OPERATOR: Greetings. Welcome to the AMD Third Quarter 2026 Conference Call. At this time, all participants are in a listen-only mode. A question-and-answer session will follow the formal presentation. If anyone should require operator assistance during the conference, please press star-0 on your telephone keypad. And please note that this conference is being recorded. I will now turn the conference over to Matt Ramsay, VP, Financial Strategy and IR. Thank you, Matt. You may begin.

MATT RAMSAY: Thank you, and welcome to AMD's Third Quarter 2026 Financial Results Conference Call. By now, you should have had the opportunity to review a copy of our earnings press release and the accompanying slides. If you have not had the chance to review these materials, they can be found on the Investor Relations page of AMD.com.

We will refer primarily to non-GAAP financial measures during today's call. The full non-GAAP to GAAP reconciliations are available in today's press release and slides posted on our website. Participants on today's conference call include Dr. Lisa Su, our Chair and CEO, and Jean Hu, Executive Vice President, CFO, and Treasurer. This is a live call and will be replayed via webcast on our website.

Before we begin, I would like to note that AMD will participate in the following events with the financial community this quarter-- the UBS Technology and AI Conference on Tuesday, December 1, and Barclays 24th Annual Global Technology Conference on Wednesday, December 9.

Today's discussion contains forward-looking statements based on our 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 the factors that could cause actual results to differ materially. With that, I will hand the call over to Lisa.

LISA SU: Thank you, Matt, and good afternoon to all those listening today. 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.

Turning to our segments, 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 health care 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 superintelligence 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.

Turning to our Data Center AI business, frontier labs, hyperscalers, and cloud providers expanded Instinct deployments for inferencing 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 multi-gigawatt deployments.

Turning to Client and Gaming, 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 health care 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 choose 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.

Turning to our 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 build-outs 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 of 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 up over the next few years, we expect our Embedded business to deliver strong double-digit annual revenue growth and continue outgrowing the market.

In summary, 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. Now I will turn the call over to Jean to provide additional color on our third quarter results. Jean?

JEAN HU: Thank you, Lisa, and good afternoon, everyone. 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 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 continue 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 agreement as we shipped MI450 Series GPUs. Our GAAP diluted earnings per share were $1.46, which included $1.1 billion non-cash charge in other income and expense, primarily reflecting the increase in our share price and the changes in the estimated probability of meeting the warrant conditions.

Going forward, this quarterly non-cash fair value adjustment are expect to increase as we and our customers progress under this agreement. The amounts will vary each quarter based on AMD's stock price and the other valuation assumptions. These non-cash adjustment 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 segment, starting with the Data Center. Data Center segment revenue was a record of $8.5 billion, up 97% year over year, driven by record server revenue on strong demand for our fifth-generation 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%, 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 cost weighted on PC unit demand, partially offset by a richer mix of high-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, communication, and the 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 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 backend 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 $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 a double-digit growth in our Embedded segment.

In addition, we expect the 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 expense 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 the World Labs acquisition in the fourth quarter. We're 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. With that, I'll turn it back to Matt for the Q&A session.

MATT RAMSAY: Thank you very much.

OPERATOR: Thank you. We will now be conducting a question-and-answer session. We ask that you please limit yourself to one question and one follow-up. If you would like to ask a question, please press star-1 on your telephone keypad. A confirmation tone will indicate that your line is in the question queue. You may press star-2 if you'd like to remove your question from the queue.

For participants using speaker equipment, it may be necessary to pick up your handset before pressing the star keys. One moment, please, while we poll for questions. And the first question comes from the line of Tom O'Malley with Barclays. Please proceed with your question.

TOM O'MALLEY: Hey, good evening, and thanks for taking my question. Congrats on the great results. I guess, Lisa, let's start off with the fact that you're giving some more color into 2027. Looks like revenue's growing 80%, but the Data Center component of that is growing much faster, where Data Center is going to be greater than 80% of your business next year.

So maybe, can you talk about, first and foremost, is this backed up by supply, whether that be wafers, packaging, substrate, memory? And then, if you look at the demand profile behind that, would demand outstrip that supply if you were able to get some more? Thank you.

