CORPORATE PARTICIPANTS

Moderator
Intel

Ravikumar Alluboyina
Robin.io – Head of Platform Engineering

Rachel Chu
QCT – Associate Manager, Telco Business Dev Team

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 hyper-automation for disaggregated Open RAN in cloud native 5G evolution.

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.

Today, we're pleased to welcome Rachel Chu from QCT and Ravikumar Alluboyina from Robin.io. Rachel Chu is the associate manager of product manager, and alliance manager at Quanta Cloud Technology, leading the team responsible for building NFVI OpenStack, and cloud native reference architectures with partners including Intel, Red Hat, Wind River, and Robin.io. Rachel's role in QCT also involves nurturing strategic alliances and product go-to-market activities.

Ravikumar Alluboyina is the senior technical director and head of platform engineering at Robin.io. With more than 15 years of experience in storage and server management products, he has deep technical expertise in distributed systems and scale-out systems design. At Robin.io, he's responsible for building the orchestration layer for the Robin cloud native platform, and leading the teams. He's also instrumental in creating a new product line at Robin.io, the Multi-Cluster Automation Platform.

Welcome, Rachel and Ravi, and thank you for taking the time to join us today. Ravi, over to you.

Ravikumar Alluboyina

Thanks, Lilian. Hi, everyone. This is Ravi and I'm from Robin.io, and I have here with me, Rachel. She's from QCT. So, what we are going to talk about today is the hyper-automation for deploying O-RAN in a Kubernetes environment. Cloud native is synonymous to Kubernetes these days. So, it is O-RAN and Kubernetes, we are going to focus on these two topics. And you're going to hear it firsthand from the two companies who have done it live.

Let's jump right in. So, what are we going to talk about today? We're going to come up with some-- not come up, we're going to put forward some of the pain points in 5G vRAN deployment, and where does Robin fit in, what are the offerings from Robin.io that can

handle this orchestration at scale, and where does QCT fit in with their top-of-the-line hardware and the platform offering, and the hardware management tooling that QCT has come up with their bare metal SDK, and finally, we're going to present the value proposition of the combined solution from Robin.io and QCT.

So, what's the problem statement? Let's start from the very scratch. Let's say if you want to build out a far edge site, which will host a DU, and there is a CU somewhere running in the core or a regional data center. A problem statement is, given a set of bare metal servers, so we have to build the entire RAN.

So, RAN, as you all know, would have a DU component, a CU which could be a CU-CP or a CU-UP, depending on the vendor, and an EMs maybe, right? So, this is a journey from point A to point B. To get there, there are several challenges.

First is obviously the bootstrapping of the hardware. The second, the bootstrapping of the cloud native platform, the Robin CNP platform. And then federating multiple clusters and deploying RAN stretching clusters across different sites. Now, is this a new problem? No, I mean, this has been done for a decade. There were towers in 4G, 3G, and going back, even. So, it's a very well understood problem statement. There's nothing new in it, but what is new is 5G and cloud native. At every stage, there are several problems that need to be solved. Starting from bootstrapping the operating system, figuring out the SKU checks, and creating the cluster, configuring the cluster, and federating these clusters, there are small, but a large number of such small problems.

There are a lot of small problems in this. There is config management, there's operating system install, there's server management, there's storage management, there is network function management. So, these are different domains.

And the other problem is the 5G environment is scattered, it's distributed across a large geographical area, and this presents a scale problem. So, these are very small problems, a large number of them across sites and at scale. So, this is a problem that we're going to handle, and let's see how we're going to decipher this and start proposing solutions.

If you were to divide the problem statements, if you were to bucket them, there are two main challenges here. One is around orchestration. The second is around platform. So, if we VIN the edge, and VIN the orchestration, the RAN orchestration problem is pretty much solved. So, what are the orchestration challenges? There is infrastructure orchestration, cluster orchestration, and network function management. When we say orchestration and management, it's not just deployment, it's the lifecycle management of these entities. So, that is on the orchestration side, how do we manage these sites?

