Broadcom Inc., Q1 2025 Earnings Call, Mar 06, 2025

25. 3. 7.

Operator

Welcome to the Broadcom Inc.'s First Quarter Fiscal Year 2025 Financial Results Conference.

At this time, for opening remarks and introductions, I would like to turn the call over to Ji Yoo, Head of Investor Relations of Broadcom Inc.

Ji Yoo

Thank you, Sheri, and good afternoon, everyone. Joining me on today's call are Hock Tan, President and CEO; Kirsten Spears, Chief Financial Officer; and Charlie Kawwas, President, Semiconductor Solutions Group.

Broadcom distributed a press release and financial tables after the market closed, describing our financial performance for the first quarter of fiscal year 2025. If you did not receive a copy, you may obtain the information from the investors section of Broadcom's website at broadcom.com.

This conference call is being webcast live and an audio replay of the call can be accessed for 1 year through the Investors section of Broadcom's website.

During the prepared comments, Hock and Kirsten will be providing details of our first quarter fiscal year 2025 results, guidance for our second quarter of fiscal year 2025 as well as commentary regarding the business environment. We'll take questions after the end of our prepared comments.

Please refer to our press release today and our recent filings with the SEC for information on the specific risk factors that could cause our actual results to differ materially from the forward-looking statements made on this call.

In addition to U.S. GAAP reporting, Broadcom reports certain financial measures on a non-GAAP basis. A reconciliation between GAAP and non-GAAP measures is included in the tables attached to today's press release. Comments made during today's call will primarily refer to our non-GAAP financial results.

I'll now turn the call over to Hock.

Hock Tan

Thank you, Ji. And thank you, everyone, for joining today. In our fiscal Q1 2025, total revenue was a record $14.9 billion, up 25% year-on-year, and consolidated adjusted EBITDA was a record again, $10.1 billion, up 41% year-on-year.

So let me first provide color on our semiconductor business. Q1 semiconductor revenue was $8.2 billion, up 11% year-on-year. Growth was driven by AI as AI revenue of $4.1 billion was up 77% year-on-year. We beat our guidance for AI revenue of $3.8 billion due to stronger shipments of networking solutions to hyperscalers on AI. Our hyperscale partners continue to invest aggressively in their next generation frontier models, which do require high-performance accelerators as well as AI data centers

with larger clusters.

And consistent with this, we are stepping up our R&D investment on 2 fronts. One, we're pushing the envelope of technology in creating the next generation of accelerators. We're tapping out the industry's first 2-nanometer AI XPU packaging 3.5D as we drive towards a 10,000 teraflops XPU.

Secondly, we have a view towards scaling clusters of 500,000 accelerators for hyperscale customers. We have doubled the radix capacity of the existing Tomahawk 5. And beyond this, to enable AI clusters to scale up on Ethernet towards 1 million XPUs. We have tapped out our next-generation 100-terabit Tomahawk 6 switch, running 200G studies and 1.6 terabit bandwidth. We will be delivering samples to customers within the next few months.

These R&D investments are very aligned with the road map of our 3 hyperscale customers as they each race towards 1 million XPU clusters by the end of 2027. And accordingly, we do reaffirm what we said last quarter that we expect these 3 hyperscale customers will generate a serviceable addressable market or SAM in the range of $60 billion to $90 billion in fiscal 2027. Beyond these 3 customers, we had also mentioned previously that we are deeply engaged with 2 other hyperscalers in enabling them to create their own customized AI accelerator. We are on track to take out their XPUs this year.

In the process of working with the hyperscalers, it has become very clear that while they are excellent in the software, Broadcom is the best in hardware. Working together is what optimizes via large language models. It is, therefore, no surprise to us. Since our last earnings call that 2 additional hyperscalers have selected Broadcom to develop custom accelerators to train their next-generation frontier models. So even as we have 3 hyperscale customers, we are shipping XPUs in volume today, there are now 4 more who are deeply engaged with us to create their own accelerators. And to be clear, of course, these 4 are not included in our estimated SAM of $60 billion to $90 billion in 2027.