LISA SU: Yeah, sure, Tom. Thanks for the question. So look, we have a lot of momentum going into the fourth quarter and into 2027. We do see a situation where there is very strong demand. So we see demand across the Data Center business, across our CPU business, as well as our Data Center AI business.

And when you look at the rate and pace of that growth, we've been planning for it for the last several years. So we've absolutely spent a significant amount of time with our customers, planning what their needs are in terms of their deployments and their deployment schedules, and also with our supply chain partners to ensure that we have the wafers, the substrates, the advanced packaging, the memory, all of the other components that go along with it.

It is true that demand is higher than supply, so we continue to work with our supply chain partners to increase supply. And I do believe that there is a good opportunity for us to do that, just given the strong relationships we have throughout the ecosystem. But sitting where we are today, to be able to say that the overall company will grow greater than 80%, especially the server business will approximately double from a fairly high point in 2026.

And then the Data Center AI business, we continue to have strong visibility, and we see strong demand from our customers, and we continue to see also good follow-through on all of the other aspects of the infrastructure that are needed. So overall, we're very positive about the Data Center growth into 2027.

TOM O'MALLEY: Super helpful. And then, just as a follow-up, on the Client side, you're saying down greater than 20% in the fourth quarter. That's kind of in line with what we've heard from others so far in the market. But if you look at '27, I think something that's difficult is that normally, you see seasonality down in Q1, and then you wouldn't get a recovery till Q2, Q3. So maybe talk about what the longevity of this correction looks like and why you've been able to outperform a bit up until this point and how you see that coming out of this trough into next year. Thanks, Lisa.

LISA SU: Yeah, sure, Tom. So talking a little bit about the PC market, the PC market continues to be an important market to us. I think we were signaling all along in the last couple of calls that we were expecting a subseasonal second half of the year, and that's how it's played out. Our view is this is one of those times when component prices go up, it affects overall system pricing. One would expect the consumer markets to be more sensitive to that, and so we have seen a bit of slowdown in the sell-through.

We're just managing the business very prudently, and we want to ensure that inventory stays healthy throughout the channel and throughout our customers. And so our planning assumption for 2027 is that the market will be double-digit percentage down. We will certainly focus on the segments where we can add the most value, places like enterprise, where we continue to be underrepresented, premium consumer, desktop channel. These are really good markets for us. But we are planning with the current environment for a down year for 2027, and we will work through that.

The good part about it is our supply chain is quite fungible across our PC business and our Data Center business, and so that actually does afford us the ability to move supply around as we see these signals in the market. And overall, as I said earlier, we do have more demand than we have supply, and so we're continuing to work on all of the aspects to increase supply as we go through the next few quarters.

OPERATOR: Thank you. And the next question comes from the line of Vivek Arya with Bank of America Securities. Please proceed with your question.

VIVEK ARYA: Thanks for taking my questions. Lisa, there's a lot of market debate about the financing of data centers and availability of land, power, and shell, et cetera. So kind of a two-parter. What's the level of confidence that AMD has that your anchor customers have the financing and the physical infrastructure to support-- I think it's roughly 3 gigawatts or so-- of planned AMD deployments into '27?

And then, as we look beyond, many of your competitors are investing in the ecosystem, providing residual value guarantees and such, and saying there is a lot more fungibility of their products. So does that create a competitive dynamic for AMD? So just curious to hear your perspective on '27 readiness and '28 competitive landscape.

LISA SU: Sure, Vivek. Thanks for the question. So as I said earlier, we're seeing very strong demand from our lead Data Center AI customers. I think you can see that there's strong demand for more compute, and MI455 is positioned extremely well. When we look at Venice, MI455, Helios, I think we feel extremely good about the competitiveness.

We have been planning for this ramp with our customers jointly. So we have good visibility into the 2027-- as you said, land, power, shells, all of those, the data center infrastructure-- and we are very much working hand in hand with our customers. So we know where the GPUs are going in terms of which data centers and which locations and all of that stuff.

Going into 2028 and beyond, I guess what I would say is the following. Look, I think we are seeing very strong growth overall. I think we have a very strong business, we have a very strong balance sheet, and from that standpoint, our goal is-- first and foremost, we're investing in our growth to support that, and that means investing in R&D, investing in supply chain and capacity, investing alongside our partners to ensure that we have the capacity for all of those things that we talked about earlier that are necessary for this type of ramp.

