

**Tony** [00:00:04] Welcome to Code Together, a podcast for developers by developers, where we discuss technology and trends in industry.

**Tony** [00:00:11] I'm your host Tony Mongkolsmai.

**Tony** [00:00:18] One of the main topics we've covered in this podcast has been how one API addresses some of the challenges in heterogeneous and accelerated computing. One of the most rewarding parts of hosting this podcast is meeting the innovators who are actually solving interesting problems. To that end, throughout the year, I will be talking to a variety of innovators who are part of the oneAPI for Startups program. We will be explaining the challenging problems these innovators are solving and highlighting the solutions their companies are creating. Today we are joined by Charlie Wardell, the president and CEO of Digital Cortex, a company with a platform that abstracts alternative computing devices so each task can be executed on the hardware for which it is best suited with simple API calls. Charlie has a passion for technology and has been fortunate to work with some of the brightest technologists at Teradata, Netezza, Vertica, NVIDIA, AMD, Intel, and currently holds patents related to distributed computing, text based knowledge mining, rule based evaluation of documents and linguistic processing frameworks. Welcome to the podcast, Charlie.

**Charlie** [00:01:18] Thank you so much. I appreciate being here.

**Tony** [00:01:21] So let's start off with what is Digital Cortex? So I gave the one liner of what it is, but why don't you dive in a little bit and talk about how Digital Cortex can help customers?

**Charlie** [00:01:33] Yeah, So, so back in my text analytics days, we were doing some text mining looking for topic detection in large volumes of unstructured data. There's a pretty large hedge fund that had the idea that, you know, we could find Alpha using text analytics alone. So they wanted to pursue that idea. And in order to do that, we had to take about 20 years worth of historical, unstructured data, and we had to identify topics in it, measure those topics, and they would use those in their algorithms to figure out, you know, what was meaningful and the weightings and things like that. My job was to figure out the topics because I have a text analytics background, so I put it on our distributed platform that I designed, and it worked really well. About 130 servers, about seven days of processing. And when the results came in, they basically did their regressions and all their backtesting and they said, Yeah, well we need to tweak the topic model, so can you go ahead and rerun that? Each time we ran it, it was about \$16,000, right? And they lost about a week's worth of processing time, plus about two or three days worth of backtesting time, only to tweak the model and have it run again.

**Charlie** [00:02:55] So I said, you know, I worked with some early technology with Netezza. They they used this FPGA technology to do this divide and conquer. And I've been fascinated with it since 2014. So I think I could do this with some FPGA technology. Now, this was before oneAPI, right? So, you know, we had to find some IP cores and some FPGA cards that were kind of in that genre. And I did it. I wrote all the device drivers to it and we did the text analytics for the most part distributed across, you know, a few FPGAs and I brought that 130 servers down to like two, and it went overnight, right? So they were estatic.

**Charlie** [00:03:39] And then the thought came to me that, well, you know, CPUs have reached their theoretical max. They're, you know, you can't fit any more transistors on that. Right? So everybody's using CPU's and they just throw more CPUs at the problem. But there's this whole market of XPU's that are coming out. You know, people need to start exploring that, right? So I started looking at all the players in the field and getting really up to speed on all the alternatives to FPGAs and GPUs and VPUs and NPUs and all this wonderful stuff. And I said, you know, they're really hard to use. What if I created a platform as a service that abstracted all those complicated technologies and basically create functions as a service backed by accelerated hardware devices? And that's how Digital Cortex was born. It was based on the experience of this text analytics problem I had and then understanding that, you know, there's gold in them, they're hills, right? You need to jump on this XPU bandwagon. You can get a lot of processing power and save the planet while you're at it, right? You don't need as many CPUs to solve the problem as you did with a few FPGAs.

**Tony** [00:04:53] So you started in the world of doing these text analytics. Where are you looking to go? Because I'm assuming that you're you're going to want to support things more than just text analytics that I don't know if that's your core competency or are you already branching out to have kernels that work in other spaces?

**Charlie** [00:05:10] Yeah, So that's a great question. What we're developing is a marketplace for kernels, right? So think of the Apple App Store where you can monetize apps. So I would like all the kernel developers in the world to, you know, park their kernels on our marketplace, monetize it. We'll take a small, small percentage of that transaction that goes through it and, you know, just open up the world to, you know, a kernel marketplace. But with that said, you know, we have to, you know, walk before we run. So what we're doing is we're going to be seeding the market with a certain amount of these functions as a service in the text analytics space and image enhancement. We are working with someone for image recognition, doing reverse image searches. We do things like redaction, you know, personally identifiable information off of X-rays, for example. Right. So it's those types of use cases where it's computationally intensive that that we're focused on initially.

