White Paper

Communication Service Providers (CoSPs), Mobile Network Operators (MNOs)

Intel-Lenovo Verified Reference Configuration for 5G Core with Intel® Infrastructure Power Manager on Red Hat OpenShift Container Platform

5G Core with Intel Verified Reference Configuration on Lenovo SR630 V4 with Intel® Xeon® 6

Authors

Intel:

Sarita Maini

Alex Mui

Gordon Noonan

Siva Prasad Tirulaka

Kashish Singh

Brian Meaney

Lenovo:

Mark Wallis

Lumir Dambrosio

Chadie Ghadie

Kelly Robinson

Table of Contents

Introduction

As the 5G era accelerates, service providers and enterprises are under increasing pressure to deliver ultra-reliable, low-latency networks that can scale to meet growing data demands. The 5G Core (5GC) network is the brain of next-generation mobile infrastructure, requiring high-performance computing capabilities, exceptional power efficiency, and intelligent orchestration.

This white paper explores how the **Intel® Xeon® 6 processor family**, combined with the **Lenovo ThinkSystem SR630 V4** server platform, provides a transformative solution for deploying and optimizing 5G Core networks. We delve into the technical advantages of this pairing, including performance, power efficiency, AI-enhanced network acceleration, and optimal Total Cost of Ownership (TCO).

Intel and Lenovo collaborated to maximize network performance and power efficiency with an enhanced standalone 5G Core NFVI implementation with infrastructure and technologies which delivered a **throughput of up to 92.6% of the line rate** with a demanding 5G Core data plane workload, using Intel Xeon 6780E processors (Intel® Xeon® E-cores). In addition, this verified reference configuration (VRC) solution achieved up to an **average of 36.14% CPU power saving** with a telco 24-hour traffic profile and up to **maximum 39.05% power saving** when idle (no traffic) using the Intel® Infrastructure Power Manager (IPM), with minimal impact on I/O throughput. The solution demonstrates ecosystem readiness of the hardened hardware, firmware, and software that enables end users to integrate their applications on top of a verified platform configuration to achieve high throughput, deterministic performance and power efficiency that are required for 5G transformative workloads.

As 5G networks grow in complexity and scale, the need for **automated, intelligent, and fine-grained power control** becomes critical. **Intel® Infrastructure Power Manager (IPM)** is a software-based platform that enables **orchestrated, telemetry-driven power and performance optimization** across infrastructure nodes, making it an ideal complement to **Intel® Xeon® 6 processors** in 5G Core environments.

Intel® Xeon® 6: Revolutionary Performance for 5G Core

The Intel® Xeon® 6 processor family introduces two microarchitectures:

- **Performance-core (P-core) architecture** for latency-sensitive workloads
- **Efficient-core (E-core) architecture** for throughput-optimized workloads

5G Core Use Cases Optimized by Intel® Xeon® 6 Processors with E-cores

- User Plane Functions (UPF) with ultra-low latency and high packet throughput
- Control Plane Functions (AMF, SMF) with concurrent session scale
- Real-time analytics and policy enforcement

Key Technical Features:

Feature	Benefits for 5G Core
Up to 144 E-cores per socket	High-density compute for scalable packet processing
Intel® IAA (In-Memory Analytics Accelerator)	Faster analytics for telemetry and observability
Intel® QAT (QuickAssist Technology)	Hardware acceleration for encryption/compression
PCIe 5.0 and DDR5	High-speed IO and memory bandwidth for demanding 5G workloads
Telemetry & Power Management	Real-time power tuning and performance scaling

Intel® Infrastructure Power Manager (IPM) Features

- The Intel® Infrastructure Power Manager (IPM), dynamically aligns CPU performance to real-time traffic conditions, reducing power usage while maintaining throughput and zero packet loss.
- IPM power manages the Red Hat OpenShift Container platform worker cores, by dynamically scaling P-states of the cores based on core utilization, which dynamically adjusts clocks and voltage
- IPM can also be used to statically reduce the uncore frequency, using the PM config file
- IPM changes the frequency of worker cores, based on a DPDK telemetry metric - "busyness" of cores
- Telemetry-driven orchestration enables workload-aware scheduling
- Efficient-core Xeon 6 SKUs deliver leading performance per watt as well as performance per dollar

Lenovo ThinkSystem SR630 V4

The **Lenovo ThinkSystem SR630 V4** is a 1U two-socket rack server purpose-built to harness the full potential of Xeon 6 processors. The ThinkSystem SR630 V4 powered by Intel® Xeon® 6 processors and optimized for high performance per watt, is the ideal choice to handle the intense processing, networking, and storage requirements of modern 5G networks.

