



- Patrick Moorhead: Balaji, it's great to see you. First time, Six Five Summit speaker, it's been so long. How are you doing my friend?
- Balaji Baktha: Patrick, yeah doing great. It's been... What? 12 years since we started our journey together in various forms, but thanks for inviting me. This is a great event. Really excited to be part of it.
- Patrick Moorhead: Oh, I appreciate that. Hey, we're bringing the biggest players in industry together and you are one of those. So thanks for signing up. And yes, it has been a long time. It's been 12 years since we talked about, when you did the first 64-bit ARM CPU for servers at Veloce and then later Applied Micro back when that whole thing was getting started and now look at that market now. You now have a company called Ventana. Can you tell us a little bit about the company?
- Balaji Baktha: Great. Well, yes, happy to. Ventana was founded by me and my co-founder Greg Favor. You know our backgrounds.
- Patrick Moorhead: Sure.
- Balaji Baktha: I've been doing this for 30 years in the high performance data center space. Drove lot of the successful products at Marvell and been doing that even before that. And since then, as you mentioned Veloce, Greg and I, we brought the first 64-bit ARM out to the world and we made it a data center class solution. And as we went through that experience, we've been able to work closely with lot of the leading hyperscalers and OEMs, and identify the key requirements for the next generation data center architecture.
- In particular, there are two trends that seem to be driving the go for strategy for these hyperscalers and OEMs. Number one, driving meaningful workload efficiency using hardware, software co-design approach, which fundamentally requires them to bring their workload specific enhancements and optimizations relative to the CPU itself and make it really close to the CPU through ISA extension capability.
- And number two is rightsizing compute, memory, I/O and other aspects of the silicon, and optimize it for workload specific requirements using composable architecture, disaggregated server architecture. These are the kinds of things that we've seen when we did Veloce, when I saw that, some of it could be done using the prior architecture, but a lot of it could not be done. And as we looked at it, we said, "Okay, we have to come up with a way to provide for that kind of an innovation that customers would like to do." So basically they have their own unique ways of accelerating some of these workloads. And we said, "Okay, the rigid CPU architectures of the past were not lending to that kind of requirement". And we thought it'd be a great idea to come up with something that's based on open and extensible and community based.
- And RISC-V happens to be that and we founded Ventana to bring RISC-V to data center. The company was started in 2019. We've been in stealth for a good number of years. And then we



came out of stealth towards the back end of last year. You know you were the first ones I spoke with about that and since then, we've been able to pretty much talk about the gaps that existed in the prior architectures and how RISC-V can bridge them and why Ventana is going to be the leader in bringing that to the market. So that's pretty much about Ventana and why we started Ventana.

Patrick Moorhead: No, I love it. The data center market has changed so much in the last 10 years. I mean, no longer are there these necessarily stovepipes of server storage, networking, and security. There seems to be a complete re-architecture of the enterprise. First and foremost, it's heterogeneous, right? Which is just all these different accelerators. Once Moore's law started to decline, we really needed to get on to something that was something other than, let's say an x86 CPU. And that's where all these accelerators came along now, along the course, we kind of pushed a lot of the effort over to the software folks, which has its challenges, but in the end, isn't it amazing how big CPU's and still are so important because they're so flexible.

They can run so much software out there. I know it may sound like I'm talking out of both sides of my mouth here, but we're heterogeneous, but big CPU cores still matter a lot. We've seen the history of x86. We've seen what happened with ARM. You were right at the very front of that with the first 64-bit ARM processor. But why was RISC-V? I think you hinted at it a little bit in your introduction of the company, but is there really interest in using RISC-V for servers?

Balaji Baktha: That's a great question Patrick, and some of the needs that drove why RISC-V, you actually touched on in your question itself but if you look at it, take a look at any computer architecture, any system on a motherboard, you'll see everything that's memory, I/O, storage, networking. Everything is standards based and open. CPU is the only thing that's proprietary. That's been pretty much proprietary architecture. Let me start all over again, Sorry.

Patrick Moorhead: Yeah, okay. No, it's good.

Balaji Baktha: That's a great question Patrick. If you look at any of the existing systems architecture, you'll see memory, storage, I/O, networking, all of it is standards based and you have multiple choices, implementations tend to be specific to various use cases, et cetera, but CPUs have largely remain proprietary. And when you have a rigid CPU in the middle, what you can achieve is directly proportional to just Moore's Law and the gains that you come from it. And as you mentioned, then that curve starts to flatten out. You got to go look for efficiencies somewhere else.

