

WHITEPAPER

Building an Open vRAN Ecosystem

Accelerating deployments with a replicable
reference architecture

September 14, 2020

NEC

DELLTechnologies

 **Red Hat**

ALTIOSTAR

 **Netcracker**
An NEC Company

Contents

Introduction

vRAN Reference Architecture at a Glance

vRAN Deployment Options

Hardware Configuration for vDU/vCU

vRAN On-Boarding and Service Orchestration

Conclusion

Introduction

The telecom industry is facing challenging times to efficiently accommodate the continuous increase in mobile usage and connectivity options. While we pave the way to 5G with infinite possibilities that will enrich our lives, traffic usage and patterns become greatly diversified, further complicated with a variety of connected devices, and traffic demand continues to surge. To address such demands, Communication Service Providers (CSPs) are seeking ways to flexibly design networks, dynamically scale network capacity, expand service coverage and improve the overall user experience whilst reducing the Total Cost of Ownership (TCO) and time to market.

In the mobile packet core network, Network Function Virtualization (NFV) has already proven itself as a viable solution to increase network flexibility and simplify deployment and operations to meet the evolving market requirements. How can CSPs leverage NFV technology to bring the same benefits to the Radio Access Network (RAN)?

The answer is with Open virtualized Radio Access Network (Open vRAN). By horizontally disaggregating the network functions (RU/DU/CU) with open interfaces, and vertically decoupling hardware and software with virtualization technologies, Open vRAN will free CSPs from the haunts of vendor lock-in, opening doors to a diverse, innovative selection of products. Additionally, these dynamics will enable the implementation of a more scalable, agile and intelligent network that can significantly increase the flexibility and efficiency of CSPs' deployments and operations.

In this whitepaper, we focus on the vertical aspects of Open vRAN. In order to help CSPs accelerate the adoption of this new technology, NEC, Netcracker, RedHat, Dell Technologies and AltioStar have joined forces to define, validate, and successfully build a reusable reference architecture in which vRAN can run on.

This suggested model uses the combination of Dell EMC's highly competitive open hardware as the foundation of the architecture, Redhat's Open Stack as a widely used virtualization platform, and AltioStar's field operable, O-RAN standards compliant, virtualized software. Further, Netcracker provides the end-to-end operations environment to automate VNF deployment, configuration, lifecycle management and radio resource optimization. Finally, NEC brings together the overall Open vRAN ecosystem enablement, in which it is

responsible for optimal and best-of-breed based vendor neutral model designing, system integration, and the delivery of a secure and mission critical carrier-grade network.

NEC is the only company building a Vertical Business Platform that is engaged in all telecom network layers from front-haul to back-haul, physical and virtual networks, network orchestration, to management and business layers with OSS and BSS. Given these end-to-end capabilities, extensive expertise and insightful understanding of Open vRAN benefits, NEC positions itself as a strategic, reliable partner for the rapid and seamless implementation of CSPs' next generation mobile networks.

This paper outlines Open vRAN reference architecture options, highlighting the key benefits and explaining how CSPs can cost effectively deploy and operate the validated Open vRAN architecture to meet the challenging new market requirements.

vRAN Reference Architecture at a Glance

Open vRAN is a software-intensive web scale architecture with open interfaces, disaggregating the hardware and software to allow a multivendor solution that combines the best-of-breed for each element of the network and enables the use of commercial off-the-shelf (COTS) components. It aims to realize the simplified deployment and operation of CSPs network assets.

The following illustrates Open vRAN architecture at high level.

