

## AMD 1Q 2026 Earnings

---

**JOHN:** Greetings and welcome to the AMD first quarter 2026 conference call. At this time, all participants are in a listen-only mode. A question and answer session will follow the formal presentation. If anyone should require operator assistance during the conference, please press star 0 on your telephone keypad. And please note that this conference is being recorded. I will now turn the conference over to Matt Ramsay, Vice President of Financial Strategy and IR. Thank you, Matt. You may begin.

Thank you, and welcome to AMD's first quarter 2026 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 a 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 on today's conference call are Dr. Lisa Su, our chair and CEO, and Jean Hu, Executive Vice president, CFO, and treasurer. This is a live call and will be replayed via webcast on our website. Before we begin the call, I would like to note that Jean Hu will present at the Bank of America Global TMT conference on Tuesday, June 2, in San Francisco.

Today's discussion contains forward-looking statements based on current beliefs, assumptions, and expectations. Speak only as of today and as such, involve risks and uncertainties that could cause actual results to differ materially from our current expectations. Please refer to our 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 in today. We delivered an outstanding start to the year, driven by accelerating demand for AI infrastructure across our portfolio. Growth was broad-based, with every segment increasing year over year led by 57% data center revenue growth. First quarter revenue increased 38% year over year to \$10.3 billion. Earnings grew more than 40%, and free cash flow more than tripled to a record \$2.6 billion, driven by significantly higher sales of EPYC CPUs, Instinct GPUs, and Ryzen processors.

These results mark a clear inflection in our growth trajectory and a structural shift in our business. Data center is now the primary driver of our revenue and earnings growth, and as AI adoption scales, demand is increasing not only for accelerators, but also for the high performance CPUs that power and orchestrate those workloads.

Turning to our segments, data center revenue increased 57% year over year to a record \$5.8 billion, led by strong demand for our EPYC CPUs and Instinct GPUs. In server, we delivered our fourth consecutive quarter of record server CPU revenue. Revenue increased more than 50% year over year, with sales to both cloud and enterprise customers each growing more than 50%.

Share gains accelerated year over year, reflecting the ramp of 5th Gen EPYC turned CPUs and continued strength of 4th Gen EPYC processors across a wide range of workloads. In cloud, AI was the primary driver of growth in the quarter, as every major cloud provider expanded their EPYC footprint to support a broad range of AI workloads, from general purpose compute and data processing to head nodes for accelerators and emerging agentic applications.

EPYC-powered cloud instances increased nearly 50% year over year to more than 1,600, with instances optimized for virtually every enterprise workload and expanded availability across the largest global cloud providers. In enterprise, demand accelerated with record revenue and record sell through in the quarter. We expanded our customer base with new wins across financial services, health care, industrial, and digital infrastructure companies, while also building momentum with mid-market and SMB customers.

We are well-positioned to continue gaining share as more enterprises standardize on EPYC across on-prem and hybrid environments based on our leadership performance and TCO. Looking ahead, our sixth Gen EPYC Venice processor, built on our Zen 6 architecture and 2 nanometer process technology, is designed to extend our leadership across cloud, enterprise, and AI workloads.

The Venice family spans a broad set of CPUs optimized for throughput, performance per Watt, and performance per dollar, including Verano, our first EPYC CPU, purpose-built for AI infrastructure. Across the portfolio, Venice widens our competitive advantage, delivering substantially higher performance per socket and per Watt versus competitive x86 offerings, and more than 2x throughput per socket versus leading ARM-based AI solutions.

Customer demand is very strong, with more customers validating and ramping platforms at this stage than with any prior EPYC generation, and we remain on track to launch Venice later this year. Looking more broadly, we are seeing a meaningful acceleration in customer demand, driven by the rapid scaling of AI workloads across both cloud and enterprise.

Inferencing and agentic AI are increasing the need for server CPU compute, as these workloads require additional CPU processing for orchestration, data movement, and parallel execution, in addition to serving as the head nodes for GPUs and accelerators. As a result, we are seeing both stronger near-term demand and deeper engagement with customers on long-term capacity planning.

At our Financial Analyst Day in November, we outlined a server CPU market growing at approximately 18% annually over the next three to five years. Based on the demand signals we are seeing today and the structural increase in CPU compute requirements driven by agentic AI, we now expect the server CPU TAM to grow at greater than 35% annually, reaching over \$120 billion by 2030.

In response to this demand, we are working closely with our supply chain partners to meaningfully increase our wafer and back end capacities to support this growth. As a result, we now expect server CPU revenue to grow by more than 70% year over year in the second quarter, with robust growth continuing through the second half of 2020 and into 2027 as we ramp our next generation EPYC processors.

Now turning to our data center AI business, revenue grew by a significant double-digit percentage year over year as adoption of Instinct accelerates across cloud, enterprise, sovereign, and supercomputing customers. We're seeing strong momentum as customers move from pilots to large scale production deployments, particularly in inference where our leadership memory capacity and bandwidth are key advantages. This momentum is driving deeper long-term customer engagements, including large scale multi-generation deployments.

