White Paper

Telecommunications
Network Functions Virtualization (NFV)

accelerated by intel

Open RAN Cloud Deployment with Amazon EKS Anywhere on Intel®-Based Dell PowerEdge Servers

End-to-end Open RAN deployment using Amazon Elastic Kubernetes Service (EKS) Anywhere on Intel® and Dell hardware can reduce costs and improve performance.

Authors

Adnan Boustany
Director—Open RAN Infrastructure Customer Programs, Intel Corporation

Akram Hassanien
Manager—Open RAN Infrastructure Customer Programs, Intel Corporation

Vivek Bhandula
Solutions Architect—Telco Edge Cloud, Intel Corporation

Yousef Hawwar
Principal Telecom Industry Specialist, AWS

Ramu Akula
Principal Telecom Industry Specialist, AWS

Cheng Liu
Principal Solution Architect, 5G RAN, AWS

Gokul Chandra
Senior Solution Architect, Containers, AWS

Mandeep Singh
Director—RAN Product Management, Mavenir

Piyush Srivastava
Senior RAN Specialist NPI, Mavenir

Nikunj Vaidya
Director, Software Engineering—OTEL, Dell Technologies

Rajendra Chayapathi
Technical Staff—OTEL, Dell Technologies

Executive overview

Network functions virtualization (NFV) is a vital part of communications service providers' (CSPs') efforts to reduce costs, improve performance, and simplify operations. These efforts extend to virtualizing radio-access networks (RANs), which enables running vital baseband functions as software. However, while standardizing interfaces and protocols for virtualizing RANs—Open RAN—can provide benefits that can range from lower hardware costs to expanding vendor options, the telecom industry needs an open ecosystem for RAN solutions in order to increase widespread adoption of Open RAN.

CSPs must navigate several selection decisions on their journey to Open RAN deployment in order to arrive at the right combination of pre-validated and pre-integrated platforms, infrastructure, and applications. Intel, Amazon Web Services (AWS), Dell Technologies, and Mavenir came together to help address some of this absence of guidance in the Open RAN ecosystem for these decisions (Figure 1). These companies partnered to create and validate a quick-to-certify solution to run Open RAN on the Amazon EKS Anywhere container-as-a-service (CaaS) solution. This solution incorporates industry-leading Open RAN software from Mavenir and Intel, paired with Intel and Dell hardware that provides a model Open RAN solution that can be customized to CSP-specific requirements and objectives. Example use cases for Open RAN solutions like this one range from far-edge deployments for telecom companies to edge and private network deployments for organizations in other industries.

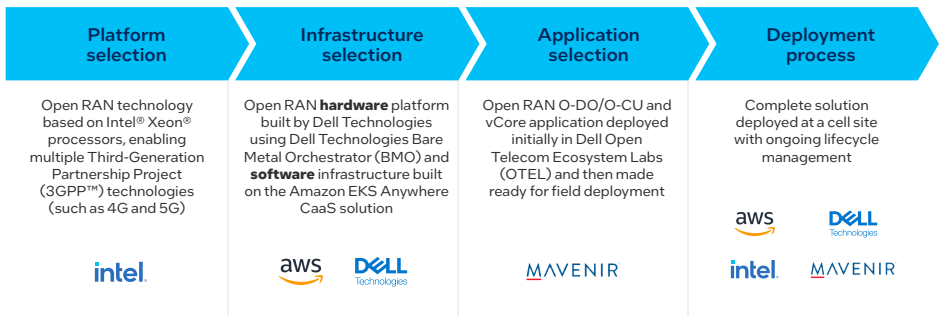

Figure 1. CSP journey to Open RAN deployment

Changes in the telco industry

Virtualizing RANs enables telco operators to run vital baseband functions such as software. Virtual RANs are part of the wider NFV efforts that many telecommunications organizations have been implementing. The virtualization of RANs enables CSPs to run RAN functions on commercial off-the-shelf (COTS) servers rather than on specialized RAN hardware. A related term, "Open RAN," refers to the standardization and openness of interfaces and protocols in a RAN. Such open standardization facilitates interoperability among different vendors' equipment, which can reduce vendor lock-in and help organizations deploy the combination of hardware and software that best meets their needs.

A Senza Fili study commissioned by Mavenir indicates that moving to Open RANs can help telco organizations save as much as 37 percent in total cost of ownership (TCO) over five years.[1] Moreover, shifting to Open RANs can free telco organizations from the proprietary hardware of any one vendor, which can help eliminate vendor lock-in.

These benefits are driving telecom organizations to move toward Open RANs. However, the industry needs to create an open ecosystem for RAN solutions in order to drive wider adoption of Open RANs. It also needs to find ways to simplify and validate those RAN solutions on COTS hardware to create a packaged offering for faster adoption.

