Hewlett Packard Enterprise

Technical white paper

Check if the document is available in the language of your choice.

VRAN 2.0 ON HPE INFRASTRUCTURE

How HPE can help change the 5G game

CONTENTS

PURPOSE OF THIS DOCUMENT

The intent of this document is to share thought leadership on virtual Radio Access Network (RAN) 2.0 from Hewlett Packard Enterprise (HPE). HPE is working extensively with many leading and large telco operators to analyze the vRAN 2.0 both from business (commercial and economic viability) as well as technical aspects. This has led to successful deployment of vRAN 2.0 solution with a number of leading telco operators.

This white paper is meant to share these experiences to help worldwide telco operators and respective hardware and software partners in this journey to vRAN 2.0.

CURRENT STATE OF AFFAIRS IN RAN

Radio Access Network (RAN) is the key piece of 3GPP cellular network infrastructure, which connects cellular devices (otherwise known as User Equipment or UE) to a core network. The key function of RAN therefore is to provide an air interface toward UE and handle all aspects associated with radio signal processing, radio resource control, processing signaling associated with subscriber attaching to the network and using its services while moving around seamlessly.

On one hand, RAN infrastructure is the ultimate edge of an operator's cellular network, which directly interfaces end subscribers and in the main part defines the overall quality of cellular service and associated subscriber experience.

On the other hand, RAN infrastructure is directly responsible for efficient utilization of radio spectrum, which is the most valuable and limited resource of any cellular operator.

It is no surprise that RAN consumes by far the largest chunk of an operator's network CAPEX and OPEX, as well as represents the biggest R&D spent of cellular technology providers and therefore is one of the most dynamic and innovation-intensive network domains.

Architecturally, RAN is a large, geographically distributed assembly of nodes called base stations. They come in all form factors—starting from outdoor macro nodes forming a blanket cellular coverage in a given area, to outdoor small cells providing point capacity in traffic hot spots, to indoor and enterprise RAN solutions improving quality and capacity inside venues.

Examples of RAN sites are illustrated as follows.

Macro tower, suburban site

Rooftop, urban site

Camouflaged rooftop, urban site

Concealed small cell, urban site

FIGURE 1. Examples of RAN site deployments

Density and capacity of RAN nodes depend on subscribers' geographical distribution and mobility patterns as well as spectrum available for an operator, with each macro base station covering an area from a few hundred feet to a few miles. The introduction of new access technologies and the release of new higher frequency spectrum bands for 5G cellular increases the density of the RAN. While the global number of base stations is pushing 10M mark, the introduction of 5G is predicted to bring a tremendous increase in RAN density, anywhere from 5X to 10X.

Given the significant amount of real-time signal processing within the RAN, as well as a need to fit within a very limited physical and power footprint across the diverse real estate of different RAN sites, existing base stations have been based mostly on specialized appliances. They come in many form factors optimized for specific deployments, however all generally include a Baseband Unit (BBU) and a Radio Unit (or RU).

- The BBU is a purpose-built appliance processing baseband signals in the digital domain, with RAN software tightly coupled with underlying hardware, which itself is based on special purpose custom ASICs, DSPs, and FPGAs.
- The RU generally executes digital/analog conversion as well as processing RF analog signal coming to/from antenna.

Example of a traditional macro base station architecture is illustrated as follows.

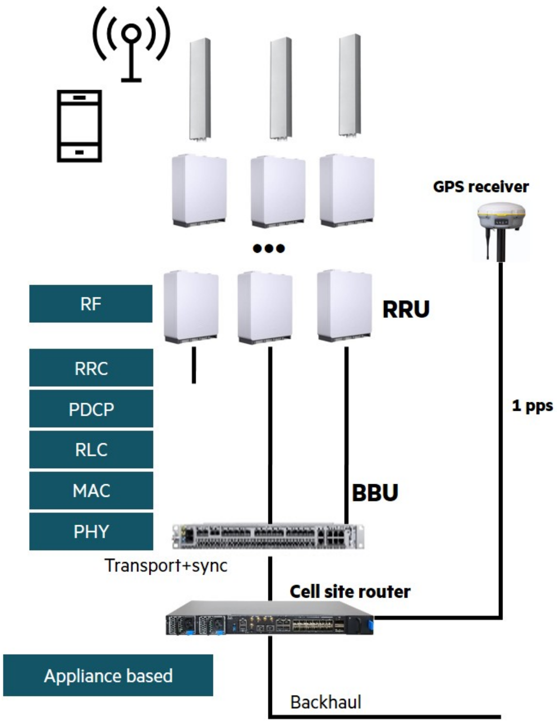

FIGURE 2. Traditional RAN base station architecture

Traditional RAN architecture is a closed ecosystem of SW tightly coupled with the underlying hardware from a single Network Equipment Provider (NEP) as a closed, proprietary appliance. The main focus of technological evolution of the traditional RAN infrastructure has been so far mainly developing more power efficient and dense appliances for BBU and RRU, and introducing deployment and mounting options for them.

Up until now, that evolution model has been fairly successful in providing operators with reliable RAN technology. However, constant traffic growth, the introduction of new generations of access technologies in combination with generally flattening revenue per subscriber has put enormous pressure on operators to constantly improve operating and capital efficiency of RAN. In many cases, traditional business and technological models of building radio access networks are no longer able to cope with that pressure.

Recent advancements in general-purpose compute as well as the number of telco driven industry initiatives that is opening up and standardizing interfaces within RAN domain poses monumental changes to the existing technological and business model of the RAN. This new paradigm of RAN, often called Cloud RAN or Virtual RAN, is based on the following technological principles:

- Fully decoupled software running on abstracted general-purpose hardware, with best-of-breed approach toward RAN solution components
- Functional components implemented as abstracted software interacting via open standardized interfaces
- General-purpose edge compute infrastructure supporting any use case and workload, with RAN workload as one of the tenants
- Cloud-native capabilities in deployment, lifecycle management, scaling, redundancy of RAN workload and underlying infrastructure

While promising fundamental improvements into RAN business and technological model, vRAN can be implemented in numerous ways each bringing many unknowns to an operator.

HPE, in partnership with leading best-in-class technology providers and leading telcos, built, deployed, and formalized learnings from a number of early vRAN deployments. Based on those learnings, HPE has created a reference technical solution which provides production grade vRAN with best-in-class total cost of ownership (TCO), verified in field.

The following sections describe main principles of RAN virtualization and key improvements it brings to traditional RAN.

INTRODUCTION AND OBJECTIVES OF RAN VIRTUALIZATION

Representing around 50% of total operator's spend on a network, RAN is very crucial segment of mobile network infrastructure. Constant evolution of existing access technologies as well as introduction of 5G NR (New Radio) puts ever-increasing pressure on RAN architecture.

Traditionally based on proprietary monolithic appliances, existing RAN architecture is reaching the limit to cope with that pressure in a cost-efficient manner. Based on experience of similar challenges in telco core networks, as well as wider technology industry, sustainable RAN evolution demands a transition to a more open, disaggregated, decoupled model with separation between software and hardware, along with introduction of cloud-native technologies into access domain.

This architectural evolution, sometimes called RAN virtualization or vRAN, targets to bring monumental improvements to traditional RAN:

- **Flexibility**, as vRAN enables best-of-breed approach toward RAN components, versus monolithic network infrastructure from a single provider in a traditional RAN model
- **Scalability**, as decoupling of software from hardware allows for independent horizontal scaling of infrastructure to address constant evolution of radio access (for example, introduction of new features, access technologies, frequency bands), versus frequent vertical upgrades and infrastructure swap projects of appliance hardware
- **Operational efficiency**, as open underlying compute platform enables flexible management and automation of vRAN infrastructure based on wider tech industry achievements in infrastructure-as-a-code and cloud-native applications management
- **Resilience to security threats**, as components of the system from one vendor can easily be exchanged for another, if the supply chain or vendor is deemed to be compromised
- **Improved security**, as open interfaces enable monitoring to detect attacks

Combination of those vRAN characteristics leads to major improvement in how operators procure, deploy, and evolve their RAN.

Side benefits of virtualizing RAN

RAN virtualization is based on deploying general-purpose compute infrastructure at the very edge of a cellular network. That same infrastructure can be used as a platform for innovation with all other edge use cases. Underlying general-purpose edge cloud will combine those use cases on the same hosting infrastructure, providing further synergies and improved return on investment into vRAN. The following figure illustrates that vision of a multitenant multipurpose telco edge cloud, running vRAN as well as other telco edge use cases on the same cloud platform.

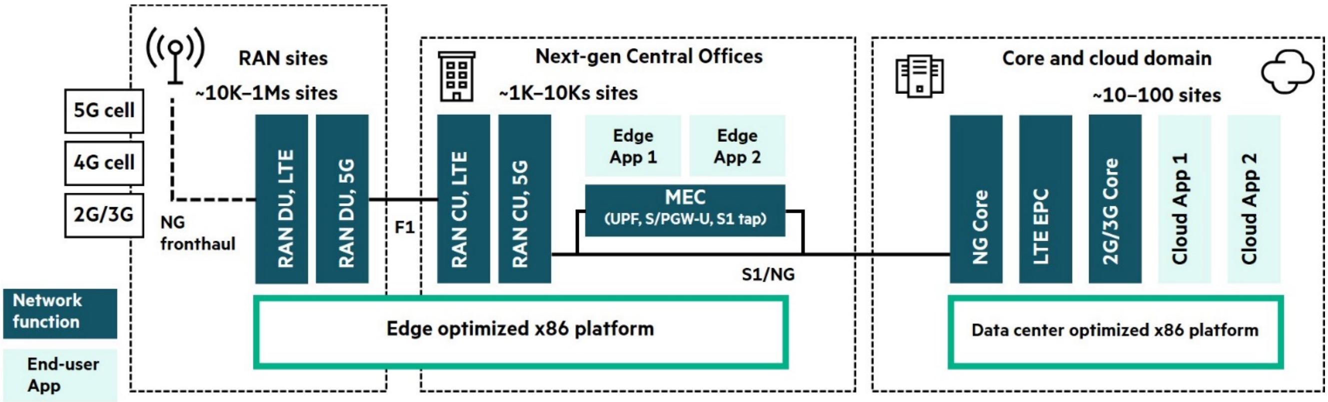

FIGURE 3. Universal edge cloud enabling multiple use cases

Examples of some of the edge use cases are as follows:

- **Private networks and cellular breakout (Private LTE/5G)**

  Private network gives telco an ability to sell a custom connectivity service to an enterprise, with ability to engineer connectivity (traffic routes, policies) around the needs of concrete enterprise use cases. Private network solution provides ability to architect the wireless network for demands of concrete location and use case, while enjoying all benefits of cellular technology—reliability, native seamless mobility over wide area, low-latency and scalability. Building private networks and providing a cellular local-breakout for enterprise traffic requires core and breakout functions to be deployed at far edge of telco network or on enterprise premises. This model requires underlying general-purpose edge cloud to host necessary functions.

