CORPORATE PARTICIPANTS

Lilian Veras
Moderator

Jeff Sharpe
Supermicro – 5G/IoT Edge AI Solutions

Jay Vincent
Intel – Senior Solution Architect

Dr. Ken Urquhart
Zscaler – Global Vice President, 5G Strategy

PRESENTATION

Lilian Veras

Welcome everyone to the Intel Network Builders webinar program. Thank you for taking the time to join us today for our presentation titled: “Delivering the Security Service Edge: High-Performance and Zero Trust”.

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, including a number of websites and documents mentioned in 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 will use the information to improve our future webinars.

Intel Network Builders Webinar Series takes place live twice a month, so check the channel to see what's upcoming and access our growing library of recorded content. In addition to the resources you see here from our partners, we also offer a comprehensive NFV and SDN training program through Intel Network Builders University. You can find the link to this program in the Attachments tab, as well as a link to the Intel Network Builders newsletter.

Intel Network Builders partners have been working to accelerate network innovation by optimizing their solutions on Intel technologies. These industry leaders are recognized in our Winners’ Circle program and Zscaler is a member, and Supermicro is a Titanium member. Learn more about our INB Winners’ Circle program by clicking on the link in the Attachments tab.

Today we are pleased to welcome Jeff Sharpe from Supermicro, Jay Vincent from Intel, and Ken Urquhart from Zscaler.

Throughout his 39-year career, Jeff Sharpe has focused on defining strategic direction for best-in-breed networking and AI training interference solutions, while building successful partnerships for new product delivery. He's currently focused on next-generation Edge and AI solutions for smart places, intelligent city infrastructure, Industry 4.0, and 5G MAC-based services. Before joining Supermicro, Jeff was responsible for leading strategies for AI Edge, telecom and platform solutions with ADLINK, Radisys, and Nortel.

Jay Vincent is a Senior Solution Architect in the Enterprise and Cloud Network division of Intel's Network Platforms Group. Jay is a 25-year veteran of designing and developing data center solutions. For the last six years, Jay has focused on delivering optimized NFV solutions to cloud service providers, and is now focused on defining the optimum compute technology for the SASE vendor POP hardware deployments.

Dr. Ken Urquhart is Global Vice-President of 5G at Zscaler. He holds three degrees in physics and executive roles at Sun Microsystems, IBM, and Microsoft. Prior to Zscaler, Ken consulted with Fortune 2000 companies on 5G, cybersecurity, and AI.

Welcome Jeff, Jay, and Ken, and thank you again for taking the time to join us today, and I will hand over to Ken to start off. Thank you.

Dr. Ken Urquhart

Thank you very much, Lilian, and welcome everyone. So we're here to talk about the Secure Edge, and you can't talk about the Secure Edge these days without talking about the telco revolution that's happening, where we're getting fiber-optic speeds with wireless technologies, whether it's low earth orbit satellites, Wi-Fi, or 5G. Just to make this very concrete, we're going to focus a little bit on 5G, but we're actually talking about a wider range of communications protocols that will change our world.

First 5G, three words: it's fast, it's live, it's massive, and fast means fiber-optic speeds. 10 gigabits per second is where they think we're going to be heading. It's 10 to 100 times faster than 4G LTE. Now, we have a little thing called network slicing, which gives you your own network independent of others.

Live, ultra-low latency, try a millisecond or less is the target. Now, speed of light gets in the way of this goal, because if you're going to haul compute bits from a device back to the cloud, do something interesting and ship them back, you're going to have latency. Speed light's the speed of light. It takes 15 milliseconds for a packet to get, at a minimum, from the West Coast to the East Coast of the United States, and that's at the absolute best conditions. You don't achieve that in practice with fiber or any other communications medium. That means we have to talk about Edge Compute. Edge Compute is the other part of 5G that makes it work. You put compute and data resources as close as possible to the user devices, and that gives you a lot of interesting new use cases.

Finally, massive. What's massive? How about a million devices per square kilometer, far more than you get today, with the same low latency, with the same high speed that doesn't fall off as you add more devices. You know, when you're sitting in the sports stadium today, says you've got five bars and you still can't get a call out? It's because the current telco network is swamped at 10,000 devices. That's going to change, and in addition to these massive, you also get lower energy usage. With 5G, you'll get about a tenth of the radio energy usage, which means now rather than having to connect devices to the grid to function, you can look at things like 10-year lifespan batteries. I'm seeing things like water pump monitors that can be buried. They have a 10-year lifespan. IoT devices that'll just keep transmitting for a very long time. No power wires. Also ultra-precise location info. How precise? How about knowing to within 10 centimeters where everything with the radio is in your factory.

