AMD Q1 2025 Earnings Call

MODERATOR: Greetings, and welcome to the AMD first 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. And as a reminder, this conference is being recorded.

It is now my pleasure to introduce to you, Matt Ramsay, Vice President, Financial Strategy and Investor Relations. Thank you, sir. You may begin.

MATT RAMSAY: Thank you. And welcome to AMD's 2025 first quarter financial results conference call. By now, you should have had the opportunity to review a copy of our earnings press release in the accompanying slides. If you have not had the chance to review these materials, they can be found on the investor relations portion of amd.com.

We will refer primarily to non-GAAP financial measures during today's call. The full non-GAAP to GAAP reconciliations are available in today's press release and the slides posted on our website. Participants on today's conference call are Dr. Lisa Su, our chair and chief executive officer, and Jean Hu, 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 take note that Mark Papermaster, our executive vice president and chief technology officer, will attend the TD Cowen TMT conference on Wednesday, May 28, and Jean, who will attend the Bank of America Global Technology Conference on Tuesday, June 3.

Today's discussion contains 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 current expectations. Please refer to the cautionary statement on our press release for more information on factors that could cause actual results to differ materially.

You will also find detailed discussion of our risk factors in our filings with the SEC, in particular AMD's most recent quarterly report on Form 10-Q and annual report on Form 10-K. And with that, I would like to hand the call over to Lisa.

LISA SU: Thank you, Matt, and good afternoon to all those listening today. We delivered an outstanding start to the year, despite the evolving dynamics related to tariffs and the regulatory environment. Growth accelerated for the fourth consecutive quarter year over year, driven by strength in our core businesses and expanding data center and AI momentum.

Revenue and EPS both exceeded consensus estimates as instinct AI accelerator, EPYC and Ryzen CPU sales grew significantly year over year. As a result, first quarter revenue increased 36% year over year to 7.4 billion as our data center and client and gaming segments both grew by a large double digit percentage.

We expanded gross margin year over year for the fifth straight quarter and increased net income by 55% driven by a higher overall percentage of data center product sales and a richer Ryzen processor mix. Despite the uncertain macroeconomic backdrop, our first quarter performance highlights the strength of our differentiated product portfolio and execution and positions us well for strong growth in 2025.

Turning to the segments, data center segment revenue increased 57% year over year to 3.7 billion. We gained server CPU share driven by the ramp of our latest 5th Gen EPYC Turin processors and sustained demand for 4th Gen EPYC. Hyperscaler demand remained strong as cloud providers expanded EPYC deployments to power their critical infrastructure and public services.

More than 30 new instances launched from Alibaba, AWS, Google, Oracle, Tencent, and others in the quarter, including the initial wave of 5th Gen EPYC Turin instances. In addition, AWS launched new FPGA accelerated instances in the quarter, powered by EPYC processors with Xilinx Virtex FPGAs that are optimized for data and compute intensive workloads like genomics, multimedia processing, network security, and cloud-based video broadcasting.

Every major cloud provider is deep in development on Turin programs, with a steady stream of public instances and internal deployments expected to ramp into production over the coming quarters. Enterprise adoption of EPYC instances was very strong in the quarter. The number of EPYC-powered cloud instances activated by Forbes 2000 enterprise customers more than doubled year over year, including new wins with internet, native streaming, transportation, financial services, and social media companies.

For example, CrowdStrike achieved major performance and cost improvements by broadly deploying EPYC instances across its multi-cloud infrastructure. At the same time, we're also actively partnering with leading application and cloud providers to deploy EPYC optimized solutions tailored for specialized industry verticals.

Siemens launched their latest software-defined vehicle solution, powered by EPYC CPUs and Radeon pro GPUs on Azure, leveraging digital twin technology to significantly speed up automotive design and validation. Oracle launched a new version of its Exadata Database platform, which is used by more than half of the Fortune Global 100.

The latest Exadata X11M has been optimized for 5th Gen EPYC processors to deliver up to 25% faster performance in transaction processing and analytics compared to the prior generation. Turning to enterprise on-prem adoption, EPYC's CPU sales grew by a large double digit percentage year over year for the seventh straight quarter, driven by new public sector winds and high volume deployments with large automotive, semiconductor, financial services, retail, energy, and technology companies.

We have built significant enterprise momentum over the last few years as our partners expanded the number of EPYC-based platforms to more than 450, and we scaled our joint go to market programs. As a result, EPYC is now deployed by all of the top 10 telecom, aerospace, and semiconductor companies-- 9 out of the top 10 automotive, 7 out of the top 10 manufacturing, and 6 out of the top 10 energy companies on the Forbes 2000.

We expect enterprise adoption to accelerate over the coming quarters, as more than 150 tern platforms become broadly available from Dell, Cisco, HPE, Lenovo, Supermicro, and others. Looking forward, we see a clear path to continued share gains as customers ramp their 5th Gen EPYC offerings that deliver unmatched performance, efficiency, and TCO across every major cloud and enterprise data center workload.

We passed key milestones in April to begin manufacturing 5th Gen EPYC at TSMC's new Arizona fab, with first production shipments expected in the second half of 2025. Longer term, we announced our Next Gen EPYC Venice processors are the lead HPC products for TSMC's 2 nanometer process node.