These RAN-on-COTS platforms consist of many components and come with many configurations and customizations, empowering CSPs to tailor their networks and meet their customers' expectations. However, with this degree of freedom, time to market with a validated solution becomes a challenge for CSPs as many components need cross-validation among solution providers. This is why Intel and AWS are working together to provide a solution that enables CSPs to get to market earlier and faster while providing quality performance from day 1. This solution provides a **pre-validated** and **pre-integrated** stack that entails a general-purpose platform (GPP), a CaaS model, and an Open RAN application. This approach is expected to simplify CSP day-0, day-1, and day-2 network launch and operations with a RAN solution based on the Amazon EKS Anywhere service. Figure 2 demonstrates how the different components from Dell Technologies, AWS, Mavenir, and Intel deployed in the lab interconnect to realize a certified Amazon EKS Anywhere Open RAN offering for CSPs.

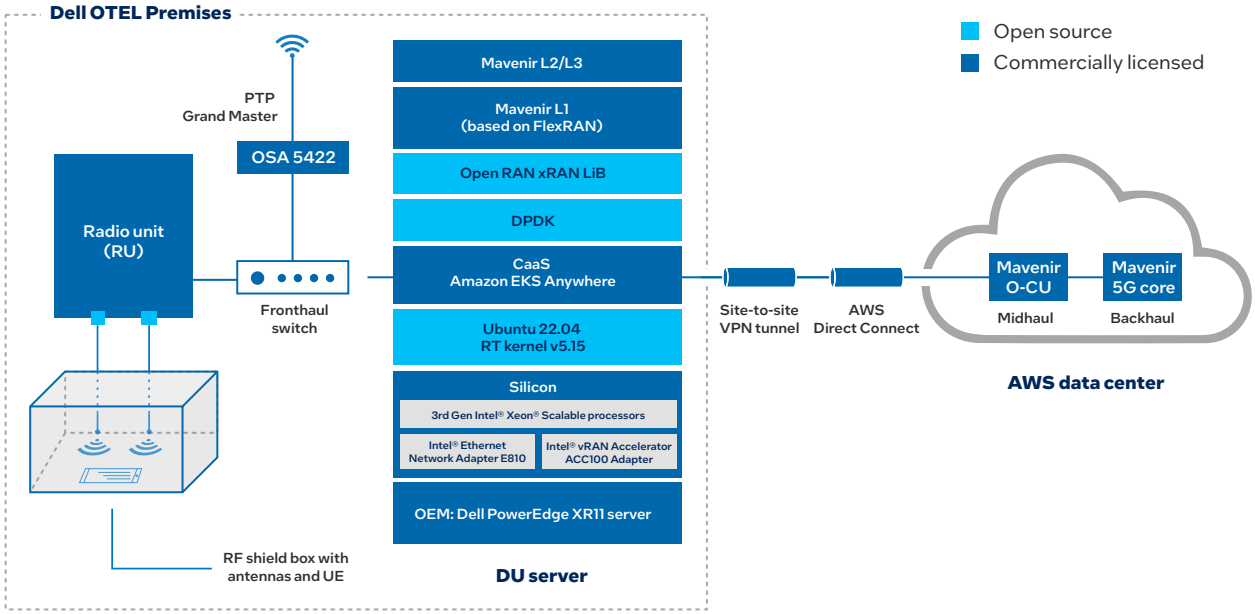

Figure 2. End-to-end architecture

Accelerating adoption

Intel, Dell Technologies, AWS, and Mavenir came together to deploy and certify an Open RAN solution on COTS OEM hardware, building a complete end-to-end virtualized telco network (RAN and core). This solution used Dell PowerEdge XR11 servers based on 3rd Gen Intel® Xeon® Scalable processors in the Dell Open Telecom Ecosystem Lab (OTEL). The software stack consisted of Amazon EKS Anywhere for the CaaS layer, Intel® FlexRAN™ reference software, and the Mavenir 5G New Radio (NR) digital unit Open RAN application, and it was validated with the radio unit (RU). This paper describes the advantages of adopting a microservices-based architecture and implementing timing and synchronization solutions. This certification exercise helped to create an end-to-end packaged offering with all the required configurations, drivers, and firmware for customers to deploy in their own environments.

Table 1 lists the different parts of the certified solution. This solution gives CSPs a new way to host their Open RAN workloads on the AWS platform, which can simplify running and managing their Open RAN networks.