Alright, so we have this ultra-high speed communications, ultra-low latency is possible, and a massive amount of devices coexisting, what are the enterprise use cases we're hearing? Well, the three of us on this talk got together and said, here's what we're hearing, and we came up with this list, in no particular order. But these are the kind of vertical use cases lit up with fast, live, massive, whether it's Smart Cities, Smart Grid, autonomous driving, energy optimization, precision asset tracking, intelligent navigation, all rendered very interestingly by the Edge. And today, you can actually see some of this in action.

Here are 5G use cases in production today, and they're color-coded. These three boxes show the high-speed lights it up today, low latency in the second box, and finally, a large number of devices. We're not there yet with a large number of devices, but we are there with the speed, we're getting there with low latency courtesy of the Edge, and that's where we've got the high-speed public hotspots down to 8K 360 virtual reality, live sports broadcasting, drone live streaming, or even robotic manipulators. And you can get these today from various suppliers in Europe, Middle East, Asia.

So where are we today? Why do we need the Edge? Let's take a moment to think about this. We're talking about high-powered compute devices you can hold in your hand. Your phone today has more power than a circa 2000, 2005 server, and the internet right now we're looking at, it's no longer just for humans. There's going to be a surge in a huge number of form factor devices, Industry 4.0, augmented and virtual reality, and these devices by their nature will have limited amounts of CPU, GPU, memory, storage, power consumption, but they do give you ultra-low latency, and mostly you get security, if they're isolated from the internet.

And let's look at that scale against the cloud today. You've got a lot of cloud choices, great tool sets to develop, build and deploy solutions, lots of compute power, memory storage, high throughput networks, all the scalability you need. And then you have any number of networks between them, And as you know, you cross networks you don't have control over, you can't verify they're secure.

It's not the old days where you own your data center, you own your leased lines, you own your devices. To get the kind of flexibility enterprises need today, you cross over someone else's network.

And so we have this story, we have devices operating in near real-time. We have a lot of compute capacity in non-real-time. And we have this dichotomy of compute power increases as you go back farther away from your device, but at the same time, latency increases. So what do you do? You need Edge.

And Edge you'll hear described a couple of ways, either Edge Compute, or Multi-Access Edge Compute is very popular when you're talking about 5G. And it comes in various flavors, and dependent on how far the server is from your workload, and how far your server is from your device. And there are three big buckets we talk about, and again, you'll hear different names for these. Currently, I hear Near Edge, meaning five to 20 milliseconds away from your device, and near being nearer to the cloud. Far Edge is defined as one to five-millisecond latency, and Deep Edge, less than one millisecond, and that's where your server is basically at the factory with your devices, or it's on the city block you're walking on. Far Edge is typically within the city, possibly within the metro area. Near Edge is close by in a suburb or in a neighboring city that's reasonably close to you.

Now, let's talk about a couple of use cases. Ultra-low latency is for the case where you need your equipment to respond very quickly in response to some compute carried out on some data transferred from the device to your workload. And one is we have a company we work with called Taqtile that does augmented reality, and they do things like no-code robotic programming, but really what they're delivering is an experience where you're looking at live images, real life, and you're doing digital overlays. And why do you need ultra-low latency for that? Well, there's a little thing in augmented reality and virtual reality called motion to photon. That's the time it takes between some object moving to the computer sending an image of it moving, to a workload saying, oh, it's moving and I'm updating my digital overlays and sending it back. Great. It turns out that if you are more than 12 milliseconds roundtripping, the person with the AR/VR goggles can get motion sickness, and about half the population is affected by it. In some people it's as low as five milliseconds, others it can be as much as 20. So you have to have the workload nearby. The Edge is critical to light this up.

The other one is Ouster Lidar. They detect objects in a volume of space, they track objects entering and exiting, and if you're trying to use that to control robots, people interactions, things like delivery robots working in the factory, and you need to know where they are in relation to the humans because you don't want them to bump into each other. Today, we run robots at two to three kilometers an hour, really slowly. So if they bump into someone, they can stop and not harm them if you know there's no humans on the aisle that you're operating the robot in, carrying parts to a production line from the supply room. There are projects going on where they'd like to run them at top speed, which means 10 miles, 15 miles an hour, and that will cause a big efficiency on getting products, but you can't run into people, so you need things like this, and you need the ultra-low latency to light it up.

