



Optimize Your NFVI Performance with Wiwynn® EP100 and Intel® Speed Select Technology – Base Frequency



1. Introduction

Communications service providers (CoSPs) are expanding their use of network function virtualization infrastructure (NFVI), which decouples software from proprietary hardware and enhances scalability with virtualized network function (VNF) software running on COTS (commercial off-the-shelf) servers.

Wiwynn® EP100¹ design is optimized for CoSPs' NFVI with flexibility, serviceability and performance. CoSPs can utilize their existing racks with flexible configuration of the 1U half-width single socket sled.

With the powerful and optimized 2nd generation Intel® Xeon® Gold 6252N processor, part of the Intel® Xeon® Scalable processor family, EP100 enables high speed packet forwarding, which is critical to NFVI. By adopting the new NFVI specialized feature called Intel® Speed Select Technology – Base Frequency (Intel® SST-BF),² EP100 boosts performance of targeted applications essential for some VNFs.³

With supports from Industrial Technology Research Institute (ITRI) and Intel, Wiwynn utilized OPNFV Yardstick/Network Services Benchmarking (NSB)⁴ framework to compare performance of Wiwynn® EP100 with Intel® SST-BF enabled and disabled in this whitepaper. The Layer 2 forwarding tests were run on Intel® Xeon® Gold 6252N processor⁵ for performance benchmark of Open vSwitch (OVS) and the Data Plane Development Kit (DPDK). This testing process not only helps to identify Wiwynn® EP100's performance under different CPU frequency configurations, but also assists CoSPs to develop a common reference set of performance benchmarks for their applications.

2. Intel® SST-BF and Performance Benchmark Tools

2.1 What Is Intel® SST-BF?

Intel® SST-BF is a new CPU feature which allows the CPU to be deployed with an asymmetric core frequency configuration. Placing key workloads on higher frequency Intel® SST-BF enabled cores enables users to boost performance of targeted applications at runtime.

Figure 1 shows both symmetric and asymmetric core frequency deployments. In a default symmetric core frequency deployment, all applications operate at the same core frequency while Intel® SST-BF enabled processor can reserve cores with higher frequency for critical applications.

Table of Contents

- 1. Introduction 1
- 2. Intel® SST-BF and Performance Benchmark Tools 1
 - 2.1 What Is Intel® SST-BF? 1
 - 2.2 Introduction of Used Performance Tools 2
- 3. Test Setup 2
 - 3.1 Configurations of the SUT – Wiwynn® EP100 3
 - 3.2 Test Traffic Setup 6
- 4. Test Results³ 6
- 5. Summary 7

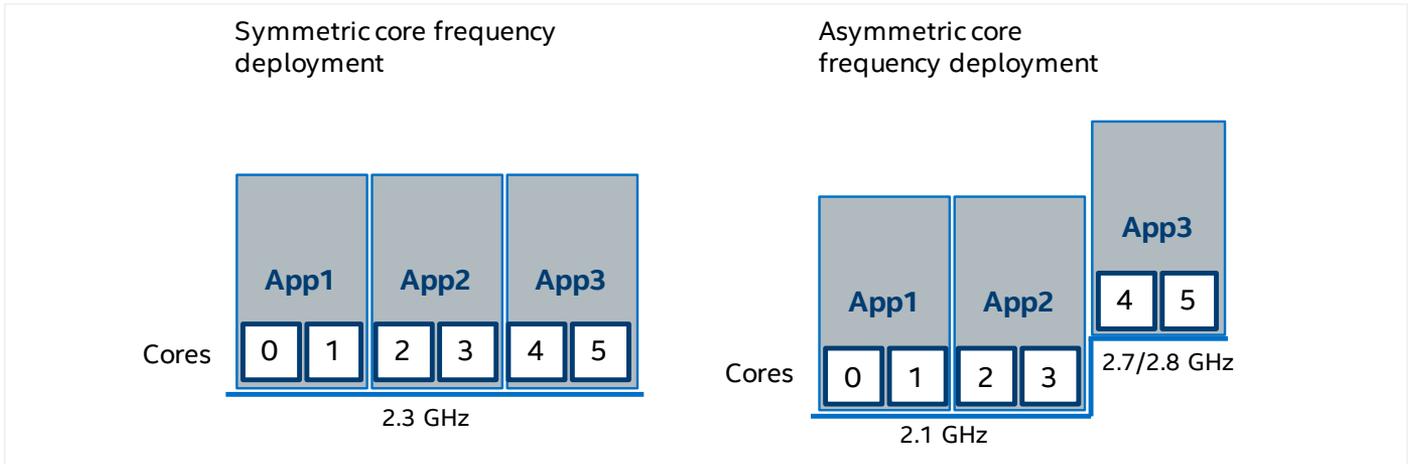


Figure 1. Core Frequency Deployment Methods

Intel® SST-BF is available on select SKUs of 2nd gen Intel® Xeon® Scalable processors, including 5218N, 6230N and 6252N, with an emphasis on network and virtualization workloads. Table 1 shows their corresponding high priority and standard priority base frequency configurations.

