JOHN: Greetings. And welcome to the AMD third quarter 2024 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, please press star 0 on your telephone keypad. And as a reminder, this conference is being recorded. It is now my pleasure to introduce to you Mitch Haws, vice president of investor relations. Thank you, Mitch. You may begin.

MITCH HAWS: Thank you, and welcome to AMD's third quarter 2024 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 the slides posted on our website. Participants on today's conference call are Dr. Lisa Su, our chair and chief executive officer, and Jean Hu, our executive vice president, chief financial officer, 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 Forrest Norrod, executive vice president and general manager, data center business solutions, will attend the UBS Annual Technology Conference on Tuesday, December 3, and Jean Hu will attend the Barclays Global TMT Conference on Thursday, December 12.

Today's discussion contains forward-looking statements based on current beliefs, assumptions, and expectations, 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. With that, I'll hand the call over to Lisa.

LISA SU: Thank you, Mitch. And good afternoon to all those listening today. We delivered strong top and bottom-line growth in the third quarter, with revenue coming in above expectations, driven by record Instinct and EPYC product sales and robust demand for our Ryzen PC processors.

Third quarter revenue increased 18% year over year to a record $6.8 billion as significantly higher data center and client processor sales more than offset declines in gaming and embedded product sales. We expanded gross margin by 2 and 1/2 percentage points and increased earnings per share by 31% year over year as data center segment revenue more than doubled.

Turning to the segments, data center segment revenue increased 122% to a record $3.5 billion. We believe we gained server CPU share in the quarter as enterprise wins accelerated, cloud providers expanded their use of EPYC CPUs across their infrastructure, and we began the initial ramp of 5th Gen EPYC processors. EPYC has become the CPU of choice for the modern data center. And our multigeneration product portfolio delivers leadership performance and significant advantages across virtually every enterprise and cloud workload.

In cloud, EPYC CPUs are deployed at scale to power many of the most important services, including Office 365, Facebook, Teams, Salesforce, SAP, Zoom, Uber, Netflix, and many more. Meta alone has deployed more than one and a half million EPYC CPUs across their global data center fleet to power their social media platforms.

Cloudflare selected Genoa-X processors with our industry-leading 3D chiplet stacking technology to power their next-generation servers that support twice as many requests per second and deliver 60% higher performance per Watt versus their prior generation. Public cloud instances increased 20% year over year to more than 950 as Microsoft, AWS, and others launched or expanded their EPYC processor-powered offerings in the quarter. EPYC instance adoption with enterprise customers also grew in the quarter, highlighted by wins with Adobe, Boeing, Micron, Nestle, Slack, Synopsys, Tata, and others.

In the enterprise, sales grew by a strong double-digit percentage year over year for the fifth straight quarter as EPYC CPU adoption accelerated and sell-through momentum grew. Dell, HPE, Lenovo, and others have expanded the number of 4th Gen EPYC platforms they offer by 50% in the last year. There are now more than 200 different EPYC solutions available that are optimized for a broad range of enterprise and edge workloads.

We are building strong momentum with large enterprise customers highlighted in the third quarter by wins with large technology, energy, financial services, and automotive companies in the quarter, including Airbus, Daimler Truck, FedEx, HSBC, Siemens, Walgreens, and others.

We launched our next generation Turin family earlier this month that delivers absolute performance and TCO leadership across both enterprise scale-up and cloud native scale-out workloads. Turin has already set more than 130 performance records for virtualization, database, AI, business applications, and energy efficiency, with the full EPYC portfolio accounting for more than 500 performance world records.

More than 130 5th Gen EPYC enterprise platforms are in development from all the leading server OEMs and ODMs. These new servers complement existing 4th Gen EPYC platforms, providing a top to bottom stack of platforms optimized for broad range of business applications.

In cloud, Google and OCI announced plans to launch 5th Gen EPYC instances early next year, and we expect broad adoption with our largest cloud customers based on the significant performance and efficiency advantages of Turin. As an example, Oracle's Turin instances delivered 35% higher performance per core, 33% faster memory speeds, and double the networking bandwidth, delivering a level of compute performance and capability that is only possible with EPYC CPUs.

Looking ahead, we are very well-positioned for continued growth and share gains based on the strength of our broad EPYC portfolio and the momentum we have built with cloud and enterprise customers. We also took a major step in the quarter to advance the x86 architecture, forming an ecosystem advisory group with Intel, several industry luminaries, and the largest cloud PC and enterprise leaders to accelerate innovation by driving consistency and compatibility across both the x86 instruction set and architectural interfaces, and ensuring we evolve x86 as a compute platform of choice for developers and customers.

Turning to our data center AI business, data center GPU revenue ramped as MI300X adoption expanded with cloud, OEM, and AI customers. Microsoft and Meta expanded their use of MI300X accelerators to power their internal workloads in the quarter. Microsoft is now using MI300X broadly for multiple Copilot services powered by the family of GPT-4 models.

Meta announced they have optimized and broadly deployed MI300X to power their inferencing infrastructure at scale, including using MI300X exclusively to serve all live traffic for the most demanding Llama 405B frontier model. We are also working closely with Meta to expand their Instinct deployments to other workloads where MI300X offers TCO advantages, including training.

