END-TO-END OPTIMIZED NETWORK FUNCTION VIRTUALIZATION DEPLOYMENT

How intelligent Network Function Virtualization (NFV) orchestration based on industry standard information models will enable carrier-grade Service Level Agreements (SLAs).

Executive Summary

Network Functions Virtualization (NFV) has been capturing the Telco industry's imagination for several years as an innovative way of improving service agility, reducing costs, and creating greater flexibility in the Telco data center. NFV's promises and benefits are clear and well understood. Intelligent, automated deployment of Virtualized Network Functions (VNFs), however, still presents complexities that the Telco industry and VNF providers must contend with.

Cyan, Brocade, Intel, and Red Hat recently collaborated with Telefónica's NFV Reference Lab in Madrid, Spain to showcase an effective NFV deployment scenario using Enhanced Platform Awareness information models and a standards-based NFV infrastructure running on Intel® Xeon® processor-based servers. The deployment showed how a Brocade* Vyatta* vRouter VNF can achieve 85X performance in an intelligent, optimized deployment over a non-optimized deployment, by taking advantage of the technologies available in the infrastructure during deployment.

This paper describes the issues that Telcos and VNF providers face and the solutions an optimized, industry standards-based deployment delivers.

Introduction

Network functions, like Evolved Packet Core (EPC), 3G wireless nodes, Broadband Network Gateways (BNG), Provider Edge (PE), routers, firewalls, etc., traditionally have been delivered on dedicated hardware appliances. But, the recent emergence of Virtualized Network Functions (VNFs) will replace this hardware-centric approach through software appliances instantiated in a carrier-grade virtualization environment on today's Intel® Xeon® processor-based off-the-shelf servers. The Telco industry understands and accepts this transformation to a Network Functions Virtualization (NFV) approach as key to making their businesses more agile, their networks more adaptable, and reducing their total costs of ownership (TCO).

While the NFV approach opens new potentials for both Telcos and the growing ecosystem of VNF appliance providers, the VNF components alone are only part of the solution. Deploying a Telco-grade VNF service presents new complexities both the Telcos and providers must deal with.

Table of Contents

Telco-Grade Deployment Needs

A Service Level Agreement (SLA) often accompanies each VNF instantiation. The SLA needs to be attained and enforced, while achieving scalability as service adoption ensues. These requirements place essential focus on how data plane workloads are handled in terms of throughput, packet loss guarantees, and latency effects. Such attributes greatly affect the VNF appliance's performance. Thus, control of network topology, VNF location, link bandwidths, and QoS guarantees, as well as awareness of the hardware capabilities in underlying server platforms, are critical in Telco.

Both Telco service providers and the appliance developers need tools and methodologies to help ensure the instantiation is launched in a manner that commits sufficient platform and network hardware resources to meet SLAs. Tools need to take into account the internal server memory topology, CPUs and I/O interfaces allocated to virtual machines, the usage of memory in "huge pages" for efficient lookups, and direct assignment of interfaces to the VM, among others. Awareness of these attributes becomes essential to help assure a given SLA in terms of performance, scalability, and predictability.[1]

Standardizing Telco NFV Deployments

Virtualized network functions are delivered by an ecosystem of many different vendors and deployed by Telcos across heterogeneous network and data center topologies and technologies. Industries deal with creating a common ground in the presence of diversity through open industry standards that bring greater predictability to the environment. Standards, such as PCI Express*, SATA, and others have enabled ecosystem growth in the PC and server industry for many years.

Such an approach is underway within the Telco space, managed by the European Telecommunications Standards Institute (ETSI), which defines a Network Functions Virtualization Infrastructure (NFVI) on which Telcos can effectively and efficiently deploy VNFs using a description language that exposes capabilities and needs.

The ETSI-NFV reference architecture defines a layered approach to VNF deployments (see Figure 1).