If you look at the data centers that evolved from bare metal to virtualized containerized, but then the next wave is software defined, which means the hyperscalers and OEMs would want to have a lot of their innovation, their proprietary capabilities, working natively with the CPU to drive some of those performance gains. So been Moore's Laws flattening out hardware software co-design, software design data center is the way forward, which means you need an open hardware platform and it starts with an open processor and an open source processor that allows you to add those capabilities while maintaining complete software compatibility through rigid base ISA compatibility. So make sure that the base ISA is rigid, capable of supporting all applications.



Patrick Moorhead: Yeah.

Balaji Baktha: And yet provide the flexibility to add extensions native to the CPU, that allows you to have those kinds of workload enhancements, which is the catalyst that drives your hardware software co-design capability is uniquely RISC-V strength.

So you've seen us dual processors. I mean, Greg has done x86 before, K6 and we've done ARM and having been in the trenches, having been through this exercise for about two decades, we know it's time for open hardware. It's time for open processor. And that's why RISC-V is the right candidate. You get the benefits of ecosystem strength that comes from an open ISA, much like your open software ecosystem that kind of really exploded with open source revolution, similar things will happen here. And that'll start with RISC-V.

And that's why RISC-V is the best thing that came along in the overall high performance space a long time, not only just high performance, it just solves all problems. One of the things that I can tell you, I learned from my colleagues at RISC-V International, that 12 billion RISC-V CPUs have already been shipped in products, 12 billion. So it's there and it's been broadly adopted across the spectrum and it's the right candidate for the next generation data center architecture.

Patrick Moorhead: Now, I appreciate the thorough answer. As we've seen out there, it's no longer a hardware game and it's no longer a "Hey, I'm a chip maker. I can do everything myself". So is the RISC-V ecosystem ready for prime time? Partners around it, software, middleware and things like that? I wasn't following the software as probably as closely I was following the hardware, but I think it took ARM at least a decade to really create something that was compelling on the... I would call it the mass market side and maybe this doesn't have to be mass market. So where is the ecosystem?

Balaji Baktha: No, again, good question. So if you look at it, for the other ISA, the one that you just talked about, it took a while because it was predominantly mobile centric for it to go to other spaces, took a while, but in our area, I mean, if you look at RISC-V, it's been around since 2012. So it's not like a brand new eyesight, it's already been around for about a decade now.

And the big principle difference between ARM and RISC-V is this, RISC-V's ecosystem is driven by the end user community, today if you talk about how many end users are there, how many members are there? It's about 2,500, getting close to 3000 and it includes who's who, all the large hyperscalers, all the large OEMs, all the semiconductor companies, all the research organizations, universities, everybody is developing rich set of software components, the building blocks to make RISC-V happen in various markets.

For Ventana's point of view, as we started Ventana, my particular goal was to make sure that we understood the software ecosystem dependency, even before we launched the company. So we spent about a good six months to a year, working with one of the largest OEMs in the valley.

Patrick Moorhead: Yeah.



Balaji Baktha: And understood what is that they would need to see happen from an ecosystem readiness point of view, before they can productize RISC-V. And we had set out a goal for ourselves that we will have all those building blocks in place even before we had our tapeout. And I'm happy to tell you, we pretty much met the goal. So in order for the kinds of applications we are going after, for them to productize, for them to realize RISC-V in production environment, all the requisite software components are here and now.

And so that's where RISC-V's ecosystem is, as it relates to the applications we are going after. In certain other cases, it'll continue to grow. I mean, it's not like RISC-V is going to be able to get there overnight. In certain other areas, it will continue to evolve. But the target applications that we've been looking at, it's pretty much here and now, and again, I told you 12,000 CPU, sorry, 12 billion CPU cores are shipped. They all run some kind of software. So it is ready to a large degree.

Patrick Moorhead: Yeah. Gosh, it's almost like we learned something the previous 15 years or something. I say that with sarcasm, of course. But as industries, we do learn and we get better at bringing up things like this. In the introduction, you talked about chiplets these days and listen, I mean, you can't go to a semiconductor conference without discussion of chiplets. I mean, can you talk a little bit more about your chiplet strategy?

Balaji Baktha: As we launched Ventana, one of the first things that we thought about is how can you enable? Who are your customers and what problems are you solving and how do you make it easy for them to solve those problems? And right from day one, it was very apparent to us that chiplet was the right way to go. Why? There are several compelling benefits. The first one is customer driven innovation in the form of workload accelerators and composable architectural solutions by rightsizing memory, I/O. These are key requirements for hyperscalers and OEMS, and chiplets are the best way to do that. Chiplets enable rapid productization of these capabilities. On the other hand, if you want to do it in Monolithic SoC, it takes three, four years to build something and you have to predict where that puck is going to be four years from now.

