

AMD Fiscal Fourth Quarter and Full Year 2023 Financial Results

JOHN: Greetings. And welcome to the AMD fourth quarter and full-year 2023 conference call. At this time, all participants are in a listen-only mode. A brief 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 as a reminder, this conference is being recorded. It is now my pleasure to introduce to you Mitch Haws, vice president, investor relations. Thank you, Mitch. You may begin.

MITCH HAWS: Thank you, John. And welcome to AMD's fourth quarter and full-year 2023 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 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 Mark Papermaster, executive vice president and chief technology officer, will attend the Bernstein Tech, Media, Telecom and Consumer 1:1 Forum on Tuesday, February 28. And Jean Hu, executive vice president, chief financial officer, and treasurer, will attend the Wolfe Research semiconductor conference on Tuesday, February 15, and the Morgan Stanley Technology, Media, and Telecom conference on March 5.

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 in today. We finished 2023 strong as data center sales accelerated significantly throughout the year despite the mixed demand environment. As a result, we delivered record data center segment annual revenue and strong top-line and bottom-line growth in the fourth quarter, driven by the ramp of Instinct AI accelerators and robust demand for EPYC server CPUs across cloud, enterprise, and AI customers.

Looking at our financial results, fourth quarter revenue increased 10% year over year to \$6.2 billion, driven by significant double-digit percentage growth in our data center and client segments. On a full-year basis, annual revenue declined 4% to \$22.7 billion as record data center and embedded segment annual revenue was offset by lower client and gaming segment revenue.

Importantly, data center and embedded segment annual revenue grew by \$1.2 billion and accounted for more than 50% of revenue in 2023 as we gain server share, launched our next generation Instinct AI accelerators, and maintained our position as the industry's largest provider of adaptive computing solutions.

Turning to the fourth quarter business results, data center segment revenue grew 38% year over year and 43% sequentially to a record \$2.3 billion. Server CPU and data center GPU sales both set quarterly and annual revenue records as sales of our data center products accelerated throughout the year. We gained server CPU revenue share in the quarter, driven by significant double-digit percentage growth in fourth gen EPYC processor revenue and demand for our third gen EPYC processor portfolio.

In cloud, while the overall demand environment remains soft, server CPU revenue increased year over year and sequentially as North American hyperscalers expanded fourth gen EPYC processor deployments to power their internal workloads and public instances. Amazon, Alibaba, Google, Microsoft, and Oracle brought more than 55 AMD-powered AI, HPC, and general purpose cloud instances into preview or general availability in the fourth quarter.

Exiting 2023, there were more than 800 EPYC CPU-based public cloud instances available. We expect this number to grow in 2024 based on the leadership, performance, efficiency, and features of our EPYC CPU portfolio.

In enterprise sales, accelerated by a significant double-digit percentage in the quarter as we built momentum with Forbes 2000 customers, we closed multiple wins with large financial, energy, automotive, retail, technology, and pharmaceutical companies, positioning us well for continued growth based on expanded production deployments planned for 2024.

A growing number of customers are adopting EPYC CPUs for inferencing workloads, where our leadership throughput performance delivers significant advantages on smaller models like Llama 7b as well as to power head nodes and large training and inference clusters.

Looking ahead, customer excitement for our upcoming Turin family of EPYC processors is very strong. Turin is a drop-in replacement for existing fourth gen EPYC platforms that extends our performance, efficiency, and TCO leadership with the addition of our next-Gen Zen 5 core, new memory expansion capabilities, and higher core counts. Internal and end customer validation work is progressing to plan, with Turin on track to deliver overall performance leadership as well as leadership on a per-core or per-watt basis across a wide range of workloads when it launches later this year.

Turning to our broader data center portfolio, our data center GPU business accelerated significantly in the quarter, with revenue exceeding our 400 million expectation, driven by a faster ramp for MI300x with AI customers. We launched our MI300 accelerator family in December with strong partner and ecosystem support from multiple large cloud providers, all the major OEMs, and many leading AI developers.

MI300x GPUs deliver leadership generative AI performance by combining our high performance CDNA 3 architecture with industry-leading memory bandwidth and capacity. Customer response to MI300 has been overwhelmingly positive. And we are aggressively ramping production to support the dozens of cloud, enterprise, and supercomputing customers deploying instinct accelerators.

In cloud, we are working closely with Microsoft, Oracle, Meta, and other large cloud customers on instinct GPU deployments, powering both their internal AI workloads and external offerings. For enterprise customers, HPE, Dell, Lenovo, Supermicro, and other server vendors are on track to launch differentiated MI300 platforms later this quarter, with strong demand from multiple enterprise customers.

In HPC supercomputing, we shipped the majority of AMD Instinct MI300A accelerators for the El Capitan supercomputer in the fourth quarter and expect to complete shipments this quarter for what is expected to be the world's fastest supercomputer when it comes online later this year. We also closed new Instinct GPU wins in the quarter, including the flagship system at the German High-Performance Computing Center HLRS, as well as what is expected to be one of the world's most powerful enterprise supercomputers for energy company Eni.

