

Turbocharging Casa vBNG Axyom™ with Intel Power Technology on 4th Generation Intel® Xeon® Scalable Processors with Dell PowerEdge

Authors

Intel

Timothy Miskell, Andrew Duignan, Ai Bee Lim, John Morgan, William Meigs, Paul Mannion

Casa

Ruinan Hu, Sudhir Krishnan, Yanhe Fan

Dell

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

Red Hat

Rony Haddad, Federico Rossi

Table of Contents

Introduction	1
Casa Systems vBNG Axyom™	2
vBNG Workload Overview	2
Intel Solution Components	2
Casa vBNG Axyom™ Test Setup	3
Hardware and Software BOM	3
BIOS Settings	4
vBNG Test Workload Topology	4
vBNG Test Results	5
Conclusion and Summary	10
<i>Appendix A – Dell PowerEdge R660 Riser Configuration</i>	10
<i>Appendix B – BIOS Settings</i>	11

Introduction

To meet the growing demand of increasing performance requirements of software-defined networking (SDN) and network functions virtualization (NFV), the 4th Generation Intel® Xeon® Scalable processors and Dell PowerEdge server offer targeted feature enhancements for networking workloads. These include support for PCIe Gen5, improved AVX-512 performance, along with power savings via Intel® Power Manager (IPM), catering to increased throughput needs while optimizing power consumption.

This whitepaper focuses on a virtualized Broadband Network Gateway (vBNG), a key workload for broadband aggregation in Virtualized Network Function (VNF) that requires not only high throughput and resiliency but also efficiency in terms of power consumption as this workload is often deployed at scale in remote central offices which are often themselves power constrained.

This document highlights the performance of the Casa vBNG Axyom™ workload on 4th Generation Intel® Xeon® Scalable processors, generating up to 433.5 Gbps supporting 40K subscribers on a single Dell PowerEdge R660 server at 0.5% packet loss. Broadband traffic loads are well understood and quite predictable- the period of the day is often not very busy (traffic) while the “busy hours” occur in the late afternoon. Leveraging Intel power technologies, the application with this knowledge released quite excellent results demonstrating up to 37% total wall power savings during off-peak conditions, as well as an average of up to 29.3% over a 24-hour period.

Under the hood of the Dell PowerEdge server for this study, key components include 4th Generation Intel® Xeon® Scalable processors, the open-source Data Plane Development Kit (DPDK), Intel® Power Manager, as well as Intel® Ethernet 800 Series Controllers.

Casa Systems vBNG Axyom™

Casa Systems vBNG Axyom™ is a virtualized Broadband Network Gateway (vBNG). For more information about Casa Systems vBNG Axyom™, visit: <https://www.casa-systems.com/solutions/bng/>

For many more details about Intel® Verified Reference Configuration (VRC), visit: <https://networkbuilders.intel.com/>

vBNG Workload Overview

A vBNG is typically deployed in a central office (CO) at edge locations close to end users and serves to aggregate data traffic from multiple home and office locations while providing access to the core of the network. The portion of the network connecting to home and office locations is referred to as the access network, while the portion of the network connecting to datacenters is referred to as the core network. Packets emanating from the access network destined for the core network are collectively referred to as upstream (US) or uplink (UL) traffic. Packets emanating from the core network destined for the access network are collectively referred to as downstream (DS) or downlink (DL) traffic.

In terms of typical IP-over-Ethernet (IPoE) traffic patterns, upstream traffic consists of Q-in-Q encapsulated IPv4 or IPv6 packets consisting of an outer service provider VLAN (S-VLAN) 802.1Q tag and an inner customer VLAN (C-VLAN) tag that allows routing to specific locations within the access network. In this case, access network subscribers connect to the BNG using either DHCPv4 or DHCPv6 protocol. Furthermore, downstream traffic typically consists of single VLAN tag-encapsulated IPv4 packets. The ratio of downstream to upstream traffic typically favors downstream traffic and oftentimes ranges from 8:1 to 9:1 for the DL:UL traffic mix. (<https://www.ncta.com/whats-new/the-asymmetric-nature-of-internet-traffic>). Due to recent developments in the increased number of subscribers working from home and the increased amount of video streams as part of the UL, the traffic mix is currently trending towards lower ratios.

Additional traffic shaping is typically performed by BNGs to implement specific Quality-of-Service (QoS) along with differentiated services (DSCP) for specific subscribers. For example, bandwidth rate limiting may be performed to differentiate specific tiers of subscribers, for example, home versus office locations, customers that may opt for higher data rates, as well as higher priority traffic classes carrying VoIP data. BNG deployments also typically include support for firewalls and ACLs to limit not only which access network locations can connect to the BNG, but also which TCP/UDP ports are allowed within the access network.

Historically BNGs have been deployed as fixed function hardware. However, with the advent of Network Function Virtualization, BNGs have been virtualized, i.e., running as a VNF orchestrated and deployed in a Virtual Machine (VM), as well as containerized, i.e., as a CNF orchestrated and deployed within a pod. Deploying virtualized or containerized appliances

allows for running the VNF/CNF on standard dual socket IA servers, for example, a Dell PowerEdge R660 server populated with 4th Generation Intel® Xeon® Scalable Processors.

Given the ever-increasing network bandwidth demands of existing subscribers, relatively fixed CAPEX and OPEX expenditures by ISPs and COSPs, narrowing margins, along with growing concerns regarding overall power consumption, it is critical to provide a repeatably deployable vBNG solution that not only delivers high performance at scale, but also in a power conscientious manner.

Intel Solution Components

The 4th Generation Intel® Xeon® Scalable processors introduce support for PCIe Gen 5, which offers double the bandwidth capabilities compared to PCIe Gen 4, going from a theoretical 256 MT/s to 512 MT/s per lane (PCI Express, 2022). This, in turn, enables the use of the Intel® Ethernet 800 Series Controllers and next-generation ethernet adapters. Specifically, E810-2CQDA2 ethernet adapters provide two ports of full 100 GbE bandwidth, offering a total of 200 GbE of aggregate throughput per NIC. On a dual socket Dell PowerEdge R660 server, it is possible to populate one E810-2CQDA2 NIC per socket, as well as an OCP form factor E810-CQDA2 on the first socket, providing for a total of 5x 100 GbE ports, for a maximum 500GbE of throughput per server.