Key Features:

- Dual-socket support for P-core and E-core Xeon 6
- Enhanced airflow and thermal design for dense edge/cloud deployments
- NVMe and PCIe Gen 5 support for high-speed storage and networking

Advantages for 5G Core

- High performance in a compact form factor for dense edge data centers
- High-speed I/O to support 25/50/100G NICs for 5G Core workloads, especially 5G UPF
- Ensures high availability and reliability with redundant power supplies and cooling fans
- Supports virtualization and containerization technologies, allowing operators to improve resource utilization

System Setup

The test setup included a Lenovo ThinkSystem SR630 V4 server as a Single Node OpenShift (SNO) cluster, as part of a standard Red Hat OpenShift Platform configuration which consisted of a single node that included the functionality of control plane node, master node and worker node.

The 2-socket server had Intel Xeon 6780E processors (Intel Xeon 6 E-cores), with two Intel E810-2CQDA2 NICs and one Intel E810-CQDA2 for OCP 3.0 NIC.

In this test setup, only one processor was used for the 5G User Plane.

The I/O throughput achieved was 92.6% of the line rate of 500 Gbps with standalone 5G Core and OCP 4.16 by utilizing 80 cores from one socket on a 2-socket server, with a Packet Size of 650 Bytes and Packet Loss of 0%. The unused 208 cores may be utilized for Control Plane network functions or other applications

IPM was used to power manage the 80 worker cores, by dynamically scaling P-states of the cores based on core utilization. The Verified Reference Configuration included BIOS, Linux kernel and NUMA optimizations. All worker cores used were NUMA node 0 cores. For the best performance with 0% packet loss, all 100 Gbps NICs were connected to NUMA node 0.

