

Qualcomm Investor Day 2026 Transcript

June 24, 2026 | New York, NY

Introduction Video

We are creating something totally new because a wave of innovation is reshaping the world around us, driven by AI that works for you. It's one of the largest shifts the technology industry has ever seen and it's the moment Qualcomm was built for.

Intelligence is moving everywhere. In autonomous vehicles, personal devices, industrial solutions, robotics, and data centers. To make this all possible, unique high performance and low power, the broadest technology portfolio and a global ecosystem with global scale. Billions of devices are being reimaged for this new generation of intelligent experiences. That's what makes Qualcomm uniquely positioned for this moment and for what comes next.

Announcer

Please welcome Senior Vice President of Investor Relations, Brett Simpson.

Brett Simpson, SVP, Investor Relations, Qualcomm Incorporated

Good afternoon, everyone, and welcome to Qualcomm's 2026 Investor Day. It's great to be here in New York, and it's great to see so many familiar faces. Now, a lot of you have been asking me recently why I joined Qualcomm and well, you know, I think it's pretty clear. I think we have a really compelling investment case. And today is an opportunity to really share with you why we're so excited about what lies ahead for Qualcomm. We've got a lot to share with you today.

Before we jump into things, I just want to say a big thanks to everyone involved from Qualcomm making this day possible. It's a huge amount of work. I really had no idea how much man hours goes into putting an event like this on. And, I just wanted to say thanks to everyone. It's really amazing. And I also wanted to say a big thanks to all the executives from Qualcomm who are here today. You've traveled a long way, and we really appreciate your support. And I also want to thank the Modular executives here today. We have Chris and Tim, and we'll hear from them a little bit later.

Now, an Investor Day wouldn't be the same without a disclaimer. And so I'd direct your attention to this slide, which contains important information regarding our use of non-GAAP financial measures and forward-looking statements, as well as the use of "Qualcomm" throughout the presentation.

So, we have a packed agenda today. Cristiano will start the day with a strategic overview. Lots of things happening at Qualcomm, which he will go through. And then we will pivot to data center. And I know you've all been waiting to hear from us on data center. Tony Pialis, who's new to Qualcomm and runs our data center division, will talk you through why we're so excited about the differentiation we can bring to the data center and why we see big

architecture changes happening with agents. Nakul will then address how Qualcomm is moving up the value chain with full-stack solutions across automotive. But more recently, industrial AI and robotics, where we see really big opportunity for Qualcomm over the next 3 to 5 years, particularly in physical AI. And then Cristiano will come back to talk about how agents drive new edge opportunities for Qualcomm, including 6G.

We're also going to share some important updates on the software side. So I'm sure you saw the announcement today with the acquisition of Modular. Akash will be our last speaker of the day. He will be bringing everything together and giving some important updates to our financial outlook. After Akash, we will have a short Q&A and then we encourage you all to join us for drinks behind the curtain in a demo room. And you'll also see some of the new products and experiences we want to share with you today.

And just one more thing. I was actually going to put my kilt on today. It's a big time in the U.S. with the World Cup. And tonight, Scotland will be taking on the mighty Brazil. And I think Scotland only needs one point to qualify here. So after six, you may not see me. I'll be going to join the Tartan Army. So with that, please join me in welcoming Qualcomm's President and Chief Executive Officer, Cristiano Amon.

Cristiano Amon, President and Chief Executive Officer, Qualcomm Incorporated

Thank you so much, Brett. Appreciate it. Thank you.

Thank you, everyone for joining us. I really appreciate it. This is an exciting day for Qualcomm. Actually, as a matter of fact, June 24th. I became CEO of Qualcomm on June 30th. So actually, I picked this week on purpose. This is exactly five years since I became CEO of this company.

Well, thank you. But this incredible company that I kind of joined as an engineer about 30 years ago. And what I'm going to tell you today, and I'll tell the reason this is special, it's going to be a little bit about the transformation.

2021: We put a strategy together. We actually had an Investor Day in 2021, in that year, and we said, what we were going to do. And I think we are at a point now that, we put everything in place to execute on the strategy, and it's time to start the new chapter of Qualcomm. And that's what my intent is: to tell you about that today. These are really exciting days. We have a lot of things, we did the best that we could to pack a lot of big pieces of information in a short period of time. Brett walked you through the schedule. We have Q&A at the end. Many of you will be tempted to run to some memory guy's earnings call starting at 4:30, but I would encourage you to reconsider and stay for the Q&A. It's going to be an awesome Q&A.

So with that, let me just get to the presentation and before we start, to each one of these chapters. So, this is where we started and I think we've been a company that has been probably the most focused semiconductor company in the mobile market. And I know, I

think the mobile market to some extent, I think, [fell] out of favor when everyone on Earth had a smartphone, but it's going to become very interesting.

But the reality is we were the most focused semiconductor company on this one market and we put together a strategy and start executing that strategy. That's what I told you. And the following five years, this is what we built. We built a diversified edge leader across multiple end markets. We built an automotive business. We built what we call an IoT, and I know there's a lot of things within this IoT segment for us, so we started to break down for you and what we call personal AI and compute, which is how wearables, virtual reality, and augmented reality evolve personal AI devices, as well industrial, networking, and robotics. And that is the company that we built: really focused on the edge to create a diversified edge leader across multiple industries. And we're now in this transition to what we're going to do in the next five years, and this is the next chapter of Qualcomm.

There's basically three pillars to it. There's actually three dimensions to the Qualcomm future. The first one, and I believe many of you came here today to see what we're going to do in the data center. We're building a data center platform. It's a comprehensive portfolio of solutions. I think you should think about it, what we have done the past few years is like a submarine strategy: we've just been executing, executing, collecting assets. And when we get to this point, we feel that we have a comprehensive portfolio to enter the next phase of the data center. And you hear about it as inference of scale, as we see this aggregation, a lot of people ask me this question and before you ask on the Q&A, I'm just going to answer. A lot of people ask, "Oh, this crowded market, is that too late?" "It is never too late for Qualcomm."

I think, because this is a market that moves very, very fast. So if you have technology leadership, there's always room for you. That's how we think of it. People can not like the company about many different things, but I have never heard anybody that says Qualcomm does not have a solid, technology capability. And I think that's what we put into work.

The second part is now that we have been building this platform of devices at the edge, we're moving to a full-stack player. So as agentic starts to transform the devices on the edge, we're going to build on those assets and it is one of the broadest portfolio semiconductors and we become a full stack player in physical AI, compute, everywhere. And that's what you're going to see in the presentation today. Actually, before I go to part number three, I wanted to break that down. There's actually one more thing.

Before the full-stack execution in automotive, industrial, networking and robotics. It's what's going to happen with our mobile business. I actually have a portion of the presentation today to tell you how you should be thinking about the future mobile edge devices that consumers are going to utilize and that's part of the second dimension of Qualcomm. And that gets to the number three.

The number three is how we're going to go from silicon to platform solutions, fully integrated platforms across hardware, software, developer ecosystem, also making Qualcomm a developer-first company.

So that is the three dimensions of this new chapter of Qualcomm and what we're going to be doing the next five years and we're going to show you some of the early wins that we have on this.

I want to summarize to you as I get off the stage, what is the Qualcomm advantage? We're always very proud of our technology and IP, but we're also a company that can build very broad relationships. That's one of the assets of the company. We have built strategic relationships across industries. We were not a player in many of the new industries that we entered. And if you look at today, Qualcomm builds very strong relationships with all of the leaders of the industries, and we build an ecosystem around it. The other part of it, and that's what when people ask about if it's late to enter the data center, you think about scale and execution. Our engineering capabilities, our operations and supply chain. It's one of the leading scale and execution machines in the semiconductor industry. When I talk about the technology advantage of Qualcomm, it's really very broad. It's really a broad portfolio and we take pride in providing leadership solutions. We have been in many areas a creator of standards. And that with everything connectivity – it was just wireless, now it's wireless and wireless including connectivity in the data center.

And that is the portfolio that we have been building. It's every form of compute. When data centers talk about disaggregated computing, this is the reality of mobile from a very long time ago. So, we have every disaggregated compute and actually we do this for different industries, including what we do on safety and industrial-grade processors. We're building a comprehensive software and developer ecosystem and you're going to hear more about that today. Sensing becomes even more important when you think about physical AI and multimedia, advanced packaging memory, and we have a very strong patent portfolio, which is a result of about a hundred billion cumulative in R&D investment. Our customer reach, and you should be thinking about those concentric circles. We started building ecosystem partnerships in mobile and you see as we go into every other industry has now expanded. And that is a unique asset of our company.

Our customers believe in Qualcomm and that is true, especially when we enter new markets and that's what we're going to be doing right now in the data center. And the last part is the execution. I'll spend a little bit on this. We're actually a proud member of the TSMC fabless ecosystem. We work with a large number of foundries and manufacturing partners and we have incredible scale. I'm just going to give you a few highlights. We consume over one million leading node wafers. It's just annual scale. We do over 75 chips tape-outs per year, of that 75 over 30 is an advanced transistor. We ship 40 billion components.

Total of 2.5 million wafers, that's a lot of scale. And we also have some things that are very unique. Some other semiconductor companies actually tell us, Cristiano, you're crazy. Sometimes we tape-out a chip. Before we get the chip back, as soon as TSMC finishes the mask, we go to production and we go to production at scale. That's the maturity of our manufacturing capabilities and we ramp completely new nodes to 100K wafers. That's the nature of the phone business within about two quarters. There's a very big ecosystem and I think that is an asset that we've also been bringing to the table. And the unique thing about Qualcomm, and I said this in Computex, is we're now present across the entire compute continuum. From sub-2 milliwatts to about 200 kilowatts. And as this gets transformed with AI and we apply our technology and we build on a developer friendly, unique software platform — I'm sure you're going to hear more about that — that creates an incredible opportunity for the company, not only for the business that we have built, but what we are going to build in the next five years.

So that's our strategy. This is what we're doing. And now what I'd like to do is start unpacking that for you. I want to bring to the stage Tony, who joined us from Alphawave and now runs our data center business and I'm sure you're going to be amazed about what we're doing and what Tony is going to share with you. Tony, please come on stage.

Tony Pialis, EVP & GM, Data Center, Qualcomm Technologies, Inc.

Hey everyone. I'm Tony Pialis, general manager of data center for Qualcomm. First of all, where's Brett Simpson? Brett, I thought I told you I never want to follow Cristiano. I mean, in terms of being a visionary in the industry and a legend, I'm so lucky to be part of his team.

So folks, one of the most common questions I've been receiving this morning as I've been meeting with many in the press is “why would I want to leave my role as founder and CEO of a semiconductor company to join Qualcomm to run their data center business?” For me, look, I've always been good at math. The equation for that has been really simple. It's all about accelerating value creation. That's what we're here to do. That's Cristiano's vision. That's why I came and that's what I'm here to explain to you today. How are we going to accelerate value for all of our shareholders and customers?

So let's get to that. Agentic AI changes the economics of compute. What does that mean? It means token counts are skyrocketing as we introduce agents. It means CPU attach rates are soaring through the roof. You can't find CPUs anywhere. They've already been bought up. So traditional infrastructure will not scale to the needs of Agentic AI. So the industry needs a paradigm shift in order to deliver this. I will explain this paradigm shift and how we will lead it. But first, let me introduce to you Dragonfly. This is our data center infrastructure to lead the industry into the next stage of AI. Let's hit it.

Dragonfly Launch Video

In less than a second with a single query, one AI agent checks 47 locations across three midwestern cities while other AI agents draft proposals, negotiate with venues, route

vendor schedules, order a lunch, pay the quarterly taxes, and much more. This is just one second with one event planner and it's what every second is about to look like at your data center a million times over. This is the moment of inference. It's the moment for Qualcomm Dragonfly. In this agentic era, every query is now a swarm. Every interaction, an avalanche. This era demands a whole new data center calculus, one where performance must meet efficiency. The calculus Qualcomm has lived and breathed for decades with billions of devices worldwide. That experience has shaped Qualcomm Dragonfly through and through. Built for the swarm. Cost efficient. Engineered for superior power efficiency and tokens per watt. Open by design, scalable by nature.

Qualcomm Dragonfly is made for this. We're ready to take you wherever you need to go as far as you need to go this second and all the seconds to follow. Qualcomm Dragonfly.

Tony Pialis, EVP & GM, Data Center, Qualcomm Technologies, Inc.

Well, Cristiano, the submarine has surfaced and here we are today in New York. All right. So, team: Agentic AI requires a new compute infrastructure. This chart illustrates why. The industry has blown through GenAI and reasoning and here we are on the cusp of deploying agents en masse. But what does that mean? It means a single agent query is generating fifty to a hundred times inference requests. We have over a million tokens being generated by a single query. Traditional compute infrastructure cannot support this scale. A paradigm shift is needed.

So first, what you see up on the screen, which by the way, I love the screen. I need to get one in my basement. What you see up on screen is a traditional GPU-based compute that's been deployed across data centers worldwide. It's been built to support both training and inference. It runs hundreds of kilowatts today.

Other competitors' solutions will be running north of 500 kilowatts. How can we deploy this and support a hundred times more inference calls than we do today? We can't. A new solution is needed to lead the industry. And so, what we in our submarine have been doing is building a world leading disaggregated compute infrastructure. Cristiano used the analogy, "Well, we already did this in mobile." That is what's needed to lead in data center. Bespoke solutions that deliver hardware acceleration for each and every function needed to deploy agentic compute. Various forms of CPUs performing specific functions, unique XPU's, some targeted to attention for prefill, some targeted for KV caching during decode. All of this blaze together using both copper and interconnect that deliver world-leading connectivity, creating one compute network that will transform the industry. Now, as I've also been asked earlier today, "How is Qualcomm positioned to win? Aren't you late?" Let me give you some background.

