DALALBYTES RESEARCH · INVESTMENT REPORT · NASDAQ: AMD

Advanced Micro Devices

The CPU War Is Won. The GPU War Is Rented.

68/100 Selective Positive. EPYC holds a record 46.2% of x86 server CPU revenue, Helios racks are running at Meta, trailing free cash flow is a record $8.4B, and the price still pays 45 times forward earnings on anchor deals carrying warrants on 320M shares. Fiscal 2029 targets: Bear $230 / Base $615 / Bull $885.

NOVEMBER 2, 2026

Reference price $504.80 (Nasdaq close, October 30, 2026). Struck on Q2 2026 reported results, one day before the Q3 2026 print of November 3, 2026. Sources: company filings and earnings releases, deal announcements, and named analyst estimates as cited. Educational purposes only, not financial advice.

1 THE VERDICT

68/100. Selective Positive. The server CPU war is effectively won, the balance sheet is net cash, and the operator is the best in the industry, but the AI GPU position is rented with warrants and equity checks, and the price still pays 45 times forward earnings printed at a peak-cycle gross margin.

What would change the score. The score rises if the Q4 2026 guide on November 3 shows Helios volume by name with Data Center tracking to more than double in 2027 and quarterly free cash flow back above $2.5B, and it falls if Helios volume slips into 2027, if a top anchor customer moves a chip generation to custom silicon, or if net cash turns to net debt.

68/100	$504.80	$824B	$230 / $615 / $885
DalalBytes score	Reference price. October 30 2026 close	Market cap at 1632M basic shares	Fiscal 2029 bear / base / bull

AMD is two companies sharing one stock. The first is a CPU franchise compounding on merit. EPYC took a record 46.2% of x86 server CPU revenue in Q1 2026, unit share reached 34.5% in Q2, lead times stretched past 30 weeks, and the 2nm Venice generation launched in July with no Intel answer on any public roadmap through 2027. The second company is an AI accelerator challenger whose seat at the frontier table is paid for in paper. The OpenAI and Meta design wins are real, and each carries a warrant on 160 million shares at a penny. The Anthropic win carries a commitment of up to $5B of AMD's own capital. Demand that a supplier has to buy is not the same asset as demand it has won, and the score keeps the two apart.

The price has done some of the work since September. The prior DalalBytes package scored AMD 72/100 at a $630.63 reference and called it priced for perfection. The stock then fell 22% from its October 6 record close of $649.42, most of the last leg on Intel's weak guide rather than any AMD news, with consensus estimates uncut through the decline. What did not improve is the quality of the earnings the multiple sits on. The 56% non-GAAP gross margin is earned at the top of a pricing cycle, the gigawatt commitments are "up to" figures rather than booked backlog, and Q2 free cash flow fell 39% sequentially as the build consumed working capital. At $504.80 the stock is a selective buy for the CPU engine and the rack option. It is not yet a full-conviction price for the GPU story.

The price basis of this report

This report uses the Friday October 30, 2026 close of $504.80 throughout, the same price carried in the DalalBytes Compute Price Reset valuation table. It is struck on Q2 2026 reported results and the Q3 2026 guide, one day before AMD reports Q3 2026 on Tuesday November 3, 2026 after the close, a date the company has announced. A post-earnings refresh of this package is scheduled for November 4, 2026.

2 THE COMPANY IN ONE PAGE

Advanced Micro Devices is a fabless semiconductor company based in Santa Clara, California, founded in 1969 and listed on Nasdaq as AMD. Dr. Lisa Su has led the company since October 2014. AMD designs its chips and manufactures them at TSMC. It reports three segments: Data Center, Client and Gaming, and Embedded. The fiscal year is the calendar year, and the latest reported quarter is Q2 2026, which ended June 27, 2026.

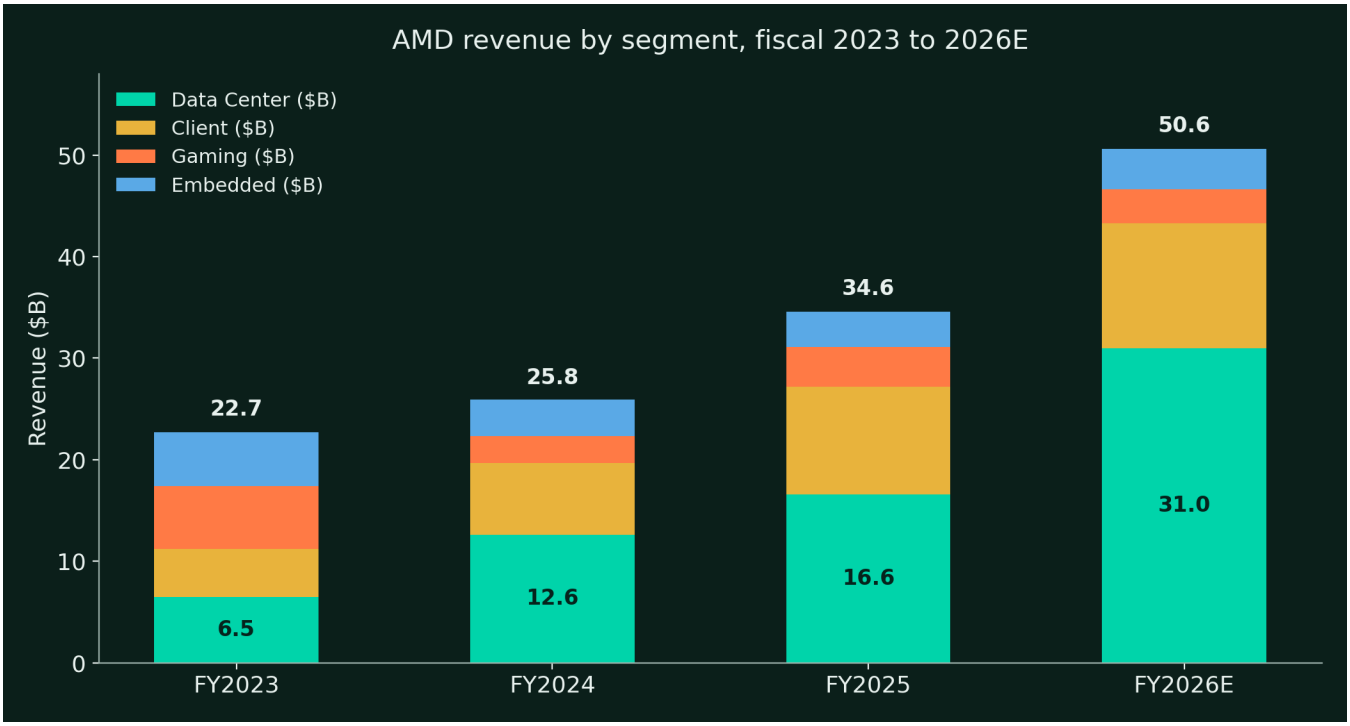
Metric	FY2024	FY2025	Q2 2026
Total revenue ($M)	25,785	34,639	11,536
Data Center revenue ($M)	12,579	16,635	6,718
GAAP diluted EPS	$1.00	$2.65	$1.38
Non-GAAP diluted EPS	$2.31	$4.17	$1.66
Free cash flow ($M)	2,405	6,735	1,560

