

AMD Financial Analyst Day

**March 5, 2020
4:00 PM EST**

Lisa Su:

Good afternoon, everyone. Thank you for joining us today for our 2020 Financial Analyst Day. We have a lot in store for you, so we appreciate everyone's time. And it's really an afternoon to take a step back and talk about the long term. We're going to talk a bit about our strategy, a lot about our technology and our technology roadmaps and product plans, and then of course talk about our financial outlook over the next 4 or 5 years.

So let me start first with just a little bit of context. Some of you are new to AMD, some of you have followed us for a while. But I like to say that we've been on a journey these last five years. I've been CEO just a little bit over five years. And I've actually shown this chart pretty often over the last few years because it is a set of guiding principles for us. First and foremost, we decided to play to our strengths. For us, it's all about high performance. High performance technologies, high performance computing, high performance products. And this is the fundamental DNA of our company. And so that's why this is our focus.

We also set a set of pretty ambitious goals. And we set out to build a culture within AMD of really focused execution, meeting our commitments to our stakeholders, employees, customers, shareholders, partners. And we believed with all of this, we would be able to create a company that would grow and grow profitability as well.

When you take a look at our strategy, actually our strategy has been very consistent. Again, this has been sort of a fundamental tenet of what we've focused on. And it's really around 3 core competencies. I like to call them our crown jewels. When you look at it, we invest in high performance graphics for growing markets like gaming, cloud gaming, console gaming, for compute and AI and virtual and augmented reality. We think graphics will continue to be one of the most important elements of the high-performance technologies going forward.

We also invest in high performance CPUs. When you think about that, think about that in our client systems. Think about that in our infrastructure and cloud environments. Everybody needs more compute. Then when you really put these two together, it's really the underlying solutions that are very differentiating. And those solutions can be at the chip level, whether you're talking about semi-custom SOCs, or they can be at the solution level, whether you're talking about platforms that put together our hardware and software,

or you're talking about partnerships where we focus on deep code design with some of the largest customers in the world. These are our 3 tenets. Investing in high performance graphics, investing in high performance compute, and bringing those together in very differentiated solutions.

Now why do we love high performance computing so much? We really believe that this is a technology that is the enabler for both the present and the future. It actually drives what can be done over the next number of years. And whether you're talking about very, very big systems like supercomputers or the cloud and the hyperscaler environments, or you're talking about new workloads like AI and big data analytics and visualization. Or you're talking about the things we enjoy like gaming and new client devices. All of these have one thing in common. They all require high performance computing and they really play to our strengths. So this is our focus from a product standpoint.

Now as we formulated our strategy and plans over the last five years, the most important decisions that we had to make, and frankly that any technology company has to make, is around those technology investments. It actually takes years to develop a new architecture and to create that foundation to build great products. And when you think about those strategic decisions, they really lead us to where we are today. So these investments include things like our Zen roadmap. With our leadership CPU roadmap, we invested in Zen. It was a big, big performance boost. We're now with Zen 2, and you're going to hear more about what's coming.

We're investing in new graphics architecture with RDNA. And RDNA is actually very unique because it spans both consoles, PCs and even mobile gaming. And it will last us, again, for the next five years. We also made a choice, and this was an important choice, to move our entire product portfolio to 7nm. And very, very aggressively. And that has really paid off for us with now best-in-class manufacturing. And when we saw some of the constraints of Moore's Law, we said, hey, there's a different way to do this, there's a better way to do this. And it involves chiplet architecture which allows you to put the best technologies on a package and in some extent, break the constraints of Moore's Law.

Now, these things actually may seem pretty obvious today, but frankly, five years ago, they weren't so obvious. And they really are the strategic decisions that lead us to today's product roadmap.

So timelines are always interesting. There are lots and lots of products if you take a look at the last few years. We first demonstrated Zen actually in August of 2016, so that seems like a long time ago. But if you look at over the last four years of products, you'll see a couple of things. You'll see consistency in the roadmap, so both on the PC side as well as on the server side, consistency with what we've been able to do on the CPU side as well as consistency in the rollout of our new graphics products. And really a cadence of product innovation across the last four or five years.

When you look at that for 2019, 2019 was actually a huge year. It was a huge year for AMD because we did introduce 7nm across our entire portfolio. And in high-end desktops, that was third generation Ryzen. In the HEDT market, that was third generation Threadripper. In mobile processors, we introduced the 7nm Ryzen 4000 series. In graphics we introduced the new NAVI products with RDNA. And in datacenter, we introduced Rome, or our second generation EPYC.

And when you look across this product set, this is performance leadership. This is performance leadership. And I know when people say that, you're like, well what does

performance leadership mean? So let me just give you a view of how we think about performance leadership. This is actually a view of first let me talk about the PC market. This is desktop and notebook performance. And what this shows is let's call it the last five years of products in the industry which show relatively incremental performance. On the desktop side, this is multi-threaded performance with Cinebench. On the notebook, side, this is productivity performance plus graphics performance. But you can see it's been relatively let's call it incremental.

When we introduced Ryzen, it changed. With first and second-generation Ryzen, we became very, very competitive. With third generation Ryzen, both on the desktop side and on the notebook side, we have changed the performance trajectory. We have changed the performance trajectory. And by the way, that's what we mean by pushing the envelope on high performance computing. And so when you look at the business results of all of that, looking first at the PC market, look. We've had very strong results in PCs. When you look at desktops today, we're over 50% share in the premium segment at many of the top global retailers. When you look at mobile platforms, we increased the number of mobile platforms in 2019 by about 70% with Ryzen. And in high end desktop, we have the best product in the industry. It is the best product in the industry. And what that has translated to in terms of business is we have consistently gained share every quarter for the last eight quarters. Eight points of share in the last eight quarters.

And now when you look at datacenter, you see the same trends. You see the same trends of incremental improvement in the industry over the last five or six years. This is looking at SPECint Rate, which is a very normalized benchmark for datacenter computing. With first generation EPYC, we became very competitive. With second generation EPYC, we've changed the industry curve. We've literally doubled the performance of our competition with the second generation of EPYC. And again, we're very excited about the progress in datacenter when you look at some of the statistics. We love the cloud. We are expanding deployments in the cloud with all of the top cloud providers. We doubled our number of cloud instances in 2019, we expect to be at over 150 in 2020. We've expanded our platforms enterprise across a number of deployments and number of OEMs, that pipeline is growing quickly. We doubled the number of platforms in 2019 in enterprise we expect to be at over 140 platforms in 2020.

And in supercomputing, and let's call this a really, really good area for us, we're winning consistently the top deployments. That's the top deployments today and the top deployments over the next three or four years. One of the things that I can say we're very, very proud of is the fact that we were selected for two of the largest DoE deployments for supercomputers with both Frontier last year at Oakridge National Labs, and just yesterday we announced with Lawrence Livermore National Labs that they have selected AMD CPUs and AMD GPUs for El Capitan, which should be the most powerful supercomputer in the world in early 2023. So lots of good progress on the datacenter side.

Now, a couple of points on graphics. Look, we're investing in graphics. We're investing in graphics. Last year we delivered our first generation of the new graphics architecture RDNA. We saw significant gains in performance per watt as well as overall performance and you're going to see us continue to invest in graphics. That has led to some very nice progress in gaming. Including, if you look at our RDNA family, the NAVI family, we're winning at 1080p and 1440p. This is a very important market for us. We are the exclusive discrete graphics provider for Macs, very important partnership for us. We are being used in the cloud in many places. Very important partnerships for us. And of course, we love game consoles. Game consoles, if you think about all of the folks that have game consoles between Microsoft and Sony systems, in this current generation we've shipped

over 150 million units since we started in 2013. So a lot of momentum there overall.

Let's just spend a few minutes on how we think about customers. In addition to products, we have majored on building very deep customer relationships. What that means is, it's beyond a typical roadmap customer/vendor relationship. It's really about how we do something special. How we co-develop, how we co-design, how we co-innovate. Microsoft and Sony are great examples of that. We love what we do in the console relationships. We've extended our Microsoft relationship to Surface and that's a very important partnership for us.

When we look at the cloud, the cloud is all about what we can do to optimize for the key workloads and you'll hear more about what we're doing with those relationships. And then what we're doing with the OEMs as well to bring out the new user experiences. So lots of focus on deepening our customer relationships and creating competitive advantage in the ecosystem.

Now what has that translated into over the last five years? It's translated to exceptional performance. If you look at where we were in 2015, we were about a \$4 billion company. We finished 2019 at about \$6.7 billion. That's 14% annual growth rate over the last few years. And this has actually come from growth across our PC business, our discrete graphics business, as well as our datacenter businesses. So really on the strength of those new products. And that growth is translated into significant margin expansion. We've expanded our margins by more than ten points in the last few years. We've improved profitability, and we've substantially strengthened the company and the balance sheet. So it's fair to say we're a much stronger company than we were a few years ago.

Now, as much fun as the last five years have been, today is really about the future, and we like to talk about what we see over the next five years. I can say for sure if you ask me or anybody on this leadership team, we are even more excited about the coming journey in terms of what we can do. And the reasons are very simple. First of all, the opportunities are larger. The impact we can make on the industry is larger. And our resources are much stronger. So if you think about those things and what we've been able to accomplish, it's really exciting to think about what we will accomplish.

So again, some of the guiding principles that we think about. First and foremost, and you're going to hear us say it probably 100 times this afternoon, maybe 101, we will stay committed to high performance computing leadership. That is our mantra. We are uniquely very, very good at it. And frankly, there are very few companies in the industry that can possibly do it. It is extraordinarily hard to stay at the bleeding edge.

We actually see even more opportunities to combine our CPU and GPU solutions. We combine them today, by the way. If you look at our solutions today in PCs and consoles, they put integrated CPUs and GPUs together. But we see that opportunity broadening and becoming more disruptive as we go over the next five years. We also will continue to prioritize very strong and predictable execution. We want to be a trusted partner for our customers. We want people to come to AMD first because they know that they can count on us. Then frankly, from a business standpoint, our aspirations are to be a best-in-class growth franchise. And we don't take that as sort of a light expression. That is our aspiration.

So let's talk about the market and what's happening in the markets. We love our markets. Our markets are big, they're growing, and they're places where it's very clear who our competition is. When you look at our TAM, it's about an \$80 billion TAM and this is let's

call it a 2023 number. We see datacenter at about \$35 billion. Datacenter includes CPUs, it includes GPUs, and it includes some telco and infrastructure. We see PCs at about \$32 billion. This is not necessarily a growing market, but it's a very large market and it's a large, good market for high performance computing. And we see gaming, and within gaming we consider consumer graphics as well as game consoles. Again, we see this to be a good market. Lots of people are gaming. More people are gaming and they all want better graphics. So it's an exciting TAM and it's an exciting opportunity for us.

Now today, we're going to spend a lot of time talking about our technology investments. When I talk to you guys, often you're like, what makes you so confident that you can continue the leadership in products? And it really is about the choices that we're making, the choices that we have made, the choices that we will make. And if you look at it, there are really a couple of key areas. First and foremost, we're going to continue to invest significantly in our cores roadmaps. That's our CPU, Zen roadmap, that's our RDNA roadmap. These are the baseline for what we have to do from a technology standpoint.

We're going to be aggressive with advanced technology. That has played to our strength, that's where we are. To do high performance compute, you have to be aggressive with advanced technology and that's in process, packaging and interconnect. We're very excited about datacenter. We think that datacenter has some very unique characteristics as new workloads come in and there's a lot of innovation to be had across cloud, enterprise and accelerated computing. And we're equally excited about PC and gaming solutions because, again, there are lot of new user experiences. So these are the areas of our technology investments.

Now my team is going to talk about this in a lot more detail, so I'm just going to give you like a small preview over the next few minutes. But it gives you an idea of what we prioritize. So industry leading CPU and GPU roadmaps. These are IP roadmaps. And David and Mark are going to go through these in more detail. The main thing that I want to say is that you can count on us to have a very strong cadence of continuing to innovate on the CPU and the GPU architecture. We can see the path. We see the path today and our teams are working on these things today.

In the area of advanced technology, I said we would be aggressive and we'll continue to be aggressive. Mark will talk more about this. In process, we are in a very, very good position, a very good position with 7nm. And we're committed to being aggressive with advanced process nodes. I think that is really part and parcel to the strategy. We're also very clear that packaging is key, that Moore's Law is slowing down, that packaging is a way to break some of those constraints. That chiplets are great. There's next generation chiplets, there's work that we can do in terms of 3D die stacking and we'll talk more about that.

And probably the area that is somewhat underappreciated when you talk about advanced technology is in interconnects. As sexy as the individual components are, how you put them together makes an incredible amount of difference. And we are uniquely positioned to really drive that interconnect architecture.

Moving onto the datacenter market, look, I said we were really excited about this area. We are really excited about this area. Forrest is going to spend quite a bit of time talking to you about it. The excitement comes from the fact that there is just insatiable demand for more compute. Everybody needs more compute, no matter where you are. And it's not just more compute, but it's different compute. And there are new workloads and new problem sets and new ways to solve the problem no matter where you're looking in this

ecosystem. Our view is that there's a huge advantage if you think about solving the problems differently. And that's where we're focused in datacenter.

Some of the key centers of datacenter leadership, again, you can expect these. CPU roadmap. Today, Zen 2 is the best CPU in the market. Second generation EPYC with Rome is the best X86 server/processor in the market. We intend to continue that with Zen 3-based Milan on track for later this year.

