WHITE PAPER

Perspective on VNF Onboarding on Intel® Architecture in an NFVI Environment

Intel's experience with interoperability testing identified a need for common configuration methods to enable seamless onboarding of VNFs on multiple NFVI platforms

Nikita Agarwal
Senior Applications Engineer

Tarek Radi
Lead Technical Program Manager

Daniel B. Shea
Senior Applications Engineer

Executive Summary

The Intel Interoperability program has been launched to make strides in simplifying and accelerating adoption of Network Functions Virtualization Infrastructures (NFVIs) and Virtualized Network Functions (VNFs). Seamless VNF onboarding processes, well established measurement methods and unified license management are some of the key factors that will enable Communication Service Providers (CSPs) to keep up with the pace of today's sky-rocketing business demands. With these improvements, CSPs can reduce VNF bring-up time, and have a consistent way of evaluating platform performance, which allows for a simpler way of comparing and selecting VNF and NFVI alternatives that are available in the market. Ultimately, CSPs will benefit by limiting their expenditure and total cost of ownership while still increasing the service capacity of their infrastructure.

In this whitepaper, we report several of the steps that Intel took in evaluating two NFVIs and several VNF products. Intel's findings show that the industry lacks a commonly agreed and unified approach to provisioning virtualized services. These conclusions seem to be in line with the findings in the "SDN Service Providers Survey" that was recently conducted by Infonetics Research[1].

Similar to the Infonetics Research survey, Intel's efforts revealed several pain points for CSPs. NFVIs are still maturing and changing rapidly, making it challenging for VNF providers to keep up the pace. VNFs are difficult to onboard and test due to the lack of common configuration methods and tools. As a result, it usually takes a long time to onboard VNFs. CSPs meet these VNF and NFVI challenges by setting up mini data centers as "playgrounds", dedicating them for an often long "proof-of-concept" (PoC) phase. Despite such an investment, PoC findings are often insufficient for fully validating real world scenarios, which require more complex topologies and much higher expenditures.

This whitepaper focuses on some initial Intel Interoperability activities and provides an overview on the evaluated NFVIs and VNFs, including test setups and results. The whitepaper calls out several findings and challenges identified during interoperability activities. Intel plans to address several of these issues in upcoming interoperability activities in 2016, such as proposing common methods of VNF configuration on multiple NFVI platforms and proposing reliable measurement methods that would enable easier and more consistent comparison of different NFVI setups.

[1] Infonetics Research SDN Service Provider Survey 7/15

Table of Contents

Introduction

Network Function Virtualization (NFV) should not become a bottleneck as Telco, Enterprise and Cloud customers grow and scale their virtualized networks. NVFIs tend to have their own deployment specifics and a different set of minimum configuration requirements. This causes the deployment and configuration of VNFs to be even more challenging.

Today, many service providers are in their early stages of switching from physical appliances to VNFs. Recent reports show that the NFV industry is expected to grow rapidly in the next few years, and the process to fully convert physical appliances to virtualized solutions may take 10 to 15 years[2]. Forecasts show that the overall NFV market value will increase by about 50% each year, and will reach approximately $11 billion by 2019[2,3]. With such a rapid growth, it is essential that standards emerge in this field to simplify deployment, configuration and management of NFVIs and VNFs.

A sharp turn in favor of virtualized solutions is already visible in the CSP market. Surveys performed in 2015 indicate that more than half of the service providers are already testing NFV solutions in their mini data centers[3]. Many others are considering doing the same. This is an understandable trend since several NFV solutions have shown that they greatly reduce capital and operational expenditures, and enable a more cost effective model of network scaling. Prior to adopting NFV solutions, one of the greatest pains of CSPs was their dependency on very expensive, inflexible, and proprietary hardware. NFV allows them to use industry-standard hardware and promises easily scalable infrastructure to meet their growing data traffic demands.[4]

Rapid growth of the NFV market cannot be achieved by sacrificing quality of CSP solutions. The CSP industry is highly dependent on the continued availability of services, thus reliability is one of the key criteria considered when CSPs select NFV business partners. Other key criteria include functionality, performance, licensing costs, and manageability. It is expected that CSPs will need strong support from NFV providers during their physical-to-virtual transition period. Efficient processes should be developed to help coordinate support activities during that transition.

