Dell Technologies Reference Architecture for Wind River Studio

Architecture Guide

H18915

Abstract

This reference architecture guides the design and creation of a production-ready Kubernetes solution for managing edge cloud infrastructure and enabling edge compute use cases such as Open RAN, vRAN, MEC, and others.

Dell Technologies

October 2021

Notes, cautions, and warnings

NOTE: A NOTE indicates important information that helps you make better use of your product.

CAUTION: A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

WARNING: A WARNING indicates a potential for property damage, personal injury, or death.

Contents

Glossary

Acronym	Definition
AR	Artificial Reality
API	Application Programming Interface
CO	Central Office
COTS	Commercial Off the Shelf
CU	Central Unit
CUPS	Control and User Plane Separation
DC	Data Center
DU	Distributed Unit
ETSI	European Telecommunications Standards Institute
FPGA	Field Programmable Gate Arrays
GPU	Graphics Processor Unit
HA	High availability
IIoT	Industrial Internet of Things
MEC	Multi-Access Edge Computing (or known as Mobile Edge Computing)
NEBS	Network Equipment Building System
NFV	Network Function Virtualization
NG-RAN	Next Generation Radio Access Network
ONIE	Open Network Install Environment
Open RAN	Open Virtualized Radio Access Network
PoP	Point of Presence
QoE	Quality of Experience
QoS	Quality of Services
RAN	Radio Access Network
RU	Rack Unit
SDN	Software-Defined Network
SDS	Software-Defined Storage
SLA	Service Level Agreement
URLLC	Ultra-Reliable Low-Latency Communication
vBBU	Virtualized Baseband Unit
vCDN	Virtualized Content Delivery Network
vCU	Virtualized Central Unit
vDU	Virtualized Distributed Unit
vEPC	Virtualized Evolved Packet Core
vRAN	Virtual Radio Access Network
VM	Virtual Machines

Acronym	Definition
VR	Virtual Reality

1

Introduction

With the evolution to 5G combined with the desire for a more competitive marketplace, service providers in the telecommunications industry are moving away from integrated, purpose-built solutions to a disaggregated model for their distributed networks. For many, this solution means moving toward Virtualized radio access network (vRAN) and Open virtualized radio access network (Open RAN) solutions based on open source and open standards including Multi-Access Edge Computing (MEC) as a use case.

Dell Technologies and Wind River have joined forces, combining our respective hardware and software expertise. Dell Technologies and Wind River synergy results in support for communications service providers (CSPs) by offering a containerized edge infrastructure that delivers the low latency, redundancy, and High availability (HA) that carrier-grade RAN workloads require today - with lower operational costs.

Together, Dell Technologies and Wind River deliver a validated reference architecture ideal for the distributed provider network. Within this reference architecture is the combination of two key attributes: unique open-source technology that meets and exceeds the performance and detailed requirements of hosting 5G and edge workloads, and a suite of high-automation tools to reduce Day 1 and Day 2 operational effort.

The joint solution between Dell Technologies and Wind River includes the market's unique ability to perform true zero-touch, fully autonomous orchestrated remote Edge sub clouds—regardless of the number of nodes. A bare-metal state moves to a fully operational online system without human intervention. This autonomous nonhuman functionality is just one example the Day 1 capabilities provide to lower operational costs for the provider and others. Single-pane-of-glass management and the ability to scale to a full cloud in a single node for low-cost use of the latest cloud-native, container runtime, and continuous integration and delivery (CI/CD) capabilities.

Dell Technologies and Wind River have been working jointly on validated solutions for the network Edge to deliver a joint vision of a new generation intelligent, programmable, and automated Edge platform.

This reference architecture describes the design and creation of a production-ready Kubernetes solution for managing Edge cloud infrastructure and enabling Edge to compute use cases such as Open RAN, vRAN, and MEC.

Topics:

Overview

As service providers begin to deploy their Edge networks to support 5G, there is much to consider. Baseline requirements are challenging and include the lowest latency, HA, and security – all deployed on a massive scale. Because of the scale and need for flexibility, it is highly challenging and costly to support Edge networks on the legacy model of fully integrated, purpose-built equipment. Service providers need the ability to deploy, update, move, and bring down workloads as needed by network demand, and it needs to happen without sending a technician to the site. Plus, disaggregation allows service providers to select the best of breed components with more competitive pricing.

Meeting baseline requirements is just the start. Service providers need to deploy and manage these networks in a cost-effective manner. Zero-touch provisioning, single-pane-of-glass host-level management, and automated updates and upgrades are examples of Day 1 and Day 2 management functionality that requires a highly distributed network.

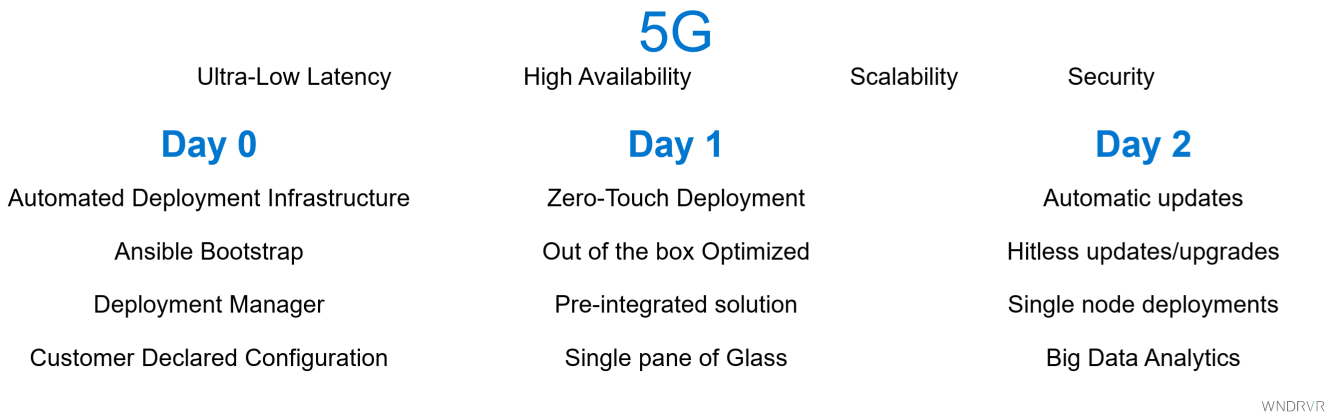

Figure 1. Representative sample of the unique requirements

Distributed Edge helps computing service providers deliver next generation 5G services faster. Edge computing allows for moving data, content, applications, and services closer to the end users or devices (delivery point), reducing round-trip delays and latency. At the same time, because data no longer needs to be sent to the central location or cloud, significant backhaul traffic is reduced, resulting in OPEX and CAPEX savings for the service providers.

Dell Technologies is working with the leading service providers on their journey to network and Edge transformations. Dell Technologies addresses many service provider problems by:

- Delivering newer infrastructure platforms to address edge requirements
- Optimizing a distributed edge cloud platform to deliver applications close to the user employing an infrastructure that spans back to the network's core
- Building a curated partner ecosystem to address specific emerging applications and Edge use cases such as Open RAN, vRAN, MEC, vCDN, and Industrial Internet of Things (IIoT)
- Contributing to several Open-source consortia focused on network transformation and the Edge

2

Reference Architecture

This reference architecture provides an overview of the key aspects of edge requirements. It provides a reference implementation of a distributed edge cloud platform to deploy distributed network services in a cloud environment at the edge.

This reference architecture focuses on Open RAN and Edge deployment (Near Edge, Far Edge, and Enterprise Edge).

Topics:

Network edge physical landscape

The distance from the Core data center on the operator's infrastructure and the services hosted at that location determines the physical location of the Network Edge.

The edge locations highlighted in the following figure are:

- RU (also Device edge or Customer Premises
- Far edge (vDU)
- Edge (vCU/MEC)
- Core Network (National/Regional office)

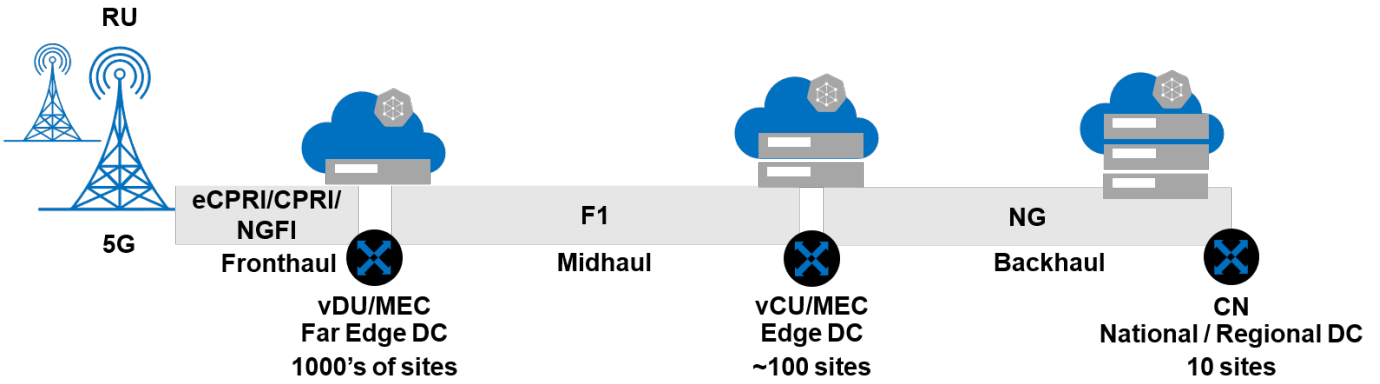

Figure 2. Edge locations

Core data center

The Core is the heart of CSP network services. It is a national/regional centralized location for the CSP's management plane and all federated control planes for the Near Edge regions. Applications such as EPC, IMS, PCRF, MANO, Analytics, and OSS/BSS, are part of this 4G/5G Core architecture.

Edge

The Edge (vCU) aggregates all types of traffic from front-haul, including fixed-line, 4G, Wi-Fi traffic to back-haul. Each Edge site manages multiple Edge sites for management, control, and data plane.

Usually, multiple infrastructure racks deploy at this site supporting various workloads such as MEC, CDN, CORD, UPF, and network slicing application functions.