A key example is our expanded strategic partnership with Meta to deploy up to 6 gigawatts of AMD Instinct GPUs spanning several product generations. Our agreement includes a custom GPU accelerator based on our MI450 architecture, co-designed to support Meta's next generation AI workloads. Shipments are on track to begin in the second half of the year, leveraging our Helios rack scale architecture, which integrates Instinct GPUs with EPYC Venice CPUs to deliver fully optimized, high performance AI infrastructure.

Together with our previously announced OpenAI partnership, these engagements position AMD as a core partner to the world's largest AI infrastructure builders, with deep co-engineering relationships and multiyear visibility into large scale deployments. More broadly, Instinct adoption continues to expand across AI-native and enterprise customers for both training and inference workloads.

Existing partners are expanding Instincts across a broader set of workloads, while a growing number of new partners are deploying production AI workloads on Instinct, highlighting the maturity of our hardware and software stack. On the software front, we continue to make strong progress with ROCm, improving performance, scalability, and enabling customers to reach production faster. In our latest MLPerf results, MI355X delivered strong competitive performance across the full suite with leadership results in multiple categories.

We also expanded day zero support for the leading open models, including the latest Google Gemma 4 family, Qwen, Kimi, and others, enabling customers to deploy new models quickly with optimized performance. To build on this momentum, we have significantly accelerated our ROCm development cadence through increased software investments and agent-based coding workflows, enabling faster performance improvements and more rapid deployment of new capabilities.

Looking ahead, customer pull for Helios is very strong, driven by our leadership performance, memory bandwidth, and scale out capacity. Helios development is progressing well with strong execution across silicon software and systems as we advance through key milestones. We have begun sampling MI450 series GPUs to lead customers and remain on track to ramp Helios production shipments in the second half of the year.

As we approach production, demand for MI450 series GPUs continues to strengthen, with lead customer forecasts now exceeding our initial plans and a growing number of new customers engaging on large scale deployments, including additional multi-gigawatt opportunities. With this expanded visibility, we have strong and increasing confidence in our ability to deliver tens of billions of dollars in annual data center AI revenue in 2027, and to exceed our long-term growth target of greater than 80% in the coming years.

I look forward to sharing more on our next generation Instinct GPUs, EPYC processors, Helios rack scale platform, and our growing customer engagements at our advancing AI event in July.

Turning to client and gaming, segment revenue increased 23% year over year to \$3.6 billion. In client, revenue grew 26% year over year to \$2.9 billion, led by strong sales of our latest Ryzen processors and continued share gains across consumer and commercial markets.

In desktop, we strengthened our Ryzen lineup, including our latest X3D processors that deliver leadership performance across gaming, content creation, and professional workloads. We also introduced the Ryzen AI 400 series and Ryzen AI Pro 400 series desktop CPUs, sending our AI PC offerings across both consumer and commercial systems.

In mobile, we delivered strong growth driven by a richer product mix as Ryzen 400 mobile PC shipments ramped and commercial adoption increased. Commercial was a key highlight in the quarter, with sell through of Ryzen Pro PCs increasing more than 50% year over year as Dell, HP, and Lenovo broadened their AMD offerings. We also closed new enterprise wins across large technology, financial services, health care, and aerospace customers.

Looking ahead, we expect demand for Ryzen CPUs to remain solid in the second quarter. However, we are planning for second half PC shipments to be lower due to higher memory and component costs. Against this backdrop, we still expect our client revenue to grow year over year and outperform the market, driven by the strength of our portfolio and expanding commercial adoption.

In gaming, revenue increased 11% year over year to \$720 million. Semi-custom revenue declined year over year as expected at this stage of the console cycle, while engagements with customers on next generation platforms remain strong. In graphics, revenue increased year over year, led by demand for our latest generation Radeon 9000 series GPUs.

We also strengthened our Radeon portfolio with updates to our FSR software that improved performance and visual quality across a broad set of gaming workloads. Similar to the PC market, we believe that second half demand in gaming will be impacted by higher memory and component costs, and we are planning the business accordingly.

Turning to our embedded segment, revenue increased 6% year over year to \$873 million, driven by strength in test, measurement and emulation, aerospace and defense, and communications, as well as increased adoption of embedded x86 products.

Design win momentum grew by a double-digit percentage year over year, with billions of dollars in new wins across markets, reflecting the continued expansion of our embedded business from a primarily FPGA-focused portfolio to a broader set of adaptive embedded x86 and semi-custom solutions, significantly expanding our TAM. Our semi-custom engagements also expanded in the quarter, as data center, communications, and other embedded customers leverage our broad IP portfolio and high performance expertise to build differentiated solutions.

In summary, our first quarter results mark a clear step up in our growth trajectory with accelerating momentum across the business. Our client business continues to outperform the market, driven by rising adoption and share gains, while an embedded design win momentum and demand are strengthening across our expanded adaptive and x86 portfolio. At the same time, our data center business is inflecting with strong demand for both EPYC and Instinct products driving significant growth.

While we are still in the early stages of the AI infrastructure cycle, the pace and scale of deployments we are seeing today reinforce both the magnitude and durability of the opportunity ahead. As inferencing and AI deployments scale, they are fundamentally increasing compute requirements, driving both larger scale accelerator deployments and significantly more CPU compute.