MI300X public cloud instance availability expanded in the quarter, with Microsoft, Oracle Cloud, and multiple AI-specialized cloud providers now offering Instinct instances with leadership performance and TCO for many of the most widely used models. Instinct cloud instance adoption is strong, with multiple startups and industry leaders adopting MI300 instances to power their models and services, including Essential AI, fireworks AI, Luma AI, and Databricks.

On the AI software front, since launching MI300 10 months ago, we have expanded functionality at every layer of the ROCm stack and increased the number of models that run out of the box on Instinct accelerators to more than one million, enabling customers to get up and running as fast as possible with maximum out-of-the-box performance.

With the release of ROCm 6.2 last quarter, MI300X inferencing performance has improved 2.4 times since launch and training performance has increased 80%. We are working closely with the growing number of marquee, cloud, and enterprise customers to fine-tune their specific inferencing workloads for MI300, with many customers seeing 30% higher performance compared to competitive offerings. And we continue to expand our work with the open source community, broadening support for key frameworks like JAX, libraries like vLLM, and hardware-agnostic compilers like Triton.

At our Advancing AI event earlier this month, we were excited to be joined by the creators and leaders of some of the most important AI software technologies who have added foundational support for ROCm into Triton, the Llama stack, SGLang, vLLM, and TensorFlow, and are working to enable broader open source community work with Instinct platforms. With this growing support from the broader AI software ecosystem and the significant advances we have made in our software stack, ROCm now provides AI developers with a truly open software alternative that has been deployed and validated at scale.

To expand our AI systems capabilities, we announced a definitive agreement to acquire ZT Systems, one of the leading providers of AI infrastructure to the world's largest hyperscale computing companies. The ZT team complements our silicon and software capabilities with critical systems expertise needed to deliver rack and cluster-level solutions. With ZT, we will be able to design and validate our next-gen AI silicon and systems in parallel, greatly accelerating time to deploy Instinct accelerators at data center scale.

Customer feedback has been very positive as the ZT acquisition enabled hyperscale customers to rapidly deploy AMD AI infrastructure at scale and provides OEMs and ODMs with optimized board and module designs for a wide range of differentiated enterprise solutions.

On the regulatory front, we made good progress as we recently passed the HSR waiting period required for US approval. We remain on track to close the acquisition in the first half of 2025. As a reminder, we plan to divest ZT's industry-leading US-based data center infrastructure-manufacturing business at the close of the transaction and are pleased that we have received significant interest from a number of parties to date.

Looking ahead, we launched our next-gen MI325X GPU earlier this month that extends our memory capacity and bandwidth advantages and delivers up to 20% higher inferencing performance compared to H200 and competitive training performance. Customer and partner interest for MI325X is high. Production shipments are planned to start this quarter, with widespread system availability from Dell, HPE, Lenovo, Supermicro, and others, starting in the first quarter of 2025.

Longer term, we have successfully accelerated our product development pace to deliver an annual cadence of new Instinct products. Our next-gen MI350 series silicon is looking very good and is on track to launch in the second half of 2025, with the largest generational increase in AI performance we have ever delivered.

Development on our MI400 series based on the CDNA Next architecture is also progressing very well towards a 2026 launch. We have built significant momentum across our data center AI business, with deployments increasing across an expanding set of cloud, enterprise, and AI customers. As a result, we now expect data center GPU revenue to exceed $5 billion in 2024, up from $4 and 1/2 billion we guided in July and our expectation of $2 billion when we started the year.

Turning to our client segment, revenue was $1.9 billion, an increase of 29% year over year, driven by strong demand for our latest generation Zen 5 notebook and desktop processors. Desktop channel sales grew by a significant double-digit percentage, led by the launch of our Ryzen 9000 Series processors that deliver leadership productivity, gaming, and content creation performance. We are seeing strength across our Ryzen desktop portfolio and are on track to launch our next-gen Ryzen 9000X3D processors in November with leadership gaming performance.

In mobile, Ryzen AI 300 series sales ramped significantly from the prior quarter as Acer, HP, Lenovo, ASUS, and others announced new consumer and commercial notebooks with leadership compute and AI performance. We made good progress expanding our presence in the commercial PC market in the quarter, closing multiple large deals with AstraZeneca, Bayer, Mazda, Shell, Volkswagen, and other enterprise customers.

We also launched our Ryzen AI Pro 300 series family, the first CPUs with enterprise-class security, manageability, and AI capabilities for Copilot+ PCs. HP and Lenovo are on track to more than triple the number of Ryzen AI Pro platforms they offer in 2024. And we expect to have more than 100 Ryzen AI Pro commercial platforms in market next year, positioning us well for share gains as businesses refresh the hundreds of millions of Windows 10 PCs that will no longer receive Microsoft technical support starting in 2025.

Now, turning to our gaming segment. Revenue declined 69% year over year to $462 million. Semi-custom sales declined as Microsoft and Sony reduced channel inventory. Sony announced the PS5 Pro with significant increases in graphics and ray-tracing performance and AI-driven upscaling, featuring a new AMD semi-custom SoC that extends our multigenerational partnership.

In gaming graphics, revenue declined year over year as we prepare for a transition to our next-gen Radeon GPUs based on our RDNA 4 architecture. In addition to a strong increase in gaming performance, RDNA 4 delivers significantly higher ray-tracing performance and adds new capabilities. We are on track to launch the first RDNA 4 GPUs in early 2025.