Now, I've told you about why the Edge is important, why it's going to change how we do things, the interesting use cases it lights up for us, but it also means we need a new generation of servers, something that's more nimble, something that can operate sometimes without air conditioning. Sometimes it has to be able to be extremely energy efficient, because you're not plugging it into a high powerline. You get what you can in an installation location.

So I'm going to turn things over to Jeff Sharpe, who's going to tell you about some of the exciting new optimizations Supermicro is doing so you've got Edge everywhere, the way you need it. Jeff?

Jeff Sharpe

Hey, thanks a lot, Ken. Hey, I'm going to go back one slide. So we were recently in Barcelona at Mobile World Congress, and Ouster and Taqtile was part of our booth, and a big part of our messaging around Edge-to-Cloud, and real life experience. So with Taqtile, we had our HoloLens glasses, and we were manipulating a robot, and we also had what the person was looking at up on a pretty large display, and there was a wow factor. I mean, looking at programming a robot, making that robot do something, if you've never done it before, it literally took five to 10 minutes to train because of all the on-screen displays that were happening. For low latency, we noticed, however, that there was a delay within our download speeds. We were using 25-Meg, I believe was our download speeds, and originally in the

first few hours were using the cloud to be able to showcase that AR. It was just too high of a latency to really create that great user experience. Once we put it on the Edge, it was amazing. We're talking microseconds of non-delay. And same thing with Ouster. With Ouster, it's around that lidar technology, non-personalized technology, but you can't do a lot of this within the cloud environment because of those latencies. So, again, it's really critical in those types of environments.

So if we step back, and we look at the market today, we live in a pretty exciting time, and luckily, we're getting out of this COVID environment, and we're starting to get out and about. However, over the past two years, myself included, we've been working from home quite a bit, and we're starting to learn Edge technologies is really required not just from working from home, but also all of those different types of applications that are out there. We know that the wireless industry, the MNOs, are gearing up for a higher throughput, better quality of experience for end users. We also know that compute technologies are gaining a lot of strength by providing more and more workloads within a single element. And Supermicro's really focused on more and more virtual cores, more and more throughput for that improved user experience and capabilities within that Edge.

We also know that at the Edge, it's not just about the Edge, it's not just about computing at the Edge, but think about all those elements that are connected from Edge-to-Cloud, and we're looking at more Edge devices, called Nano-clouds, more and more Edge devices that are out there handling the load. And the big part of this is including a lot of the artificial intelligence inferencing and training at the Edge. We're seeing more and more of that in all the market verticals. And if you think about all this, especially with a multitude of Edge devices, and a lot of those Edge devices are communicating from Edge-to-Edge, Edge-to-Cloud, Cloud-to-Cloud, and so forth, Zero Trust networking is so critical, security is so critical within all of this. And things like videos, content delivery, they're needing for that higher bandwidth that's out there as well.

So with all that said, if you again, think about the networking of devices that are out there, Ken talked a little bit about Far Edge, Near Edge, and Edge-based products, very similar to this. So from a Supermicro perspective, we look at it from mainly that compute mechanism, not really the sensors that we produce, but mainly that compute mechanism that those sensors plug into. We look at the data stores, the learning, it could be images, it could be videos, it could be analytics at the Edge that we're storing. There's a multitude of applications that are running really at the Edge. Not just one, it could be five or six inferencing items going simultaneously, and at the Edge and the network, you're having network capabilities, higher throughput capabilities, switching, soft switching, all the way into the cloud, and the key thing to this is all of these are connected, all of these items are working in parallel like a well-oiled machine, but they have to be secure. They have to have that Zero Trust environment to operate effectively.

So as we're seeing the market pretty much explode from a technology perspective, you've got a lot of different applications that require higher workloads, higher throughput capabilities. Think about deep packet inspection, bump-in-the-wire, looking at packet routing, looking at packet inspection, you don't want delays within that communications function from the Edge to the cloud. Policy enforcement, DDoS, next-gen firewall, all the privacy and security items that are required, again, you don't want to have any delays within your network. Especially when it comes to things such as autonomous vehicles, things such as augmented reality, you really don't want to have a lot of these bump-in-the-wires that are slowing things down. And then of course, as Ken said, as we're migrating into higher-speed networks such as 5G, we're getting into better utilization of the spectrum from slicing mesh networks. DAS, actually putting up a whole network of wireless technologies. We're seeing more and more enterprise 5G and enterprise LTE that's being implemented. And again, not to belabor the point, but a lot of this is being driven by some of those low latency services.

And then if you think of the growth of networks, not just in the cloud, you hear a lot about cloud-based technologies, high-end data center, but also at the Edge, you're seeing more and more areas of putting data centers on the sidewalk, putting data centers within a closet space, or outside up on the telephone pole. So you're seeing more and more workloads being driven by this exponential growth.

