White Paper

5G FlexCore 2.0 User Plane Function with Intel® Infrastructure Power Manager on Red Hat OpenShift Container Platform using 4th Gen Intel® Xeon® Scalable Processors and Dell PowerEdge R760

Dell PowerEdge R760 server powered with 4th Gen Intel® Xeon® processors and 5G FlexCore 2.0 User Plane Function workload using Intel® Infrastructure Power Manager on Red Hat OpenShift Container Platform delivers outstanding throughput and impressive power savings.

Authors

Sarita Maini

Siva Prasad Tirulaka

Gordon Noonan

Kashish Singh

Suresh Raam

Introduction

FlexCore 2.0 User Plane Function is a 5G reference design that enables the virtualization of the 5G Core User Plane Function (UPF) and runs on a Network Function Virtualization Infrastructure (NFVI). It facilitates the evaluation and benchmarking of a virtualized 5G core mobile network data plane.

This FlexCore 2.0 5G User Plane Function (UPF) solution was deployed on a dual-socket Dell PowerEdge R760 server powered by network-optimized 4th Gen Intel® Xeon® Scalable processors using a Red Hat OpenShift Container Platform multi-node cluster. This optimized hardware and software resulted in a 5G UPF data throughput of up to **948 Gbps**, which was **94.8% of the line rate of 1000 Gbps**, with **0% packet loss**. In addition, leveraging the Intel® Infrastructure Power Manager (IPM) software with a realistic 24-hour traffic profile achieved up to a **maximum of 36% power savings** and an **average of 23% power savings** without impacting I/O throughput or packet loss.

The excellent performance and power benchmarks were achieved by using Dell PowerEdge R760 with Intel® Xeon® Platinum 8470N processors, Intel® Ethernet 800 Series Network Adapters, Intel® Infrastructure Power Manager, and Intel's FlexCore 2.0 UPF in a Red Hat OpenShift Container Platform infrastructure.

The impressive I/O throughput combined with lowering energy costs, provides the opportunity for a significant competitive advantage and a lower total cost of ownership to 5G Core Communication Service Providers.

This white paper is a product of an ongoing partnership between Intel, Dell and Red Hat to accelerate the deployment of 5G core by enabling verified product readiness to achieve the lowest total cost of ownership for performance per watt and performance per dollar.

FlexCore 2.0 User Plane Function with Intel® Infrastructure Power Manager

As commercial-off-the-shelf servers continue to become more powerful, it makes them capable of deploying commercial network solutions at a much lower cost than dedicated hardware. The 4th Gen Intel® Xeon® Scalable processors have a high core count, impressive instructions per clock (IPC), increased cache size, PCIe Gen 4 bandwidth and faster DDR5-4800 memory with higher throughput, lower latency, deterministic performance, and high performance per watt. Intel Xeon N SKUs are designed for network workloads, including 5G UPF.

The IPM reference software reduces run-time power consumption by dynamically matching CPU performance and energy consumption to real-time traffic levels, delivering power savings while maintaining key network performance metrics, such as throughput and packet loss.

The FlexCore 2.0 5G User Plane Function workload-optimized solution is designed to minimize the challenges of infrastructure deployment and optimization for the best performance with a balanced I/O across sockets for core-bound and I/O-bound workloads. It uses Intel® Xeon® Scalable processors which incorporate unique features designed especially for virtualized network workloads, leading to impressive performance gains compared to systems based on prior Intel processor generations.

FlexCore 2.0 UPF optimizes the performance of 5G core, using a data plane development kit (DPDK), Vector Packet Processing (VPP), Non-Uniform Memory Access (NUMA) optimizations and huge pages. It leverages the Dynamic Device Personalization (DDP) capability of Intel Network Interface Cards (NICs) to improve performance and reduce core utilization by offloading packet forwarding, packet processing and load balancing to Intel E810-2CQDA2 NICs. The usage of the ICE COMMS DDP package resulted in an **increase in I/O throughput of up to 20%.**

The Dell PowerEdge R760 Rack Server coupled with Red Hat OpenShift Container Platform (RH OCP) delivers the ultimate performance and versatility, using optimal configurations available to meet the demanding specification requirements for 5G Core. This flexible 2U server uses less space and is power efficient.

