Intel Corporation

Network Composability for Communication Service Providers (CSPs)

CORPORATE PARTICIPANTS

Brie Hilliard
Moderator

Babu Peddu
Intel Corporation – Product Marketing Manager, Barefoot Switch Division

Tim Harrison
Extreme Networks – Director of Product Marketing, Service Provider and Portfolio, Office of the CTO – Though Leadership

PRESENTATION

Brie Hilliard

Welcome, everyone. I'm Brie Hilliard, webinar director for the Barefoot Switch Division at Intel Corporation, and your host for today's webinar. Thank you for taking the time to join us today for our webinar titled “Network Composability for Communication Service Providers”.

Before we get started, I want to point out some of the features of the BrightTALK tool that may improve your experience. There's a Questions tab below your viewer. I encourage our live audience to please ask questions at any time. Our presenters will hold answering them until the end of the presentation. Below your viewing screen, you will also find an Attachments tab with additional documentation and reference materials which pertain to this presentation.

Finally, at the end of the presentation, please take the time to provide feedback using the Rating tab. We value your thoughts and we'll use the information to improve our future webinars.

Today, we are pleased to welcome Babu Peddu from Intel and Tim Harrison from Extreme Networks. Babu Peddu is a product marketing manager and the Barefoot Switch Division at Intel where he's responsible for marketing and technology planning of the Intel Tofino Intelligent Fabric processors. Prior to this role, he was a solution and product marketing manager for the communication service providers vertical at Cisco, where he led the go-to-market and launch activities for Cisco NFV, ASR9K Edge routers, and security integration into the Cisco service provider portfolio. Babu also worked at Ruckus Wireless and launched the industry's first Wi-Fi 6 access point. Overall, he has more than 23 years’ experience in the telecommunications industry with technologies encompassing switches, routers, virtualization, and security.

Tim Harrison is the Director of Product Marketing for the Service Provider portfolio and an adjunct thought leader in the Office of the CTO at Extreme Networks. He's also the co-host of the Inflection Points podcast from Extreme. For decades, Tim worked in the internet access, managed data center, and managed service provider industries, delivering large-scale multi-vendor service provider networks and services. He shifted to take on product development and management duties to deliver some of the most innovative dedicated private cloud services available in Canada. Before joining Extreme Networks, he branched out into the value-added reseller space where he managed a national sales engineering and architect team, supported the technology journeys of great customers, and developed new businesses across multiple verticals. Over the years, Tim has attempted to learn to paint, play the Chapman Stick (look that one up!), sing opera, write novels, tread the boards as a Shakespearean actor, be a rock star, produce a podcast, and start a historical computer museum. He has learned as much or more from his failures as he has from any successes. Tim lives in the artistic and theatrical city of Stratford, Ontario, Canada with his wife, his young son, two Bengal cats, and 22 drum kits in his basement studio.

Welcome, Babu and Tim, and thank you for taking the time to join us today. Babu, I'll hand it over to you to start off.

Babu Peddu

Thank you, Brie. Hello, everyone. Good morning, good afternoon, and good evening, depending on where you are located. The agenda for today is I will introduce you to Extreme’s 9920 Network Packet Broker solution powered with Tofino 2. We will also cover the

networking challenges the telco industry is facing today. I will also introduce to you and walk you over the Intel Data Center of the Future, Intel Intelligent Fabric vision. We will also cover Intel Intelligent Fabric processors, and we will end the first part of the webinar with Intel Tofino product family.

OK, the Extreme 9920 enables cloud native visibility solution with Network Packet Broker application. It provides a composable data pipeline that allows maximum flexibility for future enhancements of the networks. Extreme Networks is one of the first customers to deploy Intel Tofino 2 programmable Ethernet Switch ASICs. The Extreme 9920 Network Packet Broker solution is deployed in large 5G mobile networks of many service providers. Extreme projects are huge opportunities for the solution over the next few years. We will cover a little in depth of the solution in the second part of the webinar.