AMD is uniquely positioned to lead in this next phase of AI with leadership products across high performance server CPUs and AI accelerators, and the ability to optimize them together as fully integrated rack scale solutions. We have a world-class supply chain and are making significant investments to expand capacity and execute at scale. With the momentum we are seeing across the business and the expanding market opportunity, we see a clear path to exceed our long-term financial targets, including delivering more than \$20 in EPS over the strategic time frame.

Now, I will turn the call over to Jean to provide additional color on our first quarter results. Jeanne.

**JEAN HU:**

Thank you, Lisa, and good afternoon, everyone. I'll start with a review of our first quarter financial results, and then provide our current outlook for the second quarter of fiscal 2026. We are pleased with our outstanding first quarter results, delivering accelerated revenue growth and the earnings expansion driven by strong execution and operating leverage.

First quarter revenue was \$10.3 billion, exceeding the high end of our guidance, growing 38% year over year, driven by strong growth in the data center and client and gaming segments, and the return to growth in the embedded segment. Revenue was flat sequentially, with continued growth in the data center segment, offset by seasonality in the client and gaming segment and the embedded segment.

Gross margin was 55%, up 170 basis points versus a year ago, driven by a favorable product mix, including a higher data center revenue contribution. Operating expenses were \$3.1 billion, an increase of 42% a year over year as we continue to invest in R&D to support our AI roadmap and the long-term growth opportunities and go to market activities.

As the business scales, operating income grew faster than top line revenue. Operating income was \$2.5 billion, representing a 25% operating margin. Taxes, interest, and other, resulting in a net expense of approximately 275 million. For the quarter, diluted earnings per share was \$1.37, up 43% year over year, underscoring the significant operating leverage in our model as we scale.

Now, turning to our reportable segment, starting with the data center segment. Revenue was a record \$5.8 billion, up 57% year over year, and 7% sequentially, driven by strong demand for EPYC processors and continued ramp of Instinct GPUs. Data center segment operating income was \$1.6 billion, or 28% of revenue, compared to \$932 million, or 25% a year ago. On the gaming segment, revenue was \$3.6 billion, up 23% year over year. On a sequential basis, revenue was down 9%, consistent with seasonality.

The client business revenue was \$2.9 billion, up 26% year over year, driven by strong demand for our latest Ryzen processors, favorable product mix, and continued share gains across consumer and commercial markets. Sequentially, client revenue was down 7% due to seasonality. The gaming business revenue was \$720 million, up 11% year over year, primarily driven by higher demand for Radeon GPUs, partially offset by lower semi-custom revenue. Sequentially, gaming revenue was down 15%, consistent with our expectations.

In addition, as Lisa mentioned earlier, we expect second half demand in gaming to be impacted by higher memory and component costs. We now expect second half gaming revenue to decline more than 20% compared to the first half. Client on gaming segment operating income was \$575 million, or 16% of revenue, compared to \$496 million, or 17% a year ago.

Embedded segment revenue was \$873 million, up 6% year over year as demand strengthened across several end markets. Sequentially, embedded revenue was seasonally down 8%. Embedded segment operating income was a \$338 million, or 39% of revenue, compared to \$328 million, or 40% a year ago.

Turning to the balance sheet and the cash flow, during the quarter, we generated \$3 billion in cash from continuing operations and a record \$2.6 billion in free cash flow, or 25% of revenue, demonstrating the cash generating power of our business model. Inventory was roughly flat at \$8 billion. At the end of the quarter, cash, cash equivalents, and short-term investments were \$12.3 billion. In the quarter, we repurchased 1.1 million shares and returned 221 million to shareholders. We ended the quarter with \$9.2 billion authorization remaining under our share repurchase program.

Now turning to our second quarter 2026 outlook, we expect revenue to be approximately \$11.2 billion plus or minus \$300 million. At the midpoint of our guidance, revenue is expected to be up 46% year over year, driven by a very strong growth in our data center segment, growth in our client and gaming segment, and double-digit growth in our embedded segment.

Sequentially, we expect revenue to be up approximately 9% driven by double-digit growth in both our data center and the embedded segments, and the modest growth in our client and gaming segment. In addition, we expect second quarter non-GAAP gross margin to be approximately 56%, non-GAAP operating expenses to be approximately \$3.3 billion, non-GAAP other income and expense to be a gain of approximately \$60 million, non-GAAP effective tax rate to be 13%, and the diluted share count is expected to be approximately 1.66 billion shares.

In closing, the first quarter of 2026 was an outstanding quarter for AMD, reflecting strong momentum across the business with accelerated revenue and earnings expansion. We are very well-positioned to build on the momentum as we scale our data center business, expand margins, drive continued earnings growth and long-term shareholder value creation. With that, I'll turn it back to Matt for the Q&A session.

**MATT RAMSAY:** Thank you, Jean. Operator, we're ready to start the Q&A session now. I would ask the callers to limit yourself to one question and one brief follow up, but please go ahead and pull for questions. 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 queue. You may press star 2 if you'd like to remove your question from the queue. For participants using speaker equipment, it may be necessary to pick up your handset before pressing the star keys. 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 Buchalter with TD Cowen. Please proceed with your question.

**JOSHUA BUCHALTER:** Hey, guys. Congrats on the results, and thanks for taking my question. Actually going to start with CPUs, which hasn't happened in a bit. It hasn't been that long since you announced the \$60 billion server CPU TAM for 2030 at the Analyst Day, and it's very quickly doubled. Agentic AI has obviously gotten a lot of attention in recent months, but would be helpful to hear your thoughts on how this TAM is inflecting and changing so meaningfully in such a short amount of time.