The processors tested in the 2-socket server were Intel® Xeon® Platinum 8470N with 52 physical cores, 104 logical cores per NUMA node with hyper-threading enabled, 208 logical cores per server, a maximum turbo frequency of 3.70 GHz, PCIe Gen 5.0, 80 PCIe lanes, 300W TDP per processor and 8 channels of DDR5-4800 MT/s memory.

Dell PowerEdge rack servers help to build a modern infrastructure that minimizes challenges and drives business success. Design optimized for a mix of demanding workloads with high database and analytics capabilities, the Dell R760 provides performance and versatility for demanding applications.

Red Hat® OpenShift® Container Platform is a foundation for building and scaling containerized applications. It is a cloud-based Kubernetes platform and comes with an automatic install to get Kubernetes up and running as quickly as possible. The OS that gets installed on the nodes when the cluster is created is an immutable OS named "Red Hat Enterprise Linux CoreOS" (RHCOS). A cluster has a web console which includes an operator hub. Operators can be used to enable automated installation and configuration of applications.

Hardware:

- Dell PowerEdge R760 server with Intel® Xeon® Platinum 8470N processors, as a worker node in a multi-node Red Hat OpenShift Container Platform cluster
- 4 x Intel® Ethernet Network Adapter E810-2CQDA2
- 2 x Intel® Ethernet Network Adapter E810-CQDA2
- Traffic Generator Servers with 5 x Intel® Ethernet Network Adapter E810-2CQDA2
- 3 control plane node servers for Red Hat OpenShift Container Platform cluster
- OC Client Server for OpenShift command-line interface (CLI)

Software:

- Red Hat OpenShift Container Platform
- Red Hat Enterprise Linux CoreOS
- Intel Infrastructure Power Manager
- Node Feature Discovery Operator
- SR-IOV Network Operator
- Intel Ethernet Operator
- Intel ICE DDP COMMS Package
- Trex Traffic Generator
- OpenShift command-line interface (CLI)

Application:

5G FlexCore 2.0 User Plane Function

Note:

The Intel® Verified Reference Configuration for 5G FlexCore 2.0 UPF with Red Hat OpenShift Container Platform on Intel® Xeon® Scalable Processors using Intel® Infrastructure Power Manager has hardware and software configuration details, and step-by-step instructions for deploying and benchmarking FlexCore 2.0 UPF and Intel® Infrastructure Power Manager.

Red Hat OpenShift Container Platform

The following figure shows the benchmarking setup for the 5G UPF test configuration.

