Arm Q2 FYE24 Results

Wednesday, 8th November 2023

Operator: Good day and welcome to the Arm Second Quarter of the Financial Year Ending 2024 Earnings Conference Call. At this time, all participants are in a listen-only mode. After the speaker's presentation, there will be a question and answer session.

To ask a question, you will need to press star one, one on your telephone. You will then hear an automated message advising your hand is raised. To withdraw your question, please press star one, one again. Please be advised that today's conference is being recorded. I would now like to hand the conference over to your speaker today, Ian Thornton, Head of Investor Relations. Please go ahead.

Introduction

Ian Thornton

Head of Investor Relations, Arm Holdings Plc

Thank you, Abigail. Good morning, good afternoon, everybody. My name is Ian Thornton, and I'm the Head of Investor Relations at Arm. I'd like to welcome you to our Earnings Conference Call for the Second Quarter of the Fiscal Year Ending March 31st, 2024. I'm joined today by Rene Haas, the Chief Executive Officer of Arm, and Jason Child, Arm's Chief Financial Officer.

Hopefully, you will all have downloaded and read the shareholder letter. If not, it is available on the Arm Investor Relations website at investors.arm.com. As the shareholder letter provides a rich update on our strategic progress in the quarter, we will dispense with the prepared remarks from the CEO and CFO, and instead focus on Q&A.

Before we begin, I'd like to remind everyone that during the course of this conference call, Arm will discuss forecasts, targets, and other forward-looking information regarding the company and its financial results. While these statements represent our best current judgment about future results and performance as of today, our actual results are subject to many risks and uncertainties that could cause actual results to differ materially from what we expect.

In addition to any risks that we highlight during this call, important risk factors that may affect our future results and performance are described in our registration statement on Form F-1 filed with the SEC on September 14, 2023. Arm assumes no obligation to update any forward-looking statements which speak only as of the date they are made.

In addition, we will refer to non-GAAP financial measures during the discussion, reconciliations of certain of these non-GAAP financial measures to their most directly comparable GAAP financial measures, and a discussion of certain projected non-GAAP financial measures that we are not able to reconcile without unreasonable efforts and supplemental financial information can be found in the shareholder letter that we released earlier today. The shareholder letter and other earnings-related materials will be available on our website at investors.arm.com. And I'll now hand you over to Rene, who will make a brief opening statement before we go to your questions.

Opening Statement

Rene Haas

CEO, Arm Holdings Plc

Thank you, Ian. And as Ian mentioned, we have given you the shareholder letter in an attempt to minimize the opening remarks by myself and Jason, but I can't resist; I'll just start with a few comments to kick off.

IPO

We are very pleased, following the IPO process, to kick off our very first quarter as a public company, and the quarter was excellent.

Revenue

We had record revenue, really fueled by demand for all Arm products, which has driven our licensing numbers up over 100% year on year. This is largely driven by what I would consider as an AI R&D super cycle, where people are investing more and more in new technologies to take advantage of the huge opportunity going forward.

Royalty

On the royalty side, slightly down year on year. However, the new businesses that we have emphasized in terms of our new growth strategy into the cloud and automotive were up approximately 20%, and the financial results relative to profitability were excellent.

Summary

So in summary, very, very pleased about the first quarter, and very, very pleased about the results we've shown as the first of our many quarters going forward as a public company. So with that, I will turn it over, I suppose, to Abigail to queue up the questions.

Q&A

Operator: Thank you. At this time, we'll conduct the question and answer session. We ask that you limit yourself to two questions today. As a reminder, to ask a question, you will need to press star one, one on your telephone and wait for your name to be announced. To withdraw your question, please press star one, one again. One moment, our first question. Our first question comes from Toshiya Hari with Goldman Sachs. Your line is open.

Toshiya Hari (Goldman Sachs): Hi, good afternoon. Thank you so much for taking the question. I had two questions, maybe one for Rene. I think your Royalty business was up mid-single digits or 5% sequentially, and I think units were down about 5% sequentially. So can you speak to what drove your revenue there? Is it chip ASPs? Is it royalty rates or is it a combination of both? I think during the IPO, you guys had talked extensively about the transition from v8 to v9. So I'm guessing that was one of the bigger drivers. But if you can provide a little bit of context, what drove your royalty business on a sequential basis that would be helpful.

Rene Haas: Sure. Yeah, thanks for the question. And you're right. It is largely driven by the transition to v9 accelerating, particularly across the smartphone segment. Additionally, as we had mentioned earlier in our discussion with analysts, we're seeing growth now across the

Automotive and Cloud Infrastructure business. And those have different royalty rates than our Smartphone business does. So as a result, what you're seeing is that even with units down, the overall numbers are actually up in terms of revenue.

Toshiya Hari: Great. And then as my follow-up; during the IPO, you had shared with us that roughly, I think it was 97% of estimated royalties under contract in fiscal 2025 kind of being locked in from a royalty rate perspective, 81% for fiscal 2026. So I was hoping now that a couple of months has gone by, if you can provide an update on those numbers. Thank you so much.

Rene Haas: Yeah, again, thanks for the question. And I would say we're about at the same level in terms of where we are in terms of mile markers towards progress. We're still confident in terms of the numbers that we had talked about in the past. But more importantly, everything's tracking as we would expect at this point in time. So those numbers are still unchanged.

Toshiya Hari: Thank you.

Operator: One moment for our next question. Our next question comes from Ambrish Srivastava with BMO Capital Markets. Your line is open.

Ambrish Srivastava (BMO Capital Markets): Hi. Thank you very much. My first question is, if I look at the fiscal year guide, given you had such a big upside on the licensing side versus what we were modeling for. What's the mix embedded in the guide between royalties and licensing?