Orchestration itself is not enough, because at the end of the day, the DU will need to run on some hardware, sitting at the far edge, and there is a platform, cloud native platform, that is going to sit on that hardware. So, somebody needs to manage the hardware and the compute resources on that machine, bare metal server, or group of servers, network, storage, and also provide the observability. There is platform. In the platform, there's a bifurcation that you need a proper hardware to bring up DU, as well as a cloud native platform that will orchestrate the function at the edge, and an end-to-end orchestrator which is orchestrating all three components or four components of RAN-- CU, DU, EMS, RIC, real-time RIC and non-real time RIC, or maybe other applications. So, this is the orchestration and the platform.

Now, that's where we come in. So, who are we? Robin.io is purpose-built for deploying and managing the lifecycle of 4G, 5G, and enterprise applications on Kubernetes. We are a platform based on Kubernetes, and we are the world's first deployment of cloud native, true cloud native IT stack in production. We are headquartered in San Jose, and we have offices around the world, and we have a very highly differentiated technology. We have filed around 72 patents, and more than half of them are awarded. We do have a lot of marquee customers, both in the enterprise segment, as well as the telco segment. And along with the customers, we have strong partnerships at the infrastructure level, and the application level, which is the network function level. We are featured in many articles, and we were the winners of some of the telco awards as well.

Let's see what Robin brings to the table. So, six years ago, when Kubernetes was just entering into the enterprise market, there was a big gap in Kubernetes which is primarily around storage. Kubernetes did not have a built-in storage stack. It was very, very good at

compute orchestration, but there was no storage stack. So, this is the time when we started investing in building a storage stack for Kubernetes.

What does it do? It is purpose-built for certain application types, which are data heavy application types-- big data, AI/ML use cases. If you look at a telco, you will have applications like Hadoop, Elasticsearch, Spark, AI/ML applications. This storage is purpose-built for them. It is application-aware storage. It can run with any Kubernetes distribution. Kubernetes is fragmented today and it is going to be fragmented. There are 20 distributions out there including clouds, there will be 20 more going forward, and Robin cloud native storage can operate with any distribution of Kubernetes.

So, this is the storage piece. This storage piece is predominantly towards a core, solving the core part of the RAN equation. Now, after some time, we have identified that storage itself is not enough. We need a carrier-grade Kubernetes for running these telco network functions. This solution, we call it CNP, cloud native platform. This is a batteries-included solution for running telco workloads. We'll get into how we do that, but this is the solution for the edge, telco edge.

The third solution that we offer is called Multi Data Center Automation Platform. This is more towards orchestrating hundreds of thousands of these far edge sites. So, we're going to focus today on the platform and the multi-site orchestrator, MDCAP.

Let's dig a little deeper into the platform. So, what does the platform look like? The platform is based on Kubernetes, Kubernetes is a phenomenal compute orchestrator, and we have an enterprise grade storage stack with all the features that you expect from the standard arrays-- snapshots, clones, QoS, backup, replication, thin provisioning, and encryption. I mean, you pick anything in the storage world, we have it. It's a distributable volume stack. It's a scale-out stack.

Alongside storage for telco applications, we have also built a big networking stack with a lot of integrations with Calico, OVS, SR-IOV, DPDK, dual stack. These are the ones that are supported, fully supported today, and there's a lot of them getting added as we speak. As we onboard more and more network functions, they come up with different requirements. You know, MAC VLAN is another example. So, we're onboarding all of those network plugins. So, we have storage, network, and compute, and the workflow manager is the one that has the brains of orchestrating compute, storage, and networking. It offers a one-click experience for deploying network functions and managing the lifecycle. If something were to go wrong, which is disk failure, node failure, NIC card failure, all of that is handled by this application workflow manager, and if you want to snapshot your application, create a DR copy, create a backup, all of that is handled by workflow manager.

Now, what does it have? It has an advanced scheduler which is NUMA-aware, it has affinity and anti-affinity policies, and a lot of primitives that will enable us to deploy these functions.