Venice silicon is in our labs and performing well with bring up and validation progressing to plan to support a 2026 launch. Turning to our data center AI business, revenue increased by a significant double digit percentage year over year, as MI325X shipments ramped to support new enterprise and cloud deployments.

More than 35 MI300 series platforms are in production from all the leading service providers, supporting the expanding number of instinct GPU deployments with cloud, enterprise, and AI customers. Several hyperscalers expanded their use of instinct accelerators to cover an increasing range of generative AI search, ranking, and recommendation use cases.

We also added multiple tier one cloud and enterprise customers in the quarter, including one of the largest frontier model developers that is now using instinct GPUs to serve a significant portion of their daily inference traffic. The depth and breadth of our customer engagements continues to expand as breakthroughs in large scale AI models like OpenAI's O3 and DeepSeek's R1 drive increased demand for traditional inferencing and increasingly as a critical part of pre-training.

The industry-leading memory capacity and bandwidth of our instinct portfolio is ideally suited for these workloads, and we are actively working with multiple customers to scale instinct from single node deployments to distributed inferencing clusters. Training engagements also ramped in the quarter, as multiple tier one hyperscale AI and enterprise customers scaled instinct GPU clusters to train internal and next gen frontier models.

In parallel, we're making meaningful progress with sovereign AI deployments as countries expand investments to establish domestic nation scale AI infrastructure. In February, we announced a strategic partnership with G42 to build one of France's most powerful AI compute facilities powered by instinct accelerators.

On the AI software front, we significantly accelerated our release cadence in the first quarter, shifting from quarterly ROCm updates to delivering ready to deploy training and inferencing containers on a bi-weekly basis that include performance optimizations and support for the latest libraries, kernels, and algorithms.

We expanded our open source community enablement in the quarter, making significantly more instant compute infrastructure available to enable developers to automatically build, test, and deploy updates to ROCm code nightly. As a result, more than two million models on Hugging Face now run out of the box on AMD.

We're also enabling an increasing number of models to launch with day zero support for instant accelerators, including Meta's Llama 4, Google's Gemma 3, and DeepSeek's R1 models that were released in the first quarter. Beyond launch, we are delivering regular software updates that increase performance for new models. For example, in the weeks following the launch of DeepSeek's R1 model, we introduced ROCm optimizations that enabled MI300 to deliver leadership inferencing throughput.

We released ROCm 6.4 in the quarter, with major upgrades that increased training and inferencing performance across popular AI frameworks like PyTorch, Jax and VLLM. The release also adds multiple ease of use features, including new cluster management tools that simplify the scaling and optimization of large scale instant deployments.

Turning to our AI solutions capabilities. Earlier this quarter, we completed our acquisition of ZT systems, adding world class systems design expertise to complement our silicon and software leadership. With ZT, we can provide ready to deploy rack level AI solutions based on industry standards built with AMD CPUs, GPUs, and networking, reducing deployment time for hyperscalers and accelerating time to market for OEM and ODM partners.

The team is fully engaged in already co-designing with key customers on rack level designs optimized for our upcoming MI400 series and with customers and OEM partners to accelerate time to market for our MI350 series. We have received significant interest in ZT's manufacturing business and expect to announce a strategic partner shortly.

We began sampling our next Gen MI350 series with multiple customers in the first quarter and remain on track to begin accelerated production by mid-year. MI350 series performance is very strong, based on the advances in our cDNA for architecture. We designed cDNA for to deliver leadership performance across a wide range of AI workloads, increasing memory capacity and bandwidth 1.5x, adding support for new data types, and improving network efficiency to deliver 35x higher throughput and performance compared to MI300x.

Customer interest in the MI350 series is very strong, setting the stage for broad deployment in the second half of this year. As one example, we are partnering with Oracle to deploy a large scale cluster powered by MI355X accelerators, 5th gen EPYC Turin processors and Polara 400AINCs.

This multi-billion dollar initiative highlights the expanding AMD and OCI partnership and the growing demand for AMD instinct to power the next wave of large scale AI infrastructure. Looking ahead, our MI400 series development remains on track to launch next year. The MI400 series is designed to deliver leadership performance for both inferencing and training, scaling seamlessly from single servers to full data center deployments.

Early customer feedback has been very positive, marking a major step forward in our instinct roadmap and significantly expanding our AI accelerator Tam, as customers plan broader instinct deployments to power a larger share of their AI infrastructure. I'm looking forward to sharing more details on the MI350 series, future MI400 rack scale solutions, and the growing customer adoption of our instinct platforms that are advancing AI events on June 12.

Turning to our client and gaming segment, segment revenue increased 28% year over year to 2.9 billion. Client revenue grew 68% year over year, marking our fifth consecutive quarter of revenue share gains. We delivered record client CPU, ASP, driven by a richer mix of high end desktop and mobile Ryzen processors.

Desktop channel sellout increased by more than 50% year over year. We set new sellout records in multiple regions as our latest generation Ryzen processors became the CPU of choice for gamers, topping best seller lists at leading global e-tailers. To build on this momentum, we extended our desktop CPU portfolio with the launch of our 16 core Ryzen 9950 X3D processor that delivers significantly higher gaming and productivity performance than the competition.