To help ensure portable and deterministic performance across a NFV-based service deployment and operation, the infrastructure must expose the relevant NFVI attributes up through the delivery stack. Likewise, the VNF's information models, describing its resource requirements and those for the services being launched, are key to enabling the provisioning layers to make intelligent and optimal deployment decisions. This Enhanced Platform Awareness (EPA) capability in the NFVI allows the orchestration platform to intelligently deploy well-designed VNFs onto the appropriate underlying infrastructure, and it helps ensure correct allocation of the resources for an end-to-end VNF service scenario.

Carrier-Grade VNF Deployment in Practice

To demonstrate the concepts described above, Intel, Telefónica, Cyan, Brocade, and Red Hat collaborated to implement a complete ETSI-NFV end-to-end service deployment solution stack using the TOSCA description language. The proof of concept (PoC), which was hosted by Telefónica in their NFV Reference Lab in Madrid, Spain, included the components listed in Table 1 (see also Figure 2). The solution stack was then tested using these components in a VNF routing service scenario.

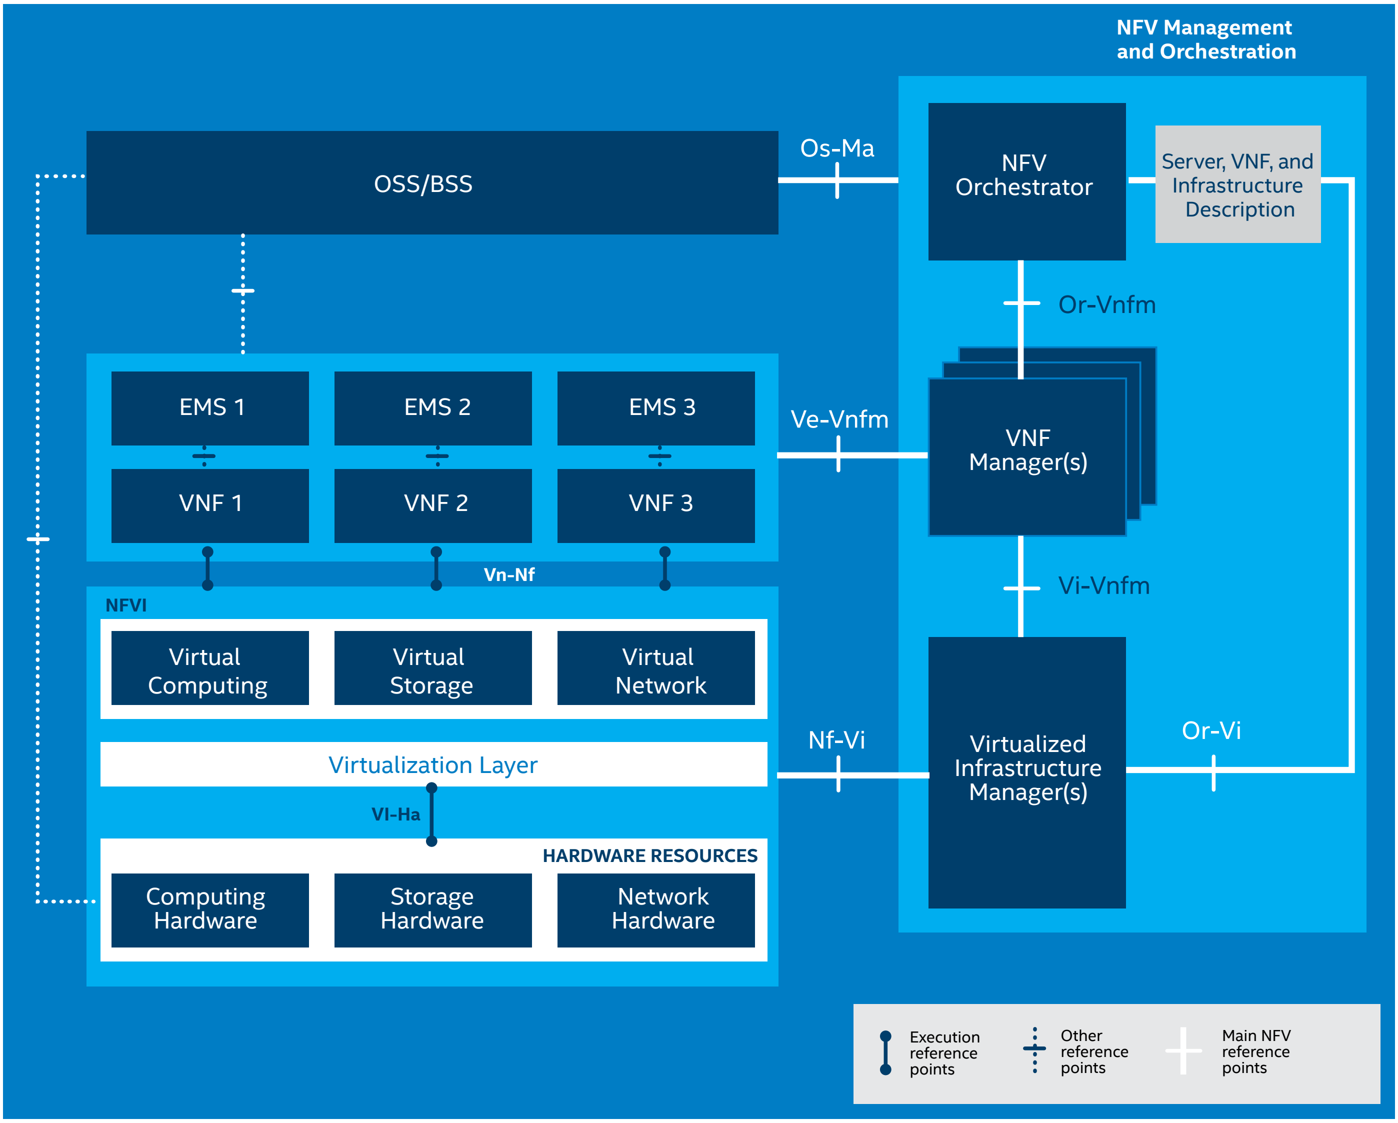

Figure 1. ETSI End to End NFV Architecture.

VNF Routing Service Overview

A VNF forwarding graph was routed using Brocade Vyatta vRouters as VNFs. A three-node network forwarding topology was designed to achieve 40 Gbps throughput between the ingress and egress points at routers A and C (Figure 3). The PoC intended to showcase the following:

