

Intel-Dell Verified Reference Configuration for Virtualized Radio Access Networks on Red Hat OpenShift Container Platform



Authors

Intel

Ai Bee Lim, Timothy Miskell, Sam Deip, John Morgan, William Meigs

Dell

Nic Lemieux, Komal Bhowad, Ryan Putman, Allan Paulino, Suresh Raam

Red Hat

Rony Haddad, Federico Rossi, Brent Roswell

Table of Contents

Introduction	1
Intel VRC for vRAN	1
Partner Solution	2
Solution Software.....	3
Performance.....	3
FlexRAN™ Reference Software Benchmarks.....	4
Summary	6
Appendix.....	7

Introduction

Intel Verified Reference Configuration (VRC) are a family of workload-optimized, infrastructure solutions, based on the Intel® Xeon® Scalable processor family targeting today's complex workloads. This document describes a reference implementation for the 4th Gen Intel® Xeon® Scalable Processor with Intel® vRAN Boost, which includes a selection of Dell PowerEdge XR servers and Red Hat OpenShift Container Platform.

When network operators, service providers, cloud service providers, or enterprise infrastructure companies choose a Reference architecture for Virtual Radio Access Network (vRAN) deployment based on a 4th Generation Intel® Xeon® Scalable Processor with Intel® vRAN Boost, they should be able to deploy various network-function-virtualized applications more securely and easily than ever before.

The reference intimation helps end users simplify design choices by bundling hardware and software pieces together while making high performance more predictable. End users will spend less time, effort, and expense evaluating hardware and software options.

This document features a workload-optimized stack tuned to take full advantage of the Intel® architecture foundation.

Intel VRC for vRAN

This workload-optimized VRC is designed to minimize the challenges of infrastructure deployments and optimization efforts, ensuring the best performance for low-latency vRAN use cases. It defines the software and hardware reference stacks and includes the FlexRAN™ reference software running as a proxy workload to showcase the performance and latencies that can be achieved with this solution.

Partner Solution

Dell has built upon the success of the previous generation XR11 product with a generational improvement by the introduction of the PowerEdge XR5610. In addition, Dell has debuted the XR8000r offering configurable sled options to support the Edge network location.

XR5610

Like its predecessor, the XR5610 is a short-depth ruggedized, single socket, 1U monolithic server, designed for the Edge and Telecom workloads. Its rugged design accommodates military and defense deployments and retail AI including video monitoring, IoT device aggregation, and PoS analytics.

Improvements to the XR5610 include:

- A CPU upgrade to the recently announced 4th Generation Intel® Xeon® Scalable Processor, up to 32 cores.
- Support for the new 4th Gen Intel® Xeon® Scalable processors with Intel® vRAN Boost, an embedded vRAN accelerator within the CPU.
- Doubling of the memory bandwidth with the upgrade from DDR4 to DDR5.
- Higher performance I/O capabilities with the upgrade from PCIe Gen 4 to Gen 5.
- Dry inputs, common in remote environments to gain some insights into edge enclosure conditions, such as door open alarms, moisture detection, and more.
- Support for multiple accelerators, such as GPUs, O-RAN L1 Accelerators, and storage options including SAS, SATA or NVMe.

Figure 1: Power Edge XR5610 1U Server



XR8000r

The XR8000r is composed of a 2U, short depth, 400mm class Chassis with options to choose from 1U or 2U half-width hot-swappable Compute Sleds with up to 4 nodes per chassis. The XR8000r supports 3 sled configurations designed for flexible deployments. These can be 4 x 1U sleds, 2 x 1U and 1 x 2U sleds or 2 x 2U sleds. The Chassis also supports 2 PSU slots that can accommodate up to 5 power capacities, with both 120/240 AC and -48 VDC input powers supported.

Figure 2: Power Edge XR8000r 2U Chassis



The 1U (XR8610t) and 2U (XR8620t) Compute Sleds are based on Intel's 4th Generation Intel® Xeon® Scalable Processor, up to 32 cores, with support for both Sapphire

Rapids SP and Edge Enhanced (EE) Intel® vRAN Boost processors. Both sled types have 8 x RDIMM slots and support for 2 x M.2 NVMe boot devices with optional RAID1 support, 2 optional 25GbE LAN-on-Motherboard (LoM) ports and 8 Dry Contact Sensors through an RJ-45 connector. The 1U Compute Sleds adds support for one x16 FHHL (Full Height Half Length) Slot (PCIe Gen5).

