Nov. 3, 2026

AMD Third Quarter 2026 Financial Results
Prepared Remarks

These remarks contain forward-looking statements based on current beliefs, assumptions, and expectations, including forward-looking statements regarding financial projections, business and industry trends, that speak only as of today and as such, involve risks and uncertainties that could cause actual results to differ materially from our current expectations. Please refer to the cautionary statement in our press release for more information on factors that could cause actual results to differ materially.

Dr. Lisa Su, AMD Chair and CEO:

The third quarter was the largest and most profitable in our history as we began ramping our next-generation data center products.

Revenue increased 45% year over year to $13.4 billion driven by record EPYC and Instinct sales and strong Embedded growth.

Earnings per share grew even faster, increasing 74% to $2.09.

We executed very well in the quarter, with Helios production ramping as planned and EPYC adoption expanding across cloud, enterprise and superintelligence workloads.

Data Center revenue growth has accelerated sequentially every quarter this year, with the segment contributing 64% of total revenue in the third quarter, up from 47% just one year ago.

We expect that acceleration to continue in the fourth quarter, driven by increased demand for inference and agentic workloads and the ramp of our newest products.

Data Center Segment

Data Center revenue grew 97% year over year to a record $8.5 billion.

In server, both cloud and enterprise sales more than doubled year over year.

We significantly exceeded our previous outlook of more than 80% growth as we gained revenue share.

Customers deployed EPYC across more applications, with agentic workloads contributing significant growth.

In cloud, hyperscalers deployed substantial volumes of EPYC servers in the quarter to run more internal services and expand their public-cloud offerings.

“Turin” is now the most widely deployed EPYC generation in the public cloud, with more than 540 instance types available globally from the largest providers.

Cloud providers also added significant EPYC capacity in the quarter for personal and enterprise agents, which are quickly becoming the next killer app for CPUs in the cloud.

Meta’s Muse and OpenAI’s dots illustrate how quickly personal agents are growing, with more than 10 million people activating Muse agents in the first six weeks.

Enterprises are moving fast as well, putting agents to work across software development, supply chain and IT operations.

These agents work continuously and in parallel across applications and enterprise systems, adding to accelerator demand while creating a substantial increase in CPU demand.

Adoption is just starting, and we expect demand to grow sharply as agents are deployed more broadly and take on more complex work.

In head nodes, we won major deployments with several of the largest model developers and cloud providers in the quarter, which will ramp over the coming quarters.

EPYC’s leadership performance and mature ecosystem – proven at scale – have made it the host CPU of choice for custom accelerators.

As one example, OpenAI selected “Turin” EPYC processors as the host CPUs for its first custom accelerator platforms.

In enterprise, EPYC server sell-through was a record, with strong demand from new and existing customers.

We added dozens of first-time EPYC customers in the quarter, including large financial services, communications, retail and healthcare companies.

Existing customers are also making EPYC their primary compute platform, extending adoption from specific workloads to a much larger share of their overall infrastructure.

We began revenue shipments of EPYC “Venice” processors in the third quarter, and the first OEM systems are on track for availability this quarter.

Early demand is the strongest we have seen for any new EPYC generation.

Our OEM and ODM partners have committed more than 240 “Venice” platforms, 20% more than “Turin,” and cloud providers are deploying “Venice” this quarter ahead of instance rollouts planned for early 2027.

We recently launched Helion, a “Venice” reference design that enables turnkey EPYC rack-scale systems for agentic infrastructure.

Customers can configure Helion for their workloads, from all-CPU racks for hyperscalers to enterprise systems combining “Venice” and Instinct GPUs that can run agents and models together.

In the all-CPU configuration, a single Helion rack has up to 144 “Venice” processors with nearly 74,000 threads, delivering the highest agent density in the industry.

Every major server OEM and ODM plans to offer Helion, with the first systems shipping this quarter.

Looking more broadly, demand for server compute is accelerating across cloud, enterprise, agentic and super intelligence infrastructure.

EPYC is the industry’s broadest portfolio spanning all of these markets, with per-socket leadership for cloud and agents, per-core leadership for enterprise and head nodes, and memory and I/O options tuned to each workload.