On AI software development, we made significant progress expanding the ecosystem of AI developers working on AMD platforms with the release of our ROCm 6 software suite. The ROCm 6 stack significantly increases performance in key generative AI workloads, adds expanded support and optimizations for additional frameworks and libraries, and simplifies the overall developer experience.

The additional functionality and optimizations of ROCm 6 and the growing volume of contributions from the open-source AI software community are enabling multiple large hyperscale and enterprise customers to rapidly bring up their most advanced large language models on AMD Instinct accelerators.

For example, we are very pleased to see how quickly Microsoft was able to bring up GPT-4 on MI300x in their production environment and roll out Azure private previews of new MI300 instances aligned with the MI300x launch. At the same time, our partnership with Hugging Face, the leading open platform for the AI community, now enables hundreds of thousands of AI models to run out of the box on AMD GPUs, and we are extending that collaboration to our other platforms.

Looking ahead, our prior guidance was for data center GPU revenue to be flattish from Q4 to Q1 and exceed \$2 billion for 2024. Based on the strong customer pull and expanded engagements, we now expect data center GPU revenue to grow sequentially in the first quarter and exceed \$3 and 1/2 billion in 2024. We have also made significant progress with our supply chain partners and have secured additional capacity to support upside demand.

Turning to our client segment, revenue was \$1.5 billion, an increase of 62% year over year and flat sequentially. We launched our latest generation Ryzen 8000 series notebooks and desktop processors in January, including our Ryzen 8040 mobile series that combine leadership compute performance and energy efficiency with an updated NPU that delivers up to 60% more AI performance compared to our prior generation that was already industry-leading. Acer, ASUS, HP, Lenovo, MSI, and other large PC OEMs will all offer notebooks powered by our Ryzen 8000 series processors, with the first systems expected to go on sale in February.

To further our leadership in AI PCs, we launched our Ryzen 8000 G-series processors earlier this month, which are the industry's first desktop CPUs with an integrated AI engine. Millions of AI PCs powered by Ryzen processors have shipped to date, and Ryzen CPUs power more than 90% of AI-enabled PCs currently in market.

Our work with Microsoft and our PC ecosystem partners to enable the next generation of PCs expanded significantly in the quarter. We are aggressively driving our Ryzen AI CPU roadmap to extend our AI leadership, including our next-gen Strix processors that are expected to deliver more than 3x the AI performance of our Ryzen 7040 series processors.

Strix combines our next-gen Zen 5 core with enhanced RDNA graphics and an updated Ryzen AI engine to significantly increase the performance, energy efficiency, and AI capabilities of PCs. Customer momentum for Strix is strong, with the first notebooks on track to launch later this year.

Looking at 2024, we are planning for the PC TAM to grow modestly year on year, weighted towards the second half, as AI PCs ramp. We continue to see strong growth opportunities for our client business as we ramp our current products, extend our AI PC leadership, and launch our next wave of Zen 5 CPUs.

Now, turning to our gaming segment. Revenue declined 17% year over year and 9% sequentially to \$1.4 billion as lower semi-custom revenue was partially offset by increased sales of Radeon GPUs. Semi-custom SoC sales declined in line with our projections in the quarter. Going forward, we now expect annual revenue to decline by a significant double-digit percentage year over year as supply caught up with demand in 2023, and we entered the fifth year of what has been a very strong console cycle.

In gaming graphics, revenue grew both year over year and sequentially, driven by strong demand in the channel for both our Radeon 6000 and Radeon 7000 series GPUs. We expanded our Radeon 7000 GPU series with the launch of new RX 7600 XT series enthusiast desktop GPUs earlier this month that offer leadership price performance for 1080p gaming. We also launched new open-source FidelityFX Super Resolution 3 software that can deliver significantly higher gaming frame rates on both GPUs and APUs.

Turning to our embedded segment, revenue decreased 24% year over year and 15% sequentially to \$1.1 billion as customers focused on reducing their inventory levels. We expanded our embedded portfolio in the quarter with new leadership solutions for key markets. We launched new Versal Prime adaptive SoCs for the aerospace, test and measurement, health care, and communications markets that deliver industry-first support for DDR5 memory and increased DSP capability compared to our prior generation.

In automotive, we launched new Versal SoC solutions that bring industry-leading AI compute capabilities and advanced safety and security features to next-generation vehicles. We also launched Ryzen-embedded processors with unmatched performance and features for industrial automation, machine vision, robotics, and edge server applications.

Looking at 2024, we expect overall embedded demand will remain soft through the first half of the year as customers continue to focus on normalizing their inventory levels. Longer term, we're very confident in the growth trajectory of our embedded business as our expanded product portfolio drove more than 10 billion of design wins in 2023, an increase of more than 25% compared to 2022.

In summary, I'm very pleased with our fourth quarter and full year results. For 2024, we expect the demand environment to remain mixed, with strong growth in our data center and client segments offset by declines in our embedded and gaming segments. Against this backdrop, we believe we will deliver strong annual revenue growth and expand gross margin driven by the strength of our Instinct, EPYC, and Ryzen product portfolios.

Taking a step back, we believe AI is a once-in-a-generation transition that will reshape virtually every portion of the computing market, starting in the data center and then expanding into PCs and across multiple embedded markets. We have built excellent customer traction based on the strength of our multiyear AI hardware and software roadmaps, and we see clear opportunities to drive our next wave of growth as we deliver leadership AI solutions across our portfolio.

