CORPORATE PARTICIPANTS

Xiaojun (Shawn) Li
Intel – Sales Director, Next Wave OEM & eODM

Kapil Sood
Intel – Principal Engineer & Network Security Architect

Chandresh Ruparel
Intel – Director, 5G/Wireless Core

PRESENTATION

Shawn Li

Welcome, everyone, to the Intel Network Builders Insights Series. I am Shawn Li, Sales Director Next Wave OEM and eODM Network and Communications Sales Organization at Intel Corporation, and I am your host for today's webinar. Thank you for taking the time to join us today for our webinar titled "Paradigm Shift in Edge to Core Security with 5G".

Before we get started, I want to point out some of the features of the BrightTALK tool that may improve your experience. There is a Questions tab below your viewer. I encourage our live audience to please ask questions at any time. Our presenter will hold answering them until the end of presentations. Below your Viewer screen you will also find the Attachments tab with additional documentations 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.

Intel Network Builders Insights Series takes place live every month, so please check the channel to see what is upcoming and access our growing library of recorded content. In addition to the resources you see here, we also offer a comprehensive NFV and 5G training program through the 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.

Today we are pleased to welcome Kapil Sood and Chandresh Ruparel from Intel. Kapil Sood is a principal engineer and a security architect for Intel Network and Edge Group, driving platform security technologies and research, and setting strategies' architecture direction for Intel Network Cloud, 5G, and Intelligent Edge business groups. Kapil has 25 years of technology leadership experience, spanning mobile and the cloud crypto systems, Intel's CPU and platform architect and start-ups. Kapil has helped define security specifications and ETSI, NFV, and authored IEEE 802.11 standards. Kapil has a MS, MBA, and BS, with more than 80 patents issued, publications and open source contributions.

Chandresh Ruparel is a business unit director of 5G and wireless core infrastructure in the Intel Network and Edge Group. In his 25 years of professional experience, he has held engineering development, product line management, product marketing, and business development roles, and successfully developed profitable networking, media, and server businesses. His organization is currently responsible for developing innovative solutions in close collaboration with the partner ecosystem and the communication service providers to deliver successful transition and the promise of 5G and the network cloudification.

Welcome Kapil and Chandresh, and thank you for taking the time to join us today. Chandresh, I will hand it over to you to start off. Thank you.

Chandresh Ruparel

Thank you, Shawn, and good morning, good afternoon, good evening, wherever you are. Thank you for joining, and taking the time. As we begin this webinar, let me without ado go over to the very first topic here.

First of all, we are living in a world that's way different now in terms of security attacks. If you look at just the last decade, and especially last year in particular, the sophistication of cyber-attacks, the frequency, and the impact is steadily increasing. In fact, I was just reading an article through SDxCentral's newsfeed that highlighted the Trend Micro update on Cyber Risk Index. So, it's a new report that highlights a survey of 3,600 global organizations, and about 86% of them expect some sort of cyber-attack over the next 12 months. So, we are going from a question of if to when. And this-- interestingly, it's not just that they are expecting an attack, but there are surveys done to confirm that when clients and customers see that the company or business they're dealing with has had a data breach, their confidence in that organization's ability to secure their data is obviously impacted, and they tend to shy away from doing business with them. So, there is material impact to individual organization's brand, their ability to do business, and expand the scope of their market. Cyber-attacks are now considered the third highest global risk as per World Economic Forum. No surprise. You just have to look at the current newsfeed and you realize the significance of the cyber-attacks.

Now, let's look at this in the context of 5G infrastructure, and the multi-dimensional transition that communication service providers are going through. We're talking about transition from single to multi-cloud domain. We're talking about the NFV to cloud native implementation, from 4G to a very different architecture for the core in 5G.

And with that in mind, let's look at the security implications. So, if you look at how security was addressed in the 4G core, you had perimeter devices, whether it's firewall, security devices that protect from denial of service attacks, or just security gateways at the periphery, and you had fewer vendors addressing the core infrastructure. So, if you have a single core vendor, the risk is relatively smaller, there were proprietary interfaces, and then you had hardware security modules provided by the vertical vendor community. So, almost a walled garden.

