Jeffrey Kvaal - *Vice President, Head of Investor Relations, Arm Holdings plc*

Thank you. We know that we've given you a lot of information today, and so you have a lot to absorb. But this time, this next half hour or so, 45 minutes is for you to ask your questions. So, if you'd like to stick your hand up and then when the microphone comes down to you, say who you are and who you work for and then ask your question.

Andrew Gardiner - *Analyst, Citi*

Question: Thank you. And good afternoon, everybody. Andrew Gardiner from Citi. I have a question on some of the numbers you just presented, Jason, particularly the trajectory of the CPU revenue over the next five years. You gave sort of, I think, the sampling dates or rather Mohamed gave us the sampling dates of sort of this year for the first gen, next year for the second gen. But in terms of the revenue, you're talking about that $1 billion. I presume that's all first gen largely in fiscal 2028. Okay. And then how do we - in terms of the second and then third generation, what's your visibility into that in terms of the customer engagements? We've seen SoftBank NRE in particular flowing through the licensing line. So, what's the level of commitment from customers that's giving you the confidence in that ramp to $15 billion at fiscal 2031?

Jason E. Child - *Executive Vice President & Chief Financial Officer, Arm Holdings plc*

Answer: Mohamed, do you want to talk about the customer side and I'll share some of the numbers.

Mohamed Awad - *Executive Vice President, Cloud AI Business Unit, Arm Holdings plc*

Answer: Sure. Sure. So, bottom line is that we've got committed customers through the first two generations and we're currently in definition phase for the third generation with customers. And so, you heard Santosh earlier today talk about a multi-generation partnership. So, there's that with Meta. But there's also other customers who are committed to both generations at this point. And just to put a fine point on your earlier question, sampling now, production by the end of this year for gen one.

Jason E. Child - *Executive Vice President & Chief Financial Officer, Arm Holdings plc*

Answer: Yeah. And then just on the numbers, expect it to be, as we said, somewhere in the line of sight to $1 billion in 2028, we expect it to kind of roughly double again the next year and then double again. And so, that's slightly more than double maybe even the year after that, as we get more customers who are deploying, we are past some of the memory shortages or at least further down the road on some of the memory shortages. And then also because, as I mentioned a bunch of times, the later generations will also be at much higher ASPs as well. And that's consistent with what we're seeing across all of our partners and their plans for core counts and plans for just more performant larger, more expensive chips. The second thing I would add is, you mentioned SoftBank. The SoftBank, that doesn't have anything to do with this. The work we've been doing with SoftBank, as I think we've mentioned on past calls, that's really to address potential demand for Stargate on the accelerated compute side. And as I said in the presentation, we don't have any updates to that project yet, but that any NRE or whatever, all that's separate from the silicon that we announced today.

Ross Seymore - *Analyst, Deutsche Bank*

Question: Perfect. Ross Seymore from Deutsche Bank. Thanks, guys, for the great presentation. Perhaps for Mohamed, but maybe even for Rene or Jason as well. Mohamed, you talked about the competitive advantages of the AGI CPU versus x86. What about versus other Arm offerings? NVIDIA has its own obviously and they've talked about going merchant. And somewhat related, how are you so confident that there's no cannibalization? You ship the chip. NVIDIA doesn't ship the chip. Doesn't that have some impact accretive as it might be on the business model?

Mohamed Awad - *Executive Vice President, Cloud AI Business Unit, Arm Holdings plc*

Answer: Yeah. So, what we see and when you think about the product space and frankly customers that are going to go off and build their own silicon at that scale, if they're doing it, they're doing it because they're trying to generate some specific type of system-level advantage. Meaning, the investment required to go off and build these things are typically tied to a broader system design that they're going off and targeting. And that's true about all of our customers. If they were just looking for a product which was really what I would say applicable to a broader sort of cloud space or specific - a broader use case, then they may not go off and do that. And frankly, that's why you've heard us say over and over again that the market was underserved. That sort of highly optimized focus on efficiency, focus on performance, focus on scale, and not compromise by adding any sort of additional things which tie it back to a particular system or a particular accelerator, et cetera, that's really what we're focused on.