- Data Center was 29% of revenue in fiscal 2023 and 58% in Q2 2026.
- Basic shares outstanding were 1632M at June 27, 2026.
- The company pays no dividend.
- AMD is the only merchant supplier fielding a full x86 CPU line, an AI GPU line, and data center networking in one portfolio.
- Every figure in this report traces to a filing, an earnings release, a deal announcement, or a named third-party source cited in section 17.

3 REVENUE ARCHITECTURE

Segment revenue ($M)	FY2023	FY2024	FY2025	H1 2026
Data Center	6,496	12,579	16,635	12,493
Client	4,651	7,054	10,640	5,947

Segment revenue ($M)	FY2023	FY2024	FY2025	H1 2026
Gaming	6,212	2,595	3,910	1,499
Embedded	5,321	3,557	3,454	1,850
Total	22,680	25,785	34,639	21,789



Revenue by segment. Fiscal 2026E is a DalalBytes estimate built from first-half actuals, the Q3 guide midpoint, and a Q4 step-up, against consensus revenue of $51.0B. Source: AMD results releases.

- Data Center revenue grew 94% in fiscal 2024 and 32% in fiscal 2025.
- Client revenue more than doubled from its fiscal 2023 trough of $4.7B to $10.6B in fiscal 2025.
- Gaming revenue fell from a $6.2B console-cycle peak in fiscal 2023 to $2.6B in fiscal 2024.
- Gaming recovered to $3.9B in fiscal 2025 on Radeon demand and firmer semi-custom sales.
- Embedded revenue digested from $5.3B in fiscal 2023 to $3.5B in fiscal 2025 as Xilinx customers worked through inventory.
- Embedded returned to growth in 2026 and rose 19% year over year in Q2.
- The Q3 2026 guide is revenue of about $13.0B plus or minus $300M at a non-GAAP gross margin of about 56%.
- Consensus for fiscal 2026 is revenue of $51.0B and non-GAAP EPS of $7.62 (S&P Global, November 2, 2026).

