Event ID: 138103384223
Event Name: Alphawave IP Group PLC Capital Markets Day
Event Date: 2024-06-04T07:30:00 UTC

C: Jose Cano;Alphawave IP Group PLC;Global Head of Investor Relations
C: Tony Pialis;Alphawave IP Group PLC;President, Chief Executive Officer, Executive Director
C: Tony Carusone;Alphawave IP Group PLC;Chief Technology Officer
C: Jonathan Rogers;Alphawave IP Group PLC;Senior Vice President - Engineering
C: Charlie Roach;Alphawave IP Group PLC;Chief Revenue Officer
C: Mohit Gupta;Alphawave IP Group PLC;Senior Vice President, General Manager - Custom Silicon
C: Babak Samimi;Alphawave IP Group PLC;Senior Vice President, General Manager - Connectivity Products
C: Rahul Mathur;Alphawave IP Group PLC;Chief Financial Officer
P: Jonathan Menon;Jefferies LLC;Analyst
P: Unidentified Participant;;
P: Rob Sanders;Deutsche Bank;Analyst
P: Simon Coles;Barclays Bank PLC;Analyst
P: Harvey Robinson;Tammy Aliberum;Analyst

+++ presentation

Jose Cano^ Good morning everyone, and welcome and thank you for making the time for joining us today. As you can see on the agenda for today we have a busy morning ahead with presentations from the management team. There will be a break just before 10 hopefully and there'll be plenty of time for questions at the end of all the presentations.

And before I hand over to Tony, our CEO, I just need to cover three items. First of all I'm not aware there's any fire drills today so if the alarm goes off we're going to have to exit the building. Second, a recording of the event will be available on our website as usual. And finally I need to draw your attention to the safe harbor statement as today's presentations and the answers to your questions may contain forward-looking statements.

And with that I'll hand over to Tony.

Tony Pialis^ Thank you Jose. Hello everyone, great to be here with you. I'm Tony Pialis, Co-founder CEO of Alphawave SEMI. I am exceedingly excited to be here with you today to walk you through the amazing accomplishments that we've achieved over the last 12 months. So given that, let's dive right in. Alphawave is the world's leader in delivering high performance connectivity and compute.

Our solutions power the industry in terms of AI and data center infrastructure. How do we do it? Well we deliver Silicon IP, Custom Silicon, Connectivity Products, and today I'm very excited to announce Chiplets that will power the next generation of AI.

Before I go into this, let me walk you through a little bit of our history. So we were founded in 2017 as a

pure place Silicon IP company, very much like ARM, we took the company public in 2021. We added on a custom silicon and a connectivity products organization. And here we stand today in 2024 at the cost of delivering a vertically integrated semiconductor company that's going to power the next generation of AI.

Let's look at a little bit at our KPIs over the years. I am exceedingly proud of the organization in their steadied and predictable delivery over the last seven years. We've grown our end customers up to a 100. We have grown our employee base to 830 at the end of last year and are on a trajectory to exceed 900 this year.

Our cumulative bookings is on pace to exceed over $1 billion by Q2 of this year. Our revenue, we finished last year at $322 million, and I expect this year to land anywhere in the range of $355 million to $365 million. We are this year delivering on the vertically integrated business model that we spoke of at our last capital markets day by delivering both Silicon IP, Custom Silicon and now connectivity products.

So how do we win in the industry? We deliver the highest performance connectivity that is powering AI and data center infrastructure. Our custom silicon business has been transformed to a higher margin business that is now incorporating our silicon IP that is leveraging leading compute from major players such as ARM, to power the next generation of AI products.

Our connectivity products are being used in the fastest Opto-Electronics that are building our next generation AI fabrics. And now while the industry is talking about chiplets, we've actually been doing something about it. I am excited to announce today that we will be deploying the world's first portfolio of scalable, interchangeable, and customizable compute and connectivity chiplets.

These chiplets will be leveraging ARM's latest Neoverse cores for compute and will be powered by Alphawave's leading connectivity silicon IP. They will accelerate the deployment of next generation AI and will help propagate Moore's Law into the future.

Let's look at how data has evolved over the years. So 2017 was a key inflection point for the world. That's where the machines took over in terms of producing more data than humanity. Since then we've entered an age of exponential data growth. What does this mean? It means from 2013 to -- 2023 to 2025, the amount of data in this world will be doubling. How do we manage this much data?

The only way forward is to deploy tremendous amounts of compute and leverage artificial intelligence to analyze and manage this data and create reason out of all of this chaos. There's no better example of how compute and AI working together can manage significant amounts of data than chat GPT.

Over a span of six months, chat GPT scaled from version 3.5 to 4.0 by training on over 2 trillion parameters. Leveraging huge amounts of additional hardware, it was able to improve its LSAT score from the 37 percentile to the 88 percentile. This just illustrates how powerful AI can be when coupled with significant amounts of data.

Next let's look at the difficulty that hyperscalers face when they deploy on scale AI. There is no doubt what NVIDIA has accomplished is absolutely remarkable. They have grown and dominated the AI

industry since 2015. However, what the chart on the right illustrates is that there is no one single piece of hardware that's optimal for all of our large language models. Some language models need intensive compute. Others need higher memory bandwidth, others require deep memory capacity, there is no one size that fits all.

What's even more surprising than all this talk about compute is that the real bottleneck for deploying AI at scale is connectivity. Connectivity is right now choking the amount of processing that can happen in our data centers. What you see on the left is for various large language models, more than 40% of the time that data sits within our data centers is being wasted as data is being moved from processor to processor across our networks.

So how do we solve this? How do we scale the next generation of AI? It's obviously going to be a combination of deploying customizable compute targeted to specific language models as well as including blazingly fast connectivity to avoid the wastage on our networks.

Now, as we continue to invest into GPUs to power our networks, we've hit a wall. Fundamentally, we are unable to manufacture silicon larger than we are today. So how do we increase the amount of compute if we cannot make our chips any bigger? And then you couple that with the fact that it costs tens of thousands of dollars per GPU.

So hyperscalers like Microsoft and Google are investing anywhere between $4 billion to $10 billion just in GPUs to power their data centers today. How do we scale this forward and deploy AI at a greater impact and across more data centers? It's going to need a combination of higher performance compute. We need to innovate our architectures to be able to deliver compute that fits within existing silicon manufacturing capabilities.

We also need to tremendously lower the cost of deploying this compute and do so in a thoughtful manner that will reduce our carbon footprint and the impact of AI on our environment. So how do we get to a broader deployment of AI? Let's look at the evolution over the years. Back in 2010, AI was software that ran on server grade processors like Intel and AMD deployed. Relatively small language models were used. By 2017, NVIDIA had introduced the GPU to the AI industry.

The GPU is able to run a tremendous amount of processes in parallel, which allowed us to use much larger -- much more significant and relevant language models that we're enjoying today with chat and with Google's Gemini. But here we are now in 2024 and a new ARMs race is underway where hyperscalers are now building custom silicon that is optimized for their specific data sets, for their language models. And now as we move forward, these hyperscalers are developing novel architectures that are optimizing each joule of energy and focused on delivering optimum performance for it.

So stepping back, how do we in Alphawave deliver and win with these hyperscalers? It's by delivering leading-edge connectivity in silicon IP. That's standards like PCIe Express Gen 7, CXL 3.0, delivering the world's fastest 224 gig service and deploying UCIe and HBM, which are chiplet connectivity interfaces. Our custom silicon business is powered by both our winning silicon IP portfolio and in partnership with ARM to deliver the world's most advanced custom silicon devices, leveraging 2.5D and 3D packaging.

Also, with the announcement of our chiplet portfolio that leverages ARM's latest Neoverse cores for compute and our silicon IP on connectivity, we will accelerate the deployment of AI, increasing the cadence that we can innovate as we power our data centers. And finally, our connectivity products are being used in the world's fastest Opto-Electronics applications, building the next generation AI fabrics.

So with all of these great advancements in the industry, we're still limited by the size that a single piece of silicon can be manufactured. So how do we continue to scale the amount of compute? How do we deploy next generation devices?

The industry knows it's through an architectural change called chiplets. And what we've done at Alphawave is assemble a portfolio of compute and connectivity chiplets that leverage ARM's Neoverse cores and our leading edge silicon IP portfolio to be able to create a scalable, interchangeable, and customizable suite of building blocks that hyperscalers, that semiconductor companies can use moving forward in order to deploy next generation devices.

And what impact does this have on the industry and to us as a company? Well, when I was here 12 months ago, since then, our total addressable market has more than doubled. We are now looking at almost a $40 billion addressable market by 2027. What has given rise to that?

It's twofold. The first is customizable AI silicon. There's been an explosion via hyperscalers, via nations, all developing their own custom silicon optimized for their own datasets. The second is the birth of the chiplet revolution. And so when you look at the right, this year, our revenue is going to be dominated by our existing licensed NRE and custom silicon business.

However, over time, what you will see is our legacy silicon will contribute less. It will be replaced by the newer design wins that we've secured in 2023 and beyond, delivering higher margins and higher revenue scale, because that is focused on the AI and data infrastructure industry. You couple in our Opto-electronics connectivity products that enter production this year and our chiplet portfolio, we are on a trajectory to meet and exceed our $1 billion revenue guidance by 2027 with the scale and the capacity that will drive continued growth beyond this decade.

In addition to investing into our product portfolio, we've also been investing into our leadership and governance. Over the last 12 months, we've added seasoned executives like Rahul Mathur, our Chief Financial Officer, and Charlie Roach, our Chief Revenue Officer. Both of whom have led and been executives in publicly traded semiconductor companies. Their pedigree has allowed them to hit the ground running and immediately contribute to our business.

In addition, we've also began a realignment in our corporate governance to better focus our leadership on our core markets of AI and data center. In addition, we've also been working on strengthening our capital structure and strengthening our balance sheet. By simplifying our debt, it is allowing us more operational flexibility to better grow the company and meet the requirements of our customer base.

Finally, we've also been focused on our WiseWave equity. We are in the process of selling that equity. We expect that to happen this year. Obviously, this would also further strengthen our balance sheet. Let me leave you with some key takeaways. Alphawave is situated in the heart of this AI revolution. How did

we get there?

It's by delivering the world's highest performance connectivity IPs, by productizing them in custom silicon and the world's fastest connectivity products, and now deploying them in this suite of chiplets that will only further accelerate Moore's Law and increase the deployment of AI across our data centers.

We're accomplishing this through our disciplined investment into R&D, which will help deliver high margin business and high margin products for our shareholders. Ultimately, we are here to help drive this AI revolution.

With that, I'll pass it to our Chief Technical Officer, Tony Carusone, to walk you through some megatrends happening in the AI and connectivity space. Tony?

Tony Carusone^ I'm Tony Carusone. I've been doing research and development on silicon solutions for high speed connectivity for over 20 years as an academic, publishing over 100 papers on the topic and publishing two of the bestselling textbooks of all time on integrated circuit design, and an industry working with small companies and some of the world's largest technology companies on connectivity solutions, and for the last 2.5 years as Chief Technology Officer at Alphawave SEMI.

I'm sure over the last year or two you've heard a lot about how AI is transforming digital infrastructure in the data centers, but what I want to do is paint you a picture of Alphawave role in that transformation. Traditional data center networks are scaling in part in response to increasing traffic due to AI applications. We see bandwidth demands increasing throughout.

These are the front-end networks that interconnect traditional compute servers in data centers. The traffic that they carry is relatively short-lived and dynamic in response to incoming requests from the outside network. As a result, networks have evolved to have multiple layers of hierarchy and to combine optical and electrical links to allow for flexible and redundant networking to carry the traffic in any different way, switching packet by packet.

So the evolution of these front-end networks is now accelerating and diversifying with AI deployment in data centers. So by that -- what I mean is that we're seeing a proliferation of optical connectivity, in some cases displacing electrical links over some shorter distances, and we're also seeing an increasing emphasis on low power consumption because we're reaching some thermal limits in the data center.

So whereas the bandwidth demands on those traditional front-end networks are continuing to grow, we're seeing even more growth being driven by the deployment of new AI hardware. So clusters of AI acceleration processors sit inside the traditional front-end networks as pictured here. So behind some of the servers are clusters of AI acceleration processors and these are interconnected to each other using a separate back-end ML network.

That back-end network is intended to allow those AI acceleration processors to be able to access each other's memory seamlessly and so they can work together on a large AI training or inference task sharing the same set of parameters in a shared memory pool. So that traffic on that back-end network is absolutely critical and therefore it's isolated from the vagaries of traffic on the front-end network which

are subject to all the incoming requests and ups and downs of front-end traffic.

On the back-end network, low latency and high speed are absolutely critical. So even if one or a few links are subject to lost packets or higher latency, that can bottleneck the entire AI process. It can really lead to a lot of really expensive AI hardware being sitting there underutilized.

So as a result, these back-end networks are evolving separately with unique requirements and because of the very rapid growth there, they're increasingly driving the industry's R&D agenda. So these AI clusters are scaling in the new AI ARMs race from hundreds to thousands to tens of thousands and even beyond AI acceleration processors working in parallel all on the same tasks. Such scaling is happening in two dimensions. We talk about scale up, scale up refers to more and more resources being added to a single compute server. So in this case what that means is more and more AI acceleration processors being added to a single server.

At the same time you've got scale out. That's a way to add more and more machines in parallel working on the same job interconnected by these back-end networks and in some cases through the front-end networks. So we've got scaling along these two dimensions to much larger AI clusters and at the same time each individual AI acceleration processor is becoming more and more powerful through the use of chiplets which I'll talk a little bit more about later. So you've got this sort of tsunami of influences causing a real explosive growth in the connectivity demands on these networks and that's why you see them increasingly driving our R&D agenda.

So to understand where Alphawave sits inside all this AI hardware, we're going to take a little look inside these compute servers. The traditional compute server actually has a lot of the same parts as your PC at home or gaming PC. There is a CPU there. There may be multiple CPUs because this is a server after all and that's connected to a network interface that allows it to talk to other servers in the network. It's also connected to its own storage.

So typically those are PCIe express links as the name suggests. It's for peripheral component interconnect or PCIe and then of course there's a bunch of memory there. So that's a traditional compute server. What's different with an AI compute node? Still a lot of the same parts. You've still got a CPU there that's orchestrating all the work that's being done that's communicating to other servers in the network via front-end network interface and that pulls memory from bulk storage and -- it pulls data from bulk storage in its own memory cache.

And all those links to the front-end network and local storage are carried traditionally over PCIe express links. But now you've added behind that a large array of GPUs, graphics processing units, although of course these aren't really graphics processors anymore. These are very large processors that are tailored to the computations required for AI training and inference.

So in some cases these are referred to as TPUs for tensor processing units or NPUs for neural processing units. So we're just going to often refer to them as XPUs in general, but the bottom line is these are AI, massive AI accelerator systems in package increasingly built out of multiple chiplets.

So those XPUs there have their own memory stores that are storing all the parameters required for the

AI computations and are being accessed frequently. And those XPUs talk to each other through the back-end network that we talked about. So there's a separate back-end network interface there. All that traffic that flows over the back-end network has to do so seamlessly and continuously and so again we need very robust reliable links providing very high bandwidth low latency connectivity.

