VIVEK ARYA: And welcome, everyone, to day one of the BofA Tech Conference. I'm Vivek Arya from the semiconductor, semi cap equipment team. Really great to see you all here and really delighted and honored to have the team from AMD join us this morning, Jean Hu, CFO, and Matt Ramsay, the Head of Investor Relations, and prior, from our tribe of proud sell-side semiconductor analysts. So really glad to have both of you here.

What I'll do is I'll go through my list of questions. But please feel free to raise your hand if you have any question that you would like to bring up. But a really warm welcome to you, Jean and Matt. Really glad that you're here with us.

MATT RAMSAY: Thank you.

VIVEK ARYA: And maybe, as a start, Jean, give us a state of the union, the way you see it-- a lot of macro crosscurrents, and then when AMD itself has kind of been on this very interesting journey in the road towards AI. So what do you find exciting right now? And then we can go through some of the nitty-gritty of each of the segments.

JEAN HU: Yeah. So Vivek, thank you so much for having us. It's great to be here. And thank you all for joining us today. Maybe I'll start to just take a step back to set the stage about the AMD. When you think about it, 2024 for AMD was the really transformation year. During the year, we really made significant progress of building our high-performance computing platform, which includes CPUs, GPUs, and embedded processing.

So specifically, if you think about the AI market, what we have done is, since the launch of MI300, which is literally December 2023, the first year of the launch, we ramped the revenue to exceed $5 billion. We made a very significant progress with our hardware roadmap, which is an annual cadence, introduced MI325 in December 2024, and also accelerated the MI350.

We are actually launching next week. So hopefully all of you can join in to listen to our MI355 launch on June 12. And most importantly, we also made significant progress in software, with ROCm really maturing. AMD Instinct and AMD software are now powering some of the most complex and the most compelling AI models with our-- help customers at scale.

So that's really exciting. On the CPU side, we actually have the best product portfolio with both our server CPU and the client CPU. It's very exciting to build that foundation. Now, when you think about the-- we're looking to 2025. We started 2025 with a very strong Q1 financial performance. Revenue was up 36% year over year, driven by data center was up 57%, and client gaming was up 28%, so really exciting. And the earnings per share, which I really focused on, was up 55%. So exciting time.

When we look at the Q2, that's, Vivek, what you're saying is-- there are a lot of noises and a lot of uncertainties. And one of the things that really impact AMD is the export license requirement for MI308. When you think about this, we were expecting the first half data center revenue to be flattish with the second half of 2024 as we were going through the product transition from MI325 to MI350.

But because of the export license requirement, we got impacted by $700 million revenue in Q2 alone with the data center GPU business. So it was quite significant. Of course, in China, because of the DeepSeek, you did see the surge of the demand for GPUs during that period of time. And what we feel really good about is we actually guided Q2 at $7.4 billion. That's the middle point, which will be like 27% year over year increase, despite of this $700 million revenue impact.

So it's really because our core business's strengths. What we're seeing is continued momentum with our CPU server business. And on the client business side, it has been performing really well. The sell-through is really strong. And we have a much higher, richer product mix versus traditionally because of the product portfolio.

So the other thing I'm actually pleased is our gaming business. We went through this really deep correction. Now, the inventory normalized. And we're actually seeing customers on the gaming console side. They actually start to build for the holiday season in Q2 and Q3. So client gaming business performed really well. We actually-- in Q2, we do expect client gaming business to have high teens sequentially, which, of course, offset the high teens decline on the data center side due to the impact.

Overall, when you think about the first half for AMD, we actually expect revenue to be up 30% despite of this export control and the tariff uncertainties year over year. And of course, we're driving earnings expansion much faster than revenue growth.

What's most exciting is actually, the second half, is we're launching MI350. We are on track with our execution for MI400 generation launch next year. Overall, not only we're adding new customers, but existing customers, we're covering more applications, more models, from both the inferencing side-- you will also see or expect us to run more training models with our customers.

So overall, very exciting. We're confident about our CPU platforms. We do believe the product portfolio innovation and the technology strengths will help us continue to gain share. So overall, I would say 2024 was the year to build the foundation. And 2025 really is an inflection for us to drive top-line revenue growth, not only for 2025, but beyond. And again, that earnings expansion-- so it is exciting time for AMD.