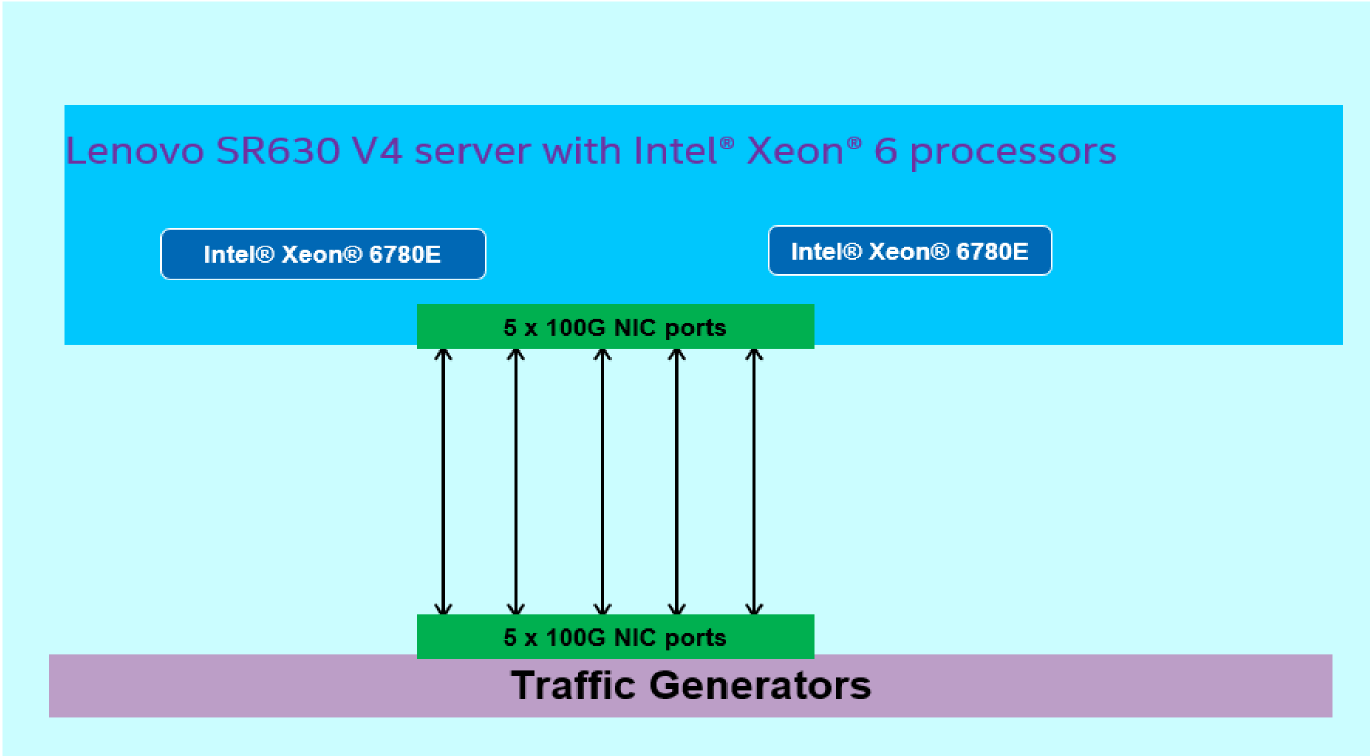

Figure 1. System Setup

Hardware and Software Details

Hardware Settings	
Processor	2 x Intel® Xeon® 6780E processor, 2.20Ghz, 144C
System	Lenovo ThinkSystem SR630 V4
Baseboard	Lenovo SB27B70074
Chassis	Lenovo Rack Mount Chassis
CPU Model	Intel(R) Xeon(R) 6780E
Architecture	x86_64
Microarchitecture	SRF_SP
L3 Cache	108 MiB
Cores per Socket	144
Sockets	2
Hyperthreading	N/A
CPUs	288
Intel Turbo Boost	Enabled
Base Frequency	2.3GHz
All-core Maximum Frequency	3.0GHz
Maximum Frequency	2.3GHz
NUMA Nodes	2
Prefetchers	L2 HW: Enabled, L2 Adj: Enabled, DCU HW: Enabled, DCU IP: Enabled, DCU NP: Enabled
Accelerators Available [used]	DLB 4 [0], DSA 4 [0], IAA 4 [0], QAT 4 [0], vRAN Boost 0 [0]
Installed Memory	512GB (16x32GB DDR5 6400MT/s [6400MT/s])
Hugepagesize	1048576 kB
Transparent Huge Pages	always
Automatic NUMA Balancing	Enabled
Network Interface Cards	2 x Intel® Ethernet Network Adapter E810-2CQDA2, 1 x Intel® Ethernet Network Adapter E810-CQDA2 for OCP 3.0
Disk	1x 447.1G Micron_7450_MTFD

Software Settings	
BIOS	IHE107B-1.10

Software Settings	
Microcode	0x3000270
Scaling Governor	performance
Scaling Driver	acpi-cpufreq
OS - Red Hat Enterprise Linux CoreOS	Red Hat Enterprise Linux CoreOS 416.94.202408132101-0
Kernel	5.14.0-427.30.1.el9_4.x86_64
Microcode	0x3000270
Drivers	5.14.0-427.31.1.el9_4.x86_64 In-tree kernel driver (ice and iavf)
DPDK	23.03.0
Intel® Infrastructure Power Manager	N/A
Red Hat OCP SR-IOV Network Operator	4.16.0-202410300306
Node Feature Discovery Operator	4.16.0-202410251436
Traffic Generator version	TRex v3.0
Run Method: ex. cold (fresh-boot), warm (post-boot after few back to back iterations)	Warm (5 iteration)
Packet Size (B)	650B
Base Core Frequency - MHz	2300
Max Frequency - MHz UFreq[CMPT:IO:IO] – Compute Die, IO Die 0, IO Die 1	2200:2200:2200