And then we are selectively investing with our lead partners where we see it's helpful to have, for example, the data center infrastructure identified way in advance, and that's multiple years in advance. And I think that's the right thing to do to ensure that we are jointly ensuring that there's the right infrastructure as our new chips become available. But we are being selective and strategic in how we do that, with the first principle being investing in our growth.

VIVEK ARYA: Got it. And for my follow-up, Lisa, one more on the competitive landscape. Many of your frontier lab customers have been working on their custom ASICs, and I'm curious, how do they conceptually make the choice between using ASICs versus AMD products? What determines their allocation? Is it cost per token, power efficiency, software flexibility, time to deployment? Just what are the puts and takes in that decision-making, and are there workloads where using an AMD GPU works better than the alternatives that they have? Thanks.

LISA SU: Sure, Vivek. So look, I think we feel extremely good about where we're positioned in terms of our performance and our total cost of ownership that we're delivering with Helios. That is fundamentally the most important factor when we think about deployment scale and deployments going forward. I think there was a significant investment that we've made in the rack-scale systems and ensuring that Helios really ramped nicely.

Third quarter was the first quarter that we shipped MI455 and Helios racks for revenue. I think this is a very major milestone. We have a big step-up in Q4. We'll continue to ramp throughout 2027. And the key is for us to continue to show that TCO advantage across workloads.

It is true that from a general-purpose GPU standpoint, that we have more fungibility as you think about various workloads that occur. There has been also very strong co-design. So beyond the MI450 Series, we're working on MI500 and MI600 as well to ensure that our technology gets even more, let's call it, tuned to what the frontier labs need. So I think those are the key factors in terms of deployment choices.

OPERATOR: Thank you. And the next question comes from the line of Joshua Buchalter with TD Cowen. Please proceed with your question.

JOSHUA BUCHALTER: Hey, guys. Thanks for taking my question, and congrats on the progress. I was wondering if you could give us insight into how we should be thinking about HBM costs impacting your model as the Instinct business ramps. I think investors have been asking how your revenue per gigawatt contribution in the last year or so-- your outlook has been so consistent in that $15 to $20 billion per gigawatt range-- given what's going on with HBM pricing.

Is that because early builds are fairly program-specific and pricing was agreed upon a long time ago? And again, as we think about that business longer term, do you have the ability to pass on these higher input costs even to your big three customers? Thank you.

LISA SU: Sure, Josh. So the way to think about it is, as we ramp, we have had, let's call it, long-term agreements with some of our suppliers as well as some of our customers. But given where the memory market is-- and not just the memory market, by the way, other component costs are also seeing some inflationary pressures as there's more CapEx that's put into the overall supply chain environment-- we are also working with our customers to ensure that those costs are shared as we go forward.

So from my perspective, back to the question that Vivek had, this is all about TCO and ensuring that from a tokens-per-dollar standpoint, that we continue to offer very strong advantages for choosing AMD Instinct and our overall Helios system. And even with some of the inflationary pressures that are out there, we certainly believe that we are offering leading TCO for the AI workloads.

JOSHUA BUCHALTER: Thank you, Lisa. Then maybe for Jean, it was good to see OpEx come in light in the third and fourth quarter. It's been a sticking point for investors for a little bit. But I guess, given how steep the revenue growth is expected to be in 2027, we're flying a little blind, given you're at more than 2x your target model.

How should we think about OpEx growth in '27? And I guess, while I'm in the neighborhood, do you think you can hold your target gross margin range-- that 55% floor-- as we move through the year and Helios continues to ramp? Thank you.

JEAN HU: Yeah, Josh, thank you for the question. On the OpEx, as you said, we have been investing aggressively. What you are seeing is the investment are translating to accelerating revenue growth we're seeing in 2026. We're actually quite pleased with the 2026, our operating leverage. We're expanding revenue much faster than OpEx growth, which driving the operating margin increase. And if you think about the Q4 guide, we're guiding very significant revenue growth, more than OpEx.

That dynamic, we believe, will carry into 2027. Lisa talked about it-- that we expect the whole company growing more than 80%. We should expect OpEx to grow less than half of the top-line revenue growth rate. So that will drive a very significant operating margin expansion in 2027.