The dynamics are completely different as we go into the 5G SA Core. First of all, the architecture is a service-based architecture. You have the 5G control plane on a single service-based interface bus, which means a single network function that gets attacked exposes the entire core. Not only that, you have web-based APIs, cloud native infrastructure, microservices deployment, you have multi-domain deployments. It's a highly distributed edge. So, now you are going beyond the perimeters of telco infrastructure and deploying solutions in areas where you don't have necessarily the confirmation on their security profile. So, it's unsecure locations that you may be deploying your solutions in. And as telcos, so communication service providers, implement new services with network slices, and they look for monetization in newer areas, it's almost like you have an enmeshed solution from edge to core. What this brings out is that the walled garden approach and perimeter approach is no more sufficient. Every device, every element in the network has to be a security device. And very important is that you have to have a consistent security framework across edge, core, and cloud.

Imagine the delay in response, the risk associated with a fragmented security strategy. First of all, not only you cannot respond to new threats that you see in the market, your response to an attack is significantly reduced-- it's much more delayed when you are struggling with multiple security frameworks across these domains. So, it's very important that you look at technologies that provide holistic solutions across each of these domains.

If we double click on the 5G core infrastructure and the use cases, some of these are fairly obvious. And this is by no means a comprehensive list, but it's to illustrate the areas where security becomes extremely important. So, we already talked about the service-based interface. Secure communication between network functions is crucial. And in this context, I'll just like to highlight that, you know, technologies, for example, that provide highly secure private key or key management is extremely important. Your security here when you use technologies like TLS is highly dependent on the security of the private key. So, your security is as good as, you know, how well you protect it. If the private key is exposed, the entire core is vulnerable, because that's really, you know, what you're using to communicate between the network functions.

You also need to ensure that you have secure UE authentication, secure converged edge, federated multi-tenancy content delivery security. You have to ensure that your communication between networks is secure, your communication with vertical applications in the core is highly secure. And in addition to that, you know, several regulatory requirements that you have to comply with, for example, lawful

intercept. Whether it is… no matter where the location is where you are introducing the LI channel, you have to ensure that the computing of that compliance mechanism is highly secure.

In billing and charging applications, if you are auditing CDRs, you have to ensure that they are being opened in a highly secure environment and there is no data breach while those audits are happening. And these are just a few examples of the type of security concerns that the communication service providers are dealing with as they transition to 5G.

When we look at the underlying technologies that address it-- and again, I'll bring back the idea of a consistent security framework that you need across multiple domains-- the key technologies that address these use cases, you want to find not only the technologies that are highly secure, have the smallest attack surface, but also address multiple of the underlying requirements. So, for example, we talked about key management as part of the core control plane. In addition to that, you need trusted execution environment for lawful intercept, for billing records et cetera. Then you need crypto acceleration for edge to core communication. You want to ensure that even in the control plane, your crypto implementation is done in a highly efficient and performant way.

You want to make sure that not only you are protecting data, but also IP and the application core that's involved in-- especially in a multi-tenant environment.

So, when you look at all of these capabilities, you know, Intel has been at the forefront of this-- of security implementation for a very long time and the approach that Intel brings to the table is holistic. All the way from essential security capabilities that you need at a platform level to workload security, regulatory compliance, as well as, you know, implementations and references that show how you can implement it as you go from virtual machines to microservices.

And the approach that is crucial here is a combination of hardware and software that ensures high integrity, smallest attack surface approach that has capabilities, the underlying capabilities that address multitudes of the essentials to address the use cases that 5G core requires.

At this point, I will transition to Kapil, my colleague here, to really dive into the specific technologies and how you can address these threats and these use cases with a holistic security approach that can be consistent across edge, core, and cloud.

Kapil, over to you.

Kapil Sood

Thank you, Chandresh. And if we stay on this foil for just one second, what we wanted to highlight here is the consistency like Chandresh said. Consistency is critical and so is the ability for the service providers and the software ecosystem to write new applications that can scale across from the edge to the core to the cloud, because we are seeing that evolution in software. We are also seeing requirements that are, basically, spanning, you know, across geos. We are looking at requirements that are coming from identity management and privacy, and we are looking to address them in a very consistent manner. And that's what we will walk you through.

As we also look at the Intel security portfolio, we focus on two important things. We call that the Foundational Security. The second one is Workload Security. And they both have to do with, you know, ensuring that the infrastructure is as secure as we can make it, that you can provably ascertain that the infrastructure that you're running is what you intended to run, and that the workloads that are performing are performing at the highest performance with the highest security.