Rene Haas: Yeah. Well, for the – so if you kind of unpack our expectations for the back half of this year, the next two quarters, we're expecting royalties will flip to positive, I'd call it, single-digit growth in Q3. And then by Q4, we expect to see double-digit growth on a percentage basis.

Licensing, we expect to continue to be strong. I do expect our assumptions on Q3, we do have some lumpiness with our licensing business, especially with ASC 606. So we do have some, as we always have, some large deals that are in play. As of right now, I think versus what we thought a quarter ago, I think there's going to be a little more falling into Q4 versus Q3. So I think, with our guidance for Q3, our expectations are to be somewhere in that kind of, call it, 0-10% growth on a year on year basis. But we expect pretty significant growth because we do expect some pretty big license deals coming in Q4.

In terms of the, I guess, the mix of revenue split, hard to say at this point. It's going to be, I'd say, yeah, closer to 50-50 then, but that really depends on how strong the royalty recovery is in Q4. There's all sorts of industry reports. And I think if you look at most of the guidance, as well as we have a pretty easy comp from a year ago, it could get maybe closer to 60% of total, but we'll see.

Ambrish Srivastava: Got it. No, I think you gave enough details. And my second question is a little bit longer term. On AI, you guys have been pretty detailed about giving us a percent. I think you said, correct me if I'm wrong, 43% of royalties are driven by AI. I just wanted to understand going forward, what's going to be the driver? And I'm assuming that majority of the 43% is on the edge. So, as we look forward, is it going to be more, as we have seen with Grace Hopper, which obviously volumes are very small, is it going to be more infrastructure-

driven i.e. something like Hopper or the data center hyperscalers or is it more going to be more of the same, more on the edge on the mobile side? Thank you.

Rene Haas: Yeah. So, this is a very, very fast-moving market relative to the models that are being released that are almost on a daily basis, combined with just how quickly some of these agents are moving across different devices. So, when you think about, for example, the endpoints, a PC or a smartphone that could be running a ChatGPT agent or Microsoft Copilot, just a quarter or two ago, we may not be classifying them as devices that were running AI. So, our expectation is that, increasingly, all of the devices that exist in the overall value chain from the cloud to the endpoint, and the endpoint can be the smallest sensor with a compute engine, will need some level of AI capability, which is why our licensing activity has been as strong as it is. People are looking to add as much capability in terms of compute to capture the workloads that are being developed. And in some cases, it's really a function of making sure you have enough compute capacity to run the model when you don't even know yet what the model is. So, I think we are in a very interesting time relative to how this overall market is going to play out. To specifically answer your question, whether it's the endpoint or the cloud. Both. And I think it's going to be a rapid acceleration across the next few years, where a few years from now we won't talk about the percentage of devices that have AI in them, it will be table stakes that they all do.

Ambrish Srivastava: Got it. Thank you, Rene. Good luck.

Rene Haas: Thanks.

Operator: One moment for our next question. Our next question comes from Vivek Arya with Bank of America. Your line is open.

Vivek Arya (Bank of America): Thanks for taking my question. Rene, for my first one, I'm curious. You had the IPO two months ago and the process started before that. What have been the big changes in your macro and industry assumptions, positive or negative, since the team went through that process? Any color by end market, geography?

And specifically, what I'm trying to get to is that if you look at the way you were thinking about royalty revenues in December and the next few quarters, have they changed in any way, positive or negative, given any potential changes in your macro assumptions?

Rene Haas: Yeah. So thanks for the question. We haven't changed anything in our models that we're talking about publicly relative to the years out forecast in terms of what any assumptions are relative to the numbers. But going back to the commentary that I made on the previous question, I do think what we're seeing from a macro standpoint is people figuring out across every end device that's being built – and again, that end device can be a smartphone, it can be a base station, it can be a laptop – people are figuring out how to make sure they have enough compute capability to take advantage of these applications and models and agents that are being introduced almost daily. So from the perspective of have we changed our models, not anything we're talking about publicly; but what I can say and feel, and again, you see it relative to the licensing activity being as strong as it is, there is absolutely a rush to ensure that there is enough compute capacity in the end devices.

One of the enemies of growth in our business is getting to good enough from a compute standpoint. And we are nowhere close to good enough. And that ends up meaning a drive for R&D to figure out just how to handle all these new capabilities.

Vivek Arya: Thank you. And for my follow-up, there's recently been excitement about the combination of Windows and Arm. I know there have been previous attempts which right, were not as successful. I'm curious, Rene, how do you think about the potential for Windows to succeed on Arm-based devices? Is that a tangible factor for 2024? Is that a factor for 2025 and beyond? Just give us your perspective on how successful it can be and what is different this time versus the prior attempts that Windows has had in dealing, interacting with Arm technology.

Rene Haas: Yeah. The Windows and Arm ecosystem is one that I have a personal history with, having been there from the very, very beginning. And we have come a long, long way from that point relative to readiness of the application ecosystem, readiness of developers, native applications. So I think from a software standpoint, everything is now in place for the next growth cycle.

One major ecosystem not called Windows has moved over 100%. And I think what they've proved is that there's amazing battery life, amazing performance, and amazing application compatibility across a number of different dimensions. You can run Windows on that alternate ecosystem and get really, really good performance. So I think we are on the cusp of getting over this hill. I feel very, very good about the growth projections for Windows on Arm.

Vivek Arya: Thank you.

Operator: One moment. Our next question. Our next question comes from Charles Shi with Needham and Company. Your line is open.