- An end-to-end NFV service delivery stack with the relevant NFV intelligence built in at each layer—from the information model to the VNF, the NFV Orchestrator, the VIM, and the NFVI—is required for an optimal VNF service chain deployment.
- The exposure of the performance-enabling NFVI attributes in the VNF descriptor and the importance of a good VNF design to define the VNF's (Vyatta vRouter) and service's requirements are crucial toward enabling an effective service deployment.
- The use of Industry-standard, open, and extensible information models, such as TOSCA, and suitable VNF formats are critical for the ecosystem of VNF vendors to deliver their services into this new end-to-end architecture.
- High-performance, well-architected servers and networking components are necessary to provide the performance required for the deployment of Telco-grade VNFs.

Optimized versus Non-optimized

To contrast the differences in non-optimized and NFV-based optimized Telco deployments, two separate system environments were constructed (see Figure 4):

- A NFV-ready NFVI pool with a Telefónica-developed, NFV-ready Virtualized Infrastructure Manager (VIM) implementing the requisite EPA, plus a Cyan NFV-Orchestrator, supporting advanced VNF deployment using enhanced NFV information models and the TOSCA description language.

- A non-optimized, computing infrastructure pool with the same Telefónica VIM connected to the same Cyan NFV-Orchestrator, but without an enhanced information model as the basis for deployment.

