

06-May-2026

Arm Holdings Plc (ARM)

Q4 2026 Earnings Call

CORPORATE PARTICIPANTS

Jeffrey Thomas Kvaal

Vice President, Head of 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

Andrew Michael Gardiner

Analyst, Citigroup Global Markets Ltd.

Joe Quatrochi

Analyst, Wells Fargo Securities LLC

Vivek Arya

Analyst, BofA Securities, Inc.

Timm Schulze-Melander

Analyst, Rothschild & Co Global Markets Solutions Ltd.

Krish Sankar

Analyst, TD Cowen

Sebastien Naji

Analyst, William Blair & Co. LLC

Vijay Raghavan Rakesh

Analyst, Mizuho Securities USA LLC

Harlan Sur

Analyst, JPMorgan Securities LLC

Lee Simpson

Analyst, Morgan Stanley & Co. International Plc

MANAGEMENT DISCUSSION SECTION

Operator: Good day and thank you for standing by. Welcome to the Arm Fourth Quarter Fiscal Year 2026 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 hand the conference over to your first speaker today, Jeff Kvaal, Head of Investor Relations. Please go ahead.

Jeffrey Thomas Kvaal

Vice President, Head of Investor Relations, Arm Holdings Plc

Thank you, Sharon, and welcome everyone to our fourth quarter fiscal 2026 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 subject 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 during the call. 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 materials are available at investors.arm.com.

With that, I'll turn the call to Rene.

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

Thank you, Jeff, and welcome, everyone. Arm delivered a record quarter and record fiscal year. Revenue this quarter was \$1.49 billion, up 20%, our highest quarterly revenue quarter ever and above the midpoint of our guidance.

Licensing revenue grew 29% year-over-year to \$819 million, driven by strong demand for the Arm platform. Royalty revenue grew 11% to \$671 million, with growth across edge AI, physical AI and cloud AI, where our data center royalty has more than doubled year-over-year. That drove record non-GAAP EPS of \$0.60 even while we continued to increase investment in R&D.

For the full year, revenue reached a record \$4.92 billion, up 23% year-on-year. Royalty revenue was up \$2.61 billion, up 21% and licensing revenue was \$2.31 billion, up 25%. Non-GAAP EPS was also a record at \$1.77. Fiscal 2026 was our third consecutive year since going public of more than 20% revenue growth, demonstrating the strength of our business and the increasing relevance of Arm in the highest growth areas of compute.

This quarter was driven by key highlights, including the expansion of Arm's product strategy and our continued momentum in cloud AI. As AI is moving from human-based queries to continuous agent-driven workloads, this shift is expanding the role of the CPU. These agentic workloads require CPUs to coordinate tasks, move data, manage memory, enforce security, and orchestrate work around accelerators.

As agentic AI scales, data centers will require more than four times today's CPU capacity, creating a datacenter CPU market opportunity of more than \$100 billion by 2030. The Arm AGI CPU, which we launched at our Arm Everywhere event last quarter and is purpose built for agentic AI addresses this need directly. Our first production silicon product for the data center will deliver more than two times the performance per rack compared with x86 platforms, with the potential to reduce AI datacenter capital expenditure by up to \$10 billion per gigawatt.

Meta is our lead partner and co-developer and is working with us on a multi-generation roadmap to support personal superintelligence for more than 3 billion users. The Arm AGI CPU expands how customers can work with Arm. Customers can now deploy Arm Compute through IP, compute subsystems or silicon. One compute platform, one software ecosystem. That is unique to Arm. IP and CSS remain the foundation of our royalty growth. Silicon extends the Arm platform and gives customers another way to build AI infrastructure.

Ecosystem support has been significant. More than 50 leading companies are supporting the expansion of the Arm Compute platform into silicon, including the very biggest names in the industry. Customer response to the Arm AGI CPU has been very strong. We now have more than \$2 billion of customer demand across fiscal 2027 and fiscal 2028. This is more than double what we stated at launch. We are on track towards our forecast of \$15 billion in this business, as stated at our Arm Everywhere event, and soon the data center will be Arm's largest business. The direction is clear: customers want Arm at the center of the AI data center.