Rene Haas - *Chief Executive Officer & Director, Arm Holdings plc*

Answer: I would also add that we announced the Arm AGI CPU today. Last week, Jensen said Vera is going to be a multi-billion dollar business and he wasn't planning on selling CPUs. So, they're having discussion about channel conflict when a month ago neither products exist tells you just about how big this market is. This market is going to be very, very large. And as I mentioned earlier, I think the 4x increase of CPUs in agentic AI, we may be under-calling that number. So I think there is a very, very large market here where multiple players can play. And right now, I think the demand is higher than we think it is. And just, again, look at the fact that Jensen announced the product that six months ago he was essentially not thinking he had a CPU business.

Harlan Sur - *Analyst, J.P. Morgan*

Question: Yeah. Good afternoon. Thanks for hosting this event. Harlan Sur, J.P. Morgan. So, following up on that question, I totally agree with you. The Arm IP architecture within the data center is set up to serve so many different applications, right. You have merchant CPU, Qualcomm, NVIDIA, MediaTek bringing merchant solutions to the market. At the same time, you've got some of your cloud customers, Google, Amazon, Microsoft, Apple designing their own Arm-based architectures, right. And all of these guys are focused on accelerated compute and AI. So, I would argue with one, two, three, four, seven players in the market, not including Arm, that this basically covers the entire 120 million CPU per gigawatt TAM opportunity that you talked about. So, again, kind of help us where does Arm AGI sort of fit into the demand profile for data center GPUs sort of going forward? Is this more of a, like you said, maybe an enterprise play or some of the Tier 2 guys that don't have the capability to design their own silicon?

Rene Haas - *Chief Executive Officer & Director, Arm Holdings plc*

Answer: Yeah. I think I'd just go back to what we talked about earlier today. Two large demand drivers. One is just the need for compute and the need for power-efficient compute. So I think that starts to move the landscape, which has been accelerating already from x86 towards Arm. There are large parts of the market today that we don't serve around enterprise simply because of the legacy software there. Could that be an opportunity the next number of years? Absolutely. Then when you look at some of the applications that Mohamed showed, whether it is a head node working with a Cerebras or somewhere in the data plane where you might be working with a networking customer like an F5 or Cloudflare, I think this just gets a lot bigger relative to the total opportunity. And I think it starts with the need for compute and the need for power-efficient compute. And we would not have embarked on this business, Harlan, if we didn't think it had long-term legs. And we would not have showed a roadmap today that had multi-generations if we didn't have commitments for it. And lastly, as we talked about earlier, we had been pulled into this business. We had had customers once they saw what CSS could do and the need for power efficiency, they pulled us towards this milestone, and that's even before the agentic stuff took off.

Mohamed Awad - *Executive Vice President, Cloud AI Business Unit, Arm Holdings plc*

Answer: The only thing I would like to add to that is in the cloud specifically, you've got tens - well over 10,000 customers using Arm Neoverse specifically in the cloud that are looking to leverage that software investment on-prem. Arm Neoverse specifically. And so, this is really the only offering that sits in that category.

Kevin Cassidy - *Analyst, Rosenblatt Securities*

Question: Hi. Kevin Cassidy from Rosenblatt Securities. Maybe just to expand on that. Mohamed, you might have answered my question already. But if making your own silicon in the cloud is a good idea, why isn't it a good idea to move into the client and edge space?

Rene Haas - *Chief Executive Officer & Director, Arm Holdings plc*

Answer: Could be a good idea. We're not talking about anything, but what we talked about today, though.

Mehdi Hosseini - *Analyst, Susquehanna*

Question: Thank you. It's Mehdi Hosseini, Susquehanna International. Two questions. One for Jason. What is your underlying assumption for just blended chip prices? I'm under impression that you do benefit from higher prices, especially given the fact that your royalty rate is fixed, but economic value is going up and the chip price is a factor in it. So, what is your underlying assumption with the inflationary trend in semiconductor industry excluding memory? And then for Mohamed, just going back to the prior question, how should I compare and contrast your AGI CPU to LPU, and just maybe from a qualitatively big picture, you don't have to get into details.

