

29-Jul-2026

Arm Holdings Plc (ARM)

Q1 2027 Earnings Call

CORPORATE PARTICIPANTS

Ian Thornton

Vice President-Investor Relations, Arm Holdings Plc

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

Jason E. Child

Executive Vice President & Chief Financial Officer, Arm Holdings Plc

OTHER PARTICIPANTS

Joe Quatrochi

Analyst, Wells Fargo Securities LLC

Sebastien Naji

Analyst, William Blair & Co. LLC

Gary Mobley

Analyst, StoneX Group, Inc. (Investment Management)

Tom O'Malley

Analyst, Barclays Capital, Inc.

Vivek Arya

Analyst, BofA Securities, Inc.

Vijay Raghavan Rakesh

Analyst, Mizuho Securities USA LLC

Charles Shi

Analyst, Needham & Co. LLC

Timm Schulze-Melander

Analyst, Rothschild & Co Redburn

MANAGEMENT DISCUSSION SECTION

Operator: Good day and thank you for standing by. Welcome to the Arm First Quarter Fiscal Year 2027 Webcast and Conference Call. At this time, all participants are in a listen-only mode. After the speakers' presentation, there will be a question-and-answer session. [Operator Instructions] Please be advised that today's conference is being recorded.

I would now like to turn the call over to your first speaker today, Ian Thornton, Vice President of Investor Relations. Please go ahead.

Ian Thornton

Vice President-Investor Relations, Arm Holdings Plc

Thank you and welcome to our first quarter fiscal 2027 earnings call. On the call are Rene Haas, Arm's Chief Executive Officer; and Jason Child, Arm's Chief Financial Officer. Today's call contains forward-looking information about the company and its financial results. While these statements represent our best current judgment, our business is subjected to many risks and uncertainties that could cause actual results to differ materially. Important risk factors that may affect our business and future financial results are described in our Annual Report on Form 20-F filed with the SEC. Arm assumes no obligation to update any forward-looking statements.

We will also refer to non-GAAP financial measures. Reconciliations of these non-GAAP financial measures to the most directly comparable GAAP measures can be found in our shareholder letter, as can a discussion of certain

projected non-GAAP financial measures that we are not able to reconcile without unreasonable effort and supplemental financial information. Our earnings material are available at investors.arm.com.

And with that, I'll turn the call over to Rene.

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

Thank you, Ian, and welcome, everyone. Arm delivered a record first quarter and a strong start to fiscal 2027. Our results reflect growing demand for the Arm Compute platform as AI expands across cloud infrastructure, edge devices and physical world. Revenue reached \$1.29 billion, up 22% year-over-year, driven by record first quarter licensing and royalty revenue. Royalty revenue grew 22% to \$715 million, licensing revenue grew 23% to \$574 million, and non-GAAP EPS increased 29% to \$0.45, above the high end of our guidance.

AI is changing where and how compute happens. We're seeing that in the data center, where the transition to Arm continues to accelerate. And we're seeing it beyond the data center as AI expand into PCs, smartphones and physical AI applications. Across each of these markets, customers are increasingly standardizing on the Arm Compute platform.

These trends are the driving demand for the Arm AGI CPU. We introduced the Arm AGI CPU in March to give customers another way to deploy the ARM Compute platform. Since that time, we've made significant progress. Initial product has now been delivered to multiple customers, and we have secured the manufacturing capacity needed to support the \$1 billion opportunity we outlined last quarter across fiscal 2027 and fiscal 2028. Demand now exceeds \$2 billion as we continue to add new customers, including multiple customers in the US and China, while the overall value of our pipeline has continued to strengthen. We're also working closely with our manufacturing and supply chain partners to expand capacity. Our confidence in achieving upside to our \$1 billion opportunity for the Arm AGI CPU business has increased in the past 90 days.

This momentum is part of a much broader shift already taking place across our Neoverse business. Data center royalties more than doubled year-over-year once again as adoption of Arm Neoverse continues to expand. This pace at which Arm is becoming the CPU foundation for AI infrastructure is accelerating. Arm Neoverse shipments have now surpassed 1.5 billion cores, with most recent 500 million shipping in just the last nine months, where the first 1 billion took six years.

The world's leading AI infrastructure providers continue to validate that trend. NVIDIA has now brought Vera into production. Built on Arm, Vera delivers up to 50% higher CPU performance and two times greater energy efficiency than comparable x86 systems, and will serve as a CPU foundation for NVIDIA's next-generation AI infrastructure. Google has stated that its Arm-based Axion CPU is a core component of its AI infrastructure strategy and the host CPU for its latest TPU AI systems. AWS announced plans to deploy tens of millions of Graviton5 cores to power agentic AI workloads. Microsoft expanded Azure Cobalt 200 virtual machines built on Arm Neoverse CSS. Additionally, Qualcomm has also announced plans to enter the AI data center CPU market with its Arm-based Dragonfly C1000.