Charles Shi (Needham and Company): Hi. Good morning. Good afternoon. Thanks for letting me ask a couple of questions. Maybe the first one, I want to ask since export control, US government put out there all the rules and that they recently updated that. I wonder if you can provide a comment on whether that has any impact on Arm's business.

And specifically, since you have a really distinct business model, especially on the royalties' side, to the extent that when your customer may be put on the entity list, are you still able to collect the royalties? That's a related part of the question. Thanks.

Rene Haas: Yeah. So for starters, every time these new export rules come out, we have a team of folks in our Trade Compliance Group that go through the information in a very detailed way and trying to understand exactly how it might impact our company. I can say that we obviously would comply with any kind of export restrictions that apply to our technology or what we build.

The latest round of documentation that came back from the US, I would say not so much, in fact, probably not at all in terms of the impact there. Generally speaking, the impact to Arm is not that significant for two reasons. One, the components and pieces that we build are generally under the thresholds that have been listed by the United States government in terms of export control. And secondly, in the areas where there's a de minimis content in terms of US people working on the design because much of our technology is actually designed and developed outside the United States, in Continental Europe and the United Kingdom, we're not impacted

quite so much. So generally speaking, the last set of rules did not impact Arm, and we have broadly speaking not seen a large impact there.

To your question relative to how it works in terms of do we collect royalties if someone's on the entity list, etc., it's pretty simple. If an end product that contains our technology can't be shipped and there's no revenue to be derived, then we feel the ripple effect of that.

Again in the last quarter, no impact from that, and as we forward a forecast to the guidance that we gave for the remainder of the year, nothing that we see on the horizon that's impacted there.

Charles Shi: Thanks, Rene. Maybe a second question I want to ask is on operating margin, you provided the full-year operating expense. What do you expect? It kind of implies the fiscal Q4 margin is going to be down. I mean, operating margin is going to be down. I mean, even if I back out that one time, increasing Social Security taxes roughly $45 million, it's still down a little bit. So how should we think about what's driving that kind of year-end margin weakness? And how should we think about going into next fiscal year? I know you have a long-term 60% operating margin target, but how do we get from here to that 60% on an annual basis? Thanks.

Rene Haas: I'll let Jason take that one.

Jason Child: So, the way I'd answer it is the margin, I think at the midpoint, it's in kind of the high 20-ish percent range. Obviously, the expenses I'm giving are independent of whether we come in at the middle or high end or even low end of the range. If you assume that OpEx and you also account for the one-time impact on Social Security, which relates to the stock vesting that was tied to the IPO, that's about 600 basis points of impact.

So I think if the midpoint applies somewhere around 20-ish, 8-ish percent, if you take out that adjustment, that would put you in the kind of the mid-30s and low 30s if you're at the bottom end of the estimate, and you'd be closer to low 40s if you're at the high end of guidance or maybe high 30s. And so that's the mechanics in Q4.

Going forward, I would say, we do expect to deliver incremental margin in the, I would say in the medium term, i.e., over the next few years, that will start to approach that 60% target that we're aiming towards. However, you just saw we've actually added about 1,000 people in the last year, and most of that is because of the headcount, mostly engineers, about 85% of those heads that we added in the last year are engineers. And those folks are specifically working on the compute subsystem, and the increased kind of complexity needed with all the designs that folks bought this quarter, and are going to forecast to buy in the coming quarters. And so that will put maybe a little bit of pressure in Q4, but I still expect you to see us deliver a solid 40-plus percent overall margin for this year and then certainly for next year. And I do expect us no change in our trajectory to get to that 60% margin over the coming years.

Charles Shi: Thanks, Jason. I appreciate the color.

Jason Child: Thank you.

Operator: Our next question. Our next question comes from Chris Caso with Wolfe Research. Your line is open.

Chris Caso (Wolfe Research): Yes. Thank you. Good afternoon. The question is another one on AI and obviously a lot of discussion about AI capabilities and client devices. Can you go

into a little more detail about how Arm monetizes that? Is it from a higher per chip royalty? Is it from a better mix at your customers, maybe some higher device ASPs? How do you see that playing out over time as AI gets embedded in client devices?

Rene Haas: Broadly speaking, the way I would think about it is whenever you're running one of these AI clients or assistants or agents, it's going to require a significant uptake in terms of compute capability, both in terms of if there's an in situ accelerator and/or through the CPU complex, keeping in mind that in a client device when you run these AI agents or whenever you're running something that's going to be a copilot of some sort, nobody wants to see their battery life suddenly go down 40% in terms of everything that was involved in running the algorithms. So what that means for us in the broad sense is I expect it's going to be a higher need for more compute capacity. We'll see more advanced cores. We'll see larger cores, more v9, which in the end game should mean higher royalty rates for us. That would be our belief going forward in terms of just the mega trend.

Chris Caso: Got it. If I could just go back and, Jason, go back to some of the comments on OpEx, you spoke about them in terms of operating margin. But just as we look at modeling operating expenses as we go into next year, obviously it sounds like we should take out that one-time Social Security tax in the fourth quarter. But what do you – well, I guess what would be the path of OpEx as you go into next year, and to what extent is that dependent on the revenue stream? Are you modulating OpEx according to revenue or are you just spending where you need it?

Jason Child: Yeah. I don't – in terms of providing guidance for next year, I'm not ready to do that. I would say kind of our long-term model approach isn't really any different, so I'm not ready to kind of go provide any updates to that. But I would expect – we definitely will be growing OpEx less than revenue, and so I do expect to get incremental margin. I just can't say exactly kind of what the quantum is for next year until we get a little later into this year.

Chris Caso: All right. Fair enough. Thank you.