Mohamed Awad - *Executive Vice President, Cloud AI Business Unit, Arm Holdings plc*

Answer: Yeah. I mean, LPU is a completely different piece that's really focused on - it's an aspect of accelerated compute specifically for the - it's specifically for inference and then decode specifically. So it's not really the same category. It's really about accelerated compute, where we're - as a CPU is really kind of think of it as the control plane, the management engine, the thing that coordinates all those accelerators. Accelerators can't exist without CPUs. LPU is an accelerator.

Rene Haas - *Chief Executive Officer & Director, Arm Holdings plc*

Answer: LPUs generate tokens, CPUs distribute the tokens. That's a very different workload.

Drew Henry - *Executive Vice President, Physical AI Business Unit, Arm Holdings plc*

Answer: And on the pricing, we, of course, don't disclose exactly what the pricing is. It's going to be different by customer based on volumes and based on who it is and what their situation is. You should assume that the market - if you look at the market of what does someone charge for an equivalent chip to AGI CPU or something - like the - maybe the x86 alternatives, they're somewhere around, I think, just shy of $2,000, somewhere in that range. So think that it's like probably not too far from that. And then, when you look at next generations, I think, most of the assumptions are that those prices go up by $50, maybe $75, maybe even double, just depending on what the core count is. And so our assumptions will probably be somewhere in line with what you see the rest of market doing.

Ananda Baruah - *Analyst, Loop Capital*

Question: Thanks. Ananda Baruah with Loop Capital. Thanks for doing this. This is great. So it's going back to use case. Sounds like you guys - and this is a clarification question. Sounds like you guys are talking about TAM expansion to the space from what you're providing with your chips today. And if that's accurate, is part of what's going on is - there's an appetite for a compute that's not present in the marketplace. You guys can really help fill in that compute need and some of your existing customers can pivot some of their resources to some of the other compute, like the right compute tool for the right compute job. Maybe you fill in more of the inferencing side. Some of the compute they're using for inferencing, they can pivot back to training, all while you're also amplifying the compute TAM. That's sort of the first clarification question. And then, I guess, the second one is, sounds like you feel solid on supply chain. Maybe memory is tight, but there's no talk of wafers yet. So just on the wafer side, is there anything - it sounds like you have that ironed out, but just wanted to get a clarification on that as well. Thanks.

Rene Haas - *Chief Executive Officer & Director, Arm Holdings plc*

Answer: I'm going to take that. I would say, to answer your comment, the way I would break down the compute needs and maybe back to the LPU question is that there's a category of compute that's generating the tokens, and then, there is the orchestration and distribution and scheduling of the tokens. We think that latter category is expanding in a large way - in a very, very large way. And that is the market that AGI CPU addresses. The generation of the token market is not directly addressed by this, but somewhat indirectly in the case that the better you are at orchestrating and distributing tokens, the more tokens that you can distribute, more tokens that get generated, it's a flywheel. The more demand you see for CPU resources, which is why we think we may be conservative on the numbers. We're not going to change them. But we think we may be conservative, because they work in tandem. But to be very clear, we're sitting on the side of the workload that is managing and distributing the tokens.

Eric Hayes - *Executive Vice President, Operations, Arm Holdings plc*

Answer: Yeah. My name's Eric Hayes, I'm the Executive Vice President of Operations. So on the supply side, we have the supply locked up very well for our ramps into production. Our supply chain partners, it's not just wafers, it's packaging, substrates, assembly test. And our supply chain right now isn't really just what we're buying from, they're partners, they've been partners of Arm for a long time, and we're expanding our business with them. They've gone through - and they have actually spoken with our customers and have confidence in our ability to deliver through to our customers, and they see us as a new channel for them, where this is going to be deployed, and it's a small investment. The numbers that you've seen up here represent small numbers to our supply chain. So it's very easy for them to supply into this and to help us develop this channel for them.