A busy eye chart, but it does showcase the spiderweb effect of things that have to come out of more and more higher Edge computing technologies that are running, again, in environmental conditions that may not be the best. Offloading things like providing video services, CDN-based services more at the Edge. Connectivity, again, all of us working from home these days, we're seeing the requirements around connectivity, and sometimes Wi-Fi just doesn't do it. You need higher-speed access for VPN and private

networking. Augmentation, aggregation of data are all driving a lot of the Edge services and applications that are out there, especially around AI functions as well.

So when you're looking at an Edge product, very simple, Zero Trust, secure, make sure that anything that's connected to that network is not trusted. I have those applications and that awareness that anything connected, either a hardware device or an application, or a software device, that we understand that that has to be trusted, that I have to look at it as a non-trusted environment. We're also seeing that cloud-native approach becoming more and more prevalence, which means that I have to manage virtual machines, containers, and so forth, which also means I need to be able to heal my systems at the Edge in a self-healing manner. I have to have high availability, like the operators do, the six-nines mentality. You're now being forced and driven into, let's say, an industrial environment, a transportation environment. And then also with a multitude of open hardware, open software, that leads a higher risk for security concerns. So all of these have to be looked at not just from the IT department, but also the OT as well.

And from a Supermicro perspective, one of our major core strengths, as Ken talked about, it's all about that Edge-to-Cloud, that latency, but we have the right hardware, we have the right capabilities to go from customer premise, anything that's like a microsecond or less, millisecond or less at that customer Edge, which a lot of times is not that environmentally-friendly area. It could be outside, it could be in a heavy moisture environment, it could be a production floor that's not sterile. So having fanless boxes, fan-based boxes, outdoor boxes, immersive cooling environments, all the way into carrier offices that we have to support NEBS and hardened availability, all the way into the data centers. So with that said, we've taken a simplified approach, because Supermicro, of course, has so many SKUs, a lot of great hardware out there, from GPU-capable Xeon SP down to Atoms, and what we've done is we've created a small, medium, large, extra-large infrastructure, really based on what the customers' needs are. We recognize not every customer is exactly the same. Some may need a super high throughput, 100 Gigs sometimes within that Edge, or maybe a lower-end processor, and you're not going to run a lot of workloads, they don't need 100-Gig, they may only need a few Meg of speed. So with that said, we've created an environment on enabling that end customer to have multiple Intel nodes, 342 at the Edge. We also recognize some of these devices are going to go outside, so IP65 cabinets that we've built that can go on telephone poles and on sidewalks, sidewalk-based digital signage that's a mini data center, and different form factors, and how many GPUs you can support? How many cores do I support? Is it a single-socket, dual-socket Intel processor, Xeon SP versus Atom? We do it all, and a lot of these systems are engineered in a way that can fit in small areas, whether it's the backroom of a 7-Eleven, or even in a digital sign or a kiosk, we can support all of that.

And with that said we are driven by that Intel-grade processor that's in there. And I'm going to go ahead and turn it over to our Intel friends, and if there's any questions later on, let me know.

Jay Vincent

Alright, thanks, Jeff. So I'm going to drill down a little bit deeper into the underlying technology that enables this Edge computing that both Ken and Jeff have talked about. So Edge-to-Cloud connectivity is going through a significant change due to the ever-expanding number of users and endpoints that you see on the left side of this slide. The need for geolocated points of presence housing MEC is becoming essential for providing optimal performance reliability, and that's what you see in the middle, are these points of presence. The technology that enables this end-to-end connectivity needs to be optimized for network functions, because that's what this center section does, it processes packets. This technology consists of the CPUs, which we're all fairly aware of, but you also have to be aware of switching ASICs, network interface cards, and hardware accelerators, and a myriad of software components that allow partners like Zscaler and Supermicro to provide the highest performance solutions on the market.

Now, as we... let me just go back for a second to a slide that Jeff had up. When we look at the Edge requirements, like he mentioned, there's a lot of different processors that are involved here. So, Intel provides these processors, like if you look at the bottom-left corner here, minimal ports and throughput, this is where Adam and core processors reside, where there's not a lot of processing required, and then goes all the way up to the top-right section where you have this extra-large right here, where we have Xeon SPs, where you have multiple Xeon SPs on a platform, and could have multiple other devices within that platform.

So if we go back to my slide here, these Edge/POPs-- well, these Edge devices on the left here, where you have thick, medium, and thin, the atom and core processors, we're processing thousands of users that connect through these enterprise Edge devices.