Perhaps one of the things that may be very new today is what we're going to do in datacenter GPUs. We used to really share the architecture, so our GCN architecture was shared between consumer graphics and datacenter graphics. But when you look at the workloads going forward, there's really an opportunity to optimize. David is going to share with you our new compute roadmap for GPUs. We call it AMD CDNA. We like that. It stands for Compute DNA. And what you can expect is this a beginning of a new roadmap that's going to take our GPU compute architecture forward, particularly around HPC and machine learning. And you can expect a cadence like we've done on the CPU side with Zen on the GPU side with CDNA. And then really exciting is how we put this system solutions together and really form together with our CPU roadmap, our GPU roadmap and our new interconnect capability, the best system solutions in the industry. That includes hardware solutions as well as investment in software solutions on the platform side. So that gives you an idea of how we view datacenter and our bets in the datacenter.

So moving onto PCs and gaming, look, we are as excited about the opportunities here. It's a different market. It's a market led more by user experiences. But the world is different today. People have a different expectation of what you can do in a notebook or a desktop or a gaming system. And in the PC market, people want more performance, they want more capability, they want more portability, they want more security, and those are things that we're good at. And for gaming, you're going to hear from Rick, there are like more than 2 billion gamers in this world. And frankly, they have a lot of expectations. They want to be able to play their games anywhere, anytime on any form factor while interacting with their friends and family. And they want to do it at high resolution. And that requires also plenty of CPU and GPU capability. A lot of technology here.

And in this part of the business, what you're going to hear from me and what you're going to hear from Rick is we are committed to building the best products for PCs and gaming. And we think there's a lot of opportunity. A lot of opportunity. In the PC market, it's a large market that has lots of demands across both notebook and desktop. We've done very, very well in the desktop market particularly in the DIY market. But we are still very underrepresented in consumer and commercial systems. And we'll talk about how we become more representative in those markets.

In graphics, we have the RDNA roadmap, but the roadmap also needs a great set of products to go along with it. And as I hear all the time from gaming enthusiasts, we are committed to a top to bottom gaming portfolio and that's a multiyear, multigenerational commitment.

And then on the console side, look, we are honored to be partnered with Microsoft and Sony for their next generation consoles. It's probably the most anticipated consumer launch of 2020. And what we do with each of them is really help them power their visions for the next generation of gaming with custom SOCs. So that gives you an idea of what we have in store for PCs and gaming.

Okay, so lots of exciting technology and products that you'll hear about today, but I also want to make sure that I give you a preview of how we're thinking about driving shareholder returns. We believe we have great markets. We believe we have great products. And we also believe that we are underrepresented in the TAM. There is a lot more that we can do. And so with continued execution, what we're driving towards is best-in-class growth. What we're driving towards is making those right investments not just for today, but for the next five years. What we believe is with our continued expansion of our product portfolio we will continue to expand margins and grow profitability. And we also believe we'll generate a significant amount of cash in that timeframe. Now Devinder is going to go through much more of this towards the end of the afternoon, but I thought I would again give you a preview of what that long-term financial model looks like.

So, in 2017 at our Analyst Day, we set out a model, which at the time, some asked whether it was a bit aggressive. It was double-digit annual revenue growth. It was gross margins at about 40% to 44%. We were about 30% margin in 2015. And it was about increasing profitability and cash. And I'm happy to say, if you look at the numbers for 2019, we more or less met these goals. In some cases, we've exceeded these goals. And so we feel good about the progress. Now, as we project out for the future, we're even more excited about what we think we can accomplish. And for our long-term model that we're going to talk about today, it actually has a few very important components. The first is, we believe we can accelerate growth. With where products are positioned, with where we are positioned with our customers, we believe that growth will accelerate and we estimate for the new model, we're calling this a long-term model, think about it as four years. So let's call it like a 2023 model. We believe we can deliver approximately 20% compound annual growth rate over that timeframe.

As we expand what we do in the commercial markets, particularly in datacenter and commercial PCs, that will help expand our margins. And so we see margins increasing from where we ended 2019, approximately 43%, to greater than 50%. We think operating margins are in the mid-20s, so again, I think we're building a very balanced model where you see growth, where you see investment, margin expansion on both the gross and the operating line. And again, our goal is to generate a significant amount of cash as we grow the business. So Devinder will have much more on this, but hopefully that gives you an idea of what we're trying to achieve. And when you think about all the technology and all the products and all the markets, this is the financial model.

So let me finish up here and just state that look, we are very ambitious with where we think we can take AMD. That ambition is motivated by building the best. So that's the motto for today. It's about building the best on both the technology side as well as the business side. And that comes with leadership in our roadmaps, that comes with really execution excellence. It comes with market share gains across all of our markets and a commitment to strong shareholder returns. So with that, let me turn it over to my team so that we can give you a lot more detail on how we get there. And with that, let me introduce Mark Papermaster to the stage to talk about our technology. Mark?

Mark Papermaster:

Well thanks to all of you for joining us here today. You know, I've had the opportunity to share with you our technical approach over the last eight years on our journey back to high performance here at AMD. And when you think about that journey, we've embraced that a company. We've embraced it in our strategy and in our culture. And so it's exciting to be here today and to talk to you about our view of our status, of how we've done. And more importantly, our journey going forward, staying on that pace to deliver high performance to the market.

Lisa touched on that earlier, but it is such a fundamental about what we do in our technology strategy at AMD, that I'm going to spend just another minute talking about these workloads in the markets that we serve. Because it is an everchanging landscape across each of these segments. But there's one thing in common and that is that incredible demand for more performance, more high performance in each of these segments. In fact, it's actually an exponential growth in most of these segments. It's not a surprise to you why that is, because you see it every day. You see smart devices around you. You have them in the appliances. You've got your smart home. You see on the factory floor the telemetry. Smart functionality has been built in across the factory floor. You see it at the emerging edge of the market. When you look at how 5G will transform the market, it changes the analytics that has to occur at the edge of network, in base stations, in the closets, traditional telco closets across the industry. And it is the approaches that supercomputers have deployed to bring massive scale compute to solve the emerging workloads in AI and analytic challenges that we're facing. You look at decision making, the massive data that's going into decision making today are requiring large scale simulations.

What about content delivery? The Olympics are going to broadcast later this year in 8K. That's a tremendous demand in terms of delivery on that content. And of course, game serving, driving up the capabilities in our cloud server.

And then bring it down to the client interface level. You look at gamers, you look at content creation, you look at what we are all doing on our PC devices today, and we want more and more visualization, more clarity, more immersive experience. And of course, we want a higher capability and efficiency at that level, at that client level.

And so all of these factors come together and that is what is driving our strategy. It's our role at AMD with the products we develop to enable our customers to harness that data and put it to work. And moreover, to do it with devices that are easy to program, that have been out there for years and have an entire ecosystem around them to bring those end solutions to bear. That is what drives our strategy. And we couldn't be where we are today, as Lisa mentioned earlier, without decisions that we made several years ago. And I will have to say that it wasn't hard for us when we saw where the workloads were going to call out that strategic focus on high performance. That was a strength that we knew that we could tap on with such deep experience in products over that area. But we made a set of tough calls on investments. We fundamentally had to transform our ability to deliver that performance. So we had the building blocks but we transformed that delivery process in a way that we could be trusted to deliver each generation after successive generation. Be a trusted partner.

So that was a change in our execution model. We changed our roadmap and we changed the very way in which we put the IP blocks together through a modular approach. And that's what we'll spend a few minutes just looking at some of those accomplishments over the last five plus years.

I'll start with that reengineering of the engineering approach at AMD. Frankly, it's one of the key accomplishments. It's I'd say equal to the tough technology decisions was some of the changes we made in the delivery model at AMD. It was a shift to a culture of high performance, a culture of collaboration, and a culture of top flight engineering execution. And in the prior FAD, I shared with you some of the ways in which we are making those changes and they have in fact proven to be highly effective. We talked before about this idea of leapfrogging teams. So what is a leapfrogging team? It's where what we did on

our CPU roadmap. We were a multiple CPU line so we were split out in our focus. We consolidated on a single Zen family. Not just one product generation, but we were working on multiple product generations from the outset. And then what we do and what we still have today is we always have one generation going to market, one well along in the design, and one in the conceptual phase. So this has been fundamental to improve our ability again to be that trusted supplier to our customers.

But you know, we went much further than that. We changed the way in which we brought our IPs together. We historically had very different methodologies across the company. It was hard to collaborate when everyone is sort of rolling their own, doing it their own way. And so what we did when we adopted that modular approach, yes, it was key to hit the performance goals, that's what drove us was to hit the performance goals. But equally, it facilitated the culture change. Because for modularity, you have to architect how the pieces come together. That drives the collaboration across the company, drives that co-engineering of how we design going forward. And it frankly changed the plumbing of how we put products together forever going forward at AMD.

We also made a change in how we put our simulation together, how we verify our designs. Because we needed to accelerate how we bring products to market, so we invested in our simulation and emulation techniques to drive earlier verification of the features that we were designing in our new products. And it turned out to be quite impactful. Because what we then were able to do through this effort is to drive left. To drive, if you think about a schedule that you lay out in front of you, and you typically historically have had that feature validation of software features on top of the hardware to the right of a line, that line is when the silicon comes back and you're testing that silicon. So that feature enablement was done in the bring-up, in the lab with that silicon. And what we've fundamentally done is we've shifted that validation left. We've shifted it pre-silicon, leveraging these advanced simulation/emulation techniques that we've deployed.

So yes, it had an immediate impact in terms of then when you validate that pre-silicon, the hardware comes back now, we're completing that bring-up in hours and days versus the weeks and months of the methodologies that we had historically done. And, it actually changed the way that we even architect our solutions because it enabled a parallelization of the architecture, the hardware and software architecture of putting these solutions together. And a great example of that is Modern Standby. Modern Standby, you all have, if any of you have laptops, it's what's giving you improved performance and energy efficiency across your device. And it's complex because of the high software/hardware/firmware interactions to get that done. So we partner from architecture through delivery with Microsoft, with our AMD software team, our AMD firmware team, chip design teams, platform teams. All of them partnered together leverage these approaches and it was -- you see it. We are shipping Modern Standby on our devices today and you see it. It's working flawlessly.

An excellent example of how we changed our approach, and frankly, this is the new normal at AMD. This is our expectation. This is how we're developing all of our products going forward. And look at our execution that this has led to on 7nm. Lisa called out that that decision, that choice, had to be made some years ago. And we have delivered as promised in 2019 over the course of the year, a comprehensive portfolio. Over 20 products in the market now with 7nm. You look at what 7nm did by doubling the density, and you could do this in roughly the same power envelope. Look at our Ryzen roadmap. We doubled the number of cores. Ryzen HEDT, high end desktop, doubled the number of cores. EPYC server doubling to 64 cores per socket. And then, in the most recently announced notebook, with the Ryzen 4000, doubling to eight cores, still having that

integrated graphics capability for leadership and with all-day battery life. So 7nm successfully executed and rolled out, again, over 20 platforms. Let's do -- oh, and of course NAVI, in our graphics line, with the impressive performance per watt games leveraging 7nm, so across our portfolio at AMD.

What I'd like to do now is just dive a little deeper on our delivery of that new Zen 2 core. When you look at that architecture, as I said a moment ago, 7nm was key of course given that density and energy efficiency. But the design was fundamental in both performance and the scalability of our implementation. We leveraged the second generation Infinity architecture, taking that Infinity fabric and broadening its application in terms of actually allowing us to still provide that on chip and socket to socket and server connectivity that we had done historically with our Infinity approach. But now we leveraged it to actually implement a chiplet approach. So what it allowed us to do, you look at our server implementation, and you see eight smaller, easier to manufacturer, 7nm CPU die connected with that Infinity architecture to essential 12nm IO and memory chip. So it delivered performance, it delivered all the scalability, it delivered -- it's actually easier to leverage that performance because it created a single pneumodomain. And it provided on top of that with the performance enhancements of feeding those engines even greater than actually with a PCIE Gen 4 to double the rate that we had before. So that Infinity architecture was fundamental to putting the whole picture together, delivering new CPU performance, 15% instruction per clock gain, but adding it at the holistic level to make sure that we could deliver the core performance.

And look at that core performance. Just going -- when we introduced Zen, it was a breakthrough. 52% instruction per clock. As I said a moment ago, 15% instruction per clock at Zen 2. But the knock we had taken on that first gain is we had previously had a gap on single thread performance. And we set out with Zen 2 to eliminate that gap. And you can see on Cinebench single thread, which is a great benchmark, very representative benchmark for content creation, 21% improvement on that single thread performance. A majority of that driven by the design improvements.

We changed the branch prediction scheme to be more accurate, so you have more accuracy of the code prediction. We improved the pipeline execution, expanding the width as well as our dispatch efficiency. We doubled our floating point to a 256-wide pipe. We doubled how we feed it, so it's a true doubling of floating point. And with that double core density, it's actually 4x the floating-point capability on every Zen 2 product that we shipped. And particularly important to gamers and to server, is that latency to memory. We improved that effective latency. We doubled the L3 cache. So really, really important decisions that deliver real world performance. And that's what's driving the rapid growth, the acceptance of the Zen product family in the industry.

I want to show you just a few stats on that. The bulk of our shipments of core, this is a count of cores, so clearly the highest volume is in the client products. And you can look with that Ryzen launch in early 2017, about a year later, 50 million Zen cores shipped in the market. We bumped that performance with a key Zen Plus release and you can see that we jumped to almost 160 million devices with that second-generation Ryzen.