Table 1. Certified solution distributed unit (DU) stack

Open distributed unit application	Mavenir DU Intel FlexRAN reference software
CaaS	Amazon EKS Anywhere
Operating system (OS)	Ubuntu 22.04
Hardware	Server: Dell PowerEdge XR11 Processor: Intel Xeon Gold 6338N NIC: 100 gigabit Ethernet (GbE) Intel® Ethernet Network Adapter E810 Acceleration: Intel® vRAN Accelerator ACC100 Adapter

Example use cases for Open RAN

This solution from Intel, Dell Technologies, AWS, and Mavenir is suitable for a variety of use cases ranging from Open RAN upgrades that accommodate network upgrades to moving applications to the far network edge and supporting private data networks.

Telco Open RAN

Traditional RANs are often characterized by proprietary hardware and a lack of interoperability between different vendors' equipment. This not only reduces telco service-providers' vendor options, but it can also hamper innovation and timely deployment of new services.

A key use case involves a CSP planning to upgrade its network. By adopting a virtualized RAN architecture, the CSP can deploy a flexible and scalable network, with the ability to dynamically allocate resources based on demand. It can also reduce both capital and operational expenditures (CapEx/OpEx) thanks to the use of COTS hardware and expanded vendor options. Furthermore, the open nature of the architecture fosters innovation, allowing the service provider to quickly introduce new services and adapt to changing market needs.

Telco far edge

The "far edge" in telco refers to the extreme edge of the network, close to end users and devices. This is where latency-sensitive and bandwidth-intensive applications need to process data, and both the control plane and the data plane are deployed and managed on the same node. Single-node Amazon EKS Anywhere clusters hold significant value in telco radio sites, as they facilitate the hosting of open distributed units (O-DUs) and sometimes open centralized units (O-CUs) on a single bare-metal server, necessitating only a minimal resource footprint.

In a far-edge use case, a telco company might use Open RAN solutions to deploy edge computing capabilities closer to users. For example, consider a densely populated urban area with a high demand for data-intensive services, such as high-definition video streaming, augmented reality, and Internet of Things (IoT) applications. By deploying edge computing capabilities at the far edge of the network, the telco company can process and analyze data closer to the source, significantly reducing latency. This can dramatically improve the quality of experience for users by ensuring smoother video streaming and more responsive IoT applications.

This approach can also help reduce backhaul costs and congestion in the network core, as less data needs to be transported back and forth. Moreover, the telco company can benefit from the flexibility and scalability inherent in the technology, easily scaling resources up or down to meet demand.

Edge and private networks

Private networks, often deployed by enterprises for specific applications like Industry 4.0, smart campuses, and more, can also benefit from Open RAN. For such use cases, multi-node Amazon EKS Anywhere clusters, which consist of more than one physical host, are suitable for users who have sufficient resources and need high availability and increased compute capacity for their workloads.

Consider a mine operator that operates a large number of automated vehicles. The miner decides to deploy a private 5G network to support IoT devices, autonomous vehicles, and real-time machine monitoring systems within the mine. By using Open RAN, the mine operator can deploy a flexible and scalable private network tailored to its specific needs. The network can be dynamically adjusted to handle peak data traffic during high-production periods, ensuring uninterrupted connectivity. Lack of downtime is essential, as lack of connectivity can result in vehicle collisions and other major accidents in the mine.

The Open RAN setup can also integrate edge computing capabilities, processing data locally within the mine. This can help reduce latency, enabling real-time control of automated machinery and faster response times for critical systems.

In addition, the miner can ensure a high level of network security, essential for protecting sensitive industrial data. Overall, Open RAN can enable the manufacturer to optimize operations, improve productivity, and maintain a high level of network performance and security.

Quick-to-certify Open RAN solution: Architectural overview

Several different technologies play important parts in this certified solution. Hardware and software from AWS, Mavenir, Intel, and Dell Technologies provide a flexible, open solution for organizations seeking to adopt Open RAN.

Amazon EKS Anywhere CaaS layer

As described in the "5G Network Evolution with AWS" and "Open Radio Access Network Architecture on AWS" white papers, AWS can provide an ideal platform for hosting 5G core and Open RAN network functions. Because Open RAN has to be placed at edge and far-edge sites, it requires various options and form-factors of edge deployment, which can be met with the AWS Outposts family and Amazon EKS Anywhere.

AWS provides Amazon EKS Anywhere for CSPs who want to use and keep their existing COTS hardware when deploying and managing 5G Open RAN network functions. As a CaaS layer, Amazon EKS Anywhere runs on bare-metal servers and operates clusters on the CSP's on-premises data centers. Amazon EKS Anywhere builds on the strengths of the Amazon EKS distribution that is common across all Amazon EKS services so CSPs can have a reliable on-premises Kubernetes environment. The Amazon EKS distribution enables CSPs to build an end-to-end platform and management framework from cell sites to an AWS region cloud using a common set of AWS services and tools.