Now, when you move over into the middle here, where we have these Edge/POPs, we're actually processing hundreds of thousands of users at these locations. These MEC POPs can exist in hyperscale cloud data centers, or they can exist in managed networks by platform service providers on bare metal servers, that we would call colos. So these can be located in different locations, but all processing packets through the system. In all cases, the base compute infrastructure requires an integrated set of hardware devices to provide the best performance reliability. Consistency of the architecture is crucial. At Intel, we provide the components from the top of the rack switch ASIC to the NIC, to the CPU, to the hardware accelerator, and the software that provides world-class performance that allows the programming of this infrastructure. Intel's focus in this particular presentation is all the elements in the points of presence, these POPs.

As I said, consistency of the architecture is the key to making a reliable and performant end-to-end network. With Intel technologies like DL Boost to accelerate analytics, QuickAssist Technology, and dynamic load balancing to accelerate packet processing on the CPU cores, we also have Software Guard Extensions and Trust Domain Extensions that enable isolated memory enclaves for Zero Trust and confidential computing. Hyperscan is a software component for high-performance pattern matching, critical for firewalls and intrusion protection systems. And the Infrastructure Programmers Development Kit (IPDK) that is a common set of APIs for programming the infrastructure. These hardware and software acceleration components optimize Edge workloads running on Xeon platforms, leveraging Intel Ethernet products. The Ethernet products provide connectivity to pull all these components together. Foundational NICs for high-bandwidth connectivity. We also provide SmartNICs and IPUs, which are used for offloading common network functions from the host's CPU to reduce processing overhead on the CPU, allowing it to do more work for the application. Switches with Xeon D control plane processors and Tofino switching ASIC provide robust programming of the network switches to increase network telemetry and reduce overhead on the switch fabric. A consistent architecture reduces TCO, enables scale in diverse locations, and maintains reliable connectivity. With ecosystem partners like Zscaler that provides Zero Trust access, and Supermicro that provides robust compute and network platforms, you can achieve a lower cost solution with consistent high performance.

The Xeon scalable processor is at the heart of this consistent architecture. Intel's roadmap of Xeon SP CPUs are built to be the highest-performing CPU cores for networking and artificial intelligence. As Jeff stated, artificial intelligence is the killer app on the Edge. DL Boost technology consists of three innovations that accelerate artificial intelligence on the CPU. Our next-generation Xeon SP processor will combine these innovations, making Intel general-purpose CPUs high-performing choices for artificial intelligence. VNNI is the Vector Neural Network Instructions, which combines three different machine learning calculations into one, providing a faster processing of artificial intelligence. BF16 is called the Brain Floating Point 16-bit format, which reduces the storage requirements and increases calculation speeds. And AMX is the Advanced Matrix Extensions, which accelerates matrix processing. These features significantly improve processing for artificial intelligence on general-purpose CPUs.

So... Today's Xeon SP Gen 3 processors are a significant performance boost over Gen 2. We have increased the core counts, added higher performance memory channels, integrated persistent memory to augment last-level cache, and added additional I/O lanes. When we add in the high-performance E810 network interface cards, we can double the I/O density on Gen 3 up to 800 gigabits per second on a single server. The hero features on Gen 3 open new use cases for Edge computing. As stated previously, these Software Guard Extensions enable confidential computing. SGX allows for the creation of secure memory enclaves to store user secrets accessible only to verified users. Crypto and compression can be accelerated using QuickAssist Technology on the chipset or in plug-in cards. Crypto acceleration is also integrated in the CPU on the Gen 3 for specific instructions and common algorithms and protocols. As mentioned previously, DL Boost will enable Gen 3 with accelerated machine learning performance. The Xeon SP Gen 3 is built with networking in mind, making it the best overall CPU for the Multi-Access Edge Compute deployments.

The Edge/POP can be designed with Gen 3 to hit every size of compute demand for a small POP consisting of a partial rack to a large implementation, requiring multiple racks with ultra-high density platforms. For example, the Xeon SP Gen 3 5318 Network SKU

provides a feature called Speed Select Technology that allows you to choose the core count and TDP to meet the power and thermal constraints in a smaller remote location. The Xeon Gen 3 6338 Network SKU provides a much higher core count and throughput for a location that needs to serve a larger metropolitan location with many more users. The network SKUs are optimized for processing packets, and provide the highest performance for network workloads, which are critical to these Edge locations. The E18 NIC provides 10, 25, 50, and 100 gigabit per second interfaces based on the POP's throughput demand. The Tofino-based network switches at the top of this diagram provide a programmable network fabric to perform network functions at the switch layer, preventing the tromboning of traffic across the rack, and provides extensive telemetry for monitoring and managing a network fabric.