With reference to high throughput vBNG VNFs, the ability to support up to 500 GbE of throughput per server allows the ability to aggregate more subscribers per server with the potential to consolidate multiple deployed vBNG VNFs onto a small subset of servers thus reducing the overall total cost of ownership of the solution.

Furthermore, 4th Generation Intel® Xeon® Scalable processors introduce support for DDR5 memory. At a platform level, the server supports up to eight memory channels. The memory frequency supports up to 4800 MT/s compared to 2667 MT/s with the previous generation. The increased memory frequency in this case provides lower latency and higher bandwidth. BNGs can be memory intensive due to the need to support high throughput data rates at scale, and being able to, for example, specifically configure QoS settings, track tens of thousands of subscriber sessions, and handle copious flow table rules.

In addition, 4th Generation Intel® Xeon® Scalable processors provide enhanced support for Intel® Advanced Vector Extensions up to 512 bits (Intel® AVX-512). As an example, Data Plane Development Kit (DPDK) based applications leveraging Intel® AVX-512 can significantly parallelize the processing of network packets in a single instruction multiple data (SIMD) fashion. AVX-512 enables the CPU to act on eight times as much data in one instruction as a non-vectorized instruction, leading to greater efficiency and better performance. In this case, VNFs or CNFs deployed with Single Root I/O Virtualization (SR-IOV) can leverage the Intel® Ethernet Adaptive Virtual Function (Intel® Ethernet AVF) Poll Mode Driver (PMD), which offers support for Intel® AVX-512 since DPDK version 21.11.

(<https://networkbuilders.intel.com/solutionslibrary/intel-avx->

[512-instruction-set-for-packet-processing-technology-guide\)](#)

It is critical to provide a repeatably deployable vBNG solution that not only delivers high performance at scale, but also in a power conscientious manner.

In terms of deployed BNG VNFs, the combination of Intel® AVX-512 along with RSS or FDIR support allows applications to significantly scale performance to meet ever increasing throughput requirements. With RSS and FDIR, applications no longer need to perform software load balancing, but rather offload this operation to the underlying NIC hardware and remove potential CPU bottlenecks. Furthermore, packet processing operations that must be performed by the CPU can ideally leverage the underlying Intel® AVX-512 ISA to perform operations on multiple packets from a single instruction. With the ever-increasing number of protocols introduced both within the network core and at the edge, Intel® 800 Series Ethernet Controllers continue to evolve Dynamic Device Personalization (DDP) technology, which offers the means to customize the input parser of the NIC to recognize novel header types. Since DDP was introduced the default DDP package now provides the immediate ability to parse Q-in-Q packets. In terms of BNG VNFs, the use of either the default or the Comms DDP profiles allows the application to parse further into the packet and in term enable RSS or FDIR support for Q-in-Q packet types.

In terms of reducing the overall power consumption for a given workload, there are several tunings available on 4th Generation Intel® Xeon® Scalable Processors. The most readily available options with respect to the processor include the use of C-states, P-states, as well as the uncore frequency. The use of core level C-states allows dormant cores to enter sleep states, for example, C6, to avoid drawing excessive power. Similarly, the use of package level C-states allows an entire socket to enter a sleep state. In addition, the P-state, i.e., the frequency of a given core, can be varied to alter the overall level of power consumption. This is especially useful in the case of DPDK based workloads that use a poll-mode driver methodology and attempt to set the CPU utilization of each core at 100%. If a DPDK based BNG, with for example Intel® Speed Step Technology enabled, does not require the use of the highest possible frequency, i.e., during off-peak conditions, the P-state for the corresponding core can be reduced to conserve power. The uncore frequency in general determines the speed of the uncore which affects I/O operations occur, or memory transactions coming into the last level cache. Similarly, during off-peak conditions, the uncore frequency can also be reduced for further potential power savings on the platform.

Intel® Ethernet 800 Series Controllers offer support for up to 256 queues for virtual functions along with Receiver Side Scaling (RSS) and Flow Director (FDIR). As a result, VNFs that enable RSS can perform a Toeplitz or XOR hash on specific fields within the parsed packet headers to distribute, i.e., load balance, traffic across multiple queues to scale performance. Furthermore, VNFs can enable FDIR to steer packets to individual queues based upon specific fields within the parsed packet headers.

The following chapter provides an overview of the test configuration to test the Casa Axyom™ vBNG application on a system under test (SUT) which in this case consisted of 4th Gen Intel® Xeon Scalable Processor and Intel Network Adapters. The SUT in this test setup is Dell PowerEdge R660 server.

Figure 1 Virtualized BNG Network Topology

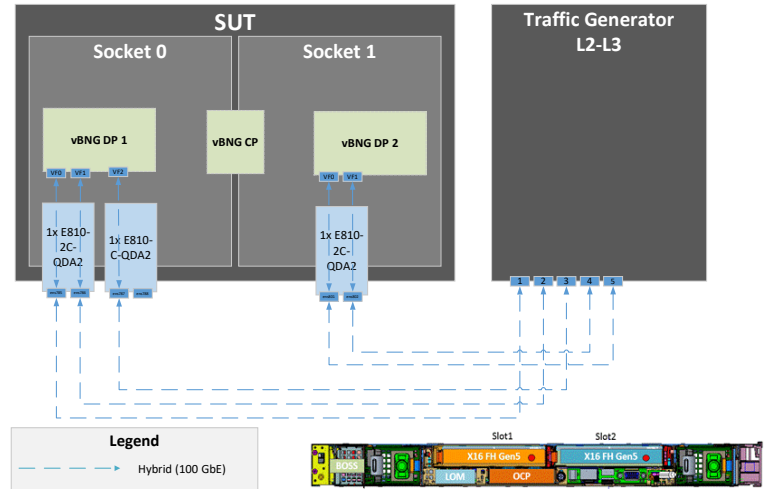


Figure 1 presents an overview of the connection setup between the SUT and the traffic generator. In this case, the traffic generator consists of an Ixia XGS12 hardware chassis capable of generating L2 to L3 traffic. Specifically, the SUT is connected back-to-back with 5x 100 GbE ports on the Ixia, and therefore the generator can simulate up to 500 Gbps of both access and network IPoE traffic.