Now, let's look at some of the key networking challenges the industry is facing today. Many enterprises have moved to the cloud in large. As a result, the service providers are also experiencing a need to move to the cloud. Bottom line, all the data traffic is moving to the cloud. We are generating data at a very faster rate than our ability to analyze, interpret, transmit, secure, and reconstruct in real-time. The analysts from IDC, they are anticipating by 2025 there will be at least 175 zettabytes of data that will be floating around in all sorts of networks-- cloud, hybrid, private, and other on-premise data centers.

So, moving to the cloud native architecture will require container-based processing, orchestration, and automation. Moreover, as the workloads are distributed across multiple cloud environments like public, private, or even hybrid, we will need to manage all this data across all these providers. The current architectures are not adequate enough to scale to cater to that. We will need more modernized architectures. AI and changing workloads will spawn need for growing network optimizations to reduce the troubleshooting time of any issue in a given network.

Security. The network perimeter is growing, and as the networks evolve, threats evolve too. The traditional methods of firewalling, distributed denial of services, advanced malware protection, intrusion detection and prevention, they are not just enough. We will need to come up with better security techniques. All these issues coupled together, combined together, are causing operational challenges for the IT administrators and also the network administrators. This is resulting in root causing to the network slowdowns. This is also increasing the capital expenditure, operational expenditure, and investments for service providers. It's time to change the way the networks are built. It's time to stop solving the future problems with older techniques. We will need networks that are smarter, agile, intelligent, that can provide increased bandwidth.

To overcome these challenges, Intel came up with the Data Center of the Future. New times, new problems, new solutions, Intel is building a Data Center of the Future. The Data Center of the Future has three categories. One is CPU for general purpose computing, number two, Intel XPUs for workload or application-specific acceleration, and three, IPU, infrastructure processing unit, to accelerate the infrastructure. The IPU is a secure programmable networking device that will increase the visibility, that will increase the performance, and reduce the overhead in the CPU. It will also allow to utilize and balance the resources between storage and compute within this architecture. All of these vectors that you'll see on either side, these critical components, are connected through Intel Intelligent Fabric that you see in the middle. So, let's talk a little bit about the Intel Intelligent Fabric.

Intel's vision to address these networking needs as part of the Data Center of the Future is the Intel Intelligent Fabric with optimized connections. Though targeted at cloud data centers, it can also be leveraged for Edge data centers, high performance computing, and communication service providers adopting cloud native technologies.

The Intel Intelligent Fabric is based on the Leaf-Spine topology. So, it brings together a comprehensive portfolio of Intel Silicon Photonics, Intel Tofino switching products, Intel IPUs, Intel Network Ethernet adapters, Intel scalable Xeon processors, Intel Agilex, Intel eASIC, Habana, and all. So, this will all provide extra visibility, added intelligence, and improves the performance. Intel is also taking leadership in bringing the industry standards together, the P4 programmability, SPDK, storage performance developer kit, EBPF, DPDK, and for the networking OS, the SONiC OS, OCP SONiC OS, and also the IPDK, infrastructure programmable development kit. These are all open source standards that Intel is planning to bring together and building a leadership role here, so that

angle is ease of use, massive bandwidth, AI-driven self-monitoring, self-analyzing, self-healing, self-remediating networks, it increases the security, it enables end-to-end optimization with in-band telemetry. It improves the density, optimizes power, reduces the costs.

Now, let's move onto the next slide. So, at the heart of the Intel Intelligent Fabric is the Intel Tofino 3 Intelligent Fabric processor. We announced this a little over a month before at the Intel Innovation event. The explosion in new services like 5G, IoT, augmented reality, virtual reality, AI/ML will require unquenchable bandwidth... demand for unquenchable bandwidth. The current data centers have outgrown and cannot scale to these new services. We cannot solve future problems with older techniques. The Tofino 3 brings intelligence through P4 programmability and AI/machine learning acceleration. This will also increase the performance up to 25.6 Tbps using a 56 or 112 gigabits per second service. This also brings extensive visibility and control using in-band telemetry and Intel Deep Insight Network Analytics Software. With IFP, it's not just the cloud data centers that will benefit, but also the Edge data centers, high performance computing platforms, communication service providers, everyone, whoever will adopt cloud technologies.