Now enabling all this AI hardware is a suite of AI connectivity IP that Alphawave is delivering to our partners. So on the front-end network that's still the purview of Ethernet. That's the technology that Alphawave was built on. We've been industry leaders in from 100 gig and now moving to 200 gigabit per second per lane. That front-end network interface is connected back to the CPUs shown here via PCI Express.

Again Alphawave has built its own in-house industry-leading PCI Express IP. Similar interfaces are connecting to the bulk storage. These can be local within the AI compute server or increasingly people are considering disaggregating them into centralized pools which further relies on high performance connectivity. Those links out to storage are in some cases being considered for using CXL instead of PCI Express. CXL is really an evolved form of PCIe that relies on the same physical layer IP that just ensures lower latency connectivity there.

Now moving over to the right you've got the XPUs that communicate to each other across large racks via the back-end network. So there's a separate back-end network interface there. That back-end network interface is typically either Ethernet or PCI Express based often with some customizations to ensure certain performance like lower latency. So for example in NVIDIA systems that's where you see InfiniBand.

The XPUs meanwhile are connected to each other densely within each compute node and back to the CPU using some internal links that are typically customized but based on the same underlying PCI Express or Ethernet physical layer IP. So that's where for example in NVIDIA systems you see NVLink play a role.

And finally the XPUs each have their own local stores of memory for storing the AI parameters. Those are in HBM memory stacks and are communicated with via HBM interfaces. Again Alphawave has the in-house developed HBM interface IP. So in summary Alphawave has the full suite of AI connectivity IP that's required to build these systems. Ethernet, PCI Express, CXL, HBM interfaces.

But there's one more component to the AI connectivity suite I haven't talked about yet and that's the dense die-to-die interfaces that are required to power chiplet based systems in package. So what I'm showing here is a complex system in package that includes multiple large dies of logic that are realized in the most advanced CMOS technologies 3 nanometer or 2 nanometer CMOS allowing them to integrate as much logic and SRAM as possible all within a single package. And these are surrounded by IO chiplets in this picture. The IO chiplets house all the connectivity IP that's required to communicate in and out of the package. For example using PCI Express or Ethernet.

Now what's shown here is representative of a large switch system in package but if this were an XPU you'd be replacing some of those IO chiplets with HBM memory stacks. And all of these chiplets are interfaced together via die-to-die interface IP that's Alphawave's got in-house. Principally UCI Express,

Universal Component Interconnect, Express Links and also in the case of memory stacks HBM memory interfaces.

So this kind of chiplet based design paradigm is absolutely essential to allowing for AI hardware to scale at the cadence that's now being demanded for the rollout -- massive rollout of AI during this AI ARMs race. So by separating the development of those large logic chips in the most advanced CMOS technologies from the IO chiplets that allows each to progress separately on its own timeline and for those solutions to be pre- developed and pre-validated independently.

It also allows people to compose variants of these systems and packages using the same pre-validated components. So combining different amounts of custom compute connectivity and memory to create these variant systems and package quickly and on a rapid cadence. So in summary Alphawave has that die-to-die connectivity IP to enable this chiplet based design paradigm and we've invested as Tony said in our own suite of chiplets to allow for this rapid evolution of AI hardware.

So it's our IP and our chiplets that are powering custom silicon throughout the data center. On the compute side we're very excited about our partnership with ARM. I'm sure you'll hear more about that after the break. ARM's processors are throughout all the hyperscalers and NVIDIA systems and so the combination of ARM processors together with our AI connectivity suite of IP is a really powerful recipe for allowing hyperscalers to develop their own custom compute processors.

But it's not just compute that's requiring custom silicon solutions in the modern data center. It's also on the networking side with a rapid scale out of AI driving such intense bandwidth increases. Cloud providers aren't sitting around waiting for standards-based solutions to emerge. They're reaching out throughout the ecosystem to look for solutions that can provide differentiation in their offerings on the networking side. And so here again it's our complete suite of AI connectivity IP that's available in all the advanced technology nodes and the major foundries together with our chiplet portfolio that's really differentiating.

And John Rogers is going to talk soon about how we're building out this 100s of IP portfolio with our R&D team. So our connectivity solutions are really battle-hardened and have been rolled out in volume deployment. This is absolutely key because a lot of the toughest challenges in rolling out connectivity they really only arise when you start rolling these things out in volume. And our connectivity solutions were architected from the beginning to allow them to evolve and adapt to emerging new challenges. And they've proven themselves able to be up to those challenges as they've been rolled out again and again.

Now later on after the break you'll hear from Babak about how he's incorporating these IPs to be able to develop our own connectivity products rapidly and reliably. And finally we're also seeing an evolution of the geographic distribution of data centers. Data centers are evolving to provide responsive AI processing closer to the endpoints, closer to the end users. And this is leading to the creation of 100s of new regional but smaller data centers.

Now this type of architecture relies on a spine of very high bandwidth connectivity and that's why we've invested heavily and positioned ourselves as leaders in coherent light technology. That's going to allow

us to bring the benefits of coherent optical connectivity to this new geographically distributed data center architecture.

So the industry is not standing still, and neither are we. We're investing in all the key technologies that are going to allow us to keep our edge going forward. Alphawave's been involved in advanced package designs practically since its inception because you have to be really to make sure these high performance connectivity solutions can work.

So we welcome the arrival of 3D silicon technologies that are really going to allow AI compute to scale beyond Moore's Law. And while we're busily working away developing products based on our latest 224 gigabit per second connectivity IP, we're also looking beyond that together with our partners to think about what's coming next because we know with the pace at which AI is developing it won't be long before people are demanding the next generation of connectivity beyond 200 gig.

And we strongly believe that coherent light is going to be a big part of that next generation connectivity picture. Coherent connectivity is really the only path that provides for ongoing bandwidth scaling potential within the data center. So you're going to see coherent light increasingly play a role inside the data center and that's why we've positioned ourselves to be leaders there as well.

So in summary we're focusing on our areas of expertise and expanding our technology leadership. We've developed a full suite of AI connectivity IP to enable custom silicon for our partners that includes PCI Express, Ethernet, and die-to-die interface IP like HBM and UCI Express, all industry leading.

We're also developing our portfolio of Alphawave chiplets to allow for scaling of AI in composable and compute -- computing networking chips. And we've invested in optics technologies that are there to meet the emerging demands of the new data center architectures.

And with that I'll turn it over to John Rogers who's a Co-founder and Senior VP of engineering to talk about R&D at Alphawave.

Jonathan Rogers^ Hi my name is John Rogers. I look after the R&D team that fuels our data center and connectivity businesses. We're a global team, we have research and development concentrations in Canada, India, and Israel. Our differentiated technical focus is really around analog mixed signal design and digital signal processing. And one of the strengths of the locations that we're in from a research and development standpoint is that each of these locations has a long pedigree and a deep talent pool of people with decades of experience in these technical sub specialities.

At Alphawave around 90% of our employees are engaged in research and development. So that's around 750 individuals. So we really are a research and development company in bringing all these new products to market. We are focused on high speed connectivity IP from a research and development standpoint. Data rates up to 224 gigabits per second. The vast bulk of our current research and development is occurring in the absolute leading edge processes.

So most of our work this year is in three nanometer and in the two nanometer gig all around processes. And we have a robust production track record that we're extremely proud of with our connectivity IPs

with products in production at very high scales of integration, hundreds of transceivers per die in seven nanometer, six nanometer, five nanometer, four nanometer.

The R&D team is closely partnered with our custom silicon business. We provide that full suite of proven differentiated IP subsystems for the AI and connectivity compute markets. As well, the Alphawave research and development team spends a lot of time on the type of problems that are involved in chiplet construction. So we've been working in 2.5D and 3D packaging around HBM and UCIe for years now. And our UCIe leaders pushing data rates up to 32 gigabits per second for this die-to-die connectivity.

Additionally, for our connectivity products, we are further enhancing our connectivity IPs with features around directly driving optics from advanced CMOS chips. And so with the very high voltage swings, for instance, that are required to directly drive lasers. And then we've also put in place a portfolio of what we term WidEye DSP features, which help us be very tolerant to interference and really allow us to have market-leading performance in these connectivity products.

IP at Alphawave is evolving. We continue to be focused on transceiver IP in the Ethernet, PCI Express, UCIe, and HBM spaces. But since we've talked a year or so ago, we've really been focused hard on extending that to pre-validated subsystems. And so for each of these transceiver types, we now have an Alphawave controller and are pre-validating a subsystem that includes both the transceiver and controller in Ethernet, PCI Express, CXL, UCIe, and HBM.

And so that allows our customers to take a fully pre-validated plug and play connectivity subsystem and slot it straight into their chips. But beyond even that, we're evolving IP to the chiplet level where we're saying not only pre-validating a single subsystem, but we're validating suites of subsystems that go together in these chiplet architectures. So for instance, working to make sure that our PCI Express plus controller and UCIe plus controller subsystems work together seamlessly to be able to transport data on these chiplets. And the same with Ethernet.

The R&D team continues to focus on differentiated connectivity IP. We deliver faster speeds by pushing data rate, lower power per bit transmitted, and lower costs by keeping the silicon size of our IPs as dense as possible. We are incredibly excited about our 224 gigabits per second AthenaCORE SerDes with its suite of WidEye DSP features, as well as our 24 gigabits per second s AresCORE UCIe transceiver with its low latency streaming controller that helps move data with minimal latency from one chip to another.

And in summary, just the R&D team at Alphawave is incredibly excited to be delivering this whole suite of intellectual property in service to all three of our businesses.

Thank you. All right, now we've got Charlie up next, who is our Chief Revenue Officer.

Jose Cano^ Thanks, John.

Charlie Roach^ And good morning, everybody. I'm Charlie Roach. I'm Chief Revenue Officer at Alphawave, and I look forward to meeting everybody during the breaks and afterwards. So what I'd like to talk to you today about is Alphawave's transformation into a connectivity company and our path to billions in revenue and billions of dollars of value. So you're going to say, why should I listen to Charlie,

right?

So I've done this before. So I ran revenue at [EnFi] for about a decade. And in that time, about a decade ago, we saw the connectivity decisions and strategies move from the very important still system vendors, such as Nokia, Sienna for coherent long haul, Cisco and Juniper for those switch chips, HP and Dell creating those server.

So what Tony CC showed you in that data center, they were making those decisions. About a decade ago, the hyperscalers said, I'm going to take that and I'm going to own that connectivity piece. I'm going to drive my own roadmaps because I have specific issues that generic solutions are not going to fix. And so what we did is in about 2016, we said we're going all in on connectivity. We sold our products that were not associated with connectivity. We came out with our own PAM4 DSP, like John Rogers just showed you all that great technology. We acquired a coherent company. Sounds familiar, doesn't it?

Alphawave acquired Banias. And then the hyperscalers said, hey, we want your IP in those connectivity products and the switches and NICs that are being developed in the ecosystem. So they forced us to become a Silicon IP company -- said we want you to sell your IP to our ecosystem vendors. So we have similar technology and DSPs running through our networks.

And then the last phase is we acquired a custom Silicon company to customize solutions for the different hyperscalers. And in the end, that connectivity technology had a value of over $10 billion to Marvell. And today, those technology products are worth over $25 billion in market evaluation. So what took EnFi over five years, Alphawave's done in less than two, is they've created a very differentiated technology of connectivity and added chiplets, which is new.

We looked forward, backwards. Now we're looking forward. The chiplet technology, a new connectivity solution that we're developing, and is really the future of die to die, then chip to chip. So how do we win?

I just talked about some of our competitors. There's some very big companies out there, Marvell, Broadcom, Cadence, Synopsys are our competitors in different areas. Well, you just heard me talk. Companies want someone who can communicate die to die. So that's millimeters apart, meters apart, rack to rack, and data center to data center. They prefer to really leverage a company.

So think about taking a flight from London to LA, you stop off in JFK, one year on Delta, one year on Lufthansa, your luggage gets lost. Good luck finding that. So they like to have consistency in who's transporting their data. Our competition really doesn't do this. They either are an IP company, or they're focused on silicon.

Alphawave, we offer great flexibility for our customers, both in technology and in doing business with us. We're very flexible. That's why they want to do business with us, and that's why we're winning at hyperscalers. So again, die to die, rack to rack, data center to data center. We solve all those issues as one connectivity company.

So entry points into the hyperscalers are not all the same. So some hyperscalers, they start engaging

with us on silicon IP. Let's say they're doing their own XPU processor. Every hyperscaler is talking today about their own development of their inference machines, AI machines. They engage with us on that silicon IP. Right?

Another hyperscaler might engage us on the connectivity products and say -- hey, we like what you're doing, the connectivity products, we like what your coherent technology allows, let's engage there. Well, you see where this goes. They engage us in one, and then they go you're doing so well here, how about your IP?

You have chiplets? Similar technology? You know, John Rogers has created this core silicon IP technology. That's great, let's use you for that. So we have four different entry points into these hyperscalers that pull each other as they connect and start developing different parts of their architecture from AI to their data center to their new inference NPU. So those are the entry points we have.

Some of these are direct, as I mentioned, maybe the hyperscalers building their own AI, they have a unique connectivity interface they want to have, but we work very closely with the ecosystem, right? You don't see the hyperscalers really building their own optical modules, and we're not going to compete with our partners there as well. Just toured Asia, first time with Alphawave, and met with a lot of the people and executives in Asia who lead in the optical and cable markets, and very excited to have us entering that market.

We just talked about our partnership with ARM on the Neoverse when our chiplets technology, we work very closely with the fabs, we're on bleeding edge technology. When John Rogers helps them develop that two nanometer by putting cutting edge connectivity, they get to test out their processes based on our technology, it's a great relationship.

As I mentioned, we've got lots of partners that are doing AI specific inferencing machines, different XPUs and connectivities. We really work with them to create specific products for the hyperscalers. Why does each hyperscaler have something different?

Why can't Alphawave make a product and sell it to all the hyperscalers? It's because they're actually solving different problems. Amazon Web Services, they invented the data center that everyone can use. They're the leader in e-commerce. They've got a very different problem set than Meta. They lead social media. Their problems are different. What Tony CC showed you was a generic architecture for a data center for an AI cluster.

Each hyperscaler actually has a little different problem. They actually have a different number of connectivities. They have a different bandwidth requirements in each step of that network. Microsoft, they serve as business. We were all editing this presentation together on our own computers. You could see each other editing it live. You did not have to sit behind everybody's one computer. They've got a very different problem than everyone else. Google, video streaming, search engine, and then NVIDIA, AI, it has its own problems.

Saying I'm doing a product, I'm engaging the hyperscalers is very generic, but what we have to do as a company is come out with base core leading edge technology that we can deploy, but then customize it

for each data center and their own issues.

Think of your iPhone as different from your tablet, as different from your laptop. They all have actually a very similar technology, but they're solving different problems. So this is an example of a Facebook campus. It's built in the middle of nowhere. Land is cheap, and it's near a power source, but they've got kind of a campus look to them.

This is a Google data center, again in the middle of nowhere. Land is cheap, but it's close to a power source, right? I mentioned Microsoft has a very different problem. They put these small data centers in and around metro areas like London and allow you to have low latency when we're editing those slides on different computers, but it looks like it's live to us. When you're updating your Facebook page, you really don't care if it updates in half a second or a quarter a second. When you're updating that Excel spreadsheet with someone else, you do. It matters.

