

AMD Announces Strategic Partnership with Meta Conference Call

MODERATOR: Greetings, and welcome to the AMD Conference Call. At this time, all participants are in a listen-only mode. A question and answer session will follow the formal presentation. If anyone should require operator assistance during the conference, please press Star-0 on your telephone keypad.

As a reminder, this conference is being recorded. I would now like to turn the conference over to your host, Matt Ramsay, Vice President, Financial Strategy and Investor Relations. Thank you. You may begin.

MATT RAMSAY: Good morning. Thank you for joining on such short notice, and welcome to our call to discuss a significant new AI partnership between AMD and Meta. By now, you should have had the opportunity to review a copy of our press release and our Form 8-K filing, discussing this announcement.

If you have not had the opportunity to review these materials, they can be found on the Investor Relations page of amd.com. Participants on today's conference call are Dr. Lisa Su, our chair and CEO, and Jean Hu, our executive vice president, CFO, and treasurer. This is a live call and will be replayed via webcast on our website.

Today's discussion may contain forward-looking statements based on our 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 expectations.

Please refer to the cautionary statement in our press release for more information on the factors that could cause actual results to differ materially. With that, I'll hand the call to Lisa.

LISA SU: Great. Thank you, Matt, and good morning, and thank you all for joining the call today. We're announcing a significant expansion of our strategic partnership with Meta, including a new multi-year, multi-generation agreement that positions AMD at the core of their next generation AI infrastructure.

2025 was a defining year for AMD, with record results across the business, and we're carrying that momentum into 2026. AI demand is accelerating rapidly, as customers scale modern AI infrastructure across both accelerated and general purpose compute.

Through the leadership, technology, and consistent execution, we have built one of the premier data center AI franchises in the industry, anchored by our differentiated Instinct GPU roadmap and our leadership EPYC CPU portfolio.

AMD is uniquely positioned to deliver high performance, energy efficient compute across the full spectrum of AI workloads. Meta has been a close partner over multiple generations, deploying millions of EPYC CPUs and hundreds of megawatts of MI300 and MI350 series GPUs across their global infrastructure.

Meta was also an early definition customer for my MI450 series, and we developed our Helios rack-scale architecture on the OCP Open rack wide standard in collaboration with Meta. Today, we are significantly expanding our relationship.

Under this agreement, Meta is expected to deploy 6 gigawatts of AMD Instinct GPUs across multiple product generations. To meet their evolving AI requirements, we are co-engineering, a custom GPU accelerator based on our MI450 architecture, optimized specifically for Meta's workloads.

Initial shipments supporting the first gigawatt deployment are scheduled to begin in the second half of 2026 and will leverage our Helios rack scale architecture with the custom MI450 based Instinct GPU and our sixth Gen EPYC CPU, codenamed Venice.

This partnership firmly establishes AMD at the center of one of the industry's most significant AI infrastructure deployments, and highlights the strength of our end-to-end platform strategy. As AI workloads scale, customers are increasingly looking for solutions, tailored to their specific architectures and performance requirements.

AMD's leadership in chiplet design and advanced packaging is a key differentiator, and enables us to rapidly leverage core building blocks of our AI platform and tailor them for the optimal compute, memory, and networking needs of specific customer workloads.

The custom AMD MI450-based GPU we are developing with Meta is a direct result of this capability, delivering workload-specific optimizations, while leveraging the MI450 platform, the Helios rack-scale system infrastructure, and our Open rack end software ecosystem, giving Meta the advantages of a custom solution, with the benefits of the broader MI450 ecosystem and GPU programmability.

As these platforms deploy at gigawatt scale, the ecosystem optimizations across ROCm, AI frameworks, and system software will extend well beyond this engagement, strengthening our broader Instinct franchise, and expanding opportunities across our entire customer base.

In addition to expanding our GPU engagement with Meta, we are further deepening our EPYC CPU partnership. We are seeing accelerating CPU compute demand, driven by the rapid scaling of AI infrastructure across model development, inferencing, data processing, and the rise of agentic AI.

