Intel Corporation

Accelerate NFVi Workloads For 5G Deployments

CORPORATE PARTICIPANTS

Lilian Veras
Moderator

Geetha Jayagopi
Intel Corporation – PLM & Strategic Planner for Wireline & NFV Business Division

Shashikiran Mahalank
HCL Technologies – Director, Product Management

Nick Davey
Juniper Networks – Product Manager, Cloud and SDN

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 “Accelerate NFVi Workloads for 5G Deployments”.

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.

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 a 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 Shashikiran Mahalank from HCL, Geetha Jayagopi from Intel, and Nick Davey from Juniper Networks.

Shashikiran Mahalank is Director of Product Management at HCL Technologies. He comes from a telecommunication background with work experience on 5G core, IoT connectivity gateway, signaling firewall, and EPC routing products. He's responsible for product management of HCL’s NFV Acceleration product offerings.

Geetha Jayagopi is a product line manager working with Intel's NFV Wireline BU to define and build network infrastructure solutions. In her current role, she is responsible for collaborating with global commercial service partners and telco customers to co-develop a roadmap and drive products and features from concept to launch in a fast-paced environment. She has over 10 years of experience working with software, networking, telecom, SDN, and NFV technologies.

Nick Davey is a product line manager working with Juniper’s Contrail business unit to define and build cloud native software-defined networking solutions. He has worked with mobile operators, wireline providers, and cloud builders to take software-defined networking and automation solutions from the lab into production networks. Nick is a quadruple JNCIE with over a decade of experience helping large enterprises and service providers design and deploy differentiated service offerings.

Welcome Shashi, Nick, and Geetha, and thank you for taking the time to join us today. Shashi, I will hand over to you to start off. Thank you.

Shashikiran Mahalank

Thank you, Lilian. Welcome, again, everybody for today's webinar around NFVi acceleration for 5G deployment.

So, before we get into the discussion of NFVi acceleration for 5G deployment, let's start with the requirements of 5G networks. With the evolution of telecommunication networks from 4G to 5G, 5G brings a new set of requirements to address three main use cases, namely enhanced mobile broadband where we have the bandwidth requirements of up to 110 Gbps data rate. We have ultra-low latency use cases where we have the requirement to support the low latency of one millisecond. And then we have the massive machine type communication use case where we expect 100 times the number of connected devices per unit area compared to LTE 4G network. Apart from that, we have the requirement to support five-nines. To meet all these 5G requirements, acceleration is the key for a successful deployment of 5G network, be it in terms of hardware infrastructure, software, and even the applications running in the 5G network.

As a part of HCL Industry Software division telecom product offering, so we offer different NFVi acceleration solutions for 5G network. We have a service acceleration solution where we offload the segmented routing header processing on top of hardware acceleration cards, and then we have the UPF acceleration solution where we offload the packet data processing rules from the UPF on top of SmartNICs, and then we have the NFVi acceleration solution where we offload the data path or the Fast Path from open virtual switch, and the Contrail vRouter, and in today's discussion we'll be focusing on Contrail vRouter acceleration.

Let us talk about the need to have this NFVi acceleration. With the billions of devices getting added every year, both connected devices, both mobile devices and IoT devices, there is a rapid explosion of data traffic over the cellular network. Along with that, now we have the 5G requirement in terms of bandwidth data at low latency. To meet all these requirements, telecommunication service providers are required to invest in terms of server infrastructure, network infrastructure, cooling infrastructure, and space.

Network function virtualization allows our telecommunication service providers to host the virtualized software on top of commercial off-the-shelf servers, instead of hosting-- instead of going with the proprietary closed loop hardwares. This allows operators to reduce their CapEx investment. However, to meet all these requirements, we have the challenges in terms of high CPU cores, in terms of achieving the low latency, in terms of achieving the higher throughput. NFVi acceleration, based on SmartNICs, is the cost-effective way to offload these compute-intensive networking functionalities from core servers to SmartNIC. HCL has partnered with Intel to come up with this acceleration solution, where we offload these compute-intensive tasks from the core servers to Intel PAC N3000 FPGA cards. So, these cards, these solutions allow operators to lower the TCO costs, along with achieving the higher return of investment.