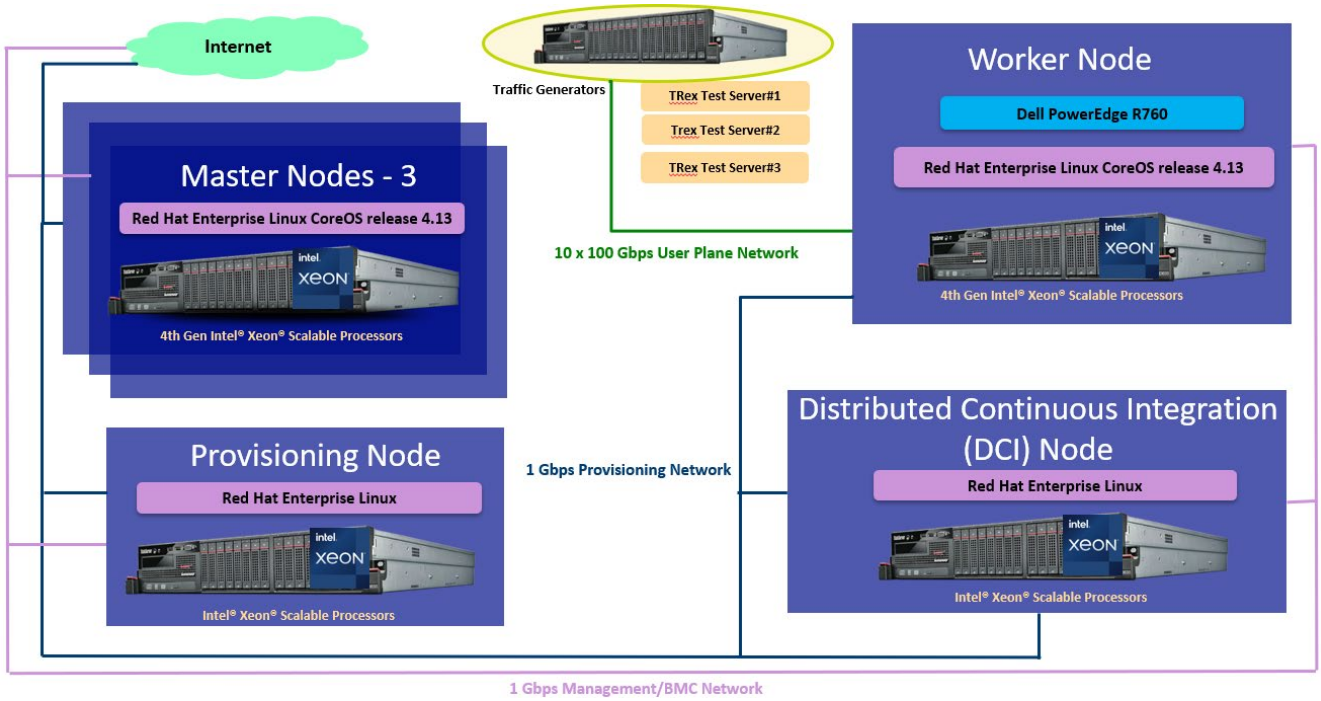

Figure 1. Red Hat OpenShift Container Platform

5G User Plane Network Test Setup

The following figure shows the benchmarking setup for the 5G UPF test configuration.