Hardware and Software BOM

Table 1 details the hardware BOM for the SUT.

Table 1: Hardware Bill of Materials

Component	Specification
Server	Dell PowerEdge R660
Central Processing Unit (CPU)	2x Intel® Xeon® Gold 6428N CPU 32 core 1.8 GHz 185W
Memory	16x 64 GiB (1024 GiB) DDR5 4000 MT/s
Network Interface Card	2x 100 GbE Intel® Ethernet Network Adapter E810-2CQDA2 1x 100 GbE Intel® Ethernet Network Adapter E810-CQDA2
LAN On Motherboard (LOM)	NetXtreme BCM5720 Gigabit Ethernet PCIe
OS Drive	2x 960GB SK hynix NVMe PE8010 RI M.2 at BOSS-N1 Controller
Storage Drive	3x 960GB KPM6XRUG960G
BIOS	1.2.1

Component	Specification
Microcode	0x2b000190

Table 2: Software Bill of Materials

Software Stack	RHEL 8.6
Host OS	4.18.0-425.19.2.el8_7.x86_64
Libvirt	8.0.0
QEMU	6.2.0
NVMe	1.0
vBNG DP Version	3.3.1
vBNG CP Version	3.3.1
DPDK	21.11
ice	1.10.1.2.2
i40e	4.18.0-425.19.2.el8_7.x86_64
iavf	DPDK 21.11
ixgbe	4.18.0-425.19.2.el8_7.x86_64

Table 2 presents the software stacks for Red Hat Enterprise Linux* (RHEL*), which is based on Verified Reference Configuration for Network Function Virtualization Infrastructure (NFVI) Forwarding Platform on Red Hat* Enterprise Linux*.

Note: Refer to Intel® Reference Architecture for NFVI Forwarding Platform on 4th Gen Intel® Xeon® Scalable Processors on Red Hat* Enterprise Linux* with vAGF Workload for details of Solution at

https://networkbuilders.intel.com/solutionslibrary/intel-reference-architecture-for-nfvi-forwarding-platform-on-4th-gen-intel-xeon-scalable-processors-on-red-hat-enterprise-linux-with-vagf-workload#utm_source=newsletter&utm_medium=email&utm_campaign=4th_Gen_Intel_Xeon_Scalable_Processors

BIOS Settings

Intel BIOS Recommendation

Intel recommends using the BIOS Settings for deterministic performance with turbo enabled to meet performance requirement for the Packet Processing throughput workload.

Refer to document BIOS Settings for Intel Wireline, Cable, Wireless and Converged Access Platform (Document ID #747130) Chapter 2.0 for detail information on the BIOS settings recommendation.

Note: Please contact your Intel Field Representative for access to documentation.

Dell BIOS Configuration

This section details Dell BIOS Configuration required to achieve optimal performance with turbo enabled for high throughput packet processing. Dell recommends using NFVI FP Optimized Turbo Profile as the workload profile.

To set NFVI FP Optimized Turbo Profile, navigate to BIOS **Setup->System Profile Settings->Workload Profile**. Reboot the system after applying the workload profile. NFVI FP Optimized Turbo Profile sets other BIOS settings necessary to meet optimized turbo performance for high throughput packet processing performance. These settings are described in the Appendix section.

In addition to setting NFVI FP Optimized Turbo Profile, set MADT Core Enumeration setting under Processor Settings to Linear.

vBNG Test Workload Topology

Table 3: Resource Allocation per vBNG CP VM

Component	Specification
vCPUs	2C4T
Memory	25 GiB
Storage	16 GiB
MGMT Interface	1
DP Interface	N/A

Table 3 presents the resource allocation for the vBNG CP VM.

Table 4: Resource Allocation per vBNG DP VM

Component	Specification
vCPUs	24C48T
Memory	50 GiB
Storage	11 GiB
MGMT Interface	1
DP Interfaces	Up to 3x 100 GbE VFs (1x VF per PF)

Table 4 presents the resource allocation for each of the vBNG DP VMs.

Table 5: Network Topology Configuration Settings

Component	Specification
Total Subscriber Counts	40K
Core Network Endpoint Count	254
Access Network Packet Type	Q-in-Q Encapsulated DHCPv4/Ipv4 IpoE
Core Network Packet Type	VLAN Encapsulated Ipv4
MGMT Interface	1
DP Port Configuration	Hybrid Mode (Access & Core)
Number of Network Topologies	5x100 GbE Networks
Number of Ports per Network Topology	1
UL:DL Line Rate Ratio	1:8
UL/DL Traffic Mix	512 B/768 B

Table 5 presents the traffic profile configured on the traffic generator as part of the benchmark. In this case, the traffic profile consists of a total of 40,000 subscribers with an UL / DL packet profile of 512 B / 768 B. Note that the packet profile represents a more real-world scenario.

Note that all ports are configured in hybrid mode, in which case each virtual function may receive and process both access and core network traffic.