- **Multiaccess Edge Compute (MEC)**

  Having custom connectivity around enterprise needs is the first step for introduction of more advanced digital services at the edge—starting from providing a hosting platform for edge-enabled applications, to providing PaaS and SaaS services to enterprise. Modern edge-enabled applications (for example, AI-based video analytics, AR edge processing, and so on) require powerful underlying compute platform, in many cases equipped with specialized accelerators (for example, GPUs or FPGAs). The vRAN edge cloud platform will be able to co-host those enterprise MEC applications with vRAN functions.

- **Network slicing and distributed user plane**

  Slicing is introducing differentiation on the services (with quality of service features as an example) and the way the network user plane/control plane and RAN is designed, managed, and associated to a network slice. This association can be dynamic, adding/removing components in different location, changing parameters, but also scaling components, including the RAN components, as number of UE (User Equipment) increases or traffic increases. Also, management of the slice can be restricted to certain users, including the management of the RAN components of the slice.

  High-reliability, low-latency communication, and enhanced mobile broadband use cases define a need for not only functional separation and independent scaling, but also optimal and dynamic placing of signaling plane and user plane functions in a network, based on operator policies and demands of a concrete use case. Dynamically placing core user plane functions (for example, UPF or PGW/SGW-U) at the edge enable flexible traffic steering across the network as well as optimal scaling of user plane in response to growing heterogeneous traffic.

While promising significant direct RAN improvements as well as side benefits, vRAN can be implemented in countless number of ways. Not every implementation of vRAN brings full value to the operator. Recognizing this fact, the next section describes alternative approaches to RAN virtualization and establishes a concept of vRAN 2.0.

RAN SPLIT OPTIONS AND DEFINITION OF VRAN

The first fundamental step of bringing more openness and flexibility into RAN architecture is disaggregating the monolithic baseband processing of a base station into separate distinct layers, communicating via well-defined interfaces. That disaggregation can be implemented in number of ways, with 3GPP TR 38.801 defining up to 8 potential splits between centralized and distributed portions of baseband.

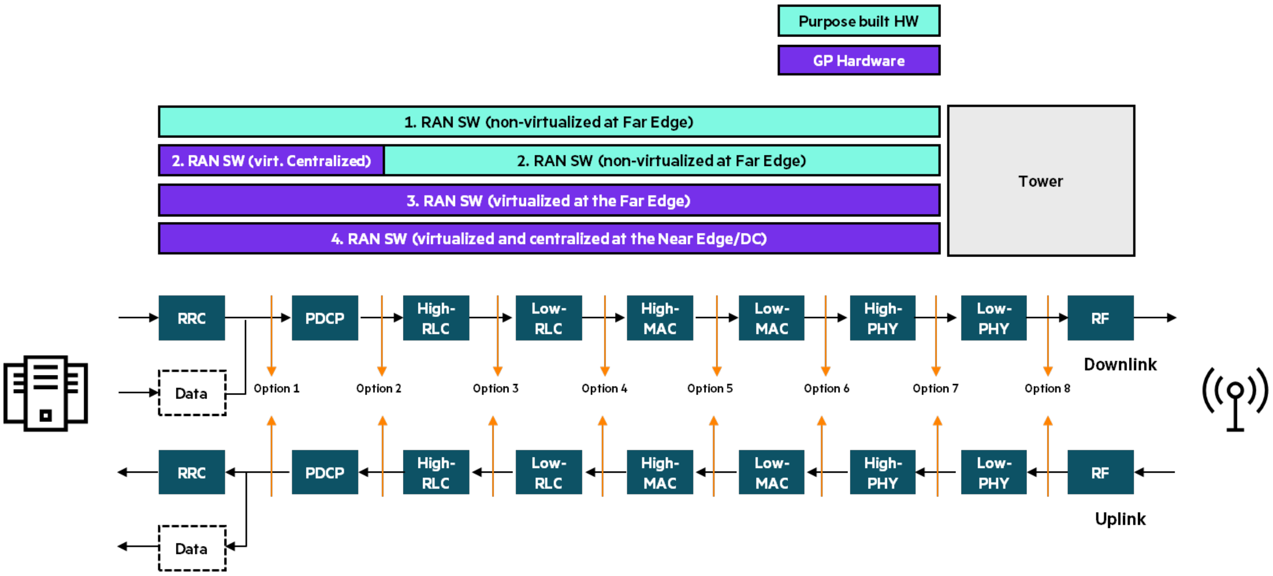

FIGURE 4. 3GPP TR 38.801 split options for baseband

Traditional RAN architecture, that is 4G, is represented by split option 8, where all baseband processing is implemented as a single monolithic appliance, and RF functions are performed by a radio unit. Alternative approaches are based on decoupling all functional elements up to a certain split point and porting them to general-purpose hardware, while implementing open standardized interfaces across each split point.

The degree of that decoupling significantly affects total economic and operational impact of vRAN compared to the traditional RAN approach. Following sub-sections describe the two most established options of vRAN.

vRAN 1.0 overview

The first attempt to virtualize baseband stack of macro RAN followed higher-level splits of baseband processing—decoupling non-real-time portions of baseband stack and porting it to general-purpose compute platform, while keeping lower-level real-time functions on dedicated appliance. These earlier vRAN architectures converged to 3GPP split option 2, with RRC and PDCP layers virtualized on general-purpose platform (vCU), and RLC-PHY kept on a dedicated appliance (DU). This approach is sometime referred to as vRAN 1.0, and overall architecture is illustrated as follows.

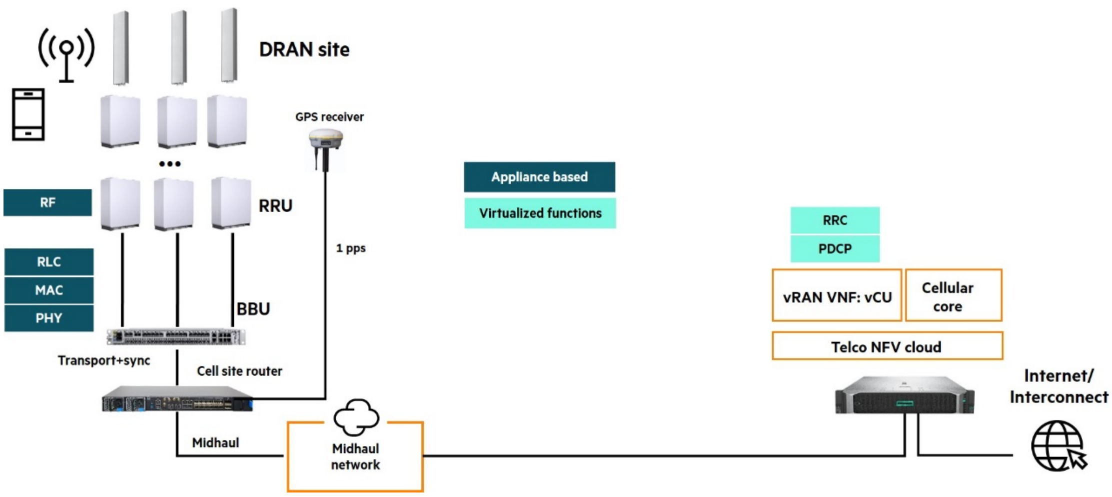

FIGURE 5. vRAN 1.0 architecture

While providing some technical benefits (for example, pooling gains in PDCP/RRC, enabling number of inter-node coordination features), vRAN 1.0 did not significantly change economics of radio access. This is mainly due to vCU processing minor portion of overall baseband (PDCP/RRC represent roughly 7–10% of overall baseband processing), therefore major chunk of processing was still executed on DU appliances not allowing for economic benefits of full hardware/software decoupling. Many operators found that vRAN 1.0 solution did not provide much (if any) economic benefits compared to traditional RAN—which is not surprising given that architecturally not much is offloaded to a general-purpose hardware.

Although this did not significantly improve the overall RAN TCO, vRAN 1.0 nevertheless was an important first venture into virtualizing portions of baseband processing stack and proving the technical feasibility of general-purpose platforms in the RAN domain. Therefore it is considered as the first stepping stone toward realizing final RAN virtualization objectives.

vRAN 2.0 overview

Unlocking the full potential of vRAN required porting of all baseband processing functions from baseband appliance to general-purpose hardware. Following 3GPP split option 7–2 for vDU/Radio and split option 2 for vDU/vCU, this concept is sometimes called vRAN 2.0 and it is the first vRAN architecture that unlocks all benefits of virtualized radio access.

The following figure illustrates the overall architecture of vRAN 2.0 base station.

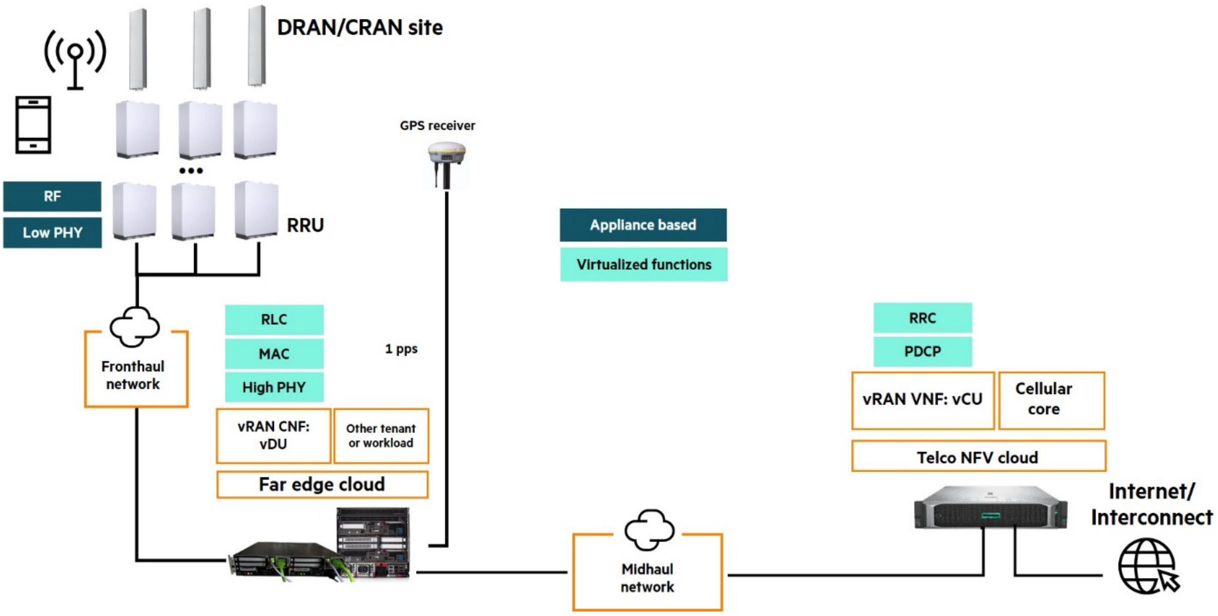

FIGURE 6. vRAN 2.0 architecture

All the baseband functions, from non-real-time PDCP and RRC to real-time RLC, MAC, and PHY layers are ported to general-purpose compute hardware. The following figure illustrates functional split of vRAN 2.0 within PHY layer, with split 7.2 defining a border between vDU and vCU functions running on general-purpose hardware and RF/Low PHY functions running on Radio Unit appliance.