So, that's platform. So, this is a glimpse of what are the real challenges in orchestrating the network function on a single server? We're not talking about a 100-node cluster. On a single server, what is the complexity? So, we are talking about dedicating CPUs for the operating system, separating the CPUs that are allocated for DU versus operating systems. We are talking about NUMA awareness here, not just for the cores and the memory, but also for the devices. It could be your network devices or the SR-IOVs. And we need multiple networks. There's Calico for internal network, there is OVS Bridge for management network, and SR-IOV virtual functions for the actual data traffic. And there is redundancy because we have to carve out two virtual functions from two different cards and assign them to a single core for redundancy. So, this is just a placement on a single node. Imagine the complexity when you're going to 100 nodes or 200 nodes, or managing 10K sites.

Now, that's complexity in deployment of the DU at the edge, but what about if you have 10,000 DUs and you want to orchestrate all of them? This is where the MDCAP comes in.

What does MDCAP offer? If you go back to the previous-- the very first slide, the problem statement, somebody has to do the dirty work, which is bootstrapping the site, and after that comes the service deployment. But bootstrapping involves a lot of steps. Installing the operating system, set up the BIOS, flash the FPGA, install the cluster, monitor, configure, license checks, a lot of these things. So, first things first, that has to be solved. That is called bare metal-as-a-service function in the MDCAP system.

The second thing, of course, is deploying RAN, which is across sites, when DU runs at the far edge, CU at the edge, maybe EMs at the core, and we need the supporting cast as well. The Hadoop cluster, Elasticsearch, log gathering, AI stacks, all of those applications need to be orchestrated, along with things like OSS and BSS systems. So, all of these can be done using different tool sets, but MDCAP provides a unified interface, a single platform that can orchestrate all of these elements. So, that is what MDCAP brings to the table. It's a single orchestrator for orchestrating anything.

How does it look like? From our experience, you have a bunch of radio sites, there are far edge sites, which are a lot of them, 10K sites. It could be a single node set up at the far edge, or three nodes or four nodes, depending on the resiliency that you want to build at the far edge. So, there are lots of DU sites, and there are some CUs which are regional centers maybe. But of course, there is core.

So, this is a typical layout, and if we were to orchestrate all of these things, there needs to be some method to madness, which is we need to construct a workflow that can talk to different pieces here, that can talk to a bare metal server, that can talk to a deployed DU, that can talk to a VCU, that can talk to a cluster at the edge, which can talk to the OSS system. So, if you combine all of these, what you get is a workflow.

So, at the core of MDCAP system, there is something called a workflow engine. So, everything is workflow driven. There are some built-in workflows, and the workflows can be added at runtime by the end user. So, the workflow that you're looking at on the screen is going to boot up the far edge clusters, it's going to boot up a MAC where the CU is running, and it is going to deploy a core data center as well, all in one shot. Obviously, it's not going to happen this way, but MDCAP is powerful enough to finish this entire pipeline, entire deployment in one shot. This is the power of MDCAP workflow management. We can define workflows for anything. Bare metal install, cluster install, cluster upgrade, site upgrade, changing the configuration of the bare metal, all of those things are handled as workflows in MDCAP. They're included as part of the platform.

Now, what does bare metal management look like? So, obviously, you have a bare metal server here. The first thing is provisioning that bare metal server, which can be core or IPC, we have bare metal profiles, firmware upgrades, storage partitioning, rate configuring. Every single detail that is required to set up a bare metal server as part of this deployment packet.

After the deployment comes the upgrades. These are day two upgrades, OS upgrades, security patches, kernel upgrades, firmware upgrades.

Then comes the monitoring. So, we need readiness, we need lightness, metrics, logs, reconciliation. There are a lot of monitoring aspects that we have to worry about for compliance. All of this is part of monitoring.

And we need admission control, which is RBAC, quotas on the system, there are schedules, and there is exclusivity, you know, lock down the system, don't allow anyone to test the system or quarantine the system. All of these policies have to be enforced on a bare metal server.

There is policy and fault management. Hardware, software, component failures need to be handled. Now, of course there is one big dashboard which shows everything in a unified pane of glass. So, imagine just for one bare metal server, we have to deal with all of these different aspects. Now, think about managing 100,000 sites, and a core cluster, which is 200 nodes or even 1,000 nodes. It's exponentially higher.