Operator: One moment for our next question. Our next question comes from Andrew Gardiner with Citi. Your line is open.

Andrew Gardiner (Citi): Thank you very much for taking the question. One on licensing to start with, clearly, you beat expectations quite handily in the quarter on that front, and I suppose this was a part of the business that during the IPO process you explained was an area where you guys had pretty good visibility. It was fairly predictable given the timing of contract renewals. So, was the beat a pull forward of demand, or are you seeing the, as you put it, Rene, the AI super cycle? Is that driving upside to the pipeline that you had had there earlier in the year?

Rene Haas: Yeah. Thank you for the question. It's a good question. I would say it was expansion of deals that we had visibility on. What we have generally pretty good visibility is when are renewals due and/or when our customer is going to be looking at uptakes of new technology. I think what we saw, broadly speaking, was the partners that we knew about that we were expecting deal closure; their appetites got bigger over the quarter, and they took more technology, so the size of the deals were larger.

Jason Child: I'd say the one thing I would add versus expectations, in the quarter, if you look at revenue certainly growing 28% is strong, but also RPO or total backlog actually grew $700 million both year on year and even sequentially quarter over quarter. And if you actually do the math on total bookings or RPO bookings, revenue plus change in RPO, you can actually see that we did over $1.1 billion in bookings in the quarter, which is the best quarter in our history. So, that definitely, to Rene's point, while we had insight into the pipeline, the size of the deal did expand and get bigger, and certainly a lot of that we think is tied to this kind of deeper investment R&D given everything that's happening in AI currently.

Andrew Gardiner: Well, and especially then you lean into my next question because in the letter you say that of that RPO, you're expecting to recognize 28% of it over the next 12 months, which given roughly where expectations were following what you guys had given us through the IPO process, it looks like you've already got two-thirds of that licensing revenue in hand. So, even if you don't sign that many more deals, it looks like you're pretty well set for the rest of the year in terms of licensing. So, is the expectation not too conservative on that front at this point?

Rene Haas: You're talking about expectations for next year or the back half of this year?

Andrew Gardiner: Well, they're 12 months forward, right? So, what you're saying about 28% of the RPO to be recognized over the next 12 months.

Jason Child: Well, we feel good about the guidance, and we did increase the targets for the back half of this year versus what we thought at IPO. And we haven't talked about next year, but certainly given the tailwinds that I think exist on the licensing side, and then now that we are seeing signs of progress on the royalty side, we're not ready to finalize the numbers, but there definitely are tailwinds.

Rene Haas: Yeah I think you're reading it correctly. To Jason's point, a billion dollars in bookings in a quarter, there were years where we didn't do that in a year, minus a few hundred million. So we are very very confident about the level of backlog we've built up and how that gets recognized over time. So we feel very confident about that.

But more importantly, from the financials, not to minimize that, it does underscore very, very strong demand for Arm technology relative to the R&D investment that people are making. We see no – in the midst of inflationary pressures, geopolitics, lots of unknowns about end markets, what we're not seeing is people pushing out deals, not making investments, staying with the current generation of technology for a cycle, none of that whatsoever. What we're seeing is as much as possible an acceleration to make sure that there's as much compute capacity in the end devices that are being built. Largely, back on the AI piece, because these models are changing so fast and being evolved so quickly, the understanding of what amount of compute capacity you need to take advantage of the capabilities that are being introduced is a bit of an unknown. What you do know is that you probably don't have enough compute in the devices that you've designed today, so adding to it is critical, which is why we saw such an expansion this last quarter.

Andrew Gardiner: Thank you very much.

Operator: One moment for our next question. Our next question comes from Harlan Sur with JPMorgan. Your line is open.

Harlan Sur (JPMorgan): Good afternoon. Thanks for taking my question. Another one on licensing. As you mentioned, there's some timing-related dynamics and revenue rec dynamics regarding licensing in the December and the March quarter. Some of the uncertainty is to be expected, especially on large deals, as you mentioned. So on the fiscal second half, is more of the uncertainty on timing of licensing deal closure or more around the revenue recognition profile of those signed deals?

Jason Child: Timing. Yeah. I mean, as Rene just mentioned a moment ago, deals certainly have the capacity to change in overall size or quantum, but that, for the most part, usually provides more upside. So in this case, it's really just about timing, and as these deals – I mean, given we did $1.1 billion in bookings in the last quarter, these are very, very large deals that require lots of complicated approvals that go to the highest levels of these organizations that can take a while, and that's hard for us to predict.

And so, it's certainly, our view on – I would not evaluate just Q3. I would evaluate Q3 plus Q4, and Q3 plus Q4 is what we took up in our guidance and feel very good about the trajectory.

Rene Haas: Yeah. So, having been with the company ten years and watched how this process works, we generally have pretty good visibility on a six-month basis. But to tell you whether something is going to close in December or January, given the fact that there may be a lot of legal language to review, it takes approvals, December is a holiday period, could be a bit out of our control. So to the level of being potentially conservative on a quarter timing, I think that's potentially the detail you're extracting here; but our confidence that the deal will actually close is quite high, given that we know what the needs are, which is why, to Jason's point, the guidance went out.

But more importantly, and I give that example of December, January as both a figurative one and a real one because that's exactly what we might be looking at here, and it makes a big difference on which side of the boundary it hits. But our degree of confidence that the technology will be needed by the end customers is quite high.