Patrick Moorhead: Yeah.

Balaji Baktha: And when you do that, you're going to oversize everything because you don't want to be wrong. And when you do that on a five or a three, assuming your memory interface changes, your I/O changes, something else changes, all of a sudden you're having to re-tapeout the whole darn thing. That's 11, 15 million dollars just for the tapeout.

If you can modularize the design, move compute to the most advanced process geometry like a five or a three, just the compute chiplet and the rest of the building blocks that tend to be analog or mix signal intense, keep them at an n-2 process node, like a 12 or a 10. When you keep it there most IPs that you need to integrate, they tend to be more robust and Silicon proven in n-2 and they yield better and they cost less. So for you to be able to do the moving parts in an n-2 process geometry and have customer specific IP incorporated into them, and then combining these two into a chip... I mean, multi chiplet based system using inexpensive packaging solution, that's what chiplet is all about.



So the benefits are really, really compelling and we thought about it from get go, and chiplets have been proven to be very good way to productize, thanks to leaders like AMDs. So if you can learn from that and make it much more open and parallel interconnect centric and minimize latencies between dies and make sure that you get best of both worlds, performance somewhat comparable to Monolithic SoC, but the flexibility that's uniquely chip centric, if you can bring those two worlds together, then you get everything that a hyperscaler and OEM cares about and that's been our founding premise.

That's been our go to market strategy from day one. And so Ventana's chiplets, 16 cores, 8 cores are all based on, the first generation's going to be based on ODSA BoW. And you've seen this announcement from Intel about UCIE, which we are really excited about.

Patrick Moorhead: Yeah.

Balaji Baktha: And we'll be one of the first ones to bring UCIE based chiplets to market, our compute chiplets to market. And so, our announcement with Intel, that was out there for everybody to see, so we are working very closely with them and to be able to bring UCIE chiplet capability to our offerings would make it even more compelling. So chiplet is way of the future and we are going to be right at the forefront of that.

Patrick Moorhead: Yeah, it's funny Balaji, probably six or seven years ago, let's just say 10 years ago, the notion of doing something in a package was negative, right? People would be like, "Oh, you couldn't fit it on the die. You were die limited. You were radical limited. And therefore there's something wrong with it". Because the previous 30 years, you take it off the die and it slowed down.

And then what happened from my point of view is Moore's law slowed down, the economics of foundries significantly changed, but most importantly, packaging became a first level citizen in the engineering team. Okay? And listen, guilty, I worked for a chip company, me say starting 20 years ago plus and we would get stuff done and we'd pass it off to the package team, okay? But now with the tens of billions of dollars spent on packaging technology, I believe you can have your cake and eat it too. Right? You can have the performance, you can hit the power, little bit more expensive, but a lot of that is made up because you're not doing the entire Silicon in bleeding edge, right?

Balaji Baktha: That's the hard part.

Patrick Moorhead: You're only doing the main... Sorry, only part of the SoC in bleeding edge and good die per wafer is higher. I mean, it just, it all ends up working and in Intel, it has been, in my opinion, the leader in moving this forward but I think as importantly, educating. Quite frankly, people like me on the arc of the possible because I was completely skeptical. So net-net, I think your intersection of Ventana and your strategy and when you're going to hit the market, makes just so much sense to me here. So Hey, take that as a compliment. I think you nailed it.

Balaji Baktha: Thank you.



Patrick Moorhead: Now, you talked a little bit about your business model before, but would you mind, for the audience, doing a double click on that? Are you going to be using chiplets mostly for the data center market or for other markets? Where are you focused and what's your business model?

Balaji Baktha: Great, so let's talk about where are we focused. Ventana is the high performance leader in the RISC-V space. We told you we've been in stealth for about two and a half, three years. And so we used that time quite effectively and wisely to gain that lead. And the target markets tend to be data center, 5G edge compute, high performance, networking, storage, security, appliances, and routers switches, and that sort of thing. And then fully self-driving automotive applications that require high performance compute and eventually clients as well but if you look at it, those are the kinds of markets we are going after, the market time for that is just huge. It's about 90 billion or so. So as you look at it, the market opportunity is quite large. And how do we go to them, I mean, how do we address what's a go to market? So for anything that's 8 cores and above, we believe chiplets are the best way to go.

Patrick Moorhead: Yeah.

Balaji Baktha: And most applications that are data center, edge compute and so forth, they tend to be in the 8 or plus space, so there chiplet is way to go. And if you have applications that are 4 cores or lower, you want to be able to do... Chiplets may not be good candidates because they add more cost to your overall bill and material cost, et cetera, Silicon cost, so the best way to go there is to license an IP, to be integrated into customer's SoC designs. So we do both, chiplets for 16 cores, 8 cores and above, and then in a typical customer socket, you could actually see up to 8 Ventana chiplets to get to the compute densities per socket that they're targeting, and as few as maybe 1. So an average could be 4 to 6. So that's kind of the chiplet base.