In mobile, AMD-based notebook sell through was very strong in the quarter. We also saw strong demand for our latest generation AI PC processors as sales ramped, increasing by more than 50% quarter on quarter. The first notebooks powered by our new high end Ryzen AI max plus and the first mainstream Ryzen AI 7 and 5 300 series processors launched to very positive reviews.

These new processors set the standard for traditional computing and graphics performance, while also delivering unmatched AI capabilities and battery life positioning Ryzen as the CPU of choice for gaming ultra thin and commercial notebooks. Demand for AMD-based commercial PCs was also very strong in the quarter.

Ryzen pro PC sell through grew more than 30% year over year, driven by new end customer wins and an 80% increase from 2024 in the number of AMD-powered commercial systems from HP, Lenovo, Dell, and ASUS. We closed multiple wins with large auto, energy, healthcare financial services, and telecom companies in the quarter.

Looking more broadly across the PC market, we remain confident we can grow client processor revenue well ahead of the market in 2025, led by expanding adoption of our desktop channel and consumer and commercial notebook portfolio, as well as a richer mix.

Turning to our gaming business results, gaming revenue decreased 30% year over year as higher Radeon graphics sales were more than offset by lower semi-custom sales. While our semi-custom SOC sales declined year over year, console channel inventories have normalized, and demand signals have strengthened for 2025.

For PC gaming, we launched our Radeon 9070 series to strong demand as our new RDNA 4 architecture delivers leadership performance for mainstream gamers. First week sellout set a record, and was more than 10x higher than our previous best Radeon launch. Demand remains very strong, and we are working closely with our board partners to replenish inventory weekly and meet the sustained demand.

We also introduced FSR 4, our first machine learning based rendering technology that delivers significantly higher frame rates and more immersive gaming experiences. FSR 4 is already enabled in over 30 games, with support expected to reach 75 titles by year end. Turning to our embedded segment, first quarter revenue decreased 3% year over year to 823 million.

Embedded demand continues to recover gradually. We expect improving demand in the test and measurement, communications, and aerospace markets will drive a return to growth in the second half of 2025. We completed initial shipments of our cost optimized parts and Ultrascale plus FPGAs and second generation versal AI edge SOCs to meet growing demand for AI at the edge.

As a part of continuing to grow our embedded X86 business, we launched our EPYC embedded 9,005 series CPUs that deliver leadership performance for networking, storage, and industrial edge applications. CISCO selected our new EPYC embedded processors for their latest high end firewall solutions, and IBM is using them to power its latest storage scale, system 6,000 for performance-intensive enterprise analytics and AI workloads.

We also released our latest Vitus AI Software suite, expanding support for the latest models and accelerating edge AI deployment across a broader range of applications, further strengthening our leadership in the rapidly emerging edge AI market.

In summary, our strong first quarter results and second quarter outlook reflect the momentum we are building across our business. While we face some headwinds from the dynamic macro and regulatory environments, including the recently announced export controls for instinct MI308X shipments to China, we believe they are more than offset by the powerful tailwinds from our leadership product portfolio.

Against this backdrop, we remain confident we can deliver strong double digit percentage revenue growth in 2025 based on accelerating share gains with our latest generation of zen 5 EPYC and Ryzen CPUs and Radeon GPUs and ramping production of our instinct MI350 series accelerators in the second half of the year to support an expanded set of customers and AI workloads.

We also expect full year growth in our semi-custom business and for our embedded business to return to year over year growth in the second half of the year driven by the reduced inventory levels and improving demand environment. To capitalize on our unprecedented growth opportunities and deliver our next major growth arc, we are expanding investments in our product and technology roadmaps, go to market initiatives, and full stack AI software and data center scale solutions capabilities.

We're also doubling down on our execution to deliver, and where possible, accelerate our industry leading roadmaps. We view the current environment as a strategic opportunity to further differentiate AMD as we deliver an expanding product portfolio that combines leadership, compute, and AI capabilities for data centers, edge PCs, and embedded end devices.

Now, I'd like to turn the call over to Jean to provide some additional color on our first quarter results. Jean.

JEAN HU: Thank you, Lisa, and good afternoon, everyone. I'll start with a review of our financial results and then provide our current outlook for the second quarter of fiscal 2025. As a reminder, for comparative purposes, our first quarter fiscal year 2025 financial statement disclosures include the combination of our client and gaming businesses into a single reportable segment to align with how we manage the business.

We continue to provide distinct revenue disclosures for our data center, client gaming, and embedded businesses. We are pleased with our record first quarter revenue of $7.4 billion, exceeding the high end of our guidance, up 36% year over year, driven by 57% revenue growth in the data center segment and a 28% revenue growth in the client and gaming segment.

Revenue declined 3% sequentially due to lower revenue in the embedded and data center segments partially offset by sequential growth in the client and gaming segment. Gross margin was 54%, up 140 basis points from a year ago. Operating expenses were 2.2 billion, an increase of 28% year over year, as we continue to invest aggressively in go to market activities and in R&D to address the significant growth opportunities ahead of us.

Operating income was 1.8 billion, representing a 24% operating margin. Taxes, interest, expenses, and other was 213 million. For the first quarter of 2025, diluted earnings per share was $0.96, an increase of 55% year over year. Now, turning to our reportable segments, starting with the data center.

