AMD gideon lens

Advanced Micro Devices, Inc. (AMD) - Gideon-Lens Equity Research

Date: 2026-04-29. Analyst: Codex ER2. Price used: $347.81. Market cap from queue snapshot: $567.1B. Fiscal year end: 52/53-week year ending the last Saturday in December.

VERDICT BLOCK

Verdict: PASS

Scorecard:

Gate	Result	Notes
Management	PASS	Su-led execution strong; governance clean but warrant dilution watch.
Moat	YELLOW	2046 test plausible, but AI accelerator moat still contested.
P/E <= 15 on confident EPS at TTM + 5 calendar years	FAIL	Base cushion -65.9%; bull also fails.
Compounder Score	Audited 46 / 100	46 / 100; Reject under the diagnostic rubric because valuation fails completely and the remaining business-quality evidence is capped by low reported ROIC, moderate acquisition-inclusive ROIIC, undisclosed AI economics, and warrant dilution uncertainty.
Business Quality Score	Audited 46 / 80	46 / 80; AMD has strong management, net cash, real franchises, and a large AI runway, but the evidence does not yet prove elite organic ROIIC or durable Data Center pricing power.
20% Return UW Score	Audited 0 / 20	0 / 20; base, bear, and bull cases do not clear the 20% hurdle at the $347.81 anchor without prohibited terminal multiple dependence.

One-line thesis (or reason to pass): AMD is a well-managed AI compute contender, but the current price requires far more durable earnings power than the reviewed base or bull cases support by the 2030-12-27 horizon.

Kill / trigger criteria:

- BUY trigger: price below roughly $119 on base-case FY+5 EPS of $7.92, before applying any additional margin of safety for the YELLOW moat gate.
- PASS triggers: Data Center FY2026 tracking below a $24B annual run-rate; Q1+Q2 Data Center revenue below $11.5B with margin below 25% and no credible H2 bridge; warrant dilution from OpenAI/Meta without matching GAAP earnings.
- Upgrade to BUY conditions: Data Center proves $50B-plus revenue with high-20s margin durability, Client/Gaming stays profitable through the cycle, Embedded resumes growth, and price falls to a P/E 15 entry on confident EPS.

Position sizing if bought: 0% today. Revisit only if price, Data Center margin evidence, and share dilution economics reset together.

Stage 0 - Time anchor

Today is 2026-04-29. AMD uses a 52/53-week fiscal year ending on the last Saturday in December; FY2025 ended 2025-12-27. The most recently reported quarter is Q4 FY2025, reported on 2026-02-03, and Q1 FY2026 has not been reported as of this run. The TTM anchor is Q1 FY2025 + Q2 FY2025 + Q3 FY2025 + Q4 FY2025, ending 2025-12-27. The forward horizon is TTM + 5 calendar years, or 2030-12-27, and every forward "FY+5" reference in this report uses exactly that five-year span.

Stage 1 - Disqualifier sweep

ROIC is yellow because GAAP returns are distorted by acquisition intangibles and AI-cycle reinvestment, so the economics must be proven rather than assumed. FCF conversion is green: FY2025 free cash flow was $5.519B versus GAAP net income of $4.335B. Net debt / EBITDA is green because AMD had $10.6B of cash, cash equivalents, and short-term investments versus $3.3B of total debt. End-markets are green, led by AI data center compute, EPYC server CPUs, and client processors, though Embedded remains mixed. Governance is green/yellow: no obvious related-party or control-structure disqualifier emerged, but large customer warrants and CEO awards require monitoring. AI disruption exposure is green/yellow: AMD is an AI beneficiary, but NVIDIA platform power and hyperscaler custom silicon pressure the moat. Pricing power is yellow because AMD has pockets of premium positioning, but semiconductors remain cyclical and competitive.

Management - one paragraph summary

Management passes. Lisa Su has an eleven-plus-year record of reshaping AMD from a financially constrained CPU challenger into a credible high-performance compute supplier. The management appendix credits the team for coherent capital allocation: heavy R&D, targeted M&A such as Xilinx and ZT Systems, and willingness to partner aggressively where AMD needs ecosystem scale. The gate is not risk-free. OpenAI and Meta warrant structures can dilute shareholders if the revenue ramp does not arrive with matching earnings power, and the one-time CEO value-creation award needs continued scrutiny. Disclosure quality is good enough for a PASS because management has been explicit about the MI308 export-control inventory charge, the Q1 FY2025 reporting change, and the scale of AI customer agreements. The management case is therefore not the reason to pass on the stock; it is one of the better parts of the thesis.

Moat - one paragraph summary

The moat gate is YELLOW. AMD has real advantages: EPYC server CPUs benefit from performance credibility and qualification cycles, Embedded carries Xilinx design-in switching costs, and Client has product and channel scale. But the blended moat is weaker than the best segment. Data Center is the decisive swing factor, and the Instinct/ROCm/rack-scale ecosystem still has to prove it can earn durable economics against NVIDIA and hyperscaler internal silicon. Embedded has the cleanest 2046 durability, but it is only 10% of FY2025 revenue. Client and Gaming add cyclical exposure and weaker pricing power. The 2046 test is plausible, not proven: a future AMD selling CPUs, accelerators, adaptive SoCs, and rack-scale systems can still make money, but the route depends on maintaining relevance in AI platforms rather than merely serving as a second source.

Stage 5 synthesis - headline numbers

The synthesis passed all 11 arithmetic tie-out checks and failed the valuation gate. Base-case TTM+5y revenue is $73.156B, operating income is $14.793B, net income is $13.000B, and EPS is $7.92. At the $347.81 price used from the queue snapshot, the required EPS for a 15x P/E is $23.19. Base-case EPS x 15 is $118.74, leaving a -65.9% cushion. The bull case also fails: bull EPS is $15.84, worth $237.67 at 15x, still -31.7% versus the current price. Bear EPS is $2.30, which shows how much downside remains if the AI ramp does not translate into durable margins.

Case	Revenue	OI	NI	EPS	EPS x 15	vs. $347.81
Bear	$47.595B	$4.185B	$3.771B	$2.30	$34.45	-90.1%
Base	$73.156B	$14.793B	$13.000B	$7.92	$118.74	-65.9%
Bull	$110.451B	$29.761B	$26.023B	$15.84	$237.67	-31.7%

The segment build is the core evidence. Data Center is modeled at $49.090B of base revenue and $13.991B of operating income, making it the dominant driver of consolidated EPS. Client and Gaming contributes $18.991B of revenue and $3.987B of operating income in base, while Embedded contributes $5.075B and $1.929B. The synthesis includes $5.114B of horizon unallocated corporate expense, which prevents the segment rollup from overstating consolidated earnings. The share count path ends at 1,642.3M diluted shares after modest buybacks and SBC issuance. This is not a close valuation miss created by conservative accounting; even the bull case does not clear the P/E 15 hurdle.

Red team - one paragraph

The red team agrees that the arithmetic survives and attacks the interpretation. Its strongest bear argument is that AMD may be priced for clean AI platform economics while actually operating as a hyperscaler second-source supplier with capped margins. The most likely optimistic Stage 5 number is Data Center TTM+5y operating income of $13.991B, because it

requires both $49.090B of revenue and a 28.5% segment margin. The monitorable exit signal is specific: FY2026 Data Center tracking below a $24B annual run-rate, or Q1+Q2 below $11.5B with margin below 25% and no second-half bridge. If Data Center moves from base toward the segment bear case, red team estimates base EPS could fall from $7.92 to roughly $3.76.

Kill / trigger criteria

- Keep as PASS while the stock trades far above the $119 base-case 15x value and while the moat gate remains YELLOW.
- Reopen only if Data Center quarterly revenue and margin prove the front-loaded ramp, especially through FY2026.
- Treat customer warrants as acceptable only if revenue conversion brings GAAP earnings, not merely headline commitments.
- Upgrade only with a much lower entry price or a revised, source-backed earnings model that clears $23.19 of confident FY+5 EPS.

Analyst flags

AMD does not disclose annual EPYC, Instinct, or Embedded product-level revenue, so segment appendices model at the reportable-segment level where required and flag missing product detail rather than fabricating it. Client and Gaming component operating income is not separately disclosed. Data Center is the key sensitivity, and the model depends on a large AI revenue ramp plus margin expansion. The moat appendix is YELLOW, not PASS, because AMD's AI platform durability is still being proven. No BUY alert was triggered because the final verdict is PASS.

Appendix A - Company Primer

Advanced Micro Devices, Inc. (AMD) - Primer

Date: 2026-04-29

Source Discipline And Time Anchor

This primer is a reader-facing orientation document for Advanced Micro Devices, Inc. It is not the controlling model source for downstream segment forecasts. Historical financial numbers come from `quant_foundation.md`, which reconciles live FMP pulls to AMD's FY2025 Form 10-K, Q1-Q3 FY2025 Forms 10-Q, and the Q4/FY2025 earnings release. The data package was present when this primer was written, but it was still an in-progress placeholder, so qualitative context here relies on the quant foundation and live primary sources rather than a completed local data package. The most important primary sources used for context are AMD's FY2025 Form 10-K and the SEC-filed Q4/FY2025 earnings release.

Stage 0 anchor: Today is 2026-04-29. AMD uses a 52/53-week fiscal year ending on the last Saturday in December; FY2025 ended 2025-12-27. The most recently reported quarter is Q4 FY2025 / fiscal year ended 2025-12-27, released 2026-02-03. Q1 FY2026 has not yet been reported as of 2026-04-29. The TTM anchor is Q1 FY2025 + Q2 FY2025 + Q3 FY2025 + Q4 FY2025, producing TTM ending 2025-12-27. The forward horizon is TTM + five calendar years, or 2030-12-27, with an exactly five-year compounding span.

Primary source links used in this primer:

- AMD FY2025 Form 10-K: https://www.sec.gov/Archives/edgar/data/2488/000000248826000018/amd-20251227.htm
- AMD Q4/FY2025 earnings release, Exhibit 99.1: https://www.sec.gov/Archives/edgar/data/2488/000000248826000014/q42025991.htm
- AMD Q1 FY2025 Form 10-Q: https://www.sec.gov/Archives/edgar/data/2488/000000248825000047/amd-20250329.htm
- AMD Q2 FY2025 Form 10-Q: https://www.sec.gov/Archives/edgar/data/2488/000000248825000108/amd-20250628.htm
- AMD Q3 FY2025 Form 10-Q: https://www.sec.gov/Archives/edgar/data/2488/000000248825000166/amd-20250927.htm

1. Business Overview

AMD is a fabless semiconductor designer focused on high-performance and adaptive computing. It designs CPUs, GPUs, AI accelerators, adaptive SoCs, FPGAs, DPUs, AI network interface cards, semi-custom system-on-chip products, chipsets, and associated software stacks. The practical customer promise is compute performance, energy efficiency, and platform flexibility across three broad environments: cloud and enterprise data centers, endpoint devices such as PCs and gaming systems, and embedded systems used in industrial, automotive, aerospace, communications, healthcare, test, measurement, broadcast, and networking markets.

The business is not a generic chip merchant. In data center, AMD sells the compute engines that hyperscale cloud companies, enterprise server OEMs, supercomputing customers, AI infrastructure buyers, and system integrators use to run workloads. EPYC server CPUs compete in general-purpose x86 server compute. Instinct GPUs and accelerators address AI training, inference, and high-performance computing. Pensando-related DPUs and AI NICs, adaptive SoCs, and FPGAs broaden the portfolio into networking, offload, and reconfigurable acceleration. The strategic direction in FY2025 moved toward rack-scale AI systems: AMD retained the design business from ZT Systems and sold the manufacturing business to Sanmina, giving AMD more systems-design capability while avoiding full ownership of a lower-margin manufacturing operation.

In client and gaming, AMD sells processors and graphics products into PCs, workstations, game consoles, add-in boards, and semi-custom designs. The customer base includes PC OEMs, independent distributors, add-in-board manufacturers, motherboard ODMs, game console customers, system integrators, retailers, and in some cases direct customer relationships. Ryzen CPUs and APUs are the core client products. Radeon GPUs and semi-custom SoCs are the core gaming products. Semi-custom console revenue is project-based and depends on the success, timing, and refresh cadence of customer platforms; client CPU/APU revenue depends more on PC unit demand, mix, share, and average selling prices.

In embedded, AMD sells adaptive and embedded compute platforms. This is the Xilinx legacy footprint inside AMD: FPGAs, adaptive SoCs, SOMs, embedded CPUs/APUs, and domain-specific tools. The customer problem is different from a PC refresh or cloud capacity build. Embedded customers often design AMD products into long-lived systems that require qualification, reliability, support, and predictable supply. The end markets include aerospace and defense, automotive, industrial automation, healthcare, communications infrastructure, test and measurement, emulation, prototyping, audio/video/broadcast, storage, networking, and edge AI. Sales can be direct or through distributors and OEM partners, with substantial design support from AMD application engineers.

AMD sells through a combination of direct sales, OEMs, ODMs, independent distributors, sales representatives, add-in-board partners, system integrators, and value-added channels. Its arrangements generally operate with product forecasts, purchase orders, and contractual agreements rather than blanket minimum-purchase commitments. This matters for quality of revenue: AMD has significant strategic customer relationships and platform design wins, but the model remains exposed to customer forecast changes, inventory digestion, platform cycles, capacity allocation, and semiconductor cyclicality.

AMD does not own leading-edge wafer fabs. It relies on third-party foundries, primarily TSMC for 7 nm and smaller microprocessor and GPU wafers, and also uses GLOBALFOUNDRIES, UMC, Samsung, and ATMP partners. This fabless model converts fixed manufacturing intensity into supply-chain dependence. It gives AMD access to leading process technology without owning fabs, but it also means gross margin, product timing, and customer service depend on external wafer, packaging, substrate, memory, and back-end capacity availability. For an investor, AMD is best understood as a high-end design, architecture, software, and customer-platform company whose physical production is externally manufactured.

2. Revenue Architecture

AMD's FY2025 / TTM revenue was $34.639 billion. The reportable segment architecture for the current analysis is Data Center, Client and Gaming, and Embedded. Beginning in Q1 FY2025, AMD combined its prior Client and Gaming segments into one reportable segment and retrospectively adjusted prior periods. Within that combined segment, AMD still discloses Client and Gaming revenue components, but the downstream modeling segment is the combined Client and Gaming segment. The standalone Gaming line must not be double-counted.

Segment	FY2025 / TTM revenue	Share of total revenue	FY2025 operating income	FY2025 operating margin
Data Center	$16.635B	48.0%	$3.603B	21.7%
Client and Gaming	$14.550B	42.0%	$2.855B	19.6%
Embedded	$3.454B	10.0%	$1.243B	36.0%
All Other	n/a	n/a	$(4.007)B	n/a
Total AMD	$34.639B	100.0%	$3.694B	10.7%

Data Center is the largest segment and the current center of gravity for the equity story. Its FY2025 revenue was $16.635 billion, up from $12.579 billion in FY2024 and $6.496 billion in FY2023. AMD attributed FY2025 Data Center growth primarily to strong demand for EPYC processors and Instinct GPU accelerators. The segment includes server CPUs, GPUs, AI accelerators, DPUs, AI NICs, FPGAs, and adaptive SoCs sold into hyperscale cloud, enterprise, supercomputing, and AI infrastructure environments. This segment has a mix of established CPU share gains and earlier-stage AI accelerator opportunity. The CPU portion benefits from x86 compatibility, platform performance, power efficiency, and data center total cost of ownership. The GPU/AI accelerator portion depends on performance, software maturity, memory and packaging supply, system availability, and customer willingness to diversify away from Nvidia-dominated AI infrastructure.

Client and Gaming generated $14.550 billion in FY2025 revenue. AMD separately disclosed $10.640 billion of Client revenue and $3.910 billion of Gaming revenue inside that combined segment. Client includes CPUs, APUs, and chipsets for desktop and notebook PCs. Gaming includes discrete GPUs and semi-custom SoCs, including console-related designs and development services. The segment grew 51% year over year in FY2025 from $9.649 billion in FY2024. AMD attributed Client strength to Ryzen desktop and mobile processor demand, unit growth, richer mix, and ASP gains. Gaming strength came from higher semi-custom revenue and Radeon GPU demand. This is a cyclical segment: PCs can swing with replacement cycles and channel inventory, while semi-custom console revenue tends to move with customer platform timing and console cycle maturity.

Embedded generated $3.454 billion in FY2025 revenue, down 3% from FY2024 and down materially from $5.321 billion in FY2023. The segment includes embedded CPUs/APUs, FPGAs, SOMs, and adaptive SoCs. Revenue pressure in FY2025

reflected mixed demand and customer inventory adjustments earlier in the year, not a collapse in the strategic rationale. Embedded is smaller than Data Center and Client and Gaming, but it is structurally important because it carries the highest segment operating margin and broadens AMD beyond commodity PC and server cycles. Embedded products often involve longer design cycles, customer-specific integration, and sticky qualification requirements.

Geographically, AMD is global and U.S.-dollar denominated. FY2025 sales by customer billing location were $11.363 billion in the United States, $7.751 billion in China including Hong Kong, $5.186 billion in Taiwan, $4.284 billion in Singapore, and $6.055 billion in other regions. International sales were 67% of FY2025 revenue, roughly consistent with 66% in FY2024. No customer represented at least 10% of consolidated net revenue in FY2025 or FY2024, although AMD had a major Client and Gaming customer representing 18% in FY2023. The revenue footprint therefore combines global exposure, China/export-control sensitivity, and some customer concentration by product cycle even where no single customer crosses the 10% disclosure threshold in the latest year.

3. Unit Economics

AMD's consolidated FY2025 gross profit was $17.152 billion on $34.639 billion of revenue, implying 50% GAAP gross margin. GAAP operating income was $3.694 billion, or 10.7% of revenue. Non-GAAP operating income was substantially higher at $7.768 billion, but the ER2 framework should treat GAAP and cash conversion carefully because acquisition-related amortization, stock-based compensation, export-control inventory charges, and tax effects are economically relevant even if some are non-cash or unusual. FY2025 free cash flow from continuing operations was $5.519 billion, calculated as $6.493 billion of continuing operating cash flow less $974 million of capital expenditures. This cash-flow profile is stronger than GAAP operating margin alone suggests because AMD is fabless and does not carry leading-edge fab capex.

A dollar of Data Center revenue turns into operating profit through product mix, wafer and packaging cost, memory/substrate availability, accelerator ramp efficiency, software investment, and customer scale. In FY2025, Data Center revenue of $16.635 billion produced $3.603 billion of segment operating income, a 21.7% margin. That margin was only slightly above Client and Gaming despite Data Center scale because AMD is investing heavily in AI and absorbed approximately $440 million of net inventory and related charges tied to U.S. government export control on AMD Instinct MI308 products. Management also highlighted higher cost of sales and higher operating expenses in Data Center. The key economic point is that Data Center can be a high-value segment, but AI accelerator ramps are not pure software economics. They require expensive silicon, advanced packaging, memory, systems validation, software enablement, and customer qualification.

A dollar of Client and Gaming revenue is more exposed to product cycles and ASP swings. FY2025 segment revenue of $14.550 billion produced $2.855 billion of operating income, a 19.6% margin. Client CPU/APU revenue benefits from scale, better mix, and desktop/mobile share gains, but it remains tied to PC demand and competition with Intel and Arm-based alternatives. Gaming revenue has two different economic profiles. Discrete GPUs compete against Nvidia and Intel with performance, software, price, channel, and brand factors. Semi-custom SoCs are more project-like: AMD receives non-recurring engineering fees and product revenue tied to customer platforms, but volumes depend on console or customer device success. As a result, this segment can show strong operating leverage in upcycles and sharp pressure when channel inventory or platform cycles turn.

A dollar of Embedded revenue has the best segment operating margin in FY2025. The segment generated $1.243 billion of operating income on $3.454 billion of revenue, or 36.0% margin. The likely drivers are design-in stickiness, differentiated adaptive compute IP, customer qualification costs that discourage switching, and a broad industrial/aerospace/communications/customer-specific base. The downside is slower cyclicality and inventory adjustments: FY2025 revenue declined 3%, and FY2024 was already below FY2023. Embedded's margin quality is high, but its growth path depends on end-market recovery, FPGA/adaptive SoC competitive positioning, and AMD's ability to attach embedded products to AI-at-the-edge use cases.

The consolidated operating model includes a large All Other burden. In FY2025, All Other operating loss was $4.007 billion, primarily amortization of acquisition-related intangibles and stock-based compensation. This accounting structure means segment operating income sums look much better than consolidated GAAP operating income. In a model, an analyst cannot simply extrapolate segment margins and ignore All Other; at minimum, stock-based compensation and acquisition amortization affect either dilution, cash compensation, or acquisition economics. AMD's FY2025 non-GAAP operating margin of 22% shows the earnings power before selected exclusions, while GAAP operating margin of 11% shows the current fully burdened income statement.

Capital intensity is moderate relative to integrated semiconductor manufacturers. Capex was $974 million in FY2025, only 2.8% of revenue, because AMD relies on external foundries and ATMP partners. But economic capital intensity is not zero.

AMD can have prepayments, long-term purchase commitments, inventory risk, packaging constraints, and working capital swings. It also must invest heavily in R&D: FY2025 R&D expense was $8.091 billion, up 25% year over year, driven by employee-related costs supporting AI strategy. In semiconductor design, R&D is the economic engine. A weaker AMD would cut R&D and fall behind; a stronger AMD must keep spending to match Nvidia, Intel, Arm ecosystem alternatives, custom silicon programs, and embedded competitors.

4. How AMD Wins Customers

AMD wins when it offers customers a better combination of performance, power efficiency, openness, total cost of ownership, platform availability, and roadmap credibility. In server CPUs, the historical win condition has been straightforward: EPYC must deliver enough performance-per-watt, core count, memory bandwidth, platform features, and price-performance advantage for cloud and enterprise customers to qualify AMD platforms alongside or instead of Intel. The more standardized the workload and the more customers care about TCO, the easier it is for AMD to win if the silicon is strong.

In AI accelerators, the win condition is harder. The chip is necessary but not sufficient. Customers need systems, networking, memory capacity and bandwidth, compiler/runtime support, model framework compatibility, libraries, debugging tools, supply availability, and confidence that the roadmap will keep pace. AMD's strategy is to pair Instinct GPUs with EPYC CPUs, Pensando networking assets, ROCm software, and systems-design capability from ZT's retained design business. The OpenAI product purchase agreement disclosed in the 10-K, involving planned deployment of 6 gigawatts of AMD GPUs beginning with MI450-series products, is strategically significant because it signals that AMD is being considered for very large AI infrastructure deployments. The investment question is whether those design wins convert into durable gross profit and ecosystem traction, not merely headline revenue.

In PCs, AMD wins through OEM design slots, product performance, battery life, integrated graphics, AI PC features, channel availability, and brand perception. The company was early to integrate neural processing units into x86 client SoCs, and its Ryzen AI roadmap is intended to keep AMD relevant as AI features move onto local devices. However, PC buyers are fragmented and price-sensitive. OEMs often dual-source across Intel and AMD. Arm-based platforms can pressure the x86 value proposition over time. Client strength is valuable but should not be underwritten as a secular monopoly story.

In gaming, AMD's win condition depends on the submarket. In semi-custom consoles, the critical factor is whether AMD secures the platform design and whether the end device sells well over a multi-year cycle. In discrete GPUs, AMD must compete with Nvidia's brand, driver reputation, software ecosystem, and performance leadership. AMD can still win at price-performance points, console-related scale, and specific segments, but discrete gaming GPU economics are less protected than embedded design-ins or server CPU share gains.

In embedded, AMD wins by helping customers design complex systems faster and with more flexibility. FPGAs and adaptive SoCs create value when the customer needs hardware acceleration, reconfigurability, low latency, reliability, and long product life. The sale often involves development tools, reference designs, IP, and application-engineer support. This is a more consultative and sticky motion than selling a standard CPU through retail or channel distribution. The tradeoff is a longer sales cycle and exposure to industrial and communications inventory corrections.

5. Competitive Landscape

AMD operates in some of the most competitive markets in technology. In Data Center CPUs, the principal competitor is Intel, with Arm-based server CPUs and internal hyperscaler silicon as additional threats. In Data Center GPUs and AI accelerators, the principal competitor is Nvidia, with a very strong proprietary software ecosystem and market position. AMD also competes against custom accelerator programs inside major cloud customers. In FPGAs and adaptive compute, AMD competes with Altera and other programmable logic providers. In Client CPUs/APUs, Intel remains the main x86 competitor, and Arm-based PC platforms are a long-term architectural risk. In discrete graphics, Nvidia is the market share leader, with Intel also participating. In Embedded, AMD competes against FPGA vendors, ASSP vendors, ASIC alternatives, embedded CPUs, DSPs, and general-purpose GPUs targeting automotive and data center applications.

The competitive factors are not just price and speed. AMD's 10-K lists performance, total cost of ownership, timely product introductions, quality, reliability, feature set, energy efficiency, form factor, selling price, adherence to standards, software and hardware compatibility, software-tool completeness, security, stability, brand recognition, and availability. These factors matter because semiconductor customers buy platforms, not isolated chips. A nominally superior chip can fail if supply is constrained, software is immature, customer qualification is slow, or the next-generation roadmap slips.

The bull case rests on AMD being a credible second source and innovation leader across CPUs and AI accelerators while maintaining a broad embedded and client base. Customers often want credible alternatives to a dominant supplier, especially in AI infrastructure where capacity, cost, and vendor concentration are major considerations. AMD's open ecosystem positioning around ROCm can help if customers prefer flexibility over Nvidia lock-in. The bear case is that Nvidia's software and system lead in AI remains too large, Intel stabilizes in CPUs, internal customer silicon absorbs incremental AI workloads, and AMD's AI revenue grows without matching margins.

6. Supply Chain, Manufacturing, And Capital Needs

AMD's fabless model is central to its economics. The company uses TSMC for production of wafers for HPC, FPGA, adaptive SoC, and leading-node microprocessor and GPU products; it relies primarily on GLOBALFOUNDRIES for certain larger-node products; and it uses UMC and Samsung for some programmable logic ICs. Wafers are sorted by foundries or sent to assembly, test, mark, and packaging partners, including Asia-Pacific ATMP partners and joint ventures related to Tongfu Microelectronics. Board-level graphics manufacturing is outsourced.

This operating model reduces capital expenditure compared with an integrated device manufacturer, but it introduces bottlenecks. Advanced AI accelerators require leading-edge wafers, high-bandwidth memory, substrates, advanced packaging, and system-level integration. If TSMC, ATMP partners, memory suppliers, substrate suppliers, or other critical vendors cannot meet AMD's needs, AMD may have to allocate scarce product among customers, miss demand, or accept higher costs. The 10-K explicitly warns that TSMC capacity at 7 nm and smaller nodes is particularly important for microprocessor and GPU products. The company also discloses risks around limited suppliers for materials, substrates, packages, memory, and other components.

The ZT Systems transaction should be read through this supply-chain and systems lens. AMD acquired ZT Systems in March 2025, retained certain design IP and employees, and sold the manufacturing business to Sanmina in October 2025. Sanmina became a preferred manufacturing partner for complex AI solutions. This structure suggests AMD wanted design and deployment expertise for hyperscale AI systems without turning itself into a contract manufacturer. If successful, that can shorten the path from GPU silicon to rack-scale customer deployment. If unsuccessful, AMD may still be judged by customers as a component supplier competing against Nvidia's more integrated systems ecosystem.

7. Financial Profile And Balance Sheet

FY2025 was a step-change year. Revenue rose 34% to $34.639 billion. GAAP gross profit rose to $17.152 billion, with gross margin at 50%. GAAP operating income nearly doubled to $3.694 billion from $1.900 billion in FY2024. GAAP net income was $4.335 billion, and diluted EPS was $2.65. The FMP/official TTM build in the quant foundation reports sum-of-quarter diluted EPS of $2.66, with the official full-year diluted EPS at $2.65 due to rounding and share-count treatment.

Cash generation was strong. Continuing operating cash flow was $6.493 billion and continuing free cash flow was $5.519 billion. Total net cash provided by operating activities including discontinued operations was higher because the ZT Manufacturing Business contributed cash flows classified as discontinued operations. AMD ended FY2025 with $10.552 billion of cash, cash equivalents, and short-term investments, against $3.300 billion of aggregate principal debt. FMP total debt including additional balance-sheet categories was $4.472 billion, but the economic balance sheet is still net cash / low leverage. That gives AMD flexibility to fund R&D, opportunistic acquisitions, supply commitments, and customer ramps.

The income statement has several quality adjustments an analyst should watch. First, stock-based compensation is material: FY2025 All Other included $1.638 billion of SBC. Second, acquisition-related intangibles amortization was $2.254 billion inside All Other, reflecting the Xilinx and other acquisition base. Third, FY2025 taxes were unusually favorable, with a GAAP income tax benefit of $103 million and an effective tax rate of negative 2.5%, helped by an IRS-related release of uncertain tax positions. A normalized future tax rate is likely higher than FY2025 GAAP. Fourth, export-control charges and reversals around MI308 affected gross margin and comparability.

For ER2 purposes, AMD is financially strong enough to pass the balance-sheet and cash-conversion filters. The harder question is not solvency. It is whether revenue growth and margins can compound enough to justify the market price at a P/E of 15 or less on FY+5 earnings. At the live FMP quote captured in the quant foundation, AMD traded at $323.21, with an indicated market capitalization of roughly $527.0 billion. That valuation demands large FY+5 EPS, not merely a good company story.

8. Management Strategy And 2025 Actions

Management's FY2025 strategy was explicitly AI-led. AMD framed AI as becoming pervasive across cloud, edge, and endpoint devices, and positioned the company as an end-to-end AI compute provider spanning CPUs, GPUs, adaptive SoCs, networking, software, and systems integration. This is a broader ambition than simply selling server CPUs. It requires AMD to improve ROCm, expand AI accelerator products, partner with hyperscalers and system providers, and deliver rack-scale platforms.

The most important 2025 actions were the ZT Systems acquisition and divestiture of the manufacturing arm, the OpenAI product purchase agreement, continued EPYC and Instinct product ramps, and expansion of AI PC and embedded AI offerings. The Q4/FY2025 earnings release highlighted Q1 FY2026 revenue guidance of approximately $9.8 billion plus or minus $300 million, with non-GAAP gross margin expected around 55%. That guidance implies continued year-over-year growth as of the latest reported outlook, although Q1 FY2026 has not been reported yet.

Management also increased investment intensity. FY2025 R&D was up 25% year over year, and MG&A was up 52%, with management attributing the latter partly to go-to-market activity supporting revenue growth. This is consistent with a company trying to win large AI infrastructure accounts and expand sales capacity. The risk is that AMD spends ahead of revenue and margins if the AI accelerator opportunity does not scale as expected. The opportunity is that those investments help AMD cross from component supplier to credible platform supplier.

9. Key Investor Questions Before Reading The Segment Appendices

First, is Data Center growth high-margin and durable, or is AMD buying AI accelerator relevance through high spending, lower margin, and customer-specific commitments? The segment nearly reached half of revenue in FY2025, but its operating margin was only 21.7%. A high-quality thesis needs evidence that Instinct/EPYC growth can scale with improving mix and software credibility.

Second, how cyclical is the FY2025 Client and Gaming rebound? The segment grew 51%, with strong client units, ASP, semi-custom, and Radeon demand. That is impressive, but PCs and consoles are historically cyclical. The analyst should separate sustainable share gains from channel restocking and platform-cycle effects.

Third, does Embedded recover while maintaining high margins? Embedded was only 10% of revenue but 16% of positive segment operating income before All Other, and it carried a 36.0% margin. If embedded stabilizes and grows, it improves AMD's mix. If it remains pressured, AMD becomes more dependent on AI and PC cyclicality.

Fourth, how should All Other be modeled? Segment operating income totaled $7.701 billion before All Other, but consolidated operating income was $3.694 billion after a $4.007 billion All Other loss. Ignoring stock compensation and acquisition amortization would overstate economic EPS. Treating all amortization as a permanent cash-like expense may understate underlying operating leverage. The synthesis appendix must handle this explicitly.

Fifth, what is the correct China/export-control discount? China including Hong Kong represented $7.751 billion of FY2025 revenue by billing location, or 22.4% of total. Export controls already affected Instinct MI308 economics in FY2025. The risk is not just lost revenue; it is product redesign, inventory charges, licensing friction, and customer hesitation in restricted markets.

10. Primer Bottom Line

AMD is a high-quality semiconductor design company with a much broader footprint than a simple PC CPU vendor. Its current economic engine is a three-part portfolio: Data Center as the largest and most strategically important growth vector, Client and Gaming as a large cyclical cash and share-gain business, and Embedded as a smaller but high-margin adaptive-compute franchise. The company has a strong balance sheet, good cash conversion, and credible product roadmaps across EPYC, Instinct, Ryzen, Radeon, adaptive SoCs, and embedded AI.

The investment challenge is valuation and execution intensity. At the Stage 0 anchor, AMD's reported TTM EPS and operating margin do not remotely explain the market value by themselves. The market is underwriting a large AI and data center earnings ramp, continued client strength, embedded stabilization, and better absorption of All Other costs. The next appendices should therefore focus less on whether AMD is a good company and more on whether each segment can credibly compound enough by 2030-12-27 to produce the required EPS at a disciplined 15x FY+5 P/E. The most likely thesis breaker is not bankruptcy or business irrelevance; it is that AMD remains an excellent competitor but does not capture enough AI system profit pool, fast enough, at high enough margin, to clear the price.

Appendix M - Management Assessment

Advanced Micro Devices, Inc. - Management Assessment

Date: 2026-04-29 Ticker: AMD Stage 0 anchor: latest reported quarter is Q4 FY2025 ended 2025-12-27; TTM is Q1-Q4 FY2025 ending 2025-12-27; forward horizon is 2030-12-27.

Source Base and Method

This appendix uses AMD's 2026 proxy statement filed March 27, 2026, AMD's FY2025 Form 10-K and Form 10-K/A, AMD's 2025 and 2026 8-K filings for the OpenAI and Meta AI infrastructure partnerships, live FMP earnings-call transcript pulls for Q1 FY2024 through Q4 FY2025, and the DR-cleared ER2 data package sections 7 and 9. The live transcript pulls were saved into `C:\Users\Admin\OneDrive\Cowork\Investing\data warehouse\transcripts\AMD\` as required by the local warehouse rule; they were not used as a stale cache. The management question is narrow: would a long-duration owner trust this team to allocate capital, retain technical advantage, communicate truthfully, and avoid governance drift through the 2030-12-27 forward horizon?

1. Who Is Running the Company

AMD is still Lisa Su's company in the economically relevant sense. Dr. Lisa T. Su has been President and CEO and a director since October 2014 and Chair since February 2022. The 2026 proxy frames her role as the leader of AMD's transformation into high-performance and adaptive computing and as the architect of AMD's current position in the AI ecosystem. That is not just public-relations language: the current EPYC server share gains, chiplet architecture, Xilinx adaptive-computing integration, Instinct data-center GPU ramp, and rack-scale AI strategy all sit within her tenure. At 2026-04-29, she has had more than eleven years as CEO and more than four years as Chair, which is enough tenure to own both the successes and the errors.

The CFO is Jean Hu, Executive Vice President, Chief Financial Officer and Treasurer since January 2023. The proxy says she leads global finance, investor relations, global corporate services, and facilities, and notes prior CFO roles at Marvell, QLogic, and Conexant. This is a credible semiconductor-finance resume, and her tenure is long enough to judge the AI-investment and buyback posture from 2023 onward, but not long enough to attribute the whole AMD turnaround to her.

The operating bench is unusually technical. Mark Papermaster, EVP and CTO, has led technology and engineering since January 2019 after joining AMD in 2011. He is central to the Zen CPU family, high-performance GPUs, modular design, and Infinity Architecture. Forrest Norrod, EVP and GM of the Data Center Solutions Business Unit since January 2023 and an AMD executive since November 2014, runs the most important segment for the forward thesis: data-center CPUs, GPUs, networking, and systems. Jack Huynh has led Computing and Graphics since April 2023 and owns the client/gaming side, including Ryzen, Radeon, NPU, and semi-custom history. Darren Grasby was appointed Chief Sales Officer and President EMEA in January 2025, giving AMD an explicit end-to-end sales leader as it moves from chip sales into hyperscale rack-scale deployments. Ava Hahn has been General Counsel and Corporate Secretary since January 2024, overseeing legal, governance, ethics, compliance, corporate responsibility, and government relations.