Customers need Arm where agentic applications run and they need Arm where accelerators scale. For example, SAP will move their core database and business application workloads to Arm, starting with AWS Graviton and expanding to the Arm AGI CPU. This represents a significant strategic shift. Cloudflare will deploy Arm across its global network to support traffic management, security and AI inference closer to users. We have also secured design wins with key network infrastructure providers, including F5 and SK Telecom.

AI infrastructure needs CPUs and accelerators working together efficiently at scale. NVIDIA, Amazon and Google are already using Arm-based CPUs as head nodes along the accelerator-based systems. Cerebrus, OpenAI, Rebellions and Positron are doing the same with the Arm AGI CPU. This momentum builds on our existing scale in the cloud.

That scale is increasingly driven by Arm Neoverse CSS and Arm-based Compute, which now represents about 50% market share with top hyperscalers. Recent announcements from key customers show that AI infrastructure is being built around Arm-based custom silicon.

At Google Cloud Next, Google announced TPU8t for training and TPU8i for inference; in both cases replacing x86 host processors with custom Arm Axion CPUs. The increased performance at 50% less power enables an 80% improvement over the previous x86 solution. AWS continues to scale its custom silicon strategy with Arm-based Graviton alongside Trainium and Nitro. While Microsoft is advancing its Arm-based strategy with Cobalt, designed to deliver high performance and energy efficient compute for Azure workloads.

And at NVIDIA GTC, NVIDIA announced Vera, the next generation Arm-based CPU built for agentic AI and building a standalone rack, integrating 256 Vera CPUs. Across the largest AI platforms, Arm-based CPU is becoming central to performance, efficiency and cloud economics.

Our opportunity is not to stop the data center. AI is moving to every device and every physical system. Phones, PCs, vehicles, factories, robots, cameras, sensors and connected devices all need efficient, secure compute with software that scales. These AI workloads will all run on Arm.

With over 350 billion chips shipped and over 22 million developers, the Arm Compute platform is the most comprehensive in history, and we are positioned to bring AI from cloud infrastructure to the edge and to the physical world through a common compute platform and ecosystem.

We enter fiscal 2027 with record results, strong customer demand and a larger opportunity ahead of us. Our strategy is clear: grow royalties through IP and CSS and add silicon as a new growth vector and scale the Arm platform across the next generation of AI workloads. With that foundation in place, our focus is execution and continuing to build the future of AI on Arm.

And with that, I will hand it over to Jason.

Jason E. Child

Executive Vice President & Chief Financial Officer, Arm Holdings Plc

Thank you, Rene. We have delivered another record quarter. Total revenue grew 20% year-on-year to \$1.49 billion, which is nearly \$250 million higher than our previous record. Our strong Q4 revenue also capped a record year. Revenue growth of 23% exceeded 20% for the third straight year and lifted fiscal 2026 revenue to \$4.9 billion.

Our revenue strength translated into record EPS for both Q4 and fiscal 2026. Royalty revenue grew 11% year-on-year to \$671 million, our highest ever figure for Q4 royalty revenue. The biggest contribution to royalty revenue growth was from cloud AI. Data center royalty revenue continues to more than double year-on-year, and we see no break in this momentum. This is primarily driven by the accelerating ramp of Arm-based server chips by all major hyperscalers, as well as increased deployments of data center networking chips, particularly DPUs and SmartNICs, where Arm has close to 100% market share.

In edge AI, our smartphone revenues continued to deliver growth despite end market weakness. This is driven by our higher royalty rates from increasing penetration of Armv9 and compute subsystems into higher end smartphones. Physical AI has also contributed to our strong royalty performance, driven by secular growth of ADAS and autonomous systems based on Arm technologies.

Turning now to licensing. License and other revenue was \$819 million, up 29% year-on-year. Growth was driven by strong demand for next generation architectures and deeper strategic engagements with key customers.

For example, we signed a long-term strategic partnership with the Indonesian government to strengthen Indonesia's capability in the development of AI technology. We also signed two next-generation CSS licenses, one to be used to develop chips for smartphones and the other for data center networking chips. These agreements reflect the continued investment by our customers in next generation Arm technology.

Of the \$819 million of license revenue, our agreement with SoftBank for technology, licensing and design services contributed \$200 million, flat with the prior quarter. As always, licensing revenue varies quarter-to-quarter due to timing and size of high value deals. So as we continue to focus on annualized contract value or ACV as it's a key indicator of the underlying licensing trend.