Intel isn't typically known for software, but we're one of the largest software development companies in the world. The software we provide is designed to optimize performance on Intel architecture. Software is a piece that ties hardware devices together to deliver high-performance compute infrastructure for software like Zscaler's. Open-source software like Hyperscan, the Data Plane Development Kit, and the high-density scalable load balancer, coupled with other open-source projects like FD.io enable high-performance execution environments that Zscaler takes advantage of for ultimate performance. Intel's open-source software and hardware accelerators are integral to the Edge platforms that Supermicro uses to create high-performance compute platforms. This integrated hardware and software creates a holistic and consistent architecture that will drive the Secure Access Service Edge and MEC in whatever Edge location it is deployed in.

So now I'll turn the time back over to Ken, that'll show us how the hardware provided by Supermicro and Intel is coupled with software from Zscaler and Intel to deliver Zero Trust network access. Ken?

Dr. Ken Urquhart

Thank you, Jay. That was great. So here's the big thing, the elephant in the room: cybersecurity. We've got these great new hardware platforms. We've got great new use cases. We're in a world with ultra-high transmission speeds, ultra-low latency lighting up the digital transformation. Over networks you can't fully assess the security posture of, or verify that they're safe, or verify that they're not already penetrated, what do you do?

This is where the Zero Trust Security Model and Zero Trust architecture approach comes in. Last year we had an Executive Order saying it's got to improve the nation's cybersecurity. Zero Trust architecture was recommended for both critical national infrastructure, government agencies, and corporations. This was amplified just in March last month in an additional statement saying we need to tighten up cybersecurity and again, moving forward, Zero Trust is the way. So what is Zero Trust? You're going to hear it as a buzzword, everybody's products are all going to say, "We use Zero Trust". To some extent, they do, they might. That's fine, but I want you to take a few minutes here and understand exactly what Zero Trust is, and at the end of this, you're going to have a very solid foundation for it, not having to read through pages of documentation.

Zero Trust says we live in a world where the breach has already happened or is inevitable, and you protect against it by constantly limiting access to only what's needed, and you look for anomalous and malicious activity everywhere. So in other words, Zero Trust really means trust nothing and no one.

Now, our Cybersecurity and Infrastructure Security Agency, CISA, has built out a description of Zero Trust architecture built on five foundational principles, and three foundational rules. Let's go over it really quickly in a way that I think is very easy to understand.

Why are we here in the IT industry? We're here to do five things. We want to connect the user from their device, over any network, to a specific application to exchange data, and we want to do it securely, and that brings us to our five pillars of Zero Trust. Connecting a user, that means identity. Do I know who you are? Are you really an employee, partner, a device that I want on my network, that I want connecting to my workloads? I have numerous ways of verifying this with identity providers such as Microsoft, Okta, Ping, and you don't have to choose just one. You can do multiple. You can use hardware keys, you can use multi-factor authentication, it's what you as an enterprise believe is enough to identify the identity of someone coming in.

Then from a device, that brings us to device. Do you trust the device? Is it something you manage? Does it have an endpoint agent on it from Microsoft, CrowdStrike, or others that'll tell you the security of it? Do you know the posture on the device? Is it something that you believe should be on your network?

Finally, network. Where are you going? Are you going to a workload in the cloud, a workload on the Edge? Are you staying on the device? Is it a destination you know about from that user, and their device?

And the application, do you trust the application? Is it having expected or anomalous behavior? Is it talking to unusual ports and protocols? Is it doing something you don't think it should be doing?

And finally, exchanging data, which is are you exchanging good things? Bad things? Are you doing content inspection? Have you secured the pipe? Is the data encrypted at in-motion and at-rest? All the questions you need to ask and answer.

And then there's the foundational principles, which is governance, the rules. Is that user and that device pair allowed to connect to what application to exchange what data? And that's really fine-grain rules, because remember, trust no one.

Automation orchestration, this is a lot of things to check and do in real-time. You need automation orchestration to make it seamless. A user should not have to do a whole bunch of gyrations to get to their workload to exchange their data, while at the same time, you're not going to allow people, who aren't allowed to be there, access.

And finally, visibility and analytics. There's a lot of infrastructure here. You have to watch what's happening at all times, log, track, trace, know that nothing funny is going on in your system.

And that's it, that’s Zero Trust in a nutshell, that’s Zero Trust in two minutes.