As deployments grow in scale and complexity, CPUs remain a strategic foundation of the compute stack, driving efficiency, orchestration, and system-level performance. EPYC is well positioned to capture outsized value in this next phase of AI expansion.

Meta is already a multi-generation EPYC customer, with EPYC processors powering the majority of core services across their global data center footprint. Building on our deep roadmap alignment, Meta will be a lead customer for our 6th Gen EPYC Venice processor at launch later this year.

We have also partnered closely on a new addition to our family, codenamed Verano, incorporating workload specific optimizations to deliver leadership performance per watt and compelling TCO.

The expansion of our partnership across GPUs and CPUs is another strong proof point that the world's most ambitious AI builders are choosing AMD Instinct and EPYC platforms as the foundation of their AI infrastructure.

I want to thank Mark and the entire team for their collaboration and partnership. We are extremely proud to work together to advance the future of AI at scale. Importantly, this engagement reflects the strong and growing demand we are seeing for MI450 series and Helios architecture.

Overall, there is significant excitement in the market for the MI450 series and Helios, and from an execution standpoint, we are making excellent progress. MI450 and Helios are currently in hardware and software validation, running the latest inference and training workloads.

We are working closely with our lead customers, supply chain, and ecosystem partners to ensure a smooth ramp. We expect to begin customer sampling shortly and remain on track to begin production shipments of both the standard MI450 series and the custom me4 50 based GPU for Meta in the second half of 2026. Now, I'll turn the call over to Jean to provide additional details on the agreement.

JEAN HU:

Thank you, Lisa. Today's announcement over 6 gigawatt agreement with Meta is another significant step in scaling our data center AI business, consistent with the ambitious plan we set out at our financial analyst day.

Let me provide some context on the financials. The Meta deployment is expected to generate data center AI revenue of significant double digit billions of dollars per gigawatt. Revenue will begin in the second half of 2026, and the ramp alongside MI450 deployment with other customers.

As part of the agreement and to strategically align the interests of both companies, AMD has issued Meta a warrant for up to 160 million shares of AMD common stock. The warrant is performance based.

The first tranche vests with initial gigawatt of shipment of AMD Instinct GPUs, with additional tranches, vesting as Meta's purchases of Instinct GPUs scaled to 6 gigawatts. Vesting is further tied to AMD, achieving certain stock price thresholds, with the final tranche vesting at a price of \$600 per share.

In addition, exercise of the warrant is tied to Meta, achieving key technical and the commercial milestones. The unique structure aligns better, and the AMD driving significant long-term revenue growth and accretive to our non-gaap earnings per share, while enabling Meta to share directly in the upside of our mutual success.

This partnership marks another significant step forward in delivering our ambitious long-term financial model, including greater than 80% of our data center AI business and generating more than \$20 in annual earnings per share within the next three to five years. Let me now turn the call back to Lisa.

LISA SU:

Thank you, Jean. Let me close by saying that today represents another major milestone for AMD's AI strategy. Through this multiyear, multi-generation agreement to deploy 6 gigawatts of AMD Instinct GPUs, we are significantly expanding our partnership with Meta.

Broadening our AI footprint and deepening our co-development and roadmap alignment with one of the world's leading AI companies that is building at massive scale. We believe that the scale of this deployment and the ecosystem benefits that it drives will further strengthen our AI platforms and expand opportunities across both existing and new customers.

The current AI infrastructure build out is one of the most significant technology investment cycles in decades. The expanded partnership with Meta, together with our previously announced partnerships with OpenAI, Oracle, and others, demonstrates the strength of our multi-generation Instinct and Helios roadmaps and firmly establishes AMD at the center of this next phase of AI growth.

From silicon to systems to software, we are executing with scale, speed, and discipline, leveraging our differentiated end-to-end portfolio and deep strategic partnerships to capture the AI opportunity, accelerate sustained data center growth, and deliver long-term shareholder value. Now, I'll turn the call back to Matt for Q&A.