Intel Interoperability Program

Through its interoperability activities, Intel aims to contribute significantly to the acceleration of the SDN/NFV ecosystem by quickly identifying many of the most crucial obstacles and then prioritizing respective solutions. The Interoperability program presented Intel with a unique opportunity to collaborate with several VNF and NFVI providers and perform joint testing to onboard different VNF software on two different NFVI platforms. Two sets of interoperability tests were performed: 1) VNF onboarding tests and 2) VNF static function chaining tests. As part of the second set of tests, traffic was passed through the VNF service function chain and then performance measurements were captured.

[2] http://www.infonetics.com/pr/2015/NFV-Market-Highlights.asp
[3] http://www.fiercetelecom.com/story/sdn-and-nfv-investments-will-account-20b-revenue-2020-study-says/2015-12-11
[4] http://www.infonetics.com/pr/2015/Routing-NFV-Packet-Optical-Survey.asp

Goals of the Intel Interoperability Program

Intel and participating NFVI and VNF providers aimed to achieve the following 6 goals:

1. Gain hands-on experience with VNF onboarding on different but relatively similar NFVI platforms.
2. Work with different NFVI platforms and VNFs to describe the process of onboarding.
3. Identify onboarding obstacles, describe them and provide conclusions with possible solutions and/or recommendations.
4. Identify basic tests to pass traffic through the VNFs and gather performance data on the NFVI setups.
5. Present the findings and conclusion to the ecosystem.
6. Identify next steps with key ecosystem partners to resolve some of the problems that hinder SDN/NFV adoption.

VNF/NFVI Details

Interoperability activities included several types of VNFs, particularly virtual load balancers, virtual routers and virtual firewalls. These VNFs were deployed on two NFVI platforms. One was an Intel® Open Network Platform Release 1.5 (Intel® ONP 1.5)[5] and the other was a Wind River® Titanium Server[6] 15.09-1. The underlying hardware for the Intel® ONP 1.5 platforms included 2 SuperMicro SuperServers 6018U TR4+[7]. As for the Wind River® Titanium Server, 2 Intel® Server Board S2600WTT[8] servers were used. Even though the hardware was slightly different, it was equivalent in terms of the key components such as CPU, memory and NICs. See Table 1 and Table 2 for more details on the hardware and software used for each NFVI setup.

After the NFVI was setup, a pool of Web Servers was deployed and configured. Then the VNFs were deployed and configured in a way such that traffic would chain from one VNF to another. Finally, test traffic was generated which consisted of consecutive HTTP requests going through a static service function chain of these VNFs before arriving at the pool of Web Servers.

HARDWARE PLATFORM	INTEL® OPEN NETWORK PLATFORM RELEASE 1.5	WIND RIVER® TITANIUM SERVER 15.09-1
Number of Servers	2 servers. (1 configured as a Controller node, and one configured as a Compute node)	2 servers, each configured as a Controller-Compute node
Server specification	· SuperMicro SuperServer 6018U-TR4+ · CPU: 2x Intel® Xeon® E5-2699 v3; 2.30GHz · HDD: 1.0 TB · RAM: 256 GB · NICs: onboard integrated 4x 1GbE ports + 1x Intel® X520-DA2	· Intel® Server Board S2600WTT · CPU: 2x Intel® Xeon® E5-2699 v3, 2.30GHz · HDD: 3 TB + 1 TB · RAM: 256 GB · NICs: 1x Intel® X520-DA2 + 1x Intel® X540-T2
Switches	· 4x Netgear GS108 (8x 1 GbE) switch · HP 5900 Series Optical Switch JC772A 48x XG(10G) + 4x 40G(QSFP)	· 2x Netgear GS108 (8x 1 GbE) switch · 1x Netgear ProSAFE JGS524NA (24x 1Gbe) switch · HP 5900 Series Optical Switch JC772A 48x XG(10G) + 4x 40G(QSFP)

Table 1. Hardware platforms used by Intel Interoperability Program

[5] http://www.intel.com/content/www/us/en/communications/network-infrastructure-open-source-open-standards.html
[6] http://www.windriver.com/products/titanium-server
[7] http://www.supermicro.com/products/system/1U/6018/SYS-6018U-TR4_.cfm
[8] http://ark.intel.com/products/82156/Intel-Server-Board-S2600WTT