Each of these companies is approaching AI infrastructure differently but they're all moving the same direction. Arm-based CPUs are becoming central to next-generation AI infrastructure. IDC reported that spending on Arm-based accelerated server platforms has nearly doubled in the past two quarters and has now surpassed x86 platforms. We're witnessing both the rapid expansion of AI infrastructure and Arm's growing role within it.

Our opportunity extends well beyond the data center. As AI moves into production, customers are increasingly focused on the economics of deploying AI at scale. Whether AI runs across the cloud, the edge, or ultimately, the physical world, efficient compute is becoming just as important as model capability, bringing AI closer to where data is created, improves performance, reduces latency, enhances privacy and lowers infrastructure costs. This is where Arm has always differentiated itself.

The result is a new generation of computing devices spanning into two distinct categories: efficient AI PCs designed for mobility and more powerful agentic systems capable of running sophisticated models locally. NVIDIA introduced RTX Spark, the first agentic PC built on Arm's Compute Subsystems, enabling sophisticated AI agents and larger AI models which can run locally. For on-the-go AI PCs, those same OEMs have continued expanding the Windows on Arm ecosystem with new Snapdragon-powered AI PCs, while Google's continued investment in AI-enabled Chromebooks is broadening access to on-device AI. As AI growth drives the transition to next generation of personal computing, Arm's opportunity continues to grow across an expanding range of AI-enabled devices.

The same economics extend into the physical world. Vehicles, robots, industrial systems and autonomous machines increasingly rely on efficient, secure and real-time Arm-based compute to sense, reason and act autonomously. NVIDIA recently expanded its physical AI platform with Cosmos 3 and the Isaac GR00T humanoid robotics platform powered by Jetson Thor, which combines an Arm-based CPU with NVIDIA Blackwell GPU.

Arm's software ecosystem continues to expand and now supports more than 22 million developers worldwide. During the quarter, Arm introduced Performix with support from Microsoft, MongoDB, Redis, and SAP, helping developers and agents analyze and optimize workloads running on Arm-based infrastructure.

We also expanded our AI developer tools, including the Arm MCP Server, which has surpassed 10,000 Docker downloads and integrates Arm's expertise into leading AI developer environments. From cloud infrastructure to PCs and physical AI, developers can build on the same Arm architecture and software ecosystem. Customers can deploy Arm through IP, compute subsystems or silicon, depending on what best fits their business. But in every case, they're building on the same Arm Compute platform, software ecosystem and developer community.

As AI becomes part of every cloud, every device and every sector, the industry is increasingly converging on a common – compute platform. We believe that convergence will define the next decade of computing. AI is changing where and how compute happens, and Arm is at the center of it.

With that, I'll turn it over to Jason.

Jason E. Child

Executive Vice President & Chief Financial Officer, Arm Holdings Plc

Thank you, Rene. We have started fiscal year 2027 with another strong quarter, delivering the highest first quarter revenue in our history. Total revenue grew 22% year-on-year to \$1.29 billion. Royalty revenue grew 22% year-on-year to \$715 million, also our highest ever figure for Q1. Once again, the largest driver of royalty growth was cloud AI. Data center royalty revenue continues to more than double year-on-year, reflecting our sustained momentum across the market. This is being driven by the continued ramp of Arm-based server chips at all the major hyperscalers, alongside increasing deployments of data center networking chips, particularly DPUs and SmartNICs, where Arm technology is deployed in nearly all leading products.

Edge AI royalty revenue continued to grow despite a soft end market in smartphones. We continued to benefit from higher royalty rates as Armv9 and compute subsystems continue to increase their penetration of the

smartphones, tablets and other consumer electronic devices. These drivers more than offset the decline in smartphone sales due to higher memory prices. Physical AI also made a strong contribution to royalty growth, supported by the continued secular expansion of ADAS and autonomous systems built on Arm technology.

Turning now to licensing, license and other revenue was \$574 million, up 23% year-on-year, also a record for Q1. Growth was driven by strong demand for next-generation architectures and deeper strategic engagements with key customers. This quarter, we signed multiple high-value agreements as existing customers renewed long-term licenses. Some of the world's largest hyperscalers, automotive and robotics companies and handset OEMs all secured access to Arm's future roadmap for their next generation of products. Of the \$574 million of license revenue, our agreement with SoftBank for technology, licensing and design services contributed \$193 million. We expect the quarterly run rate for the rest of the year to be around \$200 million.