Starting with both environment server pools empty (no VNFs deployed), the routers were deployed onto each platform through the same orchestrator. With both systems running, throughput performance was measured in real time. The NFV-optimized deployment resulted in a full line rate throughput of 23 Mpps versus 270 Kpps for the non-optimized deployment (85X faster throughput).

Information models for both scenarios were compared, showcasing the key additional attributes the end-to-end EPA delivers in the optimized NFV deployment.

Intel	Brocade	Cyan	Telefónica	Red Hat
Hardware Platforms: • Servers based on Intel® Xeon® processor E5-2680 v2 @ 2.80 GHz[3] • Intel® Open Network Platform (ONP) ingredients, including DPDK R.6[4] • 10 Gb Intel® Ethernet Converged Network Adapter X520	Virtualized Network Function: • Brocade* Vyatta* vRouter 5600 3.2 R2 • OpenFlow* switch (Brocade* NetIron* MLXe)	Orchestration: • NFV-Blue Planet* Orchestrator release 15.02	Network and Management: • Data Plane Development Kit (DPDK) R1.6-based traffic generator TIDGEN (Telefónica* I+D Generator* • Telefónica VIM* openvim R0.9	Operating Environment: • RHEL 7.0* (with patches) and QEMU-KVM version 2.0.0 (with patches)

Table 1. Solution Stack Components.[2]