Board oversight is adequate but worth monitoring because Su is both Chair and CEO. AMD's proxy discloses Nora Denzel as independent Lead Independent Director, and the board consists entirely of independent directors other than Dr. Su. Standing committees are fully independent, and the board met seven times in 2025, with non-employee directors meeting separately nine times. This is enough to avoid a governance fail. The offset is that AMD's strategy is highly CEO-centric, and the special 2026 CEO award explicitly says Su is the "principal architect" of the strategic plan. A combined Chair/CEO can be acceptable when performance is strong and the lead independent director is real; it becomes a risk if strategy needs a hard reset.

2. Tenure and Track Record

Su's track record is the strongest part of the management case. She took over AMD in 2014, before the EPYC recovery, before the Zen architecture had fully reset x86 competitiveness, before Xilinx, before Instinct became a serious data-center GPU platform, and before AMD had a credible rack-scale AI narrative. By FY2025, the data package shows revenue of $34.639 billion, GAAP operating income of $3.694 billion, and free cash flow from continuing operations of $5.519 billion. The proxy highlights that from fiscal 2020 through fiscal 2025, AMD's market capitalization increased 215%, creating about $239.3 billion of stockholder value. That is not a clean ROIC proof by itself, because semiconductor market multiples and AI enthusiasm matter, but it is a meaningful stewardship signal when paired with revenue growth and technical share gains.

The first major decision was to focus AMD's scarce engineering resources on high-performance x86 CPU leadership, embodied in Zen and EPYC. This was not just a product decision; it was a capital-allocation choice to rebuild credibility in server CPUs against Intel while using a fabless model and advanced-node access through TSMC. By FY2025, AMD reported record server CPU sales and record server revenue share, and the data package shows Data Center revenue rising to $18.370 billion in FY2025. This decision produced the cash and credibility that now fund AI acceleration.

The second decision was the Xilinx acquisition, which closed in 2022. It was expensive and stock-funded, but it diversified AMD into adaptive computing and embedded, expanded the software and customer footprint, and gave AMD a more durable portfolio than CPUs and gaming alone. The follow-through has been mixed but not broken. Embedded revenue weakened after the post-pandemic correction, yet the FY2025 call and proxy noted a return to year-over-year growth in Q4 and $17 billion of embedded design wins in 2025. The acquisition also left large goodwill and intangibles on the balance sheet, so management still needs to prove that Xilinx can be a persistent high-return contributor rather than a cyclical asset bought near a peak. Still, there is no evidence of a write-off or strategic retreat.

The third decision is the 2024-2026 AI infrastructure pivot: Instinct annual cadence, ROCm investment, ZT Systems for rack-scale systems capability, and strategic customer partnerships with OpenAI and Meta. This is the biggest live test of management. AMD completed the ZT acquisition in March 2025, retained the design capability, and sold the manufacturing business to Sanmina in October 2025. The FY2025 10-K describes this as retaining ZT design talent and IP while separating manufacturing. Strategically, that is coherent: AMD wants rack-scale design and customer integration without becoming a low-margin contract manufacturer. The risk is execution complexity. A chip company now has to deliver systems, software, supply planning, and multi-party deployment schedules at gigawatt scale.

The most visible miss is the MI308 export-control episode. In Q2 FY2025, AMD took an $800 million inventory and related charge on MI308 data-center GPU products tied to U.S. export restrictions. Licenses later allowed some shipments, and the company reversed about $360 million of charges in Q4 FY2025. This was not a fatal management failure because the regulatory environment is hard to forecast, but it is a reminder that AMD's AI ramp runs through geopolitical choke points. Management's handling was reasonably direct: it isolated the charge, explained the margin effects, and then distinguished recurring demand from the temporary China-related benefit when analysts asked.

3. Capital Allocation Behavior

The data package's five-year capital allocation table shows a company that allocates primarily through R&D, not capex or dividends. Capex was low relative to revenue: $0.301 billion in FY2021, $0.450 billion in FY2022, $0.546 billion in FY2023, $0.636 billion in FY2024, and $0.974 billion in FY2025, equal to only 1.8% to 2.8% of revenue. This is consistent with a fabless model, not underinvestment. The real investment line is R&D: $2.845 billion in FY2021, $5.005 billion in FY2022, $5.872 billion in FY2023, $6.456 billion in FY2024, and $8.091 billion in FY2025. R&D as a percent of revenue ran from 17.3% in FY2021 to 25.9% in FY2023 and 23.4% in FY2025. That pattern is exactly what management says it is doing: leaning into data-center AI, CPU, software, and systems roadmaps.

AMD does not pay a dividend, which is appropriate. A dividend would be a negative signal while the company is still trying to win AI share and server CPU share. Buybacks have been present but secondary: $1.999 billion in FY2021, $4.108 billion in FY2022, $1.412 billion in FY2023, $1.590 billion in FY2024, and $1.316 billion in FY2025. The FY2025 proxy and data package note that AMD repurchased 12.4 million shares and returned $1.3 billion to shareholders in 2025. The board approved a new $6 billion repurchase program on 2025-05-13, lifting total authorization to $14 billion, and AMD ended FY2025 with $9.4 billion of remaining authorization. This looks more like dilution management and opportunistic support than a core capital return identity. That is the right posture while the stock is volatile and strategic opportunities are large.

M&A is selective but increasingly active. The data package shows M&A cash flow of $0.822 billion in FY2022, $0.131 billion in FY2023, $0.565 billion in FY2024, and $1.760 billion in FY2025, while FY2025 10-K disclosures put total ZT purchase consideration at $4.409 billion before sale mechanics and related settlement items. The judgment on ZT is not yet in the numbers. The strategic rationale is strong because AMD needs rack-scale design credibility to compete against Nvidia's full-stack systems. The capital-allocation risk is that AI customer wins are being supported by warrants and system-level commitments whose economic quality will be hard to judge until revenue, gross margin, and share dilution are visible.

The OpenAI and Meta agreements deserve special scrutiny. In October 2025, AMD entered a product purchase agreement with OpenAI to deploy 6 gigawatts of AMD GPUs, with the first gigawatt powered by MI450-series products, and issued OpenAI a warrant for up to 160 million shares at a $0.01 exercise price. Vesting depends on purchase milestones, stock price targets, technical conditions, and commercial conditions. In February 2026, AMD and Meta announced another up-to-6-gigawatt Instinct deployment partnership, also with a performance-based warrant for up to 160 million AMD shares.

These are not normal sales agreements. They can be brilliant capital allocation if they accelerate ecosystem credibility and create high-margin, multi-generation demand. They can also dilute owners and make reported revenue quality harder to parse if economics depend heavily on equity-linked customer incentives. Management gets credit for structuring milestones, but the correct ER2 stance is to monitor warrant vesting, contra-revenue treatment, and gross-margin progression closely.

4. Compensation Alignment

AMD's compensation design is above average for alignment, with one caveat. For 2025, the proxy says about 96% of the CEO's target total direct compensation and 92% for other named executives was variable or at risk. Long-term equity was the largest component: 89% of CEO target total direct compensation and 83% for other named executives. Annual cash incentives were tied 70% to financial goals and 30% to strategic milestones. Long-term PRSUs were based on absolute and relative TSR over a three-year period plus non-GAAP EPS growth. Options and RSUs add retention and stock-price exposure. The 2025 CEO annual long-term award was 75% PRSUs and 25% stock options; the proxy reports Su's 2025 summary compensation at $55.2 million, including stock awards and option awards.

Ownership is real. As of March 19, 2026, Su beneficially owned 4,305,973 AMD shares, including 413,529 exercisable options. Jean Hu owned 148,364 shares, Forrest Norrod 424,965, Mark Papermaster 1,834,807, and current directors and executive officers as a group owned 8,107,566 shares. The CEO ownership requirement is six times salary, other named executive officers are three times salary, and the proxy states each NEO satisfied or had time remaining to satisfy the requirement. AMD prohibits hedging by employees, named executives, and directors, and pledging requires board-committee preapproval that has not been granted to date. Clawback rights cover incentive compensation tied to restated results and, more broadly, certain misconduct.

The caveat is quantum. On February 17, 2026, AMD filed an 8-K disclosing a one-time $75 million CEO Value Creation Equity Award. It is 100% performance-based PRSUs, runs through March 15, 2031, vests only if stock-price hurdles are met, and tops out at 200% of target if AMD's 90-trading-day average stock price reaches the highest hurdle, including a $600 tranche. The design is shareholder-aligned and retention-oriented. But it also confirms key-person dependence and adds another layer of equity dilution if AMD performs. This does not weaken the management gate to YELLOW because the award is performance-based and Su's track record is strong, but it increases the need for board discipline on future special awards. The 8-K states it is intended to be one-time; that promise matters.

Stock-based compensation remains a cost to monitor. The FY2025 10-K reports stock-based compensation expense of $1.638 billion in FY2025, up from $1.407 billion in FY2024 and $1.384 billion in FY2023. Against FY2025 revenue of $34.639 billion, that is about 4.7% of revenue. This is not out of control for a high-growth semiconductor company competing for AI talent, but it is meaningful. Buybacks partly offset dilution, yet the OpenAI and Meta warrants could dwarf routine employee-equity dilution if fully earned. The management assessment should not treat all equity issuance as bad, but it should force a high evidence bar that the revenue and margin AMD gets in return are worth the ownership transferred.

5. Disclosure Quality and Candor

The last eight quarters of transcript review show a management team that is promotional, but generally responsive. Analysts repeatedly pushed on the same hard topics: Instinct revenue visibility, MI300/MI350/MI450 ramps, gross-margin dilution from GPUs, China export restrictions, OpEx growth, client pull-forward risk, semi-custom cyclicality, and whether OpenAI could realistically consume multiple vendors' AI roadmaps. Su and Hu usually answered with directional numbers, timing, or explicit refusal boundaries rather than pure narrative.

Examples matter. In Q1 FY2025, analysts pressed on why data-center GPU growth appeared below some expectations and whether inventory growth was ZT-related or demand-related. Management separated the China/MI308 issue from broader Instinct demand and discussed inventory build for data-center demand. In Q2 FY2025, analysts asked about China timing and gross margins for MI products. Management did not pretend China restrictions were immaterial; it excluded MI308 from guidance and explained the charge. In Q3 FY2025, analysts questioned OpenAI's ability to engage multiple merchant GPU suppliers and ASIC suppliers. Su addressed customer engagement directly while avoiding false precision. In Q4 FY2025, analysts pushed on 2027 AI revenue expectations, supply capacity, OpEx, MI455 revenue recognition, rack-scale ramp risk, China MI308 sales, and memory procurement. Management gave specific answers on recognition when AMD ships to the rack builder, Q1 MI308 expectations, lack of expected AI supply limitation, and the need for OpEx to grow slower than revenue in 2026.

The weak point in disclosure is voluntary segment granularity for AI GPU revenue. AMD gives enough to follow broad trajectory, but it does not give a clean Instinct revenue series every quarter. Su occasionally offers modeling clues instead of a

direct number, such as telling analysts to back out China revenue to infer growth. That is defensible competitively, but less transparent than an owner would prefer given that the investment case is increasingly dominated by data-center AI. The OpenAI and Meta warrant structures also require better future disclosure: investors will need to see milestone status, accounting treatment, contra-revenue impact if any, gross margins, and whether EPS accretion claims are being met.

Overall, management's candor is good enough for PASS. The team does not hide major negatives: it disclosed the $800 million MI308 charge, the later reversal, the ZT manufacturing divestiture, and the unusual warrant terms in filings. It also names competitive threats directly in the 10-K, identifying Intel, Nvidia, Altera, Arm-based CPU suppliers, smaller accelerator companies, and customers' internal silicon efforts. The language is optimistic, sometimes highly optimistic, but not evasive.

6. Strategic Coherence

AMD's stated strategy is coherent. The FY2025 10-K says AI is shaping the next era of computing and that AMD's portfolio across cloud, edge, embedded, and end devices gives it an opportunity to become an end-to-end AI leader. It also says AI systems require full-stack innovation across compute, networking, systems architecture, and software. The investment pattern matches that claim: R&D rose to $8.091 billion in FY2025; AMD bought and then reshaped ZT Systems to retain rack-scale design capability; it continues to invest in ROCm and AI software; it has added networking through Pensando products; and it is aligning roadmaps with OpenAI and Meta.

The five-to-ten-year destination is clear: be the credible alternative full-stack AI compute supplier across CPUs, GPUs, networking, rack-scale systems, and software while retaining strong client and embedded positions. This is not a rolling sequence of unrelated pivots. It is a logical extension of AMD's historical strength in high-performance CPUs and heterogeneous compute. The risk is not incoherence; the risk is execution stretch. Nvidia has a stronger installed ecosystem, software stack, and systems motion. AMD is trying to catch up while maintaining EPYC, Ryzen, Radeon, embedded, and semi-custom cycles. Management deserves credit for identifying the right battlefield early enough to matter, but the forward assessment must stay conditional on MI450/Helios execution and software adoption.

7. Governance Flags

Governance is clean enough for a management PASS. The 2026 proxy states that fiscal 2025 had no related-person transactions considered related-party transactions under SEC rules. AMD has one share class; there is no dual-class control issue. The FY2025 10-K says management concluded internal control over financial reporting was effective as of December 27, 2025, and there was no change in the most recent quarter that materially affected controls. The proxy reports one delayed Form 4 for Philip Guido due to an administrative error on a one-time share purchase; this is not material.

Insider selling should be monitored but is not a current red flag. The FY2025 10-K disclosed that Mark Papermaster adopted a 10b5-1 trading plan on November 14, 2025 for up to 124,936 shares through November 16, 2026. That is material enough to note because he is central to the technology roadmap, but it is not a clustered senior-team exit signal by itself. The larger governance issue is structural: the company is CEO-centric, the board combined Chair and CEO, and the February 2026 special award reinforces the idea that Su's continued leadership is critical. The lead independent director structure offsets but does not eliminate succession risk.

8. Peer Comparison

Against Jensen Huang at Nvidia, Su's stewardship is excellent but second place on ecosystem execution. Nvidia has compounded from a stronger AI-software and systems position and turned CUDA, networking, and systems integration into an industry standard. AMD under Su has been a superior turnaround and challenger story, but Nvidia has allocated earlier and more completely into the AI platform layer. On transparency, Nvidia also provides enough segment and product color for investors to understand demand drivers, though both companies manage competitive disclosure carefully. For five-year stewardship, Huang ranks ahead because the capital allocation outcome has been extraordinary and the moat broader.

Against Pat Gelsinger's and Intel's recent stewardship, Su ranks far ahead. AMD stayed fabless, focused engineering resources, took server CPU share, protected the balance sheet, and avoided the capital intensity and execution drag that have hurt Intel. Intel's foundry pivot may prove strategically necessary, but the capital burden and delays make AMD's capital-light model look disciplined by comparison.

Against Hock Tan at Broadcom, the comparison is closer and different. Tan's acquisition and integration record is exceptional, with ruthless focus on cash returns and margin discipline. Su is more product-roadmap and engineering-led, with heavier R&D intensity and less immediate cash return. For a 20-year compounder, Broadcom has the cleaner capital-allocation machinery; AMD has the more asymmetric technology share-gain opportunity. Su's management ranks highly among semiconductor

CEOs because she has proven both technical judgment and strategic patience, but the AI systems pivot is still less proven than Broadcom's acquisition playbook or Nvidia's platform dominance.

9. Management Gate Verdict

Management gate: PASS.

AMD passes the management gate because Lisa Su has an eleven-plus-year record of strategic transformation, the technical bench is deep, the balance sheet is strong, compensation is mostly equity- and performance-linked, governance has no disqualifying related-party or control red flags, and capital allocation is coherent with the stated AI infrastructure strategy. The team has earned the right to be followed into the 2030-12-27 horizon. The pass is not unconditional: the OpenAI and Meta warrants, ZT integration, MI450/Helios ramp, and rising stock-based compensation create real owner-dilution and execution risks. If future disclosures show warrant-driven revenue without durable gross margin, repeated special equity awards, or slippage in rack-scale delivery masked by narrative, this gate should be cut to YELLOW. Today, the evidence supports PASS.

Primary sources used: AMD 2026 Proxy Statement filed 2026-03-27 (https://ir.amd.com/financial-information/sec-filings/content/0001193125-26-129057/d85856ddef14a.htm); AMD FY2025 Form 10-K filed 2026-02-04 (https://ir.amd.com/financial-information/sec-filings/content/0000002488-26-000018/amd-20251227.htm); AMD FY2025 Form 10-K/A (https://ir.amd.com/financial-information/sec-filings/content/0000002488-26-000021/amd-20251227.htm); AMD 8-K for CEO Value Creation Equity Award filed 2026-02-17 (https://ir.amd.com/financial-information/sec-filings/content/0000002488-26-000029/amd-20260210.htm); AMD/OpenAI 8-K and press release filed 2025-10-06 (https://ir.amd.com/financial-information/sec-filings/content/0001193125-25-230895/d28189d8k.htm); AMD/Meta 8-K press release filed 2026-02-24 (https://ir.amd.com/financial-information/sec-filings/content/0000002488-26-000045/pressreleasedatedfebruary2.htm); live FMP earning-call transcript endpoint pulls for AMD FY2024 Q1 through FY2025 Q4 saved locally on 2026-04-29.

Appendix W - Moat and 2046 Durability

Advanced Micro Devices, Inc. - Moat and 2046 Durability

Date: 2026-04-29

Stage 0 anchor: AMD's latest reported quarter is Q4 FY2025, fiscal year ended 2025-12-27. TTM equals FY2025. Segment composition and historical margins are taken from the DR-cleared `data_package.md`, which ties FMP live pulls to SEC filings and AMD primary releases. Segment deep-dive appendices were still placeholders when this appendix was written, so segment-level moat work below integrates the verified data package, the primer, AMD's FY2025 10-K, AMD investor releases, and live peer/industry context.

1. Moat taxonomy applied

AMD does not have a single software-style network effect. It sells compute platforms into ecosystems where developer tooling, validated workloads, OEM designs, hyperscaler qualifications, console platforms, and embedded design environments reinforce adoption, but those are indirect effects. The closest network-like asset is in Data Center AI: successful Instinct deployment at OpenAI, Meta, Oracle, or another hyperscaler improves ROCm validation, library maturity, model support, and procurement confidence for the next customer. That flywheel is real but young, and it remains weaker than Nvidia's CUDA/system ecosystem.

Switching costs are AMD's strongest classic moat, but they vary sharply by segment. Embedded has the clearest switching costs because FPGAs, adaptive SoCs, embedded CPUs, modules, tools, reference designs, and application-engineer support are designed into long-lived industrial, aerospace, defense, communications, automotive, healthcare, and test systems. Once qualified, the cost of redesign, recertification, supply-chain requalification, and customer engineering time is high. Data Center CPUs also have switching costs at the platform level: hyperscalers and enterprise OEMs qualify EPYC server platforms, tune workloads, validate firmware, and build procurement processes. AI accelerator switching costs arise only after the software and system stack is adopted; before that point, Nvidia is the higher switching-cost incumbent.

Scale economies are meaningful but not absolute. AMD is fabless, so it does not own wafer-fab scale, but it does have R&D scale, roadmap scale, packaging/supply-chain procurement leverage, and customer support leverage across CPUs, GPUs, adaptive compute, and semi-custom SoCs. FY2025 R&D was $8.091B, or 23.4% of revenue, and that lets AMD amortize architecture, IP, software, and validation investments across multiple markets. The limitation is that Nvidia, Intel, Arm ecosystem participants, hyperscaler custom-silicon teams, and FPGA/ASIC competitors also have scale.

Brand is useful but not independently decisive. Ryzen, EPYC, Radeon, Instinct, and Xilinx-derived adaptive compute brands help AMD get considered by OEMs, hyperscalers, developers, channel partners, and enterprise buyers. Brand mattered in AMD's 2025 client rebound, where the data package records record client CPU ASP commentary and FY2025 Client revenue of $10.640B. But semiconductor brand equity is quickly repriced by benchmarks, power efficiency, software quality, availability, and roadmap timing. AMD's brand supports trust; it does not let the company raise price without product evidence.

Regulatory barriers are not a moat for AMD. Export controls, supply-chain rules, and geopolitical licensing requirements are more often constraints than protection. FY2025 included an $800M inventory and related charge tied to U.S. export restrictions on AMD Instinct MI308 products, partly reversed by Q4 release mechanics. China including Hong Kong was $7.751B of FY2025 billing-location revenue, or 22.4% of total revenue. Regulation can reshape demand and limit competitors in specific geographies, but it is not a dependable AMD-owned advantage.

IP and patents matter, especially in x86 CPUs, GPUs, interconnect, chiplets, adaptive compute, and semi-custom SoC design, but IP alone is not sufficient. AMD's rights to ship x86 processors, its architecture know-how, and the Xilinx adaptive-compute portfolio are genuine barriers to a new entrant. The FY2025 10-K exhibit list also shows continuing patent cross-license architecture around the industry. The practical moat is not a patent wall; it is the combination of legally usable instruction-set positions, design teams, verification capability, physical implementation, software, and customer trust.

High-trust procurement is the moat most visible in the numbers and customer actions. AMD sells mission-critical components into cloud regions, AI clusters, enterprise servers, consoles, aircraft, factories, communications networks, and medical/industrial systems. Customers will not redesign around an unproven vendor just for a small price discount. This procurement trust shows up in AMD's Data Center growth to $16.635B in FY2025, OpenAI/Meta/Oracle-style multi-generation commitments, and Embedded design-win commentary of $17B in 2025 and more than $50B since Xilinx acquisition. It is strongest where qualification and roadmap credibility matter most: Embedded and Data Center.

2. Per-segment moat

Data Center is a real but contested moat. In server CPUs, AMD's moat is the EPYC performance-per-watt, core-count, x86 compatibility, TCO, and platform-roadmap position embedded in hyperscaler and enterprise server procurement. In AI accelerators, the moat is earlier-stage: AMD is trying to convert Instinct silicon, EPYC CPUs, Pensando networking, ROCm software, ZT-retained systems design, and Helios rack-scale architecture into a procurement alternative to Nvidia. The moat is durable over five years if MI350/MI450 deployments scale with improving software and gross margin; durable over ten years only if ROCm and customer co-design become normal production infrastructure; not yet proven over twenty years. It breaks if major customers deploy AMD GPUs only as a price lever against Nvidia and do not standardize material production workloads on AMD systems. The earliest signal would be large announced gigawatt deals failing to convert into recognized revenue, follow-on generations, and higher segment margins.

Client and Gaming has a thinner moat. Client CPUs/APUs benefit from x86 compatibility, Ryzen brand, OEM relationships, chiplet design, integrated graphics, and price-performance credibility. Gaming semi-custom has platform design-in stickiness once AMD wins a console or custom SoC design, but the revenue is tied to customer platform cycles. Discrete Radeon GPUs face Nvidia's stronger brand and software ecosystem. Five-year durability is acceptable because AMD's CPU share and ASP gains are visible, and FY2025 Client and Gaming revenue rose 50.8% to $14.550B. Ten-year durability is less certain because Arm-based PCs, Intel recovery, Nvidia/Intel co-developed products, and console cycles can all compress returns. The earliest break signal is not a single bad quarter; it is sustained loss of premium notebook/desktop design wins, falling ASPs despite unit growth, or next-generation console awards moving away from AMD.

Embedded is AMD's cleanest segment-level moat. It has design-in switching costs, specialized tools, long qualification cycles, application-engineer relationships, and adaptive-compute IP that is difficult to rip out after a customer ships a long-lived product. FY2025 revenue was only $3.454B, but segment operating margin was 36.0%, the highest of AMD's reportable segments. The five-year moat is strong if end-market inventory recovery continues. The ten-year moat is credible because embedded product lives and qualification cycles are long. The twenty-year moat is plausible, not guaranteed, because FPGAs and adaptive SoCs could lose share to ASICs, ASSPs, and domain-specific accelerators if customers decide reconfigurability is no longer worth the cost and power tradeoff. The earliest warning sign is design-win growth failing to translate into revenue after the normal lag, while margin falls below the mid-30s despite recovery.

The blended company moat is weaker than Embedded and weaker than the best version of Data Center. FY2025 revenue mix was 48.0% Data Center, 42.0% Client and Gaming, and 10.0% Embedded. That means almost half the company is in a high-growth but brutally contested AI/server market, more than two-fifths is in cyclical PC/gaming markets, and only one-tenth is in the clearest switching-cost franchise. AMD is a high-quality competitor, but the equity cannot be valued as if all revenue had Embedded-like durability or Nvidia-like AI ecosystem lock-in.

3. Margin persistence evidence

Consolidated GAAP operating margin does not yet prove a wide moat. The data package shows FY2021-FY2025 operating margins of 22.2%, 5.4%, 1.8%, 7.4%, and 10.7%. Revenue grew from $16.43B in FY2021 to $34.64B in FY2025, but GAAP operating income rose only from $3.65B to $3.69B because of Xilinx-related amortization, stock-based compensation, AI investment, export-control charges, and cyclical compression. That is not a failure of the business, but it is a warning against overstating current consolidated pricing power. A company with a fully mature moat would normally show more persistent operating margin through a doubling of revenue.

Segment margins are more constructive. Data Center margins were 26.8% in FY2021, 30.6% in FY2022, 19.5% in FY2023, 27.7% in FY2024, and 21.7% in FY2025. The segment absorbed AI ramp spending and MI308 export-control disruption, yet remained profitable and nearly doubled revenue from FY2023 to FY2025. Client and Gaming margins were 24.2%, 16.5%, 8.5%, 12.3%, and 19.6% across FY2021-FY2025, demonstrating cyclicality but also recovery when mix, ASP, and units improve. Embedded margins were 17.9% in pre-Xilinx FY2021, then 49.5%, 49.4%, 39.9%, and 36.0% from FY2022-FY2025. Embedded has the clearest margin persistence after the Xilinx acquisition, even through revenue digestion.

The evidence supports a narrow-to-moderate moat at the corporate level, not a fortress. AMD has enough differentiation to earn attractive segment margins in competitive markets, but the consolidated GAAP margin record remains volatile. The moat argument improves if FY2026-FY2028 Data Center revenue scales with higher gross margin and lower All Other drag. It weakens if AI revenue grows but Data Center margins stay near low-20s or below.

4. AI disruption stress test

Data Center is not a workflow that a foundation model can replicate. It is chokepoint economics around physical compute, advanced silicon, packaging, memory, networking, rack systems, firmware, compilers, workload validation, and hyperscaler procurement. AI increases demand for the segment, but it also shifts the basis of competition from CPU share gains to full AI infrastructure economics. AMD benefits if customers want an open, lower-cost, multi-vendor alternative to Nvidia and are willing to co-optimize software. AMD is exposed if AI frameworks, libraries, and developer habits remain tightly optimized around Nvidia, or if hyperscalers move incremental workloads to internal ASICs with Arm host CPUs.

Client and Gaming is platform-adjacent and partly exposed. Local AI features can raise the value of NPUs, APUs, and integrated CPU/GPU platforms, which helps AMD if Ryzen AI systems win OEM slots. But many consumer PC workflows are software experiences where the model owner, OS vendor, or cloud platform captures the economics. The hardware moat is therefore in battery life, local performance, compatibility, OEM designs, and cost, not in AI output quality itself. Discrete gaming GPUs are also vulnerable if AI-driven upscaling, integrated graphics, or cloud gaming reduce the need for midrange discrete cards. Semi-custom is protected while AMD remains designed into consoles, but the console form factor is not immune to cloud or handheld platform shifts.

Embedded is the most protected from AI disintermediation. AI at the edge can create new workloads in industrial inspection, robotics, medical imaging, communications, aerospace, and automotive systems, but the value is implemented in qualified hardware, deterministic latency, safety, reliability, and long-life support. A foundation model can help engineers design or tune systems; it cannot replace a certified adaptive SoC in a deployed industrial or aerospace platform. The main AI risk is not workflow substitution. It is that customers adopt cheaper ASICs, ASSPs, or Nvidia/Qualcomm-style edge platforms when reconfigurability is not worth the premium.

The most exposed segment is Client and Gaming, especially discrete gaming GPUs and consumer PC sockets where product cycles, price/performance, OEM allocation, and platform shifts can move quickly. The most protected segment is Embedded. Data Center sits between them: AI is the demand driver, but Nvidia's ecosystem and hyperscaler custom silicon make the moat highly contested.

5. The 2046 test

Today is 2026-04-29. Twenty years forward is 2046-04-29. A plausible 2046 AMD still sells high-performance and adaptive compute platforms, but the mix is different. Data Center remains the largest business, not as a generic GPU vendor but as a merchant and semi-custom AI compute platform supplier selling accelerator complexes, CPUs, networking/offload silicon, memory/system IP, and rack-scale reference designs to hyperscalers, sovereign AI clouds, enterprise AI infrastructure providers, and specialized HPC customers. EPYC may not exist in the same form, and x86 may be less central, but AMD's enduring role is high-end compute architecture and customer-specific platform design.

In that 2046 base case, Embedded has grown in strategic importance even if it is not the largest revenue line. Adaptive compute persists because many real-world systems need reconfigurable acceleration, long support lives, and qualification. Client PCs still exist but are a smaller share of profit, with AI endpoints, workstations, and premium local compute replacing some commodity notebook economics. Gaming semi-custom survives if AMD remains a preferred platform-design partner for consoles, handhelds, cloud-gaming appliances, or immersive devices, but standalone gaming GPU profit is not the core moat. Directionally, AMD's 2046 operating margin should be higher than FY2025 GAAP operating margin if the moat works: the company would have scaled software, systems, and adaptive-compute attach enough to absorb R&D and All Other costs. It does not need Nvidia-like margins to pass, but it should be structurally above the FY2025 10.7% GAAP level.

Why does that company still make money? The surviving moat is high-trust procurement plus R&D scale plus design-in switching costs, not a simple monopoly. AMD in 2046 must be one of a small number of vendors able to deliver merchant or semi-custom high-performance compute at the boundary of silicon, packaging, software, and systems. It makes money because customers do not want to rely on a single dominant supplier, because internal custom silicon does not cover every workload, because reconfigurable/adaptive compute remains valuable at the edge, and because validated compute platforms are hard to replace once deployed at scale.

The 2046 bear case is also specific. AMD could become a structurally lower-return second source if Nvidia keeps the dominant AI software/system profit pool, Intel stabilizes through foundry and Nvidia collaboration, Arm host CPUs and hyperscaler ASICs absorb much of incremental data center compute, and PC/gaming revenue remains cyclical without enough premium share. Embedded could also erode if ASICs and ASSPs capture high-volume edge AI designs while FPGAs remain niche. In that bear case, AMD still exists in 2046, but it is a merchant chip supplier with mid-teens or lower operating margins and weak pricing power. I would assign that bear case a material probability, roughly one-third.

6. Competitive threat inventory

Threat	Mechanism	Segment(s) affected	Earliest observable signal	Plausibility
Nvidia system/software lock-in	CUDA, networking, full racks, developer defaults, and customer productivity keep AMD as a price lever rather than a primary AI platform	Data Center	Instinct gigawatt deals do not repeat by generation; Data Center margin fails to expand despite revenue growth	High
Intel/Nvidia and Intel recovery	Intel regains CPU competitiveness and co-develops x86 CPUs with Nvidia for client/data center platforms	Data Center; Client and Gaming	EPYC share gains stall; Ryzen premium ASPs fall; OEMs shift AI PC designs toward Intel/Nvidia	Medium-high
Hyperscaler custom silicon and Arm	Large customers build internal accelerators and Arm host CPUs, shrinking merchant TAM or pressuring price	Data Center	Cloud capex growth decouples from AMD revenue growth; custom ASIC deployments take priority in AI inference	High
Console and gaming platform shift	Console cycle matures, cloud/handheld formats change, discrete GPUs lose midrange attach	Client and Gaming	Semi-custom revenue declines beyond expected cycle; Radeon share/mix deteriorates	Medium
FPGA/adaptive compute substitution	ASICs, ASSPs, edge GPUs, or domain accelerators displace reconfigurable logic in new designs	Embedded	Design wins remain large but revenue conversion weakens; segment margin falls below low-30s	Medium

7. Pricing power verification

Direct pricing power is observable in pockets, not across the entire company. The strongest evidence is in Client: management commentary in the data package cites record client CPU ASP driven by richer high-end desktop and mobile Ryzen mix, and FY2025 Client revenue grew to $10.640B. That suggests AMD can earn better mix when products are competitive. However, PC pricing power is product-cycle dependent; it can vanish if Intel or Arm alternatives improve.

Data Center pricing power is not yet fully verified. FY2025 Data Center revenue grew 32.2% to $16.635B, and Q4 Data Center operating margin reached 32.6%, but full-year margin was 21.7% after ramp costs and export-control noise. OpenAI, Meta, Oracle, and other large-scale relationships demonstrate procurement power and customer need for AMD supply, but they do not disclose price, gross margin, minimum commitment economics, or churn. Management's target to scale AI to tens of billions in annual revenue in 2027 is important, but pricing power will be verified only if segment margin expands as Instinct shipments scale.

Embedded pricing power is best inferred from margin persistence and design-in economics. FY2022-FY2025 segment margins of 49.5%, 49.4%, 39.9%, and 36.0% through a revenue downturn indicate differentiated economics. That is stronger evidence than press-release language. Still, the design-win numbers are not backlog, and the company does not disclose contract escalators or minimum purchase commitments. Embedded pricing power is good, but not immune to cyclical inventory digestion or substitution.

Contract structure is a limitation across AMD. The primer notes that AMD generally operates with forecasts, purchase orders, and contractual agreements rather than blanket minimum-purchase commitments. The OpenAI and Meta warrants align incentives, but public filings do not prove broad multi-year price escalators or take-or-pay protection. Pricing power is therefore moderate: strong where performance and design-in matter, weak where customers have dual-source leverage.

8. Moat verdict

Moat gate: YELLOW.

2046 test: YELLOW.

AMD has real moats, but they are not evenly distributed. Embedded is the clearest durable moat because switching costs and qualification support a high-margin adaptive-compute franchise. Data Center has a credible high-trust procurement and platform moat in CPUs and a developing AI systems moat, strengthened by OpenAI, Meta, Oracle, and ROCm/Helios progress, but it still sits against Nvidia's much stronger AI ecosystem and hyperscaler custom silicon. Client and Gaming is valuable but cyclical and less protected. The 2046 picture is specific and plausible: AMD can still make money as a high-performance/adaptive compute platform supplier with merchant and semi-custom roles. The gate remains Yellow rather

than Pass because consolidated GAAP margin persistence is not yet strong, the largest profit-pool opportunity is still contested, and the 2046 bear case is material rather than remote.

Source notes

- Local verified control package: `C:\Users\Admin\OneDrive\GPT\ER 2.0\Advanced Micro Devices, Inc\data_package.md`.
- Local primer orientation: `C:\Users\Admin\OneDrive\GPT\ER 2.0\Advanced Micro Devices, Inc\appendix_A_primer.md`.
- AMD FY2025 Form 10-K: https://ir.amd.com/financial-information/sec-filings/content/0000002488-26-000018/amd-20251227.htm.
- AMD/OpenAI 6GW partnership release: https://www.amd.com/en/newsroom/press-releases/2025-10-6-amd-and-openai-announce-strategic-partnership-to-d.html.
- AMD/Meta 6GW partnership release: https://ir.amd.com/news-events/press-releases/detail/1279/amd-and-meta-announce-expanded-strategic-partnership-to-deploy-6-gigawatts-of-amd-gpus.
- AMD Advancing AI 2025 release: https://www.amd.com/en/newsroom/press-releases/2025-6-12-amd-unveils-vision-for-an-open-ai-ecosystem-detai.html.
- AMD/Oracle MI450 supercluster release: https://www.amd.com/en/newsroom/press-releases/oracle-and-amd-expand-partnership-to-help-customers-ach.html.
- Intel FY2025 Form 10-K: https://www.sec.gov/Archives/edgar/data/50863/000005086326000011/intc-20251227.htm.
- Arm FY2025 Form 20-F: https://www.sec.gov/Archives/edgar/data/1973239/000197323925000016/arm-20250331.htm.
- Mercury Research summary via Tom's Hardware for CPU share context: https://www.tomshardware.com/pc-components/cpus/30-percent-of-x86-cpus-sold-are-now-made-by-amd-as-companys-market-share-grows-thanks-to-a-flagging-intel-enjoys-growth-across-all-segments-as-competition-intensifies.