“Venice” builds on that leadership, with significantly better performance per watt, performance per core, and performance per dollar than any other x86 or Arm server CPU, positioning us to gain share and significantly grow our server business as the TAM grows to approximately $220 billion by 2030.

Data Center AI

Frontier labs, hyperscalers and cloud providers expanded Instinct deployments for inference and training, driving record revenue.

We started shipping MI455X in volume in the quarter, and lead customers including Anthropic, Meta, Microsoft, OpenAI and Oracle brought their first Helios racks online quickly.

Customers are now running their latest models and workloads on Helios, and we are seeing very strong performance with the initial deployments.

We expect Helios deployments to increase substantially in the fourth quarter as production continues to ramp and HPE, Dell, Lenovo, Supermicro, and other OEM and ODM partners begin shipping systems.

We are very pleased with the overall performance we are seeing at the rack level. MI455X leads in compute FLOPs, memory capacity and scale-out bandwidth.

And as Helios comes to market, it delivers an even larger lead in memory bandwidth and higher scale-up bandwidth than the competition.

We are working closely with our manufacturing and OEM partners to scale production as customers expand their deployments.

In supercomputing, EuroHPC selected Instinct MI430X GPUs and 6th Gen EPYC processors for LUMI-AI, Europe’s next flagship supercomputer planned to come online in the second half of 2027.

MI430X delivers nearly nine times the double-precision performance of the competition, which is critical for scientific simulations.

As a result, MI430X has now been selected to power next-generation systems around the world, and we are well positioned to win additional designs as governments and research centers plan their next major investments.

In software, we launched ROCm 10, with major enhancements for production deployments at scale, including faster startup across large GPU clusters and greater resilience for long-running workloads.

ROCm 10 also adds agentic performance tuning with ROCm.AI, which profiles workloads, identifies bottlenecks, and applies and validates improvements – turning weeks of engineering optimization into hours.

Customers are also using their own agents to optimize models on our open software stack.

For example, Anthropic highlighted that one engineer used Claude to bring up their latest model on Instinct MI355X racks over a single weekend with strong performance.

Our work with frontier labs is already enabling their models to handle complex software development and performance tuning for Instinct, and as those capabilities continue to improve the benefits extend to everyone in the AMD ecosystem.

We announced two acquisitions in the quarter to strengthen our superintelligence portfolio.

Taalas brings a highly experienced engineering team and specialized silicon built for low-latency, efficient inference.

We plan to integrate their technology across our solutions portfolio to address the growing demand for low-latency superintelligence applications.

World Labs, led by co-founder and CEO Dr. Fei-Fei Li, brings a world-class research team with expertise in spatial intelligence and models that understand the physical world, giving us deeper insight into how frontier models and workloads are evolving.

Fei-Fei will join AMD as chief scientist, leading our model research and helping shape our future hardware, software and systems roadmaps.

We expect both acquisitions to close by the end of the year.

Looking ahead, our focus is on execution at scale.

The ramps of “Venice,” the MI450 series and Helios are still early, and all three are on track to grow significantly.

We now expect Data Center segment revenue growth to accelerate in the fourth quarter and full-year segment revenue to more than double from 2025.

Our visibility into 2027 also continues to strengthen.

Last quarter we said we expected annual server revenue to grow more than 70% and Data Center segment revenue to more than double for the year.

Since then, customers have expanded their EPYC and Instinct deployment plans and we have secured more supply to support that growth.

We now expect server revenue to double year over year in 2027 and data center AI revenue to grow even faster than we projected 90 days ago.

As a result, we expect Data Center to be more than 80% of our total revenue in 2027 and to drive the majority of our revenue and earnings growth.

Looking further out, customer engagements for MI500 continue to progress very well.

Interest is very high, as it is another significant leap forward in performance.

We are working very closely with multiple frontier labs to optimize MI500 for their workloads as they plan multigigawatt deployments.

Client and Gaming Segment

Segment revenue decreased 8% year over year to $3.7 billion.

In client, revenue increased 7% year over year to $2.9 billion, driven by strong sales of our Ryzen mobile processors.

Commercial is becoming a larger part of our client business, as large enterprises adopt Ryzen PRO PCs across their employee base.