**Tony** [00:06:10] Because you're providing the platform, but you're also writing kernels. I'm assuming that you'll have some type of SDK or some type of way to help people integrate their kernels into your platform.

**Charlie** [00:06:24] That's correct. Yeah. Yeah. So instantiate a class, overwrite the run method, put your kernel here and we abstract everything else. So, you know, trying to make it that simple and then we expose that function in a generic way. Right. It's the same interface to call all functions. The message is structured in a way that, you know, the function will use what it needs out of that message structure and it runs it and it returns the result. It's like ping pong, but based on hardware acceleration, it's pretty cool.

**Tony** [00:06:57] And how did the kernels know... How does the platform know where to execute these kernels? Because obviously somewhat if somebody's submitting kernels to you or you yourself are creating kernels, you're going to need some baseline to understand what performance and like going to get and where. So your platform can make kind of that right decision for the customer. How do you guys enable that workflow?

**Charlie** [00:07:17] Yeah, So initially what we're doing is we're in we do have a marketplace app framework that we're using internally and we're able to define some of

the metadata with regard to the function, right? You know, where where we think it's best suited to run and whether or not it even should be, you know, accelerated. There are some tasks that, you know, run better on CPU. So that's the initial step. But that the real answer to your question is to have something like a cost based optimizer that gets smarter and smarter as time goes by, run it, you know, capture the metrics, you know, run it again, run it again, and then make decisions based on the, you know, current resource utilization, what's available and things like that. And oneAPI does a great job at, you know, you know, selecting the best device. You know, it does a great job. And, you know, it's very easy to prioritize devices within the platform.

**Tony** [00:08:21] And as you build out this type of abstraction, my first thought is thinking of how things go to the cloud. And I know you guys have some type of cloud offering. Is it something where the customer is going to come to you and then you're going to deploy it to the cloud? Or is it something where I potentially can run this myself on-prem? What's kind of your model for getting access to the right hardware for your customers?

**Charlie** [00:08:44] Yeah. So about a month ago, you know, as we're talking to some of the VCs, they're asking us, well, what's your moat? What your moat? What's your moat? Right. And as a CEO, I'm like, you know, is it as simple as just throwing an API in front of function calls? And it really isn't. You know, the marketplace will gain the adoption, the performance and the ease of use will kind of speak for itself over time. But what I started thinking about was, well, you know, initially we were going to be a cloud. You know, you sign up, you can basically create your workflow using marketplace items and then maybe we'll make an IDE plug in so people can, you know, easily develop their kernels and upload them to the marketplace, and it is a workflow designer that that glues them together. Then it occurred to me, how are you going to compete with Amazon and how are you going to compete with Google? Who is obviously putting FPGAs in place today. So the answer to that question as well, our platform can be downloaded, right? So you can download our platform, you can install on your hardware. We hope to be in the Amazon marketplace and Google marketplace so you can basically instantiate our platform within your cloud and then we'll sniff out all of the accelerated devices on your network and basically utilize those. We also have a cloud offering, which basically you don't want to get involved in any of that, it's fully managed. You just call these API endpoints. It's that simple.

**Tony** [00:10:21] I know since you guys are part of this early access startup program with Intel that you have access to the Intel Developer Cloud, which our CTO, Greg Lavender, talked about last year at Intel Innovation. Can you talk a little bit about whether that helps you, how it might enable you to do better things and how you might utilize that in the future?

**Charlie** [00:10:41] Yeah, so the work we're doing on the DevCloud right now is to basically, you know, test the kernels in test performance and see, you know, the units of parallelism we can get with the number of work items and workgroups and things like that. So it's a big, an awesome R&D effort that DevCloud allowed us to utilize without having to stage all of this infrastructure ourselves. You know, ultimately, we'd like to be a possibility to be a front end to some of the DevCloud stuff, right? So there's no reason why this couldn't run on DevCloud itself. The DevCloud has been pretty awesome.

