Hello everyone and welcome. Thanks for joining us for the Deep Dive into AI and Productivity, a conversation about how AI on device computing and adaptive technologies are changing the way we work and how we interact with our tech. I'm Alexa Orban, Director of Marketing Communications at Cephable, and I'll be your host for today's session. We're thrilled to be joined by two incredible voices shaping the future of human computer interaction.

Rami Hassini is the Director of Project Management for Qualcomm Technologies, where he leads enterprise software and A I experiences for the Snapdragon Compute business. He has more than 25 years experience in the semiconductor industry, including 20 years at Qualcomm. After starting his career as an embedded software developer, he transitioned to a solution architect for mobile applications and now serves as a product lead for the Snapdragon Compute business. He is passionate about leading AI compute transformation and launching innovative software experiences. Alex Dunn is the founder and CEO of Cephable, and he's been leading the change in adaptive technology. Before founding Cephable, Alex held senior roles at companies like Voicify and EF Education First, specializing in applied machine learning, software architecture, and project engineering.

Today's session will cover where a I powered productivity is headed, how on-device performance opens new possibilities, and what it really means to create tools that work with people, not just for them.

Before we dive in, a quick reminder. The closed captions are available and a full transcript will be shared after the session. If you're watching via our YouTube Premiere, feel free to drop a question in the chat. We'd love to hear what you're thinking. Let's kick off with a question for both of you. If you had an AI assistant for anything outside of work, what would you want it to do? I'd love to have an AI assistant. You know, just they're so advanced in being able to kind of scrape the web and identify, you know, things that are in market and come up with great summaries. So obviously. There's a lot in this world to go explore, so I'd love kind of an A I kind of agent out there that's constantly kind of monitoring things to do activities, different cultural activities in different regions that are going on and just constantly popping up itineraries for travel, you know, to go take vacations and then go check out different world, you know, cultural events I think would be awesome. That is such a great idea. I think that could be a new product that I haven't seen break into the market yet. So definitely pin that one as a great idea. Yeah, Alex, I'd love to hear yours. Mine's probably less interesting, but

it's like the thing, the other side of productivity that's always looming in the back of my head and it's my yard work that I'm always behind on and my yard falling apart that like. I feel like I need an assistant to tell me what to do to like fix all the chaos that I've caused. I mean, being in New England, we deal with the extremes of every season and it's like every season is just destruction and I and I feel like I need, I need a next level intelligence to help me stay on top of it. Yeah. Or it turns into that even next level where it's like the movie Smart House from the 90s where it actually just does it all for you too. That that would be an ideal situation. Sign me up for that when that's in the market. Oh, definitely. If I could have, you know, somebody to schedule out when I need to cook and wash my dishes and all of that and then they did it as well. I mean, I think that's the next level of. A I for us too. Well, as

we know, A I is everywhere. And like we just talked about, I think it'll be even more everywhere in our daily lives. But I'd love to talk about what A I powered productivity actually means and how it's starting to change the way that we work. Oh, do you want to start with that?We could, we could start with Rami and then go to you. Yeah, definitely so. I think there's a lot of areas where AI productivity comes in and it comes across different user personas. So for you know the the, the traditional users obviously you've got collaboration, we've got noise cancellation and and and and you know the the, you know video background effects and being able to run those locally on the NPU allows kind of that real-time processing. But then we're also seeing a lot of things around productivity with. Auto-generating, vou know, e-mail responses, summarising text, you know, things where we can even see PowerPoint

generations happen automatically and and even some of the technology that Sethable has created with accessibility and being able to kind of monitor, use different I/O mechanisms. I think those are all powerful from a productivity perspective. And then obviously powering that with local processing allows for privacy and security. So you know it's it's amazing to see how much a I has transformed productivity and with kind of where Agentic is going, where it's going to drive more complex prompts and be able to kind of distill that into actionable insights I think is is is where the market's going from from an A I productivity perspective. Love to hear your perspective there Alex. Yeah, I think there's there's a there's a lot of ways to think. I mean machine learning AI has been sort of core to productivity since it really started to become an application like from being able to use speech recognition to dictate