To address the unique performance and latency requirements of Open RAN, Amazon EKS Anywhere provides support for customer-designated devices such as L1 accelerators and network interfaces supporting Precision Time Protocol (PTP), and it enables low-latency access to the devices with Single Root Input/Output (I/O) Virtualization (SR-IOV) container network interface (CNI) and Data Plane Development Kit (DPDK).

In order to address CSPs' requirements, AWS offers two deployment options for Amazon EKS Anywhere on COTS hardware. The first option is a multi-node cluster, which consists of more than one physical host. This option is suitable for CSPs who have sufficient resources and need high availability and increased compute capacity for their workloads.

The second option is a single-node cluster, where both the control plane and the data plane are deployed and managed on the same node. This option is suitable when a smaller deployment footprint is required, particularly in resource-constrained far-edge sites. Single-node Amazon EKS Anywhere clusters (as shown in Figure 2) hold significant value in telco radio sites, as they facilitate the hosting of O-DUs and sometimes O-CUs on a single bare-metal server, necessitating only a minimal resource footprint.

O-DUs on Amazon EKS Anywhere

To deploy an O-DU cloud-native network function (CNF) workload on Kubernetes involves integrating containerization, orchestration, and networking technologies to establish an adaptable and high-performing environment for network operations. Amazon EKS Anywhere supports key requirements for O-DU CNF deployment. Amazon EKS Anywhere enables customer-designated devices, such as Intel vRAN Acceleration ACC100 Adapter. Amazon EKS Anywhere supports Multus as a meta-CNI to enable network segmentation between the management, control, and user planes (see Figure 2).

You can use Amazon EKS Anywhere to implement O-DU CNFs as follows:

- **Kubernetes clusters** deployed and managed by Amazon EKS Anywhere can serve as the foundation for running the O-DU CNF and other services.
- **Cilium** used as the primary CNI provides network connectivity between pods within the Kubernetes cluster, enabling seamless communication among them.
- **Multus** as a meta-CNI plugin enables attaching multiple network interfaces to a single pod, allowing it to connect to different networks such as the management, data, and control planes.
- **The SR-IOV network device plugin** manages and allocates SR-IOV virtual functions (VFs) to pods, enabling high-performance, low-latency networking.
- **The real-time Linux kernel** is modified to provide more deterministic and lower latency performance, essential for real-time O-DU workloads.
- **DPDK** enables fast packet processing in user space by bypassing the kernel's networking stack, enhancing network performance for CNFs.
- **The CPU manager** kubelet feature provides better control over exclusive CPU resource allocation to containers and pods, which helps ensure optimal performance.

- **The topology manager** mitigates performance degradation caused by NUMA effects by making optimal hardware resource assignments for containers.
- **A container** runtime designed for use with Kubernetes helps ensure efficient container management.
- **CoreDNS**, a flexible and extensible DNS server, serves as a service-discovery mechanism within the Kubernetes cluster.
- **etcd**, the primary datastore for the Kubernetes control plane, helps ensure data consistency and reliability.
- **The O-DU CNF** is deployed on the Kubernetes cluster, using the networking and performance capabilities of the components mentioned above.
- **Complementary services** such as AWS Distro for Open Telemetry (ADOT) for monitoring and logging complement the CNF by providing observability, security, and operational insights.
- **Lifecycle management** is provided by Amazon EKS Anywhere for the Kubernetes cluster and workloads, including cluster upgrades, scaling, backup, and recovery.