**Tony** [00:11:17] And as you move to the public cloud, at least your public cloud solution...are there any limitations that you have running, for instance, in an AWS or GCP in terms of hardware accessibility because they do virtualize things? Is that something that

has driven you kind of to have that that on prem type solution? Is there a delta there or do you kind of get the same behavior between I run in the cloud versus I run on prem assuming that I had the same hardware?

**Charlie** [00:11:46] Yeah. So early on when I was developing this technology for this hedge fund, Amazon had this thing called I think it was called F1. It was like it was FPGA in an instance, and I can't tell you if it was bare metal or virtual. I can only imagine, you know, I'm not sure how they expose the FPGA through, you know, virtualization, but let's say they did. I had the ability to access it directly. Very expensive. It was like, you know, 60 bucks an hour at the time, just going back a few years. So I'm not sure what it is today. But that's my hope is that as you start as they start bringing more and more of these accelerated devices, they're exposed on some sort of fabric that you have access to. And, you know, be nice to, you know, just look at a one PCI bus and see what's available to you. Right. It would be nice. I'm not sure where they're going. And that's why we're offering a downloadable version.

**Tony** [00:12:46] Yeah, it's interesting because I saw that Lambda Labs is starting to offer cloud service. They do kind of AI training as a service now, so they're moving away from hey go run our platform, our buy our platform. But actually we're going to have a cloud service for that. I know obviously Intel I mentioned is doing that and Nvidia is doing that. And I think in my head as I think about it, you know, when we look at how people are building kind of these very fast accelerators and specialized hardware, it seems like it's people are moving more towards providing the accelerators as a service, kind of like you guys are already pioneering versus saying go buy a lot of hardware and then run it on prem. The cloud model really seems to make a lot of sense when it comes to... I'll say expensive compute because it's very hard to afford that and then kind of keep those machines running at full speed and utilizing them to the maximum.

**Charlie** [00:13:43] Yeah. And you know, these cloud service providers, they're they're not necessarily incentivized to make it faster. Right. They're they're really not, you know, if, you know, bringing 130 servers down to two, you know, was a good thing for them if everybody decided to do that. The way we handled it, though, was we had the FPGAs in our data center, so we had a small little colo facility and we had a, you know, a rack of servers with FPGA in it. And then we had a dedicated pipe to Amazon that extended the client VPN, VPC. So, you know, we, we looked like our servers were on their, their network and completely accessible and, you know, locked in. So that's the way we we went initially with this client who was obviously on the cloud. But I would imagine that these cloud offerings are going to have to get involved in accelerate. It's coming, right? It's absolutely coming. They need to stand them up and and make it available. I want to be the software that sits on top of it. That's all.

**Tony** [00:14:54] You've talked about how your software platform abstracts the hardware away from your end user and allows you to run kernels in the best place possible. Is there any hardware coming down the pipeline that you're exceptionally excited about?

**Charlie** [00:15:09] Yeah. So the one thing that we're talking about Digital Cortex and we're very excited about is the potential use of RISC-V. And, you know, maybe it's not considered an accelerator, but it does apply to brute force parallelism. there are some tasks that you just need. extreme parallelism across as many servers as possible. RISC-V could be a very low cost, low heat, low energy way of doing that, similar to the way Netezza created all these snippet processing units. I could see a RISC-V board doing the same thing and then just saying here you go to your compute and give us the results back.

So massive parallelism for things that are not kernel specific, right? So I can see RISC-V for that.

**Tony** [00:16:03] And that's something that interest Intel is interested in as well. Obviously people know that we have some RISC-V initiatives and although obviously Intel is heavily invested in x86, we recognize that there there are other types of compute platforms that are interesting to the market and there's different profiles for different types of hardware that makes sense. One of the things that I always have in my notes is how Dave Patterson said that the world is just going to kind of become a world of specialization because transistors can only get so small. We've got to come up with more novel ways to use our transistors in a way that makes sense power performance wise.

**Charlie** [00:16:38] Yeah, Yeah. I'm drinking that Kool-Aid. Yeah. And I, I love the fact that you guys are seeing that as well. It's just amazing when you when you see what is possible. I remember doing queries on a database when I was when I was trying to sell Netezza and Vertica and all this other MPP technologies, database technologies to customers, I would load their data, I would run the query and they would look and they would say, No, that's cached, that's not possible. They can't be that fast. I'm like, No, it's not cached. You know, I would show them and then it opens up the capability. What now? I don't need summarization tables. If I don't need summarization tables, I don't need batch jobs to populate them at night. If I only those batch jobs, I don't need this tool. Right? I could just look at my transactional data and it performs this well. Yes. So when once you understand what's possible, right. It just opens up incredible opportunity.