Appendix B1 - Data Center

Data Center - Segment Deep Dive

Company: Advanced Micro Devices, Inc. (AMD) Segment: Data Center Stage 0 anchor: TTM is Q1 FY2025 through Q4 FY2025, ending 2025-12-27. The forward horizon is exactly five calendar years later, 2030-12-27. Source hierarchy: all AMD historical dollars, growth rates, segment margins, and management commentary in this appendix come from `data_package.md`. External sources are used only for competitive context.

0. Segment Boundary And Taxonomy Note

AMD's current Data Center segment includes AI accelerators, server CPUs, GPUs/APUs, DPUs, AI NICs, FPGAs, and adaptive SoC products for data centers. The data package states that the Q1 FY2025 reporting change combined Client and Gaming into one reportable segment; it did not restate Data Center into a new boundary. Therefore Data Center is the cleanest of AMD's reportable segments for FY2023-FY2025 trend work. FY2021-FY2022 are still useful directional history from earlier segment tables, but the model anchors on the current-taxonomy FY2025/TTM control total. Data Center TTM revenue is $16.635B, TTM operating income is $3.603B, and TTM operating margin is 21.7% [data_package Section 3, Data Center table; Section 5, Data Center quarterly trend].

The major constraint is product disclosure. AMD does not disclose annual EPYC server CPU revenue, annual Instinct GPU/accelerator revenue, or annual data-center FPGA/DPU/NIC revenue separately. The data package explicitly says Data Center AI, EPYC, and Instinct remain open product-dollar gaps and instructs downstream agents not to fabricate those figures [data_package Section 4, Product and Sub-Segment Historicals; Section 10, Open Data Gaps]. This appendix therefore uses the disclosed Data Center segment as the modeled dollar line and treats EPYC, Instinct, rack-scale systems, and adaptive data-center products as qualitative drivers of that line. That is less granular than an ideal bottom-up model, but it is the only controllable approach that respects the verified data package.

1. What This Segment Actually Is

Data Center is AMD's enterprise and hyperscale compute segment. It sells processors and accelerators that sit inside cloud data centers, enterprise server rooms, AI training clusters, AI inference fleets, and data-center networking or adaptive-compute systems. The customer base includes hyperscale cloud providers, frontier AI labs, enterprises, OEM server vendors, and increasingly large infrastructure partners that build rack-scale AI systems. The segment's economic problem is simple to describe and hard to execute: customers need more compute per watt, more memory bandwidth, more networking bandwidth, and more system-level reliability every generation, while also reducing total cost of ownership versus incumbent alternatives.

The product center of gravity has two poles. The first is EPYC server CPUs. EPYC competes primarily with Intel Xeon and cloud-provider custom Arm CPUs such as AWS Graviton. Server CPUs are sold through OEMs, hyperscale direct relationships, and cloud instance rollouts. The value proposition is performance per socket, performance per watt, memory bandwidth, core count, platform longevity, and an x86 software compatibility base. CPUs are a share-gain and platform-attach business: wins compound when cloud vendors standardize instance families, enterprises certify workloads, and OEMs refresh platforms around the processor roadmap.

The second pole is Instinct AI accelerators and the related system stack. Instinct competes mainly with NVIDIA's data-center GPU platform, but the competition is no longer only chip versus chip. AI accelerators sell as systems: accelerator, memory, interconnect, networking, software libraries, compilers, model kernels, reference racks, serviceability, and customer-specific deployment support. The data package describes MI350 sampling in Q1 FY2025, volume production ahead of schedule in Q2 FY2025, and management's move toward rack-scale systems through Helios and MI450 timing [data_package Section 6, Data Center rows]. The ZT Systems acquisition and subsequent sale of the manufacturing business also matter here: AMD retained rack-scale systems design capability while divesting manufacturing operations [data_package Section 9]. That suggests AMD wants more control of system architecture without owning low-margin manufacturing capacity.

The third component is adaptive and networking-adjacent data-center silicon: FPGAs, adaptive SoCs, DPUs, and AI NICs. These products matter because data-center performance bottlenecks increasingly sit outside the accelerator die. In large AI systems, memory, networking, rack-level power, and cluster orchestration can determine realized utilization. AMD's Xilinx-derived technology gives it a credible story in adaptive compute and customer-specific acceleration. However, the data package gives no separate Data Center dollar disclosure for these products, so they should support the qualitative moat

discussion rather than drive a hidden revenue line.

Pricing is a blend of product ASP, platform attach, and system value. EPYC pricing is constrained by Intel, Arm alternatives, customer qualification cycles, and x86 platform economics. Instinct pricing has more upside when capacity is scarce and performance-per-dollar is strong, but it faces the most severe benchmark, software, and ecosystem pressure. AMD's Data Center segment therefore has both cyclical semiconductor features and platform features. It can scale very quickly when products win in cloud and AI infrastructure, but it can also see abrupt margin pressure from export controls, inventory charges, underutilization, or an inferior software stack.

2. Current Economics

TTM equals FY2025 for AMD under the Stage 0 anchor. Data Center generated $16.635B of revenue in FY2025, or 48.0% of AMD's consolidated revenue [data_package Section 1 segment table]. Operating income was $3.603B, giving a 21.7% operating margin [data_package Section 3 Data Center table]. The segment is now AMD's largest reportable segment by revenue and the main swing factor in the company-level growth thesis.

The recent quarterly path shows how volatile the economics can be even when demand is strong. Q1 FY2025 Data Center revenue was $3.674B, up 57% YoY, with 25.4% operating margin. Q2 revenue was $3.240B, up 14% YoY, but operating margin fell to negative 4.8%. Q3 revenue recovered to $4.341B, up 22% YoY, at 24.7% margin. Q4 revenue was $5.380B, up 39% YoY, at 32.6% margin [data_package Section 5, Data Center Quarterly Trend]. The annual margin therefore understates the Q4 run-rate margin but captures a real risk: export restrictions and product-specific inventory events can destroy a quarter's profit even if the five-year demand picture is healthy.

The data package attributes Q4 FY2025 Data Center growth to EPYC processors and Instinct shipments, but it does not disclose the split [data_package Section 4, Data Center rows]. It also discloses selected China MI308 revenue of approximately $390M in Q4 FY2025 and an $800M inventory and related charge in Q2 FY2025 tied to U.S. export restrictions on MI308 products [data_package Section 4; Section 9]. These two facts are important in opposite directions. The Q4 MI308 disclosure shows that export-control-affected products still generated some revenue after reserved inventory was released. The Q2 charge shows that geopolitical constraints can convert otherwise valuable inventory into a margin event.

Capital intensity at the corporate level remains low relative to foundry-integrated semiconductor peers because AMD is fabless. FY2025 capex was $0.974B, or 2.8% of revenue [data_package Section 7]. That figure is consolidated rather than segment-specific, but it frames Data Center economics: AMD does not need to fund leading-edge fabs. It must instead fund R&D, software, customer engineering, packaging relationships, supply commitments, and working capital. FY2025 R&D was $8.091B, or 23.4% of revenue [data_package Section 7]. In Data Center, reinvestment is mostly expensed through R&D and go-to-market support, not reported as segment capex.

The segment's current economics are therefore mixed. Revenue growth is excellent: FY2023 to FY2025 revenue rose from $6.496B to $16.635B, and FY2025 revenue grew 32.2% YoY [data_package Section 3]. But operating margin is not yet a straight-line scale story. FY2024 margin was 27.7%, FY2025 margin was 21.7%, and Q4 FY2025 reached 32.6% after a depressed Q2 [data_package Section 3; Section 5]. The base case should not simply extend Q4 margin to the whole horizon; the Q4 run rate is evidence of upside, not proof that five-year normalized margin is already above 30%.

3. Unit Economics

Data Center unit economics are driven by performance density, platform qualification, software maturity, and supply allocation. A server CPU sale begins with silicon cost, package cost, test cost, and wafer-supply economics. Gross profit then depends on ASP, mix, and the amount of customer-specific support required. EPYC's unit economics improve when AMD sells higher-core-count CPUs, wins premium cloud instances, and attaches platforms across multiple OEM generations. They weaken when Intel discounts Xeon, when cloud customers substitute Arm CPUs for scale-out workloads, or when enterprise qualification costs rise faster than revenue.

AI accelerator unit economics are more extreme. The chip itself uses leading-edge process capacity, expensive HBM, advanced packaging, and high-value substrates. The system around it requires networking, thermals, firmware, drivers, compilers, libraries, model support, and deployment engineering. When supply is scarce and the performance-per-dollar is accepted by customers, accelerator gross profit can be very strong. When the software stack lags, a customer may demand lower pricing to compensate for integration work. If an export-control SKU is restricted after production planning, inventory charges can overwhelm normal gross profit for that product cycle. The Q2 FY2025 negative Data Center operating margin is direct evidence that these tail risks are not theoretical [data_package Section 5; Section 9].

Operating income is then shaped by R&D leverage and systems investments. AMD must fund multiple parallel roadmaps: EPYC CPU generations, Instinct accelerators, rack-scale AI systems, ROCm software, networking/AI NICs, adaptive-compute integrations, and customer engineering. Because R&D is largely fixed before revenue arrives, revenue acceleration can create operating leverage. The CFO commentary in the data package says AMD expects revenue to expand faster than operating expense [data_package Section 6, Capital and margins rows]. That is not Data Center-specific, but Data Center is large enough and fast-growing enough that it is the likely operating-leverage engine.

The margin bridge from FY2025's 21.7% to the base-case 28.5% at the horizon has four mechanisms. First, the Q2 MI308 charge should not repeat at the same magnitude in every year. Second, rack-scale design can increase system relevance and customer lock-in, although it may also carry lower margin if the system scope expands beyond chips. Third, the EPYC installed base should support steadier CPU margins if cloud and enterprise share gains continue. Fourth, AI accelerator scale should absorb software and customer-support spending across a larger revenue base. The offset is price competition against NVIDIA and custom silicon, plus the possibility that customers use AMD primarily as a second source and demand aggressive pricing.

The segment is not structurally guaranteed to earn NVIDIA-like data-center margins. NVIDIA's FY2026 Data Center revenue scale was far larger, with full-year Data Center revenue of $193.7B and Q4 revenue of $62.3B, according to NVIDIA's February 25, 2026 release. That scale funds software, networking, systems, and ecosystem work that AMD is still trying to match. The realistic question for AMD is not whether it can equal NVIDIA's platform economics by 2030. It is whether it can earn high-20s segment operating margins while capturing enough AI accelerator and server CPU demand to become a durable second platform.

4. Competitive Dynamics

The Data Center segment faces three layers of competition: incumbent GPU platform competition, server CPU competition, and hyperscaler internal silicon.

NVIDIA is the main AI accelerator competitor. Its FY2026 Data Center revenue was $193.7B, up 68%, and Q4 FY2026 Data Center revenue was $62.3B, up 75% year over year, according to NVIDIA's investor release. That is not a peer-scale comparison; it is a platform gap. NVIDIA sells GPUs, networking, systems, software, and reference architectures into the largest AI customers. The OECD's 2025 AI infrastructure report describes GPUs as the most used chips for data-center AI tasks and cites estimates that NVIDIA has over 80% market share in AI GPU chips. The same report notes AMD as NVIDIA's most direct GPU-design rival while highlighting NVIDIA's first-mover advantage and CUDA software ecosystem. For AMD, this means the bear case is not only slower chip adoption. It is slower software adoption, lower utilization in customer deployments, and lower pricing power because customers must spend more engineering effort to use the platform.

Intel is the main x86 server CPU competitor. Intel's Q4/FY2025 release reported 2025 Data Center and AI revenue of $16.9B, up 5%, and Q4 2025 DCAI revenue of $4.7B, up 9%. That puts Intel's DCAI revenue in the same general annual scale as AMD's Data Center revenue, although Intel's segment includes a different product mix and intersegment treatment. Intel remains important because server CPU buying cycles are conservative, software compatibility matters, and Intel can defend share with platform roadmaps, price, and manufacturing messaging. AMD's EPYC moat is performance and power efficiency, but it is not a monopoly.

Hyperscaler custom silicon is the strategic substitute. AWS disclosed Trainium3 UltraServers at re:Invent 2025, describing systems with up to 144 Trainium3 chips and materially better compute performance and energy efficiency than Trainium2. Google described Ironwood as its seventh-generation TPU and highlighted inference optimization, pods scaling to 9,216 chips, and 1.77 PB of shared HBM in a superpod. Microsoft introduced Maia 200 in January 2026 as an inference accelerator built on TSMC 3nm with HBM3e and claims of improved performance per dollar. Broadcom's FY2025 results provide the merchant-enablement angle: Broadcom said Q4 AI semiconductor revenue rose 74% YoY and expected Q1 FY2026 AI semiconductor revenue to double YoY to $8.2B, driven by custom AI accelerators and Ethernet AI switches.

These custom silicon efforts do not eliminate AMD's opportunity. They confirm that hyperscalers want alternatives to paying monopoly rents for general-purpose GPU capacity. That can help AMD if customers want a merchant second source with a broad roadmap. But custom silicon also caps the bull case. If Google, Amazon, Microsoft, Meta, OpenAI, and other large customers move more inference and internal workloads to proprietary silicon, AMD's addressable opportunity becomes less like "all AI compute" and more like "external and multi-tenant workloads where a merchant accelerator has enough software support and better economics than the incumbent."

Server CPU competition has a different structure. Cloud providers may prefer custom Arm CPUs for internal scale-out workloads, but x86 remains valuable for enterprise compatibility and a broad software base. AMD can win with EPYC where customers care about socket performance, power, memory bandwidth, and migration simplicity. The server CPU market is less likely to produce a sudden 5x revenue step, but it can provide a durable base, stabilizing margin when Instinct is volatile. The base case assumes CPU share gains continue, but AI accelerators remain the main growth driver.

Pricing power is therefore conditional. AMD has pricing power when it solves a real customer bottleneck, offers supply availability, and reduces total cost of ownership. It lacks pricing power when a customer frames it as a second source, when ROCm or kernel maturity imposes integration costs, or when a hyperscaler can use internal silicon for the same workload. The correct forecast shape is rapid early growth from a small base followed by deceleration as competition, customer concentration, and software maturity become more important than raw demand.

External competitive sources used in this section: NVIDIA FY2026 release (https://investor.nvidia.com/news/press-release-details/2026/NVIDIA-Announces-Financial-Results-for-Fourth-Quarter-and-Fiscal-2026/); Intel FY2025 release (https://www.intc.com/news-events/press-releases/detail/1759/intel-reports-fourth-quarter-and-full-year-2025-financial); AWS re:Invent 2025 AI announcements (https://www.aboutamazon.com/news/aws/aws-re-invent-2025-ai-news-updates); Google Ironwood TPU post (https://blog.google/innovation-and-ai/infrastructure-and-cloud/google-cloud/ironwood-google-tpu-things-to-know/); Microsoft Maia 200 blog (https://blogs.microsoft.com/blog/2026/01/26/maia-200-the-ai-accelerator-built-for-inference/); Broadcom FY2025 release (https://investors.broadcom.com/node/63716/pdf); OECD AI infrastructure report (https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/11/competition-in-artificial-intelligence-infrastructure_69319aee/623d1874-en.pdf).

5. Moat At The Segment Level

Data Center has a real moat, but it is narrower than the market's AI excitement might imply. The moat is strongest in EPYC server CPUs, where AMD has spent multiple generations building credibility with hyperscalers and enterprise OEMs. Server CPU qualification takes time. Once a cloud instance family or enterprise server platform is built around a CPU generation, the customer does not switch casually. Performance per watt and compatibility create a repeat-purchase flywheel. That moat is not unbreakable, but it is durable enough to support a high-quality server CPU business.

The Instinct moat is earlier and more contested. Hardware roadmaps matter, and management commentary in the data package points to MI350 sampling, MI350 volume production, Helios, MI450, and large multi-year deployment anchors [data_package Section 6]. But AI accelerator moats are ecosystem moats. NVIDIA's CUDA, libraries, networking, systems, installed base, developer familiarity, and procurement mindshare are the benchmark. AMD's ROCm progress can narrow the gap, and large customers can do significant porting work if the economic payoff is attractive. Still, AMD must prove not only that it can ship competitive silicon but that customers can deploy it at high utilization without unacceptable engineering friction.

The OpenAI and Meta agreements in the data package strengthen the moat but do not eliminate execution risk. The OpenAI agreement covers a 6GW multi-year, multi-generation Instinct GPU partnership, with first MI450 deployments tied to 2H 2026, and the Meta agreement announced up to 6GW of Instinct deployments across multiple generations [data_package Section 9]. These are demand anchors, reference customers, and software-feedback loops. They also create customer concentration and milestone risk. If the deployments hit performance, power, and software targets, AMD's credibility with other customers improves. If they slip, the market may conclude that AMD is still a tactical second source rather than a strategic platform.

Adaptive compute and rack-scale design are moat extenders. Xilinx-derived technology can help AMD differentiate around customer-specific acceleration, networking, and heterogeneous compute. Retaining ZT's rack-scale design capability gives AMD a path to sell at the system architecture layer. The risk is that system-level scope can dilute margin if AMD is pulled toward integration work without commensurate pricing power. A moat built on "we will customize deeply for each hyperscaler" is valuable but not always high-margin.

Over five years, the Data Center moat is good enough for the base case but not good enough to underwrite a no-fail bull case. It protects AMD from being a pure commodity chip supplier. It does not protect AMD from NVIDIA's software lead, custom silicon substitution, export-control disruptions, or customer bargaining power. The segment-level moat verdict is therefore: credible second-platform moat, strongest in EPYC and emerging in Instinct, but still dependent on execution through MI350/MI450 and ROCm/system maturity.

6. Historical Growth Rates - Reproduced From Data Package

The table below reproduces only disclosed AMD dollars from the data package. Product lines without separate dollar disclosure remain marked as not separately disclosed. No EPYC or Instinct historical dollar revenue is back-solved.

Product / sub-segment	FY2023 rev	FY2024 rev	FY2025 / TTM rev	3-yr CAGR / clean trend	Most recent YoY	Data package cite
Data Center segment total	$6.496B	$12.579B	$16.635B	60.0% 2-yr CAGR FY2023-FY2025; 40.2% 3-yr CAGR FY2022-FY2025	FY2025 +32.2%; Q4 FY2025 +39.0%	Section 3 Data Center table; Section 5 Data Center Quarterly Trend
EPYC server CPUs	Not separately disclosed	Not separately disclosed	Not separately disclosed	[DATA GAP: no annual EPYC dollars disclosed]	Qualitative driver only; no exact EPYC YoY disclosed	Section 4 Product table, Data Center EPYC row; Section 4 growth-rate-only table
Instinct GPUs / AI accelerators	Not separately disclosed	Not separately disclosed	Not separately disclosed	[DATA GAP: no annual Instinct dollars disclosed]	Q4 includes about $390M selected MI308 China revenue; full Instinct revenue not disclosed	Section 4 Product table, Instinct row; Section 9 MI308 event
Data-center FPGAs, DPUs, AI NICs, adaptive SoCs	Not separately disclosed	Not separately disclosed	Not separately disclosed	[DATA GAP: no product dollars disclosed inside Data Center]	Not separately disclosed	Section 1 taxonomy; Section 10 open data gaps

Arithmetic tie-out on disclosed segment total: FY2023 $6.496B to FY2024 $12.579B is +93.6%, matching the data package. FY2024 $12.579B to FY2025 $16.635B is +32.2%, matching the data package. FY2023 to FY2025 two-year CAGR is `(16.635 / 6.496)^(1/2) - 1 = 60.0%`. FY2022 to FY2025 three-year CAGR is `(16.635 / 6.043)^(1/3) - 1 = 40.2%` using the FY2022 Data Center segment revenue shown in Section 3. These checks are within the 2% tolerance.

The cleanest historical read is that AMD's Data Center business has already had one extraordinary accelerator-led step-up. FY2024 revenue nearly doubled from FY2023, and FY2025 revenue grew another 32.2% even with export-control disruption. The quarterly sequence also shows demand accelerating into Q4 FY2025, when revenue reached $5.380B and operating margin reached 32.6% [data_package Section 5]. The forward model must respect that momentum, but it should not extrapolate the FY2023-FY2025 two-year CAGR blindly because the segment was moving from a smaller base and the competitive response is intensifying.

7. Management Guidance - Quoted From Data Package Section 6

The data package contains the relevant Data Center management commentary for FY2025. The following rows are reproduced as short snippets from that package, with interpretation for the segment model.

Data package row	Speaker / source	Verbatim quote	Interpretation for model
Q1 FY2025 call, 2025-05-06	Lisa Su	"We began sampling our next-gen MI350 series"	MI350 was already in customer hands early in FY2025, supporting a FY2026 ramp rather than a purely speculative product launch.
Q1 FY2025 call, 2025-05-06	Lisa Su	"actively working with multiple customers to scale Instinct"	Instinct demand was described as multi-customer scaling, not only one-off pilots.
Q2 FY2025 call, 2025-08-05	Lisa Su	"We began volume production of the MI350 series ahead of schedule"	This supports strong near-term Data Center growth assumptions after the MI308 export-control disruption.
Q2 FY2025 call, 2025-08-05	Lisa Su	"Helios is expected to deliver up to a 10x generational performance increase"	AMD is trying to compete at rack/system level, not only at accelerator card level.
Q3 FY2025 call, 2025-11-04	Lisa Su	"deploy 6 gigawatts of Instinct GPUs"	The OpenAI agreement is a specific multi-year demand anchor.
Q3 FY2025 call, 2025-11-04	Lisa Su	"first gigawatt of MI450 Series accelerators"	First OpenAI deployment timing is linked to MI450 in 2H 2026.
Q4 FY2025 call, 2026-02-03	Lisa Su	"well-positioned to grow data center segment revenue by more than 60% annually"	The FY2026 model should not assume immediate deceleration to mature-semiconductor growth.
Q4 FY2025 call, 2026-02-03	Lisa Su	"scale our AI business to tens of billions in annual revenue in 2027"	The base case must allow the AI accelerator business to become very large by 2027, while still respecting the lack of historical product-dollar disclosure.

Data package row	Speaker / source	Verbatim quote	Interpretation for model
Q4 FY2025 call, 2026-02-03	Lisa Su	"I do not believe that we will be supply limited"	The limiting variable is more likely customer adoption, software, pricing, and execution than physical supply in the stated plan.
Financial Analyst Day, 2025-11-11	AMD press release	"greater than 35% revenue CAGR"	AMD's company-level long-term target implies Data Center carries a large share of the growth burden.

The most important modeling consequence is that FY2026 and FY2027 should be high-growth years in any base case. A model that starts FY2026 at only 20% growth would contradict the Q4 FY2025 commentary unless it had a severe offsetting assumption. Conversely, a model that compounds 60%+ for all five years would ignore the competitive context, customer concentration, and natural deceleration from a much larger revenue base.

8. Forward Growth Assumptions - Explicit, Per Product Where Disclosed

Because AMD does not disclose EPYC or Instinct annual dollars, the forecast line with actual revenue control is the Data Center segment total. The product rows explain the drivers, not hidden dollar splits. This is a deliberate anti-fabrication choice.

Base Case Growth

Product / line	Historical CAGR	Latest YoY	Mgmt FY+1 guide anchor	Base CAGR TTM->TTM+5y	Justification
Data Center segment total	60.0% FY2023-FY2025 2-yr CAGR; 40.2% FY2022-FY2025 3-yr CAGR	FY2025 +32.2%; Q4 +39.0%	"more than 60% annually"; AI to "tens of billions" in 2027	24.2%	Base assumes FY2026 +60%, FY2027 +35%, FY2028 +20%, FY2029 +10%, and horizon year +3.5%. This respects the management ramp but decelerates as revenue approaches $50B and custom silicon/NVIDIA competition becomes harder.
EPYC server CPUs	No separate dollars	Qualitative share-gain driver	No product-dollar guide	Not separately modeled	EPYC supports the segment total through share gains and platform attach. It should be steadier than Instinct but cannot be given a separate revenue CAGR without fabricating a starting base.
Instinct AI accelerators / rack-scale systems	No separate annual dollars	Q4 included selected MI308 China revenue, but full Instinct not disclosed	AI to "tens of billions" in 2027; OpenAI first MI450 GW in 2H 2026	Not separately modeled	Instinct is the main early-growth driver. Base assumes major 2026-2027 ramp, then deceleration as customer deployments move from scarce-capacity buying to utilization and TCO comparisons.
Adaptive data-center silicon, DPUs, AI NICs	No separate dollars	Not disclosed	No product-dollar guide	Not separately modeled	These lines support system differentiation and attach, but the data package does not support a separate dollar model.

Base case year-by-year segment revenue path: FY2025/TTM $16.635B; FY2026 $26.616B; FY2027 $35.932B; FY2028 $43.118B; FY2029 $47.430B; TTM+5y/horizon $49.090B. The five-year CAGR is 24.2%. This is a steep deceleration from the FY2023-FY2025 period, but it still nearly triples segment revenue from the TTM base.

The base case is designed to pass two tests. First, it is not dismissive of management's stated AI ramp. FY2026 +60% and FY2027 +35% are large increases and are consistent with Data Center becoming AMD's growth engine. Second, it does not assume AMD escapes the gravity of NVIDIA, custom ASICs, and customer bargaining power. NVIDIA's FY2026 Data Center run rate and software ecosystem mean AMD can grow rapidly as a second source while still facing price and utilization constraints.

Bull Case Growth

Product / line	Historical CAGR	Latest YoY	Mgmt FY+1 guide anchor	Bull CAGR TTM->TTM+5y	Justification
Data Center segment total	60.0% FY2023-FY2025 2-yr CAGR	FY2025 +32.2%; Q4 +39.0%	Data Center "more than 60% annually"; AI "tens of billions" in 2027	37.2%	Bull assumes FY2026 +70%, FY2027 +55%, FY2028 +35%, FY2029 +22%, and horizon year +12%. This requires OpenAI/Meta ramps to validate AMD as a strategic second AI platform, plus continued EPYC share gains.
EPYC server CPUs	No separate dollars	Qualitative driver	No product-dollar guide	Not separately modeled	Bull case assumes EPYC retains high-value cloud and enterprise wins while Intel's recovery is not enough to reverse AMD share gains.
Instinct AI accelerators / rack-scale systems	No separate annual dollars	Full Instinct revenue not disclosed	Multi-GW agreements	Not separately modeled	Bull case requires MI350/MI450 performance, supply, ROCm maturity, and rack-scale deployments to make AMD more than a tactical second source.
Adaptive data-center silicon, DPUs, AI NICs	No separate dollars	Not disclosed	No product-dollar guide	Not separately modeled	Bull case assumes system-level attach helps win deployments rather than becoming low-margin customization.

Bull case year-by-year segment revenue path: FY2025/TTM $16.635B; FY2026 $28.280B; FY2027 $43.834B; FY2028 $59.176B; FY2029 $72.195B; TTM+5y/horizon $80.857B. The five-year CAGR is 37.2%. This case is not "AMD becomes NVIDIA." NVIDIA's FY2026 Data Center revenue was already far above this horizon revenue. The bull case instead says AMD becomes a highly credible second AI platform and keeps a strong server CPU base.

Bear Case Growth

Product / line	Historical CAGR	Latest YoY	Mgmt FY+1 guide anchor	Bear CAGR TTM->TTM+5y	Justification
Data Center segment total	40.2% FY2022-FY2025 3-yr CAGR	FY2025 +32.2%; Q4 +39.0%	Data Center "more than 60% annually" aspiration	12.5%	Bear assumes FY2026 +30%, FY2027 +20%, FY2028 +10%, FY2029 +5%, and horizon year 0%. This reflects a scenario where AI accelerator deployments occur but AMD remains mostly a second source and custom silicon captures much of inference growth.
EPYC server CPUs	No separate dollars	Qualitative driver	No product-dollar guide	Not separately modeled	Bear case assumes Intel stabilizes server share and cloud Arm CPUs pressure some scale-out workloads.
Instinct AI accelerators / rack-scale systems	No separate annual dollars	Q2 export charge shows downside risk	Multi-GW agreements, but execution risk	Not separately modeled	Bear case assumes deployment delays, ROCm friction, pricing concessions, or export controls limit the conversion of headline GW agreements into high-margin revenue.
Adaptive data-center silicon, DPUs, AI NICs	No separate dollars	Not disclosed	No product-dollar guide	Not separately modeled	Bear case assumes attach is useful but not enough to offset accelerator competition.

Bear case year-by-year segment revenue path: FY2025/TTM $16.635B; FY2026 $21.626B; FY2027 $25.951B; FY2028 $28.546B; FY2029 $29.973B; TTM+5y/horizon $29.973B. The five-year CAGR is 12.5%. This is still growth from the TTM base, because AMD has real demand anchors and EPYC remains competitive. The bear case is about disappointment versus AI expectations, not a collapse of the business.

9. Forward Operating Margin Assumptions

Current TTM Data Center operating margin is 21.7% [data_package Section 3]. The most recent quarter, Q4 FY2025, had 32.6% operating margin, while Q2 FY2025 had negative 4.8% margin because of the MI308/export-control charge and related disruption [data_package Section 5; Section 9]. The margin forecast should therefore normalize between depressed annual margin and strong exit-rate margin.

Case	TTM margin	TTM+5y operating margin	Mechanism
Bear	21.7%	20.5%	Pricing concessions, second-source positioning, export-control/inventory risk, and higher customer engineering costs offset revenue growth. Margin stays near FY2025 despite scale.
Base	21.7%	28.5%	MI308-type charge does not recur at the same scale, AI accelerator revenue scales, EPYC remains profitable, and operating expenses grow slower than revenue. Margin remains below Q4 FY2025's 32.6% run rate to reflect competition.
Bull	21.7%	33.0%	AMD reaches strategic-platform status with large customers, rack-scale systems improve realized value, software maturity reduces support burden, and EPYC/Instinct mix supports pricing. Margin roughly matches the strong Q4 FY2025 exit rate rather than exceeding it materially.

The base case margin expands 680 bps from FY2025. The mechanism is operating leverage plus mix, not accounting magic. Q4 FY2025 already demonstrated that a 30%+ Data Center margin is possible in a clean quarter [data_package Section 5]. But the annual FY2025 margin also proves that one product/geography disruption can erase a large amount of segment operating income. A normalized 28.5% horizon margin is therefore intentionally below the Q4 exit rate.

The bull case margin of 33.0% is demanding. It assumes AMD captures enough AI accelerator value to price above a commodity second-source level and that rack-scale design adds differentiation without dragging the company into lower-margin system integration. It also assumes EPYC remains a high-quality CPU franchise. The bull case does not assume NVIDIA-like platform economics because NVIDIA's scale, software ecosystem, and networking position remain much stronger.

The bear case margin of 20.5% is not punitive relative to FY2025, but it is a major disappointment relative to the AI thesis. In this case, AMD can grow revenue while margins fail to expand because hyperscalers use AMD to bargain against NVIDIA, custom silicon absorbs inference, ROCm integration costs remain high, and export-control risk recurs. This is a plausible downside because the most important AI customers are technically sophisticated and have the purchasing power to pressure second-source suppliers.

10. Segment Rollup - Built From Disclosed Control Total

True product-level arithmetic is not possible because AMD does not disclose EPYC, Instinct, or other Data Center product revenue. The rollup below therefore treats the disclosed Data Center segment as the only dollar-controllable revenue line and uses product drivers qualitatively. This is a [DATA GAP] caused by company disclosure, not by missing work in this appendix.

Case	TTM revenue	Revenue path assumption	TTM+5y revenue	TTM+5y margin	TTM+5y operating income	Arithmetic check
Bear	$16.635B	+30%, +20%, +10%, +5%, +0%	$29.973B	20.5%	$6.144B	$29.973B x 20.5% = $6.144B
Base	$16.635B	+60%, +35%, +20%, +10%, +3.5%	$49.090B	28.5%	$13.991B	$49.090B x 28.5% = $13.991B
Bull	$16.635B	+70%, +55%, +35%, +22%, +12%	$80.857B	33.0%	$26.683B	$80.857B x 33.0% = $26.683B

The base case is the recommended synthesis input for this segment. It reflects the data package's high-growth management commentary, the FY2025/Q4 momentum, and the OpenAI/Meta multi-year demand anchors. It also discounts the forecast for the competitive facts that cannot be ignored: NVIDIA's platform lead, custom silicon from hyperscalers, export-control risk, and customer concentration.

The probability weights used here are Bear 25%, Base 50%, Bull 25%. The probability-weighted Data Center TTM+5y revenue is $52.253B and probability-weighted operating income is $15.202B. These probability weights are not intended to override the synthesis agent's company-level scenario design; they are provided so the segment's skew is visible. The revenue distribution is positively skewed because multi-GW deployments can create large upside if execution works, but the downside is not zero because AMD already has a real $16.635B Data Center business.

11. Key Risks And Wildcards

- AI platform adoption risk - Bear/base. AMD can ship competitive accelerators and still disappoint if customers cannot deploy them at high utilization with acceptable software effort. NVIDIA's CUDA/software ecosystem is the main barrier.
- Customer concentration and milestone risk - Bear/base/bull. The OpenAI and Meta agreements are valuable anchors, but they make execution milestones more important. A delay in MI450 deployment timing or lower-than-expected utilization would hit both revenue credibility and margin.
- Custom silicon substitution - Bear. AWS Trainium, Google TPU, Microsoft Maia, and Broadcom-enabled custom XPUs can absorb large portions of training or inference demand, especially for hyperscalers' internal workloads.
- Export controls and China exposure - Bear. The data package already includes an $800M MI308 inventory and related charge in Q2 FY2025 and selected MI308 China revenue in Q4 FY2025. Future restrictions could again affect product planning and margin.
- EPYC competitive response - Base/bear. Intel's DCAI segment was growing again in 2025, and cloud Arm CPUs continue to improve. EPYC remains strong, but the CPU share-gain story should not be treated as riskless.
- Rack-scale economics - Base/bull/bear. Retaining ZT's design capability can help AMD sell systems-level value. It can also increase support burden and lower blended margin if customers force AMD into bespoke deployments without enough pricing.
- Supply chain and HBM availability - Bull/base. Management said supply should not be the limiting factor for the plan [data_package Section 6], but AI accelerators still depend on advanced nodes, HBM, packaging, substrates, and networking components. The bull case requires smooth execution across that chain.
- Inference economics - Bull/bear. If inference demand explodes and customers need a merchant alternative to NVIDIA, AMD's opportunity expands. If inference moves aggressively to internal ASICs and optimized cloud chips, AMD's accelerator TAM is smaller than headline AI compute growth suggests.

Summary Table

Case	Segment TTM+5y Revenue	Segment TTM+5y Op Margin	Segment TTM+5y Op Income	Probability
Bear	$29.973B	20.5%	$6.144B	25%
Base	$49.090B	28.5%	$13.991B	50%
Bull	$80.857B	33.0%	$26.683B	25%

Data Gaps Inserted

- [DATA GAP: no annual EPYC dollars disclosed] AMD discloses Data Center segment revenue but not EPYC revenue.
- [DATA GAP: no annual Instinct dollars disclosed] AMD discloses selected MI308 China revenue and inventory charges, but not full Instinct annual revenue.
- [DATA GAP: no product dollars disclosed inside Data Center] AMD does not disclose separate data-center FPGA, DPU, AI NIC, or adaptive SoC revenue.
- [DATA GAP: true product-level rollup unavailable] The model uses the disclosed Data Center segment control total and qualitative product drivers rather than fabricated product revenue splits.