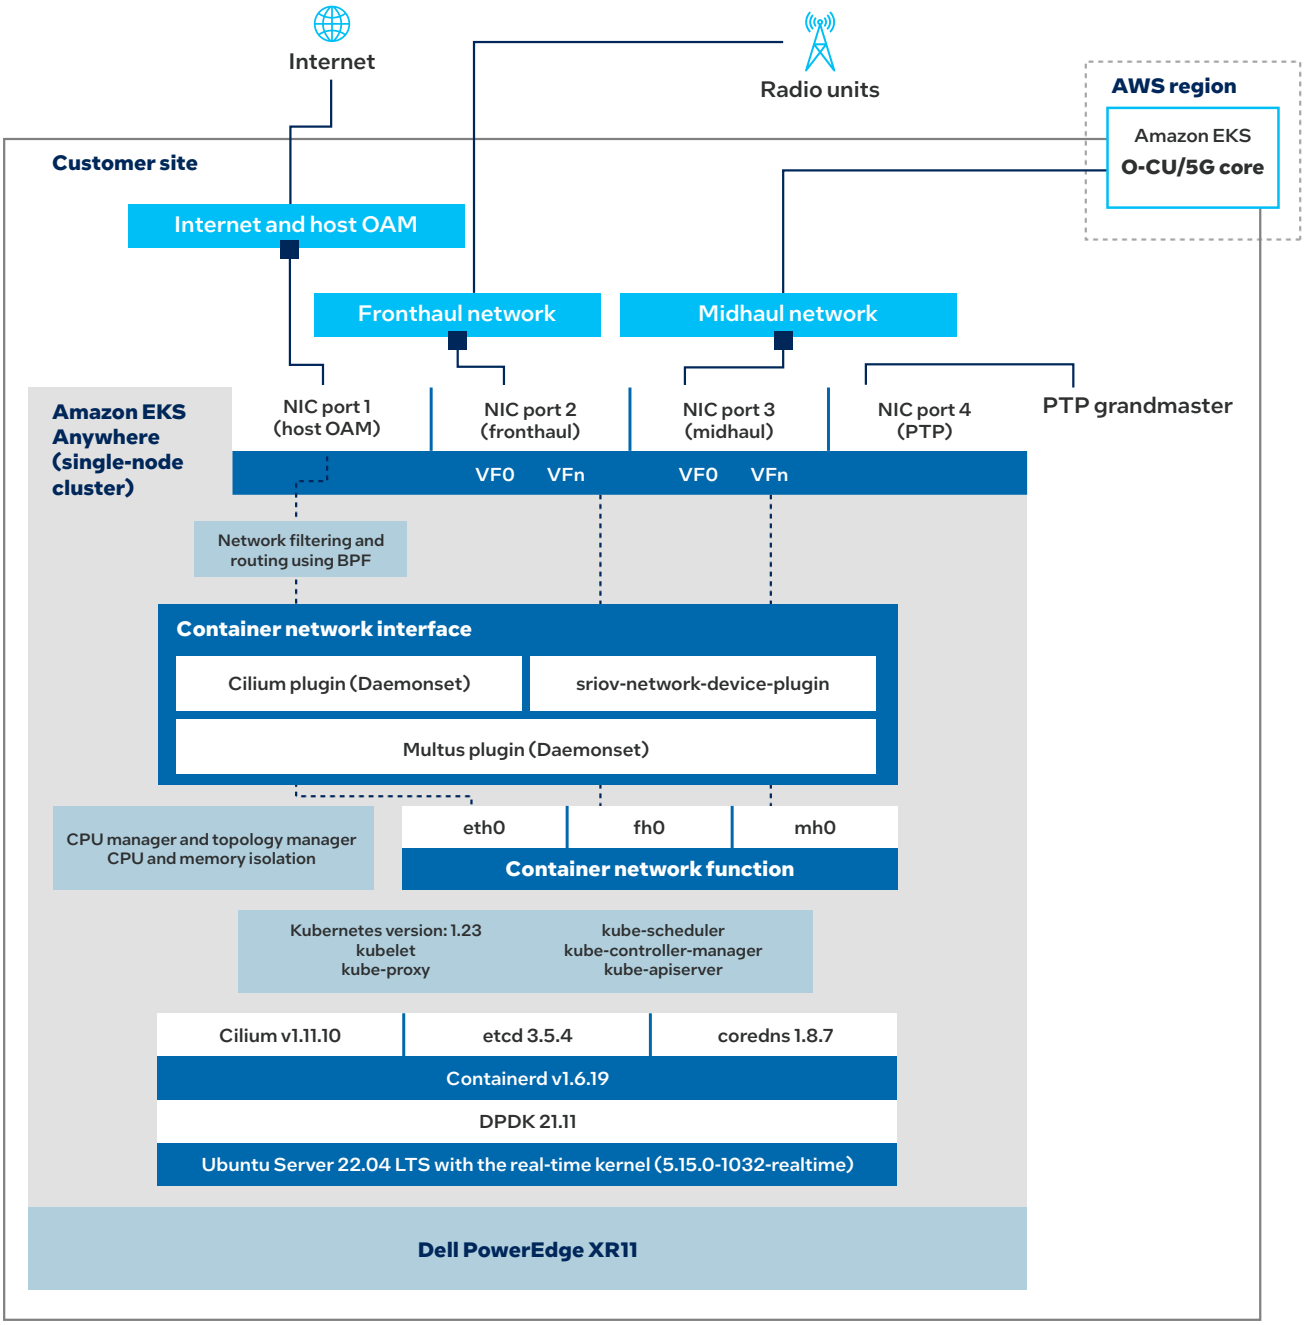

Figure 3. Architectural diagram of the Open RAN solution validated on Amazon EKS Anywhere at the time this white paper was written