text which is, you know, existed in some form or another for a while now and continues to be an area that's improved with machine learning to. I think like the the more modern Gen. Al and and just generating things conversational AI and and now into the sort of like agentic age that we're we're marching into like it's it, it really comes down to just getting things done faster. I mean one of the things that we hear a lot from especially our enterprise users is like the it's the little things that add up in the day-to-day or like we've started to kind of coin it as like digital chores like. From the little stuff like navigating menus and figuring out how to do stuff to, you know, getting started on a new project, organizing things, it's like it's all those little things that like if you need to do it from scratch, keyboard and mouse style and you need to just start writing or you need to start clicking and navigating and doing like those little things add up to take a

pretty significant amount of our day. You know, I I always sort of think back to actually from a friend of mine, Scott Hanselman at Microsoft. This is years ago. He'd shared this like number of keystrokes you have left in your life is this sort of like really ominous thing to think about all the time. It's like how many, you know, applications can you write with the number of keystrokes you have left in your life?How many letters can you write? How many emails? How many, you know, award-winning research papers does it take and. What I think is really cool about, you know, what what we're seeing with AI and just that sort of aspect of just helping you get stuff done is you're recovering a lot more of what else you can do in your day by removing some of those digital chores or at least offsetting a lot of it with with AI capabilities. And it's definitely an area where we're leaning in pretty heavily from the suffable side of things.

Yeah. And I think that's an interesting segue kind of from the machine learning side of AI that's been around. To now where we're going with the on device A I capabilities that are different because a lot of A I that we're utilizing through GBT and different plugins that people are able to use that's being sent to the cloud. So with NPUs you're able to keep a lot more A I local and on device. And I think that a lot of people don't realize too that NPUs actually help their daily tech use be even faster and smoother. So for those strokes that they have on the keyboard to all the programs that they're able to utilize. They're able to have this AI be on device. So for people that aren't as familiar with NPUs and how important on device AI will be, Rami, I'd love to hear your perspective on this. Yeah, thanks Alexa. Thanks for the the background there as well. So actually Qualcomm's been investing in the A I neural

processors since 2007 and an incredible vision our leadership has had in in this technology. I think of a neural processor as a processor unit that's custom built for AI. So AI workloads are very compute intensive, and so you really need something that you know is really specialized for those workloads. Conversely, if you look at like a CPU, those are more general purpose processors. They're supposed to take on tasks like browsing the web or watching streaming videos or streaming music. And that that ALU is really your arithmetic logic unit. That's your core engine within a CPU. And because it needs to be multi-purpose or general purpose, it's not specialized in AI workloads. And so when we built our NPU, you know it, it can handle tensor based calculations and vector based calculations. So you can see 35 to 50X faster performance on the NPU than the CPU because it is specialized for those

workloads and and by having that directly on the device. You get a lot of benefits, right?You get the benefits of running it locally from a cost perspective, privacy, security and even personalization. And so we're we're excited that you know we were the 1st to deliver an NPU for Copilot plus PC launches in 2024. We were exclusive at the time Microsoft launched Copilot plus PCs. And we're going to continue kind of driving that innovation with partners like Cephable who, you know, we're we're working closely with to kind of bring Al models onto the NPU. Yeah. And on that personal productivity side, like we were talking about before, Alex, I'd love to know how the NPUs kind of bridge to that and help the everyday user as they're using their NPUs. Yeah, I mean it's kind of like what what Rami said. It's, you know, it's like the right tool for the right job. And when you're doing Al workloads, having a tool that is built

to run AI workloads is great. I mean otherwise running on the CPU or or running it on the GPU, you know you're you're using a compute that was originally designed for something else or or general purpose, right. Like when when we started to, you know what's called just accelerate machine learning workloads on GPUs, it was because we could get a little bit closer to. You know the the types of calculations that GPUs are better at doing than a CPU is is much closer to what we do in machine learning workloads. But NPUs kind of get us all the way there. Ultimately what it means is like when you're doing machine learning it's it's you have the opportunity to make it faster like and when you do a lot of machine learning in order to you know and you're doing it all the time, it means that it's constantly faster and better on battery life and and it has like all these low-level implications that. Directly impact the the end-user experience like