And for those customers, who don't want to do their own chiplets, their own I/O hubs and their own products, we also offer ASSPs, standard products in 4 chiplet and 6 chiplet, 8 chiplet versions. So we have standard product offerings as well. So three different ways, chiplets plus a reference I/O hub design for hyperscalers to create their own version of the device, cores for those who want to do their own SoCs and for customers, predominantly OEMs, ODMs who understand the value of it, who want to build products quite rapidly, they would use our standard parts. So those are the three different go to market vehicles for Ventana.

Patrick Moorhead: Yeah, I'm glad to hear about that flexibility. Now you and I both know to prove out the chiplet you had to bill, you had to test it out with an SoC and then in between that are the different flavors that you're offering. But I spend a lot of time Balaji, with your potential customers and I will tell you that they are looking for different implementations. Now the great part though, I think is that, if you start with a chiplet mentality, flexibility, I mean, as long as you have the right interconnects is really off the scale of what you can do. And you've appropriately sized your memory controller to be able to hit all the cores and that flexibility allows a company like yours to play in so many different markets. Right?

I think people automatically, when they think, "Okay, data center, storage network, appliances, automotive, 5G edge", they might be thinking cost, cost, cost, because quite frankly, a lot of



other chip makers are doing exactly that, right? They have Monolithic designs and they're pouring literally tens and billions of dollars into designs to do that with very little leverage across cores. You mentioned AMD and while AMD isn't necessarily doing full chiplet now, they kind of kicked off and showed that you could have a flexible, I'll call it a 2D design, but AMD had money for one design and that's it.

But somehow they took that design and put that into five or six different implementations with probably one-tenth of what it would've cost me and what it did cost me in the year 2001, when I started in the chip industry. So a good timing, right? Sometimes they say, I'd rather have better timing than be good, but quite frankly, I think you've nailed both of them. I think your timing's good and I also think your roadmap looks good as well. You're the first big core in RISC-V to show up. So, from my point of view, you're the only game in town, not that you're going to be the only game in town forever, but you definitely have the jump on the market.

Balaji Baktha:

No, Patrick, thank you. You're spot on about the chiplets and us being first to market with real high performance offering the RISC-V space to enable these slew of applications and your point about reusability, building one chiplet and then the parts around that particular compute chiplet be on another chiplet, that gives it the personality to go after data center, 5G edge, open RAM, networking, security and auto... That is the biggest benefit of chiplets.

And we saw that at the very onset of founding Ventana. And so that's why chiplets has been a key. I mean, not only are we building world class processor cores, we putting them into a world class chiplet and to do that, it means you have to put together a world class fabric that connects all these CPUs together, come up with the best cash hierarchy, memory hierarchy, day to day interconnect, we've solved all these problems and Ventana, for a small company being able to do this, puts us at a unique advantageous position.

So we lead the pack by a huge margin. And our goal is to not just rest on laurels, just as we look at the strengths of RISC-V being aligned with this open hardware software centric designs, we are able to sort of continue to innovate, listen to customers. And the first generation is going to be a 5 nanometer part. And we are already working on our BT two second generation, which is a 3 nanometer part.

And continue to exploit the strengths of RISC-V, open extensible capabilities, plus the chiplet based implementation and continue to provide compelling single socket performance at unprecedented performance per watt per dollar metric, in a way that drives data center to the next growth cycle, that's what we are about. And we're doing really well, as you said, timing sometimes can be a huge advantage, we bet on it. It was like when no one thought RISC-V could go there, we bet on it and we're here and happy to hear your comments, your endorsement, and look forward to working with you and your team and with the industry at large, to make RISC-V a reality and drive it quite rapidly across many applications successfully and in so doing, make Venatana a very successful company.

Patrick Moorhead:

Listen, I love cool tech, but you know what I love better? I love industry disruption and the combination of RISC-V plus chiplets plus a unique business model is absolutely a disruptive force.



And man, I'm glad I'm part of this industry, this is fun. And listen, Balaji, I think this is absolutely the deepest that we, anybody has ever heard of Ventana Micro. And I appreciate you sharing all of this information with the audience today and thank you so much for coming on for the first time. And we'd love to have you back on the show next year.

Balaji Baktha: Great. It's been a great show and happy to be part of it and look forward to working with you in the future as well. And Patrick, thank you and thanks to your audience as well.

Patrick Moorhead: Thanks.