AMD Fiscal Second Quarter 2025 Financial Results

MODERATOR: Greetings and welcome to the AMD's second quarter 2025 conference call. At this time, all participants are in a listen-only mode. A question and answer session will follow the formal presentation. If anyone should require operator assistance, please press star 0 on your telephone keypad.

As a reminder, this conference is being recorded. And it is now my pleasure to introduce to you Matthew Ramsay, VP of investor relations and financial strategy. Thank you, sir. Please go ahead.

MATTHEW RAMSAY: Thank you, and welcome to AMD's 2025 second quarter financial results conference call. By now, you should have had the opportunity to review a copy of our earnings press release and the accompanying slides. If you have not had the chance to review these materials, they can be found on the Investor Relations page of amd.com.

We will refer primarily to non-GAAP financial measures during today's call. The full non-GAAP-to-GAAP reconciliations are available in today's press release and slides posted on our website. Participants in today's conference call are Dr. Lisa Su, our chair and chief executive officer, and Jean Hu, our executive vice president, chief financial officer, and treasurer. This is a live call and will be replayed via webcast on our website.

Before we begin, I would like to note that Jean Hu, executive vice president, chief financial officer, and treasurer, will present at Citi's 2025 Global TMT Conference on Wednesday, September 3. And Forrest Norrod, executive vice president and General manager of Data Center Solutions Business Unit, will present at the Goldman Sachs Communicopia and Technology Conference on Monday, September 8.

Today's discussion contains forward-looking statements based on our current beliefs, assumptions, and expectations, speaks only as of today, and as such involves 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 will hand the call over to Lisa.

LISA SU: Thank you, Matt. And good afternoon to all those listening today. We delivered very strong second quarter results, with revenue exceeding the midpoint of guidance as higher EPYC and Ryzen processor sales more than offset headwinds from export controls that impacted Instinct sales.

We set records for both EPYC and Ryzen CPU sales, reflecting the broad-based demand for our differentiated, high-performance data center, PC, and embedded processors. Second quarter revenue increased 32% year over year to a record $7.7 billion, and we delivered over $1 billion in free cash flow.

Excluding the $800 million inventory writedown, related to data center AI export controls, gross margin was 54%, marking our sixth consecutive quarter of year-over-year margin expansion led by a richer product mix. Turning to the segments, data center segment revenue increased 14% year-over-year to $3.2 billion.

We saw robust demand across our EPYC portfolio to power cloud and enterprise workloads and increasingly for emerging AI use cases. In particular, adoption of agentic AI is creating additional demand for general-purpose compute infrastructure, as customers quickly realize that each token generated by a GPU triggers multiple CPU-intensive tasks.

Against this backdrop, 5th-gen EPYC Turin shipments ramped significantly, and we had sustained demand for our prior generation EPYC processors. As a result, we set records for both cloud and enterprise CPU sales and delivered our 33rd consecutive quarter of year-over-year share gains.

In cloud, adoption expanded with the largest hyperscalers as they deployed EPYC to power more of their mission-critical infrastructure, services, and public cloud products. More than 100 new AMD-powered cloud instances launched in the quarter, including multiple Turin instances from Google and Oracle Cloud that deliver up to twice the performance of our previous generation, which were already the industry's highest-performing offerings.

There are now nearly 1,200 EPYC cloud instances available globally, as providers continue expanding both the breadth and regional availability of their AMD offerings. This continued expansion is accelerating enterprise adoption of EPYC in the cloud, with deployments growing significantly from the prior quarter as we closed large wins with dozens of large aerospace, streaming, financial services, retail, and energy companies.

EPYC adoption also grew with telecom customers, as providers modernize their infrastructure for next-generation networks. For example, KDDI announced plans to deploy EPYC processors to power its 5G virtualized network, and Nokia selected EPYC for its cloud platform used by service providers to build, deploy, and manage core network functions.

Turning to enterprise on-prem adoption, HPE, Dell, Lenovo and Supermicro launched 28 new Turin platforms in the quarter that delivered leadership performance, efficiency, and TCO across a wide range of enterprise workloads. EPYC enterprise deployments grew significantly from the prior quarter, supported by new wins with large technology, automotive, manufacturing, financial services, and public sector customers.

To extend our momentum with SMB and hosted IT service customers, we launched the EPYC 4005 series that combine enterprise-grade performance and features in cost-optimized platforms, purpose-built for smaller-scale deployments.

Turning to HPC, AMD now powers more than 1/3 of the world's fastest supercomputers, including El Capitan and Frontier, which retained the number 1 and number 2 spots on the latest top 500 list. We also powered 12 of the top 20 systems on the Green 500, highlighting the performance per watt advantages of EPYC and Instinct for large-scale deployments.

Looking ahead, we remain bullish on our server CPU business, driven by durable tailwinds, including growing demand for cloud and on-prem compute, sustained share gains, and the growing investments in general-purpose infrastructure required to enable AI.

Turning to our data center AI business, revenue declined year-over-year as US export restrictions effectively eliminated MI308 sales to China, and we began transitioning to our next-generation MI350 series accelerators. We made solid progress with MI300 and MI325 in the quarter, closing new wins and expanding adoption with tier-1 customers, next-generation AI cloud providers, and end users.