Now, where do we go from here? So, we talked about bare metal, we talked about RAN installation or a service deployment, O-RAN deployment across sites. If you step back, what is it that we are trying to solve? What does the system provide? So, we have to handle a physical layer lifecycle, clustering layer, network function, network service. There are MOPs, method of procedures, for out-of-band activity, this policy engine. The holy grail is closed loop automation. So, this is at a very, very high level, what end-to-end lifecycle management of the site means.

So, this is the end goal in the problem statement, the very, very last one where I want to deploy a RAN. A vendor comes in, he drops in the packages into the Artifactory. It could be Helm Charts in case of 5G and Kubernetes. We call that network function package, and it gets tested on a staging area. And this is a workflow for running the entire test suite in a staging cluster. It gets promoted to production,

and production would run the identical workflow as a test setup for deployment. No variable changes here, same identical-- whatever we have tested, we're going to deploy it in production. So, this is a typical workflow for a service deployment.

What is it? It's all intent based. The system works as an intent based system. It's based on Kubernetes. It handles multiple elements. There is network function package, topology, and location-aware. It's NUMA-aware. And all the other goodies, network function, affinity/anti-affinity policies, healing, scaling, and upgrades. So, all of this comes out of Robin's MDCAP.

Sorry. OK, to sum it all up, there were four use cases. There is RAN. So, this is a platform optimized for running RAN. And the CNP platform can host VNFs, which means virtual machine-based RAN, as well as CNFs on the same platform. In fact, they can run – we can collocate VMs and containers. This is for transition – any environment that is in transition from 4G to 5G. Now, it can be handled. MEC. Any MEC application can be rapidly deployed and managed using the platform. And core, all the complex workloads, take any complex workload that can be deployed and managed using Robin's solution.

So, this is in summary of what Robin offers for the 5G use cases. Rachel, all yours.

Rachel Chu

OK, thank you, Ravi, and thanks – thank you all for reserving your time to join to this webinar.

So, I'm Rachel Chu, I'm Business Development and Product Manager of Telco NFVI Solution in QCT based in Taiwan here. Before we directly jump into the solution introduction, please allow me to take a bit of time to share with you QCT and including who we are, where we are located, and what products we are offering to the telco market.

So, here, from this slide, you can see that we are QCT, Quanta Cloud Technology, and our parent company is Quanta Computer, established in 1988 and known as the largest laptop and server ODM manufacturer. We have strong and development team and comprehensive supply chain ecosystem partners with more than a decade experience on the design and manufacturing. And QCT is the subsidiary of Quanta Computer. We are a global data center solution provider and we combine the efficiency of the hyperscale hardware with the infrastructure software from a diversity of industry leaders to serve next generation data center design and operation challenges.

So, QCT serves cloud service providers, telecoms, and enterprises running public, hybrid, and private cloud. QCT design manufactures integrates with service cutting edge offering to our own global network.

And this is the slide we'd like to share with you, that's QCT's global footprint. We have our offices covering globally to better serve our customers. So, from here, the map, you can see that in the United States, we have one in San Jose, where we hold our annual event, and one in Seattle. In the EMEAsite , we are located in Germany. And in Asia, we have coverage in China, Japan, Korea, and the newest established in Singapore office to cover ASEAN related business. And of course, our headquarter is based in Taiwan here.

And this is the slide we would like to share with you, that slide shows how QCT has grown and facilitate the cloud adoption. So, you can see that back to 2007, we are server ODM focused, and provide the ODM direct business to the customers. And back to 2012, we start to generate, announce our own brand QCT off-the-shelf design servers. At the same time, we not only manufacture servers, but also we're building up our cloud solutions with our software-defined infrastructure solution with our ecosystem partners like VMware, Red Hat, and Microsoft to cloud service provider and enterprise market.