In the data center, we see 2024 as a start of a multiyear AI adoption cycle, with the market for data center AI accelerators growing to approximately \$400 billion in 2027. Customer deployments of our Instinct GPUs continues accelerating, with MI300 now tracking to be the fastest revenue ramp of any product in our history and positioning us well to capture a significant share over the coming years based on the strength of our multigeneration instinct GPU roadmap and open-source ROCm software strategy.

In PCs, we are focused on delivering our long-term roadmaps with leadership Ryzen AI NPU capabilities to enable differentiated experiences as Microsoft and our other software partners bring new AI capabilities to PCs starting later this year. At the same time, we are rapidly driving leadership AI compute capabilities across the full breadth of our embedded product portfolio.

This is an incredibly exciting time for the industry, and an even more exciting time for AMD, as our leadership IP, broad product portfolio, and deep customer relationships position us well to deliver significant revenue growth and earnings expansion over the next several years. Now, I'd like to turn the call over to Jean to provide some additional color on our fourth quarter and full-year financial 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 first quarter of fiscal 2024. AMD executed well in 2023 despite all the mixed market demand environment.

Delivering revenue of \$22.7 billion and earnings per share of \$2.65, we drove year-over-year revenue growth in our embedded and data center segments. In addition, we successfully launched our AMD Instinct MI300 GPUs, positioning us for a strong ramp in 2024 in the AI market.

For the fourth quarter of 2023, revenue was \$6.2 billion, growing 10% year over year as revenue growth in the data center and client segments was partially offset by the lower revenue in our embedded and gaming segments. Revenue was up 6% sequentially, primarily driven by the ramp of AMD Instinct GPUs across several leading customers, and the higher revenue from EPYC server processors, partially offset by the decline in embedded and gaming segment revenues.

Gross margin was 51%, flat year over year, with higher revenue contribution from the data center and client segments offset by lower embedded segment revenue. Operating expenses were \$1.7 billion, an increase of 8% year over year, as we invested in R&D and marketing activities to support our significant AI growth opportunities.

Operating income was \$1.4 billion, representing a 23% operating margin. Taxes, interest expense, and other was \$163 million. For the fourth quarter of 2023, diluted earnings per share was \$0.77, an increase of 12% year over year.

Now, turning to our reportable segments. Starting with the data center segment, revenue was \$2.3 billion, up 38% year over year and 43% sequentially, driven by strong growth of both AMD Instinct GPU and the fourth generation AMD EPYC CPU sales.

Data center segment operating income was \$666 million, or 29% of revenue, compared to \$444 million, or 27%, a year ago. Higher operating income was primarily due to operating leverage driven by higher revenue.

Client segment revenue was \$1.5 billion, up 62% year over year, driven by Ryzen 7000 series CPU sales. Client segment operating income was \$55 million, or 4% of revenue, compared to an operating loss of \$152 million a year ago, driven by higher revenue.

Gaming segment revenue was \$1.4 billion, down 17% year over year and 9% sequentially, due to a decrease in semi-custom revenue, partially offset by increase in Radeon GPU sales. Gaming segment operating income was \$224 million, or 16% of revenue, compared to \$266 million, or 16%, a year ago.

Embedded segment revenue was \$1.1 billion, down 24% year over year and 15% sequentially, as customers continue to work down their inventory levels. Embedded segment operating income was \$461 million, or 44% of revenue, compared to \$699 million, or 50%, a year ago.

Turning to the balance sheet and cash flow. During the quarter, we generated \$381 million in cash from operations. And the free cash flow was \$242 million. Inventory decreased sequentially by 94 million to 4.4 billion. At the end of the quarter, cash, cash equivalents, and short-term investment was strong at \$5.8 billion.

In the fourth quarter, we repurchased 2 million shares and returned 233 million to shareholders. For the year, we repurchased 10 million shares and returned 985 million to shareholders. We have 5.6 billion in remaining share repurchase authorization.

Now, turning to our first quarter of 2024 outlook. We expect revenue to be approximately \$5.4 billion plus or negative \$300 million. Sequentially, we expect data center segment revenue to be flat, with a seasonal decline in server sales offset by strong data center GPU ramp.

Embedded revenue to decline as customers continue to work down their inventory levels. Client segment revenue to decline seasonally. And in the gaming segment, as we enter the fifth year of what has been a very strong gaming cycle, and given current customer inventory levels, we expect revenue to decline by significant double-digit percentage.

Year over year, we expect data center and client segment revenues to increase by strong double-digit percentage given the strength of our product portfolio and the share gain opportunities. Embedded segment to decline, and the gaming segment revenue to decline by significant double-digit percentage.

In addition, we expect first quarter non-GAAP gross margin to be approximately 52%, non-GAAP operating expenses to be approximately \$1.73 billion, non-GAAP effective tax rate to be 13%. And a diluted share count is expected to be approximately 1.63 billion shares.

While we are not providing specific full year guidance for 2024, let me provide some color. Directionally, for the year, we expect 2024 data center and the client segment revenue to increase, driven by the strength of our product portfolio and the share gain opportunities.