Today, 7 of the top 10 model builders and AI companies use Instinct, underscoring the performance and TCO advantages of our data center AI solutions. We launched our Instinct MI350 series with industry-leading memory bandwidth and capacity and broad adoption across hyperscalers, AI companies, and OEMs.

From a competitive standpoint, MI355 matches or exceeds B200 in critical training and inference workloads and delivers comparable performance to GB200 for key workloads at significantly lower cost and complexity. For at-scale inferencing, MI355 delivers up to 40% more tokens per dollar, providing leadership performance and clear TCO advantages.

With the MI350 series, we're also expanding our system-level capabilities to support deployments powered by AMD CPUs, GPUs, and NICs. As one example, Oracle is building a 27,000-plus node AI cluster combining MI355 accelerators, 5th-gen EPYC Turin CPUs, and Pollara 400 SmartNICs.

We began volume production of the MI350 series ahead of schedule in June, and expect a steep production ramp in the second half of the year to support large-scale production deployments with multiple customers.

Our Sovereign AI engagements accelerated in the quarter, as governments around the world adopt AMD technology to build secure AI infrastructure and advance their economies. As one example, we announced a multi-billion dollar collaboration with UMaine to build AI infrastructure powered entirely on AMD CPUs, GPUs, and software.

Initial deployments are underway in key regions, with quarterly expansions planned over the coming years. In addition, we have more than 40 active engagements globally and see significant opportunities to power an increasingly larger portion of national computing centers and sovereign AI initiatives.

On the AI software front, we made significant progress this quarter increasing the performance, improving the usability, and expanding the adoption of ROCm. We announced ROCm 7 with major upgrades across every layer of the stack, delivering more than 3x higher inferencing and training performance, compared to our prior generation, and adding support for large scale training, distributed inference, and lower precision data types.

The deepened developer engagement, we introduced nightly ROCm builds and expanded access to instant compute infrastructure, including launching our first Developer Cloud that provides pre-configured containers for instant access to AMD GPUs. We also expanded native support for ROCm across key frameworks including vLLM and SGLang, enabling frontier models like Llama 4, Gemma 3, and DeepSeek R1 to launch with Day 0 AMD support.

To accelerate enterprise adoption, we introduced ROCm Enterprise AI, a full-stack platform that integrates seamlessly with existing IT infrastructure and includes everything needed for an enterprise to deploy, manage, and scale AI across their business.

Looking ahead, the development of our next-generation MI400 series is progressing rapidly. These are the most advanced GPUs we have ever built, with up to 40 petaflops of FP4 AI performance and 50% more memory, memory bandwidth, and scale out throughput than the competition.

With the MI400 series, we're bringing together everything we've learned across silicon, software, and systems to deliver Helios, a full-stack, rack-scale AI platform. Helios is purpose-built for the most demanding AI workloads, with each rack connecting up to 72 GPUs that can operate as a single massive AI accelerator.

Helios is expected to deliver up to a 10x generational performance increase for the most advanced frontier models. And we believe it will be the highest performance AI system in the world when it launches. MI400 series development is progressing well towards our planned launch in 2026, with significant interest in large-scale deployments from multiple high-profile customers.

To accelerate our development, we have invested significantly to expand our AI software and hardware capabilities both organically and inorganically with a number of acquisitions and strategic investments. We strengthened our software stack last quarter with the addition of the Brium and Lamini teams, building on our acquisitions of NOD.AI, Mipsology, and Silo AI.

On the hardware side, we added a world-class rack and data center scale design team in the second quarter with our acquisition of ZT Systems. The ZT team has integrated seamlessly, and they are actively engaging with multiple customers to accelerate deployments of our Helios solutions at scale. We also announced last quarter that Sanmina intends to acquire ZT US-based manufacturing business, becoming our lead partner for AI rack manufacturing.

Turning to the AI regulatory environment, earlier this quarter, we were notified by the Department of Commerce that it is moving forward with the review of our license applications to export MI308 to China. We appreciate the focus the Trump administration is placing on ensuring that the US technology remains central to global AI infrastructure. And we expect to resume MI308 shipments as licenses are approved, subject to end customer demand and supply chain readiness.

As our licenses are still under review, we are not including any MI308 revenue in our third-quarter guidance. Despite that, we expect revenue to grow year-over-year in the third quarter, driven by the ramp of MI350 at multiple customers.

In client and gaming, segment revenue increased 69% year-over-year to $3.6 billion, driven by record client CPU sales and strong demand for our semi-custom game console SoCs and Radeon GPUs. Client revenue increased 67% year-over-year to $2.5 billion, led by record desktop CPU sales.

Demand for our latest generation Ryzen 9000 series was strong, especially for our differentiated X3D processors. We delivered record desktop channel CPU sales, as Ryzen processors consistently topped the bestselling CPU lists at major global e-tailers throughout the quarter.

We also expanded our Zen 5 desktop portfolio with the launch of our latest Threadripper processors that feature up to 96 cores and deliver up to double the performance of the competition in many popular content creation and design workloads.

In mobile, demand for AMD-powered notebooks was strong, with sell-out growing by a large double-digit percentage year over year. We drove a richer mix of higher ASP mobile parts year-over-year as we expanded our share in the premium notebook segment, where our Ryzen AI 300 CPUs deliver leadership, performance, and value for both general-purpose and AI workloads.

