

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 half, next half hour or so, 45 minutes, is for you to ask your questions. So if you'd like to take a hand up, and then when the microphone comes around to you, say who you are and who you work for, and then ask your question.

Thank you. Good afternoon, everybody. Andrew Gardner from citi. I had 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, Muhammad gave us the sampling dates of this year for the first gen, next year for the second gen. But in terms of the revenue, you're talking about that billion dollars, I presume that's all first gen, largely in fiscal 28.

Okay. And then how do we in terms of the second and then third generation, what are you 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, you know, giving you the confidence in that ramp to 15 billion out to fiscal 31?

Muhammad, you want to talk about the customer side. And I'll share some of the numbers. Sure. So bottom line and is that we've got committed customers through the first two generations, and we're currently in definition for for the third, generation with, customers. And so, you know, you heard, you know, you heard Santosh earlier today talk about a multi-generation, partnership. So there's that, with Meta, but there's also, you know, 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 so for a Gen one.

Yeah. And it's on the, on the numbers. Yeah. I expect it to be, you know, as, 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, you know, 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 just more performant, larger, more expensive, chips. The second thing I would add is you mentioned, SoftBank, the SoftBank, that 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, updates to that

project yet, but that any NRE or whatever all that is separate, from the silicon that we announced today.

Ross Seymore from Deutsche Bank. Thanks, guys, for the great presentation. Perhaps for Mohammed, but maybe even for Rene or Jason as well. Mohammed, 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 know, 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?

Yeah. So so what we 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 you know, 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. You know, if they were just looking for a product, which was, really what I would say, applicable to a broader sort of cloud space or, you know, specific, a broader use case, and then they may not go off and do that and that, frankly, that's why we you've heard us say over and over again that the market was underserved, that sort of, you know, highly optimized, you know, focus on efficiency, focus on performance, focus on scale, and not compromise with by adding any sort of additional things which tie it back to a particular system or a particular accelerator, etc. that's really what we're what we're focused on.

I would also add that, we announced the ARM AGI CPU today, last week Jensen said Vera is going to be a multibillion dollar business, and he wasn't planning on selling the CPUs. So to have a 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 CPU is around and I 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. I 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.

Yeah. Good afternoon. Thanks for hosting this event, Harlan Sur JP Morgan. So following up on that question, I totally agree with you. You know, the Arm IP architecture within the data center has 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 GW TAM opportunity that you talked about. So again, kind of help us, what does ARM AGI sort of fit into the demand profile for data center GPUs sort of going forward? Is this more of, like you said, maybe an enterprise play or some of the tier two guys that don't have the capability to design their own silicon.

Yeah, I think I 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 that we don't serve around enterprise simply because of legacy software there. Could that be an opportunity in the next number of years? Absolutely. And then when you look at some of the applications that Mohammed 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 I we 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 have been pulled into this business. We have 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 Agentlc stuff took off.

Now, 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 that category.

Hi, Kevin Cassidy from Rosenblatt Securities. Maybe just to expand on that and maybe 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?

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

Thank you. This is 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, you're given the fact that your royalty rate is fixed, but economic value is going up and the chip prices have a as 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 to maybe from a qualitatively big picture, you don't have to go into details. Yeah. I mean, LPU is a completely different, completely different beast 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 it's not really the same category. It's really about accelerated compute where, where as a CPU is really kind of think of it as the, the control plane, the management engine, the thing that coordinates all those accelerators. Accelerators can't exist without CPUs, LPU is an accelerator LPUs generate tokens. CPUs distribute the tokens. That's a very different workload.

And on the pricing, we you know, we of course, don't disclose exactly what the pricing is. It's going to be different by customer based on volumes. And you know, based on on who it is and what what 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 a CPU or something, you know, like the maybe the x86 alternatives there 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, you know, our assumptions will probably be somewhere in line with what you see the rest of the market doing.

Thanks. Ananda Barush with Loop Capital. Thanks for doing this. This is this is great. So it's going back to use case. Sounds like you guys is in this a clarification question. It 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 right compute tool for right compute job. You know, 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 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 sounds you give that ironed out? But I just wanted to get a clarification on that as well.

Thanks. So I'm okay then. I would say to make sure of anything your comment, the way that would break down the compute needs and maybe back to the LPU question, is that there's a

category of computers at 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.

You want into this budget? Yeah. My name is Eric Hayes. I'm executive vice president of operations. So on the supply side, we have the supply locked up very well for our ramps into production. You know, 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. It's their 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 it, 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 a very easy for them to supply into this and to help us develop this channel for them.

We've worked very closely with TSMC for decades. We have relationships up and down the board, including at 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, we 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 requires a lot of focus and attention.

Thank you for hosting the day, Rene, I have a question for you. In the keynote, you mentioned that there is potentially up to 10 billion CapEx savings related to, you know, 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, you know, what's improving the GPU performance rate or, and what's kind of straight CPU?

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, 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, that math of 20 billion, I think was 10 billion.

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 is it is 60/40 where you have, you know, five big customers and everybody else, or does it look something different? And, and I don't know if this is for Rene, but, you know, the explosion and demand for tokens and CPUs, I, you know, I have to say, I haven't seen a market ramp like this in covering tech for 25 years. And it seems like, you know, there a lot of companies have been caught, offside on this. And there, you know, we talk about there's a question about memory and, what do you think is what 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? What is the gating performance? For you guys? Thank you.

Yeah. Thank you for your question. You brought up a lot of a lot of, important points there. You know, first off, every customer logo that we showed today, is a customer that we are going to have a backlog from. So these weren't evaluators, these weren't maybes. 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 on the video SAP, Cloudflare. That's a pretty big 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. You know, when we went to the numbers here. So we 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 seven and swarm a cloud, it's a little bit of a Borg mentality. The more, 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 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 the 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 to give us some very high confidence.