So Microsoft is an example here. They created DCI, okay, data center interconnect, and what it is -- is it used to be very expensive to leave a building using coherent technology. You had a standalone system, very expensive because it was made to go from London to Paris, not go 50 kilometers around London, okay?

So they created an industry standardization around that called ZR, and guess what? Facebook's like, well, that wasn't my problem -- but heck, I'll use it because I have to connect these buildings in my campus. And so Microsoft solved one of my problems. It wasn't my biggest problem, but hey, I'll use it.

So what you'll see is a hyperscaler solve one of their major issues. They get the technology proven in the ecosystem, and then you'll see other hyperscalers deploy it because it's not their biggest problem, but it does create a new value or technology chain for all the hyperscalers.

Okay, so again, why do these hyperscalers -- there's all these hot companies that you're investing in that all these hyperscalers are kicking off, and you guys are seeing a lot of interest, and the reason is they like companies that are very focused on a technology spectrum, that they know they can pick up the phone, call Tony, and he's going to pick up the phone. They know Tony. Tony knows them.

They're a director of engineering at a hyperscaler. They know Tony's going to return their text. They know Babak and Mohit and JR. JR knows them personally, right?

That's what hyperscalers get out of small to medium semiconductor companies. Flexibility in business and responsiveness. That's what you enjoy about engaging with companies as well, medium versus really large conglomerates, and you sit on the phone for 20 minutes waiting for them to pick up.

The history of this. NVIDIA was a small graphics company. Guess what hyperscalers engaged them? They became big, right? AMD stock was less than $2 dollars. They started focusing on hyperscalers. Now look at their market cap, right?

Switching companies, entire switching companies came out because they focused on the hyperscalers. Woke Cisco up, and they had to change their strategy. They now sell silicon chips and not entire systems

to the hyperscalers.

Obviously, module makers, EnoLight, Eoptolink, didn't exist a decade ago. They focused on the hyperscalers. The companies like Finisar, now Coherent, Lumentum, they had to change their strategies and compete against them. And then connectivity, right? We've got Credo as another example, and Alphawave. People focused on connectivity and focused on the hyperscalers, and we're going to be that next success story in the marketplace.

So, the path to billions in revenue, path to billions in value. Really, it's about being aggressive. We've done in two years what's been done in almost 10 in the past. We build planes, trains, or automobiles. How far do you want to go? Where do you want to take it, right? Do you need to go from London to Paris? Do you need to go just down the street in London?

We have the right technology solution for you. We have business and product customization for the hyperscalers. If you just want the engine, I'll sell you just the engine. You don't have to buy the whole automobile, or if you want it converted to electric, I'll sell you an electric car. It's the customization that we're providing to the hyperscalers.

Bottom line, we're delivering mission critical connectivity essential for compute in this AI revolution that's going on.

Mohit Gupta^ So, now in the next half of the presentations, I'll kick off with the IP chiplets and custom silicon. So, before we get in there, a little bit about myself. I joined Alphawave as part of the acquisition, which was done in September '22. I have a long history in the IP and the custom silicon business, delivered probably close to 100 chips in the market nowadays, delivered 100 plus million chips in production.

So, we carry that pedigree. I'm excited to be here. I'm seeing some familiar faces from last time, and definitely looking forward to more questions after my presentation. So, I have a few topics to cover. So, my presentation is divided into two sections. We'll talk about the IP first, and then we'll go into the chiplets and the custom silicon, which I think Tony made some exciting announcements on that front.

So, on the IP first, which is the DNA of the company. We have been in this business -- I think, from the start of the company when the company was founded. This is a fundamental technology we carry, which allows us to be a big differentiator in the market. Now, when you go in this new AI era, the silicon IP has become a fundamental to any new AI chip which is being done. So, we heard the term from Tony Carusone about AI XPUs.

So, let's understand a bit of an anatomy, I call it a little bit. I wanted to be a doctor, but never made it. So, of these XPUs, right? So, if you look at any XPU, it has a bunch of components in there. There's CPU cores, which are particularly doing the processing or managing the data. There's system IPs doing a bunch of control stuff, security stuff, memories, fabric, and then custom accelerators, which are primarily the [ALUs] and the Mac engine for a particular workload. And along this is surrounded by a bunch of blue boxes you see in the picture over here, which are the interfaces. And that's what we build at Alphawave.

This is essentially the IP portfolio we carry. So, on the north and south, you see it's written as PCIe, chip-to-chip, or Ethernet. I mean, these are used to connect these XPUs either to another chip or to another host sitting somewhere out there. So, it's a logic-to-logic chip connectivity interfaces. On the east and west, you see two types of interfaces. One is UCIe D2D, which is primarily to expand these XPUs to have more chiplets connected to them to build a bigger silicon solution.

The third you see is the memory interface, which is essentially DDRs or HBMs. Now, you see HBMs all over the map, especially what you need for the AI accelerators. These are primarily to connect these XPUs to a DRAM chip. So, we build all these interfaces today at Alphawave across various technologies.

Now, what is driving this IP market itself, right? So, the first is the workloads. And I think Tony touched in his presentation about all the various hyperscalers. And Charlie also gave a very nice thing about each hyperscaler is solving a different problem. That essentially means they have their customized workload.

So, when you go on a YouTube server, you want the best latency and the buffers that video should flow seamlessly on your screen or the laptops, right? Whereas, when you're doing a search on a Google, you just need the fastest latency for a minute, just that to pop up. It can then go slow down again, right? So, that means AI chip, which is designed for a video codec for a YouTube server can be very different what is being used for a search engine. So, this creates new design starts, which means you need more IPs, right?

Second is scaling. I mean, the process nodes just continue to scale. Why do we do that? You want better power, better power efficiencies, more compact areas, and able to drive overall performance over TCO. So, what this means is we keep driving new IP generations over and over across these process nodes. And some of you, if you have attended TSMC or Samsung events in recent times, or even Intel, they've announced angstroms, right? So, we used to joke, right? We'll soon be able to count electrons in the transistors, but we are there, almost there, right? So, we have to develop IPs in all these leading process nodes.

And last driver is the big driver is the Multi Die Integration, which is basically putting all these chiplets together in large packages. But when you put them together, they're going to talk to each other, right? So, that happens over connectivity IPs. Again, you need more IPs to talk to each other.

So, what that leads you is that this market, which is about $1.5 billion this year, is running at a nice CAGR of 21% to $2.5 billion in '27. And we have a good market hold in this position, and we continue to drive this. I mean, ultimately, it's a high margin business for us, right? Which is the licensing revenue, which comes out of this.

And then, as we have expanded the company into chiplets, custom silicon, and our own products, we are able to leverage these IPs to even generate silicon revenues, which are long term. So, it's a very important component of our business today.

Now, what are we doing over here? We have a full suite of AI silicon IP subsystems, right? What that means, I mean, these are some nice fancy eye diagrams to the investment community may not be that

relevant. But these are for engineers, a lot of blood and sweat goes in to get to these points. I mean, in fact, I brought one of the samples here of one of the test chips, there's like a three nanometer chip, which we did.

So, before we build large products, which I'll talk about later, you have to build these mini test chips to fully prove these IPs on silicon and get them out there, right? And these, I mean, of course, after the event, we can have you guys have a look at these. But these are [gizmos] we build all day long.

And we do about a dozen of test chips a year, working with leading foundries out there to keep scaling to the lower geometries. Ultimately, what do we get by doing this, we are able to reduce the risk for our customers who use our IPs in their systems and overall deliver lower cost on their engineering efforts by putting in there.

So now, if you look at what where is Alphawave on the IP side, I think I call this the team in this room from Alphawave, they've done it all on the IP side. They started at the internet era, which is early 2000. And today, we are sitting in the AI era, we have deployed these IPs across numerous applications, high performance applications, networks, switches, routers, storage devices, base stations, data centers, right? So that way, we have a proven track record of delivering these IPs across generations and generations.

So let me switch over now to about custom silicon and chiplets, how we are using these things to build more advanced silicon and being part of this whole AI era. Now, if you look at the custom silicon, I think, Tony flashed this slide, I think this thing is going to stay relevant for many years. I think if you look at it's always a decade long journey when a shift happens, right?

So the relevance of this custom workloads, driving customized design at hyperscalers, I think it's going to stay for a long period of time, at least till I can imagine working, building these chips, right? So we started with CPUs, became GPUs, massively parallel systems, and then now custom system. Now, NVIDIA is funny, right?

Even at the Computex, they announced the new relevant platform, they still call it a GPU. GPU stands for graphical processing unit, it is not for graphics. They're announcing to the world a new AI accelerator, essentially. And why they happen to be there? They just were there at the right time, right? It was not that I hope so Jensen foresee the AI era coming and they build the products around it earlier than anybody else. The GPUs were best situated to deliver for the AI market at that time, and that's why they have their leg up in the market amongst everyone over there.

Now, when you put that with the market size of this, it's been exploding, actually. It's almost a $32.5 billion market if you combine both chiplets and custom silicon in 2027, right? And there are estimates and this is a conservative estimate we have taken based on the market data and our own data. I mean, we have seen people going all the way up to $40 billion on this market or $40 billion-plus of market. And I think it will continue to grow, again going back to the same thesis of custom workloads. It's pretty much driven by one single phenomena happening in the industry.

Now, how do you win in this market? I think for people who were here last year at CMD, we talked a lot

about connectivity. And Tony announced before the meeting that what we are doing on the compute side, I think it's very essential to have both the components to win in this market. Connectivity gives the interface to talk to the other chips, but you need the ability to be able to have the performance optimized compute systems on those chips, or AI, XPUs as we call it.

And we have done a lot of work there over the last 1.5 year partnering with ARM, optimizing cores, on the PPA side with latest and greatest process nodes, optimizing latency, and very new, fresh out of the oven, being able to build leveraging ARM based newer scores compute chiplets, which will be part of our portfolio.

So when you have all this, like, what does Alphawave do, right? So I call it as a custom silicon design platform standing on five pillars. The first is process technology, right? So you hear these numbers, though gate lengths and these numbers don't match, by the way, anymore. This is -- I feel, a bit of more marketing going on, but leave that aside for a moment.

We are building on the latest and greatest process nodes. So this is different than when we talk about IP, right? So we are shipping in production at seven nanometer. We are ramping in production customer chips in five nanometer and four nanometer. We are about to tape out chips in three nanometer for our lead customers. And two nanometer, what John touched upon, is in the enablement phase, what we are building the IP test chips and the flows over there. So that's the first key component.

Second is leveraging these build IPs. I talked about it, right? You need the full suite of portfolio, 200 gig SerDes, 100 gig SerDes, PCIe, CXL, HBM, UCIe, and provide that for the custom silicon platform. Third is the most important aspect of this. It's a centerpiece, actually, of the entire strategy, which is the ecosystem. I mean, they're gone are the days when they said real men have fabs, right?

So nobody carries fabs anymore. It's actually a burden to carry a fab in a company now. So you have to work within the ecosystem, right? So there, what you get is, the first thing is ARM compute access, right? So to build these chips. So last year in September timeframe, we announced our partnership with ARM. We became part of their total design platform, where we have access to their course.

Next is leading foundries, working specifically with TSMC, Samsung, and Intel design flows. I mean, you need tools and internal engineering flows to develop these chips. And last, very important point over here is Multi-geo OSATs, right? And you hear every day, the geopolitical situation is changing. Customers do ask us, how do you manage supplies, allocations? What if that happens? China, Taiwan, you know, all those kinds of questions comes in. So you got to have a strategy to be able to take a silicon and being able to still supply to our customers, what if scenario occurs, right? So you need a Multi-geo OSATs strategy.

Fourth is the prebuilt chiplets. I think I'll touch upon it a little bit more in my next slides. But essentially, there are three components chiplets provides. Faster time to market, saving your cost, and allowing you to take a lower risk on your new design.

And last is advanced packaging, which actually, if you put it together, right, to build these complex chiplets, when you put them all together in a package, it is becoming more complex. You've got to watch

out for thermals, how power dissipation happens. Is it reliable? Is the mechanical stress can be handled in that chiplet fashion? So that requires a team and expertise, which we have invested on across all these over a period of time. And we have this fully developed platform to go work on any AI experience for our customers.

So a little bit more about ARM. So we started our journey with ARM, like I said, early last year, to be part of their total design platform, where we accessed their newest class scores, both N series and V series. And what we did is we took those scores and created a whole organization, which we call the CPU Technology Center of Excellence, which encompasses taking those ARM cores, do prototyping of those, optimizing them, being able to connect them to our own IPs, like PCIe, CXL, UCIe, Ethernet, and all those controllers, and then harden them to get the best PPA. You many times hear numbers like 3.5 gigahertz at 5 nanometer or 4 gigahertz at 2 nanometer. That's essentially the work we are doing to make sure we can get the maximum juice out of these ARM cores at each and every process node.

Next is chiplet. So I think this is a big thing for us right now. And we announced our industry's first portfolio of interchangeable, customizable, and scalable chiplets. And we put them into four main buckets. The first is IE extended chiplets. These are primarily chiplets where you are decoupling the core and the IOs. You can keep scaling down your main dies to lower geometries but have the IO functionalities being available on the side.

ARM compute chiplet, again, taking ARM newer score being available as a chiplet which can connect to an accelerator. Think of a NVIDIA Blackwell system. So they have the gray CPU talking to a Blackwell chip or a gray CPU talking to a H100 accelerator, right? So that's essentially what we are trying to mimic here.

Third are the memory chiplets. This allows you the flexibility to create SKUs on the parts. You want to build an AI accelerator where you care about bandwidth. So you put in the HBM chiplet in there. If you care about cost, then you put a DDR chiplet in there to connect to the main die.

And lastly, which is going to be a reality in future, is going to be the co-package optics chiplet. So essentially, a lot of work goes into building these. And what we are doing is we have taped out some of these already. So you will see some announcements on these very soon. And we have lead customers engaged with us leveraging this chiplet technology today.

So what does chiplet, how does it work? We talk about chiplets all day long. Chiplets are doing this, that. Let's look at an actual simulation in terms of how does it give you a time to market advantage. So when you start a design, if it was a monolithic, monolithic means a single die, large die you're building, let's say 800 square millimeter, it'll take you anywhere from 12 to 18 months to develop that chip from spec to tape out. And then you have a cycle of production, qualification, and ramping up in production.

What chiplet is able to do is compress that design cycle almost 30% to 50%, right? Because some of those components are pre-built. Because some of those components are pre-built. You don't need to spend your energy on those things. And it lowers the overall risk to be able to get it out faster. So what we are seeing is almost like 30% to 50% savings in that design cycle, leveraging the chiplets.

Second is the risk and cost, right? If you remember some of your old engineering lessons, the bigger the

chip, difficult to yield, right? So what happens is if you build a very large die, the yield will be 20%, 30% or even less sometimes in leading process nodes. Whereas if you bifurcate it into let's say two equal pieces or four equal pieces, you can get a much better yield to connect those chiplets and have gain a overall advantage. So it not only just lowers your risk, but it also lowers your cost because more yield means you can get more dies per wafer, which essentially brings down your ASP.

And lastly, I think Moore's Law is always talked about where it's headed. Yes, there are challenges to keep scaling the process every two years to deliver the performance, but we believe chiplets is the real way to extend Moore's Law, right? It allows you to pack more, leveraging the chiplet technology. So that's the third most important vector a chiplet provides when it comes to the Moore's Law.