So we do see an exciting trend here. New frontier models and techniques put unexpected pressures on AI systems. It's difficult to serve all classes of models with a single system design point. And therefore, it is hard to imagine that a general purpose accelerator can be configured and optimized across multiple frontier models. And as I mentioned before, the trend towards XPUs is a multiyear journey. So coming back to 2025, we see a steady ramp in deployment of our XPUs and networking products. In Q1, AI revenue was $4.1 billion, and we expect Q2 AI revenue to grow to $4.4 billion, which is up 44% year-on-year.

Turning to non-AI semiconductors. Revenue of $4.1 billion was down 9% sequentially on a seasonal decline in wireless. In aggregate, during Q1, the recovery in non-AI semiconductors continue to be slow. Broadband, which bottomed in Q4 of 2024, showed a double-digit sequential recovery in Q1 and is expected to be up similarly in Q2 and service providers and telcos step up spending. Server storage was down single digits sequentially in Q1 but is expected to be up high single digits sequentially in Q2.

Meanwhile, enterprise networking continues to remain flattish in the first half of fiscal '25 as customers continue to work through channel inventory. While wireless was down sequentially due to a seasonal decline, it remained flat year-on-year. In Q2, wireless is expected to be the same, flat again year-on-

year.

Resales in industrial were down double digits in Q1 and are expected to be down in Q2. So reflecting the foregoing puts and takes, we expect non-AI semiconductor revenue in Q2 to be flattish sequentially even though we are seeing bookings continue to grow year-on-year. In summary, for Q2, we expect total semiconductor revenue to grow 2% sequentially and up 17% year-on-year to $8.4 billion.

Turning now to infrastructure software segment. Q1 infrastructure software revenue of $6.7 billion was up 47% year-on-year and up 15% sequentially, exaggerated though by though by deals, which slipped from Q2 -- Q4 in Q1. Now this is the first quarter, Q1 '25 where the year-on-year comparables include VMware in both quarters.

We're seeing significant growth in the software segment for 2 reasons: One, we're converting to a footprint of large -- sorry, we're converting from a footprint of largely perpetual license to one of full subscription. And as of today, we are over 60% down; two, these perpetual licenses were only largely for compute virtualization, otherwise called vSphere. We are upselling customers to a full stack VCF, which enables the entire data center to be virtualized.

And this enables customers to create their own private cloud environment on-prem. And as of the end of Q1, approximately 70% of our largest 10,000 customers have adopted VCF. As these customers consume VCF, we do see a further opportunity for future growth. As large enterprises adopt AI, they have to run their AI workloads on their on-prem data centers, which will include both GPU servers as well as traditional CPUs.

And just as VCF virtualizes these traditional data centers using CPUs, VCF will also virtualize GPUs on a common platform and enable enterprises to import AI models to run their own data on-prem. This platform, which virtualized the GPU is called the VMware Private AI Foundation. And as of today, in collaboration with NVIDIA, we have 39 enterprise customers for the VMware Private AI Foundation.

Customer demand has been driven by our open ecosystem, superior load balancing and automation capabilities that allows them to intelligently pull and run workloads across both GPU and CPU infrastructure and leading to very reduced costs.

Moving on to Q2 outlook for software. We expect revenue of $6.5 billion, up 23% year-on-year. So in total, we're guiding Q2 consolidated revenue to be approximately $14.9 billion, up 19% year-on-year. And this -- we expect this will drive Q2 adjusted EBITDA to approximately 66% of revenue.

With that, let me turn the call over to Kirsten.

Kirsten Spears

Thank you, Hock. Let me now provide additional detail on our Q1 financial performance. From a year-on-year comparable basis, keep in mind that Q1 of fiscal 2024 was a 14-week quarter and Q1 of fiscal 2025 is a 13-week quarter.