As always, licensing revenue varies quarter-to-quarter due to timing and size of high value deals. So, we continue to focus on annualized contract value, or ACV, as a key indicator of the underlying licensing trend. ACV grew 13% year-on-year, maintaining strong momentum. This continues to be above our long-term expectation for license revenue growth.

As Rene mentioned, customer demand for Arm AGI CPU remains very strong, and this quarter, even more customers are wanting to place orders with us. We've secured the manufacturing capacity needed to support the initial \$1 billion opportunity, and we have made progress to secure additional supply, as well as optimizing our customer mix and commercial terms. Our confidence in achieving more than \$1 billion has increased in the past 90 days. We will provide an update at our Q3 results, which is when we will have better visibility of Q4 2027 and fiscal 2028.

Turning to operating expenses and profits, non-GAAP operating expense was \$733 million, up 18% year-on-year due to the ongoing R&D investment. This was about \$27 million below our guidance due to timing, as our spending plan for the year remains largely unchanged. We are expanding our engineering teams to support increasing customer demand while advancing the technologies that will underpin future growth, including next-generation architectures, compute subsystems and the Arm AGI CPU product family.

Non-GAAP operating income was \$531 million, resulting in a non-GAAP operating margin of about 41%, up 200 basis points year-on-year. Non-GAAP EPS was \$0.45, driven by both higher revenue and slightly lower OpEx than expected. More importantly, even at these elevated levels of investment, we generated free cash flow of \$665 million for the quarter and \$1.4 billion over the trailing 12 months, giving us the flexibility to continue to invest for long-term growth.

Turning now to guidance, for Q2, we expect revenue of \$1.38 billion plus or minus \$50 million. At the midpoint, this represents revenue growth of about 22% year-on-year. We expect license and other revenue to be up about 30% year-on-year and royalty revenue to be up in the low teens year-on-year. We expect our non-GAAP operating expense to be approximately \$780 million and our non-GAAP EPS to be \$0.47 plus or minus \$0.04.

Looking ahead, we continue to see strong customer demand across our business. Combined with the expansion of our technology portfolio and deepening strategic customer relationships, this gives us confidence in our long-term growth outlook.

With that, I'll turn the call back to the operator for Q&A.

QUESTION AND ANSWER SECTION

Operator: Thank you. [Operator Instructions] And our first question today comes from the line of Joe Quatrochi from Wells Fargo. Please go ahead.

Joe Quatrochi

Analyst, Wells Fargo Securities LLC

Q

Yeah, thanks for taking the questions. Maybe just first, can you give us a little bit more detail on just the confidence that's increased in terms of the upside to the \$1 billion for AGI revenue? Is there any clarity you can provide us on is it better wafer supply, customers getting access to memory? Any help there?

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

A

Yeah, I'll take that question. I think in general, it's all of the above. If we – if I look back to the call we had 90 days ago, we had committed to supply of \$1 billion. We talked about demand pipeline of \$2 billion. And we were working to secure supply for that delta between \$1 billion and \$2 billion. 90 days later, the demand picture has even gotten better, as Jason mentioned, north of \$2 billion. But our ability to secure that supply for the north of \$1 billion, our confidence has increased in the last 90 days. And I would say it's across all of those areas. Making a chip is complex relative to supply chain. You have wafers, you have substrates, you have test capacity, you got memory and all of those areas. Our confidence level in being able to secure the supply necessary has gotten better, and that's the commentary we're staying at this time.

Joe Quatrochi

Analyst, Wells Fargo Securities LLC

Q

And maybe just as a follow-up to that, I mean, on top of the extra above \$1 billion of AGI revenue that you now think you can do, how do we think about the gross margin structure relative to kind of what you were thinking about for the original \$1 billion?

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

A

Yeah, no change at this time. We said that, I think last quarter, we said it was probably going to be somewhere in the high-30% range, maybe low-40s for the first generation for, really, in Q4 of this year as well as for next year. As – I'd say over the next couple of years, we do expect to get to 50%, but that's going to basically just entail us actually bringing more of some of the work that maybe an ASIC itself helped us with the past and bring some of that work in-house. So, that'll probably take a couple of years. Those are our initial expectations. So, no real change from last quarter.

Obviously, with some of the price increases and things that have happened with the kind of working through how to digest those or what needs to be kind of – what needs to affect our pricing. So, those are some of the things that we're going to work through. And as I said in my prepared remarks, we'll provide a much more detailed update for both the revenue and margin profile before we go into Q4, so at the end of Q3.

Joe Quatrochi

Analyst, Wells Fargo Securities LLC

Q

Thank you.

Jason E. Child

Executive Vice President & Chief Financial Officer, Arm Holdings Plc

Thank you.