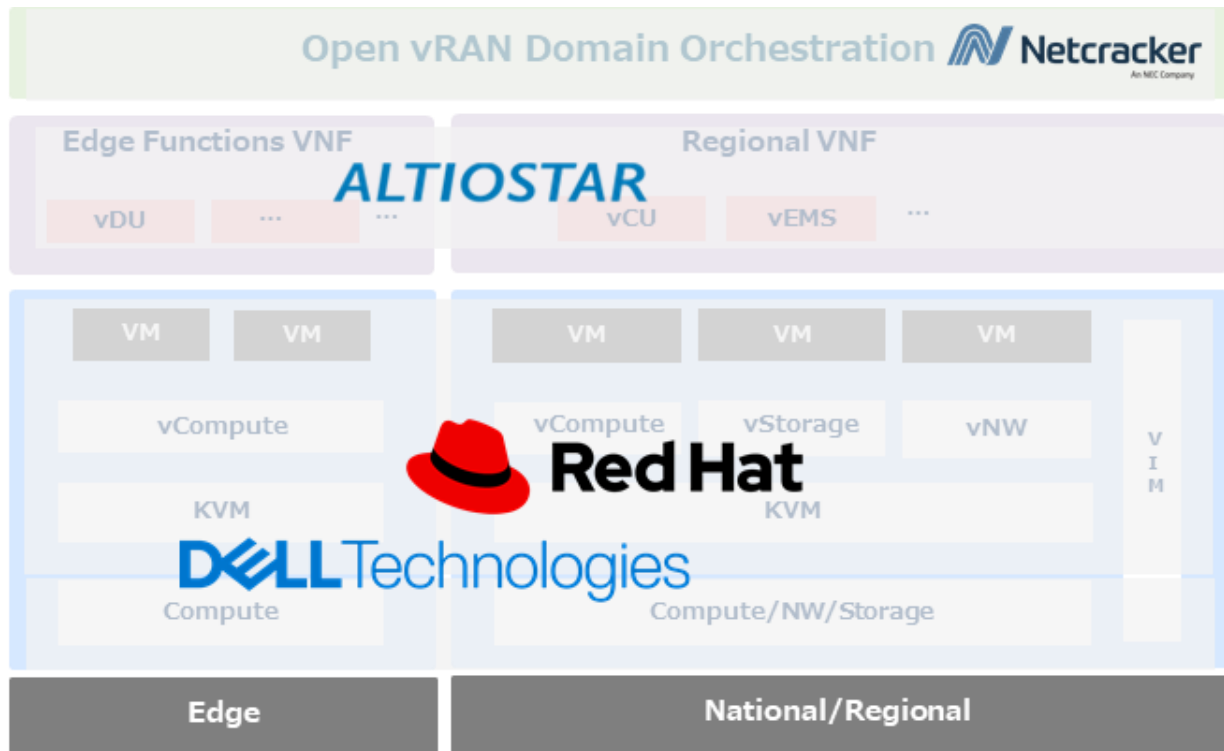


Figure 1: High Level Reference Architecture of vRAN

In this architecture, RedHat OpenStack is deployed to virtualize the resources, and the distributed compute architecture is adopted to maximize the resource usage efficiency. Intel FPGA card is implemented on Dell EMC Hardware to maximize the radio network processing whilst minimizing any latency and jitter. Altiostar brings virtualized network functions which comply with Open vRAN standards running on the infrastructure and Netcracker provides the end-to-end orchestration environment to automate VNF deployment, lifecycle management and manage overall resources.

Furthermore, NEC is responsible for integrating the diversified hardware and software components and overall project management to operationalize this reference model. Given that RAN requires physical aspects for front-haul, such as site design and installation of FPGAs, a comprehensive expertise in both physical and virtual networks is an essential requirement to execute the system integration.

Altiostar's Open vRAN, which operates based on VNFs in this architecture, consists of eNB-vCU (eNB virtual Centralized Unit), eNB-vDU (eNB virtual Distributed Unit), vEMS (virtual EMS), and third-party RRH (Remote Radio Head)/ RRU (Remote Radio Unit) & Antennas as shown in Figure 2. Each cell-site may have multiple third-Party RRHs & Antennas connected to Altiostar's eNB-vDU using an ORAN 7-2x Fronthaul interface or a single RIU over CPRI. The eNB-vDU can be located at the site, or in an edge

data center cloud, serving multiple O-RU. Beyond that, multiple eNB-vDU instances interface with a single instance of an eNB-vCU, which can be run in a centralized data center cloud. And eventually, multiple eNB-vCU instances interface with a single instance of vEMS, which can be run in a centralized data center cloud. eNB-vDU and eNB-vCU can be run in the same data center cloud.