Far Edge

Far Edge is synonymous with a microdatacenter located at a cell tower or close to a customer's premises. This location is sometimes referred to as the "Last Mile" to the subscribers. The Far Edge infrastructure has specific physical requirements for thermal/cooling, power, rack spacing, and front and rear I/O access. This location is the closest to the user and provides key services where latency, reliability, and experience are the most critical factors.

Network services deployed and used by Far Edge include vRAN (Next Generation RAN (NG-RAN)), IIoT, Private LTE, Connected Cars, AR/VR. Most of these services require ultralow latency, high network bandwidth, and real-time synchronization.

Device Edge

Device Edge refers to the network terminus that might be the radio unit at the cell tower, or the customer or enterprise premises where thousands or in sometimes, millions of devices are connected to the network and can act as sources and consumers of data. These devices include mobile devices, IoT sensors, and industrial equipment that connects either through wireless or wireline connections.

Network edge platform overview

The network edge platform architecture integrates traditionally into purpose-built equipment and contains three layers. The current model represents an evolution from legacy equipment that includes:

- **Hardware resources**—Represented by Dell EMC PowerEdge servers and PowerSwitch networking
- **Virtualization layer**—Wind River Studio
- **Virtualization layer**—Hosted containerized network functions

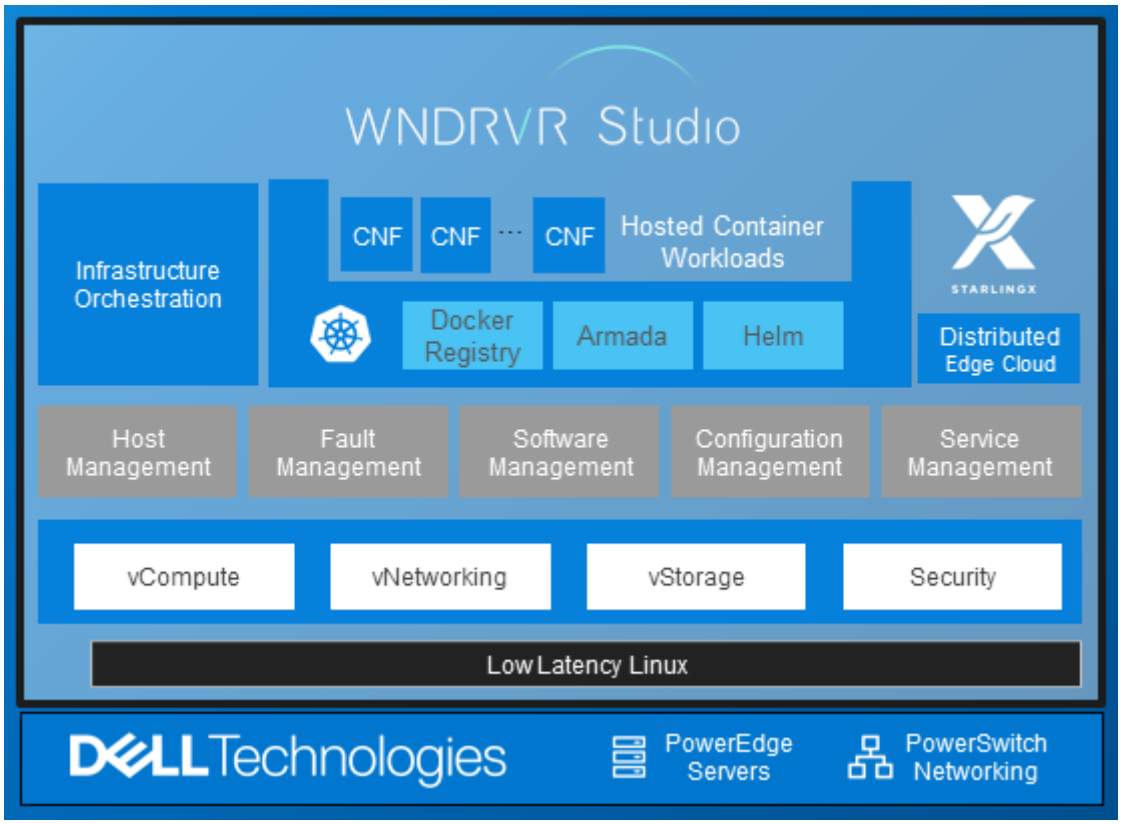

Figure 3. Network edge platform layers

Open RAN use case

Network Function Virtualization (NFV) and Software-Defined Network (SDN) are changing the game. This new virtualized, open RAN architecture provides the same benefits that the operators have seen with the deployment of virtualized services in the Core (vEPC, vIMS). Virtualization is moving to Edge locations for cloud-native-based solutions at the Edge.

Essential benefits of Open RAN

Open RAN decomposes the traditional radio stacks and maps them into discrete elements. Decomposition allows the hardware solutions to be decoupled from software implementations, enabling an eco-system of vendors. Open RAN delivers an infrastructure for more cost-effective and efficient access networks.

A genuinely open virtualized RAN helps in several areas, including:

- Using economical and ubiquitous Ethernet and or IP-based transport to support lower complexity
- Using open-source components and open standards to enable best interoperability, cost control, and future-proof deployments
- Using x86 hardware and NFV to provide elasticity of capacity
- Pooling capacity resources and licenses to cater to large sets of devices with diverse needs
- Providing a flexible platform for edge computing that is essential for a new generation of services
- Allowing dynamic orchestration of 5G slices

3

Dell Technologies PowerEdge for Edge Portfolio

Optimal compute infrastructure serves a critical role in transforming the service provider network to a highly efficient, flexible, and scalable disaggregated model. Dell Technologies offers a wide range of compute options, giving Service Providers the flexibility to design their disaggregated network architecture based on the appropriate workloads and the data center location requirements. The Dell Technologies portfolio lists the Dell Technologies products.

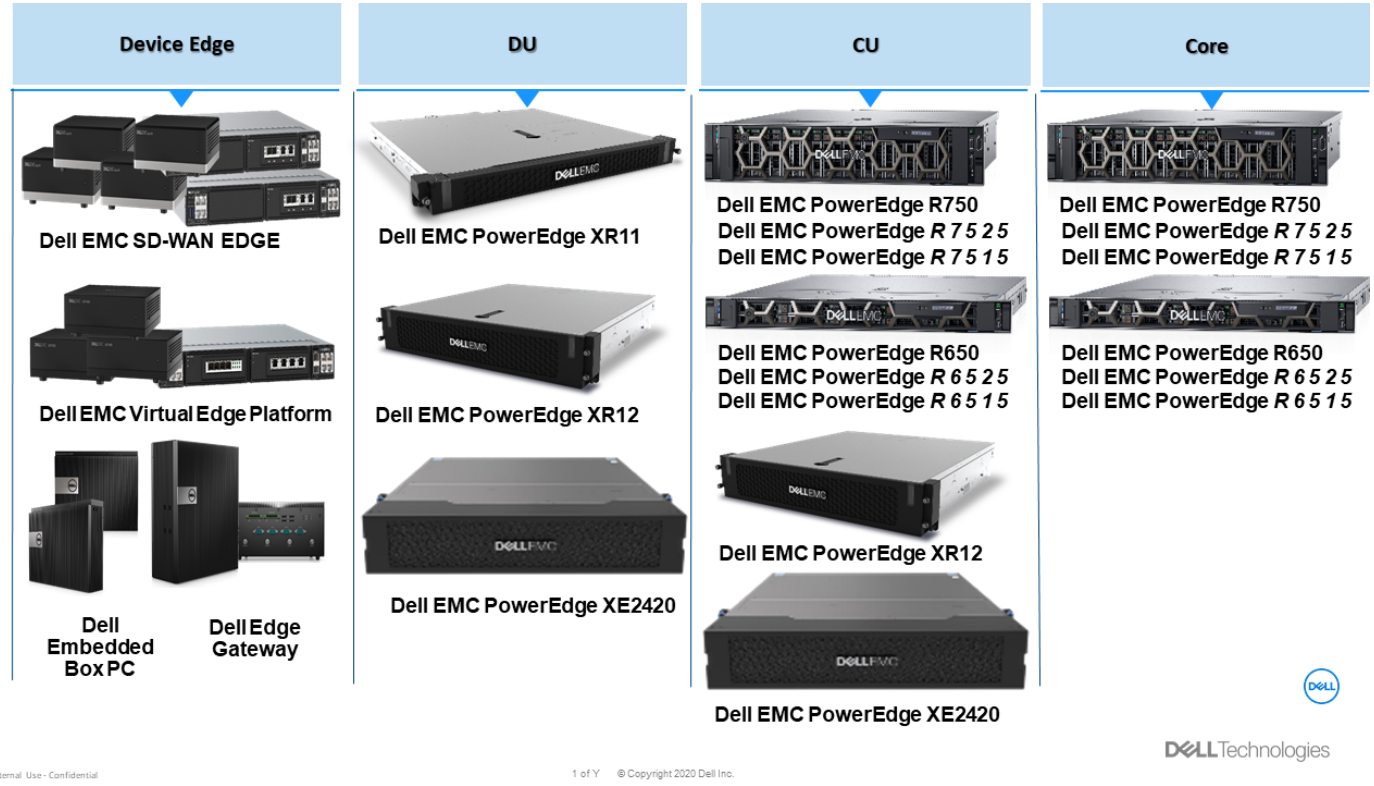

Figure 4. Dell Technologies edge portfolio

4

Wind River Studio Distributed Cloud Infrastructure

Topics:

- Open Source
- HA solution
- Scalability
- Day 0 automation
- Day 1 and 2 automation

Open Source

Wind River Studio is the first cloud-native platform for designing, dseveloping, operating, and servicing mission-critical intelligent Edge systems that require security, safety, and reliability. Studio is based on the open-source project StarlingX and is architected to deliver digital scale across the entire life cycle through a single pane of glass to accelerate transformative business outcomes.

Studio operator capabilities include an integrated cloud platform unifying infrastructure, orchestration, and analytics so that operators can deploy and manage their intelligent 5G edge networks globally.