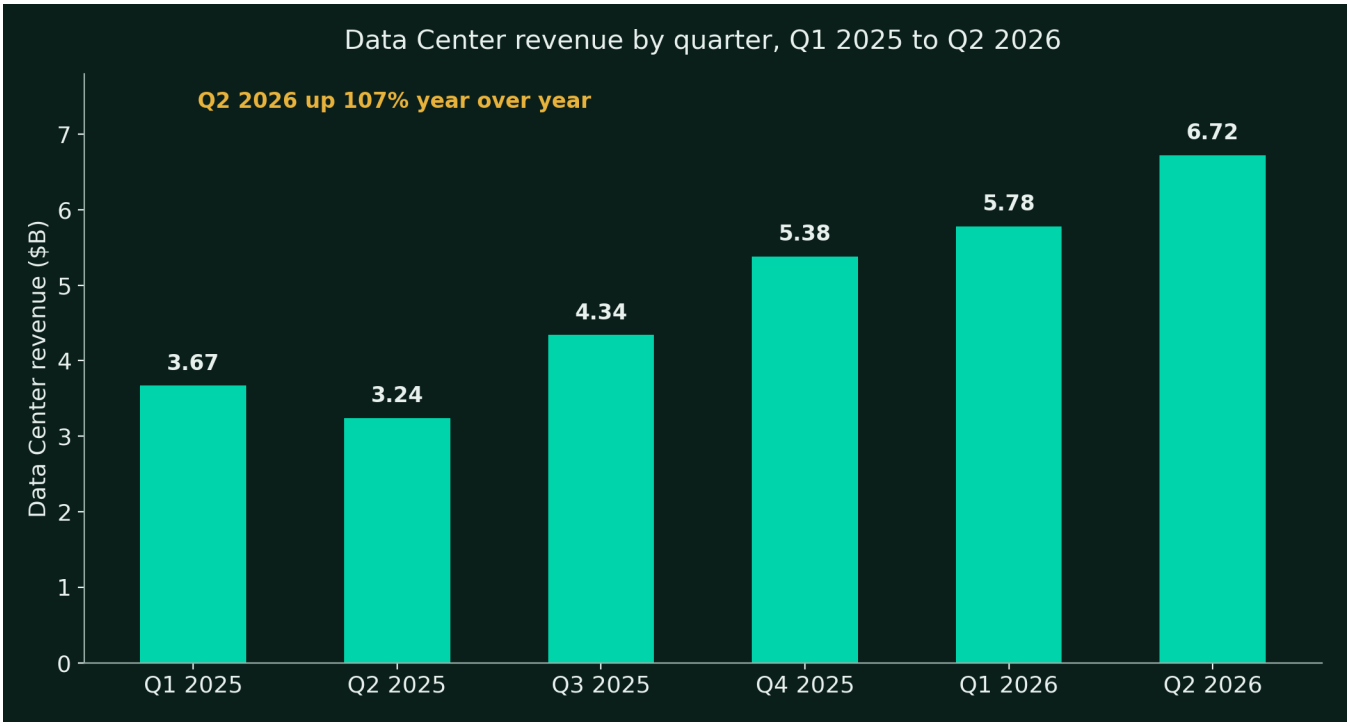
4 DATA CENTER: TWO ENGINES

EPYC server CPUs

- Mercury Research put AMD at 34.5% of x86 server CPU units in Q2 2026, up from 27.3% a year earlier and 8.9% in Q1 2021.
- AMD took a record 46.2% of x86 server CPU revenue in Q1 2026.
- Head to head against Intel Xeon SP, AMD's unit share reached 46.4% in Q2 2026, effectively parity in the premium tier.
- EPYC CPU sales grew about 70% year over year in Q2 2026 on company commentary.
- Venice, the sixth-generation EPYC on Zen 6, entered production ramp in May 2026 on TSMC's 2nm process, the first high-performance computing product to ramp on that node.
- AMD launched Venice on July 22, 2026 with up to 256 cores, 16 DDR5 memory channels at about 1.6 TB/s per socket, and PCIe Gen 6.
- AMD claims about 70% platform throughput uplift over the Turin generation. Treat the figure as a company claim pending independent tests.
- Trade reporting through 2026 described EPYC lead times past 30 weeks and price increases on selected parts.
- Management has raised its server CPU market view to about $220B by 2030, citing agentic AI demand for CPUs beside every GPU rack.
- The follow-on Verano generation adds native LPDDR5X memory support for rack-scale chassis, with Meta as the lead customer named by the company.

Instinct AI accelerators

- The MI350X and MI355X (CDNA 4, TSMC 3nm) carry 288 GB of HBM3E at 8 TB/s and have shipped since mid-2025.
- AMD launched the MI400 family on July 22, 2026, led by the MI455X with 432 GB of HBM4 and the MI430X for HPC and sovereign AI.
- The MI350P, launched in the Q2 2026 highlights, adapts CDNA 4 to a dual-slot PCIe card for standard enterprise chassis.
- AMD has never disclosed a separate Instinct revenue figure.
- Analyst estimates for AMD's 2026 AI GPU revenue run from about $8B to $12B (Cantor Fitzgerald models about $8B with upside to $10B to $12B). These are analyst estimates, not company guidance.
- Data Center segment operating income was $2.10B in Q2 2026 at a 31.3% margin, against a $155M loss in Q2 2025.
- Q2 2025 segment results carried about $800M of inventory and related charges from the MI308 export controls.
- Management guided on the Q2 2026 call for Data Center revenue to more than double in 2027.
- Management also guided for server CPU revenue growth above 80% in the second half of 2026 and above 70% in 2027.



Data Center revenue by quarter. The segment passed Client and Gaming combined in 2025 and has set a record in each of the last four quarters. Source: AMD results releases.

5 HELIOS AND THE RACK STRATEGY

- AMD acquired ZT Systems for about $4.9B in a deal that closed in March 2025.
- AMD kept roughly 1000 system design and customer-enablement engineers and divested the manufacturing arm to Sanmina in Q4 2025, a divestiture confirmed by the discontinued operations line in AMD's 2026 filings.
- Helios is the product of that team: a liquid-cooled rack with 72 MI455X accelerators, Venice EPYC CPUs, and Pensando Vulcano networking.
- Company figures put one Helios rack at about 31 TB of HBM4 and about 2.9 exaFLOPS of FP4 compute.
- Helios networking is Ethernet end to end, with UALink-over-Ethernet for scale-up at launch and Vulcano at 800G for scale-out.
- AMD launched Helios on July 22, 2026 and described it as the world's most powerful AI server rack. Treat the superlative as a company claim.
- On October 27, 2026, AMD announced that initial Helios systems are installed and being tested at Meta, with a broader deployment expected in the near future.
- The same announcement said EPYC processors power Meta's Muse AI agent and that AMD will invest $1B in Meta's equity offering.
- Microsoft will deploy Helios racks on Azure to power frontier model inference, per AMD's Q2 2026 highlights.

- Oracle is a named Helios deployment customer, with trade reporting of 50000 MI400-class GPUs from Q3 2026 in a deal reported at about $2B.
- OpenAI's first gigawatt of MI450 systems is scheduled to begin deployment in the second half of 2026.
- Meta's first gigawatt of shipments is scheduled to begin in the second half of 2026 on a custom MI450-based GPU.

6 ROCM AND THE SOFTWARE QUESTION

- ROCm is AMD's open software stack for GPU computing, now in its 7.x generation.
- AMD released ROCm.ai in 2026, an AI-native developer experience layered on the stack.
- PyTorch support is official, and the major open serving frameworks, vLLM and SGLang, treat AMD as a first-class target.
- Day-zero support for major open model releases is now routine on Instinct.
- Meta runs Instinct GPUs in production for training, ranking, and recommendation workloads, per the October 27, 2026 announcement.
- Anthropic's July 2026 partnership includes a multiyear collaboration to optimize Instinct GPUs and ROCm using Claude.
- The counter-evidence is commercial rather than technical. AMD still discloses no Instinct revenue line, so software adoption cannot be verified in dollars.
- Frontier training at the largest labs still starts on CUDA by default, and AMD is the second port, not the first.
- Third-party assessments, including SemiAnalysis, have criticized AMD's internal cluster capacity for software bring-up as a gating factor.
- The balanced read: ROCm has reached functional parity for mainstream inference serving, and the inference abstraction layer hides more of the CUDA difference every quarter. Training at hyperscale remains contested ground.

