CORPORATE PARTICIPANTS

Lilian Veras
Moderator

Roland Schröder
Ericsson – Cloud & Platform Solution Manager

Hassnaa Moustafa
Intel – Principal Engineer, Network Platform Group

PRESENTATION

Lilian Veras

Welcome, everyone, to the Intel Network Builders webinar program. Thank you for taking the time to join us today for a presentation titled; “Service Enablement at the Edge with Ericsson and Intel”. 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’ll 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 Ericsson is a Titanium partner. Learn more about our INB Winners’ Circle program by clicking on the link in the Attachments tab.

Today we're pleased to welcome Roland Schröder from Ericsson and Hassnaa Moustafa from Intel. Roland is Cloud and Platform Solution Manager at Ericsson and has over 40 years’ experience in the telecom industry, being part of the network digitalization, introduction of mobile networks, and recently the move to cloud-based deployments. He's currently working with cloud native infrastructure deployments in the telecom network.

Hassnaa Moustafa is a Principal Engineer at Intel Corporation working on 5G converged Edge products enabling IoT and AI services at the Edge. She led Edge reference solutions by Intel, enabling partners to rapidly build commercial offerings for Smart Cities, industry 4.0, connected health, connected vehicles, and immersive media. Hassnaa is passionate about working on cutting-Edge technologies that make the most significant difference in peoples’ lives. She has over 80 publications in international conferences, journals, and books, and she holds more than 25 issued patents.

Welcome, Hassnaa and Roland. Thank you, again, for joining us today, and I will hand over to Hassnaa to start off. Thank you.

Hassnaa Moustafa

Thank you. Thank you, Lilian, and thank you, Roland. I'm glad to be here today with Roland to share with Roland how Intel’s partnership with Ericsson accelerates and eases services enablement at the Edge. So, I'll give you a quick introduction first on how Intel and Ericsson see the Edge today and how 5G, and even beyond 5G, will accelerate a lot of new services at the Edge. I'll give you, also, an introduction on Intel’s Smart Edge Open framework, which accelerates services convergence, network and application services, and

accelerates, also, 5G virtualized functions, convergence with Edge services. And Ericsson also will speak about the Cloud Container Distribution and we will show you how Smart Edge Open and Ericsson CCD are partnering to enable easier services deployment at the Edge. We're happy to have feedback and questions, of course.

So, let me start first by speaking about the opportunity and the challenge with Edge services. So, as we know, and as we all know, we're now in the transition era for Edge computing from standards and concepts that have been there for a long time, and with 5G, we'll start to see realization and deployment for what we planned for several years back. So, Edge computing is becoming reality, which is good in the sense of opening new opportunities for Edge services, and also helping new players come to the business chain for providing Edge services. All this is good. It's also good that Edge services today can benefit from existing cloud native software frameworks, but this comes with some challenges. From an Edge service perspective, there is no one solution that fits all types of Edge platforms. It's not like the cloud environment. No, Edge environment is different, and it varies in Edge deployment, location, in platform capabilities, which creates sort of a challenge in terms of there is no one solution that can fit all types of Edge deployments. There is a need for diverse capabilities for the Edge solutions, for the Edge services to be agile, for deployment on any Edge location, and on different Edge platform capabilities. There's a need for, for example, seamless data ingestion at the Edge. It's not the cloud environment here. It's the Edge environment so we need seamless data ingestion. We need also AI workload optimization. When the AI workload runs at the Edge for the Edge services, it’s different from the cloud environments. New optimizations are needed. Another example, managing data storage at the Edge, again, is different from the cloud environment. So, all those new factors bring challenges for the Edge services deployment. Although there is a big opportunity, there are still technical challenges to resolve, and that's why we're here, Intel and Ericsson. We're partnering to help remove these challenges and ease Edge services.

So, how do we do? Intel offers a comprehensive set of tools that Edge solution providers can use to build their solutions, bridging the gap between the cloud and Edge. How? Through abstracting underlying hardware complexity on Intel platforms; through, also, abstracting the complexity of the network functions for solution providers who don't need to deal with the network functions complexity; and also, through offering optimization for different software, whether it’s network services software or whether it's AI workloads, offering optimization on Intel hardware.