I've been formally in the company for the last six months, but I've worked with the company for years as a supplier and partner to them. What I will tell you is this is an engineering-first company. They blazed the trail and led the world in wireless communications. Fast forward, they lead in mobile compute. Now we are leading in both automotive and PC.

When the company turns its attention to solve a new problem, we revolutionize the solution and push our way to the forefront. And folks, I am here to tell you today that is what we will do and are doing in data center. We are pushing our way to the forefront and our customers are pulling us the rest of the way in. So, how are we going to deploy our agentic infrastructure? Four simple steps.

The first one starts right now. We have our connectivity portfolio that comes from Alphawave. That is in high — finally, I've been talking about getting this into production for the last two years. I am proud to say this year we are qualified in our first leading major hyperscaler, generating meaningful revenue for the company over the next four quarters. Follow that by custom silicon. Lots of speculation and I'll have a lot more to say so you can hold your questions. We are winning in this space and we'll be delivering meaningful revenue at the end of this year, calendar year starting fiscal quarter one, 2027. Then in 2027, we are launching our third generation of AI accelerator. What makes the third generation unique is it's the industry's first near-compute AI accelerator that will transform inference. A lot more to come on that. And we're not done yet, folks. You layer on top of that in the middle of 2028 we will be launching the industry's first Oryon server class compute solution. We will launch a fleet of agentic, general purpose, and head node compute that will complete the Qualcomm infrastructure.

Now let's get to it. So, this is the best worst kept secret, or the worst best kept secret, however you want to view it, that's limiting the industry. I've been in the industry for 30 years. Compute has increased by more than 60,000 times in that span. Kudos to the amazing Qualcomm engineers for accomplishing this. But guess what? Where transformer sizes are growing 240 times over a span of two years these days, look, compute memory is only doubling in that same time span. So, what does this mean? It means there's no point packing more compute unless we solve the memory bottleneck. So, before I show you how we do that, what you see here is how a modern day XPU works.

You see that GPU off to the side. It is pulling and pushing data to a HBM stack. You have thousands to tens of thousands of wires constantly carrying data, generating tremendous amount of heat, burning significant amounts of power, moving data back and forth in order to support the scaling transformer sizes. This solution cannot keep up with the growth in AI models. So, an innovation is needed and this is what the submarine's been doing. This is what I was most pleasantly surprised to find when I joined Qualcomm. We have broken through the memory bottleneck. How have we done it? We have re-architected compute for XPUs. We have separated the AI accelerator from the XPU and what you see, we now put our XPU right under a DRAM stack. What does this mean? Very, very important. We offer all of the performance advantage of SRAM, but with the density and the memory capacity that HBM stacks offer.

And so what does this mean to the industry? It means that congestion that you saw with HBM is gone. What happens now, a great analogy I use with reporters is imagine working in the same building that you live in and so you only travel up and down. And what does that

mean for the highways and the roads that connect the suburbs to the city? Guess what? The roads are clear. So, the value this brings to the industry is lower power consumption, less heat, and that expensive road of silicon interposer that HBM solutions use are no longer needed. We can deploy multiple HBC stacks within a single compute device using standard packaging. That is a tremendous value that we deliver to the industry in terms of performance per cost advantage. Now, how do we quantify these benefits? Look off to the side. These are ultra-low latency workloads.

Things like coding workloads. HBC offers 200 times capacity per watt better solution than SRAM. There are many that are announcing SRAM based solutions today. Look what we can deliver with HBC. Now you look at the other end of the spectrum for high throughput workloads. These are the workloads that GPUs and HBM have dominated for years. With HBC, we deliver six times the bandwidth per watt versus competitor HBM based-solution. And so while the rest of the industry is now trying to deploy two unique solutions to solve these bookmarked endpoints and those in the middle, well, guess what? They're kind of stuck in the lurches. With HBC, we offer a single solution that can seamlessly span this entire sphere of workload and deliver multiple folds performance per watt and performance per dollar benefit. That is a direct TCO advantage that we offer to the industry. So, I go back to those actions asking questions.

Are you late to the game? With this kind of performance advantage, the industry is demanding and pulling us in. So, who better to introduce HBC to the industry than one of our lead partners, Microsoft? So, it is my greatest pleasure and my first but not only surprise to you today to introduce Satya Nadella, who will talk to you about how Microsoft will be deploying HBC within Azure data centers. Let's hit it.

Satya Nadella, Chairman & CEO, Microsoft [Pre-recorded]

Hello everyone, it's great to be back at Qualcomm's Investor Day. At Microsoft, we've had the opportunity to partner closely with Qualcomm across multiple waves of computing from the PC to mobile, and now AI. And across all of them we've shared a deep commitment to innovation at the systems level. Bringing together the silicon and software to deliver meaningful advances for our customers. This includes our continued collaboration to reinvent the PC for the AI era, which will only become more critical as we deliver unmetered intelligence at the edge with Windows, and we are not stopping there; in fact, with Project Solara, we are collaborating on a new platform, purpose-built for agent-first devices and it has been fantastic to see the reception since we announced it together earlier this month. And now, we are excited about your innovation in the data center, especially around high-bandwidth compute, and we look forward to building on that together. HBC implements an innovative architecture, with high memory bandwidth and integrated compute, that unlocks significant improvements in cost and performance for the next generation of AI infrastructure, and there is so much more to come. We look forward to our continued partnership as we build the next generation of computing together. Thank you so very much.

Tony Pialis, EVP & GM, Data Center, Qualcomm Technologies, Inc.

Folks, if I had a microphone rather than a lapel device, I would drop it right there. So that's our first surprise. Stay tuned. Many more on its way.

All right, how are we deploying HBC? Well, we're deploying it to target a 680 billion addressable market for us. HBC earns us the right to win a significant portion of that market over the next few years. Look on the left side of the chart to quantify Satya's comment in terms of performance benefit deploying HBC. We deliver depending on the workload anywhere from a four to 8x advantage that directly translate to TCO advantage that we deploy to our customers. I am super excited for our launch of our first HBC product in the middle of 2027. That's it. Folks, take those pictures. You're going to be seeing this to last for years to come. Our AI 250 product will introduce the first near memory compute that employs HBC and will be a complete game changer to the industry.

We're following that in 2028 with AI300, which launches our second generation of HBC. It will integrate UAL and ESON the latest scale up network fabrics. And for scale out, we will be deploying both copper and optical networks to connect AI clusters.

Now, look, as amazing as this hardware stack is, it really is just a foundation for running software. Software is where the magic is and takes a lot for me, a hardware designer, to acknowledge that. We will be deploying a full software solution stack that includes the most sophisticated orchestrators that will manage and route the traffic across a disaggregated compute cluster all the way down through frameworks and most importantly, open frameworks that will allow model developers to both develop and deploy and scale their models. We offer all the kernels and compilers that are necessary in order to optimize the latest models onto our specific hardware. While others in the industry build moats trying to protect their hardware deployments, myself and we at Qualcomm, we believe in building bridges to unite the industry.

And so my second surprise, which I know came out through an announcement, but it is still my honor to announce that today we published an announcement for Qualcomm to acquire Modular. Modular is a world leader in developing AI software solutions and who better to introduce how we will jointly transform the AI industry with our open solution to disaggregated hybrid compute than Tim Davis, co-founder and president of Modular. Tim, will you come up on stage?

Tim Davis, Co-founder & President, Modular

Thanks, Tony. And you can obviously tell between Cristiano and Tony, that's why we're so excited to be joining Qualcomm.

But hey everyone, I'm Tim, one of the co-founders of Modular and I've been building AI infrastructure for almost a decade. First at Google Brain for six years building core AI data center and edge infrastructure for mobile devices and TPUs. And then at Modular as co-

founder and president. Modular has assembled one of the best teams in the industry that have helped found, build and contribute to most of the core AI infrastructure in use today. Now, you may have read this morning that we are beyond excited that Modular is joining Qualcomm to supercharge our AI infrastructure and distribute it to the world, but you might be wondering what is Modular?

Well, I'm here to tell you. Modular is building AI's unified compute layer, a software layer that enables AI models to run on any hardware and is heterogeneous by design. For developers and enterprises, that means building once, deploying anywhere, lowering the cost of running AI at scale, and accelerating innovation into their production data center workloads. Modular is the portable alternative to Nvidia's software stack designed from day one for every AI accelerator.

Let's walk through the stack. Mojo gives developers the high-performance, low-level programming model they need without locking them into one platform. Max gives them the model and serving layer that they need without relying on Triton or TRT LLM and Modular Cloud gives enterprises the distributed serving infrastructure they need without being tied to a single silicon vendor. Together this is a full AI compute platform for the heterogeneous data center, and we are up to 50% faster when executing AI inference workloads on third party hardware and we have the numbers to prove it. After four years of R&D, Modular is rapidly coming to market with industry leading performance across many of the world's most foundational AI models.

Now importantly, our platform turns heterogeneous data center systems into multi-silicon AI token factories. In this world, enterprises, partners, and developers can use the best silicon for each workload without being locked into a single hardware stack. Because modular is heterogeneous by design, the industry can achieve lower TCO, higher performance, and greater portability across the world's compute infrastructure. We are incredibly excited that Qualcomm will help us scale our technology to data center customers everywhere, enabling broad hardware independence for the world. Chris Latner, my co-founder at Modular, will share more about our incredible future with Qualcomm later today. Back to you, Tony.

Tony Pialis, EVP & GM, Data Center, Qualcomm Technologies, Inc.

Thank you, Tim. Look, rather than preparing for this event, I've been fielding calls all morning from hyperscalers and customers asking, "Hey, how can we begin incorporating Modular's technology?" So I'm super excited about what we will be doing together.

So now let's talk about CPU technology. Qualcomm has a long lineage in leading in CPUs. They pioneered mobile compute with Snapdragon. We are winning now in both PC and automotive. The company's focus is now transitioning over to data center. And so today there's been a lot of speculation about this, but I am excited to introduce Qualcomm's C1000. It is a data center fleet of processors. These processors will run the industry's fastest cores, running greater than five gigahertz. This is more than 30% faster than any of

the competition. Coupled to that, we offer more than 250 cores to run the highest throughput workloads. Combine that with Alphawave's leading PCI Express technology delivering greater than two terabytes of I/O bandwidth.

Then add on Qualcomm's memory leadership delivering the highest performance, lowest-cost memory solutions, employing LPDDR. We have server-class RAS security embedded directly in the hardware. And finally, our CPU is also AI native. That amazing HBC technology that I walked you through for our AI inference engines couples directly as an HBC attach to accelerate AI workloads natively onto the C1000. Now, how are we deploying it through three various product lines for the C1000. The first is our Agentic CPU leveraging our HBC attach. We deliver industry-best performance. We then deploy it through general purpose CPUs running virtualized container workloads. And finally, our AI head node CPUs running and orchestrating all the traffic across disaggregated heterogeneous compute data centers. All of this targeting a \$200 billion market and that number is growing every day with each and every analyst report that gets published. So, my next surprise for you, I'd love to introduce to you Mark Zuckerberg, founder and CEO of Meta as he introduces how Meta plans to deploy the C1000 into its next generation of data centers. Hit it.

Mark Zuckerberg, Founder, Chairman & CEO, Meta [Pre-recorded]

Hey everyone. Great to be here at Investor Day with you. Meta and Qualcomm have been partners for a long time and we're doing some great work together. We first started on the Quest headsets and then we brought LLAMA to Snapdragon so people could run AI right on their phones. And today, Snapdragon is powering our AI glasses too.

Now we're bringing that partnership into our data centers. With our latest model Muse Spark, we're delivering AI to billions of people every day across our apps. The data centers, the energy, the compute to run billions of model inferences, that's what makes it all possible. Our goal is to deliver personal super intelligence to everyone in the world. And as our teams work hard to build state-of-the-art models, we need to innovate with how we get the power, we need to scale it, and make it accessible to everyone.

So that's why our work with Qualcomm is so critical. They've spent decades figuring out how to get the most performance out of every watt. They're really good at it, and now is the right time to expand this partnership. So today, I'm excited to share that we've entered a multigenerational collaboration for Qualcomm to supply CPUs to our data centers and help power our next-generation server fleet. This will help put personal super intelligence into billions of people's hands. There's a lot more to come and I'm looking forward to building together for a long time. So, thank you to Cristiano and all the teams at Qualcomm for all the work you do here.

Tony Pialis, EVP & GM, Data Center, Qualcomm Technologies, Inc.

Satya and Mark already in my presentation and folks, I'm not done yet. So let's move on to the third product line, custom silicon. There's been rampant speculation in terms of what

we're doing here. First off, I want to establish we are in custom silicon to target the highest tier of customer where we can deliver the most value add by bringing our incredible IP portfolio to play. So, our wins to date are based on both Alphawave legacy wins that are scaling into production and most interestingly in the first six months here, I am extremely excited to announce, we have won two major hyperscaler deals that will contribute meaningful revenue to Qualcomm starting at the end of this year. And so folks, how do we win in custom silicon? We work with our customers. We take their specs and we help them build their chips, whether it's in the front-end RTL design or whether we help them convert their designs into chiplet-based solutions, delivering the most advanced compute and networking solutions in the world.

And then using our manufacturing scale and know-how, we optimize their yield and enable them to deploy their bespoke solutions en masse to their data centers. I've been in this space for 30 years. The way you differentiate and win in custom silicon is through your IP portfolio. We have the world's best custom silicon IP portfolio. We have our own compute that we can optimize for our customers. We have HBC, a complete game changer in the AI industry. Add onto that Alphawave's leading electrical and optical service that kicks the butts of its competitors. We've been working in the silicon photonics and optic space for more than five years. We bring that to play in order to bring connectivity directly into compute and you couple that with Qualcomm's leading manufacturing and supply chain. This is how we've been winning in the first six months. I have not had to push my way into hyperscale customers. They've been pulling us in and when they pull us in, it gives me a chance to expand and bring the rest of my solutions to play.