But critical was Zen 2 and that acceptance that we've had and really blowing through the levels of historic captive performance of CPUs in the industry, blowing through that. And jumping in almost, just actually about 8 months, you see that we're now over 260 million Zen core family chips in the industry. So very wide acceptance. And with the strength of our roadmap, you'll see this curve do nothing but increase in its steepness going forward.

So as I wrap up, looking at that set of products and the execution that we've had just in the recent past, I'd be remiss if I didn't take a moment and talk about product security, because it is the foundation. You have to be trusted to be able to sell your compute devices. It's the bedrock of what we do. It's job one at AMD. So of course, we include a dedicated, secure processor in every device that we ship. We boot up in a trusted environment, we authenticate before we communicate and talk to any other device, and that continues strong. We continue to be very resilient in terms of the security of our CPU designs and we remain focused every day. You look at security events, we were affected with Spectre v1 and v2 like every CPU chip in the industry. We implemented a software mitigation immediately, we hardened that in Zen 2 for Spectre v1 and v2. But the resiliency of our designs means that when you look at the other side channel attacks that you've been reading sort of continual reports there, our architecture was not affected by Meltdown. Our architecture was not affected by SWOT GS. We were not affected by Zombulo 1 or Zombulo 2. So we remain incredibly focused here and we want to continue to rollout enhanced security features. Our encryption has been well accepted in the industry. Memory guard, where we have a memory across -- a memory encryption, very seamlessly implemented across our devices. The ability to then move to virtualized machines and have unique keys across unique instances in a cloud environment has been very well received. Now rolled out across the OS and hypervisor partners in the industry.

And the way that we do that there, of course, is done without the need for any code modification. So it's now -- it does not require any lift of the applications of our end users. And so you're seeing further adoption of this approach. And then going forward, again, continuing to invest in modern security going in across our devices. We will be increasing our security in those cloud environments, multi-tenant environments. So we're ready. Our leading-edge encryption abilities I just described with SEV, but growing that to add a capability to protect for what's called malicious hypervisor, somehow a bad actor can get into even a hypervisor of that cloud applications. And you'll see that coming in our roadmap soon.

And then lastly, I'm very pleased to announce that AMD has joined the Confidential Compute Consortium. We feel very proud about these advances I just described to you. But we want to join with others in the industry and work to close the final gaps to protect data throughout its entire lifecycle. So again, security will remain a bedrock and foundation of everything that we do.

Okay, let's shift gears. Let's go forward and look at really our investments. Because as excited as I am about what we've done, I couldn't be more excited about what's coming in the future. How we put our solutions together around our CPU and GPU roadmaps will be very exciting as we leverage the process and the packaging investments we're making as well as our next generation interconnect. And what's so special about what we're doing with these investments is how they come together and enable the accelerated computing and how that will drive our platforms going forward in AMD along with the software stack you'll hear next about as well as other features that are coming, from David Wang next session.

So start with roadmap. Sustained execution. We talked out how critical that is in the very foundation at AMD. So you saw the progress we made with Zen 1 and Zen 2. Zen 3, right on track. It's coming up well, it's on track for delivery late this year. And what I am really excited about is sharing about Zen 4, next generation being in the 5nm design. We're working with our foundry partner on 5nm in that same close partnership that we did for 7nm. We bring the knowhow of how to marry design and foundry technologies for high performance. And so we're continuing that same partnership and that execution for 5nm.

So look, we've called out a very clear CPU strategy. Our roadmap is stable. We're executing. We are heads down in executing and this trajectory will keep us positioned to be that trusted supplier and to meet the demanding applications going forward. You know, I have to tell you that when we designed that roadmap, we designed it assuming that like for most of my career, that we would have that continued gap versus the process capabilities of our X86 competitor. Because that's the world we had lived in. That's what we designed for and that's what we put in place. We anticipated that FinFET would help close that gap, which it did. So 14nm FinFET made very good advancements. The plot you're looking at is showing a relative server product density and relative server performance per watt versus our competition. And so you can see that we really had always learned to design with efficiency to account for some of the gaps that we've had. But what was historic was at 7nm. We did not anticipate that at 7nm that we'd actually have a leadership process capability. And so what we are doing going forward is continuing that partnership, executing very, very closely with our foundry partners, and we're assuming that our competitor will address and will come back. But once you have a gap, once you have some of the issues that may cause a delay in a new technology, it's going to take some time. So what we're doing, is we're assuming that the competitors will come back at some point, but we will keep our trajectory at the same pace of competitiveness that it's been on. And in era of Moore's Law slowing down, once you have that historic level playing field, it is about how you put the solutions together. And so it is about these innovations of integrating solutions. And that's again, where AMD will not let up.

We were at the forefront of packaging technology. You go back to 2015 when we implemented our stacked high bandwidth memories of first-generation high bandwidth memories over a silicon connectivity, silicon 2.5D packaging to a GPU, a leader in applying this approach which lowers the power and dramatically improves the GPU performance. We've continued that approach with our high-end GPU products. In our CPU, we've had excellent experience with multichip approaches and then innovated as we talked about with Zen 2 on a chiplet approach to give us tremendous configurability as well as more performance and scalability going forward.

And what I'm really excited about is sharing with you our view of the future. Because what we're working on with several of our future products is actually marrying those approaches together. Working on improving the density and also marrying the capability of what we practice on our 2.5D and 3D packaging along with chiplets. We call this our X3D approach and it plays perfectly into our modular approach that we have in AMD. So investment in packaging driving that flexibility and that density and efficiency of packaging going forward, and with it, our Infinity architecture.

So if you look at that roadmap of Infinity architecture, it was first allowing us to have some breakthrough CPU connectivity which you saw in our first generation. And with the second generation, again, that high performance chiplet implementation in the CPUs and we applied it as well on our GPU roadmap to be able to connect in either 4-way or forthcoming up to 8-way configurations of GPU connectivity. So that was a huge step forward in terms of our scalability and our roadmap.

And very exciting to share with you today what will be coming out in future products is our third generation Infinity architecture. And with that, we complete that connectivity across our CPU and our GPU roadmap. Bringing that Infinity architecture linkage and bringing a coherency across those engines.

What does that mean? It means performance. It means efficiency. It gives us unprecedented bandwidth in the industry that you'll see between those devices. It reduces the latency. But more importantly, it unifies the addressability. As you cache from the GPU into the CPU, it looks like a unified pool of memory that's available. And with a coherency, it's easier to program. The programmer is not having to manage those elements. It's very straightforward to access. What does this mean with these types of improvements? Well, look at those machine learning training that are demanding massive amounts of data, where the models are growing tremendously. They demand this type of connectivity. Look at another example, motion picture rendering that needs that combination of both CPU and GPU. So many other applications that are coming up. We're very excited about what our third generation Infinity architecture will enable.

And you don't have to look further into how this is going to play going forward than looking at accelerated computing. Because it is something that you don't have to think that far back on. Go back to 2012. That's when the gradient descent analysis sort of spawned this whole AI approach to be able to rapidly manage data. And that's where you started to see the supercomputing take advantage of GPU and CPU. That's the first wave of heterogenous computing and leveraging for both high performance computing, HPC markets, and machine learning, leveraging that heterogenous approach. But now the model sizes have exploded. The previous approach isn't good enough. It takes this combined optimized system of CPU, GPU and the ability as well to bolt on accelerators with that, that we're providing, that is the new era of computing. It's the next era of computing. It is the exascale era of computing. And I will tell you that it was a heated competition for these systems that Lisa talked about in the Department of Energy with Frontier, with Oakridge National Lab and El Capitan with Lawrence Livermore National Lab. And we love the fight at AMD. That is what we're made of. And so that fight drove us to improve our roadmap, to improve our competitiveness, to go head to head and simply beat the competition. And it is laying the foundation for our next investments for the long term at AMD. Because it drove us to the edge of what was possible.

And that's where I'd like to end my comments is on that point about the long term. Because that's what we're about. We have a deep R&D commitment at AMD. We have a deep experience at AMD. And now we've tied to that an incredible execution culture. So our core roadmaps will be relentless in computing, gain of compute as well efficiency gains. Generation after generation, we've led the industry in innovation, in modular and chiplet approach, and it'll keep us on a Moore's Law pace of performance even when Moore's Law itself, the semiconductor node alone, is tailing off. We've invested here for the future. And our execution approach, our culture, is in fact a differentiator. And it's allowing us to be a trusted supplier to Fortune 500 companies across the globe.

And lastly, our successful journey to exascale computing is driving the next wave of innovation at AMD. We will not let up on the pace of development at R&D. Thank you very much. And with that, I'm really pleased to introduce my partner in technology development, Senior Vice President of Radeon Technology Group, David Wang.

David Wang:

Thank you, Mark, for covering the exciting CPU and interconnect technologies and the future of accelerated computing. So now it's my turn to talk to you about GPU. Some of you may know I rejoined AMD 2 years ago. I've actually been spending most of my career doing GPU development. And my top goal here is to drive GPU leadership. So I'm very happy to be here to share with you our GPU technology roadmap and our journey to drive leadership in gaming, in datacenter and in accelerated computing.

Now let's dive in. Let's start with our vision. Our vision is pretty simple. We want to

drive AMD Radeon technology everywhere. At AMD, we have made tremendous progress expanding our Radeon ecosystem across PC, Macs, game consoles, now to datacenter and mobile. This is a very, very broad ecosystem spanning from cell phone to supercomputer. And to drive GPU leadership across such a broad spectrum of workload required huge focus on technology development.

So next, I'll talk to you about our technology development strategy. We have a very simple and clear strategy on process, architecture, efficiency and software. On the process side, driving aggressive adoption of advanced process technology Mark mentioned about. Because process is the foundation. But it's not enough. We also want to develop domain-specific architecture so the architecture is optimized for its workload. More on that later.

We want to drive aggressively the performance per watt and the performance per area efficiency improvement. Because the leadership on performance per watt drives higher ASP. And the improvement on performance per area will continue driving down cost. And lastly, we want to leverage our opensource software strategy to continue to expand our ecosystem.

So Mark talked about the process. I'll cover the rest in my presentation today. We want to develop domain-specific architecture because with Moore's Law slowing down, it's very, very challenging for the general-purpose architecture to achieve optimal efficiency for both gaming workloads as well as the high-performance compute workloads. And therefore, we are shifting our strategy from a GPGPU type of architecture to domain-specific architectures. That is the RDNA architecture optimized for gaming and the CDNA architecture optimized for compute that Lisa mentioned earlier. This allows us through domain-specific optimization to achieve optimal efficiency for gaming and for compute. It also means that end users don't have to pay for the performance and the features that they don't need for their application, so it's win/win.

So next I'll cover RDNA and follow by CDNA. We launched the RDNA architecture last year. It was all new architecture designed and optimized for gaming with the objectives of driving efficiency in the performance of modern gaming workload, improve power in the bandwidth efficiency, and also provide a flexible platform to implement software features such as enhanced gaming experience. And lastly, the architecture was made to be super scalable to be able to support from mobile gaming to cloud gaming.

In order to achieve these objectives, we have developed the following architecture innovations. We have a new compute unit design that has a very, very efficient pipeline for diverse and dynamic gaming workload. We have also developed a new multilevel cache that can feed data to the new compute array in an energy efficient way. And we also optimized our graphics pipeline to deliver the highest performance per clock and also improve the clock frequency. And as a result, we're able to deliver more than 50% of the performance per watt improvement from GCN to RDNA through a combination of architecture, 7nm gain, and the design optimization.

So now let's look at the young RDNA. RDNA is our new foundation of our multiyear, multigenerational gaming GPU roadmap. RDNA 2, our next generation, will continue to innovate, to raise the bar on improving performance per watt while adding advanced features such as ray tracing and variable rate shading. You're going to see RDNA 2 based products from AMD and from our partners later this year. And with our strong development pipeline, the RDNA 3 is also on the way and we'll continue to push for higher perf per watt improvement and also new features, so stay tuned.

Now let's take a closer look at how we are improving the RDNA perf-per-watt. We are leveraging the proven CPU design methodology on the Zen roadmap to make similar perf-per-watt improvements on the RDNA roadmap. We focus on 3 main areas. The micro architecture innovations to improve the perf-per-clock or the IPC in CPU terminology. And while enhancing the logic, reducing the logic complexity and associated power. And put together, we have a new physical design flow to drive the highest possible clock frequency, multi gigahertz clock frequency for our graphics engine.

With all these enhancements, our plan, our target is to drive another 50% enhancement for RDNA to RDNA 2. As you can see, we have established a very strong generational perf-per-watt improvement roadmap which will extend to RDNA 3. And this strong roadmap allows us to really drive desktop and notebook gaming leadership. And this is also a key reason why Samsung has chosen our graphics IP even for their mobile applications.

Now let's talk about ray tracing. Ray tracing is an interesting technology for gaming. However, as you know, the adoption has been slow. Mostly because of lack of content, lack of hardware, and the performance penalty when people, when gamer time is down. So we have developed all new hardware accelerated ray tracing architecture as part of RDNA 2. It is the common architecture used in the next generation game consoles. With that, we have greatly simplified the content development, that developers can develop on one platform and easily port it to the other platform. This will definitely help speeding up the adoption.

And we also provide lower level API support that gives control. That gives more controls to the developers so that they can detract more performance from the underlying hardware platforms. This will help mitigating the performance concern. You can take a look at the image that was rendered on the RDNA 2 silicon running the latest Microsoft DXL 1.1 API. That was co-architected and co-developed by AMD and Microsoft to take full advantage of the common ray tracing architecture. This is a great proof point of the benefit of Radeon everywhere. I encourage you to check it out in the demo area after the presentation.