Figure 2 Virtualized BNG Network Topology

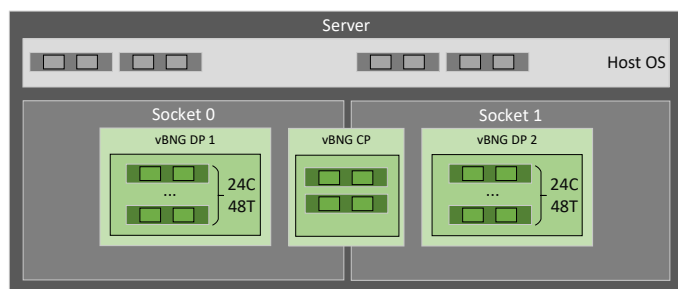


Figure 2 presents the vCPU mapping for the CP VM along with both vBNG DP VMs. Note that in this case hyper-threading (HT) is enabled on the platform, with vCPUs allocated to each of the vBNG DP VMs in HT sibling pairs. In addition, note that 4C8T are reserved for the hypervisor.

Figure 3 Virtualized BNG Network Topology

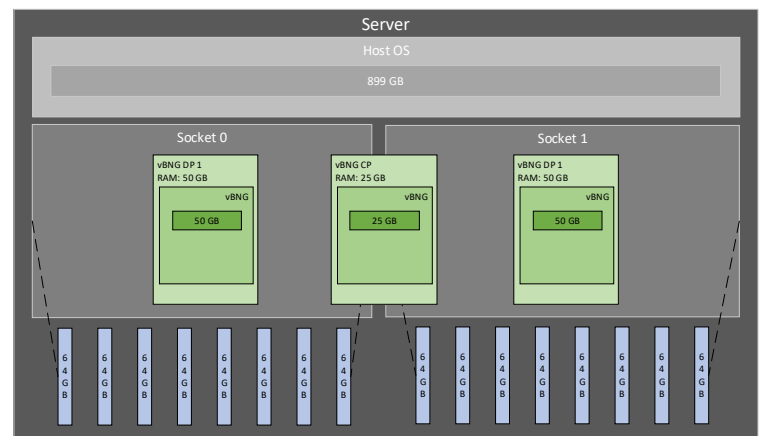


Figure 3 presents the memory resource allocation for the CP VM along with both vBNG DP VMs.

The following table presents the key performance indicators (KPIs) along with the boundary conditions for each of the corresponding KPIs where applicable.

Table 6: KPI

KPI	Units / Boundary Condition
Throughput	Gbps / RFC2544 at less than 0.5%
Packet Loss	% / Less than 0.5 %
Latency	μ s / Less than 500 μ s
Power	W /

For the benchmark a representative traffic curve captured over the course of 24 hours from a real-world production environment is applied to the Casa Systems Axyom™ vBNG VNF. In the baseline case, Intel® Speed Step Technology is enabled, with the P-state settings set to the default range of 800 MHz – 3800 MHz for each core, with the uncore frequency locked to 1.8 GHz to maximize throughput. For the power optimized case, for each hour within the course of the 24-hour traffic curve, the P-state setting is reduced such that there is no loss of performance in terms of throughput and the packet loss rate, as well as the latency, remains within the boundary conditions. Similarly, the uncore frequency is reduced such that there is no loss of performance in terms of throughput and the packet loss rate as well as the latency remain within the set boundary conditions.

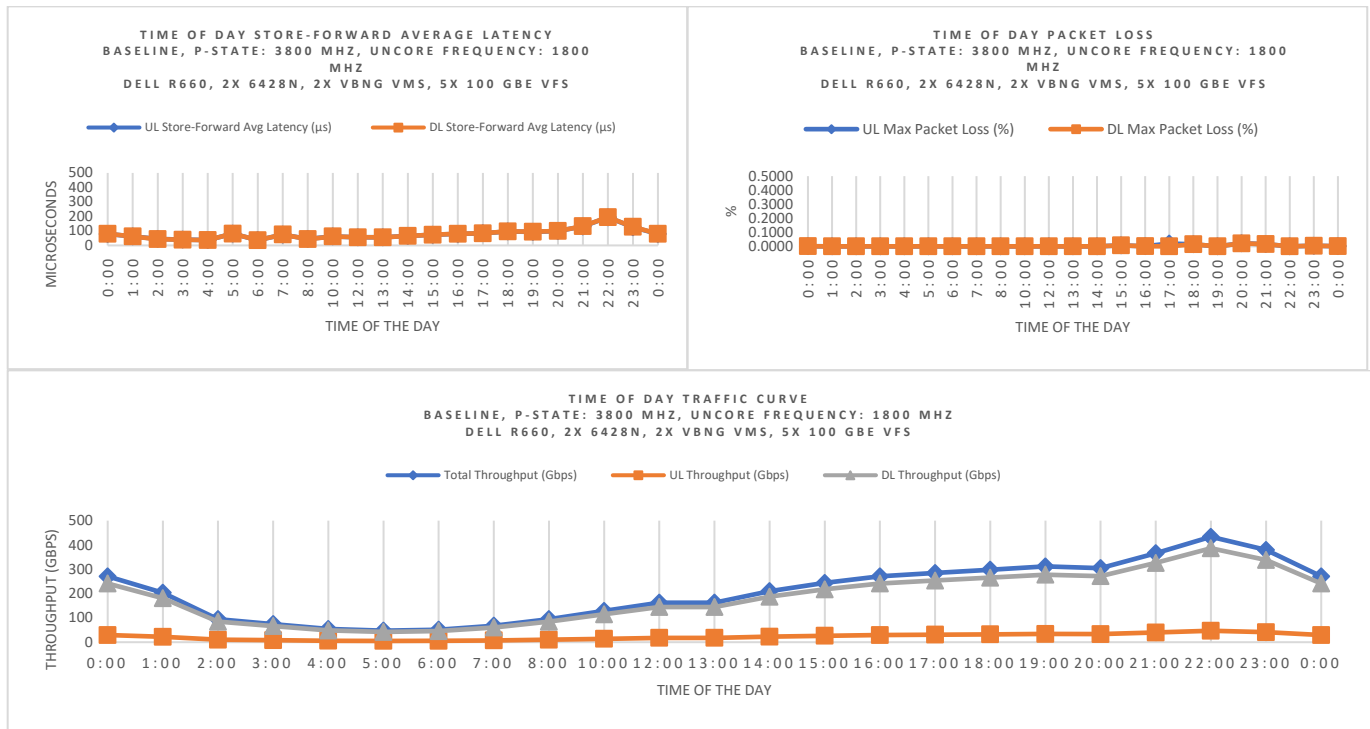
vBNG Test Results

This section presents the benchmark results for the Casa Systems Axyom™ vBNG deployed on Dell PowerEdge R660 server with the 4th Generation Intel® Xeon® Scalable processor. Figure 4 presents aggregate UL/ DL throughput for the baseline case, with a 40,000-subscriber count and a UL / DL packet profile of 512 B / 768 B. Note that the figure includes the DL throughput, the UL throughput, as well as the total throughput over the course of 24 hours. In addition, the figure below presents the store-forward average latency for the baseline