Now, I hand over to Geetha to talk about the Intel SmartNIC offerings for NFVi acceleration.

Geetha Jayagopi

Thank you, Shashi. So, I'll introduce you to the broad infrastructure acceleration portfolio. As we look to the future, Intel is rolling out additional FPGA-based IPUs and integrated ASICs as well, while creating a powerful software foundation that enables cloud operators and ecosystem vendors to build powerful cloud orchestration software. The IPU is strategic for our cloud strategy because we believe our leading IPU portfolio will provide the common infrastructure foundation that allows our cloud customers and telco customers to fully leverage the general purpose compute, XPUs, and accelerators in the heterogeneous data centers of the future. Now, there is no one-size-fits-all for where IPUs can be used. FPGA-based IPUs provide faster time-to-market for evolving standards and offer ultimate flexibility to adapt to customer needs. Cloud and comms require this flexibility. Some customers only require accelerators in our FPGA SmartNIC line, while others want to offload the infrastructure entirely to the IPU, and ASIC-based IPUs are optimized to provide the solutions at scale.

So, Intel is leading the IPU market today with customers deploying at scale in the cloud and CoSP markets, and Intel offers the products to meet tailored solutions. We also support a wide range of use cases, improving the TCO and data center utilizations.

So, why do customers want IPUs or SmartNICs to accelerate their infrastructure workloads? Some of the key motivations to using an IPU are accelerating the network infrastructure by moving some of the workload onto the IPU, which reduces the total server overhead. Workloads are localized on the IPU. Performance is increased by removing some of the PCIe latency between the host and the IPU and accelerate some specific features. IPUs are user reconfigurable and highly programmable, allowing for customization of particular features and development and deployment of the software timescales.

I'd like to highlight some intelligent infrastructure acceleration. IPU provides compute resources for programmable control. It also offers the system-level security, control, and isolation. A key benefit of the IPU is the ability to isolate a secure route of trust from the host to maintain a separation of function between the provider and the tenant. It also offers some common software frameworks. Intel will be supporting IPDK software, as well as some of the projects that support migrating hosts to the IPU to smooth out the consistent API constructs, including P4 as the network programming language. Hardware and software programmable built to customer needs. Given the broad portfolio of offerings, customers can optimize to performance and design points to meet their needs, with both hardware-level programmability and software consideration.

So, these are some of the aspects of our IPU. It's an ideal time for us to take a deep dive into the Juniper Contrail and why, actually, we would require a vRouter Contrail acceleration. Now, I'll hand it off to Nick Davey.

Nick Davey

Thank you so much, Geetha. 4G and 5G workloads have placed an incredible strain, challenge, however you want to call it, on network infrastructure. We've dramatically increased the number of endpoints that we manage through either containerization or virtualization due to our need to scale out horizontally, and even with fantastic technologies in play, like the Intel IPU and accelerated vRouter, we still need to scale this problem at the edge of the data center. This isn't something we can just solve by building bigger infrastructure.

Some of these challenges that we need to address are just the sheer scale of routing protocol adjacencies needed for 5G workloads. The majority of 5G network functions, just like their 4G cousins, I suppose, they require routing to advertise subscriber reachability, to allow for things like scale-out and load balancing, and so instead of having to scale all of these numbers of routing protocol sessions in the core of your network, or on your actual aggregation devices directly, you can offer up intelligent routing right on the edge on the SDN data plane using Contrail vRouter.