for for us it meant like from the the first thing we did was bring like more of our inferencing over the NPU, right?Like we started being able to run our computer vision models on Qualcomm's NPU so that we're using less battery power, it's faster and that unlocks all sorts of other experiences. But now we're we're bringing more and more of our machine learning workloads over to the NPU from some of the stuff that actually you had mentioned earlier like text generation, e-mail generation, summarization. You know assistants that just help you do more things and and bringing more of those generative A I capabilities over you know being able to shrink the size of models down to the point where the the limit is is being pushed on memory consumption and storage more than it is on compute. You know I think is an exciting future that you know we're we we're leaning into but ultimately it's like you know you get when you get down to it it's it's a tool

that. Just makes things faster and allows us as developers to build better user experiences. Yeah. And I think when, you know, people hear productivity, they always think about their to-do list, right?Like that's oftentimes when they're more productive, it means they're able to do more throughout their day and, you know, utilizing computers that have NPUs that are of newer tech. It allows them to do that. Like we were all just talking about before, when your computers are overloaded with too many tabs that are open or you know you have too many applications open, it slows things down. But when you're able to offload some of those things to the NPU, it allows everything to keep going. And I also think that this brings up a a big piece that we're always talking about is that security. So Rami, I'd love to hear from your perspective, you know what on device A I is allowing people to do from a security standpoint. And even from an

employer standpoint on that security too, yeah, really good question. Obviously as enterprises, you know, security and data privacy is top of mind for IT leadership. So we're seeing it come out in in multiple factors. One is endpoint security, so being able to kind of run models locally on the device to kind of assess from an anti malware, anti ransomware perspective and being able to offload those onto the NPU allows your CPU to have the full bandwidth for your employees both from a performance and a battery life. So I think typically if you look at it, most you know IT machines, they have so many agents running on there that can utilize over 10% of your CPU utilization, your battery life gets shrunk down cause you got all these background tasks for both endpoint security and data protection. So offloading those onto the NPUs is is really powerful. But even more importantly is around data protection, right?So instead of having to upload your

files to the cloud and now the processing done up there, being able to kind of locally use them whether it's in our ag. Solution or whether you're using kind of OCR to contextualize what's on the screen, that's really important for IT leadership versus the traditional way of doing it where you know maybe take a screenshot and take it to the cloud and do your processing in the cloud. Obviously no, no one wants wants that. So that's that's some of the power that the NPU brings as being you know providing that that custom you know A I processing experience to the edge that allows for kind of data privacy and security to remain local on the. Yeah, I think kind of building on that like there's there's kind of a few angles to it and and it's ease of implementation and use versus like ensuring data privacy, right. There's there's always that trade off at least in modern technology like the easiest thing is to do and implement but all like have the

sacrifice of your own data like if you want to use ChatGPT as a developer like throw a credit card down and start sending your data over to open AI's API's and. Cheapest way to do that is by letting them use it to retrain and and everything else. But then even with some of these larger models, you know running them in managed cloud infrastructure all the way down to like on Prem infrastructure, like there's significant jumps in level of effort to maintain the data privacy and security. But what what's been exciting for us is seeing like NPUs being brought to end user devices means that like we can bring these workloads to. Entirely on device, you minimize one, the number of attack surfaces when it comes like real security, which is important, right? If you're building an application that's used by, you know, 50,000 employees at your company and all of that data needs to go to either a cloud or even to a managed infrastructure setup, there's so

many more places where you need to, you know, harden and and and manage, you know, network security versus saying like, OK, we have all, you know, 50,000 people, they all have their own devices that are capable of doing all these. You know, highly intelligent automated things like the only surface then to be concerned about is the device itself, which tends to be a lot easier to manage from a security perspective versus you know, high traffic networks. But the data privacy is always the biggest one. It's it's this constant trade off of your data to ease of use. And I think what's really cool about NPUs being brought to the edge is like you don't have to sort of make that compromise anymore, yeah. When you're bringing up GBT, it really makes me think about just a I and how we think about it in a day-to-day too, because so many of us are used to utilizing that chat interface. But it's for a lot of people on the surface level for a I, that's where they