So how do we engage as a company with customers, the model in terms of our custom silicon? Now, if you look at it, there are typically we partition it into five simplistic design phases. There is spec, there's architecture, there's logic design or front-end design, physical implementation, which is the actual back-end design, and then the production silicon. I mean, this is where most of the big dollars are there, right, in terms of the cycle.

Now, if you look at a typical old school legacy ASIC business, that resides on the left side of the spectrum, where your ownership is less, but you're still delivering the end product, but you can't command margins because you have not added value to the customer that much, right?

And where we are focused is on the right side of the spectrum, where you do more what we literally call as co-developing with customers. We are part of their front-end design teams. We are working with them on architecture, spec, to define what the chip looks like, and that essentially gets us the margins and being able to drive more value with our customers.

And pretty much, I can say 100% of our AI silicon engagements today are in that space. In fact, I brought one sample over here. I know it's visible to you guys or not, but this is one of the chips we did. It's a pretty large chip, actually. It looks small, but compared to the size, it's a pretty large chip. So there is one large die, there's HBM, there's another IO chiplet over there connected over interposer, right?

So this one is already in production. We are shipping these kind of parts today, right? So it's a pretty complex chip, what we are doing in the market for our customers. And again, all these things are possible using our own technology, leveraging the ecosystem, the five pillars I talked about.

Let's talk about some design wins also, right? So this is what we have today. Already, we are shipping, but there are much bigger things we are building. We are cooking them. So hopefully, next year, when I show up here, I'll bring multiple of these parts, okay? So this is the first accelerator we are working on. It's four dies and eight HBMs, literally 12-chiplet design. It's a pretty complex design. Not every company in the world can touch this design, right?

So we are working on this with our customer. They're providing us their accelerator IP. We are doing everything apart from that on that chip, making sure what is the top-level SOC looks like, packaging, assembly, test. So our delivery to them is this full gray box, package, tested, assembled part, right? So that will be the delivery to our customer. And this is in one of the finer geometries I talked about.

Next one, ARM based stuff. So we talked about our ARM partnership. Actually, we have been working on ARM based stuff for last three years. We just came out with some stuff in the open, but this is the second generation of Neoverse chip we are doing, and again, in a leading process node, where two Neoverse chips are talking to each other with four HBMs connected to each other, right? So it's, again, a big package, almost like 60 by 80 package to manage, and this will be tipping out next year.

Another version of an accelerator we are working on, very different from this guy. It's an accelerator, too, but it's a monolithic die, right? So 800 square millimeters. Now, you can say, what about your argument about triplets over here? Well, this customer has a path. They've built so much redundancy in their chip, they don't care, right? But this tells you the customized workloads people are doing.

Very similar architectures, very similar things, but the way to implement is very different, right? So that chip is also, again, one big large die packed with six HBMs and an interposer, and will be available next year, too, right? So these are some of the big designs we have been working on for the last year or so, and there are many more. I've just picked a sample over here, and this is the future for us. This is the future for us.

Okay, to sum it up in my slides, I think there's a massive opportunity for us, right? So we have almost like a $35 billion opportunity when we combined IP, custom silicon, and chiplets, right? This is almost a 26% CAGR. I mean, if I go by my estimates, my gut feels is next year, we stand here, we'll show you even a bigger number on this, just because this market is still in early phase. Last year, the CAGR and this year has changed already, right? So when we come back, most likely this number will further blow up in the years to come. Just how the whole AI world is being driven.

We are investing early. Building these chips is -- like I said, not for faint hearts, right? We have been investing in process nodes. We have been investing in our IP to prove these things early on, so we can de-risk our customer designs. It's very important in our business.

We already have design wins, so it's not like we're just talking about it. We are fully packed to deliver these, and we're working on this for last year or so, and you will see a lot more to come over here across the worldwide customer base.

And lastly, I truly believe we are well-positioned to deliver for this market, whether it's for any complex AI silicon, could be in a different form factor. It could be an IP, it could be a chiplet, it could be entire custom silicon. We are ready for all, and we are very flexible, like Charlie indicated, right? There are companies who are rigid about the business models, but we are not. So this is one of the key things when we work with our customers. We are able to tune certain things for them to make sure it's a formula for everyone.

Thank you so much. All right, with that, I'll hand it over to Babak.

Babak Samimi^ Good morning, everyone. It's good to see some familiar faces from last year. Some of you, as you might remember, I joined Alphawave in 2022 to help build up our own connectivity products, and it's a multi-billion dollar opportunity ahead of us, really taking the best of the leading technologies

that we have and creating our own opportunities and service to market at a larger scale.

And so the connectivity products falls under my preview. As both Tony's discussed earlier this morning, right, what's happened since last time we met is the generative AI has fundamentally shifted how compute capacity is being added to these data centers. And what's happening now, discontinuity is emerging where the throughput and capacity of these GPUs are outpacing the connectivity bandwidth links.

And so as Tony mentioned, this is becoming almost a bottleneck or choke point in the network. And so it's putting a lot of focus on, if we want to get towards these accelerated computing infrastructures, if we don't deliver on the connectivity bandwidth resources, these expensive resources will be underutilized. And so this is why Alphawave is so focused in enabling the applications here and solving these problems.

And we're doing it on three fronts. On the Silicon IP business, John mentioned, right, that we have gone from the cadence of just meeting the rate speed ups to now his organization is providing a total subsystem solution. The new update now when Mohit covered was the disaggregated and scalable connectivity solutions to the chiplets.

And these two put together help solve the connectivity at the Silicon level. In my group, we're focused on bringing connectivity products together and then solve the system level capacity problem. But things are changing. Our worldview last year and this year is different. And connectivity market has been disrupted.

One area has been a catalyst to accelerate commercialization of 200 gig PAM4 -- very quickly. Our reason is to be able to balance the bandwidth gap between the GPU and the connectivity links, we need to enable 1.60 transceivers, optical modules to rise to the occasion.

The other positive thing is we need heck of a lot more optical ports. And as an industry, we're getting better at estimating the size of this opportunities. But the rule of thumb right now is for every GPU that goes in the network, there's a 3 times pulling on the optical transceivers. And I have three plus. And because the view is as multi-generational capacities increase, optics will permeate more into this back-end network that Tony CC was talking about because the rates go up. Electrically approaching things takes away from the reach. And so optics will help alleviate that.

The other thing that's emerging is this coherent light technology. Because after the disaggregation into these specialized clusters, as Charlie mentioned, the hyperscalers have co-located buildings. And so they need to spray this AI clusters across these buildings. And the reaches are 10s of kilometers. Up to 20 kilometer is an emerging application field that we anticipated and put our money where our mouth is in terms of acquiring that technology. And it's where PAM4 can't get you there in terms of reach and bandwidth. And so as you distribute AI, you have a high bandwidth connection problem, but it needs to be at low latency. And coherent light is optimized to directly address that application.

But we have to do all of this sustainably. And the volume of compute silicon that's going into the network, it's driving the industry. And we have accelerated our plans to deliver three nanometer now and then extend that to two nanometer so we can scale sustainably.

I want to leave you with one piece of learning today from if you take away anything from my presentations. In the last 20 months, we've learned the importance of the GPU as the enabling processor for generative AI. But one thing we cannot forget is there's another processing engine that's required to solve the connectivity problem, and that's the DSP. And this is where Alphawave is investing its dollars to create that connectivity processor and bring that to market.

To be a leader in that space, well, you better own all the tech key assets. And this is the unique element of Alphawave is we have all the pieces in-house. Starting from the right a high-speed connectivity technology that we talked about. We've already delivered 100 gig. We're actively working and quickly bringing to market 224. And then our eyes are on 400 gig. But doing that in successively leading generations, right? So we've done four, we're doing three and two nanometers coming.

The other piece of this market is you need to take the digital world to the optical world, to the optics. And you need to have high-precision analog capabilities in-house. This is the organization that John Rogers had built. As we go to market with our products, John's engineers and talent is part of our value proposition in the marketplace.

Because as we build, put these things together to build that SOC, to build that processor, our engineers at the forefront are working with our customers with the various optics technologies they have, be it AI/ML or this emerging differential AI/ML or in-house silicon photonics that they have to help solve the total system solution for them.

And then our firmer, provide that firmer across the application. But we can't stop there. There's two technology platforms. One shoe doesn't fit all, as Charlie mentioned. And so we have to do this at scale for both PAM4 and Coherent. Some of the other folks that are in the market for connectivity don't have these two assets. And it's critical. It's part of our differentiated offering as we go to this market.

So harmonizing our high-speed, high-precision analog with our configurable DSP engine under the umbrella of the WidEye DSP technology, this is the innovation that's powering our differentiation as we bring our products to market. We have the ability to tune for that specialized feature that Hyperscaler wants on that WidEye DSP platform.

So what are we bringing to market? We are taking those critical technology assets and putting together and products that are application-specific for the target connectivity that Hyperscaler wants. You've never seen this before, so it's the first time. We are bringing to market electrical PAM4 DSPs for active cable connections. We are bringing to market optical DSP PAM4 silicon and Coherent light for the transceiver applications at 400, 800, and 1.6T generation.

Hyperscalers want partners that can solve the broader set of their applications. What we have done quietly for the last two years is position Alphawave to bring that to market. And at the same time, do it in leading-edge technology nodes. And by owning the critical assets, we control our destiny to drive the cadence to keep up with the Hyperscalers. Having technology is not good enough if you're not there to drive the cadence and provide that optimization that they need.

This is our field of play. When we wanted to get into this space, we decided to go after the highest, biggest, end-market application to drive our revenue potential. And that was the connections inside the data center with PAM4 and anticipated Coherent light-based links for these co-located data centers and bringing to market and investing in a DSP dedicated to be a category product leader for Coherent light.

What happened in the last 20 months? Our strategy got validated. PAM4 market is bigger than what we anticipated. We'll take that. And then the disaggregation of the scale-out network and then distributing that across the data centers validated our investment in Coherent light.

And so this is why I'm saying it's very important to have these two pieces in dedicated products because these things will coexist. The consumption model for the Hyperscalers for PAM4 and Coherent will coexist as we enable 1.6T and go to 3.2T and beyond.

So how we're going to deliver on the revenue? Late this year, going into next year, we'll start seeding the revenue base that we've been talking about through our generation one dedicated products across electrical and optical PAM4 and Coherent light. Next phase, we'll scale up that revenue to a multi $100 million business as we finish our generation two investments that are happening now and bringing that to market.

So in summary, the opportunity is huge. The spending on connectivity up to electronics for connectivity is a $4 billion market and now we have all that whole market is accessible to us with the platforms we've developed. This is such a big spend and so critical to the efficiency of the Hyperscalers networks because they're paying for those expensive GPUs. They have to be idolized. CXL got created to solve that problem for memory, right? That have stranded memory.

Because the spend is so big and it's so critical to the efficiency of their network, they're paying extra attention on making sure that share where they're spending their money is more balanced in the supply chain.

And what we've done in the 3.5 and a 2 years is uniquely position Alphawave to be that trusted partner because we're delivering on the platforms that they need. We own the critical technology from analog to SerDes to digital to be able to drive the cadence, control our destiny at the cycle times and cadence they need to be able to earn that business.

And for you guys that were here last year, my takeaway was if we successfully execute to our plan next year this time, which is now, I should have earned our first purchase order and be shipping product, our connectivity products. And I'm happy to say we've done that. And we're seeding this business up, ground up new revenue stream for Alphawave, multi-$100 million. It's supported by a hyperscaler who's trusted us and we're committing to spending north of $300 million with us over successive years buying our coherent and PAM4 connectivity product.

I'll look forward to the questions and discussions later on. Thanks for listening to me. And I would like to invite our handsome CFO to provide his update.

Rahul Mathur^ Good morning, apparently, that's me. My name is Rahul Mathur. I'm the CFO of

Alphawave. I joined the company end of October, so I've been with this great management team for about seven months now. I actually just recently started my career in the semiconductor industry in 1998 at KLA instruments. So spent about eight plus years there, starting in operations and then transitioning into finance.

Over the last dozen years, I've been part of some of the great transformation success stories across our industry. So specifically, I joined Spansion about a dozen years ago. After we've gone through a pretty significant financial change we built couple of billion dollars of value there that led to a merger with Cypress Semiconductor stayed with Cypress and really worked on that transformation as well.

In 2016, I joined Rambus as CFO and spent five years at Rambus and Rambus and did many of the things that we're doing as well in terms of starting with a very strong base of silicon IP and then adding profits and creating a really compelling story much, very similar to what you've heard today.

What I wanted to do over the course of my presentation is really bring all of the different elements our previous speakers brought in into our financials to really give you a better link between our financial performance, what we report and the trends that our management talked about earlier.

What I'd like to start with is the spectacular growth that we've seen in our business. If you look at the chart on the left, this is the bookings trajectory that our company has had. On the charts on the right, it's the revenue trajectory. If you look at the [Kegres] that we've had from a bookings perspective, all told over the last couple of years, it's been about 25%. So kind of on pace or slightly better than the markets, in the market growth that we described.

If you look at the revenue trajectory, it's been 89%. Now that's something that's a combination of both organic and inorganic. One thing I'd like to point out though, is that a significant timing piece is really related to the license agreement that we have at WiseWave. As Tony talked about, that's something that's an opportunity for us from a balance sheet perspective in the coming year though we don't need any incremental inflows from any non-operating activities to fund the business and growth that we talked about earlier.

But if you take out the impact of WiseWave and how it skews our bookings and revenue what you see is a spectacular underlying growth rate and the products and technologies that we're bringing to market today and also into the future.

What you also see on the next slide is the transition that most talked about from the legacy custom silicon business that we had acquired from open five to higher margin, more non-China focused custom silicon business now and into the future.

What I wanted to do is to spend a little bit of time talking about our backlog. All the numbers that I'll walk through in the presentation today, whereas the end of last year. We ended last year with $355 million of ending backlog.

Now, as Tony mentioned earlier in our presentation, the guidance that we gave for 2024 revenue was between $345 million and $365 million. So $355 million as our ending backlog. Part of the work that

we've done over the last several quarters is to really dive into the details of our backlog, each of those transactions as well as the bookings that we've gotten over the course of this year and what we anticipate and really understanding the trends between bookings, revenue and backlog.

What I have on the charts here are what we report. So we report our business in two ways between licensing and NRE and silicon and royalties. And what you see on each of the charts for '22 versus '23 is spectacular growth in both cases. You see the strong bookings growth in license and then our RE business, you see also very strong business growth in silicon royalties. You see the same in terms of revenue growth, certainly outpacing our industry.

What you see also is the trends from an ending backlog perspective. For license and NRE, you see ending backlog trends that are similar to the growth rates that you see from a bookings and revenue perspective. For silicon and royalty, you actually see much less ending backlog at the end of 2023 and what we have in 2022.

This is exactly the transition that Mohit talked about through our acquisition of legacy open five. We had acquired a substantial amount of backlog with lower margin what I would say, China business modeling in the consumer space that we worked through over the first half of 2020 to have the three, what you'll see is in our guidance for '24, the first half our guidance for 2024, we actually expect to be less than what we printed last year because of this transition.