Consolidated revenue was $14.9 billion for the quarter, up 25% from a year ago. Gross margin was 79.1% of revenue in the quarter, better than we originally guided on higher infrastructure software revenue and more favorable semiconductor revenue mix.

Consolidated operating expenses were $2 billion, of which $1.4 billion was for R&D. Q1 operating income of $9.8 billion was up 44% from a year ago, with operating margin at 66% of revenue.

Adjusted EBITDA was a record $10.1 billion or 68% of revenue, above our guidance of 66%. This figure excludes $142 million of depreciation.

Now a review of the P&L for our 2 segments. Starting with semiconductors. Revenue for our semiconductor solutions segment was $8.2 billion and represented 55% of total revenue in the quarter. This was up 11% year-on-year. Gross margin for our semiconductor solutions segment was approximately 68%, up 70 basis points year-on-year driven by revenue mix.

Operating expenses increased 3% year-on-year to $890 million on increased investment in R&D for leading-edge AI semiconductors, resulting in semiconductor operating margin of 57%.

Now moving on to infrastructure software. Revenue for infrastructure software of $6.7 billion was 45% of total revenue and up 47% year-on-year based primarily on increased revenue from VMware. Gross margin for infrastructure software was 92.5% in the quarter compared to 88% a year ago. Operating expenses were approximately $1.1 billion in the quarter, resulting in infrastructure software operating margin of 76%. This compares to operating margin of 59% a year ago. This year-on-year improvement reflects our disciplined integration of VMware and sharp focus on deploying our VCF strategy.

Moving on to cash flow. Free cash flow in the quarter was $6 billion and represented 40% of revenue. Free cash flow as a percentage of revenue continues to be impacted by cash interest expense from debt related to the VMware acquisition and cash taxes due to the mix of U.S. taxable income, the continued delay in the reenactment of Section 174 and the impact of corporate AMT. We spent $100 million on capital expenditures.

Days sales outstanding were 30 days in the first quarter compared to 41 days a year ago. We ended the first quarter with inventory of $1.9 billion, up 8% sequentially to support revenue in future quarters. Our days of inventory on hand were 65 days in Q1 as we continue to remain disciplined on how we manage inventory across the ecosystem. We ended the first quarter with $9.3 billion of cash and $68.8 billion of gross principal debt.

During the quarter, we repaid $495 million of fixed rate debt and $7.6 billion of floating rate debt with new senior notes, commercial paper and cash on hand, reducing debt by a net $1.1 billion.

Following these actions, the weighted average coupon rate and years to maturity of our $58.8 billion in fixed rate debt is 3.8% and 7.3 years, respectively. The weighted average coupon rate and years to maturity of our $6 billion in floating rate debt is 5.4% and 3.8 years, respectively, and our $4 billion in commercial paper is at an average rate of 4.6%.

Turning to capital allocation. In Q1, we paid stockholders $2.8 billion of cash dividends based on a

quarterly common stock cash dividend of $0.59 per share. We spent $2 billion to repurchase 8.7 million AVGO shares from employees as those shares vested for withholding taxes. In Q2, we expect the non-GAAP diluted share count to be approximately 4.95 billion shares.

Now moving on to guidance. Our guidance for Q2 is for consolidated revenue of $14.9 billion, with semiconductor revenue of approximately $8.4 billion, up 17% year-on-year. We expect Q2 AI revenue of $4.4 billion, up 44% year-on-year. For non-AI semiconductors, we expect Q2 revenue of $4 billion. We expect Q2 infrastructure software revenue of approximately $6.5 billion, up 23% year-on-year.

We expect Q2 adjusted EBITDA to be about 66%. For modeling purposes, we expect Q2 consolidated gross margin to be down approximately 20 basis points sequentially on the revenue mix of infrastructure software and product mix within semiconductors. As Hock discussed earlier, we are increasing our R&D investment in leading edge AI in Q2, and accordingly, we expect adjusted EBITDA to be approximately 66%. We expect the non-GAAP tax rate for Q2 and fiscal year 2025 to be approximately 14%.