SOFTWARE INGREDIENT	INTEL® OPEN NETWORK PLATFORM 1.5	WIND RIVER® TITANIUM SERVER 15.09-1
OpenStack*	Kilo 2015.1.1	Kilo 2015.1.0
DPDK*	2.0.0	2.0.0
Open vSwitch*	Open vSwitch 2.4.9	Accelerated vSwitch* (AVS)
KVM*	QEMU-KVM 2.3.0.5.fc21	QEMU-KVM 2.2.0,C
Operating System	Fedora* 21	Wind River® Linux 6.0.0.22

Table 2. Versions of software ingredients

Interoperability Tests

NFVI platforms Setup Experience

The installation of Intel® ONP 1.5 Platform followed the steps described in the Intel® ONP 1.5 Reference Architecture Guide. The procedure started with creating the physical network topology with three separate physical networks: 1) a Tenant network which carried Data between the VNFs 2) a Management network which carried control and management traffic between the controller and compute nodes, and 3) an External/Public network which enabled the VNFs and VMs in the system to have Internet connection. Once these 3 networks were wired up, installation of Fedora* 21 was initiated, followed by updating the kernel and other specific packages. Then the network was configured and tested by simple ping commands. The controller node was then configured by running a set of scripts, which had to be edited first to specify the value of several parameters, such as the network interfaces to use. The whole operation was then repeated to setup the compute node. Even though the procedure included several manual steps (which could be simplified if there was a wizard), it was still fairly quick and straightforward.

The installation procedure for the Wind River® Titanium Server also started out by creating the physical network topology with the same three physical networks. However, one additional network was required to enable Wind River's high-availability (HA) functionality. Another difference from the Intel ONP 1.5 setup was that each Wind River® Titanium Server was configured as both a controller and compute node. The installation process was extremely user friendly and interactive. Installation of the second node was easier, but took a longer time since all the system data had to be replicated on the second node for HA support. Use of RAID and SSDs would have helped expedite the process.

Test 1: VNF Onboarding

The process of VNF onboarding was started by loading the VNF into the system using Glance or Heat (depending on the VNF's requirements). Once loaded, each VNF was booted by instantiating it from OpenStack and connecting it to all the required networks. The next step was to start the VNF and verify its basic functionalities, excluding any traffic or benchmarking considerations. Finally, the VNF was rebooted and the basic functionalities were re-examined. Table 3 captures the test results of this VNF onboarding test.

All VNFs that successfully loaded and booted on both NFVI platforms (Intel® ONP 1.5 and Wind River® Titanium Server) were considered as a 'pass'. The most common reason for VNFs failing to boot or run was the lack of support for the latest available version of OpenStack (which at the time of these tests was OpenStack Kilo).

[4] http://www.saguna.net/news-events/press-releases/akamai-and-saguna-network-aware-for-encrypted-content-optimization-solution-nominated-for-the-2015-global-mobile-awards/?utm_source=homepage?utm_medium=button-3?utm_campaign=Q1-2015

VNF	LOADING	BOOTING	RUNNING	STATUS AFTER REBOOT
Router I	Passed - Glance	Passed	Unsuccessful	Unsuccessful
Router II	Passed - Glance	Passed	Passed	Unsuccessful
Router III	Passed – Heat	Unsuccessful	N/A	N/A
Firewall	Passed – Glance	Unsuccessful	N/A	N/A
Load balancer	Passed – Glance	Passed	Passed	Passed
Web Server	Passed - Glance	Passed	Passed	Passed

Table 3. Results of VNF onboarding tests

Test 2: VNF Static Service Function Chains

First a static service function chain was setup between the VNFs (e.g. Router -> Firewall -> Load Balancer -> Web Server). Then HTTP traffic was passed and performance measurements taken. A traffic generator was created by building a script that was a wrapper around the `curl` tool. The script was designed to support consecutive HTTP requests with different transmission rates of 100 Mbps, 500 Mbps and 1000 Mbps. These HTTP requests were passed through the service function chain. In other words, they went first through the virtual router, then through the virtual firewall, then through the virtual load balancer which balanced the HTTP traffic between a virtual pool of three CentOS-based Web Servers. HTTP responses from the Web Server that handled the request were then sent back to the traffic generator. To minimize the impact on performance of the NFVI system under test, the traffic generator and receiver was deployed on a separate node, and all throughput measurements were taken on that node.