7 CLIENT, GAMING AND EMBEDDED

Client

- Client revenue was $3,062M in Q2 2026, up 23% year over year.
- Mercury Research put AMD at a record 30.3% of x86 client CPU units in Q2 2026 and 34.9% of desktop units.
- Ryzen AI parts carry the AI PC position, and Socket AM5 support now extends through 2029.
- Intel's Q3 2026 commentary on PC softness, delivered October 29, 2026, knocked AMD down 10.9% in sympathy on October 30 despite the Client segment's own growth.

Gaming

- Gaming revenue was $779M in Q2 2026, down 31% year over year on lower semi-custom sales.
- The current console generation is late cycle, and semi-custom revenue troughs before it recovers.
- Next-generation consoles are rumored to stay AMD-powered on Zen 6 and RDNA 5 with a 2027 launch window. Treat every part of that sentence as unconfirmed leaks.
- Radeon discrete GPU unit share fell from 21% in 2021 to 7% in 2025 (Jon Peddie Research), the one consumer market where AMD is losing.

Embedded

- Embedded revenue was $977M in Q2 2026, up 19% year over year, its strongest growth in over three years.
- Segment operating income was $386M in Q2 2026 at a 39.5% margin, the richest margin AMD reports.
- AMD cited more than $18B of new design wins on the Q2 2026 call.
- Versal Premium Gen 2, the Ryzen AI Embedded X100 series, and the Kria AI robotics platform all launched in 2026.
- The Xilinx acquisition price of $49B was full. The integration worked, and the segment now compounds quietly under the AI narrative.

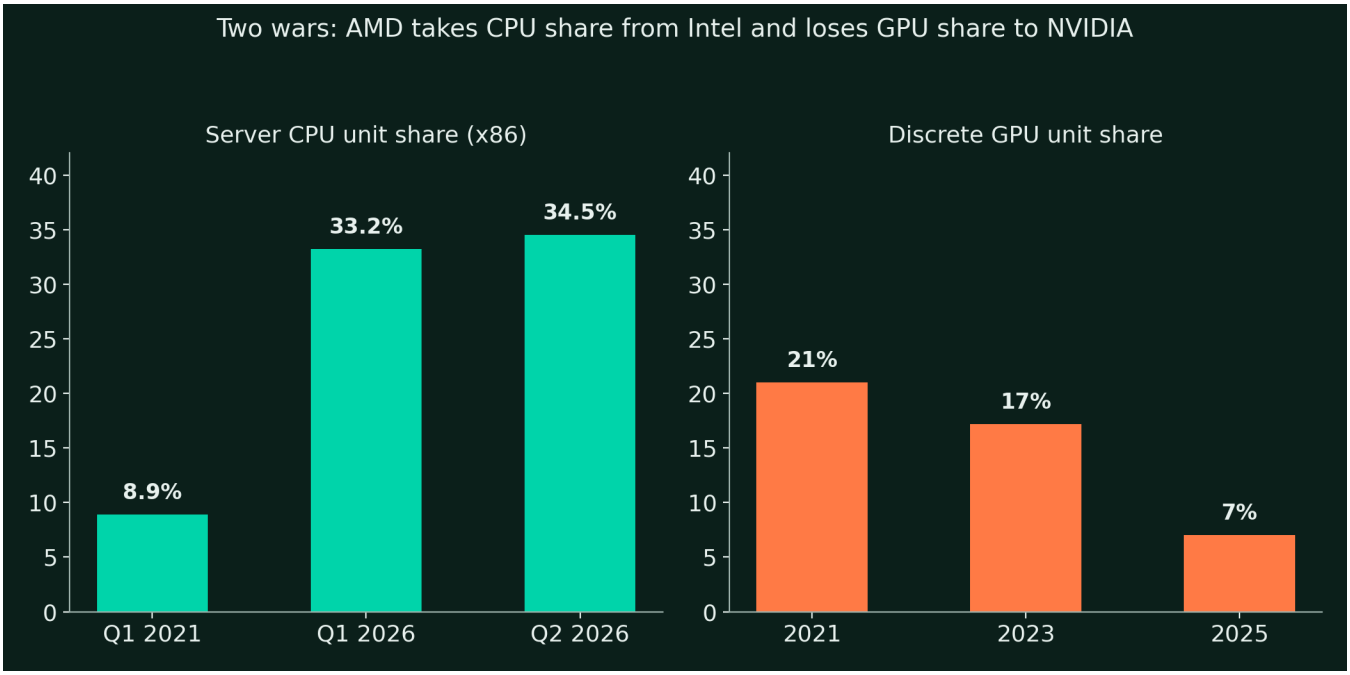
8 THE MOAT

AMD's moat is two different assets sharing a balance sheet, and grading them as one is the standard error in both directions. The EPYC moat is real and deepening. A 2nm process lead Intel cannot match on any public roadmap through 2027, 16 memory channels that make Venice the host CPU anchor of the AI rack, a record 46.2% of server CPU revenue on a third of the units, and lead times past 30 weeks with price increases are the signatures of a supplier with pricing power, not a challenger begging for sockets. Intel's available answers are a price war its foundry losses cannot fund and a process roadmap that arrives late. On CPUs, AMD no longer needs to win the argument. It needs only to keep shipping.

The Instinct position is not a moat yet. It is a window, held open by three things that can each close. The first is a genuine hardware edge in memory capacity, 432 GB of HBM4 per MI455X against a market where memory per GPU sets how large a model fits in a node. The second is the inference abstraction layer, which keeps moving workloads off CUDA-specific code. The third is the willingness of the largest buyers to fund a second source, expressed at AMD as warrants and equity checks rather than as switching costs earned through lock-in. A moat is what remains when the customer cannot leave. AMD's anchor customers can leave at every generation boundary, and two of them are building custom silicon programs of their own. The pillar grade is partial because the CPU half already passes and the GPU half, which carries the valuation, has a window where a moat would need to be.