Data center segment revenue was 3.7 billion, up 57% year over year, primarily driven by continued CPU server share gains across both cloud and enterprise customers and the strong growth of AMD instinct GPUs. On a sequential basis, data center segment revenue decreased 5%.

Data center segment operating income was 932 million or 25% of revenue, compared to 541 million or 23% a year ago. Client gaming segment revenue was 2.9 billion, up 28% year over year, driven primarily by strong customer demand for our latest generation zen five AMD Ryzen processors, partially offset by lower semi-customer revenue.

Client revenue was 2.3 billion, up 68% year over year. More than half of the growth was driven by higher ASPs from a richer mix of high end Ryzen processors. On a sequential basis, client and gaming segment revenue increased by 2%, primarily driven by stronger than seasonal performance of our client product portfolio and the increased semi-custom product revenue.

Client and gaming segment operating income was $496 million or 17% of revenue, compared to $237 million, or 10% over a year ago, driven by operating leverage on higher revenue. Embedded segment revenue was $823 million, down 3% year over year. Embedded demand continues to recover gradually.

Sequentially embedded was down 11% consistent with our expectations. Embedded segment operating income was $328 million, or 40% of revenue, compared to $342 million, or 41% a year ago. Turning to the balance sheet and the cash flow.

During the quarter, we generated a 939 million in cash from operations, and our free cash flow for the quarter was 727 million. We returned 749 million to shareholders through the repurchase of common stock and our repurchase program. We have a 4 billion remaining in our share repurchase authorization.

At the end of the quarter, cash, cash equivalents, and short term investment was 7.3 billion. Within the quarter, we reached 1.5 billion of debt and issued 950 million of commercial paper to help fund our acquisition of ZT systems, which was completed on March 31. Now, turning to our second quarter 2025 outlook.

As a reminder, in April, a new export license requirement was put in place for MI308 shipments to China-- the impact of which is included in our guidance. We expect revenue to be approximately 7.4 billion plus or minus 300 million. This includes an estimated 700 million revenue reduction as a result of the new export license requirement.

Despite this headwind, the midpoint of our guidance represents 27% year over year revenue growth. For the full year, 2025, we estimate the revenue impact due to the export license requirement to be approximately 1.5 billion. Sequentially, we expect client and gaming segment revenue to increase by double digit percentage, embedded segment revenue to be flattish, and we expect data center segment revenue to decrease due to the exclusion of MI308 revenue.

In addition, we expect second quarter non-GAAP gross margin is estimated to be 43%, inclusive of approximately $800 million in charges for inventory and related reserves. Excluding this charge, our non-GAAP gross margin would be approximately 54%. Non-GAAP operating expenses to be approximately $2.3 billion, which includes approximately $50 million in OpEx due to the addition of the ZT systems design team.

The financials for the ZT manufacturing business will be reported as discontinued operations starting in the second quarter. We expect net interest and other expenses to be 5 million due to the debt associated with the ZT system transaction. Non-GAAP effective tax rate to be 13% and the diluted share count is expected to be approximately 1.64 billion shares, which includes nine million shares related to the ZT transaction.

Looking forward, despite the ongoing macro and trade policy related uncertainties, we believe the investment we are making will position us well to address the large growth opportunities ahead, as AI expand the use of high performance computing across all our end markets.

In closing, 2025 is off to a strong start as we continue to execute on key strategic and financial goals. We delivered strong top line revenue growth, expanded gross and operating margins, and closed the key acquisition of ZT systems to expand and accelerate our data center GPU and systems roadmaps. With that, I'll turn it back to Matt for the Q&A session.

MATT RAMSAY: Thank you very much, Jean. John, can you go ahead and pull the callers for the Q&A session, please? Thank you.

JOHN: 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 question queue. You may press star 2 to remove yourself from the queue.

For participants using speaker equipment, it may be necessary to pick up your handset before pressing the star keys. In the interest of time, we ask that you please limit yourself to one question and one follow up. Thank you. One moment please, while we pull for questions.

And the first question comes from the line of Joshua Buckwalter with TD Cowen. Please proceed with your question.

JOSHUA BUCKWALTER: Hi, Dean. Thank you for taking my questions, and congrats on the results. I was hoping you could maybe expand on the drivers of upside in both the print and in particular the guide. How should we think about TCU growth by segment. And wanted to double click on client in particular. That business is up 67% year over year in the first quarter.

And there's obviously a lot of concerns on pull ins. So I was hoping you could walk through some of the drivers of the strength in Klein in particular and how you're thinking about that in Q2. Thank you.

LISA SU: OK, great. Josh Thanks for the question. Look, we were very pleased with our performance in Q1. We actually saw strength across a number of our businesses. We saw strength certainly in the client business, very strong desktop performance. We saw strength in our gaming business as well, which was really due to our strong Radeon launch. And we also saw some strength in our data center business across both stronger CPU and GPU.

So those are some of the drivers for our Q1 performance. And in particular, on your question of client performance, we've certainly looked very carefully at the ordering patterns and what customers are telling us. We have not seen a lot of tariff-related activity in that business. I would say, though, what we have seen is a real stronger mix and strength in our overall ASPs.

So the desktop channel, which is an area where we have a very strong gaming products right now, actually performed well above seasonality in Q1. And that is really the strength of the ASPs there. So that's what we saw in Q1.