AltioStar supports the following two deployment models:

1. CPRI RRU interfaced with DU via RIU: AltioStar has developed an eCPRI to CPRI protocol converter that allows using CPRI RRUs with AltioStar's vDU.
2. eCPRI RRU interfaced directly with DU: In this model the RRUs supporting eCPRI are interfaced directly with AltioStar's vDU

The two models provide tremendous flexibility to the operators in deployment by supporting both types of RRUs – CPRI and eCPRI. It may be noted that a majority of the RRUs in the market today are CPRI though the newer model are now supporting an eCPRI interface. RIU could also be used by an operator to convert their current BBU based deployment into an Open vRAN deployment by just replacing the proprietary BBU with AltioStar Open vRAN, while retaining the RRUs, thus saving majority of their CapEx investments. The solution in this whitepaper is based on CPRI+RIU architecture but also demonstrates the same architectural principles apply for direct eCPRI interface as well.




Figure 2: Overall Configuration of Open vRAN

Figure 3a illustrates AltioStar's Open vRAN architecture's functional split among eNB-vCU, eNB-vDU, and third party O-RU with 7-2x Fronthaul (O-RU). Figure 3b illustrates AltioStar's Open vRAN architecture's functional split among eNB-vCU, eNB-vDU, RIU and third-party RRH. Either the O-RU or RIU supports the lower part of the PHY (LTE L1 functionality). The eNB-vDU is a VNF (Virtual Network Function) as well as the eNB-vCU.

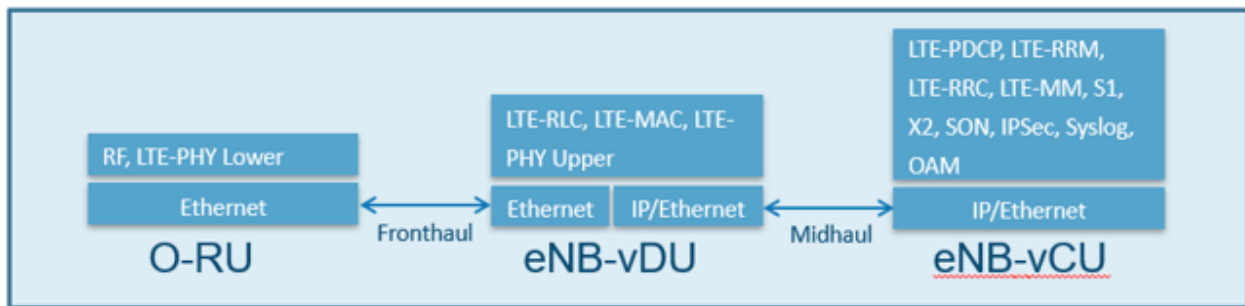


Figure 3a: Functional Split Configuration of Altiostar Open vRAN with O-RU

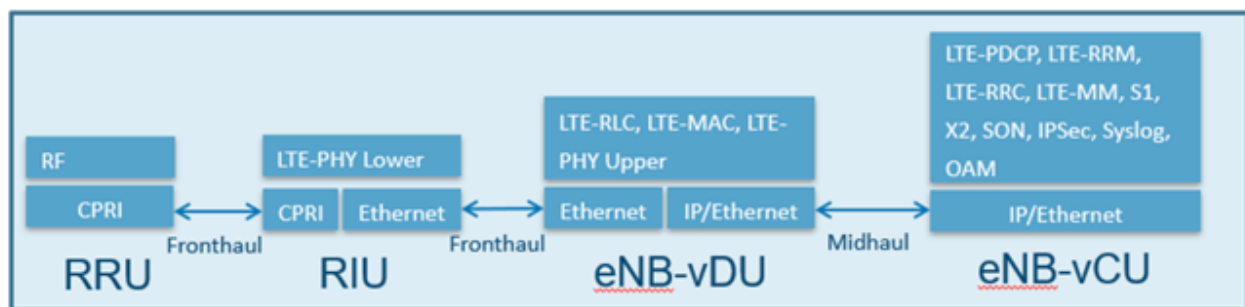


Figure 3b: Functional Split Configuration of Altiostar Open vRAN with RRH

vRAN Deployment Options

There are several Open vRAN deployment options based on the density and latency required. Depending on the site and location requirements, as well as performance of fronthaul, deployment models should be carefully considered when virtualizing the RAN function. In the Red Hat OpenStack Platform (RHOSP), Distributed Compute Nodes (DCN) for edge computing support the practice of deploying compute nodes to remote locations that might not be typical data centers.