case. As before, the maximum allowed latency is 500 μs. Note that the latency reaches a local maximum during peak throughput as expected. Furthermore, the Figure presents the

packet loss rate over the course of 24 hours for the baseline case. As before, the maximum allowed packet loss rate is 0.5%.

Figure 4 Virtualized BNG Network Topology



The following figure presents the total DRAM power, the total CPU power, as well as the total wall power over the course of 24 hours. Overlaid on the figure is a plot of the corresponding CPU frequency as well as the average uncore frequency for each hour of the day.

Note that in the baseline case the total wall power ranges from approximately 570 W during off peak hours up to approximately 654 W during peak conditions. Note also that the CPU frequency increases during off peak hours, at which point there is additional TDP headroom, and decreases during peak conditions, at which point the processor is running at maximum TDP.

Figure 5 Baseline Time of Day Power Consumption along with Core and Uncore Frequencies

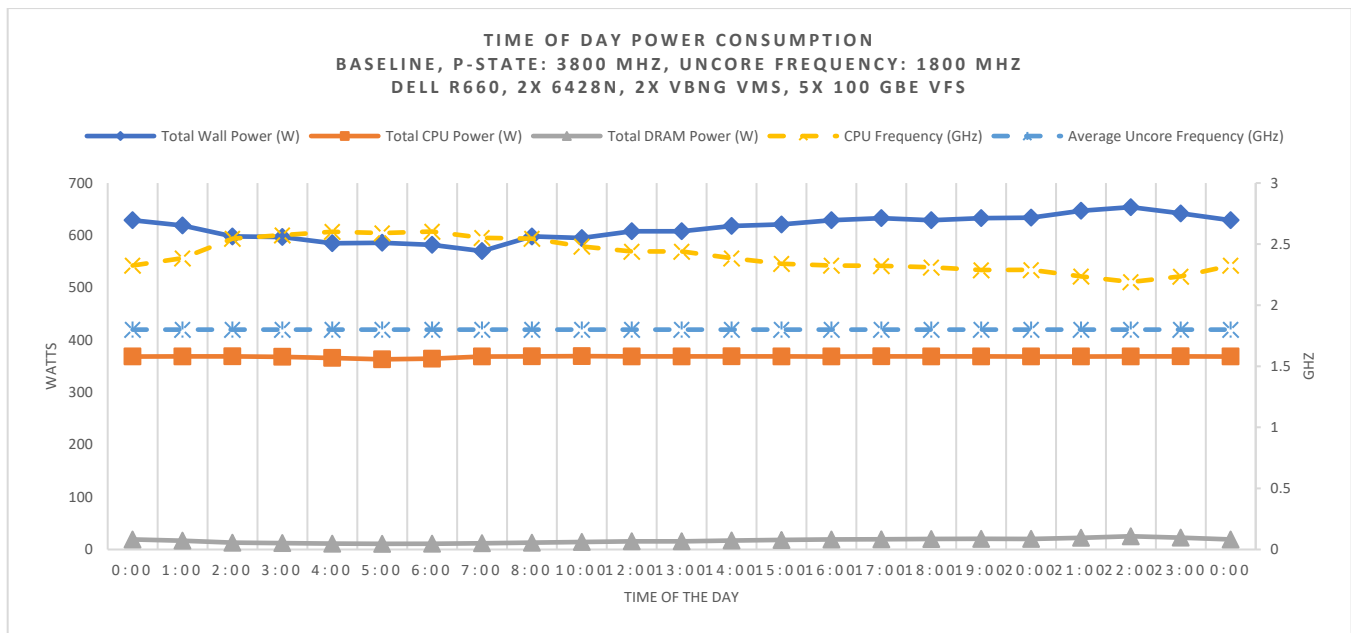


Figure 6 presents the aggregate throughput over the course of 24 hours for the case where P-state power savings alone are applied. Note in this case the throughput results presented in Figure 5 are virtually identical to the throughput results presented in Figure 1. Specifically, there is no observable loss in terms of throughput when the P-state power savings are applied. In addition, the figure includes the store-forward average latency with P-state power savings alone is applied.

Note that, like the baseline case, the latency remains below 500 μs. Furthermore, the figure presents the UL / DL packet loss rate for the case where P-state power savings are applied. As before the packet loss rate remains below 0.5 %.