Ericsson, as well, offer a telco-grade Kubernetes distribution, with functionality supporting Edge deployments. Here, again, bridging the gap between the cloud and the Edge through what Ericsson is offering here, and leveraging different optimizations from Intel on Intel hardware. So, this partnership between Ericsson CCD and Intel Smart Edge Open capabilities helps resolving some of these challenges and strengthening more the opportunity for Edge services deployment.

So, I'll give you an overview of how we see the Edge. As I mentioned, there's no one solution that fits all types of Edge deployments. As we see in this picture, we can have multiple locations for the Edge. Firstly, I'll speak about the on-premise Edge, which is triggered by the IoT, different types of services, and which can be low hanging fruit with 5G for ease of deployment and rapid deployment. We have also the network infrastructure Edge where we can have a telco network hosting Edge services, and these types of services are mainly small cities types of services, content streaming types of services. So, as we see, the location itself varies, and the services deployed on each Edge location are different types of services, looking at different verticals, whether IoT or media, or enterprise services. And accordingly, the latency itself needed and the latency requirements for each service varies per type of service at the Edge location.

So, we help all these types of diversity in terms of Edge deployment locations, in terms of platform capabilities, which differs from Edge location to another, and on-premise Edge, server or platform, wouldn't have the same capabilities of a platform sitting as access Edge, for example, of a telco. So, we offer a broad set of software building blocks from Intel, optimized for Intel architecture, and abstracting all the underlying hardware complexity here, and enabling building Edge services for diverse deployment locations, and meeting the different requirements for latency that you see here on this picture. And we work with Ericsson to see how CCD can leverage the different optimizations that we are offering.

So, what's our approach in the different building blocks we are offering in the different software optimizations we are offering? It's as simple as our approach is a cloud native microservices approach. As I mentioned earlier, we're bridging the gap between the cloud and

the Edge. So, when moving the cloud resources to the Edge, we need to follow the same approach, which is cloud native approach, microservices approach, but also modular approach; a very modular approach to build Edge services. If we look at this picture, it gives a good summary here. If we look at it from bottom up, we start by the hardware platform and Intel Edge platforms of different capabilities, and after that, on top of this platform, we have Intel Smart Edge Open. Intel Smart Edge Open is a framework abstracting-- a cloud native framework certified by the CNCF and it follows the Kubernetes type of architecture, optimized for Intel hardware, abstracting hardware complexity, allowing the network functions to work together in a convergent manner with applications, and also provides orchestration for the different microservices we are offering, and helps onboarding and management for any microservices for Edge solutions.

So, Smart Edge Open exists in different flavors. In this picture we're showing the Developer Experience Kit. The Developer Experience Kit here is focusing on the services and applications. We have other experience kits for Smart Edge Open, for private wireless types of deployments, for example, or access Edge type of deployment.

So, on top, if we go a step above here, in this picture, we see a comprehensive set of building blocks and software capabilities that we are offering. I'll give an example for AI workload optimization. We're having the OpenVINO Toolkit, which is a big set of different other building blocks like OpenCV, like Deep Learning Streamer, like OpenVINO Model Server, offering different types of optimizations for AI, for deep learning workloads on Intel architecture, and supporting several types of neural network models. We have also Edge Insights for Industrial, which provides seamless data ingestion at the Edge, and also analytics and a media pipeline, helping building Edge services for Industry 4.0 and optimized on Intel hardware. We're offering also a scalable video technology for HEVC video on Intel Xeon and optimized. It's a software optimization here, reducing the latency dramatically for encoding and decoding, and in real-time. We're offering as well a WebRTC Toolkit, which allows real-time communication, and which allows also media to be streamed over RTP and RTCP, real-time protocols.

All those are modular building blocks offered as tools and microservices on top of our platforms, and can be easily onboarded through Intel Smart Edge Open, and can be easily orchestrated. We don't stop at this. We also offer reference implementations to show how all this can work together. We build in-house reference implementations and offer it to our partners to use and test with all the underlying capabilities shown here. And we allow this to be accessible. We publish all this through Intel's Developer Catalog. If you go out to Intel Developer Catalog, you'll find several reference implementations where you can download and test, and test all the underlying capabilities I've explained. So, that's our approach to build Edge services, to build Edge solutions, how we enable our partners, how we help service providers, how we help solution providers, even ISVs. They can leverage all these capabilities to contribute to the Edge services ecosystem.