Placement of the VNFs seemed to be challenging when measuring data on both platforms. One of the constraints of the interoperability exercise was to use a minimum setup. For Intel® ONP 1.5, that consisted of two servers: one dedicated as the controller node and the other dedicated as the compute node. Here all the VNFs were placed on the compute node as depicted in Figure 2. The Wind River® Titanium Server platform used a different approach. Its minimal setup also consisted of two servers, but each acted as a controller and compute node. Consequently, the same set of VNFs were spread between both servers as shown in Figure 3.

To compare the results generated on both NFVI setups, a baseline was created, where HTTP requests were sent directly from the traffic generator to the Web Servers (avoiding the static service function chain).

One of the limitations of the `curl` tool was that it did not limit traffic perfectly. Bursts of traffic sometimes exceeded the set limit. This is the main reason why we saw some anomalies in the results shown in Figure 1 , such as 103Mbps throughput on ONP 1.5 setup when the limit was set to 100Mbps.

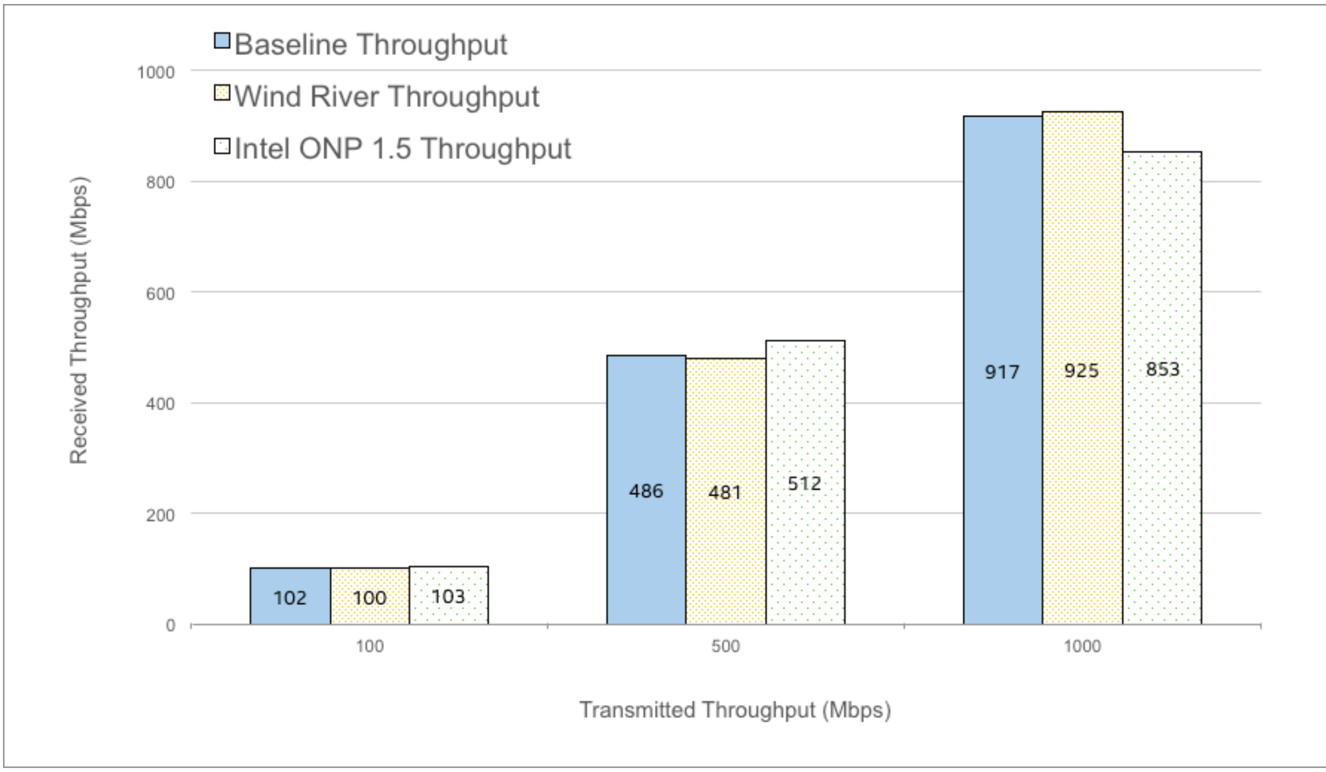

Figure 1. Throughput comparison between two NFVI platforms with the same setup of VNFs