In this whitepaper, Wiwynn compared performances of EP100 with Intel® Xeon® Gold 6252N processor with and without Intel® SST-BF enabled.

PROCESSOR MODEL	HIGH PRIORITY		STANDARD PRIORITY	
	CORES	BASE FREQUENCY	CORES	BASE FREQUENCY
Intel® Xeon® Scalable processor 6252N	8	2.8 GHz	16	2.1 GHz
Intel® Xeon® Scalable processor 6230N	6	2.7 GHz	14	2.1 GHz
Intel® Xeon® Scalable processor 5218N	4	2.7 GHz	12	2.1 GHz

Table 1. Intel® SST-BF Enabled CPU SKUs

2.2 Introduction of Used Performance Tools

With ITRI and Intel’s technical support, Wiwynn used Yardstick/NSB (Network Services Benchmarking) to run a special test VNF called Packet pROcessing eXecution engine (PROX)⁶ to compare throughput performance of EP100 with different CPU SKUs under specific packet loss rate.

NSB is part of the OPNFV Yardstick project contributed by Intel, along with contributions from industry partners, to OPNFV community. It extends the Yardstick framework for VNF characterization and benchmarking of VNFs, NFVI and network services.

The Yardstick/NSB test framework provides CoSPs with common standards and industry-accepted benchmarks for conformance to carrier-grade requirements. NSB can be used to perform deterministic and repeatable benchmarks, and presents metrics in a unified GUI.

NSB can be applied in native Linux bare metal environments, standalone virtual environments (PCI pass-through, SR-IOV, OVS-DPDK), or in OpenStack or other managed virtualized environments.

PROX, the tested VNF, is part of the OPNFV’s samplevfn project and utilizes DPDK. It implements a suite of test

cases and gathers key performance indicator (KPI) data such as packets-in, packets-dropped, packets-forwarded for comparison.

3. Test Setup

Figure 2 shows the topology of the test setup. The traffic generator (TG), System Under Test (SUT), and NSB jump node were connected to the same Ethernet management switch. Additionally, the SUT was connected to TG using two 10 Gbps network ports that made up the data plane. The NSB jump node controls TG to generate traffic and collects a variety of key performance indicators during VNF execution on the SUT.

Wiwynn® EP100 was used as the SUT in this whitepaper. It was deployed with the DPDK-accelerated Open vSwitch (OVS-DPDK) and one virtual machine with two network ports which performed PROX Layer 2 forwarding VNF (l2fwd). The hardware and software details of EP100 are described in section 3.1.

For traffic generator (TG), Wiwynn installed NSB PROX on bare metal machine and built with DPDK. The benchmark tests collected performance metrics while SUT ran with Intel® SST-BF enabled and disabled. Section 3.2 shows the test traffic setup used in the benchmark tests in detail.