In accordance with this approach, compute nodes can be deployed to remote locations, and the controller services can be more tolerant of the network issues that might arise, such as connectivity and latency. With the Red Hat distributed compute node architecture, a single overcloud heat stack spans multiple sites, with each site allocated to its own Availability Zone (AZ). The separate AZ approach allows you to target specific workloads to each site. Based on this distributed architecture, following three types of models are defined to deploy Open vRAN with more resource-efficient manner.

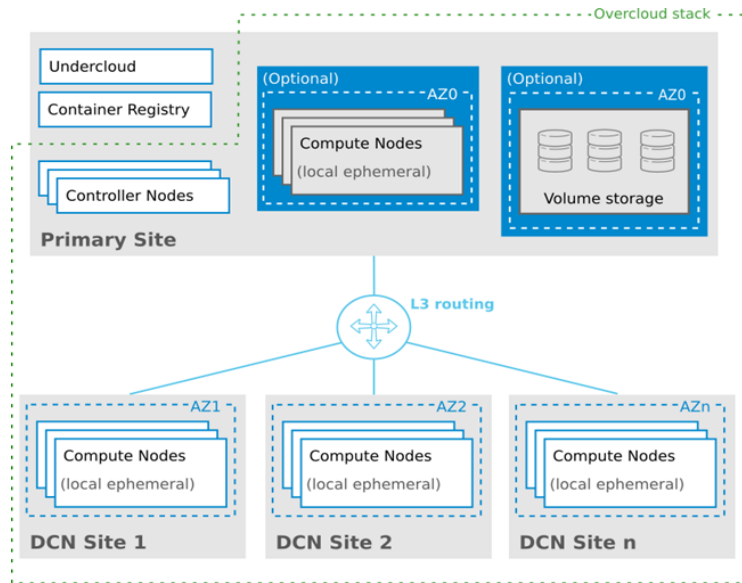


Figure 4: Concept of Distributed Compute Node

Low Density @ Tower

The model-1 uses a distributed compute node on a small-footprint server which is optimized for the deployment near to the tower. The server is expected to run a relatively small amount of vDU/vCU. Also, the computes use ephemeral storage and require QCOW2 based images for vDU/vCU.

For this configuration, the following minimum DCN RHOSP-specific network requirements per site (excludes VNF data nets) should be satisfied:

- Provisioning network (routed, DHCP Relay required)
- Internal API network (routed)
- Tenant underlay (if desired - routed)
- IPMI Network

