TIM ARCURI: Good morning. We're going to get started here. I'm Tim Arcuri. I'm the semi and semi equipment analyst here at UBS, and we are very honored to have Dr. Lisa Su with us from AMD. So good morning, Lisa.

LISA SU: Good morning. Thanks for having me.

TIM ARCURI: Great Thank you. So first, I just wanted to start by talking about the transformation that you led. I think, beginning 3 or 4 years ago, you've transformed the company from being less than 20% data center to nearly 50% this year. What have been the drivers for this transformation? Some of it has been market growth, but some of it was a decision that you made years ago to pivot the company in this direction.

LISA SU: Well, again, Tim, thanks for having me. It's great to be here with everyone. And I think, in the technology sector, it's all about making the right big bets when you look at what the inflection points. I think, over the last, let's call it 5 plus years, we've been incredibly focused on high-performance computing as a sector, knowing that as we go forward, compute would be such an important part of unlocking capability and intelligence. And then a few years ago, it became absolutely clear that this was going to be all about AI, that AI was the ultimate application of high-performance computing. And with that, the investment cycles would be there.

I think this is way before ChatGPT and large language models, but it was the idea that we could really use computing to do so much more in terms of unlocking productivity and intelligence going forward. So yes, we've pivoted, really, our R&D capabilities, both hardware, software, system integration to a significant focus on high-performance computing and AI. I think it's paid off well. Our data center business has grown very nicely, well ahead of the market, over 50% year for the last few years.

And what we see going forward is even more exciting because I think, the recognition is that computing is such an important part of the ecosystem today, that we see a very large market opportunity, as well as significant growth of our business, actually, accelerating growth from, let's call it, 50% plus over the last few years to over 60% plus as we go forward. So no question, data center is the place to be.

TIM ARCURI: Yeah. I actually wanted to ask you about that. So you had this Analyst Day recently. You gave us a new trillion dollar data center TAM by 2030. You were saying 500 billion by 2028 before, so you've upped that. But you also importantly said that you can get double-digit share that pie. You're doing $16 billion in data center this year, would put you on a 60% CAGR, as you said, that's up from the 50% CAGR over the past five years. How are you winning? And what's the crux of your competitive advantage in data center.

LISA SU: Well, I mean, when you look at what's important in the data center market, I think, the key piece is you really have to have a holistic view of the market. It is CPUs. It is GPUs. It is FPGAs. It's the possibility of doing ASICs. It's being able to integrate all of that together. And that's our unique capability. I think we are really the only semiconductor company out there that has all of this foundational IP, and we have invested way ahead of the curve in terms of some of the key enabling technologies.

We were the first to implement chiplets in high-volume production. We're now on our fifth generation of chiplets. And the reason I view these as key foundational technologies is because the one thing we know is that the workloads are going to change. There's nothing static about the computing market. There's nothing static about AI. What we see is that there's an incredible pace of innovation out there where there are new workloads, there are new models, and there are new use cases, and so you really need this entire portfolio of technology capability, which is what we have.

So we've built a incredibly strong franchise with our EPYC data center server CPU chips. We're now over 40% revenue share in that market and growing. We have a very, very strong GPU accelerator roadmap. And yes, we view that as a significant growth opportunity, the largest piece of the TAM. I think, Tim, you might remember when we originally said that the TAM was 3 or 400 million, people thought, wow, Lisa, that's really big. And now, I would say that, I think we're all believers, and the TAM is very, very large, because we're still in the early stages of this, and our differentiation is going to be offering the right solution for the right workload going forward.

TIM ARCURI: So there's been some recent news in the market that have made people think that ASICs are going to take over the cellular market, and I just wanted to get your opinion on that and the general competitive landscape in the AI world. Is our ASIC really a threat to GPUs? You've said that ASICs are going to be 20%, 25% share of the market, has anything we've heard recently changed your view on that?

LISA SU: Yeah, I actually don't think so. I think what we have said about the market is what I started with, which is the market wants the right technology for the right workload, and that is a combination of CPUs, GPUs, ASICs, and other devices. As we look at how these workloads evolve, we do see some cases where ASICs can be very valuable. I have to say that Google has done a great job with the TPU architecture over the last number of years, but it is a, let's call it, a more purpose built architecture.