MATT RAMSAY: Thank you, Lisa. Thank you, Jean. For today's Q&A session, please focus your questions on today's announcement. Given the limited time we have for questions, please do limit yourself to one question per caller. Thank you very much. Operator, you can please poll for the first question.

MODERATOR: Thank you. As a reminder, if you'd like to ask a question, please press Star-1 on your telephone keypad. A confirmation tone will indicate your line is in the question queue. You may press Star-2, if you'd like to remove your question from the queue.

For participants using speaker equipment, it may be necessary to pick up your handset before pressing the star keys. Our first question comes from the line of Joshua Buchalter with TD Cowen. Please proceed with your question.

JOSHUA BUCHALTER: Hey, guys, thanks for taking my question, and congrats on the deal. For my first one, just wanted to hit one that's gotten into my inbox a couple of times. It seems like it wouldn't be the case since Meta doesn't have a third party cloud business, like the other hyperscalers.

But could you confirm whether or not there's any overlap between this deal and the OpenAI deal you announced a few months ago? And are there any details you're able to share on how this 6 gigawatts compares to the OpenAI deal from a timing and value standpoint. Thank you.

LISA SU: Yeah, absolutely, Josh. So thanks for the question. If I give you the overall context, look, I think Meta's plans are extremely ambitious in terms of what they're trying to do, in terms of AI infrastructure build out.

We're super excited about this partnership. Meta has been a long-standing partner of AMD over the last few generations, but this agreement really takes our relationship to the next level and significantly expands on what we were doing before.

In terms of your question about whether there's any overlap, no, there's no overlap from the standpoint, that this agreement is really for Meta's workloads, But the custom GPU that we're building is optimized for Meta's workloads.

It's the first time we are doing a custom GPU of this style, using our MI450-based architecture. So I think look, both OpenAI and Meta are incredible companies. They're doing tremendous innovation, and we are super happy to be partnered with them.

And what we are really trying to do, which we've always been is, if you think about the AI space, there are lots of different workloads and workloads of different types. And what we want to do is partner deeply so that we're providing the technology that each company needs to satisfy their ambitions.

UNIDENTIFIED Thank you for all the caller there. Matt, are we doing follow-ups?

CO.REPRESENTATIVE:

MATT RAMSAY: No, we can just try to keep this to one question per caller, given we got a tight call. But thank you, Josh. Operator, the next caller, please.

MODERATOR: Thank you. Our next question comes from the line of Vivek Arya, with Bank of America. Please proceed with your question.

VIVEK ARYA: Thanks for taking my question. Lisa, I had a more fundamental question, which is where is the economic value add here? You're giving away \$30 billion roughly of value in your stock. And in return, you're getting about \$30, \$35 billion of net income, so it seems like an even swap.

So where does the value add here? So if I ask it in a different way, if the product is so good, why does AMD need to give up 10% of your equity, and is now every MI450 customer going to ask for this deal? Thank you.

LISA SU: Yeah, sure, Vivek. So let me start with some context. If you look at the context of these partnerships, every deal of this scale is very unique. As Jean said, this deal is very accretive to AMD's earnings, so it is a very good deal for AMD's shareholders.

But what we're setting out to create is a strategic partnership, where we're going much, much deeper together. And I view this as a very transformational partnership from an AMD standpoint. Meta is operating at gigawatt scale.

We are actually working with them deeply, in terms of their technology roadmap. And we're working together on hardware/software systems. And with that, it includes significant value to our overall roadmap going forward.

And I think you have to look beyond just these particular economics, which, as I said, are very accretive to our earnings. But you also have to look at what it's doing across our entire roadmap and our future roadmap going forward.

So if you look at the structure of our warrants-- in this case, again, it's a very aligned incentive structure. Meta is making a big bet on deploying at large scale for AMD, which is great.

AMD benefits from this large scale deployment which brings revenue scale, ecosystem maturity, software maturity, and assuming that we satisfy all of the purchase, the purchases, as well as the share price thresholds, AMD shareholders will benefit significantly, and Meta gets to benefit as part of that.