And maybe you could also speak to your confidence in hitting that greater than 50% share target from the Analyst Day, as your x86 competitor seems to be improving its supply and also there seems to be more momentum on the merchant and custom ARM CPU side. Thank you.

**LISA SU:**

Yeah, sure, Josh. Thanks for the question. So first of all, back to the when we think about CPU TAM, I mean, we've always said that CPUs are very critical part of data center infrastructure, and that's been where we've invested. And we saw the first signs of, let's call it AI demand, really pulling CPU demand last year. And that was the reason we updated the TAM to, let's call it the 18% category or approximately \$60 billion.

And what we've seen is all of the things that we believed in terms of agentic AI and inferencing, and all the CPU compute that is required is just happening, and it's happening at a much faster pace. So over the last, let's call it the last few months, as we've talked to our customers and we've seen how AI adoption is really unfolding, we're seeing a significant more CPU demand from really every major cloud provider as well as enterprise customers.

And the way that comes across is as AI adoption scales, you need more inferencing, as inferencing scales and you do more-- you have more agents and agentic AI, they all require CPUs for all of the orchestration and the data processing, and these other tasks. So with that, we've looked at it both bottoms up in terms of talking to customers and having them give us longer term forecasts, as well as just doing some clear workload analysis.

And yeah, I mean, it's a very exciting TAM. I think it's exciting to see your CPUs going greater than 35% to over \$120 billion. And then when you think about AMD in the context of that, I mean, CPUs are critical for so many tasks that you are seeing a lot more discussion about CPUs in the market, but we actually view it in three categories. There's general purpose compute, there's the head nodes that really support the AI accelerators, and then there are CPUs just for all of the agentic AI work.

And to do all of this, our belief is you need a broad portfolio of CPUs, and that's really what we have been focused on, is building, not just one type, but really a broader in terms of throughput optimized, power optimized, cost optimized, AI infrastructure optimized, as we've done in the Venice family. So when you put all that together, we're very excited about the larger TAM and we're also very happy with the traction that we're getting.

We're clearly feeling like we're seeing significant share gain. As we're going into our Turin portfolio that was ramped very nicely, Venice is extremely well-positioned and we're working with customers right now beyond Venice and what we're doing in those architectures. So we feel really good about the market as well as our opportunity to grow to a greater than 50% share of that market.

**JOSHUA**

**BUCHALTER:**

OK. Thank you for all the color there. I wanted to ask about the Instinct side. So in the press release, you mentioned that MI450 and Helios engagements are strengthening with customer forecasts, exceeding the expectations and the pipeline growing. You certainly have the big public OpenAI and Meta deals. Was this comment referring to those engagements upsizing versus the announced initial deployments, or was it other customers and maybe the increase on the MI450 timeline, or is it MI500 and beyond? Thank you.

**LISA SU:**

Sure, Josh. So we are very excited about MI450 and Helios. We're seeing significant customer interest in those products as well. So we have certainly talked about our large partnerships with OpenAI and Meta, and those are going really well. We appreciate the deep co-engineering that has gone on there. When we look at the totality of, let's call it based on our current visibility, how those forecasts are coming in, with all of our customers, we're actually seeing it above our initial plans that we had planned for 2027.

And I think the encouraging thing is we're seeing a breadth of customers who are now very interested in deploying at significant scale MI450 series. And those are for both training and inference workloads, although the largest deployments are for inference. And based on all of that and the scale of new customer interest, we see a path to really get to exceed our original targets of greater than 80% CAGR.

And these are really, 2027 time frame. Obviously, when we talk to customers, we're talking to them about MI355. There's a lot of good traction we're seeing there. MI450 and Helios, I think for significant large scale deployments. And then many customers are also very engaged with us on the MI500 series and all of the opportunities there. So we feel like, very, very good progress, and the key is that we're continuing to broaden and widen the scope of both customers as well as workloads.

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

**THOMAS O'MALLEY:** Hey, guys. Thanks for taking my question. Lisa, if I get your numbers correct here in the March quarter, it sounds like the server processor side of the CPU side grew over 50%. If you take it just at the word, it looks like maybe the data center GPU side actually grew in Q1. So I was curious around the cadence of this year. Previously you had talked about really a back half weighted and then more so Q4 weighted year. Could you talk about if that's changed at all?

And then the second part of the question is, as you go into 2027, clearly you're pointing out a lot of upside from the larger customers and then the ecosystem around them with new customers as well. But when you look at supply that's a major issue in the ecosystem today, could you talk about where you're concerned on supply, if you are, and then any gating factors as you look into next year, whether that be power, data center, buildouts, et cetera, or do you feel really good about the ability to grow? Thank you very much.

**LISA SU:** Yeah. OK. A lot of pieces of that question, Tom, so let me try to get through it. So first of all, on the data center segment in Q1, the server business was greater than 50% year over year, as we said in the prepared remarks. The data center AI was actually down modestly because of the China transition. We had more China revenue-- I'm sorry, sequentially, more China revenue in Q4, and it was less in Q1.