That concludes my prepared remarks. Operator, please open up the call for questions.

Operator

[Operator Instructions] And our first question will come from the line of Ben Reitzes with Melius.

Benjamin Reitzes

Thanks a lot and congrats on the results. Hock, you talked about 4 more customers coming online. Can you just talk a little bit more about the trend you're seeing? Can any of these customers be as big as the current 3? And what does it say about the custom silicon trend overall and your optimism and upside to the business long term?

Hock Tan

Very interesting question, Ben, and thanks for your kind wishes. But what we've seen is -- and by the way, these 4 are not yet customers as we define it. As I've always said, in developing and creating XPUs, we are not really the creator of those XPUs, to be honest. We enable each of those hyperscalers partners we engage with to create that chip and basically to create that compute system, call it that way. And it comprises the model, the software model, working closely with the compute engine, the XPU and the networking that ties together the clusters, those multiple XPUs as a whole to train those large frontier models.

And so -- and the fact that we create the hardware, it still has to work with the software models and algorithms of those partners of ours before it becomes fully deployable at scale, which is why we define customers in this case as those where we know they have deployed at scale and we received the production volume to enable it to run. And for that, we only have just to reiterate. The 4, I call it partners who are trying to create the same thing as the first 3 and to run their own frontier models, each of it don't have to train their own frontier models.

And as I also said, it doesn't happen overnight. To do the first chip could take -- would take typically 1.5 years, and that's very accelerated and which we could accelerate given that we essentially have a framework and a methodology that works right now and works for the 3 customers, no reason for it to not work for 4. But we still need those 4 partners to create and to develop the software, which we don't do to make it work.

And to answer your question, there's no reason why these 4 guys would not create a demand in the range of what we're seeing with the first 3 guys but probably later. It's a journey. They started it later, and so they will probably get there later.

Operator

One moment for our next question, and that will come from the line of Harlan Sur with JPMorgan.

Harlan Sur

Great job on the strong quarterly execution Hock and team. Great to see the continued momentum in the AI business here in the first half of your fiscal year and the continued broadening out of your AI ASIC customers. I know Hock last earnings, you did call out a strong ramp in the second half of the fiscal year, driven by new 3-nanometer AI accelerated programs kind of ramping.

Can you just help us either qualitatively, quantitatively profile the second half step-up relative to what the team just delivered here in the first half? Has the profile changed either favorably, less favorably versus what maybe 90 days ago? Because quite frankly, I mean, a lot has happened since last earnings, right? You've had the dynamics like DeepSeek and focus on AI model efficiency but on the flip side, you've had strong CapEx outlooks by your cloud and hyperscale customers. So any color on the second half AI profile would be helpful.

Hock Tan

Asking me to look into the minds of my customers, and I hate to tell me they don't tell you, they don't show me the entire mindset here. But why we're bidding the numbers so far in Q1 and seems to be encouraging in Q2 partly from improved networking shipments, as I indicated, to cost those XPUs in AI accelerators even in some cases, GPUs together for the hyperscalers. And that's good. And partly also, we think there is some pull-ins of shipments and acceleration, call it that way, of shipments in fiscal '25.

Harlan Sur

And on the second half, that you talked about 90 days ago, the second half 3-nanometer ramp? Is that still very much on track?

Hock Tan

Harlan, thank you. I only guide Q2, sorry. Let's not speculate on the second half.

Operator

One moment for our next question, and that will come from the line of William Stein with Truist

Securities.

William Stein

Congrats on these pretty great results. It seems from the news headlines about tariffs and about DeepSeek that there may be some disruptions, some customers and some other complementary suppliers seem to feel a bit paralyzed perhaps and have difficulty making tough decisions. Those tend to be really useful times for great companies to sort of emerge as something bigger and better than they were in the past. You've grown this company in a tremendous way over the last decade plus, and you're doing great now, especially in this AI area.

