Accelerate the Transformation to Open RAN

Acceleration technologies to optimize RAN performance

Get Started

vmware®

Summary
Introduction
The importance of acceleration technologies
VMware enables choice and flexibility
Conclusion

Summary

5G services present a significant growth opportunity for communications service providers (CSPs), but traditional radio access network (RAN) infrastructure makes it difficult to fully realize the 5G promise. Virtualizing RAN functions, or vRAN, can mitigate some challenges. Open RAN can mitigate still more challenges. With virtualization obviating the need for purpose-built hardware, acceleration options have arisen to mitigate performance losses. VMware® Telco Cloud Platform RAN, a cloud-native and RAN-optimized solution, provides CSPs a clear modernization path from legacy RAN to vRAN to open RAN. VMware Telco Cloud Platform RAN also supports various acceleration options constituting an ideal RAN solution that increases choice and flexibility for CSPs to maximize performance.

Summary
Introduction
The importance of acceleration technologies
VMware enables choice and flexibility
Conclusion

Introduction

To get the maximum benefit of 5G, CSPs must modernize their RAN. This is because traditional RAN is too inflexible in terms of both cost and time to market to serve as a viable delivery method of new 5G services. Achieving the needed agility and spurring innovation requires virtualization that eliminates reliance on proprietary hardware and enables the use of emerging industry standards for RAN interfaces to disaggregate RAN functions. In short, CSPs need an evolutionary path from traditional RAN to vRAN and then to open RAN.

To help CSPs transform their RAN, VMware has developed the VMware Telco Cloud Platform RAN. The platform enables CSPs to monetize their RAN investments by providing flexibility, scalability and increased performance. CSPs can leverage commercial off-the-shelf (COTS) hardware and choose vRAN functions from their preferred vendors.

By design, VMware Telco Cloud Platform RAN enables low latency and the potential to convert the RAN into a 5G multiservices hub. The platform serves as a common horizontal infrastructure that supports centralized automation, simplifies RAN operations and enables end-to-end visibility, root cause analysis and automatic remediation.

However, disaggregation and the elimination of ‘black box’ solutions open potential architectural requirements for accelerating RAN traffic. Primary approaches to acceleration are inline and lookaside, along with newer built-in/integrated processing. Each accelerator brings different value to CSPs in terms of performance, energy efficiency, cost and ease of deployment. VMware works closely with solution partners to help CSPs achieve their goals. The VMware Telco Cloud Platform RAN therefore enables flexibility and choice by supporting acceleration approaches.

Summary
Introduction
The importance of acceleration technologies
VMware enables choice and flexibility
Conclusion

The importance of acceleration technologies

To tap the massive potential for profitability and growth that 5G represents, CSPs must innovate and deploy new services. But they must do so cost-effectively, and the current RAN ecosystem does not make this easy. As CSPs strive to accelerate the deployment of their next-generation 5G networks, vRAN technologies play a pivotal role. Layer 1, also known as the physical layer, is the bedrock of any communication network. The responsibilities of Layer 1 range from data modulation and coding to intricate processes such as channel estimation. The processing intensity required for these operations can strain the system, causing latency and lowering data transfer rates.

The real-time functions of Layer 1, within the virtualized distributed unit, are critical given the 5G requirement for high-bandwidth low-latency networks, but commodity hardware does not typically excel in this area. vRAN/open RAN using the off-the-shelf hardware are therefore not viable options to modernize the traditional RAN with the purpose-built systems. However, methods for addressing this issue by accelerating Layer 1 functions such as FED, LDPC and others have arisen, and VMware Telco Cloud Platform RAN supports those acceleration solutions.

Accelerators are additional chips that offload Layer 1 processing from the general-purpose CPU. In vRAN solutions, inline and lookaside, including built-in/integrated processing, are a few of the primary methods of accelerating Layer 1 functions. Each make use of accelerators designed specifically for use in Layer 1 processing, and allow CSPs to take advantage of the higher connectivity speeds 5G networks can support. Inline acceleration involves a one-way flow where data passes through the accelerator and terminates at the destination without returning to the host. By contrast, lookaside acceleration involves tasks sent from the CPU to the accelerator for compute-intensive functions and then returned.