If we go to a snapshot here on the different reference implementations I spoke about, and we offer through our Intel Developer Catalog, the reference implementations, think about them as a jumpstart solution for our partners, for Edge solution providers. They cover wide industry segments, including Smart Cities and smart roads, including industry 4.0, including even the health segments. It's wide industry segments. And in this partnership with Ericsson, we were able to test three of these reference implementations mentioned here on this picture for industry 4.0 for defect detection in factories, and as a reference implementation for intelligent traffic management in Smart Cities, intersections, and third, reference implementations for telepathology or automated pathology, how lab analysis can happen in an automated way. And as you see, each reference implementation is using a part of the building blocks I showed in the previous slide. We have the opportunity to evaluate those with Ericsson CCD in this partnership that we're very excited about.

I will do a summary of what optimizations from Intel Smart Edge Open help this partnership, and was tested and evaluated by Ericsson team within the CCD. So, first, let's speak from a platform level. Through Intel Smart Edge Open, we offered resource management capabilities that Ericsson CCD used, mainly Node Feature Discovery to detect the needed resources before onboarding any microservice, and this was useful for Edge services to make sure that the latency requirements I mentioned earlier can be met if we offer the right resources. Also, there is Topology Manager to admit as a microservices based on the resources availability, and NUMA

Awareness to make sure of the colocation of memory resources with the CPU resources for workloads, which is of sensitive latency. We were able also to test and evaluate different SR-IOVs and container network interfaces from Smart Edge Open, the ones mentioned here, and I can emphasize on DPDK with SR-IOV and also Calico CNI, those were specific requirements from Ericsson CCD. So, for resource management, SR-IOVs and CNIs, that's for the infrastructure level, but from a service level, we went even a layer above and we were able to demonstrate AI workload optimization with CCD through leveraging the OpenVINO Model Server I showed in the previous slide, which was tested with Ericsson CCD, and it can offer inference-as-a-service. Meaning multiple services running AI workloads can use OpenVINO Model Server for inference-- to send inference requests to OpenVINO Model Server, which enables scalability of Edge AI workload. And we offered three of the reference implementations we had, I discussed in the previous slide, that were tested with CCD as well, and we're very excited about it because we completed the full picture. Meaning, starting from offering optimizations on our hardware platforms, up until AI workload optimization, and showing also reference implementations, leveraging all the different optimizations here. So, it was an exciting journey and still continuing, and we're happy to share the outcome here in this webinar.

And now, I'll hand off to Roland. He can walk us through more details on CCD and how this partnership also leverages the different optimizations I'm explaining. So, Roland, to you now.

Roland Schröder

Thank you. I will walk you through the Ericsson Cloud Container Distribution that has been used in this Intel-Ericsson Corporation for this specific case, but Ericsson and Intel have had many interactions and many cooperations over the years. So, this is one in the long row.

Let's start with a little bit how the world looked like and why these communication service providers are relevant for Edge deployments. So, operators all over the world have a very distributed infrastructure. I mean, one example is AT&T that has 5,000 data centers and central offices. That, of course, makes them well positioned to do an efficient Edge deployment and then be close to the place where the deployments need to be. Another benefit is that the communication service providers have a distribution network. It may be for the radio part, radio access networks. It's also they have resident fiber possibilities, and of course, this is supporting and enabling the Edge connectivity that is needed when you start distributing the functionality out in the network.

So, how does Ericsson fit in this CSP picture? Well, Ericsson is a vendor that is providing or selling or enabling CSPs around the world, and we are the leader in this market, looking at market size. So, Ericsson is a trusted partner for the CSPs for telecom solutions, either on OpenStack, using traditional VNFs, virtual network functions, or, as we're talking about now, the cloud native functionality running on a Cloud Container Distribution. So, both when looking at the solution level, and on the CCD specifically, we see continued high growth in this area. And CCD, or Cloud Container Distribution, then is also now going to be distributed around the world and being used by the different communication service providers, and this, of course, means that this combination of Ericsson CCD, Intel functionality here, and the service providers together means that we have good opportunities to support and deploy Edge functionality around the world.