VIVEK ARYA: Excellent. No, that's a great introduction. And you took away half of my question, Jean. So I have to find--

JEAN HU: OK. Now you can ask more interesting--

VIVEK ARYA: Yes. No, thank you for doing that. So let's start with the data center business. Help us understand why you thought first half would be flattish. And what gives you the confidence in this second half upside and with the launch of the MI350? Do you think it is just a product transition issue? Was there something lacking? Like, what changes in the product? And how is just your confidence around getting new customers or expanding your potential with existing customers with the MI series?

JEAN HU: Yeah. I'll start at a high level, then Matt can add. When you think about it, since we launched MI300 in December 2023, the annual cadence really accelerated our roadmap. And when you think about it, each generation of our product, we are making significant improvement. We have the competitive advantage on the inferencing side because of memory capacity, bandwidth. We are continuing to drive that advantage with each generation.

So MI300 is really the first one, MI325, which is very competitive. But MI350 is where we see the inference performance jump by 35 times. And of course, we can also support training models. So when you look at that progression of the product roadmap, typically, you go through a transition-- the significant ramp in 2024 with MI300, over $5 billion.

And then, of course, in the second half of this year, that's where we launch MI350, which we see tremendous customer attractions with not only adding new customers and existing customers. So from overall, when you have a business or market growing so fast, and you're pushing out a different generation of product, you do go through that natural transition.

MATT RAMSAY: Yeah. Vivek, I think the only thing I would add-- Jean said it well. But the things that I would add are, we were really excited as a company to be able to pull in the MI350, 355, by about a quarter. And given the capabilities jump, both on a scale out networking perspective, new data types for IPv4 and IPv6, new memory configs and bandwidth, compute capabilities of the GPUs themselves-- you pull a product in the roadmap, and it does impact some of the purchases of the prior gen as you lead up there and have a transition.

As Jean mentioned in our opening comments, Q2 was planned to be a bit more of a China-heavy quarter for us anyway as the Western customers went through that transition. And then we got some additional DeepSeek demand pull in China. And due to the export restrictions, we've had to exclude that revenue now. So it does make the shape of the year look different and the back half look like a really steep ramp.

But essentially, we were executing to what we had laid out to this community on our call in February, maybe a tiny bit better than that. And then, obviously, we had to react to the export restrictions. But what we-- the confidence levels, and now that we're getting closer and closer to launching the products, we've sampled the MI355 systems to a number of customers and gotten their feedback on it. I think we feel pretty good about where we are.

VIVEK ARYA: Got it. What is the most frequent workloads that AMD product is used for? Because there is always a question that you have the incumbent. They have all their software, and developers, and whatnot. So there is a natural case for them in a lot of public and clouds. And then you have the custom chips on the other side. And there is a case to be made for them in a lot of internal workloads. What is a AMD's workload, where you say, this is where we lead today, or this is where we hope to lead going forward?

JEAN HU: Yeah. I would say, very, very first thing is the inferencing. If you look at the AMD from MI300, we do have the advantage on the inferencing side. And if you look at the first years of RAMP, we are powering some of the most complicated models, like ChatGPT 4 and Microsoft Copilot, on the Llama side, the Meta side, and the recommendation engine, and the different things.

So priority-wise for 2024, we have been really focusing on those very complex models to support the customer into production and cover more and more models and applications. That is really, for us, from the inferencing side, it's the key advantage. And of course, we are increasingly cover the training side, cover different things.

I think, when you look at the ASIC versus merchant GPU, our view has always been, different models need different compute engine. You definitely can see, if the model is very much fixed and not changing, ASIC probably is the most cost-effective solution. But when you look at a gen AI market, the model innovation, just the pace is incredible. Everything continues to change.

So we do think general-purpose GPU will continue to be the majority of the market. That's basically how we think about it. And what we're improving each generation is not only keep the competitive advantage on the inferencing side, but like Matt said, we're adding more networking capability. We acquired ZT, which will add system solution and support the direct level, build up clusters in 2026.

VIVEK ARYA: I see.