But as we go forward, I think we see strong growth in both segments. So we guided data center Q2 up sequentially double digits, and that's double digits in both server as well as data center AI and progression as we go forward. So first on the server CPU side, we talked about growing to over 70% year over year in Q2 and that continuing into the second half of the year. And on the data center AI side, we will be ramping Helios in the second half of the year, so let's call it starting with initial volume in Q3, with a significant ramp in Q4 and then continuing to ramp in Q1.

So that's a little bit of progression. And then to your questions about customers and supply, I think I answered, Josh, the customer question. I think we have very good visibility now into the deployments that are on track for 2027. And when I say good visibility, it's visibility down to which data centers are the GPUs going to be installed in. So that's necessary just given all of the constraints out there.

We feel that there is tightness in the supply chain. There are certainly tightness in data center build outs. But we are confident in our ability to supply to the levels of growth that we're talking about and to exceed the levels of growth that we're talking about. And we're also working very closely with our customers and our partners to ensure that we have good visibility to data center power. And there is much more power that's coming online in 2027.

And so with all those things in mind, I think, again, lots of things to manage. It's a complex ramp, but we're very, very pleased with the progress on the ramp.

**MATT RAMSAY:** All right, Tom, I think you shotgunned approached the multiple questions there. So operator, maybe we can go on to the next caller, please. Thank you.

**JOHN:** Thank you. 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 couple questions. The first one is just on the EPYC competition. Lisa, you went through some of the statistics of you versus x86 and you versus ARM, but I wanted to dive a little bit deeper into that. How do you see AMD truly differentiating, especially when you're signing-- well, you see some of your competition signing up the same customers from the ARM side and the x86 competition having more supply. So I just wanted to see if you could dig a little bit deeper into how you think the market share is going to trend over time.

**LISA SU:** Sure, Ross. So look, we're very engaged with every major hyperscaler in terms of understanding their needs on the CPU side. I think we have very much wanted to, let's call it, optimize our CPU roadmap for the various workloads. I think we were early to call this AI component of CPUs. And so we've been actually optimizing very closely with those customers.

The way to think about this, Ross, is that you're going to need a broad portfolio of CPUs. Not all CPUs are the same. Frankly, you're going to need different CPUs for whether you're talking about general purpose operations or you're talking about head nodes, or you're talking about agentic AI tasks, they're going to be optimized differently. And we've thought through that and we are absolutely optimizing across the various workloads.

So from a competitive standpoint, we feel very good about where things are. And from deep relationship with the customer set, I think we feel very good about that. So from our current standpoint, I think the depth of our roadmap just expands as we go forward. And you shouldn't think about it as people are going to do one or the other. I think you're going to see people actually use x86 and ARM for many of the large hyperscalers.

And even for those who are developing their own, they're still buying lots of CPUs in the merchant market for the reason that I just stated, which is you need different CPUs for the different types of workloads, and there's very high demand at the moment.

**ROSS SEYMORE:** Thanks for that. And I guess for my follow up, maybe more for Jean, on the gross margin side of things, it's nice to see the gross margin popping up in the second quarter, guy. But I just wanted to get some trends longer term, maybe not specific numbers. But how should we think about when Helios and the Instinct side really ramps in the fourth quarter and more so next year? I could see some offsets with that carrying a below corporate average gross margin. But then everything that Lisa talked about with the EPYC side of things being significantly stronger might be more of an offset than it was in the past. So just walk us through the puts and takes of that and maybe directionally where you think gross margin goes over the next year or two.

**JEAN HU:** Yeah, Ross, thanks for the question. We are very pleased with how our gross margin is trending. It came in really strong in Q1. And also, as you mentioned, we guided the Q2 higher at the 56%. I think as we think about the second half or quarter over quarter, as you know, there are some puts and takes. I would just say from a tailwind perspective, we actually have a multiple tailwinds really are going to help our gross margin.

First is the server CPU. Lisa talked about the server CPU to expect it to grow more than 70% in Q2 and continue to be really strong in the second half. That really helps our gross margin. Secondly, in the second half for gaming actually is going to come down and our client business actually continue to go up the stack. So from client gaming segment gross margin actually is going to be also very helpful.

Embedded actually is very accretive to our gross margin. It's momentum actually is continuing in the second half. So we're really pleased with all the tailwinds we have. On the other side, MI450 will start to ramp in Q3 and the ramp significantly in Q4. That is below corporate average. So that will have different puts and takes in Q4 in the gross margin side. But when we sit here, when we look at all the positive trends we have to really offset some of the gross margin dilution from MI450 side, we actually feel really good about the setup of the gross margin for 2026.

And into next year, I think some of the tailwinds I talk about that will actually continue. That's why we feel confident about continue to drive the gross margin. We actually during our Financial Analyst Day, we outlined the long-term gross margin in the range of 55% to 58%. We think for the first year, we're making good progress there.

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

**TIMOTHY ARCURI:** Thanks a lot. I wanted to ask about units versus ASP for server CPU. If I look at the June guidance, it implies up 25% to 30% for server CPU. And Lisa, you had mentioned second half of the year. It implies that server CPU could grow like 70%, maybe a little more this year. And so I guess my question is, how much of that growth, either in June or for the year, is like units versus pricing? Are these price increases mostly captured in June or is that also helping you in the back half of the year?