Now, there's a number of different performance characteristics required by all of these workloads, and so there is really no one-size-fits-all, but the Intel IPU and the flexibility of Contrail vRouter allow you to pick and choose how workloads are connected to the data plane. We can use acceleration where we require vast scale. Like, for example, the subscriber gateways can use the accelerated data plane for increased flow and packet per second performance, allowing us to scale within our cloud infrastructure. We can also leverage flexible frameworks like Multus and Kubernetes to plumb in other types of connectivity, whether it be just kernel mode connectivity, or pass-through direct access to a network interface using SRIOV. But really, SmartNICs and IPUs are the future for 5G workloads. They allow us to have the right blend of intelligence in the server, the ability to scale horizontally in our infrastructure, while still maximizing the performance that we get out of an individual server. Remember, though, as we go down this journey, the main challenge of SmartNICs and IPUs is to get all of that intelligence from the SDN and from the network right to the port connecting the workload. So, the VM and the container have no compromise in the network services and capabilities that are exposed to it, but still have incredibly fast performance.

Contrail, in general, provides a really broad mix of options for modernizing infrastructure, and as telcos and service providers the world over have carried on down this path of virtualization and the cloudification of their infrastructure, the types of workloads that we've worked with have changed. We've moved from predominantly appliance-based deployment to predominantly virtualized telco applications, and we are now on the cusp of seeing the transformation to containerized network functions. But like I mentioned previously, none of the expectations for the workloads have changed. Even though we've decomposed our virtual machines into containers, and in a lot of cases microservices, each one of those containers or microservices still, by and large, requires the

sophisticated network connectivity, security, and services required by any 5G infrastructure, and it's the job of the SDN data plane in the cloud to implement all of those advanced features.

While Contrail has been incredibly successful in traditional telco clouds, offering up high-performance connectivity to virtualized workloads, our focus is now, and moving forward, on Kubernetes. That doesn't necessarily mean that virtual machines are going away, and I'll be the first to tell you that nothing lasts longer than legacy. We expect to see virtual machines move maybe in location, if not in format. So, we're expecting virtual machines to find a new home under Kubernetes using projects like KubeVirt to settle them into place.

The performance characteristics of these new container network functions are the exact same as their VNF brethren. So, even though we're containerizing an application, we're still expecting that particular container to achieve multiple tens of gigs of throughput, handle tens of thousands of subscribers, or maintain many routing protocol adjacencies. So, again, as we move orchestrators, the capabilities that we need to expose to the workloads don't change. In fact, Kubernetes, although introducing higher degree of dynamism into application infrastructure, really, the goal is to streamline and simplify all of the same deployment patterns that we had for the original telco clouds and 4G cloudification.

With 5G deployments, and with a focus on latency and traffic optimization, we're pushing infrastructure further out to the edge of our geographical footprints than we ever have before. In order to operate at the scale, we need efficient models. We can't just deploy a Kubernetes cluster everywhere that we're putting workloads, or else soon we'll be overwhelmed by the swarm of Kubernetes clusters that we have to manage, and so Contrail allows for increased operational scale through multi-cluster management, and our remote and edge compute architectures that allow the centralization of SDN control in, say, aggregation site or central data center, and then compute-only infrastructure to be deployed and managed from that central site wherever you need your computes. So, you can essentially stretch an island of computes on a very long wire out to the edge of your infrastructure, and again, not compromise on any of the performance or capability of the SDN data plane. By partnering with Intel and HCL, we've brought Contrail vRouter into the hardware itself, leveraging the best of both worlds from hardware acceleration and software data planes to give you an uncompromising set of features and performance.

Some of the features that we rely on in order to make our deployment successful are, of course, the advanced networking. Ultimately, what we're trying to replicate inside of the SDN are all of the advanced tools that we used to have from our traditional networks. So, what we want to approximate is every workload, every container, every VM, we want it to appear as if they were plugged into one of those ginormous, all-singing, all-dancing core routers, with the ability to apply high-scale firewall policies, load balance between a ridiculous number of endpoints, offer up thousands of routing protocol sessions, all without breaking a sweat. Contrail with the accelerated Intel and HCL data plane offers up those advanced networking capabilities to clouds and virtual machines.