Wind River Studio Cloud Platform capabilities represent the only open-source cloud infrastructure that satisfies critical requirements for deploying and managing distributed cloud networks. Key differentiators include low-latency, HA, scalability, security, ease of Day 1 operations with zero-touch deployment, out-of-the-box optimized and pre-integrated solutions, and support for single-node deployments in production at the Far Edge. Continuing ease of Day 2 operations with automatic updates, upgrades, and single-pane-of-glass management of geo-distributed edge clouds.

Wind River Studio Cloud Platform is a CNCF certified Kubernetes distribution .

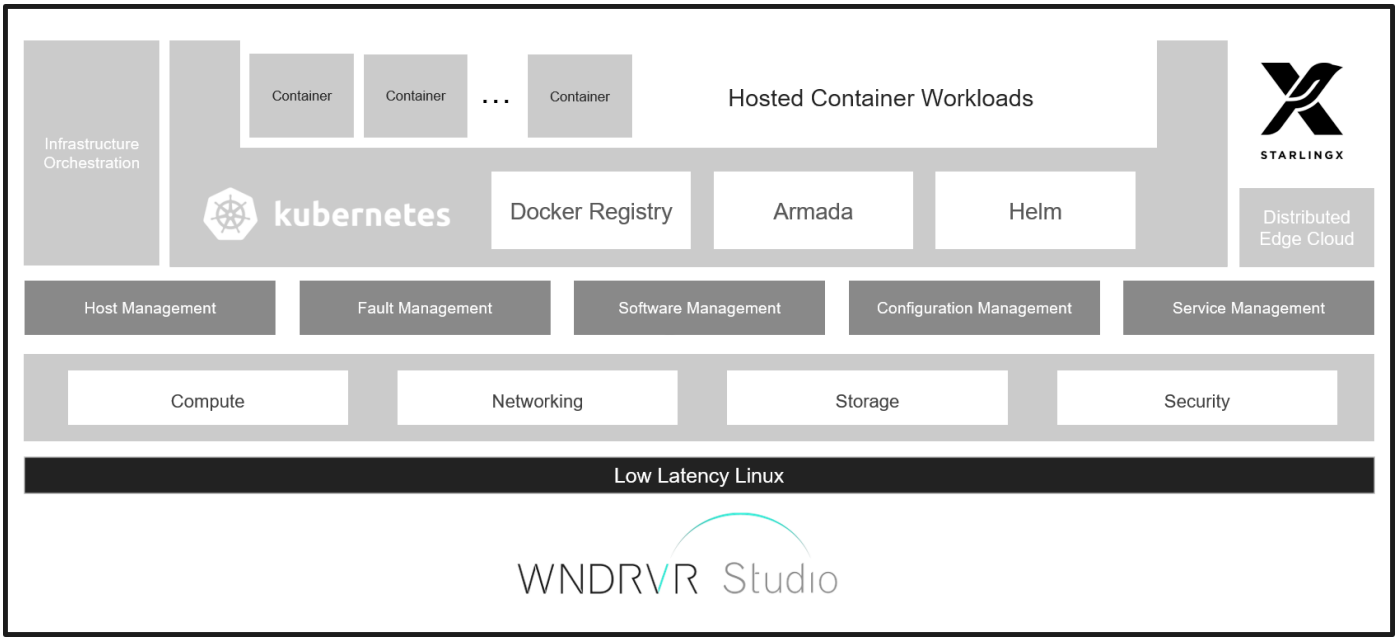

Figure 5. Wind River Studio Cloud Platform

The active community contributing to the code has produced a hardened commercial product undergoing open RAN integration and production deployment in its current version. Cloud platform practices an upstream first open-source model for StarlingX.

NOTE: The upstream first model is before the downstream production of the open-source binary distribution by Wind River Studio.

Fundamentally, StarlingX is a developed Edge cloud solution. StarlingX ensures the support for Edge-specific focused platforms, a focus on performance, acceleration, scalability both at the low and high end to ensure low CAPEX at the Edge, and the ability to manage a geo-distributed solution operationally. These features are critical aspects of ensuring this contemporary cost-effective deployed technology by service providers.

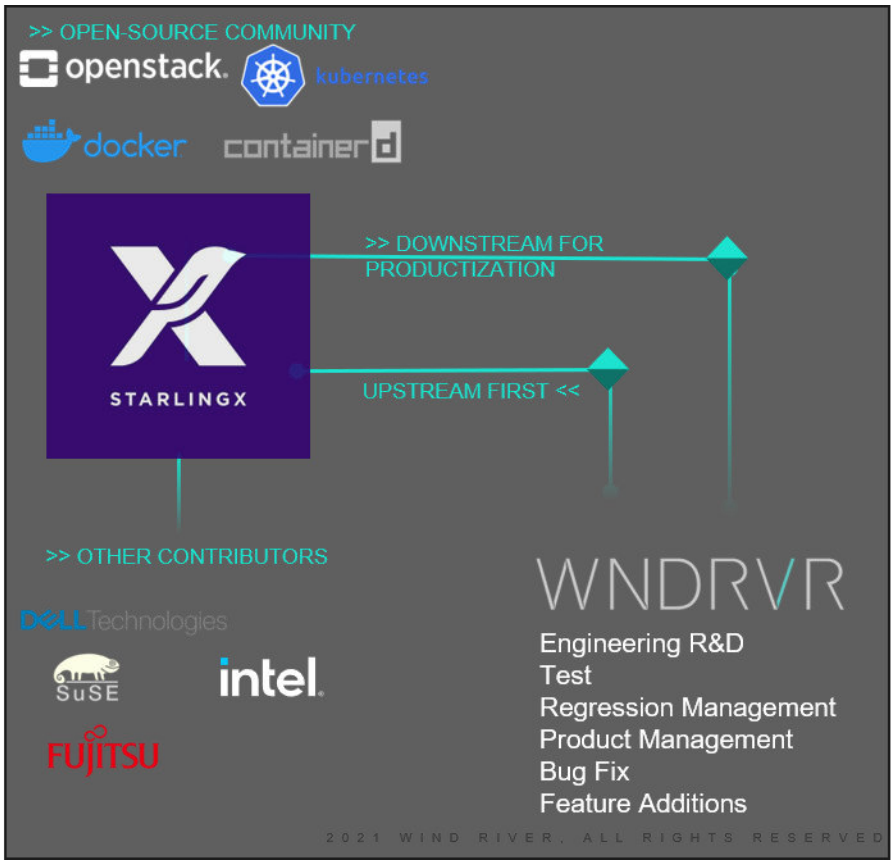

Figure 6. StarlingX integration

HA solution

The Kubernetes distribution in Wind River Studio Cloud Platform deploys in a highly available configuration with a 1:1 service model to minimize the control plane footprint in edge deployments. Application Programming Interface (API) requests are directed to activities instances by a cluster floating IP address. A DRBD-backed file system maintains persistent storage.

Service Management (SM) manages service availability and activity:

- Handles HA sparing of individual services
- Monitors host, service, and network availability
- Mitigates split-brain scenarios

Scalability

Wind River provides ultimate deployment flexibility with a range of deployment models:

- The single-node model in which the control plane (primary), worker, and storage functions are combined
- Two nodes, which is an HA model in which the control plane (primary), worker and storage functions are combined
- Two nodes plus a set of worker nodes (starting with a small number of nodes and growing as workload demands increase)
- Pair of dedicated controllers hosting the control plane (Kubernetess primary) with controller-based storage or dedicated storage nodes and up to 200 worker nodes which is the standard configuration

- Distributed Cloud configuration for infrastructure management and orchestration of thousands of geographically dispersed edge nodes

Day 0 automation

Wind River provides a fully automated deployment infrastructure driven by customer-declared configuration values.

Ansible automates the initial bootstrapping based on declared bootstrap configuration values. After bootstrapping, the Wind River deployment manager deployed by Ansible completes the deployment of the infrastructure. The deployment manager is a Kubernetes application driven by a customer-declared configuration.

Day 1 and 2 automation

For operations and maintenance, Wind River provides a rich set of capabilities through its system REST APIs. These capabilities include:

- Two-way (in-out) cluster scale (adding and removing worker nodes)
- Automated software updates
- Automated software upgrades
- Fault management feed to report infrastructure level alarms and faults
- Backup and Restore

Wind River's Distributed Cloud architecture provides single-pane-of-glass management of thousands of geo-distributed edge nodes.

5

Zero-Touch Provisioning Using Redfish

Topics:

Zero-Touch Provisioning

Zero-Touch Provisioning (ZTP) is a critical feature for service providers; it allows them to remotely install all software on every server without every server without requiring personnel. ZTP requires the Redfish firmware (version 1.2 or greater) as standardized by the Distributed Management Task Force (DMTF) to be installed on the Dell EMC servers.

NOTE: ZTP validated on Dell EMC PowerEdge R740xd and Dell EMC PowerEdge XR11

Redfish

Dell EMC PowerEdge servers offer a comprehensive range of embedded systems management functions enabled by the Integrated Dell Remote Access Controller (iDRAC) with Lifecycle Controller. These functions adhere to standard APIs, including Redfish v1.6.

iDRAC with Lifecycle Controller technology is part of a comprehensive data center solution that helps keep business-critical applications and workloads constantly available. The technology allows administrators to deploy, monitor, manage, configure, update, troubleshoot, and remediate Dell EMC servers from any location without agents. It accomplishes these tasks regardless of an operating system or a hypervisor presence or state.

Redfish supports the full range of server architectures from monolithic servers to converged infrastructure and hyperscale architecture. The Redfish data model, which defines the structure and format of data representing server status, inventory, and available operational functions, is vendor-neutral. Administrators can then create management automation scripts that can manage any Redfish, which is compliant server. Essential for the efficient operation of a heterogeneous server fleet.

To fully support the Redfish standard, the iDRAC with Lifecycle Controller supports a RESTful API as well as IPMI, SNMP, and WS-Man standard APIs. The iDRAC RESTful API builds upon the Redfish standard to provide a RESTful interface for Dell EMC value-add operations, including:

- Information about iDRAC with Lifecycle Controller out-of-band services—web server, SNMP, virtual media, SSH, Telnet, IPMI, and KVM
- Expanded storage subsystem reporting covering controllers, enclosures, and drives
- Detailed chassis information covering power supplies, temperatures, and fans
- Detailed inventory and status reporting for host network interfaces, including such details as IP address, subnet mask, and gateway for the Host operating system, which results with the iDRAC service module installed under the server operating system.