**LISA SU:** Yeah. Tim, the way I would say it is, maybe let me bring you back to Q1 for a moment. So if you look at our significant growth in the server business, it was actually-- although we were up on a year over year basis for both ASPs and units, it was actually much more unit-driven. So we are shipping more CPUs across, not just the high end Turin family, but we're actually shipping a lot of Genoa, the Zen 4 family as well.

As we go forward for Q2 and into the second half, we are guiding for a significant amount of growth. I think there's a little bit of ASP in there, but the way we're thinking about pricing, to be fair, is we are in a range where the supply chain is tight. And so there are some inflationary pressures, costs have gone up a bit, and we are sharing some of that with our customers. But we are also being very thoughtful in-- this is--

We're paying out for the long-term and that means that we are-- our goal is to ship more units and a lot more units. And so from that standpoint, you should imagine that the majority of the growth is unit-driven and the ASPs are just really to help cover some of the inflationary pressures.

**JEAN HU:** And just to add to what Lisa said, our ASP is increasing because of the mix where actually each new generation, the call counts, those are increasing. That actually drives the ASP up.

**TIMOTHY ARCURI:** Thanks a lot for that. And then I guess Lisa also, so there's a lot of new architectures that are being used from multi-tenancy all the way to low latency, and your competitor has talked about the low latency part of the market being 20% plus and they, of course, added to their portfolio there. Can you talk about how you see that part of the market? I mean, obviously you have enough business now you don't need to worry about that, probably for now. But can you talk about that? Thanks.

**LISA SU:** Yeah, sure. So look, I think what we're seeing is what we expected in the sense that as you go-- as the AI adoption continues and the volumes continue to go up and the overall market goes up, you are going to see, let's call it, different compute architectures being used because you want to get more cost optimization from that. So we expect that even in that situation, obviously the vast majority of the TAM is still going to be, let's call it, data center GPUs as the primary accelerator.

But you may choose to do optimization around inference, around low latency, around certain parts of the stack, whether it's decode versus prefill. I think that's very natural. The way we look at it is we're developing a full compute portfolio so that CPUs, that's GPUs, that's the ability to connect to all accelerators, as well as the ability to do customization for certain customers, and we've also talked about our semi-custom capabilities.

And with all of those compute capabilities in our tool chest, I think we will be able to address very effectively a large portion of this market, including the low latency portion of the market. So from our standpoint, this is a natural evolution. Now, how fast it goes depends a bit on the technology in terms of what share of the TAM these things become. But we should expect that there will be different variants, and we're well prepared to address those different variants.

**JOHN:** Thank you. And the next question comes from the line of Vivek Arya with Bank of America. Please proceed with your question.

**VIVEK ARYA:** Thanks for taking my question. Lisa, do you think agentic CPU growth is incremental or is it coming at the expense of GPUs conceptually? So if you are raising server CPU TAM, are you also implicitly raising AI TAM? So I'm interested in your perspective on what did you think a server CPU was as a percentage of AI TAM before, and what is it now with this 120 billion number?

**LISA SU:** Sure, Vivek. So the way we're thinking about it is it's largely additive to the TAM. So you should think about we need all of the accelerators to run these foundational models. And then as these agents do work, they spawn more CPU tasks. So I would say largely incremental. The key is to make sure-- what we're seeing is in these deployments, the key is to make sure the ratio of CPUs to GPUs are the right ratio.

So if you're installing a gigawatt of compute, the ratio or the percentage of CPU as part of that gigawatt will increase. Some of the conversation in the industry has been about CPU to GPU ratios. And it's very hard to call exactly, but we certainly see the movement towards where in the past the CPU to GPU ratio was primarily just as a host node in a 1 to 4 or 1 to 8 configuration. Now changing and getting closer to a 1 to 1 configuration. Or you can even imagine if you get lots and lots of agents that you could have more CPUs and GPUs.

But all in all, to answer your question, I think it's largely additive to the TAM and the key is that everyone is now planning and thinking about CPUs at the same time that they're thinking about, their accelerator deployments, which is a good thing.

**VIVEK ARYA:** And for my follow-up, Lisa, we continue to see memory prices go up. I imagine that is both a cost inflation for you, but perhaps an opportunity to take price as well. I'm curious, how is that dynamic playing out for AMD and especially for your customers? Because a greater part of their CapEx increase is really this memory inflation tax that they have to pay. So how is this dynamic playing out for you and for your customers?

And the part that I'm really interested in is that have you secured enough supply versus your other larger competitor who has disclosed a lot of prepayments and other things? So just how is this memory inflation dynamic playing out, and are you adequately supplied from that perspective?

**LISA SU:** Sure. So Vivek, let me answer the second one first. I think from a supply standpoint, we are very happy with our partnerships with the memory vendors, and we have secured enough supply to certainly meet and exceed our targets. So it is a tight memory environment, let me be clear, but I think we are very deep partnerships with the memory providers.

And then back to your comments on the inflationary pressures, I mean, look, this is something that everyone in the industry is working with. In the time of tight supply, we are seeing some cost increases on the memory side. I think we are all working through that. The way we're seeing it unfold in the market is actually on the data center side, because of the, let's call it, the demand for AI compute, I mean, people are largely focused on supply and ensuring that the supply assurance is there.