BIOS Settings	
System Configuration and Boot Management -> System Settings -> Workload Profile	Custom
System Configuration and Boot Management -> System Settings -> Power -> Power/Performance Bias	Platform Controlled
System Configuration and Boot Management -> System Settings -> Power -> Platform Controlled Type	Performance
System Configuration and Boot Management -> System Settings -> Power -> Workload Configuration	I/O Sensitive
System Configuration and Boot Management -> System Settings -> Power -> ASPM	Disabled
System Configuration and Boot Management -> System Settings -> Processors -> Turbo Mode	Enabled
System Configuration and Boot Management -> System Settings -> Processors -> Energy Efficient Turbo	Disabled
System Configuration and Boot Management -> System Settings -> Processors -> CPU P-state Control	Cooperative without Legacy
System Configuration and Boot Management -> System Settings -> Processors -> C1 Enhanced Mode	Disable
System Configuration and Boot Management -> System Settings -> Processor -> Uncore Frequency Scaling	Enabled
System Configuration and Boot Management -> System Settings -> Processors -> Intel Speed Select	Config 1
System Configuration and Boot Management -> System Settings -> Processors-> Intel Virtualization Technology	Enabled
System Configuration and Boot Management -> System Settings -> Processors -> Hardware Prefetcher	Enabled
System Configuration and Boot Management -> System Settings -> Processors -> Adjacent Cache Prefetch	Enabled
System Configuration and Boot Management -> System Settings -> Processors -> DCU Streamer Prefetcher	Enabled
System Configuration and Boot Management -> System Settings -> Processors -> DCU IP Prefetcher	Enabled
System Configuration and Boot Management -> System Settings -> Processors -> XPT Prefetcher	Enabled
System Configuration and Boot Management -> System Settings -> Processors -> UPI Prefetcher	Enabled
System Configuration and Boot Management -> System Settings -> Processors -> L1 Next Page Prefetcher	Enabled
Socket Configuration -> Advanced Power Management Configuration -> Hardware PM State Control -> Hardware P-States	Native Mode with No Legacy Support
Socket Configuration -> Advanced Power Management Configuration ->CPU – Advanced PM Tuning-> Energy Perf BIAS -> Power Performance Tuning	OS Controls EPB
Miscellaneous Configuration -> SR-IOV Support	Enabled
System Configuration and Boot Management -> System Settings -> Devices and I/O Ports -> SRIOV	Enabled
System Configuration and Boot Management -> System Settings -> Devices and I/O Ports -> Intel VT for Directed I/O (VT-d)	Enabled

5G Core Benchmarks

This section shows the benchmarks of 5G Core with Red Hat OpenShift Container Platform, using a Lenovo ThinkSystem SR630 V4 server as a Single Node OpenShift (SNO) cluster which consisted of a single node that included the functionality of control plane, master and worker node. The 2-socket server had Intel Xeon 6780E processors. Each processor had 144 cores, 144 threads, a processor base frequency of 2.3 GHz, a maximum turbo frequency of 3 GHz, PCI Express Revision 5.0, a maximum number of PCIe lanes of 88 and TDP of 320 W per processor. IPM was used to power manage the worker cores by dynamically scaling P-states of the cores based on core utilization.

I/O Throughput without and with IPM

The I/O throughput achieved was up to 92.6% of the line rate of 500 Gbps with standalone 5G Core and OCP 4.16 using 80 cores from one socket on a 2-socket server, with a Packet Size of 650 Bytes and Packet Loss of 0%.