So to your question, no, I don't think every MI450 customer is of this size and scale. Our standpoint is we are looking for defining partnerships for our AI franchise. And Meta is one of those defining partnerships.

MATT RAMSAY: Thanks, Lisa.

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

TIMOTHY ARCURI: Thanks a lot. Lisa, I also wanted to ask a question along the same lines. So it sounds like the exact same deal that OpenAI got right down to the same kind of stock prices as well. So it is a lot of revenue.

But speaking to the precedent here, would Meta have done this without this deal? So when they saw the deal from OpenAI, did they say, OK, I'll do a deal with, if I get the same deal.

So I guess speaking to the precedent, if you're going to sign with Amazon or with somebody else, does this open up the discussion that they would want the same kind of deal? Thanks.

LISA SU:

Tim, I think the way-- if I give you some background on how we came up with this, this wasn't about, hey, are we trying to do something similar with other customers? It's not like that. What it is about is Meta has been a tremendous partner for AMD over the last several generations, and we appreciate that, by the way.

I mean, they have been a big adopter of EPYC. They were an early adopter of MI300 and MI350. But without this strategic agreement, I think we would have done well. I think MI450 would have done well. But what we're looking to do is do something transformational.

And when you talk about gigawatt scale deployments and 6 gigawatts over five years, that is transformational in terms of where we see our business. And in addition to that, they're at the forefront of what's happening with models.

And model builders, they are optimizing workloads for their future. And we are optimized alongside with them. So I think if you look in that context, they are different deals, but they are a very important strategic deals in terms of the shape of AI, going forward.

And if you think about this market, there's probably only a handful of companies that are deploying at this scale. And so to have Meta today and OpenAI, as well as strategic partners that anchor AMD's AI strategy, I think is a really, really good place for us to be.

MATT RAMSAY:

OK, Lisa, thank you.

MODERATOR:

Thank you. Our next question comes from the line of Blayne Curtis with Jefferies. Please proceed with your question.

BLAYNE CURTIS:

Hey, good morning, and congrats. Lisa, I just wanted to dive in a little bit more on the economics. Should we think about this deal-- you've previously talked about a range of dollars per gigawatt, and maybe OpenAI was at the lower end of that. Is that the same kind of economics? And I'm just curious in terms of the support of these two deals, how do you think about the OpEx part of the equation?

LISA SU:

Yeah, so the way I would say it is, we've talked about-- I think Jean might have mentioned that when we think about the revenue from-- let's call it the GPUs. We're talking about something like double digit billions per gigawatt.

And so that is the range that we're talking about when you-- I'm sorry, the second part of your question?

BLAYNE CURTIS:

The OpEx considerations to support these videos.

JEAN HU:

Yeah, yeah. Blayne, on the OpEx side, because it is based on MI450, you can think about it's just another variant of MI450. So the incremental OpEx actually is quite minimal. That drives significant operating model leverage.

MATT RAMSAY: Thanks, Jean.

MODERATOR: Thank you. Our 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. In the prepared comments, obviously, this is all about the Instinct GPU roadmap, but you've also emphasized the importance that the CPU plays in these architecture deployments going forward, particularly probably around AI inferencing.

I'm curious, Lisa, if you could talk a little bit about how you see the evolving competitive landscape in GPUs, in particular, one of your key competitors, trying to push arm-based architecture, more prolifically on a standalone basis. I'm curious how you see that playing out in these architecture deployments. Thank you.

LISA SU: Sure. So, thanks for the question. I think what we've said, and what we said on the last earnings call and the last couple of earnings calls is the CPU market is absolutely on fire. I mean, there is significant demand.

It has continued to grow, and it really is a result of the AI infrastructure deployments as inferencing scales, as the agentic AI scales. And our EPYC portfolio is in an extremely good position. I mean, everything that we see, certainly with our current Zen 5 class turn products, we are very widely deployed in AI infrastructure.