MATT RAMSAY: Now, Vivek, the only thing I would add to that is, we all have to think about that this is a transformative computer science that's, what, 30 months old? So we see-- and DeepSeek was one example. And it got a lot of attention because it was in China and some claims that they made on cost. But what we see is sort of a proliferation of folks that are like DeepSeek. As the install base of GPUs, some from our competitors, some from us, get larger and larger and larger, you're going to see more folks doing software innovation on the install base.

And it is certainly-- for algorithms that settle down, it certainly makes tons of sense to do an ASIC. And we continue to think ASICs will be an important part of the market. And we've actually done some in certain instances for customers. But in a time when the algorithms are moving so quickly, another way to generate leading TCO for your customer is to take advantage of the industry software innovations over the whole depreciable life of the product.

And I think that we and our large GPU competitor probably agree on that philosophically around the programmable nature of these systems and how early we are in the software innovation of the models.

JEAN HU: Yeah. And I will say, we also made tremendous progress on the software side. You should see, or expect to see, more of our solutions with the tier 2 CSPs, because we have prioritized our top customers. But in 2025, our major focus is to broadening the customer engagement, so enterprise customers can actually go to different cloud and tier 2 CSP, get the AMD MI300, 325, and 350 to run their application models. So you will see more and more of that.

VIVEK ARYA: Got it. And just one last one on the near to medium term-- just because of the cycle time of getting a lot of these inputs done, how would you characterize the visibility, Jean? Do you think a lot of the decisions about the second half have already been made, so there shouldn't be-- like, what are the upside drivers or downside risks as you look towards the second half for MI?

JEAN HU: Yeah, thanks for the question. I think, when we engage our customers, it's always multi-generational. You're absolutely right. The lead time, the cycle time, for complicated GPU solutions are really long. So typically, you really need to plan ahead with their supply chain. And the capacity continues to be tied with CoWoS and HBM memory. You absolutely have to get your allocation for the year.

From customer engagement side, it's the same, because not only you need resembled hardware, but you also need to make sure your software really works. And to get hardware, the performance level, customers need-- the TCO, they will need. So we have been working with our customers for a long time. So it's not just recently we do feel quite good about the customer feedback, and the customer engagement, and the customer orders for the second half. In the end, we think the market opportunity actually is tremendous. It's really about how we can help a customer ramp on the software side to support their workload.

VIVEK ARYA: Got it. And how would you gauge AMD's progress so far, Jean, when it comes to things like software, the maturity of the ROCm stack, and then also networking? Because some of your peers have-- a lot of their internal solutions when it comes to scale up, proprietary or ethernet based, or they have ethernet switches and so forth. So how would you assess where AMD is from that perspective?

JEAN HU: Matt, you want to-- software is your area of focus.

MATT RAMSAY: Yeah. So Vivek, I think there's a couple of things. And you mentioned-- and you've seen this as AI has proliferated, our need to go to system-level solutions. And that's inclusive of CPU, GPU, networking, software, system-level design, cooling, power supplies. There's a lot to it.

We needed to continue to invest more in the ROCm software stack. And I think we've done that. There was a focus in 2024 of on the largest customers. And now you've seen us start to broaden out. The team's doing biweekly ROCm releases. The goal is, over time, to be the software stack of choice for open-source deployment of AI.

And that's not going to happen in a week or a month. But you saw with what AMD did from 2016 to 2019 in the server space, right? Where it was a focus on the largest of customers first and then a broadening out of the strategy as we gain scale. And I think you'll see that be something that we continue to focus on in the software side with the ROCm stack.

But the team-- the goal is really, how quickly-- how many models can you support that are performant on day 0? The ones that aren't performant on day 0, how much can you shrink the time and the friction to go from a competitor solution to an AMD solution and move the-- get that model performant in not months, not weeks, but days?

And I think, for a lot of things, we're making really good progress on that. And we've also spent a little bit more time and resource, maybe in the long tail, not yet of full support for the long tail, but on educating some of the folks in that community that might be influential in this and in the news flow as to what our progress is, how they should be measuring our progress and software, and have them know what the strategy is versus them just maybe picking up a GPU and trying to run stuff without any engagement or support from AMD, and then that experience being a little different than they expected.