But I wonder if you're seeing that sort of disruption from these dynamics that we suspect are happening based on headlines what we see from other companies? And how -- aside from adding these customers in AI, I'm sure there's other great stuff going on but should we expect some bigger changes to come from Broadcom as a result of this?

Hock Tan

You asked a -- you posed a very interesting set of issues and questions. And those are very relevant, interesting issues. The only issue -- the only problem we have at this point is, I would say it's really too early to know where we all land. I mean that's the threat, the noise of tariffs especially on chips that hasn't materialized yet, nor do we know how it will be structured. So we don't know. But we do experience and we are leaving it now is the disruption that is in a positive way, I should add a very positive disruption in semiconductors on a generative AI.

Generative AI for sure, and I said that before also at the risk of repeating here but it's -- we feel it more than ever. It's really accelerating the development of semiconductor technology, both process and packaging as well as design towards higher and higher performance accelerators and networking functionality. We're seeing that innovation that those upgrades occur on every month as we face new interesting challenges.

And when -- particularly with XPUs, we're trying -- we've been asked to optimize to frontier models of our partners, our customers as well as our hyperscale partners. And we -- it's a lot of -- I mean it's a privilege almost for us to be -- to participate in it and try to optimize. And by optimize, I mean, you look at an accelerator, you can look at it from a simple term, high level to perform to one -- to be measured not just on one single metric, which is compute capacity, how many teraflops. It's more than that. It's also tied to the fact that this is a distributed computing problem. It's not just the compute capacity of a single XPU or GPU. It's also the network bandwidth. It ties itself to the next adjacent XPU or GPU. So that has an impact.

So you're doing that, you have to balance with that. Then you decide, are you doing training or you're doing pre-filling? Post-training, fine tuning. And again, then it comes how much memory do you balance against that. And with it, how much latency you can afford, which is memory bandwidth. So you look at least 4 variables, maybe even 5 if we include in memory bandwidth, not just memory capacity when you go straight to inference.

So we have all these variables to play with. And we try to optimize it. So all this is very, very -- I mean, it's a great experience for our engineers to push the envelope on how to create all those chips. And -- so that's the biggest disruption we see right now from sheer trying to create and push the envelope on generative AI, trying to create the best hardware infrastructure to run it.

Beyond that, there are other things, too, that come into play because with AI, as I indicated, it does not just drive hardware for enterprises, it drives the way they architect their data centers. Data requirement -- keeping data private under control becomes important. So suddenly, the push of workloads towards public cloud may take a little pause as large enterprises, particularly have to take -- to recognize that you want to run AI workloads. You probably think very hard about running them on-prem.

And suddenly, push yourself towards saying, you've got to upgrade your own data centers to do and manage your own data to run it on-prem. And that's also pushing a trend that we have been seeing now over the past 12 months. Hence, my comments on VMware Private AI Foundation. This is true, especially enterprises pushing direction are quickly recognizing that how well do they run their AI workloads.

So those are trends we see today and a lot of it coming out of AI, a lot of it coming out of sensitive rules on sovereignty in cloud and data. As far as you mentioning tariffs is concerned, I think that's too early for us to figure out where do we all land. And probably maybe give it another 3, 6 months, we'll probably have a better idea where to go.

Operator

One moment for our next question, and that will come from the line of Ross Seymore with Deutsche Bank.

Ross Seymore

Hock, I want to go back to the XPU side of things. And Going from 4 new engagements, not yet named customers, 2 last quarter and 2 more today that you announced, I want to talk about going from kind of design into deployment to judge that because there is some debate about tons of design wins but the deployments actually don't happen either that they never occur or that the volume is never what is originally promised. How do you view that kind of conversion ratio? Is there a wide range around it? Or is there some way you could help us kind of understand how that works?