It's not built with the same programmability, the same model flexibility, the same capabilities to do training and inference, where our GPUs have the beauty that they are highly parallel architecture, but they're also highly programmable, and so they really allow you to innovate at an extremely fast pace. So when we look at the market, we've said that we see a place for all of these accelerators, but our view is as we go forward, especially over the next-- let's call it, 5 years or so that we'll see GPUs still be the significant majority of the market, because we are still so early in the cycle, and because software developers actually want the flexibility to innovate on different algorithms.

And with that, you're not going to a priority, what to put in your ASIC. So I think that's different. So 20% to 25% feels like the right number. I think the other thing that people should recognize is that this is absolutely a huge and growing market, and as a result, you're going to see a lot of innovation on the silicon side, as well as on the software side. And in general, I view that as a great thing, because that allows differentiation in the market.

TIM ARCURI: And if a customer came to you and wanted you to build an ASIC for them, is that something that you would do?

LISA SU: Well, the way we look at these things, Tim, it's all about, what is our secret sauce? What is our differentiation? And from our perspective, the differentiation really comes when we can take our intellectual property together with our customers intellectual property and create a case where 1 plus 1 is greater than 3. I think we are extremely good at deeply partnering with customers, and we've done that over the last 10 plus years.

We do have-- in addition to all of our standard products with CPUs and GPUs and FPGAs, we've also created a semi-custom business. I don't call that an ASIC business. And the differentiation being ASICs, you're going to do, let's call it, any chip that somebody comes and asks you to do. That's not necessarily where we shine. I think where we shine is when we can put our IP together with our customer's IP. And we have done a number of semi-custom designs that build off of our foundational capability so that customers can differentiate.

So I think our overall value proposition is-- our goal is to take all of our R&D investments, and we now have 25,000 engineers that are integrating at the bleeding edge of technology, hardware, software, system design, and really marry it with our largest customers who want to find that differentiation and work on how do we see that in the portfolio? And that could be custom system design. So we do do putting the pieces together, that could be a special SKU. We have lots of special SKUs that are optimized to given workloads, and that could be a special silicon as well. And we've done that in a number of cases across a number of markets over the last couple of years.

TIM ARCURI: Great. So I wanted to go on to another debate that's in the marketplace, and that's whether there's a bubble right now in AI. You weren't going to get away without me asking you this.

LISA SU: Well, it wasn't the first question so--

TIM ARCURI: [CHUCKLES] So can you just talk about that? I know NVIDIA went at that pretty hard on their call, so I just wanted to give you a chance to address that.

LISA SU: Yeah, absolutely. So it's curious this conversation about a bubble from my standpoint. I mean, I spend most of my time talking to the largest customers, the largest AI users out there, and there's not a concept of a bubble. What there is a concept of is-- we are let's call it 2 years into a 10-year super cycle. And that super cycle is computing allows you to unlock more and more levels of capability, more and more levels of intelligence. And that started with training being the primary use case, but that's really very quickly migrated to inference.

And now we're seeing, with all of the models out there, there is no one killer model. There's actually a number of different models that are-- let's call it, some are better in certain aspects, some are better in other aspects. Some people want to do, let's call it, fine-tuning, reinforcement learning. So with all of this capability out there, the one thing that is constant as we talk to customers is we need more compute. That at this point, if there was more compute installed, more compute capability, we would get to the answer faster.

And so yes, there is significant investment. I mean, I think all of the CapEx forecasts that have increased over the last 3 to 6 months have certainly shown that there is confidence that those investments are going to lead to better capabilities going forward. And so yeah, from a standpoint, do we see a bubble? We don't see a bubble. What we do see is very well-capitalized companies, companies that have significant resources using those resources at this point in time, because it's such a special point in time in terms of AI learning and AI capabilities.

TIM ARCURI: And I guess, just on that end, so there's a lot of talk that there's not an ROI for these CapEx dollars. People say that they're short on compute, but when you look at AI and the actual use cases, can you speak to that?

LISA SU: Yeah, absolutely. I think, again, what my view of this is the cause and effect usually takes a little bit more time than people are expecting. But what we're seeing, and I can just tell you our own case at AMD over the last 15 to 18 months. What started as, let's call it, let's try AI for our internal use cases, has now turned into significant clear productivity wins going forward. So there's no question that there is a return on investment for investment in AI.