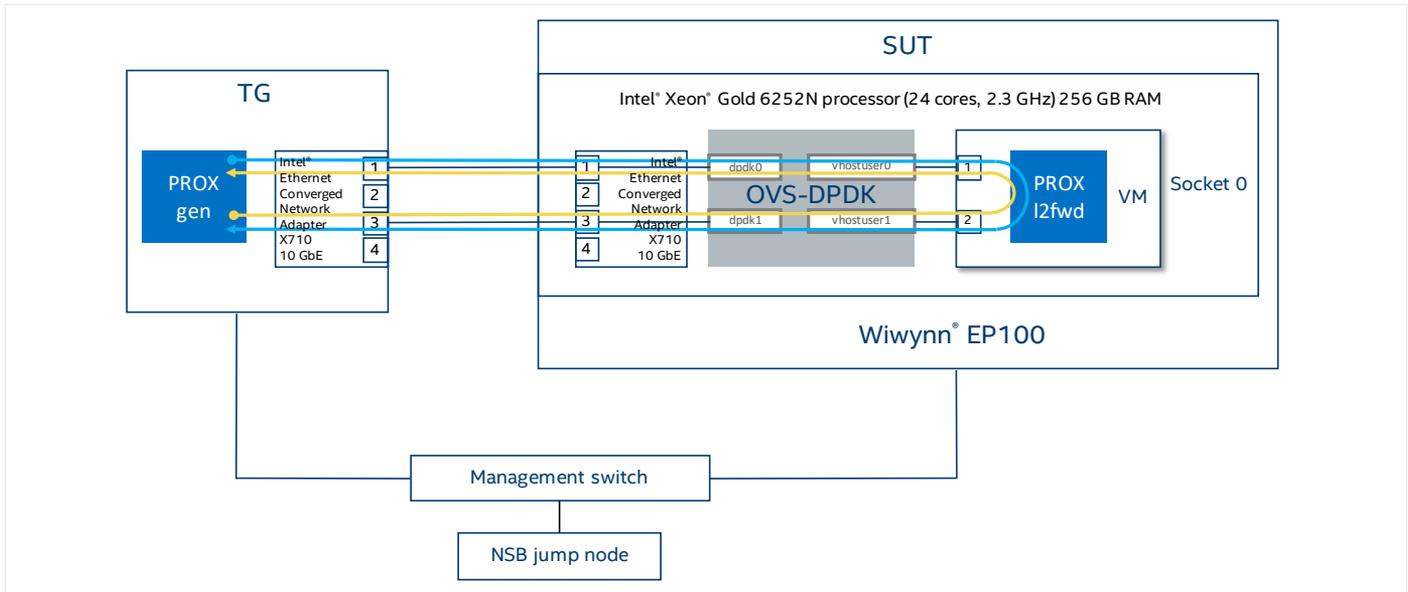


Figure 2. Test Configuration

3.1 Configurations of the SUT – Wiwynn® EP100

3.1.1 Hardware Configuration

Wiwynn® EP100 is a 3U short-depth edge system based on Nokia-led OCP OpenEDGE specification. It supports up to five 1U half-width single socket server sleds. Communication service providers can scale computing power by adding more EP100 systems for applications

ranging from base stations to regional central offices. With its flexible and efficient design, EP100 enables diverse applications requiring low latency and huge data-processing capabilities for 5G and edge computing.

In this test, one sled of EP100 was used as the SUT to test L2 forwarding performance under different frequency configurations. The 1U half-width sled’s configurations are as shown below.

	SLED HARDWARE SPECIFICATION	TEST CONFIGURATION
PROCESSOR	1S 2nd Gen Intel® Xeon® Scalable Processors	Test case 1: Intel® Xeon® Gold 6252N CPU (CPU microcode: 0x5000024), Intel® SST-BF disabled Test case 2: Intel® Xeon® Gold 6252N CPU (CPU microcode: 0x5000024), Intel® SST-BF enabled
DIMM	DDR4; up to 2933 MT/s; 8 x DIMM slots, supporting 2 x Intel® Optane™ Technology	256 GB
STORAGE	2 x U.2 hot plug SSDs	1.9 TB
	2 x onboard NVMe M.2 modules	Not used
EXPANSION SLOT	1 x PCIe 3.0 FHHL x16 AIC	Intel® Ethernet Converged Network Adapter X710-DA4 (4-port 10 Gbps SFP+)
	1 x OCP 3.0 Mezz (x16)	Not used
SYSTEM DIMENSIONS	3U; 448 x 430 x 130.6 (mm) (W x D x H), 19" rack compatible	

Table 2. Hardware Configuration of Wiwynn® EP100

3.1.2 Software Configuration

Figure 3 shows the software stack of SUT system and the software versions were as follows:

- Linux kernel: 4.20.17 (4.20 RC1 or later version is required for Intel® SST-BF function)
- Host OS: Ubuntu Server 18.04
- Virtual machine management: libvirt
- Hypervisor: KVM/QEMU
- Virtual network and connectivity between the VNFs and the physical 10Gb Ethernet ports: Open vSwitch (OVS) 2.11.0
- Packet throughput enhancement: DPDK 18.11.0 was utilized

Poll Mode Driver (PMD) cores for OVS were isolated for optimal performance. OVS was configured to use one/two/four cores, 1024 MB memory for DPDK, and single receiving and transmitting queue for each OVS port.

Four OVS rules were configured to allow all traffics between OVS port dpdk0 and vhostuser0, dpdk1 and vhostuser1.

Wiwynn defined VM with 3 CPU cores, large memory page size, dedicated CPU policy, and isolated thread policy. The detail of VM configuration was:

- 3 virtual CPUs, pinned to 3 physical cores.
- 8 GB of memory.
- CPU cores were isolated, and 4 x 1 GB huge pages were allocated.