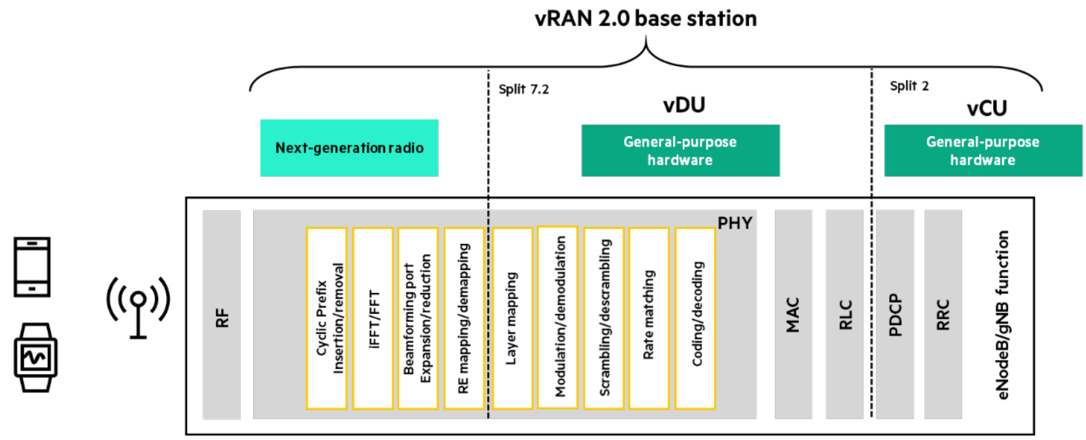

FIGURE 7. Functional splits between Radio, vDU, and vCU in vRAN 2.0

Following section illustrates lower-level design of vDU and vCU and defines main solution components of vRAN 2.0, presenting common reference architecture of vRAN 2.0.

VRAN 2.0 REFERENCE ARCHITECTURE

vRAN 2.0 reference architecture brings together number of hardware, firmware, and software components from HPE and technology partners and is illustrated in the following figure. The HPE products in vRAN 2.0 are described in detail in subsequent chapters.

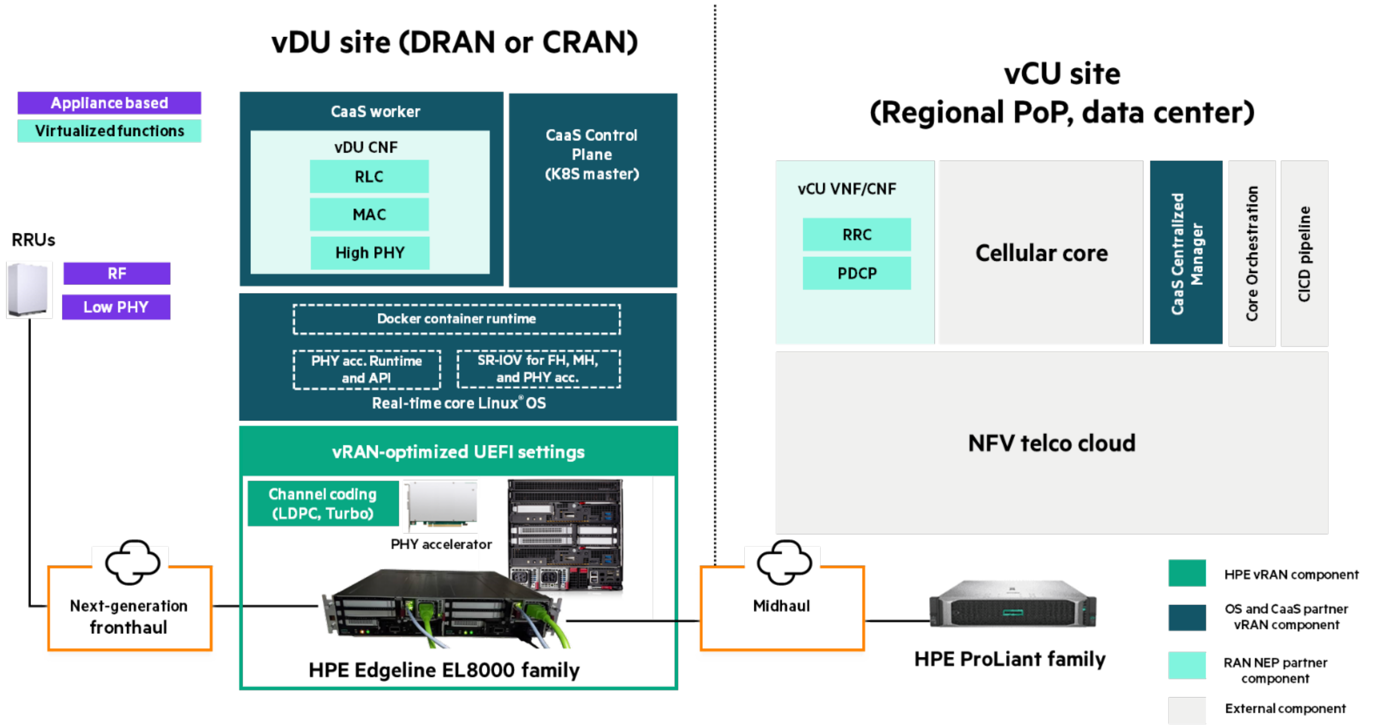

FIGURE 8. vRAN 2.0 reference architecture

vRAN 2.0 reference architecture is based on the following principles:

- RAN real-time functions (PHY, MAC, RLC) are decoupled and implemented as containerized or VM-based vDU workload on general-purpose platform
- RAN non-real-time functions are decoupled and implemented as containerized or VM-based vCU on a general-purpose platform
- vCU is deployed in a centralized facility (for example, Central Office or Regional PoP), and is deployed on top of generic NFV telco cloud
- vDU is deployed on far edge facility (for example, RAN site, far edge Central Office), and is deployed on top of far edge distributed telco cloud
- In a few exception cases, vDU could be deployed in near edge facility, this is when there is extensive fiber laid out and fronthaul latency and throughput is not an issue (~8–50 Gbps per RRU, <100 µsec latency budget)
- vDU and vCU are communicating over standard midhaul interface F1 with mild latency requirements (up to 50 ms latency tolerance)
- RRU implements rest of low PHY functions as well as RF and communicates with vDU via next-generation fronthaul with strict latency requirements (up to 100 µsec latency tolerance)
- Far edge telco cloud is based on distributed cloud-as-a-service (CaaS) platform, which provides abstraction of general-purpose compute resources for the vDU functions (and all further use cases onboarded onto the same far edge cloud), deployment, automation, and management as well as resource orchestration of the far edge—for both container based or VM-based workloads

HPE'S VRAN 2.0 SOLUTION DESCRIPTION

vRAN 2.0 components

The overall HPE's vRAN 2.0 solution follows vRAN 2.0 reference architecture and consists of two main domains:

- **Virtual Distributed Unit (vDU)**

  Functionally, vDU processes real-time portion of baseband (higher PHY per split option 7.2, MAC, and RLC layers) and it is implemented as Containerized Network Function (CNF)/Virtualized Network Function (VNF), fully decoupled from underlying hardware and running on general-purpose compute platform.

  General-purpose compute platform provides necessary compute resources to vRAN CNF and is based on Intel® Xeon® Scalable Platform with additional optional accelerator for LDPC/Turbo offloading.

  In containerized setting, vRAN CNF is implemented as an assembly of bare-metal containers, with abstraction between compute platform and CNF provided by Container Runtime environment (Docker), running on top of Linux environment with real-time core.

  In VM setting, vRAN VNF is implemented as a virtual machine with abstraction between compute platform and VNF provided by hypervisor.

  Local vDU orchestration, vRAN workload lifecycle management and underlying platform management is provided by cloud-as-a-service (CaaS) layer. CaaS is an edge-optimized platform: local instance of CaaS implements all control plane and worker node functionalities sufficient for autonomous operation of edge node, while centralized Master of Masters node (CaaS Centralized Manager) provides single pane of glass for the edge-to-core deployment. In addition to CNF/VNF and platform management, CaaS also implements single-click provisioning of edge instance and integration of vRAN production environment with rest of vRAN CI/CD pipeline.

  vDU exposes next-generation Ethernet-based fronthaul toward Radio Units and midhaul toward vCU. This implementation supports both legacy Radio Units (with external Fronthaul Gateway translating legacy CPRI into next-generation fronthaul) as well as next-generation Radio Units with native support of Ethernet-based fronthaul (for example, per ORAN fronthaul specification). In addition to that, Fronthaul Gateway serves as a synchronization source on fronthaul, ingesting reference signal from primary source (for example, GPS receiver). For the next-generation Radio Units, the role of primary synchronization source will be implemented by a fronthaul NIC.

- **Virtual Centralized Unit (vCU)**

  Functionally, vCU processes non-real-time portion of baseband (RRC and PDCP per 3GPP split option 2), implemented as a fully decoupled CNF/VNF running on general telco cloud platform.

  It exposes midhaul interface toward vDU and backhaul interface toward 5G/4G Cellular Core. As midhaul interface tolerates latencies around 50 ms, vCU can be deployed relatively far from RAN sites, in centralized sites of telco near edge (for example, Next-generation Central Offices or Regional PoPs).

HPE's vRAN 2.0 solution unlocks full potential of vRAN by addressing the following aspects:

- **Full decoupling of software from hardware**

  By following 3GPP split option 7.2, HPE's vRAN 2.0 solution architecture decouples all baseband processing from underlying hardware. That brings fundamental shift into overall RAN economy by allowing to leverage best-of-breed approach to individual solution components, as well as improves scalability and flexibility of RAN.

- **Open and future-proof framework**

  While optimized for vRAN as the first workload, the infrastructure of HPE's vRAN 2.0 solution is an open platform ready to host additional far edge use cases as they appear.

  vRAN 2.0 CaaS platform deployed at the far edge enables rapid innovation by providing general container and VM orchestration environment while integrating with rest of CICD pipeline—together with underlying HPE platform, they form a general-purpose far edge cloud hosting vRAN 2.0 as its first tenant. That future-proof investments into vRAN 2.0 environment as well as brings synergies across domains of telco business.

The above aspects bring tectonic shifts into overall RAN economics—the impact of which is described in the following RAN TCO section.

vRAN 2.0 HPE products

There are three HPE products that are very well suited for the vRAN hardware infrastructure. They are:

- HPE Edgeline EL8000 family of products (EL8000 and EL8000t)
- HPE ProLiant DL360 product
- HPE ProLiant DL380 product

The choice of which product to implement vDU on depends upon the type of site. The two types of sites are described in the following chapter.

vCU can be implemented on either the HPE ProLiant DL360 or DL380 (or any HPE ProLiant family of products currently existing in the customer's data center).