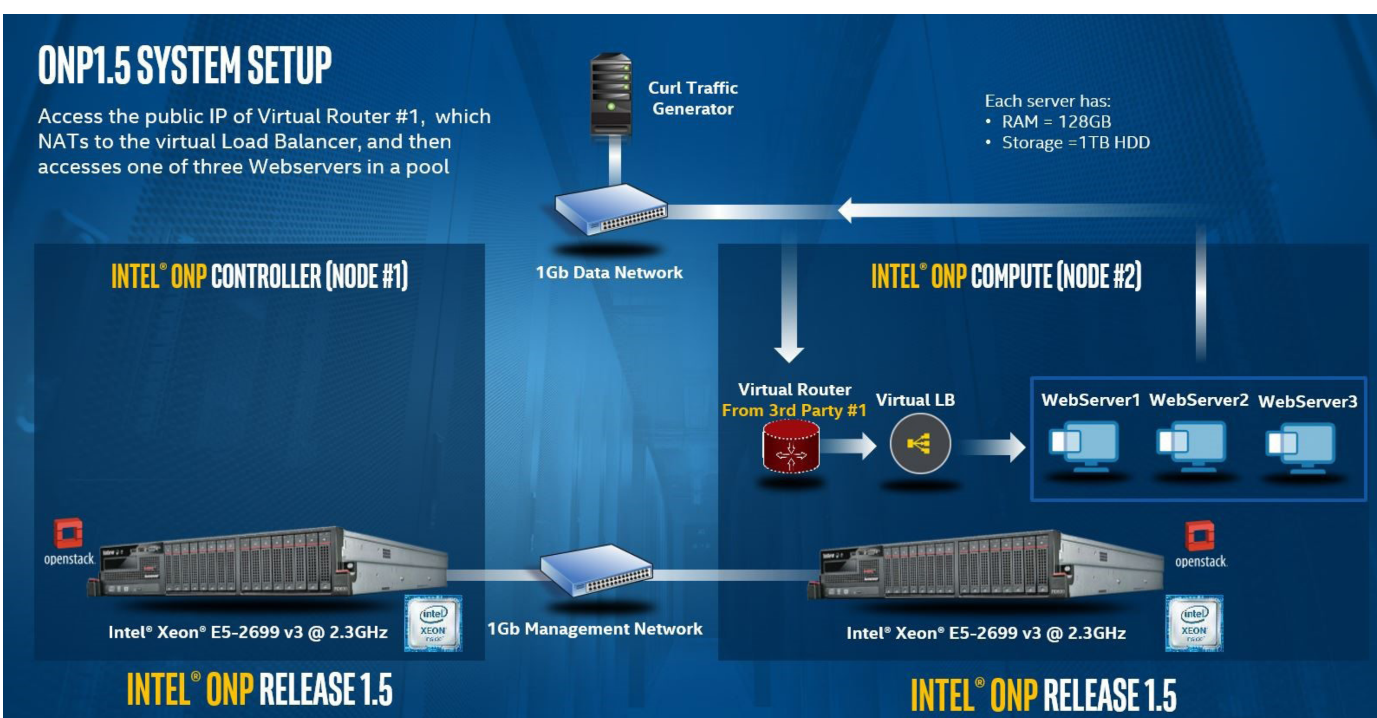

Figure 2. Intel® ONP 1.5 setup with all VNFs deployed on the compute node

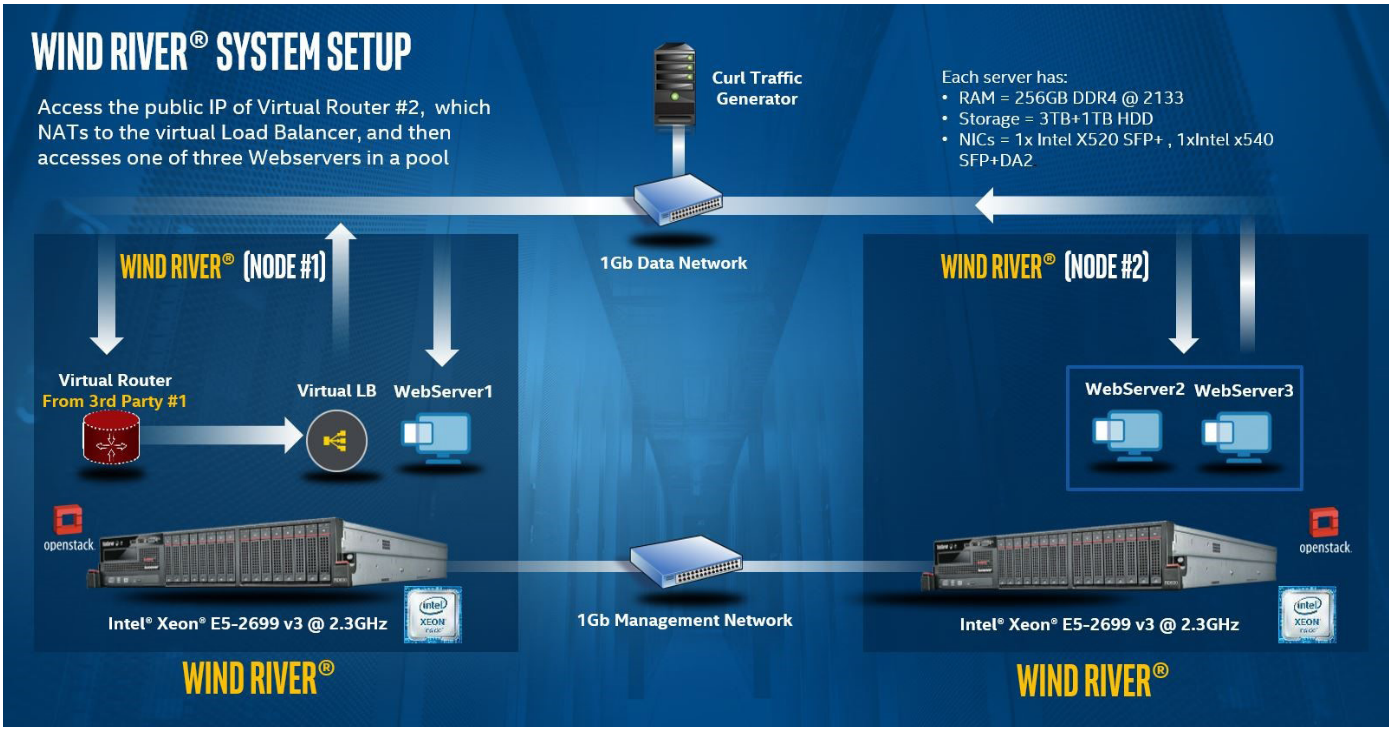

Figure 3. Wind River® Titanium Server setup, with VNFs deployed on both compute nodes

Findings and Challenges

This was the first edition of the Intel® Interoperability Program. It helped provide participating NFVI and VNF vendors with ample opportunities to learn and identify several issues related to onboarding and configuration.

The top 5 findings and challenges related specifically to NFVI onboarding and configuration were as follows:

1. **VNF-specific Configurations:** Loading VNF images required specific components. There were more tools preinstalled on the Wind River® Titanium Server system. This was useful since VNFs may require different onboarding tools (e.g. Glance vs Heat) for different NFVI. For end users like CSPs who may have different NFVI deployments, configuring the same set of VNFs across NFVI setup using different tools would be inconvenient and time consuming.

2. **VNF Configuration Dependencies:** Different VNFs required different CPU flags. Setting all necessary CPU flags was not always possible especially when the same flag had to be set with different values for different VNFs. This limited VNF deployment options, since VNFs could not co-exist on the same host. As we changed CPU flags, we found ourselves having to re-install the Intel® ONP 1.5 setup. Since Intel ONP 1.5 used DevStack, reconfiguration of the setup was necessary after a reboot. On the other hand, Wind River ® Titanium Server was more stable, and didn't require reconfiguration after reboot.

3. **Traffic Measurement:** Preparing detailed traffic test results required an additional node for traffic generation and measurement. Some differences in the two topologies made it difficult to fairly compare the results. One of the key difference was that Intel® ONP 1.5 had separate compute and controller nodes, while the Wind River® Titanium Server had both its nodes configured in redundant fashion as both compute and controller.

4. **Setup Time:** Setup of the Intel® ONP 1.5 topology took about 3 hours per node for a total of 6 hours for our minimal setup. On the other hand, setup of the Wind River® Titanium Server topology took about twice the time (~12 hours), with half of the time spent on synchronizing the two nodes.