ACV grew 22% year-on-year, maintaining strong momentum. This continues to be above our long-term expectations for license revenue growth. As Rene noted – or as Rene mentioned, customer demand for the Arm AGI CPU was very strong. We now have line of sight to more than \$2 billion of demand across fiscal 2027 and 2028. However, we are maintaining our outlook of \$1 billion while we pursue supply chain capacity and we still expect the first revenues from production chip sales to land in the fourth quarter of this fiscal year.

Turning to operating expenses and profit. Non-GAAP operating expenses were \$734 million. This was about \$10 million below our guidance and up 30% year-on-year due to strong R&D investment. These investments in R&D reflect ongoing engineering expansion to support customer demand for more Arm technology, including continued innovation in next generation architectures, compute subsystems and the recently announced Arm AGI CPU product family.

Non-GAAP operating income was \$731 million, resulting in a non-GAAP operating margin of about 49%. Non-GAAP EPS was \$0.60, driven by both higher revenue and slightly lower OpEx than expected.

Turning now to guidance. For Q1, we expect revenue of \$1.26 billion, plus or minus \$50 million. At the midpoint, this represents revenue growth of about 20% year-on-year. We expect royalty revenue and license and other revenue to both be up around 20% year-on-year. We expect our non-GAAP operating expense to be approximately \$760 million and our non-GAAP EPS to be \$0.40 plus or minus \$0.04.

The strength of demand we're seeing combined with our expanding portfolio and deepening customer engagements gives us confidence in our ability to deliver sustained long-term growth.

For those of you that missed our March Arm Everywhere AGI CPU event, we expect that by FYE31, we will be generating \$15 billion in AGI CPU revenue and \$10 billion in IP revenue for a total of \$25 billion. We expect this to translate to more than \$9 in EPS.

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

QUESTION AND ANSWER SECTION

Operator: Thank you. [Operator Instructions] We will now take our first question. And the first question comes from the line of Andrew Gardiner from Citi.

Andrew Michael Gardiner

Analyst, Citigroup Global Markets Ltd.

Q

Thank you very much for taking the question. Good afternoon all. It's only been six weeks since you held the Arm Everywhere event and announced the new AGI CPU. Somehow it feels a lot longer given the rate of change in the chip industry and in particular the AI space. You can certainly see that as well in the demand you just highlighted. You're now saying that you've got over \$2 billion of demand for the next two fiscal years.

I was just hoping that you could give us a bit more detail on how that additional demand has transpired over the last six weeks or perhaps type of customer. Is some of it greater demand with the launch customers you highlighted to us in San Francisco? Is it all new interest, new customers, applications they've got in mind?

And also, Jason, I just wanted to touch on the point you made there where you're not changing the guidance. The demand has doubled and you're now sort of seeking incremental supply from your foundry and I presume memory partners as well. What do you think you'll be able to do in terms of accessing that supply or how are you going about that? Thanks very much.

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

A

Yeah, thank you for the question, Andrew. So I'll tackle the first part in terms of where the demand is coming from, touch a little bit upon the supply situation and then let Jason address sort of the numbers aspect to it.

So one of the beauties of the product that we announced at the Arm Everywhere event is the fact that the CPU can come in a number of different flavors. But one of the more attractive options is actually buying finished racks from our partners such as Supermicro, Lenovo, ASRock. That enables customers to order and deploy quite quickly. Many customers that we've talked about are already using Arm. So whether that's through internal designs or running designs at the cloud. So a lot of the software work has been completed. So the work required in terms of bringing on new compute capacity that's Arm-based, there isn't a lot of friction to that. So you have a situation where A, software is done and B, you've got availability of a rack design that you can put the data off pretty quickly.

So to your question, it's a combination of both. It's a combination of some of the customers that we talked about during the day increasing their forecast. And there are also customers that we didn't talk about on the day who have said, hey, we are very, very interested and we're ready to deploy. How quickly can we get units?

So the number that we talked about at end of March was supply in place to support \$1 billion of demand. And that includes memory, that includes wafers, that includes packaging, that includes access to test equipment. So for the \$2 billion, we are now in the process of securing supply to support that. And the teams are working round the clock to make sure we can find the right answers for our customers. Relative to how that impacts guidance going forward, let me turn that to Jason.