Summary
Introduction
The importance of acceleration technologies
VMware enables choice and flexibility
Conclusion

Major RAN Vendors and Their L1 Strategies[1]

Vendor	Known Accelerators	Purpose-built L1 Hardware	Virtual L1 Hardware	Purpose-built L1 Software	Virtual L1 Software
Ericsson	Lookaside	Ericsson	Intel	Ericsson	Ericsson
Fujitsu	Inline	Unknown	Marvell, NVIDIA	Unknown (likely Fujitsu)	Aerial on NVIDIA; probably Fujitsu on Marvell
Mavenir	Inline, lookaside	N/A	Intel, Qualcomm	N/A	FlexRAN/Mavenir for Intel; Qualcomm
NEC	Inline	Unknown	Qualcomm	Unknown (likely NEC)	Qualcomm
Nokia	Inline	Marvell	Marvell	Nokia	Nokia
Rakuten	Inline, lookaside	N/A	Intel, Qualcomm	N/A	FlexRAN/Rakuten for Intel; Qualcomm
Samsung	Inline, lookaside	Possibly Marvell	Intel, possibly Marvell	Samsung	Samsung

(Source: companies, LightReading)

(Note: Table is not meant to be taken as a strict guide to commercial activity but highlights partners and technologies describe in interviews and press releases.)

[1] LightReading, Chip choices kickstart open RAN war between lookaside and inline, August 2023.

Summary
Introduction
The importance of acceleration technologies
VMware enables choice and flexibility
Conclusion

Inline acceleration

Inline processing is an acceleration methodology in which all operations are performed within the main data path. In this approach, specified or all data packets are directed to inline accelerators. The accelerator performs the specified computations inline in real time, removing that task from the main processor. In vRAN, inline processing can expedite Layer 1 tasks such as channel estimation LDPC, fast Fourier transform (FFT) and demodulation.

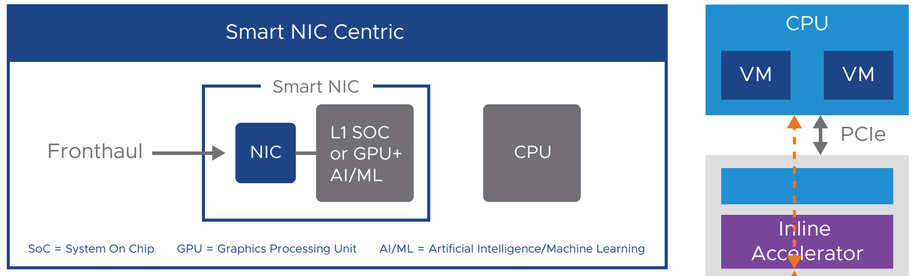

The term ‘inline,’ therefore, refers to the fact that these accelerators are within the main data path. While other accelerators may operate independently or on the side, inline accelerators are embedded directly within the data flow. They potentially interact with every data packet passing through the network, promoting swift processing and immediate data availability.

When data packets are routed through the network, inline accelerators immediately process the data, eliminating any potential lag or latency. By directly handling every packet, inline accelerators ensure quick responses and streamline operations, contributing to the overall network performance.

A sample use case for inline acceleration is FFTs—an integral process in Layer 1. FFT transforms time domain data into frequency domain data, a fundamental operation in signal processing within RAN solutions. The real-time requirements of this process mean that delays could adversely impact network performance.

An inline accelerator, positioned directly in the data path, can process FFT operations at high speeds. The unit can thereby provide instantaneous results and maintain the system’s real-time response. And it removes the FFT compute burden from the primary processor entirely.

By enabling high-speed processing and real-time responses in Layer 1 of vRAN, inline accelerators facilitate smoother migration from legacy 4G to next-generation 5G infrastructure.