Appendix B2 - Client and Gaming

Client and Gaming - Segment Deep Dive

Company: Advanced Micro Devices, Inc. Ticker: AMD Segment: Client and Gaming TTM anchor: FY2025 / Q1-Q4 FY2025 ending 2025-12-27 Forward horizon: 2030-12-27, exactly five calendar years after the TTM end date

0. Segment Boundary And Taxonomy Note

Client and Gaming is a current AMD reportable segment beginning with the Q1 FY2025 reporting change. The data package states that AMD combined the prior Client and Gaming reportable segments into one reportable segment in Q1 FY2025 and retrospectively adjusted FY2023-FY2024 presentation, while FY2021-FY2022 remain older-presentation years bridged by summing Client plus Gaming for directional continuity [data_package section 1; section 3, Client and Gaming row]. This appendix therefore treats Client and Gaming as one reportable operating segment, models the disclosed Client and Gaming revenue components separately, and does not double-count Gaming against the combined segment. Operating income is available only for the combined segment, so any component margin split below is an explicit model allocation, not a disclosed AMD fact [data_package section 10, Open Data Gaps].

1. What This Segment Actually Is

AMD's Client and Gaming segment is the PC and game-console side of the company. Client consists mainly of Ryzen desktop and notebook processors, accelerated processing units, and chipsets sold into consumer, commercial, channel, and OEM PC markets. Gaming consists of semi-custom system-on-chip products, development services, and Radeon discrete GPUs sold into game consoles, gaming PCs, and related graphics channels. The data package's product table describes Client as Ryzen desktop/mobile processors, APUs, and chipsets, and Gaming as semi-custom SoCs plus Radeon gaming GPUs [data_package section 4, Client and Gaming product rows].

The segment's customers are not a single buyer class. Client has large OEMs, commercial PC fleets, retail channel partners, system integrators, motherboard partners, and end users who choose Ryzen systems through OEM or do-it-yourself channels. Gaming has two very different businesses inside one revenue component. Semi-custom console SoCs are sold to platform owners through long design cycles, high-volume production ramps, and negotiated supply agreements. Radeon gaming GPUs are sold through board partners and OEM systems into a much more competitive and product-cycle-driven PC graphics market.

The pricing model is also mixed. Client CPUs and APUs have recurring product-cycle pricing, with high-end desktop and premium mobile products carrying materially better average selling prices than entry client products. AMD management's FY2025 commentary in the data package explicitly links the Client rebound to richer high-end desktop and mobile Ryzen mix, not just unit recovery [data_package section 6, Client and Gaming, Q1 FY2025 quote row]. That matters because the segment can grow revenue faster than PC units when AMD gains share in premium desktop, notebook, commercial, and AI PC configurations. In contrast, gaming semi-custom pricing is negotiated around console platform economics, bill-of-material targets, wafer costs, launch timing, and lifecycle volumes. Radeon discrete GPU pricing is exposed to Nvidia's product stack, channel inventory, gamer upgrade timing, and the availability of cards at suggested retail prices.

The segment is capital-light at the company level because AMD is fabless. It depends on external wafer, packaging, substrate, memory, and board ecosystems, but AMD does not own leading-edge fabs. Segment operating economics are therefore driven by gross margin, mix, R&D allocation, go-to-market spending, and the degree to which design wins amortize over a large revenue base. Client has a repeatable cadence of Zen CPU architectures, mobile APU launches, AI PC platforms, and commercial validation cycles. Gaming semi-custom has lower frequency but longer platform lives. Radeon has the weakest structural position because Nvidia dominates discrete gaming GPU share and software ecosystem mindshare, but AMD can still participate in mainstream price/performance tiers, integrated graphics, handheld gaming APUs, and console SoCs.

2. Current Economics

The current TTM base is FY2025 because the Stage 0 anchor says Q1-Q4 FY2025 ending 2025-12-27 is the latest reported TTM period. AMD's Client and Gaming segment produced $14.550B of revenue in FY2025, up 50.8% year over year from $9.649B in FY2024, and generated $2.855B of operating income, a 19.6% operating margin [data_package section 3, Client and Gaming table; source: FY2025 10-K segment table and Q4/FY2025 release]. This made the segment 42.0% of AMD's

FY2025 revenue and 37.1% of reportable-segment operating income before corporate-level unallocated items, calculated from the Data Center, Client and Gaming, and Embedded segment operating income rows in data_package section 3.

The disclosed component split was $10.640B of Client revenue and $3.910B of Gaming revenue in FY2025 [data_package section 4, Product and Sub-Segment Historicals, Client and Gaming rows]. Client was therefore 73.1% of segment revenue, while Gaming was 26.9%. This is the starting point for the forecast. It is also why applying one segment-wide growth rate would be misleading. Client is in a share-gain and mix-upcycle phase; Gaming contains a console-cycle headwind and a structurally challenged Radeon dGPU business.

The most recent quarterly trend shows momentum but also deceleration from a very easy comparison base. Segment revenue was $2.941B in Q1 FY2025, $3.621B in Q2, $4.048B in Q3, and $3.940B in Q4, with year-over-year growth of 28%, 69%, 73%, and 37%, respectively [data_package section 5, Client and Gaming Quarterly Trend]. Operating margin was 16.9%, 21.2%, 21.4%, and 18.4% across those quarters, so the FY2025 exit rate was profitable but not linearly improving [data_package section 5, Client and Gaming Quarterly Trend]. The Q4 margin dip versus Q2-Q3 is consistent with mix and gaming/semi-custom timing rather than a clean monotonic operating leverage story.

At the component level, Client revenue was $2.294B in Q1, $2.499B in Q2, $2.750B in Q3, and $3.097B in Q4, with year-over-year growth of 68%, 67%, 46%, and 34% [data_package section 5, Disclosed Client and Gaming Components]. Gaming revenue was $0.647B in Q1, $1.122B in Q2, $1.298B in Q3, and $0.843B in Q4, with year-over-year growth of -30%, 73%, 181%, and 50% [data_package section 5, Disclosed Client and Gaming Components]. Those component quarterly figures sum exactly to the reported segment revenue in each quarter, which is a useful control check and prevents double counting.

The data package does not disclose Client operating income or Gaming operating income separately. It explicitly flags that Client and Gaming operating income is disclosed only for the combined current segment, not for the components [data_package section 10, Open Data Gaps]. The appendix therefore treats the FY2025 segment margin of 19.6% as the only hard margin anchor, and uses component margin allocations only to build a transparent weighted-average forecast. The allocation that ties FY2025 is Client at roughly 22% operating margin and Gaming at roughly 13% operating margin, producing $2.855B of segment OI on the disclosed revenue mix. That allocation is analytically plausible because high-end Ryzen client CPUs carry better gross margin and semi-custom console silicon is usually lower-margin and more customer-concentrated, but it is not disclosed.

3. Unit Economics

For this segment, revenue becomes profit through product mix, architecture cadence, wafer cost, and operating leverage. AMD sells intellectual-property-rich chips rather than finished PCs or consoles. The company absorbs R&D to design CPU cores, GPU architectures, chiplets, APUs, packaging, firmware, software, platform enablement, and semi-custom SoCs. Once products are in market, gross margin depends on wafer pricing, yield, packaging cost, die size, channel pricing, competitive pressure, and whether the SKU is high-end desktop/mobile, commercial client, mainstream consumer, console semi-custom, or discrete GPU.

Client unit economics are the stronger half of the segment. Ryzen desktop and notebook processors have recurring replacement demand, a broad OEM and channel base, and an opportunity to gain share from Intel when product performance, power efficiency, or supply availability is favorable. AMD management said in Q1 FY2025 that it delivered record client CPU ASP because of richer high-end desktop and mobile Ryzen mix [data_package section 6, Client and Gaming, Q1 FY2025 row]. That is a direct unit-economics signal: growth in premium Ryzen mix can expand revenue and margin without requiring equal unit growth. The Q4 FY2025 quote that AMD believes it can grow the PC business and that focus areas are enterprise also points to commercial mix as a margin lever [data_package section 6, Client and Gaming, Q4 FY2025 rows].

Client margins can break upward if AMD sustains high desktop share, wins more premium notebooks, converts enterprise fleets, and uses AI PC features to protect ASP. They can break downward if the PC market turns into a price war, if Intel regains product leadership or floods the channel, if Arm-based Windows PCs take meaningful consumer notebook share, or if memory and component inflation forces OEMs to push down CPU platform pricing. The segment is not immune to PC unit cycles; Gartner estimated 2025 PC shipments grew 9.1% to more than 270M units, driven by Windows 11 upgrades and pre-buying ahead of memory price hikes, but Gartner also noted that many AI PC upgrades were more about future-proofing than immediate productivity value (Gartner January 2026 PC market release: https://www.gartner.com/en/newsroom/press-releases/2026-1-20-gartner-says-worldwide-pc-shipments-increased-9-point-3-percent-in-fourth-quarter-of-2025-and-9-point-1-percent-for-the-full-year). That implies some of FY2025 demand could be pulled forward.

Gaming unit economics are less clean. Semi-custom SoCs are sticky once AMD wins a console platform because the chip is designed into a multi-year hardware cycle, but the revenue curve is cyclical. Volumes peak after launch and then fade until a next-generation refresh. AMD management said semi-custom SoC annual revenue is expected to decline by a significant double-digit percentage, while the next-gen Xbox semi-custom SoC is progressing well [data_package section 6, Client and Gaming, Q4 FY2025 rows]. That combination supports a trough-and-recovery model rather than a straight-line growth model.

Radeon discrete GPU economics are the weaker link. Discrete gaming graphics cards require costly architecture, driver, software, and board-partner support, yet Nvidia has overwhelming share and stronger ecosystem gravity. Jon Peddie Research reported Q4 FY2025 total PC GPU shipments decreased 3.3% year over year and that discrete GPU penetration over the next five years is expected to be about 25% of PCs, while AMD's overall GPU share improved sequentially but the market remained difficult (JPR Q4 FY2025 PC GPU release: https://www.jonpeddie.com/news/q425-pc-gpu-shipments-decreased-by-1-3-from-last-quarter-which-is-a-prelude-to-the-creep-forward/). For AMD, the Radeon opportunity is not to match Nvidia in every high-end enthusiast tier, but to defend mainstream value, integrated graphics, handheld APUs, and console-related GPU IP.

Operating expenses matter because the same Zen and graphics engineering base supports multiple segments. Some Client and Gaming R&D benefits Data Center and Embedded, and vice versa. That makes segment OI an allocation as well as an economic measure. Still, the reported margin history is informative: Client and Gaming margin fell from 24.2% in FY2021 to 16.5% in FY2022, 8.5% in FY2023, then recovered to 12.3% in FY2024 and 19.6% in FY2025 [data_package section 3, Client and Gaming table]. The sharp FY2023 trough shows how quickly PC and gaming cyclicality can damage profit, while FY2025 shows the upside of share gains, ASP recovery, and better utilization.

4. Competitive Dynamics

The Client business competes primarily with Intel in x86 desktop and notebook CPUs, with Qualcomm and Apple representing Arm-based competition in parts of the notebook market. Intel remains the broad installed-base incumbent, with deep OEM relationships, vPro enterprise manageability, platform validation, and internal manufacturing scale. AMD's edge is execution on CPU architecture, chiplet design, performance per watt, integrated graphics in APUs, and high-end desktop credibility. The external share context supports the data package's management comments: PCWorld, citing Mercury Research figures supplied by AMD, reported that AMD desktop unit share rose from 26.9% to 36.4% in Q4 FY2025 and desktop revenue share rose from 28.0% to 42.6%, while mobile share rose from 23.8% to 26.0% and mobile revenue share from 21.6% to 24.9% (PCWorld/Mercury Research summary: https://www.pcworld.com/article/3058885/amd-surges-10-percent-in-desktop-pc-share.html). TechSpot similarly reported AMD desktop revenue share of 42.6% and mobile unit share around 26% based on Mercury Research (https://www.techspot.com/news/111297-ryzen-revenue-share-jumps-146-amd-takes-426.html). These are external competitive context figures, not AMD historical revenue figures.

The immediate Client market backdrop is a blend of cyclical recovery and structural AI PC transition. Gartner estimated AI PCs represented 31% of the worldwide PC market in 2025 and projected 55% in 2026, with x86 expected to make up a large majority of AI business laptops in 2025 (Gartner AI PC release: https://www.gartner.com/en/newsroom/press-releases/2025-08-28-gartner-says-artificial-intelligence-pcs-will-represent-31-percent-of-worldwide-pc-market-by-the-end-of-2025). That helps AMD because Ryzen AI products can compete in an upgrade category where enterprises care about Windows compatibility, performance per watt, and on-device inference readiness. It also raises competitive risk because AI PC differentiation may narrow as NPUs become table stakes and OEMs compare total platform cost.

Intel's response is the largest Client uncertainty. If Intel's process roadmap, mobile power efficiency, and platform incentives improve, AMD's recent share gains could slow. Intel can also use pricing, bundling, marketing development funds, and enterprise validation to protect share. But AMD does not need majority client CPU share to grow; it can compound from a still-minority share base if it wins premium desktop, commercial mobile, and high-ASP gaming/creator systems. The base case assumes share gains moderate, not reverse.

Gaming competition splits by product. In semi-custom console SoCs, AMD's main competitive advantage is incumbency and accumulated IP. AMD powers current major console platforms and has direct customer relationships with platform owners, but those wins are cyclical and customer-concentrated. The next console cycle can renew the business, but the timing is outside AMD's sole control. In Radeon discrete GPUs, Nvidia is the dominant competitor. JPR's add-in-board data as reported by multiple industry outlets showed Nvidia around the mid-90% range in Q4 FY2025 desktop AIB share, with AMD in the single digits; even the primary JPR release states AIB shipments were down sequentially and Nvidia gained share while AMD lost share in Q4 FY2025 (JPR AIB release: https://www.jonpeddie.com/news/q425-pc-graphics-aib-shipments-decreased-4-4-fro

m-last-quarter-to-11-million-units-with-a-cagr-to-2029-of-5-9/). That is a severe competitive constraint on Radeon margin and market share.

Substitution risk is higher in Client than it was a decade ago. Apple silicon has proven that Arm-based PCs can dominate in a vertically integrated ecosystem. Qualcomm and other Arm Windows suppliers can attack consumer notebooks, especially if software compatibility improves and battery life matters more than x86 application edge cases. Gartner's AI PC release notes Arm-based laptops gaining larger consumer share as application compatibility challenges are overcome, while x86 remains stronger in business laptops (Gartner AI PC release). AMD's defensive position is better in enterprise x86, gaming desktops, creators, and systems that benefit from CPU/GPU/NPU balance. Its weaker position is low-end consumer notebooks where price and battery claims can matter more than peak CPU performance.

Pricing power is therefore product-specific. AMD has pricing power in high-end desktop gaming CPUs when 3D V-Cache and performance leadership are clear. It has some pricing power in premium mobile if Ryzen AI platforms win OEM design slots and battery/performance reviews. It has limited pricing power in mainstream PC CPUs, because Intel can respond aggressively and OEMs multi-source. In semi-custom, pricing power is negotiated upfront and limited by console economics. In Radeon, pricing power is weak unless AMD has a clear performance-per-dollar advantage or supply scarcity favors its channel.

5. Segment-Level Moat

The Client and Gaming segment has a real but narrower moat than AMD's Data Center segment. The moat is strongest in CPU architecture execution, chiplet design, x86 compatibility, high-end desktop brand equity, console semi-custom relationships, and the engineering reuse that lets AMD apply CPU, GPU, and packaging IP across multiple end markets. It is weakest in discrete gaming GPUs, where Nvidia's hardware, software, CUDA-adjacent developer mindshare, creator workflows, ray-tracing perception, and installed base make AMD a price/performance challenger rather than a category controller.

Client's moat is not that PC CPUs are impossible to copy. It is that matching AMD's performance per watt, CPU roadmap, platform compatibility, OEM validation, and channel brand at the same time is hard. AMD's Ryzen brand has become credible with enthusiasts and increasingly credible with commercial buyers. Its 3D V-Cache products create a gaming-specific differentiation that Intel has struggled to neutralize consistently. Its AI PC roadmap gives it an NPU story at a time when OEMs are refreshing fleets for Windows 11 and future on-device AI workloads. Those advantages are durable for several years if AMD keeps executing, but they are not permanent. Intel can close gaps, Arm can erode consumer notebook share, and AI PC features can become commoditized.

Gaming's moat is two-part. Semi-custom is protected by long design cycles, platform-owner switching costs, shared software/hardware optimization, and console lifecycle inertia. Once a console uses AMD silicon, replacement does not happen mid-cycle. But the moat only protects installed design wins; it does not guarantee the next cycle. Radeon is less protected. It benefits from AMD's GPU IP, open ecosystem, integrated graphics footprint, and console/handheld relevance, but Nvidia's discrete GPU position limits AMD's ability to earn premium margins. Therefore the Gaming component has a moat in semi-custom design-in duration but not in standalone PC gaming graphics.

At the combined segment level, the moat should support mid-to-high-teens operating margins across a cycle and low-20s margins in favorable Client upcycles. It does not support assuming permanently rising margins. The FY2023 operating margin trough of 8.5% in the data package is a warning that PC correction and console timing can quickly overwhelm fixed R&D leverage [data_package section 3, Client and Gaming table]. The FY2025 margin of 19.6% shows that the segment can earn solid profits when Ryzen mix and volume are favorable [data_package section 3]. The base case therefore assumes modest margin expansion to 21.0% by TTM+5, not a straight move toward Data Center-like or software-like margins.

6. Historical Growth Rates

The table below reproduces the data package's disclosed segment and component revenue. FY2023-FY2025 are the clean current-taxonomy period. FY2021-FY2022 segment totals are directional Client-plus-Gaming sums from the older reporting presentation and are not used for clean dollar comparisons. Product-level component revenue is disclosed only for Client and Gaming revenue, not operating income [data_package section 4, Product and Sub-Segment Historicals].

Product / sub-segment	FY2023 rev	FY2024 rev	FY2025 rev / TTM	Clean FY2023-FY2025 CAGR	Most recent YoY	Data package cite
Total Client and Gaming	$10.863B	$9.649B	$14.550B	15.7%	50.8%	data_package section 3, Client and Gaming table row

Product / sub-segment	FY2023 rev	FY2024 rev	FY2025 rev / TTM	Clean FY2023-FY2025 CAGR	Most recent YoY	Data package cite
Client	$4.651B	$7.054B	$10.640B	51.2%	50.8%	data_package section 4, Client product row
Gaming	$6.212B	$2.595B	$3.910B	-20.7%	50.7%	data_package section 4, Gaming product row

The arithmetic tie-out is internally consistent. Total Client and Gaming declined from $10.863B in FY2023 to $9.649B in FY2024, a -11.2% move, and then grew to $14.550B in FY2025, a 50.8% move; those are the exact YoY values in data_package section 3. Client grew from $4.651B to $7.054B to $10.640B, producing 51.7% and 50.8% YoY growth; the last-year YoY matches the component table and explains most of the combined segment rebound [data_package section 4]. Gaming declined from $6.212B to $2.595B, then recovered to $3.910B, producing -58.2% and 50.7% YoY; this is why its two-year CAGR remains negative even though FY2025 growth was strong.

Directional longer history from the data package reinforces the cyclicality. The bridged Client and Gaming total was $12.494B in FY2021, $13.006B in FY2022, $10.863B in FY2023, $9.649B in FY2024, and $14.550B in FY2025 [data_package section 3, Client and Gaming table]. The segment's FY2025 revenue was only 11.9% above FY2022 despite the 50.8% FY2025 rebound, which shows that some of the growth was recovery from a PC and console-cycle trough rather than a new secular 50% growth baseline.

The quarterly component table also matters for forecast shape:

Quarter	Client revenue / YoY	Gaming revenue / YoY	Segment revenue / YoY	Data package cite
Q1 FY2025	$2.294B / 68%	$0.647B / -30%	$2.941B / 28%	data_package section 5, Disclosed Client and Gaming Components
Q2 FY2025	$2.499B / 67%	$1.122B / 73%	$3.621B / 69%	data_package section 5, Disclosed Client and Gaming Components
Q3 FY2025	$2.750B / 46%	$1.298B / 181%	$4.048B / 73%	data_package section 5, Disclosed Client and Gaming Components
Q4 FY2025	$3.097B / 34%	$0.843B / 50%	$3.940B / 37%	data_package section 5, Disclosed Client and Gaming Components

No product-dollar back-solving is used in this section. All cells come from disclosed data-package rows. The only data gap is margin by component: [DATA GAP: AMD does not disclose Client operating income or Gaming operating income separately; only combined Client and Gaming operating income is disclosed in data_package sections 3 and 10].

7. Management Guidance

The following guidance and commentary items come from data_package section 6. No transcript was re-fetched for AMD quotes.

Data package row	Verbatim quote	Interpretation
Q1 FY2025 call / event, 2025-05-06 / Lisa Su, Chair and CEO	"We delivered record client CPU ASP driven by a richer mix of high-end desktop and mobile Ryzen processors"	The FY2025 Client rebound was mix- and ASP-led, not merely shipment recovery. This supports a higher Client margin allocation than Gaming.
Q2 FY2025 call / event, 2025-08-05 / Lisa Su, Chair and CEO	"record desktop CPU sales"	Desktop Ryzen strength persisted beyond one quarter and supports the external market-share evidence that AMD gained share in premium desktop.
Q3 FY2025 call / event, 2025-11-04 / Lisa Su, Chair and CEO	"record quarterly sales"	The Client business was still setting records in Q3, but the quarterly YoY deceleration from 68%-67% to 46% shows that the comparison base was already getting tougher.
Q4 FY2025 call / event, 2026-02-03 / Lisa Su, Chair and CEO	"we believe we can grow our PC business"	Management expects share and mix to offset a softer or more mature PC TAM. The base case treats this as credible but not open-ended.
Q4 FY2025 call / event, 2026-02-03 / Lisa Su, Chair and CEO	"our focus areas are enterprise"	Enterprise PCs are the key forward lever because they can improve mix, visibility, and OEM stickiness relative to consumer-only share gains.

Data package row	Verbatim quote	Interpretation
Q4 FY2025 call / event, 2026-02-03 / Lisa Su, Chair and CEO	"semi-custom SoC annual revenue to decline by a significant double-digit percentage"	This is the clearest FY2026 headwind. The forecast starts Gaming with a negative year before modeling later console-cycle recovery.
Q4 FY2025 call / event, 2026-02-03 / Lisa Su, Chair and CEO	"next-gen Xbox featuring an AMD semi-custom SoC is progressing well"	Gaming can recover later in the horizon, but the timing is likely back-half weighted and not enough to erase the near-term semi-custom decline in the base case.
Q4 FY2025 call / event, 2026-02-03 / Jean Hu, CFO	"client on the business gross margin has been improving nicely"	Margin improvement is tied to Client mix and supports modest segment margin expansion if Client remains a larger share of the segment.

The main message is mixed. Management commentary is constructive on Client and explicitly negative on near-term semi-custom. It is therefore inappropriate to model all Client and Gaming revenue at a single high growth rate just because FY2025 segment revenue grew 50.8%. The correct forward setup is Client share/mix gains versus Gaming console-cycle pressure.

8. Forward Growth Assumptions

The forecast compounds from the FY2025/TTM base to 2030-12-27. The model uses two revenue lines because those are the only disclosed component revenue lines: Client and Gaming. Segment total is the sum of those lines. No separate Radeon or semi-custom dollar base is disclosed, so the Gaming forecast includes both.

Base Case Growth

Product	FY2023-FY2025 CAGR	Latest YoY	Mgmt FY+1 guide	Base CAGR TTM to TTM+5y	TTM+5y revenue	Justification
Client	51.2%	50.8% FY2025; 34% Q4 FY2025	Management believes it can grow the PC business; enterprise is a focus area	8.0%	$15.633B	Client decelerates from the exceptional FY2025 recovery to a still-healthy share-gain CAGR. This assumes AMD keeps gaining premium and enterprise share while PC units normalize after Windows 11 refresh demand.
Gaming	-20.7%	50.7% FY2025; 50% Q4 FY2025	Semi-custom SoC annual revenue to decline by a significant double-digit percentage; next-gen Xbox SoC progressing	-3.0%	$3.358B	The first horizon year absorbs console semi-custom decline, then later years recover with next-gen console contribution and handheld/integrated GPU relevance. Radeon discrete GPU pressure prevents a positive five-year CAGR in the base case.
Total Client and Gaming	15.7%	50.8% FY2025; 37% Q4 FY2025	Mixed: Client growth, semi-custom decline	5.5% implied	$18.991B	The combined segment grows because Client expands from 73% to 82% of revenue. Gaming no longer drives the segment.

The base year-by-year revenue bridge is:

Product	FY2025 / TTM	FY2026e	FY2027e	FY2028e	FY2029e	2030-12-27e
Client	$10.640B	$11.491B	$12.411B	$13.404B	$14.477B	$15.633B
Gaming	$3.910B	$3.402B	$3.130B	$3.224B	$3.321B	$3.358B
Segment total	$14.550B	$14.893B	$15.541B	$16.628B	$17.798B	$18.991B

This bridge deliberately shows deceleration rather than hiding it inside a single CAGR. Client growth steps down from 50.8% in FY2025 and 34% in Q4 FY2025 toward high-single-digit structural growth. The deceleration is defensible because 2025 included recovery from FY2023-FY2024 weakness, Windows 11 refresh demand, and strong high-end Ryzen mix. Gaming is negative in FY2026 because management gave a significant double-digit semi-custom decline guide, then stabilizes as console-cycle expectations and integrated gaming use cases improve.

Bull Case Growth

Product	FY2023-FY2025 CAGR	Latest YoY	Mgmt FY+1 guide	Bull CAGR TTM to TTM+5y	TTM+5y revenue	Justification
Client	51.2%	50.8% FY2025; 34% Q4 FY2025	PC business growth and enterprise focus	12.0%	$18.750B	Bull case assumes AMD compounds share gains into commercial PCs, AI PCs, and premium notebooks while Intel's recovery is uneven. Gartner's AI PC adoption path supports a richer market mix, and AMD captures more of that mix.
Gaming	-20.7%	50.7% FY2025; 50% Q4 FY2025	Near-term semi-custom decline but next-gen Xbox progress	4.0%	$4.757B	Bull case assumes the semi-custom decline is short-lived, next-gen console work converts before the end of the horizon, and Radeon/handheld APUs regain relevance at mainstream price points.
Total Client and Gaming	15.7%	50.8% FY2025; 37% Q4 FY2025	Mixed but improving	10.1% implied	$23.507B	Client remains the profit engine, while Gaming returns to growth rather than being a drag.

The bull bridge is:

Product	FY2025 / TTM	FY2026e	FY2027e	FY2028e	FY2029e	2030-12-27e
Client	$10.640B	$12.023B	$13.586B	$15.352B	$17.347B	$18.750B
Gaming	$3.910B	$3.519B	$3.694B	$4.063B	$4.469B	$4.757B
Segment total	$14.550B	$15.542B	$17.280B	$19.415B	$21.816B	$23.507B

This case does not assume Gaming instantly returns to the FY2023 base. Gaming's TTM+5y revenue of $4.757B is still below FY2023 Gaming revenue of $6.212B [data_package section 4, Gaming row]. The upside instead comes mostly from Client, where AMD's still-minority mobile share and enterprise opportunity provide room to compound above the PC unit market if execution remains strong.

Bear Case Growth

Product	FY2023-FY2025 CAGR	Latest YoY	Mgmt FY+1 guide	Bear CAGR TTM to TTM+5y	TTM+5y revenue	Justification
Client	51.2%	50.8% FY2025; 34% Q4 FY2025	PC business growth and enterprise focus	2.0%	$11.746B	Bear case assumes the FY2025 PC refresh pulled demand forward, Intel regains competitiveness, Arm notebooks pressure consumer mobile, and AI PC features fail to sustain ASP uplift. AMD still grows slightly because the FY2025 share base is higher.
Gaming	-20.7%	50.7% FY2025; 50% Q4 FY2025	Significant double-digit semi-custom decline	-12.0%	$2.063B	Bear case assumes console semi-custom revenue declines longer than expected and Radeon remains stuck in single-digit desktop AIB share with weak pricing. No major next-gen volume arrives before the horizon end.
Total Client and Gaming	15.7%	50.8% FY2025; 37% Q4 FY2025	Mixed with sharper Gaming drag	-1.0% implied	$13.809B	Client cannot fully offset a Gaming downturn and weaker PC replacement cycle.

The bear bridge is:

Product	FY2025 / TTM	FY2026e	FY2027e	FY2028e	FY2029e	2030-12-27e
Client	$10.640B	$10.853B	$11.070B	$11.292B	$11.518B	$11.746B
Gaming	$3.910B	$3.128B	$2.628B	$2.365B	$2.176B	$2.063B
Segment total	$14.550B	$13.981B	$13.698B	$13.657B	$13.694B	$13.809B

The bear case is not a collapse scenario; it is a normalization and mix-deterioration scenario. It matters because the FY2025 base is high after a rebound year. If the PC market weakens after the Windows 11 refresh and semi-custom enters a console trough, the segment can under-earn even with AMD retaining meaningful client share gains.

9. Forward Operating Margin Assumptions

The hard margin anchor is FY2025 Client and Gaming operating income of $2.855B on $14.550B revenue, or 19.6% [data_package section 3, Client and Gaming table]. Component margins are not disclosed [DATA GAP: AMD does not disclose Client operating income or Gaming operating income separately; data_package section 10]. For modeling transparency, the FY2025 allocation used here is Client at 22.0% operating margin and Gaming at 13.1% operating margin. That allocation ties the reported combined OI within rounding: $10.640B x 22.0% = $2.341B; $3.910B x 13.1% = $0.514B; total = $2.855B.

Base Case Margin

Product	FY2025 allocated margin	TTM+5y margin	TTM+5y revenue	TTM+5y OI	Mechanism
Client	22.0%	24.0%	$15.633B	$3.752B	Enterprise mix, premium Ryzen/AI PC mix, and operating leverage offset competitive pricing. Margin improves 200 bps from allocated FY2025.
Gaming	13.1%	7.0%	$3.358B	$0.235B	Semi-custom trough and weak Radeon pricing reduce margin. Later recovery prevents losses, but the line remains below FY2025 allocated margin.
Total Client and Gaming	19.6% reported	21.0%	$18.991B	$3.987B	Client mix rises to 82% of revenue, lifting blended segment margin despite Gaming drag.

The base margin expansion is 140 bps from the reported FY2025 segment margin. The mechanism is not generic scale leverage; it is mix. Client becomes a larger share of revenue, and Client itself shifts more toward high-end desktop, commercial PC, and AI PC configurations. Offsetting that are semi-custom decline and Radeon pricing pressure. The base case does not assume a return to FY2021 margin of 24.2% because the segment mix is different, console cycle is a headwind, and competition is intense.

Bull Case Margin

Product	FY2025 allocated margin	TTM+5y margin	TTM+5y revenue	TTM+5y OI	Mechanism
Client	22.0%	26.0%	$18.750B	$4.875B	Premium desktop leadership, enterprise penetration, and AI PC ASP protection create operating leverage.
Gaming	13.1%	16.0%	$4.757B	$0.761B	New console ramp and better Radeon/handheld mix restore fixed-cost absorption.
Total Client and Gaming	19.6% reported	24.0%	$23.507B	$5.636B	Both component margins improve, and revenue growth spreads R&D across a larger base.

The bull margin is above FY2025 by 440 bps, which requires explicit upside mechanisms. The case assumes Client earns high-20s-like economics in premium and enterprise-heavy mixes and that Gaming returns to a profitable console ramp rather than remaining in trough. It also assumes AMD does not have to buy share through aggressive discounting against Intel in client or Nvidia in Radeon.

Bear Case Margin

Product	FY2025 allocated margin	TTM+5y margin	TTM+5y revenue	TTM+5y OI	Mechanism
Client	22.0%	17.0%	$11.746B	$1.997B	PC demand weakens, Intel/Arm pricing pressure rises, and AI PC ASP uplift fades.
Gaming	13.1%	-3.0%	$2.063B	-$0.062B	Semi-custom trough extends and Radeon remains under-scale. Fixed R&D and channel support exceed gross profit on the lower revenue base.
Total Client and Gaming	19.6% reported	14.0%	$13.809B	$1.935B	Revenue mix worsens and segment under-absorbs fixed engineering costs.

The bear margin is 560 bps below FY2025. That is large but consistent with segment history: Client and Gaming margin was only 8.5% in FY2023 during the prior trough [data_package section 3]. A 14.0% margin is therefore not a worst-case depression margin; it is a mid-cycle bear outcome where Client remains profitable but Gaming and pricing pressure drag the blend.

10. Segment Rollup

The bottom-up revenue rollup ties directly to the component assumptions in sections 8 and 9.

Case	Client TTM+5y revenue	Gaming TTM+5y revenue	Segment TTM+5y revenue	Segment margin	Segment OI
Bear	$11.746B	$2.063B	$13.809B	14.0%	$1.935B
Base	$15.633B	$3.358B	$18.991B	21.0%	$3.987B
Bull	$18.750B	$4.757B	$23.507B	24.0%	$5.636B

Arithmetic tie-out:

- Bear: $11.746B + $2.063B = $13.809B. OI equals $11.746B x 17.0% plus $2.063B x -3.0% = $1.935B, a 14.0% blended margin.
- Base: $15.633B + $3.358B = $18.991B. OI equals $15.633B x 24.0% plus $3.358B x 7.0% = $3.987B, a 21.0% blended margin.
- Bull: $18.750B + $4.757B = $23.507B. OI equals $18.750B x 26.0% plus $4.757B x 16.0% = $5.636B, a 24.0% blended margin.

The base case implies a five-year segment revenue CAGR of roughly 5.5% from the $14.550B TTM base. That is much lower than FY2025's 50.8% growth because the forecast treats FY2025 as a rebound year. It is also higher than a pure PC unit growth model because AMD has share and mix levers in Client. The bull case implies roughly 10.1% CAGR and requires sustained Client share gains plus Gaming recovery. The bear case implies about -1.0% CAGR and captures the risk that Gaming declines and PC refresh demand normalizes.

11. Key Risks And Wildcards

- Bear risk - PC refresh pull-forward. Gartner's 2025 PC market release attributes much of the year's growth to Windows 11 upgrades, inventory ahead of memory price hikes, and future-proofing rather than immediate AI PC productivity. If that demand was pulled forward, Client revenue can flatten even if AMD keeps share.

- Bear risk - Intel recovery. Intel has deep OEM, commercial, and platform relationships. If Intel regains performance-per-watt competitiveness and uses pricing or platform incentives aggressively, AMD's Client share gains can slow or reverse.

- Bear risk - Arm notebooks. Gartner expects Arm to gain more consumer AI laptop share as compatibility improves. That risk is highest in consumer mobile and lower in enterprise x86 fleets, but it can still pressure AMD's notebook TAM and ASP.

- Bear/base risk - Semi-custom trough. AMD management explicitly guided semi-custom SoC annual revenue to decline by a significant double-digit percentage. The timing of next-generation console ramps is a key swing factor for Gaming revenue.

- Bear risk - Radeon under-scale. Nvidia's discrete GPU dominance limits AMD's ability to earn attractive margins in PC gaming graphics. If Radeon share remains low and channel pricing is weak, Gaming margin can fall below the model's base case.

- Bull wildcard - Enterprise PC penetration. Management's enterprise focus could be more valuable than consumer share gains because commercial fleets create validation, lifecycle, and ASP benefits. If AMD converts more enterprise designs, Client revenue and margin can beat the base case.

- Bull wildcard - AI PC mix. Gartner's AI PC forecast implies a rapid mix shift through 2026. If AMD's Ryzen AI platforms become preferred in business notebooks or premium consumer systems, the Client line can compound above PC unit growth.