The corollary of that, the larger impact that we're watching is the impact on the consumer markets. And as we said in the prepared remarks, we are expecting that there could be some demand-- the demand impact as a result of the memory price increases on things like the PC business in the second half of the year as well as the gaming business. So we're taking that into account in our overall model.

And we continue to work closely with the memory providers as well as our customers to ensure that every time we ship a CPU or GPU, that it's paired with the memory on the other side so that we don't have compute that is not being deployed.

**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, and congrats on the results. I want to stick on the topic of CPU to GPU. And as we think about the chart that you had outlined at the Analyst Day, there was obviously broken out between traditional CPUs and then the AI bucket on top of that. Obviously, I think the new forecast has a lot to do with the AI CPU expansion. I'm just curious, when you're doing a CPU in an AI workload, is there structurally a different level of ASP tied to that CPU optimized for AI relative to a general purpose service CPU? Any kind of color or help on that would be useful.

**LISA SU:** Sure, Aaron. So let me start with the broader question, the broader question regarding the way we think about the CPU TAM is, again, think about it as three categories. So there is a traditional CPUs, let's call it, general purpose CPU TAM that is increasing, but let's call it increasing at a low rate. Maybe let's call it low double-digits. Then you have your AI head node, which is connecting to accelerators, which is also growing, but it's smaller.

And then the largest piece of the growth is this agentic AI piece, which we think is really stemming from all of the agentic processes. I don't have a number that I can tell you in terms of relative ASPs, because it really depends on the workload that is being run. And what we see going forward is as core counts increase, obviously, we will see ASP increase, and that's the direction that we're going in as we go forward.

But the main point is the largest portion of this is the agentic AI-- the CPUs that are serving these agentic AI workloads in terms of the TAM increase.

**AARON RAKERS:** Yep. And as a quick follow-up, I'm curious, how do you characterize the competitive landscape as we see some of the ARM introductions in the market? Just curious of your views on the competitive landscape and server CPUs. Thank you.

**LISA SU:** Yeah, Aaron, the best way to think about the server CPU landscape is, again, number one, everyone is talking about CPUs, so that tells you how critical they are for the AI infrastructure. And I think that's a good thing. We feel like we're very well-positioned. No question, ARM is a good architecture. It has a place in the data center market. We view it as more point products relative to a portfolio where from an AMD standpoint, we've built this broad portfolio of CPUs going forward, which you're going to need for all of these different workloads.

And we have, in the Venice time frame, added an AI-optimized CPU with Verano in addition to our throughput optimized and cost optimized points. So from that standpoint, I think we're very competitive. We're continuing to innovate on architecture. We're continuing to innovate on both advanced packaging as well as all of the architectural pieces. So we feel very well-positioned going forward. And the key is that TAM is much, much larger than anybody thought, and so there's a lot of opportunity for different products to be successful in this area.

**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 guess first question was hoping to speak a bit more about client for all of calendar 26. You talk about growth, expected growth, but would love to hear your thoughts around seasonality in the second half. And I'm assuming that you are repurposing certain logic tiles from client over to the data center, and would love to better understand what the implications are for ASPs on the client side looking into the second half.

**LISA SU:** Sure. So C.J., I think the client business has performed really well for us. I think if we look at Q1, it actually was a little bit stronger than what we expected. We are seeing some mixed shift in the client business. The mix that we're seeing is the MNC or the notebook business is actually growing, especially the premium portion. We're making very good progress in the commercial PC arena with our AI PCs.

We did see desktops a little bit softer, just given desktop is a more consumer-focused market. And so in that market, it's more impacted by some of the memory pricing and the component price increases. When we look at the full year, our commentary is we are planning for some demand impact in the second half due to the memory pricing. But even in that environment, what we're focused on is ensuring that we continue to make good progress on the commercial business and continuing to focus on the premium segments of the market.

So we believe that we will continue to grow on a year over year basis for the client business compared to last year. And as it relates to ASPs, again, it's a little bit of puts and takes between notebook and desktop, but overall, I think we're feeling good about our opportunity to outperform the market in client going forward. Did I answer?

**C.J. MUSE:** That was perfect. Thank you. And then I guess a question on Instinct gross margins with compute essentially sold out. And obviously, you're building a business, so one has to be, I guess, conservative on that front. But I would think outside of passing through HBM, that given the very tight wafer environment, that this would be a place where could look to drive your Instinct margins closer to your corporate average. How are you thinking about that, either today or in the coming one, two, three years?

**JEAN HU:** C.J., at this stage, we really focus on drive the top line revenue growth on our Instinct family of products. I think on the gross margin side, you're absolutely right. It's really tied to the demand for compute is tremendous. We actually are very strategic how we think about how we work with the customers. And of course, the different customers also have different gross margin. I think over time, once we start to ramp our revenue, we'll have a lot of opportunities to improve gross margin, both on the ASP side, but also more importantly on the cost side when we scale our business.

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

**STACY RASGON:** Hi, guys. Thanks for taking my questions. For the first one, I just wanted to make sure I have the near-term AI GPU trajectory correct. So I know you said it was down sequentially in Q1 because of China. You had like 390 million of China revenue in Q4. Do the AI business in Q1 actually grow sequentially ex-China? Because it doesn't feel like it given the server outlook.