HPE's vRAN 2.0 solution hardware infrastructure unlocks full potential of vRAN by addressing following aspects:

- **Production readiness**

  On infrastructure side, HPE designed Edgeline EL8000 family from ground up to address compute needs of vRAN 2.0 framework, as well as providing a compact and ruggedized form factor for actual production deployment within RAN outdoor cabinets and RAN sheltered sites.

- **Open and future-proof framework**

  HPE Edgeline EL8000 and the HPE ProLiant family of products have best-in-class compute capabilities, which can be tuned to host any type of telco or enterprise workload—for example, UPF/PGW-U functions for network slicing and cellular breakout, AI workloads for MEC, general enterprise workloads.

- **Unique far-edge optimized compute platform**

  HPE Edgeline EL8000 family was developed from scratch addressing unique needs of vRAN and telco far edge cloud: lowest latency with single-socket and direct high bandwidth PCIe connections toward CPU, density, and modularity to address constraint spaces of RAN sites, unprecedented ruggedization allowing the platform to work within extreme temperatures of RAN sites without hit on performance and specialized climate control arrangements.

The following chapters explain in detail the configuration, value proposition, uniqueness, and benefits of each of the above platforms.

vRAN sites

There are two types of sites that the vDU can be implemented on depending upon the operator network's business needs.

Next-generation fronthaul requires sub-millisecond latency from underlying transport network, therefore vDU in most cases would be physically deployed in close proximity to Radio Unit equipment. In practice, this means that vDU platform will be hosted at RAN sites, outdoor RAN cabinets or other facilities at the very edge of cellular network. These RAN sites are classified as two types of sites—the far edge or the near edge site.

- **Far edge** sites are cell sites, requiring special environmental and physical conditions. Compute units need to be thermally suited to operate in these environments without data center cooling. They need to be shallow in depth to fit industry standard outdoor enclosures and indoor cabinets and racks.

  For these sites, HPE EL8000 or EL8000t is best suited, as it meets the environmental and physical requirements mentioned previously. Furthermore, these sites may be DRAN (Distributed RAN) or CRAN (Centralized RAN) deployments. CRAN sites require dense compute at the far edge, and HPE EL8000 is therefore suited to this deployment as opposed to EL8000t for DRAN sites.

The following figure illustrates example of vDU deployment in outdoor RAN cabinet, with HPE Edgeline EL8000t running vDU instance.

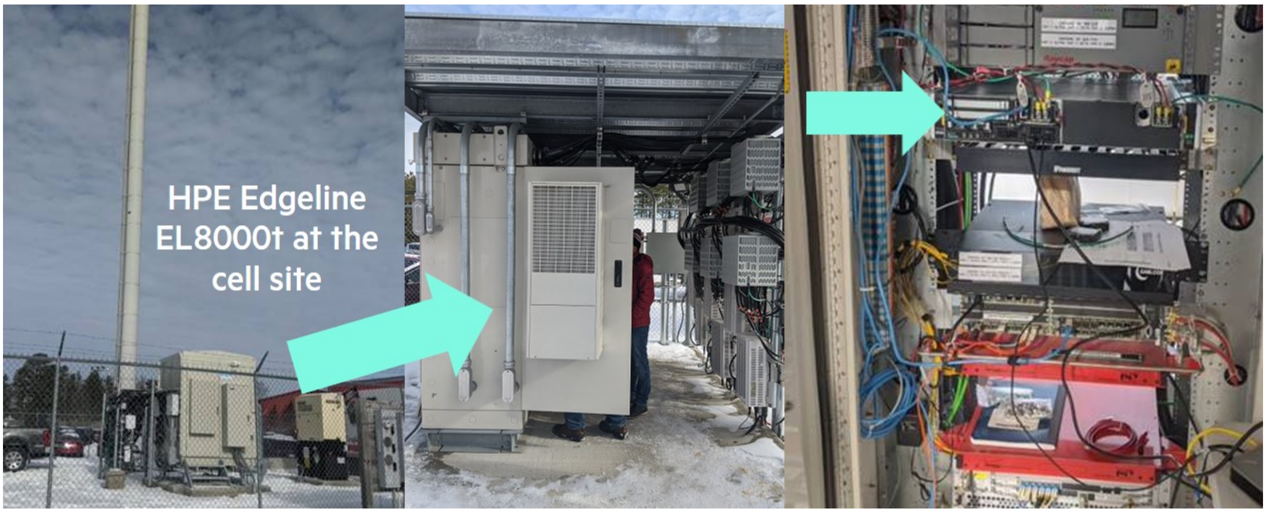

FIGURE 9. HPE Edgeline EL8000t running vDU at RAN site

- **Near edge** sites are typically small CO or regional PoP sites. These CRAN sites aggregate vDUs from ~100 RAN sites at one location. Feasibility of this deployment model is defined by last mile transport ability to provide >10 Tbps with <100 µsec latency between RAN sites and CRAN location. That limitation favors far edge CRAN deployment with operators having limited last mile networks and favors near edge CRAN deployment in operators with lots of dark fiber deployed at last mile.

  For these sites, DL360 is best suited, as it is the best-in-class dense core-to-edge infrastructure for CRAN.

Finally, as midhaul interface tolerates latencies around 50 ms, vCU can be deployed relatively far from RAN sites, in centralized sites of telco near edge (for example, Next-generation Central Offices or Regional PoPs). Depending on realities of those near edge sites, underlying telco cloud might be built on edge-optimized platforms such as HPE Edgeline family (for fairly remote COs where space, cooling, and power are limited) or on core-optimized platforms such as HPE ProLiant family (for centralized facilities).

HPE'S VRAN 2.0 SOLUTION TCO ANALYSIS

Working across multiple RAN global markets, HPE analyzed major components of RAN CAPEX and OPEX, studying potential effects of vRAN 2.0 introduction on overall RAN economy.

It was found that RAN deployments and RAN evolution initiatives across the globe are clustering into four major categories, depending on degree of RAN virtualization and deployment architecture of vRAN:

- **Traditional RAN**

  RAN is implemented on monolithic appliances, coming from a single technology vendor in a given geography. It is fully distributed, with appliance DU residing at RAN sites and processing signal for a given local base station.

- **vRAN 1.0 (PDCP/RRC split)**

  The first step of vRAN, where DU is implemented as a monolithic appliance, while CU is virtualized and deployed as a VNF in telco cloud. vRAN 1.0 is provided by a single technology vendor in a given geography and explained in previous chapters.

- **vRAN 2.0, Distributed mode**

  Following HPE's vRAN 2.0 solution reference architecture, with vDU fully distributed and physically deployed at RAN site, while vCU is centralized in telco cloud. vRAN 2.0 is an open architecture, where best-of-breed approach is applied toward selection of solution components, and explained in previous chapters.

- **vRAN 2.0, Centralized mode**

  Following HPE's vRAN 2.0 solution reference architecture, with vDU centralized in an intermediate near edge facility, aggregating processing for ~10–100 base stations. This architecture is typically evaluated by operators with significant assets in optical fiber at last mile which meets requirements of next-generation fronthaul.

The following bar graph illustrates the 5-year TCO comparison between the alternatives we have just mentioned, and the table after that provides further insight into main TCO contributing factors. All values are presented as a relative percentage points, with 5-year TCO of traditional RAN taken as 100%.

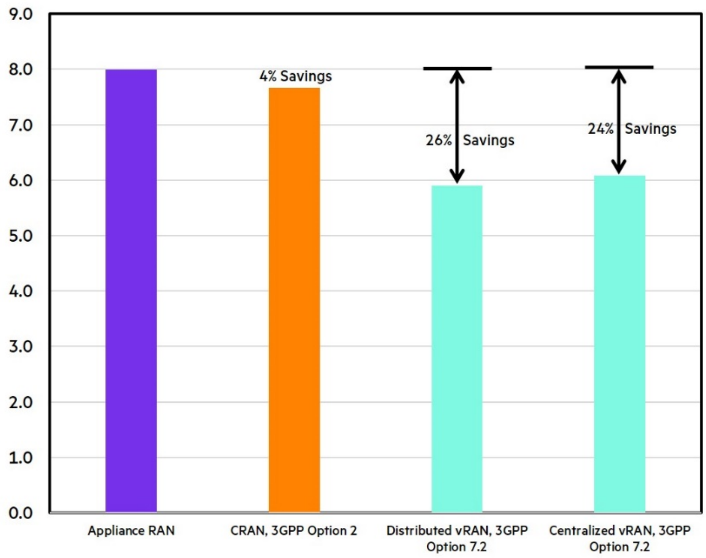

FIGURE 10. 5-year TCO of main vRAN deployment models versus traditional RAN

TABLE 1. TCO over 5 years for RAN deployment alternatives

TCO component	Appliance RAN	CRAN, 3GPP Option 2	Distributed vRAN, 3GPP Option 7.2	Centralized vRAN, 3GPP Option 7.2
CAPEX				
Base station HW	**20%**	**17%**	**13%**	**10%**
Initial deployment	6%	5%	6%	3%
HW upgrade	6%	5%	0%	0%
HW expansion	8%	7%	7%	7%
Base station SW	**6%**	**6%**	**5%**	**6%**
Rollout and tuning	**5%**	**5%**	**3%**	**1%**
Initial deployment and tuning	2.6%	2.6%	2.5%	1.0%
Post-upgrade tuning	1.6%	1.6%	0%	0%
Capacity expansion	0.8%	0.7%	0.7%	0.2%
Transport network upgrade	**0%**	**0%**	**0%**	**31%**
Backhaul upgrade	0%	0%	0%	0%
Fronthaul upgrade	0%	0%	0%	31%
Reuse of general purpose equipment	**0%**	**-1%**	**-13%**	**-10%**
OPEX				
Maintenance and health assurance	**13%**	**13%**	**10%**	**10%**
Site acquisition and rent	**50%**	**50%**	**47%**	**20%**
Electricity	**6%**	**6%**	**8%**	**8%**
Total	**100%**	**96%**	**74%**	**76%**

Following are the key observations of the model:

- vRAN 1.0 does not significantly change RAN TCO. It gains ~4% TCO over 5 years, mainly due to pooling of PDCP/RRC layers and associated savings in the overall processing needs. As PDCP/RRC layers represent insignificant portion of overall baseband processing, the pooling saving is not enough to transform RAN TCO in any significant manner.
- vRAN 2.0 brings significant TCO improvement into RAN economics, up to 26% of reduction across 5 years. Given that RAN TCO constitutes around 50% of overall operator spend on cellular network, vRAN 2.0 significantly improves overall CAPEX and OPEX of a telco operator. With fully decoupled software, vRAN 2.0 infrastructure scales horizontally and does not require vertical hardware upgrades to address evolution of network functionalities, typical for appliance RAN. That optimizes investment into RAN infrastructure, decreases truck rolls and tuning efforts associated with hardware upgrades.
- Another significant factor contributing to vRAN 2.0's superior TCO is the ability to repurpose general-purpose infrastructure for additional use cases as well as leveraging general enterprise tooling for maintenance and lifecycle management of RAN infrastructure, versus in a closed ecosystem where specialized knowledge is required for appliance operation.
- Comparative efficiency of vRAN 2.0 Distributed and Centralized mode are dependent on a balance between cost of last mile transport upgrade to comply with next-generation fronthaul requirements, versus pooling gain and potential site rent savings due to equipment centralization. This balance is unique per each RAN market, as well as each operator in a given market, therefore very sensitive to assumptions. Table 1 is based on somewhat balanced approach toward transport cost versus centralization efficiency, therefore TCO between Distributed and Centralized modes does not differ much.