So now I want to switch gears to datacenter GPU. So we mentioned about we want to develop domain-specific architecture in order to improve efficiency which is extremely critical for datacenter operations. And that's why we have designed the CDNA architecture that's optimized for datacenter compute. With objectives to enhance performance for HPC and machine learning workloads with specialized compute in the Tensor operations we added in. And to reduce the datacenter total cost of ownership through the similar type of perf-per-watt efficiency improvement that we borrowed from the RDNA methodology.

To add features that will enhance the enterprise grade RAS and the security and virtualization support, because all are critical, critical for datacenter operations. And lastly the architecture must be scalable in order to scale the performance for multi-GPU and exascale computing. Of course, leveraging our Infinity architecture that Mark mentioned about.

So this is our modern generational compute GPU architecture roadmap. We introduced our first 7nm GPU last year based on the GCN architecture. So this year, we will launch our first CDNA based product that's equipped with second generation AMD Infinity

architecture that will greatly enhance the multi-GPU connectivity. The product will be optimized for HPC and machine learning applications.

And to move forward, our next generation CDNA 2 architecture will continue to push on performance per watt. But even more importantly, you'll be equipped with our third generation Infinity architecture with CPU/GPU coherency that will extend architectural capabilities to support exascale computing. So we have very, very strong roadmap for datacenter GPU.

So now let's take a look at software, because hardware without software doesn't go too far. So let's look at our software strategy. Our datacenter software on compute is based on what we call ROCm. It's Radeon Open Compute platform. It is fully open sourced and our partners and customers, they love opensource because it enables them to innovate and to differentiate, to create their own value-add solutions without being locked into proprietary solutions. And ROCm also provides multi-platform support. People can develop in the HIP code which is platform agnostic, opensource API. And that can be compiled and run on any existing GPU. And we have also provided tools, translators if people want to convert their current code to the HIP code to preserve their investment while moving to open platform.

And lastly, the ROCm software architecture was built to scale. You scale the multi GPU performance and that's critical for exascale computing. And the software can also leverage the expanded bandwidth and the coherency from the Infinity architecture. So we have been increasing our software investment in the last few years, so let's look at our progress. So the journey started in 2018. At the time, you can tell we still had a few software components were still in the early phase of development. But we had built a solid foundation.

Last year, we put our focus on building a complete software stack to accelerate machine learning. And by working closely with Google and Facebook, AMD GPU is now officially supported by TensorFlow and PyTorch. And the combination of both represents more than 95% of the machine learning applications. So I would say we are pretty well covered there.

And this year, our plan is to develop the complete exascale software solution to cover both machine learning and extending to HPC. This is to support our supercomputer design wins. And there is no better way than working with exascale customers to make sure our software is ready for large scale deployment.

And lastly, our ROCm ecosystem is also growing. This is a partial list of our ecosystem partners in operating systems, compilers, libraries and applications. We have great community support which is important so we can offer the end to end solutions for our customers.

Now let's quickly look at our performance. You look at the chart on the left, it shows our efforts of continuous performance improvement. You can see from release 2.0 to the software release 3.0, we almost doubled the performance running on the same power. And this demonstrates the maturity of our machine learning software stack, the frameworks, the compilers and libraries. Now you can see on the right-hand side of the chart it shows our multi-GPU performance scalability. We can achieve almost linear scalability from one GPU all the way to 16 GPUs. Again, it demonstrates the maturity of our communication library and the benefit of Infinity architecture. And this scalability is critical for exascale computing.

So, to pull everything together, only AMD can provide the combined advantages of CPU plus GPU plus the opensource software. We can provide the fully integrated CPU plus GPU systems and the unified tools to make it easy for people to develop their applications. And AMD's Infinity architecture connects many CPUs and many GPUS together through enhanced bandwidth and the cache coherency. Together with our opensource software, we can demonstrate our performance advantage comparing to the competitors' combinations as you can see in the chart. This combined, advantages of accelerated computing, really enables us to drive the high-performance computing leadership. And that's a key reason why we were chosen for the supercomputer design wins.

Okay, in summary, we have developed strong GPU technology roadmap. That's based on advanced process technology, the RDNA, the CDNA architecture, the aggressive per-watt efficiency enhancement, and our opensource software. With this strong technology roadmap, it puts us on the path to leadership and we are very, very focused executing to our plan, to our commitment. This will enable the next wave of winning products. This is the end of my presentation. Thank you very much.

And with that, I'll introduce Rick Bergman to tell us the exciting business and products for the PC in the gaming business. Thank you.

Rick Bergman:

Thank you very much, David. It's quite exciting to see your passion around GPUs. I thought I'd first start off reintroducing myself. I'm actually, like David, recently joined AMD after eight years as President and CEO of Synaptics. Prior to that, I had a 30-year semiconductor background starting out with Texas Instruments and then a decade at ATI and AMD. While at ATI, I led the business that grew our GPU share to number one in the marketplace for discrete graphics. While at AMD, I led the team that developed the first APU, or accelerated processing unit, the first time an X86 core in a GPU were combined together.

But that's all in the past and it's real exciting in its own way. But what I'm really excited about now is the opportunity in front of AMD and that's what I'm going to talk about over the next little bit. So why am I so excited about it? It's really three, comes down to three things. First, and you heard this from Mark, we created this execution machine on processors. This regular cadence of leadership products coming out. Second, I can now see it's clear the same thing is happening on GPUs. David talked about our RDNA or NAVI generation of products coming out in that similar fashion. And then third, wow, we get to combine these into incredible APUs. So it's just a really, really exciting, exciting time in AMD's 50-year history.

There's a strong path forward as we look forward to where we can take this. So to my role as head of the computing and graphics group, is how can we turn that into a sustainable growth in our market share, with our profitability, with the platforms moving forward? So there's 1.5 billion PC users, active PC users out there. What an exciting user base. And they have different needs. You have some consumers that are looking for the thin and light laptop computer. Of course, we have the enthusiasts who want the absolute bleeding edge technology. And then you have content creators that view their PC as a tool. And we've kind of rewritten the rules there going from 16 to 64 cores. But at the end of the day, AMD has to pull all that together. And there's no better company that's positioned better than AMD to pull that together into solutions.

And so what is the opportunity? \$32 billion opportunity in the PC market for us, so it's a

very sizable market. We started with the Ryzen desktop 5 and 7 and now have increased to have full coverage of TAM with Ryzen products. We're focused on executing across markets. So whether it's desktop, notebook, high end desktop, commercial or consumer, we have the opportunity to lead in all of those markets.

So as Lisa mentioned, we've disrupted the market. It's always so exciting that we have the fastest solution in three distinct segments of the marketplace. Of course, our Ryzen desktop using the chip wood approach, power 7nm, we have the clear leader in desktop. There's Threadripper for the high-end desktop as well, the world's first 64-bit core processor for high end -- 64-core desktop processor for high end, allowing content developers in some cases to decrease the amount of time for example rendering by 50%. And then in notebook, we launched the mobile Ryzen third generation product back at ZES, again, setting a new level of performance for the ultra-thin notebook products.

So how does that all add up? Let's look at the market share and unit shipments. We've almost doubled our unit shipments in just 2 years. As you'd expect, our market share has followed a similar trend. But in addition, our ASPs, our average selling prices, have also increased each year. So if you think about that, we've expanded our stack for Ryzen, increased our unit shipments, and then also increased our ASPs, meaning we're moving up, we're moving into the premium segments of the marketplace, going from 30% TAM to full TAM coverage.

So what is our path forward over the next five years? Well certainly we want our growth to be long term, sustainable and built across diverse segments. We've done a great job on the desktop side. We're going to aggressively pursue the notebooks and the commercial segments as well, it's untapped potential for AMD. For commercial, we're a natural mix, as Mark touched on. We have the security, technology, the reliability and the performance to be quite successful in that particular segment. And we're looking for multigenerational support, the CPU cores based on the Zen technology. Of course, we have the GPU cores as well. But what's really exciting, again, is pulling it all together into SOCs. And we have clearly the best SOC capability and products coming to market in 2020 and I'll touch on a few of those through the course of the presentation. And of course, we need software of course to pull it all together and deliver the full potential of the hardware.

So now let's touch on the desktop leadership again. So almost any way you look at it, AMD has leadership in the key markets in the desktop area. With our desktop Ryzen products, of course you can go out to the websites, see the reviews, see the prices that we're getting, see our market share at retailers around the world. Or you can go and look at the high-end desktop type of solutions as well where we do have those 64 cores, 3 times the performance of our competition. Virtually any workload, AMD is beating the competition. But kind of don't take our word for it. Go out to the respective websites or reviews and you'll see just review after review after review talking about how AMD has clear leadership in these segments.

So now, how do we transfer that desktop leadership into notebooks? So notebooks do represent 64% of the TAM out there and certainly performance is important in the notebook business. But we want to have the entire user experience as well. So we want our customers, when they get an AMD based notebook, to love that experience from the day they go home and open up that box all the way through the lifetime of that notebook. So in the case of productivity, we've doubled the multithread performance generation over generation. In the case of responsiveness, we've really worked on our drivers, and Mark talked quite a bit about Modern Standby. But if you look at year-over-year, we

have five times the number of platforms that support Modern Standby versus prior generation. And then of course battery life, we've moved all the way up to as much as 18 hours with our new Ryzen product.

We announced at ZES our third generation AMD Ryzen mobile processor to quite a bit of fanfare and tremendous excitement at the OEMs. Of course, it's the world's first 8-core X86. It's the world's first 7nm mobile processor. Tremendous graphics, tremendous battery life. The average lifetime of a notebook is about 4.8 years, so let's just call it five years. If you go back five years and look how far we've come, we've increased the compute power by six times. We've increased the graphics by three times and the battery life of a notebook is now three to four times longer. So just tremendous progress telling you where we're going to be five years from now as well.

This is really a watershed moment for AMD. Really, it's a big moment for us when you look at the performance. Of course, we've always led in graphics for a decade, that's just expected now that we'll have the best graphics solution in the industry. Since the Zen core, our multithread performance has also been leadership as well. But what now is new, single thread performance. AMD has taken over the leadership with single thread performance. So in the 3 major categories of performance, AMD is the leader.

So what does that mean? At some point the numbers start to speak. You can see we've more than doubled our market share in the notebook space. So keep in mind we've been able to accomplish that without 7nm. This is pre-third gen. Without 7nm, without Zen 2, we've been able to more than double our market share. And now you look at the number of platforms that we have, going from 50 to in 2020, 135, 176%. It's a leading indicator where this business is going. And it's with all the top PC OEMs as well. HP, Dell, Lenovo and so forth, selecting us for premium notebooks all the way through to entry level notebooks. Clearly great progress.

The other area I mentioned that is so important to us now is the commercial business. That represents 48% of the TAM and yet great opportunity to continue that momentum and growth. Of course, the decision criteria around commercial products is a little bit different. Performance still matters. You want your workers to have snappy performance of course. Security is absolutely critical and so I'll talk to that on the next slide. You do want that battery life to keep your employees going all day long no matter where they are. And then of course manageability for deployment, imaging and management within a modern IT infrastructure.

So let's talk a minute about what it requires for security. Security is just so important in today's world. We have a 3-level approach to security. Of course, the first level is at the processor level. And there, as Mark mentioned we do have that built-in dedicated security coprocessor to protect your PC. Then as you expand at the platform level, we have features like the full encryption, the memory guard, the full encryption of memory to protect your data while it's resident in your PC or while it's being transmitted either direction as well. And then you take it up to the next level, the operating system level, we work with Microsoft and the OEMs to ensure that we have the best security at the enterprise level as well. So we support things like Think Shield from Lenovo or Sure Start from HP.

So then now, how do the numbers look here? As you can see, once again, very strong share moving forward. But again, let's look at the leading indicator. How many design wins do we have? In this case, more than double the number of commercial platforms, going from 35 to 70 plus in 2020. And again, I'll remind you, this was without the Ryzen

Pro solutions. No 7nm, no Zen 2 cores. We have that momentum with the prior products moving forward. So really expect things to pick up also in this space, again, working with top commercial OEMs like HP and Lenovo. Lenovo already announced their intention to use the Ryzen Pro in a ThinkPad product, so there will be a full lineup of ThinkPad solutions leveraging AMD processors.

So the roadmap. Of course, a great deal of interest in the roadmap. Over the last 3 years, we've had three generations of Ryzen products. The Zen Core is so critical to the foundation for the success of our products in this area. And as you heard, by the end of this year we'll have the 4th generation of Ryzen with the Zen 3 core. So it's going to be a very exciting year for AMD across all the different platforms that we participate in.

So now I'm going to shift gears. So I just talked about the PC market, now I'll talk about the gaming market. 2.5 billion gamers out there in the world. So often you have different images of who a gamer is. In some cases it's a business person playing Candy Crush sitting at a gate waiting for their flight to take off. Or it could be the enthusiast, really again, pushing, wanting the latest and greatest. We love the enthusiasts. It could be a teenager playing Minecraft at home on a PC. Or perhaps it's a 20-year old playing SIMs, also on a PC. A range of environments, different areas, different performance levels, all possibilities for Radeon technology. So where is Radeon Technology deployed? Well it's everywhere. We have an installed base of over 500 million and it's clear we're on a path over the next several years to get over 1 billion Radeon users. But what's really remarkable about this slide is who is adopting Radeon technology. Of course, you have Google up there, you have Microsoft, there's Sony, there's Apple. Of course, the PC vendors as well. These are companies that do their homework. They go out and look who's got the best technologies out there. And invariably they're coming back. When you want the best, the world's best graphics IP, best graphics solutions for gaming, it's AMD.