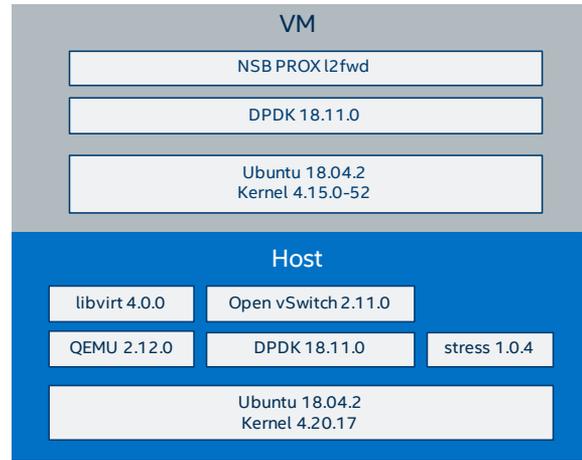


Figure 3. Software Components of SUT system

Figure 4 shows the core configurations for OVS-DPDK and VM. While Intel® SST-BF was enabled, the high priority cores were configured to OVS-DPDK and NSB PROX Layer 2 forwarding VNF in VM. These workloads can benefit from frequency scaling and have deterministic compute cycles. In fact, users can only configure the high priority cores to OVS-DPDK PMD threads to further save core resources since they are the major performance bottlenecks. The saved high priority cores can be assigned to other important workloads, such as to NSB PROX Layer 2 forwarding VNF in this test. The other cores configured as standard priority cores for workloads are not real-time or performance critical. Wiwynn used the stress 1.0.4 to generate background workloads on the standard priority cores to simulate the real environment.

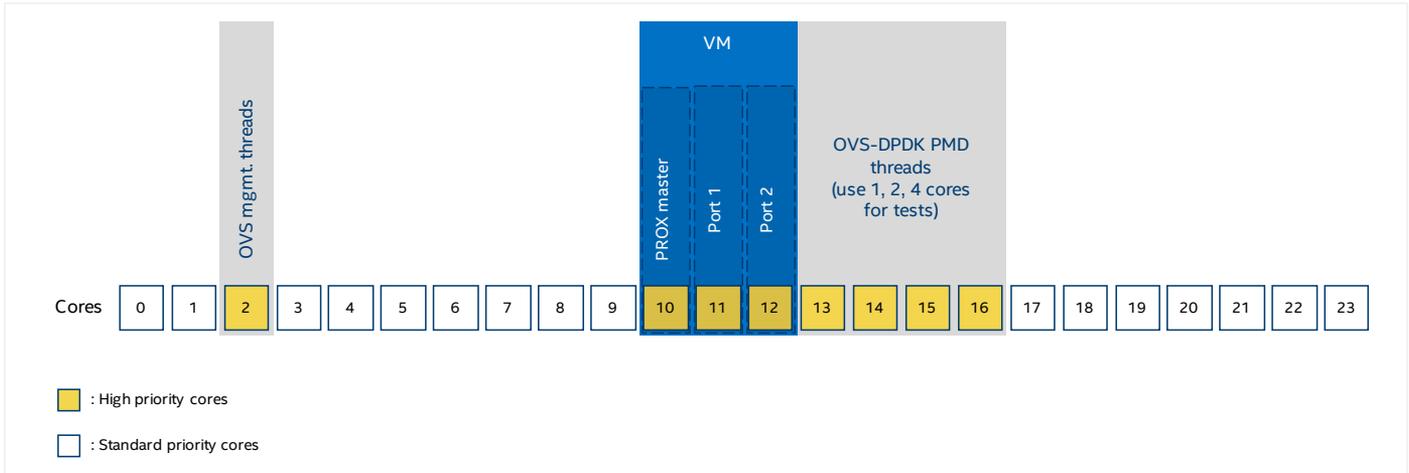


Figure 4. Core Configuration for OVS-DPDK and VM

3.1.3 Setting up Intel® SST-BF Functionality

The Intel® SST-BF functionality setup is as below:

- Linux Kernel version 4.20.17
- It requires a kernel patch that is available in the upstream Linux kernel version 4.20 and later. In

addition, to boot with the `intel_pstate` driver active is a must. Enable the Intel® SST-BF feature in the BIOS (see section 3.1.4 for the BIOS settings).