- Above analysis of vRAN 2.0 benefits does not include potential of optimized multivendor procurement policies which can be implemented by an operator toward technological suppliers, given that vRAN 2.0 enables fully open best-of-breed approach toward components. Potential economic impact of those new policies is marked by many telco operators as the number one driver for vRAN introduction, though this is not easily quantifiable.

Deeper insights into vRAN TCO modelling and present assumptions and key drivers behind the model could be provided upon request.

5G DEPLOYMENT STRATEGIES AND EVOLUTION WITH VRAN

Previous generations of mobile communication have generally addressed a single connectivity service (for example, mobile voice in 2G or smartphone mobile data in 4G), with radio interface and RAN architecture optimized for that service. Practical implication of that was somewhat similar to deployment model of rolling out those technologies by different operators, as the business model was largely similar across markets.

In contrast to that, 5G addresses a wide range of connectivity services, from low-power wide-area connectivity, to enhanced mobile broadband and low-latency mission-critical services. Therefore, 5G NR introduces unprecedented flexibility into air interface and overall RAN architecture, which allows an operator to flexibly tune 5G RAN deployment strategy around business model they choose to focus on. That implies diverse 5G RAN deployment strategies across markets and even between operators on a given market, leading to a situation where different operators start 5G NR rollout with vastly different network architectures.

While acknowledging that 5G NR deployment comes in many forms, NR network architecture is mainly defined by spectrum bands targeted for NR. Following models are evaluated by operators, as illustrated in the following figure.

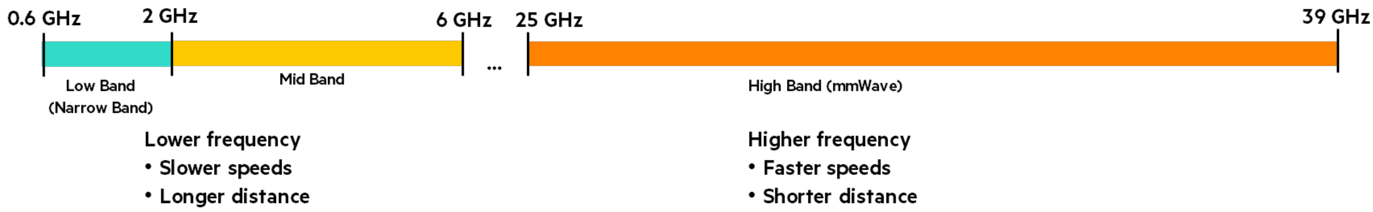

FIGURE 11. 5G NR target spectrum bands

Low Band NR

This deployment targets rollout of 5G NR in spectrum bands between 600 MHz and 2 GHz, with carrier bandwidth up to 20 MHz. Low Band NR is characterized by the following:

- Better propagation characteristics of radio signal in Low Bands for outdoor as well as superior outdoor-to-indoor signal penetration
- Spectrum availability in Low Band is limited, as low frequency bands naturally do not have much spectrum resources available and they are also leveraged by legacy 2G, 3G, and 4G services. Therefore launching 5G NR in Low Band in many cases will not lead to dramatic improvements of connectivity performance for end subscribers
- Low Band frequency does not in general allow for higher order spatial diversity and spatial multiplexing of radio signal, as it will require unpractically large size of antenna arrays. MIMO 4x4 diversity in downlink and uplink is the most commonly accepted technique

Those characteristics make Low Band particularly well suited for coverage-oriented macro network. Starting with Low Band NR is a popular choice for an operator to quickly secure wide national coverage with 5G service, claiming and promoting technological leadership in a market.

Characteristics of Low Band have following implications on RAN architecture:

- Relatively low baseband processing needs across fully distributed macro outdoor network, with a typical total antenna bandwidth per site at ~800–900 MHz
- Relatively low fronthaul throughput (<50 Gbps); low midhaul/backhaul throughput (1 Gbps)
- Macro outdoor deployment of RAN implies DU baseband processing equipment to be deployed at RAN site (for example, outdoor RAN cabinets, indoor RAN sites)

The following figure illustrates how vRAN 2.0 architecture for vDU addresses the Low Band 5G NR deployment, as well as provides some capacity estimates for that deployment.

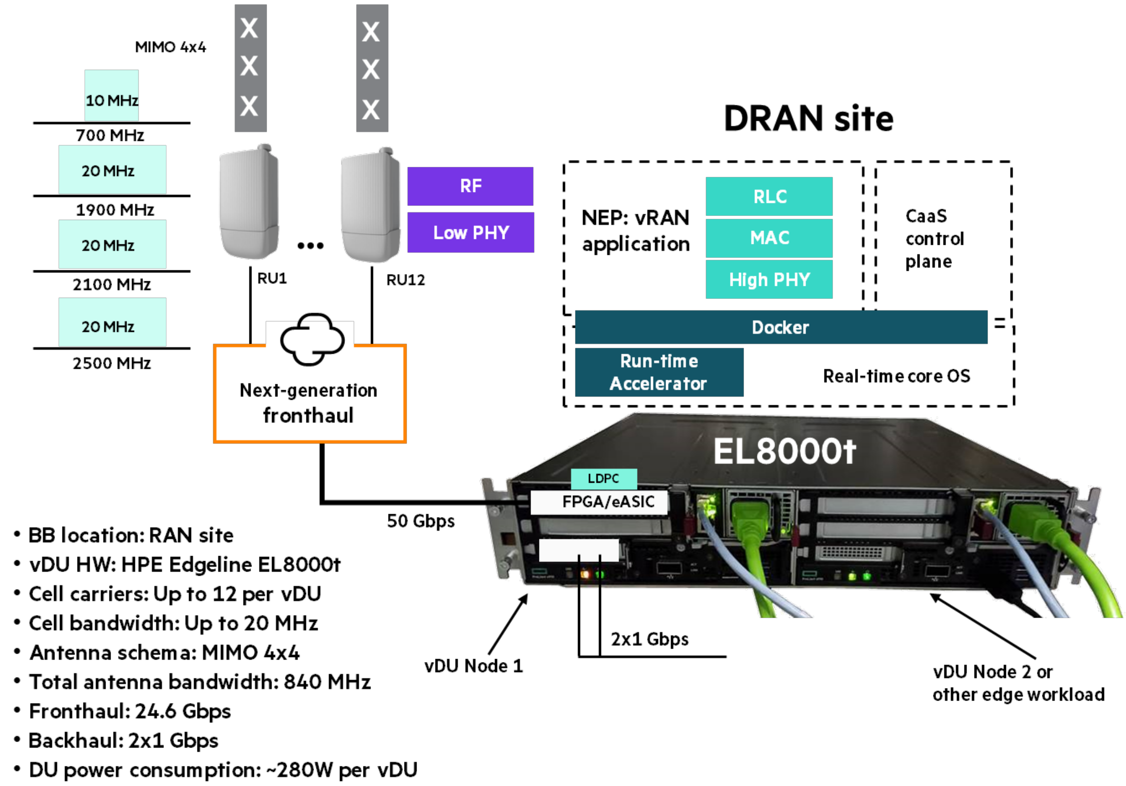

FIGURE 12. vRAN 2.0 for Low Band 5G NR deployment

Mid Band NR

This deployment targets rollout of 5G NR in spectrum bands between 2 GHz and 6 GHz, with carrier bandwidth up to 100 MHz. Mid Band NR is characterized by following:

- Average radio signal propagation characteristics, limited outdoor-to-indoor signal penetration
- Good spectrum availability as well as underutilized spectrum, as majority of spectrum bands within Mid Band are not used by legacy mobile technologies
- Higher order spatial diversity schemas are common, with 8x4 MIMO or 16x8 MIMO being commonly targeted for Mid Band. Massive MIMO techniques are being evaluated for higher frequencies of Mid Band

Those characteristics make Mid Band particularly well suited for capacity-oriented macro and small cells network. A popular use case for an operator is to offload existing LTE network in traffic hotspots and secure market leadership by providing unprecedented performance for subscribers in densest areas. Alternatively, Mid Band NR deployment can be a next step in 5G evolution after Low Band NR is deployed, by providing capacity on top of coverage where it is needed.

Addressing those deployment models, vRAN 2.0 architecture can be flexibly applied to either introduce Mid Band NR small cells to existing Low Band NR vDU, or aggregate processing of a number of Mid Band macro cells and small cells on vDU in centralized location as processing pooling starts to promise significant improvement.

The following figures illustrate how vRAN 2.0 architecture addresses those deployment models.