Summary
Introduction
The importance of acceleration technologies
VMware enables choice and flexibility
Conclusion

Lookaside acceleration

Lookaside processing is a methodology in which specified operations are performed independently from the main data path. In this approach, a CPU redirects specified computation tasks to a lookaside accelerator. The accelerator executes the computation in parallel, offloading from the main processor. In vRAN, lookaside processing can expedite Layer 1 functions such as forward error correction (FEC) coding and channel estimation.

Built-in/integrated processing are additional lookaside methodologies wherein separate accelerator cores are integrated within the CPU silicon. This allows CSPs to benefit from 5G speed and is key to successful migration to 5G network infrastructure. By understanding the functions and benefits of these technologies, RAN architects and professionals can make informed decisions while evaluating vRAN platforms for their specific network migration requirements.

The lookaside processing architecture is therefore defined by its lateral relation to the primary data path. Accelerators, specialized processing units, are not in line with the data flow in lookaside architecture. Instead, accelerators are positioned to function "on the side." Working in parallel with and independently from the main data path, the lookaside accelerators help prevent the primary CPU from being overloaded.

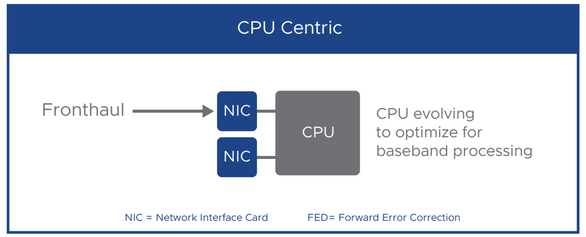

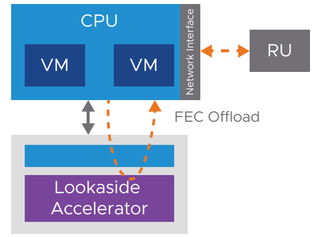

Lookaside processing is designed to address complex operations that require significant computational power. When a CPU encounters a computationally intense task, it offloads the task to the lookaside accelerator. The lookaside accelerator can then concurrently process the complex tasks while the primary CPU continues its standard operations, thereby optimizing overall system performance.

Summary
Introduction
The importance of acceleration technologies
VMware enables choice and flexibility
Conclusion

A sample use case for lookaside processing is beamforming, the signal processing technique used to control the directionality of the signal transmission or reception. It is integral to massive MIMO (multiple input, multiple output) technologies found in 5G networks. The challenge is that calculating and adjusting beam patterns in real time can significantly tax primary processors.

Leveraging lookaside accelerators to perform beamforming calculations can enhance system performance. The lookaside accelerator can independently process the complex algorithms required for beam pattern calculation and adjustment. This removes that processing load from the primary CPU, freeing it to manage other routine tasks without interruption. Separating tasks in this way helps to enable seamless operation and high-performance system function.

By shouldering the computational load off the primary CPU and optimizing network performance, lookaside processing enables a smoother transition from 4G to 5G network infrastructure. As the telecommunications industry continues to evolve, lookaside processing is a significant enabler for efficient network progression.

Summary
Introduction
The importance of acceleration technologies
VMware enables choice and flexibility
Conclusion

VMware provides choice and flexibility

VMware Telco Cloud Platform RAN helps CSPs evolve beyond traditional RAN to a virtualized and open RAN. The platform is designed to provide customers with flexible vendor and network function options. To achieve this, the platform features a horizontal design with a uniform software layer and orchestration across core, cloud and edge settings.

Enabling CSPs to modernize their RAN

Legacy RAN solutions combine purpose-built hardware with vendor-embedded software, so the key to effective modernization is to virtualize RAN functions. Doing so makes it possible, for example, for CSPs to introduce cloud-smart automation that simplifies the lifecycle management of the RAN. CSPs can easily “scale up” and “scale down” vRAN resources to accommodate variations in traffic volume. And as software becomes disaggregated from hardware, an evolution facilitated by the O-RAN alliance, there are opportunities for open standards. The open standards are referred to as open RAN.