- Enable the feature with a sample node configuration script, `sst_bf.py` (download from <https://github.com/intel/CommsPowerManagement>)

3.1.4 BIOS Settings

Table 3 shows the required BIOS settings:

PATH	BIOS SETTING	VALUE
Socket Configuration → Processor Configuration	Hyper-Threading	Disable
	Uncore Freq Scaling (UFS)	Disable
	SpeedStep (Pstates)	Enable
	Activate PBF	Enable/Disable*
	Configure PBF	Disable
	EIST PSD function	HW_ALL
	Boot performance mode	Max performance
Socket Configuration → Advanced Power Management Configuration → CPU P state control	Energy Efficient Turbo	Disable
	Turbo mode	Enable
Socket Configuration → Advanced Power Management Configuration → CPU P state control → Perf P-Limit	Perf P Limit	Disable
	Hardware P-States	Native mode with no legacy support
Socket Configuration → Advanced Power Management Configuration → Hardware PM State Control	EPP Enable	Enable
	RAPL Prioritization	Enable
Socket Configuration → Advanced Power Management Configuration → CPU C State Control	Autonomous Core C-State	Disable
	CPU C6 report	Disable
	Enhanced Halt State (C1E)	Disable
Socket Configuration → Advanced Power Management Configuration → CPU - Advanced PM Tuning → Energy Perf BIAS	Power Performance Tuning	BIOS Controls EPB
	Workload Configuration	I/O Sensitive
	ENERGY_PERF_BIAS_CFG mode	Performance
Socket Configuration → Advanced Power Management Configuration → Package C State Control	Package C State	C0/C1 state
Socket Configuration → IIO configuration	PCI-E ASPM support (Global)	Disable
Socket Configuration → IIO configuration → IOAT configuration	Relaxed Ordering	Disable
Socket Configuration → IIO configuration → Intel® VT for Directed I/O (Intel® VT-d)	Intel® VT for Directed I/O (Intel® VT-d)	Enable
Socket Configuration → IIO configuration → IIO DFX configuration	EV DFX Features	Disable
Socket Configuration → UPI configuration → UPI General Configuration	Link L0p Enable	Disable
	Link L1 Enable	Disable
Socket Configuration → Memory Configuration	Enforce POR	Disable
Socket Configuration → Memory Configuration → Memory Map	IMC Interleaving	2-way Interleave

Table 3. BIOS Settings

* The toggle of Intel® SST-BF function.

3.2 Test Traffic Setup

The test traffic generator created one flow (one IP source, one IP destination, one UDP source and one UDP destination) per 10 GbE port. The SUT is configured to have one receiving queue per port. One flow per port usually gives better performance per-core, so Wiwynn used the one flow test to measure the upper bound of throughput (packet per second).

The results reported maximum throughput with a tolerated packet loss of 0.001%. The measurement was done by binary search.

The traffic generator generated the flow at maximum throughput at the beginning and adjusted based on the loss packets rate status. When the loss packets rate was higher

than the tolerated loss rate, it decreased the throughput. If the loss packets rate was lower than the tolerated loss rate, it increased the throughput.

4. Test Results³

Figures 5, 6 and 7 show the PROX I2fwd throughput results measured at the traffic generator under different packet sizes and numbers of high priority cores for OVS-DPDK PMD threads.

The throughputs are the sum of the receiving packets monitored at the two 10 GbE network ports of the traffic generator used in the test. The theoretical maximum throughput of two 10 GbE network ports is also in the figures for reference.