In commercial PCs, Ryzen adoption accelerated as OEM consumption increased more than 25% year-over-year. We saw strong sell-through for AMD commercial notebooks with Lenovo and HP, and a significant uptick in Dell sales as they ramp availability of their AMD commercial portfolio.

We also closed new enterprise wins with Forbes 2000 pharma, tech, automotive, financial services, aerospace, and health care companies. We expect to continue growing our commercial client share, based on the strength of our product portfolio and expanded breadth of OEM offerings.

Looking more broadly, we remain confident we can continue growing client processor revenue ahead of the market over the coming quarters, driven by increased adoption of our desktop and notebook products, growing commercial momentum, and a richer product mix.

In gaming, revenue increased 73% year-over-year to $1.1 billion. Semi-custom revenue increased by a large double-digit percentage year-over-year, as console inventories inventory is normalized and our customers began preparing for the holiday season.

We announced a new, multiyear collaboration with Microsoft for custom chips that will power the next generation of Xbox devices, including consoles, PCs, and handhelds. We also deepened our collaboration with Sony through Project Amethyst, a co-engineering program that will use machine learning to power the next wave of immersive gaming experiences.

In PC gaming, demand for our latest generation Radeon 9000 series GPUs was very strong, with desktop GPUs sell-through accelerating in the quarter, as demand outpaced supply. We launched the Radeon 9600 XT, extending the performance advantages of RDNA 4 to mainstream gamers and delivering a significant uplift in gaming performance, including more than double the ray tracing of our prior generation.

As part of our end-to-end AI strategy, we introduced the Radeon AI PRO R9700 GPU for local inferencing, model fine-tuning, and other data-intensive workloads. The R9700 features more memory, full ROCm support, and multi-GPU scalability, enabling advanced AI development and deployment directly on the desktop.

Turning to our embedded segment, revenue decreased 4% year-over-year to $824 million. Demand continues recovering gradually with sell-through in the second quarter picking up, as strength in most markets was offset by a few pockets of softness in inventory reduction actions, largely with industrial customers.

We expanded our embedded portfolio with the first production shipments of Spartan UltraScale+ FPGAs that deliver leadership performance and advanced security for cost-sensitive, low-power applications.

Adoption of our Versal adaptive SoCs continues expanding in high-end applications, including next-generation Robotaxi platforms developed by Bosch in Europe, where Versal serves as a high-performance controller enabling real-time processing, security, and encryption in fully electric automated vehicles.

Looking ahead, we expect improving demand in the test and measurement, communications, and aerospace markets will drive a return to sequential growth in the second half of 2025. Longer-term design win momentum continues to build tracking ahead of this point last year and putting us on pace to surpass the record $14 billion in design wins we achieved in 2024.

In summary, demand is very strong across our product portfolio, and we are well-positioned to deliver significant growth in the second half of the year, led by the steep ramp of MI350 series accelerators and ongoing EPYC and Ryzen share gains.

Our server and PC CPU businesses are accelerating, driven by growing demand for high-performance compute, sustained share gains, the strength of our product portfolio, and expanded go-to market investments. Our embedded-in-gaming businesses are returning to growth and are well-positioned for long-term success, supported by strong design-win momentum.

And in AI, we are seeing strong adoption of our MI350 series and ROCm 7, as we deliver leadership performance and TCO advantages across a broader range of workloads and ramp deployments with an expanded set of cloud and enterprise customers.

Looking ahead, we see a clear path to scaling our AI business to 10's of billions of dollars in annual revenue. We are very excited about our next-generation MI400 series, which is another giant step forward on our roadmap and has been designed to deliver leadership performance at the chip, server, and rack levels. Customer interest for the MI400 series is very strong, and we are actively engaging with an expanding set of customers to support large-scale deployments in 2026.

We are in the early stages of an industry-wide AI transformation that will drive a step function increase in compute demand across all of our markets, positioning us for significant revenue and earnings growth over the coming years. Now, I'd like to turn the call over to Jean to provide some additional color on our second quarter results. Jean.

JEAN HU: Thank you, Lisa. And good afternoon, everyone. I'll start with a review of our financial results and then provide our outlook for the third quarter of fiscal 2025. We are pleased with our strong second-quarter financial results. We delivered record revenue of $7.7 billion, exceeding the midpoint of our guidance, up 32% year-over-year, reflecting strong momentum across our business.

Record sales of Ryzen and EPYC processors and higher semi-custom shipment more than offset the impact of the US export controls, restricting MI308 sales to China. Revenue increased 3% sequentially due to strong growth in the client and gaming segment, partially offset by the data center revenue decrease due to export controls.

Gross margin was 43%, down 10 points from 53% a year ago. The decrease was due to the $800 million inventory and related charges associated with the export restrictions. Excluding this charge, non-GAAP gross margin would have been approximately 54%.

Operating expenses were approximately $2.4 billion, and an increase of 32% year-over-year, as we continue to invest aggressively in go-to market activities for revenue growth and in R&D to capitalize on significant future AI expansion opportunities.

Operating income was $897 million, representing a 12% operating margin compared to $1.3 billion, or 22% a year ago. The decline was primarily due to the inventory and related charges. Taxes, interest expense, and other totaled $126 million. For the second quarter of 2025, diluted earnings per share were $0.48, compared to $0.69 a year ago. The inventory and related charges reduced earnings per share by approximately $0.43.