believe it's going. But as we know, there's so many more layers to a I and so I really feel like you know, with the chat interfaces, it doesn't really hit the potential of how powerful even low. local AI tools can be on your devices. So Alex, I'd love to know how an ambient tool like Cephable comes into play with changing this paradigm. Yeah, I think like outside of Cephable, like a lot of consumers don't realize how much Al machine learning is going on behind the scenes, like Rami cropping him out and putting a background like that's that's doing a pretty significant amount of machine learning and computer vision to segment you against, you know, what else is going on in the room in order to build a background likeAnd that's happening on device. You know we we do the same thing when it comes to stuff from our own computer vision models of tracking face data to running SLMS and LLMS that help actually generate stuff on device. So I think it's it's a mix of

like it's kind of always there and now it's just better and there's now more and more sort of small things that build up to to that better experience all the way to like. Yeah, you can bring an intelligent chat experience if you want to to an implementation entirely on device and it's certainly something we're always looking at with like what that interface is. But when it comes to productivity, like when we see people talking with Copilot or ChatGPT or Gemini or any of these cloud based solutions or even just any of these chat interfaces specifically, it's like. It's been a it's been an interface for people to interact with conversational AI for a long time, even before Gen. Al. But it's like so not optimized for getting certain things done. And you know what we're seeing with being able to take those same capabilities of the generative Al, but do it in context of the actual applications you're trying to use and not

have to be like, oh, I got to go like, you know, talk to my little buddy over here and then like get the answers from it and then bring it over to this environment. And instead you're just doing your job and getting things done. Like it saves you even more of that sort of small digital chore time and not introducing a new digital chore of having to go talk to your chat to then bring it over. But it also means you can contextualize a lot more, right?Like so when we do things like text generation, if I want to generate, you know, a response to an e-mail, I can have the context of what is going on in that e-mail thread be used when generating. So it's generating the most accurate response or like. if I need to generate you know code in my IDE, it should know the code around it in my style and everything else that's going on in order to provide that next-level experience. There's always going to be a use case for chat, but I think more and

more, we're going to see Gen. Al sneak its way into the ambient experience of everything else you're interacting with and other applications, the same way that regular old, good old-fashioned machine learning inferencing has been doing forYou know, at least 15 years now. Rami, do you have anything to add additionally to that?Yeah, I mean, I think the technology that that Acceptable's building is, is incredible, right?It started off with kind of the segmentation model and being able to track user gestures at 30 frames per second, which is, you know, without affecting the CPU, GPU and bringing that on the end view. Is is really incredible technology, right? Because it's all processing at real time, right?And so as the user kind of moves their head around or tilts it or kind of waves at at at the camera, it's able to not only just track the gesture, but also be able to isolate the action

and translate that into. Next steps, which is remarkable technology, right?And then being able to then expand that into different types of voice dictation. So being able to go, you know, voice to voice to text and using different ASR models and having the precision to kind of take voice and actually be able to to to dictate that with high precision is pretty phenomenal, right?And and the IT shows the rapid growth of a I development right where where things are going from when ChatGPT initially launched a couple years ago to how far you know technology's advanced across different models, right. You have traditional LM based models, but now you're going multimodal. So you could go you know text to text or text to image. And then you can use different types of modes across those to really drive different productivity experiences. The value there is employees can now leverage AI to kind of accomplish repetitive tasks,

right?So when you look at, you know, code assistant or code Gen., you know, no one really wants to comment code. No one wants to go refactor code, right?That's not the that's not the most exciting job for a developer, right?So now you have a code assistant that can go drive. That level of repetitive tasks. So your developers can focus more on innovation, what products we want to build, what do we want to innovate, how do we architect it and then let the development kind of happen kind of in parallel. So that's kind of where we see things kind of evolving is that employees can focus on the innovation and let the repetitive task kind of be handled by by kind of AI, you know, AI experiences on on the platform. I think code refactoring is probably my least favorite digital chore. It really does feel like it. Yeah, but it it brings up a good point where, you know, when we're used to on our phones or maybe some of our smart home devices where we're