Embedded segment revenue to decline, and the gaming segment revenue to decline by significant double-digit percentage. We expect to expand gross margin in 2024 and continue to invest to address the large AI opportunities while driving operating model leverage to deliver earnings per share growth faster than top line revenue growth.

In closing, we delivered solid financial results in 2023, further strengthening our product portfolio and establishing ourselves as a leading provider of data center GPUs for AI. We are very well-positioned to build on this momentum and deliver strong financial performance in 2024 and beyond. With that, I'll turn it back to Mitch for the Q&A session.

MITCH HAWS: Thank you, Jean. John, we're happy to 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 queue. You may press star 2 if you'd like to remove a question from the queue.

For participants using speaker equipment, it may be necessary to pick up your handset before pressing the star keys. We ask that you please limit yourself to one question and one follow up. Thank you. One moment please, while we pull for questions. And the first question comes from the line of Aaron Rakers from Wells Fargo. Please proceed with your question.

AARON RAKERS: Yeah. Thanks for taking the question. Just kind of framing the outlook and the guidance for this calendar first quarter, I guess the first question is, can you help us-- on a relative basis, the \$400 million of data center GPU revenue that you expected in Q4, what did that ultimately fell out to be? And then on the guidance into 1Q, can you help us appreciate what seasonal is defined as, as we think about the server business into the 1Q guide?

LISA SU: Sure, Aaron. Let me start and then see if Jean has something to add. So relative to the data center GPU business, we were very pleased with the performance that we saw in the fourth quarter. It was always going to be a very back end quarter, weighted, as we were ramping the product.

And we saw MI300A, our HPC product, actually ramp very well. And then we saw MI300x, the AI product, actually exceed our expectations based on the strong customer demand, the way the qualifications went, and then the manufacture ramp. So we were over \$400 million for that business in the fourth quarter.

And then going into the first quarter, as we look at the business, server seasonality, call it something around-- let's call it high single digit, low double digit. There are also some other pieces of the data center business. I think the key piece of it is we had originally expected the ramp to be a little bit more shallow of our MI300X. And what we're seeing now is, yeah, the supply chain is operating really well and the customer demand is strong. And so we will see MI300X increase as we go into the first quarter. And things are going relatively well.

JEAN HU: Yeah. I mean, I'll give you some color, Aaron. I'll give you some color about the client seasonality and others. So client is very similar to server. Typically, Q1 is high single digit to low double digit. That's consistent with the past.

On the embedded side, it's very consistent with what we said in the past and consistent with what you see in the industry, is embedded business is going through a bottoming process. And we think Q1, it will have a low double-digit sequential decline. That's embedded.

On the gaming side, Lisa mentioned during her prepared remarks, is we have the late stage of product cycle in the year five of gaming console. But at the same time, we also have an inventory at the customers. So the combination of those impact, we expect the Q1 gaming sequential declines probably more than 30%. So hopefully, that helps you a little bit.

AARON RAKERS: Yeah. Very helpful, Jean. And as a quick follow up, I'm just curious, the traditional server demand that you see-- I know when we look at server CPU, shipments are down, north of 20% year over year. Are you seeing any signals-- or how are you thinking about a recovery in that traditional, call it, non-AI general purpose server market as we move through '24?

LISA SU: Sure, Aaron. So look, I think-- I agree with your characterization of the 2023 demand, although we did see some strong progress in the second half of the year, especially as customers in cloud and enterprise adopted our Genoa and our Zen 4 family.

So going into 2024, I would say the traditional server market is probably still mixed, especially into the first half of the year. There's still some cloud optimization going on as well as enterprise being a little bit cautious.

That being the case, though, we also see opportunities for us to continue to grow share in the traditional server business. I think our portfolio is extremely strong. The adoption of Genoa and Bergamo, as well as our new Siena product lines, are getting a lot of traction. And then we also see Turin, our Zen 5 product, coming in the second half of the year. So even in a mixed demand environment, I think we're bullish on what traditional server CPUs can do in 2024.

JOHN: 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 give us a little bit of sense in terms of the milepost that you're marching toward on this \$400 billion TAM that you have for 2027. For example, do you think you can gain share at a rate that's similar to the rate that you gained share for server CPU? Or I guess maybe that's a different way. Is it reasonable to look at your consumer GPU share of 20-plus percent-- is that a reasonable bogey, or do you have aspirations higher than that perhaps?

LISA SU: Yeah. Thanks, Tim, for the question. I would say a couple of things. First of all, we're really pleased with the progress that we've made in our data center GPU business. I think the ramp that we've seen, the customer traction that we've seen, even in the last few months, I think has been great. And that gives us a lot of confidence in the ramp of this business. I think the beauty of the AI market here is it's growing so quickly that I think we have both the market dynamic as well as our ability to gain share in that framework.

The point I will make is, our customer engagements right now are all quite strategic. Dozens of customers with multigenerational conversations. So as excited as we are about the ramp of MI300-- and frankly, there's a lot to do in 2024-- we're also very excited about the opportunities to extend that into the next couple of years, out into that, '25, '26, '27 time frame. So I think we see a lot of growth. I think it's a little early to make market share projections. But I would say, it's a significant growth driver given the market demand as well as our own product capabilities.

