White Paper

Empowering Telecom Operators to Deploy vRAN on Cloud and Edge Infrastructure

Sponsored by: VMware

Patrick Filkins Gary Chen
October 2021

IN THIS WHITE PAPER

This white paper discusses the reliance on virtual network functions alongside the emerging use of cloud-native network functions across communications service providers (SP). The paper also provides a deeper dive into virtual radio access network (vRAN) initiatives and how they intersect with the broader reliance on telco cloud infrastructure. In focus, we discuss the ways communications SPs can deploy network applications across a common telco cloud infrastructure layer, which could include the use of physical and virtual infrastructure and container-based applications.

EXECUTIVE SUMMARY

While most leading communications SPs have deployed 5G infrastructure to drive new revenue opportunities, recent business results indicate they continue to face common challenges including service commoditization, heightened competition, and an inability to rapidly innovate to meet shifting market needs. As such, IDC generally considers most communications SPs as nonoptimally designed to capture the full potential of new revenue streams as part of the 5G era.

As a result, borrowing from the IT and network blueprints of their hyperscale cloud providers (HCPs), many communications SPs are forging a new path toward cloud-native 5G networking to enhance service agility, improve operational metrics and, in general, transform internally and deliver even more value to both consumers and enterprises. Indeed, IDC maintains that for communications SPs to truly capture all potential value from 5G, they will need to explore new ways to architect, distribute, and deploy 5G network functions across a cloud-native environment.

"5G is not just about a new radio, or core — it's also about fundamentally changing the way telecom infrastructure is designed and deployed to act as an enabler, which means using cloud native as a foundation to accelerate innovation to customers." — Global North American tier 1 communications SP

In retrospect, most leading communications SPs have already deployed virtual network functions (VNFs) to marginally lower costs, create agility, and drive service evolution in some form or fashion. While adoption has been gradual for many communications SPs, embracing software-defined networking (SDN) and deploying network functions virtualization (NFV) form factors are now evolving to include the use of cloud-native network functions (CNFs) in the 5G era. While this is important, it is only part of the story. As noted previously, not only have HCPs migrated to cloud-native functions but the cloud infrastructure those functions run on has also evolved to support the ability to rapidly introduce new services and address specific customer needs.

As such, with the adoption of virtual network functions and cloud-native network functions increasing in telco environments, IDC expects demand for best-of-breed telco cloud infrastructure (e.g., network functions virtualization infrastructure [NFVI]) to rise in tandem.

Cloud-Native Workloads Take Center Stage in the 5G Era, Enabling Communications SPs Modern Ways to Deploy Services

While network functions virtualization has been deployed for some time in the telecom space, the use of cloud-native solutions marks a dramatic shift in how telecom networks are deployed, managed, and resourced (see Figure 1). Like other markets, cloud-native architectures can potentially unlock a range of benefits for telecom operators beyond virtualization, including:

- **Service agility.** Rapidly provision, source, and deploy infrastructure to support customized mobile services.
- **DevOps model.** Focus on a DevOps model that optimizes software development.
- **Management.** Benefit from the ability to easily manage infrastructure.
- **Lower TCO.** Potentially lower costs are enabled by standardization related to containers and orchestration.
- **Reliability.** Cloud-native applications can support built-in self-healing, creating more reliable systems.
- **Reduced vendor lock-in.** Cloud-native applications can generally run on multiple hardware platforms, reducing the potential for vendor lock-in.
- **Ability to dynamically scale for usage.** Apps can be scaled only when they are needed, improving resource utilization and cost profiles.

FIGURE 1

Carrier Mobile Edge Cloud Evolution Largely Driven by 5G and Cloud-Native Network Functions