Rene Haas - *Chief Executive Officer & Director, Arm Holdings plc*

Answer: We've worked very closely with TSMC for decades. We have relationships up and down the board, including C.C. Wei's level. We talked to C.C. very early on on this development, making him aware of our strategy, our intent to get in this market, both in terms of a strategic partner as well as making sure we had security of supply. So most of the people on this panel have all worked in this business at one point in our careers. So it's something that we know it requires a lot of focus and attention.

Unidentified speaker

Question: Rene, I have a question for you. In the keynote, you mentioned that there is potentially up to $10 billion CapEx savings related to better performance of the CPU. Can you help us figure out whether it's more focused on the head node performance or the data prep funnel? Like, how do you guys conceptualize that between what's improving the GPU performance, right, and what's kind of straight CPU?

Rene Haas - *Chief Executive Officer & Director, Arm Holdings plc*

Answer: Yeah. Without giving you the entire formula, the way we thought about it was at the highest level, you're talking about roughly $50 billion of spend on a gigawatt data center. And with CPUs growing from 30 million cores to 120 million cores, and then, we calculate out what the power consumption looks like, we save about half versus x86 - that gets us roughly to that math of $20 billion - I think it was $10 billion.

Mark Lipacis - *Analyst, Evercore ISI*

Question: Hi. Mark Lipacis from Evercore ISI. Thank you very much for the presentations. The two questions, if I may. When you look out - you talk about a kind of a TAM or an opportunity going out five years. Is it - can you talk about the diversity of your clients, your customers as you go out? Is it a 60/40, where you have five big customers and everybody else, or does it look something different? And I know this is for Rene, but the explosion in demand for tokens and CPUs, I have to say, I haven't seen a market ramp like this in covering tech for 25 years. And it seems like a lot of companies have been caught offsides on this. And we talk about - there is a question about memory. And what do you think gates the ramp for you guys? Is it going to be limitations in memory? Is it going to be limitations in the industry's ability to ramp data centers? Like, what is the gating factor for you guys? Thank you.

Rene Haas - *Chief Executive Officer & Director, Arm Holdings plc*

Answer: Yeah. Thank you for your question. You brought up a lot of important points there. First off, every customer logo that we showed today is a customer that we are going to have backlog from. So these weren't evaluators, these weren't - maybe these are all customers. So to be able to stand up on the day of the launch with two very large companies in this space, Meta and OpenAI, and then, some great partners that we showed in the video, SAP, Cloudflare, that's a pretty significant milestone. So we already have a number of customers locked in, and we intend to expand that list. So I think it's going to be, A, very diverse customer base. I think, in general, whether it's turbines, whether it's access to energy, whether it's HBM, whether it's standard DRAM, we are going to see probably a supply crunch for this type of acceleration. That went into our logic when we went to the numbers here. So we didn't assume an unconstrained supply situation. We tried to be sensible about it. I do think personally, we are still in such early stages of how AI is used in large enterprises that we've got more surprises coming on the way just relative to the amount of pure compute that these AI models can subsume. Just look at the agents, right? These agents that run 24/7 and swarm a cloud. It's a little bit of a Borg mentality. The more compute you give it, the more agents that get spawned. And where does it kind of end? So, I think we are in very, very early days and I think it's typically when you see in these kind of situations, innovation starts to take over and people get very creative in terms of how to solve the bottlenecks. Good news for us. CPUs are a great tool to solve bottlenecks in all kinds of areas of a data center architecture, which gives us a high degree of confidence when we think about growing the base IP business in addition to a chip business, we think it's a really broad opportunity and we don't think the future looks like the past where you can look backwards and say, okay, you've got these guys doing chips and these guys doing IP, and I add up some of the parts and I can't get the math to work. The math never works that way when you kind of look forward because things change relative to how architectures are developed and solutions are driven. We have some core elements that give us some very high confidence. We're very power-efficient. We have a lot of software and the CPU is a really good tool to solve a lot of different jobs. And that all factored into why we went out as far as we did with the level of conviction that Jason talked about. We have a very good line of sight to where we think the future is going.