TIMOTHY ARCURI: Thanks a lot. Jean, I guess as a follow up, I know that you don't want to guide the full year, but I'm wondering if I can pin you down just a touch on maybe a milepost that you're marching to for 2024 growth is up 20% for the whole company. Is that a reasonable target? And then I guess within data center, if you just add the incremental data center GPU revenue and you assume that the server business grows a little bit, it seems like that should maybe double year over year. But I'm kind of wondering if you can give us any ranges on those numbers. Thanks.

JEAN HU: Hi, Tim. Thank you for the question. Yeah, we're not guiding a year. It's very early of the year, literally January. I think the way to think about it is, Lisa mentioned during her prepared remarks, we feel pretty good about both our data center and the client business to grow in 2024. Of course, the largest incremental revenue opportunities are going to come from data center between both the server side gaining more share and the data center GPU side with the significant ramp of our MI300. I think that's how we think about it. We do have a headwind from gaming segment. We do think, year over year, we will see very significant double-digit decline on the gaming segment. And at the same time, embedded is going through the bottoming process. We do think the second half will see the recovery. So those are the puts and takes. I can talk about it.

JOHN: And the next question comes from the line of Matt Ramsay with TD Cowen. Please proceed with your question.

MATT RAMSAY: Thank you very much. Good afternoon. Lisa, I wanted to ask. I mean, there's been so much focus and scrutiny, as there should be, on this really exciting progression with MI300. And I mean, we've progressed over the last six months from, I think, some doubts in the investment community as to software and your ability to ramp the product. And now, you've proven that you're ramping it with, I think you said, dozens of customers right across different end markets.

So what I'm interested in hearing a little bit more about-- and you guys have been open about what some of the forward programs in your traditional server business look like from a roadmap perspective. I'd be interested to hear how you're thinking about the roadmap in your MI accelerator family. Is it they're going to continue to be parts that are CPU and GPU together, or is it primarily a GPU-only roadmap? What kind of cadence are you thinking about? I'd just be-- any kind of color you can give us on some of the forward roadmap trajectory for that program would be really helpful. Thanks.

LISA SU: Yeah. Sure, Matt. So I appreciate the comments. I think the traction that we're getting with the MI300 family is really strong. I think what's benefited us here is our use of chiplet technologies, which has given us the ability to have both the APU version as well as the GPU version. And we continue to use that to differentiate ourselves, and that's how we get our memory bandwidth and memory capacity advantages.

As we go forward, you can imagine, like we did in the EPYC time frame, we plan multiple generations in sequence. That's the way we're planning the roadmap. One of the things I will note about the AI accelerator market is the demand for compute is so high that we are seeing an acceleration of the roadmap generations here, and we are similarly planning acceleration of our roadmap.

I would say that we'll talk more about the overall roadmap beyond MI300 as we get into later this year. But you can be assured that we're working very closely with our customers to have a very competitive roadmap for both training an inference that will come out over the next couple of years.

MATT RAMSAY: No, thank you for that, Lisa. Just as a follow up, I guess one of the questions that I've been getting a lot in different forums is, with respect to the \$400 billion TAM that you guys have laid out for 2027, maybe you could give us a little look under the hood as to, I guess, the-- I've gotten hundreds versions of the same question, which is how the heck did you come up with that number?

So if you could give us a little bit more in terms of-- are we talking about systems and accelerator cards? Are we talking about just the silicon? Are we talking about full servers? And what kind of unit assumptions-- any kind of thing that you can give us on market sizing, or what gives you the visibility so early into this generative AI trend to give a precise number three years out, that would be really, really helpful. Thank you.

LISA SU: Sure. Well, Matt, I don't know how precise it is, but I think we said approximately \$400 billion. But I think what we need to look at is growth rate and how do we get to those growth rates. I think we expect units to grow sort of substantial double-digit percentage. But you should also expect that content is going to grow. So if you think about how important memory and memory capacity is as we go forward, you can imagine that we'll see acceleration there and just the overall content as we go to more advanced technology nodes. So there's some ASP uplift in there.

And then what we also do is, we're planning longer term roadmaps with our customers in terms of how they're thinking about the size of training clusters, the number of training clusters, and then the fact that we believe inference is actually going to exceed training as we go into the next couple of years, just given-- as more enterprises adopt. So I think as we look at all of those pieces, I think we feel good that the growth rate is going to be significant and sustained over the next few years.

In terms of what's in that TAM, it really is accelerator TAM. So within accelerators, there are certainly GPUs, and there will also be some ASICs that are-- or other accelerators that are in that TAM. As we think about the different types of models, from smaller models to fine-tuning of models, to very-- the largest large language models, I think you're going to need different silicon for those different use cases. But from our standpoint, GPUs are still going to be the compute element of choice when you're talking about training and inferencing on the largest language models.

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 think you talked about the MI300 cloud workloads being split between the more customer-facing workloads versus internal. Can you talk about how you see the breakdown of that? And how is your ecosystem progressing? This is a brand new chip. It seems impressive that you're able to support a broad range of customer-facing workloads in cloud.