And moreover, we are growing with our customers, so back to 2016, telco industrial initial discussion how to build up the network infrastructure from the legacy equipment, or we call it proprietary appliance to the open architecture with an open interface. And which can avoid vendor lock-in and drive the digital transformation to the cloud-based network infrastructure, and providing the service in agility and intelligent way.

So, start from then, QCT designed our own – our workload-driven systems to fulfil various industry telecom, HPC, and infrastructure. So, from that, we also work with our valuable partners, Robin, to build up the telco solutions.

And this is our QCT telco end-to-end product profile. There are three product lines you could see here in the slide. From the right hand side here, we have the IronCloud, and the middle side here, you have the IronEdge, and the left hand side has the IronBox.

So, on the bottom, you can see that QCT has a comprehensive end-to-end 5G server portfolio that extends from data center to the edge site. And QCT's serves provide – servers provide a powerful compute, performance, and low latency data processing for Robin to fulfil the demanding requirement of the 5G use case, as Ravi just mentioned, including a vRAN or 5G, 4G, uCPE, or 5G core, or CDN use case etc from the original data center to the far edge site.

And in the IronCloud product line, IronCloud is QCT's network function virtualization infrastructure product line for telecom operators. Under strong partnership and cooperation with the NFVI ecosystem partners, QCT's IronCloud is built to decouple numerous functions, such as 4G EPC, 5G core, or a base-bound unit, virtual radio access network software, content delivery network, and MEC, multi-access edge computing service from the dedicated hardware and allow easy migration from legacy network infrastructure to the virtual machine, VM, and containers, reducing operators' cost and also in the service deployment.

And for the middle side, IronEdge here, it's the product line QCT built for supporting the OpenRAN in different scenarios including a centralized RAN, distributed RAN, and the server design supports the capacity required by vRAN workloads with O-RAN function split option 7.2, time synchronization, and virtualization.

In addition, the server leaves the flexibility for future expansion including a MEC application as well. So, all these key features, make QCT Iron Edge series product line perfect for OpenRAN.

And the left hand side, IronBox is a uCPE solution product line. We are offering the software-defined WAN solution with our partners to reduce the multi-protocol label switching usage and create a dedicated line between branches and headquarter for the network communication.

And here, with our valuable partner, Robin, QCT built up IronCloud-Robin Cloud Platform Solution. So, here from this slide, we show that the easiestNFVI architecture and highlight the scope of the IronCloud-Robin Cloud Platform Solution, so you can see that the IronCloud-Robin solution covers from the physical hardware and also the NFVI layer and the MANO-layer. So, in the NFVI layer, we are adopting Robin Cloud Native Platform, which is mentioned by Ravi as well, the CNP platform. And for the MANO part, we adopt Robin Multi Data Center Automation Platform, the management orchestration layer.

So, with this solution, the telecom operator can provide multiple virtual network function in an agility way from the physical to the virtual function, and then with management and orchestration software.

Then what is the use cases when we are talking about the IronCloud solutions? So, in previous successful case why QCT and Robin co-work with, is that there is APAC Tier 1 operators, and the use case they are adopting the solution including the QCT server and Robin platform and MANO MDCAP solutions, they are running containers platform on top of QCT's telecom optimized servers for 4G and also 5G virtual radio access network, and content delivery network service.

And with the Robin.io cloud platform and QCT's NUMA-balanced design server, and also we do the BIOS performance tuning, the IronCloud solution could provide a carrier grade infrastructure, which can offer the high availability Kubernetes cluster and auto-deployment and end-to-end orchestration from the data center to the edge site.

And for Robin's software on QCT's servers can orchestrate and manage the application with the network intense requirements such as OpenRAN, CU/DU, as well as the edge and multi-access edge computing application. Robin MDCAP can deliver bare metal to service orchestration at scale by using QCT's bare metal API, which I will mention in the later slide to automate infrastructure delivery and lifecycle management at hundreds and thousands of nodes, including deployment, configurations, and accelerators, and real-time operating system and network interface required.

Here, let me take one example use case to showcase QCT's end-to-end product portfolio in a use case. So, here is... we take the 4G and 5G vRAN use cases as an example. So, here you could see we come – from the bottom, you can see the servers multiple different kinds of SKU servers is why we recommend QCT different server systems in different locations.