Workload security is a very rapidly increasing paradigm. And like Chandresh said, protecting keys and user data and the owner's data, you know, things like CDRs, things like billing records, and charging, other components, they need to be protected, so does your user data. And we will walk through the details of Intel SGX and how it provides a consistency and the security that we need for these workloads.

So, let's move onto the first topic, which is Intel Platform Firmware Resilience. And as Chandresh very nicely highlighted, 5G is a highly distributed system. It's new sets of deployments that are going to happen, that are already happening indeed. And as you know, they are

happening all the way across from the edge to the core to the cloud, and they are happening in a very distributed fashion, even including multi-cloud.

So, every infrastructure today is critical infrastructure. Now, that is important because a single breach in any of these areas or any of these deployed locations could jeopardize and expose the entire network. So, we put a fair amount of emphasis on ensuring that the infrastructure, all the way from the silicon up, is highly protected and we do that with a technology called Intel Firmware Resilience, which basically ensures that we get infrastructure security at the highest level today.

So, this is basically illustrating the various features that we have in the Intel Platform Firmware Resilience feature, and the most important thing is that we take the entire platform, all the components on the platform into the security boundary. We make sure that we are able to authenticate the devices that are actually connected to the platform, and the firmware that runs on the platform to ensure that the infrastructure owner knows what supply chain components are coming in, and are they authorized to be deployed in your deployments.

So, the foundational element in the Platform Firmware Resilience, is the platform Root-of-Trust. And in addition to it, we also provide platform firmware resiliency, which means that in case of corruption in the platform, or maybe an installation or a reboot that didn't happen properly that it falls back to a good known version of the firmware. We cover the CPU. We cover the BIOS. Platform firmware itself. Things like the various Flash descriptors, the BMC firmware, the integrated NIC firmware, things that are installed at the foundational start of a platform has to be authenticated, and that's what this does.

It's also important to note that Intel has worked very closely with NIST on various standards and regulatory environment developments and specifications. And this platform firmware resiliency feature addresses the NIST SP800-193 specification requirements. So, by doing so, we have ensured that we follow other industry standards as well to make sure we protect the supply chain, we make sure that the devices and the firmware that are connected to the platform are authenticated to the CPU, and the BIOS itself is authenticated, and that can be tested as well.

So, now, let's move on to talk about Confidential Computing and Workload Security. And that is something that we have been hearing a lot more of, and Intel has been at the forefront and leader in this.

What is Confidential Computing and why do we need to really protect the core and data in use? So, confidential computing, essentially is a paradigm, is a new emerging, in fact, rapidly emerging paradigm where software, ISVs, infrastructure owners, everybody wants to provide protection for the user data, they want to protect the user or the execution environment in which those software are running, so that as to not expose anything that is being installed on your deployments, or on the cloud, or edge, or anywhere else.

So, what the confidential computing designs to protect or aims to protect is the malicious insiders with escalated privilege attacks. Now, as we start to see this, you will start realizing that now you actually start getting into the attackers that weren't previously considered attackers, right. Those who are, you know, insiders who are considered insiders because the infrastructure was secure with the perimeter security, and it was owned by a single entity. But now with the distributed architecture of 5G, that the deployments will happen on infrastructures that are not owned or operated by-- necessarily by the person or by the company who's providing the services. So, insiders are an important part of that protection scheme.

So, are the prevalence of hackers. I think Chandresh illustrated that very well earlier that we are seeing more and more attacks. There are bugs that attackers, hackers look to exploit, and as they exploit a bug in the operating system, they very rapidly go into the workload which is running in the user space, and they can eventually compromise the data, the execution modes, and various other things that can go wrong.

And then there is the emerging case scenarios around the third parties that are accessing data without owners' consent. And we see that predominantly happening-- everybody sees that predominantly happening. And so, these are the three big vectors for confidential computing.

And what we also see is that our applications, in the cloud, or even in infrastructures are constantly under attack. And they are under attack from the various bugs in the operating system, rom the hypervisor, from other applications. There is an emerging threat vector

that we are looking at is around the physical security on the platform, so to prevent scopes and other components from injecting or reading data coming out of the system buses. Host administrators or the VM administrators. So, there is a whole suite of attackers that has now been considered a security threat that they weren't in the past.

So, confidential computing is a paradigm, is a new computing-- is the new way of computing that we expect is going to change how people protect their code and data.