We're very power efficient, we have a lot of software, and a 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 that Jason talked about. We have a very good line of sight to where we think the future is going.

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

Yeah, I can start. And then Mohammed can fill in the dots. One of the 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 and customer 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 on 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 and 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 a minimum, and you can maybe fill in the detail on the Meta deal, but that's a lot of how those things go.

Yeah. I mean, I guess what I would say is that, you know, obviously the cloud business is, is on a has 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 had already had programs that were already going. And so, you know, we were in there trying to figure out how to move the next one over. You know, and we were positioning our, our IP and our CSS, etc. and then obviously, you know, like, oh, you're developing CSS, that's taking us a little bit further. And of course things kind of, you know, continued, continued from there, you know, 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. You know, 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, you know, focusing on the biggest consumers of compute and then helping them address the fact that the market was underserved.

You know, I mean, everybody can do everything there are 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.

And, you know, the only thing and I would add to that is that if you think about how we think about our go to market going forward, it is the next largest group of consumers of compute, right? We're not going to every mom and shop pop shop because, you know, there's support requirements around that and etc. we're really focused on those big hits.

Okay. Thank you. Lou Miscioscia with Daiwa capital markets. So, one question has to do more with China. Maybe I'll break it down into two parts. So Rene, the comment that, you know, 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? In the second part, I think, is to Drew Henry, you said that you were just in China, just, curious as to what you see there doing with robotics and auto, and are they materially further advanced than us here in the U.S.?

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-Chinese customer would take the Ami. High 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 would observe that. But right now we don't see that. So no, no customers today, but no reason we can't have any.

And then to your question about what's going on in robotics, I was in China last week was one of those was one of those trips where it's it's 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 as you could possibly see in the amount of time that, 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 speed, historically just very famous at the speed at which they move things. And, and there's a lot that's happening. They're 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, in, in Europe, and Europe, as you know, for quite some time as missed a lot of the waves that have been these, these, these, these technology where it's competing waves, it right around the University of Zurich, right around the ETH University of Zurich, is 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, in the US. So there's, there's there's quality and there's quantity and those things are very different. So the thing that that I'm recognizing in the areas where there's there's quantity, where there's lots of manufacturing capacity that might be made available at very inexpensive prices, folks. People are filling that with, with quantity. But the quality is where the where the real value of the market is going to be. And that to me is happening actually globally. And so I'll spend as much time in Europe and in the U.S. as I will in China as we work more in this particular, in this particular space. The last thing I'll say, and that's and what that quality is, is about where computing investments are being made. And that's the key thing about robotics is about where the investments in computing are. And that's being made globally in the kind of low level infrastructure side of things, things actuators and motors and controllers and things like that, where manufacturing is, is cheap. There's an awful lot that's happening in, in the, in the low cost manufacturing worlds, which had been a lot in China, has now being a diversified has people figured out, listen, we've got to diversify how manufacturing is happening. So it's a it's as much happening in in in other places today.

The last point I'll, I'll make about it is the interest that's going on in robotics to me 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 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 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, as, as scale, a kind of to scale hit. The thing that's common

across all of that, though, is the compute platforms. And the compute platforms are going to be as high performance and has and as high capability as you can possibly provide for the next decade or more in those compute platforms that, that, that I talked about it and I find that interest is is global.

Hi, Charles Shi from Needham in the company. Maybe I want to, abstract all the discussion around, AGI, CPU, all things you guys are doing on the silicon strategy. And, 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 through down that, that path of 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 a sourcing chip design. I mean, the Hyperscaler customers of yours. But now it looks like at 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 you looks like you're going back to offering to off the shelf chips.

So going a little bit the opposite of the trend that has been happening. So my question is are the maybe the two fold, one why this is happening now? Because the trend has been one way, but yeah, it's seems to be going back and two, your customers, especially the Hyperscaler customers, who elect to buy the off the shelf chips are from you guys. How do you make sure they stay with you on the silicon business? Instead of maybe they will even maybe they're going down the same path as many of the other guys have have gone through, going from the merchant to silicon, going to ASIC, going to maybe even COT Customer owned tooling.. That's that's a long question. But hopefully we can get some color.

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? What do I mean by that? Will 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 maturely take market share from, 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, that, 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 is 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.

Thanks. David Gibson from MST financial. Two questions on the commitments from your customers. Are they minimums? They volume in 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 it's we haven't even produced the chip yet, but wondering how you think about the life of the product itself. Like, is it a five year cycle replacement for Meta in five years to replace it with generation 6, 7 whatever? Wondering how you conceptualize the life of these chips and how long they last for.

Yeah, I'll answer the first part. Mohamed, you can comment on the on the second part, I'll give you my view, to your question of how solid are the commitments from the customers and how have we baked that into our forecast? Jason would not allow us to show any numbers that he didn't feel confident in. So the commitments on forecast and our confidence levels are pretty high relative to support 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.

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, you know, amortization and CapEx costs and, you know, 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, you know, where we are today. But I think that's dynamic. We'll see. We'll see how it goes. I mean we're certainly introducing new products. They're going to be more efficient, more performant. We expect them to pick those up as they they happen.

I don't have any more questions. Do we want to wrap up? Okay. If there are no further questions, I wonderful set of questions. Obviously, you guys know your 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 the business is going. I think the last analyst day we did was when we went public in September 23.

So that'll be like August of 23. 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 two and a half years later. So that should give you the confidence that as a management team, when we, when we sign up to a projection and share it with an audience like this, we're serious about it.