So, from the right hand side, date center part. In a Central Cloud area, we will recommend the standard depth 1U/2U system for the high density compute network performance to provide the 5G core and the data center application services.

And in Aggregation – and also Aggregation Site and Far Edge site, we will strongly recommend to have the short-depth with front access design servers for a CU/DU, radio access network service. This is because that in edge site, most cabinets are small and less than 500-millimeter depth, and the operation and maintenance can only be accessed by the front for the easy maintenance and operation. So, in order to make the server fit in edge site cabinet, we design our servers, QCT edge servers with less than 450 millimeters and compliance with the NEBS Level 3 to have the wide temperature operations.

And moreover, the Robin Cloud Native Platform, NFVI and MANO MDCAP software stack has been validated into the solution package from the data center to the edge site, especially for the vRAN use cases.

And while we are mentioning the vRAN use case, there is a global community called Telecom Infra Project, TIP, it's one of the key communities to put companies and organizations working together to accelerate the development and deployment of open disaggregated and standard-based technology solutions that could deliver the high quality connectivity to the telecom operators. And QCT also joined the TIP project. And here from this slide, you can see that QCT's design servers have already got an award and compliance ready ribbon for the CU/DU including the top one – the top one is the indoor CU/DU servers with the 1U and 2U system.

And the bottom one is the outdoor edge servers. It can be wall mounted design to support wide operation temperature.

And these three systems have been included in TIP Exchange Marketplace. It delivers excellent expandability and high computing capability, making the space and performance requirements of telco deployments. So, if you have any further questions or would like to learn more detail of these servers, you could also go to the TIP Exchange website to learn more.

And let's take a further deep dive into the solution stack architectures. The left hand side shows that the reference architecture from the physical layer, virtualization layer, application layer, and the MANO layer.

The physical layer, we select two best-fit system SKU, which you could refer to the right upper side here. The right upper side here, we select two models, and these two models support the latest Intel 3rd Generation platform, we call that Ice Lake platform, and they are NUMA-balanced design servers. 1U system is a master node in Kubernetes platform and 2U systems are as a worker node in Kubernetes platform. And in 2U systems, it supports up to 10 PCIe expansion slots and also supports 24 NVMe or flash devices.

In order to the easy maintenance and most commodities in the servers are designed with toolless including – toolless design including the hard drive, SSD, PCU, PCIe, and OCP to make sure that operators and customers can be easy maintenance.

And moreover, we also have our security enhancement features with the chassis intrusion mechanism. So, while the system cover being opened by different kinds of urgent items, the system will trigger alerts and have the log event to notify the administrators to let them know that the system has been opened.

And one of the key value while we're building up the solutions with Robin is that we select the best-fit servers and switches to have the pre-integration and pre-validation to make sure the reliability of the infrastructure and optimized performance. So, here we show that we adopt 1U system as a cloud native platform master node, and 2U system as the worker node. In the way – 1U drives the cost effective, and 2U system offering the compute and network performance optimized for the resources pool.

So, we got the two recommended reference configuration in different Intel Generation platforms. The one is the latest one. Here, we show that's the latest Intel 3rd Generation platform, Ice Lake platform, support the PCIe Gen 4. And the other one is the previous generation, it's the Skylake and Cascade Lake platform. It supports the PCIe Gen 3. So, while telecom operators are doing the migration

from the 2nd Generation to the 3rd Generation CPU, different kinds of CPU platforms, we have done the interoperability on Robin CNP MDCAP as well.

And previously, we are talking about the management and orchestration, so according to previous engagement with the different kinds of telecom operators and customers, we learned that while managing cloud native infrastructures, telecom operators could leverage MANO platform with NFVI interface to get further information of the virtualization resource pool. However, the management's lack of the physical layer server part, which with the challenge and the pain points we got from the telecom operators, QCT and Robin we discussed further on that, and we can joint develop a solution, we called that QCT Bare Metal Software Development Kit Tool, and it will be integrated as a bare metal profile in Robin MDCAP. With the tool, it provides the features with bare metal introspection which can discover the system detail spec, including the CPU, memory, disk, NIC. And after discovering the spec, it can offer the pre-defined SKU validation to make sure the system config is same as which SKU previously telecom operators have been defined. And also, the BIOS and BMC patch update, backup and restore are significant during the telecom operator doing the operation and the maintenance. It's also being included in QCT bare metal profile tool.