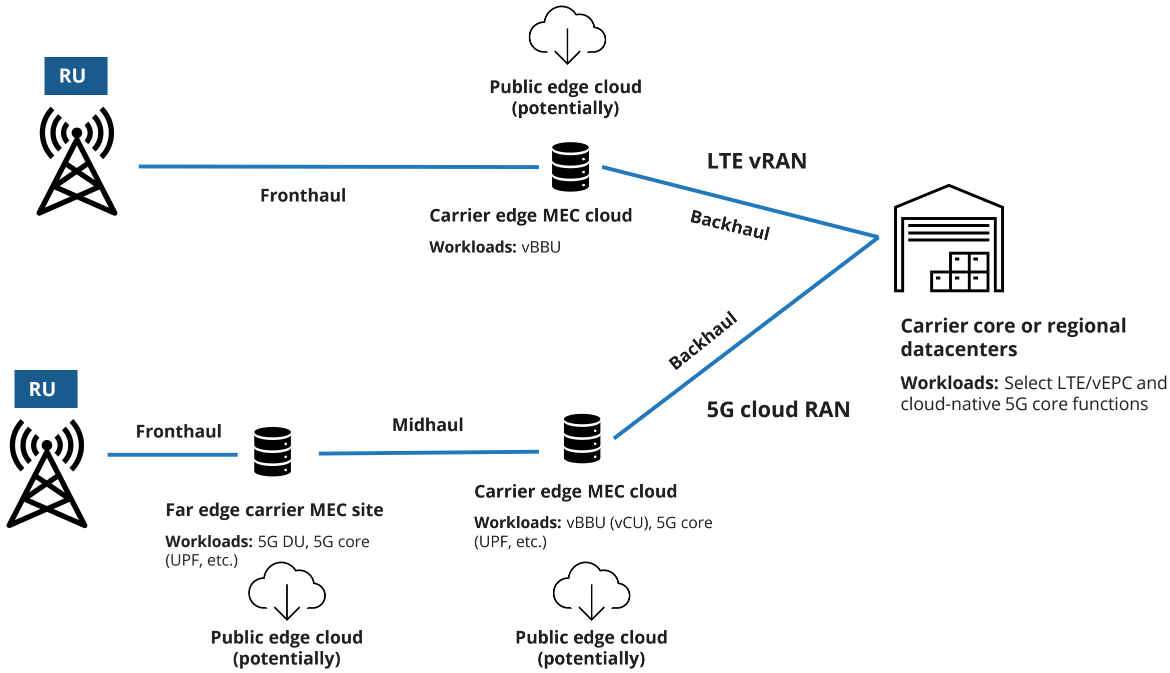

Source: IDC, 2021

Like virtualization in the 4G/LTE era, the mobile core domain is emerging as the early proving ground for container-based solutions. In focus, the SA 5G core envisions select 5G network functions deployed as containers able to dynamically be called upon and distributed across carrier core and edge sites using a standardized software approach. As such, the SA 5G core is the logical starting point to design telecom networks for cloud-native environments; however, IDC expects CNFs to see usage in other network domains over time, including the RAN.

Beyond the mobile core, communications SPs are evaluating how to migrate radio access network (RAN) workloads to a virtualized, cloud-based architecture called vRAN. In contrast to RAN workloads, which traditionally run on tightly integrated hardware platforms, vRAN represents the decoupling and tactical redistribution of RAN functions on the network. In theory, this decoupling and redistribution could reduce capex, deliver service agility, and enable better customer experience. However, vRAN also introduces significant complexity, control, and operational challenges.

While IDC believes commercial vRAN will reach mainstream adoption in the 2020s, vRAN best practices remain in development. For example, most communications SPs understand the benefits associated with virtualization; however, operationalizing a mix of CNFs and VNFs on a cloud-based infrastructure remains a challenge because of the lack of in-house experience with cloud-native, container-based solutions.

As such, vRAN considerations include how to not only deploy network functions but also deploy the underlying cloud infrastructure. Even with those challenges, IDC data shows that spending on vRAN architectures is expected to grow significantly through 2025 (see Figure 2), often in conjunction with the use of carrier multi-access edge compute (MEC) sites.