What is the return on investment for enterprises? It is more productivity. It's building better products. It's being able to actually serve your customers in a way that is more intuitive than you have today. And if you look at today's AI, as much progress as we've made over the last couple of years, we're still not at the point where we're fully exploiting the potential of AI. So we're seeing actually a lot more effort over the last 3 to 6 months on the use of agents, and how we make sure that AI not only suggests answers in a copilot fashion, but actually gets to a place where it can actually do lots of productive work.

And that is flowing through-- we're seeing that across multiple customers. We're seeing that across the largest hyperscale customers. We're seeing that across the large enterprises that are using AI. And I still say that we are in the very, very early innings of seeing that payoff. So as we talk to the largest enterprise customers, I think, every conversation is, Lisa, how can you help us? How can we learn faster so that we can take advantage of the technology?

So I think the return on investment certainly will be there. I think the debate is perhaps more around the largest foundational model companies and whether there's return on investment there. But again, my view is that there's not going to be one best something, or there going to be multiple models that are best optimized for use cases. And the secret sauce is really in how you integrate it so that customers can take advantage of the technology as smoothly and as easily as possible.

TIM ARCURI: Thank you. So another point is that you're moving from being a silicon company to being a systems company. And a big piece of that was your acquisition of ZT, and your partnership now with Sanmina. So can you actually speak to that. And you're a bit of a fast follower in building these racks and these systems, so do you think that you've learned from some of the growing pains that your peer had?

LISA SU: Well, I think if you take a step back and come to why are we doing this integration? The reason we're doing this integration is the time to useful capability, the time that it takes for our customers to bring up this really complex infrastructure is super critical to make as fast as possible. So the full stack solution is a way for us to help customers get to, let's call it, productive compute capacity. And we're very happy with our acquisition of ZT. I think it's one of the smoothest acquisitions integrations that I've seen.

And what we've been able to do is really take, let's call it, best-in-class system design and combine it with our best-in-class hardware and software capability to come up with very strong full stack solution. We're super excited about MI450 series and the Helios rack that will come to market in 2026. I do think we've learned. I think we've learned as an industry. We're always going to learn that putting together these complex, rack-level systems is hard. There's nothing easy about it, but there are certainly ways that you can derisk and ensure that you can go as fast as possible.

I think the key elements for us in our strategy, when we think about rack-level solutions is as important as it is to have that reference design capability, it's also really important to have an open ecosystem. And that open ecosystem means that we have an open rack architecture, which, together, we've developed with Meta, which I think has taken a lot of the best practices out there in the industry.

We're working with all of the key suppliers within the rack to ensure that, again, we learn how to bring these up as fast as possible. And then frankly, the ZT team has brought 1,000 plus really, skilled engineers to the capability. So I think we feel really good about our rack-level solutions. I think the feedback that we've been getting on the Helios rack has been fantastic. I think people see that we've made really, smart engineering decisions to ensure that we're able to bring these systems up as smoothly as possible.

TIM ARCURI: Great Thank you. One thing I also hear is that you're fighting a battle on multiple fronts. You're fighting Intel and AMD and PC. You're fighting NVIDIA, and you're fighting ASICs, and you're not that large of a company yet. So when you think about prioritizing development, do you feel like you're having to disinvest in certain areas and invest in others?

LISA SU: Well, actually, I think you're actually pointing out one of our strengths. So I think one of our strengths is the fact that we have a really, really capable and efficient R&D engine. I give Mark Papermaster and the team a lot of credit for that. We've built an execution engine. We've done 5 generations of server CPUs right on time, at performance, best in class, and the way we develop is we actually develop foundational capabilities that bring all of these computing elements together. So CPUs, GPUs, FPGAs. I actually think this is one of our strengths.

We're not religious about the world is going to be taken over by X, because I can tell you for sure, I do not believe the world is going to be taken over by X. I think you're going to need the right compute for the right workload, and that is our strength, and I think we've developed a R&D engine that knows how to execute that. Now, there's no question that AI sits above all of this. And so all of the innovation that we're doing in AI, all of the software investments that we're making in AI are there to ensure that it works across the entire portfolio.

TIM ARCURI: Great. Can we talk about the deal with OpenAI? You offered them 10% of the company with warrants. There's various strike prices at each tranche. How did the deal come together, and how does it change your engagement with the other customers?

LISA SU: Well, first of all, we're very pleased, excited, happy with the OpenAI deal and partnership to give you some idea of how it came together. It really came together over the last couple of years. We've always been working with them as one of the leading foundational model companies to understand where do they think model evolution was going, because that's so critical in determining our long-term roadmap.