As we're building out this infrastructure, the security is front of mind for everyone. Whether it's just good infrastructure hygiene, or whether you're offering a multi-tenant service and need to isolate your attendance from one another, the ability to both observe and enforce firewall policy-- or observe traffic and enforce firewall policies is key to any infrastructure. That infrastructure really needs to span between all of our types of managed infrastructure because, like I mentioned, nothing lasts longer than legacy, so we're going to see not only deployments of OpenStack continue and grow in the years to come, but we're going to see a migration of VMs into Kubernetes, and regardless of where our workloads land, we need to offer up the same set of sophisticated tools.

If we are deploying large macro infrastructure, we are going to have the need for multi-cluster connectivity, and at the heart of Contrail is a set of old and battle-hardened tools, BGP and overlay MPLS encapsulation, and by combining just BGP route advertisements to allow clusters to export all of their route context, and by layering MPLS on top of that, we get a really fantastic architecture that allows us to efficiently route multi-tenant traffic across a network core that doesn't require a great deal of configuration. And then finally, everything that we build is built upon a foundation of open source and upstream tools. So, Intel, Juniper, and HCL all know how to go with the flow. Where there's a solution in community for us to support some of the sophisticated ambitions of ours, we use those tools, whether it be upstream projects like KubeVirt, a fantastic innovation like Multus and multi-network detachments, or simply the Linux

Foundation projects like Tungsten Fabric, we extend the open source tools to build a solution that's available to all of our customers and developed with the full transparency of the community.

What this all gives us is an ability to abstract away the complexities of our networks. We can build logical topologies in our software-defined networks. We can slice and dice those topologies to fit the needs of our applications, enforcing firewall policies between individual workloads in the same subnet, and using network policies to control the flow of traffic between all of our virtual networks. We can redirect and steer traffic through service instances, whether they be containers or VMs, using service chaining. And with this set of tools, we can help operators either modernize their existing telco clouds to leverage the best of containerized solutions for 5G, or we can help operators who are still running legacy appliance-based infrastructures rebuild the same logic and services that they're used to in a much more flexible fashion.

Finally, I want to talk about just the state of where we are with offloads and IPUs. Today, we have a phenomenal architecture working with Intel and HCL, leveraging DPDK and SRIOV, to connect workloads directly to the accelerated-- or the network acceleration complex. In the future, we'll continue to work with our partners to develop full offload engines. This moves a lot of the sophistication that still is up on the operating system down to dedicated compute complexes into the IPUs and SmartNICs themselves and will offer a truly independent experience with the SDN data plane implemented-- or data plane and control plane implemented completely in the NIC. This is still a future that we're all working towards, but seeing the phenomenal progress that Intel, HCL, and Juniper have achieved over the past year, I have no doubt that we're marching steadily towards this fantastic future.

That's it for me for now. I'll hang around and watch the questions in the chat, but yes.

Shashikiran Mahalank

OK, thanks Nick. I've proceeded to HCL Tungsten Fabric vRouter Offload solution. HCL TF vRouter Offload solution allows operators to offload the Fast Path from Contrail vRouter to Intel SmartNIC. So, in this slide, we cover the architecture of HCL vRouter Offload solution, where we have the Contrail vRouter agent, which provisions the flow rules received from Contrail SDN controller to Intel SmartNIC using DPDK-based APIs via vRouter DPDK component, also known as Slow Path. And then we have the SmartNIC, which provides the interface towards the VMs or the VNFs running on the host.

Yes, talking about the architecture, when we receive the first packet from the network, our solution running on the Intel SmartNIC performs the Flow Table Lookup. So, if you find the matching flow provisioned by the vRouter, we forward the packet directly to the respective VNF or VM running on the host, also known as Fast Path, and if no flow matching rule is found, we forward the packet to the vRouter. The vRouter performs the Flow Table Lookup and it provisions the corresponding flow to the Intel SmartNIC, and any subsequent packets related to the same flow will be directly routed to the VNF running on the VMs. So, in this way, we completely offload the Fast Path from Contrail vRouter on top of Intel SmartNIC.