So, Intelligent key benefit vectors. Key benefits fall into three major buckets or benefit vectors—intelligence, performance, visibility and control. So, under intelligence, we have a fully customizable P4 programmable pipeline. Intelligent packet processing for accelerating AI/ML workloads. Expandable table and buffer sizes with Intel FPGAs. Enhanced security with Intel Software Guard Extension with Intel Trust Domain Extensions. Performance, we have multiple form factors, 6.4, 4.8, 25.6 Tbps total throughput. It comes with 112 or 56G SerDes for high speed and easy migration. It provides high speed Intel Silicon Photonics. It provides power-optimized hyperscaler use cases for Intel Tofino Intelligent Fabric processors. Visibility and control provides enhanced congestion control. It identifies delays or hotspots with real-time in-band telemetry. It analyzes the packet flows with Deep Insight Network Analytics Software. It provides remedy using AI or Deep Insight reports. It performs traffic monitoring and steering for enhanced security and reliable transport. It provides an increased in-band telemetry data available with the Intel IPUs and Ethernet network adapters.

So, I wanted to present you the Intel Tofino product family roadmap slide. So, we have three form factors here. The first generation Tofino comes with 16 nanometers with up to 1.8 to 6.4 Tbps, has 24 gigabits SerDes, and then a second form factor is Tofino 2 with seven nanometers, provides speed up to 4.8 to 12.8 Tbps, it has a modular chip design and has 56G SerDes. The third form factor is Tofino 3 with seven nanometers, which is built on the same architecture as Tofino 2. It comes from 6.4 to 25.6 Tbps, has the same modular chip design, and SerDes is 112G and 56G. It provides intelligence with P4 programmable AI/ML acceleration and really highly secure. Performance, it provides up to 25.6 Tbps throughput with IFP, 112G/56G SerDes for Tofino 3. Visibility and control, as we spoke earlier, edgeo-to-cloud real telemetry, that is one of the features, provides enhanced congestion control, and also enables self-healing network capabilities. The networks will be able to remediate any issues automatically.

So, with that being said, folks, I will pass on the baton to my partner, Tim. Tim, please go ahead.

Tim Harrison

Thanks very much, Babu, and thanks so much, Brie, for getting this all set up. Good morning, good afternoon, and good evening, everyone. Forgive my long intro. I'll skip further details and jump right in.

First of all, thank you for spending some of your valuable time with us today. We really appreciate it. I loved one of the things that Babu said earlier-- well, all of the things about Babu said earlier, but in particular, as he was speaking, he said networks solving problems for the future can't use techniques from the past. I think that's really insightful. So, let's talk about how some of that has changed over the past little while.

As we all know, 5G for communication service providers has become the standard and the goal going forward, and that rollout was actually going fairly well, we were seeing some good transition into that, and it was really starting to blossom, and then along came COVID. So, we all know what COVID has done to our lives and to the business models of service providers. So, what does this mean for Extreme in the new world, and what has changed culturally and socially that changes the nature of 5G deployments and other technologies requiring visibility?

Well, the thing that we've discovered is that service providers need to be moving faster than their customers. Networks must adjust for new use cases as they appear. Operators need to be able to see changes coming and to react before those changes impact their customers or degrade their services. We've all heard the standard 5G use cases, and we feel that those are definitely starting to develop, but I believe that the killer apps for 5G are still to come. How does a service provider build a network for the long-term now, and have the flexibility to change when something changes our lives again? Back to what Babu said earlier about networks solving problems for the future, can't use techniques from the past. So, what is the story that we're telling to help address this?

Well, we are telling the story of composability. There are some different concepts of composability out there. Some believe that it's primarily resource management. We see it slightly differently. So, we're designing everything with our forthcoming solutions to be composable. Composability really enables a service provider to adapt to changing use cases by defining how they wish to use their infrastructure. The power to change how packets are processed through that data pipeline. The flexibility to enable, migrate, and disable features with the click of a button. The depth of visibility anywhere within the network. The freedom to use the infrastructure to deliver new services without rip and replace of the infrastructure. And all of this is enabled by the power of cloud native technologies. That means that a service provider can change the nature or recompose their services based on the use cases that they see identified through our visibility tools, actioned through our automation, delivered, managed, and assured through our platforms, and enabled by the ability to scale anywhere. Look at this as the leveling up of the hardware abstraction layer to a network abstraction layer that turns the network into a platform onto which you can then compose applications. We don't want to just focus on the technologies underlying the network, but really on composing what is service delivery on top of the network. And this is really important for 5G. As we see those use cases build, we want to make sure that service providers are building a network for the long-term now, and have the flexibility to change when something changes our lives again. This is why composability is vital for the service provider network at scale.