Mavenir Open RAN software

The Mavenir Open RAN solution presents a groundbreaking methodology for constructing and managing cloud-native telecommunication networks. It uses a microservices-oriented architecture and open interfaces. A fully containerized, cloud-native, Open RAN solution disaggregates the DU and centralized unit (CU) that can be implemented on any cloud-infrastructure COTS hardware. This disaggregation facilitates the integration of various components, including open radio unit (O-RU), COTS hardware, CaaS, and service-management and orchestration (SMO) layers with RAN CU/DU network functions (NFs). These disaggregated components allow CSPs to combine best-of-breed components from multiple vendors, which promotes innovation and minimizes vendor dependency.

The integration activity for this Open RAN solution involves Mavenir RAN network functions (O-DU, O-CU, and management), onboarding, and overall platform readiness, which includes Intel FlexRAN reference software validation on Amazon EKS Anywhere, an optimized OS, and the real-time kernel. The integration results in a successful demonstration of end-to-end calls with commercial-user equipment (UE), hence proving the viability of the open ecosystem and the benefits of the open standard interface.

Mavenir's Open RAN solution makes use of Intel FlexRAN reference software and Intel silicon-based COTS servers for 5G RAN O-DU, O-CU, management, and PTP workloads. The Open RAN solution validation focuses on the latest real-time kernel adaptation, OS host tunings, and Kubernetes enablement to validate its performance after deploying the Amazon EKS Anywhere cluster and performing Kubernetes conformance testing.

The Mavenir Open RAN solution can be deployed over any best-in-class cloud application-modernization platform, and it can deliver a unique and differentiated distributed cloud-edge 5G solution. The solution allows CSPs to build and deploy NFs across the highly distributed cloud that extends from on-premises infrastructure to edge locations, supporting varied use cases as per network requirements. The Mavenir solution, managed by cloud operations teams, can provide a single point of governance to CSPs, giving them a birds-eye view of the deployed edge fleet and their health. The deployment can be automated centrally by making use of cloud CI/CD services, with the capability to manage thousands of edges remotely—which can enable CSPs to focus on developing solutions rather than managing infrastructure.

To realize the full potential of 5G, Mavenir's solution implements cloud-based technologies and adheres to strong design principles and architecture. With the AWS and Amazon EKS Anywhere platform delivered on the Intel FlexRAN reference architecture, Mavenir's 5G products can now be quickly deployed in a Kubernetes upstream-compatible environment on COTS hardware. The flexible kernel and network configuration allows for standardized deployment of Mavenir's Open RAN DU/CU products across different clouds or on-premises.

Intel technologies

Intel software and hardware form the backbone of this solution. The Intel FlexRAN reference architecture provides the foundation for cloud-enabled telecom network functions. Intel processors, network adapters, and the Intel vRAN accelerator provide high performance for those network functions.

Intel is driving the transformation to virtualized RAN by developing processors, NICs, and hardware accelerators to run heavy RAN workloads that require high compute and processing capabilities while reducing TCO. 5G Open RAN powered by Intel technology drives and enables flexibility, performance, and the intelligent-edge capabilities needed for CSPs to accelerate collaboration and innovation across solution providers.

Intel FlexRAN reference software (release 22.07) was installed and validated on Amazon EKS Anywhere on a PowerEdge XR11 DU server in this certification project.

Intel FlexRAN reference architecture

The Intel FlexRAN reference architecture includes multiple generations of the Intel Xeon processor family. This family features new and evolving instruction sets that enable efficient implementation of L1 processing, in addition to L2 and L3 workloads. As openness is one of the key benefits of Open RAN, Intel FlexRAN reference software enables ISVs like Mavenir to build a RAN software stack on a GPP platform such as Intel Xeon processors and to develop their own differentiated solutions tailored to operators' needs by making use of the Intel FlexRAN reference software stack. After deploying the Amazon EKS Anywhere cluster and performing Kubernetes conformance testing, the Intel FlexRAN reference software validation took place focusing on OS host tunings (Ubuntu release 22.04 for this project) and Kubernetes enablement. This focus enabled performance validation of the Intel FlexRAN reference software, the native CPU manager for resource allocation to pods, the SR-IOV device plugin, the operator, the forward error correction (FEC) operator, and the CNI in the Amazon EKS Anywhere environment. This Intel FlexRAN reference software validation enables platforms for real-time RAN workloads.