Jason E. Child

Executive Vice President & Chief Financial Officer, Arm Holdings Plc

A

Yeah. Thanks, Andrew. So in terms of the expectation, I think we said back at the Arm Everywhere event that you should assume probably \$90-ish million or so of revenue, just shy of \$100 million in Q4. At this point, we're not changing that target. As we get deeper into the year, we will provide some kind of indications on how things are going with the supply chain. And then of course, in Q3, we'll give you a much firmer estimate of what we expect to deliver in Q4. And then maybe some indication of what we think FY 2028 will look like as well.

Andrew Michael Gardiner

Analyst, Citigroup Global Markets Ltd.

Q

Thank you, guys. Very clear.

Operator: Thank you.

Jason E. Child

Executive Vice President & Chief Financial Officer, Arm Holdings Plc

A

Thanks, Andrew.

Operator: We will now go [ph] to the next (00:19:07) question. And the question 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 first, just as we think about, I guess, like the royalty rate growth for 1Q and then just thinking about the full year, is there any help that you can provide, just kind of the puts and takes there? Clearly, data center, very strong, accelerating. But then how do you kind of think about consumer electronics, smartphones, et cetera?

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

A

I'll [indiscernible] (00:19:43). Thanks for the question, Joe. So in terms of Q4, as we said before the quarter, we had a bit of a tough comp in that we had a particularly strong ramp of MediaTek [ph] about (00:19:58) 400 a year ago, more so than what we expected this year. And so as a result, you saw a bit of a slowdown in royalty revenue. As indicated by our guidance, we're expecting that to get back to the kind of 20% range by Q1.

So I would say within – the assumptions within our expectations are we will probably continue to see unit growth, I think, actually slip to negative for the mobile market in this last quarter. We're going to continue to see very flattish, maybe slightly negative numbers for the overall market. The lower end of the market is probably going to be where most of that impact is. So it doesn't have too much impact on us. And then, of course, any sort of

negative impact we do see, we expect that to be more than offset by demand in cloud AI or specifically in the data center.

I would say on the data center, there's certainly been a lot of positive announcements from some of our partners, certainly AWS and Google in particular about kind of their deployments and either accelerating some of those deployments. And now I – you could say that the three largest GPU providers across NVIDIA with Vera or Grace, and then of course with now Google pairing TPUs with Axion to the latest Arm-based chip. And then of course, Trainium, all paired with the latest version of Arm-based Graviton. All of those partners are now on Arm and we're expecting that that should provide continued growth and upside throughout the year.

Hard to say exactly what that pace of growth will look like since we – we're not going to announce their numbers that they haven't announced, but we do expect to see continued upside. So any of the weakness that we see in memory and mobile, we do expect that to be largely offset plus, I would say, on the cloud side.

And then outside of that the other category of automotive continues to be growing strong and we don't see any changes there. We continue to gain share and continue to grow kind of in that double-digit growth rate. And so overall, feel pretty good about the royalty outlook for the balance of the year.

Operator: Thank you. We will now take the next question. And the next question comes from Vivek Arya from Bank of America Securities. Please go ahead.

Vivek Arya

Analyst, BofA Securities, Inc.



Thanks for taking my question. Rene, yesterday when AMD spoke about the CPU market, they gave a 2030 TAM of \$120 billion, slightly more than \$100 billion that you had given before. But then they also suggested that they expect to maintain 50% share. I imagine Intel will stay relevant in that timeframe also. And then there are all these captive programs, right, Graviton and Axion and Vera and so forth. So I'm just trying to understand what is Arm's natural kind of niche in the market? Like which of these parties conceptually can you take share from to grow your business? Thank you.

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc



Yeah, thank you for the question. Yeah, when we talked about \$100 billion TAM on March 24 at the Arm Everywhere day, I think we were the first company to talk about numbers in that magnitude, and we had a bunch of questions about it. Now, it's sort of nice to see the rest of the market catching up and going higher than the number.

Could the number be \$120 billion out in that timeframe? Certainly. We are seeing literally not only an explosion of CPU demand, but one of the areas that we're seeing growth in terms of CPU is number of cores per CPU. Many of these agents want to run independent jobs or flows or a batch on a specific CPU core. So again, the Arm AGI CPU is 136 cores, which is much larger than many of the competitive offerings.

In the go forward space, could I see a world of 256 cores, 512 cores? Absolutely. It's a great place for Arm because in a very, very high core count design, what really matters is efficiency per core. And that's where we are – we're world class.