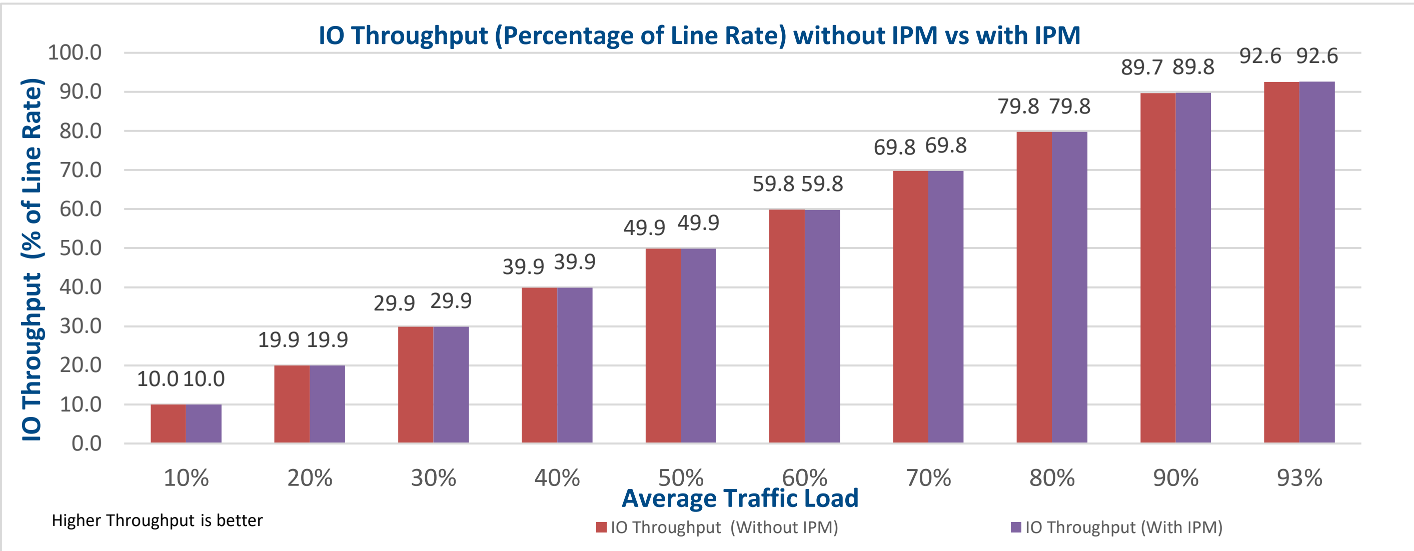

Figure 2. I/O Throughput (% of Line Rate) without and with IPM

Power Usage with 24-hour Traffic Profile

The results of the 5G Core benchmarks in a Red Hat OCP worker node with and without IPM, using a realistic telco 24-hour traffic profile, showed an average 36.14% power savings with IPM, with a packet loss of 0% and a packet size of 650 Bytes. Using IPM which leverages the DPDK telemetry cores "busyness" to determine actual 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 a reduction of power consumption in real time.