talking to them to get things done or, you know, we talk to our device to, you know, be able to send messages or communicate. But there hasn't been that open up in many cases until probably the past year where you can tell your computer what to do and you can use different. Inputs beyond the keyboard and mouse. So we're starting to see that, but I think it really can improve workplace efficiency. And Rami, you were kind of touching on that a little bit. So I'd love to, you know, expand on that thought or even Alex, you can jump in to talk about this too, because I really think it can make the workplace more efficient if you're able to just talk to your computer to get done what you'd like to. Exactly. Yeah, you're you're right. Because like you know speech is the guickest form of input, right. And and if you know where we kind of see things going with with kind of agentic workloads. So agentic workloads is where where we kind of see services kind of

happening in the background. So imagine a service that's constantly monitoring your your e-mail and identifying which ones you typically respond to prioritizes those. Can actually pre-draft the response as personalized with my tone, not just a general response, but you know with with the way I would respond to it. And then in that maybe another service that's running in the background, scraping the web, identifying product launches from competitors and then doing a competitive analysis on what was launched versus our current product portfolio and and repositioning our go to market strategy. So all these agents that are running in the background to have voice be that input that kind of guides these agents and these services that are. Running in the background will be super important, right. And and then pairing that with an NPU, you're going to need a dedicated processor to handle, you know these multiple agents that are kind of your digital employees

in a way running in the background so that you're not overloading your CPUGPU, which were kind of built for more traditional tasks of browsing and you know, you know Excel and and Office type of workloads, right. So that's kind of where we see the efficiency really take off in in the next year or so. Yeah. And I'm even seeing, you know, your Al org chart. I've seen this across the board on, you know, on small nimble teams they have, you know, their digital COO, which I I definitely think that's where we can see these things coming into play. But that point of they're already doing these things for you, you're not prompting them to make that change. Alex, anything to add?Yeah, I mean I've been in specifically the the voice space within the sort of A I machine learning side of things for like. Eight years now, almost exclusively from Voicify to all my private research that led to Cephable and obviously Cephable voice is like the number one input that

we see people using and then, you know, face and buttons and everything else on top of that. I think that like Rami made a really great point. Like it's it's the fastest way to communicate. It's also like a a way that most people communicate naturally. But it's always been one of those very elusive challenges forever in technology on like, how do I process speech to true intention to get stuff done?You know, speech recognition has been around for for like probably the within 10 years since the beginning of modern computing. Like the speech recognition, Speech Recognition Institute's been around since I believe the 80s, if not earlier. Like early forms of speech recognition were commonplace pretty quickly, but. We've never really gotten to this point, you know, where I can just tell my computer what to do and it figures out how to do it. And it's not because we can't recognize the speech. It's because how do we take the context of what you're

thinking and what your intention is to actually get stuff done. And like, that's the sort of front that we're pushing now with Cephable is like, yeah, we can do speech recognition all on your device and all personalized to you and it gets more accurate with you and you can dictate and you can give commands. But it's really that like. Context of what you're trying to do and then starting to execute on those things for you with basically little agents that are running behind the scenes or workflows that are being automated for you from very simple stuff. Like you know, if I'm in a, you know, PowerPoint presentation, I say go to the next slide, it should figure out how to go to the next slide to. I needed to go start writing an e-mail for me in response to that that message I got last night and just have a go do that. With the right type of intention, with your tone and the way you write and everything else. It's like, you know, that's that's where we're trying to push

the frontier. It's like voice as an input is great, but it it really becomes powerful when you have the right output from understanding the context and intention. And that's sort of, you know, that's that's the secret sauce and magic behind what we're doing in the A I space. Such a good point. And and yeah, I love that. I just envision it, right. Those little agents that are able to get done what you need to do. Because when we do come onto our computers, often there are those digital tours, there's that long to do list. And if you could just sit there and be like, OK, what could I offload? And to be able to have that, I mean, it's a game changer for so many people. So when we're talking about the reach of this, I just love some examples in industries that could really benefit from the productivity tools that we've talked about today and also utilizing NPUs and we could kind of dive into some industries that most people might not