Eight of the 10 largest banks in the world and seven of the 10 largest healthcare and automotive companies now use Ryzen PRO PCs.

OEMs continue expanding their AMD-based commercial lineups, from everyday business PCs to high-end workstations.

HP, Lenovo and others launched the first commercial notebooks and small form factor desktops featuring our newest Ryzen AI Max PRO 400 Series processors.

These systems deliver workstation-class performance and run superintelligence models with up to 300 billion parameters locally.

Looking at the market, PC sell-through slowed late in the quarter as higher memory and component costs weighed on demand.

We are managing the business accordingly, aligning sell-in with sell-through to keep customer and channel inventory healthy.

In 2027, we are planning for PC market unit shipments to decline by a double-digit percentage.

In this environment, we expect to outperform the market as we deliver our Ryzen roadmap and grow our commercial footprint.

In gaming, revenue declined 40% year over year to $773 million, primarily due to lower semi-custom sales as expected at this stage of the console cycle.

Gaming graphics revenue declined year over year while Radeon GPU channel sell-through increased sequentially.

Beyond gaming, Radeon AI PRO revenue grew significantly from the prior quarter as more developers chose AMD to run and fine-tune models locally.

We expect strong growth for Radeon AI PRO to continue, driven by our leadership performance-per-dollar and an open software stack optimized for the latest open-source models.

Embedded Segment

Revenue increased 29% year over year to $1.1 billion, the fourth consecutive quarter of accelerating growth.

Demand was very strong across all of our end markets, led by strong growth with aerospace and defense, automotive, communications and industrial customers.

Data center buildouts are also driving higher demand for our embedded CPUs and adaptive SoCs in control plane and optical networking applications.

We remain on track to close more than $18 billion in design wins in 2026, our third consecutive record year.

Adaptive design wins have increased every year since we acquired Xilinx, with embedded x86 and semi-custom wins growing even faster.

As these wins ramp over the next few years, we expect our embedded business to deliver strong double-digit annual revenue growth and continue outgrowing the market.

Close

We delivered an excellent third quarter and expect an even stronger fourth quarter as Data Center and Embedded segment growth more than offsets lower client CPU sales.

We enter 2027 with the strongest product portfolio in our history and significant customer momentum.

We have worked across our supply chain to secure the wafers, substrates, advanced packaging, memory and other key components needed to support much higher growth next year.

Based on our current visibility, we now expect revenue to grow more than 80% in 2027, which is more than twice the long-term financial model we outlined at our financial analyst day last November.

Data Center will be the largest contributor, driven by expanding Helios deployments and continued EPYC share gains.

We also expect Embedded growth to accelerate next year driven by adoption of our adaptive and x86 solutions.

Beyond 2027, we are increasing R&D investment across our silicon, software and systems roadmaps and working closely with our largest customers to support their ambitious infrastructure plans for the next three to five years.

We are also deepening our collaboration with our supply chain partners and investing alongside them to secure the long-term capacity our growth will require.

Taken together, these give us strong confidence in our ability to sustain exceptional revenue growth for years to come… with earnings growing even faster.

Jean Hu, AMD Executive Vice President, CFO and Treasurer:

I'll start with a review of our financial results and then provide our outlook for the fourth quarter of fiscal 2026.

We delivered an outstanding third quarter.

Revenue increased 45% year over year to a record $13.4 billion, exceeding the high end of our guidance, driven by our Data Center segment, which nearly doubled year over year.

Revenue increased 16% sequentially, driven by strong growth in the Data Center and Embedded segments, partially offset by a decline in the Client and Gaming segment.

Our Data Center segment accounted for approximately 64% of total revenue, up from 58% last quarter, underscoring the continued shift in our business mix and its increasing contribution to AMD's growth.

Gross margin was 56%, up 190 basis points year over year, reflecting a richer product mix.

Operating expenses were $3.6 billion, an increase of 29% year over year, as we continued to invest in R&D to expand our data center silicon, systems and software capabilities.

Operating income grew 75% year over year to a record $3.9 billion, significantly outpacing our revenue growth. Operating margin was 29%, expanding more than 5 percentage points year over year, underscoring the operating leverage in the business as we scale.