And so the TAM in this business is also sizable, it's a \$12 billion market and it's growing. RDNA as you heard from David really forms the foundation for powering the next decade of gaming. It's going to enable a full stack of solutions from AMD. And of course, every gamer knows it's not just about the hardware, you also have to have great software as well. So our first solution, the NAVI, we call it our NAVI 1X solution, using RDNA, we introduced last year. It was a very ambitious introduction. Of course, a new ground-up architecture in a leading-edge process like 7nm, key features like GDDR6, and we also led the industry with PCIe 4.0. Very bold architecture.

We targeted the largest audience in the PC market. And it was really around 2 distinct segments. So first, the 1080p market, that resolution, that represents about 60% of the PC gamers. And then the 1440p market as well which is the majority of the balance of the gamers out there. Two different product lines, 5600 and the Radeon 5700.

Another thing to note here is the breadth of solutions. Just like I mentioned previously on Radeon, these are all the leading graphics add-in board vendors. And of course, it also happens that several of them also support us on the motherboard side as well. We have 3 times in this generation the number of SKUs that we had in the prior generation due to the strong pull and interest for our NAVI 1X products. And as further validation, the Radeon 5700 XG won the GPU of the Year from the PC Gamer. And so the importance of this at the end is, how well do you run games? And more importantly, modern games or recently released games? So this is about a half dozen of the really critical games out there looking at our performance. As you can see, it's quite good, which means very fluid game play and also shows actually we're a little better than the competition as well. Because we targeted the right area with the right architecture with their right solution.

And so I've mentioned software a few times, I just want to reemphasize the importance of software in this marketplace. We do put a great deal of emphasis on it because it really allows the performance and the features that come out of our solutions. So in December we released our latest software, we call it Adrenalin. It enabled a new graphics user interface making it easier for gamers to game. Brand new hot titles come out, we have to be there on day-0 with optimized drivers that work flawlessly. That is our commitment. And then there's features. So software gives us a touchpoint with the gaming community. And so we hear when they're looking for certain additional features or maybe we innovate and come up with features. And one of those examples was Radeon Image Sharpening which just allows a better visual experience on games.

And so now the shift. Well what's next of course is the big question. Of course, that will be our NAVI 2X family of products. We'll introduce those products at the end of this year. But clearly, we're stepping up and we're targeting the enthusiast class performance as well. Which means we have to have great performance at 4K resolution. Remember I talked about 1080 and 1440, the next big step is 4K resolution. That means uncompromised performance, very fluid play, and of course we have to have just stunning visuals as well. As David talked about the importance of ray tracing to show the shadows and reflections to be able to game at a very fluid, consistent rate. So we'll add additional features like variable rate shading as well to give you that performance uplift so critical in this segment of the market. And we'll take that 2X stack across over time across the entire top to bottom lineup for graphics.

So now we shift into what does our GPU roadmap look like? This probably looks very familiar to the CPU roadmap I just showed. Remember at the very beginning of my presentation I said we're getting that same clockwork, that same execution machine on the GPUs like we've had on the CPUs. So we're going to have cadence, that drumbeat. So yes, as I mentioned, NAVI 2X starting later this year and then continuing forward. The development team is off working on that next generation, creating those even better performance, better set of visuals. Because it demonstrates our commitment to be in this market. To not only be in it, but to actually win in this market as well.

So now let's talk another key segment of the gaming market which is the game consoles. 10-plus year relationship with Sony and Microsoft. As Lisa mentioned, having shipped over 150 million units of the current generation. But we're on the cusp of the next generation. So the next generation consoles will use our latest Zen technology, use our latest RDNA technology as well, really creating this very immersive experience with ray tracing, 3D audio, fast load times to really excite a new generation of gamers. So whether it's a Radeon mobile processor, NAVI 2X, or one of these next generation consoles, clearly 2020 is going to be a very, very exciting year for games.

So now to wrap it all up, you've heard about our great processors, you've heard about our great graphics, incredible SOC capability. We'll have more APUs this year than ever in the history of AMD. Gives us the opportunity to have leadership across all segments in PC and gaming. Remember, combined, that's a \$44 billion market and we have the opportunity to lead in every segment. As you walk the hallways here in AMD or any site worldwide, it's absolutely clear that our employees are maniacally focused on executing to the roadmaps that I showed today. There's really no reason at all that we can't seize the opportunity for sustained growth and success over the next decade at AMD. Thank you very much. So thank you and now we will take a brief break and we'll come back with Forrest talking about our datacenter solutions.

[break]

Forrest Norrod:

Good afternoon. Welcome back. I'm incredibly pleased to be here to share with you an update on the progress we've made so far in the data center and the future that we see ahead of us. And for us at AMD, when we think about our mission to deliver high-performance computing, there is perhaps no market for which that is more important than the data center market. Because the data center market of today is a market that is seeing continuous innovation and continuous disruption with new applications that weren't even dreamed of a few years ago springing into fruition all the time. With machine intelligence, of course, with the advent of webscale applications, with artificial intelligence, it is truly an era in the data center of continuous change and an endless need for compute power, an endless need for high-performance computing, which is a perfect match for our mission.

And it's also a very large and interesting market. As Lisa said earlier, it's a \$35 billion market in 2023 with a number of rapidly growing segments that we believe are a perfect fit for our technology and our direction. And so I want to talk to you a little bit about how we're going to continue to attack this market, but let me first reflect on where we began this journey; this journey back to leadership in the data center. We embarked on this several years ago. We actually talked about this and our ambition to bring innovation back to the data center at our Financial Analyst Day back in 2015.

The first major milestone on this journey was when we introduced the first-generation AMD EPYC processor in 2017, which truly was the first competitive X86 server CPU to challenge our competitor in quite some time. And when we introduced that processor, that first EPYC processor, codename Naples, the processor that with its incorporation of the Zen core began our tour of Italy and a return to the data center, we put up a roadmap showing what we were going to deliver from 2017 all the way out to 2020.

One of the reasons that we did that is because we knew that to be considered for the data center, much less to be a leader, we had to be not just a provider of high-performance components; we had to be a reliable partner. Somebody that end customers could count on to be there and to always be there with each generation, new products, new innovations, maintaining the value of the investment that they would have to put into any new entrant into the data center.

And so we put this roadmap up. And I'm incredibly pleased and proud to say that the teams have delivered with metronomic regularity on this roadmap with Naples in 2017, Rome of course last year, and Milan on track to ship this year, as Lisa mentioned earlier. So the execution is there, and the execution with the second-generation EPYC has delivered an amazing part.

This is the highest performing X86 processor ever. And it's not close. When we introduced Rome, it was twice the performance of the competitive X86 processor, which enables our customers to in many cases drive 25% to 50% lower total cost of ownership for providing the data center services to their internal and external customers, which allows them to transform their data centers, their IT, their operations. And as Mark mentioned earlier, the Zen 2 core at the heart of Rome also incorporates not just high performance, but the high security features that are so critical to secure the data center, to secure our customers' valuable IT assets. And so Rome has been a revolution in the data

center.

And on the performance side, I'm incredibly proud to say that we have garnered over 140 world records so far. We're the ultimate in performance at a wide range of workloads from enterprise IT, through high-performance computing, through cloud infrastructure, Rome delivers superior performance, delivers world-record performance.

If you take a closer look at that and you look at the processors that are available for the vast majority of service today, and that means 2-socket servers, and you look at what we've delivered with Rome versus what the competitor has to offer today, including, including the processor refreshes that they just introduced a few weeks ago -- or few days ago -- we clearly have incredible performance lead; double the performance of the competitive offerings inclusive of their most recent refreshes in the 2-socket market. And the same story in 1-socket where the unique value of Rome, and indeed Naples before that, was to first establish the no compromise 1-socket market, where by not placing any artificial limitations on the availability or features of a single socket processor, we have enabled customers who need all the reliability and the resiliency of an enterprise-grade server processor to address -- to scale their processor needs to fit their application, and hence, drive yet further TCO value.

Now, Rick talked earlier about the various aspects of performance in the client market, and indeed, these performance metrics that I've shown you so far demonstrate the throughput superiority, the multi-threaded superiority of AMD EPYC over the competitive parts. And that holds true even when you look at lower core counts and can compare core to core. So in the heart of the market where most enterprises buy in the 16-core processor segment, if you take a look at our performance there, again, inclusive of the two brand new parts that our competitor just introduced, AMD EPYC has a performance lead and an unmatched performance-per-dollar advantage. So almost 3.5 times the performance per dollar of the most recently introduced 16-core part from our competitor; a truly remarkable achievement.

So with that, I want to reflect on where we stand today. So today, we're in a place where we have demonstrated repeatedly predictable execution to the market, quite frankly, highlighted and thrown into high relief by uneven execution by others in this segment. We've demonstrated leadership performance on multiple dimensions, and we've demonstrated strong ecosystem support where we have the OEMs, the software vendors and the IHVs that provide the rest of the ecosystem to compose an entire solution all-embracing EPYC.

And so that's a great foundation. That's a great place to be. And so we're taking that foundation, and our imperative, our mission now for the entire team is to continue accelerating growth of the AMD EPYC product portfolio. And so our mission is to broaden the deployment, sort of already begun first with Naples and then more rapidly with Rome, across enterprise, cloud and HPC. It's to work with our customers to continuously unlock all of the power and performance of our solutions by working with them to optimize their workloads and tune their software to get every last drop of performance out.

It's to continue ramping our field and customer support organizations as our customer set grows and to support optimization and deployments. And then that's all in service of course of growing our market share. Now on the market share, the next significant milestone is one that we are going to hit next quarter when, as promised, we believe we will achieve double digit server CPU market share in Q2 of this year.

Now, when we look at what we've got in more detail across each one of these segments, we have built a product, we built a portfolio that really can demonstrate leadership performance in the workloads that matter. And the markets that matter, broadly speaking, are these three: the cloud market, which now constitutes, quite frankly, about half of the overall market for server processors; the enterprise IT market, which is a little north of 35%; and then of course the HPC market, which constitutes most of the remainder.

With the first-generation EPYC, taking that first-generation Zen core, we were able to produce a product, Naples, that effectively addressed about 60% of the workloads across those broad markets. The vast majority of the HPC workloads, much of the cloud workloads and a good proportion of the enterprise IT workloads. And for them we had leadership performance and a great solution.

With Rome and with the work that we've done with our ecosystem partners and end customers, we've expanded that footprint substantially. And we see over 80% of the workloads across these segments are addressed best -- are addressed in a leadership fashion with second-generation EPYC Rome.

And so to give you a few examples in each one of those segments, on enterprise, we have a great performance story. We had with Naples superior virtualization capability and superior TCO. With Rome we've taken that to the next level; up to 50% TCO savings over the alternative. But beyond that, in many other application areas, in Java performance, in database performance we're providing nearly half -- or sorry, nearly twice the performance of the competitive solutions. Even on things such as SAP, we've demonstrated world records and we've demonstrated demonstrably superior performance.

You couple that with the embrace of the OEM partners who have greatly expanded their portfolio of solutions, their portfolio of platforms that are made available to the end customers, growing from 22 platforms available in 2017 with the first-generation EPYC processors to over 140 expected this year. We've got a winning solution for enterprise IT, and we've got that broad support across the enterprise ecosystem both on the OEM side, on the software side and the IHV side. Our mission: keep driving, keep driving.

On the cloud computing side, it has been incredibly gratifying to watch the growth here as well. In 2018, our first instances were launched with some of the major cloud providers. We had 18 public instances available, public instances or services available on first-generation EPYC systems. And now this year we expect over 150 public services and instances available so that customers around the world who have embraced the cloud computing paradigm can do so now with AMD EPYC and see the full performance of Rome.

In order to support that, we have developed and deployed customized versions of the EPYC processors, uniquely tweaking in many cases to ensure that the environments and the TCO is maximized for our cloud customers. And they in turn are using the fact that we can support 60% more virtual machines with each being at the same performance level or better, that we can support better Docker performance, that we have much better performance per watt, which is critically important for any cloud provider operating at scale. And then we've got the memory bandwidth to unlock all of that performance and feed the beast, as Mark wanted to say, to keep all of those cores humming and to keep that performance available to end users. With those advantages, it's no wonder that the world's leading clouds in the U.S., in Asia and in Europe have embraced AMD EPYC.

A few examples. Google's largest general purpose -- actually largest instance of any type is available on second-generation Rome. We've delivered more than 25% better TCO for Twitter. And Microsoft has a great set of instances in both general purpose, as well as purpose-specific instance types, including examples that are demonstrating over 60% higher SQL performance for companies that are embracing a move to the cloud. So a great place for us so far and a great place for us to continue to grow with our cloud customers.

On high-performance computing, of course, this is a real area of pride for us at AMD, as Lisa mentioned earlier. We have a part and a roadmap that has demonstrated leadership in HPC applications by virtue of a substantial lead in floating point performance, and demonstrated leads in performance on commercial applications with up to 72% higher structural analysis, 95% higher computational fluid analysis than that available on our competitor. Very importantly as well, and one of the things that's led to a number of major supercomputing design wins with the large meteorological research associations and institutions around the world, 120% faster weather forecasting performance.

And with the Frontier and El Capitan announcements, it is a real source of pride for the AMD team to be powering the world's fastest exascale systems and to be contributing to supercomputing and HPC clusters around the world from industrial applications to research to defense applications as well.