Some of the key features of HCL vRouter Offload solutions are we completely offload the Fast Path, which leads to 100% of core savings for Fast Path processing. We do support different overlay protocols, namely VXLAN and MPLS or UDP. We support IPv6 and VLAN as part of underlay protocols. We do support LACP and LAG protocol-- LAG for link aggregation. We support NAT. We do support mirroring where we mirror the flow to the different destinations of the VMs, which can be used for service assurance. We support multi-queues where each VF can support up to 16 queues. And we support PCIe pass-through and VDP-based interfaces towards the VM running on the host.

Talking about the value proposition of HCL VRouter solution, so we say 100% costs for Fast Path. So, here no costs are used for Fast Path processing. We support... the other advantage is the increase in throughput. We achieve up to six times increase in throughput for smaller packets here. We achieve five times lower latency than the software. And with our solution, we'll be able to achieve 100% line rate compared to software, where software is able to achieve 50% of the line rate we achieve with acceleration. Our solution supports OpenStack-based orchestration and the drivers required to support offload solutions are readily upstreamed to the TF community.

OK, to conclude, so here we summarize and call for action. So, our solution allows operators to achieve a higher TCO savings in terms of CapEx and OpEx. With TCS HCL VRouter Offload solution, operators will be able to achieve higher throughput with lower latency, achieve 100% cost savings for Fast Path processing. Since our solution is based on Intel FPGA SmartNIC, it allows the same SmartNIC to be repurposed for any future new use cases and also extends the current offload solution on the same SmartNIC.

For any demos, or any deep dive discussions, yes, we are happy to assist you, and you can reach out to us at NFVi_Acceleration@hcl.com.

Over to you, Lilian.

Lilian Veras

Thank you, excellent. Thank you for sharing such great information. We do have a few questions coming in from the audience, while you were presenting, so let's get started on our Q&A section.

Let's see the first question we have here. What are the applications or workloads that benefit most from accelerated data planes?

Nick Davey

I can either take that or, Shashi, if you want to grab it, that's cool.

Shashikiran Mahalank

Yep.

Nick Davey

Sorry.

Shashikiran Mahalank

Yes, sorry. Yes, you can-- yes, so we do support the applications or workloads starting from access network to core, or even to the end applications hosted on the network. So, our solution can be leveraged for any of these use cases.

Nick Davey

I think the clear benefit, though, is for applications that have high PPS requirements or have very strenuous latency requirements. I think any high-performance streaming analytics application, as well as on the subscriber side of things, whenever any of the subscriber management, like a VBNG or the data plane of a CUPS deployment, or even like a PGW workload or legacy GGSN, those types of workloads where we have loads of subscriber traffic coming out of a virtual machine, where we have a subscriber anchored to an endpoint, those are the type of instances that are difficult to scale horizontally. You can't seamlessly scale them easily without revealing to your subscribers that you're doing something to your network. So, those types of applications benefit greatly from acceleration because they allow us to go big, but in the cloud, so we can scale vertically.

Lilian Veras

Awesome, thank you. Thanks to you both. Another question we have here, what challenges need to be overcome in the VNF ecosystem in order for acceleration to become widely deployed?

Nick Davey

I mean, I'll take a stab at this, and then maybe hand it over to the HCL folks again. But yes, some of the challenges we've seen over the past couple years have just been the fragmentation in the ecosystem as multiple IPUs or multiple SmartNICs have used different

frameworks and had completely different implementations. So, the portability of a data plane or application between SmartNICs and IPUs has been incredibly challenging. We're seeing Intel leading the way in this space around standardization efforts for DPDK and P4 to give folks a common language that we can use to build offload implementations. And then I think the other work is just to engage with the VNF and CNF providers to ensure that there are representative interfaces within their applications so that they can make full use of offload and acceleration. Shashi, you probably have a better idea of that coming from the data plane implementation side.