Turning to our embedded segment, third quarter revenue decreased 25% year over year to $927 million. Embedded demand continues recovering gradually, led by strength in test and emulation, offset by ongoing softness in the industrial market.

Momentum continues building for our differentiated Versal family of adaptive SoCs, led by strong demand for our Versal Premium VP1902, which is the world's largest adaptive SoC and FPGA that is powering multiple platforms for all three of the largest EDA vendors.

Our Versal portfolio is also being adopted broadly across multiple aerospace customers. To build on this momentum, we taped out Telluride last quarter, the first product in our 2nd Gen Versal family that delivers up to 10x more compute and enables AI application acceleration on a single chip.

Design win momentum is very strong across our portfolio, tracking to grow more than 20% year over year in 2024 and positioning us well to grow our embedded business faster than the overall market in the coming years. In summary, the business accelerated in the third quarter, and we expect strong demand for our Instinct, EPYC, and Ryzen processors to result in another quarter of significant year over year growth.

Taking a step back, this month marks my 10th anniversary as AMD's CEO. In the last 10 years, we have successfully completed multiple arcs, first by turning the company around and setting the solid financial and operational foundation required for sustained growth, and then by transforming AMD into the high performance and adaptive computing leader.

While I'm incredibly proud of what we have accomplished, I'm even more excited about the unprecedented growth opportunities in front of us. Looking out over the next several years, we see significant growth opportunities across our data center, client, and embedded businesses, driven by the nearly insatiable demand for more compute. Each of these opportunities is amplified exponentially by the rapid adoption of AI, which is enabling new experiences that will make high performance computing an even more essential part of our daily lives.

In the data center alone, we expect the AI accelerator TAM will grow at more than 60% annually to $500 billion in 2028. To put that in context, this is roughly equivalent to annual sales for the entire semiconductor industry in 2023. Beyond the data center, we are adding leadership AI capabilities across our product portfolio and partnering deeply with the broad ecosystem of partners to deliver differentiated AI solutions at scale.

This is an incredibly exciting time for AMD as the breadth of our technology and product portfolios, combined with our deep customer relationships and diversity of markets we address, provide us with a unique opportunity as we execute our next arc and make AMD the end-to-end AI leader. Now, I'd like to turn the call over to Jean to provide some 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 current outlook for the fourth quarter of fiscal 2024. We're very pleased with our strong third quarter financial results.

On a year over year basis, data center segment revenue more than doubled, and client segment revenue grew 29%. We expanded gross margin by 250 basis point and drove earnings per share growth of 31%.

For the third quarter of 2024, revenue was $6.8 billion, up 18% year over year. And the revenue growth in our data center and client segment was partially offset by lower revenue in our gaming and embedded segment. Revenue increased 17% sequentially, primarily driven by growth in our data center and client segment. Gross margin was 54%, up 250 basis point year over year, primarily driven by higher data center segment revenue.

Operating expenses were $1.96 billion, an increase of 15% year over year, as we continue to invest in R&D and go-to-market activities. Operating income was $1.7 billion, representing a 25% operating margin. Taxes, interest expense, and other was $211 million. Diluted earnings per share was $0.92, an increase of 31% year over year and 33% sequentially.

Now turning to our reportable segments, starting with the data center. Data center delivered record quarterly segment revenue of $3.5 billion, up 122%, nearly a $2 billion increase year over year, an increase of 25% sequentially. Growth in revenue was led primarily by the strong ramp of AMD Instinct GPU shipment and growth in AMD EPYC CPU sales. The data center segment accounted for 52% of total revenue in the third quarter.

Data center segment operating income was $1 billion, or 29% of revenue, compared to $306 million, or 19%, a year ago. Data center segment operating income more than tripled compared to the prior year, driven by higher revenue and operating leverage.

Client segment revenue was $1.9 billion, up 29% year over year and 26% sequentially, driven primarily by strong demand for our Zen 5 AMD Ryzen processors. Client segment operating income was $276 million, or 15% of revenue, compared to our operating income of $140 million a year ago, or 10% of revenue, primarily driven by higher revenue, partially offset by higher operating expenses.

Gaming segment revenue was $462 million, down 69% year over year and 29% sequentially, primarily due to a decrease in semi-custom revenue. Gaming segment operating income was $12 million, or 2% of revenue, compared to $208 million, or 14%, a year ago.

Embedded segment revenue was $927 million, down 25% year over year, as customers continue to normalize their inventory levels. Revenue increased 8% sequentially as demand improved in several end markets. Embedded segment operating income was $372 million, or 40% of revenue, compared to $612 million, or 49%, a year ago.

Turning to the balance sheet and the cash flow. During the quarter, we generated $628 million in cash from operations. And the free cash flow was $496 million. Excluding certain nonrecurring payments related to acquisitions from operating cash flows, our free cash flow was $619 million.

Inventory increases sequentially by $383 million to $5.4 billion, primarily to support the continued ramp of data center segment products. At the end of the quarter, cash, cash equivalents, and short-term investment was $4.5 billion. In the third quarter, we returned $250 million of cash to shareholders, repurchasing 1.8 million shares. And we have 4.9 billion of authorization remaining.

Now, turning to our fourth quarter of 2024 outlook. We expect revenue to be approximately $7.5 billion, plus or minus $300 million, up 22% year over year, driven by strong growth in our data center and client segments, more than offset decline in the gaming and embedded segments.