The final product line I will walk you through is our connectivity. This is the third bottleneck in the industry. The first was memory. We solved that with HBC. The second was the software stack. Modular, Tim and Chris will help us solve that. The third is connectivity and this is near and dear to me. This is where hyperscalers, as they deploy customers of compute, as AI compute doubles, so does connectivity every two years. And then you have this transition from copper cables to optical solutions with new low latency scale up and scale out fabrics. This is a race to the forefront and we have the technology pieces needed to win.

We have everything you need to scale from the millimeter of connectivity all the way through to tens of kilometers from our leading die-to-die technology to our co-packaged optics interfaces that bring the world's fastest lower power connectivity right next to our compute. Add onto that, leading PAM4 electrical and optical SerDes to drive scale up and scale out networks. Today in production at 224 Gbps, soon we'll be in production with 448 Gbps. And finally, coherent-lite. Today it drives optical connectivity across campuses. But when PAM4 runs out of steam, coherent-lite will be connecting compute clusters within the data center. So with all these pieces, I'm excited to announce our connectivity portfolio. We are already in production with our first generation of 800 Gbps electrical and optical DSPs, including our first generation of coherent-lite. By the end of this year, we will

be in production with our second generation of electrical and optical solutions deploying 224 Gbps solutions.

These products are anchored with the lead hyperscaler win already. And then looking forward to 2028, where we'll be bringing in our third generation of connectivity solutions. This will be based on 448 connectivity and we'll also deploy our next generation of coherent-lite solutions. So folks, stepping back now, we have developed a transformational infrastructure that is already winning in the industry. Four product lines, each of them already anchored with multiple customer wins and a pipeline that'll blow your heads in terms of accumulated value. Incredible metrics. Look at that. Up to eight times better tokens per watt per second than traditional GPUs. Greater than 200 times memory capacity compared to SRAM solutions. Six times memory bandwidth per watt and for our CPUs greater performance than 2X than our competition. All of this is direct TCO advantage. So why are we entering now? Because we have the performance that the industry needs.

So folks, final takeaways I want you to remember. Tokens per watt replaces FLOPS. The race has changed. Embedded solutions, embedded providers, they're playing the old game. There's a new game in town and it's all about delivering agentic-first rack scale platforms that delivers the world's best TCO. That is what we've been building in our submarine. Add on to that four product lines that we've already anchored with hyperscaler wins. And now with today's announcement, we have the world's best software stack that will build bridges across all the hardware of the industry. And finally, and I will leave the numbers to our man, Akash. I am very, very proud to announce within the first six months on the job, we will deliver multiple billions of revenue starting fiscal '27. And for those of you that aren't aware, that means starting this calendar year. And so I am excited now to introduce my colleague Nakul, but before it comes on stage, I have one last person I want you to hear from.

It's Tareq Amin, the CEO of HUMAIN. He's a visionary in the industry and he has been our first data center customer. Let's hear what Tareq has to say. Let's hit it and thank you.

Tareq Amin, CEO, HUMAIN [Pre-recorded]

Congratulations to Cristiano and the entire Qualcomm team on this bold milestone. AI is no longer a technology trend. It is becoming the operating system for every industry, every economy, and every society. I believe the next decade belongs to inference. Billions of agents. Trillions of interaction. Continuous intelligence operating across devices, enterprises, and government. Through our collaboration, HUMAIN and Qualcomm are deploying the next generation AI infrastructure by combining Qualcomm breakthroughs in semiconductor innovation with HUMAIN full-stack AI capabilities, from infrastructure, cloud platforms, foundation models, and sovereign AI systems. Success will not be measured by peak performance alone. It will be measured by performance-per-watt, performance-per-dollar, performance-per-outcome. This is where Qualcomm brings something extraordinary. A really fundamentally different approach to AI compute that

challenges really conventional assumptions about power consumption. Years from now, we'll look back on this moment as the beginning of a new era for AI.

Nakul Duggal, EVP & Group GM of Automotive, Industrial and Embedded IoT, and Robotics, Qualcomm Technologies, Inc.

Good afternoon, ladies and gentlemen, and a very big round of applause for Tony first of all. Welcome. Dragonfly is here.

I run Qualcomm's Automotive, Industrial, and Robotics businesses, and as someone who's been with the company for over 30 years, there has not been a more exciting time to be at Qualcomm. Each of these three businesses are quite different. They're quite unique, and they need different strategies. But over the years, to diversify Qualcomm, we've had to build new muscles that strengthen over time. We see the next several years belonging to physical AI, and massive transformative change that physical AI is going to drive in industrial and enterprises, with robotics becoming a key catalyst. You've seen what we've done with Auto. I'm going to give you a sense over the next 30 minutes of how we are preparing for physical AI.

Physical AI is the next great computing wave. It doesn't run in the cloud, it runs on the edge. It's going to run in factories, in warehouses, in retail, in hospitality, in hospitals, and robots are going to be a very important part of physical AI. Automotive is the first example, and you've seen what we've done with the Automotive business, Industrial, and Embedded. I gave you an update on this space in November 2024, and we've been preparing, and I'll give you a sense as to how this business is progressing. And robotics is a space that we entered only in the last nine months or so, and we have expanded very quickly into multiple environments. We are finding strategies to move up the stack. One thing that is very unique about Qualcomm is that while these are three unique businesses, there is one single common theme, one IP roadmap, one product foundation, and one physical AI platform.

If you look at the way physical AI is now moving into our lives, you've seen human-facing AI, you've started to now see machine-facing AI, and, ultimately, embodied AI. These three layers are highly interwoven, and they compound over time. And you will start to see how this plays out now as I walk through Automotive, Industrial, and Robotics. Hopefully, this will all become clear. But they are all underpinned by the same technology layers that create vast automation capability. Human-facing AI has changed our interaction layer, first with chatbots and digital assistance, but now with body cameras and with XR glasses. And you are changing the physical space interaction layer between humans and the devices, the products that they own. Instrumented AI, machine AI, is about putting AI that is embedded into sensors, embedded into cameras. And it mostly comes down to that sensor, that endpoint, being situationally aware. The real economic unlock is, however, physical AI and embodied AI. And what we see here is the ability for devices to perceive, to reason, and to actuate, with the goal of being able to complete a physical task. And that

evolution is just starting. We find ourselves at this inflection point, and, as this matures, we see a massive edge content upgrade cycle hit us.

This content shift across the edge is expected to be pretty significant across Automotive, across Industrial, and across Robotics. In Automotive, over the next seven years, you will see 500 million vehicles produced that will have AI cockpits and that will have anything from L2 to L4 autonomy. This was not the case if you look at what was getting deployed, what was getting shipped, over the last five years. 50 billion IoT endpoints by 2035. And over a million robots will get deployed globally. What is today a \$300 billion addressable market is going to become over a trillion dollars within the next decade. Right now, while we see Automotive and Industrial as very large TAMs, we expect this to invert where Robotics will actually become very large. This is the market that we are going to lead.

Our ambition is simple. We need to own the solution, the silicon, the software, and the stack for physical AI. We will build full-stack platforms where there is a wide space. Our operating mantra—and you've seen this—is we will win in Automotive, we will disrupt Industrial, and we will define Robotics. Let's get started.

We introduced our first generation of automotive products, especially compute products, 10 years ago. Today, we are one of the largest automotive, compute, and advanced connectivity players globally. Five generations of compute silicon delivered in 10 years. Today, from the first silicon to the start of production of the vehicle, we have brought that timeline down to 15 months. This is as fast as the consumer product life cycle. We now have over 500 million Snapdragon cars on the road. 90 million cockpits powered, and we only entered the cockpit business in 2016. We've launched 450 new car models since 2021, which is two new models every week for the last five years. The Snapdragon Digital Chassis is the underpinning of vehicle compute and connectivity globally. There is not a modern vehicle that is built without the Snapdragon Digital Chassis. We will exit fiscal '26 at \$6 billion in annualized revenue, and this is after delivering 23 consecutive quarters of double-digit, year-over-year growth. We built a \$65 billion design win pipeline. Our content value, from Gen 3 to Gen 5, has uplifted 8 times. We are engaged with over 70 automakers, and over 100 Tier 1s and Tier 2s globally. This is what true diversification looks like.

We are a systems company, and we have shaped the automotive industry and its platform architecture across hardware, across software, across compute, across AI, and across opportunities like ADAS. And we will do this consistently, over multiple generations. We are, ladies and gentlemen, on track to become the largest automotive semiconductor supplier across all pure-play automotive.

If you look to the right, we have built leadership in the cockpit because of all of the value, all of the IP, all of the differentiation that Qualcomm brings, and all of the breadth of access that we have to so many ecosystems. As we start to see AI come into our lives, AI is transcending the traditional domain architecture of what a cockpit is, and what ADAS is. So we designed for this. Gen 5 was designed, keeping in mind that we can't really be

traditional in our way of thinking. We designed for a mixed-criticality fabric. That means the customer can run cockpit applications and ADAS applications as they feel. They can run them separately and they can run them together. As we build more powerful chips, those chips necessitate customers to figure out how their architectures change. And this is creating tremendous optionality for customers to figure out how to design the next generation of vehicles.

With AI now, you can process any sensor input, across any specific vehicle domain. You don't have to tie the physical hardware to a specific domain. We are running 30 billion parameter models on the cockpit today. Commercial. Concurrently, we can run L2 to L4 stacks. And then, as you see on the left, what we used to call the software-defined vehicle, or SDV, last year or the year before, this has now become an AI-defined vehicle because we can now run agents, directly on top of the SDV, that get access to the vehicle context. We now have use cases where a car drives into a parking lot, sees a QR code, scans it, and pays for it. And that's an agent. This is how quickly the AI-defined vehicle is moving.

ADAS was a new space for us about three years ago. We didn't really have any customers to talk about either in the SoC space or in the stack space. We are now at 25 OEMs. The open-platform strategy that we have built with ADAS has allowed us to be able to provide tremendous optionality to customers. This is a complicated business. It's about safety, it's about cost sensitivity, and it's about which part of the world you're deploying in. We have a dozen different stack partners that we have engaged with, and we are building our own stack as well. Now, Snapdragon provides you with the best performance-per-watt-per-dollar capability across the industry, and that is why customers are moving to our platforms. In parallel, we have built the Snapdragon right pilot stack ourselves. We debuted this last year with BMW, and we are now validated in 60 countries. Stellantis is the latest OEM who has picked not only our Ride Pilot stack, but the entire Snapdragon Digital Chassis, which we will deploy starting SOP '28.

As Automotive is evolving, we see a tremendous amount of new growth opportunities. Robotaxis are things that you might have questioned, "When will robotaxis be real?" They're starting to happen. We do expect that, by the end of this decade, we will start to see these scale. And our strategy with robotaxis is actually very straightforward. Tony talked about HBC. We will actually build accelerators that will connect our SoCs and HBC Gen 2 to provide that same tech to our automotive customers. And we are planning to do that in the '28 timeframe.

We are also seeing this very interesting transition around token generators inside the car. As automakers are putting in so much of compute and so much of memory to be able to run models locally, we are receiving requests to see an offline mode. Could those be part of a federated use case for token acceleration? We will use the same exact HBC capability for token acceleration in the car.

The other area that we are seeing a lot of interest in is AI and ML use cases for the car, for the powertrain, for the drivetrain, and for battery management. There are so many domains in the vehicle that need local machine learning processing. We acquired a company called Edge Impulse a year ago, and we are actually running their MLOps locally in the vehicle. That allows us to use the Snapdragon NPU for local AI compute. Same exact chassis, no big difference; you just add more local processing capability. And, finally, we are starting to see a buildup in satellite connectivity needs. We've had a tremendous telematics and connectivity portfolio over the years. Customers are now looking to add satellite to that as well.

So, here is why we keep winning. We operate the full system architecture of the car. We support a global footprint. We are multi-generational in our silicon roadmap. We go across every tier. We go across every domain. We are compatible across generations, and we allow customers to plan a decade against our roadmap. We are building an ADAS L2+ stack ourselves because we know that this technology is going to be standardized across every vehicle. We have built the deepest and widest software and AI stack in the industry. We partnered with Google. We've integrated with Gemini. We partner with every global digital ecosystem, and we bring in OEM-preferred ecosystems. We've built years of safety expertise. We are a smartphone company, but we have made this transition. We have diversified. We now build safety as part of every single chip, every single piece of software. And even our stack and our tools are all safety-grade. And we've built tremendous supply chain complexity and resilience because we have become one of one in the automotive space. We understand supply complexity, capacity complexity, regional and geopolitical complexity, and we've been able to scale this business up really well. So big thank you to the Qualcomm team who's been driving this for years.

Automotive is a playbook for diversification for the company, and we've gained leadership, we've gained tremendous scale, and we've built a multi-generational strategy. And, as AI is upon us, we find ourselves very well prepared for that transition. Now, we've learned a lot about how to diversify through our Automotive experience. And I'll share with you what we have done in the last 18 months in the Industrial and Embedded spaces.

Before we get started—if you look to the slide—you know, the OT, or the operational plane, in any industry, in any enterprise, has traditionally never had to really do any processing at the edge. It was always about capturing information and sending that information to the cloud. It was mostly deterministic, it was mostly static, and the concept was to send data to the cloud and the processing happened in the IT layer. Now, as you start to see AI come into the picture, it's really the same concept. You have data at the edge and you have enough information at the edge to be able to process and get to a specific outcome, whether it's to detect an anomaly or whether it's to extract specific analytics. The operational technology plane is where these endpoints are getting deployed, and that OT plane is getting re-architected. Even brownfield settings are now being re-architected because you can add an intelligent AI aggregator. This re-architecture creates a once-in-a-generation opportunity for the entire OT plane to become more intelligent, with more data,

smarter models, better insights, and more intelligent endpoints. So, we started to build our roadmap for this specific space. If you look at what we have built with Dragonwing, over the last 18 months, you'll see a variety of different solutions that are very vertical focused. So we've dived very deep into what our vertical customers need, and we have built solutions based upon that. We have a silicon roadmap that addresses connectivity, camera, commercial processors, and industrial processors. These power everything from AI boxes, connected industrial gateways, edge appliances, industrial PCs, payment terminals, smart home appliances, drones, and body cameras. And we've picked three vertical categories that we focus on: industrial, commercial and mobility. And these are further segmented across 12 verticals. So we are building for every device class, every connectivity standard, and every stack layer.