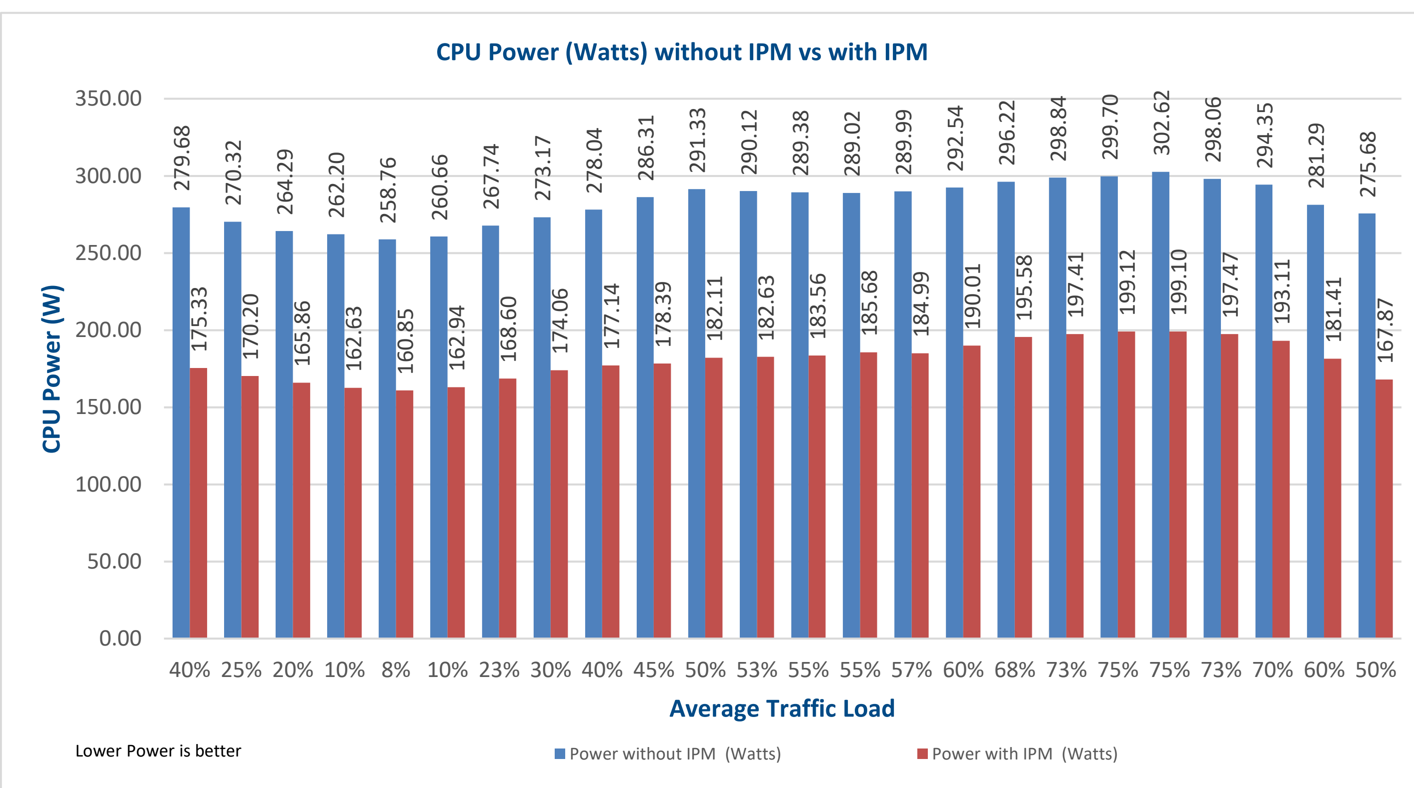

Figure 3. CPU Power Savings with 24-hour traffic profile using IPM

CPU Frequency Scaling with IPM

The CPU frequency for each worker code was dynamically increased or decreased by IPM, depending on the core utilization which changed based on the traffic load percentage, resulting in a low CPU frequency across most of the traffic load throughout the 24-hour period and bringing impressive power savings. The functionality of IPM is based on this dynamic scaling of P-states, based on a DPDK telemetry metric called "busyness" of cores.

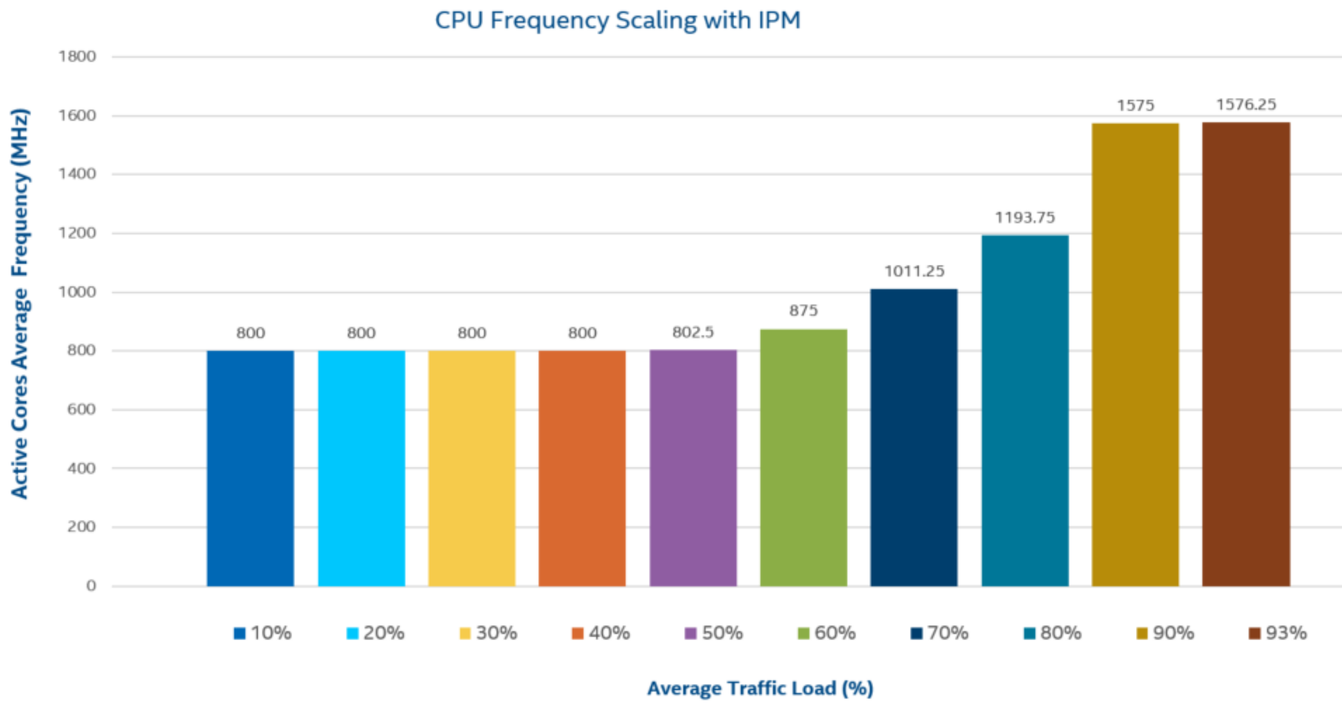

Figure 4. CPU Dynamic Scaling of P-States with IPM

CPU Utilization

The average core utilization of all the worker cores is shown with and without IPM. The core utilization is higher as the traffic load increases, and the core frequency is scaled up or down by IPM based on usage. DPDK telemetry is used to determine the CPU load.