And then to your question about the guide for Q2, as Jean mentioned, we do have the new export control limitation on MI308. So we have taken out that revenue, which is a $700 million headwind in Q2. But with that, we have a strong outlook given the strength across the rest of our businesses.

So we continue to see strength in client going into the second quarter. Again, the desktop business continues to perform above typical seasonality. We're also seeing the beginning of the commercial ramp, which is a place where we have traditionally been quite underrepresented. We see continued strength in gaming, I would say, much better than typical seasonality.

That is really our AIB business with the Radeon products are ramping. As well as consoles have now drained all of their inventory. And so they are starting their ramp into the year. And from a data center side, we see sequential growth on the CPU side. We see the GPU right on track minus the China export controls. And so for all of those reasons, we're pleased with where the performance of the business is right now.

JEAN HU: I'll just add one point to what Lisa just said on the client business. We have really strong performance in Q1. Sequentially, client revenue is largely flattish versus Q4. When you look behind it, our unit actually declined a double digit. So the revenue flattish is largely driven by the ASP increases sequentially due to the richer mix Lisa just mentioned.

JOSHUA BUCKWALTER: Thank you for all that color. To follow up, I wanted to ask about how the ex-China X308 instinct family performed in the quarter and how you're thinking about the back half of the year. I think you mentioned in the prepared remarks significant double digit year over year.

Could you maybe provide some color on how instinct did in the first quarter, how you're thinking about the first half ahead of the 350 ramp in next month? Thank you.

LISA SU: Sure. So on the instinct ramp, I would say the a two, one performance of data center GPU was in line with maybe a little bit better than expected. I think the key point that we've always said about the instinct ramp is very excited about the MI350 launch. We're right on track for that launching mid-year.

I would say, customer interest has been very high. So from a competitiveness standpoint, we feel really good about where it's positioned. Overall, I think one of the advantages that we have with the MI350 launch is that from a systems overall environment, it's actually very similar to the MI300. So we believe it's going to ramp fast.

And we already have a couple of deals that have been announced, including a very important relationship with Oracle in terms of the MI350 series for a number of joint customers. So we're excited about the overall AI business. I think we continue to see strength there. I know there are some uncertainties as it relates to tariffs and other things.

But this is one of those areas where from an infrastructure standpoint, there continues to be investment in AI infrastructure. And so with that, we would expect strong growth into the second half of the year.

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

TIMOTHY: Thanks a lot. Lisa, you said that data center GPU grew significant double digit, but it was like $600 million last March. So I would think that I think a lot of us thought it was going to be like 1.7 to 1.75. So is that the wrong way to interpret that? Because it seems like it more than it went up triple digits at least. So can you help us there?

And also, I'm curious, the additional 800 million that sort of has to come out of from the ban, does that all come out in September, or is there some remnants of that have to come out in the fourth quarter as well?

LISA SU: So again, what I would say is the data center GPU business did perform very well in the first quarter. I think we have to go back and look at what you had for first quarter 2024. But overall standpoint, it performed right where we would expect. Relative to your conversation, as to where does it come out? I would say the vast majority comes out in the September quarter.

So think about Jean mentioned 1.5 billion. You would see the majority of it in Q2 and Q3 with very little in Q4. So we had always expected that the fourth quarter because it would be very focused on the MI350 family would be non-China revenue. And that's how it was planned.

TIMOTHY: Got it. And then if Jean, just on the inventory, it was up a lot. Is that just due to ZT or is there something else happening there? Thanks.

JEAN HU: Oh, on the inventory side, we built some inventory primarily to support the very strong client and the server ramp and also the second half data center GPU ramp. As you probably know, the lead time is really long. To build for the Q3, Q4 ramp, we really need to start wafers right now. That's why the inventory has increased.

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

HARLAN: Hey, good afternoon. Thanks for taking my question. I know there's been a lot of focus on your upcoming MI350 series, Lisa, but MI400 next year is where you potentially close the competitive gap in a big way, right? You're bringing frontier class model training performance GPU, interact scale, solution.

More and more of the challenges has been standing up these rock scale platforms power, cooling, footprint, networking, connectivity, telemetry, et cetera, right? Lots of well-telegraphed issues with standing up these rock scale architectures. So as you've shared your MI400 rock scale solution architecture with customers, like what is the AMD team doing to potentially address the ease of these deployments with the MI400? And just in general, what's been the overall feedback been like on MI400?

LISA SU: Yeah, Harlan, thank you for the question. I think look, we're excited about the MI350 series launch that's coming up, but we are extremely excited as well about the MI400 series and the roadmap there. I think we've been very active with customers on our roadmap. As you know, this is one of those areas where you absolutely have to be planning many quarters in advance for that.

One of the primary reasons we acquired ZT systems was exactly to address this rack scale architecture. And so from that standpoint, the closing of the ZT acquisition has been very timely. What we're doing right now is, together with our ZT design team, as well as our customers design teams and our own systems design capability, really actively planning what those rack scale systems are going to look like.

I would say, the MI400 series enthusiasm from customers is high, and there's a lot of activities that are going on right now to ensure that we do in fact learn from some of the-- let's call it some of the challenges that have occurred with some of the recent deployments.