LISA SU: Yeah. Sure, Joe. So yes, look, we are really happy with how the MI300 has come up. And we've now deployed. And working with a number of customers, what we have seen is certainly ROCm 6 has been very important as well as the direct optimization with our top cloud customers.

We always said that the best way of optimizing the software is working directly on the most important workloads. And we've seen performance come up nicely, which is what we expect, frankly, with the GPU capabilities that we would have to do some level of optimization. But that optimization has gone well.

I think to your broader question, the way I look at this is there are lots of opportunities for us to work directly with large customers, both on the cloud side as well as on the enterprise side, who have specific training and inferencing workloads. Our job is to make it as easy as possible for them. And so our entire tool chain, all of our-- the overall ROCm suite has really gone through significant progress over the last six to nine months.

And then we're also getting some nice support from the open-source community. So the work that Hugging Face is doing is tremendous in terms of just real-time optimization on our hardware. Our partnership with OpenAI on Triton and our work across a number of these open-source models has helped us actually make very rapid progress.

JOE MOORE: Great. And for my follow up, I guess, a lot of the forecasting of your business that I'm hearing is coming from supply chain, and we're sort of hearing, AMD's building x in Asia. I guess how would you ask us to think about that? Are you looking at being sold out for the year, and so the supply chain would be closer to revenue? Are you building for the best case scenario? Just I worry about sometimes expectations when people hear these supply chain numbers, and I'm just curious how you would bridge the gap.

LISA SU: Yeah. So I mean, Joe, I think we updated our revenue expectations this quarter from our original number of \$2 billion to \$3 and 1/2 billion to try to give some bounding on some of the discussion out there. The way to think about the \$3 and 1/2 billion is, these are customers that we're working with who have given us firm commitments on what they need. As you know, the lead times on these products are quite long. So it's important to have those forecasts in early. And we have a strong order book. So that gives us good confidence to exceed the \$3 and 1/2 billion.

From a supply chain standpoint, our goal is always to build more supply. And so from that standpoint, we have also worked with our supply chain partners and secured significant capacity. Think about it as first half capacity is tight and more comes on in the second half of the year. But we've certainly made more progress there. So we do have more supply, and we're going to continue to work with our customers on their deployments. And we'll update that number as we go through the year.

JOHN: And the next question comes from the line of Toshiya Hari with Goldman Sachs. Please proceed with your question.

TOSHIYA HARI: Hi. Thank you for taking the question. I had one on the MI300 as well, Lisa. I guess, first of all, how should we think about the quarterly trajectory beyond Q1? You talked about Q1 being up sequentially. Is it fair to assume a straight line as we progress through the balance of the year, or is it more second half skewed? How should we think about that?

And I guess more importantly, some of the cloud potential customers that have yet to officially sign up for sign off on the MI300-- I guess, what's the sticking point? Is it just a function of time, and you just need a little bit more time to go back and forth and tweak things? Or is there a software concern? I guess, what's holding them back at this point?

LISA SU: Yeah. Yeah, Toshiya. Thanks for the question. So first, on the MI300 trajectory, I think you would expect that revenue should increase every quarter from now through the end of the year, but it will be a bit more second-half weighted. And part of that is just customers, as they're finishing up their qualifications in their lines as well as how our supply chain is ramping. So yeah, it should increase each quarter, but be a bit more second-half weighted.

And then to your comment about customers, look, we're engaged with all of the large customers. These are all folks that know us really well given our deep relationships in EPYC. I think people just have different adoption cycles as they consider what they're trying to do in their roadmap. But I view this as still very, very early innings for us in this space.

And I think a question was asked earlier. I think the key is this is not just about an MI300 conversation, but it really is about our long-term multigenerational roadmap. And so that's the context on which we're working with our largest customers as well as-- as you know, there's a lot of demand coming from folks that are more AI-centric and not necessarily the typical cloud customers, but more enterprise, or, let's call it, AI-specific companies that we're also very well-engaged with.

TOSHIYA HARI: Got it. That's super helpful. And then as my follow up, maybe one for Jean on the gross margin side. You're guiding Q1 to 52%.

I'm curious. Again, I'm sure you're not going to give quantitative guidance beyond Q1. But how to think about the trajectory for Q2 and beyond-- I'm pretty sure you're working through some kinks as it pertains to the Instinct ramp. Hopefully, that improves over time.

So that should be a tailwind. FPGAs, perhaps the second half turns for the better. And you've got server CPU volume growth throughout the year. So it feels like you've got multiple tailwinds as we think about gross margin progression on a sequential basis, but what are the potential headwinds as we move throughout 2024? Thank you.

JEAN HU: Yeah, Toshiya. Thank you for the question. Yeah, you're absolutely right. We have some puts and takes that impact our gross margin. We guided the Q1 120 basis point higher than Q4 sequentially primarily because the higher data center contribution actually more than offset the decline of embedded business in Q1.

Going forward, the way to think about it is, as you said, is the major driver is going to be data center business is going to grow much faster than other segment. That makes change that will help us to expand the gross margin nicely. I think you also are spot on on the embedded coming back in second half, which will be a tailwind with the data center GPU.