Going back to gross margin, we are actually very happy about where we are with our gross margin. In Q3, we are ramping MI450. Data Center AI business is doing very well-- actually, better than our expectation when we got into Q3 quarter. And in Q4, we're guiding sequentially revenue growing 50% for Data Center business. It's actually going to be the quarter Data Center AI business surpasses server business. So we all know there's a margin a little bit dilutive to corporate average, but the gross margin dollars are increasing very significantly.

I think, looking ahead over 2027, we actually think our Data Center business will be more than 80% of total company revenue. So the mix of Data Center business is going to drive the company's gross margin significantly. The way to think about it is, Data Center AI will continue to ramp through 2027, and the server business is expected to double. So quarter over quarter, there may be variations because the mix of those business will be a little bit different. In addition, Embedded business is strengthening in 2027. I think that will also help our gross margin.

Overall, when you put the whole business model together, what you should expect us to do is significant top-line revenue growth, significant gross profit expansion, and OpEx growing less than revenue growth, which will result in very significant operating margin expansion and the profitability expansion.

OPERATOR: Thank you. And the next question comes from the line of Timothy Arcuri with UBS. Please proceed with your question.

TIMOTHY ARCURI: Thanks a lot. Lisa, I'm wondering if you can break down a little bit how Data Center is going to look like next year. You did revise server CPU up from up 70% to up more than 100%, but you're not being super specific about Data Center AI. So is the right way to think about it that the GPU business will still be bigger in dollars than the server CPU business but not much? Is that the right way to think about the split? And then how to think about the shape of the year? I know that the GPU business is lumpy, but I imagine it's probably pretty ratable to the year. Is that kind of fair to think? Thanks.

LISA SU: Yeah, Tim, maybe let me break down a couple of things. So first of all, we think 2027 is going to be a very strong momentum year for both the CPU business as well as the Data Center AI business. And when we gave you some of the markers, you should just consider them markers. So we've said the overall company will grow greater than 80%, and for the CPU business, we will about double, and for the Data Center AI business, it will significantly more than double.

And from that standpoint, you should expect that the Data Center AI business continues to be-- as Jean just said, Q4 is the first quarter where the Data Center AI business will cross over in terms of size compared to the CPU business. And we would expect both businesses to grow nicely throughout the year. So hopefully, that gives you a bit of shape as we go through the year.

TIMOTHY ARCURI: Thanks a lot. And then, Jean, I guess, can you help us on the warrants? How are they going to layer in through the year? Any sort of guardrails you can give us on how to roll that into our model through the year and into 2027? Thanks.

JEAN HU: Yeah, thanks for the question. I think the simple way to think about it is, there are three parts for the warrant. The first one is, we report net revenue, the Q3 revenue, Q4 guide. And whenever we talk about the revenue, it's always net revenue, which already incorporated the warrant revenue reduction piece, which is determined at the grant date fair value. So that one you should not worry about it. It's always incorporated in our revenue whenever we talk about the revenue.

The second piece is, probably you're asking, is the share count. The way to think about it is, the share count only come into diluted share once all the vesting conditions and the exercising conditions are met by our customers. That means the shipment need to be there, and some commercial conditions need to be met. We are going to provide the guidance in 2027 to give you clear guidance when we believe the share vesting and exercising conditions are met. That is something we'll provide the transparency once we get into 2027.

The last piece is, it's a quarterly non-cash charge, which is a remeasurement of fair value. That actually does not impact our non-GAAP financial results and free cash flow. So on the share count, just wait for us to provide you clear guidance when we get there.

OPERATOR: Thank you. And the next question comes from the line of CJ Muse with Cantor Fitzgerald. Please proceed with your question.

CJ MUSE: Yes, good afternoon. Thank you for taking the question. I guess, first question-- on the server CPU side, you've obviously taken your outlook higher. I'm curious how we should be thinking about units versus mix, including higher core count, as well as higher ASPs associated with Venice. Is there a way to break that down to help us understand that 100% growth?

LISA SU: Sure, CJ. So on the server CPU side, when we talk about 100% growth, you should imagine that we are definitely shipping significantly more units in that. Of course, some of those units have higher core counts as we go into Venice. But I would say the growth in revenue is mostly in units and core count, and less so in ASPs, although there's an ASP uplift as we increase core count.