We expect revenue to be up approximately 10% sequentially, driven primarily by growth across data center, client, and gaming segments. In addition, we expect fourth quarter non-GAAP gross margin to be approximately 54%, non-GAAP operating expenses to be approximately $2.05 billion, non-GAAP other net income to be $17 million, non-GAAP effective tax rate to be 13%. And the diluted share count is expected to be approximately 1.64 billion shares.

In closing, we're pleased with our strong execution in the third quarter. We delivered record revenue along with strong year-over-year expansion in gross margin and earnings per share growth. Looking ahead, we are very well-positioned to deliver another record quarter of revenue in the fourth quarter, driven by continued momentum in our data center and client segments.

Importantly, we are making the strategic investment and position AMD as the end-to-end AI infrastructure leader and to drive long-term profitable growth. With that, I'll turn it back to Mitch for the Q&A session.

MITCH HAWS: Thank you, Jean. John, let's pull the audience for questions.

JOHN: Thank you, Mitch. We will now be conducting a question and answer session. 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 to remove yourself 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 Toshiya Hari with Goldman Sachs. Please proceed with your question.

TOSHIYA HARI: Hi. Good afternoon. Thank you so much for taking the question. My first one is on the data center GPU business. Lisa, you took up your '24 outlook by $500 million. Curious what drove the change there.

And more importantly, as you look forward into calendar '25, I doubt you're going to give us quantitative guidance on this call, but, conceptually, how are you thinking about growth in your Instinct business between your large cloud customers and your enterprise customers? And if you can speak to the possibility of adding new customers within cloud, again, specific to Instinct, that would be really helpful.

LISA SU: Sure, Toshiya. Thank you for the question. So first of all, we had a very strong quarter for the data center overall in Q3, and especially for the Instinct product portfolio. We actually completed some important customer milestones, and we were able to ramp a bit above our initial expectations.

So data center GPU was very strong in the third quarter, and we raised overall guidance for the year from exceeding $4.5 billion to exceeding $5 billion based on the completion of some of those customer milestones. So we feel good about the trajectory as we go through the end of this year.

And then to your overall question about 2025, at a high level, look, we feel very good about the market. From everything that we see, talking to customers, there's still significant investment in trying to build out the infrastructure required across all of the AI workloads. And then within that, our product portfolio is getting stronger with the annual cadence launching 325 later this quarter and 355 in the second half of next year.

And then in terms of customers, our customer engagements are actually broadening quite well. And it's broadening in two ways. So certainly, cloud, our largest cloud customers are broadening the set of workloads that they're running on AMD Instinct. And we're also very engaged with a number of large cloud and enterprise customers that are actively working with us on optimizing their workloads. And we would expect those would be good opportunities for us over the next couple of quarters as well.

TOSHIYA HARI: Great. And then a quick follow up, maybe one for Jean, on gross margin. You're guiding Q4 gross margin essentially flat sequentially. If you can walk through the puts and takes there, that would be really helpful.

And then again, as you look forward into '25, you're kind of speaking to continued data center growth. One would think the embedded business hopefully begins to recover. And then within some of your businesses, like server CPU, I would expect enterprise to grow perhaps faster off a low base. So I do think you have a lot going for you from a gross margin standpoint. But are those points valid? And what should we be thinking about in terms of potential headwinds as well? Thank you.

JEAN HU: Yeah. Yeah. Thank you for that question. First, we are very pleased with our Q3 gross margin performance. We delivered 53.6%. And we are guiding approximately 54% in Q4. In general, when you look at 2024, our gross margin improvement has been primarily driven by the mix, especially data center business continue to be the strong growth driver of our business. It's accounting for more than 50% of our revenue mix. That helps us to improve gross margin.

Going into 2025, we're not going to guide specifically. But as you mentioned, there are puts and takes that will help us going forward. I will say first is, going forward, the largest growth driver is our data center business, both the CPU side, the GPU side. And you are absolutely right, we see our enterprise server business continue to expand. That will be a tailwind on gross margin side.

And secondly, embedded business is recovering, and gradually, but that will also help us with our gross margin. At the same time, we are also seeing our client business expansion nicely. Client business today is more focused on consumer side, which tend to be below corporate average. That's something headwind will deal with.

And lastly, I will say, our team has done a great job to continue to improve operational efficiency. When we scale the company next year, you can see we're going to benefit from economies of scale to continue to drive operational efficiency to improve gross margin.

JOHN: And the next question comes from the line of Aaron Rakers with Wells Fargo. Please proceed with your question.

AARON RAKERS: Yeah. Thanks for taking the question. I'll do two as well. I guess maybe building on the latter point that you just made, I'm curious of how you would characterize your supply chain. Having evolved-- now talking about $5-plus billion, how do we think about what you've been doing on the supply side given the lead times, the production cycles of these MI300 and 325 GPUs as we look out into 2025? And I have another quick one.

LISA SU: Yeah. Sure, Aaron. So, look, I've been very happy with how our supply chain has ramped over the last number of quarters. Clearly, it's a tight supply environment. But we've done a great job getting-- ensuring that we have capacity across the entire supply chain. Again, that was part of the reason for the higher revenue in the third quarter around our Instinct business, just both customer demand as well as supply chain improvement.

And going into the next few quarters, going into 2025, I think we expect that the environment will continue to be tight, but we've also planned for significant growth going into 2025. And so we feel good about our overall supply chain capability.