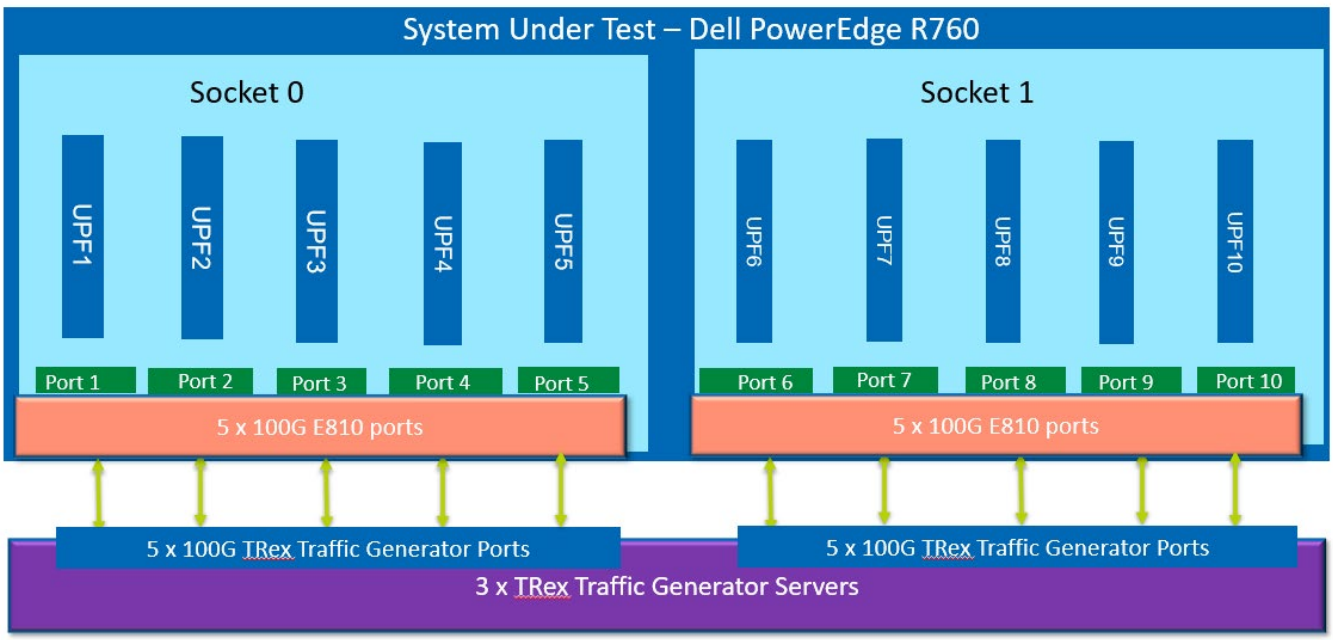

Figure 2. 5G User Plane Network Test Setup

5G FlexCore 2.0 UPF Benchmarks

This section shows the benchmarks of 5G FlexCore 2.0 UPF with a Red Hat OpenShift Container Platform cluster, on a 2-socket Dell PowerEdge R760 as a cluster worker node. Each socket in the Dell PowerEdge R760 server was populated with Intel® Xeon® Platinum 8470N processors. The TDP per processor was 300 Watts, resulting in a TDP of 600 Watts per 2-socket server. The throughput and power benchmarks using Intel® Infrastructure Power Manager with the IPM agent running are compared with benchmarks without the IPM agent.

I/O Throughput with and without IPM

The Dell PowerEdge R760 server with two 4th Gen Intel® Xeon® 8470N processors, a maximum 1000 Gbps Line Rate (10 x 100G Intel® Ethernet Network Adapter E810-2CQDA2 ports), 650 Bytes Packet Size for both Uplink (UL) and Downlink (DL) and a traffic ratio of UL/DL 1:3 with 0% Packet Loss with IPM and without IPM, showed up to 948 Gbps throughput, which was 94.8% of the line rate.

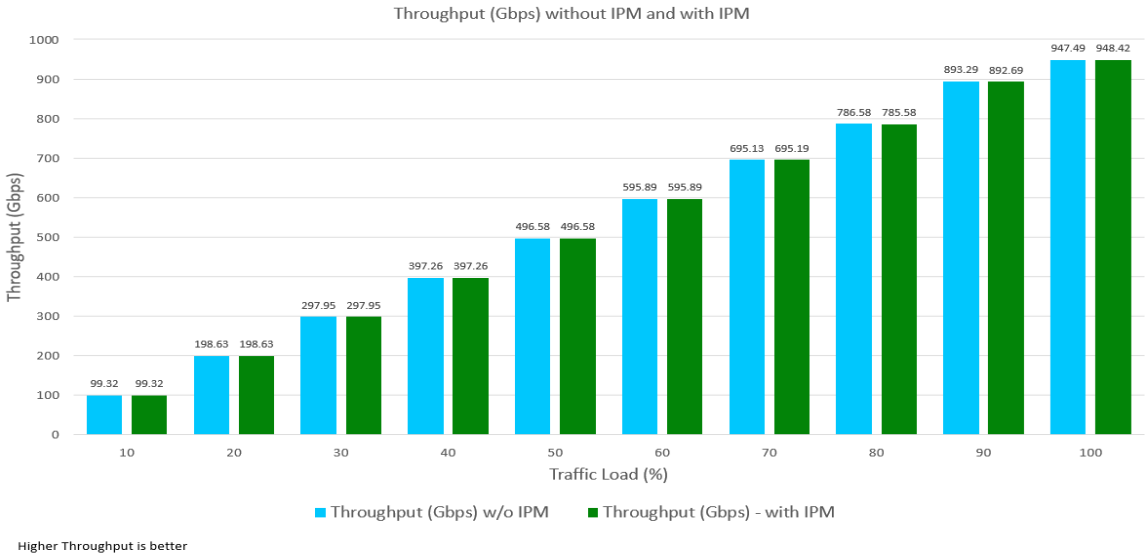

Figure 3. FlexCore 2.0 UPF I/O Throughput (Gbps) with and without IPM[1,2]

I/O Throughput with 24-hour Traffic Profile

The I/O throughput of the 5G FlexCore 2.0 UPF solution using a realistic 24-hour traffic profile with a packet loss of 0% and a packet size of 650 Bytes, demonstrated no degradation in I/O throughput and yet impressive power savings when IPM was used.

