

S1 EP36 - Silicon Customization for a 3nm World

Monday, May 08, 2023 · 12:35

Kevin O'Buckley, Senior Vice President, Hardware Engineering, Compute and Custom Business Unit, and podcast host Chris Banuelos discuss silicon customization, advanced technologies and their increasing importance for data infrastructure on this week's episode. Learn more about the factors driving customization and what is required. Hear Kevin discuss how Marvell is enabling customers to deploy advanced AI/ML applications for hyperscale clouds and for 5G radio access networks (RAN) in the most space and energy efficient fashion. Be sure to also read the latest press release on Marvell demonstrating the industry's first 3nm data infrastructure silicon, mentioned in this podcast: <https://bit.ly/3n7YIYI> Kevin

Speakers

Kevin O'Buckley

Senior Vice President, Hardware Engineering, Compute and Custom Business Unit

Host

Christopher Banuelos

Senior Manager of Global Social Media Marketing

C Christopher Banuelos 00:04

Welcome to the Marvell Essential Technology Podcast. I'm your host, Chris Banuelos. On today's episode join me and Kevin O'Buckley, Senior Vice President of the hardware engineering, compute and custom business unit, discussing silicon customization and its increasing importance. Learn more about the factors driving customization, and what's required, as well as what Marvell is doing as a partner in the custom space. Be sure to also check out Marvell's recent announcement on demonstrating the industry's first three nanometer infrastructure silicon, link to the press release will be in the description of this podcast episode. Kevin, it's great to have you on today's podcast. Thank you again for joining. Wanted to start today's conversation with the type of work you're doing here at Marvell, can you give our audience just a brief description?

K Kevin O'Buckley 00:56

I'm responsible for Marvell's custom silicon business. So it's the business that takes Marvell's technology portfolio, and the building blocks of all the products that we develop as Marvell, from networking, switching products to storage products to compute and security products, and offers that same technology platform to our customers who are looking not for a standard product, but to leverage those building blocks of our technology, with customization to serve a very specific need that they have that's not otherwise met by a standard product.

C Christopher Banuelos 01:31

Kevin, I've heard you say previously, that silicon customization is increasingly important. Why is that?

K Kevin O'Buckley 01:38

Yeah, that's a super question. And I think for to answer that, you take a step back and look at the trends in the technology industry in general. And maybe the the most obvious example of that is a transition that's happened in our industry over the last few years, away from general purpose computing products and general purpose networking products, to customize products that are really driven by hyperscale level customers. I'll give you a simple example. You know, if you ask someone even just five years ago, Hey, what is cloud computing, I keep hearing about this, a euphemism you will hear often as well, cloud computing is just you're gonna use someone else's computer instead of your own computer. And to a certain extent in their early iterations. That was true. You know, cloud computing was using a general purpose Intel or AMD processor that was located in the hyperscale data center. But as the business scaled up, what our hyperscale companies in the industry found is, those general purpose computers actually didn't scale effectively for the diversity of workloads that the entire technology industry needs. So they found very quickly that they could drive specific customization in those products to better address the requirements of scaling at the hyperscale. So things like moving from a general purpose load store processor architecture, to a custom silicon that's uniquely designed to process certain database applications, or video process applications or offload of network transport off of general purpose processors and under special processors with security to do so. All of these are new demands that were only required in our industry, as hyperscalers really became the epicenter of how we as consumers and as companies interface with technology.

C Christopher Banuelos 03:33

For those out there who don't know, what's required for customization?

K Kevin O'Buckley 03:38

Yeah, customization, at the high level really is two things. Effective customization is two things. The first is unquestionably technology. We at Marvell, are a technology company. Our customers are technology companies, and they're looking for the latest and greatest advanced technologies. In the semiconductor industry. They're looking for advanced node silicon over the last couple of years. That's meant things like 5nm technology and advanced IP design using that technology. It's meant advanced packaging technology for massive large scale projects, including multi chip projects, and in the not too distant future, as, as we'll be announcing, or are announcing now at Marvell, we're moving that platform to 3nm. So first and foremost, what's required for customization in the infrastructure space is the most advanced technology in the in the industry. The second of two things that's required for customization really describes less a specific piece of technology and more a model for collaboration. When we're doing silicone customization, it's enabling a product for a very specific customer in a very specific application. And that means starting at the very early product development phase and continuing all the way through the Design and qualification phase of the product. Two companies need to act as one seamless virtual development team. And that means we have to design the architecture together, we have to work together on the specific implementation of that architecture. And we have to work together and getting that product qualified successfully and ramped to market as quickly as possible. It's a series of aligned interests, but it also requires a lot of expertise, technology and behaviors to pull off successfully.

C Christopher Banuelos 05:32

Kevin, what is Marvell doing as a partner? And what are some of Marvell's technology investments?

