This IDC Technology Spotlight discusses the overall trends, potential benefits, and challenges communications service providers will face as they increasingly adopt containers and virtual machines to optimize telco cloud applications such as vRAN.

Optimizing Telco Workloads on Cloud Infrastructure: vRAN in Focus

June 2021

Written by: Patrick Filkins, Senior Research Analyst, IoT and Mobile Network Infrastructure, and Gary Chen, Research Director, Software-Defined Compute

Introduction

Communications service providers (SPs) have largely embraced network virtualization as a means to lower costs, create agility, and drive service evolution. While adoption has been gradual for many communications SPs embracing software-defined network (SDN) and deploying network functions virtualization (NFV), form factors are evolving to include the use of cloud-native network functions (CNFs) in the 5G era. As such, with adoption of virtual network functions (VNFs) and cloud-native network functions (CNFs) increasing, communications SPs will also need to prioritize telco cloud infrastructure (e.g., network functions virtualization infrastructure [NFVI]) to reap the most benefits from a cloud-centric approach.

For example, the 5G mobile core will deploy as a set of cloud-native, container-based microservices on day one. 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 not only deploying network functions but also how to deploy the underlying cloud infrastructure. IDC data shows that spending on vRAN architectures is expected to grow significantly through 2025 (see Figure 1).

AT A GLANCE

KEY TAKEAWAYS

- » Communications SPs are embracing telco virtualization as part of the 5G era, which includes moving into container-based applications.
- » A vRAN strategy is emerging as a viable approach for communications SPs to lower costs and improve service agility.
- » vRAN considerations increasingly include whether to deploy RAN functions on VMs, containers in VMs, or containers on bare metal.

FIGURE 1: ***Worldwide vRAN Spending, 2020–2025***

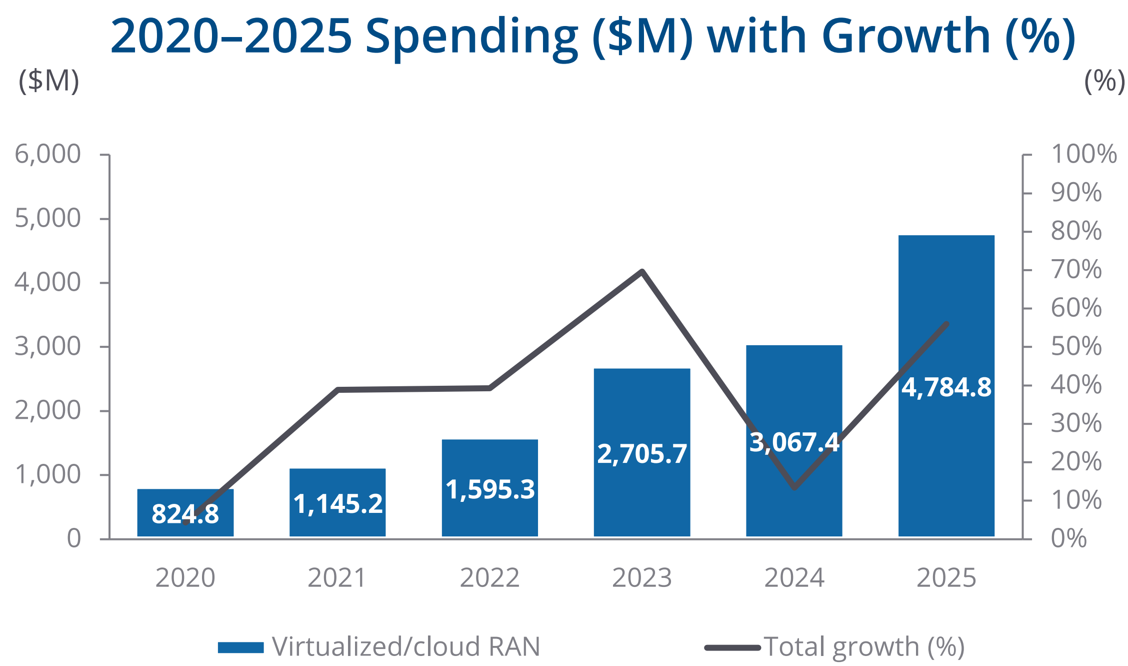

Note: vRAN started with vBBU and baseband pooling initiatives and will evolve to include splitting the baseband into central units (CUs) and distributed units (DUs) deployed at the carrier edge. This split will drive more interest in running vRAN or container-based RAN apps on either virtual machines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