Let me give you an example of what we have done in the vision space. So we believe that vision is a major unlock in industrial. We've had tremendous expertise in the company from our smartphone heritage and from our automotive heritage, and we are building industrial machine vision, we are building robotics, and we are building surveillance as three additional layers of capability. Now, vision AI is a major unlock because the physical world is best processed through the lens of a camera frame. We understand lighting conditions and how to improve them. We can reconstruct scenes in 3D. We can semantically annotate these scenes. We can feed them to a VLA so that it can tell us what the artifacts of interest are. We can even predict what the next scene needs to look like in a situational awareness scenario. And so we've built an entire video AI stack from camera chips and edge AI boxes to on-prem appliances and a full video AI service. And we are deploying this across every vertical: retail, small and medium businesses, smart cities, and venues. Any use case. All verticals. We believe video intelligence is a major edge play.

The other area that we had to spend a lot of time on was figuring out how we can become much more developer-centric, much more developer-first, as Cristiano mentioned in his remarks. And there were two problems. How do you simplify access to the product? And how do you accelerate the journey from prototyping to commercialization? As Tim mentioned in his remarks, Modular is going to be a significant game changer for Qualcomm because we will now be able to write and serve models faster. We will also be able to build application-specific acceleration in the libraries that Modular is going to bring, which will make our entire AI stack much stronger. Over the last 18 months, we made three additional acquisitions to be very developer-centric. One was Arduino. Arduino brought us 33 million developers and a massive global footprint—completely open source—which allowed us to be able to get access to pretty much every vertical out there. We acquired Edge Impulse. That allowed us to be able to get model training and tuning, and containerized development of models at the edge. And we acquired Foundries.io, which allowed us to be able to manage industrial-grade Linux. So developers now have the ability to prototype rapidly on Arduino and Dragonwing, and we can scale these projects across standardized, System-on-One modules or chip onboard capabilities across all of these various use cases that I mentioned.

Last October, we launched Arduino UNO Q, which was the first Dragonwing product on the Arduino ecosystem, to tremendous success really across every vertical, every type of developer, and prosumer market. We are able to launch VENTUNO Q in August this summer. 40 TOPS of AI, an 8-core, 12 cameras, a safety island built-in, and built-in AI models—all part of this overall ecosystem that we talked about. It'll focus on Industrial, Consumer, Embedded, and, of course, Robotics. This runs full upstream Linux. We've not been a company that has been on that path. We are now a full upstream Linux company, and these will all run out-of-the-box. So all of the goodness of the Qualcomm platform available in full upstream Linux. We are also working on agentic development. So, you can essentially take these development boards and code directly with Claude, with Codex, with Cursor, and you can start wipe coding. And you can buy these on Amazon. Qualcomm is changing, and we are becoming massively developer-centric.

So while we were pulling together this product portfolio, this developer centricity, and this video AI, we were also, in parallel, building our focus on verticals. We have 12 different verticals across three major industry types, and I'll give you a sense as to what we are doing in three of them. Retail, energy and utilities, and oil and gas. But we are building blueprints so that they are repeatable, so we understand what the OT blueprint is that we need to be able to replicate, and we are leveraging the entire Dragonwing portfolio for this. We've even built a solutions engineering team for Dragonwing. Let me step through a few of these quickly.

In retail, the store is evolving. It is continuously sensing; it is making decisions by itself. It has to act because stores are now hybrid. Some stores in the evening are almost not manned. They need to have more autonomy. There are AI cameras that are getting added for loss prevention, for shelf intelligence, electronic shelf labels, RFID for dynamic pricing, AMRs and robots for automatic restocking, for removing products that have expired, and for cleanups. And we are partnering with key partners in retail, like Vusion Group, to be able to drive this expansion. Similarly, in energy and utilities, Qualcomm addresses every step of the value chain: generation, transmission, and distribution—the entire grid. We go after sensors, industrial gateways, fixed cameras, and meters. In our hometown, San Diego, we work with San Diego Gas & Electric to help mitigate wildfires, with autonomous drone inspection to execute automatic power shutoff. Schneider has been a great partner of ours. We've built with them the capability to be able to deploy industrial gateways as part of their substation concentrators. In oil and gas, we have built solutions for upstream, midstream, and downstream capabilities. We are proud to have worked very closely with Aramco and with Schlumberger to bring connectivity and edge computing to very high-risk environments in vital industries. From autonomous drones to well-drilling operations and monitoring the safety of workers, there is a tremendous amount of complexity in these industries. And the results are showing: our indirect revenue is up 77% from '24 to '26. We have tens of thousands of unique customers, and we are working to grow thousands more. Over 200 hardware and tech solutions, more than 35 leading distributors, and 45 global GSIs. Our partners span all verticals, and we address every step of the value chain. The channel has become the multiplier of what we are building.

To summarize, AI is, and will continue to, re-architecting the operational plane, and that disruption is going to create this upgrade cycle across billions of endpoints, which is a massive market opportunity for us. And, in the last 18 months, we have rebuilt our entire product portfolio, our developer platform, and a vertical go-to-market. We have purpose-built silicon for software and AI. We have a clear prototype-to-commercialization path across 38,000 customers. And this is a repeatable, full-stack blueprint across industries.

This business has really helped us in understanding very deeply, as we take on new challenges like Robotics, how to diversify the company. And, really, all of these build up on each other. All these learnings compound. The Automotive and Industrial businesses have created a lot of focus on how to ramp the teams to be able to go after Robotics. We understand how to do safety. We understand how to get into vertical-specific markets. And Robotics requires four key building blocks: computing at the edge, connecting the edge, enabling high-performance AI hardware, and orchestrating intelligence systems. And all of this builds on our collective experience, knowledge, and investments from various Automotive and Industrial initiatives.

Now, robotics is where embodied AI gets physical. So the objective is to perform human tasks, which would include mobility or emotion, perception and reasoning, and actuating or manipulating in the physical space that is around you. So these are systems that have to sense, they have to think, and they have to act. This is, in our mind, at least a trillion-dollar opportunity over the next decade. And it requires a very broad set of technologies, products, and experiences that no real, general-purpose chip maker has today.

Before we get into what we are building, I want to describe to you what robotics will do. And, really, you can think of this as a time continuum. So embodied AI implementations will encounter many tasks that will have varying degrees of complexity across the mobility, the actuation, and the intelligence domains. Starting with inspection, where mobility is the underlying skill, tasks are going to include: reporting status, visually documenting, measuring, and surveying the real world. Next is transportation, and the movement of goods, tools, packages, and even people. After that comes the interaction with the physical world. Core skills to start, and then finer, more precise, more dexterous skills, like pick-and-play, sorting, assembly, and insertions. Some can be single-chart; some can be longer-horizon. This builds up to multi-agent fleets: teams of robots that are working in partnership, in coordination with each other across multiple use cases. And then, finally, robotics comes to the home with consumer interaction, which will require tremendously high levels of safety and testing trials. And each of these steps adds more and more unique value. We have silicon shipping in every single tier today. And let me share with you what we are doing in this space.

So as we have done in pretty much every market that we get into, we always look at a full-stack approach. We want to be able to make sure that we capture value across the stack. We want to be able to make sure that we can drive the pace of the inflection point if a

market is ready to mature. We focus on six layers: the compute, the next generation operating system, the simulation, the data pyramid, the data flywheel, the models, and the hardware reference design. And as we have learned from our Automotive experience, we capture disproportionate content value, and we build tremendous systems expertise, whenever we take a full stack approach. A robot is not one computer, it's three computers. They work in concert and they are an in-tandem hierarchy. System 2 is the reasoning brain. It's the cerebrum. It's the heavy, mixed-critical AI workloads that require deliberative thinking. System 1 is the action layer. It's the layer that plans the motion. System 0 is executing the motion. It is your reflex system. It's the millisecond control. It's the highly decentralized part, what you call the nervous system. And the main takeaway is this is a heterogeneous compute problem statement, something that we understand a few things about. And this requires you to have this optimal balance of distributed thinking, planning, and real-time reflexes that require the appropriate compute engine. We've solved similar challenges in the history of the company. Most recently, with Flex and Automotive, which is, in our mind, another association of physical AI. And, to my knowledge, we are the only company that is architecting across all three domains.

I'm going to get a little bit technical. The hierarchical compute architecture comprises of the thinking brain—with central compute—of several limbs and joints—with their own local compute and split-second motor control—and intelligent sensing for end effectors. We are building this entire system from brain to fingertip. The Dragonwing IQ10 is our central compute SoC. It's purpose-built robotic silicon, and it is already commercial. The perception IP allows us to visualize the world around us across multiple, context modalities: the motion control IP for trajectory and balance, the actuation and control IP at the server motor control level, wireless and wired IP for time sensitive networking, and, obviously, always-on sensing. We've taken this chassis mindset that we adopt from our learnings in Automotive, where we start off with a specific area that we are good at, and, over time, we expand to get to a system-level focus. This platform-level thinking allows us to think about the embodiment, which is a unique differentiator that Qualcomm has.

One more, maybe last, complicated slide. Second to last. Let me orient you to this slide. If you look to the robot to the left—and I used this as an example when Cristiano and I were discussing this—think about the concept of a robot picking a jug of water and pouring it into a cup. As it is doing that, the weight and the shape—assume it's a paper cup—the weight and the shape of the cup is going to change. And that requires the robot's hand to sense, and to adjust its grip pressure. In real time. This is decentralized. This is exactly what you will feel as a human. That is the complexity of a robot. When we talk about bringing a robot into the home, just think about what level of complexity you are talking about introducing technology into. And that is why this is a longer gestation cycle time period. Three systems are active at the same time. The brain controller, System 2, is highly perceptive. It is aware of the kinematic shape of the environment. It knows its degrees of freedom. It knows the range of motion. It's responsible for balance control. It can identify the cup. It can identify the water jar. The body controller, System 1, controls the limbs. It controls the movement of the hand that goes to the jar, it knows where the cup is, and it

actually takes that motion. It coordinates that movement. System 0 is actually able to sense the grip pressure, the tactile feedback, the temperature, the moisture, and the weight, which allows it to act reflexively. Across these different systems, we have multiple, real-time local loops that run within a system, and we have slower loops that are running across systems. And we are building embodiments across all these three different systems. We are also building a full software and application stack, complete with raw support and SDKs for manipulation, to write to various types of sensors and effectors. We will ship sample applications, and we can place robotic arms, office scout applications, AMR for navigation, and follow-me applications. And this is open to every developer ecosystem, including the Arduino ecosystem that we just enabled.

The other aspect of robotics development is the simulation, data, and training model flywheel. We are building this environment in-house. We are building, as you can see on the left, a simulation platform, where, before a robot ever touches the real world, you have to train it in a virtual world. It has to be aware of the physics, the sensors, and the rendering as to how that will take place in the real world. This saves you tremendously in terms of the physical involvement of trial and error. Then we have the data pyramid. That's the fuel for these systems. We combine real-world data that Qualcomm has access to, synthetic data that we generate ourselves, and a lot of open-source data. And then, on the right, we train the foundation model. This is a single model that can take multimodal input—like vision, like depth, like touch, and natural language—and it generalizes across use cases. We train these models with simulators, behavioral cloning and teleoperations, and reinforcement learning. So the workflow is end-to-end. We build the hardware, we generate the data, we develop the models, and we deploy them into the customer environment.

We announced our IQ10 reference design at COMPUTEX in June, and this is purpose-built silicon, which is shipping today. We also have IQ9 and IQ8, for simpler embodiments, and we are already designed into the NEURA robots, which you can see outside in the demo area. With NEURA, we offer a complete reference design that powers the NEURA MAiRA cognitive robot arm, as well as the 4NE1 humanoid. These robots are trained in the NEURA Gym, with the robotics foundation model, and they run the Neuraverse application platform. This is a full-stack. Silicon solutions with a key customer, a key partner, in less than six months.

Today, we are powering every type of embodiment, and several are shipping already. We have over 100 engagements, spanning the entire robotic stack, with companies like NEURA, Figure, and KUKA. We are working with several drone OEMs and many AI sensor and embodiment partners. With physical AI upon us, IQ10, the robotics reference design and the end-to-end solution stack, ensures that customers and partners design with us. We are taking the same approach that has allowed us to scale very quickly in other businesses. Robotics is already a reality at Qualcomm, and we are very excited to be powering this next generation of physical AI, where we believe we are very well positioned.

To conclude, a few takeaways. Automotive—I hope you are all believers—is now a track record. We've had 23 consecutive year-over-year, double-digit quarters of growth. We don't expect to let you down anytime soon. \$65 billion in the design win pipeline. We have delivered, we are still accelerating, and we are on track to becoming the largest automotive semi-player globally. We are now a category leader in every domain we enter, and that's not easy to do. Industrial and Embedded IoT is now scaling. 18 months in, we have built a product portfolio with Dragonwing. We've built developer muscle with three acquisitions—four, with Modular—and a full-vertical stack that goes from silicon to solutions. And Robotics is happening now. It's already shipping. Dragonwing IQ10, IQ9, and IQ8 are all in production. Partners are integrating them into every environment, from humanoids to quadrupeds, from cognitive arms to AMRs and drones. Three industries, one IP foundation, and one physical AI platform. Thank you very much. And, before I turn it over to Cristiano, I would like to play a video from one of our partners, David, from NEURA. Thank you.