CJ MUSE: Great. Thanks. And I guess, as a follow-up, just to button down 1Q seasonality, based on what you said tonight, I'm interpreting that we should see kind of normal Client and Gaming down seasonally in Q1, kind of ratable Data Center above December. And I guess, as part of that, any sort of vision for how to think about Embedded?

JEAN HU: For the Embedded business, typically there is seasonality. We do see overall, when we look at 2027 for the whole year, Embedded business is actually going to grow significantly and also strengthening through the year. But typically Q1, actually there are some seasonality for Embedded business.

OPERATOR: Thank you. And the next question comes from the line of Stacy Rasgon with Bernstein Research. Please proceed with your question.

STACY RASGON: Hi, guys. Thanks for taking my questions. I was wondering-- I know you said that you hit crossover for GPU versus CPU in Q4 within the context of that 50% guide. I was wondering if you could help us with the magnitude of that crossover. Is it just barely there, or is it significantly higher? I'm just trying to get the sizing in my head in terms of the entrance right into next year.

JEAN HU: Yeah, hi, Stacy. Maybe I'll give you some color. So the way to think about it is, if you think about the MI450 ramp, it's a very significant step-up sequentially from Q3 to Q4. And that really, compared to server side-- server is growing really nicely, strongly, but the ramp for Data Center AI surpassed the server in Q4. That is a very significant step-up. I think the way-- server business, of course, is very strong, but the Data Center AI will be much higher than server business in Q4.

STACY RASGON: In terms of dollars?

JEAN HU: Yeah, in terms of dollars. Yeah.

STACY RASGON: OK. Great. And thanks. And for my follow-up, Jean, there were some questions earlier about gross margins next year. I wasn't exactly quite sure how to interpret your answer in terms of the gross margins next year relative, I guess, to the Q4 rate. Do you think you could give us some color, some highlights on a good entrance rate and a good exit rate for gross margins in 2027?

JEAN HU: I think that--

STACY RASGON: Given the ramps of the different products and everything that are happening.

JEAN HU: Yeah, Stacy, the reason we're not able to provide you quarter-over-quarter-- it's really because when you think about the data center mix, it's a primary driver in 2027. When you have a Data Center business, revenue is more than 80% of the total company business. And then, within Data Center, the interesting part is you have a Data Center AI business continue to ramp significantly.

On the server side, Lisa mentioned that we're going to double in 2027. Of course, we also know the supply-constrained environment, and Lisa also mentioned that we are working with our suppliers to ensure we get more supplies. So the mix between those two business actually determines what the company's gross margin will be.

Right now, I actually don't have that visibility. We hope both business are going to do great, and we know they're both doing great. Q4 actually is a good benchmark to think about it. It is because Q4 Data Center AI business is much higher than CPU business. So I think that gives you some idea at that benchmark. Then, in 2027, different quarter could be different. There could be variation there.

OPERATOR: Thank you. And the next question comes from the line of Joe Moore with Morgan Stanley. Please proceed with your question.

JOE MOORE: Great. Thank you. We've heard quite a bit of positive feedback on the Helion racks, and I wonder if you could talk a little bit about that. In the CPU-only version of that, what's different for these agentic workloads versus a traditional rack? And then, do you anticipate an attach rate of the GPU blade in a lot of these?

LISA SU: Sure, Joe. Yeah, that there's quite a bit of positive excitement around the Helion racks. And I think the main point there is, these agentic racks, they really want to be optimized for the environment that they're in. So if they are in, let's call it, hyperscaler or large cloud racks, you're going to have them in all CPU configurations. It's going to be quite dense. Venice is, from a performance-per-socket standpoint, the best in the industry, and so that will translate into a very strong agentic rack if you fully populate it with CPUs.

But in enterprise areas, you will also see the opportunity to mix CPUs and GPUs, and it just allows people to really optimize for the workloads that they're in. And it is-- we have contributed to OCP, so it's a fully open rack, and I think that is viewed as a positive. And from that standpoint, we've gotten very positive feedback from both OEMs and ODMs, as well as customers.

JOE MOORE: OK, thank you. And if I could follow up, the 100% growth in server-- do you think you're meeting demand at that level, or do you think you still have a shortage condition through all of next year?