We are at a very early stage of ramp. We are improving testing time yield and continue to expand gross margin. And we expect it to be accretive to corporate average. So those are all the tailwinds coming in the second half.

I would say, the headwinds aside continue to be in the first half, where we see embedded business-- not only Q1, we see sequential decline. Q2 probably are going to be sequentially flattish versus Q1. That is a headwind for us because it does have a very nice gross margin. But overall, we feel pretty good about the trajectory of the gross margin improvement, especially second half.

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

ROSS SEYMORE: Hi. Thanks for let me ask a question. I wanted to get into the competitive environment, first on the Instinct side of things, how that's going-- it doesn't seem to be slowing down your ramp whatsoever-- but then also on the straight server CPU side of things. Lisa, you said you're gaining share in that area. But as we think about future roadmaps, pricing incentives, those sorts of things, any meaningful change in the competitive environment that you're seeing through the 2024?

LISA SU: Sure, Ross. So look, I think the environment for us is always competitive. So I think that has not changed. If I look at the Instinct side, I think we have-- I think we've shown that MI300 and our roadmap are actually very competitive. There are some places where-- let's call it-- it's more even in the training environment. But as we look at the inferencing environment, we think we have significant advantages, and that's shown through in some of our customer work. So we think for both training and inference, we will continue to be very competitive.

And then as you go into the CPU side, again, from our view, with each generation of EPYC, we've continued to gain share. I think we exited the fourth quarter at record share for AMD. And we're still quite underrepresented in enterprise. So I think there's an opportunity for us to continue to gain share as we go through 2024.

From a competitive standpoint, what we see is Zen 4 is extremely competitive right now with Genoa, Bergamo, Siena. And as we go into turn, we're deep in the design in cycle for Zen 5 and Turin, and we feel very good about how we're positioned.

ROSS SEYMORE: Thanks for that. I guess as my follow up, on the data center side, another theme that's been pervasive throughout 2023, at least, was the GPU side crowding out the CPU side. You mentioned that there was still a little bit of cloud digestion going on within your EPYC business, but where do you see that standing? I know you're going to gain share, et cetera, but you guys clearly benefit from the Instinct side of the data center GPU side, but what about on the CPU side of things? Is that headwind now behind us, or is it still an issue?

LISA SU: I think we expect the CPU business, from a market standpoint, to grow, Ross, as we go into 2024. I think the rate and pace of growth will depend a little bit on the macro and just overall CapEx trends. But from our standpoint, we are starting to see some of our larger customers plan their refresh cycles.

There's a lot of, let's call it, older equipment that has yet to be refreshed. And the value proposition for refresh is so strong because the energy efficiency and the footprint of the newer generations are so much better than the four or five-year-old infrastructure that we do see that refresh cycle happening as we get into 2024. I think the exact timing, we'll have to understand more as the market evolves as we go through the year.

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 question. So first one, Lisa, you gave us a \$2 billion-plus number for MI before. Now, you have raised it to over \$3 and 1/2 billion. And I'm curious, what drove the change? Was it incremental demand signals? Was it supply? And can you supply more if, let's say, demand is \$4 or \$5 or \$6 billion? What is the limitation?

And related to that, on the competition side, your competitor will launch their B100 later in the year. Do you think that will change the competitive landscape in any way?

LISA SU: Yes. Sure, Vivek. So I think what we said is, as we went from \$2 to \$3 and 1/2 billion, it really is mostly customer demand signals. So as orders have come on the books, and as we've seen programs move from, let's call it, pilot programs into full manufacturing programs, we have updated the revenue forecast.

As I said earlier, from a supply standpoint, we are planning for success. And so we worked closely with our supply chain partners to ensure that we can ship more than \$3 and 1/2 billion, substantially more, depending on what customer demand is as we go into the second half of the year.

And then in terms of, again, roadmaps, as I said, we're very focused on a competitive roadmap, what the next generations are beyond MI300. So I do believe that we have a strong roadmap in place and continue to work with our customers to adopt our roadmap as quickly as possible.

VIVEK ARYA: All right. And a longer term question, Lisa. If I look at the success that AMD has enjoyed, it's many factors. But a few of them included your early adoption of chiplets and then the strong partnerships you have had with TSMC.

But now, we are seeing your x86 competitor, Intel, also adopt chiplet or tile technology, as they call it. And then I think recently, the manufacturing update they gave, they said they are two years ahead in terms of incorporating gate all around and backside power delivery. So let's assume they are right and they have either caught up to TSMC or maybe they are ahead. What impact does that have on AMD in the medium to longer term?

LISA SU: Yeah. Sure, Vivek. Look, we're always looking at what's next, right? So on the chiplet technology, I mean, we're on the fourth generation of the chiplet technologies. I think we've learned a lot about how to optimize performance there. We are very aggressive with our adoption of leading edge technology as it's needed. But I think those are only a few of the pieces.

We're also focused on continuing to innovate on architecture and design. So I think the longer term question that you ask is-- I think we're expecting that the competition is going to be on a similar process technology. And even in that case, I think we feel like we have a very strong roadmap going forward, and we'll continue to drive both the CPU and the GPU roadmap very aggressively.