Now turning to our reportable segment, starting with the data center. Data center segment revenue was $3.2 billion, up 14% year-over-year, driven by strong EPYC CPU revenue and the share gains across both cloud and enterprise customers. On a sequential basis, data center revenue decreased 12% due to the impact of the export controls on MI308.

The data center segment operating loss was $155 million, compared to operating income of $743 million a year ago, or 26% of revenue. The loss was primarily due to the inventory and related charges. Client and gaming segment revenue was $3.6 billion, up 69% year-over-year and 20% sequentially, driven by record client CPU sales and strong demand for our PC and console gaming products.

In the client business, revenue was a record of $2.5 billion, up 67% year over year, driven by record sales of our Ryzen desktop CPUs and richer product mix. Gaming revenue rose to $1.1 billion, up 73% year-over-year, reflecting strong demand for our newly launched gaming GPUs and higher semi-custom revenue as inventory has now normalized and customers prepare for the holiday season.

Client and gaming segment operating income was $767 million, or 21% of revenue, compared to $166 million, or 8% a year ago, driven by richer client product mix and operating leverage on higher revenue.

Embedded segment revenue was $824 million, down 4% year-over-year and flat sequentially, as embedded and market demand remains mixed. Embedded segment operating income was $275 million, or 33% of revenue, compared to $345 million, or 40% a year ago. The decline in operating income was primarily due to product mix.

Before I review the balance sheet and the cash flow, as a reminder, we closed the acquisition of ZT Systems early in the second quarter. As we had announced our intent to divest ZT manufacturing business, the financial results of the ZT manufacturing business are reported separately in our financial statements as discontinued operations and are excluded from our non-GAAP financials.

Subsequently, in May, we entered into an agreement with the Sanmina Corporation to sell the ZT manufacturing business for $3 billion in cash and stock, inclusive of contingent payments. The transaction is expected to close near the end of 2025, subject to regulatory approvals and customary closing conditions.

Turning to the balance sheet and cash flow, during the quarter, we generated $1.5 billion in cash from operating activities of continuing operations, and free cash flow was a record of $1.2 billion. We returned $478 million to shareholders through share repurchases, resulting in $1.2 billion in share repurchases for the first half of 2025.

In May, our board of directors approved an additional $6 billion authorization. Exiting the quarter, we have $9.5 billion remaining under our share repurchase program. At the end of the quarter, cash, cash equivalents and short-term investments were $5.9 billion. Our long-term debt was $3.2 billion. During the quarter, we paid down $950 million of commercial paper used to finance the ZT system acquisition clause.

Now turning to our third quarter 2025 outlook, please note that our third quarter outlook does not include any revenue from AMD Instinct MI308 shipment to China, as our license applications are currently under review by US government. For the third quarter of 2025, we expect revenue to be approximately $8.7 billion, plus or minus $300 million.

The midpoint of our guidance represents approximately 28% year-over-year revenue growth, driven by strong double-digit growth in our client and gaming and data center segments. Sequentially, we expect revenue to grow by approximately 13%, driven by strong double-digit growth in the data center segment with the ramp of our AMD Instinct MI350 series GPU products.

Modest growth in our client gaming segment, with client revenue increasing and gaming revenue to be flattish. And our embedded segment revenue to return to growth. In addition, we expect third-quarter non-GAAP gross margin to be approximately 54%, and we expect non-GAAP operating expenses to be approximately $2.55 billion.

We expect net interest and other expenses to be, again, over approximately $10 million. We expect our non-GAAP effective tax rate to be 13% and diluted share count is expected to be approximately 1.63 billion shares.

In closing, we executed very well in the first half of the year, delivering record revenue and building strong momentum for growth in the second half. The strategic investment we are making position us to capitalize on the expanding AI opportunities across all our end markets, driving sustainable long-term revenue growth and earnings expansion for compelling value creation. With that, I'll turn it back to Matt for the Q&A session.

MATTHEW RAMSAY: Operator, will you please pull the audience for questions? Thank you.

MODERATOR: Yes, thank you, Matt. 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 to remove yourself from the queue.

For participants using speaker equipment, it may be necessary to pick up the 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 poll for questions. And the first question comes from the line of Thomas O'Malley with Barclays. Please proceed with your question.

THOMAS O'MALLEY: Hey, thank you for taking my questions. And I appreciate it. Lisa, I want to start on the client business. So you had previously laid out a second-half outlook that was roughly flattish with the first half, as you were kind of protecting against some pull-forward.

So first, do you think that your Q2 results included some pull-forwards and the second half should still be flattish? And longer-term, after the Intel commentary regarding 18A maybe what that means is a knee-jerk reaction just right away for AMD longer-term in terms share and ASPs?

LISA SU: Sure, Tom, thanks for the question. So first of all, our client and gaming segment, and particularly our client business, just performed very well in the first half of the year. I think if you look at the entire first half, it was up 68% year-over-year.

I think if you look underneath that, what we're seeing is strength in every part of our client business. So we saw very strong sales in our desktop channel area. We have a leadership product there, best gaming GPUs with our X3D GPUs. We've had strong rise in AI adoption. As well in the first half of the year, we see that in sell-through.