So, how do we look at the design principles for the Ericsson infrastructure and for the CCD? As it says on the top left here, 5G core has been initially the driver for introduction to cloud native and the CCD, but now we see that we have multiple applications coming also for operation maintenance, for business support systems, as well as specific Edge applications. So, you could say it started with 5G, and that being also the need for distribution networks, the Edge enabler functionality was there from the beginning. But we are now evolving to new applications as well.

One aspect of Edge is the large number of deployments. If you have a few big, for example, central office deployments in the telco world, you may have manual handling of, for example, software upgrades. But if you have thousands of Edge sites, then you really need to have a high degree of automation to handle, for example, the lifecycle management of the software in these Edge deployments. So, the integrated automation is a key thing when we do our development here as well.

We are open, and open also means that we apply an Upstream First functionality, or Upstream First philosophy. So, for long we have been focusing on being open, and open means that we're using the best-of-breed open source components, which are then selected to be telecom compliant. And in doing so, we say that we would like to use the best-of-breed, use the open source, and we do Upstream First. So, when we see a need for some functionality that we may not find in the open source, we then go and propose a new functionality to the community, and once they have accepted that, we also include that into our CCD solution. And this, of course, means that we have to be very active in the different standardizations. So, in the number of these open source projects, for example, IPv6 dual stack, Metal Kubed, Cluster API, MetalLB, Multus, and so on, Ericsson is very active and proposing improvements that we think are needed for the telco world.

Following open source also means that we are very active in the Cloud Native Computing Foundation, the CNCF, and we are certified in that one, and this is a way of securing the interoperability and the portability of our solution, and possibility into work with other components and other, for example, 3PP application vendors. The community has the first feature evolution and CCD then is matching this evolution by having frequent upgrades, securing that we are using a supported Kubernetes release. As you know, they have three releases per year and a certain 12-month support period. And also, to make sure that we are having the latest security patches. Security is, of course as usual, a very important part, and here we can see that we may have some regulatory telecom aspects, and also telecom operator aspects of security that we maybe don't find in all container-as-a-service deployments. So, Ericsson is working very hard to provide the required security functionality. It could be to have secured interfaces using TLS. We are hardening all our images and all our software, making sure that they are closing all attack surfaces that it might have. We also do nightly checks or security using a number of tools, and also benchmarking our level of security using things like CIS-CAT and so on. The final principle on this page is the workloads that CCD, of course, should not only support Ericsson workloads and telco detail workloads. We should be workload agnostic. As such, we could handle various workloads on our CCD. And by having the CNCF compliance, we have a controlled way of enabling onboarding.

So, we talked about the Cloud Native Computing Foundation, the CNCF, and here we see lots of small boxes and that is also showing a little bit about the selection that is needed to be done when we can pick our components for the CCD. And some of these are maybe more for the IT industry, or more for the telecom, so we need to look at what are the telecom requirements and then choose from this landscape. And we have chosen functionality needed for, for example, the Intel Corporation. We have things that have been mentioned to have high throughput, to have SR-IOV, to have the Multus in the container network interface. Also mentioned before, we have the Node Feature Discovery, that is part of the Kubernetes release. So, with this Node Feature Discovery done, we could see the hardware and the software features that are offered by the different nodes, and then, as was shown before, we could enable and optimize the deployment of different applications, making sure that the applications get on the node or the server that is actually fulfilling all the requirements that the application might have.

If you have a good eye, and a big screen, you can see the small Ericsson sign under the platform, so that indicates that Ericsson has certification, CNCF certified distribution. In this, we also have the Intel part with the Smart Edge Open, which is also then listed in this landscape as being something that CNCF has looked at and said, yes, this is fulfilling the CNCF requirements and way of designing.

Looking a little bit inside the Cloud Container Distribution, we could see some details of what has been selected from this big landscape that we saw on the previous slide. And this has been selected from all the options to make sure that we have functionality that supports all the tool functionalities in the Smart Edge Open toolbox to see that they could get what they need. For example, this container networking interface. If you see in the CNI box, the green box, we have a number of connectivity functionalities, and then we see SR-IOV, we see Multus, and we see Calico. That is part of what the Intel applications are needing.