However, the shift to virtualized and open RAN calls attention to the need for acceleration technology. Specifically, using off-the-shelf chipsets rather than purpose-built silicon can lead to reduced performance. As such, virtualization can require acceleration technologies, such as inline, lookaside, and built-in/integrated acceleration to match the performance of custom silicon. The choice between different acceleration technologies often depends on factors such as latency sensitivity, the desire for fewer data hops and specific function requirements in the context of open RAN. VMware Telco Cloud Platform RAN supports multiple approaches, maintaining a vendor-neutral and technology-neutral stance.

The abstraction layer provided by virtualization separates the hardware from logical units exposed to applications. The key role VMware Telco Cloud Platform RAN plays is to expose hardware capabilities such as hardware acceleration to the vRAN functions sitting atop the hypervisor layer, thereby enabling alignment with industry and operator requirements. For example, the platform enables the ability to manage multiple accelerator cards in one server (single socket vs. dual socket) and effectively load balance/distribute traffic between them. This allows RAN vendors to tap into RAN-centric infrastructure/server capabilities even though there is a hypervisor layer between the underpinning infrastructure and the applications.

Summary
Introduction
The importance of acceleration technologies
VMware enables choice and flexibility
Conclusion

Delivering RAN-optimized performance

Because 5G services require low latency and high-performing vRAN functions, VMware Telco Cloud Platform RAN has undergone extensive integration work with key RAN vendors. This integration serves to maximize performance and improve resource utilization, meeting and exceeding stringent requirements inherent to RAN systems. Moreover, the flexible and intelligent orchestration and automation allow for deployment of vRAN functions at the best locations to serve their functional purposes, resulting in increased performance without linearly scaling the entire RAN.

Providing a horizontal platform, automation, and service assurance

A horizontal platform

VMware Telco Cloud Platform RAN enables CSPs to virtualize multi-vendor RAN functions on a horizontal platform specifically optimized for RAN systems that has been hardened through strenuous testing and integration work with key RAN vendors. In this way, CSPs can manage the full lifecycle of multi-vendor vRAN functions and services deployed across the RAN with consistent, automated operations. In addition, CSPs could gain significant advantages from deploying vRAN and edge workloads on the same platform. By running vRAN functions alongside edge workloads, such as virtualized cell site routers or RAN Intelligent Controllers (RICs) as an example, CSPs can reduce hardware footprint, leading to the total cost of ownership (TCO) reduction, including energy saving and lower carbon emission.

Cloud-smart automation

5G RANs tend to be widely distributed, fast moving and complex. They defy manual management processes and make deep automation capabilities vital. Powered by VMware Telco Cloud Automation, VMware Telco Cloud Platform RAN supports consistent, cloud-smart automation and management across dispersed RAN sites. VMware Telco Cloud Automation provides lifecycle management and orchestration, from server provisioning and vRAN functions instantiation to service automation, such as automated provisioning of accelerator cards.

Service assurance

To meet service assurance needs, VMware Telco Cloud Service Assurance empowers CSPs to simplify the operations of their RAN and services. VMware Telco Cloud Service Assurance provides capabilities—such as the ability to monitor the health of accelerator card(s) and gather energy metrics—that enable deeper visibility, root cause analysis and remediation in the RAN.

Summary
Introduction
The importance of acceleration technologies
VMware enables choice and flexibility
Conclusion