When we were looking at what should the MI400 series look like? What would really make it special? How do we differentiate long term? Clearly, one of our key strengths has been our memory architecture that's enabled by chiplets and all that. And a lot of that came from talking to our largest customers, OpenAI being one, but a number of our other large partners, Microsoft, Meta, Oracle, et cetera, also contributed to those thoughts.

And when we thought about, where do you want to go going forward? This is all about going big and not necessarily the typical way that technology evolves is sometimes, hey, we do smaller partnerships here and there. In AI, it's all about really bringing together hardware, software, co-optimization and co-design, and that's what we've really put together with this OpenAI partnership. I think, we view it as a way to ensure that we are highly developing with one of the largest model companies in the world.

The key here is that-- with the current structure of our 6-gigawatt partnership, it's a win-win on both sides. So on one hand, we get significant scale with this, if you think about it as each gigawatt is deployed that's significant scale to AMD. That's double-digit billion of revenue. And it's also an opportunity for OpenAI to be very invested in our technology success as well, because there are a number of commercial, as well as technology milestones.

Very much a win-win, very highly accretive to our portfolio, and as it relates to other customers, I think the idea of having a very optimized roadmap is a good thing, and we view it as-- again, as much as we love OpenAI, we also deal with the entire set of customers out there from AI-native standpoint, as well as the largest hyperscalers, and we're seeing great traction with the roadmap.

TIM ARCURI: And are you any more engaged? Have you had any more conversations lately that you might not have had you not announced that deal?

LISA SU: I believe that it has given people a view of AMD's capabilities. I think, we always had good conversations, but I think, the idea of just how competitive the MI400 series roadmap is, what we have going forward, has certainly been helped since we announced the OpenAI deal.

TIM ARCURI: And do you worry about customer concentration? Can you speak a little bit about breadth? If you look out in your forecast, how broad will your customer base be?

LISA SU: Yeah look, our view is-- we are a general purpose supplier in the sense that OpenAI is a great partner and we very, very much believe in their success and their roadmap, but we are highly engaged across all of the largest hyperscalers out there. And from a customer concentration standpoint, the key point is this is a big multi-generational, multi-gigawatt partnership. We have a number of others that are at similar scale, similarly, multi-generation.

And the truth is, compute is a premium. This is one of the areas where there are so few companies that can offer this capability. I'd like to believe that in addition to great technology, we focus on our customer success, so it's about total cost of ownership, ensuring that there's significant differentiation, and also ensuring that we're very flexible in how people want to operate in terms of the overall ecosystem. So from that standpoint, I don't worry about customer concentration. I view this similarly when-- if I give you the example of where we were in the server CPU market when we started with the hyperscale accounts, they didn't all start on day one, at the same time.

The different hyperscalers went large at different points in time, and that's the same thing that we're going to see in the AI-accelerator roadmap. We're seeing a very similar pattern in terms of how we engage with customers and how customers view AMD as really a long-term partner, especially since there's this recognition that in addition to the GPU roadmap, the CPU roadmap, the networking roadmap, the overall capabilities are very attractive.

TIM ARCURI: Great. Well, we've made it 23 minutes, and we haven't talked about CPU yet. So maybe, we can talk about that. So demand is obviously very strong in both PC and server, so maybe, we can talk about that. We keep hearing about hyperscalers asking for supply, and we keep hearing about long-term contracts, particularly on the server side. So can you just talk about that? And just talk about this the supply-demand environment.

LISA SU: Yeah, absolutely. The last, I would say, several months has been a very interesting story around the CPU world. We are really happy and proud of our partnerships on the CPU side. I think there was this narrative last year that somehow GPUs were going to take over the world, and the refresh cycles for CPUs would lengthen, and you wouldn't have as much, let's call it, market momentum. I think what we started seeing at the beginning of this year is actually, significant refresh cycle starting, so that was very positive, but more interesting is over the last 3 months, what we've seen is really a significant uptick in CPU demand.

And when you look underneath that it's not just refresh cycles. I mean, there's no question that there were some refresh cycles that were, let's call it, delayed as a result of some of the AI CapEx spending. But a lot of that is being caught up now, and what we're also seeing is that as AI moves to more inferencing, and there's more work being done, and things like the agent workloads are starting, they're spawning more general purpose CPU needs. Because if you think about it, if you have, let's call it 1,000 agents or 1,000 virtual employees, they need to operate on some data set. They need to operate on some computing capability, and that requires general purpose CPUs.