In the Kubernetes, we have this IPv4, IPv6, the CPU pinning or the NUMA Awareness that was also mentioned that we could make sure that we can allocate workloads to a specific NUMA optimizing that work. And it's also in the Kubernetes part that the new Node Feature Discovery is located. We also have other parts you'll see. Next to the CNI, you'll see a CSI, container storage. Here, we support multiple options. What is relevant for the more distributed Edge part is the host storage where we use the CEPH functionality for that one,

meaning that we don't have any dedicated storage node. We could distribute the storage over the servers, the few servers, in an Edge deployment. We also have monitoring, which is of course an important thing, that via VictoriaMetrics, we can extract information about the CCD, and what was going on in that one, and it could be, for example, presented using Grafana to visualize what we are collecting.

Then, of course, we need to handle the fault management. We have the lifecycle management that we said was very important to do the optimization, and on the top we also have the identity and access management, which is part of the security, how you could access the CCD.

So, when we look at applications and how we onboard that, it's been proven that it requires quite a detailed understanding of the requirement for what the application needs from the CCD. So, we see that there is a certain degree of integration and certification of the functionality that needs to be done, to see that the performance is correct and so on.

To the right, to sum up this one, we have the CNCF certification. So, it's something that is aligning to CNCF, securing the interoperability. We had the Upstream First principle to only use open source functionality. We say it's 5G ready, and that was maybe the first application that was used. Within that one, we see a lot of the connectivity aspects that are supported. And we had the hyperconverged storage, which was then very optimized when it comes to small deployments. But also remember that it's not only 5G applications, but a lot of other applications that is using the same CCD. So, it's workload agnostic. Lifecycle management, important to remember that this is a key thing both for the initial deployment, but also for upgrades and the handling of the software. And that it's an industrialized container-as-a-service, fulfilling the telco grade, handling all the fault and performance management that is required. It's secure with hardened images, and having continuous integration, continuous deployment of the Ericsson generic microservices that are used within the CCD. And of course, logging is another important part of this.

So, we have said that Ericsson is very much for openness and for fulfilling the intention of having interoperability, and we have a long-term commitment. We are active in many different standardization bodies, Anuket is one of them. Anuket is coming as a merge of the OPNFV and the Cloud infrastructure Telco Taskforce, the CNTT. And that was created then to create a common understanding of new capabilities for infrastructures across the telecom industry. So, it was a way to, or is a way to bring things together, to get standardization of the interfaces, for example, to get interoperability.

To the left, in the red part, you could see a little bit about the baseline for how things are worked through. So, you start with a reference model and create a reference architecture for both, one for OpenStack and one for Kubernetes, and this lists the high-level components and their interaction. From this, we could specify a reference implementation, more details on how this should be done, and also then to look into how do we have a reference conformance to this one, what is required to be within this. From these specifications in red, we come to implementation, where we have actual reference implementation code, and in the green box we say that to fulfill the intentions, to make sure that they are aligned with their reference architecture, we have a test framework. And then you could deploy your code and you could run the test and you get a batch for this one. And as I was mentioning, there is one reference architecture or one for OpenStack and two for Kubernetes. So, they have a separate path for this one.

Anuket takes lots of different input from various sources. So, OpenStack, Kubernetes, are two key ones, of course, and we have also ETSI and other ones, and we process this, and we create these reference architectures. So, it uses the inputs from standardization bodies, from various operators and vendors, and in this way, then we align the way of how we could implement this in an interoperable way. Both Intel and Ericsson are active members in Anuket, and supporting and working with it to achieve these goals.

So, this is a very short information and presentation, and in the material, you could also then see. If you would like to know more about this, on the Ericsson side, you could go and have further information about the Cloud Container Distribution, but also general cloud infrastructure, and how the 5G and the dual-mode 5G core fits into all of this. There’s a good, interesting link for Edge computing and how Edge computing and the 5G together then pave the way for a success. And also, there is a link on cloud native, and that cloud native is now happening.