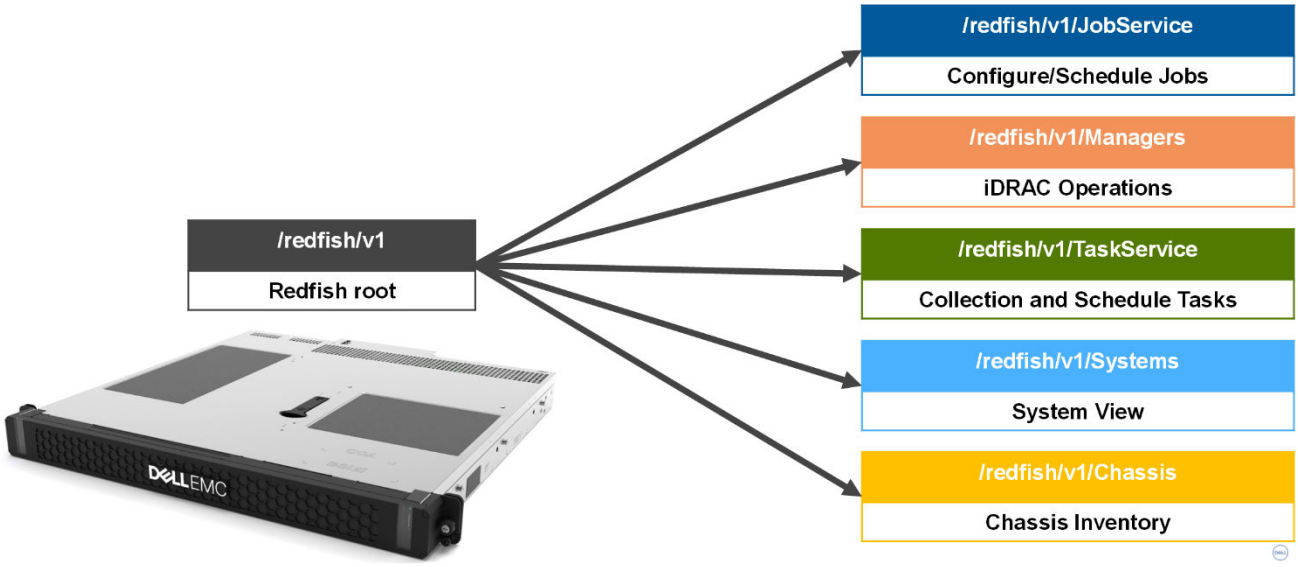

Figure 7. Redfish configuration

Wind River Studio ZTP implementation

The Wind River Studio cloud platform distributes a single ISO file, making installation easy. Redfish platform management services allow installing this ISO on all nodes using standard JSON files and RESTful APIs. After installing, connecting, and powering the hardware, an engineer thousands of miles away can complete the entire software installation using familiar automated tools.

The cloud platform user documentation Cloud Platform Distributed Cloud Configuration and Cloud Platform Installation Overview describes how to use the Redfish-enabled Dell EMC servers. The automated ZTP deployment has 3 phases: installation of the ISO, remote bootstrap, and deployment of subclouds (edge nodes) from the Central Cloud (at the Regional or National Data Center). You can run the procedures that are outlined in this section by using RESTful APIs from existing IT infrastructure, from a MANO, or directly on the central cloud.

Full ZTP of a system controller in the central region with hundreds of vRAN sites (subclouds) support Wind River Studio Conductor in a distributed cloud deployment.

Installation

The cloud platform creates a modified ISO image matching the edge cloud's networking configuration. Next, the customized ISO initiates a Redfish remote media installation when the edge cloud and the network connect to the iDRAC and OAM interfaces. The edge cloud's networking configuration embeds in the ISO installed from Redfish remote media. When the remote node boots connectivity is enabled without manual interaction which permits the second step, bootstrapping, to initialize.

Bootstrap

Similar to the installation, bootstrap values reside in a remote YAML file. The following parameters are defined and based on an Ansible Playbook to run and initialize the edge cloud:

- System Mode [Simplex, Duplex, or Standard]
- Name
- Description
- Location
- Networking
- Registry
- Security and Certificates

Deployment

The Wind River Studio Cloud Platform Deployment Manager provides a data-driven method for configuring the platform components of a StarlingX-based cloud platform installation. This implementation eases the adoption of StarlingX systems by automated deployment systems (CI/CD pipelines). By providing a data-driven method of defining each system installation, the end user is no longer required to manually interact with the system through the system CLI, UI, or directly with the System API install the system initially.

The current scope of the Deployment Manager is to install the initial system only (so-called Day-1 operations). The Deployment Manager consumes a YAML-based deployment configuration file provided by an end user or automation framework. It attempts to reconcile the system state to match the wanted state defined in the supplied deployment configuration file. When each host completes reconciling, it transitions to the unlocked/enabled state. When each Host has reached the wanted state, the system is ready to deploy an application workload.

When the system has reached the wanted system state, the Deployment Manager no longer accepts further changes to the configuration. The end user must interact via one of the accepted user interface methods (that is the system CLI, UI, or the system REST API) to modify the system.

In the future, the Deployment Manager will support postinstallation operations (that is so-called Day-2 operations). End users will modify the system configuration by supplying an updated deployment configuration rather than interacting with existing system interfaces.

The end user must supply a deployment configuration model that conforms to the supported system definition schema. The schema defines a set of Kubernetes Custom Resource Definitions (CRD) instances. It provides documentation for each attribute and validation rule that conforms to the OpenAPI v3 schema validation specification. The complete schema definition is stored in this repository and found in the `config/crds` directory. The CRD instances are automatically generated based on annotations added directly to the source code found under `pkg/apis`.

A complete system deployment configuration is composed of several Kubernetes Custom Resource (CR) instances. Each CR conforms to a CRD instance defined under the `config/crds` directory. For example, a system deployment may be composed of several instances of the following CRD types:

- System
- Platform Network
- Data Network
- Host Profile
- Host

It is possible to move common Host attributes into a Host profile definition and reuse that definition from many Host resources to streamline defining many Host records. Similarly, it is possible to define multiple layers of HostProfile resources so that attributes common to multiple HostProfile resources can be grouped into a common HostProfile resource and reused by other HostProfile resources. A Host resource can inherit from a HostProfile but still provide overrides for individual attributes that may be Host-specific.

When the Deployment Manager prepares to configure a Host resource, it first resolves the final Host attributes by merging the hierarchy of HostProfile resources related to that particular Host and applies any Host-specific overrides. The final Host attributes are validated and processed. See the HostProfile schema documentation for more information about individual attributes during the resolution of the HostProfile hierarchy.

6

Far Edge Reference Architecture

Topics:

Wind River Studio Cloud Platform architecture

The Wind River Studio Cloud Platform is scalable from a single server to thousands of servers, all manageable from the data center through a single pane of glass in a distributed cloud architecture. SAN and Ceph are popular storage options.

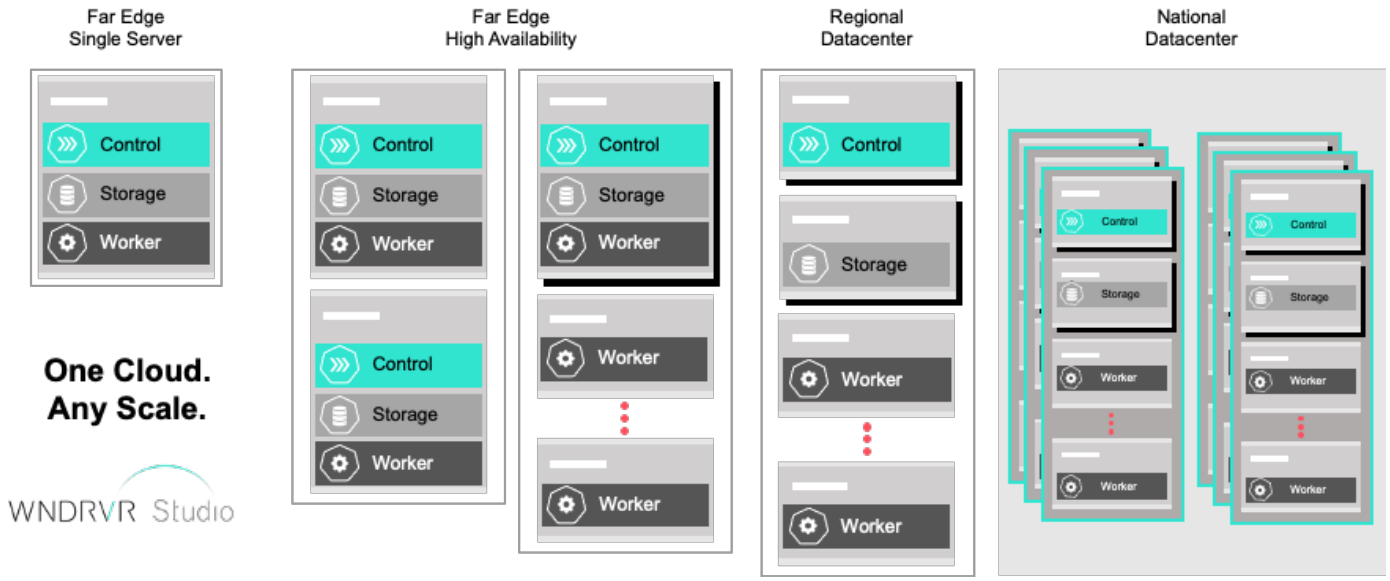

Figure 8. Wind River Studio Cloud Platform architecture

Wind River Studio Cloud Platform generalized deployment configuration

The Wind River Studio Cloud Platform architecture supports various hosts, networks, and networking hardware in different configurations.

This section uses the Standard with Storage Cluster on dedicated Storage Nodes deployment option as a reference configuration to introduce some common concepts of a Cloud Platform deployment. Regional or national data centers use the complete configuration with all the elements shown in Figure 9 Logical dedicated deployment options.