Victor Chiu - *Analyst, Raymond James*

Question: Thank you. This is Victor Chiu from Raymond James. Could you provide some color, I guess, around the genesis of this initiative? And how the discussions with your launch customers took form? I'm just kind of curious if the hyperscalers approached some of your current customers, first to kind of develop a solution which they weren't able to do, and then came to you to develop the chip, afterwards or if it was vice versa. And just wondering how they came to the conclusion that partnering with you would be the better alternative than doing it with, someone that already has the infrastructure in place and the scale in place to produce a chip that meets their efficiency needs?

Rene Haas - *Chief Executive Officer & Director, Arm Holdings plc*

Answer: Yeah. I can start and then Mohamed can fill in the dots. One of the things that makes Arm incredibly unique relative to where we sit in the ecosystem is because the Arm compute platform is where all the software starts. We have a lot of conversations about our roadmap with customers who aren't our direct customers. And what do I mean by that? Take Microsoft Windows for a moment. We spend probably more time than anyone in the Arm ecosystem talking to the Microsoft Windows Group about where Windows is going. So, we can understand exactly what are the things that we need to put in our architecture, so people can provide solutions to drive for that ecosystem. So with that overhang, we're having conversations all the time with people like a Meta or an SAP who are going to be users of our product somewhere in the value chain. And it's in the course of those discussions when customers are looking at the art of the possible, that the discussions take place from a creative standpoint of, gosh, instead of if I go down path X, could I go down path Y, which is a bit of the Meta deal. Mohamed, you can maybe fill in the detail on that, but that's a lot of how those things go.

Mohamed Awad - *Executive Vice President, Cloud AI Business Unit, Arm Holdings plc*

Answer: Yeah. I mean, I guess what I would say is that, obviously, the cloud business is on a - is a great upward trajectory. But it wasn't too long ago that we were not. There wasn't a whole lot of traction for us. And when I think about where we were focused, we were focused on the greatest consumers of compute. And so, you can count the greatest consumers of compute on one hand. Right. And if you look back all of the other hyperscalers were already on Arm. They'd already had programs, they're already going. And so, we were in there trying to figure out how to move the next one over, and we were positioning our IP and our CSS, et cetera, and then obviously they were like, oh! you're developing CSS. That's taking us a little bit further. And of course, things kind of continued from there. Their point to us was like, hey, we really like the offering. We wish we could get this thing in-house. We wish we could, but we don't. We are not in a position to go off and build a CPU. The market is underserved. Where do we get this, right? So, it was really about, focusing on the biggest consumers of compute and then helping them address the fact that the market was underserved.

Rene Haas - *Chief Executive Officer & Director, Arm Holdings plc*

Answer: Yeah. I mean, everybody can do everything. There are lots of customers that we have who are building Arm-based silicon today and they have incredibly capable design teams, yet they buy off-the-shelf Arm parts from people who are building chips as part of their business.

Mohamed Awad - *Executive Vice President, Cloud AI Business Unit, Arm Holdings plc*

Answer: And the only thing I would add to that is that if you think about how we think about our go-to-market strategy going forward, it is the next largest group of consumers of compute. Right. We're not going to every mom-and-pop shop because there are support requirements around that, et cetera. We're really focused on those big hits.

Lou Miscioscia - *Analyst, Daiwa Capital Markets*

Question: Okay. Thank you. Lou Miscioscia with Daiwa Capital Markets. So my question has to do more with China. Maybe I'll break it down into two parts. So, Rene, the comment that, the logos you had are real customers. Just wondering what the opportunity is in China. Is it possibly as big as the US given everything they're doing over there? Is there any restrictions or have you not been able to go in there? And the second part, I think, is to you, Henry, you had said that you were just in China. Just curious as to what you see they're doing with robotics and auto. And are they materially further advanced than us here in the US?

Rene Haas - *Chief Executive Officer & Director, Arm Holdings plc*

Answer: Yeah. What I would say with China, as far as the Arm AGI CPU, we have no customers to announce today. But there's no reason why we should not see market adoption in China. At the highest level, there's very good product fit relative to the very reasons that a non-China customer would take the Arm AGI CPU are the same reasons why a customer in China might. We just don't have anybody to talk about today. The big question always relative to export control and how would that kick in? From everything that we understand about the export control rules, there's nothing about this product that would be in tripwire with that. Obviously, should things change, we'd observe that. But right now, we don't see that. So no customers today, but no reason we can't have any.