9 COMPETITION, QUANTIFIED

Competitor	Where they meet AMD	The number that matters
NVIDIA	AI accelerators, networking, gaming GPUs	About 75% of AI accelerator revenue in 2026 estimates and 93% of discrete GPU units in 2025
Intel	Server and client CPUs	65.5% of x86 server units in Q2 2026 but falling revenue share, and a weak Q4 2026 guide after a strong Q3
Broadcom	Custom AI silicon design	AI revenue of $16.7B in Q3 fiscal 2026 and a $73B AI backlog
Arm licensees	Server CPUs from above	Arm designs took 13.6% of server CPU shipments in Q2 2026
Hyperscaler silicon	TPU, Trainium, MTIA, Maia	Custom ASIC units pass merchant GPU units in 2027 on JPMorgan shipment forecasts



Two wars, opposite records. Left: AMD's x86 server CPU unit share (Mercury Research). Right: AMD's discrete GPU unit share (Jon Peddie Research). Sources as named.

- NVIDIA's data center revenue share across the NVIDIA, AMD, and Intel peer set rose from 25% in Q1 2021 to 86% in Q4 2025 (Bloomberg via Visual Capitalist).
- AMD's share of that same set has held a 7% to 9% band since 2023, after a 17% peak in late 2022.
- Intel's Q3 2026 print showed data center and AI revenue up 78% year over year, growth that rides on Xeon host CPUs inside other vendors' GPU racks.
- Intel then guided Q4 2026 well below consensus, and the guide, not the quarter, repriced the whole CPU complex on October 30, 2026.

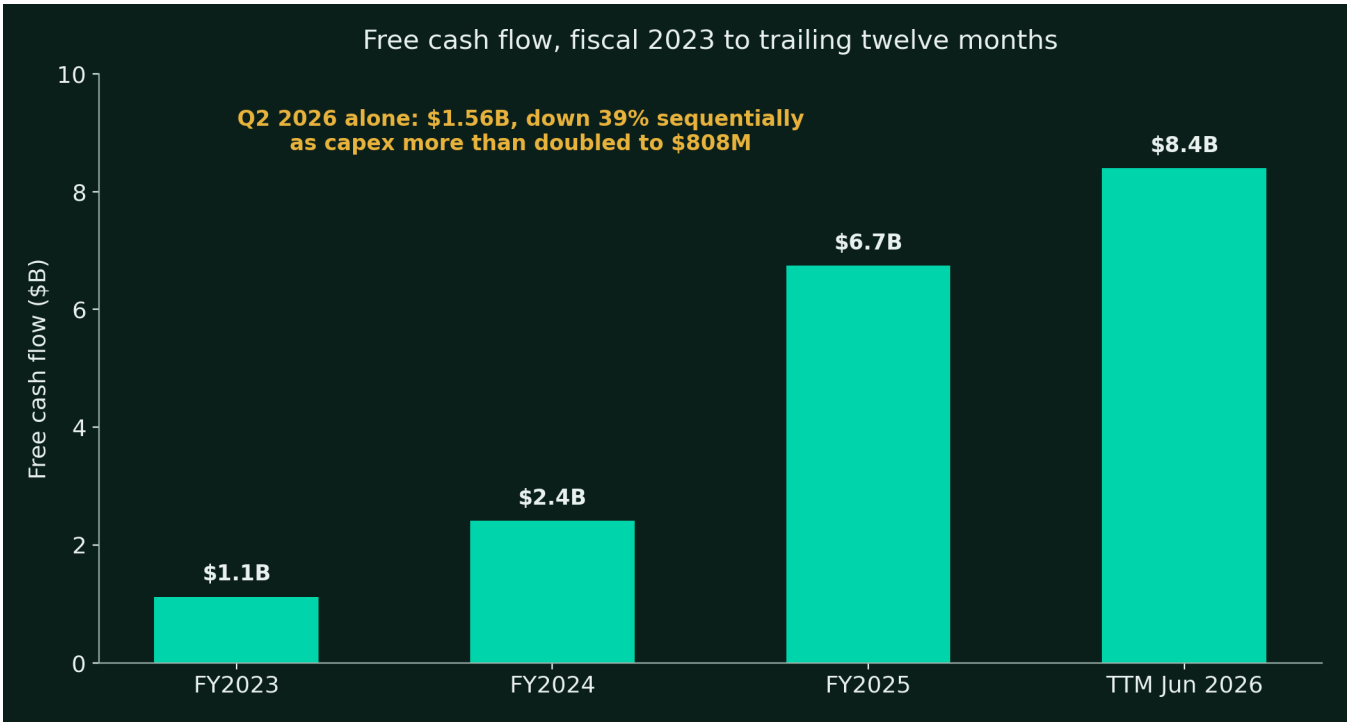
- Google now sells TPU capacity externally, with Anthropic committed to up to one million TPU chips. Every TPU gigawatt is a gigawatt no merchant vendor bids on.
- Meta's three-vendor structure is the template: its own MTIA silicon for the base load, AMD for a custom inference layer, and NVIDIA for the frontier.
- AMD's EPYC socket sits inside many custom-silicon racks as the host CPU, which keeps AMD in the tent even where its GPUs are not.