Harlan Sur: Perfect. And then maybe mid to longer term, the step-up in royalty rates over the next few years is, in large part, driven by the adoption of your total compute or compute subsystem solutions, where you're not only delivering more CPU or MCU cores to your customers but also integrating some of the key subsystem blocks like bus architecture, cache memory management, memory controllers, security, etc., saves your customers significant engineering, design, and validation costs, and in turn, you guys get a higher royalty rate. TCS has been very successfully adopted by several of your large mobile customers. Can you guys just give us a sense on the traction of driving more subsystem solutions into your automotive, industrial, PC, data center, compute customers, and any way to quantify the momentum there?

Rene Haas: So we just announced for our Infrastructure business, our CSS Partner program, where we are engaged with people like TSMC, Cadence Synopsys, Intel, etc., to more rapidly accelerate partners who want to move into this solution space. That's really been driven by the fact that the demand for this has been higher than we expected.

And if you just step back and think about, well, why would this be of such high interest to end customers, from – if a customer is designing a, an SoC that had one microcontroller core into it or two, handing IP to that customer, they then develop their chip, they put their IP around it, and then ultimately developing end product. That model works very well for certain

segments. But if you're trying to build something that is for a laptop or a cloud infrastructure or a 5G switch, and you're putting down 16 CPUs, 30 CPUs, 100 CPUs, and you're trying to incorporate the fabric, and you're trying to incorporate the cache memory interfaces, and you're also trying to build a chip that used to take you 16 weeks from TSMC to now 26 weeks, you've got 10 weeks added to your cycle time, and now the subsystem part that you have to integrate is really hard.

So, particularly in the cases where many of these subsystems are exactly what we just described, they are the compute, they are the block that Arms delivers, it actually makes a ton of sense for partners to look for us to provide that. That applies extremely well, as I said, to the markets I mentioned, including automotive, ADAS, including mobile. So we are over-subscribed on this, and again, on the guidance standpoint, not going to change any words that we provided in terms of what the overall future looks like, but it is, I think, a huge value driver for our end customers. So we see this direction of travel only increasing.

Harlan Sur: Absolutely. Thank you.

Rene Haas: Yep.

Operator: One moment for our next question. Our next question comes from Ross Seymore with Deutsche Bank. Your line is open.

Ross Seymore (Deutsche Bank): Hi, guys. Thanks for asking the question. For my first question, I just wanted to get into the implied December, and, well, actual December guide and implied March guide. It looks like you're missing the Street a little bit in December, but then beating it in March. Is that just the lumpiness of the licensing you've mentioned a bunch of times? And I guess more precisely, what's the general expectation on the royalty side of things, especially in the March, considering that there's lots of moving parts cyclically right now, but seasonally, that doesn't tend to be the best of quarters for your mobile business, the smartphones, etc. So just the puts and takes on those would be helpful.

Rene Haas: Sure. So, yeah, on the licensing side, exactly as we described. Our six-month visibility is very, very good. Our month-to-week visibility is a little fuzzier. And as a result, we're going to err on the side of caution and not overstep, but make sure we deliver on what we say we're going to do. And as I said, we're extremely confident in the deals that we've identified and the need for the technology.

I'll let Jason comment a bit more in terms of the direction of travel on royalties. But broadly, we've seen three quarters of sequential growth. We have a lot of strong indicators from partners that we are out of the trough and climbing out of the trough relative to the direction of travel, the slope of the curve, I'll let Jason speak to that. But generally speaking, our indicators are pretty strong as far as that market goes. And as I said, in the other markets where we continue to grow and gain share in cloud and automotive, our confidence level is quite good.

Jason Child: Hey, Ross, on the royalty side, what I would say is in this most recent quarter, we did see positive sequential growth return. And if you look at some of our largest partners, they've seen the same. Our guidance, which is in part, we are looking at some of the industry reports as well as also looking at some of the forecast from our partners.

And I think we're forecasting something pretty similar to what others are saying, and that is we're expecting to see somewhere in the probably high, kind of mid to high-single digit sequential growth in the next, I would say, next two quarters, each of the next two quarters.

And so, when you factor that into the downturn that really kind of took hold last year, that means you're going to see year on year growth in royalties get back to, I'd say, positive single-digits in Q3, and then I'd say get definitely kind of well into the double-digit growth by Q4.

And then we'll see from there, but obviously the comps certainly are easier as well in the first half of next year. So I think we have a good setup, but as long as this kind of recovery that us and our partners are seeing continues to come to fruition, it should be a great setup.

Ross Seymore: Thanks for that, Jason and Rene. I guess for my second question, this has been a bit of a rolling correction. You just talked about some of the dynamics coming out the other side, thankfully, but some of the other markets, automotive, industrial, broad-based ones, seeming like they are just rolling over now to the downside. What's the impact to Arm if some of your more client businesses improve? I realize you have a bigger exposure to those, but as far as implied royalty revenue rates, those sorts of things, if the automotive and the industrial IoT side of things weaken, can you make up for that with the mobile side of things, the client side, or are there trade-offs that we all need to appreciate?

Rene Haas: Our expectations are the combination of increasing v9 products that actually are starting to shift, because we're still – on the royalty side, we're still in the relatively early days of v9 shipments. Think of it as being probably somewhere in the like 10% or so of our royalties are shipping with v9. So since really a lot of these designs were sold over the past couple of years, those products are really just starting to come to market now. So that's certainly going to be a tailwind to growth, both as we see recovery in units, but then also as we see these higher rates flow through.

There's also a lot of – there's some strong growth drivers certainly happening in the infrastructure business where certainly while the industry had been a bit slow in this last year, things seem to be picking up, and then obviously with all of the different hyperscaler AI efforts, there's a lot going on there. And so, we're going to continue to gain share on the infrastructure side, specifically on the cloud compute side.