ORAN or CPRI RU

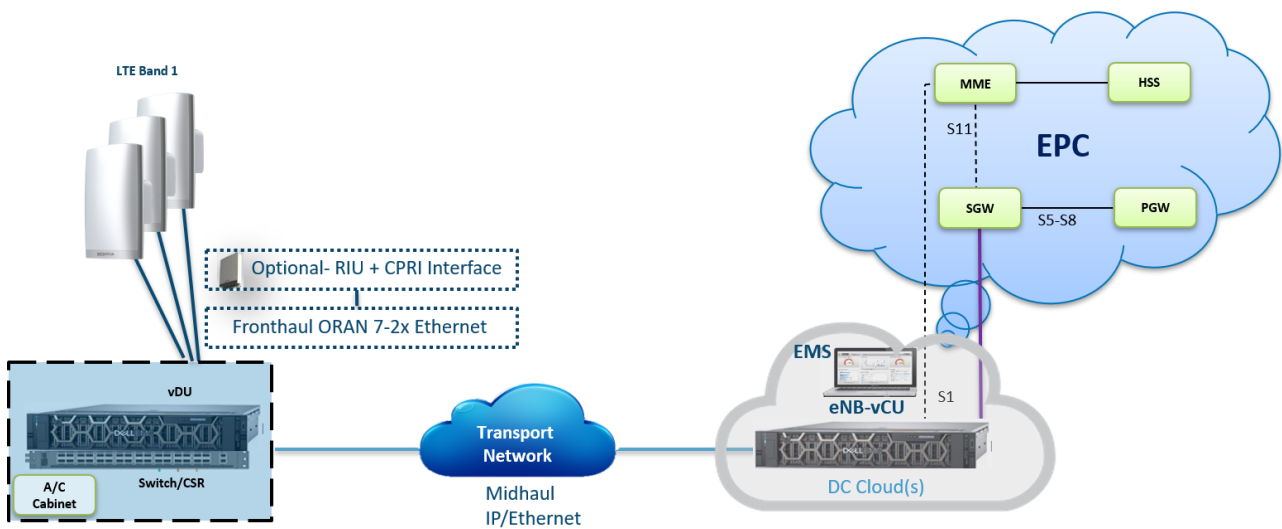


Figure 5: Low density @Tower

High Density @ Tower

The model-2 uses a distributed compute node server which includes an Intel N3000 Vista Creek FPGA card to offload mobile fronthaul processing. These servers are configured to run enough amount of vDU/vCU to accommodate high density cell sectors at the tower site. Also, computes use ephemeral storage and requires QCOW2 as based images for vDU/vCU.

For this configuration, the following minimum DCN RHOSP specific network requirements per site (excludes VNF data nets) should be satisfied:

- Provisioning network (routed, DHCP Relay required)
- Internal API network (routed)
- Tenant underlay (if desired - routed)
- IPMI network

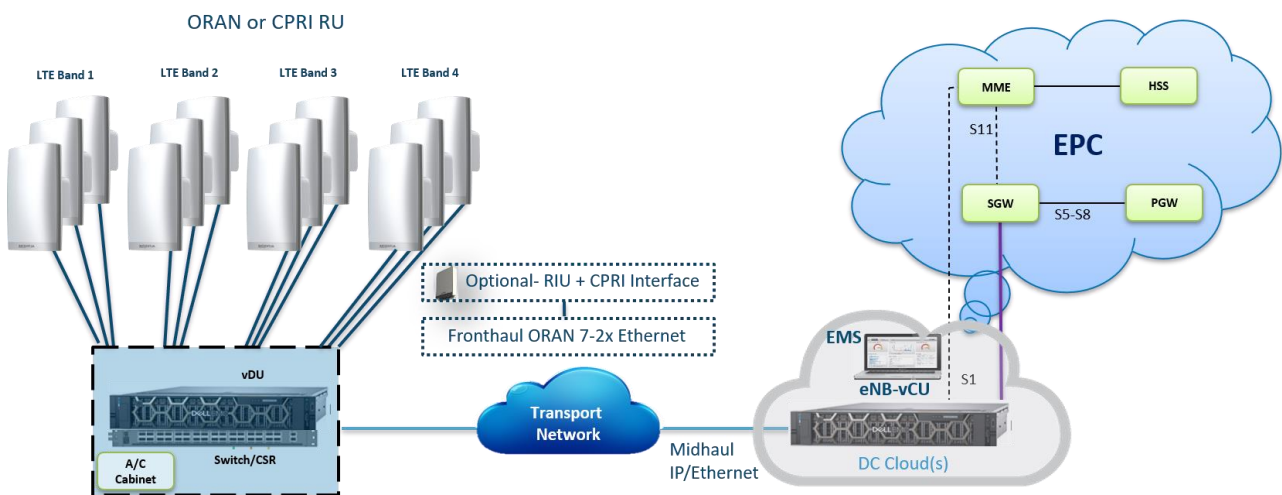


Figure 6: High density @Tower

High Density @ Edge Cloud

The model-3 uses server with large foot print with Intel N3000 Vista Creek FPGA cards and will host vDU/vCU. Multiple compute nodes will be deployed based on sector requirements for vCU and vDU. vCU compute nodes do not require FPGA cards and can support oversubscription to maximize the resource efficiency.