10 THE FINANCIAL STATEMENTS

GAAP results ($M)	FY2023	FY2024	FY2025	Q1 2026	Q2 2026
Revenue	22,680	25,785	34,639	10,253	11,536
Gross profit	11,402	13,671	18,183	5,416	6,203
Gross margin	50%	53%	50%	53%	54%
Operating income	401	2,086	3,741	1,476	1,990
Net income	854	1,641	4,335	1,383	2,297
Diluted EPS	$0.53	$1.00	$2.65	$0.84	$1.38
Non-GAAP gross margin	52%	53%	52%	55%	56%
Non-GAAP diluted EPS	$1.38	$2.31	$4.17	$1.37	$1.66

- Fiscal 2025 GAAP gross margin of 50% carries the MI308 charges. The non-GAAP line at 52% is the cleaner read on that year.
- Q2 2025 GAAP gross margin printed 40% with about $800M of MI308 inventory and related charges inside it.
- Stock-based compensation ran $503M in Q2 2026. Acquisition-related amortization ran $544M in the same quarter.
- The GAAP to non-GAAP gap is structural while Xilinx amortization runs, and it is roughly $1B a quarter of earnings quality to keep in view.
- Q2 2026 GAAP net income includes $483M of gains on long-term investments, per the reconciliation in the results release. Headline GAAP EPS flatters the run rate.
- Non-GAAP operating margin reached 27% in Q2 2026, up from 12% a year earlier.
- Segment operating income in Q2 2026: Data Center $2,103M, Client and Gaming $582M, Embedded $386M.
- Inventories rose to $8,468M at June 27, 2026 from $8,045M at the end of Q1. Accounts receivable rose to $7,281M.

- Free cash flow was $1.12B in fiscal 2023, $2.41B in fiscal 2024, $6.74B in fiscal 2025, and $8.40B across the trailing twelve months.
- Q1 2026 free cash flow set a record at $2.57B. Q2 2026 fell 39% sequentially to $1.56B as capex more than doubled to $808M.
- First-half 2026 operating cash flow was $5.3B against $1.2B of capex, per the Q2 2026 Form 10-Q.



Free cash flow by fiscal year and trailing twelve months, with the Q2 2026 sequential drop annotated. Source: AMD filings.

11 ROADMAP VERSUS DELIVERY AUDIT

Promise	Set when	Status at November 2, 2026
Annual Instinct cadence, MI350 then MI400	Advancing AI 2025	Delivered. The MI400 family launched July 22, 2026
Venice EPYC on TSMC 2nm in 2026	2025 product disclosures	Delivered. Production ramp began May 2026 and the family launched in July
Helios rack systems shipping in Q3 2026	2025 OCP disclosure, repeated on the Q2 2026 call	On track in test form. First systems are running at Meta as of October 27, 2026
OpenAI first gigawatt of MI450 in H2 2026	October 6, 2025	Pending. No shipment confirmation yet

Promise	Set when	Status at November 2, 2026
Meta first gigawatt shipments begin in H2 2026	February 24, 2026	Pending. Systems installed and in test at Meta
Anthropic first gigawatt in H1 2027	July 22, 2026	Pending
Data Center revenue more than doubles in 2027	August 4, 2026	Pending. This is the load-bearing guide for the multiple
Server CPU revenue up more than 80% in H2 2026	August 4, 2026	Pending. Scored first on the November 3 print
ROCm parity for mainstream inference	Multi-year claim	Partially delivered. Serving parity is evidenced, frontier training is not
More than $20 of annual non-GAAP EPS in the long-term model	Financial Analyst Day, November 2025	Pending. Fiscal 2026 consensus is $7.62

The pattern across a decade of this audit is consistent: AMD delivers silicon on the date it names, and the misses in its history are commercial, not engineering misses. Venice ramped on 2nm before any rival server part, the MI400 family arrived on the annual cadence, and Helios exists as physical systems at a named customer rather than as a slide. That record is why the market extends AMD credit on the next promise.

The promises have changed character, though. The 2025 promises were products, and products can be inspected. The 2026 promises are gigawatts, and a gigawatt converts to revenue only through a customer's construction schedule, a memory supply chain that is sold out industry-wide, and a software stack that has to hold up at a scale AMD has never operated. The audit therefore grades the roadmap in two halves. Everything AMD controls directly is delivered or on schedule. Everything that requires a customer's data center to exist on time is pending, and the pending half now carries the valuation. November 3 is the first scoring date, because the Q4 guide is where Helios volume either appears as a number or moves into 2027 as a narrative.

12 CAPITAL STRUCTURE, BUYBACKS AND WARRANTS

- Cash, cash equivalents, and short-term investments were $13,111M at June 27, 2026, a record.
- Total debt was $3,226M at June 27, 2026. The balance sheet carried about $9.9B of net cash.
- On August 17, 2026, AMD closed a $4.75B senior notes offering, the largest bond deal in its history.
- The notes are $1.25B at 4.600% due 2029, $1.50B at 5.000% due 2031, $1.00B at 5.250% due 2033, and $1.00B at 5.500% due 2036.
- Pro forma gross debt is about $8.0B against the $13.1B liquidity base. AMD remains net cash, with the new notes adding roughly $240M of annual interest.
- Buyback authorization stands at roughly $10B, including $6B added in May 2025.