AARON RAKERS: Yeah. And then as a quick follow-up, when I look at the sequential guidance and revenue at the midpoint-- let's call it about $680 million up sequentially-- could you help us frame how much of that is driven by client versus the data center piece of the business? Thank you.

LISA SU: Sure, Aaron. So certainly, if you look at the sequential guide, the largest contributor is the data center business. And that is now such a large piece of our business. It's over 50% of our business in Q3, and it will continue to grow in Q4.

Client segment, we also expect to perform well. I think we've done very well with our launches this year, both on the desktop Zen 5 launches as well as the notebook AI PC launches. So we expect growth there. And then the other segments, on a sequential basis, gaming and embedded would be more modest.

JOHN: And the next question comes from the line of Ross Seymore with Deutsche Bank. Please proceed with your question.

ROSS SEYMORE: Hey, guys. Thanks for letting me ask a question. And Lisa, congrats on the 10-year anniversary. Jean, when you gave the segment guidance, kind of playing off the last question, you didn't mention embedded as being up. And then, Lisa, you just talked about it being up less. Just curious. I know it's a kind of a muted recovery, but what's happening in embedded business implied in your fourth quarter guidance? And how are you thinking about 2025 just directionally?

LISA SU: Sure, Ross. Let me start, and then maybe Jean can add. Look, we've seen some improvement in the embedded business. You saw some improvement in the sequential in Q3. We expect some modest improvement in Q4.

What we're seeing is really a mix across the different subsegments in embedded. There are some segments that are stronger. So our test and emulation business actually did very well. We're ramping our new Versal platform there. Aerospace and defense did well, also continues to be relatively strong. Communications has not seen much recovery. So I would say that that's still fairly muted. And industrial is also a little bit on the softer side.

So between all of those, we expect a little bit of growth into Q4 and, let's call it, modest growth in 2025. But we're planning for it to be a little bit mixed amongst the segments. Maybe, Jean, if you want to add.

JEAN HU: No, I think you covered.

LISA SU: OK.

ROSS SEYMORE: Great. And I guess for my follow up, just on the EPYC business, you've done incredibly well there, Lisa. Have you seen any loosening up of the crowding out effect with the-- well, in your case, the Instinct side taking demand away from the EPYC side? Doesn't seem to have slowed you down much this year. You're still up, by my math, maybe 33% or so. And how are you thinking about that kind of CPU versus GPU dynamic from a customer spend perspective as we look into '25?

LISA SU: Yeah. Sure. Let me talk a little bit about the data center CPU business. Look, we've been extremely pleased with the progress there. I think the market environment has certainly gotten better over the last couple of quarters. We've seen some of the large cloud customers now adding to their data center capacity and refreshes. We've seen enterprises also start with some of their modernization activities.

Within that sort of market environment, our product portfolio has done extremely well. So here in the third quarter, the Zen 4 portfolio between Genoa and Bergamo was very strong. We saw the beginning ramp of our Zen 5 Turin capabilities. And then we also actually saw pretty good demand even on Zen 3 or our Milan family, just given the performance price ratios there. So overall, we have a very strong top-to-bottom stack, and we do see some strength in the overall server market, which adds to some of the AI opportunity that we have.

JOHN: And the next question comes from the line of Ben Reitzes with Melius Research. Please proceed with your question.

BEN REITZES: Yeah. If I could stick to the pattern of two, I want to follow up a little bit on servers, and thanks for the question. Going into next year, it seems like you're poised for more share gains. And I was just wondering. There is a perception that GPUs are cannibalizing the servers, but what are you seeing specifically in Turin?

With this momentum that you're talking about, is it the right CPU for a GPU? Are hyperscalers in particular liking the fact you can consolidate and create more room for AI gear? What is the particular catalyst that you're seeing out there with this upgrade? Thanks.

LISA SU: Sure, Ben. So I think when I look at Turin and the environment going into 2025, yeah, Turin is actually very well-optimized for the broad set of server workloads or traditional CPU workloads, including both scale-up and scale-out workloads. So I think that is very positive. We believe that there will be strong AI content on CPUs, and that helps Turin as well. It's also an important piece of the head node of any GPU configuration.

So overall, I think those are all good catalysts for us. And primarily, we're at this point where we're expanding the workloads in general as we are working with our largest cloud customers. And we're making strong progress in enterprise, which has been very important for us.

The enterprise sale takes longer. There's a significant amount of POC work and just making sure that CIOs are familiar with our offerings. But I think between our Zen 4, Zen 5 offerings now, we really have a very broad portfolio that satisfies the vast majority of the traditional CPU workloads.

BEN REITZES: OK. Thanks. And I just want to follow up on the PC market. You guys keep outperforming, but there's just general concerns of consumer weakness. And I was just wondering what you're seeing in the 4Q. And is there a risk of more than seasonal decline in the 1Q, or do you feel like things will keep going unabated in the PC market for you guys particularly?

LISA SU: Well, we see a few things in the PC market. So our business tends to be more consumer-weighted. So the second half of the year is usually stronger than the first half of the year. And this year, that's added to the fact that we have a couple of product launches. So we launched our desktop products with Zen 5 and our Ryzen 9000 Series. And we also launched our AI PC next-generation Ryzen AI 300 product.