David Rigger, CEO, NEURA ROBOTICS [Pre-recorded]

Hello everyone. My name is David Rigger and I'm Founder and CEO of NEURA Robotics.

What we do is we build robots and enable them to have the cognitive abilities to see, hear, feel, think, and react fully autonomously on all kinds of physical tasks, like humans do. The benefit of working with Qualcomm is that we're giving the robot more than just the brain. What I mean by that is today we're seeing that physical AI is mainly seen as a vision-language action model, but it's actually much more than that. It's a little bit similar to swimming. You can't learn swimming by just your brain and vision. So it means you cannot just watch a video and think you are a swimmer. You learn to swim by simply going into the water, trying it out yourself, and then actually training your memory effect of your muscles. You're training your reflexes and nervous system how to breathe, how to move your body, and how to actually stay above the water. These tasks always require more than just vision. They need the feel of touch, they need to hear, and combine all that to build the foundation model, which can actually do all the physical tasks on this planet. We built the physical AI platform, which we called the Neuraverse. This is a deployment platform where everyone in the world can actually train and contribute to make this one brain smarter. What you can expect for all the partnership is NEURA and Qualcomm are basically setting a new standard in physical AI, enabling every robot on the planet, and every human on the planet, to actually train the robots and enable them for all kind of physical tasks on our platform, Neuraverse.

Cristiano Amon, President and Chief Executive Officer, Qualcomm Incorporated

All right. So, I got the last part of the presentation, I think, before Akash will walk you through the financials.

And, I think, before I start, what I'm going to tell you is what is unique about Qualcomm. I know we have a limited amount of time, but there's a lot of new vectors in technology, and, hopefully, you'll be able to see that. It's not only about one solution in the data center, but it's also about what we're doing Automotive, what we're thinking about in Industrial, which

is a whole different industry, and this field of Robotics. You need to have the breadth of semiconductors technology that we have. And I think that's an opportunity for Qualcomm.

With that, I'm going to talk to you about the future of mobile edge devices. I'm going to try to unpack a couple of different trends that are going to happen, as we think about the role of agents. I said it in the keynote at COMPUTEX: the event of agents and orchestrators was a very significant milestone that actually provided clarity on how those devices are going to evolve. And the industry tends to think in binary terms. It's all of a sudden, everything is going to stop, and it's going to go do the other thing. No, they're going to coexist. But devices are going to have different types of use cases, and that's what I'm going to try to unpack in this presentation.

The first one I want to talk to you about is what we all have been used to with mobile: the smartphone is at the center of your digital life, and everything is around that smartphone. The app store becomes the control point of the OS, and the apps understand human intentions, and everything is around the smartphone. Even other devices, they're just an extension of the smartphone. That's not the case anymore. Actually, every AI company, every foundational model company, now talks to us. The devices are the end points for agents. That's where humans are, and the agent is at the center. It's not about the phone at the center anymore. The agent, for the agentic experience—once you understand human intentions—is at the center, and devices are just endpoints of the agent. And the purpose of my presentation, right now, is to tell you how the device is going to change because of everything that has happened.

Let's start talking about the user experience. So those devices have been built for the human as the user. So the workflow is based on humans going to an app and doing things at human speed. But now the device has the orchestrator and the agents, which are going to do other things on behalf of the human. So the agent is now going to operate the device, and we're starting to see that right now. If you ask me, "Where's the epicenter of new agentic use cases?" It's actually happening in China right now. And you started to see the agents go to your device, operate the device for you, and go to the web for the agentic experience. That tells you that the device now has two use cases.

It was interesting, I said this before, and I'm going to repeat it. When people want a computer to run OpenClaw and they have the computer running OpenClaw. Once you start having that experience with you—not just the amount of software developers that exist in the world, but the six billion people that have smartphones—when you started to use agentic experience as part of your experience and interaction with the device, you're not going to carry the computer. It's all going to happen on the same device. And the device is going to have two users because there's two different workflows: there's you and there's agents. So, that's one big change.

The other big change is perception in sensing, and that's what is actually changing the device in itself and enabling different endpoints, like personal AI devices. So if we now

have the computer that interacts with us, the way we interact with each other, then the context—that we are inserted in, especially as we think about how we communicate with audio, how we speak, how we listen, how we see, and then how we are integrated into our surroundings—becomes important. So you now have a lot of sensing data that needs to be part of you know the processing of those devices, and that is enabling also a different class of device. That's why I told you about what happened to wearables, and what happened to augmented reality, mixed reality, and virtual reality. The transition is very important. Because the reason we've been very focused on glasses is because glasses are close to our senses—they are close to our eyes, to our mouths, to our ears—and those devices, wearables, are an extension of the smartphone. When the smartphone is in the center for the agentic experience, they actually become an endpoint for the agent. And the stuff that you actually wear becomes very interesting. And especially when you think about glasses. Just as a side comment: at Microsoft BUILD, Satya announced Project Solara, a badge with a camera. It's a new class of devices. We have 40 different designs today, with some of the largest AI and model companies in the world, thinking about those new form factors. And there's one form factor that we know is going to get scale: it's glasses. And the use case is "see what I see" and "hear what I hear". And I'm going to come back to the use case because I'm also going to talk about another change that's going to happen on devices at the edge. But now what I'd like you to do is this. New companies are looking at this big change. The devices now are endpoints. The barriers to entry, for OSs and app stores, for new experiences are no longer the same. You're going to see a lot of excitement. I want to start by showing a video from Amazon.

Panos Panay, SVP, Amazon [Pre-recorded]

We're in the middle of a fundamental shift in computing right now, and what's most exciting about this next wave of AI is what it unlocks for people, for their creativity, their curiosity—the things only humans can do. Now, as technology continues to fade into the background, the customer must stay at the center. That shift creates new requirements, not only for the devices already in people's homes, but for entirely new devices built for AI-first experiences. At Amazon, we're building for this future with Alexa, creating experiences that work seamlessly for a customer when they are in the home as well as when they're on the go. Qualcomm is one of our critical partners in expanding both the capabilities and reach of these experiences, so AI can seamlessly move with people throughout their daily lives. This relationship is about building what comes next together.

Cristiano Amon, President and Chief Executive Officer, Qualcomm Incorporated

So there's a lot of exciting things coming. And, by the way, I'm actually so grateful. I have a relationship with Panos that spans decades. He's an incredible, individual visionary, and I'm very grateful for the partnership.

So the other thing we talk about is the orchestrator, and, as you heard today, those agents, they generate demand for a lot of tokens. The reason a lot of the hyperscalers just see a wall in front of them of compute demand is because the economics of AI are fundamentally changing. Agents and orchestrators, that's why I said that OpenClaw was

an incredible milestone. They're redefining the architecture and economics of AI. Not only creating an entry point for Qualcomm into the data center, but actually creating a fundamental change in the architecture of compute that touches all the devices on the edge. And it's an order of magnitude increase. If you look at how we started, with conversational to now agents, you see the order of magnitude increase. The projection is 40X, the increase in annual token demand between 2026 and 2030.

Now, I'm going to show an example to you, and I'll tell you what you're going to see. We've been doing a number of these things. And you can try it yourself, you can try the different prompts, and you're going to see that it changes. Sometimes you get 10%, 20%, 30%, 40%, or 50%. I just picked a very simple example for you to see today. We've been showing this. So what you're going to see, right now, is that we've got two computers. Those computers have orchestrators. You give them a prompt, a complex prompt. You ask them to do some research, design a webpage, and put up the results. One we're going to be using is just Claude. We're using OPUS 4.6, 100% on the cloud. You see the thing working. The other one uses smart routing. You use some models that are locally installed into the machine, and some others in the cloud, and you get those things to work. And what you see, at the very end, is what hybrid AI really means because you will get to exactly the same outcome that you want. But now, when you think about things, like the mix of experts, when you think about different kinds of models, you can actually see how the architecture of AI is evolving. The reason you saw some foundational model companies saying, "I am going to give up on doing video creation." It's very obvious because you can use your compute capacity to monetize tokens of higher value. And that's what we're starting to see. Actually, when Microsoft said at BUILD that they are now creating unmetered intelligence, people really understood what an AI PC is, and that this is happening everywhere.

And the beauty of these devices, like phones, is that it's actually happening on the same device with the different users. And you can see a lot of useless debates. I have seen over the years about, "Is this edge or is this is cloud?" It's actually the wrong conversation to say, "I have something that I need to do on the cloud. Can I do it on the edge and vice versa?" That's the wrong approach. Things that are going to be done in the cloud are going to be done in the cloud. The growth of the cloud is incredible, but the edge now also becomes computing that is going to generate tokens. And, I think, that's how the industry is naturally going to evolve. And the message here to you is like what's happened in the data center. Inference is actually becoming disaggregated and distributed everywhere. It's just a new form of compute. For example, when you look at the architecture, you're going to have the cloud data center, which has hundreds of megawatts. This diagram, I'll bring it back to you. You'll understand that. I'm going to overlay two diagrams that are going to make it very interesting. But you look at the data center, with hundreds of megawatts—two gigawatts—and you're going to have regional data centers with tens of megawatts. And on-prem, there's a lot. If you look at results of some of the server companies like that, there's a lot of movement for on-prem right now. It's not against the cloud. They're both growing, and you also see that as you have more of a hybrid AI. You're going to see more computing happening on PCs and devices at the edge. So inference becomes distributed. Just don't

take my word for it. We have an incredible partnership with Google, and I want you to hear from Google about that.

Rick Osterlow, SVP, Platforms and Devices, Google [Pre-recorded]

Hello, everyone. I'm Rick Osterlow, Senior Vice President of Platforms and Devices at Google. And I'm incredibly excited to talk about our longstanding partnership with Qualcomm. As the industry shifts towards agentic workflows, we're moving beyond simple responses to offer true digital agents. These are proactive, multi-step systems that seamlessly anticipate user needs across your entire ecosystem of devices. To bring this to life, our teams are focused on bringing a shared, full-stack vision. We're combining Google's advanced Gemini models and Android system-level intelligence with Snapdragon silicon. Together, we're innovating side by side to scale on-device AI and Gemini Intelligence to the most advanced mobile devices. With AI, we're redefining next-generation automotive experiences, and pushing into new frontiers with wearables, like XR glasses and intelligent eyewear. And this innovation is also at the heart of our brand-new GoogleBooks effort to revolutionize the laptop experience. True agentic workloads require what we call distributed intelligence. By efficiently balancing processing between the cloud and on the device edge, we can deliver seamless experiences that are private, instantaneous, and personalized. Our deep engineering collaboration with Qualcomm ensures that the broader ecosystem and our OEM partners can scale these innovations rapidly. Gemini Intelligence will elevate the Android ecosystem, making it smarter, more proactive, and more helpful than ever before. Cristiano, congratulations on Investor Day. Thank you for our outstanding partnership, and I'm really excited about our road ahead.

Cristiano Amon, President and Chief Executive Officer, Qualcomm Incorporated

So big thank you to Rick, for the partnership and the confidence of what we're going to do together. So I have to pick up the pace now and I'll get to my last part of this presentation, which is about 6G.

So we have been designing 6G for this AI era, and I'm going to now show the role of 6G in this conversation we're having, and I'm going to go fast. So I just thought about new classes of devices, like glasses. So the goal, of the connectivity of 6G, is to transform all of us into walking cameras in this world. I need a very fast, high-definition video. The opposite of what we did with 5G, which enabled streaming of high-definition video. We're going to do that for uplink and across the cell site so everyone has the ability to stream high-definition video for "see what I see." That's very important context for agentic experiences. That's what we're going to do with connectivity. I'll go through the details. There's a lot of improvement on the connectivity side, but that's only one third of the story. The other part of the story is the computing part. And the computing is going to be required because the 6G infrastructure is no longer dedicated equipment for communications. The network of 6G—actually, you need to be thinking about how that was a network that was designed for voice. It's now transporting bits, and it's going to also generate tokens. And the reason that's important is because of the premise of how 6G is going to deal with radio frequency. It's going to look at radio, and, for instance, physical AI, and you're going to need compute.

Look at the architecture. It's kind of the same as the **distributed**. You're going to have a big data center, and you're going to have a regional data center: that's the core of the network. You have to have an edge data center. You're going to have the cell site, and then devices at the edge. And that's why we're also building our data center solution to be scalable. Because when you think about 6G, it becomes one sovereign AI workload, and some operators will actually be selling token generation machine capacity, like a CSP, for AI companies for this distributed compute. And, once you have all this compute, it brings the next part of 6G, which is sensing. Because every single RF is going to be treated as a radar, and you will use models trained on that RF performance and the RF radio characteristics. You're going to be able to sense everything. Drone detection is top of mind. It becomes a critical infrastructure. Everything that moves and flies becomes an important context for models, and that's part of the perception.

So with that, I'm going to summarize what's happening with devices at the edge right now. They're evolving for agentic experiences. They're going to have more than one user, and we're building the next architecture of those devices. You now need a CPU for the orchestrators. The orchestrator is going to operate your device for you. I think everybody now understands that the CPU matters, and it's important. You need to have a different architecture on inference because you have to have very high-performance, low-power inference, even when you are not using the device as a human. And the high-bandwidth compute that we're doing for the data center; we're building a version of that as a co-processor for phones. There's a complete change of the perception and sensor, as well as a new modem, and it's not just about just the phone. It's different classes of devices. And that is also bringing other people to the space.

So I'm going to finish this part of the presentation by asking Hark, a new company also founded by Brad, to show you what they're doing.

Brett Adcock, Hark Labs [Pre-recorded]

I founded Hark because today's AI is just not good enough, and nobody's taking real advantage of it. It constantly forgets things about you, and it runs on devices built a long time ago. So, at Hark, we're an AI lab building the world's most advanced personal intelligence. I think the next AI platform is something that truly knows you, that can see, hear, and act in the world with you. And you don't get there by bolting AI onto existing devices. You do it all together. The models, the hardware, and the interface is one product. That's what we're building at Hark, an intelligence that things like you and sometimes ahead of you. We're grateful to have excellent partners like Qualcomm who are helping us realize some incredible ideas. We're excited to release the Hark platform this summer, and then the next generation of consumer devices after.