A

Operator: Thank you. We will now take the next question, and the question comes from the line of Sebastien Naji from William Blair. Please go ahead.

Sebastien Naji

Analyst, William Blair & Co. LLC

Thank you and good afternoon. My question is on the smartphone market. Because of the much higher memory handset costs, handset OEMs are absorbing meaningfully more BOM inflation. They're starting to raise prices. I think given that smartphone royalties remain a significant portion of your overall royalty base, can you maybe just talk a little bit about what you're seeing in terms of unit demand and mix and how this impacts your royalty revenue outlook for the fiscal year?

Q

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

Yeah. So, I'll take the first part of the question and I'll let Jason then talk about some of the details around the numbers. Generally speaking, we have been somewhat isolated from the negative growth of the smartphone market because of the fact that we have moved the vast majority of customers to v9, in some cases, CSS, and in some cases, the second version of CSS. And what that's all contributed to has been an increase in royalty growth. So whereas the smartphone market has been going – project to be down by some level of double digits, we are projecting a double-digit growth in royalty in smartphone market.

A

Vis-à-vis exactly how that translates to the forward guidance and forecast, I'll let Jason go in with some more color.

Jason E. Child

Executive Vice President & Chief Financial Officer, Arm Holdings Plc

Yeah, I think there's obviously others that are reporting this week and key partners that we'll learn more from. But when we look at our forecast, we mostly try to look at the industry projections across all – the entire industry. IDC is, of course, one of things that we look at. And then of course, we update for partner mix where we have slightly different royalties based on whether it's CSS or v9 or v8, et cetera. So when we add all that together, we have seen some incremental slowdown versus what was expected at the beginning of the year. And I think the new piece was initially the expectation was just going to affect the lower end of the market, and now, we are seeing all parts of the market, even some upper and mid-tier being affected. And so, that's going to be an incremental for us.

A

And so as a result, I would say I think going into the year, we probably thought we were – I think we said last quarter that we're expecting somewhere around 20% year-on-year for the next few years in royalties, including this year. Right now, if I had to guess, that's probably somewhere closer to the high teens right now, but hard to say until we learn more about the next couple of quarters. I'd say right now, the – we do think royalties will come down a bit in this next quarter. We guided to kind of the low to mid-teens for Q2 and we'll give you a better update next quarter.

But I think the one thing I would make sure you don't lose sight of is the good news is the overperformance that we've been seeing at cloud AI continues to help offset that. And so while there is weakness on the smartphone side, the overperformance on the cloud AI business continues to accelerate. And that's the piece that gives us confidence in full year and next year's numbers. But again, we'll give you more updates over the next couple quarters.

Sebastien Naji

Analyst, William Blair & Co. LLC

Great. Okay. That's very helpful. I appreciate all the color.

Q

Operator: Thank you. [Operator Instructions] We will now go to the next question, and the next question today comes from the line of Gary Mobley from StoneX Group. Please go ahead.

Gary Mobley

Analyst, StoneX Group, Inc. (Investment Management)

Hi, everybody. Thanks for taking my question. You highlighted maybe a little bit of upside to the demand profile for the next two years, \$2 billion plus, I think, to paraphrase correctly. But what about the demand pipeline for the out years, fiscal year 2029 through 2031? And I assume the \$15 billion in projected AGI revenue expected for 2031 contemplates all the supply chain-related headwinds, that's a supply served revenue. Correct?

Q

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

Yeah. I'll let Jason comment further in terms of any forward-looking comments he wants to make. But I'll say, what has changed since March is that, to your point, clearly, supply is an issue across a number of different factors. The outlook that we gave at that time was based upon our view of the supply picture.

A

On the flip side, we did say at that event that the CPU TAM, we thought was about \$100 billion in those outer years. And going back, you may recall that prior to that, most folks were talking about a number of around [ph] \$50 billion to \$60 billion (00:35:34). So when we talked about the \$100 billion number in March, there were some surprise around it and there was a lot of back and forth in terms of justifying that number. Since that time, a number of my peers have talked about numbers quite a bit bigger than that, in some cases, up to \$200 billion. I think all of that is really being driven by the increase of overall compute capacity, which will be inference-based, and all of that inference-based compute is going to largely be running agentic workloads, and the agentic workloads are essentially capacity constrained in terms of throughput by the number of CPUs you have. So as inference demand goes up, which is clearly is, and as agentic demand goes up, as it clearly is, that means an increase in CPU demand. So, we're not changing the numbers at all but there are a lot of indicators that the numbers that we've talked about back in March relative to our view of the TAM may have been conservative.

Jason, do you want to add anything to that?