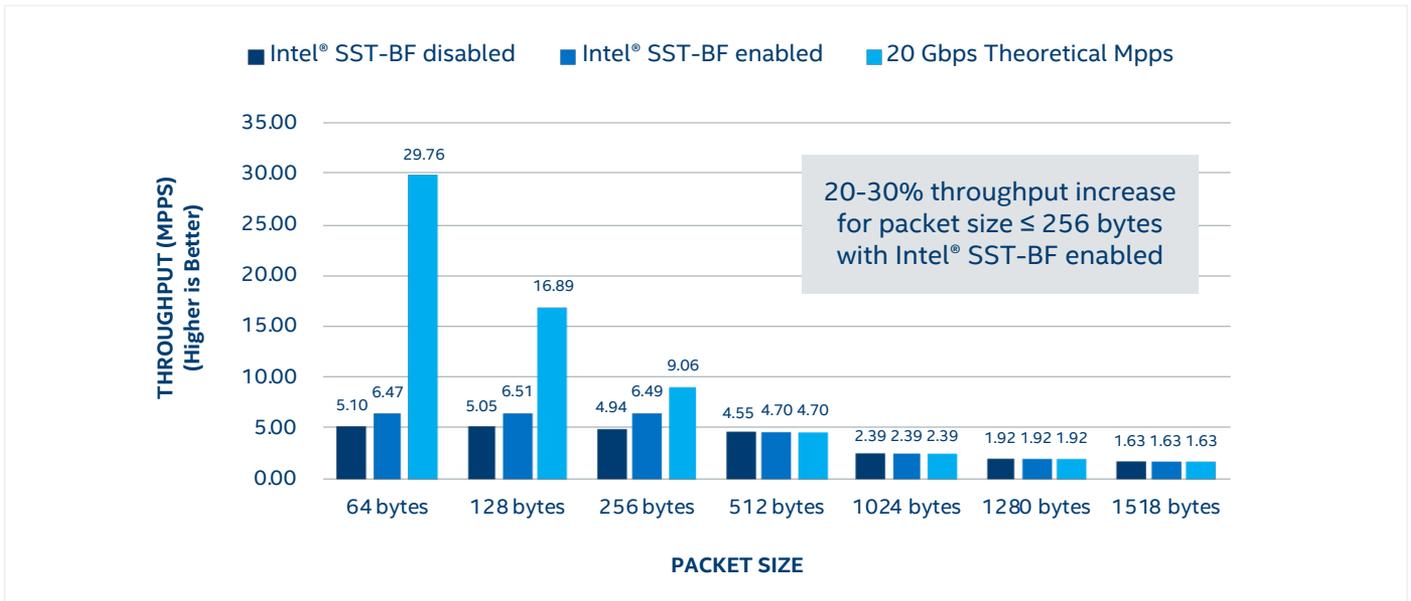


Figure 5. PROX I2fwd throughput results while 4 high priority cores were pinned to OVS

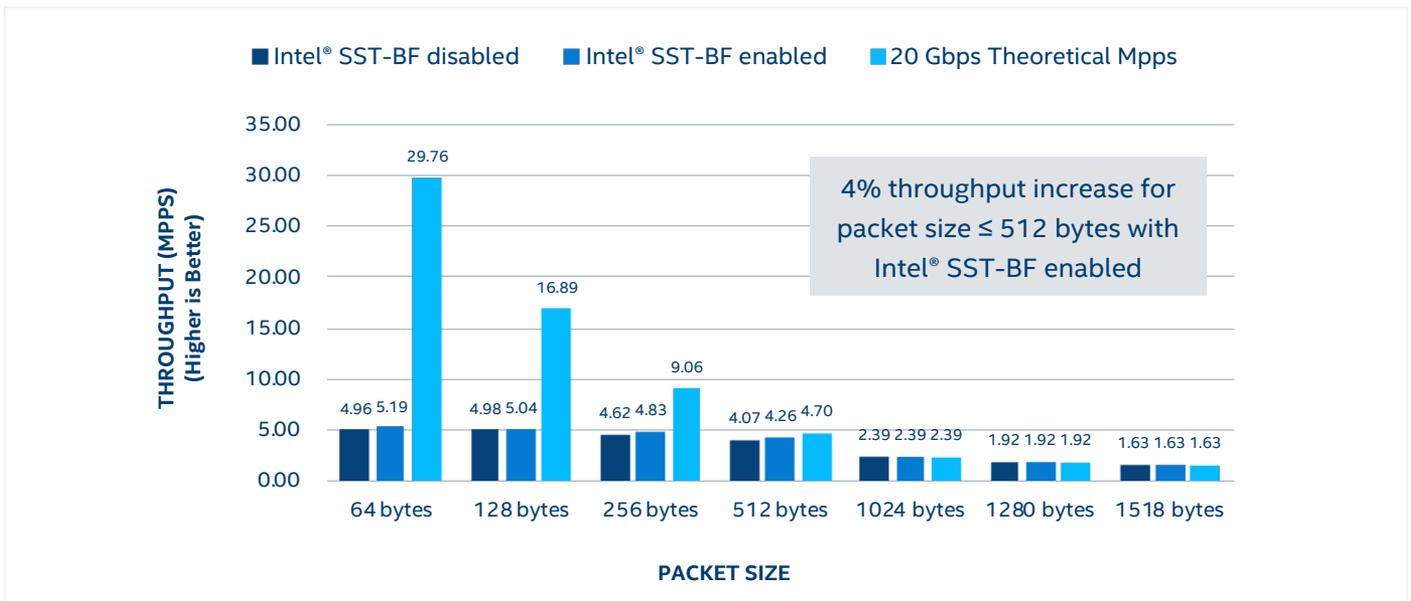


Figure 6. PROX I2fwd throughput results while 2 high priority cores were pinned to OVS