The following figure shows a logical view of the reference deployment:

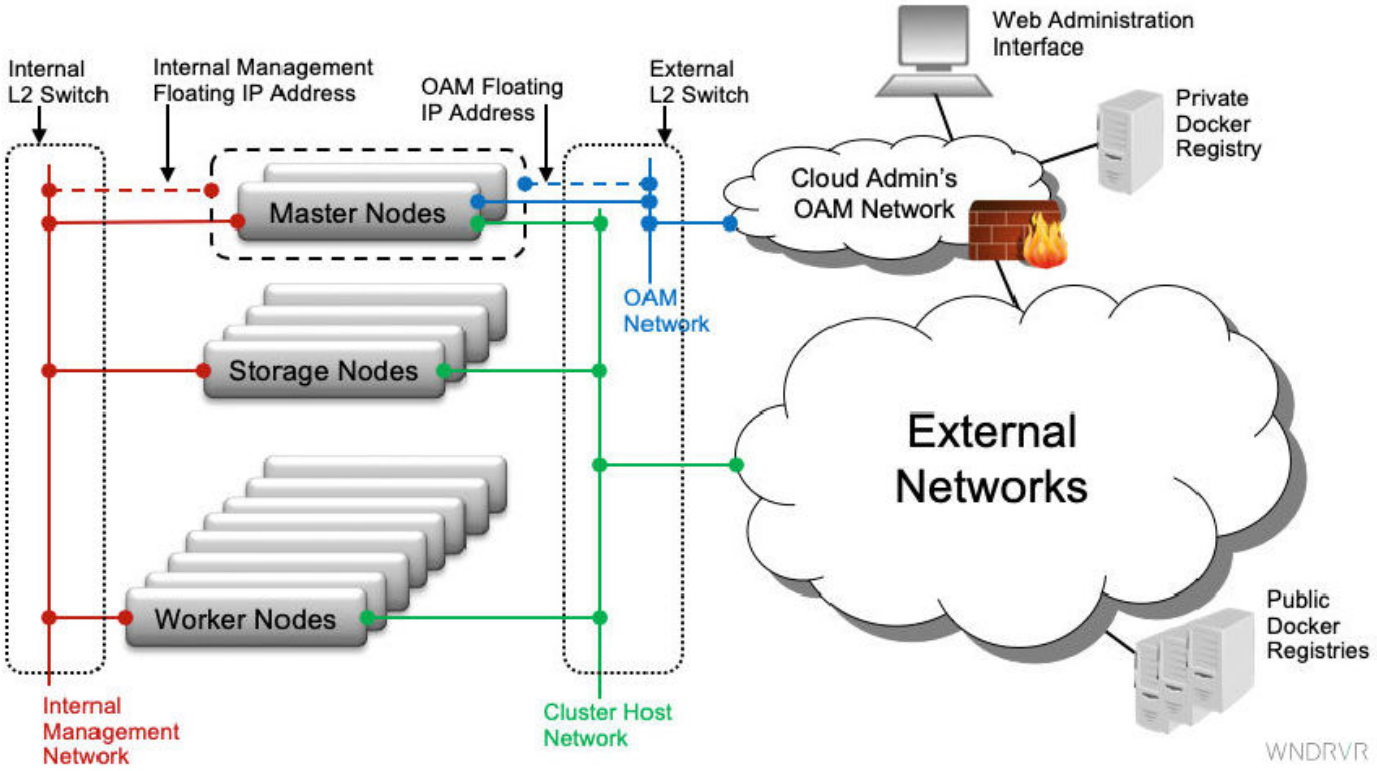

Figure 9. Logical dedicated deployment options

The Figure 9 Logical dedicated deployment options reference architecture illustrates the Standard with the Dedicated Storage Nodes deployment option. The controller, storage, and worker functionalities are deployed on separate physical hosts allowing controller nodes, storage nodes, and worker nodes to scale independently.

As shown in Wind River Studio Cloud architecture, the deployment is often scaled down to one or two servers for Edge applications.

Far Edge vRAN HA

The smallest HA Wind River cloud solution (two servers) is ideal for vRAN applications with six 9's reliability.

Ideally, each of the two servers is identical (including PCIe cards), but the Central Unit (CU) system (if present) implements reduced functionality for cost reduction. In this case, upon failure and replacement of the Distributed Unit (DU) server, the workload and system design must allow for reduced functionality.

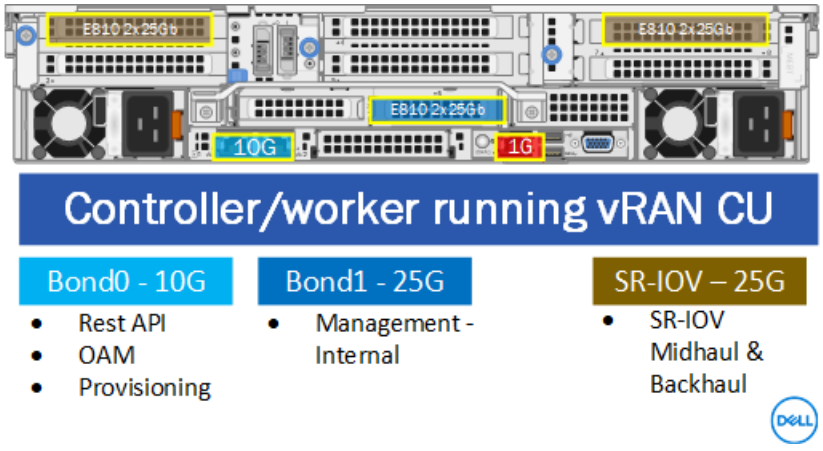

Figure 10. CU server configuration

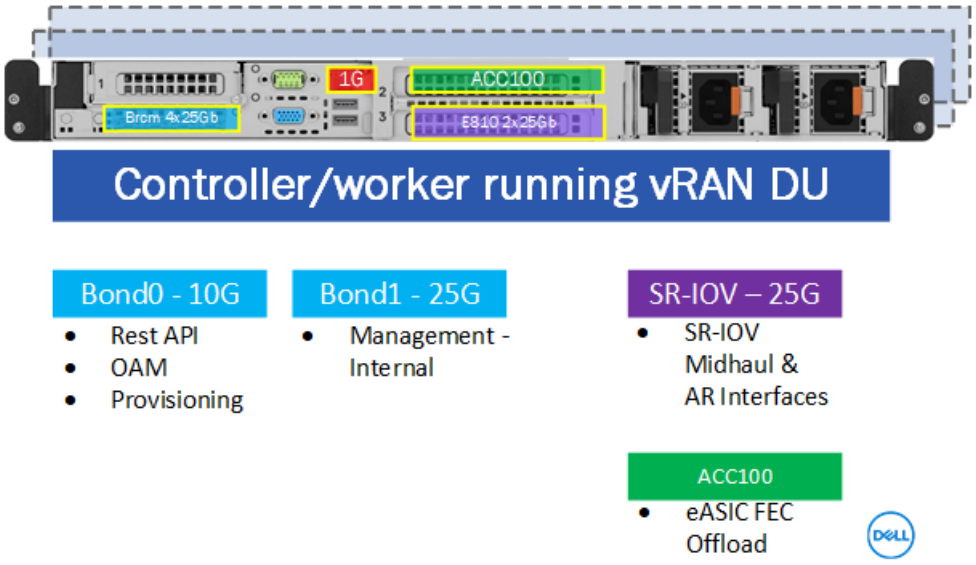

Figure 11. DU server configuration

Far Edge vRAN Single-Server

One of the essential Far Edge features of Wind River Studio is that it can scale down to a single server and still perform intense, Far Edge functions such as vRAN DU. This configuration features up to 70 percent less CAPEX than competitive solutions.

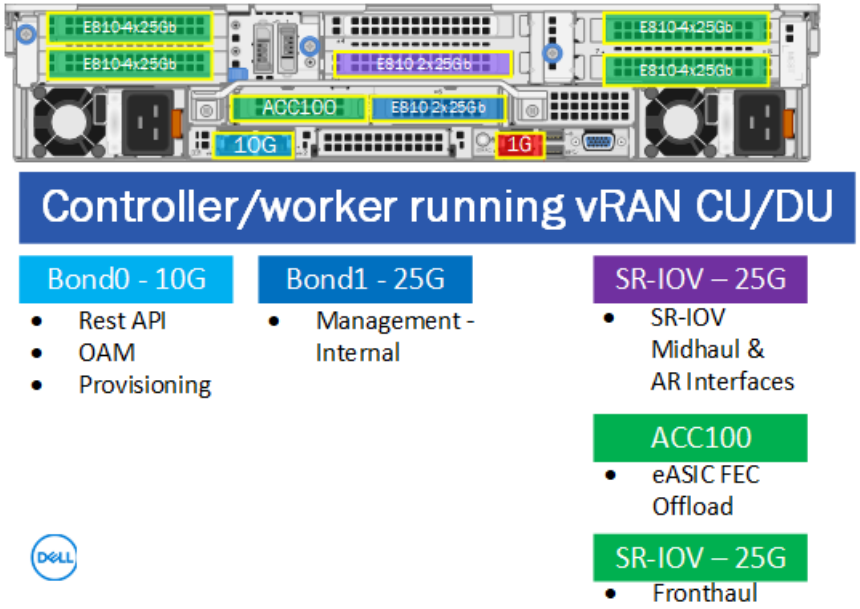

Figure 12. CU/DU configuration

Physical network topology

The following figures show typical physical network connections for a vRAN use case:

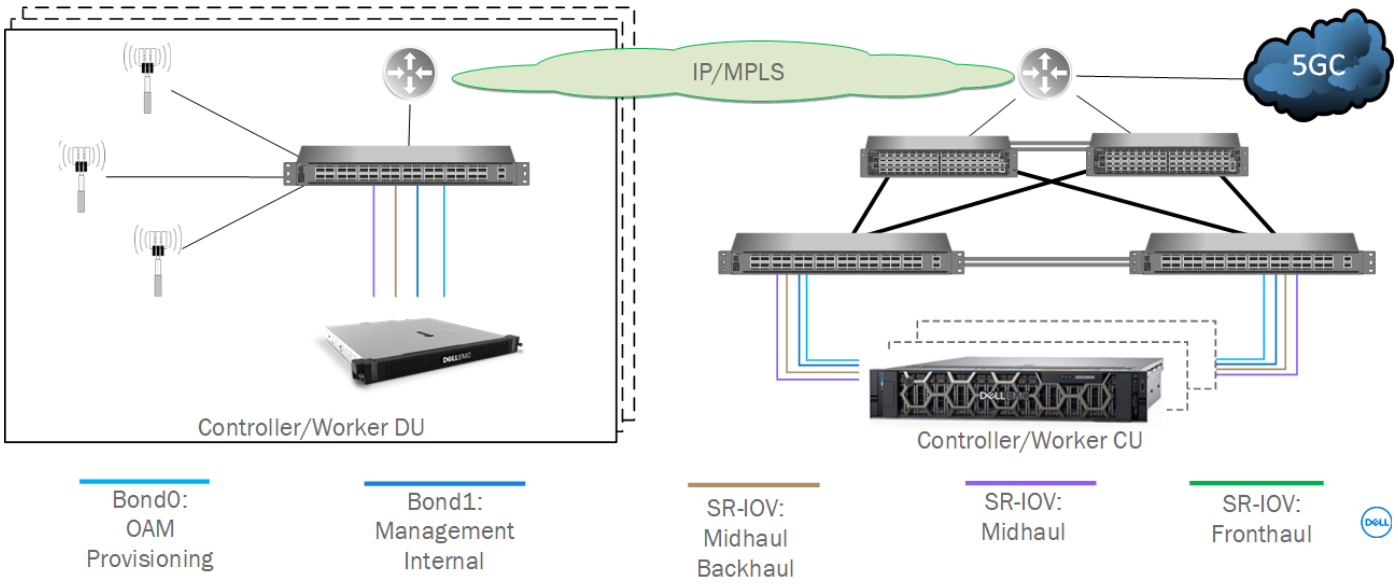

Figure 13. IP/MPLS configuration

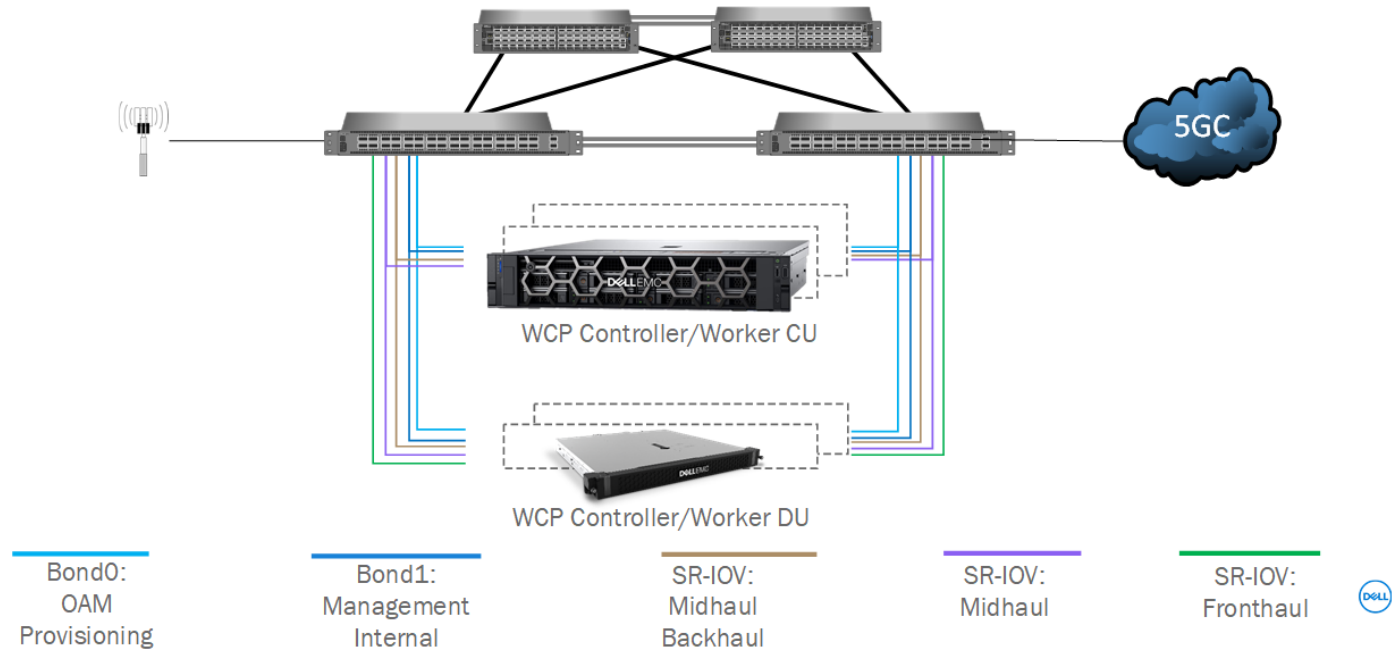

Figure 14. Single CU/DU site configuration

Logical networking

All-in-one Simplex configuration (single server)

The All-in-one Simplex deployment configuration provides a scaled-down cloud platform that combines controller, storage, and worker functionality on a single nonredundant host. Deploy the simplex configuration at the Far Edge.

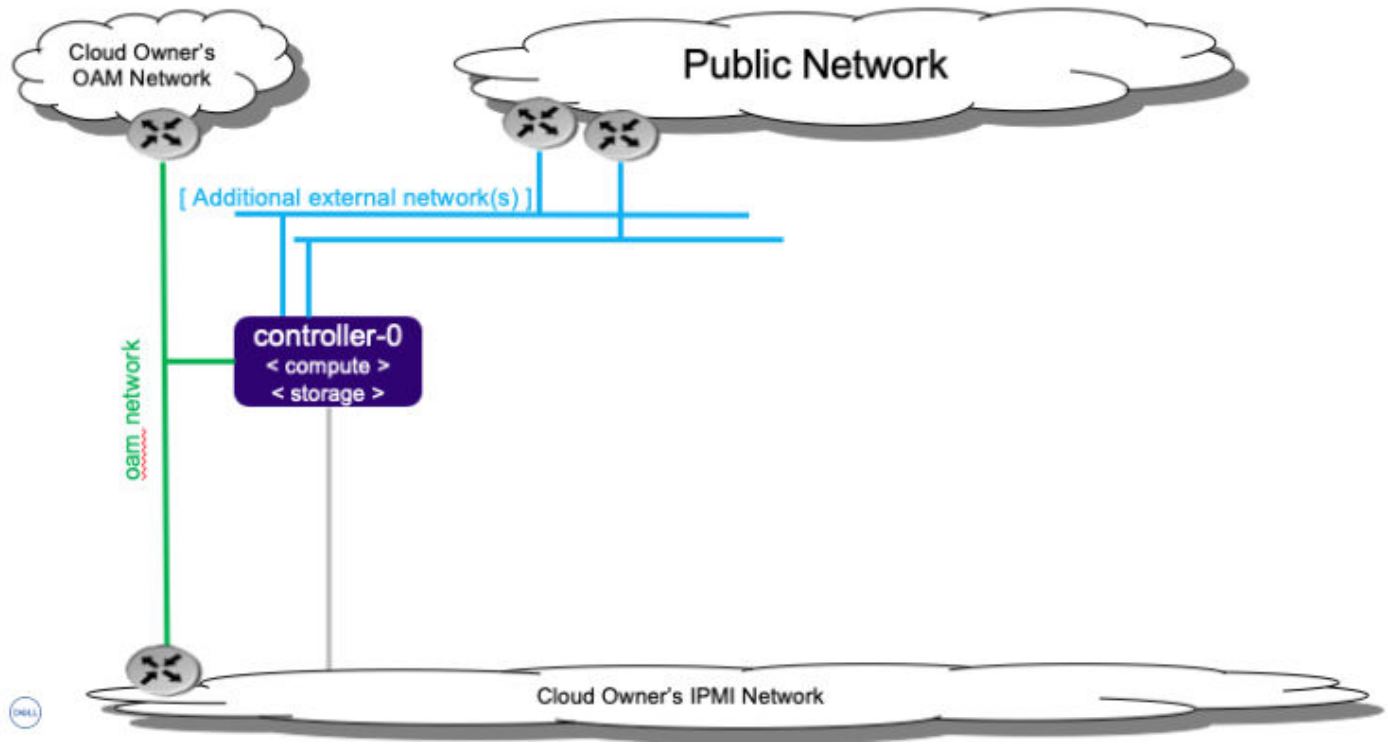

Figure 15. All-in-one Simplex single-server configuration

All-in-one Simplex configuration (dual-server HA)

Wind River Studio Cloud Platform Duplex provides a scaled-down HA cloud platform deployment option that combines controller, storage, and worker functionality on a redundant pair of hosts. The duplex configuration deploys at the Edge.