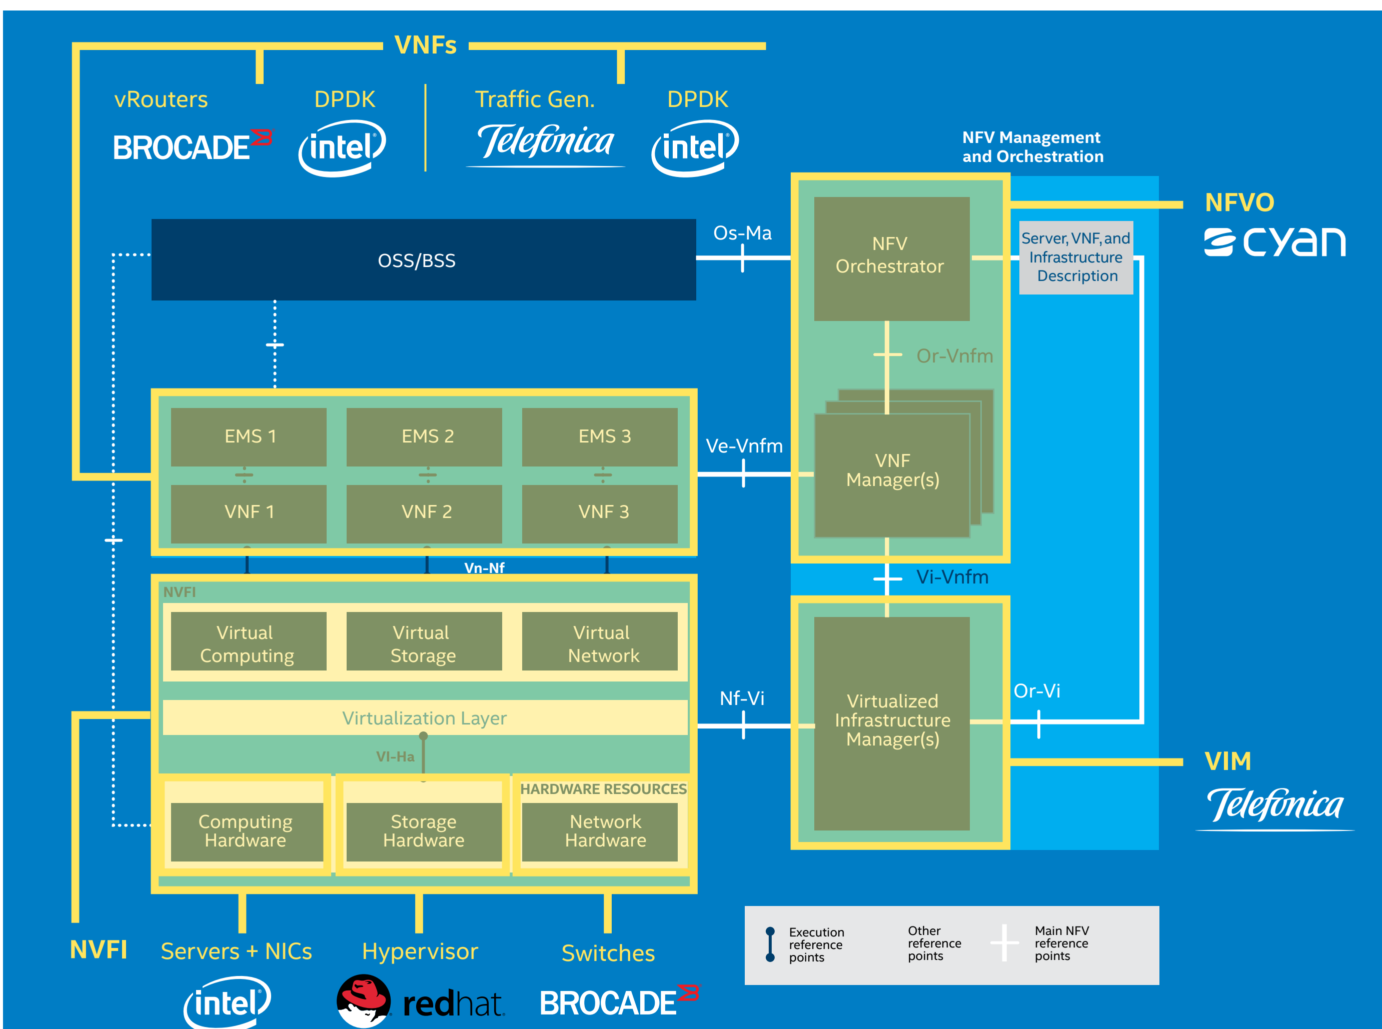

Figure 2. Partners and System Component Contributions.

Results

Non-optimized Deployment

The Brocade vRouter was deployed through an "unaware" NFVI stack. It could only achieve a throughput of 270 Kpps, instead of its designed 23 Mpps (40 Gbps @ 192 byte packet size). See Figure 5.

Optimized NFV Deployment

The optimized NFV deployment uses TOSCA and well-designed VNF models. The information in the models allowed the Planet Blue orchestrator to intelligently deploy the vRouter configuration through the Telefónica VIM. The result was full line rate performance of 23 Mpps (40 Gbps @ 192 Bytes) as shown in Figure 5. Similar line rate performance for larger packet sizes was achieved as shown in Figure 6. The vastly improved performance was largely attributed to the following:

- PCIe* pass through enabled: The NIC was directly connected to the vRouter, bypassing the vSwitch. A non-PCIe-pass through path to the VNF limits the throughput rate with an acceptable packet loss, while the PCIe pass through mode allows unhindered communication.
- NUMA affinity: vCPUs were intelligently allocated from an appropriate CPU socket with local memory, potentially accelerating CPU performance for the vRouter. Without NUMA affinity, a CPU could have been assigned from any socket, possibly not directly attached to the NICs, and might also use a non-local memory bus.
- CPU pinning: vCPUs allocated to the vRouter were dedicated to the VNF, helping to ensure dedicated processor resources to the VNF. Without pinning, CPUs could be shared or dynamically rescheduled, limiting determinism and, thus, performance.
- 1G huge page setup. Huge pages greatly improve achievable performance for the DPDK. It also leverages the recent advances in server IOTBL and VTd architectures, especially for small packet sizes.