**Tony** [00:17:43] ASo your kernels...We talked about image enhancement, image conversion. AI obviously is one of the most talked about uses of accelerators right now. It used to be a probably things like rendering, like things that you would see out of Pixar and how we're getting cool things out of computer graphics now. It's definitely AI: Generative AI, Stable Diffusion, ChatGPT. What are you guys thinking about in that space? Because obviously you guys want to be a platform for accelerated compute that is the popular cutting edge of accelerated compute. How are you guys thinking about that space?

**Charlie** [00:18:20] Yeah, so we're giving a lot of thought to that space. One of the things that we've noticed on the FPGAs is the ability to do Inferencing of Tensor models or, you know, PyTorch models on an FPGA. So inferencing is a low hanging fruit. So now we start thinking about the marketplace and saying, okay, so we can, you know, in our marketplace, we can allow a person upload their model and put an API in front of inferencing, these very specific FPGA accelerators that do inferencing. As far as the market is concerned, I'm not so much interested in the creation of the models as so much. I am interested in the, you know, lambda calls to that, that our hardware accelerated. We think I was talking to our COO this morning that people are going to want to create their own models by extending, the base ChatGPTs of the world and everybody's going to have their own. They're going to have their own domain specific they're going to be tuned. They're going to be happy with it, and they're going to need a place for it to run. I want to be the place where they run it, not necessarily where they create it.

**Tony** [00:19:33] And as part of that, whenever you're building any type of solution, there's obviously the training part that the hardware matrix multiply, convolution, which is we know is hard or expensive. And that's what most people think about when they think about the AI, or at least people who are doing it that's what they think about because that's the most compute intensive. But there's a lot of work that needs to go into making an AI solution successful. And with something like ChatGPT, there's a lot of text, there's a lot of input that

needs to be done and translated into a way that the model can actually learn. Is that a space that you guys are targeting as well? Because it seems like something that would make sense if, if people are trying to do kind of text analysis, I need to take that text and somehow figure out how to feed it into a model.

**Charlie** [00:20:22] Yeah. Eventually will be full force in that space. Right now what we're doing is we're taking an existing model and we're basically enhancing it with domain specific text. Right? So, you know, you think about a particular vertical, I don't know, hotel and hospitality, you know, grab their base models from ChatGPT, for example, and then, you know supplemented with the host hotel and hospitality data set that you know you have internally. That's a much easier lift than creating the base model. It's just a much easier lift. So we facilitate that. We already do. We already facilitate the ability to go ahead and enhance the existing models that are in place. But creation of the models, you know, that's an area that is on our roadmap. It's just not, not the initial one.

**Tony** [00:21:17] Yeah. And that's a very tough space too. It's you can see that a lot of these large language models that people are using come from very few players in the market that tend to do a lot of these things. It's a very specialized space.

**Charlie** [00:21:30] You're getting into like, you know, massive, massive parallelism across, you know, many, many GPU's use. And, you know, it's just not something that we can tackle, right? So we're going to leverage what you've tackled.

**Tony** [00:21:45] And if you had to pick the most interesting end user use case of what people have done, with your platform or service, is there a particular one that stands out in your mind that is really interesting and novel use of your technology?

**Charlie** [00:22:02] You know, we want it to be just generic and pipeline acceleration, right? And what's starting to be interesting and, you know, we've had a few people approach us with regard to natural language understanding, and it's different than the with different approaches then, you know, ChatGPT it's more and trying to understand the contextual nature of symbols. A stop sign is a symbol to stop, right? That's language. So we're starting to see people move into this direction of trying to create these ontologies of concepts and relationships and utilize that as its base of knowledge. You know, autonomous vehicles have a language as well, right? You know, the lanes in a road, you know, the signs, the traffic lights, those are all language, right? If you can model that language contextually, you could do some amazing things.

**Charlie** [00:23:08] The problem with it is that in some of these cases, the parsing of a sentence is so extraordinary that it takes about a second per sentence. So the question is, can you accelerate that? What can you do to accelerate that? And we're working very closely with a few firms now, one specifically that has created this incredible method of... They're doing part of speech tagging on steroids. It's pretty amazing and storing all of those symbols and relationships. And we're looking for the hotspots where we can accelerate that and be part of a different movement from ChatGPT, which is still in the AI space. But it's a natural language understanding, which is not necessarily based on neural networks. So that's interesting.