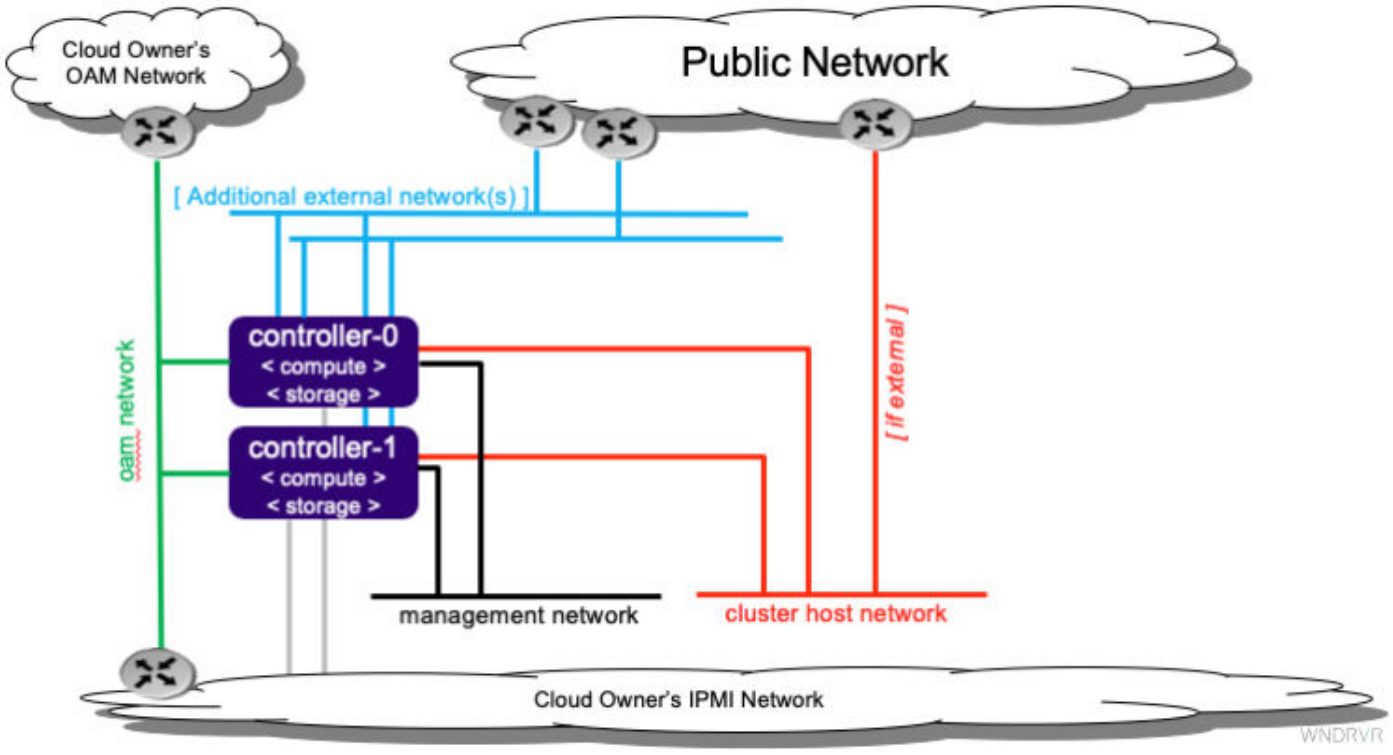

Figure 16. Dual-server HA configuration

Standard configuration with controller storage

The Wind River Studio Cloud Platform supports a small-scale deployment option using a small Ceph cluster as a back end for Kubernetes Persistent Volume Claims deployed on the controller nodes instead of dedicated storage nodes. SAN storage deployed at the Near Edge is also a popular option.

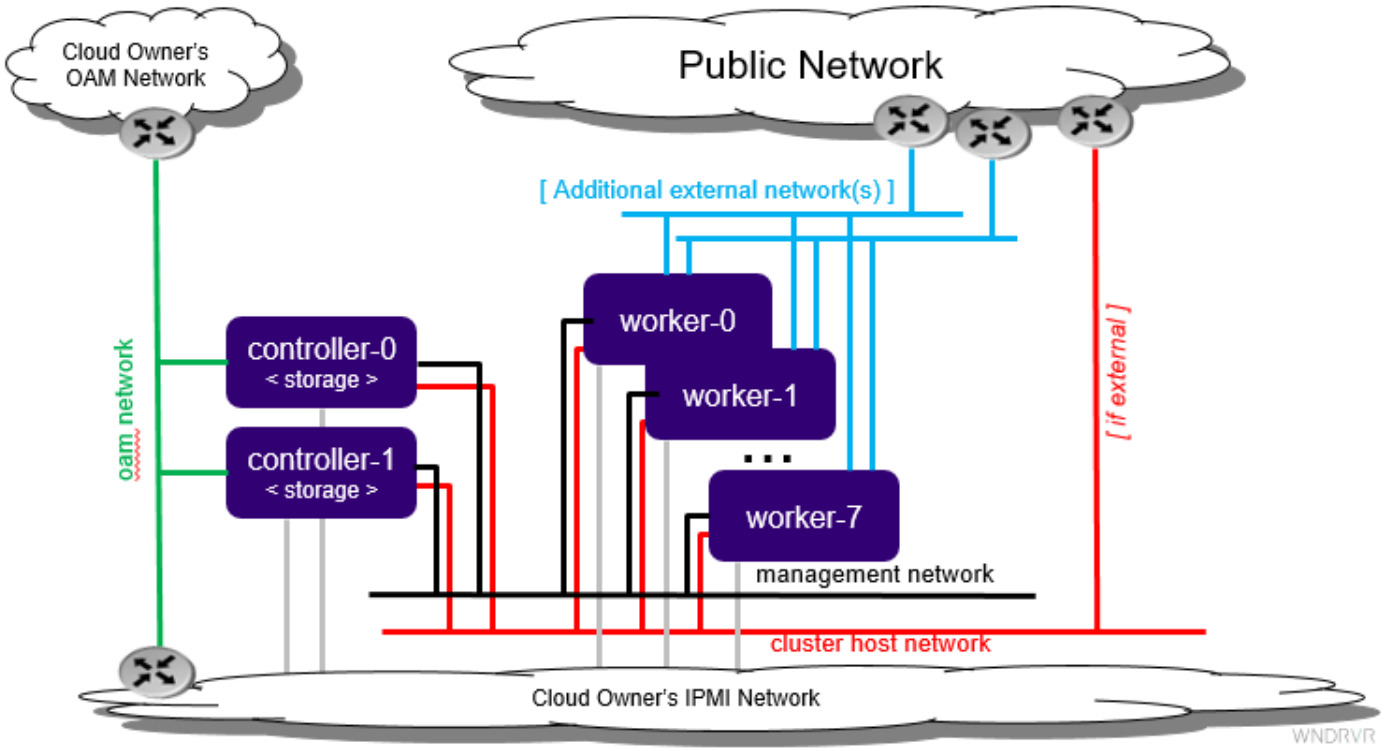

Figure 17. Standard configuration with controller storage configuration

Standard configuration with dedicated storage configuration

Wind River Studio Cloud Platform deployment with dedicated storage nodes provides the highest capacity (single region), performance, and scalability. Near Edge and larger data centers deploy this option.

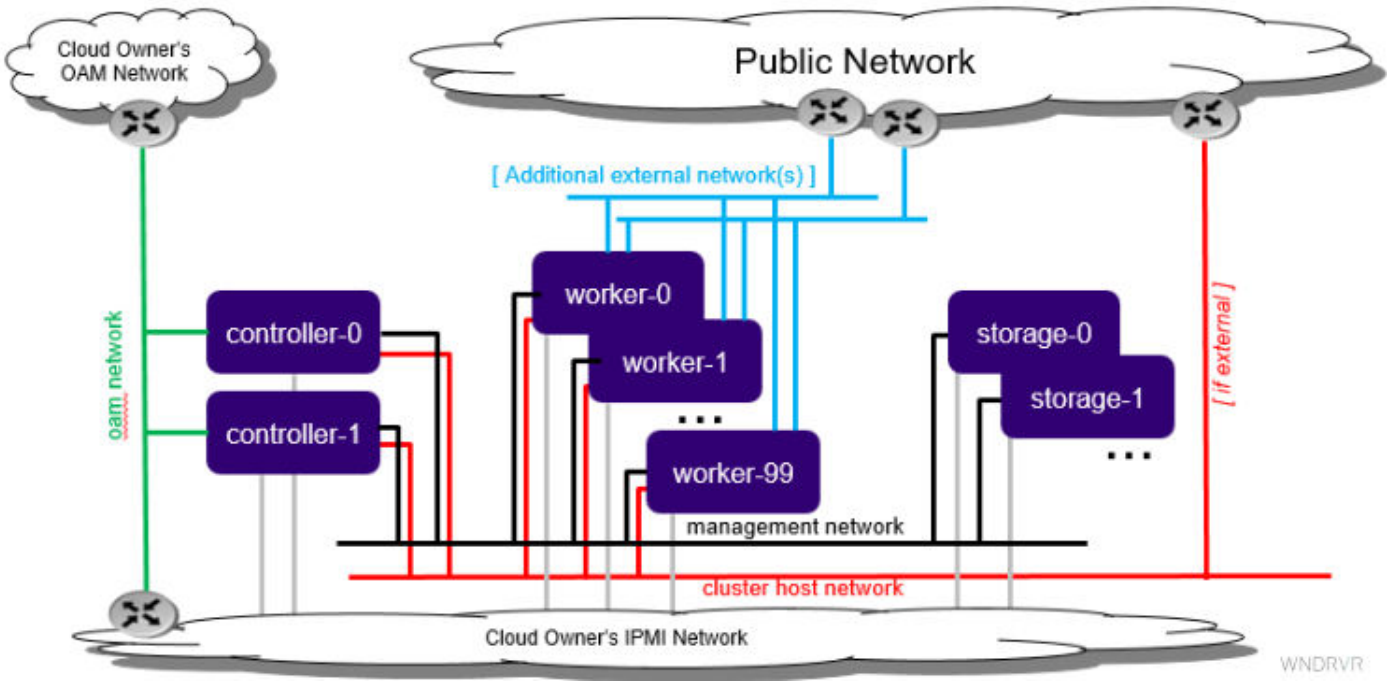

Figure 18. Standard configuration with dedicated storage configuration

Multi-Tenancy

Multi-Tenancy is compliant with optional OpenStack running Virtual Machines (VM). The Kubernetes community is still working out the details of Multi-Tenancy.

Far Edge reference components

Networking

Networking components include:

- Dell EMC Z9264-ON - Spine
- Dell EMC S5232F-ON/S5048-ON - Leaf
- Dell EMC S4048-ON/S4148T-ON - Management

NOTE: Optional NEBS model validation in the future.

Servers and switches

The following tables describe hardware and firmware versions that are only validated for the listed servers and switches.

Table 1. Dell EMC PowerEdge XR11/Dell EMC PowerEdge XR12 servers configuration

Model	System info	Firmware version
Dell EMC PowerEdge XR11/Dell EMC PowerEdge XR12	1 x Intel® Xeon® Gold 6338N	
	64 Gb 3200MT/s RDIMMs	
	2 x 480 Gb SSD SATA Mix Use	
	Intel E810-XXV DP 10/25 Gb – SFP28	20.0.18
	Intel ACC100	

Table 1. Dell EMC PowerEdge XR11/Dell EMC PowerEdge XR12 servers configuration (continued)

Model	System info	Firmware version
	Broadcom QP 25 Gb	21.80.16.95
	BIOS	1.0.2
	iDRAC9	4.40.35.0

NOTE: Dell EMC PowerEdge R650 and Dell EMC PowerEdge R750 to-be validated.