Drew Henry - *Executive Vice President, Physical AI Business Unit, Arm Holdings plc*

Answer: And then to your question about what's going on in robotics, I was in China last week. And one of those trips where it's a series of meetings all day. The next morning, get on a plane, fly somewhere, series of meetings all day. Next morning, get on a plane, fly somewhere. It's one of those kind of trips. I saw just about as many companies you could possibly see in the amount of time that I was there. And it was incredibly impressive. The thing that I've taken away is that, of course, China continues to move at China's speed, historically just very famous at the speed at which they move things. And there's a lot that's happening there, rumored to be in the area of 200 robotics companies in China alone. But the more interesting thing that I've noticed is the global development in robotics. For instance, Europe, as you know, for quite some time has missed a lot of the waves that have been these technology waves. Right around the University of Zurich, right around the ETH University of Zurich. There are a bunch of robotics startups. One of the top robotics companies is in Germany. Across the US, we've got a number of robotics companies in the US. So, there's quality and there's quantity, and those things are very different. So, the thing that I'm recognizing in the areas where there's quantity, where there's lots of manufacturing capacity that might be made available at very inexpensive prices, people are filling that with quantity. But the quality is where the real value of the market is going to be, and that to me is happening globally. And so, I'll spend as much time in Europe and in the US as I will in China as we work more in this particular space. The last thing I'll say is that quality is about where computing investments are being made. That's the key thing about robotics: where the investments in computing are, and those investments are being made globally. In the kind of lower level infrastructure side of things, things like actuators, motors, and controllers, and things like that where manufacturing is cheap. There's an awful lot that's happening in the low-cost manufacturing world, which had been a lot in China but is now diversifying, as people have figured out, listen, we've got to diversify how manufacturing is happening. So it's as much happening in other places today. The last point I'll make about it is the interest that's going on in robotics right now is there's an awful lot of cost in the kind of low-level actuator infrastructure that exists. Because people are trying to figure out what is the kind of materials you want to buy, you want to build stuff with, and so, there's lots of machines that are getting built. The machines are expensive to build, and then you're not putting a lot of volume through them. So each component is expensive today. But as that market begins to consolidate around more, and more, and more common parts, we're going to see a collapsing of that cost. That's going to happen. And the thing that's interesting to me is that in much of the robotics space today, a significant portion of the cost is actually in the mechanicals and actuation systems. And that is going to go through a dramatic cost decrease as economies of scale hit. The thing that's common across all of that, though, is the compute platforms and the CPU platforms are going to be as high-performance and as high capability as you can possibly provide for the next decade or more in those compute areas that I talked about and I find that interest is global.

Charles Shi - *Analyst, Needham & Company*

Question: Hi. Charles Shi from Needham & Company. Maybe I want to abstract all the discussion around AGI CPU, all things you guys are doing on the silicon strategy. To put that a little bit into the perspective of where things may be going a little bit even longer term. So the industry has gone down that path from the merchant model, companies buying off-the-shelf, and then more like ASIC model, and then going to like more recently customer-owned tooling, basically in-sourcing chip design, I mean, the hyperscaler customers of yours. But now it looks like this whole launch of silicon, you guys are, well, by the way, you guys have been a big enabler of that path I just described. But now looks like you're going back to offering off-the-shelf chips. So going a little bit opposite of the trend that has been happening. So my questions are the maybe the twofold. One, why this is happening now because the trend has been one way, but it seems to be going back. And two, your customers, especially the hyperscaler customers who like to buy the off-the-shelf chips from you guys, how

do you make sure they stay with you on the silicon business instead of maybe even maybe they're going down the same path as many of the other guys have gone through. Going from the merchant silicon to ASIC, going to maybe COT, customer-owned tooling. That's a long question, but hopefully we can get some color. Yeah.