So I think the combination of those two have given us, let's call it, a stronger than the normal second half of the year. I would expect there will be seasonality going into the first half. That's typical in our business. I think the main point is, this is the strongest PC portfolio we've had in our history, I think, across desktop and notebook.

And as we go into 2025, I think there's generally some optimism about the PC market, let's call it maybe growing mid-single digits. And within that, we have the AI PC catalyst as well as some of the Windows 10 end of support coming in 2025 as catalysts. So I think we feel good about the PC market in 2025, but we would expect some level of seasonality going into the first half of the year.

JOHN: And the next question comes from the line of Joshua Buchalter with TD Cowen and Company. Please proceed with your question.

JOSHUA BUCHALTER: Hey, guys. Thanks for taking my question. I'll keep it to one since I'm new here. On the data center GPU side, I realize everyone wants new customer announcements in 2025 guidance, but we're not going to get that, certainly not today.

That said, maybe you could speak to how much of a runway you think exists at existing customers and how much of a catalyst MI325X on CDNA3 and then MI355 on the new architecture can be to growing revenue and workload breadth internally and externally at these existing customers as we think about growth into next year at the data center GPU side. Thank you.

LISA SU: Yeah. Sure, Joshua. Thanks for the question. So look, maybe let me start again with this notion of we're very pleased with the progress that we've made in the data center GPU business. I mean, if you think about when we started the year, we were just launching MI300. We had talked about perhaps $2 billion of revenue in 2024.

And as we've gone through the last few quarters, what we've done is we've successfully completed a number of customer milestones. And those customer milestones include things like ensuring that we are at scale in data centers, meeting all of the reliability requirements, and so on, and so forth, as well as optimizing on specific workloads and ensuring that we have very, very performant out-of-the-box performance.

So I think we have gained a lot of confidence over the last couple of quarters just seeing how the customers have ramped. I know everyone would like it to go faster. We think it's actually going really well. And being able to talk about exceeding $5 billion of revenue in 2024, I think we feel really good about that.

Going into 2025, we feel great about the market. The market continues to be the place where there's significant CapEx investments. We feel great about our product portfolio. It is getting stronger with everything that we've learned off of the MI300 ramp.

And then in particular on the software side, we have greatly increased our customer support, customer engagements, out-of-box performance, open-source ecosystem, all of those components that are necessary to ensure that customers can run at scale, at performance, at a great TCO.

So all of those together, I think we have a good opportunity to grow at our current customers with the number of workloads. You heard Meta talking at our event about expanding from inference on their large language models with Llama 3.1 to some training workloads. Microsoft has also been a very, very great partner with that.

And then you should assume that we're working with all of the large customers out there. And many of them are very deep engagements with us to continue to optimize their software to our hardware. So those are the opportunities in front of us.

JOSHUA BUCHALTER: Thank you. And congrats on an amazing decade at AMD.

LISA SU: Thank you.

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

TIMOTHY ARCURI: Hi. Thanks. I had a quick one and then a more extensive question. So the first one is I wanted to ask about the September actuals for data center GPU. It seems like it was in the $1.5 billion range, and that would put December in the $2 billion range. Is that about right?

LISA SU: So that's a pretty granular question, Timothy. But maybe let me help you with this. We actually did better in the data center GPU business relative to our initial expectations. So you would imagine that the business was actually greater than $1.5 billion. I mean, we're actually seeing now our GPU business really approaching the scale of our CPU business.

TIMOTHY ARCURI: Great. And then just more broadly on the shape of next year, I mean, you see the market with these big customers. How do you see the shape of revenue off of that C Q4 number? Can you continue to grow off that level? Or do you see some risk maybe of a pause next year?

And I ask because MI355 is dropping into the existing infrastructure. And I think you did say you'll have a liquid-cooled option. But you do hear about some customers wanting rack scale. So I wonder if you worry a bit about there could be a pause before you get to rack scale in 2026. Thanks.

LISA SU: Yeah. I guess I would put it like this. I mean, look, this business is-- there's a lot of activity going on, and we're spending time with customers as they're building out their data centers. Some preferred air-cooled environments, liquid-cooled environments. Some want rack scale. Some are perfectly happy to go into their existing data center infrastructure. So there are lots of different variants.

What I would say about 2025 is we feel very good about the growth opportunities. I would say that it might be lumpy. In general, these are large customer acquisitions, and it's not always predictable exactly which quarters you would expect the significant buildout. But I think, overall, we feel good about the trajectory into 2025. I don't know if that's helpful.

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

JOE MOORE: Great. Thank you. I wonder if you could talk about the $5 billion of AI revenue, how that breaks down between training and inference, if you're able to assess that, and just-- I know you started off with most of the traction in inference, but you've seen some traction on the training side. Can you just talk about what that split looks like going forward as well?

LISA SU: Sure, Joe. So certainly, from the $5 billion that we're talking about, the early traction has been primarily with inference. Just given the strength of the product portfolio, MI300 is very, very well optimized for inference, given the memory capacity and memory bandwidth capabilities. But we have had some training adoption, and we expect that that will continue to grow as we go through the next few quarters. And so as we, let's call it, fast-forward a year, I would say we would have a fairly balanced portfolio between training and inference.

JOE MOORE: Great. Thank you.

JOHN: 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. I had two. So, Lisa, for the first one, how do you address this investor argument that MI is off to a great start, but spec-wise remains one year behind the industry leader? You're shipping something comparable to Hopper while they are starting to ship Blackwell next year. When you are at MI350, they will be on Blackwell Ultra or Rubin. So how do you see AMD closing that gap? And can you really gain share until that gap is closed?

