What an amazing day. Days like this don't happen often in a company's history, and for me personally, it's just an exciting moment to be participating in it.

I've had an opportunity to participate in a lot of new technologies being brought into various different markets over my career. This one ranks at the top for me personally.

As many of you may know, I spent quite a bit of time in the early days at Silicon Graphics, actually building out compute platforms, where we had to build our own silicon on these little tiny wafers that were like this big. Then moving on from there to NVIDIA and taping out and launching a whole bunch of 500-square-millimeter die platforms for all the compute GPUs when I was doing what I was doing for Jensen.

So I've been doing stuff in this industry for a very, very long time, and it's been fun to participate in all those other kinds of programs. But this one in particular is a big one for us.

Rene walked through the history of where things started when we went private. I joined Arm shortly after that, and joined Arm specifically to help get started the infrastructure business that's now this Cloud AI business.

As a matter of fact, I remember Mohammed and I one day standing in front of a whiteboard trying to figure out how we help customers build stuff faster, and that's when we first came up with the early ideas that we called the virtual SoC, which has now become CSS.

But that journey over that time, from when we were starting that to now, was not just about figuring out how to get new products into the marketplace or start the businesses. It was also about transforming the company.

For the last number of years, I've been involved in taking our physical design group — a physical design group that would do very, very low-level silicon optimizations for the top foundries across the world — and we would work with them to help build and make sure that they had the most optimized silicon in the wafers that they were putting through the foundries. Then we worked very closely with our partners to make sure that they could get up to the clock rates and speeds like Chris was talking about.

That group eventually transformed into the group that Steve Halter runs, which is our silicon group.

At the same time, I'm working on how it is that we take our entire software ecosystem and get our software ecosystem optimized, because Arm, if we're going to be participating in CSS and doing silicon and stuff like that, we also have to participate and help people optimize their silicon platforms. So that's what I've been helping with for quite some time.

Rene asked me just recently to start this physical design group for us, and I told him I was delighted to do so because I think that physical design is going to be the most interesting computing platform in the history of computing.

Let me get into that for a second.

[00:41:05] Physical design, or a physical AI, to kind of understand what I mean by that. Physical AI is where AI is embodied in a machine, where that machine then, AI embodied, where the machine then has to sense, decide, and act safely in the physical world to be able to take action.

I like to think of it this way: the key metric of this particular industry is the latency, the time between a photon hitting a sensor and an actuator actually firing. If you're traveling 65 miles an hour down the road and something is detected by a car that's operating autonomously, that moment from when that sensor detects something to when the braking system or steering system fires is the key attribute.

It's an incredibly complicated computing problem. It is, I think, going to become one of the most profound computing markets in our industry and in history.

To understand a bit about the TAM here compared to what Chris showed and compared to what Mo has talked about today, we think about the Physical AI TAM as being relatively small compared to some of these others — of course, about $25 billion a year or so, moving up to about $50 billion a year.

At $25 billion a year, it's really not hard to do the math on that. The Physical AI space today is principally automotive platforms beginning to move to more autonomous platforms. 100 million-ish transportation vehicles are sold each year, so you can kind of get a sense for that market.

Industrial robotics, which is kind of the robotics space that's really generating value today in the world, is only about 500,000 units a year. It's not a big opportunity today, and that will grow over the next five years or so. It's growing as a result of content growing — compute content growing — inside both of these devices.

The compute content is going from traditional automotive platforms that you see today into much more autonomous platforms, that have started to show up with robotaxis and the like. But the units aren't increasing dramatically, the compute content is. You're seeing a large increase in total compute content, which is driving this doubling of the TAM for this particular time.

But I had to beg and plead with the team: please let me show something a little bit past 2031. This is where I think the hockey stick happens for this particular market.

I think this market will grow in 2031 and don't ask me what year Drew will this thing hit a TAM of $200 billion, and while I can't tell you exactly what year this thing is going but I absolutely believe to my core that it is going to happen. It's going to happen. This hockey stick is going to go up high.

You've heard this from others, and this is really as a result of physical AI being embodied in robotics, in robotics and robotics platforms, and humanoid robotics platforms in particular. And its amazing how fast this will continue to grow. But I believe this is what's going to happen.

Frankly, I think the $200 billion undercalls it. I think it will be a trillion-dollar TAM sometime in that timeframe. But you've got to get ready for it.

Of course, this is platforms that span across a whole bunch of different types of form factors. In the transportation space, it's everything from traditional cars — which are becoming more autonomous now — and eventually, of course, autonomous robotaxis or the likes that you're seeing. If you're just cruising around here in San Francisco and you see them now all over the place.

But of course, it also moves into autonomous trucking, autonomous heavy machinery, and the like.

I'm working on programs right now with customers where we're talking literally about how it is you deploy an entire fleet of heavy equipment into a construction project, and that heavy equipment actually goes in and does all the work that needs to be done autonomously, which would be quite remarkable.