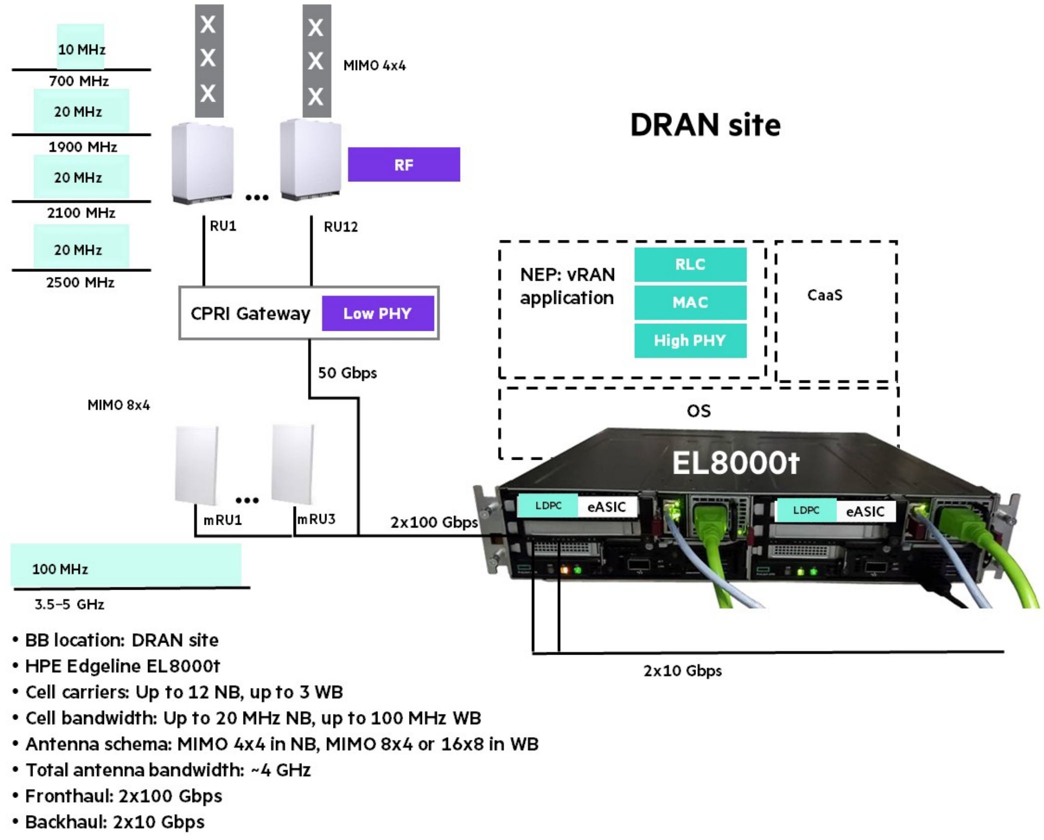

FIGURE 13. Adding Mid Band NR small cells to existing Low Band deployment

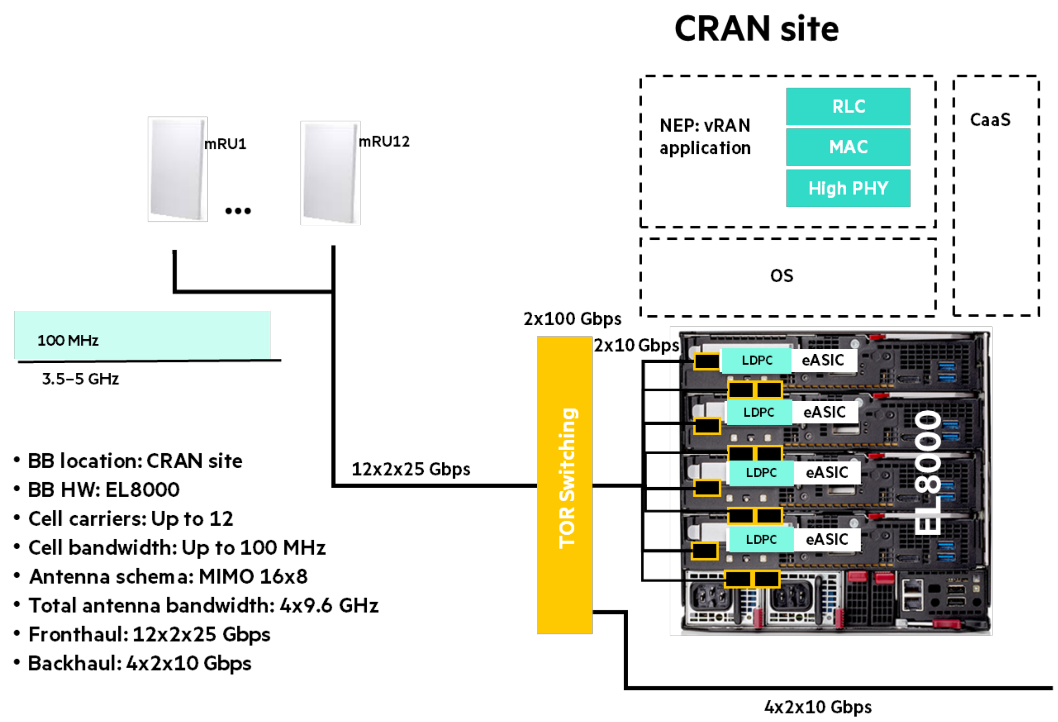

FIGURE 14. Introducing stand-alone Mid Band NR small cells with baseband centralization

mmWave

This deployment targets rollout of 5G NR in spectrum bands between 25 GHz and 39 GHz, with carrier bandwidth up to 800 MHz. mmWave NR is characterized by following:

- Extremely weak signal propagation characteristics, which means utilizing the mmWave spectrum even for small cell NR is unpractical without beamforming and massive MIMO techniques
- Higher frequencies make massive antenna arrays feasible—with massive MIMO techniques and active beamforming running on 64x64 arrays
- Extremely high baseband processing requirements (combination of wide spectrum and advanced antenna techniques)

Feasibility of vRAN 2.0 architecture based on 3GPP split option 7.2 for mmWave is under study. With total antenna bandwidth reaching dozens of GHz, baseband processing and potential fronthaul requirements explode, likely requiring new architectural approaches and next-generation element base for vRAN infrastructure added to vRAN 2.0 road map.

As seen from previous examples, vRAN 2.0 virtualization enables operators to experiment with deployment model and flexibly adjust network architecture and deployment strategy as 5G business models matures and evolves.

HOW HPE CAN FACILITATE 5G VRAN ROLLOUT

HPE Edgeline EL8000 family: Telco edge product

HPE designed Edgeline EL8000 is an edge platform that is specifically designed from ground up to address unique challenges of telco edge. HPE Edgeline EL8000 was built using the following principles:

- **Best-in-class performance with the smallest footprint**

  Top-of-line Intel Xeon Scalable processors, most powerful GPUs, and FPGA accelerators in a quarter of standard server footprint
- **Low-latency design for real-time processing**

  HPE EL8000 family is based on single-socket design and flexible firmware, specifically optimized to minimize latency of communication between server components (for example, between 5G accelerators and CPU)
- **Uncompromised operation under extreme conditions**

  Compliant with and going beyond NEBS Level 3 and GR-3108 Class 1 requirements with -5°C to 55°C continuous operation, operating fully loaded (with accelerators and such) and without hit on performance
- **Optimized designed for easy and reliable operations**

  Modular hot-pluggable architecture with built-in redundancy of all mission-critical elements
- **Industry-leading built-in security**

  Based on iLO 5 technology, HPE EL8000 platform delivers Silicon Root of Trust and plethora of other security features essential at the edge of the network

Following sections provide in-depth overview of product options and features of HPE Edgeline EL8000 family.

HPE Edgeline EL8000t, compact 2U chassis for DRAN deployment

HPE Edgeline EL8000t is a compact platform featuring up to two compute cartridges in a shallow depth chassis. The platform is optimized for extremely space limited environments of RAN outdoor cabinets and compact indoor racks of RAN sites.

That design mostly targets fully distributed deployment of vRAN at macro RAN sites, sometimes called DRAN (as explained in the previous section).

HPE Edgeline EL8000t allows deployment of up to 2 e910t compute cartridges in a single chassis, each featuring single-socket top-line Intel Xeon Scalable processor, built-in high-throughput networking and optional additional accelerators (for example, FPGA, GPU).

The following figure illustrates the HPE EL8000t.

FIGURE 15. HPE EL8000t chassis front view

HPE Edgeline EL8000, dense 5U chassis for CRAN deployment

HPE Edgeline EL8000 is an extremely dense compute platform, featuring up to 4 single-socket servers in a quarter of a standard server footprint.

This design mostly targets high-density deployment of vRAN with centralized baseband processing, sometimes called CRAN (please refer to the previous chapters for details).

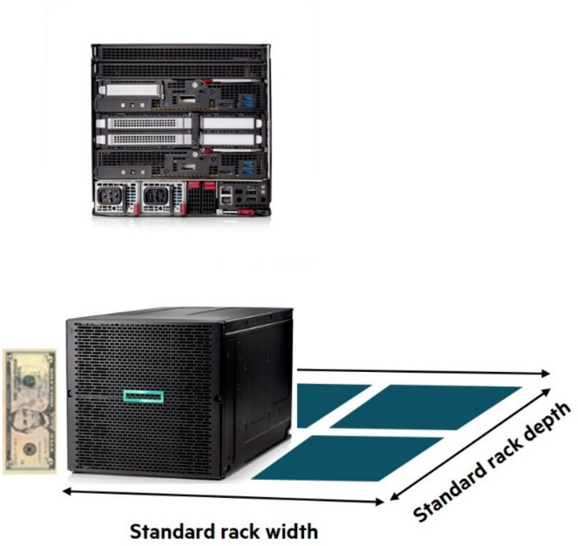

FIGURE 16. HPE EL8000 chassis occupies quarter of a standard rack

HPE Edgeline EL8000 allows deployment of up to four e910 compute cartridges in a single chassis, each featuring single-socket top-of-line Intel Xeon Scalable processor, built-in high-throughput networking and optional additional accelerators (for example, FPGA, GPU).

Why HPE Edgeline EL8000 family for vRAN 2.0

HPE's vRAN 2.0 solution offering is the industry unique field-proven solution for virtualizing full stack of Radio Access Network and fully unlocking business value of vRAN to cellular operators. HPE brings to the table a number of technology, operational and business advantages differentiating the vRAN 2.0 offering.

Flexibility and adaptability

HPE's vRAN 2.0 solution reference architecture easily adapts to any vRAN deployment scenario of 5G NR—from fully distributed macro outdoor in narrow band to distributed and centralized deployment of small cells and macro outdoor in Mid Band.

That flexibility is possible by leveraging industry unique HPE Edgeline EL8000 family and HPE ProLiant DL360 as underlying universal compute platforms. Flexibility and modularity of the HPE EL8000 family make them fit into all edge environments from remote RAN outdoor cabinets to indoor sheltered shallow depth racks, operate without specialized climate control arrangements as well as provide unprecedented density of compute resources required for vRAN 2.0 baseband processing. Flexibility of the DL360 products make them the best-in-class CO/PoP near edge solution platform.

Solutions versus component provider

HPE is working with several NEP, OS, and other partners to provide the best HW infrastructure solution for vRAN 2.0. HPE is not going to be an OEM of vRAN 2.0 solution, meaning we will not be integrating the vDU, vCU, and RU to provide a complete solution. However, we are working very closely with our partners to not only collaborate but lead the RAN HW infrastructure that is best suited for vRAN 2.0 and beyond.

While designing our servers, we work with both hardware and software partners to determine the best-in-class components considering at all times the vRAN 2.0 requirements. The following are some examples, and by no means exhaustive.

- Using the best CPU, NIC, and HW acceleration techniques to support all aspects of vRAN processing
- Selecting components and designing servers that not only meet the performance but are optimized for power consumption
- Conducting thermal tests with all components and in some cases the appropriate vRAN workload running on the server
- Providing software and firmware capabilities that allow for easy deployments and upgrades in the field
- Influencing the road map of component providers/partners in anticipation of future needs

Staging and management

Open and fully disaggregated, the vRAN 2.0 architecture is a collection of a number of inter-dependent hardware, firmware, and software components. However, during field rollout of vRAN 2.0 the scope of the activities which could be executed at the RAN site at a Day 0 are limited because of a number of factors:

- **Limited skill set of site personnel**: Site technician may have enough skills for basic scripted provisioning actions (such as assigning the IP address, loading a certificate, and so on), however, they may not be able to install and configure the entire ecosystem of software and firmware components locally.
- **Limited throughput for O&M traffic on midhaul network**: While remote loading of software and firmware components via centralized automation engine is feasible, the last mile link in many cases leaves very limited room for Operations and Maintenance (O&M) and Out Of Band (OOB) traffic—and in peak hours might not have more than ~100 Kbps bandwidth.
- **Lack of physical security at the site** limits the possibility for automated network-based node discovery for the initial L2/L3 connectivity establishment, using technologies like Stateless Address Autoconfiguration (SLAAC), Dynamic Host Configuration Protocol (DHCP), automated Preboot Execution Environment (PXE) boot, and so on. Static network configuration is the preferred way of network provisioning.