- Bull wildcard - Console and handheld ecosystem. A successful next-gen Xbox semi-custom design and continued growth in handheld gaming APUs could turn Gaming from a drag into a modest growth contributor by the end of the horizon.

Summary Table

Case	Segment TTM+5y Revenue	Segment TTM+5y Op Margin	Segment TTM+5y Op Income	Probability
Bear	$13.809B	14.0%	$1.935B	25%
Base	$18.991B	21.0%	$3.987B	50%
Bull	$23.507B	24.0%	$5.636B	25%

Sources Used

- AMD historical segment and component data, management commentary, taxonomy, and data gaps: `data_package.md` in this ER2 workdir, especially sections 1, 3, 4, 5, 6, and 10.
- Gartner worldwide PC shipment release, January 2026: https://www.gartner.com/en/newsroom/press-releases/2026-1-20-gartner-says-worldwide-pc-shipments-increased-9-point-3-percent-in-fourth-quarter-of-2025-and-9-point-1-percent-for-the-full-year
- Gartner AI PC release, August 2025: https://www.gartner.com/en/newsroom/press-releases/2025-08-28-gartner-says-artificial-intelligence-pcs-will-represent-31-percent-of-worldwide-pc-market-by-the-end-of-2025
- Jon Peddie Research Q4 FY2025 PC GPU release: https://www.jonpeddie.com/news/q425-pc-gpu-shipments-decreased-by-1-3-from-last-quarter-which-is-a-prelude-to-the-creep-forward/
- Jon Peddie Research Q4 FY2025 AIB release: https://www.jonpeddie.com/news/q425-pc-graphics-aib-shipments-decreased-4-4-from-last-quarter-to-11-million-units-with-a-cagr-to-2029-of-5-9/
- PCWorld summary of Mercury Research Q4 FY2025 AMD share figures: https://www.pcworld.com/article/3058885/amd-surges-10-percent-in-desktop-pc-share.html
- TechSpot summary of Mercury Research Q4 FY2025 AMD share figures: https://www.techspot.com/news/111297-ryzen-revenue-share-jumps-146-amd-takes-426.html

Appendix B3 - Embedded

Embedded - Segment Deep Dive

Stage 0 anchor: today is 2026-04-29. AMD's fiscal year ends on the last Saturday in December. The latest reported quarter is Q4 FY2025, ended 2025-12-27, and the TTM base for this appendix is FY2025. The forward endpoint used throughout is 2030-12-27, exactly five calendar years after the TTM end date.

0. Segment Boundary And Taxonomy Note

Embedded is a current AMD reportable segment and includes embedded CPUs, GPUs, APUs, FPGAs, system-on-modules, and adaptive SoC products [data_package Section 1, current segment table; data_package Section 4, Embedded product row]. The Q1 FY2025 taxonomy change combined Client and Gaming into a single reportable segment, but the data package does not identify a comparable Embedded restatement inside the last three years. Embedded dollar comparisons should still treat FY2021 cautiously because FY2021 precedes the Xilinx acquisition and shows only $0.246B of Embedded revenue; FY2022-FY2025 are more meaningful for the post-Xilinx business, and FY2023-FY2025 are the cleanest three-year current-period history [data_package Section 3, Embedded table; data_package Section 10, taxonomy notes].

1. What This Segment Actually Is

AMD Embedded is the inherited Xilinx adaptive-computing franchise plus AMD embedded processors and related system-level products. The customer is typically an OEM, tier-one supplier, network equipment vendor, aerospace and defense prime, industrial automation vendor, medical device maker, broadcast/pro-AV vendor, or cloud/communications infrastructure supplier that needs semiconductors inside long-lived systems. The end product is not a consumer PC or a data center accelerator card sold into a fast replacement cycle. It is more often a board, controller, sensor-fusion system, industrial PC, robot, network appliance, ADAS subsystem, image-processing pipeline, radar or RF system, medical imaging unit, or mission-critical embedded platform.

The segment's product set is broad. The data package says Embedded includes embedded CPUs, GPUs, APUs, FPGAs, system-on-modules, and adaptive SoC products [data_package Section 1, reporting structure; data_package Section 4, Embedded product row]. AMD's embedded solutions page describes the portfolio as adaptive SoCs, FPGAs, and CPUs supporting industrial, automotive, healthcare, space, and other edge-to-cloud applications. It also lists Adaptive SoCs and FPGAs, Embedded Processors, Embedded+, SOMs, and accelerator cards as portfolio categories [AMD Embedded Computing Solutions, lines 477-482 and 596-603]. This matters because Embedded is not a single chip family. It is a portfolio of products that sit at different price points, power envelopes, lifecycles, and software-tool dependencies.

The adaptive SoC and FPGA portion is the heart of the moat. A field-programmable gate array gives the customer reconfigurable hardware logic. An adaptive SoC adds hard processor cores, programmable logic, hard IP, and software-programmable engines in a more integrated platform. AMD's Versal page describes general-purpose Versal devices as combining programmable logic, embedded Arm application and real-time CPU cores, programmable network-on-chip, high-speed serial transceivers, programmable I/O, and hard IP [AMD Versal Adaptive SoCs, lines 515-519]. Specialized Versal products add AI Engines for machine learning or DSP, RF-ADCs/RF-DACs, expanded DSP hard IP, or high-bandwidth memory [AMD Versal Adaptive SoCs, lines 551-553]. These features explain why a customer may use an AMD adaptive SoC rather than a generic microcontroller, CPU, GPU, or ASIC: the hardware can be tuned to the sensor, communications, safety, latency, and power constraints of the application.

The pricing model is semiconductor product revenue, not usage-based software. AMD sells chips, modules, boards, tools, IP, and support into long design cycles. Revenue recognition depends on product shipments and channel consumption, while commercial momentum is often visible earlier through design wins. The data package explicitly says Embedded design wins are not revenue and should be treated as a commercial indicator rather than backlog. The Q4 FY2025 transcript disclosed $17B of design wins in 2025 and more than $50B since acquiring Xilinx, but the data package cautions that conversion timing and margin profile are not disclosed [data_package Section 4, growth-rate-only table; data_package Section 6, Embedded commentary; data_package Section 10, downstream modeling guardrails].

Go-to-market therefore looks different from AMD's Client or Gaming segments. Embedded customers care about product longevity, qualification, functional safety, board support, tool-chain support, power, deterministic latency, security, and the confidence that a design can remain available for many years. A customer switching a qualified embedded design from one FPGA or adaptive SoC vendor to another can require re-architecture, tool migration, firmware and HDL work, validation,

safety documentation, and customer requalification. Those costs can be larger than the chip price delta, which gives successful sockets sticky revenue once the product is in production. The same dynamic creates a lag: design wins may not become material revenue for years, and revenue can decline sharply when inventory digestion hits distributors or when a large communications, industrial, or automotive program transitions.

2. Current Economics: FY2025 / TTM

The FY2025/TTM base is $3.454B of Embedded revenue and $1.243B of segment operating income, equal to a 36.0% operating margin [data_package Section 3, Embedded table; quant_foundation P4, TTM per-segment construction]. Embedded was 10.0% of AMD FY2025 revenue and 16.1% of disclosed reportable-segment operating income before All Other corporate expense, despite being the smallest of the three reportable segments by revenue; its $1.243B of segment operating income also equaled 33.6% of AMD's $3.694B consolidated GAAP operating income, though segment and consolidated operating income are not directly additive because of All Other corporate expense [data_package Section 1, current segment table; data_package Section 2, consolidated historicals; data_package Section 3, segment historicals]. That is the central economic point: Embedded is small, but structurally high-margin relative to AMD's consolidated GAAP operating margin and relative to the more cyclical Client and Gaming business.

Revenue declined 2.9% in FY2025 after falling 33.2% in FY2024. The three-year post-Xilinx revenue path is volatile: FY2022 was $4.552B, FY2023 rose to $5.321B, FY2024 fell to $3.557B, and FY2025 slipped to $3.454B [data_package Section 3, Embedded table]. The decline was not matched by a complete margin collapse, but the margin compression was material. Operating margin was 49.5% in FY2022, 49.4% in FY2023, 39.9% in FY2024, and 36.0% in FY2025 [data_package Section 3, Embedded table]. The segment still produced more than one-third operating margin in a downcycle, which is strong evidence that the product portfolio has pricing power and a high gross-margin content base; however, the 1,340 bps decline from FY2023 to FY2025 also shows that the segment has meaningful fixed cost and mix leverage.

The quarterly trend improved late in FY2025. Embedded revenue was $0.823B in Q1, $0.824B in Q2, $0.857B in Q3, and $0.950B in Q4. Year-over-year growth was -3%, -4%, -8%, and +3%, respectively [data_package Section 5, Embedded quarterly trend]. Operating margin was 39.9% in Q1, 33.4% in Q2, 33.0% in Q3, and 37.6% in Q4 [data_package Section 5, Embedded quarterly trend]. The Q4 positive growth and margin rebound support management's commentary that demand was recovering, but the year as a whole still ended below FY2024.

There is no disclosed revenue split between Embedded FPGAs, adaptive SoCs, embedded CPUs, GPUs/APUs, SOMs, and accelerator cards. The data package says "No finer product dollars found in this run" for Embedded [data_package Section 4, Embedded product row]. This appendix therefore models the segment as one disclosed product line: "Embedded portfolio." That is not ideal, because the underlying businesses likely have different growth and margin profiles. Adaptive SoCs for automotive AI edge and aerospace may have very different cycle timing than industrial embedded CPUs or older FPGA families. But inventing a mix would create false precision. The correct treatment is to use the disclosed segment dollars and operating income, then discuss the qualitative subdrivers separately.

[DATA GAP: AMD does not disclose Embedded revenue by FPGAs, adaptive SoCs, embedded CPUs/APUs/GPUs, SOMs, accelerator cards, end market, customer type, or geography in the data package.]

Capital intensity is low at the segment manufacturing level because AMD is fabless. The data package provides company-level capex of $0.974B in FY2025, equal to 2.8% of revenue, and R&D of $8.091B, equal to 23.4% of revenue [data_package Section 7, capital allocation history]. Segment-level capex is not disclosed. Embedded reinvestment intensity is better thought of as engineering, software-tool, IP, validation, field application engineer, and product-longevity expense rather than owned fabrication capex. The segment depends on advanced foundry, packaging, and supply-chain partners, but its operating income is more sensitive to product mix, volume absorption, and R&D leverage than to AMD-owned plant spending.

3. Unit Economics

Embedded revenue turns into profit through a combination of high-value silicon content, long design lives, software and IP lock-in, and a relatively low need for direct selling expense once a product is qualified. The data package does not disclose gross margin by segment, but the operating-margin history is enough to infer that Embedded is one of AMD's highest-return pools. A segment can only sustain 49% operating margins in FY2022-FY2023 if gross margins are high and operating expense can be spread over a specialized product base [data_package Section 3, Embedded table]. Even after the FY2024-FY2025 revenue decline, Embedded produced a 36.0% FY2025 operating margin [data_package Section 3, Embedded table].

The main revenue-unit driver is the design win. A customer evaluates an AMD device, commits engineering resources, integrates the product into its hardware and software architecture, validates it, and then moves through production. A win can convert into revenue over a long time, sometimes over the life of an industrial, aerospace, medical, or automotive platform. That makes the conversion curve slow but sticky. The Q4 FY2025 call disclosed $17B of design wins in 2025 and more than $50B since the Xilinx acquisition [data_package Section 6, Embedded commentary]. The data package is explicit that these are not backlog or revenue. They are evidence that the funnel is alive after the inventory correction, not a basis for a dollar-by-dollar forecast.

The operating-cost structure has several major drivers. First, platform engineering and EDA/tool development are required to keep Vivado, Vitis, board support packages, IP libraries, compilers, and safety/security collateral relevant. AMD's Versal page highlights Vivado and Vitis as the core hardware and software development environments that reduce development time and help developers achieve performance per watt [AMD Versal Adaptive SoCs, lines 593-597]. Those tools are not optional in the economic model; they are part of why customers can use AMD adaptive compute, and they are part of the fixed cost that must be leveraged over revenue.

Second, product qualification and lifecycle support are meaningful. Embedded customers often require extended availability, safety certification, security documentation, industrial temperature support, radiation or ruggedization characteristics, or automotive qualification. AMD's Versal AI Edge page refers to safety and security requirements including ISO 26262, IEC 61508, and IEC 62443, and describes automotive-qualified Versal AI Edge adaptive SoCs [AMD Versal AI Edge Series, lines 484, 510, and 532-534]. Those capabilities are expensive to build and maintain, but they support pricing power and switching costs.

Third, mix matters. If near-term recovery comes mostly from older industrial and communications inventory replenishment, margins may recover but probably not to the FY2022-FY2023 high. If recovery is driven by higher-value Versal AI Edge, RF, automotive ADAS, aerospace, and high-end adaptive SoC sockets, margins can recover more. If recovery comes from lower-priced embedded processors or competitive mid-range FPGAs, operating margin may remain in the mid-30s even with revenue growth.

Fourth, cyclicality matters. Embedded is less consumer-cycle exposed than Gaming, but it is not non-cyclical. The FY2024 revenue decline of 33.2% was too large to dismiss as a rounding issue; it reflects inventory digestion and end-market softness in long-cycle markets [data_package Section 3, Embedded table; data_package Section 6, management commentary]. When distributors and OEMs hold excess inventory, shipments can fall below end demand for several quarters. Conversely, once inventory normalizes, shipments can grow faster than end demand for a time. This makes one-year growth an unreliable signal for long-term normalized growth.

The margin breakpoints are therefore clear. Margins move up when design-win conversion lifts revenue, higher-value adaptive SoCs mix up, channel inventory normalizes, and fixed engineering expense is absorbed over a larger revenue base. Margins move down when volume falls, older products mix higher, price competition rises in mid-range FPGAs, or AMD must increase field and software investment to defend sockets against Altera, Lattice, Microchip, Achronix, ASICs, GPUs, and customer-specific silicon.

4. Competitive Dynamics

Embedded competes in a concentrated but segmented programmable-logic and embedded-compute market. The direct FPGA/adaptive-compute competitors are Altera, now majority-owned by Silver Lake with Intel retaining 49%; Lattice Semiconductor in low-power and small-form-factor FPGAs; Microchip in secure, low-power, industrial, aerospace, and defense FPGA families; and smaller players including Achronix, Efinix, QuickLogic, Gowin, Anlogic, Flex Logix, and ASIC/ASSP alternatives. In embedded CPUs and industrial PCs, AMD also faces Intel's edge and embedded CPU portfolio, Arm-based SoCs, NXP, Qualcomm, Nvidia, TI, Renesas, and specialized edge-AI accelerators.

AMD's strongest competitive position is in high-performance adaptive SoCs and FPGAs where legacy Xilinx tools, IP, customer history, and product breadth matter. AMD's embedded portfolio spans Versal, Zynq, FPGAs, Ryzen Embedded, EPYC Embedded, Embedded+, SOMs, and accelerator cards [AMD Embedded Computing Solutions, lines 596-667]. That breadth lets AMD offer a customer several routes: pure programmable logic, Arm-plus-FPGA adaptive SoC, x86 embedded CPU, x86 plus Versal board-level architecture, or an ecosystem module. This is valuable when customers are trying to consolidate vendors or use a common tool chain across product generations.

Altera is the most direct high-end and mid-range competitor. Intel launched Altera as a standalone FPGA operation in 2024, and Altera announced in September 2025 that Silver Lake had completed its acquisition of a 51% stake while Intel retained

49% [Altera press release, lines 498-502]. Altera's technology roadmap is relevant because it targets many of the same use cases. Its 2024 developer-conference materials described Agilex 3 FPGAs for embedded and intelligent-edge applications, with higher integration, security, compact packages, and densities of 25K to 135K logic elements [Intel/Altera release, lines 32-39]. Altera also emphasized Agilex 5 SoC FPGAs, AI tool flows, and ecosystem support for embedded applications [Intel/Altera Embedded World release, lines 178-204 and 231-239]. The private-equity majority ownership could sharpen execution and commercial focus, but separation from Intel may also create transition risk. For AMD, the threat is not that Altera destroys pricing overnight; it is that Altera becomes more nimble in mid-range and high-end sockets where customers value a second source or a focused FPGA supplier.

Lattice competes more strongly in low-power FPGAs, companion chips, security, control logic, board management, and smaller edge devices. Lattice's FY2025 release shows that its communications and computing revenue reached 64% of Q4 revenue and 56% of FY2025 revenue, while industrial and automotive represented 30% of Q4 revenue and 37% of FY2025 revenue [Lattice FY2025 release, lines 1686-1730]. Management also said normalized industrial and automotive channel inventory should let revenue better track consumption in 2026, while communications and computing demand remained strong [Lattice FY2025 release, line 401]. Lattice is not a substitute for AMD's highest-end Versal devices in many applications, but it can pressure lower-power sockets and companion-chip opportunities, especially where power, cost, board area, and simplicity matter more than raw compute.

Microchip is a credible competitor in low-power, secure, industrial, aerospace and defense, and long-lifecycle FPGAs. It does not need to beat AMD across the entire high-performance adaptive SoC stack to take sockets; it only needs to be good enough in applications where security, reliability, radiation tolerance, power efficiency, and supply assurance dominate. Smaller vendors and regional Chinese FPGA suppliers can also matter in cost-sensitive or geopolitically constrained markets, though they generally lack AMD's high-end breadth and software ecosystem.

The broader substitute threat is ASICs and dedicated edge-AI accelerators. If an application reaches high enough volume and the algorithm stabilizes, a customer may build an ASIC or choose a dedicated NPU/AI SoC rather than keep an FPGA. This is a recurring risk in programmable logic: FPGAs are attractive for flexibility, time-to-market, and changing standards; ASICs are attractive for unit cost and power at scale. AMD's defense is to keep adaptive SoCs relevant for workloads where algorithms, sensors, standards, or safety requirements change faster than ASIC cycles, and where one programmable device can replace several fixed-function chips.

External market forecasts support a mid-to-high-single-digit or high-single-digit structural growth backdrop, but not a hypergrowth baseline for the entire segment. Lucintel forecasts the FPGA market reaching $17.4B by 2030 at a 9.5% CAGR from 2024 to 2030, with drivers including ADAS, IoT, automotive, and AI [Lucintel FPGA market report, lines 76-80 and 199-204]. Microchip USA cites a similar 2026-to-2031 FPGA market forecast of $11.02B to $17.23B, or about 9.35% CAGR [Microchip USA FPGA market availability, lines 6-10]. These are market-research figures rather than AMD-specific guidance, and they include data center and computing FPGAs as well as embedded markets. They are useful as a sanity check: AMD Embedded can grow above market if design wins convert and Versal AI Edge gains share, but a base case well above low-teens growth would need much stronger evidence than the data package provides.

5. Moat At The Segment Level

The Embedded moat is real but narrower than AMD's company-level AI narrative. It rests on five assets: installed base, tool chain, product breadth, certification/lifecycle support, and customer design-in switching costs.

Installed base matters because Xilinx was an embedded and programmable-logic standard for decades. Engineers trained on Xilinx devices, IP blocks, Vivado, Vitis, and Zynq/Versal architectures are less likely to switch unless the alternative offers a decisive cost, power, or performance advantage. The tool chain matters because FPGA/adaptive SoC selection is not just a chip-purchase decision. Customers build HDL, software, IP integration, timing closure flows, simulation, board support, safety collateral, and debug workflows around the vendor's tools. AMD's product pages emphasize Vivado and Vitis as development environments for adaptive SoCs and FPGAs, enabling different developer profiles from AI scientists to hardware designers [AMD Versal Adaptive SoCs, lines 593-597]. The tools do not make AMD invulnerable, but they raise switching costs.

Product breadth matters because AMD can span high-end Versal, Zynq adaptive SoCs, standalone FPGAs, Ryzen Embedded, EPYC Embedded, and Embedded+ architectures [AMD Embedded Computing Solutions, lines 604-663]. A customer with a robotics, medical, automotive, industrial, or communications platform can choose from multiple AMD architectures while staying inside a similar vendor support model. Embedded+ is especially interesting because it packages Ryzen Embedded processors with Versal AI Edge adaptive SoCs for sensor fusion, AI inferencing, industrial networking, control, and visualization [AMD Versal AI Edge Series, lines 553-555]. This is a practical response to edge AI systems that

require CPU software compatibility, programmable I/O, deterministic preprocessing, and AI inference in one board-level design.

Certification and lifecycle support matter because embedded customers often have long production lives. The value proposition is not simply peak TOPS or logic density; it is the confidence that the part, tools, and support will exist through the product life. AMD's Versal AI Edge safety/security positioning for ISO 26262, IEC 61508, and IEC 62443 is directly relevant for automotive and industrial sockets [AMD Versal AI Edge Series, lines 484 and 510]. Once a customer qualifies a safety-critical design, it has a high threshold for change.

The moat is not absolute. Altera can match AMD in many technical domains and is now a focused FPGA company with new ownership. Lattice can win where low power and simplicity beat high-end integration. Microchip can win where security, reliability, and long lifecycle matter more than performance. ASICs can displace FPGAs in high-volume stable workloads. Edge-AI NPUs can capture applications where the workload is standardized and software ecosystems mature. The moat also does not prevent cyclicality; Embedded revenue fell from $5.321B in FY2023 to $3.454B in FY2025 [data_package Section 3, Embedded table]. A true moat should protect normalized returns over a cycle, not prevent every downcycle.

My conclusion is that Embedded has a good segment-level moat in high-performance adaptive SoCs and long-lifecycle programmable logic, but the moat supports margins more than it guarantees growth. The right five-year model should assume recovery and design-win conversion, not a straight return to FY2023 revenue and 49% operating margin. A durable base case needs to recognize that FY2025 margin was 36.0%, not 49%, and that the market is competitive even if switching costs are high.

6. Historical Growth Rates Reproduced From Data Package

AMD discloses Embedded only at the segment level in the data package. There are no disclosed product-dollar rows for Embedded FPGAs, adaptive SoCs, embedded CPUs/APUs/GPUs, SOMs, or accelerator cards [data_package Section 4, Embedded product row; data_package Section 10, open data gaps]. The table below therefore uses the finest disclosed line: Embedded portfolio. FY2022 is included as the 3-year CAGR base because TTM/FY2025 is the anchor and FY2022 to FY2025 is the last completed three-year period; FY2023-FY2025 are also shown because they are the cleanest current-period dollars.

Product / sub-segment	FY2022 rev	FY2023 rev	FY2024 rev	FY2025 / TTM rev	3-yr CAGR FY2022-FY2025	Most recent YoY	Data package cite
Embedded portfolio: embedded CPUs, GPUs, APUs, FPGAs, SOMs, adaptive SoCs	$4.552B	$5.321B	$3.557B	$3.454B	-8.8%	-2.9%	FY2022-FY2025 revenue and YoY from data_package Section 3, Embedded table; product boundary from data_package Section 4, Embedded row

Operating margin history for the same disclosed line:

Product / sub-segment	FY2022 OI	FY2023 OI	FY2024 OI	FY2025 / TTM OI	FY2025 / TTM margin	Data package cite
Embedded portfolio	$2.252B	$2.628B	$1.421B	$1.243B	36.0%	OI and margin from data_package Section 3, Embedded table; TTM tie to quant_foundation P4

Arithmetic tie-out:

Check	Calculation	Result
FY2023 YoY	$4.552B x 1.169	$5.321B, matches FY2023 within rounding
FY2024 YoY	$5.321B x 0.668	$3.555B, matches FY2024 $3.557B within 0.1%
FY2025 YoY	$3.557B x 0.971	$3.454B, matches FY2025 within rounding
FY2022-FY2025 CAGR	($3.454B / $4.552B)^(1/3) - 1	-8.8%

[DATA GAP: No Embedded product-level revenue or margin split is disclosed in the data package, so this section cannot present separate historical rows for FPGAs, adaptive SoCs, embedded CPUs/APUs/GPUs, SOMs, or accelerator cards.]

7. Management Guidance Quoted From Data Package Section 6

The relevant Embedded commentary in the data package points to a gradual recovery, stronger second-half demand, and a large design-win funnel. It does not provide a segment-dollar FY2026 guide.

Data package row	Verbatim quote	Interpretation
Q1 FY2025 call / event, 2025-05-06, Lisa Su	"Embedded demand continues to recover gradually"	Management framed early FY2025 as a recovery period rather than a growth period. This fits the Q1 revenue decline of 3% YoY and supports using Q1-Q3 softness as inventory-cycle context, not as a permanent run-rate collapse.
Q1 FY2025 call / event, 2025-05-06, Lisa Su	"drive a return to growth in the second half of 2025"	The expected 2H recovery was directionally realized in Q4, where Embedded returned to +3% YoY growth [data_package Section 5, Embedded quarterly trend].
Q3 FY2025 call / event, 2025-11-04, Lisa Su	"demand environment strengthened across multiple markets"	The improvement was broad enough that management did not frame it as one customer or one end market. That supports a base case with multi-market recovery, but not an aggressive product-specific ramp because no end-market dollars are disclosed.
Q4 FY2025 call / event, 2026-02-03, Lisa Su	"demand strengthened across several end markets"	Q4 confirmed recovery language after three negative YoY quarters. It is a near-term positive, but still qualitative.
Q4 FY2025 call / event, 2026-02-03, Lisa Su	"We closed $17 billion in design wins in 2025"	This is the strongest commercial funnel indicator. The model treats design wins as long-cycle conversion potential, not as backlog, because the data package explicitly warns against treating them as revenue.
Q4 FY2025 call / event, 2026-02-03, Lisa Su	"more than $50 billion of embedded designs since acquiring Xilinx"	The cumulative pipeline is much larger than annual revenue, but conversion timing, cancellation risk, content per program, and margin mix are not disclosed.
Q4 FY2025 call / event, 2026-02-03, Lisa Su	"return to growth in our Embedded segment"	This supports a positive FY2026 growth assumption, but there is no disclosed segment-dollar guide.

[DATA GAP: No explicit FY2026 Embedded revenue dollars, Embedded CAGR target, product-level growth target, design-win conversion schedule, or Embedded product margin guide is disclosed in the data package.]

8. Forward Growth Assumptions: Explicit, Per Disclosed Line

Because AMD discloses only one Embedded dollar line, the forward model uses one product/sub-segment line: Embedded portfolio. The qualitative subdrivers are adaptive SoCs/FPGAs, embedded processors, Embedded+, SOMs, and accelerator cards, but they are not separately dollar-modeled. The base, bull, and bear cases compound from FY2025/TTM revenue of $3.454B to the TTM+5y endpoint of 2030-12-27.

Base case: Embedded returns to healthy but not explosive growth as industrial, communications, automotive, aerospace/defense, and healthcare demand normalizes and design wins convert. The base case does not assume direct conversion of the $17B 2025 design-win figure into revenue; it uses that disclosure as evidence that the design funnel can support above-GDP growth after a downcycle.

Product	3-yr historical CAGR	Latest YoY	Mgmt FY+1 guide	Base CAGR TTM-to-TTM+5y	Base TTM+5y revenue	Justification
Embedded portfolio	-8.8% FY2022-FY2025	-2.9% FY2025; Q4 +3% YoY	Qualitative "return to growth"; no dollar guide [data_package Section 6, Embedded commentary]	8.0%	$5.075B	The base case assumes a cyclical recovery from two down years plus moderate design-win conversion. It stays below the most aggressive edge-AI narratives and roughly in line with external FPGA market forecasts around 9%-10% CAGR, because AMD has no product-dollar guide and competition remains strong.

Base year-by-year bridge:

Period ending	Revenue	YoY / growth assumption	Rationale
FY2025 / TTM 2025-12-27	$3.454B	actual	Data package control total.
TTM+1	$3.730B	8.0%	Recovery year after Q4 FY2025 turned positive, supported by management's return-to-growth language.
TTM+2	$4.029B	8.0%	Design wins begin converting, offset by normal competitive price pressure.
TTM+3	$4.351B	8.0%	Versal AI Edge, industrial automation, aerospace/defense, communications, and embedded processors contribute.
TTM+4	$4.699B	8.0%	Growth normalizes near the external FPGA market growth range.
TTM+5 / 2030-12-27	$5.075B	8.0%	Segment remains below FY2023 margin peak but returns near FY2023 revenue scale.

Bull case: Embedded converts a larger share of design wins, wins more AI-edge and automotive sockets, and benefits from customer preference for programmable hardware as edge models, sensors, and safety standards change.

Product	3-yr historical CAGR	Latest YoY	Mgmt FY+1 guide	Bull CAGR TTM-to-TTM+5y	Bull TTM+5y revenue	Justification
Embedded portfolio	-8.8% FY2022-FY2025	-2.9% FY2025; Q4 +3% YoY	Qualitative recovery plus $17B 2025 design wins and more than $50B cumulative since Xilinx [data_package Section 6]	12.0%	$6.087B	The bull case assumes that design wins convert faster and that Versal AI Edge / Embedded+ take share in ADAS, robotics, industrial vision, aerospace, and medical imaging. It is above broad FPGA market forecasts but not outlandish because the FY2025 base is depressed and FY2023 revenue was already $5.321B.

Bull year-by-year bridge:

Period ending	Revenue	YoY / growth assumption	Rationale
FY2025 / TTM 2025-12-27	$3.454B	actual	Data package control total.
TTM+1	$3.869B	12.0%	Sharper rebound as inventory digestion ends and design-win conversion begins.
TTM+2	$4.333B	12.0%	AI-enabled embedded applications and automotive/industrial programs ramp.
TTM+3	$4.853B	12.0%	Higher-value Versal mix grows and Embedded+ improves system-level attach.
TTM+4	$5.435B	12.0%	AMD regains and exceeds FY2023 revenue level.
TTM+5 / 2030-12-27	$6.087B	12.0%	Segment reaches a new high but still below any mechanical design-win conversion math.

Bear case: The recovery remains shallow, Altera and Lattice pressure sockets, industrial and communications demand stays uneven, and design wins convert slowly or at lower content.

Product	3-yr historical CAGR	Latest YoY	Mgmt FY+1 guide	Bear CAGR TTM-to-TTM+5y	Bear TTM+5y revenue	Justification
Embedded portfolio	-8.8% FY2022-FY2025	-2.9% FY2025; Q4 +3% YoY	Qualitative "return to growth"; no dollar guide	2.0%	$3.813B	The bear case allows a low-growth recovery from a depressed base but assumes design wins take longer to convert, mid-range competition intensifies, and ASIC/edge-AI substitutes limit pricing and content gains. It does not assume another FY2024-style collapse because Q4 FY2025 returned to growth and the design-win funnel is real.

Bear year-by-year bridge:

Period ending	Revenue	YoY / growth assumption	Rationale
FY2025 / TTM 2025-12-27	$3.454B	actual	Data package control total.
TTM+1	$3.523B	2.0%	Modest recovery after Q4 positive growth, but no broad acceleration.
TTM+2	$3.594B	2.0%	Customers remain cautious; conversion is slow.
TTM+3	$3.666B	2.0%	Competitive pressure offsets edge-AI content.
TTM+4	$3.739B	2.0%	Segment remains below FY2022-FY2023 levels.
TTM+5 / 2030-12-27	$3.813B	2.0%	AMD keeps sticky sockets but does not regain prior peak.

Thesis-consistency check: The base case moves from -2.9% FY2025 growth and +3% Q4 growth to an 8% five-year CAGR. This is not a sharp deceleration problem; it is a recovery assumption from a trough. The bull case at 12% is deliberately above external FPGA market forecasts because it assumes share/content gains from a depressed base, while the bear case at 2% assumes the large design-win funnel converts slowly. No case treats the $17B 2025 design-win number as backlog or divides it mechanically into revenue.

9. Forward Operating Margin Assumptions

The FY2025 / TTM Embedded operating margin base is 36.0%, calculated from $1.243B of operating income on $3.454B of revenue [data_package Section 3, Embedded table; quant_foundation P4]. Margin history is high but volatile: 49.5% in FY2022, 49.4% in FY2023, 39.9% in FY2024, and 36.0% in FY2025 [data_package Section 3, Embedded table]. The five-year margin cases should therefore distinguish between cyclical recovery and full reversion. I do not assume a return to 49% in the base case because FY2025 showed margin erosion, competition has sharpened, and the product mix of future design wins is not disclosed.

Product	Current operating margin	Bear TTM+5y margin	Base TTM+5y margin	Bull TTM+5y margin	Mechanism
Embedded portfolio	36.0%	32.0%	38.0%	42.0%	Bear: weak volume absorption, pricing pressure, and higher support/tool investment. Base: modest scale leverage and mix improvement from recovery. Bull: stronger Versal/adaptive SoC mix and higher revenue absorption, but still below FY2022-FY2023 peak margins.

Base margin moves up 200 bps from 36.0% to 38.0%. The mechanism is operating leverage from revenue recovering from $3.454B to $5.075B and a mix shift toward higher-value adaptive SoC and edge-AI products. The base case still remains below FY2024's 39.9% and far below FY2022-FY2023's roughly 49% because the model does not have product-level margin evidence and because competitive pressure in mid-range FPGAs is real.

Bull margin moves up 600 bps to 42.0%. This assumes design wins convert into high-value products rather than low-end price-sensitive sockets, and that AMD's tool chain and certification support allow pricing power in automotive ADAS, aerospace/defense, industrial robotics, healthcare imaging, and communications. The margin is still conservative versus FY2022-FY2023, because those years may have benefited from post-Xilinx acquisition mix, supply tightness, and peak volume absorption that should not be treated as the steady-state norm.

Bear margin falls 400 bps to 32.0%. This assumes revenue growth is too weak to absorb the fixed engineering and support base, while Altera, Lattice, Microchip, and ASIC/NPU substitutes pressure price and mix. Bear margin does not collapse into the teens because sticky embedded sockets, software/tool switching costs, and long product lives should protect some economics even if growth disappoints.

The margin cases are segment-level because product margins are not disclosed. A revenue-weighted product margin calculation is not possible without inventing subsegment revenue. If AMD later discloses FPGAs versus embedded processors or Versal versus legacy FPGA dollars, this section should be rebuilt with product-level revenue and margin assumptions.

10. Segment Rollup Built Bottoms-Up

Since the data package discloses only one Embedded revenue line, the bottoms-up rollup equals the single disclosed product/sub-segment line. Arithmetic is explicit below.

Case	Product / sub-segment	TTM base revenue	CAGR TTM-to-TTM+5y	TTM+5y revenue	TTM+5y operating margin	TTM+5y operating income
Bear	Embedded portfolio	$3.454B	2.0%	$3.813B	32.0%	$1.220B
Base	Embedded portfolio	$3.454B	8.0%	$5.075B	38.0%	$1.929B
Bull	Embedded portfolio	$3.454B	12.0%	$6.087B	42.0%	$2.556B

Rollup checks:

Case	Product revenue sum	Segment revenue stated	Difference	Product OI sum	Segment OI stated	Difference
Bear	$3.813B	$3.813B	0.0%	$1.220B	$1.220B	0.0%
Base	$5.075B	$5.075B	0.0%	$1.929B	$1.929B	0.0%
Bull	$6.087B	$6.087B	0.0%	$2.556B	$2.556B	0.0%

The base case implies Embedded grows from 10.0% of AMD FY2025 revenue to a larger absolute profit contributor if consolidated AMD growth is led by Data Center. It does not make Embedded the company thesis. Instead, it treats Embedded as a high-margin stabilizer and optionality pool: smaller than Data Center, less cyclical than Gaming, and more profitable than the consolidated business on a segment operating basis. The segment's value in the total AMD model should come from high operating income per revenue dollar and long-cycle design-win durability, not from pretending it will drive the majority of company growth.

11. Key Risks And Wildcards

Bear risk: design wins fail to convert at expected timing. The $17B 2025 design-win disclosure and more than $50B cumulative post-Xilinx figure sound large, but the data package warns they are not revenue or backlog [data_package Section 6, Embedded commentary; Section 10, downstream modeling guardrails]. A design win can be delayed, reduced, lost, or converted at lower content than expected. This is the most important modeling risk because it could make the base case too optimistic.