So on an underlying basis, as the chart showed previously and see in the numbers now, the underlying growth of the business that we have going for is nothing short of spectacular. And you're seeing that transition come through in terms of our financials.

What I also have in the bottom half of this chart is a little more detail on how all of these things get accounted for in our financials over the course of the product life cycle. This is the same lifecycle that most of the showed in terms of what we work on our customers with. And what you see is when we recognize a booking versus when we recognize revenue, when we recognize expense and how all that works is different depending on the different businesses that we have.

So there's not always the same correlation between the timing of a booking, the timing of revenue and when we receive cash. In the early parts of this cycle where we list architecture and product specifications, that's predominantly internal investment from us as a company.

These are the 740, 750 engineers that work for Tony and John making these world leading products that enable the changes that we all experience in our lives into our future. Now that's a combination of headcount, but in some cases, it's also large CapEx amounts, for example, from mass sets. Fourth, the business at the Babak is running, right? Where -- what we have to do is we spend money on mass sets for these test chips that are coming out into the future. And that's in the early part in terms of internal investment.

Now around the time of a design start is when we receive the order from the customer. And in many cases, we actually receive cash upfront associated with that design. Now Mohit and Babak assures me, we'll get more and more cash upfront in those transactions. But this is part of the work we're doing from

an infrastructure perspective of how we look at our contracts and how we structure them to make it easier for us to run our company and also for you to understand what we are doing. What happens over the next 12 months to 18 months, is then we'll recognize NRE revenue over that development period.

One of the changes that you saw from a guidance perspective for us reflects this change. Because what happens in our deals where we have both licensing as well as ASIC revenue. We recognize revenue on a percentage of completion basis.

Now if you look at the total amount of the deal and the expenses associated that tends to be heavily skewed towards the end of the development period because of the cost of the masks. Masks can be millions to [10 million] or more for each individual product. So from a total cost of a project perspective, it tends to dwarf the overall cost of the projects, which pushes revenue recognition later into that process.

Now for the mass sets we work on with customers, we actually get paid for those mass sets in advance of our us paying our foundry partners. Mohit, Babak, thank you again.

But the timing of that revenue and the timing of the cash flows then can be different. And that's what you're seeing in some of the trends in terms of what happens to our financials. So again, a tape-out will get the cash from asset for the tape-out, then what happens is then were released to production, and then we'll recognize silicon revenue over the course of that production period.

So those trend -- transactions could be $50 million or more sometimes less, sometimes more depending on the size of the design win, depending on the longevity of the contract, which sometimes happens based on what the end market looks like as well. And then we'll recognize that silicon revenue.

Now one thing to note is in the $355 million of backlog that we reported at the end of last year, that tail of production revenue was not included. That tends to be recorded almost on a transactional basis, right? We'll get an order from our partner.

We'll release that order to our foundry partners as those products shipped, then we'll recognize the revenue will also collect the cash in terms of what's there. So the revenue recognition profile can change fairly significantly depending on where we are in the business cycle, also what it means for us from a business perspective. But what we look at is cash.

One of the ways that I was trained is that P&Ls opinion, the cash is true. And so this is something that we're actively focused on making sure that our company is properly capitalized, and we have the cash on hand to execute on the growth that we've talked about earlier today.

What I wanted to do now is to spend a little more time specifically on elements of our P&L. So the investment that we're making in these products is what drives the P&L at this point. What you saw from a revenue perspective year-over-year is a 74% increase.

Now, again, that included acquisitions that we brought in towards the tail end of '22. But whether it's organic or inorganic investment, our growth is nothing short of spectacular. What you see is a gross

margin perspective in '23 of about 51%.

Going forward, I expect that gross margins to be roughly flat because what you'll see is you'll see growth in all aspects of our business. The silicon IP business has phenomenal gross margins in the 97%, 98% range. What happens, though is that then you have products, products have typically between a 50% and 60% gross margin. And then you have our custom silicon business.

The legacy custom silicon business that we acquired has lower gross margins kind of in the 20%, 30% range. As we move up the value change, you'll see a transformation in those custom silicon margins between the three elements of our business, it should converge in the low 50% range for the next several years.

As we grow our product business out into the future, I'd expect to see those gross margins continue to increase. What you see then is an adjusted EBITDA perspective of 34%, about $63 million in 2023. That's predominantly because of the substantial investment that we're making in products.

So we spend $132 million in R&D, sales and marketing and G&A, again, almost 90% of our employees are engineers. What we also did is capitalized another $53 million in term -- excuse me or so in terms of R&D capitalization.

Now the capitalization that we do of R&D for future products. So that investment in future products where we don't have revenue as yet. And from an IFRS perspective, what we're asked to do is to put that actually onto our balance sheet and we'd only start to amortize that in future years as those products are launched.

What we tried to do is to be very transparent in terms of what our spend is. So that you see it, but it also gives you an idea of the rough balance of spend that we have that's supporting existing R&D as well as future product.

One thing that we're also doing is working on our infrastructure. This is a business that has grown exceptionally well, both organically and inorganically. And right now we being the royal wave mainly [Chris] is operating for different ERP systems, right? That creates times, it creates complexity, it doesn't allow us to give the information we need to our management team, to our investors, and that's part of the work that Tony talked about earlier in his presentation, is making the investment not just in good looking new executives, but also in the infrastructure that we need to go run our business, right.

So that we can accelerate from a decision-making perspective. So we can make those trade-offs very quickly in terms of where we want to make our investments in providing the information both inside and outside the company as we continue to make sure that we're investing in the right places.

Here's a little more information specifically about that R&D investment and where we are from an employee perspective, a lot of this information is covered. I think as John Rogers mentioned, we have design centers in Canada as well as in Israel, in India and in California, which some of us call home.

So certainly, we have a industry leading design team that's very focused on making sure that we continue

to extend our technology leadership in these critical elements. Now, I mentioned, we capitalized $54.5 million of development costs again for future products. So that spend that's not included in our existing P&L. That is predominantly still cash that we're monitoring.

On the next slide, I'll talk a little bit again about our balance sheet and cash flows. As the slide indicates, this is a key priority for us as a company and certainly as a finance organization. Because we have to make sure that we continue to have the cash that we need on hand to fund the products that are going to change our lives going up the future.

From a cash perspective, we ended last year with about a little over $100 million of cash, not a lot of money for many of you in the audience, but for our company, fairly adequate. What we saw was a drawdown from '22 to '23 as we make that investment.

Going forward over the course of '24, what I expect because of the trajectory of our business and the investments we're making in our products, I'd expect to generally use cash in the first half of the year and then to generate cash back again in the second half of the year. But it's really going to depend on the timing of the engagements and the deals that we have and how our business flows over the course of the year. Net-net, over the course of the year, I would expect that we would use cash, but we'll certainly have the cash on hand that provides us the flexibility that we need.

One thing that we're also doing is very closely monitoring what we're doing from a debt perspective and from a covenant perspective and what we want to do and what I'll tell you is we have a fantastic relationship with our existing lending community. Actually half of them are represented in the room right now. So we continue to have very active conversations on what we can do from a capital structure perspective.

What I'd suggest is that you should expect us to do something in that timeframe. It's probably measured more in and months than it is in terms of quarters. But as I said earlier, we're confident that we have the cash on hand that we need right now to continue to execute on the growth and we do not need external resources from a funding perspective in order to meet our covenants or anything else that we need to do.

There are some additional information that we private here in terms of changes in working capital and expect to generate cash from working capital over the course of this year, but expect to continue to spend in terms of the capital expenses that we have for a company as we invest in our infrastructure and the products that we need.

And again, talk to this change in R&D capitalization was just from the cash flow statement what we have analyzed actually $54.5 million.

What I want to talk about next is just the guidance that we had issued for '24 that we had initiated in April as well as the guidance that will be stable out for 2025. Now what you saw for 2024 in terms of the guidance that we put out, reflects the timing that I had in the slides, right. I think historically, what our company had done is looked at the bookings flow and the timing of our bookings and translated that directly into revenue.

As I've talked to you earlier, depending on our different sources of our business, the timing can be different, right, between when we recognize bookings, when we get cash and revenue recognition. So the revenue forecast that you see here in 2024 is based on a bottoms-up analysis, right. We did a bottoms up analysis of the 20 -- of the ending '23 backlog of $355 million.

At the beginning of this year, we expect roughly half of that backlog should be revenue for us in 2024. If you look at the bookings that we had in Q1, you'd say that that adds another about 10% of what we have in terms of our revenue estimates.

So from a risk or opportunity perspective, I think it's roughly 10% higher or lower in terms of midpoint that we had. And I think what's really going to drive it is just the timing of the custom silicon gauges that we have over the course of the next several quarters. And I think that's going to determine where we end up in our range for 2024.

I talked about the gross margins we'd have for '24 as well as the OpEx. The OpEx percentage, roughly 30% shows that continued investment in R&D. This is something that is critical for our company. It's critical for companies in our space.

My experience is that in our industry, typically, people have two or three folks that they have from a competitive perspective. If you're number one or number two, number one and number two generate that 90% of the market share typically command a 50% to 60% gross margins that you expect.

For our founders, if you're number three year or worse, you're left with the remaining 10% of the share in your gross margins are in the 20%, 30%. So making sure that we hit the investments we need to invest in these growth businesses this year is critical for us to hit our growth out into the future.

Again, we continue to expect to see CapEx of roughly 10% and expect to have capitalized $50 million to $60 million of R&D. What you see from '24 to '25 is that growth to continue to kick up, right? So from about $355 million to $450 million. And that trajectory, again is based on the bottoms-up analysis of where we see bookings, where we have our existing backlog and also the engagements that we have with our customers.

Now, of course, our business is dynamic. The world is dynamic. My experience in doing this for the short period of time I've been doing it is that it usually ends up better than what you anticipated, but not in the way that you thought.

Any given Sunday to the American phrase, things are going to be a little bit worse, things are going to be a little bit better. What I look at is the overall trend, and this is also why I've looked at numbers from an annual basis because otherwise, if you look at half-year or quarterly, any given period may be a little bit better, maybe a little bit worse but what we want to look at is making sure that we're following the trends.

We're growing as fast or faster than the markets that we described that we're doing so profitably and that we're also making sure that we make the investments we need to continue to be competitive and

be good partners for our customers, right. And that also that we have the cash and liquidity on hand to continue to invest.

With that, what I'll do is I'll pass it over to Tony for concluding remarks before we go into Q&A.

Tony Pialis^ Thank you, Rahul. Thank you, everyone today for listening to our talks. I ended my portion of the talk stating that Alphawave is at the heart of the AI revolution. And hopefully after hearing the knowledge and the insights from the executive team, you can understand why we've earned that spot.

It's the planes, trains, and automobiles that Charlie referenced. There is no other supplier in the industry that's able to service the hyperscaler and the ecosystem around them and give them all the technology building blocks that we're able to deliver. Whether it's helping the hyperscalers, build their custom silicon via our silicon IP portfolio, via our custom silicon business, and now through the launch of the industry's first portfolio of chiplets, we solve their silicon problems.

And with the addition and the investment into buybacks, connectivity problem -- connectivity products, we're helping them now connect the GPUs, connect the servers, distribute data within a data center. And as Charlie's photo shows, connect data across corn fields, across their campuses.

So there's no one else in the industry that's able to solve all of their problems the way we do. And that is why I am so excited to be here. And for those of you that were here last year, I think you can now see the increased scale and the increased footprint that our investments are yielding. Those investments will ultimately power the AI revolution.

So I'm an engineer. I love mathematics. I really understand things when they're described in an equation. I think there's no better equation than the one shown at the bottom here. Our investments will yield our path to $1 billion of revenue by 2027, but it will not end there. I assure you we will continue to scale and drive ongoing growth into the billions through the end of this decade and onwards.

And so with that, I thank you very much for your attention. We will be opening the floor to Q&A. You'll have access to an entire executive team. I want to make sure that they get the stage time that they deserve because they're the engine, okay, easy for me to be up here and be the face, but they are the engine driving the revenue and solving our customers' problems. So thank you very much.

+++ q-and-a

Jonathan Menon^ Hi. Thanks for those presentations. It's Jonathan from Jefferies. Maybe I'll just I've got quite a few. So I'll come back again later perhaps, but I'll start with maybe the financial questions and then come to the other questions here.

Rahul, I'm just a little bit unsure why the gross margins are sort of flattish over a sort of a three year period in what you're showing '23, '24, '25 at about 50%. Because you are showing the but your bookings on the custom silicon, which presumably the China driven low margin phenomenon, you're showing that coming down from $100 million to $60 million, something like that.

You should be compensating for that in -- with higher NRE and presumably some of the higher custom silicon chips are from the sort of new business that you're getting the connectivity you're confirming will ship this year. So presumably, we'll ship more next year at a higher margin and further. So when you add all that together, I don't see why the gross margin should not go up and you should be at the flattish 50% gross margin?

Rahul Mathur^ Sure, and thank you very much for the question and I think as I said earlier, I want to make sure it's and people on the webcast heard there's an opportunity to submit questions via the webcast or the portal as well for those of you who aren't in the room.

In terms of the gross margin trajectory over this couple of years, as you described, it's really related to the mix that I talked about earlier. Is that we continue to see solid growth in just our silicon IP business, which tends to be a very high margin percentage, 96%, 97%, 98% range.

Now exactly, as you described, you'll see less of our lower margin custom silicon business, which had kind of a 20%, 30% gross margin as we've worked through or canceled part of that backlog over the last year and a half. Now that's going to be replaced by more custom silicon business. That custom silicon business though won't yet be at a 50 or higher percentage range. So it's still going to be lower than what we see from a corporate average perspective. right.

And then as you see the connectivity products ramp, those should be in the 50%, 60% gross margin range. But in the time periods that we're talking about, it's still going to be a relatively smaller portion of our revenue and so the trend that you're talking about over the medium term is that gross margins stay roughly flat as we go through these transitions.

Now as we talk about the longer term in our path to $1 billion or so run rate in '27, then I'd start to expect to see gross margin start to increase above the 50% into 60% and beyond because you'll see a larger contribution from products that are higher than kind of that 50%. But we're not yet there from a transition perspective.

Now what I'll tell you also, as I said earlier, any given quarter, any given reporting period, you'll see it sometimes it will be a little bit higher, a little bit lower just depending on the trend of what our customers are shipping. But really on the medium term basis, it's just because of the trend of kind of those three different profiles and their how they fit into.

Jonathan Menon^ And then just on the 2025 outlook of the $450 million of revenue, that one of Tony's chart showed that silicon revenue coming on top. Can you give us an idea of how much you are assuming of that revenue and that number? Is there assumption of something? And what kind of visibility do you have on that sitting today?

I mean, is it -- you could probably paint a picture where we had $100 million or $150 million all could be at $30 million because things don't ramp as you expected, but probably ramp. So what is the level of confidence in that number, in that visibility? And what are you factoring in into that $450 million.

Rahul Mathur^ So one of the things I mentioned earlier is that we ended '23 with $355 million of

backlog. As I mentioned, roughly half of that we expect to recognize in 2024. So natural assumption is that the other half you'd expect in '25. So that's the first basis in terms of the [$450 million].

Another [Peter], as Mohit talked about for silicon is that that's not something that we report in backlog, but we have a design win and we also have expectations and estimates from our customers in terms of what they predict to ship this year and next year as well. So that's part of our forecast.