FIGURE 2

Worldwide vRAN Forecast, 2020-2025

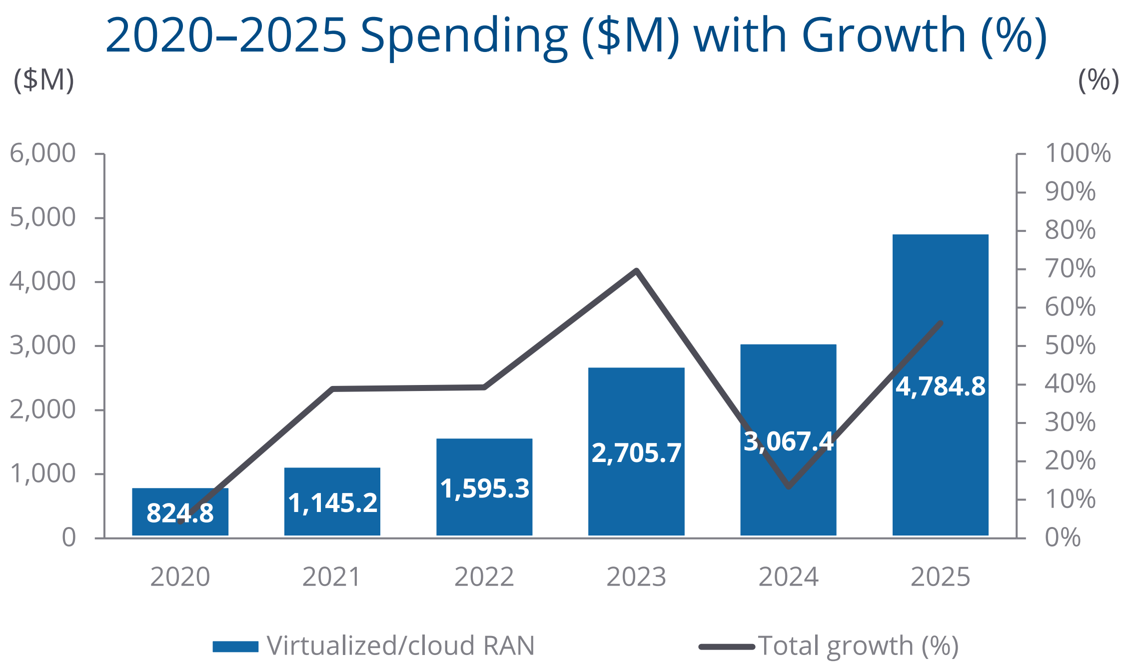

Note: vRAN started with vBBU and baseband pooling initiatives and will evolve to include splitting the baseband into central units (CUs) and distributed units (DUs) deployed at the carrier edge. This split will drive more interest in running vRAN or container-based RAN apps on either virtual machines or even bare metal infrastructure.

Source: IDC, 2021

vRAN and Open RAN: Defining the Opportunity for Containers and Virtual Machines in the RAN

While initial vRAN deployments are underway, in conjunction, communications SPs are also evaluating Open RAN solutions. With interest in both vRAN and Open RAN accelerating, it is important to understand how these approaches are defined and how they will impact and ultimately run on top of a telco cloud foundation.

> *"Telco virtualization has taken time, but we are beginning to reap the benefits of using cloud and virtualization at scale. While some network domains, such as the mobile core, lend themselves to virtualization from day one, we expect vRAN to play an integral role in our cloud-native road map over time as well. — Global European tier 1 communications SP*

Virtualized RAN

Although virtualization is well understood in the context of IT and network infrastructure, vRAN is about more than just moving from hardware to software. vRAN starts by decoupling the traditional carrier basestation, which consists of a virtual baseband processing unit (e.g., vBBU) and the radio unit (RU). The vBBU is pooled in an edge location with other vBBUs, while the radio hardware (RU) remains at the cell site. In a vRAN model, the vBBU is further disaggregated into the central unit (CU) and the distributed unit (DU). Control plane and user plane functions make up a CU, while select latency-sensitive or real-time processing functions can be deployed at the edge in the DU. Disaggregating these network functions enables them to be deployed in multiple configurations. The expected result is a RAN architecture better suited to deliver customized service levels aligned with specific applications beyond what today's mobile network can support. Further, vRAN functions, while disaggregated for optimal deployment, are generally delivered from a single vendor.

Open RAN

Open RAN incorporates all the cloud-based aspects of either LTE or 5G vRAN but relies on standardized interfaces between the CU, the DU, and the RU, which could enable communications SPs to mix and match vendor solutions across the mobile network. The decoupling of the baseband and radio is key to enabling this approach; however, standardizing both vertical and horizontal RAN interfaces enables multivendor compatibility.

RAN Intelligent Controller

Theoretically, communications SPs could mix and match CUs, DUs, and RUs to produce a best-of-breed solution. However, this approach, particularly in the context of 5G, remains largely in the proof-of-concept (POC) stage. Secondarily, Open RAN envisions the use of the RAN Intelligent Controller (RIC). The RIC allows external applications, called xApps, to control parts of the LTE and 5G network faster than current control applications. For example, the RIC, in conjunction with xApps and machine learning, could theoretically improve RAN traffic steering across different radio access technologies.