And in addition, we've had strong enterprise sell-through as we brought that forward. So to your question of how much is pull-forward, we don't think a whole lot of that is. When we look at the sell-through patterns, the end user consumption is actually quite strong for client in terms of going into the second half of the year.

As we said in our Q3 guide, the primary driver of our Q3 guide is a very strong data center driven by MI350 ramping. We are expecting some growth in the client business, so I wouldn't say it'll be flat to the first half, but we're planning for it to be a little bit less than seasonal, just given some of the uncertainties out there.

But the client business is performing extremely well for us. We believe we are gaining share in all the right places. So if you look at the numbers in the first quarter, and it'll show through in the second quarter as well, a lot of the uplift in revenue is in ASPs. And that is basically, we're selling up the stack on the strength of our portfolio.

And I think we're still quite underrepresented in the enterprise portion of the business. That is where we have increased our go-to market resources and focus. And we're seeing nice traction there, especially with the portfolios that we have from HP and Lenovo in enterprise PCs. And now we're adding Dell as well as it's ramping here, started in the second quarter. We'll ramp more in the second half of the year.

So I think all of those are tailwinds for our client business, beyond the second half of '25, but really, into the next number of quarters, as we think about the portfolio and the opportunities for us.

THOMAS O'MALLEY: Super helpful. And then secondly, I was hopeful you could provide us a little more color on China. So the guide doesn't include MI308. But perhaps you could comment on when you get approval if the supply chain is ready, what's currently in inventory, and maybe compare what you think the contribution will look like versus the $700 million in Q2 and the $800 million for the second half you spoke about in April.

LISA SU: Sure, Tom. So, yeah, let me answer some of the questions on China. I'm sure that there are some questions. Look, we're very pleased with the progress that's been made with the administration over the last couple of months. We've been working very closely with the administration. I think the focus here on ensuring that US technology gets utilized throughout the world is something that we certainly support and very much want to contribute to.

China's an important market for us. Given the timing of licenses, we have a number of licenses that are under review now. We are working with the Department of Commerce to get those reviewed. We do expect that once those licenses are approved, we will start MI308 shipments.

In terms of the supply chain, most of our inventory was not in finished goods, so it was work in process, and it will take us a couple of quarters to run through that. The exact timing of revenue and contribution will depend a bit on when the licenses are actually granted.

But overall, I think this is a better position than we were 90 days ago. And we certainly view China as a market that we would like to service with MI308. And we're working closely with the administration to do that.

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

VIVEK ARYA: Thank you for taking my question. If we look into 2026, that's when I think the Sovereign opportunity could get quite meaningful for AMD. What is the right way to size that? What does this JV structure mean with some of the contracts that you have signed? And would you consider this incremental to the kind of growth rate that you're seeing with your current MI business, or would this be instead of?

So just if you could give us a way to size, what is that incremental opportunity from a sovereign customers when it comes to '26 , Lisa. Is it dependent on MI400? In which case, it might be more backup dated, et cetera. So just some ways to think about Sovereign for AMD next year.

LISA SU: Yeah, absolutely, Vivek. Thanks for the question. So look, we're really excited about the overall AI opportunity for us with MI355 and the MI400 series as we go through the back half of this year and into 2026. I think, there's a very large opportunity with, let's call it hyperscalers, some of the leading AI companies, as well as Sovereign. I think Sovereign is additive to that, for sure.

From the standpoint of what to expect, there are also some regulatory things that need to be worked through on the Sovereign side. But again, we're working closely with the administration as they go through the various regulatory decisions that need to be made.

But from my perspective, I think the fact that countries want their own Sovereign computing capability is very, very clear. I think we see that all over the world. The humane opportunity that you're referring to that we announced with the Kingdom of Saudi Arabia, I think is a great example of where together with their ambitions, our technology--

I think you heard from Tareq he was at our event-- saying that would start with MI355, that we would expect that that would continue on. I think what's attractive about our offering is our open ecosystem. And I think that really resonates with the Sovereign community. But to your original question, I think it's an additive opportunity, and it's one that we believe will continue to be very important for us going forward with both MI355 as well as the MI400 series.

VIVEK ARYA: Got it. And for my follow-up, I wanted to ask about gross margins for your MI product. So I understand in the early days, right, it has been dilutive. What kind of sales level is required for it to start becoming additive to margins?

And let's say if I fast-forward to Q4 and assume that your Q4 sales are growing year-over-year roughly the same rate as Q3 sales, should we expect gross margins to stay at these Q3 levels, or are there other plus/minus drivers we should think about in terms of gross margins as you go into Q4? Thank you.

JEAN HU: Yeah, Vivek, thank you for the question. The gross margin of our MI product, we said it's a little bit below corporate average. I think at this point, our priority is really to address the large, fast-growing revenue opportunities we have and provide customers a better TCO to really expand our presence in the marketplace.

I think the way to think about our gross margin, there are different dynamics, different customers, different generations. But also, our operations team has been continuing to really drive operational efficiency to improve MI family's gross margin. That has been ongoing.

So it's not necessarily really tied to, OK, the revenue level each quarter, but you should think about it as a trend in the longer term. It should improve continually going forward. Overall, the way I think about it is gross margin dollars. This is one of the fastest-growing market opportunities. For any financial metrics, gross margin dollars is what we try to grab as much as we can. Hopefully, that helped your question.