As far as the market share numbers, AMD has 50%, Intel has 50% and we have 50%. So you add up to some crazy number. All I would say is as follows: we see a very, very strong direction with all of the Arm hyperscalers who use our technology today. As Jason mentioned, whether it's NVIDIA, whether it's Amazon, whether it's Google. The very largest and most prevalent accelerators by volume is the TPU, it's Trainium and it's Rubin, today Blackwell going to Rubin. Those all connect to Arm, and increasingly they are going to be 100% Arm. So we feel very, very good about the market share there.

We also talked about a lot of customers such as Cloudflare, such as Meta, such as SAP, SK Telecom, OpenAI, customers who simply aren't going to design their own Arm-based CPU. They've – either for CapEx reasons or engineering reasons. For whatever the reasons are, they're not going to go off and make the spend there. So we think it's a market that we can play in in a very large way. And I think even indicators of AWS selling Graviton to outside partners, it's kind of an indication that there's just huge, huge demand for Arm-based capacity.

So we think we're going to play alongside our partners in this space. And we also think the opportunity is very, very large for both. And I'm actually confident that by the end of the decade, I believe the largest market share by CPU type will be Arm.

Vivek Arya

Analyst, BofA Securities, Inc.

Q

Thanks, Rene.

Operator: Thank you. Our next question today comes from the line of Timm Schulze-Melander from Rothschild & Co. Please go ahead.

Timm Schulze-Melander

Analyst, Rothschild & Co Global Markets Solutions Ltd.

Q

Yeah, hi. Thanks very much for taking my question. So Rene, maybe just to start with you and to key off that CPU TAM commentary you just made there. Just want to check that I heard you right that you anticipate 100% attach rate of Arm CPU with those accelerators you mentioned. And then maybe just looking forward from an OpEx perspective, as you get into that merchant market, as your products attach to some of your partners' products, do you have any undertakings in terms of operating expenses, in terms of in-market customer support? And then I had a quick follow-up for Jason.

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

A

Yeah, thank you for the question, Tim. Yeah. So to clarify my comment, my expectation is that for the Trainium platform over time, TPUs over time and NVIDIA's accelerators over time, I believe that the vast majority of the market share there will be Arm. NVIDIA's there essentially, and we are starting to see that happen with Graviton already over the last number of quarters and the announcement that Google made at Google Next with the TPU8t and TPU8i, the training and inference chips. So that trend is well underway.

And the reason for it, as stated, is that by getting much better performance in the same power envelope, the overall performance of the platform is greatly improved. Google is talking about an 80% improvement in terms of the overall performance. So it's really numbers like that and the advantages that customers see in terms of embracing the platform that gives us very, very high confidence that that trend should continue.

If I understood your question correctly on the OpEx side or customer support issues, et cetera, et cetera, relative to Arm, Arm-based server racks being bought by customers, what I would expect, it would be a relationship that we would typically have with our partners, with the ODMs, where they're building finished racks and designing the customers response for the software. We're not delivering the application software, but anything regarding the low-level code, the firmware, the boot ROM, that's all in us. So we'll have customer support ready for that. And then of course, if there's any issues regarding the hardware, we will manage that all ourselves.

Relative to the OpEx associated with that inside the company, I'll let Jason bridge on to that. But that's all baked in in terms of when we talk about the growth rate of head count. And we've taken all of that into consideration relative to what a customer support platform solution looks like.

Jason E. Child

Executive Vice President & Chief Financial Officer, Arm Holdings Plc

A

Yeah. In terms of the OpEx for support, that was baked in the numbers that we shared at Arm Everywhere and embedded within our long-term guidance, and so already accounted for.

Timm Schulze-Melander

Analyst, Rothschild & Co Global Markets Solutions Ltd.

Q

Got it. And maybe the quick follow-up, just looking at FY 2027 for the year, could – Jason, could you just give us some help in terms of the shape of the – of how you expect royalty revenues to kind of flow through the year and maybe some commentary around the OpEx as well would be really helpful. Thank you.

Jason E. Child

Executive Vice President & Chief Financial Officer, Arm Holdings Plc

A

Yeah. So I expect that the royalty growth is going to be in the roughly 20-ish percent for the year. And so across each of the quarters, plus or minus depending on the quarter, but it will all be relatively close to the 20-ish percent range. And that's really for both license and royalty. The license will be just like the last, I guess, the last three years, will be a little bit back end weighted. I think it's about 60% probably second half versus 40% front half.