Hock Tan

Well, Ross, interesting question. I'll take the opportunity to say, the way we look at design win is probably very different from the way many of our peers look at it out there. Number one, to begin with, we believe design win when we know our product is produced in scale -- at scale and is actually deployed, literally deployed in production. So that takes a long lead time because from taping out, getting in the product, it takes a year easily from the product in the hands of our partner to when it goes into scale production, it will take 6 months to a year is our experience, that we've seen, number one.

And number two, I mean, producing and deploying 5,000 XPUs, that's a joke. That's not real production in our view. And so we also limit ourselves in selecting partners to people who really need that large volume. You need that large volume from our viewpoint in scale right now, in mostly training, training of large language models, frontier models in the continuing trajectory. So we eliminate ourselves to how many customers or how many potential customers that exist out there, Ross, and we tend to be very selective who we pick to begin with.

So when we say design, it really is at scale. It's not something that starts in 6 months and die in a year and die again. Basically, it's a selection of customers. It's just the way we run our ASIC business in general for the last 15 years. We pick and choose the customers because we know this and we do multiyear road maps with these customers because we know these customers are sustainable. I'll put it bluntly. We don't do it for start-ups.

Operator

And one moment for our next question, and that will come from the line of Stacy Rasgon with Bernstein Research.

Stacy Rasgon

I wanted to go to the 3 customers that you do have in volume today. And what I wanted to ask was, is there any concern about some of the new regulations or the AI diffusion rules that are going to get put in place supposedly in May impacting any of those design wins or shipments. It sounds like you think all 3 of those are still on at this point but anything you could tell us about where is that new regulations or AI diffusion rules impacting any of those wins would be helpful.

Hock Tan

Thank you. In this era or this current era of geopolitical tensions and fairly dramatic actions all around by governments, there's always some concern at the back of everybody's mind. But to answer your question directly, no, we don't have any concerns.

Stacy Rasgon

Got it. So none of those are going into China or to Chinese customers then?

Hock Tan

No comment. Are you trying to [indiscernible] who they are?

Operator

One moment for our next question, and that will come from the line of Vivek Arya with Bank of America.

Vivek Arya

Hock, whenever you have described your AI opportunity, you've always emphasized the training workload. But the perception is that the AI market could be dominated by the inference workload,

especially with these new reasoning model. So what happens to your opportunity and share if the mix moves more towards inference. Does it create a bigger TAM for you than the $60 billion to $90 billion? Does it keep it the same but there is a different mix of product? Or does it more inference heavy market favor a GPU over an XPU?

Hock Tan

That's a good -- interesting question. By the way, I never -- and I do talk a lot about training. We do our XPUs also focus on inference as a separate product line. They do. And that's why I can say the architecture of those chips are very different from the architecture of the training chips. And so it's a combination of those 2, I should add, that adds up to this $60 billion to $90 billion. So if I have not been clear, I do apologize, it's a combination of both. But having said that, the larger part of the dollars come from training, not inference within the service, the SAM that we have talked about so far.

Operator

One moment for our next question and that will come from the line of Harsh Kumar with Piper Sandler.

Harsh Kumar

Thanks Broadcom team and again, great execution. Just Hock, had a quick question. We've been hearing that almost all of the large clusters that are 100,000 plus, they're all going to Ethernet. I was wondering if you could help us understand the importance of when the customer is making a selection, choosing between a guy that has the best switch ASIC such as you versus a guy that might have the compute there, can you talk about what the customer is thinking and what are the final points that they want to hit upon when they make that selection for the NIC cards?

Hock Tan

Okay. No, it's -- yes, it's down to -- in the case of the hyperscalers now very much so, is very driven by performance. And its performance, what you're mentioning on connecting, scaling up and scaling out those AI accelerators, be they XPU or GPU among hyperscalers. In most cases, among those hyperscalers, we engage with when it comes to connecting those clusters. They are very driven by performance. I mean if you are in a race to really get the best performance out of your hardware as you train and continue to train your frontier models, that matters more than anything else.