Shashikiran Mahalank

Yes, so like you explained, Nick, some of the use cases, namely the data plane use cases, for example, packet core or the UPF, so these are the use cases where we really see the advantage of using acceleration there, and we already see the challenges the VNF vendors have to achieve their required throughput, and in terms of the requirement in terms of low latency there.

Nick Davey

I will say that the VNF vendors, the operators, and the network infrastructure providers, we all recognize the challenge that 5G in particular places on our infrastructure. What we tend to have in these waves of infrastructure deployment is a great sense of urgency to get infrastructure deployed and into production, and sometimes that masks over some of the work that remains to be done. I think, in this case, some of what we need to double back on and make sure we have a good solid ecosystem around is the interface support in the VNF and CNF ecosystem, making sure that vendors are anticipating some kind of SDN data plane, some kind of SmartNIC. If we all know that we have this problem, we should be planning for the solution together. I'm optimistic that industry forums like CNTT, like the Linux Foundation Networking, they're acting as communities of interest, bringing multiple parties from industry together under this common set of goals, where we can work together on some of these bigger problems. Not that they're slowing down deployment. I think 5G deployments are marching steadily onward, but there's definitely opportunities for optimization here.

Geetha, any comments around just that common interface for IP programming around DPDK or P4?

Geetha Jayagopi

I mean, in the sense, like when things are getting ready?

Nick Davey

Just about the strength of that approach, perhaps? Like why that matters? I don't think I did a great job explaining it.

Geetha Jayagopi

No, I think you did a great job. So, the common infrastructure, we're mentioning it's really important, and based on what you said on the CNTT and other communities, I think that that's really required, and Intel is actually marching towards working that and our IPUs, which the next generation, upcoming ones, will be a better fit on all that. So, yes, you did a great job, Nick. Thanks on that.

Nick Davey

I listen well every once in a while.

Lilian Veras

That's great, thanks. We do have another question here. A member from the audience asking if it's possible to use the legacy network architectures to host high-performance network services.

Nick Davey

Ooh, OK, yes. If you don't mind, I'm going to jump in and snag that one.

Yes, it is absolutely possible. With enough propulsion, pigs can fly, but sticking the landing can be tricky. The issue with using traditional network architectures to deliver these types of services is that you're asking a lot out of your edge devices. So, whether it be like the leafs in a leaf-spine network, or the aggregation or cell site router in your actual points of presence, if you have 100 containers running on a server-- well, that's a bit farfetched, but say we've got like 40 containers running on a server, and we've got two servers. Now, we've got 80 BGP sessions or 80 VFD sessions that we need to run from that router or from that switch to those servers, that problem will just scale up linearly. As we scale out our infrastructure, as we add more containers and VNFs, we're going to see further load being placed on edge devices, and it's not only just a straight-up scaling problem. You also run into convergence issues, and all kinds of stability issues when you have that much scale and churn in a network. So, it's just a risky proposition. You also have a gargantuan automation challenge. You have to programmatically, I hope, manage the network connectivity to all of the end devices using a set of traditional network automation frameworks.

If you have network automation expertise, yes, that's a tractable problem. You can solve that with a little bit of Ansible, some XML RPC, and a lot of time, but it's still a very cumbersome approach to deploying modern infrastructure. And that's to say nothing of the scale that things like multi-tenancy and network slicing are going to introduce, because as we have to start scaling out copies of parts of our infrastructure to enable a network slice for our customer, well, then that scaling problem just grows linearly, again, and it's tied to our revenue. So, it's a scaling problem that we want to have. That means we need to address it in an optimal way.

So, yes, you can definitely do this. If you have high-scale infrastructure, if you have routers and switches that can take thousands of routing protocol sessions, and believe me, they do exist, then yes, with a lot of automation, and a lot of time, you can make this work with classic architectures. But with SDN solutions matured and hardened from the first round of 4G cloudification, with operators the world over having a load of expertise in this space, having learned what works and doesn't work, with mature software solutions and hardware solutions from Intel, HCL, and Juniper, just to name a few, there is a better way to deliver high-scale infrastructure. You don't need to use just the network devices to deliver the networking. The better place for that now in modern infrastructures is inside of the servers themselves.