So a new area for us, a new area of focus has really been to expand our reach beyond the traditional areas of cloud, enterprise and HPC and to dive deeper into the realm of the telco and infrastructure products that they themselves are undergoing a revolution. Many of these systems that are powering today's networks have moved or are moving from being based on proprietary hardware, on proprietary ASICs, to running on software on industry standard servers. And the thing that's enabling that is the incredible power of today's servers, and that in turn gives those customers who have embraced that paradigm the opportunity to add significant agility at significantly better cost to their businesses.

EPYC, particularly the second-generation EPYC is uniquely capable in this area with much higher, double the I/O bandwidth per link with much higher bandwidth capability in terms of more links, more memory bandwidth. We demonstrate leadership, networking data plan performance and truly are the perfect fit for the emerging 5G telco infrastructure.

Nokia, one of the leaders in telco infrastructure, recently demonstrated why they have embraced AMD Rome, because they see twice the performance on their 5G infrastructure based on Rome than they see using a competitive -- our competitive CPU. That's the type of performance that Rome provides. And so we've got a great part, we've got a great start in all of those markets, and we're just going to keep going.

I'm very pleased to say that the third-generation AMD EPYC part, codename Milan, we expect to continue demonstrating performance, leadership in both throughput and low thread count applications. And we expect to be shipping it later this year as promised. And with Milan, we will open up the aperture even further. I talked to you about with Rome, we've got better than 80% coverage on the workloads across HPC, enterprise and cloud computing. With Milan and the work that we're doing with our ecosystem partners, we believe we'll have leadership performance in virtually 100% of the workloads that power today's data center, which is incredibly exciting and we believe continues our progress in every segment.

But our tour of Italy doesn't end there. So Milan -- Rome is an incredible part. Milan will be fantastic. But Mark and the team are completely committed to that CPU leadership that he talked about before, and the Zen 4 core will show up in our next next-generation server part, codename Genoa, that you'll see in 5nm and that we announced yesterday will be the CPU that's at the heart of the El Capitan system. So, the journey will continue. The leadership products in server will continue.

But beyond server, the other market that we've got to talk about is GPU. Data center GPU for us is a rapidly growing market that we believe by 2023 will constitute an \$11 billion market opportunity for us. And in data center GPU, there's really a number of different application areas or segments that where it is crucial, it is vital. First one in virtualization and cloud gaming. Many customers are making the move to do virtual desktop, virtual gaming in the cloud, unlocking a new set of economics, unlocking a new set of collaboration, unlocking a new set or new points for TCO optimization.

Beyond virtualization and cloud gaming, of course the most actively discussed application of GPUs has got to be machine intelligence. And we are committed to driving our roadmap aggressively for machine intelligence. And then of course high-performance computing where the top end supercomputers have been the province of GPU-accelerated computing for the last decade.

Our start in this market has already been made. We have the world's first 7nm GPUs for the data center. We are the only company to implement industry standard hardware-based virtualization that allows the resources of GPU to be scaled up or down, depending on the needs of a specific application, and that's critical for almost any application in the cloud. We're taking the same approach to developing a multi-generational roadmap -- not developing. Not just developing; deploying, delivering a multi-generational roadmap for GPUs that we took on the CPU side.

And when we look at the innovation that's happening in each one of these areas, we think that innovation is best fueled by an open software environment, by an open software community. And AMD has made a strong commitment to having completely open source tool chains from drivers to libraries to enable each one of these segments.

In GPU virtualization, it was incredibly gratifying to us earlier this week to have Microsoft Azure introduce its latest virtualization instance, the NVv4, which provides accelerated VDI solutions via the cloud. And with, by virtue of our hardware virtualization technology, scales up and down and delivers superior economics than any other alternative. That same technology is also critically important to providing the scalability and the long-term economics that are so important to cloud gaming to ensure that that segment grows as well.

In machine intelligence, look; the applications of MI are just incredible. Anybody that's -- you all have seen it. The revolution in natural language processing, in image recognition, in recommendation engines, and even in industrial automation and robotics has been incredible. The types of capabilities that we're seeing now in these systems are things that people would not have dreamed of a decade ago.

And we're seeing applications spring up, applications for artificial intelligence spring up everywhere. One of the interesting ones that was discussed the other day is using machine intelligence in concert with HPC to steer scientific research and traditional HPC computing and by doing so to get what was characterized as a force multiplier deployed.

So, we believe that machine intelligence applications are going to continue to grow, but what's going to grow even faster is the insatiable need for performance for these applications. Because as researchers refine their algorithms, as they explore the limits of what machine intelligence can do, they're driving an exponentially increasing demand for performance by virtue of larger models, longer run times, more data going into all of these models.

And that looks a lot, quite frankly, like HPC, demands as well. Mark showed earlier a portion of this data on HPC, but just to put a finer point on it, we have seen on the supercomputing side a 10,000x increase in performance over the last 15 years. And if you look at just this year to the exascale era, we need to do another 10x increase in performance to meet that exascale barrier, to break that exascale barrier. Similar things happening on machine compute.

You look at the model sizes of some of the most interesting work that's being done out there where you're seeing the most interesting results, model sizes have grown 300x in the last 3 years. And so for both of these segments, we see this unending, to use Mark's word, relentless demand for increased performance, and the only way to meet those performance needs is with accelerated solutions. And so our roadmap is committed to providing that performance for those applications. And David showed this earlier. We're already shipping our first 7nm GPU for the data center, our MI50 product based on the GCN. Later this year, we'll be introducing our first CDNA product also based on 7nm and optimized for HPC and MI applications to be a highly efficient, scalable product to meet those needs. CDNA2 will be coming soon to usher in the exascale era and to provide that 10x uplift in performance at the system level.

And it's not just about performance. It's about making sure that that performance is usable. And this was touched on earlier, but I want to turn up the contrast, because look; accelerated systems have now been around for about a decade. But what we have done is we have bolted on accelerators onto a server system architecture that quite frankly was created for webscale applications and databases. And so the CPUs and the GPUs are isolated from one another. They don't work well together. And although the performance is there, it's difficult to reach. It's difficult to program effectively applications to fully unlock the performance in a system with this topology.

The first step in addressing this we're taking with our CDN architecture to allow better scalability amongst the GPUs and coherency amongst the GPUs to allow up to 8 GPUs to scale more efficiently and work more efficiently with one another. But with the CDNA2 architecture, we get something truly special where we extend that Infinity Architecture, in this case Infinity Architecture 3, to couple the CPU and the GPUs together into one unified data view. Where this not only provides additional performance, much more importantly it allows programmers to stop worrying about the explicit management of data movement, explicit management of pre-emption, the explicit management of a bunch of functions that they haven't had to worry about on CPUs for many, many years and that are acting as barriers to embracing accelerated computing. So we're super excited by this unified accelerated computing architecture that we're bringing to bear in CDNA2.

And it's that architecture, quite frankly, that led to the exascale systems that Lisa and Mark both mentioned earlier. First with Frontier, which will be a 1.5 exaFLOP system deployed next year, and it is capable of outperforming the top 100 supercomputers on today's supercomputer list combined. It'll be deployed in the middle of next year, supported by AMD CPUs and AMD GPUs connected together in a coherent way.

The only thing that's more exciting than Frontier and what can be done with Frontier is of course El Capitan. Announced yesterday in an event with Lawrence Livermore National Laboratory and HPE, it will produce over 2 exaFLOPs of performance, expected to be more powerful than the top 200 supercomputers on today's list. And it'll be supported by AMD CPUs and GPUs shipping in 2022.

Those two systems to us are a source of tremendous pride and a validation to our assertion that we're making tremendous progress on our goal to be the new data center leader is coming true. And I think that -- I hope you see, and I think that the industry is seeing, that AMD has produced leadership products with leadership execution and we have a leadership roadmap that's going to continue on to the future. Delivering the best performance available on all sorts of applications across a wide range of workloads and leading the way to the future of accelerated computing by defining the architecture of tomorrow's accelerated system. And so I think with that, I think we are well on our way to demonstrating to any that question, we are the new data center leader.

And with that, I'd like to turn the stage over to my good friend Devinder.

Devinder Kumar:

Thank you, Forrest. You know, it's been 3 years since we had our last Financial Analyst Day in 2017, and I can tell you it's been quite an exciting journey. You've heard from Lisa and my colleagues about our plans to accelerate momentum, execute product and technology roadmaps and the long-term growth that's ahead of us. I will share with you the financials since 2017 in terms of our progress of what we have done since 2017, and then the next phase of our financial journey, our long-term financial model and our capital allocation strategy. So let's get started.

Here is exactly the priorities we laid out in 2017: grow revenue from the base that Lisa showed you earlier, just about over \$4 billion; gross margin expansion; operating expense discipline to return to profitability. And from a standpoint of what we have done, we have focused on the financials to grow our revenue, expand the gross margin steadily over the last few years, exhibited the OpEx discipline to continue to increase our profitability over the last few years. Let's take a look at the details.

On revenue, we've made great progress. We've had increased revenue of \$2.5 billion from where we were in 2016 to 2019; 56% growth in revenue from 2016 to 2019. It's been driven by the new leadership products that you heard my colleagues talk about: the AMD Ryzen, the AMD Radeon and the AMD EPYC products that have been introduced since that time.

Let's turn to margins and OpEx. If you look at the left, gross margin coming off of 2016, 31%. And we've showed steady increase, accelerating gross margin improvement, 12 percentage point improvement in gross margin in just 3 years. And since the Financial Analyst Day in 2017 when we met, every quarter since then we've had year-over-year increase in gross margin. Every quarter since we met the last time, gross margin has gone up when you look back at year-over-year gross margin expansion.

We made the OpEx investments needed for R&D and in particular what you heard a lot about today, the multi-generation roadmap, whether you go from first generation or second generation to third and now onto the fourth generation of products. In 2017 and 2018, our focus was R&D. In 2019, in addition to the investments in R&D, we accelerated our investments in go to market. And we had lot of go-to-market activities in 2019, and that's why you see the \$2.1 billion OpEx in 2019.

With that, if you take revenue, margin and OpEx, let's look at growth in profitability. On the left -- actually I love this chart. First of all, when you look at the operating margin coming off of 1% in 2016, we've had 11-point increase in operating margins in the 3 years from 1% to 12%. And then on the right, EPS. We've made very good progress on the bottom line. Earnings per share has gone up. We've had growth in EPS for the last 3 years, and with that the P&L has gotten better.

But let's turn our attention and look at the balance sheet, which I know a lot of you are asking about when we met in 2017 about AMD's balance sheet. If you look on the left of this chart, \$1.8 billion of debt has gone down by \$1.2 billion in the 3-year period. It's less than \$600 million as we ended 2019. We've ended 2019 with \$1.5 billion of cash, which turns net cash positive for the first time in a long time to be about \$900 million as we ended 2019.

The gross leverage target that we set in 2017 was to get under 2x. We've exceeded that target. The gross leverage is 0.5x in particular because the EBITDA of the company as we ended 2019 on a trailing 12-month basis was more than \$1 billion. And that led to a very strong foundation of a balance sheet as we go on with the investments needed for the next few years in terms of everything you've heard from my colleagues this afternoon.

On the next slide, let me summarize what we said in 2017. 2020 long-term target model is exactly what we said in 2017 -- double digit growth in revenue. Double digit growth in revenue, and as we end 2019, we've had about 16% compound annual growth rate in revenue. Gross margin, we had said we'll get from the low 30s to 40% to 44%. We ended 2019 at the upper end, at the upper end of that range that we laid out in 2017.

EPS, strong growth every year. Strong growth every year. We've had solid momentum, solid financial momentum for the last 3 years, which brings us to today, which brings us today. And let me show you the priorities from a financial standpoint for our next 4 years.

You've heard a lot today from my colleagues. Mark and David shared with you the multi-generation CPU and GPU technology roadmap. You heard about architecture. You heard about a product roadmap from Rick and Forrest and their plans to grow market share and grow revenue in many different areas. All of this sets us up for continued success. All of this sets us up for continued success, and for the next 4 years it's all about growth. It's all about growth. You saw Lisa put out the chart earlier about 20% compound annual growth rate on revenue for the next 4 years. It's about growth.

Now with growth, we also want to focus on continued margin expansion. Continued margin expansion to go further from where we are today. And that will continue to be a focus. We want to further increase operating margin and increase profitability. And finally, we want to generate significant amount of cash. It really sets up for a very exciting 4 years, given where we have come from the last 3 years in terms of everything we have done.

But it starts -- it starts, as you've heard from all of us, with our market opportunities. We are playing in markets today where the TAM is \$79 billion. You've heard about the data center from all of us. \$35 billion TAM in the 2023 timeframe. PCs, large market, \$32 billion TAM. And then the gaming, which is a combination of consumer graphics and our game console business, \$12 billion TAM. The opportunities ahead of us are pretty large. The opportunities ahead of us are pretty large. We execute and we grow the revenue, we can improve our financials in a significant manner.

Let me show you the long-term target model. Revenue growth. 20% compound annual growth rate, and a lot of it from increasing market share in the markets we already play in with the products you heard from today. Gross margin greater than 50%. Higher ASPs, higher margin products, premium products that we're introducing in premium markets driving the gross margin higher.

We invest in OpEx at 26% to 27% of revenue, prioritizing R&D and also go-to-market activities. And we expect to double the operating margin. You just saw me show you the 2019 operating margin at 12%, and with the mid-20s percent in this time frame, we doubled the operating margins from where we were in 2019. And we expect to generate free cash flow margins greater than 15% and generate significant cash over the long term. Significant cash. This long-term model projects a very exciting period for us from a financial standpoint for the next 4 years.