OK, so Intel SGX is Intel's Software Guard Extensions. Intel Software Guard Extensions are one of the premier ways of doing confidential computing, and we are happy to share that we have had Intel SGX on our platforms for a while now, and very recently we introduced it for the first time on our Xeon Scalable processors, on the 3rd Generation Xeon Scalable processors. Now, let's look at how SGX actually protects your code and data.

So, as we saw in the previous foil, your application is constantly under attack from the operating system, hypervisor, remote administrators, anybody that can get access to the system. And so, it is critical to have the most sensitive parts of your application, for instance, keys, private keys, and the code that operates on those private keys in a highly secured environment. We call that SGX Enclave.

And what you see there on the right hand side in the diagram with a lock on top, is the SGX Enclave, which is running, let's say, your key management application, like a PKCS 11 cryptographic library. Now, we all know that those cryptographic libraries process keys that are used for generating private keys and for processing the key management functions. So, we want to protect those and put them in the enclave.

Now, when that code, when that key management code is running on the enclave, it is running in a specialized environment on the Intel SGX, and Intel's CPU prevents any of these other attackers from looking into the code or data, or jumping into arbitrary parts of the code and start executing, or even dump memory.

So, it is a highly secured environment. Intel SGX has been-- it has been out there in our products for about four or five years now, and it has been very thoroughly vetted by the industry. We are, basically, providing-- Intel is providing a number of reference architecture and libraries for use of adoption of Intel SGX, and we hope to address a lot of the attacks with Intel SGX.

OK, so this basically gives a view of how Intel SGX operates and why it is important. So, Intel SGX is now available on all Intel SB servers as well, and it will be a capability that will be now consistent across from the various networks, whether they are deployed in the edge, cloud, or core, or enterprise. One of the advantages of Intel SGX is that it runs bare metal workloads that can scale as the application scales. So, Intel SGX allows applications, whether they are multi-threaded or individual processes to be able to run on bare metal or in VMs. SGX provides protection in virtualized environments, which means that you can run Intel SGX inside a VM as well.

Intel SGX gives customers-- the developer community granular controls over how they can actually manage their code and data, which means that Intel SGX, basically, allows that enclave where the ISVs can actually run the code of their choice inside the enclave. And there are a number of POCs and reference code that Intel has. There is a lot on the user space, open source community, and all that can be leveraged. Intel SGX, like I said, has been proven with the real world deployments and has been battle hardened, which means that we've had great opportunity to have, you know, the academia, the researchers look at it, provide us feedback, and we have enhanced Intel SGX, you know, as a result of that.

Intel SGX has a very unique feature, which is called Attestation. And this attestation feature actually scales for the cloud and edge. So, what attestation means is that as a developer, as an owner of that code, you want to ensure that that is the right enclave that is running on the Intel CPU on the platform that you intended it to run on.

And so, SGX has a capability where it can provide a platform level attestation, so the Intel CPU will sign the SGX code and data for the customer and deliver the attestation code, so that the ISV or the owner of that software can ensure that that is the right enclave that they intended to run. So, once that trust has been established, then the proper keys, billing records, or anything else can be activated or deployed in that enclave.

As we talked about the trust boundary, SGX provides the smallest attack surface, which means that the smallest of your security libraries, or the most critical of your security functions, all the way to the entire containers can be running inside the SGX enclave. So, the trust boundary scales from how little you want it to be, to how big you want it to be, and that is a differentiation where customers are in control of how much and how they want to protect in the cloud. So, by doing components or operations inside the enclave, everything else outside of the enclave is outside of your trust boundary, which means that those components could be... even if those components are compromised, like your operating system, they will not be able to read any memory or execute any code inside the enclave at arbitrary points.

OK, so let's take an example of how Intel SGX is protecting and securing the 5G core with a service mesh. So, service mesh is also an emerging paradigm and it is one of the ways of providing a very horizontal, consistent software architecture and that allows the 5G functions to scale. And service mesh deployments are synonymous with containers and microservices. It is a mechanism where, you know, it basically allows the various 5G functions to communicate with each other in a very consistent manner, in a secure manner. And therefore, this is an important piece that should be protected with Intel SGX.

So, what we are showing here in this reference application is on the left hand side, we have a Centralized Key Server. So, this centralized key server could be a key vault, it could be, you know, a key HSM, it could be those sensitive key servers where you have-- the COSP has their most pricy fits.