The vision aims to enable SPs to maximize the full benefits of an open virtualized RAN by providing them with the freedom to choose their RAN vendors based on their unique requirements and priorities during their RAN modernization journey. Whatever the approach, both vRAN or vRAN underpinning an Open RAN implementation will likely require a foundational telco cloud infrastructure for maximum efficacy, including the ability to deploy container-based RAN applications on virtual machines (VMs) or bare metal.

Telco Cloud Initiatives Poised to Expand, Fueled Using Both Public and Private Telco Cloud Infrastructure for 5G Workloads

As noted previously, with the broader availability and use of VNFs and CNFs as an adjunct to physical networks, IDC expects requisite demand and overall need for a telco cloud abstraction layer designed to facilitate the deployment, distribution, and real-time enablement of telco services across the telco footprint. Said differently, communications SPs are slowly evolving to mimic the approaches taken by HCPs. Regardless of the approach, IDC expects the market for telco cloud infrastructure, or NFVI, to grow rapidly through 2025 to $8 billion (see Figure 3).

FIGURE 3

Worldwide Telecom Network Functions Virtualization Software (NFVI and VNF/CNF) Revenue Snapshot

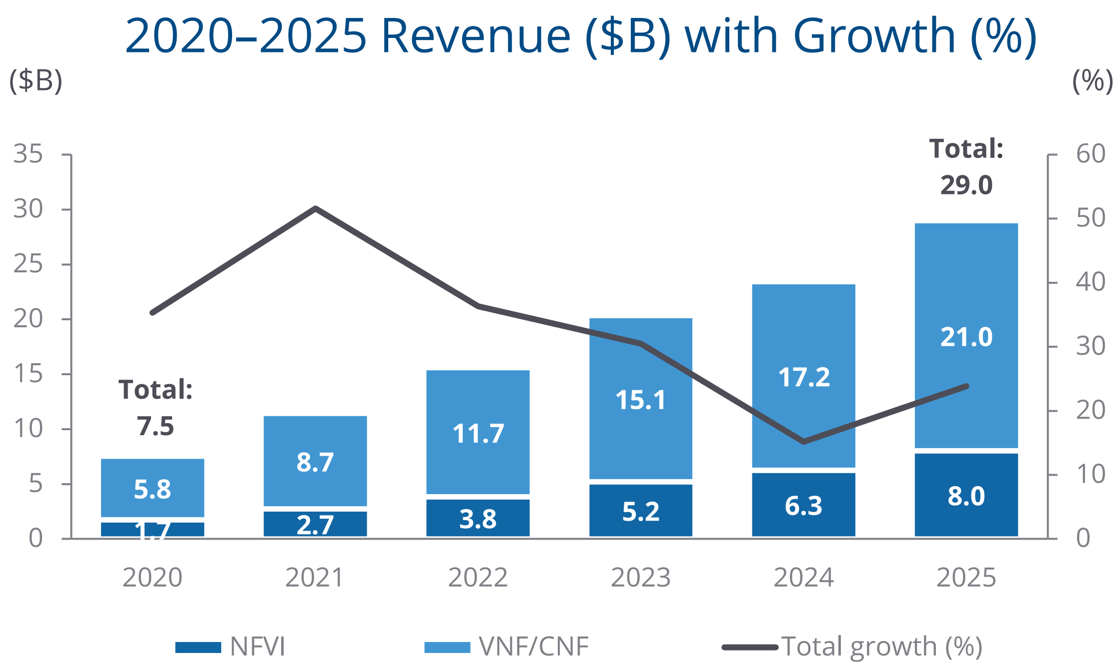

Selected Segment Growth Rate	Total Market CAGR
▲ NFVI CAGR 35.9% ▲ VNF/CNF CAGR 29.3%	30.9%

Note: Chart legend should be read from left to right.

Source: IDC, 2021

In tandem, many communications SPs are evaluating the use of public and private cloud infrastructure for network workloads. Leveraging HCPs for network workloads remains a challenging venture, but one that a few tier 1s and new market entrants have built into their 5G network rollout plans. While in the early days, the use of cloud-native network workloads paves the way for more deployment models going forward. In this instance, while the public IaaS remains the foundational hardware element, there is potential to run virtual machines, leveraging the hypervisor, as a standardized abstraction layer spanning both the telco and IaaS framework. The goal here is to standardize or create a more horizontal framework for running 5G network applications.