Let me show you a little bit on the revenue mix over the next 4 years. In 2019, we ended with \$6.7 billion of revenue. And you've heard us say the data center revenue, CPU and GPU combined, is about 15% of the revenue. With the compound annual growth rate that we have for the company of 20%, the overall revenue of the company gets bigger.

The PCs and gaming part of it grows mid-teen percentage, but we expect that in the long-term time frame in the long-term model, data center becomes greater than 30% of the revenue in the 2023 timeframe. The data center, everything you heard Forrest talk about, driving the revenue to be greater than 30% of our revenue which is overall higher from where we were in 2019.

In addition to revenue, we are also continuing to focus on gross margin, margin expansion. So let's talk about that. Where does the margin expansion come from? We have 2019 43% gross margin. The high-end gaming is slightly accretive, and the margins will improve as we build out our gaming portfolio. The margins improve as we build out our gaming portfolio. PC products are above corporate average and a significant contributor to margin growth in the long-term model.

And then data center margins, which are well above corporate average, are the largest contributor with the growth in the data center revenue becoming 30%, or greater than 30% of the overall revenue. In combination, this drives our gross margin higher, and we would like to do for the next 4 years exactly what we have done for the last 3 years.

The tax rate, let me just cover that. We expect that sometime during the long-term target model, given our consistent profitability, the tax rate on a non-GAAP basis will move up to approximately 15%. The long-term cash tax rate stays at about 3%, similar to what we had in 2019 and what we are guiding to in 2020, fundamentally due to the fact that we have about \$6.7 billion of net operating losses that carry forward and allow us to pay at a lower tax rate, even though the non-GAAP tax rate is 15%. And those NOLs protect the approximately 3% cash tax rate through the period of the long-term model.

Let me move on to the capital allocation strategy. OpEx investment comes first. As you heard me say earlier, we will continue to invest in R&D and go-to-market acceleration. We think about many things beyond that. We have to fund the revenue growth in terms of the working capital needed. And then on the shareholder return side of it, we will consider shareholder return vehicles, including limiting share dilution and strategic initiatives. And lastly, building on the credit rating progress that we've made over the last few years, we want to have a goal of achieving investment-grade rating from a rating standpoint.

In a nutshell, our priorities are: invest in the business; drive growth; and deliver shareholder returns. Invest in the business, drive growth and return to the shareholders.

So from an overall standpoint, in addition to the business momentum, the product momentum, the revenue momentum, here is a summary of the financial momentum. We have delivered great products and established great financial momentum. We have been laser focused on execution and on all the roadmaps you heard about today. We have significant opportunities ahead of us with the almost \$80 billion TAM.

We are still in the early innings of market share growth across PCs, gaming and data center. And that's where you see the ongoing market share gains as one of the drivers of financial momentum. We want to accelerate the financial momentum, continuing to drive gross margin expansion, continuing to increase profitability and generate significant cash in the time frame of the long-term target model. And finally, we want to deliver strong returns to our shareholders. Thank you very much.

So it's a pleasure to invite Lisa back on stage for closing remarks before the Q&A.

Lisa Su: Thank you, Devinder.

Devinder Kumar: Thank you.

Lisa Su: All right. How's everyone doing? So look, I hope you've enjoyed the last couple of hours and got a feel for the excitement that we have in the products and the technology and the business. And I'm just going to spend just a couple minutes and just summarize in a few key takeaways.

Hopefully it's very, very clear: we are committed to leadership in high-performance computing. And that's across data center, that's across PCs and that's across gaming. And I hope it's also really clear that we are assuming the competition is going to be very, very strong. We have big competitors, and we respect them a lot. But at the end of the day, we're playing our game. And we know that if we execute our roadmaps, that we will see the growth that is exciting us. And that growth is in a set of great markets that we are under-represented in today. And so we come back to our ambitious view of delivering the best, and that's the best in technology and also best in class in terms of overall growth. So those are really the key takeaways.

Now, I think we're going to turn it over to Q&A. I'm sure there are some questions. So let me have the team come back up, and I think we'll reset the stage for Q&A. Just give us a couple of minutes here.

Okay. So look, before begin the Q&A, we spent the entire afternoon really focused on the long term because that's what this is about. It's about our strategy and our roadmaps. But I do want to address probably a topic that's on the minds of many of you, which is a little bit about what's happening in the short term.

Obviously, there's a lot of volatility in the markets with the coronavirus, and we want to make some comments about that as well. Our first priority of course, like all of our colleagues, is to ensure the health and safety of our employees and our partners and our customers, and so that is our focus. And we have taken steps to minimize potential exposure at our global sites and with travel like most of our peer companies have.

From a business standpoint, it is a very dynamic situation, so let me give you some color to kind of give you a view of what's going on. From an overall supply chain standpoint, our supply chain, our supply chain is primarily focused in China, Malaysia, as well as Taiwan, and I would say it's a very robust supply chain. So we have taken a number of actions to ensure that we have continuity in that supply chain. And based on what we see today, we're actually back to near normal supply capacity in our supply chain. So that is something that we continue to be very focused on.

We're also monitoring our customers since a lot of our customers have supply chains that are very dependent on China and some of those operations. And we did see some disruptions certainly through Chinese New Year and in the month of February. There's a lot of progress being made. I would say all of us in the ecosystem are trying to return those operations to as normal as possible, and we expect that to continue over the next couple of weeks. I'm sorry, over the next coming weeks.

Now let me turn to the demand standpoint. I think from the demand standpoint, again, this is a very fluid situation, so there are lots of puts and takes. What we have seen is outside of China, the overall demand has actually been about what we expected for the first quarter. In China, we have seen some reduction in consumer demand, particularly in the offline channel networks, and those I think will continue for some time. We have also seen some other puts and takes where the demand for infrastructure has increased beyond what we had expected originally.

And so with all of that, we had guided the first quarter at our first quarter earnings call at \$1.8 billion plus or minus \$50 million. We are not updating that as of now. Our best visibility is that the impact in the first quarter will be modest, but we'll keep watching that. And perhaps we'll be in the lower half of the range, but still within the range of our original guidance. You also saw from Devinder that for the rest of 2020, we are standing where our first quarter -- our 2020 guidance remains unchanged. And we see a very exciting growth path over the 2020 year.

So with that, let me turn it over to questions from the audience.

Ruth Cotter: Great. Thank you, Lisa. So we have microphones in the room. We have Laura here in the middle and [Saski] and Jason. So if you wouldn't mind putting your hand up if you have a question and we'll take questions in the room. Aaron?

Aaron Rakers: Yes, thank you. Aaron Rakers with Wells Fargo. I guess I want to unpack the model a little bit more. One of the things that I'm a bit surprised by is that you've got this big push in the GPU side, particularly the data center side. How do we as kind of analysts start to think about modeling that out? Have you thought about separating that out from a segmentation perspective? And what's embedded in your expectations as well with regard to the margin profile of the gaming SoCs as Microsoft and Sony come on late this year into 2021?

Devinder Kumar: Yes, I think on the segment piece of it, we report the segments as we do with the CG and EESC. We have provided additional color where we feel it's helpful, especially on the data center piece of it which is a combination of CPUs and GPUs. If you look at what I presented, we talked about PC gaming, in totality, mid-teens growth over the time frame of the long-term model.

And then on the data center side of it, it's higher growth and that's why it's becoming 30% of revenue on an overall standpoint. And if you look at the numbers and do the math, I

think you can get to the numbers in terms of how much is in data center, CPUs and GPUs. CPUs generally, if you look at it from a viewpoint of the market, is essentially flat. GPUs is where the growth is on the data center side of it. And that's where I would leave it.

Ruth Cotter: Ross?

Ross Seymore: Thanks. Ross Seymore from Deutsche Bank. I wanted to stick on that data center side of things, and maybe Devinder what you just said, a little clarification. Going from the 15% of sales to 30% of sales, a little color on how you think that growth will be driven between the GPU and the CPU side, the server CPU.

And then maybe a related follow-on is on the market share within servers as a whole, I think Forrest you talked about getting to the double digits in the second quarter, hitting your target there. What sort of target should we think of as being next? And is the market size itself the full server market, or are you still kind of judging the market as two thirds of what people describe it as in its entirety? Thank you.

Lisa Su: Yes. So, lots of different questions there, Ross. Let me try to take a couple of them and then maybe Forrest will respond as well.

Look, when we think about the model, let me just take a step back and say that we are designing a model that has a number of different outcomes can get there. And so that is obviously a lot of growth in data center, but also a lot of growth on the PC and gaming side. Within data center, clearly from a dollar value growth rate, a sort of growth number, the data center CPUs will be the larger number. From an overall growth rate, given we're starting from a low base on the GPU side, the growth rate will be higher, and as well as the market growth rate is higher.

Relative to market share goals, I think our view is that our product portfolio, whether you're talking about the data center or PCs or gaming, really supports a very strong market share over the next number of years. We're not putting a new market share target, but what I would say is that we believe that our product portfolio is strong and can certainly meet our previous market share capability, as well as beyond that, across a number of years. Forrest, did I touch --

Forrest Norrod: No, I think you hit it well, Lisa. I think on that last point, we're certainly not done when we hit double digit market share. Our imperative is to continue to grow. And as Lisa mentioned, our ambition is not to stop anytime soon. Keep it rolling.

Ruth Cotter: Laura, Tony here in the front row -- front.

Unidentified Participant: My question is on the supercomputer wins, specifically El Capitan and Oak Ridge. The -- I thought was interesting is you won both the CPU and the GPU at El Capitan. And this is a 3 year out, so they're making decisions 3 years out on obviously a forward roadmap at 5 and I'm presuming CDNA. So, what was the reason why they chose you? I can understand the CPU side, and the GPU was a little surprising. Is that because of this GPU/CPU interconnect and that that was what -- that was the trigger to win the combined both sockets? And then does this -- is this a precursor to hyperscale design wins, and will that translate to hyperscale CPU/GPU combined wins? So that's my question.

Forrest Norrod: Maybe I'll take a first whack at that and maybe ask David and Mark as well to weigh in there. Look, I think it was immensely gratifying to be chosen as the CPU and GPU

provider for both of those. One of the things that is -- makes it most gratifying is these are extremely rigorous evaluations that are done. And so they look at the technical characteristics of the proposed solution. They also, quite frankly, look at your execution capability and track record quite closely, and I think because you are procuring things that have not yet finished design.

And so I think David and the team did a marvelous job on the GPU side, developing an architecture that we think can scale up quite a bit, and maybe he'll talk about that in just a second. But then the ability to put them together. The ability that we talked about throughout Mark, David's and my presentation to be able to have a unified accelerated computing architecture with the CPU and the GPUs working seamlessly together and dramatically simplifying the programming model with something that I think was tremendously attractive to DOE and to the National Labs.

David Wang: Yes, I think we didn't show this perf-per-watt progression on the CDNA side, but you can imagine whatever we have done and what we're planning to do on the RDNA 1, 2, 3, that trajectory will be happening on the CDNA side as well. So we are aggressively enhancing the performance per watt on the type of operations that the data center and the HPC customers care about. And we'll be scaling our technology very aggressively through the CDNA architecture migration. So I will say besides the obviously the Infinity Architecture that provides all the benefits of cache coherency, programmability, unified memory, I think the GPU performance on its own it also a key factor of winning the deal.

Mark Papermaster: And lastly, I would add we have really listened carefully to the requirements from the Department of Energy and their goal to have this system be optimized for both HPC and MI. And listened to them and presenting to them how our software stack could really -- and an open source approach optimized across a CPU and GPU and be optimized with the libraries we would provide and then enable them to enhance it even further was also a key element.

Forrest Norrod: The one thing I would add to that is that to add on the back to answer the second part of your question was we do think that that architecture is the right architecture for machine intelligence as well. That necessarily when you look at those demands, we're going to have to move to clusters of accelerated systems, and this is the right architecture for those large scale machine intelligence applications as well.

Ruth Cotter: Great. Saski, at the very back of the room.

Trip Chowdhry: Thank you. Trip Chowdhry with Global Equities Research. A phenomenal presentation, lot of learning. I was wondering, if you look at two industries -- the semiconductor industry, the chip industry and the software industry -- for the semiconductor industry, the catalyst in data centers is machine learning models are getting bigger. But when you look into the software industry, they are going with distributed training, transfer learning, and their motto is extract as much power from your existing CPUs and GPUs and delaying the purchase. How do you see this evolving? Would it be a growth linear or would it be a step function? Thank you.

Forrest Norrod: You want to answer that one?

Mark Papermaster: Well, I'll just start from a workload standpoint. I think you're going to see the algorithms are still changing, and so this is dynamic. It's one of the reasons that you saw across our presentations that our whole approach is scalability. Our CPU roadmap runs unabated. Our GPU roadmap has a relentless unabated growth as well. And then it's how you put it

together. So there's no question that certain workload applications will remain GPU only, will remain CPU only.

But what you're seeing in these supercomputer applications is that -- it's the leading indicator of where the industry is going. Often, if you go back through history, the problems that first needed supercomputer class, think about the old supercomputers, you're actually performing those operations on your phone today. So it is indeed, from my standpoint and what I'm hearing from CTO peers in the industry, it is indeed the leading indicator of the approaches being used on many of these analytic and machine learning workloads.

Ruth Cotter: Thank you, Mark. Jason, Tim here. Thank you.

Timothy Arcuri: Tim Arcuri, UBS. Hi, thanks. Devinder, I actually had two. So first of all, I guess I'm a little surprised that OpEx is coming down so much as a percent of revenue from last year to 2023, given that you have all these new architectures and initiatives that you had to support, A. And B, if you compare say to like in video, they spend 28% of revenue on OpEx and they're supporting a single architecture. So can you sort of talk to how you're going to keep OpEx so low and yet grow revenue so much? Thanks.

