Hello everyone. My name is Jason Child. I'm CFO. I think I've probably met all of you, and probably multiple times, so it's great to see you again.

Very exciting day. Exciting to walk you through some of the numbers. I imagine many of you have been waiting to see what it all adds up to.

What I'm going to do today is make sure that the business case and the financial model behind the Arm AGI CPU is very, very clear.

There are really three points to this.

First, demand from new and existing customers is allowing us to materially expand our opportunity through selling chips.

Second, our existing IP business continues to have strong underlying growth drivers, and the chip business is compounding, and not displacing, the IP business.

Third, when you put those together, the combined model has much larger revenue, profit, and EPS potential by FY31.

All right. So, introducing the first phase of Arm's market expansion.

Let me start with the framework. Customer demand for Arm-created chips and the size of the opportunity has led us to explore the chip market over the past three years. After exploring the market and our own capabilities, we're introducing, as Rene and the team have talked about, the Arm AGI CPU.

Furthermore, the Arm IP business is going from strength to strength. Royalty revenue growth and license revenue growth both have multiple long-term structural growth drivers, which we expect to continue for years to come.

Finally, the financial consequences by FY31, that the combined model is materially accretive to revenue, gross profit, operating profit, and EPS. Importantly, much of the investment is already in the business, so the additional chip gross profit has meaningful drop-through to earnings.

Arm's opportunity is huge, and we're growing into the largest market in history — one that is just getting bigger and bigger. It's over $500 billion today, and we think this grows to more than $1.5 trillion in FYE31.

Just to be clear, this is just semiconductor logic — CPUs and XPUs. No memory, no optical, just chips where you might find Arm technology either today or in the future.

Breaking that down a little bit more, Cloud AI includes cloud compute, both CPU and XPU, enterprise compute and networking, supercomputers, and so on. This is a $330 billion market today, growing at over 30% a year to around $1.2 trillion in FYE31.

Edge AI includes the chips that go into smartphones, consumer electronics, and IOT devices. This is a $180 billion market, which we expect to grow at around a 7% CAGR over the next five years to $250 billion.

Physical AI, which includes automotive applications and robotics, this is a $25 billion chips market today, doubling in size over the same period.

You've heard Drew talk about how the autonomous vehicles and robotics opportunity could be larger than both Cloud AI and Edge AI. However, our current view is that the inflection point likely happens after fiscal 2031. It could be earlier.

This adds up to over $1.5 trillion.

[01:32:43] Double-clicking on the Cloud AI market, we have greater than $100 billion of Cloud, AI and Enterprise data center silicon, and $55 billion for wired and wireless networking. The remaining $1 trillion includes data center accelerator chips. We will come back to that opportunity another day.

Today we are just focused on the CPUs, and there is $455 billion of non-accelerator CPU TAM that is an Arm sweet spot.

I want to direct particular attention to the $100 billion-plus in the data center CPU TAM, as that is the market we are addressing today with the Arm AGI CPU.

These numbers may be higher than you've seen before. However, we have visibility into our customers' roadmaps. We are confident that the demands that inference and agentic AI will put on CPU performance will drive healthy volume growth and ASP increases.

We're expanding our opportunity in three ways.

Firstly, with our IP and CSS offering, the only customers we can address are the large hyperscalers who want to build their own chips. Our maximum revenue is limited to a fraction of the chip value. Assuming a 100% market share at a 10% CSS royalty rate, this adds up to $2.4 billion.

Supplying a complete chip allows us to address the full value of this market at $24 billion.

Secondly, as only the largest cloud service providers we are building their own chips, we are expanding the opportunity by offering products that all data center companies can use — from the largest hyperscalers to neoclouds to telcos and enterprises.

Thirdly, over the next five years, we expect the size of the market to increase significantly, driven by the expansion of CPUs used for inference and agentic AI.

Together, this takes our total opportunity from $2.4 billion of possible royalty revenue this year to more than $100 billion of revenue in FYE31.

Not all of this will be captured by chips. We still will be offering our customers CPU IP and compute subsystems, and this will help us maximize our total revenue from the data center by allowing our customers to choose the right solution for them.

[01:35:16] This slide makes the economics, and thus our rationale for entering the market, more tangible with a simplified hypothetical example using an illustrative $1,000 chip price.

CPU IP has about a 5% royalty rate, so for every $1,000 of chip sales, Arm receives about $50 in royalty revenue. This is at 100% gross profit, so you get $50 of gross profit dollars.

Compute subsystems have about twice the royalty rate, so it generates about $100 in royalty revenue and gross profit. For or a chip.