Jason E. Child

Executive Vice President & Chief Financial Officer, Arm Holdings Plc

Yeah, yeah. I think in general, as Rene said, yeah, the TAM at the back when we did the event, it – when we talked about to \$100 billion, we said \$100 billion plus. Obviously, it's gone up to, I think now, the most recent estimate is \$220 billion sitting on – from a variety of sources. So I think our expectation is if you flow through the same kind of market share that we were expecting at \$100 billion, could it be much higher if, in fact, the market is at \$200 billion plus? Well, certainly could be. The constraint over the next couple of years, I would say, the – for

A

us, FYE 2027 or calendar 2026 or FYE 2028, which is, I guess, calendar 2027, quite seems pretty tight. And so, we're – maybe there's an upside to get to this \$2 billion or \$2 billion plus and we'll provide updates on that over the next couple of quarters.

But to go beyond that, the supply chain is going to loosen or at least expand quite a bit as we get into calendar 2028 and 2029, or for us, FYE 2029 and 2030. Right now, if you look at wafer capacity and memory capacity, you probably know as well as I do, I think there's estimates that the capacity is going up somewhere between 70% to 100%, depending on your assumptions on wafer and some of the different partners that can help there and even more so on the memory side. So our expectations are, there certainly is potential to beat the number that we give externally. Certainly, internally, we do have higher targets. But until we start to see the supply chain capacity come online, we're mostly focused on just trying to get to the \$2 billion number next year and we'll provide updates, of course, later as we learn more.

Operator: Thank you. Our next question today comes from the line of Tom O'Malley from Barclays. Please go ahead.

Tom O'Malley

Analyst, Barclays Capital, Inc.

Hey, guys. Thanks for taking my question. Last week, AMD hosted an Analyst Day and talked about a \$220 billion TAM by the end of the decade. And kind of underneath that, the largest contributor was agentic AI that kind of laid out traditional CPU to head nodes and then agentic applications. So I was curious, when you look at your silicon business longer term, where do you see that fitting in? Is that going to cover all three of those buckets in your early wins and your early pipeline? Is there one way that that's leaning or another? Just would like to get a flavor of what you foresee for that silicon business?

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

Yeah. And I'm sorry, could you could you repeat those three categories, so I make sure I've got the definitions right?

Tom O'Malley

Analyst, Barclays Capital, Inc.

Yeah, traditional server and then you have head nodes and then you have agentic applications.

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

Got it. Yeah. Our Arm AGI CPU is going to play in all three. When we talked about the customers that we had signed up back at the Arm Everywhere event, we had folks like Cerebras and OpenAI, which were largely around head node type applications. We also talked about Meta and Cloudflare, which are around agentic applications, but also general-purpose server. We've also just talked about OCI, Oracle. So, the Arm AGI CPU is a very good fit for all of those, to be quite candid, which is why we are very, very optimistic about the demand. And we did talk about quite a number of customers that day and we were very specific about inside the cloud, the agentic workloads. It's a pretty broad term because the agents are running through the head node but they're also running through the general-purpose racks that sit inside the data center. So, short answer is those three categories that you defined, as described earlier, the Arm AGI CPU is a great fit for all three and we have customers in all three.

Operator: Thank you. Your next question comes from the line of Vivek Arya from Bank of America. Please go ahead.

Vivek Arya

Analyst, BofA Securities, Inc.

Q

Thanks for taking my question. I had a near and a longer-term question on your AGI CPU. Near term, I'm very curious, what is preventing Arm from securing supply for just \$1 billion more? I mean, it's a \$50 billion market, right, and there are a number of foundries who can make it. So, I'm just curious what is preventing Arm from getting that extra \$1 billion a year from now?

And then longer term, Rene, when I look at the three players that have Arm-based CPUs in AI, whether it's NVIDIA or Amazon or Google, they each have their proprietary accelerators also, which are often co-designed with the CPU. So don't you think that restricts Arm's opportunity, given that you have a CPU only offering, or is your intention to add other things to that CPU-only offering over time? Thank you.

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

A

Yeah. So, I'll take both parts of that question, and Jason, you can add on. I'm glad you think \$1 billion is not a big number. There are a lot of folks just trying to secure extra supply up to \$50 million. It is a very, very tight market across everything, whether it's memory, whether it's test equipment, whether it's substrates, whether it's adjacency wafers. It's a very, very, very tight world. I think even my compatriots who are in the CPU business for a living and have their own fabs have not been able to supply the demand. So, we're – we are more optimistic than we were in 90 days ago, which is great. Demand is increased but we feel better about the supply side.

On the proprietary accelerators, if I understood your question correctly, a couple of proof points. So, that is an area that we can certainly play in. First off, at Google, they have used Axion as an interface into their TPUs away from x86. So clearly, there's a space for Arm to exist with custom accelerators.