Below that, we have a centralized server, attestation services, the Intel DCAP library. So, like I mentioned, Intel SGX-- Trust and Intel SGX Enclave is through attestation and this is the attestation server that allows the owner to ensure that the enclaves and the other software that's running inside the enclave is actually running on Intel CPU. It is signed by the Intel CPU, and you can verify that attested code in the key server-- in your attestation server.

Moving to the right hand side, we see a number of 5G functions, the AMF, SMF, PCF. It could be other functions as well. But what this is illustrating is that these functions communicate with each other over the 5G core service based interfaces, and these interfaces are TLS protected. And TLS is the predominant Transport Layer Security, which provides a secure channel between the various network functions, like the AMF, and the SMF, and the PCF to communicate with each other.

So, what we are showing here is the purple boxes, let's say, the SMF and the PCF have to communicate. They will create a protected TLS channel and this will be done usually through mutual authentication, which means that the SMF authenticates the PCF, and the PCF authenticates the SMF, and then they create a security channel over which they can communicate.

Now, as we heard earlier, TLS requires private keys, and private keys are critical in-- for signing operations, as well as for key decrypt operations in TLS. So, since the TLS operations are now handled within the service mesh, what we are doing now is protecting the service mesh with Intel SGX. So, we have a set of libraries and we are enabling with Istio and with Envoy, which are components of the service mesh to take advantage of Intel SGX, where these components can use standard interfaces, standard cryptographic interfaces to execute those critical, sensitive, you know, key management functions inside the SGX Enclave.

The most critical part of this is that the keys are never in the clear. So, on the left hand side, the centralized key server will securely deliver a key into the SGX Enclave. Let's say, this is the STOD signing key, and this key never leaves the enclave. It never leaves the key server on the left hand side in an unprotected fashion, and the one that's inside the enclave is decrypted and it's used only inside the enclave. So, we are providing key security at data protection, at rest, in transit, and in execution.

So, this is an important use case. We believe it is prevalent and it will be applicable across the various deployments, and it can be ubiquitous wherever we want to run our 5G applications.

Now, this is an example. It can be extended to other usages around data privacy, edge, analytics. It can also be applicable to billing records, LI, and others. So, these capabilities, the basic building blocks that Intel is providing here are broadly applicable.

OK, so security is extremely important and critical for Intel, and so is performance. We have happily shared that we have substantial improvements in our crypto performance, and Intel has a new, enhanced instruction set on our Intel Scalable processors, which offer,

you know, performance in the orders of leaps and bounds from the prior generations, all the way from OpenSSL RSA operations to the TLS encrypted connections per second, and for AES-GCM, including CRC.

Now, let me walk you through what these actually mean and why they are important.

So, RSA operations are handled and are required in TLS for key handshake. So, the operations that we discussed in the previous SGX, or where the mTLS keys are established, those keys are established using the RSA encryption and signing operations. These are highly expensive operations in terms of CPU cycles, so if their performance is increased, it increases the overall performance of all the mTLS connections that are being established on that platform and all the other servers.

Now, SGX can utilize all these cryptographic enhancements inside the enclave. So, running SGX on our Intel Scalable Xeon processors with these instruction enhancements gives you SGX security and the performance that comes from our enhanced instruction set.

Connections per second in TLS are extremely important, because-- excuse me-- because they allow your servers to scale to the extent that you want them to scale. So, the more number of connections per second that we can perform, the better it is for the customers to scale their existing software and their infrastructure. So, 4.2 times is a significant improvement from where we were previously.

Excuse me. IPsec GCM-- AES-GCM, so AES is an encryption algorithm, which is actually used for encrypting data at high speeds. And what we have seen is that with the industry moving to AES, predominantly the different modes of AES as well, that GCM mode, the Galois/Counter Mode is of prime criticality in TLS as well as in current infrastructures that are being defined and deployed for 5G.

We have listed some performance numbers here. Feel free to please check those out on our web. We also have details on Intel SGX.

And with that, I would like to hand over control to my friend, Chandresh.

Chandresh Ruparel

Thank you, Kapil, for a fairly detailed overview of both crypto acceleration and SGX capabilities. Again, just to summarize, you know, we all understand-- I think all we have to do is look at newsbites streaming through, and you can't escape the fact that the frequency and impact of cyber-attacks is continuing to increase. And it does impact customer confidence with both immediate and long-term impact, especially as the carriers go to a move from 4G to 5G core, and implement cloud native, a multi-vendor strategy where you want to bring best of breed capabilities into your network to address monetizable services, and bring more operational efficiency with cloud native deployment. The security exposure, obviously, is higher. The attack surface has increased significantly. And a highly secure hardware and software framework that can be consistently applied to edge, core, and cloud, is crucial.