think of too. Ram, if you want to kick that off. Yeah, definitely. We we work with the top enterprises across the globe and we actually see it as universal across every industry. Every every IT lead that we're talking to is looking about how do they build out their AI strategy, right. And and you know what, what might be the AI strategy for, you know, the financial institutes might be a little bit different than what insurance or manufacturing is doing. But they all see this opportunity again to kind of automate repetitive workflow to drive innovation. And some of it might be universal, right?Some universal solutions are like an IT help desk, right, which is grounded on your set of, you know, documentation, whether it be, you know, your your back end around how you open up tickets or or reset passwords or different types of experiences that allow you to offload some of that that kind of the, you know, the initial. IT requests to almost like a

rag bot in a way, right. And others are more industry specific, right. So think of like a like an agent that can do your financial planning or your tax agent, right. So we're seeing it across different industries and and and and and we're seeing like a lot of universal aspects, which is exciting, right, because I think every enterprise that is driving down this path is, is, is, is really transforming the way their employees will work. The way their employees will interact both from a collaboration perspective as well as a productivity perspective, everything from code assistance to RAG to kind of drive IT help desks. To, you know, to the way we were describing about e-mail generation, right. And we all work in, you know, global environments, right. So we all have peers that work in different parts of the globe. So when, you know, when when we start our day, we might already have an inbox full of emails from different regions, right. So

to have that all, you know, pre-sorted, prioritized, responded is, you know, would would allow us to really accelerate our our productivity when we start our day. So hand it over to Youngs. Yeah, I think the the every industries is definitely something we see too. And it's it's one of those things that you know when you when you look at the sales side, people don't like to hear and they're like, OK, well like what industry is like Staffable's AI gonna be the most impactful for?And it's like all of them. I mean we have customers in just about every industry. I think one of the more surprising ones that we've been seeing growing. I guess two of them is manufacturing and in like where hands free computing actually it's like so obvious when you think about it, but it's like in a very different type of interaction and the other is in education where you know AI adoption is is sometimes very quickly adopted and

sometimes there's a lot of backlash around it, especially because early versions of generative AI were quickly adopted by students to. Call it plagiarizing, cheat, what have you. And so there's this interesting sort of back and forth push back. But where we see it creating a lot more value in the education space really is in inclusive education, supporting students with disabilities and being able to sort of level the playing field in the the education sector too. But certainly creating value for just about anyone from like note taking assistance to just easier communication or adaptive communication capabilities. Yeah. This brings up a really good point about, you know, making workplaces more efficient, but it's also making them more accessible to everyone and just from from AI tools like we had talked about throughout this entire webinar. So Alex, I'd love for you to build off that thought and and where

we're seeing this in the workplace. You know, I think the most interesting thing about you know where where Sappable is and and came from obviously like we were built as an assistive technology originally and continue to build for people with disabilities and. Sort of fell into, you know, this quote UN quote curb cut effect where we started to just see more and more people using the capabilities purely for productivity. And of course now we're building more capabilities specifically around that from text generation and everything else. But there's this, you know, more more of a of an approach to universal design and what I think A I specifically has the capability of doing is like. Truly being an equalizer when it comes to the workplace and education, like being able to have something that understands you, that is personalized to you, that adapts to you and allows you as an individual to then communicate to your own technology and to others in a way

that just works best for you is that's that's what accessibility is for, right? And so I think, you know, while we've been building Staffable and and even the early research, you know, for the last five years specifically for individuals with disabilities. This this curb cut effect is has been really cool to to experience directly. And you know, it's it's also been interesting just to see how many different innovations globally have come from that same type of thing, right from the first capacitive touch screens that now live in everyone's phones and tablets and PCs being built by, you know, an individual and invented by an individual to support their child with a disability to be able to use a computer without being able to use a physical keyboard to. Closed captions to literally the curb cut, which was created for people that were using wheelchairs but benefits everyone from delivery people to strollers to just not smashing your toes when you're trying to cross the