Intel Xeon Scalable processors

Intel Xeon Scalable processors feature intelligence built into the underlying silicon platform (platform-level optimizations) that enables operators to deliver with high-performance, power-efficient virtualized RAN. In addition, each new generation of Intel processors has demonstrated frequency-scaling headroom, new performance levels, and technology that includes the integration of acceleration into the CPU. Simultaneous 2G, 4G, and 5G RAN software pipelines can run inline within the GPP, eliminating the need for specialized external components. 4th Gen Intel Xeon Scalable processors can deliver up to 2x the capacity using approximately 20 percent less power compared to the previous generation.[2] They also provide higher memory bandwidth compared to the previous generation, with I/O capabilities that meet the highest capacity requirements for next-generation Open RAN deployments. Along with the processor, Intel Open RAN acceleration (implemented in a PCIe add-in card) is required to handle 5G L1 forward error correction (FEC) compute-intensive workloads. With 4th Gen Intel Xeon Scalable processors (with Intel vRAN Boost), the Open RAN acceleration is now integrated directly into the CPU.

4th Gen Intel Xeon Scalable processors (with Intel vRAN Boost) eliminate the need for external accelerator cards, which can help reduce the complexity and bill-of-materials (BoM) cost of Open RAN deployments for CSPs. Virtualized RAN based on Intel Xeon processors enables CSPs to achieve cost savings through cloud economies of scale while optimizing their networks' performance-per-watt using cloud-native capabilities such as:

- **Simplified network upgrading and testing**—A microservices-based approach enables upgrading components independently of each other at a granular level.
- **Converged services**—End-to-end virtualization makes it possible to run RAN, edge, and core workloads on a common platform and also enables automation and workload rebalancing. You can monitor telemetry data and resource utilization, and then rebalance workloads dynamically.
- **Energy savings**—Make use of power-management technologies to adjust frequency (P-states) and put cores into various sleep states (C-states) to minimize power use, even under high-load conditions.
- **Baseband pooling in the DU**—In periods of low activity, you can consolidate baseband traffic onto fewer processing cores.
- **Processor-core reuse**—Apply unused cores to non-traffic workloads, such as administration and management (OAM) functions.
- **End-to-end network slicing**—Ensure that service-level agreements (SLAs) can be deployed on demand for mission-critical services such as Ultra-Reliable Low Latency Communications (URLCC).
- **Artificial intelligence (AI) and machine learning (ML)**—Dynamically reconfigure networks to save costs, get more value from infrastructure, and support new revenue streams.

Dell Open Telecom Ecosystem Labs (OTEL)

This engagement was carried out in Dell Technologies OTEL, a collaboration platform, facilitating Intel, Dell Technologies, AWS, and Mavenir to build and qualify the Amazon EKS Anywhere solution on Dell PowerEdge servers. OTEL is a state-of-the-art 5G labs facility that offers a collaborative platform for CSPs, network equipment providers (NEPs), and system integrators to build, qualify, and showcase open 5G solutions for CSP customers. OTEL offers various services, ranging from joint partner/industry certification to customer lab validation services, implementing continuous engineering with partners to drive the common technical innovation deeper and iteratively improving outcomes and business value for the CSP customers. The goal is to facilitate CSPs to accelerate the adoption of solutions in their production environments. Dell Technologies OTEL has embarked on its transformation journey by offering various services like the solution-validation capability, which provides comprehensive, end-to-end validation, lifecycle management, and enhanced service support. Highlights of OTEL's capabilities include:

- Test and certify in state-of-the-art virtual and physical labs and flagship facilities in Round Rock, Texas, and Cork, Ireland
- Build and test new, customer-centric solutions in production-like, vendor-neutral environments
- Drive industry-standard compliance and simplify integration and interoperability
- Facilitate experimentation and learning with access to cutting edge resources
- Reduce risk and develop certified, secure solutions that capture 5G's potential for CSPs

For the end customers/partners, the OTEL platform offers:

- Reduced costs of innovation and integration
- Reduced risk of deploying new technology
- Optimized solutions to fit their needs