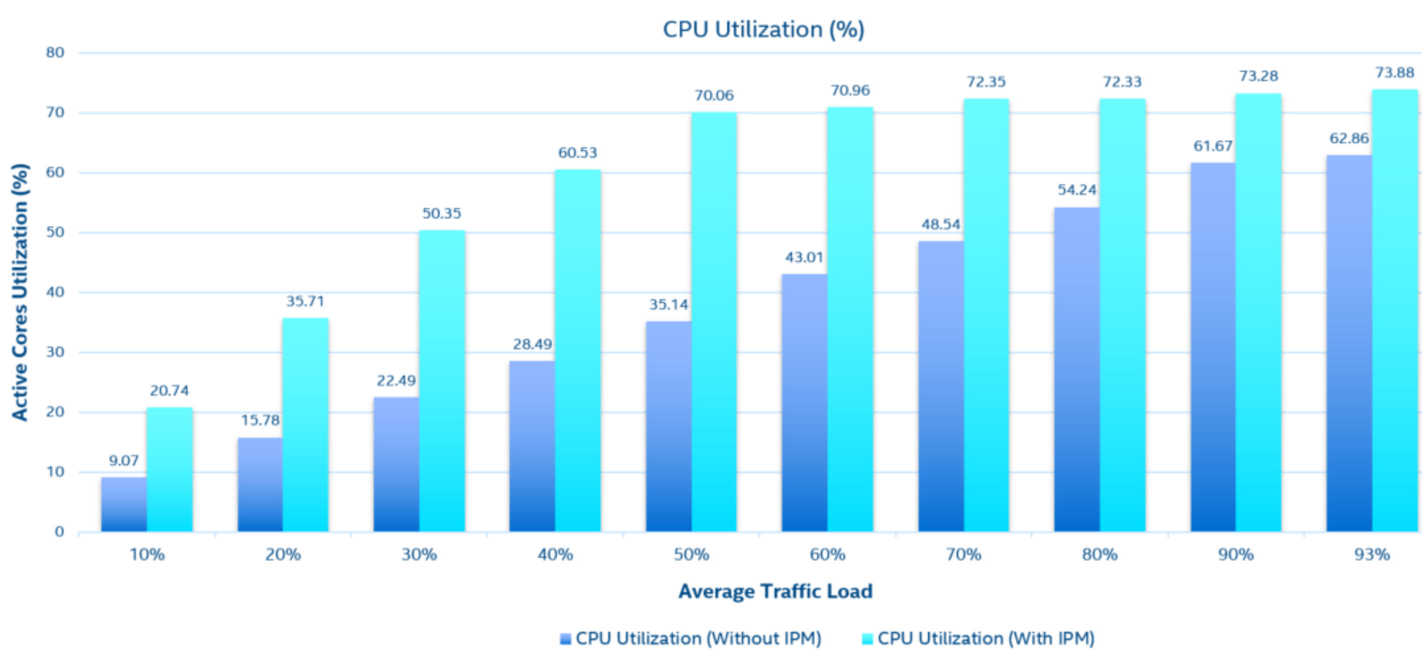

Figure 5. CPU Utilization

Summary

This Intel Verified Reference Configuration defines the hardware, software, and BIOS settings required for a 5G Core workload on a Lenovo SR630 V4 server with Intel Xeon 6 processors with E-cores. The collaboration between Intel and Lenovo enables CoSPs to quickly integrate their 5G solution on top of a verified platform configuration to achieve impressive throughput and power efficiency.

The Lenovo SR630 V4 server with Intel Xeon 6780E processors and Intel Network Adapters achieved up to 92.6% of the line rate at 650-byte packets with zero packet loss, regardless of Intel Infrastructure Power Manager (IPM) usage. With IPM enabled, power savings reached a maximum of up to 39% and averaged 36.14% without impacting throughput or packet loss. By dynamically adjusting CPU P-states and CPU frequency, based on DPDK telemetry, IPM ensures optimal efficiency for next-generation 5G deployments.

Notices & Disclaimers

0725/SM/PDF ♻ Please Recycle 366330-001US