If your question was around the Arm AGI CPU specifically, there has been some announcements made by NVIDIA about NVLink Fusion. And it's a very interesting platform approach where they're offering a mix-and-match, where you could take a Vera CPU and plug it into a proprietary accelerator, or you could take a CPU made by someone else in the Arm family and connect it into a Rubin accelerator. So, we don't have anything to announce around that today, but there are absolutely paths for Arm to connect to custom accelerators, whether it's through a self-hosted design done internally and/or connecting to something – through something like NVLink Fusion.

Operator: Thank you. Our next question today comes from the line of Vijay Rakesh from Mizuho. Please go ahead.

Vijay Raghavan Rakesh

Analyst, Mizuho Securities USA LLC

Q

Yeah, hey, Rene and Jason. Good to hear you've got some capacity on the agentic AI CPU side. I was wondering, when you look at your agentic AI CPUs, are they mostly in the 128 core? I know core count is important. So, are you looking at – when do you expect the next – the 192 or 256 core CPUs coming out and if any color on the basic side as well?

And for Jason, do you expect to like start breaking out the silicon side of revenues, the agent CPU side? Because that seems to be becoming bigger and bigger, I guess. Thanks.

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

A

Yeah, so thank you for the question. On the product side, nothing to tell you specifically today about the roadmap and where we're going with that. So, the Arm AGI CPU today is a 128-core based design. There are instances in the market using Arm that are better – or actually greater number of cores, and that's Graviton5, which is 192 cores. Certainly, the direction of travel is more cores. And the reason for that is for running agentic workloads, more cores is a better outcome because the software overhead is simpler. You can run virtual machines or virtual jobs on single cores and they're going to be much more efficient in terms of throughput. So while I don't have anything to talk about today in terms of the roadmap and the core count, it's not a stretch to think that the next-generation designs are going to have more and more cores, because that's clearly the direction of travel.

And I'll let Jason comment on the second part of the question.

Jason E. Child

Executive Vice President & Chief Financial Officer, Arm Holdings Plc

A

Yeah, on the revenue breakout, well, first, as soon as we start shipping, which will be at the end of this year, that's the first milestone, and then once it becomes at least 10% of revenue, we'll break it out separately as a third line, so separate from license, separate from royalties [indiscernible] (00:46:01) revenue. So, I would expect that based on our forecast that that should be broken out then in FYE 2028.

Operator: Thank you. Your next question today comes from the line of Charles Shi from Needham & Company. Please go ahead.

Charles Shi

Analyst, Needham & Co. LLC

Q

Hi. Thanks for taking my question. Jason, I think if I hear you correctly, you are talking about maybe royalty revenue growth this year instead of 20%, probably going to be high teens because of all the things happening in the smartphone incremental weakness, et cetera. Anything we should be looking at at the moment, let's say, on the licensing side that can show up with a little bit of upside to offset with the incremental weakness on the royalty side? And maybe on that topic, what's the early view on maybe FYE 2029, which is largely CY 2028, and do you think a 20% royalty is still a good number? And I want to get some thoughts there. Thank you.

Jason E. Child

Executive Vice President & Chief Financial Officer, Arm Holdings Plc

A

Yeah. In terms of guidance, for Q2, we did lower, I think, maybe what we – what our expectations were a couple of quarters ago. We didn't guide to it but I think our expectations were they were going to be a little higher. And so, we're now guiding to [ph] 30% (00:47:31) royalty growth in Q2, but then, we also increased our license growth. So overall, on a combined basis, we're actually ahead of where we had expected to be. I would expect similar trends for the rest of the year, and that is any softness or weakness that we experienced in royalties will be at, I would say, at least as much license revenue to overcome that. So, I don't really expect overall to be any sort of a slowdown.

Now, the thing that could change is we continue to see the cloud AI business overperform, and in particular, we're really seeing strong deployments from some of our partners like Google with our Axion as it's [ph] growing

(00:48:16) with all the new TPU deployments. Certainly, Amazon, and then certainly, with Vera at NVIDIA, Amazon with Graviton, and then we have other hyperscalers who are also exceeding their plans, so – and increasing their trajectory and their velocity. So, there is certainly some possible upsides on royalties throughout the year. But if I had to say right now, I would not expect there to be weakness across the overall business. It just might be some weakness in the smartphone side of the business. And we'll just have to wait and see how the cloud AI business is able to overcome all of it or just some of it. We'll have to update you as we go later into the year.