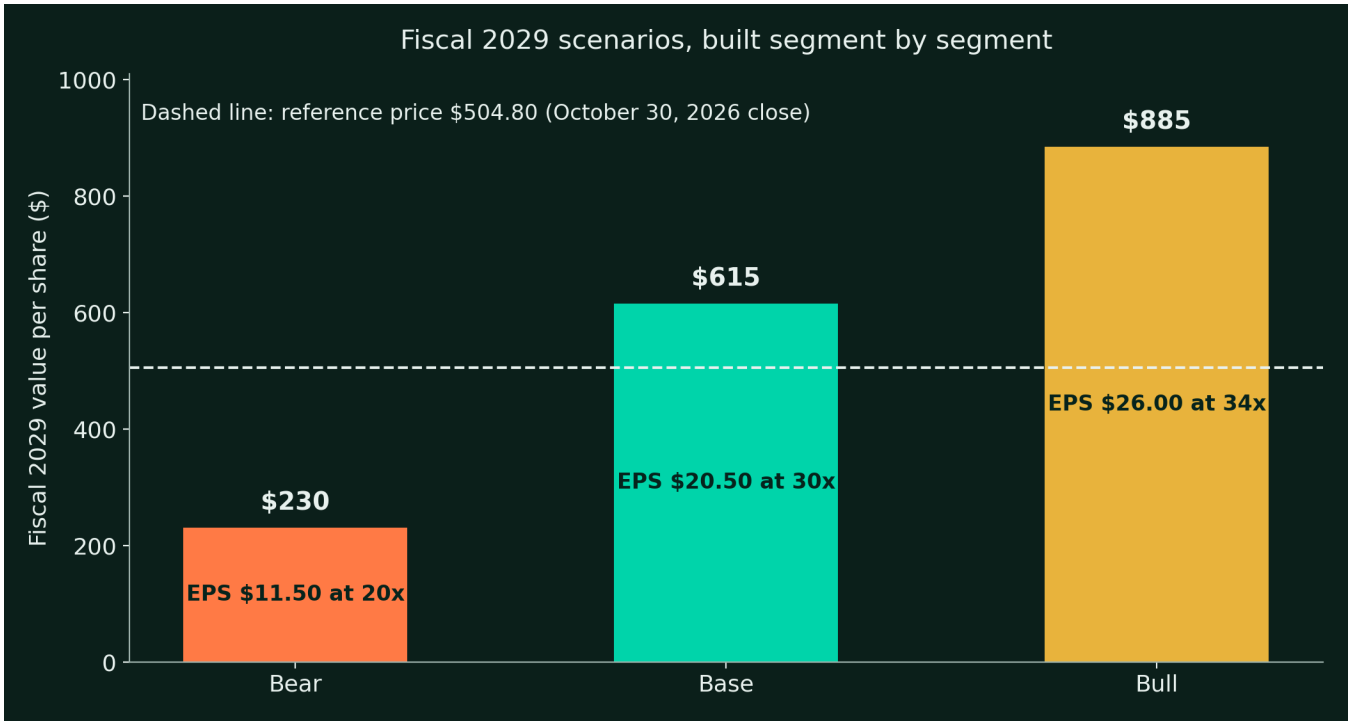
- Repurchases in the first half of 2026 were $221M, with zero open-market repurchases in Q2 2026, per the Form 10-Q. Growth capex is outranking shareholder return at the top of the cycle.
- The OpenAI warrant covers up to 160M shares at a $0.01 exercise price, vesting in tranches as deployments scale from 1 gigawatt to 6, with vesting also tied to AMD share-price targets and OpenAI technical and commercial milestones.
- The Meta warrant covers up to 160M shares at $0.01 on the same pattern, with stock-price thresholds that escalate to $600 and an exercise window running through February 2031.
- Together the two warrants cover 320M shares, about 19.6% of the 1632M basic share count, if every tranche vests.
- The Q2 2026 Form 10-Q states both warrants are classified as liabilities until conditions for equity classification are met, and that no tranche had vested at June 27, 2026.
- AMD has committed up to $5B of equity investment in Anthropic, expected to fund through fiscal 2028 subject to contingencies.
- AMD will invest $1B in Meta's October 2026 equity offering, announced October 27, 2026.
- Trade reporting in 2025 described a $300M AMD guarantee behind a Goldman Sachs loan to the neocloud Crusoe for chip purchases. Customer financing of this kind mirrors the circular structures now common across the AI supply chain.
- AMD announced the acquisition of Taalas, a Toronto AI inference silicon specialist, on August 6, 2026, and closed the MEXT memory-technology acquisition in Q2 2026.

13 VALUATION AND FISCAL 2029 SCENARIOS

- Market cap at the reference price is about $824B on 1632M basic shares.
- Trailing non-GAAP EPS is $5.76 across Q3 2025 to Q2 2026. The reference price pays about 88 times that figure.
- Fiscal 2026 consensus non-GAAP EPS is $7.62. The reference price pays about 66 times it.
- The forward multiple on next-twelve-months consensus is 45.0 times (StockAnalysis snapshot, November 2, 2026).
- Trailing free cash flow of $8.40B puts the price at about 98 times free cash flow.
- Consensus revenue for fiscal 2026 is $51.0B, which puts the price near 16 times forward sales.
- The 54-analyst consensus target is $646.74 on average and $641 at the median, with a range of $430 to $1,250 (S&P Global via StockAnalysis).

Fiscal 2029 build ($B)	Bear	Base	Bull
Data Center, EPYC CPUs	34	40	46
Data Center, Instinct GPUs	21	59	84

Fiscal 2029 build ($B)	Bear	Base	Bull
Data Center, networking and other	3	5	6
Client	12.5	15.5	18.5
Gaming	3.6	5.2	6.8
Embedded	4.6	6.0	7.5
Total revenue	79	131	169
Non-GAAP EPS	$11.50	$20.50	$26.00
Exit multiple	20x	30x	34x
Value per share	$230	$615	$885



Fiscal 2029 scenarios against the reference price. The base case delivers management's long-term EPS ambition roughly on time and still returns about 22% over three years. Source: DalalBytes estimates.

The skeptic case

The base case's hardest-working assumption is Instinct revenue reaching about $59B by fiscal 2029, from a 2026 base the company does not disclose and analysts estimate at $8B to $12B. That path needs the 14 gigawatts of "up to" commitments to convert at high rates on schedule, HBM4 supply to hold through a single primary memory partner, and rack revenue not to dilute the 56% gross margin line it is guided at today. The evidence that breaks it is specific: a Q4 2026 guide on November 3 that reaches its number on chips while Helios language moves into 2027, or any anchor customer shifting a generation to its own custom silicon.

- The bear case is a conversion failure, not a demand collapse. EPYC still compounds, the gigawatts arrive at a fraction of announced scale, and the multiple compresses to 20 times on $11.50 of earnings.
- The base case pays the 2027 Data Center doubling guide in full, then decelerates. It lands near management's own long-term EPS ambition of more than $20.
- The bull case requires AMD to bend the accelerator share line up out of its 7% to 9% band toward the low teens, with warrant dilution already inside the share count.
- One discipline runs through all three cases. The 56% gross margin is treated as a cycle-peak print that mix must defend, not a floor that price will raise, consistent with the DalalBytes Compute Price Reset analysis of where this cycle sits.