Cristiano Amon, President and Chief Executive Officer, Qualcomm Incorporated

The obvious question that everybody's asking, "What is the timing?" It's hard to predict the timing right now. And I think the mobile market is dealing with the uncertainty of the memory situation. But it's very interesting. We are incredibly encouraged with the design

activity and the number of new entrants. When I talk about China, I am in a situation right now where I don't know who the mobile customers are anymore because there are OEMs, but every single AI foundational model company building agents is also a customer. The surface area is tremendous. There's 6 billion phones, 2 billion personal AI devices, 2 billion PCs, and 500 million cars. And when you think about this new form factor, I'll just give you an idea of how you should think about glasses. We're just at the beginning of this. 600 million glasses are shipped annually. There's less than 1% market penetration for smart glasses. We're building a very small reference design, and you can build this to any glasses. You can see, you can have the ability to access an agent for audio, multimodal, or premium display, and you can see the smart E-BOM. So that is a great opportunity, and that's what's also going to happen on the mobile devices business for Qualcomm. So I hope you saw that there are a lot of interesting trends in technology. And, as I get to the end of my presentation—and I think you're all eager to see what Akash has in store for you—I am going to talk about this number three pillar that we've been talking about in this presentation. About the next chapter of Qualcomm.

From silicon to platform solutions, and to really building a fully integrated platform for hardware and software, while also changing the company to a developer-first mindset. And, you saw, we're doing this with all the new business we build on the edge, and we're going to be doing this for everything we have on the compute continuum. That's why this acquisition of Modular is so significant for Qualcomm because it also builds on the pillars of the Qualcomm advantage. Not only the technology, but the focus on deep customer partnerships, the creation of an ecosystem, and the scale. We have proved that we can partner across the industry, and it's never the role of one company to innovate, though that creates an incredible opportunity.

Before I bring Chris up here and he's going to talk to you for a few minutes, I am going to say—and I may be a little bit aggressive saying this—but I'm going to say I'm willing to bet that not everyone here will understand what we're trying to do. I think some of you will understand, and this is not a negative comment in any shape or form. It just takes a while to understand how we've been thinking about this. We've been starting this journey with Modular for more than a year, and it took a lot, I think, for me to convince Chris and his team—and he'll share his story—but I think we may have an Android moment here. And, maybe, I'm going to be so bold as to say that maybe there's even a Linux moment. I don't know, but I think we have something good. And I think we have momentum. As AI goes everywhere, as compute becomes distributed, as you have every single endpoint become an endpoint for agents during inference, and you have an industry that wants an open ecosystem. Maybe that's what Qualcomm can do: support everyone.

With that, I would like to bring to this stage a legend, Chris Lattner. I thank the founder and CEO of Modular, who will tell you about what we're going to do. Chris, please come on stage.

Chris Lattner, Co-Founder and CEO, Modular

Thank you, Cristiano. All right. Well, thank you, Cristiano. By way of introduction, I spent my early career building today's software platforms. This includes the compiler technology that runs every phone—no matter whose vendor is in your pocket—the data centers that span all the hyperscalers, I built the Swift programming language at Apple, and I also built the software stack that powers Google's amazing TPU AI Scale Platform. Now, I decided to join Qualcomm because I found that the team recognizes something. They recognize the opportunity of AI today. They have the ambition to do something big, but also, they've made all of the investments already that put them in a perfect position to do something about it. Today feels familiar to me. It feels a lot like back at Apple, when it was about to take off. Back then, there were a lot of doubters. People did not really understand what was going on, but we had all done the formative work already. And so, all we had to do was get people to see it through the products and their lives. So let me walk you through what I see today.

So it turns out that compute has fundamentally changed. You've heard about that a lot today. It's no longer about a single chip. Compute today is a large-scale data center distributed systems problem. We all need to program diverse AI accelerators from multiple different vendors. We need to get the best performance, the best TCO, and we need usability because doing all this is harder than it's ever been before. Now, the world is still struggling to get individual systems to compete with the industry leader, but that's where Modular comes in. At Modular, we spent the last four and a half years building a novel platform that actually scales. All with the goal from the beginning of unifying the industry and opening a new chapter for accelerated compute. Now, we built this platform to scale, across a full spectrum starting from the data center, but then going all the way down to the edge.

And so this is why I'm so excited that Modular is joining Qualcomm. We're bringing together the perfect combination of scalable hardware and scalable software. This is joined with a shared ambition from both teams to make a world that is better for everyone with AI. And I got to tell you a little bit about where we're going. So, now let us remember that everybody wants tokens, but they also want amazing economics. Now, most people don't really want to know how it works. The systems, the software components, all the different pieces that go into this, are amazing, and, for nerds like me, I love it. It's great. I think I'm in fellow company here, but a lot of people just want a solution. And so, together, we're lifting the Qualcomm AI silicon business. The silicon no longer. Now, it's about solutions. Full it-just-works solutions and an AI solution business is far more valuable. Now, this end-to-end solution approach starts in the data center, of course. That's our core focus, but it won't end there. This is the first software platform that was designed from the beginning to unify edge and data center, utilizing diverse accelerated compute in all the crazy form factors that are pervasive in our lives.

Now, we've been using a lot of operating systems over the course of the last decade, but this platform will grow into a full operating system, built natively distributed, natively accelerated, and agentially native by design. And we are not just building this for us. We're building an open developer platform, enabling AI developers, AI researchers, and

app developers to innovate like never before. Together as an industry, we understand that we need to scale incredible amounts of compute. We know that this will require many different organizations to come together to make it possible. The result of doing this is an incredible opportunity for everyone, all of our partners, millions of developers, and, of course, all the consumers will benefit from AI products in our lives. Now, throughout my career, I've had the privilege to drive open standards. I've built several very large-scale, open-source communities. I've been part of multiple waves of compute, and I feel that today Qualcomm is really quite well positioned to be the best in the industry to lift the entire world, and, as such, we're committed to this being an open platform. Qualcomm is obviously an incredible hardware company, but today we're going further. We're now a full-stack vertical AI solution company. I couldn't be more excited to build together. So thank you.

Cristiano Amon, President and Chief Executive Officer, Qualcomm Incorporated

Thank you.

Hope you share our enthusiasm about this change in the company, this opportunity to change in the industry, what we're going to do together, and we're super happy to have Chris and the rest of Modular as part of the Qualcomm family.

So, I was about to end here, but I thought maybe there's one more thing.

Hopefully you are entertained, but we have one more thing, and that's just building on this. We're also very happy to announce a very strategic partnership that we're making with Hugging Face. And I'll tell you about this partnership. Qualcomm and Hugging Face started a very unique collaboration because they share exactly the same vision we presented to you today. For the data center, Dragonfly, Hugging Face will be very focused on demand creation from Qualcomm Dragonfly silicon. Both their inference and storage services will map to all the Qualcomm Dragonfly products. There will be agentic model onboarding, combining Hugging Face, their 16 million developers, and what we're doing with Modular, across the whole family of Qualcomm chipsets—from Snapdragon to Dragonwing and Dragonfly. All models are going to be onboarded on Qualcomm technology platforms using an agent that is going to handle the setup, optimization, and deployment with zero manual integration work. And then, just as you heard from Chris, we will build on an end-to-end agentic AI with distributed intelligence. And that distributed AI framework is when agents can operate similarly across the entire compute continuum, leveraging Qualcomm technologies models and cooperating with everyone. And I want you to hear from Clement.

Clem Delangue, Co-Founder and CEO, Hugging Face [pre-recorded]

Hi everyone. I'm Clem, Co-Founder and CEO of Hugging Face. If you haven't noticed recently, something big is happening in AI. More and more, the world is running on open source and local models. And for good reasons. They're way more affordable than the big LLM APIs, more customizable by companies. And because they run on your own device, they're private by design. Your data stays yours. Today, over 16 million AI builders create

this future in the open on Hugging Face. And that's why I couldn't be more excited to announce a new collaboration with Qualcomm. Together, we're going to make open models easy to run everywhere, from a device in your hand to a full rack in the data center. Snapdragon, Dragonwing, and Qualcomm's Dragonfly cloud, all powered by the open-source community. You'll be able to take any model—big or small—deploy it, optimize it on any Qualcomm platform, with agents running on device and orchestrating across the cloud. We'll also offer Hugging Face pro subscriptions to many developers using Qualcomm platforms. Local, private, and affordable for everyone. That's the future of AI we want, and we can't wait to build it together. Thank you very much.

Cristiano Amon, President and Chief Executive Officer, Qualcomm Incorporated

So that's it. I think I got to the end of the presentation. Hopefully we gave you an opportunity to understand, I think, what Qualcomm's going to do within the next five years and I'm going to summarize it to you.

Data center will add a meaningful new vector of growth. And I think what you saw is, when we originally talked about this, we talked about how we're building a data center portfolio, and we expect revenue to be in fiscal '28. Then, we get more traction, and we move it to fiscal '27. We get more traction, and then we get to fiscal '26. And you heard from Tony, we're just starting.

Automotive, Industrial, and Robotics will extend Qualcomm in the next frontier of physical AI. We'll build the platform. We have the market scale, and we're executing on all the technology trends. Agentic AI at scale will drive an upgrade cycle across edge devices. They are going to be machines, they are going to generate tokens, and they are going to be interacting not only with the users, but they're also going to be interacting with agents, and that's going to happen across the entire industry. The token economics will make distributed inference inevitable. We're excited about the growth in the cloud. That's what creates the opportunity for us to enter into this aggregate, and that will continue. We're just at the beginning of that, but everything will become, you know, an AI computer, and, I think, that is going to fundamentally change and create a massive opportunity for us across the compute continuum. 6G will be foundational infrastructure for the age of AI. And, if you haven't forgot, we have that asset, too. And we're going to be expanding beyond silicon to a full-stack software platform with the most industry-friendly, I think, platform. And, at the end of the day, I think, we have seen in our industry, open horizontal systems will win. And, I think, that's kind of our bet.

So with that, thank you so much for listening to our presentation. And now, I think the main attraction of today, our CFO, Akash.

Akash Palkhiwala, EVP, Chief Financial Officer and Chief Operating Officer, Qualcomm Incorporated

All right. Good afternoon, New York. It's incredible to be here. Lots of familiar faces. Great to see all of our investor friends here in the room as well. It looks like all of you decided to

stay here rather than go to an earnings call. That's a great decision. That's a great decision. We're going to make it worth your while. This is the climax of the show. So, we have closed the doors now. You're stuck here. You'll have to listen to the rest of what I have to say. Just kidding aside, you heard Cristiano, Tony, Nakul, talk through all the great stuff we are doing across our businesses. And so my job is now to try to wrap it up in a financial framework and so let's just get to it.

Through my presentation today, I'll try to address these key areas. Revenue and EPS are going to grow much faster than what we had told you before. We're going to see with diversification and growing into Data Center, the mix of businesses will change radically versus our previous estimate. Our operating scale, it's going to be a key differentiator for us going forward, and you heard a little bit about that in the various presentations. And then capital return, capital allocation, remains consistent with what we've told you before.

But before I go through all of this, let me just quickly address how Qualcomm has changed over the years. We obviously started off by inventing 3G. We led 4G, we led 5G, but today we have changed. We are more of a computing company than we are a connectivity company. We're still the best involved in connectivity, mind you, but we are a computing company today. So for a lot of investors who've known Qualcomm from a long time ago, it's important to make the switch that we are a computing leader that also happens to be best-in-class in connectivity.

The second change in Qualcomm, the second transformation, started when Cristiano became CEO five years ago. We went from being a smartphone company to all edge devices: auto, personal AI, networking, industrial, PC — all of these devices we are leader in now. The third transformation starts now. We're going to go from being a devices company to being a cloud and device company. And as we change, as we transform, it completely redefines what's in front of the company and how you should think about the company going forward.

Let's start with some financial lookbacks. This is the last five years. What has happened to Qualcomm overall? We've doubled revenue during this period, \$44B. During this period as well, we tripled EPS, and this performance and track record validates our growth strategy. It also sets the platform for what we're going to do going forward. Within this period, what did QCT do? QCT far exceeded the performance of overall Qualcomm, much more than 2X growth in revenue. We also had double-digit CAGRs in each of our revenue streams. I'll highlight Auto: 44% CAGR over this period. Within handsets, within Android handsets, we grew at a CAGR of 12%. This is a market that's perceived as mature, but during this period we grew because content increased, and the mix across tiers got stronger. We expect this trend to continue.

If you think about what happened to earnings during the same period, we grew four times, 2X faster than revenue and this is because of operating leverage and really investing and growing and diversifying across the businesses. This slide outlines the businesses we are

in today. We're in Licensing, Android Handsets, Automotive, IoT, and now in Data Center. Across these businesses, we address \$1.7T of TAM, and as we diversify more, as we launch new products, a very large portion of this becomes addressable to us. Over the next few slides, I will focus on three key areas in terms of financials. I'll talk about Data Center, I'll talk about Automotive, and I'll talk about IoT.

But let's first start with the numbers. I'm sure you're waiting for the revenue forecast. Last time, when we were here 18 months ago, we set a target of \$22B in non-handset revenue. This is Auto and IoT: \$8B in Auto, \$14B in IoT. And I remember a lot of investors said, "Wow, that's a very aggressive target. Are you going to be able to meet the target?" 18 months later, we are here again and we are very happy to say that we are revising the target. Our fiscal '29 revenue target is now \$40B. Just to repeat. Let me just repeat this. Same here. Last time, 18 months ago, we said \$22B. Now we're saying \$40B. This is nearly 2X increase in the target for revenues in fiscal '29. What this also means is the four-year CAGR from '25 to '29 is 40%. Very strong growth as a result of the diversification efforts. We have an incredible opportunity in front of us.