"Public cloud providers are an interesting way to deploy 5G network workloads and accelerate time to market for 5G services. Cloud providers already play a role for many telco IT workloads, so it's natural as the industry moves to cloud-native network functions that they could become a telecom network enabler as well. A likely outcome is we will run a mix of both dedicated and public workloads across the telco environment, which will require a rich set of resources to manage, orchestrate, and support." — North America tier 1 communications SP

VMs and Containers: Evolution of Software-Defined Compute

While communications SPs embraced hypervisor-based virtualization of the network and RAN later than other industries, virtual machines are widely deployed today, and the benefits are well known. Server virtualization allowed communications SPs to move away from proprietary boxes and use cheaper commodity hardware. Consolidation on a common and shared hardware layer also raised utilization and increased efficiency. Finally, the software-defined nature of VMs allowed for a greater level of management agility.

Containers and Kubernetes are emerging as a new compute model for vRAN as applications are shifting to a cloud-native microservices architecture. Containers are highly application centric, offering a more efficient way to package and ship applications as well as codifying operational knowledge for these applications. This enables modern cloud-style automation and orchestration patterns with Kubernetes, which can speed up software deployment, increase scalability, and reduce risk as the pace of changes and deployments increases. Figure 4 shows IDC's forecasts for VM and containers.

FIGURE 4

Virtual Machine and Container Growth, 2016-2024

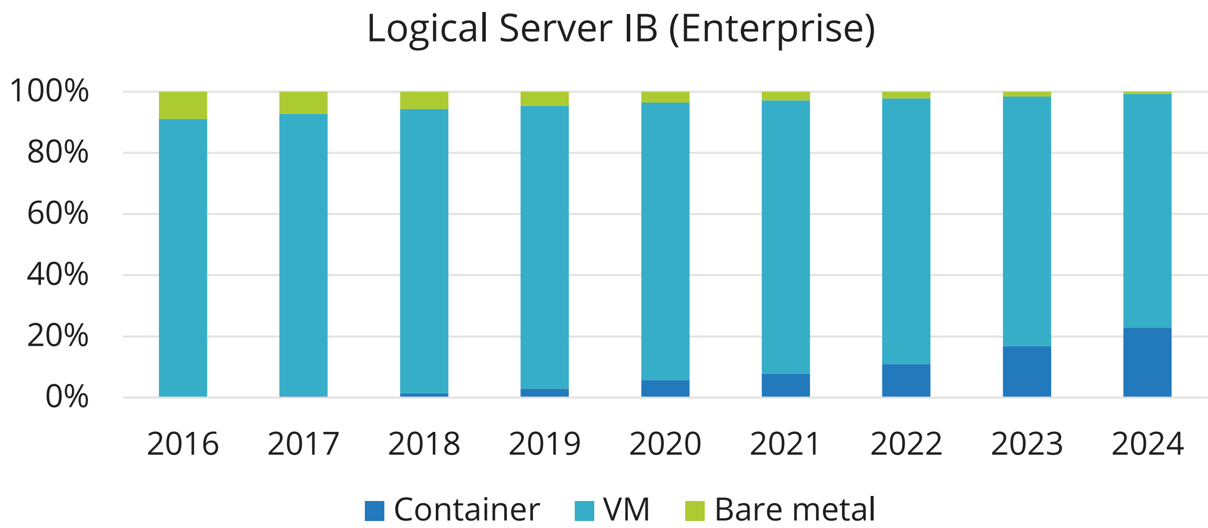

Source: IDC, 2021

Containers on Bare Metal and VMs

From a runtime and execution point of view, containers are implemented at the operating system (OS) level. Since container runtimes operate at the OS level, they can be deployed in either VMs or a bare metal server. There are many considerations when deciding to deploy containers in a VM or on bare metal.

Interest in containers on bare metal is typically driven by:

- **Cost.** Elimination of the hypervisor could save on software licensing. However, the cost savings may not be that simple to calculate. Today, virtualization is the standard system to provision and manage servers. Eliminating virtualization would require customers to acquire and implement different software to do those tasks or possibly live without certain management functions. Bare metal management tools are not as plentiful today as most of the past two decades have focused on VM management.
- **Performance.** Eliminating the hypervisor could increase performance by removing the overhead of virtualization. Virtualization overhead has decreased significantly over the years, especially as it became accelerated in silicon. CPU and memory performance overhead are today quite low and don't impact most workloads. Eliminating the duplication of multiple operating systems in VMs would recover some memory. However, operating systems are becoming more optimized for container use cases. Instead of a traditional fat OS, an OS for a container deployment only needs to do one task, which is to run a container and single application. Container-optimized OSs eliminate much of the unneeded functions, reducing image size and resource usage.