Helping navigate around real world limitations, HPE offers tailored services to enable smooth vRAN 2.0 rollout and operations via staging environments. The staging services from HPE can be customized to a particular deployment and operational practices of a cellular operator and can come in multiple flavors:

- **Full initial preprovisioning** of systems, integrating telco operator site planning and RAN planning information into staging process (for example, IP addresses, OS, and application configuration files, image repositories and CICD pipeline)
- **Partial preprovisioning** with loading the system with common images and components shared across all sites (for example, vRAN custom UEFI settings, OS image, CaaS, and application images), followed up by remote provisioning of site-specific information (IP addresses, configuration files for OS, CaaS, and application)

In addition to custom services around initial provisioning, HPE has the capability to support Day 1 operations—including support services tailored for far edge deployment and edge spare parts management.

Security

Traditional centralized deployments of general-purpose infrastructure supporting telco functions are located in relatively secure physical locations (data centers, regional PoPs), which provide relatively safe trusted zones for infrastructure they host, as well as tooling around security (DMZs, firewalls, secure anchors of trust, and so on).

Taking general-purpose infrastructure to the physically non-secure far edge location (for example, RAN sites) significantly expands the risk, with hardware and software stack potentially directly exposed to a malicious agent.

So far, traditional appliances of telco edge (for example, RAN appliances) were relatively robust to attacks on its hardware and software stacks, as they were closed, single purpose and proprietary. Therefore, many attacks on cellular network last mile are limited to attempts of hijacking radio interface (for example, jamming control channels on air interface, intercepting subscriber attach procedures, presenting fake base stations as legitimate ones, and so on).

With general-purpose compute hosting open software stack deployed at the RAN site, the attacker has unprecedented opportunity to attack the infrastructure itself (the following figure illustrates some of the examples).

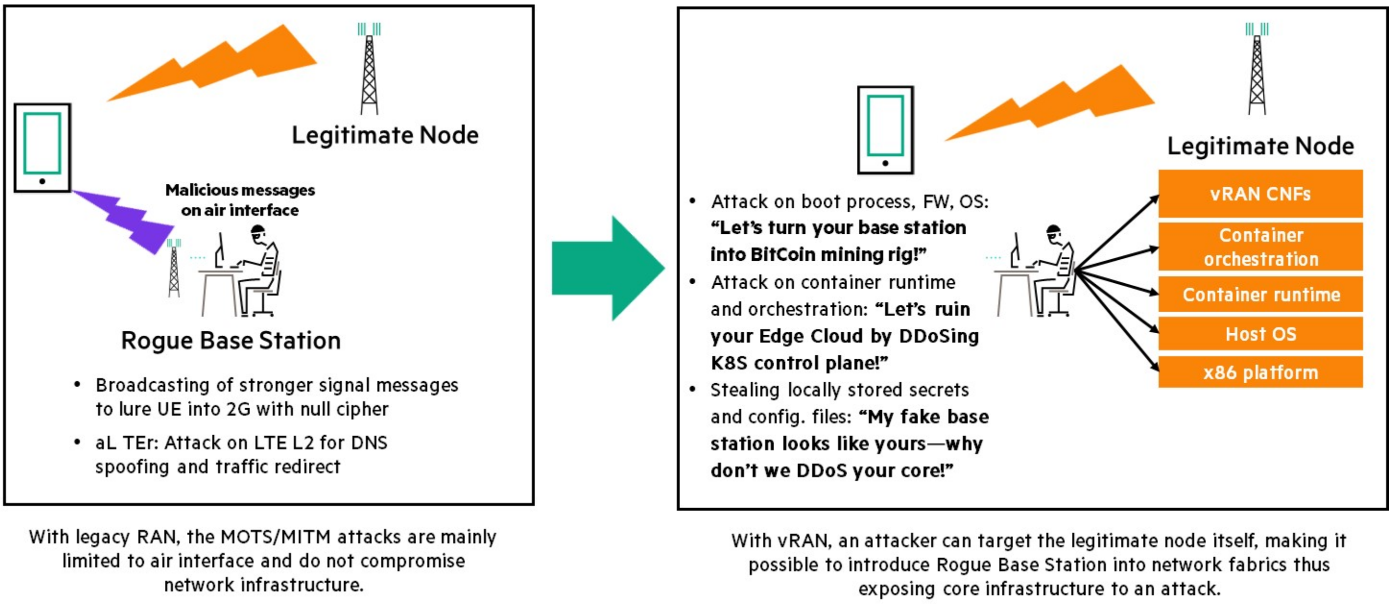

FIGURE 17. New attack vectors with vRAN

Addressing those challenges, HPE brought in number of key security features into design of HPE Edgeline and HPE ProLiant products.

- **iLO 5 Silicon Root of Trust**: Deployed at a non-secure location; the infrastructure can't be considered trusted by default, therefore software stack deployed at that infrastructure can't anchor its security mechanism to such infrastructure. HPE Silicon Root of Trust is a unique HPE technology, which guarantees authenticity of HPE hardware and firmware stack throughout the lifecycle of the system: From factory, to supply chain, to last mile logistics, to production site to decommissioning as shown in the following figure.

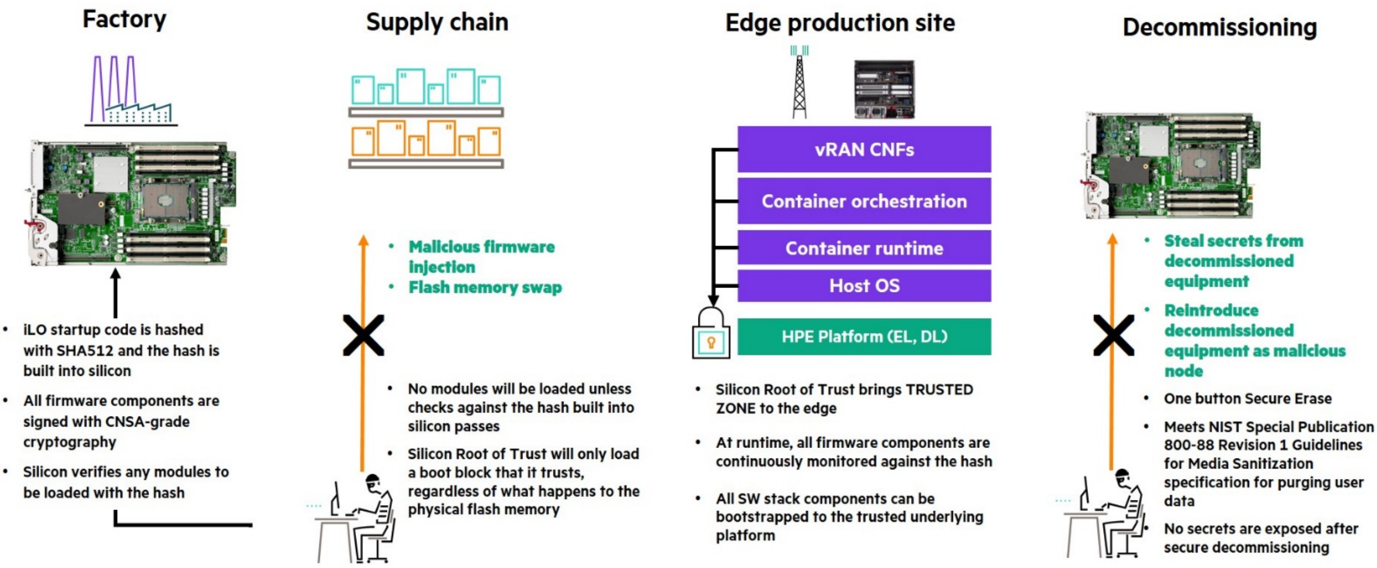

FIGURE 18. HPE Silicon Root of Trust helps to secure entire lifecycle of far edge system

- **Secure Erase:** Decommissioning of legacy equipment is in many cases an overlooked security loophole. Taking infrastructure nodes off air and dumping them at non-secure warehouse or even landfill provides an attacker a straightforward way to acquire intimate knowledge about security framework as well as critical secrets. General-purpose infrastructure and software solution stack deployed at far edge site will store many secrets—security tokens, passwords, as well as less critical information about security practices which can be used to engineer future attacks.

  To make sure that the system is decommissioned safely (or in case of emergency action on suspicious node), HPE introduces remote capability to Secure Erase all information and settings from the system. This unique capability complies with NIST Special Publication 800-88 Rev. 1 Guidelines for Media Sanitization, and makes sure no secrets are leaked through unsafe decommissioning practices.
- **UEFI security:** HPE products support a list of UEFI security features, such as TPM 2.0, secure boot, ability to enable/disable physical interfaces to limit unauthorized access at far edge site, TLS (https) boot, and so on.

Thermal advantages

Many operators consider environmental requirements of NEBS Level 3 as a reference for what the far edge system should be able to tolerate.

NEBS Level 3 is a good place to start evaluating edge equipment capabilities, but it does not necessarily address the entire range of environments at far edge on practice. At RAN sites, the equipment needs to be ready for deployment in extremely tight spaces with limited airflow and operational temperatures frequently go beyond thresholds defined by NEBS Level 3. Keeping requirements of real RAN sites in mind, HPE designed Edgeline EL8000 systems to withstand tougher-than-NEBS environments without affecting performance or long-term reliability.

The following table illustrates key area where HPE Edgeline EL8000 goes above and beyond formal compliance with NEBS Level 3.

TABLE 2. HPE Edgeline EL8000 and EL8000t operating temperatures

Criteria	NEBS Level 3 requirement	HPE Edgeline EL8000/EL8000t
Normal Operating Temperature, min °C	5°C	-5°C
Normal Operating Temperature, max °C	40°C	55°C
Short-term Operating Temperature, min °C	-5°C	-5°C
Short-term Operating Temperature, max °C	55°C	65°C
Cold start, °C	0°C	-5°C
Continuous Operating Temp (max)	Not defined	55°C

As it is seen from the above, HPE Edgeline EL8000 family is able to continuously operate up to 55°C ambient temperature without affecting performance or longevity of components, which represents 15°C improvement compared to NEBS Level 3 requirement. Together with 10°C improvement on minimum operational temperature, that represents 50% improvement of normal temperature operational range versus typical NEBS-compliant equipment.

With the widest tolerance to high and low ambient temperatures on a market, HPE Edgeline EL8000 family makes it possible to deploy vRAN 2.0 baseband processing where it is required (RAN outdoor cabinets, rooftop, attic, or basement RAN sites) without any need to upgrade HVAC equipment eliminating costs associated with it.

Compactness and density

For far edge locations, HPE Edgeline EL8000 family has two form-factors: EL8000t is the most compact vDU platform particularly suited for far edge DRAN deployment with vDU deployed at RAN sites, and HPE EL8000 is the densest vDU platform addressing CRAN deployments with maximum compute-per-square-foot.