And what we also see is continued growth in our silicon IP business, and that's a high-margin IP business that we were talking about. And so you start to see that delivery. Then on top of that, the new end products, right? So in 2025, could 10% of that revenue would be part of the new products that we've talked about? I think so, absolutely, it could be more. It could be less depending on the timing of shipments. And then you see the other portions of our business continue to grow very nicely from that growth trajectory in terms of what's there.

What I talked about from a risk or opportunity perspective is really the timing, particularly of custom silicon in terms of what those orders look like, what the ramp looks like, what engagements with customers look like as well in terms of what's there. And that's something that's harder for us to kind of predict, but one of the things we did is exactly that bottom-up forecast, right?

What we've done within the company from a process perspective is we've created communication between the finance and accounting team, between the sales team, between operations and engineering and of engagement by engagement that says not just what order do we have, but what's our delivery schedule associated with that and then based on our delivery schedule, what's our revenue recognition profile and that what gives us a little more confidence in terms of that forecast.

That's also why I think the forecast was trimmed about kind of from the $500 million to the $450 million to reflect just that trajectory and the timing of revenue recognition. Now, of course, there's still a lot of time left in this year, and one might argue you have all of next year still ahead of you, right. So there's still time for things to get better or for risks to evidence themselves, either from a customer perspective or macro perspective.

And sometimes those are going to be tailwind, sometimes those will be headwinds, and that's tough for us to predict which is why we read your reports in terms of what's going to happen in terms of our business. But those are some of the things that we've done to help us give us confidence, in terms of the guidance that we have for '24 as well as '25, and also some of the trends that we see in our business.

Jonathan Menon^ So just on the 10%, sorry, just to understand, clearly, the 10% that you said in 2025 is the connectivity products or is that the custom silicon shipment revenue?

Rahul Mathur^ I think that could be more predominant in the connectivity product, kind of you'd see more from a custom silicon.

Jonathan Menon^ So my question was actually on the custom silicon revenue shipment revenue. So what is -- are you assuming anything of that in the $450 million?

Rahul Mathur^ Yeah, absolutely. We'd expect to see substantial growth in custom silicon, over the course of '24 and '25. And I'm sorry, Babak, when I said 10%, I mean [50] in terms of what's there, but that's what you start to see in terms of the trends.

But again, and there's one part of making sure that we invest in products to get the design wins. The other part is the ramp and when they ship and that tends to be more of a industry's timing perspective than a company specific.

Jonathan Menon^ Just one last question for me. On the on the product side, you've sort of that a bit on the compute silicon side versus last year, sort of where you are more connectivity and in its compute as well. On -- and a lot of it seems to be coming from the ARM relationship.

ARMs, if you see ARM's direction of travel, it's trying to gain. I mean that the story is a bit like yours. I mean, we are trying to offer more and more to our customers, so that our customer virtually has to do nothing at all, but by our new verse and compute subsystem.

So does that -- is that a risk in the longer term that your partner is trying to eventually take some of your values, the window in between where you can add that, that that service of being the in between who put it all together, but eventually you're going to because they're going to they're raising their royalty rates in order to take as much value for themselves as possible. Does that sort of edge you out in the in the longer term?

Rahul Mathur^ Mohit, why don't you take that question?

Mohit Gupta^ So when you look at -- I mean of course, they're fundamentally still an IP company out there right now, what they do the large chunk of revenue. If you look at $500 million coming from licensing a $1.5 billion coming from royalties, right?

So the partnership we have put in place, I think it serves two purposes. One is when you build these custom chips, you are integrating those ARM cores inside those chips, right? So it expands our Horizon and the market we can so to build those custom silicon chips, right.

Second is the which is a very differentiated element to build chiplet and integrate those ARM compute chiplets as part of the overall solution for our customers. So at least my ARM of opinion, I don't think our intention to go into both of these because that will definitely put them as a very different company what they are today, right?

So that partnership, they do see us as a very complementary to each other, right? So the work we are doing with them over the last, like I said, we formally announced this over a year. So but we have been working with them for last three years, right? And we are shipping ARM-based chips actually even from a legacy perspective. So we're not newcomers to the ARM ecosystem and where it stands, I think ARM ecosystem is just going to drive.

Of course, the penetration in the market right now is pretty much the market we focus on is almost 100% right now, we don't see any other architectures in there. So for us, I think it's a great vantage point

what we have delivered with them. And I don't think there will be an intersection point with them the way we are structured and where they are headed.

Unidentified Participant^ Question on your I mean, you've talked a lot more about your AI chips this time than it did last year, and I need that. (inaudible) And when you look at where who is winning there at the moment, it's virtually entirely dominated by [1.6]. And so it means that the cost of your customers buying this silicon from you are either going to be in a customized part of the data center, which you did talk about as well or that you think that a media is going to lose a lot of share over the next few years which anybody's guess of that, especially given the software lock-in, et cetera, that developing. So how do you see that play out essentially because these pretty important to your revenue that how this metrics plays out in that part of the market? That's my first question. I've got one more.

Rahul Mathur^ Charlie, why don't you take this?

Charlie Roach^ Alright. I think what we're seeing in the marketplace is the -- that market a market is continuing to grow at a very fast rate. But what you're seeing from the hyperscalers is they've all announced there are all developing their own silicon that is going to compete with that ecosystem.

And what they're doing is they're announcing is when they build out their data centers and their AI, they're going to still say, hey, if you want to work on and in video processor, you can go specifically to an end video processor. If you want to work on our internally developed processors, you can do that, and the pricing might be different for that compute.

So what you're seeing play out and our strategy has really played out with the hyperscaler strategy. Is there going to give you an option of working on a AMD-based processor, their own grab a ton processor or and in video processor, and they're going to let the market dictate. Do we buy more and video processors, do build more of our graviton processors or do we buy met more AMD processors.

So that's really what we're betting on is that there is going to be an increased market overall, but the hyperscalers are spending hundreds of billions of dollars developing their own solutions as well. And so I would just go back to that, and that's what we're really relying on is that those hyperscalers and ecosystem is going to be successful. In video, it's still continuing to grow because AI is in its very beginning of deployments in the marketplace. So I think there's going to be plenty of room for everybody.

Unidentified Participant^ Then just actually a follow-up to that would be the connectivity itself in Media's connectivity is quite different from other people's connectivity, InfiniBand as well as CO -- you don't there is CXL versus there is envy link. So it I mean, is this going to be a problem as the other players ramp up? Or is these connectivities technology is going to be an impediment for other people to break into this market?

Rahul Mathur^ Tony, why don't you take this one?

Tony Carusone^ I mean, I think there is a proliferation of standards being developed to address this challenge you're talking about. I don't think we're in a sustainable situation here where you've just got

one supplier supplying so much of the market there.

In addition to what Charlie talked about, it also just point out that all these hyperscalers have their own workloads right rather than just offering these platforms for the open market, they've got their own workloads that they know best, and that's really where they can leverage custom silicon solutions.

So what they want is these customizable silicon solutions mixing and matching connectivity, IP compute, IP memory and so on, but also sort of standardized connectivity solutions. So that they can have a component supplier ecosystem, that's more diverse and so that they're not locked in.

So that's. Now that some of that is playing out right now over the course of one to two years, you've seen sort of the initially proliferation of standards. And then I think we're turning the corner now and seeing it starting to converge. People are coming together to agree upon alternatives to the NVIDIA ecosystem to enable this sort of combination of ecosystem partners that's really going to lead AI scale sustainably.

Unidentified Participant^ -- These ASICS that you are developing at this point? I mean, do you start -- Have you started shipping them already in revenue terms? And or because you've I think there was one slide which showed second half this year, you're going to start seeing some revenue. So how has the revenue already started coming in?

Tony Pialis^ So we have two types of ASICS, that will be shipping. The first is Mohit custom silicon business. So the new types of wins that we've been securing since we've integrated opened five. Right now will be and the first generation will be entering production late this year, early next year. So revenue will follow thereafter.

From Babak's connectivity products, which are the pleasingly fast optoelectronics, the power, the next generation of AI fabrics. Those revenues as he stated will also start to come in right at the end of this year, first on PAM4 then on coherent and start to ramp through '25, '26 and but you get to '27 contributing a significant portion to the overall corporate revenue.

And I think that also obviously dovetails into the whole margin discussion. By '27, these higher margin products will be shipping and high enough volumes to provide that gross margin uplift to the overall business.

Rob Sanders^ Yeah, here, Rob Sanders, Deutsche Bank. I have a question for Babak. What I didn't really get from your presentation was some data on your relative position against Marvell in PAM4 and optical DSP. So when you talk to customers, where do you score ahead of Marvell, where are you behind? Because I think we all agree there's a huge opportunity, but I didn't really get the sense of what alpha ways unique value proposition is in the optoelectronic space. Thanks.

Rahul Mathur^ So we're behind our market share. And that's going to correct itself with our sponsor hyperscaler, who is working very actively engaged with us to give us that revenue gain. And then we highlighted one area with the wide IDSP platform that we're developing to be able to compete effectively. And other few other things are also in the hopper. But the share gain is our number one

disconnect, and we've accelerated our cadence.

We will be announcing our next generation faster than you probably would have anticipated, and this is where we're going to start catching up on our competition by surprise in terms of the agility with which we pivoted to -- we were taking our Generation one to production this year and move very fast follow the turnaround time from our Gen two start to finish is eye-opening, and so we're going to close that gap.

Rob Sanders^ Great. Thanks for that. And a question for Mohit. Can you just confirm how many hyperscalers have already you've designed in or at least are working with on custom silicon for AI related products. I'm just not quite clear on what the number is that you've already locked in?

Mohit Gupta^ Yeah. So we don't publicly disclose like which and how many hyperscalers. But what I can tell you is we are shipping like either direct or indirect way. So many times we are working with customers whom we have built chips for them, and they are shipping to hyperscalers. So hyperscalers, like Charlie indicated, right? They are all unique beast in themselves in terms of engagement, style, and touch points. And we see both kinds of engagements coming in.

On direct side, yes, we are already in there in many of the hyperscalers in production where they have put our technology out there, both North America and A-Pac hyperscalers out there.

Rob Sanders^ Thanks. And just my last question would just be on the situation. I think we all agree this is big opportunity, but you are this is a question for a whole lot of questions. And you see you're playing a big boy's game against massive companies, right? And yet you're small and you're capital constrained, are you may not agree with that? I would argue you're capital constrained. So why not just raise some equity capital quickly to shore up the balance sheet? Because I mean, I assume that if customers are looking at you, they are wanting to know that you all well capitalized as company. Thank you.

Rahul Mathur^ Sure. Thank you for the question. I would say we're smaller, but just as mighty. The other thing I would tell you is that in some of these competitors, the phrase that we use even as a management team is that they have to be a mile wide, but an entity, right? And so what you saw in the presentation earlier is the specific focus where we can be a mile deep, right, and really having a holistic solution for our customer and then their customer as well. And it's a different dynamic, right?

When you're in one of the larger companies like the ones you're describing because from a scale perspective, yes, they may have more number of people or engineers, but they have the same number of people, do they have the same focus and the same expertise and these critical areas that are going to drive the decisions, right, in terms of what's there.

And we don't have that burden. We don't have that burden of having to be that mile wide, right. We can focus on specifically what our customers need in order to do that. From a balance sheet perspective, this is something we actively manage on a daily weekly basis in terms of what's there and what we're doing from a cash perspective and everything we said see that we have enough. Now all things equal, it's better to be richer, right, in terms of what's there. What you should imagine is that from a financial perspective, we're perpetually evaluating all sorts of different options.

And it's interesting because I had a conversation with one of our investors earlier, and we'll have a conversation with an investor that says, oh, you should absolutely choose door number A, and the very next conversation will have a different investor will say, no, no, it's door B and the next conversation, we'll have another investor will say, no, no, it's door C, right.

And so what we'll do, we'll depend on what's in front of us and what we see there because each of the different avenues that you could have to do something from a capital perspective as opportunities has risks has benefits, right? And your equity is one option where we're trading now, I think equity is very expensive, right?

There is a great support that we have with our existing partners from a debt perspective, I mentioned half of them, half of our debt partners are in our room right now here, and they've had the misfortune of meeting me multiple times over the last couple of quarters or so.

So there's a lot of different things that we can do and it's something that we're actively considering. I think if we were to do something that would be measured again in months, not quarters on. But we see the same numbers that everyone else does and sometimes even a few layers and down below.

Tony Pialis^ If you know, Charlie referred to the various in the playbook with the hyperscalers to engage flexible, nimble companies with us, right. And I don't know about Mohit, but I haven't had the full Rahul in a customer engagement that justify our size, right? Because it's in the playbook. They know they can stand you up. And if you're providing the technology and the cadence that you need. So there hasn't been a concern and it's just part of the engagement model.

Simon Coles^ Hi, Simon from Barclays. Two quick ones. I think so the times obviously got much, much bigger, and I guess ChatGPT took off just after the last CMD. So we knew that was coming. But I just wanted to dive in on the custom silicon parts, $15 billion now $28 billion$, 30 billion. How much is the new opportunity with the compute that you're now disclosing that you can offer to customers and how much was just the general AI market and connectivity market increasing driving the change in the custom silicon?

Rahul Mathur^ Mohit?

Mohit Gupta^ Yeah. So if you look at it, like I said, the custom silicon market when people projected for last year, they were looking at 2027 to be around probably little less than $25 billion, right? And we took an estimate of $28 billion, and there are estimates around even $30 billion plus.

So give or take wherever that number is $27 billion, $28 billion, I would say the compute is probably roughly about 30% of that market, right, which we were not addressing at that time. And that's it suddenly becomes, but it's not certain for sure, but becomes part of our addressable TAM now at this point.

Still, I mean, there are a lot of architectures, like I said, it’s all driven by customer workload, right? So when you talk to a customer, everybody wants choices, everybody want what options in there. So this is I

see it as one more pillar in our portfolio, like I showed the five pillars of our custom silicon platform. It is one of those portfolios which will allow us to get more and do more with our customer base.

Simon Coles^ Great. Thank you. And then another one, I guess everyone tends to focus on the hyperscalers when they think about their business and now to over 100 customers. Mohit, I think you sort of indicated that some of those customers are shifting to hyperscalers. So indirectly, that's how the sterling to hyperscalers, but it would be great to get some more color on who those other customers outside the hyperscalers really are where they're focusing, what they're trying to achieve. Just to get a bit more color about the broad customer base that you have.

Rahul Mathur^ Yes, go for it.

Mohit Gupta^ Okay. So if you look at the customers right, we should in totality, right? The 100 plus customers primarily there between IP and custom silicon and of course, Babak is working with some leading hyperscaler on the product side. But if you take the vast majority of those customers when it comes to the IT side, they are companies who are building traditional networks, right, switches, routers, enterprise, enterprise-class storage devices, right? Even servers. So that's and 5G base stations. So we are shipping our IPs into those type of, those categories of customers, right.

And we are custom silicon sits today. I mean, if you go back 20 years, custom silicon was sitting primarily into the networking domain. That's how Broadcom and Marvell, the LSI logic of the world were formed at that time, right? So over the years that custom silicon has evolved, and AI is the big boost right now in the custom silicon business,

There is still a custom silicon market which exists for non AI chips as well, right, which are switches, routers, enterprise-class SSDs, and we do serve those. It just that the growth we are seeing on the AI segment is just phenomenal, right? And everybody you talk to the just talk about that. But there are still markets we are playing in for those traditional if you want to call it an ASIC market.