Source: IDC's Worldwide Telecom Network Functions Virtualization Software (VNF and NFVI) Forecast, 2021–2025 (IDC #US47662821)

Further, cloud infrastructure is evolving to include reliance on orchestrated, container-based systems that are ideal for meeting the needs of cloud-native microservices. In some cases, container-based systems can greatly improve scalability and manageability and speed up software development and deployment. Yet container-based systems also require a very different approach and skills for infrastructure management, application architecture, and software deployment. Success will be dependent on not just technology but also people and process transformation. To get the maximum benefit from containers, customers must pivot to a cloud operating model, DevOps organizational structures, microservices application architectures, and agile development methodologies.

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.

Virtualized RAN. While 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 base station, 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. 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.

Whatever the approach, both solutions will rely heavily on a foundational telco cloud infrastructure for maximum efficacy, including the ability to deploy container-based RAN applications on virtual machines (VMs) or bare metal.

Containers and Kubernetes

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 increase.

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. 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 (OS)-level technology that sandboxes individual applications. Because container runtimes operate at the OS level, they can be deployed either in VMs or on a bare metal server. 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. Containers and VMs used together provide dual layers of isolation, improving overall isolation strength.
- **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.
- **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. 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 and various hardware accelerators in servers.

  Virtualization greatly increased server utilization rates, and containers can raise those rates further. Most enterprises will run a mix of large-scale and small-scale apps and will still need to carve up a physical server into smaller pieces. Additionally, 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 level 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 will benefit from resilient infrastructure underneath because server failures can still cause chaos for Kubernetes. Additionally, while Kubernetes can increase pod counts easily, virtualization is needed to increase the size of nodes or provision new nodes.

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. Additionally, the distributed nature of Kubernetes makes 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 Kubernetes to run on smaller edge hardware. Another approach for even more resource-constrained edges separates the Kubernetes control plane and data plane. 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 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.

Considering VMware's Telco Cloud Portfolio, Including Telco Cloud Platform RAN, as the Foundation for vRAN and Open RAN

VMware is a leading provider of enterprise and telco cloud infrastructure and operations solutions. It's 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 a number of VMware products.

Telco Cloud Platform RAN

In light of 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

As a whole, 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.

Telco Cloud Infrastructure

Considering infrastructure, VMware offers two options to address 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 communications SPs ready to move to cloud-native, container-based workloads in the 5G core, vRAN, or other domains over time.

Telco Cloud Automation

Telco Cloud Automation delivers cloud-first multilayer automation designed to automate and orchestrate telco applications with integrated life-cycle management automation across a virtual or cloud-native infrastructure.

Telco Cloud Operations

This solution 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 a 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 to 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 the provider is in its evolution.

VMware Commercial Projects Showcase Scale and Cloud-Native Prowess

Vodafone selected VMware to deliver its Telco Cloud portfolio across Vodafone's expansive operations in Europe, the Middle East, Africa, and Australia, providing the cloud infrastructure foundation for nearly 50% of Vodafone's global traffic. Vodafone's project represents one of the most cohesive and far-reaching global telco cloud efforts to date.

Key statistics are as follows:

- VMware's Telco Cloud supports nearly 1,500 VNFs across Vodafone's global footprint, including voice, data, and service platforms.
- TCO of core network functions decreased by nearly one-half.
- VNF design, build, test, and deploy time was reduced significantly.