Auto is an area where we've also seen pretty strong gains. Certainly the market, I think our expectations won't be as strong as it's been over the last year or so. I think their inventory levels have probably caught up a bit more, and so at least from – I'd say from an auto inventory, not a chip inventory perspective.

And then our expectations are, I think, similar to yours. We don't expect IoT to be a big growth driver in the near term. I think certainly could be further down the road as we see what AI does for edge computing and whatnot, but for right now, probably not a strong growth driver for next year. So that's I'd say kind of how we're looking at it on a lot by lot basis. And we'll certainly let you guys know as we learn more and progress throughout the year.

Ross Seymore: Thank you.

Operator: One moment for our next question. Our next question comes from Pierre Ferragu with New Street Research. Your line is open.

Pierre Ferragu (New Street Research): Hey, thanks a lot for taking my question. I'd like to come back to OpEx. I mean, you've given a lot of clarity on numbers. Thanks a lot for that. But my question is probably a bit more generic. So if we say like Q4 to Q4, your OpEx has increased by – will have increased by about like 20% – 20% or so, I'd like to better understand like the operational drivers of that increase and in general how your OpEx is increasing. Maybe I have like too much of a simplified view of your model, but to me, you spend a lot of OpEx on developing products. And then after that, you license these products. And then after some time, these licensed products end up into the products of your client. And so I would have expected OpEx to grow very early in the process and not really at that point in time when actually your licensing activity is very, very rapid, and the next stage is for your clients to actually integrate your IP in their product. So there is probably an element of the model that I'm missing, and I'd love if you could help me better understand that.

And also one thing that could help is to give us a sense of that increase in OpEx, is that almost exclusively product development, like R&D or actually there are a lot of like business development and managing client relationships within the start of these new licensing program that maybe we are missing the way we understand it?

Jason Child: Thanks for the question. So first, yeah, so if you look back over the last year, we've added about 17% increase in headcount, about a 1,000 people. 85% of those heads are in R&D. And so, while there may be some G&A and some sales, whatever we're talking about, it's 15% is everything that's non-R&D, and our total R&D as a headcount and a percentage of total OpEx runs about 80%, so we're, I'd say, pretty consistent in terms of our overweighting towards building capabilities, future capabilities. And most of this R&D to your point, when we're – the R&D teams are creating designs and that we're then of course selling, but we're constantly working on the next evolution. Right now, what's been happening over this last year and the reason why there's been such a significant amount of hiring is because as we started selling the compute subsystem capabilities because that's certainly what customers have clearly been and partners have been looking for from us, that's required us to build a solutions engineering team, which is a bit of a new muscle for us. And so that team went from, I'd say very, very small a year ago to now about a 1,000 people. And so that's been a big area of hiring and a lot of the contracts and a lot of the royalties that we're going to get from those hires, we're going to start seeing really not until we get into our fiscal year in 2026, as we talked a bit about during the roadshow. So it is to some extent, some forward investment, and but there is contracts. I think someone earlier asked the question that we had something like 81% of our royalty contracts were already signed for financial year 2026. So the vast majority of the benefit we're going to get from them has been signed but will not ship until we get into late 2025 and early 2026.

Rene Haas: Yeah, and as far as the product development cycle goes, your question is a very good one relative to how to think about product development and cycle times. We put out products very, very frequently. The mobile and PC world, they need to see new CPUs every single year. They need to see new GPUs every single year. So we are developing the next generation product and releasing something literally on an annual beat. The hyperscaler market is probably every two years, but we also do performance scores and efficiency scores on a bit of a tick-tock basis. And then automotive cores are probably anywhere between two to three years. So our people are always working.

And I would also say that, one of the things we did during the SoftBank years to invest is we actually got out of a number of commodity businesses that we were in, such as video IP, display IP, where we were highly undifferentiated, and used those resources to develop a Neoverse CPU roadmap, and to develop a Automotive AE roadmap. AE is a Automotive Enhanced, including functional safety.

So there is a constant treadmill of products and CPUs and GPUs and NPUs that are being developed. And then to Jason's point, when we start to put those into subsystems, that's a new output and a deliverable. So the IP group provides those IP cores to the Solutions Engineering Group, which is then essentially the group responsible for stitching them together as a subsystem.

So it is an ongoing engineering flywheel that does not abate. And as I said, relative to the broader market demand, we're nowhere close to good enough. People want smaller, faster, better all the time, which is what we're working on.

Pierre Ferragu: Thanks, that's very, very clear. Like this subsystem engineering really answers the question I had. Maybe one quick follow-up, if I may. If I look at where you're guiding and how it compares to like the sellside consensus, you're kind of quite significantly higher on OpEx for Q4. And so my question here is maybe a bit provocative, but does that mean that the sellside analyst didn't listen to you carefully enough during the IPO process and mis-modeled a bit OpEx in the near term or does that mean that today, compared to three months ago, you've actually built up OpEx faster than what you were thinking three months ago?

Jason Child: I can answer that. The short version is the assumptions back when we had Analyst Day in early August was for a lower stock price than we actually ended up issuing most of the equity at. And also I would say higher social security taxes, especially in the UK, where they're about 2x higher than they are in the US. And when we flow through the increased stock price and then we flow through the increased taxes, that's the driver. So, it's not an ongoing cost driver. That's why our expectations are still to deliver somewhere in the 40% range of non-GAAP operating margin in the near term in this year. And we still have a long-term target that will be getting to the 60% range. There's no change to those targets. It's the short-term aspects of dealing with some of the IPO-related costs which just ended up being a little higher than we had previously forecasted.

Pierre Ferragu: Thank you, very clear.

Jason Child: Thank you.