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

TIMOTHY ARCURI: Thanks a lot. So my question is on data center GPU. You did say that June is up year-over-year. So it sounds like it's maybe a little more than $1 billion. And you use words like "strong ramp" into the back half of the year. Can you give us just any color on what that means? Can you get to, say, $7 billion for the year? And can you give us some maybe milepost on what you're assuming for Q3? That would be great.

LISA SU: Yeah, thanks, Tim, for the question. I think what we said in the prepared remarks is that we are seeing a strong ramp from Q2 into Q3. MI355, we actually started production in June, so we had some shipments in the month of June, but it really is ramping as we go through this quarter in the third quarter.

So in terms of guideposts, we said it would grow year-over-year from last year. And that I think is a strong ramp. And then we would expect it to grow into the fourth quarter as well. The demand, I should say, what we're seeing from customers is, I think really positive around MI355. The way I would contrast it with maybe the MI300 ramp.

I think MI300 started with perhaps some smaller deployments. I think what we're seeing with MI355 is very competitive versus the B200, GB200 family of products. I think there's a strong desire to really use us at scale. MI355 is very strong for inferencing. We're also working with a number of customers on training.

And this is also an opportunity for us to build into the MI400 series as we go into 2026. So we're bullish on MI355 and where the AI opportunity is for us. And I think we're right on track to what we expected to be as we were going through the development of the roadmap.

TIMOTHY ARCURI: Thanks a lot. And then, Lisa, just on that point, also, you did talk about a new Developer Cloud. So obviously, you're beginning to lease back some of the capacity that you're selling into the clouds and the NIO clouds.

Is that going to be a material portion of the revenue you're going to recognize for MI355 in the back half of the year? Can you just talk about that and maybe how to think about how much demand that's going to stimulate and what the ultimate goal is for that cloud? Thanks.

LISA SU: Yeah. So there are a couple of things in that question. So let me answer. So the Developer Cloud is simply we want to make it super easy for developers to get on AMD Instinct GPUs. One could say-- again, if we look back at the MI300 family, we were very focused on the largest hyperscalers and the largest customers.

But there's a lot of interest in our GPUs across a number of customers who just wanted easier access. So by ensuring that Developer Cloud is there, that it has, ready-to-deploy containers, you can run training and inference easily, you don't necessarily have to make longer-term commitments, I think that's the purpose of the Developer Cloud.

I don't think it adds meaningfully to revenue in the second half of the year, but it certainly adds to customers getting experience with AMD. I think the larger revenue opportunities for us are really with large customer adoption as they ramp to larger deployments. And we're very actively trying to get those deployments up and running as soon as possible.

One of the things, just as a reminder, that the MI355 is-- given that it's a similar infrastructure to MI300, we actually think it's going to ramp very quickly and very well for customers. And I think that's one of the attractive portions of it as well.

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

ROSS SEYMORE: Hi. Thanks for letting me ask a question. Lisa, I want to go back to the Instinct side of things, and the MI355 ramp. Looks like the second half is going to ramp really significantly. You said it's going to be up year-over-year in the third quarter.

I believe a quarter ago, you said roughly the same thing in the MI308 is out of both numbers. So that shouldn't really matter. And I guess it would be upside. But I just wondered how have things changed from a quarter ago as far as the MI350 family adoption, especially because you launched a little bit early? Is the growth a little bit more than you would have expected a quarter ago, about the same, or a little worse? Just any sort of color on that would be helpful.

LISA SU: Yeah Ross, thanks for the question. I think the main thing I would say is I think the adoption is a bit faster than we might have expected. Again, whenever you launch a product, we want to make sure that we go through the full validation and all of that with our customers. I think there's a lot of interest, broad-based interest in MI355. So I feel like over the last 90 days, I think we've had significant sort of new customer interest. And that's certainly positive. I will say--

ROSS SEYMORE: I guess--

LISA SU: Oh, I'm sorry, Ross, I was just going to add, our engagements are-- I think the other piece is I think there's also a lot of excitement around MI400 and what we can do with the Helios rack. So there are a number of customers who, based on the strong roadmap that we're showing, want to get familiar and really work with us earlier in the life cycle, which, I think, again, positive.

ROSS SEYMORE: Great. Thank you for that. I guess is my follow-up, an earlier question you talked about a little bit below seasonality in the second half of the year for your client business. It seems like there's just-- I don't even know if seasonality is a framework that matters.

But how are you thinking about that for the second half as a whole for client and then gaming? It was just up a huge amount sequentially. In the second quarter, you described a little bit of what you're expecting there. But how do you think about seasonality for the second half in its entirety in the gaming side as well?

LISA SU: Yeah. Let me try, and then Jean might add if you want a little bit more color. So the way to think about it is, we do expect some sequential growth in client as we go into the third quarter. I would say single digit-type growth. We continue to see good traction for our products in that portfolio.

On the gaming side, I would call it flattish to Q2. And we're coming off of such a strong Q2 that I think flattish is actually to be expected. As we go into the fourth quarter, the dynamics that we would see is, we would see that the console business would actually be down substantially. So think about it as down strong double digits.

The customers usually build for the holiday season, sort of before that. And then that will be completed by the fourth quarter. So we would expect as the client and gaming segment that the segment would probably be down in the fourth quarter. So hopefully, that helps. Jean, did you want to add to that?