What do we do? Zscaler’s whole thing, spent more than a decade worrying about Zero Trust security and providing an easy-to-consume unobtrusive way, where you get a great user experience coupled with reliability of system, high availability and high scalability, and you do not get that on day one. There are a lot of new entrants into the field saying, “We do Zero Trust”. That's great. Do you have 10 years of experience providing ultra-high availability, ultra-high scalability, and reliability in a way that is unobtrusive to your end users? We do this with 150 data centers around the world providing Near Edge five to 20 milliseconds latency to a lot of our customers. But that's not just one number to think of, and all powered by Supermicro and Intel, and great features that you've heard. We’re to making use of all of them. But that allows us to process 250 billion or more requests per day, blocking 100 million threats per day, doing 175,000 unique security updates as new threats are found and identified, because when we find it for one customer, we found it for all of them, and we update everyone within moments. And that's what you're getting with the Zero Trust architecture, and because we're already on the Edge, you get that with high bandwidth, high throughput, low latency, low overhead, and that's what you need in this world of the Edge.

So that's it for us, and thank you very much for listening. I'm sure you have questions, please type them into the question area, we'd love to hear from you. We'd love for you to connect with us on LinkedIn. Ask your questions now or later. Lilian?

Lilian Veras

Thank you. Thank you to the three of you for such a great presentation and sharing very insightful information with us all. We do have a few questions that have come in while you were presenting, so let's get started on those. The first question I have here is for Jay, I guess. What is unique about Intel architecture that makes it ideal for Secure Service Edge?

Jay Vincent

Oh, yeah, that's a good question. Probably most people aren't really thinking about Intel architecture when they're talking about Service Edge, but we've designed-- like I said earlier, consistency of the architecture is really critical to making this reliable and performant, and so when you tie all of the components that we have, from the CPU to the network interface card, to the switching fabric, it creates a complete underlying architecture that makes it easy to program the infrastructure. And so, Jeff might want to comment on this, but

having these underlying components in the core hardware that provides the service, really provides consistency that enables this performance. So, yes, that's what that architecture provides.

Jeff Sharpe

And if you think about Intel, too, they're not just about that CPU that's embedded within the motherboard, or a socket. As Jay was saying earlier, we have NIC cards, we have network interface controllers and cards, we have hardware accelerators. There's also an overhead of great software solutions that's optimized for those silicon-based solutions on x86. So I think it's a pretty cool roadmap and future with Supermicro and Intel, combining all of those, and again, in a small, medium, large infrastructure, meeting what those demands are of the network, meeting the demands of customers, and we have the ability to provide those options of, again, high-end Xeon SP multiple sockets, multiple areas, all the way down to an Atom or a, i7 core-based device, so very cool. And not to mention, with Ice Lake D coming out, we just introduced Ice Lake D, and those are heavily positioned for industrial environments. So you've got that high-end processor that is optimized for industrialized use. So again, a pretty cool area, more and more cores, additional hardware accelerators and NIC cards, all running in a nice network of information.

Lilian Veras

Thank you. A member of the audience is saying that the telco providing his Edge services says that they are already secure. Why does he need Zscaler, or any Zero Trust solution?

Dr. Ken Urquhart

Now, that's an excellent question. This is Ken Urquhart. Telco providers are signaling something very important, which is they believe they're doing and providing you a secured network. Meaning their network is hardened against attack and hacking. At the same time, you have a responsibility as a customer to secure your workloads. Meaning they provide the secure network, you take care of the security of your workloads, in most cases. So you'll need an additional layer of security that we, for example, can provide. And that's really about it.

Jeff Sharpe

Hey Ken, let me—

Dr. Ken Urquhart

Remember that shared security model.

Jeff Sharpe

Let me add to that too, Ken. If you think of... a great example is this webinar. So you have Jeff Sharpe and Ken and Jay, the BrightTALK and those Intel folks. We're across the country, we may even be across the world, and we're connected to our local provider, and all of this could be going through wireless networks. It could be going through landline networks. It could be going through VPNs and whatnot within our own environment. But think about all the devices, think about all the applications that are running simultaneously that could add risk to that network. Absolutely, the MNOs and the telcos have provided a great infrastructure for that security, but if you're not secure coming into the network, then that becomes an issue. So I think when you're looking at Edge-to-Cloud with a multitude of networks that are combined, that network of devices and core devices, you may lose that environment, which is why you need Zero Trust.

Lilian Veras

Perfect, thank you. Another question we have here: How can a cloud solution like Zscaler provide protection at the 5G Edge?