A $1,000 chip revenue model delivers about $500 of gross profit dollars. So while IP and CSS are extremely attractive on gross profit margin, the chip model produces more gross profit dollars per chip. This could be an order of magnitude higher than the IP model.

That is why we are pursuing all three models.

It expands our market to include customers that were not interested in an IP model, gives our current customers choice, and for Arm it creates a much larger profit opportunity.

We're not going to force any of our existing customers to migrate to this new model. We are welcoming customers to stay on their current IP/CSS model should they choose. Should they decide to embrace the silicon model, however, this chart illustrates how that decision would be significantly positive for gross profit dollars.

All right, so how are we landing the new chip business?

This business is based on customer demand from multiple companies and across hyperscalers and large enterprises. These are companies who prefer to buy a chip from us over building one from our IP.

We believe that Arm is uniquely positioned to build a CPU for the data center. If you look at all the Arm-based CPU chips, most of the technology already comes from Arm.

Because we have such strong initial demand, we have been able to quickly turn customer interest into actual business. As you see, we already have multiple customers lined up, and

we have line of sight to more than $1 billion in chip demand over the next two years. The vast majority will fall into FYE28.

Our biggest challenge is not finding customers who want our chips. It's actually memory shortages limiting our customers' ability to deploy our chips.

We expect material revenue from the Arm AGI CPU starting in FYE28, with an exponential ramp to around $15 billion in FYE31.

The first driver is increasing demand from new customers, naturally increasing chip volumes. We also expect increasing chip volumes, as well as rising complexity, to lift ASP significantly by FYE31.

Turning now to our IP business, starting with royalties.

Arm's royalty revenue has multiple secular growth drivers. The end markets to which our technology is being deployed are growing. We are gaining share. Our customers deploy more Arm-based chips. Increasing complexity is driving up core count, especially in the data center and high-end automotive chips, which leads to higher royalty per chip. And our most advanced technology commands a higher royalty rate.

Over the past five years, royalty revenue has grown at about a 14% CAGR. This has accelerated to over 20% in the past two years as Armv9 and CSS have started to ramp.

Looking forward, we expect that royalty revenue CAGR will be 20% over the next five years.

It might not reach 20% every year, as there will still be the occasional market downtick or inventory correction. Our licensing revenue overperformance in the past few years lends confidence to these future royalty growth rates.

We have discussed for a long time our efforts to add more value to our customers and to be compensated for that value. AI is a significant tailwind to our journey, as customers are pressed for both more compute power and faster time to market.

We all know we've essentially doubled the royalty rates from Armv8 to Armv9, and again to CSS. We could also charge a significantly higher royalty rate for each Armv9 generation, I am sorry a slightly higher royalty rate for each Armv9 generation. And a higher royalty rate for each CSS generation.

[01:39:47] The increase in complexity of the chips of today and tomorrow is also contributing to our royalty revenue growth.

In the data center, for example, AI — including agentic AI — is driving our customers to design an increasing number of cores into their chips. Tracking both customer increasing

core count over recent years and our customers' planned increase in future core count, we can see the number of Arm cores per chip increasing by about 20% per year.

The significant annual increases in cores and rising price per core as customers take more Arm technology is a big part of our confidence in continued robust growth in the data center royalty business.

One of the questions we often get is around our visibility into the future revenue trajectory. I think this tells the story.

Many of the contracts that underpin our royalty revenue forecast are already signed, and the royalty rates are already agreed in contract. We have delivered the technology. Our customers are building the chips and, in many cases, are already shipping the chips in high volume.

Looking out over years 2027 to 2031, 70% of revenues we're forecasting to collect are already covered with royalty rates set in contract. Even by fiscal 2031, the contracted base is still around 60%.

The remaining 40% is almost all with existing customers who we are confident will want access to the next generation of Arm technology, typically at higher royalty rates, higher core counts, and higher volumes than they do today.

We expect continued Armv9 adoption across all edge devices, from smartphones to smart glasses to smartwatches — pretty much every device with a screen. We expect a further boost from CSS adoption, not just in premium smartphones, but across all personal AI computing platforms, including in the PC space.

The combination of higher royalty rates from next-generation Armv9 and next-generation CSS will deliver outsized royalty revenue growth.

You can also see that we have a very sizable 65% of our forecasted royalties based on rates that are already under contract through FYE31.

Royalty contract coverage tends to be lower in edge devices than for Cloud AI and Physical AI devices due to very fast design cycles in consumer electronics.