But then of course it's moving into robotics. These robotics platforms will be everything from humanoid form factors for doing all whole other bunch or type of projects, but then surgical robotics, and things used in industrial applications, warehousing, logistics and the like, security platforms, medical delivery via drones, food delivery, and all this kind of stuff.

All these things are happening. This is happening. The key thing across all of it is that compute is key to it.

This is what we've been getting ready for.

Similarly to what Chris was talking about for these businesses, for the physical AI business similarly we've been preparing for this by getting our platforms ready for helping people become more vertically oriented.

For the longest time, Arm of course just provided the individual components of IP, and generation to generation to generation we improved the capabilities of those platforms, even now through to v9. But it's not sufficient to just be able to offer all this stuff as bags of capabilities. You need to be able to help people build stuff faster.

That was the key thing we saw early on in the infrastructure business: the vertical OEMs that wanted to go off and build Arm platforms similar to what AWS was doing they wanted to get to market faster. They wanted to get their own silicon into the market. How do you do that? You do that by actually enabling them to be able to do that themselves.

So our taking, so instead of saying, "Here's all this stuff, try to figure out how to build it yourself," we come in and say, "Listen, we know it's hard to do this stuff, so we're going to curate it for you. We're going to design it in the compute subsystem. We're going to make it something that you're going to be able to build much more quickly into your platforms because you're going vertical."

I think we're in a decade or more of more and more companies going vertical in the platforms that they're building, for a whole bunch of different reasons that we can talk about sometime.

This move from individual bits of IP into curated designs — where we've designed the subsystem itself — has been a big, big, big move for us, and one that's worked out quite successfully.

What it means for our business is a move from v8-type architectures into v9-type architectures which Chris described the benefit of that eloquently. But that has resulted – we put a ton of investment into that and hence a lot more value into it, that itself has helped in a doubling of the rate that we collected on those types of platforms.

Then when you move into CSS, it's the doubling of doubling. We're building on top of v9, and we then again increase the amount of value that we bring to our customers. Of course, in exchange for that value that we bring to them, we get value back, and that's increased royalty rates.

To understand what it means to go faster time to market, or to reduce the amount of labor necessary to build stuff, this is hundreds of millions of dollars of value to our customers. Hundreds of millions of dollars of value.

We've done the analysis on it. You can't show up to your customers and say, "I'm increasing the royalty rates," without also showing up and saying, "Here is my financial analysis that shows you what the value is that we're bringing to you when you get a product to market faster, when you actually reduce the amount of time it takes for you to build it."

Things like being able to provide a compute subsystem so that they can actually tape out a chip and take that to market without making any changes to that chip — that's a huge economic value to them. That's the success that we've had with doing these platforms, like our CSS platforms.

These CSS platforms are now designed and optimized in the physical design space. They are designed and optimized now to be able to support this world of physical design space. This is the work that we've been doing.

As the world transitions over time, you have to understand where we're coming from, and this is where there are significant investments that we're making.

The world today — the vast majority of vehicles on the planet are traditional automotive platforms — and Arm has been in this space for a very, very long time. We created the Physical AI business unit just a few months ago, but we've been involved in this business area and these market areas for quite some time — decades.

So for instance, you might be surprised to know that in the last 12 months, just into the physical AI space, we've shipped over 2 billion Arm devices into our ecosystem to support this space. That's because there is an awful lot of compute that goes into these platforms.

Traditional automotive spaces goes everywhere from the actuation systems that manage braking, acceleration, and the monitoring of a bunch of different devices inside the vehicle, all the way through to the central compute that does things like keeping you in lane, helping you monitor speeds, and those sorts of functions, to the more sophisticated platforms that are becoming more autonomous.

That's all moving now to new levels of compute inside these autonomous vehicles. The big things that are changing are the architectures in these platforms. These architectures are evolving to become much more complicated, and again, with that more complication is more value that we bring into it.

An autonomous vehicle has substantially more Arm value that's brought into them than even traditional automotive platforms. But because we invest so heavily in those traditional automotive platforms and help them advance, all that technology into those platforms then moves into the autonomy platforms as well.

What's happened in these autonomous car platforms is you're seeing an increasing amount of compute being put into them. Increasing amount compute is built on a of Arm. Again, our CSS platforms are very enabling of this, and our latest platform is called the Zena CSS.

The big thing now is that as you move from this progression from autonomous – automotive platforms with ADAS capabilities into autonomous vehicle platforms — where the vehicle is moving itself — when the vehicle is moving itself, that's a big investment not just in the platform that moves the vehicle around, that manages the autonomy where decisions have to be made in milliseconds or faster, but you also have to improve the information that's displayed inside the vehicle.

Because if you're not driving, what are you doing? You want to know where you're going, you want to see mapping systems, you want displays inside these vehicles, and in addition entertainment that's going to be provided to you.