5. **OpenStack Dashboard:** The OpenStack dashboard was more advanced in the Wind River ® Titanium Server setup and provided more usability and configuration options.

The top 5 findings and challenges related specifically to VNF onboarding and configuration were as follows:

1. **OpenStack Version:** Occasionally we found that several VNF providers were not supporting the latest versions of OpenStack, which limited the variety of VNF implementations we could use in our NFVI setups.

2. **VNF Configuration Complexity:** Each VNF had very specific requirements which made the process of onboarding very slow and error prone. A solution to this issue is planned to be addressed during the next round of Intel Interoperability activities. Intel will work towards proposing common configuration methods for similar types of VNFs on different NFVIs.

3. **VNF Settings:** There are some VNF settings, such as specific CPU flags, which must be set only during the installation process. This does not seem to be an effective approach in a virtual environment.

4. **Licenses:** Handling licenses in a virtual environment is complicated. Most of the VNF licenses were pinned to the system's UUID which usually changed after system reconfiguration and restart. As a result, it was occasionally not possible to reuse the same VNF license after a VNF was terminated and redeployed on the same NFVI.

5. **Limited Debugging Tools:** There were no standard ways of debugging VNFs during the onboarding process. The Industry seems to lack robust tools for tracing VNF boot progress. Due to the lack of logs and other convenient methods of capturing errors, multiple live debugging sessions with VNF vendors were required, which was very time consuming.

Call to Action - ISVs and the Ecosystem

1. **Implementing Standard Benchmarking:** There is currently no standard way of performing benchmarking among various VNFs and NFVIs. This is very problematic as it means CSPs cannot compare performances of VNFs and NFVIs off the shelf, but instead need to build PoCs to see how a handful of these might perform. Ecosystem collaboration is required to identify a standard set of tools and models for measuring performance and benchmarking in an SDN/NFV topology.

2. **Providing L3 Plugins:** Several VNFs only came with an L2 Plugin. L3 VNF plugins are recommended for complete functionality, and easier manageability.

3. **Avoiding dependency on CPU Flags:** ISVs building VNFs should avoid requiring specific CPU flags to be set on the host, since that limits flexibility in deploying different VNFs on the same host.

4. **Easily adaptable NFVI Configurations:** Simple configuration changes in an NFVI setup (such as updating network settings or adding CPU flags) would end up requiring a rebuilding of the node. Dedicated efforts are required to address issues causing instability in the infrastructure and enable better scalability as solutions grow and become more complex.

Intel® Interoperability Next Steps

Solutions for the challenges and issues discussed in the previous section will be systematically addressed in upcoming Intel Interoperability activities in collaboration with VNF and NVFI vendors. The first goal of the next round of Interoperability activities will be to propose a common configuration approach for VNFs of the same type. This objective is crucial not only for VNF providers, but also for CSPs. Through availability of common configuration methods, it will be much easier to onboard VNFs with a minimal set of common working configuration commands. This in turn will enable quick execution of tests of various VNFs on several different NFVI platforms and thereby accelerate adoption and deployment for CSPs as they advance in their virtualization journey.

Acronyms

ATIS	Alliance for Telecommunications Industry Solutions
AVS	Accelerated Virtual Switch
CPU	Central Processing Unit
DPDK	Data Plane Development Kit
ETSI	European Telecommunications Standards Institute
GbE	Gigabit Ethernet
HA	High Availability
HDD	Hard Disk Drive
ICT	Information and Communications Technology
Intel® ONP	Intel® Open Network Platform
ISG	Industry Specification Group
ISV	Independent Software Vendor
KVM	Kernel-based Virtual Machine
NFV	Network Function Virtualization
NFVI	NFV Infrastructure
NIC	Network Interface Card
OEM	Original Equipment Manufacturer
PoC	Proof of Concept
QEMU	Quick Emulator
RAID	Redundant Array of Independent Disks
RAM	Random Access Memory
SDN	Software-Defined Networks
SDND PAE	Software Defined Networking Division, Platform Application Engineering group
SSD	Solid State Data
TEM	Telecommunications Equipment Manufacturer
UUID	Universally Unique Identifier
VNF	Virtual Network Function
vSwitch	Virtual Switch