Operator: One moment for our next question. Our next question comes from John DiFucci with Guggenheim Securities. Your line is open.

John DiFucci (Guggenheim Securities): Thank you, and thanks for taking my question. My first question has to do with the related party revenue, which was flattish year over year versus the rest of the revenue was up just almost 40%. Can you provide more color around the license versus royalty mix for the related party business and how we should think about that going forward?

Jason Child: Sure. Thanks for the question, John. So related party revenue is Arm China, where we have a couple hundred customers in China that are all aggregated and treated as one

based on how the joint venture was set up. You should think of – so China did – it still grew, I call it in the low-single digits, but as a percentage of total, it fell from about 25% in the most recent period down to about 20%. And that's just because the rest of the world just grew so much faster.

In terms of the mix of license versus royalty, think of it as being pretty close to 50-50. A little over 50% of license, a little less than 50% is royalty. But all of that revenue for Arm China is treated as other because of the fact that the way the Arm China joint venture is structured.

John DiFucci: Got it, okay. Great. Thanks, Jason. That's really helpful. Actually, how about going forward, how should we be thinking about that mix going forward for Arm China?

Jason Child: I think the mix is pretty similar. They historically – at least that's the color we have right now. They've historically been I'd say closer to 50-50 than the rest of the world, which is closer to 40-60, 40% license, 60% royalty. In terms of – we'll certainly update you if we see a change, but that's our expectation, I'd say, for at least the back half of this year.

John DiFucci: Okay, great. And second question. When we think about the guidance for next quarter and the year, obviously the macro backdraft has an effect, and you guys have your own crystal ball, I guess we all do, and it has an effect on both license and royalties. But the royalty part is really something you probably have less control over and visibility into that timing, even though there were a lot of questions here on license, which probably you have more visibility there.

Jason, you did hit on the royalty visibility in one of the questions, but I just want to make sure I understand what's implied in guidance in regards to units – and I know there's a lot of other things that affect royalty revenue, most of which are going in the right direction. We can all see industry analysts' forecasts for units, and you guys have intimate relationships with your customers, so you can have even better visibility into that. But I'm just curious, when you look at industry analysts' estimates on units, so are you assuming about the same, a little below, or even perhaps a little even better, or do you just see things a little bit better than industry analyst estimates?

Jason Child: Sure. Okay, so here's what I would say. So in this quarter we just reported we were minus 5% on royalty. If I compare that to our three best comps, at least closest in terms of mix, you could look at MediaTek, you could look at Qualcomm, and you could look at TSMC. Those guys were all between minus 11 to minus 24%, as they just reported the last few weeks. So we had stronger growth than those guys, primarily because of our share gains that we're seeing in infra and auto, and then also in some part because – or in part because of the Armv9 adoption. So we are expecting that – and I think those guys as well as counterpoints from the other industry analysts, all seem to be kind of triangulating around flipping to positive sequential growth in this last quarter, and then in the current quarter, expecting to see that get kind of in the high-single digits approaching double-digit range.

And that's pretty consistent with what we're expecting, and then same for Q4. So from everything I can see – and we do get paid royalties on 7 plus billion chips per quarter, so we do see quite a bit from a bunch of folks. As far as I can tell, we're kind of all triangulating relatively similar impacts. I think the difference probably are our v9 rates, maybe the piece to that is why we typically grow a little bit faster than maybe some of the others. Does that answer your question?

John DiFucci: Yeah, that's really helpful. Thank you very much, Jason.

Jason Child: Great.

Operator: One moment for our next question. Our next question comes from Matt Ramsay with Cowen. Your line is open.

Matt Ramsay (Cowen): Yes. Thank you very much, guys. I guess, Rene, my first question is just the phenomenon we've seen with Gen AI computing in the last 12 months or so. I just wanted to get your take on what it may mean for Arm. Is it – do you view it – I mean, maybe the balance of it as a positive catalyst in that, it can pull forward a bunch of Arm server designs that might be hosts for accelerators in a data center and things like that, or do you have concern that maybe it limits or shrinks the CPU TAM that you can grow into organically? How do you think that puts and takes of that are? Thanks.

Rene Haas: Yeah. No, Matt, thank you for the question. I think it's broadly a positive, and the reasons for that are as more and more of these LLMs are being used in cloud data centers for training, obviously, any accelerator or GPU needs a CPU, that's table stakes. Then when you drill down one level deeper in terms of what is the type of CPU that you need, you need something that's energy efficient, you need something that's very, very low power, something that you can customize in terms of an overall system, then it gets very interesting because when you look at TCO, but more importantly, total system power, by actually developing a custom Arm-based SoC that can interface with one of these accelerators, you're actually able to get a very, very high degree of customization relative to power efficiency and throughput. One of the biggest power requirements or beasts that – from these applications is actually in terms of feeding the engine, which is memory bandwidth. So if you can design something that's actually custom that can interface into the accelerators, it can be hugely beneficial. I think, as you know, there's a lot of work going on inside the community today developing custom chips that are Arm-based, so I think that's only a net positive. Then when you layer on top of it the leading player, leading actor in that field on accelerators is obviously NVIDIA. NVIDIA has been doing a big benefit in terms of making the CUDA drivers for the A100 and H100s available on Arm. So whether using a standard product from an Ampere or building a custom chip based upon Arm-based cores, I think, this Generative AI workloads being pushed onto AI clouds is a tailwind for Arm. We're pretty excited by it.

Matt Ramsay: Thanks, Rene, for that. As my follow-up, I think in another question, Jason, you spoke a little bit about Arm China. Given all the things that are going on regulatory-wide, I wanted to step back a bit. And you guys control the IP that gets given to Arm China for them to then do things with a license into China and the royalties come back out the other end.