And next one, we also add the hardware security enhancement features with the Intel Ice Lake to enhance the security product.

And the tool is building up with all Redfish interface and it can provide a northbound API service to either a VM platform or MANO platform to be integrated. We believe that the single pane of glass could offer a better user experience to the operators and have the monitoring and management all covered from hardware, NFVI, CNF aspect.

So, to sum up that QCT and Robin joint development on the IronCloud Cloud Native Platform solution, so QCT has strengthened our role as a 5G ecosystem enabler these past few years. We work with industry leading partners like Robin to drive the fundamental change in a way telecom operators plan, deploy, manage their infrastructure.

So, it also supports the innovative NFVI and the virtual network function technology the product and solution to realize the truly software-centric network infrastructure covering for the fixed line and mobile network on the unified cloud platform.

And the solution solves the challenge from complexity in multiple layer integration management platform, QCT and Robin, the solution offers the pre-integrated validated solution with optimized performance and centralized management platform for the telecom operator and multiple various use cases, including vRAN, 4G, 5G core, EPC , and the CDN as well.

And last but not least, Robin and QCT team have released our solution brief in Intel Network Builder Platform for who are interested in the solution products, please feel free to download it. And we also attach the solution brief and QCT product portfolio in the download document area. You could also refer to it to download the marketing collaterals to learn more.

And thank you for your time to join to this webinar. If there is any further questions regarding to the solution, please feel free to contact Ravi and me to learn further information , and welcome to browse QCT's and Robin's website to learn our products and solutions.

So, Lillian, let me hand over to you to see if there is any questions regarding our webinar information. Thank you.

Lilian Veras

Thank you, Rachel, and Ravi, such a wonderful presentation. We do have a few questions that have come in while you were presenting, so let's get started on the questions.

First question I have here is, "How QCT design optimize edge server for vRAN?"

Rachel Chu

So, the question is for me, right? So, from QCT's side, we join our – we joined the community including O-RAN and also Telecom Infra Project as I mentioned before, and while we... we also engage with our partners and customers, we collect a lot of feedback from the direct customers and from the community we also get the feedback and also the collection of observation. And based on that, we

designed our edge servers with a front access design and the NEBS Level 3 in order to put it up to our in the cabinet system, and also the outdoor environment.

So, this is how we work with. Thank you.

Lilian Veras

Thank you, Rachel, that's great. Ravi, I have a question for you too. Someone asking for Robin.io. "Is there a provision to restore config and detect drift in the configuration?"

Ravikumar Alluboyina

Yes, so it's a very good question. In some of our discussions with QCT, this did come up, because the... I mean, just to give you a background. The Redfish API is not complete in terms of detecting the config drift or presenting the entire config. There's a good feedback from the QCT team on that, on what is required for preserving the configuration of the bare metal server. It could be BIOS or firmware. So, this is when we have designed this concept of configuration drift, which is every version of the... we constantly monitor the system and we pick up the configuration. When there is a change discovered, the change could be based on a trigger or based on a periodic monitoring.

So, every version is preserved and we could show the differences. And the reason is, you know, the – there could be different... there could be changes made by the user or the Redfish API is upgraded or BIOS is upgraded, all of these changes would come in, so it's really crucial that we have this config drift.

It is built into the platform. Short answer.

Lilian Veras

Awesome, Ravi, thank you so much. There's another question here for you. "How are different versions of BIOS and firmware, Redfish APIs handled in server bootstrap process?"