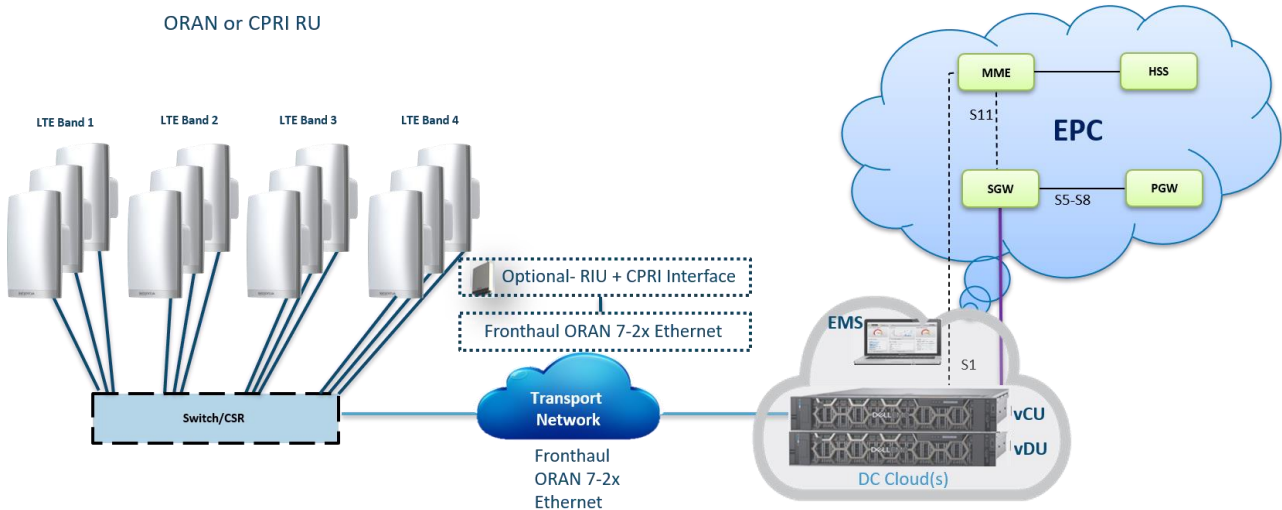


Figure 7: High density @ Edge Cloud

The RHOSP Cluster deployed in these models can support additional DCN sites as needed. Currently Red Hat supports RHOSP 13 deployments with 300 compute nodes. Red Hat is actively working on expanding the number of compute nodes that can be supported from a single director node. As mentioned above that is 300 today, however 500 have been successfully testing internally and we expect official support for 500 soon. The overall target is up to 1000 compute nodes. There are other factors such as failure domains, update/upgrade maintenance window times and such which should be factored into planning for scale. Smaller deployments with up to 100 Compute nodes per cluster are generally more flexible and easier/quicker to manage.

Hardware Configuration for vDU/vCU

The following are the two examples of server for the compute to be deployed. Even though this does not mandate server specification to on-board Open vRAN, the intent is to show a set of re-usable example which is already tested to on-board Open vRAN.

Dell EMC PowerEdge R640 server

The Dell EMC PowerEdge R640 is the ideal dual-socket, 1U platform for dense scale-out cloud computing. The scalable business architecture of the Dell EMC PowerEdge R640 is designed to maximize application performance and provide the flexibility to optimize configurations based on the application and use case.

With the Dell EMC PowerEdge R640 you can create an NVMe cache pool and use either 2.5" or 3.5" drives for data storage. Combined with up to 24 DIMM's, 12 of which can be NVDIMM's, you have the resources to create the optimum configuration to maximize application performance in only a 1U chassis. This can simplify and speed-up deployments of the Red Hat OpenStack Platform.

Dell EMC PowerEdge R740 servers

Maximize your application performance with the Dell EMC PowerEdge R740 server, which brings you the perfect balance of accelerator cards, storage and compute resources in a 2U, 2-socket platform.


With a wide range of GPU and FPGA options, the Dell EMC PowerEdge R740 has the versatility to adapt to virtually any application and provides the optimum platform for VDI deployments. The Dell EMC PowerEdge R740 offers up to 16 x 2.5" or 8 x 3.5" drives and iDRAC9, so you can scale to meet demands and simplify the entire IT lifecycle.

Dell EMC PowerEdge R740xd servers

The Dell EMC PowerEdge R740xd delivers a perfect balance between storage scalability and performance. The 2U two-socket platform is ideal for software defined storage. The R740xd versatility is highlighted with the ability to mix any drive type to create the optimum configuration of SSD and HDD for either performance, capacity or both.