Now, you are going to have some variances based on the domain in which you are implementing these, but those are... you know, you want to keep them relatively minor. You want to have the smallest attack surface implementation and capabilities that address trusted execution environment, that address key management, crypto acceleration et cetera across the board. And Intel does provide a highly secure hardware-rooted solution stack that allows the smallest attack surface. And it provides a holistic framework that protects data at rest, inflight, and while processing.

And the key technologies that we went over today include the Security Guard Extension, SGX, Intel Platform Firmware Resilience, PFR, and the instruction set in Xeon Scalable processor for crypto acceleration.

So, with that, I would like to pass it back to Shawn for Q&A.

Shawn Li

OK, thank you Chandresh and Kapil. We've got some of the questions coming in, appreciate. Let’s start with the first question. “There are other technologies in the market that provide VM application and memory accelerations. How are they different from what Intel is offering?”

Chandresh Ruparel

So, that's a great question. As I mentioned earlier, I think the focus from Intel's standpoint has been how do you provide the smallest attack surface. It has to be a highly secure framework. Because every element in the 5G infrastructure now needs to be a security device, and if you leave vulnerabilities at hypervisor level or guest OS or application level, the hackers will take advantage of it. They are looking for those. So, I mean, it's pretty obvious that they will look at vulnerabilities and try to exploit it.

So, you want a solution that is-- that has the smallest attack surface, is able to address critical requirements such as key management and trusted execution environment for VM, application, memory, crypto isolation. And these capabilities also have to be battle-tested, meaning you have to have some experience in the field working with partners, working with carriers to ensure that there is robustness and maturity to it, and that's exactly what Intel is delivering with its solutions and security framework.

Shawn Li

Great. Thank you, Chandresh. Another question. "How does confidential computing address regulatory and industry standards?" Kapil, you are on mute.

Kapil Sood

Thank you, Shawn. Yes, I was on mute. As we talked about earlier, confidential computing is a paradigm that is fast evolving and with the cloud that is being deployed, you know, confidential computing allows applications to run in an isolated environment, protected from the infrastructure. So, regulatory and industry standards are extremely important, right.

So, the first... first of all, you know, when our customers want to deploy a commercial product that has to meet certain regulatory requirements, it is of primary importance that they use the right ingredients to build that, because they can get those assurances, get those regulatory approvals, and they can, you know, scale that deployment across the entire area. So, regulatory requirement is prime and we, you know, within Intel, we have done our due diligence to ensure that with the critical and code-- critical code and data that can be protected inside the enclave can address the capabilities and the ingredients that are required to build these regulatory approved products.

Now, Shawn, I also wanted to emphasize that we have, Intel has, worked with the industry very closely over the last decade or more. We worked with NIST, we also worked at CNFV, we have worked at other forums to make sure that the right requirements and the industry requirements are understood, that our technologies, you know, make sure that we can help our customers address those requirements when and as they become important.

So, regulatory requirements are becoming important. Industry standards are critical. And I want to actually also emphasize, SGX as an example of confidential computing, based on the requirements that we've seen from NIST and from the various standards/bodies that it addresses, those-- or it helps address those requirements by providing that smallest attack surface for the enclave. And with Intel PFR, which is the Platform Firmware Resilience, that addresses the NIST requirements for ensuring that the supply chain remains protected, that the infrastructure owner can attest and verify the devices that are going into the platform and how the firmware is being authenticated.

So, we take a holistic view of compliance. We try to make sure that all our securities are helping our customers in the best way possible to address those needs.

Shawn Li

Great. Thank you, Kapil. And another question. "Telcos have traditionally addressed security with... perimeter security devices and hardware security modules, are they obsoleted now?"

Chandresh Ruparel

So, just to confirm, is the question, are they obsoleted? Is that correct?

Shawn Li

Yes.

Chandresh Ruparel

Is that the question? OK.

Shawn Li

That's correct.