Dr. Ken Urquhart

Another great question. Remember, with Zscalar, we provide a cloud service, but that is our control plane. Meaning that's where policy is enforced, that's where your rules about who accesses what are entered and enforced, and that's where we watch for threats continuously across all devices, all workloads. But at the same time, when you're talking about Edge, you say, well, but my data's got to go from an Edge server sitting somewhere out there directly to a device, and that's absolutely true. We provide the control infrastructure and set up those secured connections between authenticated and identified devices and users and workloads, then we get out of the way. So you still can have that ultra-low latency while benefiting from a global control infrastructure we provide. So yes, and even today, for example, on Amazon AWS, you can select us to secure your workloads, your devices connected to those workloads, over Amazon Wavelength. We're in the catalog, check us out.

Lilian Veras

That's great. Thank you, Ken. Another question we have here: Can I try Zscaler's Zero Trust in my Edge environment now?

Dr. Ken Urquhart

Yes, you can. You can go to our website, check it out, you can ask for access, you can talk to our people, we're happy to help you out.

Lilian Veras

Great.

Dr. Ken Urquhart

And there's more where that came from as Edge use cases expand. Right now, as I said, high speed, low latency is just starting, and then there'll be the scalability out when you'll want to attach 10,000, 100,000 IoT devices in your factories. We'll be ready for you.

Lilian Veras

That's perfect. Is there a limitation on specific Intel SKUs that Supermicro and Zscaler does not support?

Jeff Sharpe

At least for right now, most if not all of our x86-based SKUs support Zscaler, again, from small Atoms all the way up to Xeon SP processors across the different generations of silicon. However, we do make recommendations when we do sit down with our customer base, and we're trying to solve a security solution, we'll make recommendations just based on the earlier testing and certification that we've done with Zscaler, and lessons learned on the market. So we want to make sure the right systems, along with the right Zscaler attributes are fulfilling what those customer needs are. I would say 99.99999% of all of our systems will actually work with Zscaler.

Ken Urquhart

Yes, and Jeff, I want to add to that one thing people don't, I think, realize about Supermicro is you work with a lot of companies globally, and you see a lot of use cases, and you're not just-- when you buy from the Supermicro-- this is because we are a satisfied customer-- you're not just buying a bunch of servers. You're buying an expertise in outfitting, customizing the servers to what you need them to do, and they will tailor-fit you. You're not going into a big box shoe store and buying shoes. You're going to a custom tailor who can provide specific shoes for you at volume, just like a big box store, and that I think is a really interesting aspect of Supermicro people just don't realize.

Jeff Sharpe

And also from a risk perspective, too, Ken, to add to that, is that when customers approach us on not just our products, but how do I implement a Zscaler-based solution as well, there could be other ISVs that they're interested in, other software components they're interested in, and some of the benefits of Supermicro was they've done their R&D on our hardware, they've done their R&D and development all on Supermicro, so it de-risks that, and we have this business-to-business program that enables that matchmaking that may occur as well for not just security, but for other AI-based use cases or other types of software requirements around CDN and content delivery and IP delivery of different video components. So it's pretty awesome.

Dr. Ken Urquhart

Absolutely.

Jay Vincent

Just to add a little bit more to that. This partnership that we have with Zscaler or Supermicro at Intel is a trifecta of delivering value to the end user, and while Zscaler is the software component-- correct me if I'm wrong, Ken, but that that's all running on Intel architecture underlying, and Supermicro providing that consistent architecture underneath is really a key value statement for Zscaler that provides the reliability, the consistency, and the performance that really makes it shine for the end user.

Dr. Ken Urquhart

There's a reason both you two companies power our Edge. Said very well. And the other one is, of course, there's all these different form factors, as Jeff had said. Whether you're running with a hyperscaler over something like Amazon Wavelength, or Azure, or if you want to run in the factory containerized on your own servers, we have a wide variety of solutions to fit that, just like Supermicro provides a wide variety of solutions depending on your environmental power needs, and your compute power, based on Intel. So there's-- and when you talk to one of us, you're talking to all of us.

Lilian Veras

That is awesome. Those were the questions we had. I would like to thank the three of you again, Jeff, Jay, and Ken for such an insightful presentation. I would like to ask our audience to please not forget to give our team a rating for the live recording so that we may continuously improve the quality of our webinars. Thank you again for joining us today, and I'll see you next time, and this concludes our webcast.

Dr. Ken Urquhart

Thank you, everyone.

Jay Vincent

Thank you, thanks.

Jeff Sharpe

Thank you, everyone.