I guess what I want to get a little bit more granular on, given the dynamic environment we're in, is just what kind of visibility do you actually have through the structure of Arm China into the forward licensing trends, and then, from an audit perspective, the royalties that are coming out on a quarterly basis. Just like the level of visibility you have to the operations that go on within that organization as IP goes in and royalties come back out.

Rene Haas: Yeah, so Matt, I'll let Jason comment on the audit component and also the integrity of the information that comes back. But just a couple of things I wanted to note on Arm China so that you and the rest of the group can understand. First off, from a delivery standpoint, when Arm China signs a contract with a PRC customer, the IP actually goes from Arm Limited

directly to the customer. It doesn't go into Arm China, so they are not a holder, if you will, of the product. The product is essentially downloaded directly from our servers to the end customer.

Secondly, for many of the high-value designs, particularly, whether we're working in the networking space or the cloud space or automotive, they're generally working with our latest edge technology. And because of that, there is a lot of interaction between the customers in China, the Arm China salespeople, and Arm Limited marketing and engineering. So we have really, really good visibility in terms of when these large strategic deals are being consummated because generally speaking, everything around demand creation and the technical interactions between the engineers at the customer and the engineers at Arm is something where you have complete visibility into.

So we have a very, very good idea of when large deals will close in China just by the nature of the relationship between our engineers, the partners' engineers, and the sales folks for Arm China. I'll let Jason talk about audit and things of that nature.

Jason Child: Yeah, so the Arm China customers run through the same process that all the rest of our customers run through, and that is, there's royalty audits that are conducted after the fact. And again, that's no different than any other of our customers globally. And occasionally, there's findings and we work through those findings and then get recoveries or adjustments, and that process works well.

Second, we also have audits from independent auditors. So Deloitte Touche does independent audits of Arm China, and they go through the same set of audit requirements that the rest of Arm and SoftBank goes through for that matter. So all of that work is done in parallel, and I would say the integrity information and the responsiveness and all that is really the same for Arm China as it is for all the other regions and parts of the organization that we work with.

Matt Ramsay: Thank you, guys. Really helpful.

Rene Haas: Thank you.

Operator: One moment for our next question. Our last question comes from Sara Russo with Bernstein. Your line is open.

Sara Russo (Bernstein): Great. Thanks for taking my question. Hello, Rene, Jason. So it's been about three years since you launched Flexible Access, more than three years, and about three years since Total Access was launched, and the letter gave some helpful details around slight increase in ACV. Just wondering, for renewing customers, are you seeing any trend on increasing IP adoption because they're into an all-you-can-eat-to-subscribe mode? And for those renewing customers then after they've been in Total Access for a while, are you seeing any increasing spend from those Total Access customers?

Rene Haas: Yeah. So Total Access, as you recall, about three years ago that we rolled it out, and we've seen a few things happen as the program has launched. One is customers that were initial adopters of it, when they've gone up to the next cycle they've actually taken a larger consumption, either more tape-outs and/or more IP. Secondly, the program has worked extremely well from the standpoint of it reinforced what we kind of believed going in, and that is the churn rate for our large partners is pretty small, if not zero. The mean is around 16 years. The median is 19 years of just the longevity of the relationships that these partners

have with Arm. It's allowed our FAEs to be much more involved and engaged in terms of pull-through of other IP.

And then as mentioned earlier, particularly with everything going on with AI and such, it's really, really moved the dial. So ATA has been everything we hoped it would be, and probably a bit more. It's removed a lot of churn in terms of sales cycle, and for the partners, it's very, very easy for them because they do subscriptions around EDA tools. They've rather – they know they're going to spend money with Arm. As mentioned, many of these customers have been with Arm 10 years, 15 years, 20 years. So, it's been a pretty natural evolution on that.

So the program, I would say, has exceeded expectations. I would hope over time that we would get the vast majority of all our partners on this. I think we will because a) there is very little churn to our business, and b) it puts all the resources in the right place in terms of having people accelerate the tapeout of chips.

Sara Russo: That's great. And maybe just a quick follow-up. As part of that program, because you're working with customers slightly differently, are you getting more visibility into customer design program such that it gives you more confidence on forecasting royalties, and what you expect to see from a royalties perspective than maybe what you got in the more traditional licensing models?

Rene Haas: Yeah, I think one of the things that was a byproduct of that, and I would say combination of industry trends/Total Access, is compute subsystems because once we started to get involved with partners more deeply, we started to understand exactly what their tapeout schedules were, exactly what they were trying to use from a process standpoint, what libraries they were using. We were suddenly in a completely different domain relative to how we were interacting with partners in terms of schedules. So what it's done for us is I think it's accelerated subsystem engagements, and at the same time, our understanding and visibility of customer programs is at a level that we've never had before.

Sara Russo: That's great. Thank you very much.

Rene Haas: Sure.

Operator: Thank you. That concludes the question and answer session. At this time, I would like to turn the call back to Rene Haas, CEO, for closing remarks.

Rene Haas: Okay. Thank you, Abigail, and on behalf of myself, Jason, and Ian, I'd like to thank everyone for their excellent, thoughtful questions. This was the first time around for us in terms of as a trio doing this. We'll get better each time, but thankfully, we had a very good quarter to come off on and talk about, which made the job a bit easier.

As mentioned before, we're very, very excited about the prospects going forward, very, very excited about the opportunity, and look forward to continuing to engage with you all. Thank you so much.

Operator: Thank you for your participation in today's conference. This does conclude the program. You may now disconnect.

[END OF TRANSCRIPT]