So, speaking of scale, when we talk about scale, and when we talk about things like infinite distribution, we really mean at scale, and in so many ways, infinitely distributed. Think about 5G as a model for incredible scale, not only in bandwidth that's usable by the consumer, but also from an infrastructure standpoint. So, we have a picture here of a tower on top of a mountain, think of the truck rolls when something goes wrong, think of the management complexity of having network devices spread throughout the landscape that are either at Extreme's scale or distant, but also the ability to scale the resources to manage and to configure and to compose that network. So, we're traditionally focused on horizontal and vertical scale. We think about this as more of a circumference of scale where everything is distributed significantly out rather than centralized deep in the heart of data centers. So, we want to make sure that when we talk about scale, we understand the nature of scale, specifically in 5G and cloud networks.

And speaking of cloud, service providers are at different points along the cloud path. The key is that if a service provider does not have a strategy for cloud or cloud native technologies, they're in trouble. We want to be able to help service providers understand the value of cloud and cloud native technologies, and for Extreme, who has often been known as an enterprise business, when we talk about cloud, we mean more than just our enterprise wired and wireless management platforms. Those are absolutely key, and they have been huge values for Extreme, but what we see is that the cloud native technologies underlying those are the real key to composability and consumption at scale, and in an infinitely distributed world. Now, think about the 5G proof points. We've seen AT&T leverage Azure for 5G. We've seen Bell Canada going to GCP to manage their virtual infrastructure, Microsoft Azure for Operators, for example, and even the AWS enterprise cellular or enterprise 5G announcements that we've just seen, cloud makes a huge difference, and people are starting to understand the value that they can gain from cloud native technologies. We've certainly seen that within our own organization, as we've pivoted to more of a software standpoint to allow our customers to actually start using these technologies and to change the way that they do business. So, I think you're going to see some amazing new solutions from Extreme built on cloud native technologies for the service provider, and we're going to talk about one of those today.

So, today, we're going to focus on visibility. So, without the ability to know what's actually happening in your network at any given time, it's extremely hard to make effective decisions on service delivery. For Extreme, visibility is about having the ability to help you make decisions, the best decisions possible, by getting network data consumed, processed, and off to analytics tools at speed, and at scale.

When we look at the key ways that we can help you today, we feel that we deliver an industry-leading flexibility. So, this is our composability story. The ability to compose not only the software, but also the Intel Tofino 2 chipset that we use using P4. To enable CSPs to scale anywhere through automation and simplified management. Again, when we look at that circumference of scale, and we look at how infrastructure is going to be in any place necessary to provide services, you need to be able to manage that, and managing it individually as you might do today on certain network management apps is not going to be the way to do it. We need to be able to do it in a better way to allow you to scale not only your network, but your resources to manage that network. And the power of visibility across highly distributed physical and virtual environments is vital, as we look at composing the applications. That means we need to be able to push visibility to wherever you need it in the network, and that means that it may live inside of virtual environments. That may be in private cloud, it may be just in a virtualization stack, where there's a MEC stack somewhere at the base of a cell tower. We need to be able to get visibility into there, and that takes multiple types of form factors from the physical to the virtual.

So, what are the key tenets, if you will, of our network visibility solution today? Well, we have at the top layer, if you will, Extreme Visibility Manager. This is a centralized management tool for all of our visibility solutions. That includes legacy tools that we have, legacy platforms that we have, that have been great tools and resources within major service provider networks today, and also what we have going forward. We don't want to strand anybody because the technology has been working for a period of time, and they still need to get value out of that. So, we continue to be able to manage not only existing platforms, but platforms going forward.