So we actually have a view that the CPU market actually will substantially grow over the next 4 or 5 years, as we see the AI usage really spawn more traditional computing applications. So it is certainly a good thing to see. We love seeing that. I think it's one of the reasons that we're so passionate about the overall roadmap being important in terms of all of the capabilities, and we see the CPU business as a great business going forward.

TIM ARCURI: And you've gained a bunch share in data center, do you think that in server has your lead at all shrunk? Do you think that you'll continue to gain share?

LISA SU: We do. We're in a very fortunate place right now, where we are a trusted partner on the CPU side, especially for the largest hyperscalers, and the conversations are such like, how can we work together to build, let's call it, the best-in-class roadmap going forward? I think, as great as our fifth generation Turin is, we're super excited about our next generation, Venice CPUs. We think that extends our leadership going forward, and that extends as we go into the next generation as well.

So I think we have a very strong franchise there, and the key is we're a trusted partner going forward. We're also quite underrepresented in the enterprise space, but I've seen that also as a significant growth opportunity for us. The largest enterprises are all looking for help as to how they modernize their data centers and how they make their choices, and we're very happy to be part of that conversation.

TIM ARCURI: Great. One thing that I was quite surprised about from the Analyst Day was that you have pretty strong share aspiration gains in client, actually. You think you can be more than 40% share in client, can you just talk about that?

LISA SU: Yeah. So the client PC business has been a place for us that it's not a market that is necessarily growing by leaps and bounds, but it is a important market. It is a market that has a very good customer-facing capability for us, and we've grown extremely well over the last couple of years. I think we've really streamlined our roadmap. I think we have put it as an AI-first roadmap, and that has been appreciated. We're now, let's call it, mid to high 20s share.

And as we go forward, we see that only growing. There are areas where, I think, we are already best in class, when you look at things like, the desktop gaming market. This is an area where we've had historically a lot of success desktop channel market going forward. And we're continuing to grow in premium notebooks. Let's call it, the most valuable part of the PC, Tam is where the product does matter in the premium segments, and that's where we're actually gaining the most share because our products are superior.

TIM ARCURI: Do you worry that because memory prices have gone up so much, do you worry about some of this specking in PC, or do you worry that it hurts the market at all, that it hurts demand?

LISA SU: Yeah. I mean, we're certainly watching, Tim, the commodities. There's no question that as the market has gotten tighter, some of the commodities like memory have become tighter, and we certainly are watching for that. I don't think it's a major perturbance to the market. I think it might be a minor perturbance, and we're watching that closely.

TIM ARCURI: Great. And maybe, we can talk about some of the bottlenecks that you're worried about over that 2030 forecast you gave. Are there things that you're worried about like, HBM or CoWos or what is something that keeps you up at night that could constrain your growth?

LISA SU: Well. The great thing about the semiconductor market is I think we're used to expanding and expanding quickly. So if you put aside, very temporal things, what are the most important things? It is advanced technology access to the most advanced wafers. It's high bandwidth memory, Packaging, costs these elements. We have built a very, very strong supply chain over the last couple of years. We have deep partnerships with TSMC, all of the memory vendors all of the packaging vendors.

And I think, we feel very confident that we can achieve our growth rates. I think, the industry, as a whole, is very much around ensuring that we do satisfy all of the demand that's out there. The other area that we're watching very closely is power and how data center power is coming online, not just in the United States, but across the world. I will say that this administration has really activated a lot of the power build out, so we're seeing things moving faster. We're seeing that there is a desire to put more power on as quickly as possible, trying to get rid of some of the bureaucracy around that, and I think those are all good things.

We're also looking at power outside of the United States, and so there are lots of opportunities we didn't get to talk about, sovereign AI and a lot of the nation state, investments that are happening there, which we think are another adder on top of it. So I would summarize it to Tim as it is, there are lots and lots of things on the radar screen, but the most important thing is that everyone in the ecosystem recognizes how important the enablement of this computing technology is, and so we're all working together to do that.

TIM ARCURI: Great. Well, we're out of time. Thank you, again, Lisa.

LISA SU: Wonderful. Thank you so much.