Turning now to royalty revenue in the Cloud AI business, I've touched on the drivers of our expectation in rapid cloud growth already: rapid growth in the market, ongoing share gains, rising numbers of cores per chip, and delivering greater value per core create a powerful compounding story.

Our confidence is bolstered by the contracts that cover 85% of our expected royalties over the next five years.

We expect our healthy royalty growth in Physical AI to continue as cars, particularly autonomous and ADAS, continue to adopt more sophisticated silicon for the digital cockpit and driver assistance. We also anticipate continued share gains in this sector.

Our confidence is very high given the long lead times in automotive. Ninety-five percent of our royalties are under contract through FYE31.

As Drew explained, we are very excited about our opportunity in robotics. Much of this opportunity lies beyond FYE31 and thus is not captured in these figures.

Over the next five years, we expect that Cloud AI will be our fastest-growing revenue driver, even without the contribution from AGI CPU chips. When you add in chip revenue, it will surpass Edge AI in FYE30 and become by far the majority of our revenue in five years' time.

Finally, to licensing as you know, this has been growing well ahead of our expectations. At the IPO, we said that we would grow low single digits, which we quickly lifted to mid to high single digits, and it has more recently been growing over 20% per year.

This has been driven by a combination of the AI cycle, more customers getting access to Arm technology, subscription licenses and compute subsystem agreements, and the expansion of our licensing and design service agreement with SoftBank.

This license growth is the basis for the royalty commitment that you saw on the prior slide. We think all these drivers will continue, with the SoftBank licensing growing around high single digits, and with the AI cycle continuing to provide the majority of growth through more demand for next-generation CPU IP and compute subsystems at higher royalty rates.

Of course, the strong license revenue growth should lead to higher royalty revenue growth in the years and decades to come.

[01:44:58] As I mentioned right at the start, the Arm AGI CPU business, the royalty streams, and the licensing revenue all compound on top of each other. The chip business is targeting customers who either don't have the internal resources or don't have the desire to develop their own chips.

We do not expect the chip business to displace the IP business. Though some customers may ultimately choose to switch, it is accretive to earnings power as we previously discussed.

With around $15 billion of revenue from the chip business expected in FYE31, and another $10 billion of IP revenue, we are forecasting very strong revenue growth from the combined business over the next few years.

The good news is that we have already done a significant part of the heavy lifting when it comes to hiring the engineers needed to hit our plan.

If you've been following our financials for the past few years, you will know that we've already ramped R&D to support our product roadmaps.

Increasing R&D, combined with good execution, creates a virtuous cycle of new products driving revenue growth. From here, we are forecasting mid-teens OpEx growth through FYE31. Most of the incremental spending is our R&D investment in new technologies.

We expect our revenue by FYE31 to have grown by more than 2.5 times faster than our non-GAAP total costs. So as revenue and gross profit scale, particularly in chips, much of that incremental gross profit can drop through. That is the operating leverage in our model.

Our focus here is on the long-term earnings power of the company.

Before we get there, yes, we recognize that you have interest in the near term as well, and so we are affirming the Q4 guidance that we issued in February.

Back to FYE31: we see two meaningful profit engines.

First, we expect the IP business to reach about $10 billion of revenue, achieve a 99% gross profit margin, and deliver over 65% non-GAAP operating margin. We are today increasing our operating margin target by 500 basis points from our previous long-term target of 60%.

Second, we expect the Arm AGI CPU business to reach about $15 billion of revenue, with a gross margin of at least 50% and a non-GAAP operating margin of over 30%.

Putting those together, we have a consolidated business with $25 billion of revenue, industry-leading blended gross profit and operating margins, and more than $9 of non-GAAP EPS power in FYE31.

This is not a story of choosing between IP and chips. It's a story of combining a very high-margin IP model with a large, fast-growing, and accretive chip business.

Let me close on the three points that I started with.

First, customer demand is allowing us to materially expand our opportunity through selling chips. We already have line of sight to over $1 billion of demand from some of the companies that you met today, and are forecasting $15 billion of incremental revenue.

Second, our existing IP business continues to have strong underlying growth drivers, with the chip business compounding and not cannibalizing the IP business. We expect this to deliver around $10 billion of revenue in FYE31.

Third, by FYE31 the combined model is significantly accretive to revenue, gross profit dollars, and operating profit dollars, with more than $9 of EPS power. Because much of that investment is already in the base, the incremental economics are very attractive.

All right. With that, we're now going to conclude my presentation, and we're going to make a quick transition to start the Q&A session.

[01:54:58] Q&A session