The second is the Composable Network Operating System. So, this is really a key element of the strategy. We need to be able to compose not only the application in microservice form, but we also need to be able to be focused on customization of that platform and the ability to deliver rapidly. So, using those cloud native technologies that we talked about, microservices inside of containers for example, this allows us to be more effective and more efficient with the platforms that we deliver for the composable network and visibility. And the Extreme Network Visibility Platform, so this is some hardware, some software, and some chip that is built for a cloud native world. So, literally with the operating system, and with the tools that we use on top of that platform, this is a cloud native application, and the application that runs on top of that platform that we're talking about today is our Network Packet Broker application, and this is very key. This is where Intel and Extreme are attached at the hip. This platform is fundamentally built using the Intel Tofino 2. Now, we are in love with the Tofino 2. It allows us to do things that our competition cannot do and cannot catch up to us with. We think we've leapfrogged them with the use of the Intel Tofino 2 chip.

Extreme Visibility Manager, so often, when we talk about service providers, we think of CLI junkies who loves to get in there and type and spend time scripting things. That's true, and I come from that world, and some days I miss scripting those types of things. However, when we think about what the things that people need to accomplish are, we find that the CLI is just the end result, and we don't want to focus on-- sorry, the end step, we don't want to focus on the end step. We want to focus on the end result. So, how can we make it easier for somebody who is administering one of these visibility networks or configurations or deployments? We want to give them the easiest, most user-centric experience as we possibly can, leveraging the tools that we have. It's very complex to manage a type of environment like this without having data about that environment as well. So, we have streaming platform telemetry. As Babu mentioned, there's lots of telemetry that comes from the chip, and comes from the architecture itself, that we can feed back into the application, allowing the application to give a constant update on the readiness and power and availability of the visibility platform.

The key value prop here, this improves productivity, simply just by lowering the time it takes to build the infrastructure rule sets and to manage the applications that are delivered out onto the network, and that reduces operating costs. It also reduces risk. So, when you have a bunch of humans doing individual tasks, there's an increased risk. When you have tools that allow more intuitive access to configuration, there's less opportunity for mistakes, there's less opportunity for somebody to miss something, and that reduces risk, which overall reduces costs. We think that's really important to understand.

So, let me show you the fruits of our labors from a platform perspective. This is the Extreme 9920 that Babu mentioned previously. So, this platform is fully composable. This is from the chip level, all the way up to the operating system, all the way up to the management from Extreme Management Center. It's also all built on cloud native technologies. Extreme Visibility Manager has been built on cloud

native technologies. Therefore, it's composable all the way down to the chip. Always, we're focused on the most powerful packet visibility possible and we've worked very hard to make sure that we are providing the best slicing and dicing of those packets possible, and making sure that we have incredible performance, and that Intel Tofino 2 chip gives us the ability to do more with the platform and to do it at speed. So, very important for us.

And as part of the overall cloud native visibility, there is obviously the Extreme 9920 in the center, but we have Extreme Visibility Manager. We also have Extreme Session Director, which is useful in LTE and 4G/5G networks. We have the SLX 9140/9240, which are key collectors and aggregators inside of the network, but again, as we talked about physical and virtual, the key is that we do have virtual packet broker applications that can be orchestrated out to those virtualized environments. All of this, of course, is built leveraging our BGP-EVPN configurations, on top of Edge infrastructure, and is supported by powerful automation. We have our Extreme Fabric Automation tool, which is sometimes actually built into the platform on certain switches, but allows you to very easily create your configurations to interconnect all of your devices together, allowing you to have that visibility network instantiated on top of either your existing network or instantiating a visibility network that is separate from your production infrastructure.

Just to show you where that fits in our product line, so we have our IP Fabric tools and our IP Fabric switches. Those are very powerful Edge and MEC and data center switching architectures. We also have our border routing platforms, also our IP transit, but key to this is, of course, our cloud native visibility quadrant at the bottom you see there. So, this is a very holistic architecture. There are tools that fit each of the different types of use cases, some of them blend across. We also have our trusted delivery initiative where we have some very protected risk mitigation, factory to production, and full lifecycle management type of switches for those that are concerned about supply chain and nefarious activity there.