JOHN: And the next question comes from the line of Harsh Kumar with Piper Sandler. Please proceed with your question.

HARSH KUMAR: Yeah. Hey. Thanks for letting me ask a question, guys. I had two questions. Let me start off with the accelerator side.

The question we get a lot from our customers is they want to understand the value proposition of the MI300. So, Lisa, I was hoping you could give us some understanding of price versus power comparison or compute power. And then today, are you seeing your customers that are buying the MI300, are they primarily buying it for inferencing today, or are they using it primarily for tooling? And maybe for Jean, Jean, do you think is it possible for MI300 to finish the year at a run rate of about \$1.5 billion?

LISA SU: OK, Harsh. So let me start. So your question about the value proposition for MI300. Again, customers are using it for different reasons. But presume that there is a performance per dollar benefit to using AMD. So that's one piece of it.

The other piece of it, though, is we intrinsically have more memory bandwidth and memory capacity on MI300x compared to the competition. And what that means is, for large language models that are many tens of billions of parameters, you could potentially do the workload in fewer GPUs. So it's a substantial system savings and allows you to do much more work within the same system.

In terms of what customers are using MI300 for today, I would say there are a number of customers using it for large language model inferencing, and there are also customers that are using it for training. So I think the whole point is being a strong partner. When you put these AI systems in place, they are sometimes mixed-use systems. So they would be used for both training and inference.

MITCH HAWS: John, we have time for two more questions.

JEAN HU: Yeah. Harsh, let me answer your question about the MI300. Exiting Q4 2024, is it possible to get to \$1.5? It is possible, right? Because, Lisa mentioned earlier, we'll see sequential increase each quarter and the more back end loaded in second half.

And we do have a supply of more than \$3.5 billion. And of course, we'll continue to make progress with our customers. So the mass-- yeah, it's possible. But right now, we are really looking at-- focus on the execution of the current \$3.5 billion plus.

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 questions. For the first one, you had talked about that you had expected a more shallow ramp of the MI300, and it's clearly doing better than that. So some of the upside, I guess, in the near term, is that being pulled forward from the second half, or is this like a step up, or is it more of a step up in demand, both in the first half and the second half relative to what you were seeing before? How do I interpret that shallow comment that you made?

LISA SU: Sure, Stacy. I don't think it's a pull forward of demand. I think what it is, is we wanted to see how long it would take for customers to fully qualify and get their workload's performance. So that depends a lot on the actual engineering work that's done.

And now that we're, let's call it, a quarter later, we've seen that it's gone really well. So it's actually gone a bit better than our original forecast. And as a result, we've seen stronger demand signals. And customers are gaining confidence in their ability to deploy a significant number of MI300 this year.

STACY RASGON: Got it. Thank you. For my follow up, I want to ask Matt's question a little more directly. You didn't quite answer it, that the \$400 billion number that you've got out there, is that just silicon and chips, or is there hardware, and servers, and stuff like in that number as well? What's in that number?

LISA SU: Yeah, I thought I had answered it. But yes, I'll answer it again. It is accelerator chips. It is not systems. So think of it as GPUs. ASICs, that will be there. So those types of things. But obviously, it includes memory and other things that are packaged together with the GPUs.

JEAN HU: Yeah, memory will be quite significant, right? So memory is a big portion of it too.

JOHN: And the final question comes from the line of Chris Danley with Citi. Please proceed with your question.

CHRIS DANLEY: Hey. Thanks for squeezing me in, team. I guess a question for Lisa. As MI300 revenue ramps, how do you see the customer concentration, let's say, a year or two from now? Do you think you'll have one or two customers that are in double digits, or one or two that are half the revenue, or do you think it'll be totally fragmented?

LISA SU: I don't think it'll be one or two that are half the revenue, Chris. I think we're building this as a-- really, we're happy to see the broad adoption. As always, with of the large cloud partners, we might see one or two that are higher than others. But I don't think you see the type of concentration that you mentioned.

CHRIS DANLEY: Great. And then just a follow up on somebody else's question on Intel's roadmap versus TSMC. So I'm sure you're intimately familiar with TSMC's manufacturing roadmap. And we've all seen Intel open up the kimono on what they expect to happen over the next couple of years. I mean, do you think Intel is going to close the gap somewhat with your foundry over the next couple of years, or do you think they'll be able to maintain the lead?

LISA SU: Look, I feel very good about our partnership with TSMC. They continue to execute extremely well. We'll see what happens over the next few years. But I'd like to reemphasize what I said earlier. Even in the case of process parity, we feel very good about our architectural roadmap and all of the other things that we add as we look at our entire portfolio of CPUs, GPUs, DPUs, adaptive SoCs, and put them together to solve problems. I think we feel really good about what we can do with our customers.

So we're always going to be paying attention to the process race. But I think we feel very good about our strategy and how do we continue to push the envelope on the computing roadmaps.

JOHN: OK. That is the end of the question and answer session. I would like to turn the floor back over to the AMD team for any closing comments.

MITCH HAWS: Great, John. That concludes today's call. Thank you to everyone for joining us today.

JOHN: And this concludes today's teleconference. You may disconnect your lines at this time. Thank you for your participation.