**Charlie** [00:24:04] The other area of interest we have is, is in health and medical. I would love to be diagnostic, right? I would love to, you know, not necessarily call it, but assist and say, hey, doc, you might want to look over here on this MRI. You know something's here. You might want to check that out. To be able to do things like that are very exciting to me.

Reverse image searching is another incredible opportunity, especially in a space of child trafficking. Right. So, you know, it's sad to say, but, you know, there are pictures of hotel rooms where people have been exploited. And you know what? If instead of looking for the missing child, you're looking for a particular room and then you can tell where that room was and what hotel, what time this was taken, things like that. You just more in giving evidence towards, you know, trying to track missing children. So, you know, it's going to require massive amount of work to get the images in the feed and stuff like that. But I would like to be part of accelerating that search. And we're working with a company that that does reverse image search right now.

**Tony** [00:25:25] That's pretty interesting. A lot of diverse use cases that actually potentially solve problems that we care about in the world. Probably we don't think about very much. So that's pretty cool, actually. One question I have when I think about kind of all the large amount of compute that you guys potentially are going to enable. One of the biggest challenges that we see nowadays as we build out scale systems is how do I get the data there, How do I get all of the information I need from wherever I've got it stored into my compute? You guys are kind of a compute service. How do you guys connect up to the large data source to actually access all this information to compute on.

**Charlie** [00:26:08] Yeah. Data locality. Right. And so you know this this this was what made Teradata and Netezza so incredible is that they had, I'll use Netezza as an example because they know it really well. Netezza is a database appliance that had these cards that went vertically into a rack and these cards, they supported a hard drive, not an SSD but a real hard drive and it had a CPU on it, It had an Ethernet port on it and had an FPGA and it had RAM. Right. So what they did was they took their data when you loaded your data and they sharded it across as many of these cards as they can cram into a single rack. And I think they had about 100 or 106 of these they called them snippet processing units (SPU). And so if you had 100 of these and you had a billion records, right, you would record one on the SPU one and record two on SPU two and paper thinly sliced. And then when I launched my query, all 100 of them go simultaneously against its local data set and says, I have an answer, I have an answer, and then then the head node would reduce it all into a result set. Unbelievable performance. It was just brute force. It was unbelievable. They optimized it with indexes later on, but it was unbelievable. Queries that that that took you know, 46 hours were running in 11 minutes in my experiments. It was just just incredible. So I understand data locality but not all you know when you're dealing in the world of parallelism based on how complex your computation is, sometimes the latency of moving the data to the node is is small in comparison to the processing time. So in the example of the hedge fund, what we did was we went to their S3 bucket, we pulled a packet of data distributed across our cluster of FPGAs and each one ran autonomously and independently on its local data source. But there was that initial move the data over, divide it up, run it locally and then put it back. So that's the approach that we have been taking.

**Charlie** [00:28:31] We're not like Hadoop style where we're going to have, you know, our own little database on every single node. We may look into that, but we're looking at some advanced accelerators, storage and retrieval stuff of fabric of storage and retrieval that could potentially be centralized and provide low latency. So we're looking into that. But I would like to be able to have, you know, data locality by customer and have their own little repository so that we don't have to worry about the sharding.

**Tony** [00:30:39] So one of the things that many developers who would be listening to this would be interested in is how you went from working on a problem as a developer to

deciding I want to make the leap and I want to actually create a company and found a startup. Can you talk a little bit about what that journey has been like for you, however you want to talk about it, whether it's why you decided to do it or the challenges that you're facing? Maybe all of it.