Devinder Kumar: I think it starts with really some of the things that we talked about. Revenue's growing on a compounded basis at 20%. And some of the things that Mark talked about in terms of the efficiencies in terms of the engineering helps from that standpoint. And we look overall at the investments needed, and our view is in the -- out in that time frame, not right away. I mean, we ended at 31% in 2019, but out in the time frame in 2023 where the revenue has grown, we can manage it within the 26% to 27% OpEx rate.

Ruth Cotter: Matt.

Matthew Ramsay: Thank you very much, everybody. Just wanted to say the balance sheet stuff, Devinder, congrats.

I had one question and then I guess a clarification. The clarification bit was, Mark, on your slides, I think in the past you guys had talked about Zen 3 being on 7-plus, and you talked today about it being on 7. Maybe you could just clarify if that's a nomenclature change or if that's an actual change in architecture.

And then backing up to the question, I think if you look through the long-term model there, 30% of revenue would be, I don't know, \$4 billion, \$4.5 billion in data center. Maybe you could talk about how much you think of that is cloud versus enterprise versus HPC, which seems to have a lot of momentum. And then you talked about wireless infrastructure today for the first time. Maybe you could just break that, say, \$4 billion, \$4.5 billion revenue down and just how you're thinking about growth in those segments. Thank you.

Mark Papermaster: Matt, the clarification is in fact nomenclature as you said. So we work very closely with TSMC, and as you see how their public roadmap on 7nm evolved, there was at one time a 7nm Plus. And what is often happens is these new full-node changes like 7nm, what happens is some of the enhancements actually get folded in the base roadmap. So 7nm actually has encompassed in that nomenclature several different grades, shall we say, of its development. And so we matched up with TSMC nomenclature.

Forrest Norrod: Then with regard to the second question. I think, look, we think that that split in the

market is stabilizing that we do think that cloud currently constitutes about 30% of the overall market, enterprise about -- sorry, about 50%. Enterprise of about 35%, HPC about 15%. When we look out in that time frame, we don't see it dramatically changing. There will probably be a few percentage point shifting. And then certainly our ambition is to participate in all of those segments quite strongly. So ideally, we would be looking at a relatively balanced spread across all three of those businesses.

Ruth Cotter: Jason, over here. Thank you.

Mitch Steves: Thanks. Mitch Steves from RBC. So I really have two questions. The first one is really just regarding your competitor. Every day there's a new leak in terms of what the specs are for both you guys and your competitors. So, maybe you guys could provide some sort of high level view on what you guys expect to happen over the next couple years, because a long-term roadmap's obviously going to have a lot of competition in it.

And then secondly, I'm sure you're getting a lot of questions on this already, but you have a 20% target over the next several years. Just want to confirm that's more of like a long-term CAGR and it's not going to be a 2021, number. So what I mean specifically by that is probably 2021, based on my math, has got to be higher than that and it would decelerate. Just want to make sure that's correct.

Lisa Su: Want to do the first one and I'll do the second one?

Mark Papermaster: Well, just on the roadmap, of course there's always speculation and leaks, and we don't comment on the speculations. What we really do is really is what I alluded to earlier. We -- and what Lisa said explicitly. We of course have formidable competitors, and we're fighters here at AMD. So what we do is focus on getting the most competitive roadmap that we possibly can and really listening to our customers to make sure that what we're developing will address the workloads that they have coming at them.

Lisa Su: Yes. And Mitch, to your comment about the overall compounded annual growth rate, look, it's fair to say that 20% is a very strong number when you look across 4 years. If you just do the math, it's more than double the size of what the company was in 2019. So I would say it'll depend on how exactly things ebb and flow, but we would be very, very pleased, and that would be lots of market share gain to achieve the CAGR of 20%.

Ruth Cotter: Saski, at the back here.

Unidentified Participant: Hi. The question is about some of the developments in the startup space, and there's a lot of money going into new developments for semiconductors. And also some of your customers are talking about their own silicon. Maybe can you talk a bit about how you see that? And is this a new fight you have to take on, or is this something that we should be thinking about more deeply, or how do you think about that dynamic?

Ruth Cotter: Mark?

Mark Papermaster: Yes, I'll start just from a -- just take you back on even the historic view of the industry. There's always been a need for specialized devices. There's been specialized ASICs out there or starting often in an FPGA implementation. This is how new approaches, new workloads, are typically implemented. But when you look at them, they're typically targeted at a more narrow set of workloads. And this won't diminish at all the need for these very high performance, easy-to-program GPUs and CPUs we have out there. Code that's out there already leveraging these approaches is massive.

And you also commented on some of the larger companies developing their own silicon. Again, in a market where there's an insatiable growth of compute demand, there's room for these tailored solutions. And so it's an ecosystem, and we're very, very confident on our growth as we shared with you today. And there's plenty of room for tailored solutions in the industry. Some of them will have staying power; others won't. But there will always be a need for these easy-to-program, general purpose solutions that stay on the competitive path that we set out and are implementing here at AMD.

Ruth Cotter: Jason, Gus here.

Gus Richard: Your R&D efficiency's really impressive. And you sort of highlighted a couple areas where you're getting there: emulation, simulation, concurrent software/hardware design and modulatory. Everybody else is doing sort of the same thing. And I'm just wondering if you could give a little more color as to how you're getting there and how you're getting these products out so fast relative to your competitors with fewer dollars.

Mark Papermaster: Well, I'll start and --

Lisa Su: Are you trying to get Mark a raise?

Mark Papermaster: A short answer is necessity is the mother of invention. So just speaking candidly, when you look at the turnaround we were facing, and we really couldn't have the kind of investment we would have liked to have had, the things that I talked about aren't something that we're trying to do now. What I shared with you over my comments was looking back at what we implemented and allowed us to deliver these high-performance products to market. So it's really a credit to the team and how they responded to the challenges we faced in bringing AMD back to high performance. And the good news is they love it, and so we don't see any lack of that kind of innovation of how to improve going forward.

Ruth Cotter: Great. Laura.

Nathan Brookwood: Thank you. Nathan Brookwood, Insight 64. You were very explicit when you talked about Genoa being based on 5nm technology. But you were less specific when you talked about the advanced versions of RDNA and CDNA being on an advanced node. Can you give us some color on that advanced node and why you are characterizing the one very explicitly and the other kind of very generically?

Lisa Su: So, maybe let me take that. I think when you look at our processor roadmap, again, we have the win from El Capitan. It is a 5nm roadmap to get some of the performance efficiencies that we need to get. As it relates to RDNA and CDNA, David has a lot of things in the hopper. And we will talk more about what node and what architecture and all that stuff as we get closer to product.

Ruth Cotter: Great. Jason.

Blayne Curtis: Thanks. Blayne Curtis, Barclays. So just wanted to ask on the data center GPU market, maybe you can just talk about that split. So the CDNA, what exactly is different versus in graphics? Is it just you're adding acceleration, or is there something more fundamentally different? And then just from a competitive landscape-wise, unlike in CPU, your competitors should have 7nm this year. Can you just describe where you see the differentiation in the data center market for your product?

David Wang: Sure. I think as I mentioned in my presentation, the big part of that CDNA optimization is indeed adding some more higher density math for HPC, and for some of the what we call the tensor ops, the matrix acceleration ops that are needed for HPC and machine learning acceleration. So I will say that's a very big part of innovation on the CDNA.

And of course, we say we are optimizing the architecture also for compute. So that does mean we are putting less focus on other operations that are less important for high-performance computing. So those are the possibility and the opportunities for us to reclaim the silicon area to enhance the compute capabilities.

I think the other big one, though, is what Mark talk about Infinity Architecture. As we -- someone mentioned about explosion of AI, the machine learning, the size of the training. That is very, very important that the size of a training really grows. We need a lot more GPUs. They all need to be interconnected together with a high bandwidth efficiency. So that second-generation for Infinity Architecture that Mark talk about really allow us to interconnect multiple CDNA-based GPUs at a much, much higher bandwidth and the flexible topology to allow this large training model to be performed in a much more energy efficient way. So I would say those are the kind of key elements.

Ruth Cotter: Laura?

Unidentified Participant: You mentioned go-to-market investments, and it strikes me that that's a good idea because you wouldn't want something as mundane as salespeople to get in the way of everything you're trying to achieve. So maybe you could flesh that out a little bit, what kind of hiring expectations you have in what areas. And maybe a sense for how flexible you can be and what you're trying to build towards, or are you trying to do it as you go?

Lisa Su: Yes, maybe let me take that. So we are at a place where go-to-market investments are very, very important for us. Darren Grasby has our overall worldwide sales organization, and the focus is on customer-facing commercial, both for client as well as data center and enterprise, as well as sort of field application engineering at top hyperscalers to help them optimize our capability. So, lots of focus in this area. You will continue to see that be an area that we invest. And I think it's one of those places, again, that we can scale. As the company scales, we scale our capabilities there as well.

Ruth Cotter: Jason.

Shane Rau: Shane Rau with IDC. Forrest, my question's regarding data center GPU. Remember the three application segments you highlighted -- a virtualization, MI and HPC. Could you map your existing product line, and which product lines serve those three applications? And if then you could allow a quick follow-up.

Forrest Norrod: Yes, sure. So right now we have, if you look at the virtualization, we're a cloud gaming space. I think we've got some of our MI parts in current GCN-based MI parts in the virtualization space and in the cloud gaming space. We certainly see looking forward the CDNA parts, particularly starting with the ones later this year, are going to build on the wins that we've already had with our existing architecture and focus those in the MI and HPC segment. So on the virtualization side, it'll be a mixture of parts going forward.

Shane Rau: If you'll allow me to map into more product names as in you've got the V340 series, you've got business in cloud with Stadia and you have Radeon Instinct. And I'm trying to map those into those applications that you highlighted earlier and get a sense for where

you stand now, given where your strategy that you just highlighted is going to take you as far as server GPU.

- Forrest Norrod: You know, I don't think we get into the details of the product mapping particularly if they're embedded in our customers' end products. So I think I'll stand by what I said a moment ago, which is we've got some of the existing MI products in the virtualization space based on GCN architecture, and we certainly see CDNA driving heavily into MI and HPC going forward.
- Shane Rau: And then if you'll allow, under MI, would you segment the MI in a training versus inferencing way? And do you consider both of them under the TAM of MI?
- Forrest Norrod: I think we definitely consider both of them under the TAM of accelerated MI.
- Ruth Cotter: Saski, over here.
- Harsh Kumar: Hi. Harsh Kumar, Piper Sandler. So you talked about CDNA and unified data. Is that in Milan or the generation after that? And then also, my understanding from this is that the data will appear in the same rate to a CPU and a GPU. Is that roughly accurate? And how big a deal is that to your end customers? And when your competitor comes out with a GPU, how easy or hard is it for them to emulate something like this?
- Forrest Norrod: A number of questions there. Maybe I'll take the first one and then leave the others to David. All that we've said so far on the supercomputers is that Frontier is based on a future custom EPYC processor, and then El Capitan is based on Genoa. I think that's all that we've said thus far on the processor side.
- David Wang: I think on the -- given our CDNA2, that is how we are coherency, that allow us to have the data in its preferred location. Instead of needs to be replicated, the GPU or CPU are copied to each other.
- A very simple example is CPUs doing pre-processing on a set of a data, and that needs to be processed again by GPU and then have to pass back to CPU for post-processing. Imagine now we have a unified memory that's cache coherent. You can just keep that data in HBM memory for the GPU, and you can just do the whole processing without moving data back and forth between CPU and the GPU. And that's tremendously helpful both from the programming point of view, as well as from performance point of view. As far as our competitor-wise, it depends on how they work out the coherency scheme with whoever their partner is.
- Ruth Cotter: Great. I think we have time for one last question. Laura.
- Joseph Moore: Hi. Well, I actually had two, but I don't know how that works. But just one nit question. Why is the free cash flow margin and operating margin 10 points, as much as 10 points apart? Because I would think for you it's pretty comparable.
- And then I just -- you also mentioned M&A as a potential something to think about. Just any context on what type of M&A you might think about. What -- is that tuck-in type of stuff or something bigger? Seems like you have a pretty full portfolio already.
- Devinder Kumar: Let me take the first one. So the first one, if you look at the operating margin at approximately mid-20s, as we call it, you do have the 3%, approximately 3% cash tax rate. You have CapEx investments. You have interest expenses. And then obviously as

you're growing the business, there's some amount of money needed to fund the growth with the significant CAGR that we showed from a revenue standpoint. And these are approximate numbers, so yes, you're right; there's a difference between that. But those are the three or four factors that drive to the difference between operating margin and free cash flow margin.

Lisa Su: Yes, I wouldn't expect though that, Joe, there should be a 10-point difference. So, I think that's just the way we built the model at this point.

And then look, to your second question, hopefully what you've seen this afternoon is like we're really excited about our organic growth opportunities. There's a lot of growth. There's a lot of technology. There's a lot of market opportunity. And we see that as really a great business model. Now, the fact is the company's a lot stronger today than it was a few years ago, and so that balance sheet is also something that we're proud of. And we're always going to look at what are good priorities for us to continue that growth, and that's the way we think about strategic M&A.

Ruth Cotter: Great. Well, thank you, everybody, for joining us today and to everybody for tuning in on the webcast. And for those of you here in person with us today, we're going to retire to a demo area out in the front lobby and you can spend some time also with the presenters out there. Thank you.

Lisa Su: Fantastic. Thank you.