I/O is the area where virtualization overhead can make the most impact, though this has significantly improved over time. Acceleration technologies like SR-IOV can help nearly eliminate this overhead, though at the expense of some virtualization features depending on the platform. In the future, more acceleration technologies such as Smart NICs or data processing units (DPUs) may further offload and accelerate network and storage I/O.

The performance of containers on bare metal may be attractive to customers with highly performance-sensitive workloads or for the few applications that never were virtualized for one reason or another:

- **Simplification of the stack.** Elimination of the hypervisor would decrease the number of layers in the stack and simplify it. While taking out layers in the stack would certainly make it smaller, whether that would simplify it is more debatable. A typical operating stack has grown over time, instead of shrinking. Many of the layers added over time are important for abstraction and management, which are needed as the world has moved to complex constructs like fully automated clouds. Many customers feel that removing the hypervisor could make things like troubleshooting containerized applications easier, but this is often a double-edged sword. In some cases, that may be easier by removing another variable. But in other cases, the virtualization layer may provide insight and metrics to actually help troubleshooting. Today, most enterprises have more instrumentation at the virtual layer than the physical layer.

While many perceive that containers are the next generation of virtualization that replaces the hypervisor, they are actually very different technologies and work well together; virtualization is a hardware-based abstraction that partitions a physical server and containers are an operating system-level technology that sandboxes individual applications. A quintessential example of the use of virtualization and containers together is the public cloud, where nearly all containers run inside virtual machines for reasons of security, manageability, and efficiency. Virtualization is the default in the enterprise datacenter today, which means the majority of on-premises containers deployed are in VMs. Most containers today run in VMs for a variety of reasons:

- **Security and isolation.** Container boundaries are not as secure as VM boundaries, especially for multitenant scenarios. This is a major reason why public cloud containers are run in VMs. However, outside the public cloud, there are still very good reasons to have strong isolation as communications SPs may have strict requirements to segregate different groups, users, and applications, often for compliance reasons. The stronger VM boundaries can also limit the blast radius if a container or container host is breached. Containers and VMs used together provide dual layers of isolation, improving overall isolation strength. Layered security is a best practice that can provide extra protection in case of failure, software flaws, or misconfiguration and is often required in many types of deployments.
- **Manageability.** Containers and Kubernetes were developed to be better application deployment and management solutions and do not address management of the underlying infrastructure. Kubernetes does not address the underlying virtual or physical infrastructure but expects the user to present a robust infrastructure on which it can operate. Running containers on bare metal would require the user to manage and provision bare metal infrastructure in an agile, cloudlike way that is difficult to do today without virtualization. Other problems that virtualization solved, such as being able to perform nondisruptive hardware maintenance, rear their head again with bare metal.
- **Performance and utilization efficiency.** While containers on bare metal may seem attractive to customers from a performance standpoint, there are trade-offs. For most customers, only a small percentage of applications would benefit from the removal of any hypervisor overhead, which has been greatly reduced over the years. The primary factor for the majority of workloads is manageability, not performance, and virtualization more than delivers on manageability benefits at the expense of a small level of performance. Enterprises should keep in mind that

virtualization overhead continues to drive toward near zero with virtualization acceleration in CPU silicon, software optimizations, and various hardware accelerators in servers.

- **Server utilization.** Virtualization greatly increased server utilization rates, and containers can raise those rates further still. Most communications SPs will run a mix of large-scale and small-scale apps and will still need to carve up a physical server into smaller pieces. Being able to provision a Kubernetes node only by physical server increments is not very granular, and for most clusters, that would be too large of a node to effectively utilize. In addition, communications SPs may run a mix of various operating systems and multiple versions/patch levels. Containers all share the same host OS, so all containers on a bare metal server must be for the same OS version. Being able to mix and match different operating systems with virtualization provides more flexibility.

  In addition, for RAN workloads, which can be transaction time and latency sensitive, the combination of containers running on virtual machines offers more granular and robust tools to manage performance by allowing providers to configure and set policies at both the hardware and the application level.