Figure 3: XR8610t 1U Compute Sled



The 2U Compute Sled builds upon the foundation of the 1U Sled and adds support for an additional two x16 FHHL slots.

These 2 Sled configurations can create both dense compute and dense I/O configurations.

Figure 4: XR8620t 2U Compute Sled



The 2U Sled also provides the ability to accommodate GPU-optimized workloads.

This sledded architecture is designed for deployment into traditional Edge and Cell Site Environments, complementing or replacing current hardware and allowing for the reuse of existing infrastructure.

Dell Hardware and Firmware Details

Details for vRAN solution is based on the hardware configuration which showcases the best combination of latest Intel® CPU technology coupled with Intel® platform technologies, Intel® Ethernet and Intel® acceleration technologies. These technologies are integrated on the motherboard to deliver best-in-class vRAN performance with the low latency requirements using the Data Plane Development Kit (DPDK) and Baseband Device (Bbdev).

Table 1 shows the Platform Hardware Configuration and Table 2 provides the Platform Firmware details.

Table 1: Platform Hardware Configuration

Hardware	Description
Processors	20-core or 32-core, 4 th Gen Intel® Xeon® Scalable processors with Intel® vRAN Boost

Hardware	Description
DRAM	32GB Dual Rank DDR5 4800MHz 1 DIMM per channel Total Memory 256GB
Network Interface Card	Quad Port 25GbE Intel® Ethernet Network Adapter E810-XXVDA4TG1 (Gen 4 x8)
Storage	2x 960GB SSD NVMe solution as boot device

Table 2: Platform Firmware

System	Components	Version
XR8620t	BIOS	1.1.3
	iDRAC	6.10.89.00 (Build 17)
	CPLD	1.1.2
XR5610	BIOS	1.1.4
	iDRAC	6.10.25.00 (Build 05)
	CPLD	1.0.1

Notes:

- For XR8620t drivers and firmware updates go to: <https://www.dell.com/support/home/en-us/product-support/product/poweredge-xr8620t/drivers>
- For XR5610 drivers and firmware updates go to: <https://www.dell.com/support/home/en-us/product-support/product/poweredge-xr5610/drivers>

Intel BIOS Recommendation

Intel recommends using the BIOS settings for Max Performance with low latency configuration to meet the optimized deterministic performance requirements for the vRAN reference implementation.

Refer to document BIOS Settings for FlexRAN™ reference architecture based on Intel® Xeon® Processors (#640685) for information on the BIOS settings.

Dell BIOS Config

This section describes Dell BIOS Configuration required to achieve optimal performance with low power consumption for FlexRAN™ implementation. Dell recommends using Telco Optimized Profile as the workload profile.

To set Telco Optimized Profile, navigate to BIOS Setup->System Profile Settings->Workload Profile. Reboot the system after applying the workload profile. Telco Optimized Profile sets other BIOS settings necessary to meet optimized performance. These settings are described in the Appendix section.

Additional BIOS settings to be manually set after applying Telco Optimized Profile

- Set MADT Core Enumeration to Linear (Location – Processor Settings -> MADT Core Enumeration)
- Set SR-IOV Global Enable to Enabled (Location – Integrated Devices -> SR-IOV Global Enable)
- Select SR-IOV on individual NICs under use (Location – Device Settings-> (Select NIC) ->Device Level Configuration->Virtualization Mode)

Notes:

- Reboot is required after applying the above settings for changes to take effect.
- In cases where Telco Optimized Profile is set, and BIOS is upgraded to a newer version, follow the steps below for updates in Telco Optimized Profile to take effect.
 - Load BIOS defaults, and
 - Re-apply Telco Optimized Profile

Solution Software

Solution deployment is on Red Hat OpenShift 4.12 with the following FlexRAN™ Reference Software revision.

Table 3: Solution Software Version

Ingredient	SW Version Details
DPDK	22.11
FlexRAN™	23.07
Red Hat OCP	4.12
RCOS Kernel	Kernel 4.18 Realtime


Installation Guide for RHOC

Red Hat provides documentation on many methods of installation supported. Intel collaborated with Red Hat to release the FlexRAN™ Reference Software Cloud-Native Setup for Red Hat OpenShift Container Platform 4.12 Deployment Guide. Please refer to Document (#779745) for details.

Note: Contact your Intel Field Representative to access this document.