Bear risk: Altera becomes more aggressive after separation from Intel. Altera now has focused ownership and a clear FPGA mandate, and its Agilex 3/5 roadmap targets embedded, intelligent edge, AI inference, industrial IoT, automotive, and communications use cases [Altera press release, lines 498-502; Intel/Altera release, lines 32-39]. If Altera gains share in mid-range or high-end sockets, AMD's growth and margin recovery could be capped.

Bear risk: low-power FPGA sockets migrate to Lattice or Microchip. AMD does not need to win every low-power application, but if edge AI and industrial automation volume shifts toward small, low-power, lower-cost devices, Lattice and Microchip can take attractive share. Lattice's end-market mix shows real exposure to communications/computing and industrial/automotive, overlapping with AMD's embedded opportunity set [Lattice FY2025 release, lines 1686-1730].

Bear risk: ASICs and dedicated edge-AI processors substitute for programmable logic. In high-volume stable applications, ASICs can deliver better unit cost and power. In AI edge, dedicated NPUs may become easier to program and good enough for vision or sensor-fusion workloads. AMD's best defense is flexibility, long lifecycle, and heterogeneous integration, but the substitute risk is real.

Base risk: industrial and communications cycles remain uneven. Embedded may recover from the FY2024-FY2025 inventory correction, but these markets can still swing with customer capex, distributor inventory, macro conditions, and regional demand. The base case assumes normalization, not a boom.

Bull wildcard: Versal AI Edge and Embedded+ become the default architecture for complex edge-AI systems. AMD's Versal AI Edge positioning spans sensor ingress, preprocessing, AI inference, real-time control, safety, and security [AMD Versal AI Edge Series, lines 532-537 and 553-555]. If customers prefer a programmable heterogeneous platform rather than a fixed-function NPU, AMD can grow above the broad FPGA market.

Bull wildcard: automotive ADAS and aerospace/defense programs scale faster than expected. AMD disclosed Subaru's selection of Versal AI Edge Series Gen 2 for next-generation EyeSight ADAS, with AMD adaptive SoC technology already used in current EyeSight-equipped vehicles [AMD Versal Gen 2 press release, lines 471-473]. A few large automotive or defense platforms can support multi-year revenue if they reach production at scale.

Bull wildcard: AMD's broader company portfolio helps Embedded win. AMD can bring x86 CPUs, GPUs, DPUs/NICs, AI accelerators, FPGAs, adaptive SoCs, and embedded processors to the same strategic customer. That does not automatically translate to Embedded revenue, but it can improve executive access and system-level credibility in edge-to-cloud architectures.

Summary Table

Case	Segment TTM+5y Revenue	Segment TTM+5y Op Margin	Segment TTM+5y Op Income	Probability
Bear	$3.813B	32.0%	$1.220B	25%
Base	$5.075B	38.0%	$1.929B	55%
Bull	$6.087B	42.0%	$2.556B	20%

Source Notes

AMD historical revenue, operating income, operating margin, quarterly trends, and management quotes in this appendix come from `data_package.md` as required. External sources are used only for competitive context, product architecture, and market backdrop: AMD Embedded Computing Solutions, AMD Versal Adaptive SoCs, AMD Versal AI Edge Series, AMD Versal Series Gen 2 press release, Altera/Silver Lake press release, Intel/Altera embedded FPGA releases, Lattice FY2025 release, Lucintel FPGA market report, and Microchip USA FPGA market availability note.

Appendix S - Synthesis (Combined Earnings Model)

Advanced Micro Devices, Inc. - Synthesis (Combined Earnings Model)

PATH_RECEIVED: C:\Users\Admin\OneDrive\GPT\ER 2.0\Advanced Micro Devices, Inc\appendix_synthesis.md

This appendix is the combined AMD earnings model for the ER2 Production Mode run. It uses the orchestrator anchors without recomputing them, reproduces the reviewed segment outputs, adds the required consolidated operating expense bridge, builds net income and EPS, and tests the result against the P/E <= 15 rule at the orchestrator-provided current price of $347.81. The file is intentionally calculation-heavy. Qualitative language is included only where it explains why a mechanical arithmetic result should or should not be trusted.

The important conclusion is direct: the base case does not come close to clearing the P/E 15 gate. The base case is not a weak operating forecast. It takes AMD from $34.639B of FY2025 / TTM revenue to $73.156B at the 2030-12-27 horizon and from $3.694B of reported GAAP operating income to $14.793B of modeled operating income. Even so, after corporate expense, normalized tax, and an explicit share-count path, base EPS is $7.92. At 15x, that is $118.74 per share, or 65.9% below the current price. Required horizon EPS at the current price is $23.19. The bull case also fails at 15x, although by a smaller amount, because even $15.84 of modeled EPS is worth only $237.67 at 15x.

1. Anchors

Anchor	Value
Today	2026-04-29
FY end	Last Saturday in December
TTM anchor end date	2025-12-27
TTM quarters	Q1 FY2025 + Q2 FY2025 + Q3 FY2025 + Q4 FY2025
Latest completed FY label	FY2025
Horizon date	2030-12-27
Compounding span	Exactly 5.0 calendar years from TTM end
Current price	$347.81
TTM diluted share count	1,636M

The TTM base is FY2025 because the data package states that AMD's latest reported quarter is Q4 FY2025 and that Q1 FY2026 has not been reported as of the package date. The model therefore does not create a separate in-progress fiscal-year adjustment. Historical tie-outs use FY2025 as both the TTM anchor and latest completed fiscal-year sanity check.

The consolidated FY2025 control totals from `data_package.md` are revenue of $34.639B, GAAP operating income of $3.694B, net income of $4.335B, diluted EPS of $2.65, and diluted shares of 1,636M. Segment revenue totals are Data Center $16.635B, Client and Gaming $14.550B, and Embedded $3.454B. Segment operating income totals are Data Center $3.603B, Client and Gaming $2.855B, and Embedded $1.243B. The segment operating income sum is $7.701B, which is materially above consolidated GAAP operating income. The difference is unallocated corporate expense, discussed below, and it must be carried forward or the model would overstate earnings.

2. Segment-Level Forward Assumptions

The segment horizon values below reproduce the reviewed segment appendices. No horizon case is recomputed from a different base. The only calculations in this section are summations and margin checks.

Segment	TTM revenue	Bear TTM+5y rev	Base TTM+5y rev	Bull TTM+5y rev	Bear OI margin	Base OI margin	Bull OI margin	Source
Data Center	$16.635B	$29.973B	$49.090B	$80.857B	20.5%	28.5%	33.0%	`appendix_B_data_center.md`, sections 8-10
Client and Gaming	$14.550B	$13.809B	$18.991B	$23.507B	14.0%	21.0%	24.0%	`appendix_B_client_gaming.md`, sections 8-10

Segment	TTM revenue	Bear TTM+5y rev	Base TTM+5y rev	Bull TTM+5y rev	Bear OI margin	Base OI margin	Bull OI margin	Source
Embedded	$3.454B	$3.813B	$5.075B	$6.087B	32.0%	38.0%	42.0%	`appendix_B_embedded.md`, sections 8-10

Data Center is the dominant swing factor. The base case nearly triples segment revenue over five years. It uses a year-by-year path of +60%, +35%, +20%, +10%, and +3.5%, ending at $49.090B. This is aggressive in absolute dollars but still below the bull case because the reviewed segment appendix deliberately discounts Nvidia's ecosystem lead, customer custom silicon, export-control risk, and the bargaining power of hyperscale customers. The appendix also states that AMD does not disclose EPYC or Instinct annual revenue, so the entire Data Center segment is the dollar control line.

Client and Gaming is modeled at the component level for revenue because AMD discloses Client and Gaming component revenue, but operating income remains a combined-segment construct because AMD does not disclose Client operating income and Gaming operating income separately. The base case ends at $18.991B of revenue and $3.987B of operating income. It assumes Client continues to grow from the FY2025 rebound but Gaming remains a drag through the console-cycle trough.

Embedded is a single disclosed dollar line. The base case grows from $3.454B to $5.075B at an 8.0% five-year CAGR and uses a 38.0% operating margin. That is a recovery case, not a direct conversion of the $17B 2025 design-win figure or the more-than-$50B cumulative design-win commentary into revenue. The Embedded appendix explicitly treats design wins as evidence of funnel depth, not booked revenue.

3. Combined Revenue Build

There are no revenue eliminations or corporate revenue adjustments in the data package. The three reportable segments sum to consolidated FY2025 revenue, and the same structure is used at the horizon.

Line	Bear	Base	Bull
Data Center TTM+5y revenue	$29.973B	$49.090B	$80.857B
Client and Gaming TTM+5y revenue	$13.809B	$18.991B	$23.507B
Embedded TTM+5y revenue	$3.813B	$5.075B	$6.087B
Sum of segment TTM+5y revenue	$47.595B	$73.156B	$110.451B
Corporate eliminations	$0.000B	$0.000B	$0.000B
Combined TTM+5y revenue	$47.595B	$73.156B	$110.451B
Implied TTM to TTM+5y CAGR	6.6%	16.1%	26.1%

The base case revenue CAGR of 16.1% is driven almost entirely by Data Center. Client and Gaming contributes moderate growth, while Embedded provides high-margin recovery but is too small to determine the consolidated revenue path. This matters for verdict discipline: if the Data Center ramp is delayed, the base case earnings cushion does not merely narrow. It can disappear because the consolidated model does not have another segment large enough to replace that profit pool.

4. Combined Operating Income Build

The data package and QF derived calculations show that FY2025 segment operating income sums to $7.701B, while consolidated GAAP operating income is $3.694B. The $4.007B difference is treated as unallocated corporate expense. This includes corporate-level operating cost that is not allocated to the three segment appendices. It is not optional. Leaving it out would transform segment operating income into consolidated operating income and materially overstate AMD's earnings power.

The forward corporate expense assumption is TTM unallocated corporate expense grown at 5.0% per year for five years. The horizon corporate expense is therefore $4.007B x 1.05^5 = $5.114B. This is lower growth than base revenue growth because the model assumes corporate leverage, but it is not flat because AMD must keep funding data-center AI software, systems work, customer engineering, legal/compliance, and equity-compensation infrastructure. The management appendix supports this posture: AMD allocates primarily through R&D rather than capex, and FY2025 R&D was $8.091B, or 23.4% of revenue.

Line	Bear	Base	Bull
Sum of segment TTM+5y OI	$9.299B	$19.907B	$34.875B
Unallocated corporate expense	($5.114B)	($5.114B)	($5.114B)

Line	Bear	Base	Bull
Combined TTM+5y OI	$4.185B	$14.793B	$29.761B
Implied OI margin	8.8%	20.2%	26.9%

The base case consolidated operating margin of 20.2% is not simply the revenue-weighted segment margin. It is the segment OI sum after corporate expense. This distinction is crucial. The base segment OI sum is $19.907B, or 27.2% of revenue. After unallocated corporate expense, consolidated operating margin is 20.2%. The $5.114B corporate line consumes 7.0 points of revenue margin at the base horizon.

The bull case reaches a 26.9% consolidated operating margin. That is possible only if Data Center becomes a much larger profit center while corporate cost does not scale proportionally. The bear case remains profitable but unimpressive: $4.185B of operating income on $47.595B of revenue is only 8.8%, below FY2021 AMD operating margin and below what a durable compounder would need to justify the current price.

5. Net Income and EPS Build

The net income bridge uses a normalized 13.0% tax rate from the data package's ROIC and AMD non-GAAP tax-rate convention. FY2025 GAAP tax was a benefit and is not used because it would overstate normalized operating economics. Net interest is modeled as $0.150B of income at the horizon. AMD had net cash of $7.330B at FY2025, so a modest net interest income assumption is reasonable, but it is intentionally small because it is not the thesis driver. Other income is set to zero.

Line	Bear	Base	Bull
Combined TTM+5y OI	$4.185B	$14.793B	$29.761B
Net interest income / (expense)	$0.150B	$0.150B	$0.150B
Other income	$0.000B	$0.000B	$0.000B
Pre-tax income	$4.335B	$14.943B	$29.911B
Tax rate assumed	13.0%	13.0%	13.0%
Net income	$3.771B	$13.000B	$26.023B
Diluted share count at TTM+5y	1,642.3M	1,642.3M	1,642.3M
EPS at TTM+5y	$2.30	$7.92	$15.84

The base EPS calculation is $13.000B of net income divided by 1,642.3M diluted shares, which equals $7.92. The bull EPS calculation is $26.023B divided by 1,642.3M, or $15.84. These are GAAP-oriented operating-model outputs. They do not add back stock-based compensation, amortization, customer warrant economics, or other non-GAAP items to chase management's strategic-target EPS framework. That is intentional. ER2's P/E 15 test is a common-share earnings-power test, not a press-release non-GAAP target test.

6. Share Count Path

The starting diluted share count is 1,636M from the data package. The buyback assumption is $1.316B annually, equal to FY2025 repurchase dollars and consistent with the idea that AMD's buyback is a dilution-management tool rather than a dominant capital-return identity. The management appendix notes that buybacks have been present but secondary and that AMD has a large authorization. It also warns that customer warrants tied to OpenAI and Meta can become a major dilution source if milestones are met. This base share path includes routine SBC offset, not full customer-warrant vesting.

SBC issuance is estimated from FY2025 stock-based compensation of $1.638B divided by the orchestrator current price of $347.81, or 4.7M shares per year. Buyback average price starts at $347.81 and rises 5.0% per year, which makes repurchases less effective over time. This is a conservative mechanical approach because it does not assume AMD can repurchase large amounts of stock at depressed prices.

Year	Starting share count	Buyback $	Avg price assumed	Shares retired	SBC issuance	Ending share count
TTM+1	1,636.0M	$1.316B	$347.81	3.8M	4.7M	1,636.9M
TTM+2	1,636.9M	$1.316B	$365.20	3.6M	4.7M	1,638.0M
TTM+3	1,638.0M	$1.316B	$383.46	3.4M	4.7M	1,639.3M
TTM+4	1,639.3M	$1.316B	$402.63	3.3M	4.7M	1,640.8M
TTM+5	1,640.8M	$1.316B	$422.77	3.1M	4.7M	1,642.3M

The five-year path retires 17.2M shares and issues 23.5M SBC-equivalent shares, increasing diluted shares by 6.3M to 1,642.3M. This is a small change relative to AMD's current share base, but it matters because the company is already far below the required EPS for the P/E 15 gate. The model does not assume that repurchases "maintain EPS" automatically. Buybacks at a high stock price do little to reduce share count, while SBC and potential customer warrants create a real dilution overhang.

7. P/E 15 Test

Required EPS for a 15x P/E at $347.81 is $23.19. The base case is the controlling case for the ER2 verdict. If base EPS x 15 is below the current price, the stock fails on price even if the bull case is attractive in operating terms.

Case	EPS at TTM+5y	EPS x 15	Current price	Cushion	Pass?
Bear	$2.30	$34.45	$347.81	-90.1%	FAIL
Base	$7.92	$118.74	$347.81	-65.9%	FAIL
Bull	$15.84	$237.67	$347.81	-31.7%	FAIL

The P/E 15 gate is therefore FAIL. This is not a close fail caused by a rounding convention or a conservative net-interest line. To pass the gate, base EPS would need to be $23.19, or about 2.9x the modeled $7.92. Even if the base case used the starting 1,636M shares instead of the explicit 1,642.3M ending share count, base EPS would be only $7.95. Even if corporate expense were held flat at $4.007B instead of grown to $5.114B, base EPS would increase by roughly $0.59 after tax. Those changes do not alter the verdict.

8. Reflexivity Check

The TTM and FY2025 reconstruction ties because TTM equals FY2025 under Stage 0. Segment revenue sums to $34.639B, exactly matching consolidated revenue. Segment OI sums to $7.701B. After the $4.007B unallocated corporate expense line, consolidated OI is $3.694B, matching the data package. This is a critical control check because the forward model compounds from the segment base but must report consolidated earnings after corporate expense.

Historical tie-out	Segment sum / reconstructed	Consolidated reported	Difference	Result
TTM revenue	$34.639B	$34.639B	0.0%	PASS
TTM operating income after corporate expense	$3.694B	$3.694B	0.0%	PASS
FY2025 revenue	$34.639B	$34.639B	0.0%	PASS
FY2025 operating income after corporate expense	$3.694B	$3.694B	0.0%	PASS

The reflexivity risk is not historical arithmetic. It is forward interpretation. AMD's segment appendices are strong enough to build a plausible high-growth model, but the current share price already discounts something much closer to management's aspirational non-GAAP EPS target than to this GAAP-oriented base case. The moat appendix supports that caution: moat gate YELLOW and 2046 test YELLOW. AMD has real assets, but not a fortress-like right to compound earnings at any price.

9. In-Progress FY Reconciliation

No separate in-progress FY reconciliation is required because Stage 0 uses FY2025 / TTM as the latest completed fiscal year and Q1 FY2026 had not been reported in the data package. There are no YTD FY2026 actuals to reconcile. However, the Q4 FY2025 management commentary and Q1 FY2026 guide provide a useful guardrail. AMD guided Q1 FY2026 revenue to approximately $9.8B plus or minus $0.3B, which annualizes to $39.2B before seasonality and quarter mix. That is above the FY2025 TTM revenue base and consistent with a model that begins with a higher forward year, but it is below the model's base TTM+1 revenue path of $45.239B.

The base TTM+1 revenue bridge is Data Center $26.616B, Client and Gaming $14.893B, and Embedded $3.730B, for $45.239B. The reason this is materially above FY2025 is not hidden in the synthesis model: Data Center is explicitly modeled at +60% in the first forward year based on management's "more than 60% annually" Data Center commentary and the OpenAI / Meta multi-year demand anchors. This is an explained acceleration, not an arithmetic accident. It should still be monitored carefully because it creates a high burden of evidence in FY2026 and FY2027.

10. Sensitivity

The sensitivity table holds all other base assumptions constant and varies one driver at a time. For the Data Center revenue sensitivity, the base Data Center five-year CAGR is 24.2%, and the table shows 22.2% and 26.2%. For Client and Gaming margin, the base segment margin is 21.0%, and the table shows 19.0% and 23.0%. For share count, the table shows the base ending share count plus or minus 5%. For tax, the base normalized rate is 13.0%, and the table shows 10.0% and 16.0%.

Sensitivity	EPS at TTM+5y	Delta vs base
Data Center CAGR -200 bps	$7.35	($0.57)
Base case	$7.92	$0.00
Data Center CAGR +200 bps	$8.54	$0.63
Client and Gaming OI margin -200 bps	$7.71	($0.20)
Client and Gaming OI margin +200 bps	$8.12	$0.20
Share count -5%	$8.33	$0.42
Share count +5%	$7.54	($0.38)
Tax rate -300 bps	$8.19	$0.27
Tax rate +300 bps	$7.64	($0.27)

The most important variable is Data Center revenue scale, followed by share count only if customer warrants begin to vest in size. Client and Gaming margin and tax rate matter, but neither can rescue the P/E 15 test. A 200 bps Data Center CAGR improvement adds only about $0.63 of EPS in this one-variable table. That is meaningful for earnings revisions but not nearly enough to close the $15.27 EPS gap between base EPS of $7.92 and required EPS of $23.19. The verdict therefore depends less on small modeling refinements and more on whether AMD can deliver an earnings level far above this base model while avoiding large warrant dilution and preserving Data Center margins.

11. Combined Earnings Summary

Case	Revenue	OI	OI margin	NI	EPS	Probability	EPS x 15	vs. $347.81
Bear	$47.595B	$4.185B	8.8%	$3.771B	$2.30	30%	$34.45	-90.1%
Base	$73.156B	$14.793B	20.2%	$13.000B	$7.92	50%	$118.74	-65.9%
Bull	$110.451B	$29.761B	26.9%	$26.023B	$15.84	20%	$237.67	-31.7%
Probability-weighted EPS	n/a	n/a	n/a	n/a	$7.82	100%	$117.24	-66.3%

The probability weights are Bear 30%, Base 50%, and Bull 20%. They are not uniform because the disqualifier sweep is mixed: management passes, but moat and 2046 durability are YELLOW, Data Center execution risk is high, and the current price already requires an exceptional earnings outcome. The bull case is possible, but it is not the correct center of gravity for a price-discipline framework.

Model Control Notes

Three control choices deserve explicit documentation because they prevent common AMD modeling errors. First, the model does not add back stock-based compensation. SBC is both an expense and a dilution source, and the share-count path already uses FY2025 SBC to estimate annual issuance. Adding back SBC to net income while also ignoring the resulting shares would double count economics in favor of the equity holder. AMD may report non-GAAP earnings that exclude SBC, amortization, and other items, but this synthesis is designed to answer whether the common stock clears a disciplined P/E test on durable earnings power.

Second, the model does not capitalize customer warrants or treat them as free demand creation. The OpenAI and Meta warrant structures can create enormous strategic value if they validate Instinct at scale, but they can also transfer meaningful ownership to customers. Because the exact milestone vesting, accounting treatment, and revenue quality cannot be known from the reviewed appendices alone, the base case excludes full warrant dilution and flags the warrants as a monitoring item. This choice is favorable to AMD compared with a fully diluted customer-warrant case, so the valuation failure is not caused by an overly punitive warrant assumption.

Third, the model keeps Data Center product disclosure discipline. The appendices make clear that AMD does not disclose annual Instinct revenue or annual EPYC revenue. The synthesis therefore uses the disclosed Data Center segment as the operating control total. This avoids the appearance of precision that would come from inventing AI accelerator dollars, then separately assigning margins to an undisclosed product line. The cost of this discipline is that the model cannot isolate which part of Data Center creates or destroys the margin bridge. That is an analytical limitation, but it is preferable to false product-level confidence.

Tie-Out Checks

- Reported-financial tie-out: PASS. TTM equals FY2025 under Stage 0. Segment revenue sums to $34.639B versus consolidated $34.639B. Segment OI of $7.701B less unallocated corporate expense of $4.007B equals consolidated OI of $3.694B. Latest completed FY2025 ties on the same basis.

- NI / shares = EPS: PASS. Bear $3.771B / 1,642.3M = $2.30. Base $13.000B / 1,642.3M = $7.92. Bull $26.023B / 1,642.3M = $15.84. Differences are rounding only.

- Segment sum tie-out at horizon: PASS. Horizon revenue equals segment sums in all cases: Bear $29.973B + $13.809B + $3.813B = $47.595B; Base $49.090B + $18.991B + $5.075B = $73.156B; Bull $80.857B + $23.507B + $6.087B = $110.451B. Revenue eliminations are zero.

- Product-to-segment rollup: PASS. Data Center has no disclosed product-dollar split and uses the segment control total as the dollar line. Client and Gaming explicitly ties Client plus Gaming to segment revenue in all cases. Embedded is a single product/sub-segment line and ties exactly.

- Weighted margin check: PASS. Data Center and Embedded have no product-dollar margin split, so the segment margin is the controllable disclosed line. Client and Gaming weighted product margins tie: Bear OI $1.935B on $13.809B is 14.0%; Base OI $3.987B on $18.991B is 21.0%; Bull OI $5.636B on $23.507B is 24.0%.

- Share count path: PASS. Starting diluted shares are 1,636.0M. Across five years, routine SBC adds 23.5M shares and buybacks retire 17.2M shares. Ending diluted shares are 1,642.3M. The path covers exactly five years, not six.

- Growth plausibility: PASS. Data Center base CAGR of 24.2% is below 1.2x its disclosed three-year historical rate of 40.2% and is supported by management's Data Center growth commentary. Client base CAGR of 8.0% is far below the FY2023-FY2025 rebound rate. Gaming base CAGR is -3.0%. Embedded base CAGR of 8.0% is above its negative three-year historical rate but is justified as recovery from a downcycle and is within the external FPGA-market range cited in the Embedded appendix.

- Historical arithmetic: PASS. The five-year CAGR calculations tie within 1% to the segment horizon values. Examples: Data Center base $16.635B x 1.242^5 = $49.162B versus stated $49.090B; Client base $10.640B x 1.08^5 = $15.634B versus stated $15.633B; Gaming base $3.910B x 0.97^5 = $3.358B; Embedded base $3.454B x 1.08^5 = $5.075B. Rounding differences are immaterial.

- Recency anchor: PASS with explicit acceleration explanation. A strict comparison of the base TTM+1 path of $45.239B to the FY2025 TTM revenue base of $34.639B shows +30.6%, above the normal 10% guardrail. This is not hidden model drift; it is driven by the explicit Data Center +60% first-forward-year assumption from the reviewed segment appendix and management's Data Center ramp commentary. Q1 FY2026 guide of about $9.8B annualizes to $39.2B, so early FY2026 must show clear acceleration for the base path to remain credible.

- P/E 15 arithmetic: PASS as arithmetic, FAIL as valuation gate. Required EPS is $347.81 / 15 = $23.19. Bear, base, and bull EPS x 15 are $34.45, $118.74, and $237.67. All are below the current price, so the P/E 15 gate is FAIL with a base-case cushion of -65.9%.

- Probability-weighted summary: PASS. Scenario probabilities sum to 100%. Probability-weighted EPS is 30% x $2.30 + 50% x $7.92 + 20% x $15.84 = $7.82. EPS x 15 is $117.24, or 66.3% below the current price.

Verdict

The combined earnings model is internally tied and reproduces the reviewed segment values, but it does not support a BUY under ER2 price discipline. AMD can have a strong management team, real Data Center upside, and a plausible five-year revenue doubling and still fail at the current price. The problem is the required earnings level. At $347.81, the P/E 15 bar requires $23.19 of horizon EPS. The base model produces $7.92, and even the bull case produces only $15.84.

The correct synthesis posture is therefore FAIL on valuation, not because AMD lacks a growth story, but because the current price demands a far more extreme and cleaner earnings outcome than the bottom-up segment model supports. The strongest trigger that could change this would be evidence that Data Center revenue is converting into durable consolidated operating margin far above the base case without major warrant dilution. The clearest kill criteria are Data Center growth below the FY2026 ramp path, recurring export-control or customer-incentive margin leakage, and visible warrant dilution without corresponding GAAP earnings power.

Appendix CS - Compounder Score

Advanced Micro Devices, Inc. - Diagnostic Compounder Score

PATH_RECEIVED: C:\Users\Admin\OneDrive\GPT\ER 2.0\Advanced Micro Devices, Inc\appendix_compounder_score.md

1. Executive Verdict

This appendix scores Advanced Micro Devices, Inc. under the ER2 Compounder Score rubric. The score is diagnostic only. It does not override the completed ER2 verdict, which was PASS on valuation discipline. The evidence supports AMD as a capable, strategically important semiconductor competitor with excellent management, a strong balance sheet, real product franchises, and a large AI/data-center runway. The evidence does not support classifying the stock as investable under a 20% annualized return underwriting standard at the ER2 price anchor of $347.81.

The total diagnostic score is **54 / 100**, split between **Business Quality 54 / 80** and **20% Return Underwriting 0 / 20**. The score-implied classification is **Reject** under this rubric. That classification is not a statement that AMD is a bad company. It is a statement that the current evidence does not show a high-ROIIC, durable, conservatively priced compounder at the starting valuation used in the completed ER2 report.

The main reason for the cap is simple: AMD's business narrative is better than its consolidated reported return evidence. FY2025 revenue was $34.639B and official continuing free cash flow was $5.519B, but GAAP operating income was only $3.694B and the data package's reported ROIC was 5.8%. The five-year reported ROIC series was 73.6%, 3.9%, 0.7%, 3.1%, and 5.8% from FY2021 through FY2025. That series is distorted by Xilinx, amortization, AI investment, and cyclicality, but the rubric does not allow AMD to earn high ROIC points from a future story alone.

The valuation block receives no points because the 20% case requires either enormous NOPAT growth from a low current base, a terminal multiple that remains extremely high, or both. The synthesis base case modeled FY2030 / TTM+5 revenue of $73.156B, operating income of $14.793B, net income of $13.000B, and EPS of $7.92. Even the bull case produced EPS of $15.84. At the $347.81 price used in the completed report, the base case was worth only $118.74 at 15x EPS, and the bull case was worth $237.67. The 20% return case therefore requires an outcome far beyond the reviewed base case and still above the reviewed bull case.

Source confidence is **Medium**. The completed ER2 packet has strong local evidence for FY2025 consolidated financials, segment control totals, management, moat, and the five-year earnings model. Confidence is not High because product-level Data Center revenue is not disclosed, Instinct and EPYC economics cannot be separated, customer warrant dilution is still uncertain, and current market data was intentionally not refreshed because this is a backfill against the completed ER2 report.

2. Data Source Map

Scoring input	Source used	Confidence	Notes
Price / market cap	`main_narrative.md`, `appendix_synthesis.md`, `quant_foundation.md`	Medium	Completed ER2 report used $347.81 and $567.1B market cap from queue snapshot; quant file also retained an earlier FMP quote at $323.21 and $527.0B. The score follows the completed report anchor.
Enterprise value	Calculated from completed report market cap and data package net cash	Medium	EV approx. $559.8B = $567.1B market cap less $7.330B net cash.
Income statement	`quant_foundation.md`, `data_package.md`, raw FMP annual income	High	FY2025 revenue $34.639B, operating income $3.694B, net income $4.335B.
Balance sheet	`data_package.md`, raw FMP balance sheet	High	Cash + short-term investments $10.552B, total debt $3.222B for leverage, goodwill + intangibles $41.831B.
Cash flow	`quant_foundation.md`, `data_package.md`, raw FMP cash flow	High	FY2025 official continuing CFO $6.493B and FCF $5.519B; raw FMP including discontinued operations was higher.
Share count	`quant_foundation.md`, `appendix_synthesis.md`	High	FY2025 diluted shares 1,636M; horizon model ends at 1,642.3M.

Scoring input	Source used	Confidence	Notes
SBC	`appendix_synthesis.md`, `appendix_management.md`, raw FMP cash flow	High	FY2025 SBC $1.638B, or 4.7% of revenue and 29.7% of continuing FCF.
Segment data	`quant_foundation.md`, `data_package.md`, segment appendices	High	Data Center $16.635B revenue / $3.603B OI; Client and Gaming $14.550B / $2.855B; Embedded $3.454B / $1.243B.
Management / incentives	`appendix_management.md`	High	Management gate PASS; proxy and filings reviewed locally.
Moat / durability	`appendix_moat.md`, segment appendices	Medium	Moat gate YELLOW and 2046 test YELLOW; strongest moat is Embedded, most important growth moat is still contested Data Center AI.
TAM / market size	`data_package.md`, segment appendices	Medium	AI/data-center opportunity is large, but management TAM and partner announcements are not treated as booked revenue.
Competitors	`appendix_moat.md`, segment appendices	Medium	Nvidia, Intel, hyperscaler custom silicon, Arm, Altera, Lattice, Microchip, and ASIC/ASSP substitutes.

3. Key Financial and Quality Metrics

Metric	Value	Source / formula
Analysis date	2026-04-29	Completed ER2 packet date
TTM end	2025-12-27	Stage 0 anchor
Current price used	$347.81	Completed ER2 report
Market cap used	$567.1B	Completed ER2 report
Enterprise value used	$559.8B	Market cap less net cash
FY2025 revenue	$34.639B	Quant foundation / data package
FY2025 operating income	$3.694B	Quant foundation / data package
FY2025 normalized NOPAT	$3.214B	$3.694B x 87%
FY2025 official continuing FCF	$5.519B	QF/data package
FY2025 diluted shares	1,636M	QF/data package
5-year revenue CAGR	28.8%	FY2020 $9.763B to FY2025 $34.639B
5-year NOPAT CAGR	21.9%	FY2020 normalized NOPAT $1.191B to FY2025 $3.214B
5-year FCF CAGR	48.0%	FY2020 FCF $0.777B to FY2025 official continuing FCF $5.519B
Reported ROIC 5-year median	3.9%	Data package reported ROIC series
3-year acquisition-inclusive ROIIC	17.4%	ΔNOPAT FY2022-FY2025 / capex + working capital investment + acquisitions, FY2023-FY2025
5-year acquisition-inclusive ROIIC	12.4%	ΔNOPAT FY2020-FY2025 / capex + working capital investment + acquisitions, FY2021-FY2025
FY2025 FCF / NOPAT	171.7%	$5.519B / $3.214B
5-year FCF / NOPAT average	204.8%	FY2021-FY2025, using official FY2025 continuing FCF
FY2025 SBC / revenue	4.7%	$1.638B / $34.639B
FY2025 SBC / FCF	29.7%	$1.638B / $5.519B
Net debt / adjusted EBITDA	-0.86x	Data package
EV / NOPAT	174.2x	$559.8B / $3.214B
FCF yield	1.0%	$5.519B / $567.1B
Dividend yield	0.0%	No common dividend
Net buyback yield	0.2%	FY2025 net repurchase dollars approx. $1.031B / $567.1B

The ROIIC calculations are deliberately conservative. Working-capital cash outflows were treated as reinvestment, and acquisitions were included. This matters for AMD because the last five years include Xilinx, ZT Systems, large R&D growth, and major balance-sheet intangibles. If one excludes goodwill and treats some acquired intangible amortization as non-economic, AMD's apparent return profile improves. The full compounder rubric, however, requires acquisition-inclusive return evidence because acquired capital and customer warrant economics are real common-shareholder costs.

4. Rubric Scorecard

Category	Score	Max	Rationale
Incremental Returns on Capital	14	25	Cash conversion is strong and ROIIC is positive, but reported ROIC including acquisition capital is low and current Data Center AI returns are not yet separately proven.
Reinvestment Runway	16	20	AI/data-center runway is large, AMD has shown ability to fund R&D at scale, and Embedded design wins add depth; score is capped by contested market structure and margin uncertainty.
Durability of Moat	9	15	AMD has real moats in Embedded, EPYC, x86, adaptive compute, and procurement trust, but company-level moat is only YELLOW and Data Center AI moat is still forming against Nvidia and custom silicon.
Accounting and Financial Quality	8	10	Financials tie to primary sources, cash conversion is strong, and balance sheet is net cash; score is capped for SBC, customer warrants, and acquisition-intangible complexity.
Capital Allocation Discipline	7	10	R&D-first capital allocation is coherent and management is strong; score is capped by unproven ZT/AI systems returns, Xilinx goodwill, buybacks at high prices, and potential warrant dilution.
Price and 20% Return Underwriting	0	20	Base, bear, and even reviewed bull cases fail the 20% return standard without terminal multiple expansion.
Total	**54**	**100**	Score-implied classification: Reject.

Subscore Detail

Rubric item	Score	Max	Evidence
1A Normalized ROIC is high	2	5	FY2025 reported ROIC was 5.8%; excluding goodwill would look better, but acquisition-inclusive capital is real.
1B ROIIC is high	4	8	3-year acquisition-inclusive ROIIC estimated at 17.4%; 5-year at 12.4%.
1C ROIIC is not deteriorating sharply	2	3	FY2023-FY2025 recovery is visible, but the five-year record is volatile.
1D Unit economics support reported returns	2	3	Embedded margins and Client mix support differentiation; Data Center AI unit economics remain undisclosed.
1E Cash conversion is real	3	3	FY2025 FCF / NOPAT was 171.7%; CFO and FCF are strong versus earnings.
1F ROIC is not an accounting artifact	1	3	Reported ROIC is not inflated by tiny capital; it is low because acquisition capital and intangibles matter.
2A Large remaining market opportunity	5	5	AI accelerators, server CPUs, client AI PCs, and embedded adaptive compute provide major runway.