So, to get more information on powering our visibility solutions together with Intel, we have some resources on our solution/visibility page. There's also a fantastic blog written by my co-worker, and also my question gatherer today, Joerg Bonarius, and you can get more in-depth architecture information there on how it is that this platform can be designed to fit into your architecture, and that architecture can be used to make sure that you're making more money, you're making sure that you're addressing your solutions at scale, with infinite distribution, and in a consumer-centric way, so that you're focused on what it is you're delivering for customers, and the Intel platforms have been absolutely fantastic. We have the ability to do things that others cannot do, because we focused on working with Intel on this. And I also want to make sure that you understand the nature of the composability here. This Network Packet Broker application that runs on top of the 9920 is the same packet broker application that can be imposed and moved into different other platforms. The 9920 is a perfectly enabled solution to be able to process those packets at scale at extreme speed. Thank you for Intel's Tofino 2 chipset, and the pun there. But this gives us the performance in large scale environments, cloud service providers, and communication service providers, and we have the ability to feed that through the rest of our visibility platform and to be able to manage it at scale through our Extreme Visibility Manager.

So, I want to thank you all for listening, and I will turn this back over to Brie who will take us through to our questions.

Brie Hilliard

Great, thank you both so much for that excellent presentation. Tim just shared several links, and I just want to direct everyone to the Attachments tab that's right below your screen. You'll be able to find all of those links and attachments easily accessible right there.

Now, let's get into the Q&A. The first question I have is for Babu. So, Babu-- let's see, bear with me one second. Babu, is the Tofino 3 IFP targeted only at data centers? Can it provide benefit in any other deployments?

Babu Peddu

The Tofino 3 is not just for cloud data centers that it will benefit, but also the Edge data centers, high performance computing, and communication service providers, all who are adopting cloud native technologies, so all of them can benefit out of it.

Brie Hilliard

Great, thanks so much. Another question for you. Why do you call the Intel Tofino 3 an Intelligent Fabric processor? Is it really just a switch?

Babu Peddu

Well, it is a processor that goes into your networking switch. So, the Intel Tofino 3 is based on the P4 programmability, and accelerates packet processing up to 25.6 Tbps. It also optimizes power consumption for hyperscale use cases. So, these are the use cases that it can solve. So, bottom line, it is a processor that goes into a switch.

Brie Hilliard

Great, thanks very much. Now, the next question I have is for Tim. Tim, is the Extreme Networks visibility solution targeted at service providers only?

Tim Harrison

Great question. Actually, no. So, if you are an enterprise and you are looking for visibility in your platform as well, by all means, this is absolutely for you. The service provider focus was where we started originally. So, the legacy platforms that came in with Extreme's acquisition of Brocade were very much service provider directed but not necessarily focused. What we think in the terms of our new platform, the 9920 and the Network Packet Broker application, they can be delivered anywhere. It's really about how much value you're going to get from your data, and we think that large enterprise, even some medium enterprise, as well as service providers can get significant value from this platform.

Brie Hilliard

Great, thanks so much. Let's see, Babu, one more for you. When will the Tofino 3 IFP be available?

Babu Peddu

The customer samples will be available in Q2 2022 next year.

Brie Hilliard

Great. And, Tim, what environments benefit from the 9920 in the NPD application?

Tim Harrison

Well, obviously 5G environments benefit directly. We also see opportunities at MEC sites, Sso where service providers are delivering multi-axis Edge computes. So, where there may be third-party applications running and they need to understand the value of the service that they're providing in the MEC environment, that they need to understand maybe how to bill their MEC customers, which is also an important part of the services that they provide. We see, when you're looking inside of large enterprises, there's obviously the security aspect of things, being able to understand the traffic flows within your network. So, maybe the 9920 being fed by some virtual packet brokers make some significant sense there. In distributed environments, you may just need to get that data somewhere so it can be filtered and processed. So, you may not want to put a 9920 at the Edge, but having the virtual packet broker as an element maybe of your remote branch or even as part of an architecture for healthcare, we think there's also a great opportunity for people to leverage this for healthcare, for security, also for regulations. Understand that the technologies being used, that are based on certain regulations, are doing what they need to do. So, I think there's a broad swath of industry that can definitely benefit from this.

Brie Hilliard

Excellent, thank you. That's a really great insight. And, Babu, another one for you. What's new with Tofino 3, apart from the higher throughput you mentioned?