The following presents the deployment of Red Hat OpenShift via a hub and managed cluster with Zero Touch Provisioning (ZTP).

Figure 5: Deployment Configuration



Performance

This chapter aims to verify the performance metrics of the vRAN reference implementation ensuring that there is no anomaly seen. Refer to information in this section to ensure that

the performance baseline for the platform is as expected.

FlexRAN™ reference software is another key component of the reference implementation for vRAN reference design which provides hardware acceleration for the FEC baseband processing in layer one application. As such, reference implementation for vRAN reference implementation, should meet the performance published in the Performance Report for FlexRAN™.

Performance Baseline

This section includes information on a few applications that are required to be executed after the Platform is configured as per BOM, BIOS configuration, and Software Stack, as described in the earlier section.

The output provides a performance baseline on expected latency performance, memory bandwidth and jitter seen on the system.

Memory Latency Checker (MLC)

The first application is the Memory Latency Checker which can be downloaded from <https://software.intel.com/en-us/articles/intelr-memory-latency-checker>

Download the latest version and execute this application, unzip the tarball package and go into Linux* folder and execute `./mlc` or `./mlc_avx512`.

Table 4: MLC Data

Idle Latencies for Sequential Access	NUMA Node 0
NUMA Node 0	117.8 ns

Read-Write Ratio	Peak Injection Memory Bandwidth (MBPS)
ALL Reads	235535.8
3:1 Reads-Writes	217400.1
2:1 Reads-Writes	217379.9
1:1 Reads-Writes	205371.0
Stream-triad like	211210.3

Inject Delay	Latency (ns)	Bandwidth (MBPS)
0	208.18	234643.3
2	208.30	234696.7
8	206.10	233499.5
15	201.72	231812.2
50	201.09	233740.9
100	146.10	159943.3
200	128.23	76507.3
300	125.03	52909.9
400	123.93	40262.7
500	123.36	32662.0
700	122.53	23708.6
1000	122.51	16863.3
1300	121.73	13140.7
1700	121.65	10200.1
2500	121.20	7137.8
3500	121.45	5633.7

5000	121.68	3854.3
9000	121.08	2374.9
20000	121.09	1364.4

Cache-to-cache Transfer Latency	Latency (ns)
Local Socket L2→L2 HIT Latency	83.3
Local Socket L2→L2 HITM Latency	84.3

Cyclictest

For vRAN configurations, the worker node must demonstrate the system latency for the wake-up time of the threads to be approximate to 10 us and should not exceed 20us for at least 12 hours.

Testing conducted based on Table 3 solution software version, results for 16 threads is summarized in Table 5 below.


Table 5: Test Results of Cyclictest

Tool	Test Duration (Hours)	Status
Cyclictest	1	Pass
Cyclictest	2	Pass
Cyclictest	4	Pass
Cyclictest	6	Pass
Cyclictest	8	Pass
Cyclictest	12	Pass

FlexRAN™ Reference Software Benchmarks

FlexRAN™ reference software is being used as the proxy workload to ensure that the performance and latency of the Device Under Test (DUT) are as expected to our internal testing. There are numerous test cases to emulate the RAN workloads supporting narrow band, massive MIMO with different number of cells in timer mode or O-RAN Mode testing.

Table 6: FlexRAN™ Testing Methods






Table 6 shows the testing methods supported by the Intel FlexRAN™ reference software, timer mode and O-RAN mode.

The test case for 6389 and 18220 are tested under timer mode on the system described and the results are tabulated in Table 7. Pass results are established by comparing the published results.

Table 7: FlexRAN™ Software Results

System	6 Cells Massive MIMO Test Case 6389	18 Cells Narrowband 20MHz Test Case 18220
XR8620t	Pass	Pass
XR5610	Pass	Pass

Summary

The Intel virtualized Radio Access Network(vRAN) Reference implementation, Intel VRC based on 4th Gen Intel® Xeon® Scalable processors with Intel® vRAN Boost, combined with architectural improvements, feature enhancements, high memory, and IO bandwidth, is a tremendous performance and scalability advantage in today's network environments.

These processors are optimized for network and RAN workloads. The integrated FEC Accelerator offload and Intel® Ethernet E810 Network Controllers offer multiple hundredths of Gigabit per second for front haul and back haul traffic throughput.