So I think we've made a lot of progress. But there's a lot of work to go. I mean, we've done things like Silo AI, where we brought in an acquit hire that has some vertical-specific software for certain industries. And I think you'll hear more about that. It's a bit of a weird time to have this conversation, with our launch event next Thursday. So anyway, stay tuned.

VIVEK ARYA: OK. As you move to rack scale next year, what is the content lift that AMD can get? Or is it that a lot of the incremental content comes from partners and others? Just, is there a simple way, Jean, to look at what kind of content lift you can get as you migrate from more chips and boards towards more rack-scale systems?

JEAN HU: Yeah. I think AMD's approach has always been embracing the ecosystem. And when you think about this, we do have our GPUs. And also, on the networking side, we acquired a company called Pensando. They have the programmable DPUs, which is a great asset for us from the networking side.

And then, with the ZT acquisition, we have 1,200 very talented engineers. They are like power management, like thermal system solution engineers. They will help us to design reference, design the cluster, and the system-level design. But we work with our partners. On the networking side, we work with Arista, Cisco, Broadcom. And also, we work with other suppliers.

So for us, it is about to provide our customers the system solution. But we get benefit by selling GPUs, CPUs, and DPUs. But our customers also get benefit by selling their networking gears. So I would say, each generation, the content is increasing, especially for us. If you look at MI350, we're still just selling GPUs. But when we get to MI400, not only GPUs and DPUs, and also the head node of CPUs.

So our content will increase very, very significantly. And at the cluster level, if we can support really very, very large clusters, that will help us tremendously. But our partners will get benefit by selling ethernet switches or other things. Yeah.

VIVEK ARYA: Significant as in 30%, 40%, 50%?

JEAN HU: No, we have not quantified any of that yet.

MATT RAMSAY: Yeah. Vivek, the only other thing I would add is, the ZT deal, that we did both ends of it. So we brought the ZT design team in-house and then recently announced the planned divestiture of the manufacturing co. That whole ZT design team, under consultant contracts, was working alongside of the AMD design teams on the MI400 series the whole time. The deal was pending. And now they're part of the company.

A unique focus for them, having been part of a manufacturing company, is design for manufacturability. When they were part of ZT as an integrated company, if they didn't design systems that could then be manufactured and sold, they had no revenue. So we're really, I think, pleased with some expertise that they brought into the company, not just on system-level design, but system-level design with an emphasis on manufacturability and time to market. So I think you'll hear the team talk a bit more about that next week.

VIVEK ARYA: Got it. One of the questions-- we'll come back to the data center side. But one question that has come up, Jean, as you well know, on the client side, is, just given all the trade and tariff issues, the risk of pull-ins and so forth, and AMD has had such significant strength in your PC business. So how are you assessing the risk? I know you have already guided to subsidy in the second half. But was that more out of conservatism, or is that something that you're truly seeing?

JEAN HU: Yeah. Thanks for the question. We absolutely have seen our client business performing extremely well. But we have not seen any meaningful pull-ins and pull-forward from our customers. When we really look at our client performance, for instance, in Q1, client-- actually, PC business is up 68%. But 43% is because ASP increased.

And that is not to like ASP, it's more, AMD has been introducing the latest-generation product. So the mix is at the high end of the stack we play. That has been the key driver of our Q1 performance. Frankly, the client discussion has been ongoing for several quarters because the performance has been really good.

If I look at the trend, what happened in Q1 actually was quite consistent with prior quarters. It is the ASP increase. It is the richer product mix. So Q2, we actually continue to see the momentum. Especially right now, we are in the middle of the quarter. We have not seen a slowdown of the sell-throughs. So it's very good in Q2. We do think, sequentially, it will be better than seasonality.

For the second half, it's just-- the macro uncertainties just are a lot. So we cannot predict what's going to happen. Every day, there's some news. From that perspective, We're trying to be really mindful of the macro environment and make sure we are conservative, thinking about the second half. We didn't mention, it's subseasonal. But it's really driven by the uncertainties in the macro environment and the tariffs.

VIVEK ARYA: Right.

JEAN HU: Yeah.

MATT RAMSAY: Vivek, I think, if the client market from a unit perspective performed, quote unquote, "normally," I think we would have a really good back half of the year in our client business. But given the uncertainties that are out there, I think it's just prudent at this time, given how strong the business was in the first half of the year, to take that approach.