And then on OpEx, I know we initially had thought maybe the OpEx would step up a little more into Q – from Q4 into Q1, as we just guided to. We're expecting that to be a little less than we thought before. And now I would expect that the license – or I'm sorry that the OpEx growth is going to grow sequentially every quarter by, call it, a few percent and the overall expense will still kind of be – we'll show incremental margin improvement throughout the year and we will be growing expenses less than revenue by the time we end the year, and it will be building kind of slowly throughout the year. But we will be getting back to delivering incremental margin by end of year, like we – like where we were at a couple years ago before we started investments [indiscernible] (00:31:20).

Timm Schulze-Melander

Analyst, Rothschild & Co Global Markets Solutions Ltd.

Q

Okay, thanks. Great. Very helpful. Thank you.

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

Krish Sankar

Analyst, TD Cowen

Q

Yeah, hi, thanks for taking my question. Rene, I had a question on the CPU to GPU ratio. It seems like with inference, there's a compelling case for AGI CPU. But I'm wondering, what is your assumption on when we get to 1 to 1, and do you expect the CPU to GPU ratio to cross 1 to 1? And also, if you can differentiate between the opportunity in the head node versus the host node, that'll be very helpful. Thank you.

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

A

Yeah, so I think it's a kind of a complex question in terms of thinking about today's world and statically how things look. The way I think the – think about it is that while the ratios may not go to more CPUs than GPUs from a chip standpoint, they probably will from a core count standpoint.

And what do I mean by that? The way to think about Blackwell and Rubin and some of these large accelerators is that they're pretty much reticle limited. Meaning that the size of the chips is already limited by the amount of area that a [ph] mass (00:32:36) can print. So it's not like you're going to get many, many more GPUs. And then one could argue how efficient those CPUs are as they consume all that silicon.

On the flip side, CPUs today, the Arm AGI CPU, for example, it's 136 CPU cores. Vera, that's 88. As I mentioned earlier, could I see those core counts doubling or quadrupling over the next number of years? Absolutely. Does that mean that, oh, the ratio of chips stay the same if one chip has 500 cores and it used to have 136 cores? So clearly, the ratios, they're going to – are going to change from a CPU core count, maybe not a chip count.

Where we'll see the growth, in my opinion, is not so much in the head node to a GPU architecture because it's a little bit fixed given the way the GPU is architected and how it feeds the CPU. But will you see many, many more CPUs inside a data hall, dedicated racks of CPUs that are doing agentic orchestration and scheduling and management? 100%. You simply just have to look at NVIDIA announcing a dedicated Vera rack, 256 Vera CPU chips, 88 cores per chip in a 200-kilowatt liquid cooled rack. That is designed to sit in a data center adjacent to a Vera Rubin system. And that's simply because of the size of the system, it's liquid cooled. So imagine a world where you had scores of Vera Rubin racks. Now you may actually have a Vera rack in between or two Vera racks. So that changes the ratios completely.

I think one thing we know for sure is that we probably have undercalled the CPU demand in terms of the transition here. We talked about a 4x increase. We could get our heads around a bigger number than that. But ratios are a tough way to look at it, just given the math that I just described.

Krish Sankar

Analyst, TD Cowen

Q

Thank you, Rene. Very helpful.

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

Sebastien Naji

Analyst, William Blair & Co. LLC

Q

Yeah, thank you. It's great to see the strong demand for your AGI CPU. One of the questions that is top of mind for many investors is how this shift in your business model is impacting some of your existing IP customers who also sell Arm-based CPUs in some form? Could you maybe just give us a sense of how your larger customers

have reacted since your announcement in March and how you think about managing that potential tension between your product and IP businesses?

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

A

Yeah, thank you for the question. It's a super important question. One of the things that we wanted to ensure relative to the overall strategy of selling silicon is that we had the support of the ecosystem. Now for us, the ecosystem represents a lot of different partners. The ecosystem are the people who build Arm-based chips like Samsung or TSMC, that ecosystem is EDA partners like Synopsys and Cadence; the huge amount of people who work in the software space in the Linux world and everything around Kubernetes containers, for example; and then of course, our licensees, AWS or Microsoft or Google or NVIDIA, in this case, all of them have products.