Appendix

Table A1: User Adjustable BIOS Settings set by Telco Optimized Profile

BIOS knob	Setting	Location in BIOS Setup
AC Power Recovery	Last	System Security->AC Power Recovery
AMP Prefetch	Enabled	Processor Settings-> AMP Prefetch
APS Rocketing	Disabled	System Profile Settings->APS Rocketing
AVX ICCP Pre-Grant Level	512 Heavy	Processor Settings->AVX ICCP Pre-Grant Level
AVX ICCP Pre-Grant License	Enabled	Processor Settings->AVX ICCP Pre-Grant License
AVX P1	Level 2	Processor Settings->AVX P1
C1E	Disabled	System Profile Settings->C1E
CPU C1 Auto Demotion	Disabled	System Profile Settings->CPU C1 Auto Demotion
CPU C1 Auto unDemotion	Disabled	System Profile Settings->CPU C1 Auto unDemotion
CPU Power Management	System DBPM (TELCO)	System Profile Settings->CPU Power Management
C-States (Processor C6 or CPU C6 Report)	Enabled	System Profile Settings->C-States
Custom Uncore Frequency	1.6GHz	System Profile Settings->Custom Uncore Frequency
Dynamic L1	Disabled	System Profile Settings->Dynamic L1
Energy Efficient Policy (ENERGY_PERF_BIAS_CFG mode)	Performance	System Profile Settings->Energy Efficient Policy
Energy Efficient Turbo	Disabled	System Profile Settings->Energy Efficient Turbo
GPSS Timer	0us	System Profile Settings->GPSS Timer
Homeless Prefetch	Enabled	Processor Settings->Homeless Prefetch
LLC Prefetch	Disabled	Processor Settings->LLC Prefetch
Logical Processor (Hyper-Threading)	Enabled	Processor Settings->Logical Processor
Memory Patrol Scrub (Patrol Scrubbing)	Disabled	System Profile Settings->Memory Patrol Scrub
Monitor/Mwait	Enabled	System Profile Settings->Monitor/Mwait
OS ACPI Cx	OS Cx C2	System Profile Settings->OS ACPI Cx
Package C-States	Disabled	System Profile Settings->Package C-States
PCI ASPM L1 Link Power Management	Disabled	System Profile Settings->PCI ASPM L1 Link Power Management
Scalability	Disabled	System Profile Settings->Scalability
System Profile	Custom	System Profile Settings->System Profile
Turbo Boost (Turbo Mode)	Enabled	System Profile Settings->Turbo Boost
Uncore Frequency (Uncore frequency scaling)	Maximum	System Profile Settings->Uncore Frequency

BIOS knob	Setting	Location in BIOS Setup
Uncore Frequency RAPL	Disabled	Processor Settings->Uncore Frequency RAPL
Virtualization Technology (VMX)	Enabled	Processor Settings->Virtualization Technology
Workload Configuration	I/O Sensitive	System Profile Settings->Workload Configuration
X2APIC Mode (XAPIC)	Enabled	Processor Settings->X2APIC Mode

Table A2: Non-adjustable BIOS settings set by Telco Optimized Profile

BIOS knob	Setting
Boot Performance Mode	Max. Performance
EIST PSD Function	HW_ALL
EPP enable	Disabled
HardwarePM Interrupt	Disabled
Hardware P-States	Native with no Legacy Support
Intel SpeedStep (Pstates) Technology	Enabled
Memory Configuration	8-way interleave
Memory Paging Policy (Page Policy)	Closed
Memory POR & Memory Population POR (Enforce POR)	Enabled
Native ASPM	Disabled
PCIE AER Error Handling – PCIE Correctable Errors	Disabled
PCIE ECRC generation and checking	Disabled
Power Performance Tuning	BIOS Controls EPB
UMA Based Clustering Status	Quadrant
Virtual NUMA (MCC)	Disabled



Notices & Disclaimers

The information in this publication is provided "as is." Dell Inc. makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any software described in this publication requires an applicable software license.

This document may contain certain words that are not consistent with Dell's current language guidelines. Dell plans to update the document over subsequent future releases to revise these words accordingly.

This document may contain language from third party content that is not under Dell's control and is not consistent with Dell's current guidelines for Dell's own content. When such third-party content is updated by the relevant third parties, this document will be revised accordingly.

Copyright © 2023 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC, and other trademarks are trademarks.

of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

Learn more at www.Intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others. 0823/ABL/Wipro/PDF  Please Recycle 787009-002US