LISA SU: Yeah. Yeah. In fact, I actually don't see that. So maybe let me state it in another way. I think MI300, when we launched it, was behind H100. H100 was in the market for a much longer time. And we have, with our accelerated roadmap, actually closed a good part of that gap. I think MI325 is a great product. It's going to compete very well with H200, and the MI350 series will compete very well with Blackwell.

In the overarching view of the world, frankly, the market continues to be constrained, particularly in the newer product generations. It takes a long time to go from, let's call it, shipping your first samples to actually ramping in volume production workloads. And I think one of the advantages that we have with our portfolio is that from a data center retrofit standpoint, it's actually a much easier ramp just given the infrastructure is the same.

So look, there are lots of opportunities across the set of AI workloads. We think this roadmap is actually strengthening over time, and that's the feedback that we're getting from our customers.

VIVEK ARYA: And for my follow up, Lisa, back to the question on the PC market. So it was up-- or the client was up 26% sequentially. I think you're guiding to some growth in Q4 as well. How do you see the state of the channel? And I ask that just because sell-through in Q3 was not that great. And I think Intel had guided down their Q3 to be flat or down.

So just how much of the growth that you are seeing is because of ASPs? How much of this is units? And as you look out to 2025, do you think the ASP strength that we have seen this year, can that mix benefit continue for you in '25 also?

LISA SU: Yeah. In fact, the way I would say it is, our client business has a few factors that may be slightly different from the overall market. So let me start with our desktop channel. I think our share is very high in the desktop channel. And we've done very well there. So we saw strength across that portfolio in Q3. We actually saw some of our highest sell-through. So it was a strong sell-through quarter. And obviously, it shows in the revenue results.

I think on the notebook standpoint, our business tends to be more heavily consumer-weighted. And so that is more of a second half story. There is good momentum around AI PCs. I mean, we've looked at some of the activation rates of our newest Ryzen AI 300 processors, and the activation rates are good. It's still very early in the AI PC cycle, but we are seeing some good momentum there.

So I think we'll see both units up, and the ASPs will depend a little bit on the mix between consumer and enterprise, and desktop and notebooks. But my earlier comments about the client segment, I think it is an opportunity for us. We are underrepresented in the client segment. And we see an opportunity to grow both in consumer as well as enterprise.

JOHN: And the next question comes from the line of Harlan Sur with JP Morgan. Please proceed with your question.

HARLAN SUR: Good afternoon. Thanks for taking my question. Lisa, on your core EPYC business, you did start to see recovery in enterprise last quarter. It was broad-based across many verticals.

You've also got a growing share dynamics as well, right? It sounds like that trend continued in the September quarter. Do you anticipate continued follow-through in terms of quarter on quarter growth this quarter for enterprise? And then maybe what sort of enterprise and general cloud demand trends are you seeing out of China?

LISA SU: Sure. So, Harlan, thanks for the question. We did see positive growth momentum here in the third quarter in enterprise. We are getting broader adoption. And we're seeing that growth both in on-prem deployments, as well as cloud third-party deployments. And so we're very happy with the adoption of our cloud instances.

When people are doing migrations, they're migrating from on-prem to cloud. They're migrating to AMD. Or they're migrating from older cloud instances to newer cloud instances. We're seeing migrations to AMD. So I think that helps us in both the enterprise segment as well as the cloud segment.

As we go forward into the fourth quarter, I think we are expecting another quarter of growth on a sequential basis for our server business and strength in both enterprise and cloud. And I think as we think about our opportunities going forward, the enterprise business is a place where we have been underrepresented.

I think our portfolio has strengthened. So the platforms that are being offered by the OEMs have broadened with not just Zen 4, but also with the Zen 5 portfolio launch. And so I think those are opportunities for us in 2025.

JEAN HU: Yeah. On your China question, we are underrepresented in China market in the server CPU side. So as Lisa said, that's another opportunity for us to continue to gain share.

HARLAN SUR: Perfect Thank you.

JOHN: 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 question. For the first one, I wanted to go back to what you said about the size of the data center GPU was in accordance, you said it was approaching the size of your compute business, which would put it around, what, under $1.7 billion, maybe a little more?

And number one, is that right? And if that is right, it implies that at $5 billion for the year, you'd actually be down in Q4. So it's probably got to be $5.2 or $5.3 billion for the full year just to be flat sequentially and more than that to get growth. So I guess the question number one, is that the size of the business that you're talking about in Q3? And is that the kind of level of growth you're thinking into Q4? How should we be thinking about that?

LISA SU: Right, Stacy. So first, a couple of things. You have to remember that in our data center segment, we have some other revenue that is not CPUs and GPUs. We have some FPGAs and other things.

But the question earlier was, was the revenue $1.5? And I said that it was greater than $1.5. So take that as a fundamental. And then as we talked about, we didn't guide in exact number for the data center GPU. We said exceed $5 billion. And so you should assume that it exceeds $5 billion.

STACY RASGON: OK. Thanks. For my follow up, you talked a little bit about lumpiness next year. Is there any kind of seasonality or something that I ought to be thinking about specifically in the Q1 for data center GPU? I know you have the MI300, I guess, ramping into production at that point. But should I be thinking about a seasonal-- seasonality into Q1, given the general statement on lumpiness?