Tony Pialis^ No, there's an interesting stat. By 2027, I believe 2/3 of the networking and storage equipment are going to land and hyperscaler data centers. And so the whole reason that we focus so much, and I think this is listed in Charlie's slide is because there are 2/3 of the market and there's only a handful of them. And so it's important to integrate yourself deeply with them, which is why we're focused on the planes, trains and automobiles and solving all of their form of data transport and now compute.

Charlie Roach^ I think to your point, we are also working with leading larger semiconductor companies. And just remember where I came from, I was a customer of Alphawave licensing the technology, I built my previous products and then I joined kind of the driver on connectivity products, but we do absolutely continue to engage and have business with the major semiconductor players.

Tony Pialis^ And most of that silicon eventually will make its way into the hyperscaler data centers.

Harvey Robinson^ Hi. It's Harvey Robinson from [Tammy Aliberum] Three, I think related questions about share of value. Just to help us with how those sort of your market opportunities evolving. Could

you just give us feel for how much of a piece of silicon with chiplets in it that chiplets accounts for what percentage of the value. And the second thing really is within that, what's your share? I mean, obviously, you're suggesting it is growing but gives a feel where you are now where you're going to.

And then coming finally, back to the ARM ecosystem, obviously mentioned a number of times today looking at their press releases and on total designed, I didn't specifically call it out, but I believe there's lots of numbers, lots of people in that. Can you give us a feel for how important you are to ARM? Thank you.

Tony Pialis^ So, Mohit, why don't you take the first question, which is what is the value of chiplet? Of the chips themselves within a final product.

Mohit Gupta^ So if you look at -- so today, as if you the way we look at the chiplet market, right, it's first of all brand new market. It never existed until a few years ago. And even now it is still call it is in the nascent phase where the known good die, iIf you go out in a market to sell other triplet is still early.

So that market, as I will still suspect few years out from where we stand today. What we see is to be able to use those chiplets as part of our custom silicon offering, right? So that's where I would say about 20% of the business is what I would expect where we can leverage some of those chiplets into our custom silicon business, right.

I think coming to maybe '26, '27 timeframe, I do expect as the packaging and test equipment get formalized. There will be a market where we will be able to sell known good dies. But again, time will tell. There is a standard data over a particular protocol, which is called UCI, Universal Connector Interface, right, over which and each click and talk to each other. But there's lots of work to be done there. So that you can buy a chiplet from vendor A, a chiplet from vendor B, put it in a package and they will work. I think it's going to take a couple more years to get into that market and project that revenue. Right now, it will be all in custom --

Tony Pialis^ So regarding your second question, what share of the chiplet ecosystem or chiplet market we win? There is no chiplet market, which is why I announced before bringing to the world into the industry, the first portfolio of computing connectivity chiplets. So we will own 100% of it given that no one else has done it yet. It's inevitable. It's the only way we can scale compute.

And so by being at the forefront, yeah, look, we get first mover advantage. And because there's not an ecosystem, as Mohit said, for us, this is a vehicle to win our in the custom silicon space. This is how we can go up against Broadcom and win. We will not need to charge tens or hundreds of millions of dollars of NRE because we've already pre-built these scalable, interchangeable, and customizable chiplets.

And your third question was? -- stay tuned over the next 24 hours. Okay. Stay tuned. That's -- that should answer that piece. Yeah.

Harvey Robinson^ Hi. What may be follow-up question for you is on the hyperscaler customers on the custom silicon side. When I look at your competition, it's everyone broadly knows which hyperscaler Broadcom is making stuff for, which hyperscaler Mediatech is making stuff for these days and there are a

new entrant into this market, probably even after you announced your entry.

And who does -- do you see do it forward? And also to explore and their multiples are reflecting a more direct connection to that. So why don't we just be more open and give us a few names, and I'm sure your share price will reflect that, that information quite nicely. What I mean, how come they can do it and you can't do it.

Tony Pialis^ Our agreements don't allow us to do it otherwise would be a pure announcing. Look we're providing these chiplets, these custom A6, these connectivity products into these hyperscalers. Let's see maybe next year, the SME buyback had set a goal.

I look forward to having a goal where I could stand up here and start pointing to specific hyperscalers, specific types of engagements and say that's powered by Alphawave. But look, it's a secretive industry. Apple doesn't announce who their suppliers are, nor do typically any of the hyperscalers.

A lot of the information that you publish, and others publish comes from interviews with direct people rather than through press announcements. But yeah, we'll set that as a goal for next year to be able to provide more concrete examples of how we are at the heart of the AI revolution and where we're contributing.

Harvey Robinson^ Well, just some leaky interviews like that.

Tony Pialis^ I do not condone that. All of our employees honor their confidentiality agreements.

Harvey Robinson^ Just on the chiplet, I mean because again, there was three and now you've sort of got a chiplet block way, prominently, which I think is a change. If I'm not mistaken. So that refresh referring to the '26, '27 standard chiplet that you've talked about?

Mohit Gupta^ Eventually we'll get there. But right now, like I said, it does and that's how I presented. If you see my presentation, I didn't say chiplets one presentation or custom silicon. I see chiplets today is a big enabler for custom silicon. It's a inherent part of that strategy. But eventually, as the chiplet market will mature, I do see it can stand up on its own leg.

Harvey Robinson^ And one last question to Babak. I mean, while things have changed. The $300 million hasn't changed in the last two years or so has it actually increased, or you can't announce anything or are we still with the one customer and the $300 million agreement?

Rahul Mathur^ Well, we have the one customer, which you noticed has a greater than $300 million. But we're actively working with other hyperscalers and are in various stages, but some might be with our next generation that will be into intersect point. Right. But the overall revenue opportunity for the products that we're working on is multi-hundred million dollars. But even with the one hyperscaler, it's greater than $300 million.

Harvey Robinson^ -- As a chip when they you'll self now. I mean, you have to test them and take them to market in high-volume production. And this has not been this is not the easiest thing to do with these

multiple chiplets on the same dye, et cetera. So I mean, can you talk a little bit about where you are because you will have to help your customers go to market with these with the products as such really. I have a couple of small follow-ups.

Mohit Gupta^ Okay. So no, great question, actually. So we have actually for one of our first chiplets is in the fab. It's coming out later this year. We are ramping up our test equipment, working with test vendors, how you're going to test it. We are working with us at how you're going to integrated and different platforms. If you are familiar with the buzzwords like [cowas], info, you have heard like TSMC capacities are full.

So we are looking at alternate dose at all for right, working with the likes of AAC. Intel has e-mail -- Gil cetera, or sorry, [Amkor] has another equal and process node. There is a new chiplet packaging, how silicon box they're just focusing on using glass panel substrates to do it.

So yes, we are working on those as part of this journey. So like I said, the first is value extension chiplets then comes in down computer-based chiplets and then the memory and then eventually the Optec will follow.

Tony Pialis^ But we're following the same pattern as we did with the connectivity products, connectivity products draft behind our silicon IP. Our chiplet strategy track behind our custom silicon business where Mohit already shipping more than a dozen chiplet based custom silicon A6 to the market. And so for us, a lot of that infrastructure is already in place given the business that we're already running.

Mohit Gupta^ And it's unique actually, I must tell you like I was with Tony, and I were in Asia last week and we listed quite a few customers when we show our triplet portfolio days open up. Nobody has shown that portfolio and how we are able to do. So I mean, time will -- we will come back and show you the actual products. But we have big engagements sitting out there to engage just leveraging our triplet portfolios to build bigger customers.

Harvey Robinson^ (inaudible) Where is it in terms of adoption by, I mean, clearly one hyperscaler wants to use your product, but where is it in? Because it I mean, it's not technology, which is used today in the data center, where is it in terms of adoption of the technology, the data center today?

Rahul Mathur^ So there's two hyperscalers in particular, who are very interested in coherent light technology. And our plan is we will be shipping for revenue next year.

Harvey Robinson^ And all the software and the whole ecosystem is ready for that already?

Rahul Mathur^ Yeah. That's being worked on the comment you made about ecosystem. It's absolutely true. The number -- we have multiple calls on a weekly basis with hyperscaler plus the supply chain ecosystem partners are solving things like the optics, the tuned optics for those applications for coherent light. So it's here and now we're like living it.

Mohit Gupta^ And it goes further. Tony, do you want to give an update across the broader industry for coherent edge.

Tony Carusone^ Again, the coherent light is really the only way that bandwidth is going to be able to continue to scale inside the data center and more and more people are seeing that. And that's why you're seeing standards evolve around coherent Lite. And I think you're going to see new products come to market from system vendors as well that are going to rely on coherent light technologies to support that back-end networks increasingly and provide that path to the fast cadence, they're looking to upgrade those networks.

Tony Pialis^ So validation is standardization of which over the last 12 months there's been significant IEEE, which is the governing body for virtually all forms of connectivity that have standardized now behind coherent Lite because they see the momentum that's gaining there.

Tony Carusone^ And the other thing is we have contributed to that standardization by providing some innovation around forward error correction and things like that, where we're providing the test case model that others will test against potential interoperability because we're just kind of a first mover advantage there.

Mohit Gupta^ And we are very active in pushing that standard forward and people are very open to our input on that and our influence on that because they recognize that we had data to offer just being technology leaders, they were ahead and being able to do that.

Harvey Robinson^ I have just two quick questions on the first would be, and you might -- you might say to me, wait and see mark on this one but with respect to ARM, everyone likes to say they're an ARM partner, basically, I know I can you just give some indication of the level of partnership in areas like is there like a sequence of ARM partners and where do you stand on that?

Mohit Gupta^ Okay. So what -- let's see how much I can say. So we are pretty unique in what we have done, right. If you look at all the ARM partners out there, we are not just one in there for the name or the logo. First of all we have bought an IP and a silicon company. They don't even have any other partner like that.

What we are working with ARM is not just technology IP and just build chips around it, right. We are actually you will hear a lot more in a very short duration. We are actually going to make some major investments to build leveraging ARM cores to do those chips. So it is a much different than what other partners and ecosystem as even the total design, if you're referring to, we are taking a jump out of that leak into going very differently from that.

Harvey Robinson^ Great. Thank you for that. And the second question is up. Thank you for being here in London. I know it's a lot of the sort of center of everything is elsewhere in the world for this world. In terms of incentivization and obviously, as a company, there's been some sort of speed bumps over the last few years, et cetera.

Can you talk a little bit about incentivization and keeping key people in place and maybe equity participation that sort of thing?

Tony Pialis^ Sure. Let me take that one. First off, what motivates engineers, first and foremost is quality of work, all right? What we do is probably the hardest form of electrical and Computer Engineering. It's the sexy part of our industry. And so by staying at the forefront of new technologies, new data rates, new processing nodes, incorporating compute in our DSPs, incorporating compute now via chiplets. I mean, this is a sandbox for electrical and Computer Engineering.

Obviously, beyond that, you need to compensate people for all of their hard work. We do so. We're global, which means we're not -- look many of my competitors, most of their engineering teams are housed in the Bay Area, right? It is extremely competitive there. We've taken a more expansive footprint with our R&D being located in the Bay Area, but also in Canada, also in Israel, also in India. That's helped us as well, right, bring this quality of engineering and work globally.

I have two engineering teams abroad. And then finally, look, I still call us the best funded publicly start-up in the world. And so equity is another key form of incentive. If you look at the valuation of our competitors like a sterilizer which has one product and you see all the various products on the screen now.

I think once the investment community it realizes the value that we bring, I see tremendous upside. And so do our employees see tremendous upside. So they ascribe a huge amount of value to their equity based compensation.

Harvey Robinson^ Is somebody I part of your TAM. And then second question role, if you are blessed with a green [bond] from a allocation -- capital allocation strategy standpoint, would you be recommending buying back stock or maybe pursuing further acquisitions or giving R&D teammates raises or how would you deal with that?

Tony Pialis^ So I'll take the sovereign AI quickly? It is huge part of our business, okay? Everyone, every nation is building their own AI because they view it as a national paradigm. That's part of the TAM explosion from last year to this year. It's part of the reason. It's not just hyperscalers building their own AI, but its nation's building their own AI. So definitely, and I see this continuing to scale because now you have the Middle East investing a lot of in this space. Okay? And we see a lot more action there. And I think all of us in this room acknowledge they can deploy capital when they want to.

So I think there's a lot more upwards movement that could occur due to Sovereign AI.

Tony Carusone^ In terms of the question about capital allocation, I think we're blessed to begin with right in terms of having what I think is an attractive debt instrument with very supportive partners, in terms of what's there, if we had more liquidity, it would give us a little more flexibility just from a timing perspective. And so it's something where you can spend a little less time in terms of the active management in terms of the daily weekly monthly work that you kind of have to do.

I've worked at multiple companies where we've done some of the transactions that you've talked about. What I look at today is that are there ways for us to accelerate our growth? What are there ways for us to make the risk more assured. And those are the investments that we're making. And I think we're very attractively priced.

I would say you can read our RNS reports and see that I and I've put my money where my mouth is right in terms of what's there. So I think there's a phenomenal opportunity that's exciting for each of us on the stage and then our employees see it as well.

My experience is that and often reputations and valuations can be trailing indicators, since so just as we execute what I expect to see is a substantial improvement in terms of that value. And that creates another form of equity for us as well, not just for employees, but for other things that we can do. But the simple answer to your question is that what I would choose to invest in is either additional products or more investments that we can make that would make us more, feel more assured about hitting the financial targets that we have.

We have all the products needed to succeed here. Rahul, now that we do not need to invest further to add more products. Everything is already within our wave as it stands today.

Harvey Robinson^ And quick question on the hyper scale engagements. Could you briefly describe the relationship that you guys have all the level of collaboration with companies like Marvell in those engagements with the larger hyperscalers, for example, Microsoft.

Tony Pialis^ Babak, do you want to take that one?

Babak Samimi^ Yeah. So engagement level is the primary on two fronts, is a very active engineering level where we even take some of our brain trust out of John's organization and Tony himself personally involved in driving the engagements on the products that we have today and then also our roadmap in terms of inflows in our roadmap. And we're actively engaged on both of those fronts with the hyperscalers.

And then Mohit maybe provide some comments in terms of wins beyond just the connectivity products that his organizations driving?

Mohit Gupta^ Yes. So first is on the IP side, right? As you saw, the scaling is happening the process nodes, but you still keep developing higher speed connectivity, IP, right? So now gone are the days when industry body, they're deciding the new specs and what IPs you need now you see, okay, what are the hyperscalers pushing in the market? That's where the wind flows. So we are actually very closely working with many architecture teams inside the hyperscalers. Okay. When is the next 400 gig coming, 400 gig and [uncertainities]. What does that 400 gig so these will entail. So those are a level of discussions we have with those hyperscalers to drive some of the technology aspects of it.

Then comes the further silicon aspect of it where chiplets comes as a big enabler for us at hyperscalers, right? So we are definitely seeing that will be another interesting point for us to work with many more hyperscalers where we can actually come up with these customized chiplets, which they can introduce in their own silicon because they're -- if you look at hyperscalers, are also getting vertical, right? They're almost like semiconductor companies themselves, but they're doing it for themselves, right? So that's where chiplets is a great intersection point for us. Also on top of custom silicon to talk to them.