And as we go into our Zen 6 family with Venice and the addition of Verano, we actually see our workload coverage increasing across a number of our largest customers. So I think it's great, that there's so much interest around CPUs.

I think that is something that we've always believed, that you need all different types of compute. And we feel we're very, very well positioned from a competitive standpoint.

MODERATOR: Thank you. Our next question comes from the line of Tom O'Malley with Barclays. Please proceed with your question.

TOM O'MALLEY: Hey, thanks for taking my question, and congrats on the deal. Lisa, I wanted to dive in on the custom commentary from the release, and then also in your answer to one of the questions earlier. What does it mean specifically about a custom design?

Is it just a different flavor of a tape out? Is the system architecture going to look a little bit different? And then in the future, are you going to be doing more custom-style tape outs and/or systems with other customers that come on board. Just a little bit curious what that means specifically. Thank you.

LISA SU: Yeah. Tom, I think this is actually a pretty interesting new thing that we're doing here. We've always believed that when you look at AI infrastructure, there's no one chip that does it all, especially when you look across training and inference and big models and small models and different workloads that you're trying to optimize.

So what's unique about this deal is we started with the workload first. We didn't start with the chip. We started with the workload. What is most important to Meta for their future workloads, their highest volume workloads?

And then we worked back from that with our chiplet architecture. What's unique about our chiplet architecture is we have all the building block pieces, but you can put them together and configure them in different ways to give you different performance and system characteristics.

And so what we've done together with Meta is we're using our chiplet architecture to come up with a new variant based on the MI450 architecture. As Jean said, it's highly leveraging the base capability, which gives us both development scale, as well as, frankly, a lot of leverage, as we're bringing up these technologies.

But it's not just a chip optimization. I think it's chip level, board level, system level. And that comes together in a solution that I think gives the best of both worlds, which is something that is highly tuned to their workload, but takes advantage of the entire infrastructure and supply chain, and everything else that we are developing for the base MI450 architecture.

So to your question, do we expect to do more? I would expect that for high-volume workloads there, there will be benefits to doing something that is let's call it more customized in a GPU format.

Note that we're not doing full ASIC, which, again, usually takes a lot more time, but we're doing something based off of our foundational architecture. So I think it's a good opportunity to expand our portfolio, as customers are increasing their volumes.

JEAN HU: Hey, Tom, just add to what Lisa said that there's no additional tape out needed for this customer chip.

MODERATOR: Thank you. Our next question comes from the line of Antoine Chkaiban with New Street Research. Please proceed with your question.

ANTOINE CHKAIBAN: Yes. Hi. Thank you very much for the question. Actually, I'd like to follow-up on the prior question. I'm wondering, given this is a custom deployment, how much of Meta's software investment in rack end is truly transferable versus custom?

And if I ask the question differently, does this deal really create a self-reinforcing ecosystem that makes subsequent AMD deployments progressively easier, as I imagine custom hardware also means custom software. Thank you.

LISA SU: Antoine, actually, it's extremely leverageable. So I don't know if I would say 100%, but 95-plus percent is the way you should think about it. Because all of the underlying software is using the MI450 architecture.

So if you think about all of the work that we have to do in terms of the base libraries, the kernel optimizations, all of that stuff is highly leverageable to the rest of the AMD ecosystem. And again, this is a GPU.

So from a GPU standpoint, it's already highly programmable. So Meta has already been a very strong partner with us on the software ecosystem. If you look at all of their work with PyTorch and the open ecosystem, and I expect that we will continue to work closely together, as we scale these deployments to gigawatts plus.

ANTOINE CHKAIBAN: Thank you.

MODERATOR: Thank you. Our next question comes from the line of Mark Lipacis with Evercore ISI. Please proceed with your question.

MARK LIPACIS: Hi. Thanks for taking my question. Lisa, can you help us understand what kind of visibility do you have under this deal? I'm trying to understand what hard orders do you have right now for the near term.

And is there a take or pay element on Meta's part here, or can they elect to opt out? Can you just help us understand how hard are the orders? I appreciate you anticipate getting the full 6 gigawatts, but if you could help us understand the near-term visibility and the longer term visibility. Thank you.

