

Hi, I'm Sharbani, and I'm one of the over 2,100 teammates that we have here at Arm obsessing every day about software.

I feel like Drew and Chris and Mohamed kind of did my job for me already, but it's probably a good time to underscore: you need software to unlock the power of incredible hardware. You can't have hardware without great software, and that's why the breadth, depth, and value that we provide to our partners through our software ecosystem is our advantage.

Our ecosystem is massive. Rene already called out that we have 22-plus million developers building on Arm. Twenty-two million. I'm pretty sure that makes this the biggest software ecosystem there is.

Now when I think about this breadth, it covers our full platform. I wanted to pull all of these logos together that you've seen. We covered each of these kinds of individually, and this is just a representation of the various companies, the various partners that we work with. But it goes from cloud to edge. It goes from the phones that are in your pockets to the laptops that some of you are working on right now to run amazing reports talking about all of the awesomeness that we are bringing forth here today.

It's helping to power the agents that you're going to use to help you write your next reports, to help you figure out what your next strategies are going to be. It's in the cars. It's in the robots. It's everywhere. This ecosystem is absolutely massive.

It's not just about it being the most ubiquitous compute platform. We also have incredible depth. We go up and down the stack. If you look, you can see we've got the best operating systems all the way up to your favorite applications, like Santosh talked about earlier today.

When it comes to depth, the depth of our software ecosystem is our durable advantage. We've been at this for a while. We've been doing this for 15 years. We have over 1,300 open-source and upstream projects that we work on, and we have over 50,000 partners — 50,000 companies that we collaborate with. Twenty-two million developers building on Arm, 50,000 companies — that's a lot of elbow grease. We've been working really hard, right Drew?

And like I said, we work up and down the stack. We start off at firmware, doing kernel-level optimizations. We're working up through open source, through all the operating systems, all the way up to the applications. We're incredibly strategic about how we think about where we need to optimize software.

That's a lot of software. That's a lot of territory to have to cover. So we have to think about how we're going to have the absolute best leverage and the best ROI on our investments in each of these projects.

That means a lot of times we're thinking about how we have optimizations at the operating-system level or, like Chris was talking about, our Kleidi optimizations. We want to be thinking about how we're not just working with partners, but also how we're putting things out into the ecosystem, helping our partners with their proprietary stacks, so the software optimizations that we provide just run seamlessly from cloud to edge and work behind the scenes most of the time like magic.

When I think about what we want to obsess over when we're trying to drive the depth of these optimizations, I'm going to tell you all the same thing I tell my team every day: we obsess over three KPIs — performance per watt, making it as easy as possible for developers to build on Arm, and making it as fast as possible for them to deliver innovation into the market through those applications that each of you love to use every day.

These advantages only compound with the AI economy, with our agentic future.

[01:23:07] How many of y'all are using agents today? Going to hope a lot more after this session as well.

I'm going to let you in on a hot tip: I like to think about agents like apprentices. Each of you are great artists in what you're trying to do, and an apprentice can only help you amplify and provide all of the additional support you need to have your next great innovation.

To do that, they have to do a lot of experimentation. They're running all the time. They've got to be running as fast as possible. That means we're going to have to swap in and out a lot of different types of software.

So when you send your agent to go and run off and solve something for you, you want it to be able to have good judgment and be able to quickly swap out different models or different cloud providers so that you're not just blowing the bank on your latest experiment, or if it's something that's absolutely mission critical.

The good news is, the agents are going to deliver the work to your fingertips — to your phone, to your laptop — because when you think about our ubiquitous platform, that's all running on Arm. The agents are going to be delivering the value where Arm is actually serving you the best.

Then the software ecosystem comes into play because our software has been completely optimized across the board from cloud to edge, so these agents can really

easily go and pick and choose software that's going to work on those devices. They don't have to think about it. It's just going to work. It's going to run great on Arm.

These advantages only compound. So it's breadth, depth, and I want to talk a little bit more about the value.

We provide our partners with incredible value by helping them gain horizontal leverage across the Arm platform. I want to jump into an example with Meta. Seems appropriate — we heard a lot from Santosh, from Paul today.

Meta was facing a challenge, and we've been working with Meta for years, not just on the hardware side but also on the software side, and this is something that's near and dear to my heart as well.

You think about Meta: Meta is famous for having a single stack, a monorepo, where you have a developer who is pushing something new, building one of your favorite applications and a new feature, and whenever they push something out into the internal code base, everybody's able to go and pick that up. It's got to meet an incredibly high bar, but you have to have horizontal leverage within the company so anybody can pick up the best-in-class, the latest and greatest, because they're always at the cutting edge.

This means that when we're partnering with Meta on software, we have to be at our absolute A game.

Now they had a challenge because they wanted to unlock the value that we have from some of our new architecture and from our CPUs. Chris talked a little bit about Kleidi and our SME2 capabilities. Now that's incredible performance and incredible perf per watt that we're driving for.

[01:26:13] But when you think about how models are actually made, you're training massive models in data centers, right? You've got to use a ton of compute. You can have billions of parameters, and you've got to squish that down and put it onto a device. Again, how am I going to fit that in my pocket so I can go be checking my favorite update on Instagram or sending my family great videos on WhatsApp?

So, what we knew we needed to do was jump in and help them with optimizations through their AI frameworks.

Many of you may be familiar with PyTorch. It has been the fastest-growing framework. I have a little bit of bias — I used to work in frameworks at another company — so I love all frameworks, they're all great. But I can see the power of needing to take those AI frameworks, which can become quite large, and the models are larger, the frameworks

themselves are large, and we needed to be able to help Meta optimize once and deploy many, to fit these models onto the devices that we're all obsessing over.

So, what we ended up doing is we worked with them through PyTorch. Arm actually has an active seat on the board. We help out with the Technical Advisory Board, and we help direct both PyTorch in open source as well as work directly with Meta.

Then we also provided incredible contributions to ExecuTorch, which some of you may have seen went 1.0 last October. That means, again, it's not only benefiting Meta and all of the developers inside Meta, but it's also benefiting open source altogether.

So now when you're going out and some of you might be playing around with PyTorch models, you don't have to go through a whole long path to deploy it on Arm. You can actually go through ExecuTorch, you can go faster, it can be easier for you, and you can still get the best perf per watt that you get on the Arm ubiquitous compute platform.

So that's a repeatable performance gain and efficiency for Meta, and that's also something where we've been able to take those lessons and extend them more broadly into our ecosystem.

Again, we have breadth — a massive, massive ecosystem. We have the depth, and we're strategic about how we help up and down the stack. We work directly with our partners, both on their proprietary stacks and to help them in open source, so that we can have these repeatable gains that compound and will compound for years to come.

All right. I think next Jason's going to come up and give us a little bit of an overview.