Figure 6 Store and Forward Latency, Percentage Packet Loss, and Time of Day Traffic Curve with P-State Power Savings

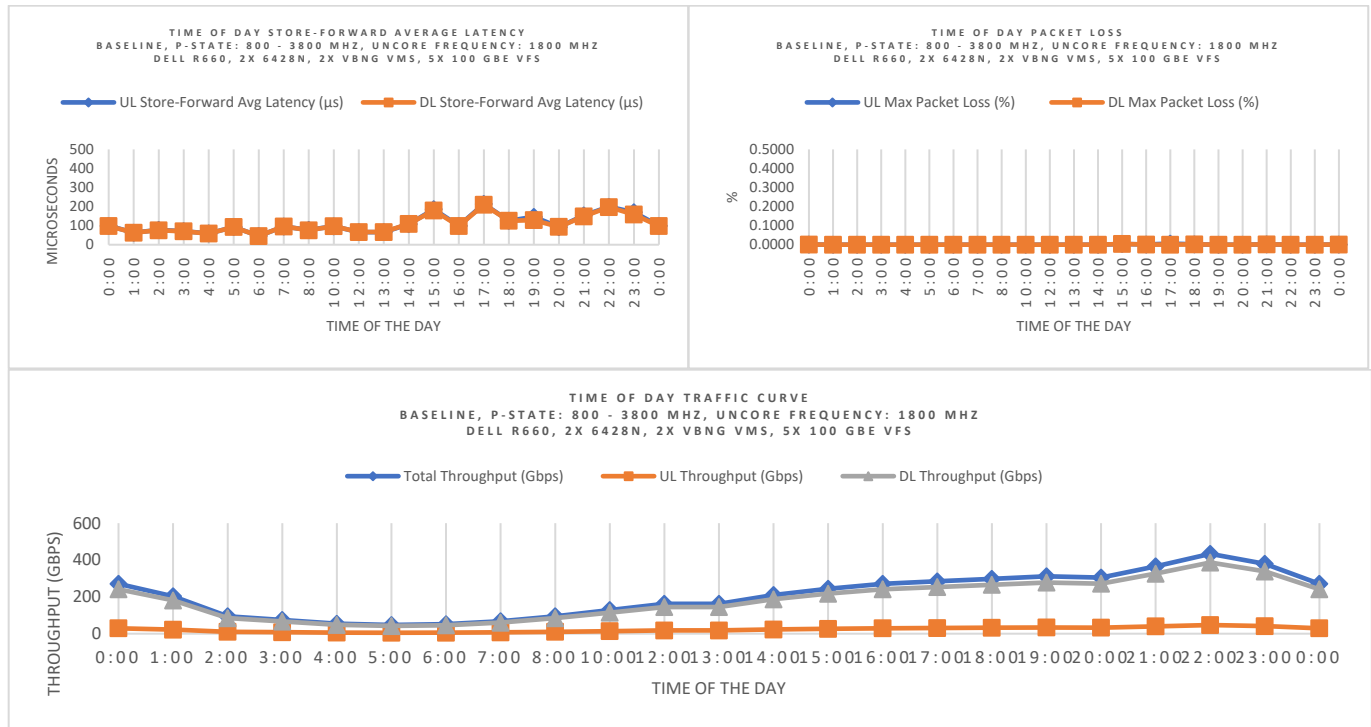


Figure 7 presents the total DRAM power, total CPU power, as well as the total wall power consumption over the course of 24 hours with P-state power savings applied. As before, the figure includes an overlay of the CPU frequency as well as the uncore frequency. In this case the total wall power ranges from approximately 407 W during off-peak conditions up to approximately 592 W during peak hours.

Figure 7 Time of Day Power Consumption with Core and Uncore Frequencies with P-State Power Savings