On the Intel side, we see what we heard earlier, links to the Smart Edge Open Developer Experience Kit, to the Private Wireless Experience Kit, and the Open Building Blocks, and which I found interesting, a little bit more details on the different reference implementations. So, by looking at the defect detection, the ITM, and telepathology, you could learn more about what the feature is, and also learn quite a bit on how it was implemented.

So, to sum this up, Ericsson CCD is the leading Kubernetes distribution in the telco community. We can deploy it both on top of an infrastructure-as-a-service to deploy it inside a virtual machine, but we also have received an interest for bare metal deployments is increasing a lot. So, the CCD can be deployed in both of these deployment options. Ericsson drives a number of initiatives to standardize the Kubernetes infrastructure, to make sure that we can have the interoperability, we can have the vendor mix, and match for the different communication CSPs.

Ericsson CCD supports the required resource management. So, we talked, for example, about the Node Feature Detection, the AI workload optimization was also mentioned, and the connectivity throughput with the high throughput using, for example, SR-IOV and DPDK to speed things up. This has been verified in the reference implementation demonstrations of these three features. So, the strategic collaboration between Ericsson CCD and the Intel Smart Edge Open enables an accelerated path for Edge innovations, and it's done by simplifying and optimizing application development for several players from ISVs and CSPs.

So, with that, let's move over to some questions.

Lilian Veras

Awesome, thank you, Roland, and thank you, Hassnaa, for sharing such great information with us. Great presentation. We do have a few questions that have come in while you were presenting. So, let's start with those. First question I have here is, how do you see Edge computing acceleration with 5G and beyond?

Hassnaa Moustafa

Thank you, Lilian. Interesting question. I can start and, Roland, please feel free to add. So, yes, as we mentioned earlier, Edge computing has been around for a long time, standards and concept. 5G, and of course, beyond 5G is triggering Edge computing true services to happen, as simple as offering the reliable connectivity, reliable connectivities to have an Edge platform sitting anywhere, that can be in the factories, that can be in a clinic, and can be surely connected in a reliable way, similar to a cloud environment where everything is reliable. So, reliable connectivity is something very important helping Edge computing services to be there at the Edge, and with the required service level agreements. Also, what we have in 5G, like network slicing, can ensure that every service can have the network resources needed, and again, adding to the reliability, adding to the service level agreements that can be met.

One more thing, also private 5G. Private 5G is helping a lot because we're talking with our partners, we're talking with the ecosystem, we have a segment like the health segment, for example, industry also in factories, they are looking at private 5G for several types of reasons. So, private 5G is helping and accelerating Edge services also to happen. So, that's how we see Edge services, starting to see the true deployment with 5G and, of course, beyond. So, that's from my side. I don't know, Roland, if you'd like to add anything here.

Lilian Veras

Roland, would you like to chime in?

Roland Schröder

Sorry, mute as well. So, it was well summarized, I think, that we have all these possibilities with 5G and reliable connectivity is, of course, a key part in this one. And to add on the connectivity part, also, the predictable latency or the low latency that we could achieve by having these Edge deployments is a key thing in this. For example, when it comes to augmented reality, virtual reality, and robotics

within the industry, or maybe gaming for the private user. So, 5G gives a good platform for doing these Edge deployments, and then making sure that the different applications get the characteristics they need.

Lilian Veras

All right, that's great. Another question we have here is, how the microservices approach can accelerate the interactive services like AR and VR?

Hassnaa Moustafa

Yes, that's an interesting question, also. Yes, AR and VR, if we think of AR and VR bottom up, how the service is built, it's mainly like there is a component which is media, images, video, and streaming, and also like, interactive interfaces and real-time. If we think of it, it's a composition of different services that existed before, but interacting together, taking a new form to create a new service. So, why do we start from scratch, for example, rather than leveraging what we have in real-time video streaming, for example, what we have in media overlay, what we have in getting information from sensors and leveraging them to produce a new type of frame, which existed in gaming already? So, if we think of AR/VR, microservices are crucial here in order to scale the deployment of these services, and in order also to have faster time-to-market. Rather than inventing new solutions, new real-time protocols, or new media streaming protocols, we can leverage all what we have from different types of services, and how to leverage and consume and compose, it's microservices, using microservices. I'll give an example. Something like WebRTC, web real-time communication, when it started early 2000s, it started for video conferencing applications. Now it's being used for media streaming, for live TV, for example. So, a potential in AR/VR, a big potential in AR/VR applications, and so that's how we look at AR/VR. It's a new type of service, that's true, but it's a composition of different technologies that existed in different verticals. So, again, microservices are crucial here, yes.