So we went to them early on this and we explained to them what we were doing. We explained to them why we were doing it, and we explained to them why it was beneficial to the Arm ecosystem that we do it. And that's largely because the more software that's written and optimized for Arm makes everyone stronger. And we asked for their support of the strategy at the Arm Everywhere day, and every single partner we asked said yes.

We had probably over 50 different partners, the names I just gave you, all of them. In some cases, they provided a quote. In some cases they provided references to other partners, and then in some cases they did a video for us, which we used at the event. So I think by the fact that everyone was asked and everyone said yes and how can I help, I think is about as good an endorsement I could give for that. So we're deeply grateful for it. We don't take it for granted. We're very, very appreciative of it. And we think, again, it lifts all boats.

And probably the last thing I would say on this is that the primary reason we did this was that our customers asked for it. And at the end of the day, we are responding to customer demand in a market, and we see this today because we're sold out and we've got people looking for more products. There's demand for these products. And at the end of the day, customer demand speaks volumes.

Sebastien Naji

Analyst, William Blair & Co. LLC

Q

Very clear. Thank you.

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

Vijay Raghavan Rakesh

Analyst, Mizuho Securities USA LLC

Q

Yeah, hi. Thanks. Great quarter and guide here. Rene, just a quick question, as you look at the demand on agentic AI CPUs, and obviously the CSPs seem to be ramping up on that as well, how is your data center royalty revenues growing, I guess, as you look at fiscal 2027 over last year? I'm sure that's picking up, but any color on that?

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

A

The royalties associated with customers building chips based on Neoverse, those royalties have doubled year-on-year. And I'm looking at Jason here. I expect they're going to double year-on-year again in this year. So to answer that question – the answer to that statement is yes. So that business is incredibly strong. And when we talked last

quarter in February and I said that this was going to be our largest business, I was speaking about it in the context of just how big the royalties were.

So now when we add the Arm AGI CPU business, we have two extremely strong sources of revenue that I think the best way to think about it is don't cannibalize each other. They're going to run in tandem with each other. We're going to have a strong business around the ARM AGI CPU. We talked about \$15 billion by fiscal year 2031. And then we also have an IP business that we expect to double to \$10 billion, and that IP business will largely be driven by datacenter [ph] revenue (00:39:39).

Vijay Raghavan Rakesh

Analyst, Mizuho Securities USA LLC

Q

Got it. Thanks. And then on the licensing revenue side, should we still think of it as a high-single digit percent growth fiscal 2027, 2028, or is that trending higher too with Armv9 and other stuff?

Jason E. Child

Executive Vice President & Chief Financial Officer, Arm Holdings Plc

A

Well, for this year, I think we said you should expect license revenue to be more in the 20-ish percent range. Over time and I would expect, I think, the long-term target is probably in the high-single digit, low-double digit range. But it's hard to say. We've seen this kind of AI investment supercycle or whatever folks are calling it has now gone on for three years. So who knows how much longer it goes, but it's at least going to happen for the next year. And beyond that, hard to say. But I would expect at least 10% year-on-year growth for the long term as probably being the floor that we can see for right now.

Vijay Raghavan Rakesh

Analyst, Mizuho Securities USA LLC

Q

Great. Fantastic. Thank you.

Operator: Thank you. We will now take the next question. And the question comes from the line of Harlan Sur from JPMorgan. Please go ahead.

Harlan Sur

Analyst, JPMorgan Securities LLC

Q

Hey, good afternoon. Thanks for taking my question and congrats on the expanded AGI demand profile. On the billion dollars of revenues for AGI across 2027 and 2028, most of it obviously is going to be in 2028. I think at the event, you told us gross margin profile on these first gen products is about 30% plus.

I know it's embedded your total cost structure, but Jason, can you just give us a rough sense on the OpEx attributed to supporting your chip business this year and next year? I know it's early revenues, but I'm just trying to figure out like the scale of the chip business cost structure. And as you drive your chip revenues higher, like what – when can you drive accretion to the earnings power? I know the calendar 2030 endpoint, I think, is roughly \$3 billion in OpEx for the chip business. But where is it this year and next year?