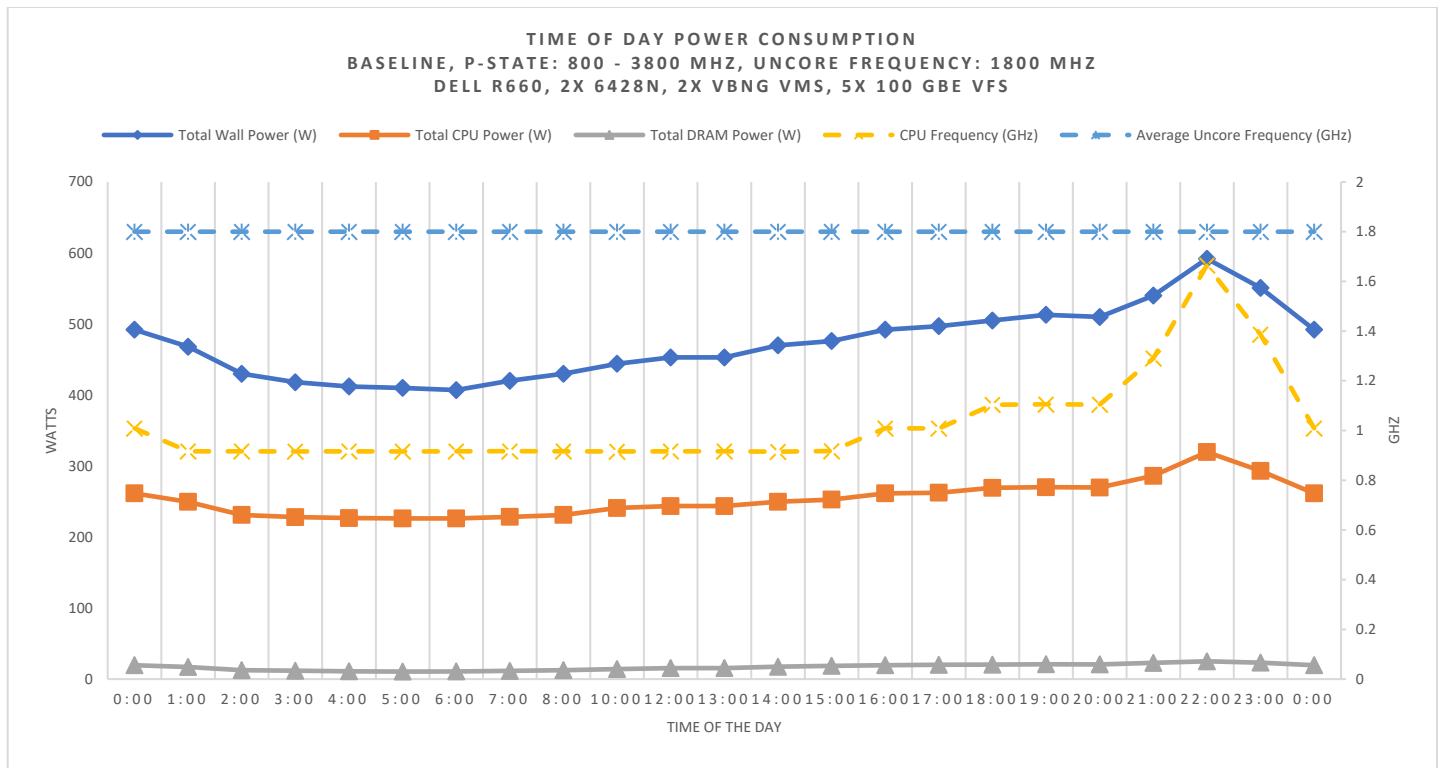


Figure 8 presents the aggregate throughput over the course of 24 hours for the case where P-state power savings and uncore frequency power savings are applied. As before, note that the throughput results presented in Figure 8 are virtually identical to the throughput results presented in Figure 4. Specifically, there is no observable loss in terms of throughput when P-state and uncore frequency power savings are applied. In addition, the figure presents the store-forward average latency with P-state power savings and uncore frequency power savings applied.

Note that, like the baseline case, the latency remains below 500 μs. Moreover, the figure presents the UL /DL packet loss rate for the case where P-state power savings and uncore frequency power savings are applied. As before the packet loss rate remains below 0.5 %.

Figure 8 Time of Day Traffic Curve with P-State and Uncore Power Savings




Figure 9 presents the total DRAM power, total CPU power, as well as the total wall power consumption with both P-state and uncore frequency power savings applied. As before, the plot includes an overlay of the CPU frequency as well as the uncore frequency over the course of 24 hours. In this case the total wall power ranges from approximately 368 W during off-peak conditions up to approximately 557 W during peak hours.

Figure 9 Time of Day Power Consumption with P-State and Uncore Power Savings


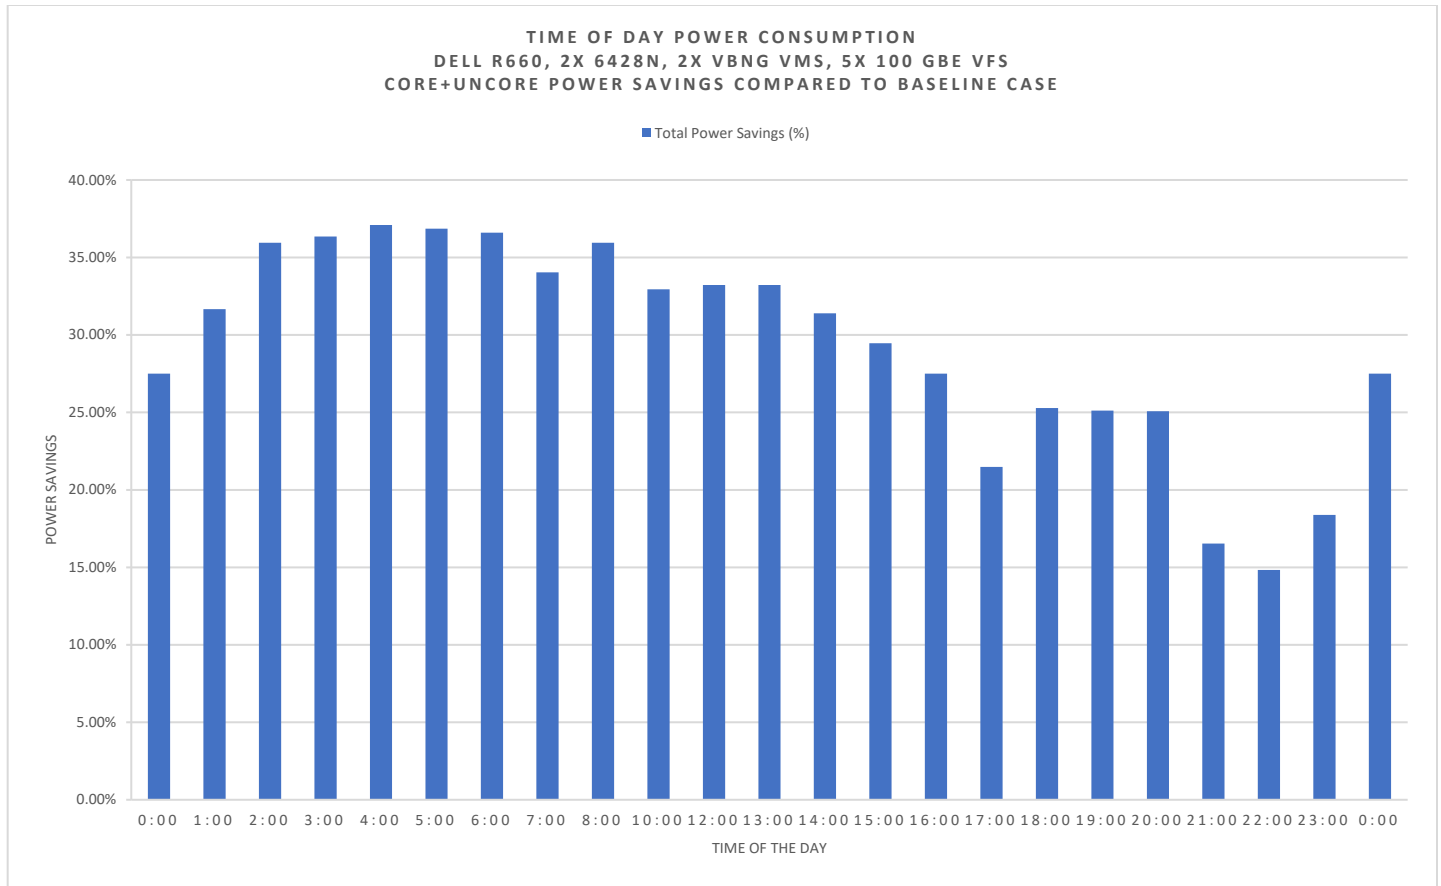


Figure 10 presents the total wall power savings for the case where P-state power savings and uncore frequency power savings are applied relative to the baseline case. In this case, the total wall power savings reaches a maximum of up to approximately 37.09% during off-peak conditions. Furthermore, with P-state power savings and uncore frequency power savings, the Casa Systems Axyom™ vBNG deployed on 4th Generation Intel® Xeon® Scalable Processors can achieve on average 29.3% total wall power savings over the course of 24 hours.