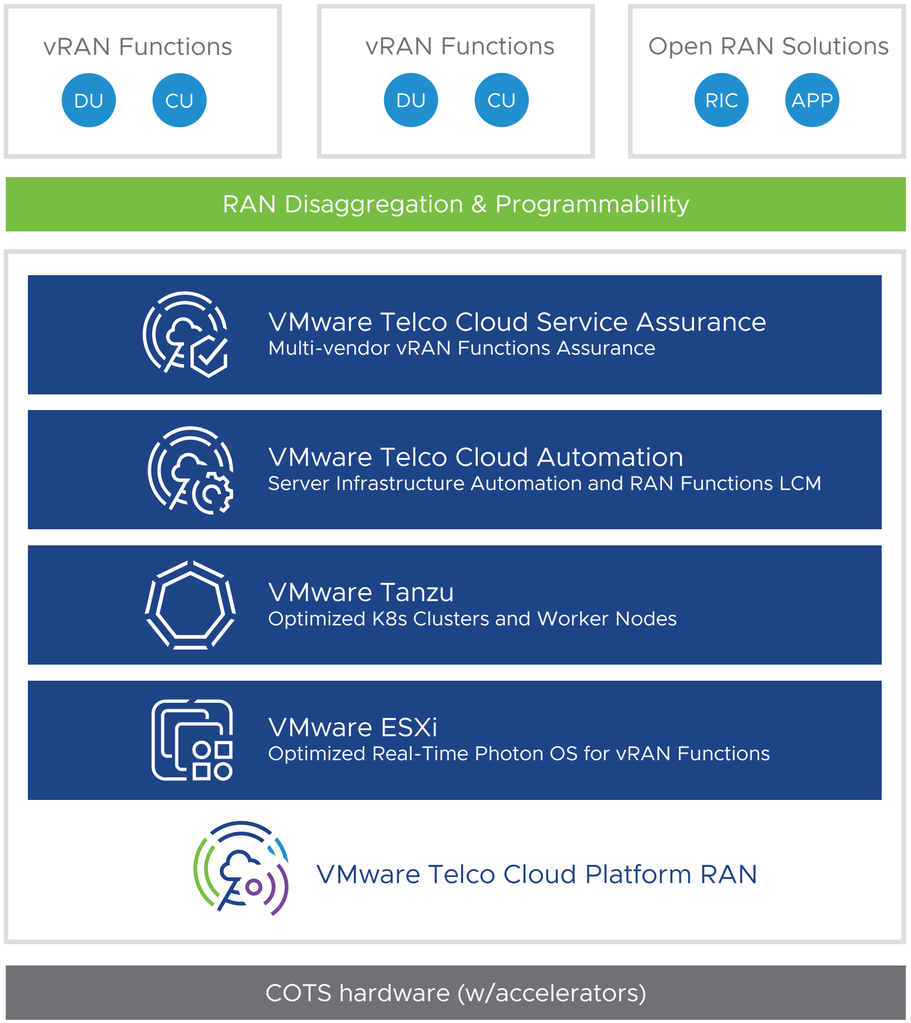

With the right combination of VMware products and virtualization technologies, CSPs can reduce their costs while increasing their operational efficiency.

Enabling design and architectural flexibility and choice

VMware Telco Cloud Platform RAN provides a wide range of choices and flexibility. The platform provides consistent and automated operations for multi-vendor RAN environments by supporting different hardware and compute options through the tight integration work with various RAN vendors. VMware Telco Cloud Platform RAN provides further choice by supporting both virtual network functions and cloud-native network functions that have been certified with multiple vendors. These capabilities make VMware Telco Cloud Platform RAN the optimal choice for RAN modernization.

Summary
Introduction
The importance of acceleration technologies
VMware enables choice and flexibility
Conclusion

Conclusion

The potential for growth driven by 5G technology is massive. But at times such growth is hindered by existing RAN infrastructure. Open RAN can address infrastructure issues through virtualization. In particular, the rise of acceleration technologies can compensate for potential performance losses when moving away from purpose-built hardware.

VMware Telco Cloud Platform RAN enables currently available acceleration technologies, including scenarios in which both inline and lookaside processing are in use, along with newer built-in/integrated acceleration. This gives CSPs the ability to choose which best fits their needs, allows RAN vendors to tap into RAN-centric infrastructure/server capabilities and ensures a seamless transition from legacy 4G to next-generation 5G networks. Additionally, VMware focuses on specific nuances in the operation of different technologies and provides a comprehensive view of the entire network through its unique orchestration and observability platform. This support makes VMware Telco Cloud Platform RAN the optimal choice for RAN modernization and the journey to profitable RAN monetization with 5G.

Get Started Today

Accelerate the transformation to Open RAN

LEARN MORE

Join us online:

vmware®