LISA SU: Sure, Mark. So I think, as we said, we are assigning a long-term strategic agreement that goes across five years for 6 gigawatts. The first gigawatt is committed, and we'll start shipments in the second half of 2026.

As you know, the supply chain, overall, is tight, and so we're planning it very tightly together for their data center builds in terms of which data centers these things are going in. But I think we have very good near-term visibility.

The important thing, I want to mention, again, Mark, is we had already a very strong relationship with Meta. But what this agreement does is it really does take our near-term work together to the next level.

And that has been very, very positive. And I really feel that it is a big addition to what we see in terms of MI450 adoption to include the custom GPU optimized for metas workloads. So in terms of long term we are actively working on beyond MI450 as well.

As you know, we've already been well into the development of MI500 and beyond. And I expect with each generation, we can get even more optimized, as we learn more about their workloads, and they learn more about our architecture.

MARK LIPACIS: Very helpful. Thank you.

LISA SU: Thanks.

MODERATOR: Thank you. Our next question comes from the line of Harsh Kumar with Piper Sandler. Please proceed with your question.

HARSH KUMAR: Yeah, Lisa, thank you for letting me ask a question. I was curious about how the deal was run, especially the custom GPU part. This is not the norm for you guys. So was this a bake off against some other people, or did Meta approach you specifically with an idea for custom chip? I'd be curious how this played out.

LISA SU: Harsh, I think the best way to say it is we're always in deep discussions with Meta. They are a very close partner. What has been a very good about this engagement is Mark and I have spent a good amount of time together on how we really align our roadmap with their next generation infrastructure.

I think Mark is extremely ambitious with what he wants to do with Meta, in terms of what they're doing with their models. I think they have a set of requirements. He has a great team that works with him, and we have been working closely with them on how can we both optimize, as well as expand our relationship.

So from that standpoint, we started with workload first, what are you trying to accomplish. And we came up with ideas for how we could even enhance the base standard product of MI450 to broaden our workload adoption.

So it was a very collaborative effort and one that I think talks about the types of bets that companies are making. We are making certainly a big bet on Meta, and I think Meta is making a big bet on AMD.

HARSH KUMAR: Thank you so much.

MATT RAMSAY: Operator, I think we have time for just one more question before we end today's call. Thank you.

MODERATOR: Thank you. Our last question today will come from the line of Srini Pajjuri with RBC Capital Markets. Please proceed with your question.

SRINI PAJJURI: Thank you. Good morning, and thanks for squeezing me in. Lisa, just want to clarify-- you gave us some targets at the analyst day last year, \$20 plus EPS. Just curious if those numbers already included the Meta deal.

And then you also talked about potential other customers, and obviously, Meta is one of them. I'm just curious, the \$20 number includes any other future deals as well. Thank you.

LISA SU: Yeah, so going back to our financial model, I think we put some very ambitious goals out there in terms of our revenue, as well as our EPS and the \$20 EPS. When we came up with that number in November, I wouldn't say we had this Meta deal specifically baked in.

This was still very, very much in the works. This does expand our relationship with Meta, which is a great thing. And when we look at these financial models, we have a broad set of customers and a broad set of customers that we're actively engaging.

So I think what this should give you is clearer visibility into how we intend to both achieve and exceed our financial model, having a great strategic partnership, like Meta, as well as some of the other partnerships that we have talked about.

And you should imagine that the overall interest in MI450 is very high. And in addition to the partnerships that we're talking about right now, there are a number of other strategic partnerships that are underway. And I feel very good about our trajectory towards that long-term financial model.

MATT RAMSAY: All right, thank you very much, everyone. I think this is the end of the Q&A period. We really appreciate everybody jumping on short notice. It's an exciting day for AMD. And thank you for your interest and your time. Operator, please go ahead and close the call. Thank you.

MODERATOR:

Thank you. This concludes today's conference call. You may disconnect your lines at this time. Thank you for your participation.