Chandresh Ruparel

So, yes, I'll take that. The answer is no, they are not obsoleted. It's necessary but not sufficient. What we are seeing now is that because of the web-like infrastructure, you know, web based APIs, you have HTTP/2, you know, being used for communication between network functions et cetera, and the highly distributed, you know, edge with locations that are not necessarily secured by the telco infrastructure. It's very important that we treat each element as a security device, because any one of them can be exposed and with it, it can expose the infrastructure. And so, the point is that the perimeter security requirements around firewalls, security gateways, depending on the location and application being deployed will be required, the implementation models are changing, but it's not sufficient. The idea is that make sure that every element has a solid security framework so you minimize the risk, right, with the smallest attack surface. That's the point, so it doesn't obsolete. The security requirements have increased because of the bigger attack surface.

Shawn Li

Thank you. Thank you, Chandresh. Another question from the audience. "How does Intel SGX integrate in Docker and Kubernetes? Also, does it support environments like AWS, GCP, if underlying hardwares is Intel Xeon?"

Kapil Sood

Good question and very pertinent as well. So, in short, yes. Intel SGX integrates with Docker and Kubernetes. A little bit of explanation there is that as we talked in our use case that we shared about 5G service mesh, the underlying infrastructure for orchestration is Kubernetes, and Docker containers are being deployed. So, we are working with very rich ecosystem partners-- a rich system of ecosystem partners to integrate SGX into their software, so whether it's service mesh, Kubernetes deployments, Dockers, yes, Intel SGX integrates in multiple ways as well. So, as we talked about earlier, containers can use small enclaves to just protect the keys, just like we talked about an example on service mesh and how the Envoy can actually-- or the STOD agent can actually protect their keys for signing inside the enclave.

The other model is where the entire container can run inside the SGX Enclave. So, here, this is important because on our latest Scalable processor, the maximum available Intel memory for SGX-- memory for Intel SGX is one terabyte, which means that a lot of larger portions of your software code can run inside the SGX enclave. There is no platform limitation to do that. So, you can scale, you can run containers that are even larger inside the SGX enclave at a high performance.

Now, we also work with our partners on SGX. So, we have worked with ISVs, our partners in OSVs, OEMs, and our CSP and COSP partners, and the ecosystem to make sure that SGX is deployable and ready when it becomes available on Intel Xeon servers and when those are deployed in the field.

And we do that in a number of ways. We, basically, have reference architectures, we upstream our drivers, we make sure that all the software is open source that the customers can use, so we do expect broad availability. There have always been-- there have recently been announcements with some of our CSP partners on Intel SGX, and we are also looking to scale this, you know, to all the other partners and customers.

Shawn Li

OK, thank you. Thank you, Kapil. And this is our last question. “Have you extended the security supervision, monitoring, and analytics to 5G NDL related to NF applications context and business logic separations, and stored it as structured and unstructured data and the supervision analysis through the support for thematics, as well as the MA traffic 3 IGG and N3 IGG through the RG and ATSSS functions?” This is kind of a long question, and thank you in advance. Thank you.

Kapil Sood

OK, so, yes. I, unfortunately, don’t understand a lot of these acronyms, so pardon me for that. But SGX can be utilized in a number of different applications and scenarios. So, analytics is a prime example. We have illustrated and worked with partners to make sure that we can do, you know, multi-party analytics in a confidential environment.

Now, what that use case means is that parties that won't necessarily share data with each other, say, they are hospitals or they are doing some sensitive analysis for, you know, for any problem they want to solve and they don’t want to share data with each other, they could actually use Intel SGX to pool their data into the enclave, and the enclave runs the analytics on that data without exposing that data.

So, this is called Privacy Preserving Analytics, and it is actually one of the predominant use cases of Intel SGX that our customers are using, especially in the healthcare industry, and it can be used in other areas as well.

So, that’s another use case, you know, for using SGX.

Unfortunately, I couldn’t get to the other acronyms, so my… sorry, I couldn’t answer that completely. Chandresh, do you know any more on that?

Chandresh Ruparel

I think I’ll pass on that. I think there is also a contextual thing that is not necessarily coming out in the question completely on those acronyms, so I’ll leave it at that. Maybe we will respond offline.

Shawn Li

Also, we will have contact information and the audience could send the questions to our speakers. And thank you.

Chandresh Ruparel

We would like to thank everyone.

Shawn Li

Yes, that is the last question for today, and thank you for joining us today. Please do not forget to give our team a rating for the live recording so that we may continuously improve the quality of our webinars, and be sure to join us next time, Wednesday, September 1st at 8 a.m. Pacific for the Intel IPUs Fundamental Role In Your Cloud Strategy.

Thank you again for joining us today. This concludes our webcast.