street. So I think, you know, in the form of digital technology, AI is as a whole, you know, very much a part of that inclusivity and accessibility journey. Yeah. That's such a great point. And I know you were just talking about so many industries and just people who can really be impacted by the power of AI. So Rami, I'd love to know as everything evolves in the AI space, where do you really see these tools going in the next few years? And what do you think the most exciting up and coming new AI tool is? Yeah, thanks, Alexei. So First off, yeah, I want to echo what Alex mentioned. You know, that was one of my favorite projects I've worked on with was getting Cephable to run on the NPU just because of the reach it has across accessibility and enabling, you know, employees to kind of have, you know, be able to, you know, take full maximize the full use of the device even if. They had certain disabilities to to really be able to

maximize productivity across across the entire you know estate of employees at A at a company. So I want to First off thanks thank the Acceptable team for allowing that partnership to to to expand from that perspective. We were super excited to to to help enable that. As far as where we see things going forward, you know we love the fact that you know the NPU on these devices is is driving innovation. We've seen things from the ability to drive training content. So things that would have taken 3 weeks to go create create a video tutorial, whether it be for, you know, team onboarding or customer training videos for new products, things that would have taken, you know, weeks or even a month to go create can be done in this in a matter of seconds, right?You click through your workflow, your video gets auto-generated and it can get translated in over 20 languages, right?So that's. A pretty powerful tool that again, no one

wants to create the script and and and and and do those repetitive tasks. They want to drive innovation around building the message and and and building the strategy around driving that message to to their end channel. So that that was one exciting technology that just that just got leveraged on the NPU. Another is. The ability to, and this is more of an agentic workflow, the ability to say, hey, I want to go compare my product against the competition and and generate a competitive analysis. And So what you want to do is build a Knowledge Graph of files that are on your device with files that are, you know, in a private cloud and and the public environment and actually be able to have this agentic AI workflow leverage your entire file system to go build out a PowerPoint slide deck. It really drives a competitive analysis and actually can can can create visual dashboards of of of of you know that kind of addresses your prompts. And so

we're seeing the technology curve go higher because our NPU brings that performance and models are getting more compressed and more powerful. And so we're kind of at that inflection point where we're really seeing you know adoption curves around Agentic A I on on device. That's awesome. Yeah. Alex, I'd love to hear some of your closing thoughts. I mean, I'm where where technology is going in the next few years. Do I just spoil our like 2 year road map or I, you know, I think the the main thing that we see is in, you know, I don't want to belabor the point that we've talked about this this whole time, but. It's that just I I want to tell my computer what to do and it just gets the job done right. It should it should be an extension of me the same way that my own team members are an extension of of me and and I'm an extension of them to be able to just get things done. I I I

think that's where you know we're seeing

that type of enablement head. It's that it's the agentic side of things but it's there's a lot of complexities and and nuances within that that are going to take some time to develop. I think we're gonna start to see more standards coming coming out too around this stuff. You know, MCP being one of them around how you actually can expose modalities and rag capabilities and how developers can start to enable their existing applications and data to be more easily navigable and usable with a I and I think that. That's a longer Rd. to get to where it can become that much of A standard that everyone's sort of implementing it the same way that, you know, we saw an initial uptake from like restful APIs being developed by developers building applications that allows other folks to interact with it. Like, I think we're going to start to see more of that when it comes to enabling an outside, you know, generative AI and

and Agentic AI to be able to go interact with my data, my business logic, my information in a way that makes sense. So I think that's where things are going to be heading and you know we're very much excited to be a part of that journey over the next few years. Yeah. Well, thank you both for joining today to chat about this. To summarize, I really think what we've been saying is AI for productivity, it's not just about the to-do list, it's not just about asking AI chat to get something done. It's really about being that extension of you and just. Being able to communicate to your computer what you want and it's already doing it for you. And I really think that that's the next phase. So when it comes to productivity, it's not just about checking off your to-do list, it's also about how you can become more of an innovator. And I think we all would love that additional time to really be able to think and let our computers be able to check off some things from that probably

never ending to-do list for for many of us. So I I appreciate the insights on. At all ends. And I think it also just opens up really the capabilities of on-device AI and NPUs. So we appreciate you both today and we're looking forward to more conversations like this. Awesome. Thanks so much, Rami, for taking the time and joining us. Yeah, definitely. Thanks, Alex. And thanks, Alex. That's been, it's been a great discussion. Appreciate the time.