And then in terms of 2029 – I'm sorry, 2028 and 2029, yeah, I think the 20%-plus royalty growth that we provided, I think that's very much still intact. The things that are happening right now, the only real weakness in royalties is really on the smartphone side, in particular, to some of the memory issues there. And while they're certainly having a tough year this year, I think most of the partners that we talked to seem to think that there's going to be stronger recovery in the back half of the year. And so, I would expect that the year-on-year growth rates for next year should be kind of back in line as we start to lap some of these challenges. And then, of course, as we continue to see the AI deployments continue to grow at the level they have been, that we should be able to, again, be in that 20% range next year and beyond.

Operator: Thank you. Our next question today comes from the line of Krish Sankar from TD Cowen. Please go ahead.

Q

Hi. Thanks for taking my question. This is [ph] Steven (00:50:14) calling on behalf of Krish. Rene, I had a question on low-cost AI models. Just with the recent cycle of headlines around low-cost AI, whether from China or elsewhere, do you see the proliferation of low-cost AI being a net positive or negative for your AGI CPU franchise given your market positioning? And also, if you can also walk us through how low-cost AI might be a [indiscernible] (00:50:38) for your IP business, like for both data center and the edge side of things? Thank you.

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

A

Yeah, thank you for the question. I'm going to presume that when you say low-cost AI, you're referring to open source models, open weights, et cetera. If the world moves to something that's more open source based, AKA open weights, for us, I think it's somewhere between a net neutral to a net positive. And I start with net neutral from the perspective of we're somewhat agnostic relative to whether it's a close frontier model or an open source model. Those models all require CPUs, they all require the IP that we deliver for people to build custom chips. So whether it's an opportunity for Arm AGI CPU or the IP or compute subsystems, we're going to be required independent of what those models are. Just frankly, just given the nature of the workloads that need to get run, they're kind of independent of the hardware underneath it. You need CPUs to run agents, whether that's open source or closed source.

There's a very interesting argument about to be made relative to the potential of those open source models being much more differentiated, smaller, more efficient that could run in different edge footprints. If that were to happen, it'd be a great thing for Arm, just given the footprint that we have in smaller edge devices. Right now, the most sophisticated models that are the close models all are literally 100% cloud-based. The RTX Spark that was announced is a very interesting product where you could potentially run some smaller models locally.

I think open source has traditionally shown that whether it's in this area or any other software domain, widely [ph] brightens (00:52:44) the choice relative to how you run different things. In other words, open source leads to higher innovation and differentiation. So, I think it's either neutral to a positive depending on where it goes at the edges, but probably most important thing to remember about Arm, and it applies here as well, is we're pretty agnostic to the top layer. All of that will need to run on our CPUs and our IP.

Operator: Thank you. We will now take our final question for today, and our final question comes from the line of Timm Schulze-Melander from Rothschild & Co. Please go ahead.

Timm Schulze-Melander

Analyst, Rothschild & Co Redburn

Q

Yeah, great. Thanks for taking my questions. I had two. One was a more kind of operational question maybe for Jason. Just talking about the timing and the sequential quarter-to-quarter cadence of R&D, you talked about how it come in below your expectations. And I just wondered maybe if you could talk about how that kind of comes to pass. Is it kind of a milestone-based thing and maybe they're not met and so the cost slips a quarter or two? Is it a capitalization issue? Just how do – how does – what are the things that dictate the timing of R&D cost in the P&L? And then I had a quick follow-up on the CPU.

Jason E. Child

Executive Vice President & Chief Financial Officer, Arm Holdings Plc

A

Sure. So with OpEx, we did come in a little bit lower. When it comes to OpEx, you have to break out R&D,. You have to break it down into really two buckets. So you have the people or developer costs, and then you have the tools that they're using, which is typically emulation, cloud spend, et cetera. And so in this last quarter, we did come in a little bit below and that would be more, not on the people side because that was pretty much right in line, I would say it was really more on kind of the tools that were utilized and necessary versus what our forecast was. So I think for the most part, you should think about kind of banking that savings that we delivered in Q1. And then going forward, our estimates, I think, are pretty similar, maybe a little bit lower. So overall, again, you can [ph] bank (00:54:54) the savings from Q1 and then assume that for us, they are just at a high level. It's going to grow by kind of mid-single digit percentage quarter-on-quarter, the overall OpEx. And the things that can move, it's less around the total number of engineers because that number is one that we're – we can forecast and we're pretty consistent with. It's really just about what is the tool utilization that's necessary. And so certainly, with things like emulation and cloud spend, that's going to flex a little bit based on what are we learning on the latest developments? And do we need to increase some of our testing or are we kind of done. And so as a result, we're getting better and better at forecasting this. But it does move around a little bit because we've really only been doing this now for the last year or two. And so, I would expect our variance probably to get smaller going forward.