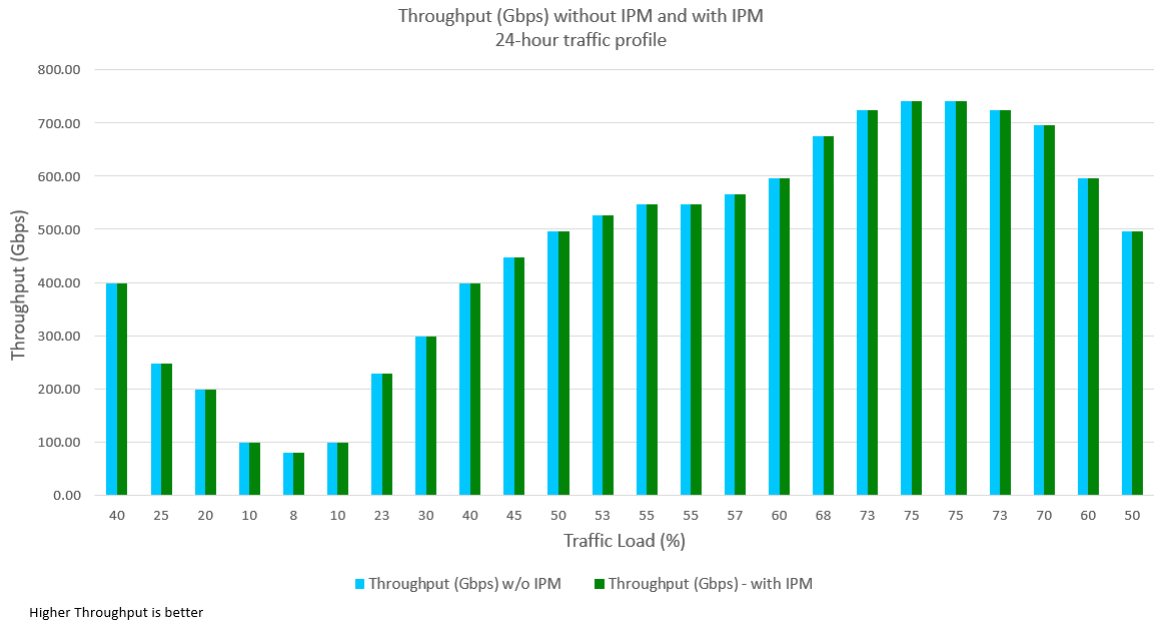

Figure 4. FlexCore 2.0 UPF I/O Throughput (Gbps) using 24-hour traffic profile with and without IPM[1,2]

Power Usage with 24-hour Traffic Profile

The results of the 5G FlexCore 2.0 UPF benchmarks in a RH OCP worker node with and without IPM, using a realistic 75% traffic profile, showed up to a **maximum of 36% power savings and an average 23% power savings with IPM**, with a packet loss of 0% and a packet size of 650 Bytes. Without IPM, the power usage was ~100% even in an idle state with no traffic. Using IPM which leverages the DPDK telemetry logical cores "busyness" to determine CPU load for the monitored active UPF worker cores, the CPU frequency was downshifted or upshifted by IPM, depending on the core utilization which changed based on the traffic load percentage, resulting in much lower power usage.