The Dell EMC PowerEdge R740xd is the platform of choice for software defined storage and is the foundation of Red Hat Ceph Storage for this Architecture Guide - Version 1.0

Physical Topology



vDU Server Specification



Table 1: Recommended configuration for vDU

	Description
Model	Dell EMC PowerEdge R740 XL (2U) Server
CPU	2 x Cascade Lake 6248, 2.4GHZ, 20 Core
Memory	12 x 16GB
Disk	2 x 480 SSD, Boss Raid Controller: RAID 1
Network Interfaces	25G (dual port) for OVS & OpenStack control plane

	CPU 0: 25G (dual port) for SR-IOV CPU 1: 25G (dual port) for SR-IOV
FGPA	CPU 0: Intel N3000 Vista Creek (2 x 25G) CPU 1: Intel N3000 Vista Creek (2 x 25G)

vCU Server Specification



Table 2: Recommended configuration for vCU

	Description
Model	Dell EMC PowerEdge R740 XL (2U) Server
CPU	2 x Cascade Lake 6248, 2.4GHZ, 20 Core
Memory	12 x 16GB
Disk	2 x 480 SSD, Boss Controller: RAID 1
Network Interfaces	25G (dual port) for OVS & OpenStack control plane CPU 0: 25G (dual port) for SR-IOV CPU 1: 25G (dual port) for SR-IOV

Intel FPGA Programmable Acceleration Card (Intel FPGA PAC) N3000 Card

The Open vRAN vDU node needs to be able to perform a large amount of Digital Signal Processing (DSP) in real time. To accomplish this the vDU node will need to offload the DSP to a hardware device. The Intel PAC N3000 card has been selected and certified by Altiosstar for this task.

The Intel® Programmable Acceleration Card (Intel® PAC) N3000 accelerates network traffic for up to 100Gbps to support low latency, high bandwidth 5G applications. It allows vendors to create custom tailored accelerated virtualized solutions for Open vRAN and Core network workloads and achieve faster time to market (TTM) with support of industry standard orchestration and open-source tools. Intel is accelerating NFV adoption by enabling ecosystem partners such as Telecom Equipment Manufacturer (TEM), Virtual Network Function (VNF) vendors, system integrators, and Telecommunications companies to bring scalable and high-performance solutions to market.



Figure 8: Intel® PAC N3000 Card

Orchestration and Management of Open vRAN

Netcracker Open vRAN Domain Orchestration comprises a suite of orchestration and OSS systems, together with advanced analytics, to automate the end-to-end service lifecycle from planning and design to activation and assurance across the entire domain. Orchestration systems include Network Orchestration and VNF Management for vCU/vDU resources and Service Orchestration for the entire physical and virtual Open vRAN domain. OSS systems include Configuration Management, Active Resource Inventory, Fault and Performance Management as well as Planning and Design Management that, together with the orchestration systems, enable end-to-end automation for day 0 and ongoing operations.

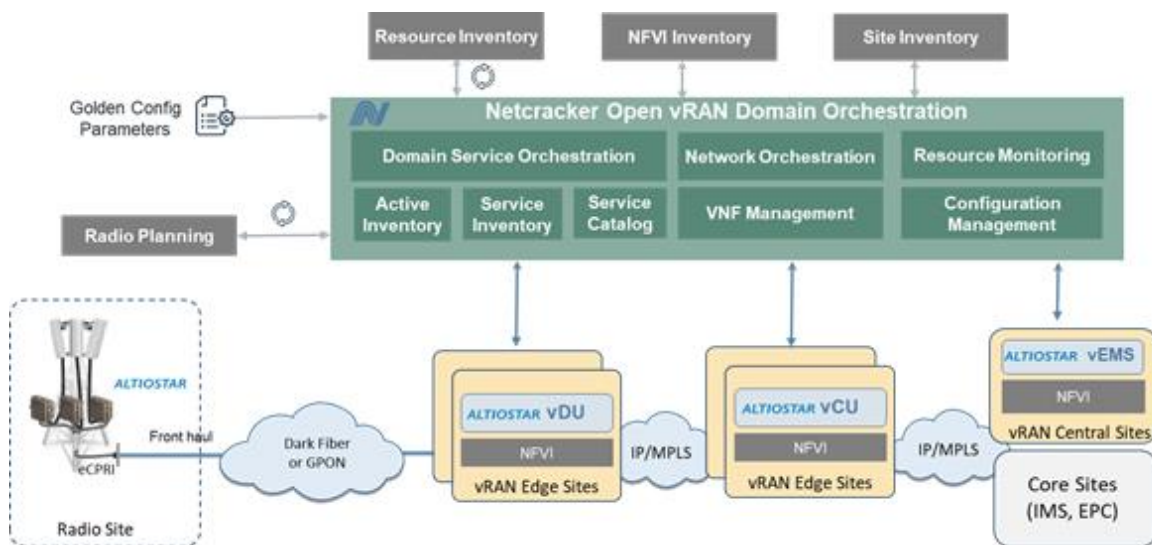


Figure 9: Open vRAN Domain Orchestration

Open vRAN Domain Orchestration can be divided into three stages:

- RAN domain long-term financial and technical planning and design
- Open vRAN activation and configuration of physical and virtual network functions
- Open vRAN lifecycle management with continuous tuning and optimization

In this paper we focus on the orchestration and automation aspects of Open vRAN activation and lifecycle management as part of the reference architecture.

Deployment and activation of Open vRAN components is automated using Netcracker Network Orchestration and VNF Manager (ETSI NFV MANO compliant). The MANO system provides the optimal placement of virtual resources at the vCU and vDU edge sites. Virtual resources are automatically instantiated by the MANO, which also performs centralized license management for multivendor software. Each network element (physical and virtual) is then configured in a central manner with specific parameters using Netcracker Configuration Management. Netcracker Service Orchestration brings together the physical and virtual assets to enable automated service provisioning and service lifecycle management across the Open vRAN domain. Built into the solution is the optional capability to manually validate and accept each operational step.

The procedure to instantiate and configure vCU/vDU systems is as follows:

1. When the radio site installation is complete and the O-RU is powered up, a notification is sent to Netcracker Open vRAN Domain Orchestration (DO) through AltioStar's EMS (EMS).
2. Open vRAN DO initiates the eNB (vDU/vCU) deployment process on a particular site/NFVI and interacts with the edge cloud VIM and VNF Manager to instantiate vCU and vDU VNFs, apply day 0 configuration and perform additional NFVI/network configuration.
3. Open vRAN DO then creates a self-commissioning profile for the eNB in the EMS.
4. The deployed vCU requests initialization through the EMS, triggering the Open vRAN DO to start the eNB service activation process with design and template-based configurations.
5. The EMS requests Open vRAN DO for eNB license information and imports it to the EMS.
6. The EMS requests the Open vRAN DO for eNB configuration, which is provided to the EMS to prepare the eNB configuration data. The EMS notifies Open vRAN DO about progress of the self-commissioning process.

After activation, Open vRAN lifecycle management uses the orchestration systems to automate operations after initial instantiation and includes scaling, healing, upgrades, termination and continuous optimization. Active Resource Inventory maintains real-time inventory and topology data that can be used by the orchestration and assurance systems to perform automated lifecycle events and optimization. Assurance systems collect events, metrics and telemetry across the vRAN network, triggering closed-loop control with orchestration and enabling automated root cause analysis and auto healing. Configuration Management is used to continuously compare eNB configurations with the golden configurations. Netcracker's AI-driven Advanced Analytics is also used to further optimize the network through trending and forecasting and by using KPIs and business data to identify optimizations that will yield positive economic benefits.

Conclusion

In this paper, we outlined the reference architecture as a reusable model to help CSPs accelerate the adoption of this new technology. We also highlighted the benefits for CSPs to cost-effectively transform their RAN and simplify operations for the deployment of new innovative services. We believe that this validated Open vRAN reference model, with NEC's RF expertise and telco system integration capabilities, will help CSPs to expand and evolve their networks at a faster pace to accommodate surging traffic growth and improve their user experience while reducing the CAPEX and OPEX.

Disclaimer:

This whitepaper is issued for information only. It does not constitute an official or agreed position of NEC, Netcracker, RedHat, Dell Technologies and AltioStar. The views expressed are entirely those of the author(s). We decline all responsibility for any errors and any loss or damage resulting from use of the contents of this White Paper. Furthermore, Information on this documentation contains forward-looking statements regarding estimations, forecasts, targets and plans. The White Paper is made based on information currently available and certain assumptions considered reasonable as of the date of this material. These determinations and assumptions are inherently subjective and uncertain and are not guarantees of future performance, and actual operating results may differ substantially due to a number of factors.