And then I look at what's maybe suggested for Q2. Are you thinking GPUs and servers grow similar rate sequentially? Because it would probably put GPUs in Q2 below the overall level that you were at Q4, which seems low to me. I'm just trying to tie all that out. Could you help me with that, please?

**JEAN HU:** Yeah. So I think, Stacy, appreciate the question. I think if you look at the Q1, we didn't mention data center AI was down modestly sequentially, primarily due to lower China revenue in the quarter. I think on your second question regarding Q2, you're right. Both data center AI and the server will grow double-digit in Q2.

**STACY RASGON:** Yeah, but you didn't answer my question. In Q1, did it grow sequentially ex the China step down I guess is what I'm asking.

**JEAN HU:** The China for our business in Q1, it's not material. So I think I will repeat what I just said, is that, yeah, the China revenue in Q1 is not material.

**STACY RASGON:** OK. OK. So you don't want to answer-- you don't want to-- OK. Second question. OpEx for spending, but it continues to blow past the targets. You give an OpEx guide and then it blows through it and then you guide higher. So again, I'm not bothered by the spending. I'm just wondering why is the OpEx been so hard to forecast, and how should we think about OpEx through the rest of the year given the revenue growth?

**JEAN HU:** Yeah. Thanks, Stacy, for that question. I think the most important thing is given the tremendous market opportunities we have, we actually are investing aggressively. If you look at the past several quarters, we really leaning in investing, but all the AI investments are driving the revenue momentum. So if you look at the Q1, revenue was 38% up, but then Q2, it was 4-- we guided 46% up.

The investments are driving the revenue momentum. Some of the OpEx increase, of course, is tied to the revenue. When you look at our beat on the revenue side versus our guidance, we did beat on the revenue side. So that impacts a little bit. But also at the same time, we have a lot of customer engagement with our data center AI business. We do continue to make sure we have the resources to support our different customers.

**MATT RAMSAY:** Thank you very much. Operator, I think we have time for one more caller on the call. Thank you very-- thank you.

**JOHN:** Thank you. Our final question comes from the line of Blayne Curtis with Jefferies. Please proceed with your question.

**BLAYNE CURTIS:** Hey. Well, thanks for squeezing me in. Lisa, I just want to go back to the supply side. There was a lot of story about your competitor restarting 7 nanometer. I'm just curious, as you look at that landscape, which is quite robust through the end of a decade, do you think that the older products will stay around longer? And is there a way to think about the implications for gross margin in such a strong market? Is that actually a negative?

**LISA SU:** Actually, Blayne, I don't think we see the older products hanging around longer in our case. I think it might be company-specific stuff. In our case, we actually see-- first of all, Turin is very strong. We actually crossed over 50% of our revenue being Turin in this quarter. Genoa is very strong. We're still shipping some Milan, but I would say that's come down over time.

So in general, people want to use the more-- the newer products because they're just more efficient in every aspect, from performance, from cost structure, from power standpoint. So that's what we're seeing. By the way, I should also mention, in addition to what we're seeing in the cloud segment of server, we're seeing really nice strong pickup in enterprise. And there as well, we're seeing our newer products do very well.

So from our standpoint, it is all about ensuring that we ship what the customer needs. And in this case, it typically is our newer products, and we expect that to continue. As we transition into Venice later this year, we will expect Turin and Genoa to continue shipping, but there's a lot of goodness in going to the new products.

And on the supply chain side, I know there's been a lot of discussion about how tight the supply chain is. The supply chain is tight, I would definitely say that, but I also think, this is an area where we excel. We have very deep relationships across the supply chain, on the wafer side, on the back end capacity side, and we are seeing meaningful improvements in that. And as our customers come to us with more demand, we are getting more supply.

And the good thing about this is we're now talking about 27 CPU demand. We're talking about 28 CPU demand. And so that allows us to just plan much better as we go forward.

**BLAYNE CURTIS:** And then just a quick one for Jean, I'm just curious to follow-up on Stacy's question on OpEx. I guess I was a little surprised that SG&A is outpacing R&D. I was just curious is that startup costs? I mean, because in a strong market, you wouldn't think you would have to discount or have a big sales effort. So I'm just curious for the year how you think about R&D growth versus SG&A.

**JEAN HU:** I think for the year, you should expect us to grow R&D much faster than SG&A. I think in the past few quarters, we have been really building our go-to-market machine, and we have been investing more in sales and marketing side. But going forward, you should expect the year over year growth. R&D will grow faster than SG&A growth.

**LISA SU:** And if I just add to that, Blayne, the places that we invest-- Jean is absolutely right. We're investing in R&D ahead of sales and marketing. But the places that we're investing in sales and marketing are paying off. So the investments are going into enterprise servers. They're going into commercial PCs. They're going into mid-market, small and medium business. These are places where AMD traditionally didn't invest.

But now that we have a much broader portfolio, both on the server CPU and on the commercial PC side, it makes sense for us to invest because that's the very best part of those markets.

**MATT RAMSAY:** All right. Thank you very much, everybody for joining and your interest in AMD. John, you can go ahead and close the call now. Thanks.

**JOHN:** Thank you. 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. Please disconnect your lines, and have a wonderful day.