Table 2. Switches

Component	Operating system	Version
S4148T-ON (ToR)	OS10	10.5.2.7
Z9264-ON (Spine)	OS10	10.5.2.7
S5232F-ON/S5248-ON (Leaf)	OS10	10.5.2.7

Wind River Studio software requirements

The following table lists the software requirements:

Table 3. Wind River Studio software versions

Component	Version
Wind River Studio Cloud Platform	21.05
Wind River OpenStack (optional)	21.07
Wind River Studio Analytics(optional)	21.06
Wind River Conductors (optional)	21.05

7

Appendix

Topics:

Bill of materials (BOM)

Table 4. The following table lists the bill of materials for the Dell EMC PowerEdge R750

Dell EMC PowerEdge R750 - controller/ worker CU	Option description	Quantity
	Dell EMC PowerEdge R750 server	1
	Dell EMC PowerEdge R7500 motherboard	1
	Performance optimized	1
	32 GB RDIMM, 3200MT/s, dual rank	16
	C1, no RAID for HDDs/SSDs (mixed drive types allowed)	1
	Intel Xeon Gold 6330N 2.2G, 28C/56T, 11.2GT/s, 42M cache, turbo, HT (165 W) DDR4-2666	2
	Trusted Platform Module 2.0 V3	1
	Chassis with up to 8x2.5" drives	1
	PERC H345 with rear load bracket	1
	BOSS controller card + with 2 M.2 Sticks 480 GB (RAID 1),FH	1
	1.6 TB SSD SAS Mix Use 12Gbps 512e 2.5in Hot-plug AG drive, 3 DWPD,	2
	Intel X710 dual port 10 GbE SFP+, OCP NIC 3.0	1
	Intel E810-XXV dual port 10/25 GbE SFP28 Adapter, PCIe full height	2
	Dual, Hot-Plug, fully redundant power supply (1+1), 1400W, mixed mode	1
	Riser config 1, 6x8, 2x16 slots	1

Table 5. The following table lists the bill of materials for the Dell EMC PowerEdge XR12

Dell EMC PowerEdge XR12 - controller/ worker CU/DU	Option description	Quantity

Table 5. The following table lists the bill of materials for the Dell EMC PowerEdge XR12 (continued)

	Front Port Access/Rear Port Access with onboard Broadcom 57504 Quad Port 25 GbE	1
	Trusted Platform Module 2.0 V3	1
	Front Port Access Chassis with 4 SATA drivers (Onboard SATA)	1
	Intel Xeon Gold 6338N 2.2 G, 32C/48T, 11.2GT/s, 48M Cache, Turbo, HT (185 W) DDR4-3200	1
	High-Performance Chassis Thermal Configuration	1
	Performance optimized	1
	8 GB RDIMM, 3200MT/s, dual rank	8
	C20, No RAID with embedded SATA for HDDs or SSDs (mixed drive types allowed)	1
	480 GB SSD SATA Mix Use 6Gbps 512 2.5in Hot-plug AG drive	2
	1.6 TB Enterprise NVMe Mixed Use AG drive U.2 Gen4 with carrier	1
	Telco BIOS setting	1
	Front Port Access, Single, Hot-plug, Power Supply (1+0), 1400W, mixed mode	1
	Front Port Access Riser Config 0, 1 x 8, 2 x16 slots	1
	iDRAC9 Datacenter 15G with OpenManage enterprise advanced	1
	Broadcom 57504 Quad Ethernet 57504 4x25 GbE	1
	Intel E810-XXV Dual Port 10/25 GbE SFP28 Adapter, PCIe low profile	1
	Intel ACC100 (Full Height)	1
	ReadyRails Sliding Rails for 2/4-post racks without cable management arm or strain relief bar	1

Table 6. The following table lists the bill of materials for the Dell EMC PowerEdge XR11

Dell EMC PowerEdge XR11 - controller/ worker DU	Option description	Quantity
	Front Port Access/Rear Port Access with onboard Broadcom 57504 Quad Port 25 GbE	1
	Trusted Platform Module 2.0 V3	1
	Front Port Access Chassis with 4 SATA drivers (Onboard SATA)	1

Table 6. The following table lists the bill of materials for the Dell EMC PowerEdge XR11 (continued)

	Intel Xeon Gold 6338N 2.2G, 32C/48T, 11.2GT/s, 48M cache, turbo, HT (185W) DDR4-3200	1
	High-Performance Chassis Thermal Configuration	1
	Performance optimized	1
	8 GB RDIMM, 3200MT/s, dual rank	8
	C20, No RAID with Embedded SATA for HDDs or SSDs (mixed drive types allowed)	1
	480 GB SSD SATA Mix Use 6Gbps 512 2.5in Hot-plug AG drive	2
	1.6 TB Enterprise NVMe Mixed Use AG drive U.2 Gen4 with carrier	1
	Telco BIOS setting	1
	Front Port Access, Single, Hot-plug, Power Supply (1+0), 1400W, Mixed Mode	1
	Front Port Access Riser Config 0, 1 x 8, 2 x16 slots	1
	iDRAC9 Datacenter 15G with OpenManage enterprise advanced	1
	Broadcom 57504 Quad Ethernet 57504 4x25 GbE	1
	Intel E810-XXV Dual Port 10/25 GbE SFP28 Adapter, PCIe low profile	1
	Intel ACC100 (full height)	1
	ReadyRails Sliding Rails for 2/4-post Racks without cable management arm or strain relief Bar	1

Dell Technologies portfolio

Dell EMC PowerEdge XR11 servers

Dell EMC PowerEdge XR11 is a specialty edge server that is engineered to deliver powerful performance for harsh environments. It is a single-socket, 1U, short-depth, front/rear accessible server designed to support demanding edge applications such as streaming analytics, manufacturing logistics, and 5G cell processing applications.

Dell EMC PowerEdge XR12 servers

Dell EMC PowerEdge XR12 is a specialty edge server that is engineered to deliver powerful performance for harsh environments. It is a single-socket, 2U, short-depth, front/rear accessible server designed to support demanding edge applications such as streaming analytics, manufacturing logistics, and 5G cell processing applications.

Dell EMC PowerEdge R750 servers

The Dell EMC PowerEdge R750 server is a hyper-dense, two-socket, 2U rack server. The R750 server is the ideal dual-socket, 2U platform for dense scale-out cloud computing. The scalable business architecture of the R750 server maximizes application performance and provides the flexibility to optimize configurations based on the application and use case.

Dell EMC PowerEdge R650 servers

The Dell EMC PowerEdge R650 server is a hyperdense, two-socket, 1U rack server. The R650 server is the ideal dual-socket, 1U platform for dense scale-out cloud computing. The scalable business architecture of the R650 server is designed to maximize application performance and provide the flexibility to optimize configurations based on the application and use case.

Dell EMC PowerEdge XE2420 servers

The Dell EMC PowerEdge XE2420 is a specialty edge server that is engineered to deliver powerful performance for harsh environments. It is a dual-socket, 2U, short-depth, front-accessible server designed to support demanding edge applications such as streaming analytics, manufacturing logistics, and 5G cell processing applications.

Dell EMC Open Networking

Open Networking is a core element of networking strategy and mission from Dell Technologies. Open Networking separates the hardware from the operating system, giving you the choice of picking the operating system that best fits your unique network infrastructure needs. Open Networking uses standards-based open-source building blocks.

In the Dell EMC Networking portfolio, any switch model with an "-ON" suffix, such as the Dell EMC Networking Z9264-ON and the Dell EMC Networking S5248-ON switches, has Open Network Install Environment (ONIE) enabled.

Dell EMC OpenManage Enterprise

Dell EMC OpenManage Enterprise is an intuitive infrastructure management console that allows IT staff to discover, deploy, update, and monitor Dell EMC PowerEdge servers. It also enables IT administrators to view and make changes to other equipment in the data center infrastructure, including chassis, storage, and network switches on the enterprise network. OpenManage Enterprise helps users to:

- **Simplify**—OpenManage Enterprise brings the ability to handle a wide range of IT administration tasks into a single, intuitive systems management solution, reducing complexity.
- **Unify**—OpenManage Enterprise is a one-to-many systems management console, built to scale: With a single instance of OpenManage Enterprise, IT Administrators can manage up to 8,000 devices regardless of form factors, such as Dell EMC PowerEdge rack-, tower-, or modular servers, or PowerVault MD and ME storage systems, or third-party devices. Have more than 8,000 devices in your infrastructure? Just add additional instances of OpenManage Enterprise.
- **Automate**—OpenManage Enterprise helps to boost IT Admin productivity by automating tasks throughout the server life cycle. For example, OpenManage Enterprise can speed server discovery and deployment; streamline BIOS and firmware update processes, and produce customized reports.
- **Secure**—Security is a top priority with Dell EMC OpenManage solutions, including the OpenManage Enterprise console. For example, OpenManage Enterprise can detect drift from a user-defined configuration template, alert users, and remediate misconfigurations based on presetup policies.

8

References

Dell EMC PowerEdge XR11 Rack Server

https://www.dell.com/en-us/work/shop/povw/poweredge-xr11

Dell EMC PowerEdge XR12 Rack Server

https://www.dell.com/en-us/work/shop/povw/poweredge-xr12

Dell EMC PowerEdge R750 Rack Server

https://www.dell.com/en-us/work/shop/povw/poweredge-r750

Dell EMC PowerEdge R650 Rack Server

https://www.dell.com/en-us/work/shop/povw/poweredge-r650

Dell EMC PowerEdge XE2420

https://www.dell.com/en-us/work/shop/povw/poweredge-xe2420/techspecs

Dell EMC PowerSwitch S-Series Switches

https://www.dell.com/en-us/work/shop/povw/networking-s-series-25-100gbe

Dell EMC PowerSwitch Z-Series

https://www.dell.com/en-us/work/shop/povw/networking-z-series

Dell EMC OpenManage Enterprise

https://www.delltechnologies.com/en-us/solutions/openmanage/enterprise.htm

Wind River Studio Cloud Platform

https://www.windriver.com/products/cloud-platform/