LISA SU: I wasn't implying something about seasonality of the data center GPU business. I was implying more that if you think about the evolution of the business, it depends quite a bit on a specific number of customers. So these are large customers that drive deployments.

For example, the third quarter was a bit higher than we expected. That was driven by some additional customer demand, and we may see that type of lumpiness. So that was what I was implying. And we'll have to see how things evolve as we get into 2025.

JOHN: We have time for two more questions. The next question comes from the line of Harsh Kumar with Piper Sandler. Please proceed with your question.

HARSH KUMAR: Yeah. Hey, Lisa. First of all, congratulations on your 10-year anniversary. I just looked at the 10-year chart. The stock is around $3. So a heck of a job here. And also, congratulations on $5 billion in instant revenues.

So I wanted to ask the first question. This is one we get from investors all the time. In the coming year, let's say 2025, your key competitor will take most of the TAM off the AI market, the GPU market. Rough count, they'll take in something like $50-$60 billion. You'll get another $5 to $10 billion, call it.

So the question is, what do you think is the major hindrance? You've got chip-level compatibility. So does it boil down to the fact that you're just earlier in the game, you've been doing this just 12 months in a serious manner, or is there still a rack-level disparity? If you could just help us think about what the hindrances are to you becoming a major player here.

LISA SU: Yeah. Sure, Harsh. Thanks for the question and for the comment. Maybe let me say, I view them as opportunities.

If you remember, Harsh, and I think you do, our EPYC ramp from Zen 1, Zen 2, Zen 3, Zen 4, we had extremely good product even back in the Rome days. But it does take time to ensure that there is trust built, there is familiarity with the product set.

Although we're both GPUs, there are some differences, obviously, in the software environment. And people want to get comfortable with the workload ramp. So from a ramp standpoint, I'm actually very positive on the ramp rate. It's the fastest product ramp that I've seen overall.

And my view is, this is a multigenerational journey. We've always said that. We feel very good about the progress. I think next year is going to be about expanding both customer set as well as workload. And as we get into the MI400 series, we think it's an exceptional product. So all in all, the ramp is going well. And we will continue to earn the trust and the partnership of these large customers.

What I will say is customers are very, very open to AMD. And we see that everywhere we go. Everyone is giving us a very fair shot at earning their business, and that's what we intend to do.

HARSH KUMAR: Very helpful. Lisa, for my follow up-- and maybe Jean can help out here-- it's kind of understood that your gross margin for MI300, 325 will be below corporate margins. Could you help us maybe think of some framework on how those gross margins for Instinct might get to parity with your corporate business? I know you probably won't give us a revenue level or a time frame, but maybe you could help us frame in some manner that we could try and understand it.

JEAN HU: Yeah. We are very pleased with our overall revenue ramp of our data center GPU business. And our team not only support the revenue ramp and continue to improve the gross margin. Overall, it's below corporate average. And when you think about it, going into next year, of course, our top priority right now is really focusing on to address customer demand and to provide the TCO benefit, really increase significantly our market presence, and drive substantial revenue growth.

On the gross margin side, once we continue to ramp the revenue, we do think we'll have the opportunity to continue to improve gross margin. When you think about it, this is a data center business. Over in the longer term, it tends to be better than corporate average. We'll take some time to get there.

But when you look at our data center segment performance, we more than doubled the revenue year over year, but we tripled the operating income year over year. So I think that's how we are thinking about it, is really to drive the long-term growth and get market presence, at the same time drive a gross margin up.

JOHN: And our final question comes from the line of Thomas O'Malley with Barclays. Please proceed with your question.

THOMAS O'MALLEY: Hey, guys. Thanks for letting me on. I really appreciate it. Lisa, I just wanted to ask on mix in the MI300 platform into Q4. Obviously, you're launching the 325 series. Should we expect a significant contribution of the new product in Q4? Or is this a situation in which it's launching late in the year, and most of the impact is into Q1? I guess that's the first one.

And then the second one is just, if you look into next year, could you talk about where the end markets are in embedded for Xilinx? You gave some color on comm and a little on industrial there. But just where we're at in terms of starting Q4 and into Q1, are there any that are outperforming, underperforming? Any color there would be helpful. And thanks for sticking me in.

LISA SU: Yeah. Absolutely, Tom. So back to your first question. In terms of the mix in Q4, we would expect the majority of the mix will still be on my MI300. MI325 is going into production late in the quarter, and it will be more of a first quarter ramp. So that question.

And then in terms of the embedded business going to 2025, I would say that the trends that we see are similar to what I mentioned for the Q3, Q4 timing. We do expect that some of the markets will recover. We're expecting that it'll be a gradual recovery. So we see that.

And the strength that we see is in some of the test and emulation segment. We see some strength, certainly, in aerospace and defense. There was a little bit of recovery in automotive that we started to see. We're still waiting for comms and industrial. So we'll get more visibility on that as we go through the end of this year. But that's what we see currently.

THOMAS O'MALLEY: Thanks, Lisa.

JOHN: And at this time, we have reached the end of the question and answer session. And I would like to turn the floor back over to Mitch Haws for any closing comments.

MITCH HAWS: That concludes today's call. Thanks to all of you for joining us today.

JOHN: And thank you. That does conclude today's teleconference. We thank you for your participation. You may disconnect your lines at this time.