K Kevin O'Buckley 05:37

Yeah, to talk about that, let's use it as an example of something that, you know, is pretty visible in our industry today. And that's the the very rapid development and deployment of these natural language models that we hear about things like ChatGPT, and other infrastructure products. And enabling those types of machine learning shifts is extraordinarily complex. So we're making investments in things like memory access for absolutely unbelievably large scale training models with an incredibly large number of parameters to enable those, those language models to be more and more sophisticated over time. In addition to the memory investments that we're making, there's specialized investments in compute capability, contrasting historical investments, and compute, which has been about taking, you know, ALUs and typical math actions like like additions and others and doing them as quickly as possible. Machine learning compute drives, special math

functions, like things called multiply accumulate function. So we're making technology investments in 5nm and 3nm technology, to actually harden those math functions in as energy efficient and space efficient fashion as possible, so that our customers can develop in partnership with us the most advanced silicon solution using those memory and compute building blocks.

C Christopher Banuelos 07:06

Kevin, can you talk about some of the projects that Marvell is working on in the custom silicon space?

K Kevin O'Buckley 07:11

This is this is one of the most exciting and frustrating aspects of my job in the custom silicon space at Marvell, I'm fundamentally working on projects that will be branded appropriately as my customers so so I can't actually a podcast like this, say, hey, you know that product out there, you know, we develop that in partnership with a company that's for my customers to announce, but what I can talk about are some of the very specific applications that those custom silicon products are doing. So for example, we're working on and are shipping now products that deliver 10s of terabits per second of data in high end switching and routing. In both the enterprise and the data center space. In 5G and 6G wireless, we're developing custom products from silicon right at the antennas that you see when you're driving down the highway, all the way down through the value stream to the complex processing that happens in in base station and in centralized data centers. And in data center compute, we're doing things like the most advanced server class processors using Arm compatible instruction sets and custom fabrics, all the way through to you know, some of the most exciting things that people are talking about the industry today, like AI and machine learning (ML) applications that are being deployed at hyperscale clouds. All of those are examples of projects that we're both developing today, and shipping today, to our customers.

C Christopher Banuelos 08:37

Kevin, this might be a good time for us to transition our conversation into the recent announcement. Let's talk about that.

K Kevin O'Buckley 08:45

Yeah, so our recent announcement that we made was about bringing Marvell's technology platform to the most advanced semiconductor technology in the industry. And for us, that's 3nm today. We've had an extraordinary success bringing our 5nm based products to market over the last couple of years and are ramping and production in many cases across many applications. And our announcement recently is about bringing that entire technology platform to the new node, that 3nm technology node, which enables a whole new class of power and performance. And also as a companion to that a whole new level of integration capability with next level multi package. So we're doing things like developing new long reach, short reach, and ultra short reach interconnect technologies that allow us to scale compute and networking and storage and security applications. From very small scale edge requirements all the way up through the highest performance hyperscale level compute applications with single designs able scale up in multi chip packaging to allow a customer to develop a single chip that can be deployed in a wide variety of performance requirements in an incredibly cost effective and energy efficient fashion.

C Christopher Banuelos 10:13

And what are customers excited about? Or what should they be excited about?

K Kevin O'Buckley 10:17

Yeah, that's a super one. I think that the the buzzword and customization now in in our industry is scalability. And we're making some really exciting technology investments to enable scalability. And the importance of scalability is about you know, our customers from hyperscale, down to the edge all have common requirements for moving data in a network or computing data, you know, in small scale at the edge, and then massive hyperscale in the data center, and scalability in our investments around scalability allow our customers to develop a singular building block, maybe scaled for the edge. And through technologies like multi chip

packaging, and advanced chip, the chip and dilated eye interconnect technologies enable them to scale that building block up to massive hyperscale with a single design, and a different package designs that allow them to scale that design up from small to large. So it's an incredible, exciting innovation. And we've made tremendous investments in multi chip packaging, and interconnect IP to enable our solutions to scale to our customers requirements.

C Christopher Banuelos 11:26

Kevin, this has been a great conversation so far. Just wanted to end with what's next for Marvell?

K Kevin O'Buckley 11:31

Chris, this has been a blast for me to thank you. You know, I think as this has been a great opportunity for me to sort of express some of my enthusiasm for custom silicon. And I really am excited about it. Just as an example. You know, today, custom silicon is maybe meeting 25% or so of our of our customers requirements is represented in the products that we're shipping. And as we discussed earlier, the importance of this is just growing and growing and growing. So I'm really enthusiastic to be part of a team and a company that's making the investments in leading edge technologies that will let us take that 25% up and up and up and meet the needs of our customers.

C Christopher Banuelos 12:12

Kevin, thanks again for joining today's episode. Looking forward to hearing more from you and your team.

K Kevin O'Buckley 12:18

Likewise, Chris, it was a blast. Thank you.

C Christopher Banuelos 12:22

Thank you for listening to the Marvell Essential Technology Podcast. As always, please feel free to visit our website to learn more. And we'll see you on the next episode.



To deliver the data infrastructure technology that connects the world, we're building solutions on the most powerful foundation: our partnerships with our customers. Trusted by the world's leading technology companies for 25 years, we move, store, process and secure the world's data with semiconductor solutions designed for our customers' current needs and future ambitions. Through a process of deep collaboration and transparency, we're ultimately changing the way tomorrow's enterprise, cloud, automotive, and carrier architectures transform—for the better.

Copyright © 2023 Marvell. All rights reserved. Marvell and the Marvell logo are trademarks of Marvell or its affiliates. Please visit www.marvell.com for a complete list of Marvell trademarks. Other names and brands may be claimed as the property of others.