Babak Samimi^ One of the other things I've seen since the whole market AI boom happened is interesting phenomenon where we are being asked by the hyperscaler to engage in conversations of what we're doing next because they see what we have and their as their mapping out last week, Tony was in the Bay Area, engage in those conversations and they were -- we weren't asking for the meeting the meeting request, it's coming to us there.

Harvey Robinson^ (inaudible) hyperscalers and it's interesting that this came out. I think early in the presentation and you there to talk about one thing. And then we always we layout our full portfolio and it's amazing where they see intersections with other needs they have in other parts of our portfolio. So I think that synergies are really particularly amazing with the hyperscalers because they have needs in all these areas.

Tony Pialis^ And Charlie has driven this strategy specifically since he's landed here. Because as he stated and I have reiterated, we are the only one in the world, they can deliver all four forms of these solutions. And hyperscalers and our partners that ultimately sell into the hyperscalers need at least one, if not more of these product types.

Harvey Robinson^ Then just one quick clarifying question. This initial purchase order that you I spoke about, it relates to the hyperscaler partnership that was announced earlier?

Tony Pialis^ Yes.

Harvey Robinson^ Okay. Another more market related question. We saw this announcement from this Korean AI company, SAP on a couple of weeks ago, and they are I chip and the partnership that they have announced with you guys, is that a market segment where you guys feel there could be additional potential in those like smaller, almost us top up financed companies. I mean a lot of capital at the moment is a flowing into the space also from the BC side?

Mohit Gupta^ So is -- you saw that? Yes, they had to do a public disclosure on that. So yes, they are one of the custom silicon partners for us. I mean that we are building the chips for them. And Tony talked about it, right? So rainy, I mean, that's a big thing and they are trying to address certain market of them sold. Of course, they have global ambitions. So we look at each custom silicon opportunity, which comes our door and wait it out ourselves.

It's not like a general purpose shelf, that or hey, I want to build this chip. Can you do it for us? We look at what is in there, what is the value of that? what value we bring to the table and overall, what is the future silicon potential? Because ultimately, that's what you guys are looking for, right, how we can ship long-term silicon revenues.

So we read it out. So we definitely are ready. And we and I'll kind of go about their promise for the future where they come the heritage of that company comes back from SK Telecom, right? They used to be part of SK Telecom and they were carved out from there. So we definitely are keen to see how they will ramp up the volumes on that AI accelerate.

Harvey Robinson^ And then the last question on my end, could you share your thoughts on the recent

announcement from ARM that they also want to build AI chips?

Tony Pialis^ So I think ARM has a tremendous opportunity ahead of it. I have read similar. I don't think that announcement came from them, but I think someone referenced discussions that happened with an ARM employee. I think for the largest IP provider in the world to transition to selling chips, it is a little bit of an inventors dilemma, right? Because you'll be cannibalizing your own business because they supply the vast majority of the AI space with their own IP. So it's not clear to me whether they are committed to making this transition and it's not clear to me how they would make this transition.

Yeah. So what I would say is that to me, a really interesting strategy that might exist for ARM can right would be to partner up with someone that has complementary IP like connectivity and to enable them to build chiplets that way, meaning we, would bring to the market and they would continue to collect royalties, which fuels their existing business.

So I could see someone like us helping to monetize their leadership space and help propagate chiplets. ione of the tag lines that Mohit and I are really pushing in the industry is we truly see chiplets as the next form of silicon IP silicon IP was and is today a basic building block in semiconductors. But in the future, chiplet will be that same basic building block.

And so I think I know for a fact ARM executive see this and they see a partnership is going to be critical for them to continue to stay at the forefront of supplying these building blocks. And so hopefully in the next 24 hours or so, there will be another more broader announcement made that will illustrate that this close partnership with them to bring about the next form of building block in the industry, which is moving beyond just silicon IP, but also into chiplets.

Rahul Mathur^ Changing tack slightly, obviously, I think you've given a fairly strong steer that the WiseWave investment is likely to be realized in a relatively short period of time. And also, I think just thinking about the shift that's going on within the business in a sense, a business that was cleaner of China becomes much more easy to see the progress that you're making based on the -- on all of the things that you've talked about today.

So from that point of view, I wonder if you could just talk a little bit more about thoughts about that and also remind us what's embedded in the guidance for 2025 in terms of thoughts about where China is and what you're incorporating with respect to WiseWave in your numbers.

Tony Pialis^ Got it. So I'll handle the first question and then Rahul, in terms of guidance, our I'll pass it over to you. Yeah, so look, why we are driving towards our core focus areas, which is AI, data center and digital infrastructure. Obviously, a WiseWave was an early strategy to tried to penetrate China and do so with an onshore entity, obviously, geopolitics and headwinds, as Rahul referenced, it faced us. And so we've adjusted our strategy.

And so yes, we've stated in private meetings and we've stated publicly, our goal is to exit or exit our equity ownership in that partnership this year. It doesn't mean we're completely exiting China. China is a market that all of our players continue to sell in, but you have to sell and thoughtfully and obviously performing a significant amount of diligence to make sure that we stay on very cleanly on the right side

of export and compliance?

Rahul, do you want to take on the second question?

Rahul Mathur^ Yes, absolutely. So from a guidance perspective, it's one of the reasons I was prescriptive in one of the core just charts that I presented. That specifically talks about our trends from a bookings and revenue perspective and called out what portion of that was related to WiseWave for a couple of reasons.

One, I look at that as a great transaction for the company, right in terms of getting bookings and revenue and momentum that help us invest in R&D. And but as I mentioned earlier, we've delivered everything that we needed to as part of our SLA with WiseWave.

So that initial SLA with them was done at the end of last year. Now we still may have transactions with them for specific products or purchases or initiatives, but it's different than what we had historically. And so I would expect our percentage of revenue associated with China to go down pretty substantially year-over-year because of that transition and the completion of delivery for the Escalade for WiseWave. \

But as Tony and Charlie also intimated, that's still a big market, making sure that we're approaching it the right way with the right partners and targeting business that's more higher -- that's higher margin more aligned to our strategic goals as one of the key transitions for us. So I would expect that portion of our business that goes in China, the trend down because of rates.

Harvey Robinson^ Thank you. And it's a technical question, and it will be an answer to this really. I mean, when we think about tape-out for a specific chip, there's a process to do that, how does that actually work with chiplets?

Mohit Gupta^ Okay. It's just Chiplets is just mini as if you want to call it, right. So the way chiplet integration gets done. So you take a traditional process where you take out the design database, which engineers develop. You give it to a foundry, upload the database, they convert digital files into actual physical myself, a chip comes out and then it goes through an assembly and testing process where you put it on a package. Now that's where the magic happens with the chiplet.

So there are multiple chips which comes in audio and DRAM.s from Micron, Samsung, SK Hynix, take those HBMD ramps and at assembly point, you are assembling multiple pieces together on a package, right? Now to do that correctly, you have to do a lot more work upfront in your design cycle itself, right? So it's not just magically those chips show up and they will connect and talk to each other.

And since a monolithic die has been broken down into different pieces coming from different parts of the world with different process nodes and geometries and more, they have their own characteristics. The thermals are very different, the reliability is very different.

Just a simple example in a very common language the correct the distress on the chip like it can create a crack when you're assembling these things, right? So D-Ram stack is multiple layers of D-Ram versus a if you look at razor thin wafer, I'll show you some of the chips I have, but it's very thin, very thin.

So a lot goes in. That's why TSMC has been talking a lot about their cohorts and in info, and they're investing heavily into those, not just alone their fabs, but in the assembly and test houses. And all over the world, even outside foundries there's so much investment happening. That's a big area where this chiplet thing gets fueled up as an ecosystem.

Tony Pialis^ So this is why the complexity is so great building these advanced chips. It's why hyperscalers continue to work with custom silicon providers because they don't have that capability in-house. And this just illustrates the value of the open five acquisition that we were able to bring this capability and redirect it in less than two years and now be able to serve as hyperscalers and leading semiconductor companies the same way, Broadcom and Marvell does.

Harvey Robinson^ Thank you. Things are working on from active electrical cables through to 1.6 terabit code coherent light multiple custom -- pieces of custom silicon for many different customers. In addition, to selling IP for five or six different interconnect technologies at a variety of different speeds is it's astounding, and you've got 700 engineers. Can you actually get all of this done or are you going to end up focusing on the three or four or five or six or seven things that really gain a lot of traction and then push forward on those?

Tony Pialis^ So I'll quickly answer and then I'll pass it to John Rogers, the time in, given he's the execution engine for all of us. So yeah, we're committed to doing it all. Okay. This isn't speculative product development. There's pull from customers for all of it, but there's a development strategy feeds all the business units that John is has developed and we probably should patent at some point. John?

Jonathan Rogers^ I mean, I talk about a couple of different things on that front. One is from a sort of our research and development group we focus on IPs that are valuable to at least two, ideally, all three of our businesses. And so we get a tremendous amount of re-use across them and a lot of synergy, frankly, what we learn from customers and say the custom product area, we can rapidly translate back into our IP offerings in many cases.

I think the other thing that underlies that is I'm Tony and myself have been in the IP space for a long time, have also run product groups in the past. And so we come at design in a much more scalable way than I think you would see at a lot of companies.

We have a real platform approach to doing things like transceiver design, where we have a -- we allow for significant innovation, but in a very modular way within a larger system design. And so we are able to rapidly improve designs from one generation to another without tearing apart all the software and different sign-off and system infrastructures across them.

This allows us to get innovation out to the market a lot more quickly than a lot of engineering teams, I think in the industry do. And so we have that real focus on platform design and in fact, are addressing many of these activities from the same baseline engineering platforms. And so that's a big part of that as well.

Tony Pialis^ I can give you a specific example of how these synergies translate to direct value to the end

customer, which was your question, how do we differentiate ourselves? Because our DNA is an IP development. We get measured by how much real estate we consume on silicon. So we are very compact Okay. Now we are leaders in terms of being able to deliver the most compact connectivity solutions in the world.

John's delivered his designs into switches that incorporate hundreds of these. So you've got to make it form, fitting. Value for the end customers, when Babak productize is this technology, we now have the world's smallest DSP processors as well.

And so our customers ask us, are you sure. Everything that you say is actually in this chip because when they look at competitors were significantly smaller, the value for us obviously is less silicon we consume, the more margin we make and the end value for the end customers, we can pass on some of that savings to them.

And so part of this R&D engine that John has built has information and learnings flowing across all the different business units, having learnings from our customers flowing into our products, learnings from our products flowing back to our customers. It's very synergistic. And I think the only way to continue to add to scale this form of business, delivering planes, trains, and automobiles.

Harvey Robinson^ I think my friend here when we gave some slides, we will talk with some customers we call music. Does that die size correct or is there a mistake or there's no less correct?

Tony Pialis^ And it comes directly back from the IP lineage, okay. They're being forced to build things as compact as you can and being measured by it.

Harvey Robinson^ And on the connectivity piece. So your insertion to coherent light is at the 1.6 terabit?

Tony Pialis^ We can serve that application with the first generation product we have in the multi-die configuration.

(inaudible) We'll do 1.6, (multiple speakers)

So even the 1.6 T implementations today outside of [Hanger] multi diabetes. So the die currencies, 800 gig, and you have multiple of those. And so hence, again, the die size becomes important but from how the customers or the optical transceiver suppliers, we'll package it. We can do 400, 800 or 1.6 T. beyond 800 gig as multi-die configurations. Thank you.

Unidentified Participant^ First one, but your last CMD and given the history of the custom design team, you made a few references to risk five lots of discussion of the relationship with ARM today that you're continuing to work with the risk five, are you or your customers focusing more on at this point? Any particular reason for this and how are you thinking about risk five going forward?

Tony Pialis^ Mohit, why don't you take that as one?

Mohit Gupta^ So it's a -- like I said, there are particular markets we are focused on, right? And there's a

natural pull for ARM being. There is this dominance of ARM in those markets. And that's why we have strengthened our partnership. We are not taking a very polarized position honestly that okay, we are just going to do ARM and not do risk five.

We do have some risk five base design, but from a large portion of our business, we are getting pulled into doing more and more of ARM base design, and that's what we are doubling down on right now. So we continue to see, and I look at what our customers say in custom silicon, I mean, the name of the business is custom silicon, right? So you have to hear what customers are saying.

Unidentified Participant^ Perfect. And the last one I got. Just can you talk about the competitive environment we said is becoming recognized as an important AI enabler. Have there been new competitors entering the space? How have incumbents reacted to your continued growth.

Rahul Mathur^ Let me tried to take that one new players in the service space? No, I don't think there's been any new players in years. The reason for that being the technology that goes into service these days is its revolutionary, the Wide IDSP that John referenced. It has taken us seven years to develop and refine. And this is after John me, Tony C have been doing this for our entire careers, okay.

This is our third business delivering service IP. So yes, there's not really many people with that kind of caliber that can enter a space and immediately penetrate. And now as the industry shifts to coherent, that just increases the complexity by an order of magnitude.

So is a very, very, very high barrier of entry for any new players in this space. And that's why you've seen such a convergence, right? You go talk to any hyperscaler in this industry and you talk about high-performance connectivity, they will say Half wave, Broadcom, Marvell, that's about it.

Unidentified Participant^ And I'm just wondering about the Board membership, what you're looking for is obviously some of the guys that have exited had backgrounds in telecom and semiconductor, but you've mentioned you're looking for specific skill sets. Maybe you can elaborate on that and just the progress on that recruitment?

Tony Pialis^ Sure. Look, a huge responsibility of the board is obviously to provide governance, but also help drive the planning and how we move forward as a business. So when we went public, we were focused on silicon IP and telecom was a meaningful space for us.

Obviously, as we've transitioned the company and vertically integrated and given the birth of chat and 80 plus percent of our pipeline being focused on AI. This is a key market for us. This is the key market for us, and our time is now to capture.

So as a result, we are refocusing the Board to better leverage skill set in this specific space. So that individuals that have worked within the data center space, that have delivered semiconductors to the industry.

Those skill sets, that learning, that understanding of cadence from silicon development through production through shipping and volume that will help us plan and create the long-term business plans

that will enable us to execute on this opportunity.

Unidentified Participant^ Thank you. And just one other quick follow up separate subject with can you talk a little bit about outerwear is positioning [inference]?

Tony Pialis^ Mohit, you're probably well positioned.

Mohit Gupta^ Yeah. So I mean, when you say inference, I mean AI inference, right, the I mean, pretty much all the accelerators you see the ones I talked about, they are addressing the inference market training market is much smaller, which many of the hyperscalers they're doing the chips themselves, right? The big volume sits inference. So all the chips we talked about where we are focused is actually around the inference market. A lot of those accelerators Our going in solving does influence stuff.

Tony Pialis^ And just so everyone understands and in video is clearly the dominant player in the training market, okay. The inference market is much more open for new players, new opportunities, and that is where hyperscalers are actively investing. That's where software and AI is also coming in trying to solve local language models, local datasets.

All right. I think that might be it. And so with that, thank you very much everyone. We will be standing around. So feel free to come chat with us. If you have any questions, you'd like to ask privately or in a smaller form but thank you very much. Hopefully, you found this useful. Thank you.