- **Reliability, availability, and scalability (RAS).** Kubernetes provides many new and enhanced application RAS capabilities, but it enhances rather than replaces hypervisor RAS features, which operate at an infrastructure level. Kubernetes focuses on container orchestration and thus provides RAS from an application point of view, such as making sure the application is always running and scaling the number of instances. Virtualization provides infrastructure-based RAS with features such as server-based high availability, nondisruptive maintenance, and live migration. Certain containerized workloads and the Kubernetes control plane can still benefit from resilient infrastructure underneath as server failures can still cause chaos for Kubernetes. In addition, while Kubernetes can increase pod counts easily, virtualization would still be needed to increase the size of nodes or provision new nodes.
- **Storage and networking subsystems.** Storage and networking are key elements to the scalability and performance of workloads. Storage in particular is a notable issue for containers and Kubernetes, which were originally designed for stateless apps. However, many key stateful apps are still needed, and many customers want to containerize them, driving the need for data persistence, for which Kubernetes has recently been improving support. Hypervisors already have mature storage and software-defined networking subsystems and interfaces, and by running containers on a hypervisor, users can ease the integration of these elements in their container systems. Virtualization also provides an abstraction from NICs and HBAs, which would be directly exposed to the container host in bare metal. For example, new firmware revisions for adapters would be directly exposed to the container host OS and the container in a bare metal implementation, which often caused compatibility problems for bare metal OS implementations. Kubernetes drivers for storage and networking on bare metal have generally been more limited in availability and functionality, but this has improved somewhat over time.

TRENDS: CONTAINERS AND VIRTUAL MACHINES SET TO PLAY A ROLE IN THE TELCO CLOUD

VM and Container Coexistence

While containers are the future primary compute model for telco applications and network functions, VM-based applications will continue to exist for the foreseeable future, creating a coexistence paradigm. Communications SPs will likely have both VM and container-based functions during this transition and will desire the flexibility to accommodate both models as different application vendors may offer their software as one or the other. This leads to important management considerations to prevent siloing and many users will look to solutions that can provide a unified platform or management control plane across VMs and containers.

Containers and Kubernetes for the Edge

The efficient, self-contained, and standardized packaging capability of containers makes them easily distributable and portable. In addition, the distributed nature of Kubernetes make it ideal for edge computing, on which vRAN will rely. Kubernetes is evolving in several ways to accommodate different edge scenarios. One approach is shrinking the entire Kubernetes footprint to allow it to run on smaller edge hardware. Another approach, for even more resource-constrained edges, separates the Kubernetes control and data planes. The data plane runs at the edge, while the control plane can be in a centralized datacenter or the cloud.

Kubernetes Everywhere

With the industry standardizing on Kubernetes as the core container platform, there are increasingly many ways to deploy Kubernetes, from customer-managed traditional software to varying cloud models. Customers that do not want to set up and manage Kubernetes can easily get a variety of hosted Kubernetes services in the cloud or even on premises in a remotely managed model. These services are also increasingly offering higher levels of management and automation options that can offload even more management burdens from the customer. The wide availability of and standardization on Kubernetes allow customers to easily deploy containers where they are needed, in the management model desired, and increase the portability of workloads across environments.

VMWARE PROFILE

VMware is a leading provider of enterprise and telco cloud infrastructure and operations solutions. Its end-to-end portfolio provides a foundation for applications to run uninterrupted across multiple clouds, networks, or devices. IDC views VMware's horizontal approach to telco cloud enablement as a key requirement for the 5G era. Further, VMware's overall solution is delivered as a platform whereby solutions can be delivered as an integrated end-to-end offering or consumed as features where needed. The totality of the solution is broad, combining several VMware products.

Telco Cloud Platform RAN

Considering rising interest in vRAN and Open RAN, VMware now offers telco cloud platform RAN (TCP RAN), its RAN-optimized platform designed to act as the common platform for virtual and cloud-native RAN network functions. As such, TCP RAN is an extension of VMware's existing telco cloud portfolio, enabling communications SPs to further scale telco cloud infrastructure and operations. Along with acting as the common cloud platform for RAN applications (e.g., CU/DU) or other 5G custom applications, TCP RAN offers the following benefits:

- Designed to support the high-performance needs of RAN cloud-native workloads
- Integrated security
- Open RAN support for future migration
- Consistent operations from the RAN to edge to core
- Cloud-first automation
- Programmable resource provisioning for customized needs

VMware's approach supports communications SPs that want to build horizontal, cloud-native control into the RAN. What started in the communications SP core is now being brought to the RAN. Said differently, as communications SPs migrate to vRAN and Open RAN over the next decade, it is logical to take a cloud-based approach as opposed to continuing with a siloed, domain-centric strategy.