In the third quarter, we began recognizing the accounting impact of warrants issued under our previously announced customer agreements as we shipped MI450 Series GPUs. Our GAAP diluted earnings per share were $1.46, which included a $1.1 billion noncash charge in other income and expense, primarily reflecting the increase in our share price and changes in the estimated probability of meeting the warrant conditions.

Going forward, these quarterly noncash fair value adjustments are expected to increase as we and our customers progress under these agreements. The amounts will vary each quarter based on AMD's stock price and other valuation assumptions. These noncash adjustments will not affect our non-GAAP results or free cash flow.

Our non-GAAP diluted earnings per share were $2.09.

Now turning to our reportable segments.

Starting with Data Center …

Data Center segment revenue was a record $8.5 billion, up 97% year over year, driven by record server revenue on strong demand for our 5th Gen EPYC processors and by record data center AI revenue, highlighted by the launch of our AMD Instinct MI450 Series GPUs.

On a sequential basis, Data Center revenue increased 27 percent, primarily driven by initial MI450 Series sales and strong server revenue growth.

The Data Center segment operating income was $2.9 billion, up 168% year over year, and operating margin increased to 34% compared to 25% a year ago, demonstrating strong operating leverage of data center business.

Client and Gaming segment revenue was $3.7 billion, down 8% year over year due to lower semi-custom and graphics sales, partially offset by higher demand for the latest generation Ryzen mobile processors.

In the Client business, revenue was $2.9 billion, up 7% year over year, driven by higher sales of Ryzen mobile processors. Sequentially, Client revenue was down 4% as higher memory costs weighed on PC unit demand, partially offset by a richer mix of higher-end processors.

Gaming revenue was $773 million, down 40% year over year due to lower semi-custom revenue and lower graphics demand. Sequentially, Gaming revenue was essentially flat.

Client and Gaming segment operating income was $573 million, or 15% of revenue, compared to $867 million, or 21%, a year ago.

Embedded segment revenue was $1.1 billion, up 29% year over year and 14% sequentially, as demand strengthened across several end markets, led by growth in aerospace and defense, automotive, communications, and industrial end markets.

Embedded segment operating income was $453 million, or 41% of revenue, compared to $283 million, or 33%, a year ago.

Turning to the Balance Sheet and Cash Flow:

During the quarter, inventory increased approximately $3 billion sequentially to $11.5 billion to fulfill our strong data center demand. We generated $1.2 billion in cash from operating activities of continuing operations. Free cash flow was $554 million, reflecting increased working capital investment and ongoing capital expenditures to secure equipment and capacity for the rapid growth of our Data Center business, including back-end capacity dedicated to support our Data Center demand. We expect to continue to invest in inventory and capacity to support our strong growth in 2027.

At the end of the quarter, cash, cash equivalents and short-term investments were $15.8 billion. Our total debt was $7.1 billion, including the $4.75 billion senior notes offering we completed in August.

Now turning to our fourth-quarter 2026 outlook:

We expect revenue to be approximately $17 billion plus or minus 3%. The midpoint of our guidance represents approximately 66% year-over-year revenue growth primarily driven by the acceleration of our Data Center segment.

Sequentially, we expect revenue to grow by approximately 27%, driven by:

- A sharp increase in our Data Center segment of approximately 50% growth.
- A decline in our Client and Gaming segment of greater than 20%.
- And double-digit growth in our Embedded segment.

In addition, we expect fourth quarter non-GAAP gross margin to be approximately 55%, and we expect non-GAAP operating expenses to be approximately $3.75 billion.

We expect non-GAAP net interest and other expenses to be a gain of approximately $4 million.

We expect our non-GAAP effective tax rate to be 13% and diluted share count is expected to be approximately 1.67 billion shares including the estimated impact of closing the Taalas and World Labs acquisitions in the fourth quarter. We are not expecting our fourth quarter share count to be impacted by customer warrants.

In closing, the third quarter was an outstanding quarter for AMD, with record revenue, significant operating margin expansion and earnings growth, reflecting disciplined execution across the business. As our Data Center business continues to scale, we are very well positioned to drive increased profitability and long-term shareholder value.