HARLAN: Thanks for that. And then I continue to be impressed. I mean, seven consecutive quarters of strong year over year growth in your EPYC enterprise and on prem traction. You have high 30s, low 40s type share of the overall server market and enterprise and on prem. Your share is probably in the sort of low 20% range but significant share momentum. Can you just remind us like what has the AMD team done?

What have you put in place go to MarketWise to drive this strong tailwind here in what has been a very, very tough market segment to crack?

LISA SU: I think there are a couple of things, Harlan. First of all, the strength of the product cannot be undersold. At this moment with 5th Gen EPYC, the overall cloud adoption has been fantastic. And then on the enterprise side, we've really broadened the product portfolio for Turin, that includes, let's call it, low core count up through the highest core count and frequency ranges. So that's very helpful.

But probably the largest impact has been in go to market. In the go to market space, we have added significant headcount and capability to address end users directly. And with the use cases, I think some of the things that we talked about across industries, we're actually learning from each deployment and replicating that across many of the Industrial partners.

So overall, I think it's been a strong effort on enterprise, and we're really still in the very early stages of that. I would say we're still quite underrepresented in enterprise, but with the platform coverage and the processor coverage, I think we feel good about the opportunities.

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

AARON RAKERS: Yeah. Thanks for taking the question. Going back to the data center business and particularly the GPU business, I think last quarter you had alluded to the fact that you'd expected the data center revenue to be roughly flat in the first half of the year. I guess if we were to take out the 700 million impact from China, what would the expectation still be flat for the year? Is that a fair assumption?

JEAN HU: So, Aaron, so you're right. Last time we did mention the first half data center GPU is flattish versus second half. The way to think about what Lisa mentioned is the 1.5 billion impact largely will be in Q2 and Q3. And so when you take out $700 million in Q2 and majority in Q3, that is what the impact in Q2 and Q3.

But remember what Lisa mentioned is we do see second half weighted. As we launch MI355, we will see significant ramp. Year over year, we see strong double digit growth over our data center business and the GPU business also.

AARON RAKERS: And then as a quick follow up, kind of thinking about the gross margin, obviously, this quarter's guidance is reflective of the charge that you're taking. Should we assume that in the back half with mix attributes to be considered that you would see a return to that 54 plus percent gross margin in the second half of the year? Is that a fair assessment?

JEAN HU: Yeah. Aaron, thank you for the question. There are a few takes on the gross margin. If you think about the Q2, excluding 800 million charge related to the MI308, our gross margin actually is around 54%. So at the company level, the mix is less favorable because the client and the gaming business is growing sequentially.

But we do have a few drivers to drive the gross margin up. First, as I mentioned earlier, if you look at our client business, the gross margin has been improving because of the richer mix of our latest generation product portfolio. That really helps. And also secondly, within data center, when we expand the enterprise market share, we do see gross margin improvement.

Of course, in addition, MI308 data center GPU gross margin is on the low end of our data center GPU margin. So that also helps us. Overall when we think about the second half, we actually think the gross margin will improve slightly because data center continue to be very strong growth driver-- number one growth driver second half versus the first half, which will be partially offset by continued strength on the client and the gaming side. Hope that answers your question.

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

THOMAS O'MALLEY: Hey, Lisa, and others. Thanks for taking my question. I really appreciate it. And Jean, thanks for the helpful answer there. I just wanted to understand your view on system-based architectures and whether you feel like you have what you need right now. Obviously, you will think 1.0 is coming out.

You can use third party providers to do the interconnect. ZT system does do a lot for you in terms of the system architecture. But from the interconnect side, do you think that you need more? Is that something that you're going to do internally? Look externally. Just want to understand where you think the portfolio is today and whether you can address system-based architectures with what you have today.

LISA SU: Sure, Tom. Absolutely. I think we feel like we have all the pieces required as well as deep partnerships in the ecosystem. And I consider it a system-level optimization between CPU and PPU and networking capability, RAC scale architecture. I think all of those pieces are things that we are investing in.

And we're also partnering with others in the industry who are offering these capabilities. I think when we look at the architectures that our customers want, our customers are really asking for one, that we have a reference architecture that works, but also that we work with them as they want to interchange various pieces, particularly on the networking side.

I think there are a couple of different solutions out there. And we are very much focused on ensuring that we interoperate across the spectrum.

THOMAS O'MALLEY: Helpful. And then if we look at the full year, I mean, we'll get the units with the filing. But it looks like there's some material share gains here in the first quarter. When you look at the full year, just to level set us on share gains versus market growth, could you maybe talk about what you see the client business growing as a base level?

And then just obviously, it's difficult to predict where share will go. But just any comments on what you're seeing thus far. Is a couple points of shares kind of what you're seeing in the first quarter as well? We'll get a little more later, but mostly just on the market growth for 2025. Thank you.

LISA SU: Sure, Tom. So if you're asking about share in the client business, I think that was the conversation. Look, we are very pleased with our client business performance over the last couple of quarters. I think we are seeing unit growth, particularly in desktop, but where we're seeing probably the most growth is overall revenue share.

And so we're gaining share in the right places, which is in high-end, notebook and commercial as well as in desktop overall. So from that standpoint, that's where we think we're going. As we go through the year, I know there's a good amount of conversation about what happens in the macro and what happens with tariffs. And does that change things going forward?