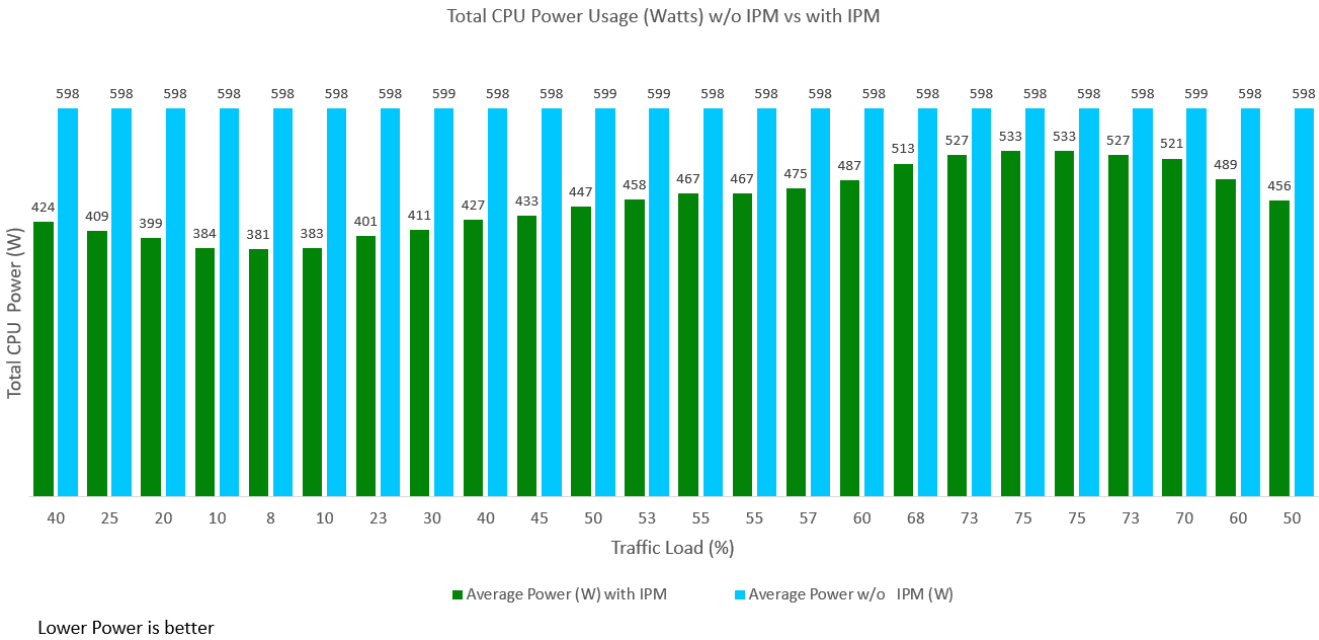

Figure 5. Server Power Savings with 24-hour traffic profile using IPM[1,2]

Summary

Network-optimized Intel Xeon Scalable processors help Communication Service Providers to manage and optimize 5G Core UPF data traffic, with energy-efficient, fast I/O throughput and high density of connections per server. The processors are designed for high throughput and low latency for network workloads from 5G core to the edge.

The Dell PowerEdge R760 mainstream 2U dual-socket stack server is a powerful and dependable server which brings PCIe Gen 5, DDR5, and integrated acceleration to the Dell PowerEdge server line.

The outstanding 5G FlexCore UPF performance of up to **948 Gbps** which was **94.8% of the maximum 1000 Gbps line** rate, on a single Dell PowerEdge R760 server with Intel® Xeon® Platinum 8470N network-optimized processors, in addition to meeting and exceeding 5G UPF high-speed requirements, also enables Telco companies to reduce the footprint of number of servers needed in 5G Core and Edge 5G platform solutions.

The advantage of Intel® Ethernet Controller 800 Series Dynamic Device Personalization for Telecommunications (COMMS DDP) was that it dramatically increased the bandwidth for packet throughput, by offloading packet processing and load balancing to the Network Interface Cards (NICs), resulting in lower core utilization, which was used to increase the I/O throughput by up to 20%. The COMMS DDP software provides the benefit of lower Total Cost of Ownership (TOC) which reduces the number servers needed for 5G Core workloads.

The usage of the Intel Infrastructure Power Management tool with FlexCore 2.0 UPF running on Intel 4th Gen Xeon Scalable processors demonstrated a **maximum of 36% power savings and an average 23% power savings**, using a realistic 24-hour traffic profile. The power consumption was reduced with no impact on performance metrics such as I/O throughput or packet loss.