Timm Schulze-Melander

Analyst, Rothschild & Co Redburn

Q

Got it, got it. And then just to circle back on the AGI CPU conversion of customer engagement and interest, and as you've described, being able to kind of bring that across the Rubicon and book it as a firm order, just wanted to understand and just kind of circle back, just confirm what we're talking about here, so we're talking about maybe early calendar 2027 as you come into the December quarter report that you should have more visibility. And I guess one of the things I'm just thinking about is, is there a concern or is it possible or probable that you get into a kind of a golden screw type situation that you've got 99% covered, but because there's tightness, you're missing one small component that could kind of stymie that or is it really not quite that vulnerable to one single sort of small component to be missing? Thank you.

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

A

If I [indiscernible] (00:56:57) answer your question is as we get closer to fulfilling the demand for the product, is there still some risk at the 11th hour that one critical component that we haven't secured could tip the whole thing over. No, there's not a concern of that. We are very confident in terms of the partners we're working with. We have a very good understanding of their capabilities and their upside capabilities. So, there isn't one single point of failure that, to your point, of the 1% that could cause an issue.

Operator: Thank you. This concludes today's Q&A. I will now hand back to Rene for closing remarks.

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

Thank you. The quarter for us was – is really a tale of two stories. One is the core business itself. The IP business continues to perform extremely well, a record first quarter, and we are guiding ahead of expectations, as Jason had mentioned, and our results were beyond expectations. The growth is really being driven by the data center. Arm continues to increase its market share. And we have many key customers now shipping Neoverse cores, which as I mentioned before, have now surpassed 1.5 billion.

But we also now have our new business, the Arm AGI CPU business, which we just announced back in March. The confidence level we have in the success of that product only continues to increase. And as we talked about, the demand is greater than what we shared 90 days ago, and at the same time, our confidence to fulfill that demand has also increased. So all in all, very proud of the results that we've achieved, a great quarter for the company, and thank you all for your interest and questions.

Operator: Thank you. This concludes today's conference call. Thank you for participating. You may now disconnect.

Disclaimer

The information herein is based on sources we believe to be reliable but is not guaranteed by us and does not purport to be a complete or error-free statement or summary of the available data. As such, we do not warrant, endorse or guarantee the completeness, accuracy, integrity, or timeliness of the information. You must evaluate, and bear all risks associated with, the use of any information provided hereunder, including any reliance on the accuracy, completeness, safety or usefulness of such information. This information is not intended to be used as the primary basis of investment decisions. It should not be construed as advice designed to meet the particular investment needs of any investor. This report is published solely for information purposes, and is not to be construed as financial or other advice or as an offer to sell or the solicitation of an offer to buy any security in any state where such an offer or solicitation would be illegal. Any information expressed herein on this date is subject to change without notice. Any opinions or assertions contained in this information do not represent the opinions or beliefs of FactSet CallStreet, LLC. FactSet CallStreet, LLC, or one or more of its employees, including the writer of this report, may have a position in any of the securities discussed herein.

THE INFORMATION PROVIDED TO YOU HEREUNDER IS PROVIDED "AS IS," AND TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, FactSet CallStreet, LLC AND ITS LICENSORS, BUSINESS ASSOCIATES AND SUPPLIERS DISCLAIM ALL WARRANTIES WITH RESPECT TO THE SAME, EXPRESS, IMPLIED AND STATUTORY, INCLUDING WITHOUT LIMITATION ANY IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, ACCURACY, COMPLETENESS, AND NON-INFRINGEMENT. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, NEITHER FACTSET CALLSTREET, LLC NOR ITS OFFICERS, MEMBERS, DIRECTORS, PARTNERS, AFFILIATES, BUSINESS ASSOCIATES, LICENSORS OR SUPPLIERS WILL BE LIABLE FOR ANY INDIRECT, INCIDENTAL, SPECIAL, CONSEQUENTIAL OR PUNITIVE DAMAGES, INCLUDING WITHOUT LIMITATION DAMAGES FOR LOST PROFITS OR REVENUES, GOODWILL, WORK STOPPAGE, SECURITY BREACHES, VIRUSES, COMPUTER FAILURE OR MALFUNCTION, USE, DATA OR OTHER INTANGIBLE LOSSES OR COMMERCIAL DAMAGES, EVEN IF ANY OF SUCH PARTIES IS ADVISED OF THE POSSIBILITY OF SUCH LOSSES, ARISING UNDER OR IN CONNECTION WITH THE INFORMATION PROVIDED HEREIN OR ANY OTHER SUBJECT MATTER HEREOF.

The contents and appearance of this report are Copyrighted FactSet CallStreet, LLC 2026 CallStreet and FactSet CallStreet, LLC are trademarks and service marks of FactSet CallStreet, LLC. All other trademarks mentioned are trademarks of their respective companies. All rights reserved.