The following figure illustrates typical DRAN sites—a deployment model addressed by HPE EL8000t.

Indoor 2 post rack, shallow depth

Indoor shallow depth, no rear access

Outdoor shallow depth, no rear access. Mounting only on front plane (no rear supports)

FIGURE 19. RAN site environments where vRAN 2.0 baseband will be deployed

The spaces are characterized by the following mandatory requirements, addressed by HPE EL8000t.

- **Shallow depth of the cabinets**, with back-to-front space measured at ~23". With ~3" usually reserved for safe cable management and ~3" at the back for air exhaust, it leaves no more than 17" for equipment.

  HPE Edgeline EL8000t and EL8000 are just 16.9" deep.

- **Limited mechanical support**, with best case scenario of 2 post rack (see above, indoor rack) and worst case scenario of just a frontal plane available for mounting without any rear support column and no rear access inside cabinet in general (see above, outdoor shallow depth cabinet).

  HPE Edgeline EL8000t was specifically designed with optimized weight and center of gravity as well as flexible mounting kit to mount reliably into cabinets with 2 post or only frontal plane while passing NEBS Level 3 seismic tests in that assembly.

- **Uncontrolled and limited airflow**: Majority of industry standard outdoor cabinets do not enforce controlled airflow between frontal and rear compartments of the cabinet (for example, no dummy plates in unoccupied rack space to control separation of cold air at front and hot air at the back of the cabinet). Together with significant density of equipment hosted in those cabinets, that requires unparalleled tolerance to high operational temperatures of x86 equipment introduced into those cabinets, going well beyond NEBS Level 3 in some cases.

  HPE Edgeline EL8000t and EL8000 were specifically designed to fit those environments thermally, tolerating anywhere between -5°C and 55°C continuously without hit on performance or longevity of the system, as explained in the previous chapter.

Another typical deployment environment is CRAN, where vDU locations are somewhat more generous with space and power—edge COs and larger RAN sites. While still requiring shallow depth systems and ability to operate under NEBS Level 3 ambient temperatures, the main criteria for this model of deployment is density of CRAN system.

Addressing the need for the most dense vRAN vDU solution, HPE Edgeline EL8000—a half-depth and half-width chassis, hosting up to 4 separate nodes in a quarter of a standard server footprint.

Putting two HPE EL8000 systems side by side in a standard rack would allow to bring up to 8 vDU nodes in 5U space, while fitting into 2-post racks or shallow depth cabinet, as illustrated by the following Figure 20.

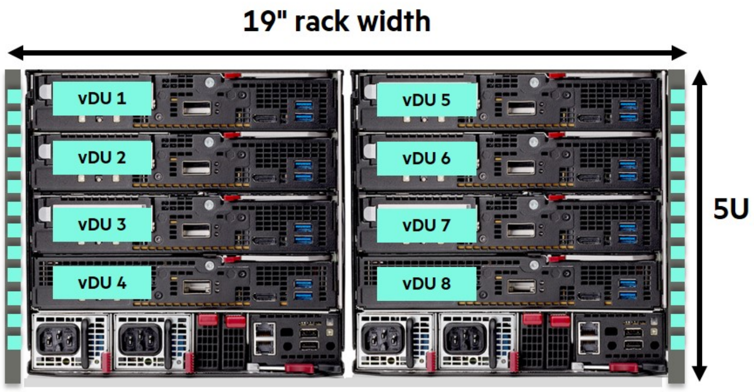

FIGURE 20. 8–10 vDUs in 5U space, based on two HPE EL8000 chassis or five DL360 servers

HPE ProLiant DL360 for vRAN 2.0

The HPE ProLiant DL360 Gen10 is a highly flexible, performance-driven two-socket server that features Intel®'s Scalable family of processors in a densely packed 1U form factor. The HPE ProLiant DL360 Gen10 server supports the Intel Xeon Scalable processor family with up to 28 cores, plus 2933 MT/s and DDR4 RAM up to 3.0 TB.

Additionally, the DL360 Gen10 features a verity of storage configuration options such as a 4 LFF drive, 8 SFF + Optical Disk Drive, 8 SFF + 2 SFF, and 10 SFF NVMe.

This design mostly targets high-density deployment of vRAN with centralized baseband processing, sometimes called CRAN (please refer to the previous chapters for details).

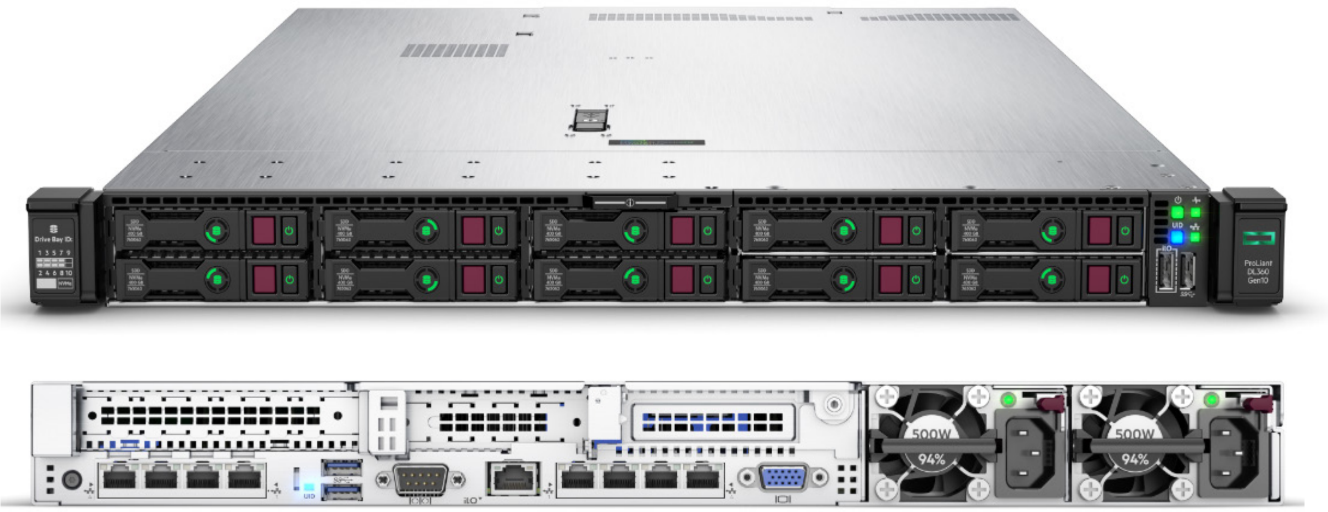

FIGURE 21. 8–10 DL360 Gen10 chassis occupies a single standard rack unit

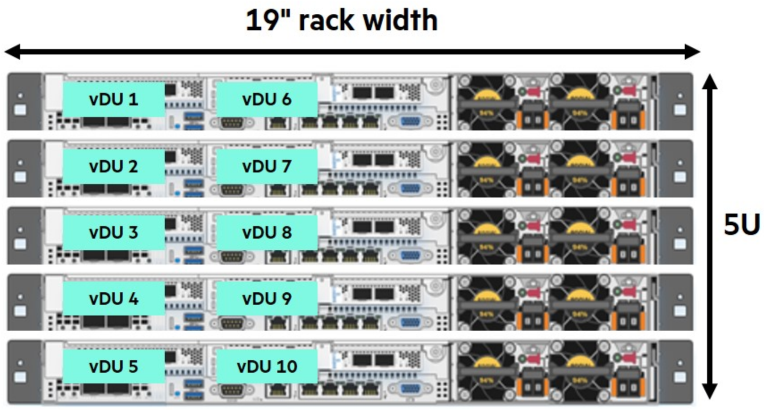

FIGURE 22. 10 vDUs in 5U space, based on five DL360 servers

The secure 2P HPE ProLiant DL360 Gen10 server delivers exceptional flexibility and unmatched expandability packed in a dense 1U rack design with high-throughput networking and optional additional accelerators (for example, FPGA, GPU).

While addressing two distinct deployment models of CRAN and DRAN, HPE Edgeline and ProLiant products share the same underlying platform, iLO 5 management, firmware/UEFI platform and hence are deployed, integrated, and managed in a consistent way, allowing to build a homogeneous far and near edge cloud across CRAN and DRAN locations.

CONCLUSION AND SUMMARY

It is inevitable that telco operators will drive a more open and flexible RAN architecture by disaggregating the monolithic baseband processing of a base station into separate distinct layers. It is just a matter of timing, some operators are ahead of others based on their market economics. Irrespective of when this happens, HPE is ready with the appropriate TCO analysis, solution analysis with partners, and ultimately the right HW infrastructure to help them migrate to this new architecture.

APPENDIX A: ACRONYMS

Acronym	Definition	Acronym	Definition
3GPP	3rd Generation Partnership Project	NEBS	Network Equipment Building Systems
5G NR	5G New Radio	NEP	Network Equipment Provider
AI	Artificial Intelligence	NFV	Network Functions Virtualization
API	Application Program Interface	NgCO	Next-generation Central Office
AR	Augmented Reality	NGMN	Next-Generation Mobile Network
ASIC	Application Specific Integrated Circuit	NIC	Network Interface Card
BBU	Baseband Unit	O&M	Operations and Management
CAPEX	Capital Expenditure	OPEX	Operating Expenses
CaaS	Cloud as a service	ORAN	Open RAN
CICD	Continuous Integration/Continuous Development	PDCP	Packet Data Convergence Protocol
CNF	Containerized Network Function	PDN	Packet Data Network
CPRI	Common Public Radio Interface	PoP	Point of Presence
CPU	Central Processing Unit	PXE	Preboot Execution Environment
CRAN	Centralized RAN	QoE	Quality of Experience
CU	Centralized Unit	QoS	Quality of Service
DDP	Dynamic Device Personalization	RAN	Radio Access Network
DHCP	Dynamic Host Configuration Protocol	RE	Resource Element
DMZ	De-Militarized Zone	RF	Radio Frequency
DPDK	Data Plane Development Kit	RLC	Radio Link Control
DRAN	Distributed RAN	RRC	Radio Resource Control
DU	Distributed Unit	RRU	Remote Radio Unit
eCPRI	Evolved Common Public Radio Interface	RU	Radio Unit
EL	Edgeline	SGW	Serving Gateway
FPGA	Field Programmable Gate Array	SLAAC	Stateless Address Autoconfiguration
GPP	General Purpose Processing	SoC	System on a Chip
GPU	Graphics Processing Unit	SW	Software
HVAC	Heating, Ventilation, and Air Conditioning	TCO	Total Cost of Ownership
HW	Hardware	TLS	Transport Layer Security
MIMO	Multiple Input Multiple Output		

APPENDIX B: REFERENCES

For more information on relevant products, please use the following references:

- hpe.com/info/edgeline
- hpe.com/us/en/servers/proliant-servers.html
- hpe.com/us/en/solutions/infrastructure-security.html

LEARN MORE AT

hpe.com/dsp/infrastructure

Get updates

Hewlett Packard Enterprise

a50001874ENW, July 2020, Rev. 1