With 5G Core software using DPDK-based applications without power savings technologies, CPU power consumption is always at a maximum. With IPM using DPDK telemetry to compute the core "busyness", every monitored core has the optimal power at any time. Therefore, in 5G applications and workloads based on DPDK, usage of the Intel Infrastructure Power Management tool becomes very important for real-time power consumption and meeting sustainability goals.

ADDENDUM: System Configuration

iDRAC Settings

System Information	
System Model Name	PowerEdge R760
System BIOS Version	System BIOS Version
System Manufacturer	Dell Inc.
System CPLD Version	1.0.5
BIOS Version	1.4.4
iDRAC Version	iDRAC9
iDRAC Firmware Version	6.10.80.00

BIOS Settings

Processor Settings

	Current Value
Logical Processor	Enabled
CPU Interconnect Speed	Maximum data rate
Virtualization Technology	Enabled
Kernel DMA Protection	Disabled
Directory Mode	Enabled
Adjacent Cache Line Prefetch	Enabled
Hardware Prefetcher	Enabled
DCU Streamer Prefetcher	Enabled
DCU IP Prefetcher	Enabled
Sub NUMA Cluster	Disabled
MADT Core Enumeration	Linear
UMA Based Clustering Status	Quadrant
UPI Prefetch	Enabled
XPT Prefetch	Enabled
LLC Prefetch	Enabled
Dead Line LLC Alloc	Enabled
Directory AtoS	Disabled
AVX P1	Normal
Dynamic SST-Performance Profile	Disabled
SST-Performance Profile	Operating Point 3 | P1: 2.1 GHz, TDP:300w, Core Count:52
Intel SST-BF	Disabled
Intel SST-CP	Disabled
x2APIC Mode	Enabled
AVX ICCP Pre-Grant License	Disabled

Dell Controlled Turbo

Number of Cores per Processor	All
CPU Physical Address Limit	Enabled
AMP Prefetch	Disabled
Homeless Prefetch	Enabled
Uncore Frequency RAPL	Enabled
Processor Core Speed	2.10 GHz
Processor Bus Speed	16 GT/s
Local Machine Check Exception	Enabled
CPU Crash Log Support	Disabled
Family-Model-Stepping	6-8F-8
Brand	Intel(R) Xeon(R) Platinum 8470N
Level 2 Cache	52x2 MB
Level 3 Cache	99840 KB
Number of Cores	52
Microcode	0x2B0004B1
Family-Model-Stepping	6-8F-8
Brand	Intel(R) Xeon(R) Platinum 8470N
Level 2 Cache	52x2 MB
Level 3 Cache	99840 KB
Number of Cores	52
Microcode	0x2B0004B1

System Profile Settings

	Current Value
System Profile	Custom
CPU Power Management	OS DBPM
Memory Frequency	Maximum Performance
Turbo Boost	Enabled
Energy Efficient Turbo	Disabled
C1E	Enabled
C-States	Enabled
Memory Patrol Scrub	Standard
Memory Refresh Rate	1x
Uncore Frequency	Dynamic
Dynamic Load Line Switch	Enabled
Energy Efficient Policy	Performance
Monitor/Mwait	Enabled
Workload Profile	Not Configured
CPU Interconnect Bus Link Power Management	Disabled
PCI ASPM L1 Link Power Management	Disabled
Workload Configuration	IO Sensitive