So the basic first thing they go for is proven. That's a proven piece of hardware. It's a proven system, subsystem in our case, that makes it work. And in that case, we tend to have a big advantage because I mean, networking RS, switching and routing RS for the last 10 years at least. And the fact that it's AI just makes it more interesting for our engineers to work on. And -- but it's basically based on proven technology and experience in pushing that -- and pushing the envelope on going from 800 gigabit per second bandwidth to 1.6, and moving on 3.2, which is exactly why we keep stepping up the rate of investment in coming up with our products where we take Tomahawk 5. We doubled the radix to deal with just 1 hyperscaler because they want high radix to create larger clusters while running bandwidth

that are smaller but that doesn't stop us from moving ahead to the next generation of Tomahawk 6.

And I would say we're even planning Tomahawk 7 and 8 right now and we're speeding up the rate of development. And it's all largely for that few guys, by the way. So we're making a lot of investment for very few customers hopefully with very large served available markets. That's -- if nothing else, that's the big bet we are placing.

Operator

One moment for our next question, and that will come from the line of Timothy Arcuri with UBS.

Timothy Arcuri

Hock, in the past, you have mentioned XPU units growing from about 2 million last year to about 7 million you said in the 2027, 2028 time frame. My question is, do these 4 new customers, do they add to that 7 million unit number? I know in the past, you sort of talked about an ASP of 20,000 by then. So those -- the first 3 customers are clearly a subset of that 7 million units. So do these new 4 engagements drive that 7 million higher? Or do they just fill in to get to that 7 million?

Hock Tan

Thanks, Tim, for asking that. To clarify, as I -- I thought I made it clear in my comments. No. The market we are talking about, when you translate the unit is only among the 3 customers we have today. The other 4 we talk about engagement partners. We don't consider that as customers yet, and therefore are not in our served available market.

Timothy Arcuri

Okay. So they would add to that number.

Operator

One moment for our next question, and that will come from the line of C.J. Muse with Cantor Fitzgerald.

Christopher Muse

I guess, Hock, to follow up on your prepared remarks and comments earlier around optimization with your best hardware and hyperscalers with their great software. I'm curious how you're expanding your portfolio now to 6 mega scale kind of frontier models will enable you to and won't blush, share tremendous information but at the same time, a world where these 6 truly want to differentiate. So obviously, the goal for all of these players is exaflops per second per dollar of CapEx per watt. And I guess, to what degree are you aiding them in this efforts? And where does maybe the Chinese wall kind of start where they want to differentiate and not share with you kind of some of the work that you're doing?

Hock Tan

We only provide very base basic fundamental technology in semiconductors to enable these guys to

use what we have and optimize it to their own particular models and algorithms that relate to those models. That's it. That's all we do. So that's the level of -- a lot of that optimization we do for each of them. And as I mentioned earlier, there are maybe 5 degrees of freedom that we do. And we play with that. And so even if there is 5 degrees of freedom, there's only so much we can do at that point. But it is, and how they -- and basically how we optimize it, it's all tied to the partner telling us how they want to do it. So there's always so much we also have visibility on.

But it's what we do now is what the XPU model is, share optimization, translating to performance but also power, that's very important, how they play it. It's not just cost, though power translates into total cost of ownership eventually. It's how design it empower and how we balance it in terms of the size of the cluster and whether they use it for training, pre-training, post-training, inference, test time scaling, all of them have their own characteristics. And that's the advantage of doing that XPU and working closely with them to create that stuff.

Now as far as your question on a China and all that, frankly, I don't have any opinion on that at all. To us, it's a technical game.

Operator

One moment for our next question, and that will come from the line of Christopher Rolland with Susquehanna.

Christopher Rolland

And this one's maybe for Hock and for Kirsten. I'd love to know just because you have kind of the complete connectivity portfolio how you see new greenfield scale-up opportunities playing out here between could be optical or copper or really anything and what additive this could be for your company?