LISA SU: Well, Joe, for us to grow that amount, we're meeting a lot of demand. And when you look at-- when we talked about the TAM growing, let's call it, about 50% CAGR over the next three or four years, we saw growth in three areas-- so traditional cloud and enterprise, agentic CPUs and agentic workloads, and head nodes. We believe we're gaining share in each one of those segments as we give you that forecast of roughly 100%.

I do believe there's more demand, and we are actively working with our supply chain partners to satisfy additional demand. Obviously, we also have to make sure that these squared sets are there and DDR5 and all of the other components are there. But my view is, we are taking a big step-up in terms of the demand forecast from our customers.

I'm excited about the fact that our demand is broad-based, so it is really virtually across every customer in terms of their expansion of their footprint with AMD, and so that's a positive. And we're going to continue to work on increasing supply to meet the additional demand as it comes to fruition.

MATT RAMSAY: John, I'm actually going to ask you to have one more caller join us, please. I think we have time for one more questioner. Thank you.

OPERATOR: Thank you. And the final question comes from the line of Jim Schneider with Goldman Sachs. Please proceed with your question. Hello, Jim. Your line is now live. OK, our final question comes from the line of Harsh Kumar with BMO Capital Markets. Please proceed with your question.

HARSH KUMAR: Yeah, hey, Lisa. Question is on Helios. So you're one of two companies shipping complete racks. Can I ask you about what your share aspirations is? And the reason why I'm asking is you've done a wonderful job in servers, taking share from basically a very difficult competitor. And also, the second part of that question is, do you have aspirations to address the enterprise market with the Helios rack and the setup? And then I have a follow-up.

LISA SU: Sure, Harsh. So look, we're very happy with the progress that we've made with Helios. I think this is a significant-- the rack-scale solutions are significant R&D efforts. I think Helios has ramped very nicely. We saw a good initial quarter in Q3. We're ramping nicely into Q4 and into 2027.

Our aspirations are to be a very significant compute supplier to our largest customers, both frontier labs as well as hyperscalers. And from that standpoint, the focus on Helios, MI450 ramp is very much on our minds for Q4 and 2027. But we're in very active dialogues with our largest customers on MI500 and MI600 and beyond. And our aspirations are to continue to have a larger and larger piece of their compute needs. And from everything that we see, demand continues to go up. So I think that's the overall view on Helios. Was there another part of that question?

JEAN HU: Enterprise.

LISA SU: Oh, enterprise.

HARSH KUMAR: Yeah, about the enterprise. Yeah.

LISA SU: Sure. Yes, certainly. I think we continue to see-- although enterprise is developing, let's call it, a little bit slower than the largest frontier labs and hyperscalers, we do believe that the enterprise business is quite important. We have made a lot of progress in enterprise CPUs.

Many of our enterprise customers are actually asking for our advice as to how to deploy AI and how to ensure that they're deploying it as efficiently as possible when you think about both on-prem as well as cloud deployments. So we are expanding our portfolio. We've introduced some enterprise-specific GPUs, and that will be a place where we would like to grow our footprint to, again, be a very strategic supplier to our largest enterprise customers.

HARSH KUMAR: Understood. And my follow-up is on the Client side. Historically, you've done a great job managing this business through situations like this. What would be your share-keep strategy if your chief competitor was to cut ASPs? Would you try to match them, or would you have to just gauge the situation as it goes? Just be curious about how you're going to manage this.

LISA SU: Yeah, I think our view is, PCs are an important market for us in the standpoint of-- I'm a big believer in end-to-end AI deployment, so from cloud to edge to client, and we have a unique portfolio in that fashion. That being the case, I think we're always quite disciplined in how we manage.

Our focus has really been on the highest value and most profitable segments of the PC market. Enterprise, premium consumer, desktop channel-- these are places that we do very well. We do recognize that there's a good amount of volatility in the low end of the market, and that's the place where we have the least exposure, frankly, to that portion of the market.

MATT RAMSAY: All right, thank you, John, for organizing the Q&A. Thank you, everybody, for joining our call. We really appreciate you spending some time with AMD, and thank you very much for your questions and your attention. John, go ahead and close the call. Thank you very much.

OPERATOR: Thank you, Matt. And ladies and gentlemen, thank you for your participation. This does conclude today's teleconference. You may disconnect your lines at this time, and have a wonderful day.