Jason E. Child

Executive Vice President & Chief Financial Officer, Arm Holdings Plc

A

Yeah, I would – it's a good question. I would say, yeah, the revenue split for 2027, 2028, something like \$90-ish million for Q4 2027 and then \$910 million or whatever for 2028. That's kind of what we laid out five, six weeks

ago. And as said, we have demand above that. But for right now, let's just assume that's the number until we work through some of the wafer and memory shortage issues.

In terms of the OpEx, yeah, so the OpEx that's in our plan and what we've guided to, that does include the support related to that business. If you kind of just look at that business on a standalone basis, we see it as a business that because a lot of the cost of development is really – probably the most expensive part of developing the chip is really the compute die, which really is effectively kind of the CSS, so we're able to kind of leverage that. That automatically makes this business much more profitable as a standalone chip business given that we get kind of that work as part of our IP business already.

So the incremental costs and OpEx we are having to add to the chip business is not that significant. It's a team that's in the probably in the dozens of people, not hundreds. And so you can assume that it's operating profit positive next year. And when you go out to, say, 2031, I think we said the IP business probably gets to something like a 65% operating margin business or EBITDA margin business and that the chip business is probably in the 35-ish percent range. And that's where we expect to get to.

How quickly we get there is probably going to be a function of how fast the revenue grows. And at least at a \$15 billion rate, I think those are the right numbers based on everything we see right now.

Operator: Thank you. We will now take our final question for today. And the final question comes from the line of Lee Simpson from Morgan Stanley. Please go ahead.

Lee Simpson

Analyst, Morgan Stanley & Co. International Plc

Q

Great. Thanks so much for squeezing me in at the end here. I just wanted to go back to the ratio of CPUs to GPUs, particularly as it relates to orchestration. And I think, Rene, you mentioned it's better to look at this really from a core count perspective, particularly when handling these agentic sessions. I'm just trying to get a sense from a bottom-up perspective, are we looking at this as each agentic flow requires one core? Or could we look at this differently and say that there is an average number of Arm instructions that are needed to orchestrate each token generated by the accelerator? Thanks.

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

A

Yeah. Thank you, Lee. You've gone into deep math there. I think the latter is a little too complicated to think about it. The – maybe a more straightforward way to think about it is each of these agents are running a batch or running a job themselves. There is certainly a level of complexity in terms of the way the branch prediction, coding is handled and essentially the way you would code the example.

But if you just think about the nature of an asynchronous workload for an agent, it runs a job, it does some scheduling, it stops, it waits, it pauses. It's actually pretty good for a single core design to handle that, as opposed to having multi-cores, having to run that all together in unison. It's going to be more power efficient if you run it through a single core. And the more cores you have in theory, the more batches you can run.

So our viewpoint is very much one of more cores is better. And that's why I think you're going to see increasingly larger core counts in these CPU chips. So you'll see more CPUs, cores being shipped. You may not see three extra chips, but the chips will be more expensive, which is why when people look at \$100 billion TAM in five years or \$125 billion or whatever the number is, largely it's going to be driven by the fact that these CPU chips are going

to have lots of CPU cores which will drive ASP up. But I think it's a core – it's a per core batch job, not a – not multiple instructions across multiple cores.

Operator: Thank you. I will now hand the call back for closing remarks.

Rene Anthony Andrada Haas

Chief Executive Officer & Director, Arm Holdings Plc

Thank you and thank you for all the questions. This was an amazing quarter, \$1.5 billion, approximately in revenue. Not long ago is what Arm did in a single year. And can do nearly \$5 billion in revenue a few years after the public offering of the company when we projected around this time we will be doing about \$4.5 billion is a great testament to all the work done by Arm's employees, customers and partners.

What we are seeing, to be clear, is unprecedented compute demand, and that we are at the center of that demand growth. The best way to think about Arm's growth trajectory, particularly now that we've announced the Arm AGI CPU is that we have two growth vectors that will drive this. The ARM AGI chip, as we talked about, with \$1 billion of demand in March 2024. We're now seeing over \$2 billion of demand over the next two years, 2x what we talked about. And we're on track to achieve \$15 billion of revenue by fiscal 2031.

And our business around IP, the Neoverse IP with CSS has doubled year-on-year. We're projecting it to double again year-on-year and we are seeing adoption and acceleration with AWS, Google, NVIDIA and Microsoft. Both of these vectors represent a structural growth for Arm that is very, very strong and very sustainable.

With that, thank you again for all your questions and we appreciate all your interest in Arm.

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.