14 RISKS

The first risk is that the demand is partly purchased. Two customers anchor the entire forward Instinct story, and both were paid in warrants that dilute existing holders by up to a fifth of the company if every tranche vests. Add up to $5B committed to Anthropic's equity, $1B into Meta's offering, and a reported guarantee behind a neocloud's chip loan, and a pattern forms that the industry now shares with NVIDIA and Broadcom: suppliers financing the demand they report. None of this makes the deployments fake. Meta's racks are physically running. It does mean the quality of AMD's forward demand is lower than the quality of its forward supply, and the gap is where a repricing would start.

The second risk is physical. Every MI455X carries 432 GB of HBM4 into a memory market that is sold out industry-wide, and AMD has named Samsung as its primary HBM4 partner under a 2026 memorandum, which puts the flagship on the memory supplier that has trailed SK hynix on HBM qualification. TSMC packaging capacity is contested by NVIDIA's Rubin ramp and Broadcom's XPU programs at the same time. A 14-gigawatt conversion schedule is, in the end, a memory allocation schedule.

The third risk is China and the policy stack. China has historically been roughly a fifth of AMD's revenue, mostly client, gaming, and EPYC. The April 2025 license requirement on the MI308 produced about $800M of charges, licenses later permitted about $390M of MI308 shipments in Q4 2025, and the Q1 2026 guide included about $100M of MI308 sales. The MI450 class remains restricted. The exposure runs in both directions: licenses can widen the served market, and escalation can remove the CPU and client lines that currently flow freely.

The fourth risk is the one the score already prices. At 45 times forward earnings on a gross margin earned at the top of a pricing cycle, AMD has no cushion for a merely good outcome. The base case in this report

delivers management's own long-term earnings ambition and returns about 22% over three years. Anything less than the ambition, delivered late, meets a multiple with further to fall than the earnings do. Key-person concentration sits underneath all of it: the strategy, the customer trust, and the roadmap credibility are substantially one person's record, and no public succession discussion exists.

15 THE STORY

In October 2014 AMD was worth about $2B and was a plausible bankruptcy candidate. Lisa Su's response was a sequence of unfashionable decisions executed in order: rebuild the CPU core and accept years of low share while it matured, sell the company's manufacturing past and go all-in on TSMC, buy Xilinx for the embedded cash flows the AI story would later need, and decline to chase NVIDIA's training fortress head-on while building the one portfolio, CPU plus GPU plus networking plus rack integration, that a buyer could substitute as a system. The Zen architecture compounded quietly for five years before the market repriced it. By September 2026 the company was worth $1T, the fourth US chipmaker to reach the mark.

The story's second act is now running, and it is a different genre. The first act was a technology recovery, graded on products. The second is a capacity and credibility buildout, graded on gigawatts, and AMD is playing it against an opponent with 75% of the accelerator revenue pool and against its own customers' custom silicon at the same time. Its designated role in the industry's script is the second source, the supplier whose existence disciplines NVIDIA's pricing. That role is worth a great deal. It is not worth any price, and it converts to a first-source economics only if the share line in section 9 bends. The next four quarters, starting with the November 3 guide, are where the story either changes genre or confirms the role.

16 PRICE BASIS AND TECHNICALS

- Reference price $504.80, the Friday October 30, 2026 Nasdaq close.
- The stock fell 10.9% on October 30, 2026 in sympathy with Intel's weak Q4 guide and PC commentary.
- The October decline was 17.5% for the month, against a record close of $649.42 on October 6, 2026. The reference price sits about 22% below that record.
- No consensus estimate cuts were found through the October decline.
- AMD reports Q3 2026 on Tuesday November 3, 2026 after the close, with a call at 5:00 PM ET. The company has announced the date.
- The Q3 bar is the company's own guide: revenue of about $13.0B plus or minus $300M and a non-GAAP gross margin of about 56%.
- Compiled consensus sits almost exactly on the guide, at revenue of about $13.06B to $13.09B and adjusted EPS of about $1.95.
- The consensus path implies a Q4 2026 revenue guide in the region of $16.2B. Treat that figure as the street's working expectation, not a company number.
- This package supersedes the September 28, 2026 DalalBytes package, which scored AMD 72/100 at a $630.63 reference price and used a 35 times exit multiple in its base case.

- A post-earnings refresh of this package is scheduled for November 4, 2026.

17 SOURCES AND DISCLOSURES

- AMD Form 10-K for fiscal 2025 and Form 10-Q for Q2 2026, including segment results, balance sheet detail, and the warrant classification note.
- AMD results releases: Q4 and full year 2025 (February 3, 2026) and Q2 2026 (August 4, 2026), including the GAAP to non-GAAP reconciliations and the Q3 2026 outlook.
- AMD deal announcements: OpenAI partnership (October 6, 2025), Meta expanded partnership (February 24, 2026), Anthropic partnership (July 22, 2026), and Helios systems at Meta (October 27, 2026).
- AMD Form 8-K and prospectus supplement for the $4.75B senior notes offering (August 2026).
- Mercury Research server and client CPU share for Q1 and Q2 2026, as reported in trade press.
- Jon Peddie Research discrete GPU unit share, and Bloomberg data via Visual Capitalist for data center revenue share, both as compiled in the DalalBytes Compute Share Wars article.
- S&P Global analyst consensus via StockAnalysis, snapshot of November 2, 2026.
- DalalBytes Compute Price Reset (November 1, 2026) for the cycle and margin framing, and the DalalBytes AMD Q3 2026 earnings preview scorecard for the November 3 bar.
- Disclosures: the author is long AMD. This report is research and opinion, not investment advice.