Rene Haas - *Chief Executive Officer & Director, Arm Holdings plc*

Answer: Yeah. Yeah, it was a long question, Charles, but I think I got the gist of what you're saying. So let me let me try to play it back a little bit. I think what we see for the Arm architecture is an opportunity to have a significant share growth that isn't in conflict with what you are saying. So what do I mean by that? Well, hyperscalers still continue to do their own chips in-house to get the kind of efficiencies you're talking about. Yes. Do we think the TAM for Arm expands because more and more compute is needed in the data center? We do. Because we don't think the hyperscalers will be able to do every single chip for every single application. Maybe they can, maybe they can't. But also more broadly, we think the constraint on power in general really brings our power efficiency advantage to the table where we think we can now materially take market share from x86. So, I think you've got a expansion of the TAM of other applications. I don't think the hyperscalers, I said, need to stop for us to be successful. I think it's a large market for both. And I think, as I said, I think the demand for Arm will grow because A, power efficiency really, really matters. And all of the work we've done around software is really paying dividends. And just look at the comments that Paul made from Meta this morning of that porting exercise. It went from a can we do this to, yes, let's do this because the power efficiency - I think his words were, it's too large to ignore and I think you are going to start to see that in areas that have been x86.

David Gibson - *Analyst, MST Financial*

Question: Thanks. David Gibson from MST Financial. Two questions on the commitments from your customers. Are there minimums? Are volume and price set? Can you give us a sense of how flexible those commitments are and hence the risk to your revenues going forward? That's the first one. The second one is going to be a tougher one perhaps. I know we haven't even produced the chip yet, but wondering how you think about the life of the product itself? Is it a five-year cycle replacement for Meta in five years to replace it with generation six or seven whatever. Wondering how you conceptualize the life of these chips and how long they'll last for? Thank you.

Rene Haas - *Chief Executive Officer & Director, Arm Holdings plc*

Answer: Yeah, I'll answer the first part, and maybe Mohamed you can comment on the second part, I'll give my view. To your question of how solid are the commitments from the customers and how would we bake that into our forecast, Jason would not allow us to show any numbers that he didn't feel confident. And so the commitments on forecasting our confidence levels are pretty high relative to supporting the numbers that we showed. So it's not like, oh my gosh, we think we've got a lot of work to go off and do. Is there upside opportunity? There always is. But we took a very conservative view as far as the numbers go. I would imagine and Mohamed, you can chime in here, too. I think this probably looks like a five-year replacement cycle, even though we're going to be bringing our products faster than every five years.

Mohamed Awad - *Executive Vice President, Cloud AI Business Unit, Arm Holdings plc*

Answer: Yeah. I mean, I think five years is kind of what we typically see, though. Those numbers are changing now, right? Because I mean, the challenge is, is that, listen those numbers are driven by amortization, and CapEx costs and how do you depreciate it? And when you're talking about numbers this big and frankly, this goes to whether or not you even build your own silicon. The calculus for what it means to go replace silicon or build silicon changes pretty dramatically. And so five years is a good, where we are today but I think that's dynamic. We'll see how it goes. I mean, we're certainly introducing new products that are going to be more efficient, more performant. We expect them to pick those up as they happen.

Rene Haas - *Chief Executive Officer & Director, Arm Holdings plc*

Do we want to wrap up? Okay. If there are no further questions. Wonderful set of questions. Obviously, you guys do your homework. But as Jason so eloquently put, we have extreme confidence in two businesses that compound on top of each other that produce a very game-changing result for us in five years. We didn't take the announcement lightly relative to what it meant to the ecosystem, what it meant to the partnership, which is why you saw some of the videos that you saw. And at the same time, we don't take lightly providing you numbers that look quite different than how you've thought about the company before. But we have been very meticulous in terms of looking at where we think the trajectory of the business is going. I think the last Analyst Day we did was when we went public in September 2023 - something like August 2023, and we showed numbers that we had high confidence in that showed a number of royalty agreements under contract. We beat all of that as we stand here today 2.5 years later. So that should give you the confidence that as a management team, when we sign up to a projection and share with an audience like this, we're serious about it. So thank you very much.