Rubric item	Score	Max	Evidence
2B Proven ability to reinvest at scale	4	5	R&D rose to $8.091B in FY2025; management has executed EPYC, Ryzen, Xilinx integration, and Instinct ramp.
2C Organic growth opportunity	3	4	Data Center and Client growth are visible, but Instinct and EPYC product dollars are not disclosed.
2D Repeatable expansion playbook	2	3	Chiplet, x86, and adaptive-compute playbooks are real; rack-scale AI is newer.
2E Market structure supports long growth	1	2	Market grows, but Nvidia, Intel, Arm, ASICs, and hyperscaler silicon keep structure tough.
2F Growth does not require ROIC collapse	1	1	Base case assumes operating leverage and FCF generation; not enough to rescue valuation.
3A Clear source of advantage	3	4	Procurement trust, x86, chiplets, EPYC, Xilinx adaptive compute, and R&D scale.
3B Pricing power	1	3	Pricing power is visible in pockets, not yet consolidated.
3C Competitive evidence	2	3	Share gains and partnerships are real; Nvidia and custom silicon remain high threats.
3D Customer stickiness	1	2	Embedded and server qualification help; hyperscalers retain bargaining power.
3E Resilience through stress	1	2	AMD survived export controls and cycles, but consolidated margins remain volatile.
3F Low disruption risk	1	1	AI is more demand driver than direct substitution risk, though competitive disruption remains.
4A Clean revenue and earnings quality	2	2	Primary-source tie-outs passed.
4B Cash flow supports earnings	2	2	Cash conversion is strong.
4C Conservative SBC and dilution	1	2	SBC is meaningful and customer warrants could dilute owners.
4D Balance sheet strength	2	2	Net cash balance sheet.
4E Transparent disclosures	1	1	Consolidated and segment disclosure are adequate.
4F No governance/accounting red flags	0	1	No fatal red flags, but warrant/accounting complexity and late Form 4 keep this from a clean full point.
5A Reinvestment priority is rational	2	2	R&D-first posture is appropriate.
5B Acquisition discipline	1	2	Xilinx and ZT are strategically coherent but not fully proven on returns.
5C Buyback discipline	1	2	Buybacks offset dilution but are not compelling at high valuation.
5D Incentives align with value creation	1	2	Pay is performance-linked, but special CEO award and warrant dilution need monitoring.
5E Owner mentality	1	1	Lisa Su and executives own meaningful shares.
5F Evidence from past decisions	1	1	Long-term turnaround record is strong.
6A Base-case IRR >= 20%	0	6	Base case is negative under conservative exit multiple.
6B Conservative exit multiple	0	4	A lower exit multiple destroys return; flat/current multiple would be extreme.
6C Bear case is acceptable	0	3	Bear case has severe downside.
6D Reverse DCF is beatable	0	3	Required NOPAT CAGR is too high under a sane exit multiple.

Rubric item	Score	Max	Evidence
6E Multiple does not price perfection	0	2	Starting valuation already prices exceptional execution.
6F Cash returns help	0	2	No dividend and buybacks are too small relative to market cap.

5. ROIIC and Reinvestment Analysis

AMD's reported return profile is the central tension in this score. The business is strategically better than the reported ROIC line suggests, but the compounder rubric requires evidence that incremental capital is earning high returns. The evidence is not strong enough to award high points.

NOPAT was calculated as GAAP operating income multiplied by one minus a 13% normalized tax rate. The data package uses this normalized tax rate because FY2025 GAAP tax was a benefit and would overstate normalized economics. On that basis, FY2020 NOPAT was about $1.191B, FY2022 was about $1.100B, and FY2025 was about $3.214B. The five-year NOPAT CAGR from FY2020 to FY2025 was 21.9%, which is attractive. The problem is the denominator: AMD also absorbed Xilinx, ZT, higher R&D, and working-capital investment.

For the three-year ROIIC estimate, FY2022 to FY2025 NOPAT increased by roughly $2.114B. Reinvestment over FY2023-FY2025 was estimated as capex of $2.156B, working-capital investment of $7.525B, and acquisitions of $2.456B. That gives acquisition-inclusive ROIIC of about 17.4%. For the five-year estimate, FY2020 to FY2025 NOPAT increased by roughly $2.023B, while FY2021-FY2025 reinvestment was about $16.330B, producing about 12.4% ROIIC. These are positive but not dream-compounder numbers.

The score is not harsher because AMD is fabless and R&D-heavy. Capex alone understates reinvestment because AMD's real capital budget is people, design, verification, software, and systems engineering. FY2025 R&D was $8.091B, or 23.4% of revenue. That R&D creates valuable intangible assets if AMD's roadmaps win, but it is still a real economic investment. A score that gives AMD full credit for R&D-created future profit before the Data Center AI margin proof arrives would violate the rubric.

Cash conversion is the strongest quantitative offset. FY2025 official continuing CFO was $6.493B and official continuing FCF was $5.519B, compared with normalized NOPAT of $3.214B. The data package also notes that raw FMP FCF including discontinued operations was higher, but the official continuing-operations number is the right quality metric because ZT manufacturing divestiture mechanics created noise. Even using the lower official FCF, cash conversion is excellent.

6. Runway Analysis

AMD's runway score is high. Data Center is already 48.0% of FY2025 revenue and has the largest growth vector. The completed synthesis base case grows Data Center revenue from $16.635B to $49.090B by 2030-12-27. That is aggressive, but it is not unsupported: the data package cites Q4 FY2025 management commentary that AMD was well positioned to grow Data Center revenue by more than 60% annually and scale AI to tens of billions in annual revenue in 2027. OpenAI and Meta announced up-to-6GW Instinct deployment partnerships, and AMD is pushing from chips toward rack-scale systems through ZT design capability and Helios.

The runway score is capped because the most important runway is also the least proven at product-level economics. AMD does not disclose annual Instinct revenue or EPYC revenue. Data Center segment operating margin was 21.7% in FY2025, with Q4 at 32.6%, but the full-year number still reflects ramp cost, export-control noise, and mix uncertainty. The model needs Data Center not merely to grow, but to grow at attractive incremental margins while avoiding heavy customer-incentive dilution.

Client and Gaming adds a secondary runway. FY2025 Client and Gaming revenue rose 50.8% to $14.550B, with Client strength supported by richer Ryzen mix and enterprise-PC opportunity. Gaming remains cyclical because semi-custom console revenue can decline sharply and discrete GPUs compete against Nvidia's ecosystem. Embedded adds the highest-quality but smallest runway. FY2025 Embedded revenue was only $3.454B, but operating margin was 36.0%, and the Q4 transcript cited $17B of 2025 design wins and more than $50B since Xilinx. Those design wins are not backlog, but they support a credible recovery path.

7. Moat Analysis

The moat score follows the completed moat appendix: **Moat gate YELLOW** and **2046 test YELLOW**. AMD has real moats, but they are uneven. Embedded is the clearest segment-level moat because qualified adaptive-compute and FPGA designs are sticky, long-lived, and supported by tools, IP, and customer engineering. Data Center has a credible moat in EPYC CPUs and a developing procurement alternative moat in AI accelerators and systems. Client and Gaming has product-cycle advantages, brand, x86 compatibility, and console design-in value, but it is more cyclical and less protected.

The company-level moat is not a fortress. The largest forward profit pool is Data Center AI, where Nvidia has a stronger software ecosystem, networking position, rack-scale platform, and developer habit. Hyperscaler custom silicon and Arm-based host CPUs can also pressure merchant silicon economics. AMD may still win an attractive share of the profit pool as a second platform and customer co-design partner, but the evidence does not show Nvidia-like lock-in.

Pricing power is moderate. Embedded margin persistence is good evidence. Client ASP commentary is supportive. Data Center pricing power remains unverified because AMD does not disclose AI accelerator price, margin, customer concentration economics, or contract structures. The OpenAI and Meta relationships prove customer interest and strategic relevance, but they do not by themselves prove high-return, non-dilutive, recurring economics.

8. Accounting and Capital Allocation Analysis

Accounting and financial quality are above average. The quant foundation passed 8/8 reconciliation checks. Consolidated FY2025 revenue, operating income, segment revenue, segment operating income, cash flow, and balance-sheet liquidity were tied to primary sources and live FMP pulls. The balance sheet is strong: cash plus short-term investments were $10.552B, total debt for leverage was $3.222B, and net debt / adjusted EBITDA was -0.86x.

The main accounting-quality deductions are not fraud concerns. They are economic interpretation concerns. SBC was $1.638B in FY2025, or 4.7% of revenue and 29.7% of official continuing FCF. SBC is a real cost under this rubric. The OpenAI and Meta warrants could become much larger than routine employee dilution if milestones are met. The completed synthesis deliberately did not assume full warrant dilution, which is favorable to AMD; therefore the valuation failure cannot be blamed on an excessively punitive dilution case.

Capital allocation is coherent but not fully proven. Lisa Su's long-term track record is strong, and the management appendix rated the management gate PASS. AMD allocates primarily through R&D rather than capex or dividends, which is appropriate for a fabless high-performance compute challenger. The Xilinx acquisition broadened AMD into adaptive compute and Embedded, but it also created large goodwill and intangible assets. ZT Systems is strategically coherent because AMD retained rack-scale design capability and sold the manufacturing business, but the return on that move remains to be proven through Data Center revenue, margin, and customer renewal evidence.

Buybacks do not materially help the score. FY2025 buybacks were $1.316B, while the market cap used in the completed report was $567.1B. The synthesis models buybacks as a dilution-management tool: routine SBC adds 23.5M shares over five years and buybacks retire 17.2M shares, leaving ending diluted shares slightly higher at 1,642.3M. Repurchases at a high stock price do not create enough per-share value to offset the valuation problem.

9. 20% Return Underwriting

The 20% return underwriting score is **0 / 20**. The base case does not work, the bear case is unacceptable, the reverse DCF is not beatable under a conservative exit multiple, and cash returns are negligible.

Case	Horizon NOPAT / value basis	Exit assumption	Ending value per share	Implied IRR from $347.81
Bear	$3.641B NOPAT	20x EV / NOPAT	~$48.8	-32.4%
Base	$12.870B NOPAT	20x EV / NOPAT	~$160.9	-14.3%
Bull	$25.892B NOPAT	20x EV / NOPAT	~$319.5	-1.7%
Base P/E 15 cross-check	$7.92 EPS	15x EPS	$118.74	-19.3%
Bull P/E 15 cross-check	$15.84 EPS	15x EPS	$237.67	-7.3%

The 20x EV / NOPAT exit assumption is already generous for a semiconductor company whose moat gate is YELLOW and whose largest profit pool is contested. It is far below the starting EV / NOPAT of roughly 174x, but that starting multiple is the problem. If the exit multiple were held at today's approximate EV / NOPAT multiple, a 20% return would be mechanically easier, but that would be the exact terminal multiple dependence the rubric prohibits.

The reverse DCF makes the problem explicit. Starting EV is approximately $559.8B and FY2025 normalized NOPAT is $3.214B. To earn 20% annually for five years with a 20x exit EV / NOPAT multiple, required ending NOPAT would be roughly $69.7B, implying an approximate **85% NOPAT CAGR**. The completed bull case's horizon NOPAT was only about $25.9B. With a 30x exit multiple, required NOPAT CAGR would still be roughly 68%. The evidence does not support that as a base case.

The base-case maximum buy price for a 20% IRR, using the 20x EV / NOPAT exit framework, is approximately **$64.70 per share**. This is not a target price. It is a diagnostic output showing how far the 20% hurdle is from the completed report's $347.81 price anchor. Using the synthesis's stricter 15x EPS value of $118.74 would imply a maximum buy price of roughly $47.70 for a 20% five-year return.

10. Hard Veto Checklist

Hard veto	Status	Evidence
ROIIC below cost of capital with no credible path	Watch	Estimated 5-year acquisition-inclusive ROIIC is only 12.4%, but runway and margin recovery create a credible improvement path.
Reported ROIC high only because capital is tiny	Pass	Reported ROIC is low, not artificially high; capital base includes acquisition intangibles.
Growth requires persistent external financing	Pass	Net cash balance sheet and strong FCF.
Share count rises materially without value creation	Watch	Routine share count is controlled, but OpenAI and Meta warrants could be large.
Large acquisitions without disclosed returns	Watch	Xilinx and ZT are strategic, but acquisition-inclusive returns are not clearly proven.
FCF structurally lags earnings	Pass	FCF exceeds normalized NOPAT.
Permanent impairment risk from concentration/regulation/cyclicality	Watch	Export controls, China exposure, AI customer bargaining power, and PC/gaming cyclicality are material.
20% return case requires terminal multiple expansion	Fail	20% case does not work under conservative exit multiples.
Needs flawless execution to justify today's price	Fail	Current price requires a far above-base earnings outcome.
Cannot explain moat in one paragraph	Pass	Moat can be explained as high-trust compute procurement, x86/chiplet/adaptive-compute IP, R&D scale, and design-in switching costs.
Accounting/governance issues impair trust	Watch	No fatal trust issue, but warrants, SBC, and special equity awards require monitoring.
Balance sheet risk could wipe out equity	Pass	Net cash and low leverage.

There are two valuation hard vetoes: the 20% return case requires an inadmissible multiple assumption, and the current price requires near-flawless execution relative to the completed ER2 model. Those vetoes are enough to prevent an "Investable now" classification regardless of business-quality positives.

11. Final Decision

Score-implied classification: Reject.

The evidence supports AMD as a high-quality challenger with a large runway, but it does not support the diagnostic compounder ideal at the completed report's price. Business Quality is respectable but below the 65 / 80 threshold because current reported ROIC is low, acquisition-inclusive ROIIC is not high enough, Data Center AI unit economics are not separately disclosed, and the moat gate is YELLOW. Valuation / 20% Return Underwriting is 0 / 20 because the base case, bear case, and even a reviewed bull case do not clear a 20% return hurdle without relying on terminal multiple support.

Compounder Score note: 54 / 100; Reject under the diagnostic rubric because business-quality evidence is good but not elite and the 20% return case fails at the completed ER2 price anchor.

Business Quality Score note: 54 / 80; runway, management, balance sheet, and cash conversion are strong, but reported ROIC, acquisition-inclusive ROIIC, Data Center AI margin disclosure, and moat durability cap the score.

20% Return UW Score note: 0 / 20; base-case IRR is negative under a conservative 20x EV / NOPAT exit multiple, and the 20% case requires an implausibly high NOPAT CAGR or terminal multiple dependence.

12. Machine-Readable JSON


```
{
  "company": "Advanced Micro Devices, Inc.",
  "ticker": "AMD",
  "exchange": "NASDAQ",
  "analysis_date": "2026-04-29",
  "currency": "USD",
  "current_price": 347.81,
  "market_cap": 567100000000,
  "enterprise_value": 559770000000,
  "business_quality_score": 54,
  "business_quality_score_max": 80,
  "valuation_underwriting_score": 0,
  "valuation_underwriting_score_max": 20,
  "total_score": 54,
  "total_score_max": 100,
  "classification": "Reject",
  "clears_20pct_hurdle": false,
  "base_case_irr": -0.143,
  "bear_case_irr": -0.324,
  "bull_case_irr": -0.017,
  "max_buy_price_for_20pct_irr": 64.7,
  "scorecard_notes": {
    "compounder_score": "54 / 100; Reject under the diagnostic rubric because business-quality evidence is good but not
    "business_quality_score": "54 / 80; runway, management, balance sheet, and cash conversion are strong, but reported
    "twenty_percent_return_underwriting_score": "0 / 20; base-case IRR is negative under a conservative 20x EV / NOPAT
  },
  "category_scores": {
    "incremental_returns_on_capital": {"score": 14, "max": 25},
    "reinvestment_runway": {"score": 16, "max": 20},
    "durability_of_moat": {"score": 9, "max": 15},
    "accounting_and_financial_quality": {"score": 8, "max": 10},
    "capital_allocation": {"score": 7, "max": 10},
    "price_and_20pct_underwriting": {"score": 0, "max": 20}
  },
  "threshold_tests": {
    "total_score_gte_80": false,
    "business_quality_score_gte_65_of_80": false,
    "valuation_score_gte_15_of_20": false,
    "roiic_plus_runway_gte_35_of_45": false,
    "base_case_irr_gte_20": false,
    "bear_case_irr_gte_8_to_10": false,
    "no_hard_vetoes": false,
    "does_not_require_multiple_expansion": false
  },
  "key_metrics": {
    "adjusted_roic_5yr_median": 0.039,
    "roiic_3yr": 0.174,
    "roiic_5yr": 0.124,
    "acquisition_inclusive_roiic_5yr": 0.124,
    "reinvestment_rate_5yr_avg": null,
    "revenue_cagr_5yr": 0.288,
    "nopat_cagr_5yr": 0.219,
    "fcf_cagr_5yr": 0.48,
    "fcf_conversion_5yr": 2.048,
    "share_count_cagr_5yr": 0.063,
    "sbc_to_revenue_latest": 0.047,
    "sbc_to_fcf_latest": 0.297,
    "net_debt_to_ebitda": -0.86,
    "ev_to_nopat": 174.2,
    "fcf_yield": 0.0097
  },
  "hard_vetoes": [
    "The 20% return case requires terminal multiple support or NOPAT growth far above the reviewed bull case.",
    "The current price requires near-flawless Data Center execution and margin expansion relative to the completed ER2
  ],
  "top_positive_evidence": [
    "FY2025 official continuing free cash flow was $5.519B versus normalized NOPAT of $3.214B.",
    "Data Center revenue reached $16.635B in FY2025 and management cited more than 60% annual Data Center growth potent
    "Embedded produced a 36.0% FY2025 operating margin despite a revenue downturn.",
    "Management gate was PASS, with Lisa Su's long-term strategic record and a net cash balance sheet."
  ],
  "top_negative_evidence": [
```

```
    "Reported FY2025 ROIC was only 5.8% and five-year median reported ROIC was 3.9%.",
    "Estimated acquisition-inclusive 5-year ROIIC was only 12.4%.",
    "Moat gate and 2046 test were YELLOW, not PASS.",
    "At $347.81, base-case EPS x 15 was only $118.74 and bull-case EPS x 15 was $237.67."
  ],
  "data_gaps": [
    "AMD does not disclose annual Instinct revenue or annual EPYC revenue.",
    "Data Center AI accelerator gross margin and product-level ROIIC are not separately disclosed.",
    "OpenAI and Meta warrant vesting, accounting treatment, and full dilution impact remain uncertain.",
    "Client and Gaming operating income is not disclosed separately by Client and Gaming components.",
    "This backfill used the completed ER2 price anchor rather than refreshing live market data."
  ],
  "source_confidence": "Medium"
}
```

Appendix CSA - Scoring Audit

Advanced Micro Devices, Inc. - Scoring Audit

PATH_RECEIVED: C:\Users\Admin\OneDrive\GPT\ER 2.0\Advanced Micro Devices, Inc.\appendix_scoring_audit.md

1. Original vs Audited Scores

This audit reviews the diagnostic Compounder Score only. It is not an ER2 gate and does not change the already completed ER2 valuation discipline conclusion. The original scorer was appropriately skeptical on valuation: Valuation / 20% Return Underwriting was 0 / 20, base-case IRR was negative, bear-case IRR was deeply negative, and the reverse DCF required a level of NOPAT growth far above the reviewed bull case. That part is preserved.

The main audit issue is that the Business Quality score still gives too much credit to runway, capital allocation, and accounting quality while the evidence base itself repeatedly states that the most important profit pool is not disclosed at the product level, the customer-warrant economics are unresolved, and acquisition-inclusive return evidence is not yet dream-compounder quality. The original report says reported FY2025 ROIC was only about 5.8%, five-year median ROIC was 3.9%, 3-year acquisition-inclusive ROIIC was 17.4%, and 5-year acquisition-inclusive ROIIC was 12.4%. Those numbers are positive but do not support a mid-50s diagnostic score when the 20% return case is already fully broken.

Item	Original	Audited	Notes
Business Quality Score / 80	54	46	Reduced for management-claim runway reliance, warrant dilution uncertainty, acquisition-return proof gaps, and weak reported ROIC.
Valuation / 20% Return UW Score / 20	0	0	Preserved. Base, bear, and bull cases do not clear 20%; the reverse DCF is not beatable under a sane exit multiple.
Total Score / 100	54	46	Diagnostic score remains Reject.
Classification	Reject	Reject	Score below 70 and hard valuation veto remains.

2. Score Reductions and Reasons

Incremental Returns on Capital is reduced from 14 / 25 to 12 / 25. The original scorer did calculate ROIIC, so Category 1B is not zero. However, 4 / 8 for "ROIIC is high" is generous when the actual five-year acquisition-inclusive ROIIC is only 12.4% and the three-year figure is 17.4%. That is below the >20% long-duration compounder standard and is not cleanly separated between organic return, acquired capital, and AI ramp investment. The audit preserves some credit for positive improvement and strong cash generation, but reduces 1B by 1 point and 1E by 1 point because cash conversion was tested with a five-year average that is flattered by low NOPAT and a one-year continuing-operations FCF bridge. FCF is real, but it does not prove high ROIIC.

Reinvestment Runway is reduced from 16 / 20 to 12 / 20, the largest reduction. The runway is real, but several points rely mainly on management claims and strategic customer announcements. The data package cites management statements about Data Center growing more than 60% annually and AI scaling to tens of billions in annual revenue, but the files also say AMD does not disclose annual Instinct revenue, EPYC revenue, AI accelerator gross margin, customer concentration economics, or contract structures. Under CSA rules, a qualitative runway score that relies mainly on management TAM or ramp claims cannot receive near-full credit. Data Center may prove highly valuable, but the audit cannot score it like a proven high-return flywheel yet.

Durability of Moat is reduced from 9 / 15 to 8 / 15. The moat appendix rates both the moat gate and 2046 test YELLOW. Embedded has strong design-in switching costs, but the largest forward profit pool is Data Center AI, where Nvidia has a stronger software ecosystem, rack-scale platform, networking position, and developer habit. Hyperscaler custom silicon and Arm host CPUs are also meaningful threats. The original moat score was already capped, but one more point is removed because Data Center pricing power remains explicitly unverified.

Accounting and Financial Quality is reduced from 8 / 10 to 7 / 10. Primary-source tie-outs passed, the balance sheet is net cash, and continuing FCF is strong. The reduction is for unresolved customer-warrant accounting and dilution, meaningful FY2025 SBC of $1.638B, and product-level disclosure gaps that matter because the investment case increasingly depends on

AI accelerator economics. A clean revenue tie-out does not fully solve revenue-quality risk when equity-linked customer incentives could become economically large.

Capital Allocation is reduced from 7 / 10 to 7 / 10, unchanged at the category total after review. The original deductions were fair: Xilinx and ZT are strategically coherent but not fully proven on returns; buybacks mostly offset dilution; and incentives are mostly aligned but include a large special CEO award and potential customer-warrant dilution. The audit considered cutting this category, but the management appendix provides enough evidence of long-term strategic execution, net cash conservatism, and coherent R&D-first allocation to preserve 7 points.

Valuation / 20% Return Underwriting remains 0 / 20. At the $347.81 price anchor, the original scorer estimated bear, base, and bull IRRs of -32.4%, -14.3%, and -1.7% using a 20x EV / NOPAT exit multiple. The reverse DCF calculation requires roughly $69.7B of ending NOPAT, or about 85% NOPAT CAGR, to earn 20% annually for five years at 20x exit EV / NOPAT. The completed bull case reaches only about $25.9B of horizon NOPAT. There is no credible way to award valuation points without terminal multiple dependence.

3. Unsupported Claims

The evidence does not support near-full runway credit. The most important unsupported bridge is from "large AI demand and customer partnerships" to "high-return, durable, non-dilutive Data Center economics." OpenAI and Meta commitments show strategic relevance, but the reviewed files do not disclose realized Instinct revenue, customer gross margin, warrant contra-revenue treatment, milestone probability, or full dilution impact.

The evidence does not support treating Data Center AI pricing power as proven. The moat appendix explicitly says Data Center pricing power remains unverified because AMD does not disclose AI accelerator price, margin, customer concentration economics, or contract structures. That claim should cap moat and runway scores.

The evidence does not support a high acquisition-discipline score beyond "strategically coherent." The data package says FY2025 10-K disclosures put total ZT purchase consideration at $4.409B before sale mechanics and related settlement items, while the scorer's acquisition-inclusive ROIIC denominator uses only $2.456B of acquisitions for FY2023-FY2025 and $16.330B of reinvestment for FY2021-FY2025. The direction is not necessarily wrong, but it is not enough to prove acquisition-inclusive returns are high.

4. Calculation Errors

No fatal arithmetic error was found in the original score. The original JSON and markdown agree on 54 / 100, 54 / 80 Business Quality, 0 / 20 Valuation Underwriting, base IRR -14.3%, bear IRR -32.4%, bull IRR -1.7%, and $64.70 maximum buy price for a 20% IRR.

There are calculation limitations rather than clear breaks. The 5-year FCF conversion figure of 204.8% is arithmetically plausible, but it is not a clean measure of compounder quality because AMD's NOPAT base is depressed by amortization, acquisition integration, AI investment, and cyclical margin troughs. The ROIIC calculation is present, but the files do not fully reconcile acquisition cash flow, ZT gross consideration, disposal proceeds, customer warrants, and capitalized R&D into one acquisition-inclusive return bridge. That weakness lowers confidence and caps scores.

5. Missing Data

The most important missing evidence is product-level Data Center economics. AMD does not disclose annual Instinct revenue, annual EPYC revenue, AI accelerator gross margin, customer concentration economics, or contract structure. The completed synthesis correctly uses the disclosed Data Center segment as the control total rather than inventing product dollars. That conservatism is good modeling practice, but the missing data must reduce the diagnostic score.

The second missing evidence item is warrant economics. The synthesis explicitly does not assume full OpenAI and Meta warrant dilution, and says that choice is favorable to AMD. Because the valuation score is already zero even without full warrant dilution, this does not alter the 20% failure. It does reduce confidence in business-quality and accounting-quality scoring because the equity transferred to customers could be economically material.

The third missing evidence item is a fully organic ROIIC bridge. AMD's business is R&D-heavy and acquisition-shaped. The scorer calculated acquisition-inclusive ROIIC, but the audit would need a cleaner split of organic NOPAT growth, acquired NOPAT, working-capital investment, capitalized R&D, ZT purchase economics, and disposal proceeds before awarding higher incremental-return points.

6. Hard-Veto and Threshold Test Audit

The original hard-veto analysis is directionally correct. The 20% return case requires terminal multiple support or NOPAT growth far above the reviewed bull case. Current price requires near-flawless Data Center execution, margin expansion, and limited warrant dilution. Those are hard valuation vetoes under the diagnostic rubric.

The threshold tests also remain failed. Total score is below 80, Business Quality is below 65 / 80, Valuation Underwriting is below 15 / 20, base-case IRR is below 20%, bear-case IRR is not acceptable, and the 20% case does not work without terminal multiple dependence. The audited classification therefore cannot improve above Reject.

7. Final 20% Hurdle Pass/Fail

Final 20% hurdle result: Fail. The base case produces negative IRR under the original 20x EV / NOPAT exit framework, and the synthesis P/E 15 cross-check is even harsher. At $347.81, required horizon EPS at 15x is $23.19. The base case produces $7.92 and the bull case produces $15.84. The required earnings outcome is far beyond the reviewed base case and still above the reviewed bull case.

8. Dream Business / Dream Investment Classification

Audited classification: neither. AMD may become a better business than current consolidated reported ROIC suggests, and management deserves credit for the long transformation. But the audited evidence does not yet prove a dream business because reported ROIC is low, acquisition-inclusive ROIIC is moderate, Data Center AI unit economics are undisclosed, moat durability is YELLOW, and customer-warrant economics are unresolved. It is also not a dream investment because the 20% return underwriting score is zero.

9. Stage 7 Scorecard Notes

Compounder Score note: 46 / 100; Reject under the diagnostic rubric because valuation fails completely and the remaining business-quality evidence is capped by low reported ROIC, moderate acquisition-inclusive ROIIC, undisclosed AI economics, and warrant dilution uncertainty.

Business Quality Score note: 46 / 80; AMD has strong management, net cash, real franchises, and a large AI runway, but the evidence does not yet prove elite organic ROIIC or durable Data Center pricing power.

20% Return UW Score note: 0 / 20; base, bear, and bull cases do not clear the 20% hurdle at the $347.81 anchor without prohibited terminal multiple dependence.

10. Auditor JSON

See `compounder_score_audit.json` in this folder for the machine-readable audit.

Appendix RT - Red Team

Advanced Micro Devices, Inc. - Red Team

Stage 0 anchor: TTM ends 2025-12-27. The forward horizon is 2030-12-27, exactly five calendar years later. This red team attacks the Stage 5 synthesis even though the valuation verdict already fails: at $347.81, AMD needs $23.19 of horizon EPS to clear a 15x P/E gate, while the base case produces $7.92 and the bull case produces $15.84. The purpose here is not to relitigate arithmetic. The synthesis passed its 11 tie-outs. The purpose is to identify the assumption most likely to make even the already-failing base case too optimistic, and to specify the observable signal that should end any residual "wait for proof" thesis.

1. The most fragile assumption

The most fragile assumption is the Data Center segment revenue ramp from $16.635B in FY2025/TTM to $49.090B at the 2030-12-27 horizon, with the first two forward years doing the heavy lifting. In the Data Center appendix, the base case is not a smooth 24.2% CAGR mechanically applied across five years; it is a staged path of $16.635B, $26.616B, $35.932B, $43.118B, $47.430B, and $49.090B. That means the model requires +60% in FY2026 and +35% in FY2027 before decelerating. The synthesis correctly discloses this, but the red team should not let the disclosure become validation. This is the assumption that can break first because it depends on AI accelerator deployment timing, ROCm maturity, hyperscaler utilization, customer concentration, export-control stability, and pricing discipline all moving in the right direction at once.

The fragility is not that AMD lacks demand anchors. The data package documents OpenAI and Meta multi-gigawatt arrangements, MI350/MI450 product timing, and management's ambition for Data Center growth. AMD's official FY2025 release also confirms that Q4 Data Center revenue reached $5.4B, up 39% year over year, and FY2025 Data Center revenue was $16.6B, up 32% year over year. Those are real facts, not hype. The problem is the gap between "strong and accelerating" and the model's need for $26.616B in the next forward year. The Q4 run rate annualizes to about $21.6B before seasonality. The base case needs materially more than that, and it needs the ramp to occur while the segment avoids recurring export controls, customer engineering cost leakage, and second-source pricing pressure.

The Stage 5 sensitivity table shows that a 200 bps reduction in Data Center CAGR lowers horizon EPS by $0.57. That table is useful for small-error testing, but it understates the red-team problem because the plausible downside is not a 200 bps miss. The segment appendix's own bear case is $29.973B of Data Center revenue and 20.5% segment operating margin at the horizon, producing $6.144B of Data Center operating income. The base case is $49.090B of revenue and 28.5% margin, producing $13.991B of Data Center operating income. The difference is $7.847B of segment operating income. Applying the synthesis tax rate of 13% and the base diluted share count of 1,642.3M, that gap is roughly $4.16 of EPS. If Data Center lands at the segment bear case while Client/Gaming and Embedded remain at their base cases, base EPS falls from $7.92 to roughly $3.76. That is the break: the base case is not merely valuation-insufficient; it is highly exposed to one segment's two-year acceleration.

The product-level disclosure gap makes this worse. AMD does not disclose annual EPYC server CPU revenue or annual Instinct accelerator revenue. The segment appendix deliberately avoids fabricating those dollars, which is good modeling discipline, but it leaves the model unable to separately test whether growth is coming from durable CPU share, bursty accelerator shipments, China-specific product sales, or early rack-scale deployments. The Q4 FY2025 release notes selected MI308 China revenue and a reversal of previously reserved inventory charges; the FY2025 release also says full-year results included net inventory and related charges from U.S. export controls on MI308. A clean Q4 margin does not prove clean steady-state economics.

2. Most-likely-to-be-optimistic Stage 5 number

The Stage 5 number most likely to be revised down is base-case Data Center TTM+5y operating income of $13.991B. It is more vulnerable than base revenue alone because it embeds both the $49.090B revenue endpoint and 28.5% segment operating margin. If revenue ramps but margins do not, the earnings result still misses. If margins expand but deployment timing slips, the revenue base is too small. The model requires scale and value capture together.

The synthesis is careful not to assume Nvidia-like economics. It keeps base Data Center margin below the Q4 FY2025 exit-rate margin and below the bull case. That is appropriate, but the red-team objection is that AMD's AI accelerator position may be structurally second-source for longer than the model allows. Nvidia's FY2026 release shows GAAP gross margin around 75% in Q4 and 71.1% for the fiscal year, supported by a full hardware, networking, software, and developer

ecosystem. AMD can grow rapidly as a credible alternative and still face enough customer bargaining power to cap margins. Hyperscalers buying AMD for supply diversification do not have to let AMD earn platform economics. They can use AMD to reduce Nvidia dependence, bargain on price, push custom silicon, and demand customer-specific engineering support.

The base Data Center margin expands from 21.7% in FY2025 to 28.5% at the horizon. That 680 bps improvement is not impossible. It is the right number to red-team because it assumes the ugly parts of the AI ramp get amortized away: ROCm support, rack validation, customer engineering, product-transition inventory, export compliance, and pricing concessions. The segment appendix itself states that Q2 FY2025 Data Center margin was negative because of the MI308/export-control charge and that Q4 was strong. A normalized five-year margin should not be inferred from the best quarter when the business is exposed to product/geography discontinuities and highly concentrated customers.

The share count is a secondary optimistic number. The base synthesis ends at 1,642.3M diluted shares, only 6.3M above the starting 1,636M. That path includes routine stock-based compensation offset by routine repurchases, but it does not include full vesting of customer warrants tied to OpenAI and Meta. The data package states that AMD issued OpenAI a warrant for up to 160M shares and Meta a performance-based warrant for up to 160M shares. Those warrants should not simply be added without considering the revenue that would accompany milestone achievement. But the red-team point is that the base EPS denominator is not stress-loaded for the exact contracts that are supposed to validate the Data Center upside. If 320M incremental shares ultimately vest and are not economically offset by much higher GAAP earnings, a $13.0B net income base would translate to roughly $6.63 of EPS rather than $7.92. In other words, the customer proof that helps revenue may also tax per-share economics.

3. Strongest bear case in its strongest form

The short case is that AMD is priced as if it will become the clean second platform for AI compute, but its actual economics may look more like a strategic supplier used by hyperscalers to discipline Nvidia and diversify supply. The OpenAI and Meta agreements validate demand, but they also concentrate the investment case around a few technically sophisticated buyers with the power to demand pricing, support, and warrant economics. Over the next six to eight quarters, AMD can report impressive Data Center revenue growth while segment margins stall in the low-20s because ROCm friction, rack-scale integration, export controls, and customer-specific engineering absorb much of the scale benefit. If Data Center growth tracks below the FY2026 +60% path or Data Center margin fails to move above 25% as revenue ramps, the model should be rewritten around second-source economics, not platform economics.

This bear case affects Data Center first and consolidated EPS second. Client and Gaming can be fine and still not matter enough. Embedded can recover and still not matter enough. The synthesis itself says Data Center is the dominant swing factor. The current price requires earnings power far above the base model; if the base model is also too optimistic, the investment case has no ER2 margin of safety at all.

4. Cross-examination of management commentary

First, management's Data Center growth language is directionally useful but too easy to over-map into the model. The data package quotes management as saying Data Center revenue can grow "more than 60% annually" and AI can scale to "tens of billions" in annual revenue in 2027. Those are the right phrases to test, not to adopt. The base model does adopt the first-year +60% path, then assumes the ramp remains strong enough for +35% in the second year. The skeptical read is that management is describing a strategic objective around a product cycle and customer pipeline, not guaranteeing high-margin revenue conversion. If the commentary is wrong, the first evidence will be the quarterly Data Center revenue bridge: not enough sequential growth from the $5.4B Q4 FY2025 base, or revenue growth accompanied by margin compression.

Second, management said it did not expect to be "supply limited." That removes one excuse but creates another risk. If supply is not the bottleneck, then slower-than-modeled revenue growth would point to customer qualification timing, software readiness, demand conversion, or pricing. That would be worse than a temporary supply shortage because it would question the strength of AMD's platform pull. The first visible tell would be commentary shifting from deployment scale to pilots, optimization, customer timing, or digestion.