JEAN HU: No, I think you covered it.

MODERATOR: And the next question comes from the line of Joshua Buckhalter with TD Cowen. Please proceed with your question.

JOSHUA BUCHALTER: Hey, guys, thank you for taking my question. I wanted to ask about lead times on the Instinct family, both for the MI350 family and MI400. As you move into larger cluster sizes, which it sounds like you're doing at least with Oracle on 350, and then endeavor to do more so on 400, how much visibility and lead time do you need from your customers?

Because I would imagine the lead time for your parts is measured in months. But on the infrastructure side, in particular for the larger-scale deployments, those things are measured in years at this point. So maybe you could speak to the visibility specifically on the 400. Thank you.

LISA SU: Yes, sure, Josh. So, yes, our lead times are long, given all of the processing steps that we have to go through. Think about it as somewhere between 8/9 months, that type of thing. We have a very, very strong supply chain. We've been preparing for these ramps of both MI350 series and MI400 series. And that preparation is ongoing. So we feel like we have a very strong supply chain there.

In terms of visibility with customers, we're absolutely working with customers very closely on near term MI350 series deployments, getting those deployments up as quickly as possible. Again, one of the things about the MI350 series that is good is that it can go into existing data centers, just given the platform that is that it is in. So we have been certainly working with our customers there.

And then for the MI400 series, there are lots and lots of details in of full rack scale design implementation. And we're actively working with the largest customers right now on just ensuring that our Helios rack is fully compatible with their data center build out as we go into 2026.

So that visibility is important. I think that co-development, co-engineering is important as we get into the rack scale architecture. And the team that we've brought in has been extremely, extremely helpful in terms of both internal platform build out, as well as ensuring that we're working closely with our customers on their data center needs.

JOSHUA BUCHALTER: Thank you for that, Lisa. And maybe for Jean, I wanted to follow up on Vivek's question earlier on gross margins. So if we add back the charges in 2Q and then your gross margin implied in the guidance is roughly flat sequentially, and that's despite what's implied to be data center GPUs up meaningfully, sequentially.

It doesn't seem like consoles falling off in the third quarter. Can you maybe talk to the underlying drivers of how you're able to keep the flat gross margins despite what sounds like still margin-dilutive data center GPUs up significantly within the mix? Thank you.

JEAN HU: Yeah, Josh, thanks for the question. Yeah, we are guiding our Q3 gross margin around 54% and Q2 , you're right, excluding the $800 million charge. It was close to 54%. I think the gaming business actually is quite high. So the mix actually is unfavorable.

But we do have some tailwinds we have been really driving. First is we have been expanding our server business, which has really nice gross margin. And on the client business side, we are expanding the commercial PC business. That really helps us to drive the gross margin up.

In addition, we do have a really strong operational team. They are driving the gross margin improvement just from an operational efficiency perspective across the board. So overall, our objective is to continue to improve gross margin. Despite MI350 very strong ramp in Q3, we are able to continue to drive the margin up.

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

JOE MOORE: Great. Thank you. You used this language before of the kind of 10's of billions opportunity around MI400. Can you talk about the time frame when that might occur, and not to pin you down too much, and what would help you get to that level sooner rather than later? Should we think of that 2027 realistic outcome that you could be looking at $20 billion-plus? Just a little bit more color on that 10's of billions comment.

LISA SU: Yeah. Maybe without being specific, Joe, I can give you the way I look at it and back to this notion of, are we incrementally more confident? I think we're seeing a lot of positive signs in our AI customer adoption. I think the strength of the MI350 series, the very positive feedback that we're getting on MI400 from customers, the work that we're doing in terms of ensuring that we are fully ready for large-scale deployments of not just inference but training.

I think when we get to 10's of billions of dollars, we're talking about significant gigawatt-scale type deployments. And those would be important for us to get there. And we're certainly, I think, engaged with all the right customers that can enable that type of ramp. But I won't necessarily speculate on the exact time other than to say certainly, that would be our set of aspirations.

JOE MOORE: Great that's helpful. Thank you. And then, as these workloads evolve-- you've talked about inference and training as different opportunities for AMD. Are you seeing those start to come together? It seems like with inference, the reasoning models are requiring much higher complexity. Is rack scale more important to the inference market than you thought it might be? Just any color around how that complexity of inference is impacting you guys.

LISA SU: Yeah, I think that's absolutely true, Joe. With the proliferation of models, I think what we're seeing is GPUs continue to be the computing of choice as you think about all the models that are out there. And then as you go into distributed inference and some of the newer techniques, we are seeing the importance of the scale-up and scale-out architecture, which we are, investing in.

But I think the overarching thing is, I think we have a very competitive roadmap across the next couple of generations. I think that has now gotten strong customer validation. We're getting a lot of feedback from customers on where they would like to see us continue to add more resources and add more focus. So that is very helpful. And the key is to be a full-scale solution provider for these large customer deployments. And that's what we're working on.

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

AARON RAKERS: Yeah, thanks for taking the question. I do have a follow-up as well. I guess the first question is, when we look at the data center guide, Jean, you alluded to double-digit sequential growth. Obviously, the MI355 series kind of ramping. I'm curious, how could we conceptualize what you're expecting in the server side, and where do you think your market share is today in traditional enterprise servers outside of cloud?