Figure 10 Percentage Time of Day Power Savings for P-State and Uncore Power Savings Compared to Baseline Case




Summary

This whitepaper demonstrates the performance of the Casa vBNG Axyom™ on 4th Generation Intel® Xeon® Scalable processors with Intel® 800 Series Ethernet Controllers and the potential power savings. In this case, the Casa vBNG Axyom™ can achieve an average of 29.3% total wall power savings over the course of 24 hours when benchmarked with a representative real world traffic curve. Specifically, with both P-state and uncore frequency power savings applied, the overall power consumption from a single Dell PowerEdge R660 server is dramatically reduced, which in turn implies significant savings in terms of CAPEX and OPEX.

Appendix A – Dell PowerEdge R660 Riser Configuration

The Dell PowerEdge R660 server might have different Risers Configuration, the configuration used in this study is the Riser Configuration 3 (R1P + R4P) consisting of two Full Height Half Length PCIe slots as shown in the picture below.



The following steps were performed:

1. Install Intel Chapman Beach E810-2CQDA2 in Slot 1 and Slot 2
2. Update the Dell PowerEdge R660 server to the latest BIOS, CPLD, and iDRAC firmware available.
3. Bifurcate both Slot 1 and Slot 2. Go to BIOS Settings > Integrated Devices > Slot Bifurcation > Select “x8” for both Slot 1 & 2

Appendix B – BIOS Settings

Table A1: BIOS settings set by NFVI FP Optimized Turbo Profile

BIOS knob	Setting	Location in BIOS Setup
AC Power Recovery	Last	System Security->AC Power Recovery
AMP Prefetch	Disabled	Processor Settings-> AMP Prefetch
AVX ICCP Pre-Grant Level	N/A	Processor Settings->AVX ICCP Pre-Grant Level
AVX ICCP Pre-Grant License	Disabled	Processor Settings->AVX ICCP Pre-Grant License
AVX PI	Normal	Processor Settings->AVX PI
CIE	Enabled	System Profile Settings->CIE
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
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
Homeless Prefetch	Enabled	Processor Settings->Homeless Prefetch
Intel SST-CP (RAPL Prioritization)	Disabled	Processor Settings->Intel SST-CP
LLC Prefetch	Enabled	Processor Settings->LLC Prefetch
Logical Processor (Hyper-Threading)	Enabled	Processor Settings->Logical Processor
Memory Patrol Scrub (Patrol Scrubbing)	Standard	System Profile Settings->Memory Patrol Scrub
Monitor/Mwait	Enabled	System Profile Settings->Monitor/Mwait
Memory Refresh Rate (Memory DIMM Refresh Rate)	1x	System Profile Settings->Memory Refresh Rate
OS ACPI Cx	OS Cx C2	System Profile Settings->OS ACPI Cx
PCI ASPM L1 Link Power Management	Enabled	System Profile Settings->PCI ASPM L1 Link Power Management
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
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 NFVI FP Optimized Turbo Profile

BIOS knob	Setting
APS Rocketing	Disabled
Boot Performance Mode	Max. Performance
CPU C1 Auto Demotion	Disabled
CPU C1 Auto unDemotion	Disabled

BIOS knob	Setting
Dynamic L1	Disabled
EIST PSD Function	HW_ALL
EPP enable	Disabled
GPSS Timer	0us
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
Package C-States	Disabled
PCIe AER Error Handling – PCIe Correctable Errors	Disabled
PCIe ECRC generation and checking	Disabled
Power Performance Tuning	BIOS Controls EPB
Scalability	Disabled
UMA Based Clustering Status	Quadrant
Virtual NUMA (MCC)	Disabled



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. 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

CONFIG1 (BASELINE): Test by Intel as of 05/21/2022. 1-node, 2x Intel® Xeon® Platinum 6428N Processor, 32 core HT On Turbo On, Total Memory 1024GB (16 slots/ 64GB/ 4800MT/s), BIOS 1.2.1 (ucode: 0x2b000190), 2x Intel® E810-2CQDA2, 1x Intel® E810-CQDA2, Red Hat Enterprise Linux 8.7, 4.18.0-425.19.2.el8_7.x86_64, gcc (GCC) 8.5.0 20210514 (Red Hat 8.5.0-16), Casa Systems Axyom™ vBNG CP Ver 3.3.1, Casa Systems Axyom™ vBNG DP: 3.3.1. score=54.18 Gbps, 585 W at 512 B / 768 B mix with 40K subs with P-state 800 - 3800 MHz, uncore 1800 MHz

CONFIG2 (VARY P-STATES & UNCORE FREQ.): Test by Intel as of 05/21/2022. 1-node, 2x Intel® Xeon® Platinum 6428N Processor, 32 core HT On Turbo On, Total Memory 1024GB (16 slots/ 64GB/ 4800MT/s), BIOS 1.2.1 (ucode: 0x2b000190), 2x Intel® E810-2CQDA2, 1x Intel® E810-CQDA2, Red Hat Enterprise Linux 8.7, 4.18.0-425.19.2.el8_7.x86_64, gcc (GCC) 8.5.0 20210514 (Red Hat 8.5.0-16), Casa Systems Axyom™ vBNG CP Ver 3.3.1, Casa Systems Axyom™ vBNG DP: 3.3.1. score=54.18 Gbps, 368 W at 512 B / 768 B mix with 40K subs with P-state 800 MHz, uncore 800 MHz

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.

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. 1116/TM/Wipro/PDF

Please Recycle 793452-001US