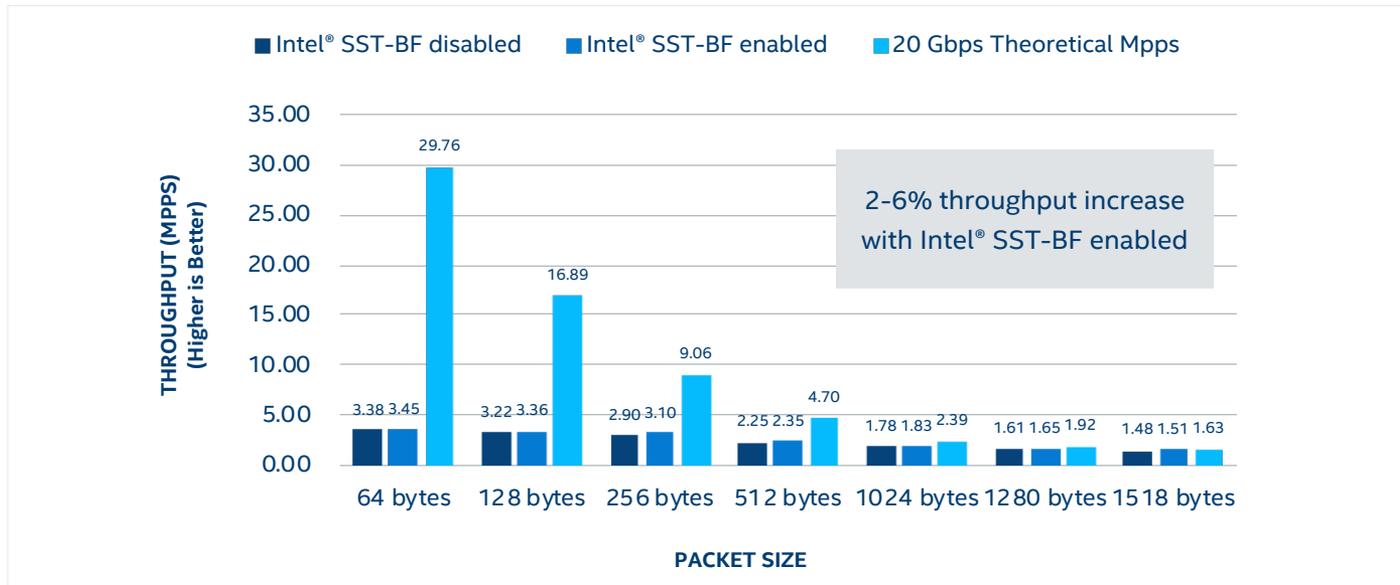


Figure 7. PROX I2fwd throughput results while 1 high priority core was pinned to OVS

The test results showed that there were different throughput improvements with Intel® SST-BF enabled on Wiwynn® EP100. The performance improvements of assigning one or two high priority cores for OVS-DPDK PMD threads were not significant compared to the case of assigning four cores. In addition, packet size was also a key factor in performance.

In the case of allocating four high priority cores for OVS-DPDK PMD threads, the throughput increased by 20-30% when the packet size was equal or smaller than 256 bytes. However, the performance improvement is not significant when the packet size is equal or larger than 512bytes. This is because the benefit of Intel® SST-BF is to assign high priority cores to heavy loading and mission critical workloads, which are OVS-DPDK PMD threads and PROX I2fwd in this test case, to guarantee their resource and performance. From the test results, the loading of OVS-DPDK PMD threads and PROX I2fwd are only heavy when the packet size was equal or smaller than 256 bytes. And starting from 512 bytes, the test reached line rate.

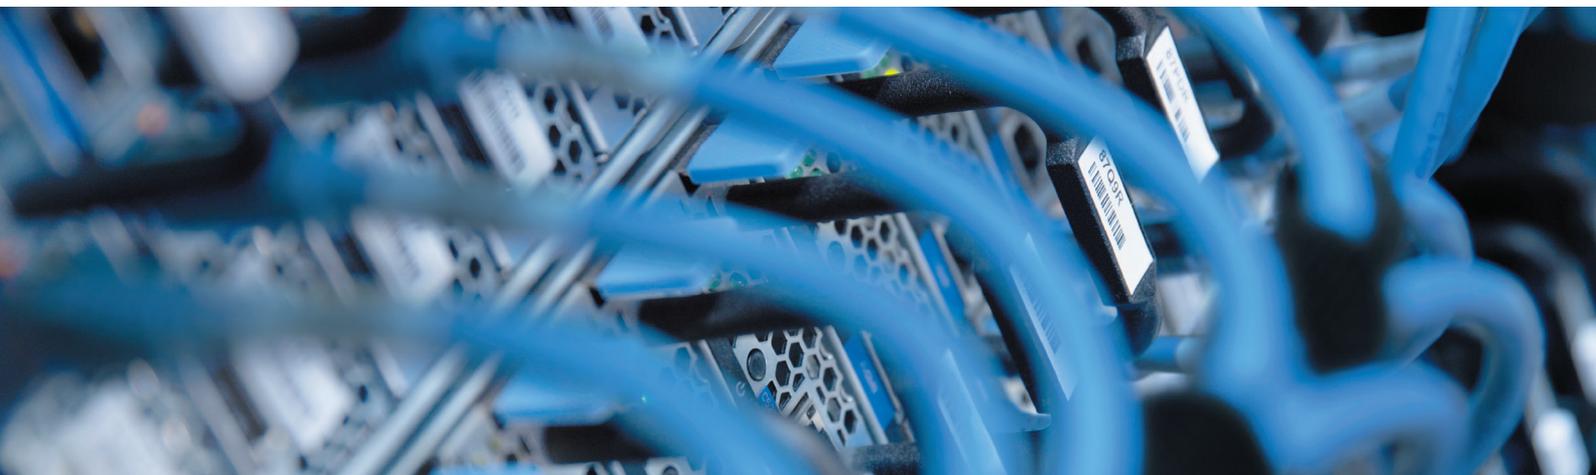
The test results also brought another issue of what's the best CPU cores configuration for OVS-DPDK PMD threads. More cores assigned to OVS-DPDK PMD threads can bring better network performance but sacrifices the resource of other workloads. From the test, the packet size of VNFs is a good indicator for OVS cores setup to achieve the performance