JEAN HU: Yeah, I think when you look at our data center business, we do have a strong double-digit growth. Both server and the MI, both sides are growing sequentially. Of course, MI ramp is the most significant one. As far as server market share, we do think we continue to drive the market share up compared to Q1. Third party has not published report yet, but we feel really good about the Q2 market share increase versus Q1.

LISA SU: And Aaron, if I just add to that, one of the things that-- it's important for people to understand is, in some of the cloud CapEx numbers that have come out that have been quite positive, that is not only a GPU statement, but there's actually significant CPU CapEx in there as well.

We've started to see more robust forecasts going out a number of quarters on the server CPU side, because all of that AI content really requires traditional CPUs as well. So we're very bullish on the opportunity in servers. I think the team has really executed extremely well.

If you look at our portfolio now, Turin Genoa, our very well-adopted broadening workloads. Enterprise adoption is also increasing. So I think all of those are positive for the server opportunity in the second half of '25, as well as going into '26 and beyond.

AARON RAKERS: Yeah. And then as a follow-up, I'm thinking about the China, the MI308 opportunity. When we do see a license-- I think you alluded to this earlier-- it's going to take a little bit of time to ramp and get the supply chain to satisfy the demand. But I'm curious, the $800 million writedown that you had taken, is there no kind of finished inventory there? Does that come back? Do you have any reversal aspects of that once a license gets approved?

JEAN HU: Let me start first, then Lisa can add. The first on the $800 million. Majority of them are WIPs We really don't have on the shelf finished good we can ship immediately. So we do need to take time if we get a license.

AARON RAKERS: Thank you.

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

CJ MUSE: Yeah, good afternoon. Thank you for taking the question. I guess, Lisa, I was hoping you could level-set us on the Instinct ramp into 2026. How are you thinking about the timing of the handoff 350 to 400? How are you thinking about Helios contributions. And I guess very importantly, from a customer contribution perspective, how you might see that evolve from traditional hyperscalers to perhaps more sovereign and NIO cloud within the mix?

LISA SU: Sure, CJ. So certainly, as second half this year, it's all about MI355 ramp into first half of next year, I think the MI400 series development is right on track. The development of the Helios platform is also right on track. We would expect significant revenue contribution from Helios in 2026.

And then relative to the contribution of the various things, hyperscalers, versus NIO clouds versus Sovereign, I think it's a little early to really talk about the different pieces other than to say, you would expect that hyperscalers and let's call it NIO clouds that would be working for other large AI natives, maybe significant pieces of the initial ramp. And then Sovereign may come a little bit later in time, just given the timing of when different build-outs would happen. So hopefully, that gives you some color.

CJ MUSE: Very helpful. And then, Jean, I guess a question for you, with the sale of ZT for $3 million in cash and stock, and you only have $3 billion of debt outstanding, how are you thinking about the use of proceeds? Is there saving for a rainy day or bolt on acquisitions, perhaps more aggressive share buyback? How are you thinking about it today?

JEAN HU: Yeah, thanks for the question. Our business model actually generates a lot of free cash flow. As you see in Q2, free cash flow generation was $1.2 billion. So if we close the ZT sale, and we'll get more cash. Overall, our capital allocation principle continues to be the first phase investing, especially with the tremendous opportunities ahead of us.

And then we will continue to return cash to shareholders. We did a $1.2 billion repurchase in the first half of the year. We are committed to continue to return cash to shareholders through share repurchase.

MATTHEW RAMSAY: Operator, I think we have time for one more caller, please.

MODERATOR: The final question comes from the line of Ben Reitzes with Melius Research. Please proceed with your question.

BEN REITZES: Hey, thanks for squeezing me in here. I wanted to clarify a little bit on the $1 billion increase in sequential sales. It would seem like it's coming from GPUs, primarily. I was wondering if you could back that. And that's with nothing in China.

And if the answer to the prior question that GPUs are over a billion, that kind of puts you at a $2 billion run rate. And I was just wondering if that was accurate in terms of thinking. And then I have just a very quick follow-up. Thanks.

JEAN HU: Hi, Ben, thanks for the question. If you look at the sequential revenue increase, as I mentioned during the prepared remarks, we see data center strong double-digit increase, which includes both GPU and CPU. But GPU definitely drives the largest incremental amount increase. We also mentioned the client actually is going to increase sequentially.

In addition, the embedded business will return to sequential growth. So multiple businesses contributed to sequential increase, but the majority of increase really driven by MI355's strong ramp.

BEN REITZES: Great. And then if indeed that gets you pretty close to a couple billion dollars, if the MI300 comes in, do you see it at the same run rate that you exited and then you have the ability to get at that $800 mill run rate right away? Or do you think it'll take several quarters to ramp when you get the license? Thanks.

LISA SU: Yeah, Ben, it will take some time to ramp, just given-- particularly, today, we're sitting already in early August so I don't think you would see a lot of it in Q3. But certainly, as licenses would be approved, we would schedule that, and it would take a little while to ramp.

MATTHEW RAMSAY: All right, operator. Thank you very much. We appreciate everybody that joined the call today. And we just like to end the call now. Thank you.

MODERATOR: Yes, ladies and gentlemen, that does conclude today's teleconference. We thank you for your participation. You may disconnect your lines at this time.