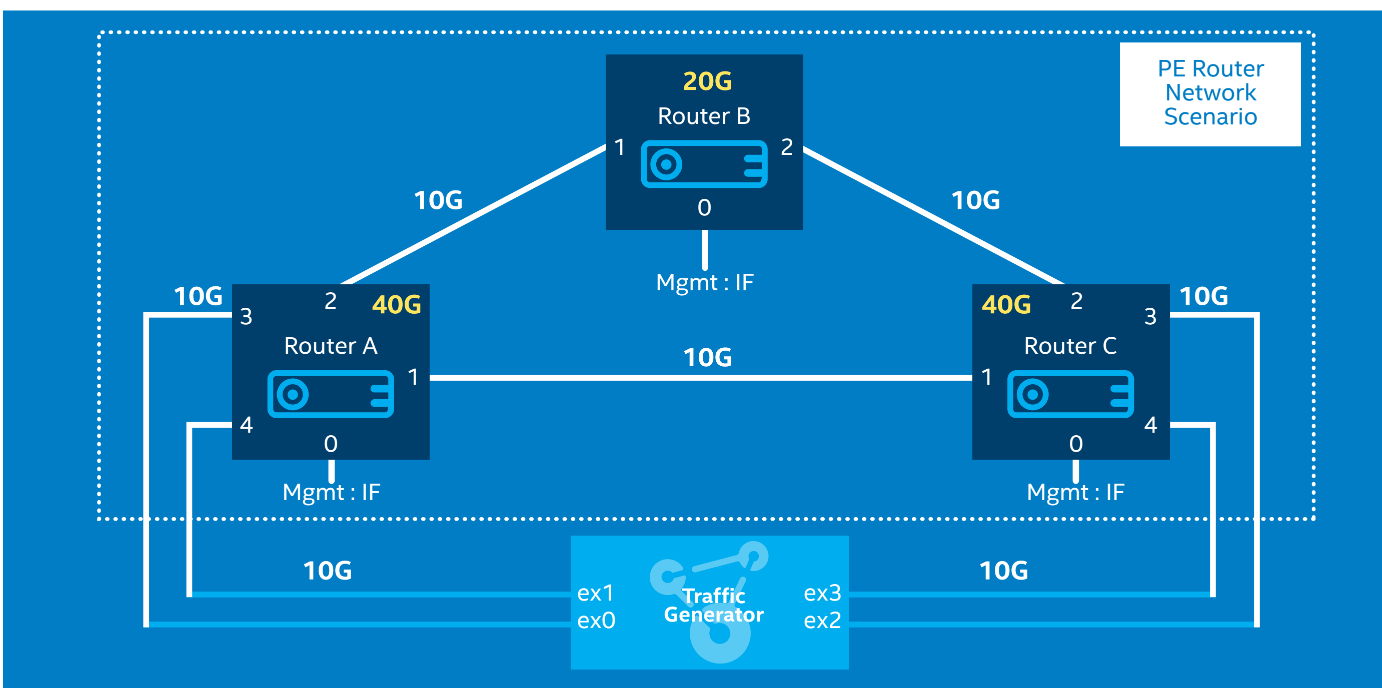

Figure 3. PE VNF Routing Service Chain.

Conclusions

Performing intelligent, Enhanced Platform Awareness-based NFV deployments onto an ETSI-NFV infrastructure with TOSCA and extended information models can significantly improve VNF performance to meet the requirements of an NFV-based Telco data center. Using the enhanced models in the PoC, the Brocade vRouter was able to achieve the designed performance by implementing PCIe pass through, NUMA affinity, CPU pinning, and huge pages during deployment. In summary the PoC highlighted the following,

- End-to-end NFV deployments will enable the significant performance improvements necessary for Telco-grade services performance.
- A well-designed VNF, with EPA attributes defined for the appliance, will help expose appliance and service requirements up through the stack. The VNF developer must design these attributes into the VNF model.
- Standard and open information models enable an open VNF ecosystem and the transition from hardware-based appliances to software-based NFV.
- The standardization of the NVF service information model as well as the availability of open source components, such as DPDK, Open Stack, and optimized KVM, are key components toward unleashing the promise of open NFV solutions that can leverage best of breed cloud open source technologies and optimized, high-performance servers. Intel, Red Hat, Cyan, and Telefónica will continue to work to enable VIMs, such as OpenStack, with these critical NFV EPA enhancements.[5]