Okay. So now I'm going to start with Data Center and talk through the Data Center financial forecast. As Tony outlined, an incredibly exciting product portfolio. It's built on the basis of technology leadership, and this is by far our largest growth opportunity. In terms of revenue ramp timing, we have revenue today in fiscal '26 from the connectivity products we acquired through Alphawave. As we get to fiscal '27, that's only three months away, we will start ramping custom silicon revenue in first quarter of fiscal '27.

In second half of fiscal '27, we'll ramp AI accelerator revenue and then second half of fiscal '28, we will ramp CPU revenue. All these revenue streams will layer on top of each other, similar to what we did in Auto. We started with connectivity, cockpit, ADAS, same thing in Data Center and that's how our revenue will scale. Across all of these markets, there's a \$1T TAM opportunity for us. So, we are incredibly excited as we launch these new products, we'll be able to access a very large portion of that TAM.

So, this is our financial forecast for the Data Center business. Let's start with fiscal '27. We're targeting \$5B of revenue in fiscal '27 and let me highlight a few key things. We will have two hyperscaler customers that are at global scale that will drive at least \$1B of revenue within the year. This is not a concentrated revenue stream. We have diversification within the customer base. We also expect custom silicon gross margin to be slightly below our overall Qualcomm gross margin, but it'll be accretive at the operating margin level. So this is a very attractive financial business for us. In terms of AI accelerator and CPU, we'll be investing ahead of revenue ramp in fiscal '27.

Let's talk about fiscal '29 now. We're targeting \$15B of revenue in Data Center in fiscal '29. And this revenue, what gives us the confidence in being able to achieve this revenue, is the diversity of products, diversity of customers, and the fact that we are talking about multi-generations across our customer base. So this is the forecast that we're putting out for

fiscal '29. But let's talk about the opportunity beyond that. As we discussed earlier, it's a \$1T market cap — \$1T TAM opportunity for us, and we have an incredible product portfolio. Across all of this, we are targeting greater than 5% share in five to seven years. So this is not a question of, “What will we do in fiscal '27?” Which is very strong. It's not just a question of what we'll do fiscal '29, but the long-term opportunity for us is incredible.

So, I'll go to Automotive now. Snapdragon has become the platform of choice for the automotive industry. We are very proud of what the Auto team has built, and as Nakul discussed, that becomes the platform with which we jump into Industrial and then into Robotics. Content increase has been a key part of the story, and I often get the question on, “Hey, auto market is mature—when will you stop growing? You're growing much faster than everyone else in the industry.” The reality is that the part of the industry that we are in, the amount of content growth is tremendous. Between our third-generation product, fourth-generation product, and fifth-generation product, there's an 8X increase in content. This translates into financial growth and, really, a very long runway for us. We just launched our fifth-generation product.

As you think about the overall TAM for cars in that flat market, our SAM is going to grow—easily grow—double digits, and we are going to grow much faster than that. The drivers of that growth include digital cockpit capabilities, increased sensors, ADAS going to L2++, and then finally generative AI capabilities. And we are leading in each one of these areas.

Let's go to the automotive design win pipeline. As Nakul mentioned, \$65B. Two years ago, we were here, we're talking about \$45B design win pipeline. Now we are at \$65B. As a reminder, what design win pipeline conveys is really the cumulative revenue expected over the designs that we have won. But what's incredibly impressive about the design win pipeline is the diversification. Diversification of products: you have cockpit, connectivity, and ADAS. We have diversification customers and we have diversification within the regions. And this is a very important attribute of our pipeline. We're really winning globally across all OEMs. We're looking forward to continuing to grow this pipeline going forward.

Okay. So now I'll talk about the revenue forecast. This is the last forecast we gave 18 months ago. And we said we'll be at \$8B in '29 and we'll be greater than \$9B in '31. And at that point we had pulled in our revenue ramp by two years. So, what we're going to do now is we're going to pull it in again. By two more years. We will hit \$10B of revenue in fiscal '29 and, as Nakul mentioned, we'll be the largest automotive silicon supplier shortly. And we are not done yet. If you think about the growth vectors that remain, we're incredibly excited about them. Robotaxis, autonomy going to L4, token accelerators—a separate accelerator that we'll use HBM for—and then AI HBC 4, and AI workflows. So, a tremendous amount of vectors still remaining; we'll continue to grow this business for a very long period of time.

I'll go to IoT next. Crisitano outlined this. We're thinking about IoT in two buckets going forward. You have personal AI and compute, and industrial networking and robotics.

For personal AI and compute, agentic AI is driving an inflection point in these devices. We used to talk about glasses. We don't talk about glasses anymore. It's not just glasses. It's a bunch of different devices that come with it. For PCs, of course we are in Windows PCs, but now we're extending to Chromebooks in addition to tablets. For industrial networking and robotics, digital transformation is creating an incredible opportunity for us. So, I'll talk through both of these areas in some more detail, but here's the financial forecast. We're going to be over \$14B across these markets, which is a CAGR of 20% from where we were in 2025. We have strong growth drivers and plays in each one of these areas. Let's start with personal AI.

Personal AI is a set of devices that can see what you can see, hear what you can hear, and you can have an agentic AI conversation with it. We have the broadest technology portfolio that allows us to win here. If you think about things that are required—small form factor, low power, connectivity, great camera, sensor fusion, object tracking—we have all these technologies in-house, we integrate it into our products, and it has been adopted by all kinds of customers. You have the traditional OEMs, you have the hyperscalers, and you have new entrants, and all of them are using our chips. We have included a modest forecast in our financials for this market. If the vision that Cristiano outlined plays out and this market turns out to be much bigger, we have tremendous financial upside opportunity that is not captured into our forecast.

Let's talk about PC next. Actually, this week, two years ago, was when we got into the PC market. It's been a very short period of time, and here's our report card since then. We've launched products across all tiers, flagship to entry. We are the performance leader in every single tier. We have multiple generations of products now so we can launch current generation products and previous generation products at the same time. But when we got into the market, one of the questions was, will you be able to build a channel? Here are the metrics on the channel.

We now have not just design wins at OEMs, but very broad acceptance of applications that have been ported over to our platform, printers, peripherals, consumer channels with retailers and enterprise channel. We've built a very, very strong platform over the last two years, and we are ready to scale the volume as a result of it. Agentic AI is also driving an inflection point, and it changes the way the device is going to be used, and we are very well positioned to take advantage of that.

Rick mentioned GoogleBook. We're the lead partner for GoogleBook. We'll be launching devices with various OEMs over the next several months, and these are devices that will bring Snapdragon along with Gemini together to deliver the best agentic AI experiences. So, while there's more work to do for us in PC, we are tremendously proud of the progress that we've made, and we are set for takeoff.

Finally, Industrial, Networking, and Robotics. This is another market, as Nakul outlined, AI is accelerating the transformation. 18 months ago, when we were here, we were talking

about how micro controllers will move to microprocessors and AI. Today, I know a lot of our peers come in and talk about that same transition. This positions us extremely well for what's going to happen going forward.

Like PC, one of the criticisms that we heard is, "Are you going to have the channel that is required to drive the volume here?" So we've built that channel. Now we have hardware and technology partners, we have distributors, we have system integrators, and 38,000 unique customers, and we are ready to ramp the volume based on the products we have. We are targeting \$8B in '29 for this business as well.

Okay. So, let me summarize what we just said on the financials and add a few more metrics to it. Not \$22B anymore, \$40B in revenue outside handsets, with \$15B for data center. We are forecasting Android Handset revenues to grow modestly at 5% going forward. This assumes that the current memory environment doesn't materially change, and it also does not assume any uplift from the agentic AI conversation we just had. Both of those things would be upside to our handset forecast.

Licensing: We continue to be very happy with the way that business is stable, and it will continue to scale with 4G and 5G units globally. And then finally, operating margins, no change to the targets we've said before. We expect QCT, as we include Data Center into it, to be at 30% in the long term, and we expect QTL to continue to be at 70%. So, what does this do to the mix of businesses across QCT?

When we get to '27, Handsets will be less than half of our revenue. This is a conversation we have with investors all the time, "But you are a handset company." Here's how the mix is going to change. We'll be less than 50% in '27. In '29 with the forecast we just showed you, handset will be a third of our revenues. We'll be truly diversified across Handsets, Data Center, and Industrial IoT and Automotive.

Our capital allocation strategy really is unchanged. We're continuing to proceed down the path that we have discussed with you in the past. Our priority is investing in the business, maintaining the leadership on technology, and continuing to accelerate diversification. For M&A, we've done 35 acquisitions over the last five years, and each one of them was contributing to our growth strategy. We've picked a strategy and each acquisition comes in and says, "How do you help me accelerate something that I'm planning to do anyways?" And there's Alphawave, Modular, are great examples of what we've done. In terms of capital returns, over the last five years, we've returned \$40B to the shareholders. Over the last 10 years, we've retired 30% of our shares.

As we go forward, the strategy remains unchanged. We're going to keep increasing dividends in mid to low single digits, and we'll return most of our free cash flow to shareholders. Also, through this, we'll maintain a strong balance sheet. We'll retain financial and strategic flexibility, which is so important in our industry. Let's talk about OPEX for a second. What has happened over the last five years is, as we have diversified,

we've grown revenue at 15%, but we've grown OPEX only at 6%. This is in spite of funding all the diversification efforts. And what that did is OPEX as a percent of revenue came down from 31% to 23%.

As we go forward, we expect that to decline further to 19 to 20%. So very happy with the way we are managing the need to invest and focus on diversification, while realizing operating leverage financially. So we've talked about '27; we've talked about '29. So the question is, "Are we well positioned to continue to grow beyond that?"

So, before I wrap up, I want to talk about all these growth drivers that we have beyond '29. Data center: clearly a massive growth factor for us as we look forward. Robotics: as that becomes one of the largest markets in the long term, we are positioned to win there as well. Industrial: we'll see an upgrade cycle that'll happen over a very long period of time, positioned to win there as well. ADAS: both from ADAS and autonomy perspective, we have an opportunity in silicon and stack, and we are executing on all those opportunities as well. Personal AI, as Cristiano outlined, is a tremendous opportunity. If this becomes a scenario where everyone in the world has a couple of devices that are personal AI devices, in addition to phone and PC, we're going to have a very significant growth vector beyond '29. And then finally 6G as we outlined before.

So very excited. Our growth curve is not done in '29. We have a long-term growth opportunity, and we are focused on executing to it. Finally, I want to wrap this up with some key takeaways. We have a clear line of sight to diversified revenue base where handset will be a third of our revenues in '29. We expect non-handsets to be \$40B within QCT by '29. Our EPS target is up, greater than \$18 in that same timeframe. And then finally, when you think about long-term, we have an opportunity with the things we outlined to scale our revenue to \$100 billion.

Thank you very much for coming. Thank you for listening to us. We're very excited about what's in front of us. Thank you very much, and I'd like to invite Cristiano, Nakul, and Tony back on stage for Q&A.

Cristiano Amon, President and Chief Executive Officer, Qualcomm Incorporated

Thank you so much for staying with us. I'm glad that most of you stayed.

Tony Pialis, EVP & GM, Data Center, Qualcomm Technologies, Inc.

Before we start, Akash, given that financial performance, can I buy some more shares?

Akash Palkhiwala, EVP, Chief Financial Officer and Chief Operating Officer, Qualcomm Incorporated

Go for it. Ask Ann before you do.

Cristiano Amon, President and Chief Executive Officer, Qualcomm Incorporated

All right. Who's going to go first?

Akash Palkhiwala, EVP, Chief Financial Officer and Chief Operating Officer, Qualcomm Incorporated

Please go ahead.

Christopher Rolland, Analyst, Susquehanna

Hi, Chris Rolland, Susquehanna. Thank you so much for the day. I think data center is probably the most interesting here. The \$15B and then it sounded like more than \$50B over time. If you could talk about the linearity of this given your product releases, and then also customer deployments, I think that would be great. And your overall ability, you think, to address this with your partners as well, given supply constraints, et cetera. Thanks.

Akash Palkhiwala, EVP, Chief Financial Officer and Chief Operating Officer, Qualcomm Incorporated

So I think the best way to answer the linearity question is you have a number for fiscal '27 of \$5B, fiscal '29 of \$15B. And as we said, our product launches, the way it pans out is '29 will have the benefit of CPU accelerator and custom silicon and connectivity. So all four product launches, and then '28 CPU comes in at the end of the year. So you're going to see this ramp that happens between 5B to 10B, but it aligns with the product launches timeline across the period.

Cristiano Amon, President and Chief Executive Officer, Qualcomm Incorporated

Maybe I'll just add a few things to try to answer your question on supply chain. So as we align, we have visibility right now of \$5B in fiscal '27. For that revenue, we have secured capacity as well as secured memory. So even our customer commitments right now on the high bandwidth compute, HBC technology, we have secured memory as well. We're not a small company. We have a capacity allocation. We consume a lot of leading node wafers. I also think our suppliers are betting on Qualcomm and want Qualcomm to succeed, and I think that is reflected in the capacity commitments we have for the projected revenue we made of fiscal '27 of \$5B.

Tony Pialis, EVP & GM, Data Center, Qualcomm Technologies, Inc.

The one thing I'll just add is '29, as Akash said, is when all four product lines truly launch and that's just the beginning of the launch. So '30-'31 is when we start delivering multiple generations of these products and that trajectory is going to change from what you'll see over the next three years.

Cristiano Amon, President and Chief Executive Officer, Qualcomm Incorporated

Who's driving the microphone? I think there's a couple of hands.

Janco Venter, Arete Research

Hey guys, Janco Venter, Arete Research. I was wondering if you can zoom into the CPU commentary that you laid out for us. I mean, some of those statistics, performance

statistics versus your peers were quite compelling — the 5 GHz. And I was just wondering if you can maybe just hold my hand a bit and just tell us how you get to this, versus your competition, given how significant the TAM expansion has been over the last couple of months.