Lilian Veras

That's awesome. Thank you, Nick. One more interesting question we have here is, how did open source development change the way Intel, HCL, and Juniper collaborated on this project?

Shashikiran Mahalank

Yes, so HCL... OK, yes, HCL has been able to work closely with Intel and Juniper, where we're able to upstream the required DPDK interface to Tungsten Fabric community, which enables the larger TF users from the TF open source community to leverage the acceleration solution for their deployment.

Nick Davey

Working in a forum and a community that's built to enable cross-company collaboration, that's built to focus on the results, and I guess less of the-- less encumbered by process was a huge advantage. Having the ability to pull together multiple organizations, and then even take outside expertise from the community, is really the biggest benefit of open source. It goes back to that old notion of all ships rise with the tide. If we have the expertise in the community around a project that we're working on, it makes everybody smarter. As we bring in more advanced data plane capabilities into the Tungsten Fabric community, it also spurs other innovators to, I mean, be inspired by the work that we've done, to put it lightly, and to bring in their own advancements. So, innovation in the open drives other innovation. It's just the best way to collaborate between organizations in my opinion, and if you're curious about what we've got cooking over in Tungsten Fabric, I encourage you to check out Tungsten Fabric and the Linux Foundation Networking. You can see that at tungsten.io and learn a little bit more about the project, what we're working on, and what we have coming next.

Lilian Veras

Awesome. We do have time for one last question. What challenges do containerized network functions present compared to VNFs?

Shashikiran Mahalank

Yes, I think we do have both pros and cons here of using containers versus VMs. Yes, I mean, containers allow you to adopt this cloud native framework and leverage the implementation of microservice- based architecture. At the same time, it brings the challenges in terms of security, and that is where Juniper Contrail solutions clearly help to achieve those security aspects.

Nick Davey

And I think you've got just this notion that the... how do I put this? The transformation between VMs and containers should be more than just the repackaging of the workload itself. I know we're definitely on the first generation of container network functions, but we see this with every first generation transition, like an architecture transition, and we saw this in the transition from appliances to virtual network functions. First, you make something possible, then you make it optimized. So, I think right now with container network functions, in particular, they have been made possible through the hard work of the upstream community, the operator community, and vendors like Intel, HCL, and Juniper. We've taken a lot of the successful approaches from the first generation telco clouds, and we've brought them forward into containerized telco clouds, and that means things like modernizing DPDK to allow for high-performance offload inside of containers, modifying Kubernetes to support things like multiple interfaces for pods. I think all of the possible work has been done now, but there's definitely a lot of room to go back and optimize to make an actual cloud native network function, not just a VM that we containerize.

And by the way, I'm not throwing any shade at any organization that's modernized their VNFs into CNFs. Like I said, we've got to do-- we're going to make something possible before we make it optimized. So, everything's moving on the right track, but I'm really excited to see what's going to be possible as we start decomposing network functions into whatever microservice network function is going to be called. If anyone has cool names or buzzwords, now's the time to jump in.

And yes, I guess to Shashi's point as well, like just the scale of containers, as we start scaling out more and more and more, we have to replicate our policies, our visibility, all of our monitoring. That has to scale out to the size of the infrastructure that we're managing. The old gold standard used to be the number of servers managed by an admin. We need to start thinking in terms of the amount of workloads managed by an admin because containerization is going to absolutely explode that number.

Lilian Veras

Fabulous. Well, thank you all for this great presentation. I ask our live audience to please do not forget to give our team a rating for the live recording so we may continuously improve the quality of our webinars. Thank you, Shashi, Geetha, and Nick again for the great presentation, and this concludes our webcast. Bye for now.

Nick Davey

Thank you all, take care.

Shashikiran Mahalank

Thank you.

Geetha Jayagopi

Thank you.