VMware Radio Intelligent Controller for Open RAN

While much of Open RAN's charter is about delivering multivendor RAN networks, openness can also deliver best-of-breed deployments and further innovation in the RAN layer. VMware's Radio Intelligent Controller (RIC) modernizes the RAN to be truly open and modular, built with best-of-breed solutions from a multivendor ecosystem while providing the RAN intelligence to ensure all the solutions work harmoniously. VMware's vision and strategy are to be vendor neutral, and therefore, it has engineered its RIC to seamlessly integrate with RAN centralized units (CUs) and distributed units as well as xApps and rApps from a large ecosystem of partners. This embodiment of openness is a foundational element of Open RAN, and one the many communications SPs are exploring as they look for ways to consume RAN solutions in a new way.

In addition, VMware RIC brings programmability to the modernized RAN through the disaggregation of the control and management functions of the RAN from the underlying RAN data plane functions. This disaggregation enables decoupling of the RAN intelligence from the data plane and repositioning the intelligence into the VMware RIC. As a result, control and management of the RAN is now performed through xApps and rApps that are hosted on the VMware RIC.

VMware's vendor-neutral approach combined with VMware RIC's programmability promotes the best-of-breed RAN architecture, creating a rich and vibrant xApp and rApp ecosystem through VMware RIC's SDKs. While the SDKs elevate the activities of app developers, they also empower communications SPs to develop their own xApps and rApps, meeting their specific business and technology requirements and priorities.

Telco Cloud Infrastructure

Considering infrastructure, VMware addresses both virtualized and container-based workloads:

- **Hybrid CaaS and IaaS infrastructure.** VMware's infrastructure layer delivers the foundational cloud infrastructure (e.g., NFVI) needed to run both VNFs and CNFs.
- **Tanzu for telco (cloud-native infrastructure).** Tanzu provides container-based orchestration for communication SPs ready to move to cloud-native, container-based workloads in the 5G core, vRAN, or other domains over time.

Telco Cloud Automation

This automation product delivers a cloud-first multilayer automation from infrastructure and CaaS to network functions and services, all designed to automate and orchestrate telco applications with integrated life-cycle management automation across a virtual or cloud-native infrastructure.

Telco Cloud Operations

Telco Cloud Operations provides real-time automated assurance across both physical and virtual network functions and also includes monitoring and performance management across multivendor architectures, common in the telco space.

Overall, VMware is taking an holistic approach to telco cloud enablement. In focus, as communications SPs look to migrate from physical to virtual, and eventually to container-based applications, it is clear that telcos will need help build the hybrid, multicloud foundation that is needed. As such, VMware's experience in cloud infrastructure management, automation, and operations positions the company as a leading partner for any communications SP, regardless of where they are in their evolution.

CONCLUSION

With leading RAN vendors just now starting to offer cloud-native RAN apps, and Open RAN still a couple of years away from broad adoption across the mobile network, today's focus remains on vRAN and traditional RAN implementations. However, what is clear is that communications SPs are likely to deploy a mix of these architectures, particularly as the 5G coverage phase shifts to densification and as small cells come to the fore.

Leveraging a common horizontal NFVI platform that can simplify RAN deployments, as well as tie in adjacent network domains (such as the mobile core), is the vision NFV and NFVI have been striving to execute on all along — namely, a common cloud-native infrastructure running VNFs and CNFs that can drive operational synergies, enhance service agility, and reduce costs. While there is growing interest in containers on bare metal, there are still many challenges to deploy that at scale. The container market today is largely dominated by containers in VMs because of security and manageability and IDC expects this trend to continue.

As 5G just begins to hit its stride, it's clear that virtualized and cloud-native containerized solutions will provide the foundation for the next wave of 5G and edge-based services designed to improve existing services and help provide the agility needed for new enterprise and industrial connectivity services.

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

140 Kendrick Street
Building B
Needham, MA 02494
USA
508.872.8200
Twitter: @IDC
blogs.idc.com
www.idc.com