Third, Client and Gaming commentary is mixed and should not be treated as a hidden rescue. The data package has constructive Client commentary around enterprise focus and improving gross margin, but it also includes management's guide that semi-custom SoC annual revenue will decline by a "significant double-digit" percentage. The Client and Gaming appendix handles this by growing Client at 8% and shrinking Gaming at 3% in the base case. That is reasonable. It also means there is no credible path for Client/Gaming to offset a Data Center miss. If Client growth slows after the Windows 11 refresh and Gaming declines longer than expected, it adds pressure, but it is not the central break.

Embedded commentary has the same issue at smaller scale. AMD disclosed $17B of design wins in 2025 and more than $50B since Xilinx, but the data package correctly warns that design wins are not backlog or revenue. The Embedded base case assumes 8% growth and 38% margin. That may be fair, but it is not large enough to rescue consolidated EPS if Data Center disappoints.

5. Stage 5 tie-out reverification

The synthesis's historical tie-out is sound. TTM equals FY2025 under Stage 0. Segment revenue sums to $34.639B. Segment operating income sums to $7.701B, and after $4.007B of unallocated corporate expense, consolidated operating income ties to $3.694B. The red team does not find a historical arithmetic break.

The recency anchor is where the model carries real burden. The synthesis explicitly notes that the Q1 FY2026 revenue guide of about $9.8B annualizes to $39.2B, while the model's base TTM+1 revenue path is $45.239B. The synthesis labels this as explained acceleration rather than drift, and that is fair. But "explained acceleration" is not "verified acceleration." The first forward year has to prove it quickly. If Q1 and Q2 FY2026 do not show Data Center revenue materially above the Q4 FY2025 run-rate trend, the base path becomes stale almost immediately.

The share count path is arithmetically correct but economically incomplete for a red-team view. The model treats buybacks as dilution management, which is conservative compared with assuming heroic repurchases. However, the warrant structure tied to OpenAI and Meta is not a generic footnote. It is attached to the very customer agreements used to support the Data Center ramp. The right monitor is not just ending diluted shares; it is diluted shares plus any warrant milestones versus incremental GAAP operating income from those same customers.

The product CAGR comparison also needs skepticism. The synthesis says Data Center base CAGR of 24.2% is below historical rates and therefore passes growth plausibility. That is mathematically true, but historical CAGR is a weak comparator because the base was much smaller, the AI accelerator market was capacity constrained, and AMD's product disclosure does not separate EPYC from Instinct. The more relevant test is whether AMD can convert the MI350/MI450 roadmap into repeated, high-utilization deployments against Nvidia, custom ASICs, and internal Arm/accelerator strategies.

6. Signal that forces an exit

The exit signal is specific: by the time AMD reports the next two quarters after FY2025, abandon any residual thesis if Data Center revenue is tracking below a $24B FY2026 annual run-rate and Data Center operating margin is below 25.0%, unless management discloses signed deployment timing that mathematically bridges to the $26.616B base-case FY2026 segment revenue path. A practical threshold is Q1 plus Q2 FY2026 Data Center revenue below $11.5B combined, with no credible H2 deployment bridge, or any quarter in which Data Center revenue grows but segment margin remains below 25.0% because pricing, support cost, export-control, or inventory effects recur.

This signal invalidates the thesis because the base case depends on a front-loaded Data Center ramp. The horizon endpoint of $49.090B does not come from slow compounding; it comes from a large FY2026 and FY2027 step-up. If the first two quarters do not make the $26.616B first-year path plausible, then the model has to move toward the Data Center bear case. If revenue grows but margin does not clear 25%, the model also has to move toward second-source economics. Either outcome destroys the only segment that can plausibly move consolidated EPS close to the levels embedded in the stock, even before considering the P/E 15 failure.

7. Red-team conclusion

The thesis does not survive as a BUY candidate, which is already the synthesis verdict. The more important red-team conclusion is that even the base operating model survives only conditionally. The arithmetic is clean, but the model is dominated by a Data Center acceleration that is not yet proven at product-level disclosure granularity. AMD has excellent management, real products, credible customers, and a much stronger strategic position than the pre-Zen company. None of that gives the common shareholder a safe claim on $23.19 of horizon EPS at the current price.

The synthesis should carry the red-team narrative explicitly: AMD is not a broken company; AMD is an exceptional company whose stock is discounting cleaner AI platform economics than the evidence currently proves. The most fragile assumption is not the tax rate, not the buyback path, and not Embedded recovery. It is Data Center converting multi-gigawatt AI partnerships into nearly $50B of horizon revenue at a 28.5% segment margin without heavy warrant dilution or recurring customer-specific margin leakage. Until FY2026 quarterly Data Center revenue and margins validate that path, the correct posture is PASS, not watchful optimism.

Sources checked

- Local synthesis: `C:\Users\Admin\OneDrive\GPT\ER 2.0\Advanced Micro Devices, Inc\appendix_synthesis.md`.
- Local segment appendices: `appendix_B_data_center.md`, `appendix_B_client_gaming.md`, and `appendix_B_embedded.md`.
- Local management and moat appendices: `appendix_management.md` and `appendix_moat.md`.
- Local data package: `data_package.md`, especially Section 6 management commentary and Sections 9-10 corporate actions and modeling guardrails.
- AMD Q4/FY2025 results: https://www.amd.com/en/newsroom/press-releases/2026-2-3-amd-reports-fourth-quarter-and-full-year-2025-fina.html.
- AMD Financial Analyst Day 2025 release: https://www.amd.com/en/newsroom/press-releases/2025-11-11-amd-unveils-strategy-to-lead-the-1-trillion-compu.html.
- NVIDIA Q4/FY2026 results: https://nvidianews.nvidia.com/news/nvidia-announces-financial-results-for-fourth-quarter-and-fiscal-2026.
- Intel Q4/FY2025 results: https://www.intc.com/news-events/press-releases/detail/1759/intel-reports-fourth-quarter-and-full-year-2025-financial.

Appendix D - Data Package (raw historicals)

Advanced Micro Devices, Inc. (AMD) - Data Package

PATH_RECEIVED: C:\Users\Admin\OneDrive\GPT\ER 2.0\Advanced Micro Devices, Inc\data_package.md

Stage 0 anchor: today is 2026-04-29. AMD uses a 52/53-week fiscal year ending on the last Saturday in December; FY2025 ended 2025-12-27. Most recently reported quarter is Q4 FY2025, released 2026-02-03. Q1 FY2026 has not been reported as of this package date. TTM anchor is Q1 FY2025 through Q4 FY2025, equal to FY2025.

Live-source rule: this package used live FMP helper pulls, SEC EDGAR pages, AMD IR/newsroom releases, and FMP transcript endpoint pulls performed in this run. Local data warehouse caches were not read. FMP API keys are not included in any citation or file.

1. Reporting Structure and Taxonomy

Current reportable segments from the FY2025 10-K and Q4/FY2025 press release are Data Center, Client and Gaming, and Embedded. Data Center includes AI accelerators, server CPUs, GPUs/APUs, DPUs, AI NICs, FPGAs, and adaptive SoC products for data centers. Client and Gaming includes desktop/notebook CPUs, APUs, chipsets, discrete GPUs, semi-custom SoC products, and development services. Embedded includes embedded CPUs, GPUs, APUs, FPGAs, system-on-modules, and adaptive SoC products. Source: FY2025 10-K segment note and Q4/FY2025 press release selected corporate data.

Current segment	FY2025 revenue ($B)	% of FY2025 revenue	FY2025 operating income ($B)	Description/source
Data Center	16.635	48.0%	3.603	FY2025 10-K segment table; Q4/FY2025 release
Client and Gaming	14.550	42.0%	2.855	FY2025 10-K segment table; Q4/FY2025 release
Embedded	3.454	10.0%	1.243	FY2025 10-K segment table; Q4/FY2025 release

Taxonomy investigation: reporting changed in Q1 FY2025. AMD combined the prior Client and Gaming reportable segments into one Client and Gaming reportable segment and retrospectively adjusted 2025 presentation for FY2023-FY2024. The 2025 10-K still separately discloses Client and Gaming component revenue inside the combined segment, but only the combined Client and Gaming line is reportable segment operating income. For FY2022 and FY2021, the company had already moved to Data Center, Client, Gaming, and Embedded after the Xilinx acquisition; this package sums Client plus Gaming for directional comparison but flags FY2021-FY2022 as pre-current-reportable-segment presentation. Do not double-count Gaming against Client and Gaming.

Geographic exposure is material and international. FY2025 sales by billing location: United States $11.363B, China including Hong Kong $7.751B, Taiwan $5.186B, Singapore $4.284B, other regions $6.055B, total $34.639B. International sales were 67% of net revenue in FY2025 per AMD MD&A.

2. Consolidated Historicals - Five Years

Metric	FY2021	FY2022	FY2023	FY2024	FY2025	Source
Revenue ($B)	16.43	23.60	22.68	25.79	34.64	FMP annual income / SEC 10-K
Operating income ($B)	3.65	1.26	0.40	1.90	3.69	FMP annual income / SEC 10-K
Operating margin (%)	22.2%	5.4%	1.8%	7.4%	10.7%	calc
Net income ($B)	3.16	1.32	0.85	1.64	4.33	FMP annual income / SEC 10-K
EPS diluted ($)	2.57	0.84	0.53	1.00	2.65	FMP annual income / SEC 10-K
Diluted shares (M)	1,229	1,571	1,625	1,637	1,636	FMP annual income / SEC 10-K
Capex ($B)	0.30	0.45	0.55	0.64	0.97	FMP annual cash flow

Metric	FY2021	FY2022	FY2023	FY2024	FY2025	Source
FCF ($B)	3.22	3.12	1.12	2.40	5.52	FMP annual cash flow; FY2025 continuing FCF per AMD Q4 release
ROIC (%)	73.6%	3.9%	0.7%	3.1%	5.8%	calc - NOPAT / average invested capital

ROIC definition: NOPAT equals GAAP operating income multiplied by (1 - 13% normalized tax rate). Invested capital equals total debt plus total stockholders equity less cash and short-term investments; denominator is average beginning/end invested capital. The 13% tax rate is AMDs 2025 non-GAAP normalized tax rate used in quarterly reconciliations; actual FY2025 GAAP tax benefit is not used because it would overstate operating return economics. FY2025 ROIC calculation: NOPAT $3.694B x 87% = $3.214B; average invested capital approximately $56.177B; ROIC 5.7%.

3. Segment Historicals

Comparable caution: FY2025 10-K provides current-taxonomy segment revenue and operating income for FY2023-FY2025. FY2022 and FY2021 come from AMDs FY2022/FY2023 10-K four-segment tables and are bridged by summing Client plus Gaming for the combined current segment. Use FY2023-FY2025 for clean current-taxonomy trend work.

Data Center

Metric	FY2021	FY2022	FY2023	FY2024	FY2025	Source
Revenue ($B)	3.694	6.043	6.496	12.579	16.635	SEC 10-K segment tables; FMP segment pull check
YoY %	n/a	63.6%	7.5%	93.6%	32.2%	calc
Operating income ($B)	0.991	1.848	1.267	3.482	3.603	SEC 10-K segment tables
Operating margin (%)	26.8%	30.6%	19.5%	27.7%	21.7%	calc

Client and Gaming

Metric	FY2021	FY2022	FY2023	FY2024	FY2025	Source
Revenue ($B)	12.494	13.006	10.863	9.649	14.550	SEC 10-K segment tables; FMP segment pull check
YoY %	n/a	4.1%	-16.5%	-11.2%	50.8%	calc
Operating income ($B)	3.022	2.143	0.925	1.187	2.855	SEC 10-K segment tables
Operating margin (%)	24.2%	16.5%	8.5%	12.3%	19.6%	calc

Embedded

Metric	FY2021	FY2022	FY2023	FY2024	FY2025	Source
Revenue ($B)	0.246	4.552	5.321	3.557	3.454	SEC 10-K segment tables; FMP segment pull check
YoY %	n/a	1750.4%	16.9%	-33.2%	-2.9%	calc
Operating income ($B)	0.044	2.252	2.628	1.421	1.243	SEC 10-K segment tables
Operating margin (%)	17.9%	49.5%	49.4%	39.9%	36.0%	calc

4. Product and Sub-Segment Historicals

Segment	Product / sub-product	What it is	Dollar disclosure?	Last three years revenue + source
Data Center	EPYC server CPUs	Server microprocessors for cloud and enterprise	No separate dollars	AMD discloses Data Center segment revenue, not EPYC dollars; Q4 FY2025 press release says Data Center growth was driven by EPYC processors and Instinct shipments.
Data Center	Instinct data center GPUs / accelerators	AI accelerators, including MI300/MI350/MI400 families	No separate dollars	AMD discloses selected China MI308 revenue in Q4 FY2025 (~$390M) and inventory charges, but not full Instinct annual revenue. Treat Instinct dollars as not separately disclosed.
Client and Gaming	Client	Ryzen desktop/mobile processors, APUs, chipsets	Yes, revenue only	FY2023 $4.651B; FY2024 $7.054B; FY2025 $10.640B. Source: FY2025 10-K segment table and Q4/FY2025 press release.
Client and Gaming	Gaming	Semi-custom SoCs plus Radeon gaming GPUs	Yes, revenue only	FY2023 $6.212B; FY2024 $2.595B; FY2025 $3.910B. Source: FY2025 10-K segment table and Q4/FY2025 press release.
Embedded	Embedded CPUs, GPUs, APUs, FPGAs, SOMs, adaptive SoCs	Adaptive and embedded compute portfolio	Segment dollars only	FY2023 $5.321B; FY2024 $3.557B; FY2025 $3.454B. No finer product dollars found in this run.

Growth-rate-only / non-dollar disclosure table:

Segment	Product	Disclosure method	Last 4 quarters YoY growth / notes	Source
Data Center	EPYC CPUs	Qualitative driver, no dollars	Q4: strong demand; Q3: server share gains; Q2: robust EPYC demand offset MI308 China headwinds; Q1: server CPU share gains. No exact EPYC YoY dollar growth disclosed.	AMD Q1-Q4 FY2025 releases and transcripts
Data Center	Instinct GPUs	Qualitative driver plus selected MI308 China disclosure	Q4 included about $390M MI308 China revenue and release of reserved inventory; Q2 had $800M inventory charge from export controls; full Instinct revenue not disclosed.	AMD Q2/Q4 FY2025 releases
Client and Gaming	Client	Quarterly dollars and YoY disclosed	Q1 $2.294B +68%; Q2 $2.499B +67%; Q3 $2.750B +46%; Q4 $3.097B +34%.	AMD quarterly releases
Client and Gaming	Gaming	Quarterly dollars and YoY disclosed	Q1 $0.647B -30%; Q2 $1.122B +73%; Q3 $1.298B +181%; Q4 $0.843B +50%.	AMD quarterly releases
Embedded	Design wins	Annual design-win dollars, not revenue	Q4 transcript: $17B design wins in 2025 and over $50B since Xilinx acquisition. Treat as backlog-like commercial indicator, not revenue.	FMP Q4 FY2025 transcript

No back-solved product revenue is used. All product-dollar rows above are explicitly disclosed dollars. Data Center AI, EPYC, and Instinct remain open product-dollar gaps.

5. Most Recent Quarterly Trend

Company FY ends the last Saturday in December. Most recently completed fiscal quarter at time of writing: FY2025 Q4, covering 2025-09-28 to 2025-12-27, reported on 2026-02-03 via AMD Q4/FY2025 press release: https://www.amd.com/en/newsroom/press-releases/2026-2-3-amd-reports-fourth-quarter-and-full-year-2025-fina.html.

Data Center Quarterly Trend

Quarter	Revenue ($B)	YoY %	Operating margin (%)	Source
Q1 FY2025	3.674	57.0%	25.4%	AMD press release dated 2025-05-06; FMP quarterly segment pull
Q2 FY2025	3.240	14.0%	-4.8%	AMD press release dated 2025-08-05; FMP quarterly segment pull
Q3 FY2025	4.341	22.0%	24.7%	AMD press release dated 2025-11-04; FMP quarterly segment pull
Q4 FY2025	5.380	39.0%	32.6%	AMD press release dated 2026-02-03; FMP quarterly segment pull

Client and Gaming Quarterly Trend

Quarter	Revenue ($B)	YoY %	Operating margin (%)	Source
Q1 FY2025	2.941	28.0%	16.9%	AMD press release dated 2025-05-06; FMP quarterly segment pull
Q2 FY2025	3.621	69.0%	21.2%	AMD press release dated 2025-08-05; FMP quarterly segment pull
Q3 FY2025	4.048	73.0%	21.4%	AMD press release dated 2025-11-04; FMP quarterly segment pull
Q4 FY2025	3.940	37.0%	18.4%	AMD press release dated 2026-02-03; FMP quarterly segment pull

Embedded Quarterly Trend

Quarter	Revenue ($B)	YoY %	Operating margin (%)	Source
Q1 FY2025	0.823	-3.0%	39.9%	AMD press release dated 2025-05-06; FMP quarterly segment pull
Q2 FY2025	0.824	-4.0%	33.4%	AMD press release dated 2025-08-05; FMP quarterly segment pull
Q3 FY2025	0.857	-8.0%	33.0%	AMD press release dated 2025-11-04; FMP quarterly segment pull
Q4 FY2025	0.950	3.0%	37.6%	AMD press release dated 2026-02-03; FMP quarterly segment pull

Disclosed Client and Gaming Components

Quarter	Client revenue ($B) / YoY	Gaming revenue ($B) / YoY	Source
Q1 FY2025	2.294 / 68.0%	0.647 / -30.0%	AMD press release dated 2025-05-06
Q2 FY2025	2.499 / 67.0%	1.122 / 73.0%	AMD press release dated 2025-08-05
Q3 FY2025	2.750 / 46.0%	1.298 / 181.0%	AMD press release dated 2025-11-04
Q4 FY2025	3.097 / 34.0%	0.843 / 50.0%	AMD press release dated 2026-02-03

6. Management Commentary on Forward Growth

Transcript source control: FMP earning-call-transcript endpoint was fetched live for AMD FY2025 Q1-Q4. Returned dates were Q1 2025-05-06, Q2 2025-08-05, Q3 2025-11-04, and Q4 2026-02-03, matching the quarter chronology. Quotes below are short verbatim snippets with interpretation; use the raw transcript files for full context.

Data Center

Source	Speaker	Verbatim quote	Topic	Interpretation
Q1 FY2025 call / event, 2025-05-06	Lisa Su, Chair and CEO	"We began sampling our next-gen MI350 series"	Data Center	MI350 was already in customer sampling early in FY2025.
Q1 FY2025 call / event, 2025-05-06	Lisa Su, Chair and CEO	"actively working with multiple customers to scale Instinct"	Data Center	Management framed Instinct as moving beyond pilots into scaled deployments.
Q2 FY2025 call / event, 2025-08-05	Lisa Su, Chair and CEO	"We began volume production of the MI350 series ahead of schedule"	Data Center	Supports near-term accelerator ramp assumptions after export-control disruption.
Q2 FY2025 call / event, 2025-08-05	Lisa Su, Chair and CEO	"Helios is expected to deliver up to a 10x generational performance increase"	Data Center	AMD is pushing from chips toward rack-scale systems.
Q3 FY2025 call / event, 2025-11-04	Lisa Su, Chair and CEO	"deploy 6 gigawatts of Instinct GPUs"	Data Center	OpenAI agreement creates a specific multi-year demand anchor.
Q3 FY2025 call / event, 2025-11-04	Lisa Su, Chair and CEO	"first gigawatt of MI450 Series accelerators"	Data Center	First OpenAI deployment is tied to MI450 timing in 2H 2026.
Q4 FY2025 call / event, 2026-02-03	Lisa Su, Chair and CEO	"well-positioned to grow data center segment revenue by more than 60% annually"	Data Center	This is the most explicit segment-level forward growth aspiration.
Q4 FY2025 call / event, 2026-02-03	Lisa Su, Chair and CEO	"scale our AI business to tens of billions in annual revenue in 2027"	Data Center	Downstream models should not treat AI accelerator revenue as disclosed, but should test against this magnitude.
Q4 FY2025 call / event, 2026-02-03	Lisa Su, Chair and CEO	"I do not believe that we will be supply limited"	Data Center	Management was explicit that planned data center AI ramp is not expected to be supply constrained.
Financial Analyst Day call / event, 2025-11-11	AMD press release	"greater than 35% revenue CAGR"	Data Center	Company-level target implies data center must carry much of the growth load.

Client and Gaming

Source	Speaker	Verbatim quote	Topic	Interpretation
Q1 FY2025 call / event, 2025-05-06	Lisa Su, Chair and CEO	"We delivered record client CPU ASP driven by a richer mix of high-end desktop and mobile Ryzen processors"	Client and Gaming	Client rebound was driven by product cycle and mix, not only units.
Q2 FY2025 call / event, 2025-08-05	Lisa Su, Chair and CEO	"record desktop CPU sales"	Client and Gaming	Client strength persisted into Q2 despite macro uncertainty.
Q3 FY2025 call / event, 2025-11-04	Lisa Su, Chair and CEO	"record quarterly sales"	Client and Gaming	Ryzen growth became a full-year share-gain trend.
Q4 FY2025 call / event, 2026-02-03	Lisa Su, Chair and CEO	"we believe we can grow our PC business"	Client and Gaming	Even with a softer PC TAM, AMD expects share/mix to offset market pressure.
Q4 FY2025 call / event, 2026-02-03	Lisa Su, Chair and CEO	"our focus areas are enterprise"	Client and Gaming	Commercial PCs are a forward growth lever, not just consumer channel.
Q4 FY2025 call / event, 2026-02-03	Lisa Su, Chair and CEO	"semi-custom SoC annual revenue to decline by a significant double-digit percentage"	Client and Gaming	Console cycle is a 2026 headwind inside Gaming.
Q4 FY2025 call / event, 2026-02-03	Lisa Su, Chair and CEO	"next-gen Xbox featuring an AMD semi-custom SoC is progressing well"	Client and Gaming	Gaming semi-custom could reverse later, but likely not before 2027.
Q4 FY2025 call / event, 2026-02-03	Jean Hu, CFO	"client on the business gross margin has been improving nicely"	Client and Gaming	Client mix and commercial penetration are margin-relevant.

Embedded

Source	Speaker	Verbatim quote	Topic	Interpretation
Q1 FY2025 call / event, 2025-05-06	Lisa Su, Chair and CEO	"Embedded demand continues to recover gradually"	Embedded	Early FY2025 embedded was still inventory digestion/recovery.
Q1 FY2025 call / event, 2025-05-06	Lisa Su, Chair and CEO	"drive a return to growth in the second half of 2025"	Embedded	Management expected recovery to start in 2H 2025.
Q3 FY2025 call / event, 2025-11-04	Lisa Su, Chair and CEO	"demand environment strengthened across multiple markets"	Embedded	Sequential improvement was visible before Q4 positive YoY growth.
Q4 FY2025 call / event, 2026-02-03	Lisa Su, Chair and CEO	"demand strengthened across several end markets"	Embedded	Q4 confirmed return to YoY growth.
Q4 FY2025 call / event, 2026-02-03	Lisa Su, Chair and CEO	"We closed $17 billion in design wins in 2025"	Embedded	Design-win funnel is the main long-cycle embedded growth indicator.
Q4 FY2025 call / event, 2026-02-03	Lisa Su, Chair and CEO	"more than $50 billion of embedded designs since acquiring Xilinx"	Embedded	Embedded pipeline is larger than near-term revenue; conversion timing remains uncertain.
Q4 FY2025 call / event, 2026-02-03	Lisa Su, Chair and CEO	"return to growth in our Embedded segment"	Embedded	Company expects FY2026 embedded recovery, but no segment-dollar guide disclosed.

Capital and margins

Source	Speaker	Verbatim quote	Topic	Interpretation
Q4 FY2025 call / event, 2026-02-03	Jean Hu, CFO	"revenue to be approximately $9.8 billion plus or minus $300 million"	Capital and margins	Q1 FY2026 guide was above the Q1 FY2025 base and includes about $100M MI308 China sales.
Q4 FY2025 call / event, 2026-02-03	Jean Hu, CFO	"we do expect revenue to expand faster than operating expense"	Capital and margins	Operating leverage is part of the FY2026 earnings bridge.
Q4 FY2025 call / event, 2026-02-03	Jean Hu, CFO	"gross margin should progress issue generation"	Capital and margins	Transcript says "issue" but context indicates each generation; treat as directional margin commentary.

7. Capital Allocation History

Year	Capex ($B)	Capex / revenue	R&D ($B)	R&D / revenue	Buybacks ($B)	Dividends ($B)	M&A ($B)	Source
FY2021	0.301	1.8%	2.845	17.3%	1.999	0.000	0.000	FMP annual cash flow and income statement
FY2022	0.450	1.9%	5.005	21.2%	4.108	0.000	0.822	FMP annual cash flow and income statement
FY2023	0.546	2.4%	5.872	25.9%	1.412	0.000	0.131	FMP annual cash flow and income statement
FY2024	0.636	2.5%	6.456	25.0%	1.590	0.000	0.565	FMP annual cash flow and income statement
FY2025	0.974	2.8%	8.091	23.4%	1.316	0.000	1.760	FMP annual cash flow and income statement

Capital allocation notes: AMD does not pay a common dividend. FY2025 M&A cash flow reflects the ZT Systems acquisition and manufacturing-business sale mechanics; the FY2025 10-K discloses total ZT purchase consideration of $4.409B and sale of the manufacturing business to Sanmina in October 2025. AMD repurchased 12.4M shares and returned $1.3B to shareholders in FY2025 per Q4 transcript / Q4 release; the board also approved a new $6B repurchase authorization on 2025-05-13, raising total authorization to $14B per FY2025 10-K.

8. Balance Sheet Snapshot - FY2025

Item	$B	Source
Cash and equivalents	5.539	FMP annual balance; Q4 release cash plus ST investments tie
Short-term investments	5.013	FMP annual balance
Total debt	3.222	AMD Q4/FY2025 release selected corporate data; FMP balance-sheet debt includes additional categories
Net debt (cash)	-7.330	calc: debt less cash and ST investments
Net debt / adjusted EBITDA	-0.86x	calc; adjusted EBITDA $8.521B from Q4/FY2025 release
Goodwill + intangibles	41.831	FMP annual balance sheet
Tangible book value	21.168	calc: equity $62.999B less goodwill/intangibles

9. Recent Corporate Actions - Last 12 Months

- 2025-05-13: board approved new $6B share repurchase program, in addition to existing authorization, raising total authorization to $14B. Source: FY2025 10-K.
- 2025-08-05: Q2 FY2025 release disclosed an $800M inventory and related charge tied to U.S. export restrictions on AMD Instinct MI308 data center GPU products.
- 2025-10-06: AMD and OpenAI announced a 6GW multi-year, multi-generation Instinct GPU partnership; AMD issued OpenAI a warrant for up to 160M shares tied to milestones. Source: AMD/OpenAI press release and 8-K.
- 2025-10-27: AMD completed the sale of ZT Systems data center infrastructure manufacturing business to Sanmina; AMD retained rack-scale systems design capability. Source: FY2025 10-K / Sanmina release.
- 2025-11-11: AMD Financial Analyst Day set long-term targets including greater than 35% revenue CAGR and more than $20 non-GAAP EPS over the strategic horizon. Source: AMD Financial Analyst Day press release.
- 2026-02-03: AMD reported FY2025 revenue $34.639B, GAAP operating income $3.694B, and FY2025 free cash flow from continuing operations $5.519B. Source: Q4/FY2025 release.
- 2026-02-24: AMD and Meta announced up to 6GW of Instinct GPU deployments across multiple generations; AMD issued Meta a performance-based warrant for up to 160M shares. Source: AMD/Meta press release.

No CEO or CFO transition was found in the last 12 months. No dividend initiation was found. No separately disclosed large legal settlement was found in the fetched sources; Q3 non-GAAP reconciliation includes a $12M loss contingency on a legal matter, not material for this package.

10. Open Data Gaps and Use Instructions

- Data Center product revenue is not separately disclosed for EPYC vs Instinct; do not fabricate AI accelerator dollars. Only selected MI308 China revenue and inventory charges are disclosed.
- Instinct annual revenue is not disclosed. Management gave forward aspiration to scale AI to tens of billions in annual revenue in 2027, but historical base remains undisclosed.
- Client and Gaming operating income is disclosed only for the combined current segment, not for Client and Gaming components. Component revenue is disclosed; component margins are not.
- Embedded has segment revenue and operating income; no finer product-dollar split was found for FPGAs, adaptive SoCs, embedded CPUs, or SOMs.
- FY2021-FY2022 segment data precede the Q1 FY2025 Client/Gaming combination. Client+Gaming sums are shown for directional continuity, but FY2023-FY2025 are the clean current-taxonomy period.
- FMP FY2025 free cash flow includes discontinued operations; AMDs official continuing-operations free cash flow of $5.519B is used for FY2025 cash-flow quality and capital allocation context.
- Quarterly segment operating income was sourced from AMD selected corporate data/10-Qs; Q4 segment operating income equals the Q4/FY2025 release and 10-K annual bridge.

Downstream Modeling Guardrails

The model base should use FY2025 consolidated revenue of $34.639B, GAAP operating income of $3.694B, and the three reportable-segment revenue base of Data Center $16.635B, Client and Gaming $14.550B, and Embedded $3.454B. These are official-source control totals from AMDs FY2025 10-K and Q4/FY2025 release, and they tie to the live FMP annual income statement and FMP revenue-product-segmentation pull. If a downstream model starts from a different FY2025 segment base, the model should be treated as out of control until reconciled.

For Data Center, keep the entire segment together unless the downstream subagent explicitly creates a scenario split for server CPUs and AI accelerators. That split must be a model assumption, not a historical fact. The historical package supports statements such as "Data Center grew 32% in FY2025" and "Q4 FY2025 Data Center grew 39% YoY"; it does not support statements such as "Instinct was $XB in FY2025." The Q4 transcript and Financial Analyst Day materials support ambitious forward growth commentary, but those comments should be mapped to explicit revenue scenarios rather than inserted as historical actuals.

For Client and Gaming, revenue components can be modeled separately because AMD discloses Client and Gaming component revenue. Operating income should remain combined unless the model explicitly allocates margin using a documented assumption. The 2026 setup is mixed: Client has positive share and enterprise-PC commentary, while Gaming has a semi-custom console-cycle headwind. A model that applies one growth rate to the entire combined segment should explain why the Client tailwind and Gaming headwind are being netted that way.

For Embedded, the most reliable historical signal is segment revenue and operating income. Design wins are important but should not be treated as backlog or revenue. The Q4 transcript says $17B of 2025 design wins and more than $50B since Xilinx, but the conversion timing and margin profile are not disclosed. FY2026 growth assumptions should therefore be built from recovery in end-market demand and design-win conversion, not from direct conversion of design-win dollars into revenue.

For cash flow, use AMDs continuing-operations free cash flow when comparing FY2025 with FY2024 because the ZT manufacturing divestiture created discontinued-operations cash-flow noise. The live FMP cash-flow row is still retained in raw files, but the package uses the companys Q4/FY2025 release control number of $5.519B for FY2025 FCF. For debt, use the Q4 release total debt of $3.222B for leverage, while noting that FMP balance-sheet total debt includes additional balance-sheet categories and is higher.

Source Quality Notes

Primary-source hierarchy was applied as follows. Consolidated statements came from live FMP helper pulls and were tied to AMDs SEC-filed FY2025 10-K and Q4/FY2025 release. Segment annual data came from SEC filings, with FMP revenue-product-segmentation used as a cross-check for current additive segment revenue. Quarterly segment data came from AMD quarterly press releases and selected corporate data tables. Transcript quotes came from live FMP transcript endpoint files with returned dates checked against expected call dates. Corporate actions came from AMD IR releases, SEC filings, and Sanmina transaction confirmation where Sanmina was the counterparty announcing completion.

The highest-risk taxonomy issue is the Q1 FY2025 combination of Client and Gaming into one reportable segment. AMD provides restated FY2023-FY2024 current-taxonomy values in the FY2025 10-K, but FY2021-FY2022 remain older presentation years. This package bridges FY2021-FY2022 by summing old Client plus old Gaming only to provide trend context. It should not be used to infer that AMD managed the business in FY2021-FY2022 under the exact same current reportable-segment structure.

The highest-risk product disclosure issue is Instinct. AMDs public comments make clear that Instinct is strategically central, but the company does not disclose total Instinct revenue. The Q4 FY2025 release discloses approximately $390M of MI308 China revenue in Q4 and explains the inventory reserve release, but that number is a special item and not a proxy for total AI GPU revenue. The appropriate downstream treatment is to build Instinct explicitly as an assumption inside Data Center and mark it as model-derived.

Source Manifest

- FMP annual income: FMP stable income-statement, symbol=AMD, period=annual, limit=8; raw/data_package/fmp_income_annual.json

- FMP annual cash flow: FMP stable cash-flow-statement, symbol=AMD, period=annual, limit=8; raw/data_package/fmp_cashflow_annual.json
- FMP annual balance: FMP stable balance-sheet-statement, symbol=AMD, period=annual, limit=8; raw/data_package/fmp_balance_annual.json
- FMP segment annual: FMP stable revenue-product-segmentation, symbol=AMD, period=annual; raw/data_package/fmp_revenue_product_segmentation_annual.json
- FMP segment quarter: FMP stable revenue-product-segmentation, symbol=AMD, period=quarter; raw/data_package/fmp_revenue_product_segmentation_quarter.json
- FY2025 10-K: https://www.sec.gov/Archives/edgar/data/2488/000000248826000018/amd-20251227.htm
- FY2024 10-K: https://www.sec.gov/Archives/edgar/data/2488/000000248825000012/amd-20241228.htm
- FY2022 10-K: https://www.sec.gov/Archives/edgar/data/2488/000000248823000047/amd-20221231.htm
- Q1 PR: https://www.amd.com/en/newsroom/press-releases/2025-5-6-amd-reports-first-quarter-2025-financial-results.html
- Q2 PR: https://www.amd.com/en/newsroom/press-releases/2025-8-5-amd-reports-second-quarter-2025-financial-results.html
- Q3 PR: https://www.amd.com/en/newsroom/press-releases/2025-11-4-amd-reports-third-quarter-2025-financial-results.html
- Q4/FY PR: https://www.amd.com/en/newsroom/press-releases/2026-2-3-amd-reports-fourth-quarter-and-full-year-2025-fina.html
- Financial Analyst Day PR: https://www.amd.com/en/newsroom/press-releases/2025-11-11-amd-unveils-strategy-to-lead-the-1-trillion-compu.html
- OpenAI partnership: https://ir.amd.com/news-events/press-releases/detail/1260/amd-and-openai-announce-strategic-partnership-to-deploy-6-gigawatts-of-amd-gpus
- Meta partnership: https://ir.amd.com/news-events/press-releases/detail/1279/amd-and-meta-announce-expanded-strategic-partnership-to-deploy-6-gigawatts-of-amd-gpus
- FMP transcript FY2025 Q1: stable earning-call-transcript, symbol=AMD, year=2025, quarter=1; returned date 2025-05-06; raw/data_package/fmp_transcript_2025_q1.json
- FMP transcript FY2025 Q2: stable earning-call-transcript, symbol=AMD, year=2025, quarter=2; returned date 2025-08-05; raw/data_package/fmp_transcript_2025_q2.json
- FMP transcript FY2025 Q3: stable earning-call-transcript, symbol=AMD, year=2025, quarter=3; returned date 2025-11-04; raw/data_package/fmp_transcript_2025_q3.json
- FMP transcript FY2025 Q4: stable earning-call-transcript, symbol=AMD, year=2025, quarter=4; returned date 2026-02-03; raw/data_package/fmp_transcript_2025_q4.json

Self-Verification Log

- Largest-segment most recent quarterly growth: live AMD Q4/FY2025 press release confirms Data Center segment revenue in Q4 FY2025 was $5.4B, up 39% YoY.
- Management quote check: live FMP FY2025 Q4 transcript returned date 2026-02-03 and contains the phrase "well-positioned to grow data center segment revenue by more than 60% annually".
- Full-year revenue check: live FMP annual income-statement row for FY2025 reports revenue $34.639B, tying AMDs Q4/FY2025 press release and FY2025 10-K.