Challenges

- **Market messaging.** VMware's greatest challenge in integrating containers and VMs will be one of perception and communicating deeply technical and complex issues. VMware needs to simplify a complex topic and effectively convey to the market the differences between containers and VMs and why a hypervisor is still needed.
- **Lack of knowledge around container-based architectures.** Communications SPs were slow to transition to virtualization, and containers are another complex and transformative technology. The impacts are arguably broader than those of virtualization, which was primarily infrastructure focused. Container adoption is usually tied to application rearchitecture and transformation of the entire software development and deployment process that requires organizational change. Additionally, container, Kubernetes, and cloud-native skills are in tight supply today.
- **Refactoring and replatforming of applications.** While some organizations may be able to lift and shift existing applications to a container, many applications require refactoring. Thus, many applications may not realize the full benefits of containerization until they are rebuilt from the ground up as cloud native. Any kind of replatforming, refactoring, or rebuilding of complex applications can take a lot of time. The entire IT industry is experiencing this situation today as an overall shift to cloud and containers is taking place. Communications SPs' RAN applications typically have been slow to transform, so the limiting factor to containerizing RAN may ultimately depend on how fast RAN application vendors can adapt their software.
- **Lack of a cohesive plan among communications SPs to build a horizontal foundation of cloud-based infrastructure across multiple domains.** Communications SPs have generally lacked an end-to-end approach to telco cloud, instead opting to develop virtualization and cloudification in select telco domains (e.g., mobile core). This approach runs counter to the overarching reason for and potential benefits gained from a telco cloud, which is a horizontal, common cloud-native platform on top of which VNFs/CNFs can be deployed.
- **Lack of understanding regarding how to operationalize NFV and SDN for maximum benefit using management and orchestration (MANO) platforms.** Similar to the overarching telco cloud challenge, domain-specific MANO platforms have emerged, but they solve only specific domain challenges. To truly achieve zero-touch automation across multiple domains, communications SPs will need to look for a way to bridge those domains.

- **Ability to ensure security as VNFs/CNFs traverse multiple telco clouds or even telco edge clouds.** With any application, security remains top of mind. While communications SP networks have built-in security measures today, deploying application across a telco cloud is still a new approach, requiring a different method to single-pane management and simplified monitoring built into the telco core and edge infrastructure assets, perhaps even including in the application itself.

Conclusion

Containers are sweeping the IT industry, as customers pursue more cloudlike operations. Subsequently, in the telecom market, communications SPs are evaluating cloud-native, container-based solutions as a means to create their own telco clouds. As such, the RAN is emerging as a logical target to realizing the many benefits they bring through the emergence of vRAN and Open RAN projects. While virtualizing, containerizing, and cloud enabling the RAN are still relatively new initiatives, they are key to meeting the increasing demands of mobile, which will require increased scalability, efficiency, and manageability. VMware's leadership in virtualization and subsequent integration of the Tanzu container technology into vSphere, and VMware's broader Telco Cloud offering, provide communications SPs an end-to-end platform to support vRAN migration over the next decade.

In the telecom market, communications SPs are evaluating cloud-native, container-based solutions as a means to create their own telco clouds.

About the Analysts

Patrick Filkins, *Senior Research Analyst, IoT and Mobile Network Infrastructure*

Patrick Filkins covers IoT and Mobile Network Infrastructure. He is responsible for market and technology trends, forecasts, and competitive analysis related to the IoT network edge, gateways, and protocol strategies. Additionally, Patrick's research is focused on mobile infrastructure, fixed-mobile convergence, and the emerging 5G ecosystem worldwide.

Gary Chen, *Research Director, Software-Defined Compute*

Gary Chen is IDC's research director of Software-Defined Compute. His research focuses on server virtualization, container infrastructure and management, and cloud system software (system software used to build IaaS clouds such as OpenStack).

MESSAGE FROM THE SPONSOR

VMware powers complex digital infrastructure for more than 200 leading global communications service providers. By applying agile software and cloud methodologies to IT data centers and communications networks – including fixed, wireless wireline and cable. VMware helps CSPs increase business agility and improve operational efficiency. For more information, visit telco.vmware.com.

IDC Custom Solutions

The content in this paper was adapted from existing IDC research published on www.idc.com.

IDC Research, Inc.
140 Kendrick Street
Building B
Needham, MA 02494, USA
T 508.872.8200
F 508.935.4015
Twitter @IDC
idc-insights-community.com
www.idc.com