balance. And Yardstick/NSB test framework can assist CoSPs, system integrators and equipment providers to benchmark performance and evaluate the core configuration to speed up their NFVI deployments.

5. Summary

Wiwynn® EP100 has adopted 2nd Gen Intel® Xeon® Gold 6252N processors, part of the Intel® Xeon® Scalable processor family, and enabled Intel® SST-BF feature. In its test, Yardstick/NSB test framework was used to compare Layer 2 forwarding performance on Wiwynn® EP100 with Intel® SST-BF enabled and disabled. The test results have proved that up to 20-30% throughput increase were observed when 4 high priority cores were pinned to OVS (Intel® SST-BF enabled) and the packet size was equal or less than 256 bytes.³

By assigning high priority cores to heavy-loading and latency-critical workloads, such as OVS-DPDK PMD threads in this whitepaper, users can boost performance of targeted applications. With Yardstick/NSB test framework, users, including CoSPs, system integrators and equipment providers, can further benchmark VNF performance and adjust CPU core assignment strategy based on different workload characteristics and combinations for NFVI workload optimization.



About Wiwynn

Wiwynn® is an innovative cloud IT infrastructure provider of high quality computing and storage products, plus rack solutions for leading data centers. We aggressively invest in next generation technologies for workload optimization and best TCO (Total Cost of Ownership). As an OCP (Open Compute Project) solution provider and platinum member, Wiwynn actively participates in advanced computing and storage system designs while constantly implementing the benefits of OCP into traditional data centers.

For more information, please visit Wiwynn [website](#), [blog](#) or contact sales@wiwynn.com

Follow Wiwynn on [Facebook](#) and [LinkedIn](#) for the latest news and market trends.

About Intel® Network Builders

Intel Network Builders is an ecosystem of infrastructure, software, and technology vendors coming together with communications service providers and end users to accelerate the adoption of solutions based on network functions virtualization (NFV) and software defined networking (SDN) in telecommunications and data center networks. The program offers technical support, matchmaking, and co-marketing opportunities to help facilitate joint collaboration through to the trial and deployment of NFV and SDN solutions. Learn more at <http://networkbuilders.intel.com>.



¹ Wiwynn® EP100 video: <https://www.youtube.com/watch?v=BBDlnwziPqc>

² Intel® Speed Select Technology – Base Frequency official website at <https://www.intel.com/content/www/us/en/architecture-and-technology/speed-select-technology-article.html>

³ Testing conducted by Wiwynn as of Sept. 12, 2019. See the Test Setup section for configurations.

⁴ OPNFV Yardstick/Network Services Benchmarking (NSB): https://docs.opnfv.org/en/stable-danube/submodules/yardstick/docs/testing/user/userguide/14-nsb_installation.html

⁵ Intel® Xeon® Scalable Processor 6252N official data at <https://ark.intel.com/content/www/us/en/ark/products/193951/intel-xeon-gold-6252n-processor-35-75m-cache-2-30-ghz.html>

⁶ Packet pROcessing eXecution engine (PROX): <https://wiki.opnfv.org/pages/viewpage.action?pageId=12387840>

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more complete information visit www.intel.com/benchmarks.

Testing conducted by Wiwynn. Performance results are based on testing as of Sept. 12, 2019, and may not reflect all publicly available security updates. See configuration disclosure for details. No product or component can be absolutely secure.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at intel.com.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Notice Revision #20110804

Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

Other names and brands may be claimed as the property of others.