**Charlie** [00:31:07] Yeah. Yeah, that's that's a that's a great question. So I like to stick to my knitting and say that, you know, I'm really comfortable in the tech space. I would say that I am an introvert to the extent it is that I will look at your shoes while I'm talking to you as opposed to my own. So being a CEO of a company, you know, it just requires a different thing. Things that I didn't even think about. Now, with that said, for the prior 11 years, I was a co-founder of a text analytics company where I was the CTO, and I basically went behind the scenes and I did my things, I created my patents, I did the demos and all that other stuff and I stayed out of the business side for the most part, but I was in every Shake the Money tree meeting. Right. Because they wanted to know about the tech. So I did have a little experience in business to understand the types of questions that were asking, you know, the obvious stuff, like, you know, what is your moat? And you know what it was, you know, what is your use of funds and things like that. So taking the leap to say, you know what, after this company was sold and trying to figure out my next thing, I basically said, you know, all of the people in the tech space were really interested in the tech. And I think that if I, you know, put on a CEO hat or at least pretend to, I could basically build some confidence because not only am I the tech guy, but I'm the strategic direction of it. But with that said, we have an incredible COO who has his MBA and, you know, he is just brilliant. You know, he he basically cleans up after me and he's my executive function. So he tries to keep me in the closet and slide pizzas under the door. Let me, you know, build out the core of our framework. And then when I get the core framework to a point, I will go out and I'll become the evangelist. Right? So I'm going to be the evangelist of oneAPI/SYCL/FPGA/XPUs. I want to create as many relationships as possible with vendors and see what technologies that they're there creating so we can look at bringing it into our environment. He's responsible for, you know, all of the handshaking and business dealings and financials and stuff like that. So yeah, I am CEO, but I'm not COO, right? That that's it's different. It's different.

**Tony** [00:33:48] Is it something that you would definitely do again?

**Charlie** [00:33:51] I would, yeah, I would. And if it fails, I'm going to do it again. And if I'm going to do it again, it just it I'm an innovator. I've been doing this for, you know, 25 plus years, looking for incredible technologies. And every once in a while, every once in a while, there's this wave. And if you catch it just right, you just may you just may go for a great ride. I think the XPUs are incredible wave. And I am so on board with what one API is doing. Just imagine being that oneAPI to all of these XPU platforms that you know no vendor lock in, you know and it's it's. Still, it has a long way to go. But where the vision is so amazing that it just. Yeah, why not? Of course.

**Tony** [00:34:42] Yeah. It's actually great to see how excited you are about oneAPI because obviously my job is to be a oneAPI evangelist a lot of the time. It's funny because when you think about it, oneAPI is exposing hardware and Intel is obviously very invested in exposing hardware and accelerating hardware. One of the places where it's challenging is trying to make sure that the people who want to utilize the hardware can use it in a way that's as easy as possible.

**Charlie** [00:35:10] It will be amazing when these device manufacturers develop devices that are oneAPI compliant, right? You know, it would just I could just imagine them all



coming out and then giving access to, you know, this army of developers to utilize these XPU technologies relatively easy. So not everybody wants to get down to that SDK level. You know, moving data between buffers and stuff like that, you know, they're used to setting variables and running functions. So that's what I want to for them. You know, you go ahead and set your variables, run your functions. I, you know, we'll have the SDK, obviously in C++, Java, you know, Python, you know, eventually we're just going to make it really super easy for you to utilize these functions that are in the ecosystem of the marketplace and you can do your job and the, you know, hardcore engineers at a lower level can do theirs.

**Tony** [00:36:34] So I'll ask one last question because I think we're probably almost out of time. Where do you hope to be in five years? And also, where do you hope the technology industry goes in five years?

**Charlie** [00:36:45] I'm seeing some interesting things in technology that are very exciting. Centralized memory, you know, set the fabric of centralized, compute centralized. And I think you guys are at the forefront of a lot of that. I'm excited about that, you know, because these PCI buses are getting very, very fast. PCI five they're getting very fast. So, you know, things like data locality become less of an issue and things become a lot easier to develop. Where do I see us is I want to be that platform as a service like the MuleSoft of workflow, but where each component on your canvas is a hardware accelerated function. Right? And I want to create the composer that allows you to develop those kernels, see it in a marketplace, and let other people share it. So I would like to be that, you know, cloud offering. I'd like to be the platform like, you know, Hadoop was where you can download it freely and then businesses came out of it: Cloudera, Hortonworks and whatever. I would love to be that as well. So there is an aspect of, you know, a lot of work that we need to do. We need to recruit some, you know, volunteers on our open source in this initiative for our platform, We're going to obviously have a premium version of it, which has support and greater scalability. But I would like to be, you know Intel Inside is a great mantra, and having Digital Cortex use Intel Inside is pretty good. Runs on Digital Cortex with Intel Inside.

**Tony** [00:38:31] That's such a good tagline. We should probably end our podcast on that. I'd like to thank Charlie for joining us and talking about Digital Cortex and how oneAPI is enabling their business and also talking about his journey from engineer to company founder. I hope you'll join us next time when we talk more technology and trends in industry.