Babu Peddu

Yes, so Tofino 3 IFPs optimize power consumption for hyperscale use cases, and also accelerates the machine learning for distributed workloads. The Tofino 3 also supports real-time in-band telemetry, spanning across the edge-to-cloud spectrum.

Brie Hilliard

Sounds great. Let's see, Tim, where can people get more information on the visibility solution from Extreme Networks?

Tim Harrison

Absolutely. If you go to extremenetworks.com/solution/visibility, you will see a significant amount of information there that we've posted. There will be more to come as we continue to expand on the portfolio and find new use cases, but to start with, I think that's the greatest place to go.

Brie Hilliard

Great, thank you, and how does Extreme Networks apply the concept of composability across a highly distributed environment?

Tim Harrison

Wow, great question. So, the wonderful thing about composability, from our perspective, again, is not that it's just resource management, but that it's also the orchestration of the applications that deliver services on top of that network. So, look at a highly distributed network as a very complex management challenge. What we want to do with composability is break that management challenge down by creating the resource, the network as a resource, the network as a platform and abstraction, which removes some of the challenges of trying to manage individual devices at scale, but applies their resources to the pool. So, again, going beyond resource management. Once that resource pool is available, that abstraction is created, the ability to then just deliver services where it's necessary becomes much more easy. It's not about the individual component at scale. It's about the component being able to deliver that application to the user in that particular area. It's a different kind of vision of how the network works. Whereas many other vendors may consider their intellectual property focused on the platform to be the key driver, we think it's the application and the user's consumption of that application on top that's the driver. So, we want to provide highly distributed environments with that abstraction so that you're just focused on delivering the application and the Tofino 2 chipset and the Tofino family is really vital for that. The ability to go in and tweak how packets are processed, to be able to go and program the chip to do something different in a particular area, I think is really a game changer. I think we're going to see more of that as we go along, and I think there's going to be a lot of people scrambling to catch up when they see the power that can be unleashed by having a programmable chip like the Tofino family.

Brie Hilliard

Great, thank you so much. Just a reminder to our audience, please continue to feel free to ask questions. We have several coming in, and so we've got some time left.

The next question is for Babu, actually. This is appropriate, Tim, since you were just talking about the Tofino 2. So, is there any architectural difference between the Tofino 2 and Tofino 3?

Babu Peddu

Not really, they are built on the same architecture, except that Tofino 3 is capable of providing twice the bandwidth of its predecessor, which is Tofino 2. So, they are very similar.

Brie Hilliard

Great, thank you. Tim, is the 9920 and NPB application available now?

Tim Harrison

Yes, it absolutely is. Yeah, we launched it earlier this year. We have it in multiple sites currently running, and we couldn't be more proud of this platform. So, yes, if you're interested, please reach out, and we'll get you connected with somebody to talk about this, but yes, the application and the platform are available today.

Brie Hilliard

Excellent, and then other than NPB, what other solutions are planned for the 9920?

Tim Harrison

Ooh, good question. It depends on the use case, and actually, what we've seen as we've been deploying the 9920 is that it's been more about changing the nature of the packet rather than putting different applications on top. However, because it was originally focused on a Network Packet Broker platform. We feel that that's what the best use case is today. What you'll see going forward, I think, as we go more into our composability story, is that elements of the 9920 will start to show up across the platform. We see this as really a fundamental technology that's going to start delivering across the platform and you'll see more composability applications that roll out towards NPB sometime in the next year, I suspect.

Brie Hilliard

Sounds great. And then Babu, does the Tofino 3 have embedded NPUs and Arm cores?

Babu Peddu

Well, I will have to look into that and get back to the person who's posing this question. So, I have no answer at that time, so I'll get back.

Brie Hilliard

No worries. Well, everyone, I think that concludes our questions for today. Thank you everyone for attending, that was a really great lively discussion, and thank you to our presenters, Babu and Tim. Really appreciate you taking the time to present today. You'll note that I have on this slide up everyone's contact information. So, if we didn't get to your question today, please feel free to reach out to us directly. The presenters would be happy to answer your questions via email as well.

So, again, thank you for joining us today. Please don't forget to give our team a rating for the live recording so that we can continuously improve the quality of our webinars. Your feedback is truly helpful to us. This concludes our webcast.