We are spending quite a bit of time ensuring that we are aligning with our customers, looking at inventory levels, looking at consumption and overall sell through. And we believe that we have a good overall inventory position. And there is not, let's call it a tremendous amount of pull-ins or other things that are coming into play. And, we will continue to be very agile in how we look at that going forward.

JOHN: 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. I had two questions as well. On the first one, just near term, Lisa, did your GPU sales grew sequentially in Q1? How much was MI308 in that number? And if you look at 2025 overall, do you think GPUs can still grow despite this China headwind that you mentioned relative to the 5 billion plus you did last year?

LISA SU: Yeah, sure, Vivek. So let me answer the second question first. We absolutely believe the data center GPU will grow, and we think it will grow strong double digits. We had a plan that was second half weighted and it still is. Relative to the MI308 situation, it's certainly a headwind, but one which we think is well contained, given everything else that we have going on.

And relative to the Q1 performance of data center GPU, it was down very modestly from Q4, which is what we expected. We did see good overall demand actually in the first quarter, driven by MI325. So we had a significant adoption by a large foundational model company, which was very positive there.

And as we go forward, we expect that we will continue to broaden both customers as well as workloads within our current customers for the instinct portfolio.

JEAN HU: And, Vivek, in Q1 MI325 and MI300 are majority of our revenue.

VIVEK ARYA: Great. And then longer term, Lisa, in the past, you described, I believe, almost a $500 billion or so addressable market for AI accelerators. How much of that roughly is China, because that now seems to be somewhat restricted for US companies?

And then also related to that, how should we think about these AI diffusion rules that I think there is an implementation date that is coming up on May 15? So I'm curious what you have heard. So just the implication of China restrictions and these AI diffusion rules on thinking about the addressable opportunity for you longer term. Thank you.

LISA SU: Vivek, I think it's a good question. I think overall, it is a very dynamic market. So you will appreciate that. On the China export controls, I think we always expected that there would be some amount of, let's call it limitation on leading edge GPUs going into China. So that was factored in to our TAM expectation when we talked about $500 billion.

So I don't think that dramatically changes the TAM. But what I will say is on the AI diffusion side, we're very actively working with the government as they're thinking through these rules. And it's a very fine balance that we have to have. At the end of the day, when we look at the US AI companies, we have leading edge technology.

We want to ensure that the rest of the world can really use us as the primary platform. So I think it will be important to work through the AI diffusion rules and all of that as we think about longer term TAM. And we're certainly spending quite a bit of efforts trying to ensure that it's well understood, the importance of the overall ecosystem and having the rest of the world really adopt the US ecosystem, given our strength in leadership overall.

JOHN: And the next question comes from the line of C.J Muse with Cantor Fitzgerald. Please proceed with your question.

C.J MUSE: Yeah, good afternoon. Thank you for taking the question. I wanted to revisit your assumptions around client. If you were to just flat-line the Q1 actual, you would grow the business about 30%. You're obviously very bullish on taking share. You talked about huge tailwinds from ASPs. But curious, when you put it all together, how should we think about traditional seasonality into the second half, particularly with the potential of some pollings here in the first half?

LISA SU: First, C.J, it's a fair question. Look, we want to be very clear that our client business performance is primarily driven by the strength of the product portfolio. And it's driven by some of the desktop channel products that traditionally are not so well tracked, if you look at the IDCs of the world.

We are planning for let's call it a second half sub seasonal, given that we're off to such a strong start in the first half of the year. And that is what we're putting into our internal planning numbers. So you wouldn't see necessarily typical seasonality since the first half is better than seasonal.

That being the case, I think we feel strongly that from a consumption basis standpoint, we can see the data. So when we look at the Q1 performance, it was a very, very strong Q1 in terms of sellout and consumption for our desktop business. And as we start Q2, we're now four weeks into it. We see those patterns continuing.

So we're in an upgrade cycle right now. Gaming CPUs are usually purchased when they're gaming GPUs that come out in new cycles. And I think we're benefiting from that on both the CPU and the GPU side, which is great. We're very happy with that, and we're ramping up production to ensure that we keep the channel full.

C.J MUSE: Very helpful. And then I guess looking to next year, can you talk about 400 series and rack level solution go to market strategy? You talked about trying to obtain partners. Is there a certain number that you're targeting? And then how are you thinking about of getting through some of the learning curve challenges of getting the rack scale, working with your OEM partners such that you can deliver that ramp in 2026? Thanks so much.

LISA SU: Sure, C.J. So I think the right answer is we're getting a very early start, and that's what we have to do so that we maximize the overall learning cycle that is required for rack scale solutions. We are working very closely with a number of our hyperscale partners today to define those solutions and make sure that we're thinking about all of the various areas that could require work.

And we're also working with our OEM partners who also have, let's call it, learned quite a bit over the past couple of months and quarters as other rack scale solutions have been coming online. So I think we're doing everything to, let's call it, move ahead the learning cycle. And again, we have the benefit of the MI350 series being a relatively, let's call it, not large lift. And so the focus on the rack scale stuff is on MI400.

JOHN: And the next question comes from the line of Stacy Rasgon with Bernstein Research. Please proceed with your question.

STACY RAGSON: Hi, guys. Thanks for taking my questions. For the first one, given the China data center GPU headwinds in Q2 and Q3, do you think that GPU business actually grows year over year in Q2 and Q3? Understanding your comments for the full year on it, but do you think, given those headwinds in Q2 and Q3, it can actually grow year over year?