Ravikumar Alluboyina

Yes, good question. So, this is one of the practical challenges that we faced in one of the customer sites, because the hardware from the same vendor in different SKUs have different BIOS versions, and the same Redfish API doesn't work. So, there are different versions of the API that needs to be used for targeting different servers, different SKUs. So, this is where we come up with a concept of bare metal profiles, so a profile... a bare metal profile will encapsulate the hardware SKU, the operating system, and the workload, which means the target workload that is going to run on that bare metal, it could be DU or it could be hardware stack or a RAN, or a rig, or EMS.

So, bare metal profile will encapsulate different versions of Redfish API, BIOS operating system, and hardware.

Lilian Veras

That's great. Thanks, Ravi. We have another question here for you, Rachel. "What are the challenges the product solves?"

Rachel Chu

OK, good questions. So, regarding for the challenges and also what the product solves, we can discuss in different aspects from the interoperability, we understand that currently telecom operators are being open architecture from the hardware SKU layer to the VDI layer, to a VNF. So, with multiple vendors in different kinds of – different layer vendors, how we can work together to build up our ecosystem to have the pre-integrated and validated solution to a telecom operator before they are doing the implementation, deployment, to the production. And that would be the one key significant scenario for them, so that's why QCT and Robin co-worked together to build up a solution.

And the other one is when we mentioned about the management orchestration. Previously, the manual platform is more of a VNF application layer, and the VIM layer. And now with QCT's support and we work with Robin on the bare metal profile and QCT BMP tool, we combine it into the single pane of glass panel that can be offering to telecom operators to have the single point part to cover all the management and monitoring from hardware, VIM to the VNF. So, I think that will be two key factors to the products of the challenges.

Lilian Veras

Awesome. Thanks, Rachel. There is another member from the audience here asking Rachel, if you can mention some of the use cases for the product.

Rachel Chu

OK, regarding for the use case in the telecom operators, now, previously, it's more focusing on the VM, virtual machine part for the 4G EPC. And as Ravi just mentioned that Robin Cloud Native Platform could support either VM and also containers in the same platform. So, with that, it could mean that you can run QCT with system and Robin platform NFVI to offering the different kinds of application service for the 4G EPC core and the 5G core network, and to the radio access network as well. And the other use case, as we mentioned, it's the content delivery network is also now trending for the telecom operator to offer in the live streaming or different kind of streaming services.

So, these three might be the top use cases for the solutions.

Lilian Veras

That's great. Thank you, Rachel. We do have time for one last question. I believe this one is for you, Ravi. "How is disaster recovery handled by CNP platform?"

Ravikumar Alluboyina

Right, so DR is one of the biggest features of the CNP platform. It's actually a differentiator too... so, if you've... in the presentation, I did mention that Robin has a storage stack and it has network plugins, and it is a Kubernetes platform. So, we have control over storage, network, and compute.

So, what Robin provides is a DR facility, which can run in a different site and an application could be set up to have async replication to a different site. Which means you have site one and site two, the application is active on site one, and we can setup DR policies, so the application gets replicated to the other site asynchronously, it's a DR site, so sync is not possible. And with a click of a button, so we can boot up the application on the other site.

We could do that because we have control over the storage. The storage part of Robin takes care of the async replication, it is always differential, which means whatever has changed, only that portion is synced with the other site. It's not just the data, because to recreate an application on the other site, especially in 5G environments, in cloud native environments, application is not just one container or one pod. The makeup of an application could be 20 pods and 40 volumes, and four secrets, and 10 config maps. There are a lot of elements. So, we backup everything. We backup the data. We backup the configuration. We backup... when I say backup, copy to the other site. Data, configuration, and topology. The makeup of the application.

So, DR is a core construct in CNP. With just a click of a button, you can enable a DR policy and the application can be migrated over to a different site, and that site could be cloud as well. Thank you.

Lilian Veras

Fabulous. Thank you, Ravi.

Thanks to both of you, Ravi, and Rachel, for this great presentation. Thank you all members of the audience for joining us today and 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.

Thanks again, Ravi and Rachel. And join us next time. This will conclude our webcast.

Ravikumar Alluboyina

Thank you. Thank you, Lillian. Thank you, Rachel.

Rachel Chu

Thank you so much.