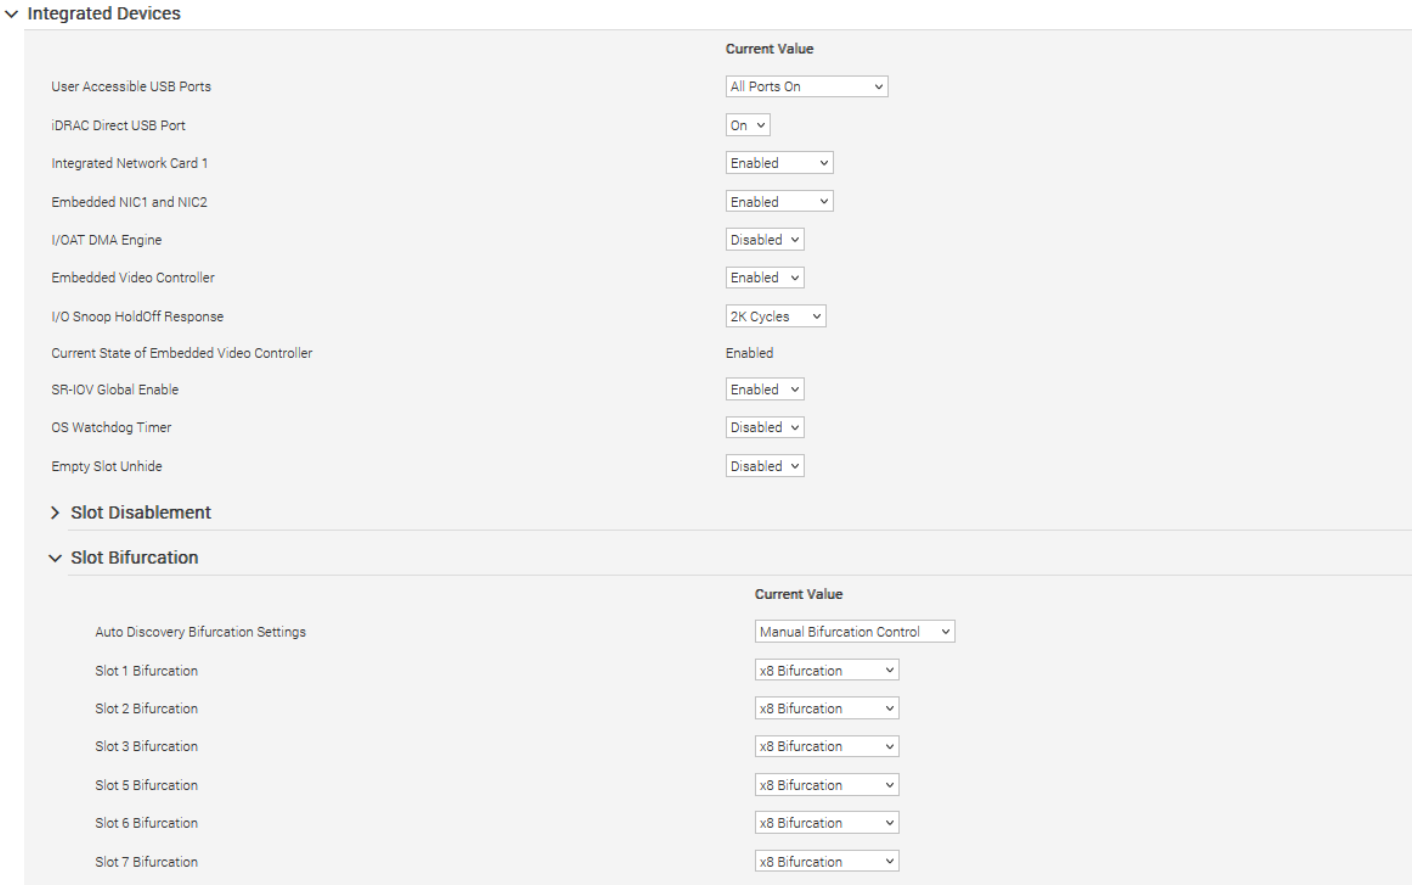

[1] **Notices & Disclaimers**

Performance varies by use, configuration, and other factors. Learn more on the Performance Index site. 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.

[2] **Configuration**

Test by Intel as of 3/17/2024. 1-node, 2x Intel® Xeon® Platinum 8470N, 52 cores, HT On, Turbo On, Total Memory 2048GB (32x64GB DDR5 4800 MT/s [4400 MT/s]), microcode 0x2b0004b1, Red Hat Enterprise Linux CoreOS 413.92.202311151359-0 (Plow), kernel 5.14.0-284.41.1.el9_2.x86_64, DDP ICE COMMS package version 1.3.45.0.

0524/SM/PDF ♻Please Recycle 820829-001US