Roland Schröder

And I mean, also in the Cloud Container Distribution we use microservices and with the same reasoning that we could have smaller, reusable software components that we could combine, and also evolve in a more independent way compared to if we have a big monolithic software. So, it enables, say, a forced evolution and forced movement in the feature enhancement.

Lilian Veras

Awesome, great. So, Roland, a question for you. How does Ericsson support CCD software developers who want to optimize your service enablement at the Edge solutions on Intel platforms?

Roland Schröder

So, I mean, CCD is a telco grade container-as-a-service distribution. That, of course, then gives a good platform for the characteristics needed for a software developer that wants to have an Edge functionality. And it supports things like we have talked about, the high throughput, for example, SR-IOV and so on. We have the Node Feature Detection, that's secure, that a software developer could say that, OK, for optimizing my service in an Edge, I would like to make sure that this is deployed on the correct type of server, that has either an acceleration function, GPU, or whatever, or, for example, have the correct NIC card for a SR-IOV interface of 100-gigabit and things like that. And then, of course, also, we have these possibilities for small deployments in the Edge, where the hosted storage is one example. And then when you deploy your software, you also need to be able to manage that site and do software upgrades and so on, and also here, we say that we have a good automated lifecycle management functionality that's secure, that once you have deployed that you could maintain your Edge deployment as well.

Lilian Veras

That's great. We do have time for one last question, and the question asks, what are the key services that can benefit from Edge computing in the next two years, and how about in the long-term?

Hassnaa Moustafa

Oh, yes, that's a nice one, yes. We can say what we see for today, for the low hanging fruit for the short-term. We see industry 4.0 as a segment which has imminent need for Edge solutions and specifically with private 5G. So, we see big potential for industry 4.0. We see also like the health segment. The health segment has a need also and a requirement for Edge services, mainly to minimize the cost for communication between different premises and also for some requirements for storage and on-premise storage when it comes to medical records, and also different geos have different requirements when it comes to storage of medical records. So, we see the industry 4.0 and we see the health in terms of managing communication and synchronization between medical records across different premises and different campuses. We see Smart Cities. Smart Cities will come progressively because we speak about vehicle-to-everything connectivity, we know that a lot of constraints are coming also from different regulatory rules that differ from geo to geo, and having autonomous vehicles, and by when. So, V2X is a part of the Smart Cities use cases, which needs Edge computing, of course. But this might take a time phase to happen, different between geos, but we see something like traffic monitoring at intersections, and traffic management applications, and even the maps applications, all those are Smart Cities applications, and we see them as kind of like midterm in terms of wide deployment for Edge services.

For a long-term, we see a lot the need for telehealth. Telehealth in the sense of remote surgery, even of course with a different type of progression here. We see also like... in terms of the public sectors, we see a lot of applications for drones for different purposes by the public sectors, and here Edge computing and Edge communication with Edge data centers is something which is important. Entertainment has a different level also of services, with cloud gaming, of course, is there, but this is bending the cloud gaming platforms themselves, and their readiness to migrate from the cloud to the Edge. So, that's how I can summarize different services we see with different timelines for availability and commercialization.

Roland Schröder

Yes, I think you covered the most of the ones I had on my list. As you said in there on the gaming, we see that already being one application that could benefit from this low latency and so on, and typically it could have quite high end user interest then driving the market. Otherwise, I think my years in the industry have proven that it's very hard to predict what the new network functionality and capabilities will generate when it comes to new functionality and new applications. I was thinking about remote surgery as well, as you mentioned, even though I may not be really happy if I was laying on the operation table with a remote surgeon. But for sure, the functionality and the network could be something for the future.

Lilian Veras

Excellent. Thanks to both of you, Hassnaa and Roland, again for such a great presentation. I'd like to also ask our audience to 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. Join us next time and this concludes our webcast. Thank you.

Hassnaa Moustafa

Thank you.