So there's going to be an awful lot of change that happens inside these vehicles that is perfectly set up for the way our business models are set.

[00:52:21] Of course, all of this moves into robotics. The same platforms that are going into autonomous vehicles are the same platforms that are going into robotics and empowering robotics.

The reason why the world and the industry is so fast-ready, so enabling and ready to move from autonomous vehicles into autonomous robotics — and why so many autonomous vehicle companies are also doing autonomous robotics investments and building those kinds of products — is because the compute is the same. It's a much more complicated actuation system.

As a matter of fact, that humanoid robotics platform is hands down the most complicated computing system ever. It is incredibly complicated. Controlling all the actuation systems, being able to assure that you're able to move a pinky appropriately at the same time that you're walking, at the same time as you might be getting instructions, at the same time that something is speaking to the robotics platform — it's an incredibly complicated platform.

So the computing has increased substantially, and therefore, as a result of that, our royalties could substantially increase. I think it's going to be a fast-growth product for us.

What's interesting, when you actually take a step back and look at it all, is that there are about four planes of computing that exist in these more autonomy platforms.

The perception-driven intelligence, which we are incredibly good at designing, has to run in robotics platforms and in automotive platforms. This is the platform that is responsible for how something moves around, how it makes autonomous decisions about how a car drives around the city, or how a humanoid robotics platform or some other type of robotics platform makes its way around a factory floor or a home.

That's a very, very, very complicated computing platform — one of the most difficult computing platforms you can build because of the real-time aspects in it. You've got to make decisions in milliseconds and nanoseconds.

Then there's a completely separate compute domain, which is the interaction-driven compute domain. That's where you're conversing. You're sitting in an autonomous car and you say, "Hey, wait, stop. I want to get a coffee right over there." The autonomous vehicle has to recognize exactly the context of that command, and then it has to communicate to the actual drive system that's managing the drive.

All of this stuff is incredibly complicated. Then it moves into how you manage the actuation system, and then into how the entire system is connected through the cloud. It's an incredibly complicated platform, and one in which we are really, really well situated.

Examples of that span the industry.

We've been working with the absolute pioneers in this space. Tesla, of course — I drive a Tesla. I drive 95% of my time now fully autonomously in my Tesla. We've been working with Tesla for a very long time.

We've seen announcements from Nuro, which is doing an autonomous platform to help companies that want autonomous vehicles go off and do that. They're building platforms using our technology.

Rivian just announced last December that they now have an AI platform that they've built for themselves — again, a company going vertical, and our ability to enable them to build that vertical platform themselves, being able to build the silicon platform they want using our technologies, v9 technologies, and our CSS technology, is what enables all this stuff.

Then robotics platforms like AGIBOT, who also build on top of the Arm platform today.

All of these companies are examples of real-life applications of people moving from v8 to v9, or moving from v9 into CSS platforms, and being able to use these systems that we provide.

[00:56:10] Interestingly, as Chris was saying, adoption in our space, the Physical AI space, adoption happens a little bit slower because it takes a little bit longer to put these platforms into the marketplace because of the way that safety requirements are imposed on these platforms.

You want to make sure these platforms are safe. So you have to do a certification to enable that. So we move along little bit slower than some other industries in adoption of these

types of technologies. But the adoption is absolutely happening. We're moving into v9 where v9 is part of the royalty collections and CSS, building on top of that doubling and doubling I referred to before, are going to again increase the growth of this particular industry for us.

This is all through the 2030s before the big, big growth happens as we move into the beginning of the next decade, which we are incredibly well positioned for.

That's why I tell people that we are just at the center of this incredibly large market opportunity that is staring right over the horizon for us. It'll be a $200 billion-per-year TAM. I think probably even more when it really hits its stride.

Like I said, I'm spending a lot of time in China looking at actuation systems and how those will drive down and become cheaper. This is how involved we are in the ecosystem.

This is going to be an incredibly fast-growing market for us when it hits that portion of the hockey stick, and we are really, really well positioned for it.

But we wouldn't be well positioned for it if we weren't making investments in software. That's my final point.

We do invest substantially in software ecosystems. As a matter of fact, we work very hard to migrate any platform on any other legacy architectures and make sure that they're optimized and moving into the Arm ecosystem.

As was described earlier, that's becoming incredibly easy for us to enable now for a whole bunch of different reasons. We have a very strong software ecosystem that we are able to rely on now that is incredibly mature, running all the platforms that we run. That means we are very very well situated for this growth into the physical AI space. Building on top of the investments we made over number of years in these platforms, from v8 to v9 to CSS and the like, the intimacy that we have with customers — customers going vertical, wanting to be able to build these kinds of platforms — we're bringing all the capabilities of our ecosystem together to enable this. But it does all require this very rich software ecosystem.

All right. I'm going to turn it over now to Mohammed, who's going to come back up and close on the opportunities now in cloud.