And if we do better than that, great. But I think, for now, the business, we have seen a few pull-forwards, as Jean said. But we've been working really hard to manage around that. Maybe a data point that will be helpful-- in the first quarter, our units were actually down more than seasonal. And the sell-through was more than selling across our client business.

So we've been trying to manage the business as best we can to protect ourselves from the trends that are happening. But there's a lot of uncertainty. And I think being prudent about it right now is probably the best approach.

VIVEK ARYA: Right. So you wouldn't be surprised to see some semblance of seasonality if macro doesn't change from where we are today.

JEAN HU: Yeah, we would be happy to see that.

MATT RAMSAY: Apply your large grains of salt to that.

VIVEK ARYA: OK, understood. One other thing that also came up in your rival's call on the x86 CPU side-- they said that some customers are going down the stack to address some of the tariff issues. I know the tariff situation keeps on changing and that the mix of AI PCs has come down. Did you also see that? Or how do you see that field of AI PCs? Like, you mentioned ASP strength. So what is driving that richer mix in ASP strength for AMD?

JEAN HU: We have not seen that. We ask our team, sales team. We have not seen that. I think the strength of our ASPs has been because the latest generation of product we have, not only we have the best GPUs for the gaming players, we also have the best CPU and APU on the AI capability side.

So overall, if you look at this, we're stressing the high end of each stack, versus, in the past, AMD tend to stay at the lower side, consumer side. But more importantly, on the commercial PC side, we were very, very much underrepresented. We are still underrepresented. But with the Dell engagement, that really helped us with more OEMs and more platforms in the market. And we're focusing on enterprise go-to-market. So with enterprise customer adoption, that also helps. The commercial PC side, it tends to have higher ASP.

So I think our strategy in PC market is-- first is to lead with the technology and the product portfolio. That has been the key driver Lisa and the team is focusing on. That really helps us to-- from market share side, we're still a small player. But for us, we really want to be profitable growth, not only just growing revenue, but to get profitability there. Yeah.

MATT RAMSAY: Yes. And the other piece, Vivek, is in the enthusiast desktop market, both for developers and for gamers. I think, for this audience, go find whatever-- Amazon or whatever e-tailer site you'd like to go find and figure out where AMD SKUs are in the top 10 on those. And I think you'll be pleased with the results. It's been a-- over the last six, seven quarters, it's been a significant share shift in the highest end of the desktop market. And the sell-through has been very strong relative to the sell-in as well.

But that's a new ASP-rich part of the TAM that AMD had not traditionally done that much business in. And the market share is sort of flipped on its head in that piece. And we feel pretty strongly that we can sustain that if we continue to innovate on those products.

VIVEK ARYA: Got it. And in the last minute or so, just wanted to quickly touch on gross margins, since we have our CFO here. So gross margins-- the one pushback we hear about is that, well, AMD's growth will be driven by a lot of these GPUs. And so far, they have been below corporate average. So as that mix shifts more to these GPUs, and then within that, the mix shifts more to rack-scale systems, so how do you look at the trajectory of gross margins from there, Jean? What are the upside and downside risks?

JEAN HU: Yeah. I think AMD's business, there are always a few puts and takes, because the mix actually really drives the gross margin. Overall, when you think about the second half, it is true. We'll see data center to be the number-one revenue driver for the second half. And of course, GPU is an important element of that, which is dilutive to corporate gross margin.

But we do have our server business enterprise, both on the server side and the commercial side, will actually help to drive the gross margin up. So we do think we will see modest gross margin improvement in the second half. Of course, it depends on mix.

Going forward, into next year and beyond, I'm actually quite confident about continued gross margin expansion, despite of data center GPU side is slightly dilutive, or dilutive to gross margin. But the overall enterprise play, both on the server side and the commercial side-- more importantly, we did not touch embedded business. Actually, we do see really important demand signals that the cyclical upturn is coming.

So we do think, next year, the embedded business will also be much stronger versus this year. That is a really high-gross-margin business for us. That will help us with the mix also.

VIVEK ARYA: Terrific. With that, Jean, Matt, thank you so much. Really appreciate the discussion. And thank you all.

MATT RAMSAY: Thank you, everybody.

JEAN HU: Thank you so much.