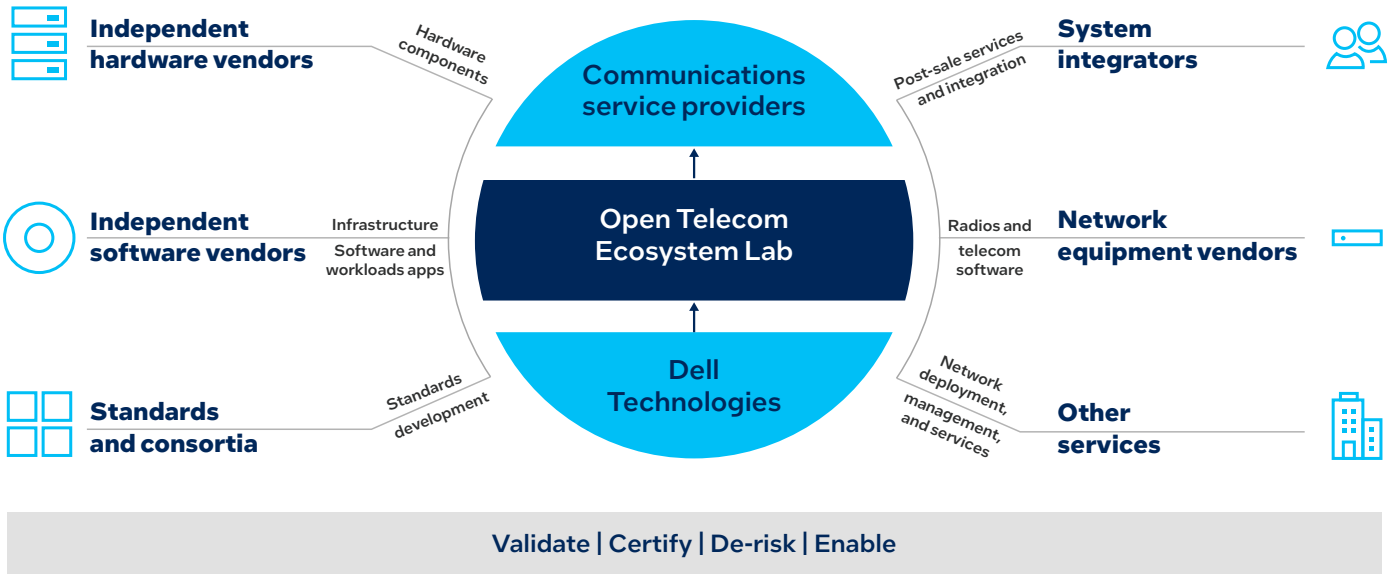

Figure 4. Overview of the Dell Technologies OTEL

OTEL offered the following benefits for this engagement:

- Joint lab with key partners: Intel, AWS, and Mavenir
- Best-in-class, high-performance Dell servers, RF chambers/RF shield boxes, test tools, and expertise
- Security-enabled remote access from anywhere, anytime

The teams went through iterative design review cycles and finalized the architecture. The OTEL team built the hybrid cloud architecture with the RAN system deployed on-premises in the OTEL and connecting to the AWS cloud, hosting Mavenir's CU and 5G core. Security-enabled remote access was provided to the Intel, AWS, and Mavenir teams, who collaborated on the integration effort.

The RAN system comprised user equipment (UE) and an antenna, which were enclosed in the RF chamber, connecting to the 5G NR RU, which in turn connected to the DU server via a fronthaul switch. The OTEL team collaborated with the respective teams and participated in the installation of Amazon EKS Anywhere, Intel FlexRAN reference architecture, and Mavenir DU. The remote access mechanisms allowed the Mavenir, Intel, and AWS teams to access the system remotely and integrate the systems end-to-end. The OTEL test tools and utilities provided remote debugging capabilities and were utilized to conduct end-to-end validations with data transfers and benchmark the performance.

Conclusion and next steps

Open RAN can unlock enormous benefits for the telco industry, ranging from lower TCO to avoidance of vendor limitation. Yet for all of its potential benefits, industry adoption of Open RAN has been held back in part by the lack of an open ecosystem to support it.

This certified Amazon EKS Anywhere solution for Open RAN from Intel, AWS, Dell Technologies, and Mavenir helps provide some guidance for navigating selection decisions in the Open RAN ecosystem. This solution incorporates industry-leading hardware and software solutions to create a model Open RAN solution that can be customized to CSP-specific requirements and objectives. To learn more, contact your Intel representative to learn how to validate your hardware stack or software. You can also contact your AWS, Dell Technologies, and Mavenir representatives for further details.

Learn more

- "O-RAN architecture on AWS"
- "The Journey to a Cloud-native, Fully Software-defined vRAN Architecture"

[1] The Senza Fili study, commissioned by Mavenir, found up to 37 percent lower five-year TCO, blended from 49 percent lower CapEx and 31 percent lower five-year OpEx. The study used "virtualized RAN" and "cloud RAN" interchangeably. Source: Monica Paolini. Senza Fili. "How much can operators save with a Cloud RAN?" Sponsored by Mavenir. 2017. mavenir.com/wp-content/uploads/2020/01/SenzaFili-Mavenir-TCO-WP.pdf.

[2] For workloads and configurations, visit www.Intel.com/PerformanceIndex: 4th Generation Intel® Xeon® Scalable Processors. Results may vary.

Printed in USA 1023/JS/PRW/PDF Please Recycle 356908-001US