LISA SU: I think you're-- let's see. Stacy, the best way to answer that question is, in Q2, it's not going to grow year over year, just given what we've said about the 700 million coming out of Q2 and how we had previously talked about the evolution. But we do believe that we'll grow year over year going forward in Q3 and Q4, certainly for us to do the full year with strong double digit growth.

STACY RAGSON: So you do think it can grow year over year in Q3? OK. For my second question, I wanted to ask about the trends in Q1. So you said it was data center GPU was down, I guess modestly in Q1 as expected. But again, if I go back to your double digit year over year comments, it couldn't have been any more than 1.4 billion in Q1 for GDP.

And it feels like it's less than that, which means it would have been down at least 20% sequentially, maybe more, which also implies that the server CPUs in Q1 were up sequentially, which is also well above seasonal, similar to clients. I guess what I'm asking is, are those trends correct? Am I modeling that correct? And I guess, what are the implications in that case of server CPUs actually up well above seasonal in Q1 given this environment?

JEAN HU: No. I think, Stacy-- this is Jean-- I think when you think about the Q1 data center performance, it's declined 5%. So it's a little bit better from a server perspective because it is declined sequentially. Same thing like data center GPU. Like Lisa mentioned earlier, it did decline.

So I think that is the overall data center performance. I think I don't know about your model, but that is how we really look at the numbers, how we think about it.

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

ROSS SEYMORE: Hi, guys. Thanks for letting me ask a couple of questions. I'm going to go to the embedded space. I know it's not the biggest one, but everything else has been addressed pretty detailed. You mentioned the second half getting up to year over year growth. It seems like that requires some significant double digit growth sequentially in both quarters, just to get the full half there. What gives you the confidence in that sort of ramp?

JEAN HU: Ross, thank you for the question. On the embedded side, we started to see a gradual recovery. I think there are signs, especially the order pattern, the book to bill ratios we see improving. In the aerospace defense and also test the measurement side, we see a very visible improvement industrial side.

The improvement is less so. Their inventory still are among different customers. But overall, the trend, the demand, pattern does improve. I think Q2, we did guide the sequentially flattish. And I think we'll start to see Q3, especially Q4, you will see year over year increase, especially in Q4.

STACY RAGSON: Great. Thanks for that. And I guess my last question would be on the OpEx side of things. You guided to the overt number for the second quarter, 2.3 billion. You said there's 50 million from ZT in there. Is that the entirety of the ZT side of things, or what should we think for full year OpEx or the second half however you want to discuss it?

JEAN HU: Ross, thank you for the question. For the ZT design team, we view it as quarterly. The incremental OpEx is about 50 million. That 2.3 billion includes everything from ZT because we closed the transaction on March 31.

I think when you look at the overall, our OpEx increase year over year, we continue to drive revenue growth to increase more than OpEx. Looking at the Q2 at the midpoint of our guidance, revenue will be increasing 27%. And we do expect the earnings per share growing much faster than the top line revenue growth. So OpEx side will be very disciplined to continue to manage it.

MATT RAMSAY: Operator, I think we have time for one more caller, please. Thank you.

JOHN: No problem. And the final question comes from the line of Joe Moore with Morgan Stanley. Please proceed with your question.

JOE MOORE: Great. Thank you. One of the things your cloud customers have been talking about is this kind of growth in inference costs, the sort of reasoning models using a lot of inference compute and some tightness. Can you talk about that from AMD's perspective? Are you seeing that in your business? Does that change the focus that you have going forward?

LISA SU: Sure, Joe. So I think overall, what we're seeing is that with these new reasoning models, the inferencing is more important. And there's also a move to more distributed inferencing. So I think that plays into our strengths. I think we have demonstrated with MI300 that we are an excellent inference solution.

And that holds true for 35 and 350 series as well. So we continue to see with our memory bandwidth and memory capacity advantages, that's a positive. I will say that as we're going into this, the number of workloads that we're seeing overall is expanding. So we're seeing both training and inferencing as important workloads that we're working on.

And our customers continue to demonstrate-- I think the desire that we're seeing probably from a trend standpoint is that there are many models that people are using today. So they're not necessarily using one model. They're actually using several different models. And so the optimizations around that are the things that we're doing with our ROCm software suite.

JOE MOORE: Great. And then just an update on your thoughts on competing with custom silicon with ASICs in the AI space. Most of your largest customers also have a custom silicon offering. So will they invest in both AMD and ASICs? And just how do they decide how to apportion that investment?

LISA SU: I mean, Joe, I view them as really two different things. I think one of the primary aspects, as we've talked about, the $500 billion TAM and the opportunities there, look, we think ASICs have a place. We happen to think GPUs have a larger piece of that because the models are changing so much.

And from our standpoint, it's really important to have competitive TCO. And people want choice to get there, especially as inference costs become so important. And we're working on trying to expand the overall inferencing capability out there. So I don't think it's an either/or. I think it's a let's get the best solutions out there, and we will certainly believe that we're very competitive in inferencing. And I think we are also becoming a much more solution for training as well.

JOHN: And ladies and gentlemen, that does conclude the question and answer session. And that also concludes today's teleconference. We thank you for your participation. You may disconnect your lines at this time.