And then, Kirsten, I think OpEx is up. Maybe just talk about where those OpEx dollars are going towards within the AI opportunity and whether they relate.

Hock Tan

Your question is very broad reaching in our portfolio. Yes, we have the advantage and a lot of the hyperscale customers we deal with, they are talking about a lot of expansion. But it's almost all greenfield, less so brownfield. It's very greenfield. It's all expansion, and it all tends to be next generation that we do it, which is very exciting. So the opportunity is very, very high. And we deploying -- I mean, we are both -- we can do it in copper. But while we see a lot of opportunity from is when you connect -- provide the networking connectivity through optical.

So there are a lot of active elements, including either multimode lasers, which are called VCSELs or edge-emitting lasers for basically single mode, and we do both. So there's a lot of opportunity just in scale up versus scale out. We used to do, we still do a lot of other protocols beyond Ethernet to consider PC Express, where we are on the leading edge of the PC Express. And the architecture on networking, switching, so to speak, we offer both. One is a very intelligent switch, which is like Jericho

family with a dumb NIC or a very smart NIC with a down switch, which is the Tomahawk. We offer both architectures as well.

So yes, we have a lot of opportunities from it. All things said and done, all this nice wide portfolio and all that adds up to probably, as I said in prior quarters, about 20% of our total AI revenue maybe going to 30%. Though last quarter, we hit almost 40% but that's not the norm. I would say typically, all those other portfolio products still add up to a nice decent amount of revenue for us.

But within the sphere of AI, they add up to, I would say, on average, be close to 30% and XPUs, the accelerators is 70%. if that's what you're driving perhaps that gives you some -- shed some light on towards where -- how one matters over the other. But we have a wide range of products in the connectivity, networking side of it. They just add up though, to that 30%.

Kirsten Spears

And then on the R&D front, as I outlined, on a consolidated basis, we spent $1.4 billion in R&D in Q1, and I stated that it would be going up in Q2. Hock clearly outlined in his script, the 2 areas where we're focusing on. Now I would tell you, as a company, we focus on R&D across all of our product lines so that we can stay competitive with next-generation product offerings. But he did lay out that we were focusing on taping out the industry's first 2-nanometer AI XPU packaged in 3D. That was one in the script, and that's an area that we're focusing on.

And then he mentioned that we've doubled the radix capacity of the existing Tomahawk size to enable our AI customers to scale up on Ethernet towards the 1 million XPUs. So I mean that's a huge focus of the company.

Operator

And one moment for our next question, and that comes from the line of Vijay Rakesh with Mizuho.

Vijay Rakesh

Just a quick question on the networking side. Just wondering how much it was up sequentially on the AI side? And any thoughts on M&A going forward, seeing a lot of headlines around the Intel products, et cetera?

Hock Tan

Okay. On the networking side, as I indicated, Q1 showed a bit of a surge but I don't expect that to be -- that mix of 60-40, 60% is compute and 40% networking to be something that is normal. I think the norm is closer to 70-30, maybe at best, 30%. And so who knows what Q2 is, we kind of see Q2 as continuing but that's just, in my mind, a temporary blip. The norm will be 70-30. And if you take it across a period of time like 6 months, a year, to answer your question.

M&A, no, I'm too busy. We're too busy doing AI and VMware at this point. We're not thinking of it at this point.

Operator

Thank you. That's all the time we have for our question-and-answer session. I would now like to turn the call back over to Ji Yoo for any closing remarks.

Ji Yoo

Thank you, Sheri. Broadcom currently plans to report its earnings for the second quarter of fiscal year 2025 after close of market on Thursday, June 5, 2025. A public webcast of Broadcom's earnings conference call will follow at 2:00 p.m. Pacific.

That will conclude our earnings call today. Thank you all for joining. Sheri, you may end the call.

Operator

Thank you. Ladies and gentlemen, thank you for participating. This concludes today's program. You may now disconnect.