Tony Pialis, EVP & GM, Data Center, Qualcomm Technologies, Inc.

Look, that's a great question. As I mentioned during my talk, this is a company founded by engineers. Everything we do is about technical innovation. The Oryon CPU core is transformational, as you've mentioned. 5 GHz is remarkable and it's not even a custom hand-built design. It's built using automated tools and it's updated and refreshed each and every year. So look, it's foundational in architecture. You cannot just stitch this type of performance in. It has to be built from the ground up. So even though it's based on mobile compute, the server-class compute has been built from the ground up to lead in terms of performance. So I've been asked, "Why are you launching in '28?" Because even in '28, we will have the industry's best performance in compute and IO capability. And then when you think of bolting on the HBC attach to integrate native AI inference workloads on top of this industry leading compute, folks, this is game over across the industry.

Cristiano Amon, President and Chief Executive Officer, Qualcomm Incorporated

Maybe I'm just going to add a few things. I'll kind of remind you a little bit of our journey. So we have been building, I think, our own CPUs. The first thing we did, we built a CPU to have an Apple compete, as we entered the PC space to create an Apple compete. This is following the Apple M series. Then we built a CPU for mobile devices. We built a CPU safety-grade for automotive. So this is the next-generation CPU that we designed.

One commentary to what you said: right now the demand for CPU is massive. So everybody has CPU chips. I've seen this in the pandemic, but if you look at what happened in the other markets with design CPU, our metrics have been very, very good from a performance perspective, from a power perspective. The feedback that we got on our CPU that Tony outlined, I want to talk about it a little bit.

We receive very specific requirements from all the hyperscalers about the CPU that they want to see in '28 when we start shipping it. And the feedback is, "This is too good to be true. When can I get the silicon?" So you should be thinking about us having the capability, understanding where the puck is going for agentic, and building a CPU for that. Right now, everybody's shipping. It almost doesn't matter — the demand is high, but we expect by that timeframe, if you have true competition, Qualcomm's going to fare very, very well in this area.

Joshua Buchalter, TD Cowen

Hi, Josh Buchalter from TD Cowen. Thanks so much for hosting the informative presentation and taking my question. I was hoping you can maybe speak to your software maturity as we think about you merging into the data center ecosystem. I think we appreciate Qualcomm's rich heritage in silicon design and manufacturing, but it's a new

venture for you guys and one where others have been inhibited by their software platforms. So could you speak to that and sort of what Modular brings specifically as we think about merging that into your roadmap? Thank you.

Cristiano Amon, President and Chief Executive Officer, Qualcomm Incorporated

Very good. Thanks, Josh, for the question. So let me break this conversation in two pieces: what we're doing right now and what we're going to do soon, right? So first we've been very focused on inference and that is the focus right now, the disaggregated inference. It doesn't mean that's the only thing we're going to be doing, but we've been very focused on inference for the disaggregated accelerator into the data center.

We also bring another interesting capability, especially when you think about open-source models. Because when models have to run on the edge, they have to run on Qualcomm. We have been embracing a lot of the industry standards; we've been supporting ExecuTorch, we've been supporting Triton, for example, and we have been working over the years — there was a purpose. I told you that we have been preparing for this, building assets. There was a purpose for AI100. The purpose of AI100 is basically to start understanding how we need to mature our software stack to the point now where we have new models in 24 hours actually running on our accelerator. So that's what we have been building. We've been building the capabilities supporting all of the different open ecosystems.

But now there's something else we're going to be doing. Because the reality is some of those platforms right now are old. I think the incumbent platform has been designed about 20 years ago, and I think a company like Modular has a very modern platform which has been probably designed for the disaggregated heterogeneous compute and is open. That's how we're going to change the conversation a bit, not only creating something that delivers higher performance and is easier for developers to use not only for us, but for the rest of the industry and actually make sure that happens on the data center as well on the edge. So those are the two vectors. We'll continue to check the box and do what everybody is doing for inference, but we want to do something much better.

Louis Miscioscia, Daiwa Capital Markets

Yeah. Thank you, Lou Miscioscia here with Daiwa Capital. So, to continue on the Modular situation, it definitely seems very interesting. Maybe you could just talk about the obstacles. Where do you think it's going to be deployed? Hyperscalers, neoclouds, enterprise, because obviously once software does get defined, obviously Nvidia has made a lot of progress in this area. It's hard to overcome, as you've seen, with Windows and other things, but there has been success stories like VMware throughout the year. So, it seems like you could have an opportunity. So, more details would be great.

Cristiano Amon, President and Chief Executive Officer, Qualcomm Incorporated

Well, I'll tell you what we see right now, and I'll talk a little bit about the vision. But first I'll answer the question: if for inference Nvidia was the only option, Nvidia would be the only

thing shipping right now — which is actually not the case. I also believe that as you think about what the industry really wants and what a lot of the industry has been developing, you see the progress of Google with TPUs as an example. You saw the acquisition that Nvidia made of Groq. So you have different architectures and I think the moat of inference is actually not as strong as it has been for training, but it also created an opportunity because you now have clusters of compute with a bunch of different hardware and you want a solution that could actually abstract that problem for developers. Now I'm going to tell what I see.

What I see is the Modular team, Chris and team — who are here and we can't wait to get them as part of Qualcomm. They develop something that is modern and actually abstract this for developers and get a lot of performance out of the hardware. And just don't take my word for it. If you actually look at what they have done, they have achieved a lot of performance working with Nvidia hardware, with AMD hardware, with CPUs and it's truly an open platform that can actually run across different types of environments and compute and also scale for the edge. That's how we started working with them. So you're always going to have this conversation, which is, somebody's going to say, "I'm just going to stay within CUDA," — which is tied to Nvidia hardware — "I'm just going to stay there and I am just going to get one ecosystem, or I'm actually going to see the benefit of having heterogeneous compute, a disaggregated compute, and see what's going to happen on the edge, which will happen regardless." Like I'm telling you, I am getting some of the major AI companies meeting with us. You saw some of those videos today saying, "I need to move all of those workloads to the edge. I have a better use for some of my tokens in the cloud." That's going to happen and that's going to bring different types of hardware.

With that, I'm going to tell you what the vision is. At the end of the day, we have a lot of customers that make things and those customers that make things, they are going to start adopting a lot of AI in inference compute. You saw that across physical AI. So, I think the customer reaction is they're actually looking for a platform that scales; scales across the edge, is open, is easy to use. And we expect that there's a lot of different customers in the cloud. They're dealing with the fact now they have three or four or five different software stacks that they have to play with. That's the bet. And as I said, we're going to be actively driving it. We're going to continue what you heard from Chris. It's going to be open. It's going to be available to everyone and we're going to see what happens.

Chris Caso, Wolfe Research

Okay. Thank you. It's Chris Caso from Wolfe Research. For the fiscal '27 data center guidance, I think it'd be helpful if you could clarify exactly what's in that guidance and I guess from the product launches you've discussed, it sounds like it's the accelerator plus maybe some of the connectivity you got from Alphawave. And then with regard to the customers, you've already announced HUMAIN as an accelerator customer. Microsoft was today and you talked about two customers. So, are those the two that are included in that?

Akash Palkhiwala, EVP, Chief Financial Officer and Chief Operating Officer, Qualcomm Incorporated

So, from a product perspective, I'd say the largest part of the revenue base will be custom silicon. As I mentioned, there are two customers who are global hyperscalers who will each be greater than \$1B. So, by far, that'll be the largest part of the revenue. There will be a portion of AI accelerator coming in towards the end of the year and then connectivity products that came from Alphawave acquisition will be also a portion of it. So, it's really those three product lines with AI accelerator really coming at the end of the year. From a customer perspective, as I said for custom silicon, the two large customers will drive the custom silicon revenue and then we have a very large customer base in connectivity that came through the acquisition and will have HUMAIN as a significant portion of it as well. So that's the base.

James Schneider, Goldman Sachs

Jim Schneider from Goldman Sachs. Thanks for taking the question. Maybe just to follow up on the prior question, can you maybe talk a little bit about how you expect the customer diversity change from fiscal '27 onward. I think you talked about three customers, you sort of indicated maintaining the largest lion share of the custom silicon. Do you have orders for all of that \$5B already to be covered with that revenue and then maybe talk about how much more diverse you expect that revenue base to get in '29 and can you actually hit the '29 targets based on the customers you have now?

Akash Palkhiwala, EVP, Chief Financial Officer and Chief Operating Officer, Qualcomm Incorporated

So let me address it in two parts. I think first is we have high confidence in our forecast and so I'll leave it at that. The way you should think about the '29 forecast is we're engaged across a variety of customers today. We are engaged across a variety of products with them and then we are talking about multi-generation. So it's a combination of those factors that gives us confidence in the fiscal '29 number.

Tony Pialis, EVP & GM, Data Center, Qualcomm Technologies, Inc.

And the one thing I will add is remember in data center, the discussions are moving from megawatts to gigawatts. And when you deploy full infrastructure, as I outlined today, a few gigawatts can get you to the '29 numbers.

Holger Mueller, Constellation Research

Holger Mueller, Constellation Research. I would be amiss not asking a Brazilian: against who Brazil will win the World Cup?

Cristiano Amon, President and Chief Executive Officer, Qualcomm Incorporated

I don't know. It's a tough one. I'm kind of encouraged that Brazil started playing bad because when they start playing bad, usually it brings down, I think it brings more humility in the team. They start playing together and they start to improve in the second half. So I'm

going to take that as probably a consolation for the performance I have seen in the beginning of the season.

Holger Mueller, Constellation Research

Perfect. To the serious question. AI tells me you're building between 250 or 500 chipsets. I know you don't make that number complex, but I estimate with all your plans for 2029 that number might easily double. How do you plan to handle the complexity, because every chipset is an adventure — not all adventures end happy — from a human, skill, capacity, risking, supply chain perspective?

Cristiano Amon, President and Chief Executive Officer, Qualcomm Incorporated

Okay. Look, I'm going to give the answer that we always give to ourselves. So we're actually called "Quality Communications." Maybe I need to — It was interesting, during the 40-year anniversary of Qualcomm, we had an event and our founder, Dr. Irwin Jacobs, he came to speak and he said to me, "You know Cristiano, I made a mistake. I think I made a mistake. I should have done Qualcomm with one "m" because then it could be communications or compute interchangeably."

But I think the story here is we actually have a quality reputation. When we talk about our customers, we get awards all the time from mobile customers about the lowest defect density. We ramp brand new chip and IP in a very fast period of time. We get over and over, messages from Apple that we have been probably one of the best quality suppliers and you saw what happened in automotive.

So, I think we look at that skill that we have developed as something that is very important. What we saw, and this is a little bit complicated. This is a little bit complicated. I'm going to try to simplify the complexity. There's often discussion about leading on design and process technology. There's often this discussion and about who's best and yield and all of this. And one of the things we learned, for example, when we do Snapdragon for mobile phones that has to ramp very, very fast. You have to design the product to a very narrow spec. I could never afford to do, for example, what Intel does and say, "I have this distribution and I'm going to sell this as a high clock speed, as the i9, this is going to be the i7 and I'm just going to bin it. I have to develop the same exact part because you never hear Samsung saying, this is a fast Galaxy or not."

And we saw what happened with this incredible demand on the data center. We hear anecdotal data from a lot of the customers about parts that are going to have rework and failure rate. We actually look at that as a vector of differentiation for Qualcomm, which is our ability to reliability. And we're not small. We ship 40 billion components every year. As I said, we do a large number of tape-outs of leading node chipsets, do all of them in parallel. You saw record dates on very complex automotive industry, which is very strict from all the gates we have have gone to. We're breaking new records from tape-out to cars. So we're going to bring all of that to the data center and this is already happening.

I wanted to use this opportunity — sorry for the long answer —but I think sometimes the world is actually a lot more simpler than it looks. Alphawave had customers. They have been licensing IP and they have engaged with customers. What we saw as soon as we closed Alphawave, that conversation accelerated because just Qualcomm had more muscle. We just had more capacity. We have a bigger supply chain and we're building to actually make commitments that people will bet a large volume on. And I think that's kind of what happened. That's why they got celebrated; a lot of the custom ASIC engagements, they had Alphawave IP on it. And I think that's what I expect to happen. I think size really matters. As I said, all of the \$5B we outlined and forecasted right now, we have wafers and we have memory committed to those.

Announcer

We have time for one final question.

Joseph Cardoso, J.P. Morgan

Hey, thank you. It's Joe Cardoso from J.P. Morgan. Maybe more of a question for Tony on the connectivity side. Nice to see the roadmap here across copper and optical solutions. How we're curious, in one of your earlier slides, you also mentioned CPO. How are you thinking about that opportunity on the connectivity side? How is Qualcomm looking to participate and maybe timeline around the roadmap there?

Tony Pialis, EVP & GM, Data Center, Qualcomm Technologies, Inc.

Thanks for that question. So, at Alphawave, we had started working on silicon photonics and co-packaged optics about five years ago. And so what the plan is right now is to initially deploy first generation of silicon photonics in our AI300 series. So that will immediately empower optical scale out, drive down power consumption dramatically because you're no longer going from copper to optics, you're going straight to photons. And look, beyond that, AI fabrics will be moving optical. So we start with scale-out in and around 2028 and then scale-up beyond that.

Akash Palkhiwala, EVP, Chief Financial Officer and Chief Operating Officer, Qualcomm Incorporated

Okay.

Cristiano Amon, President and Chief Executive Officer, Qualcomm Incorporated

That's it? All right.

Tony Pialis, EVP & GM, Data Center, Qualcomm Technologies, Inc.

I think that's it.

Cristiano Amon, President and Chief Executive Officer, Qualcomm Incorporated

Thank you so much. Thank you. Thank you for being here with us. Really appreciate it. Thank you.