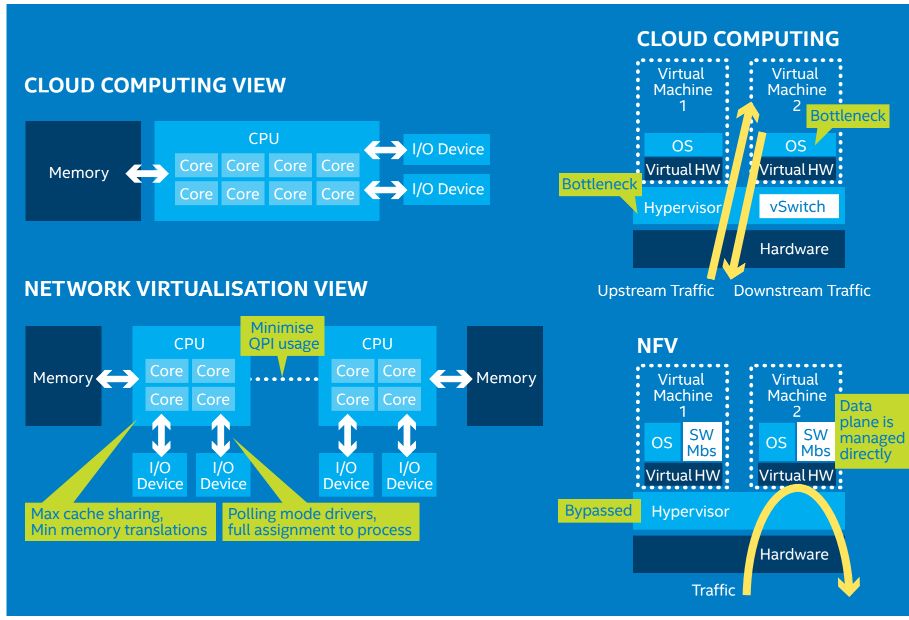

Figure 4. Cloud vs. NFV.

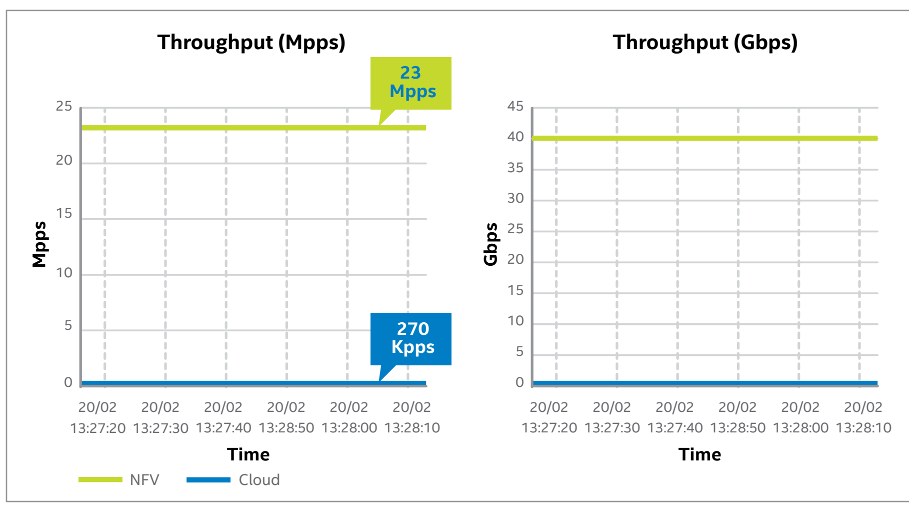

Figure 5. Performance Comparison for 192 byte frame size.

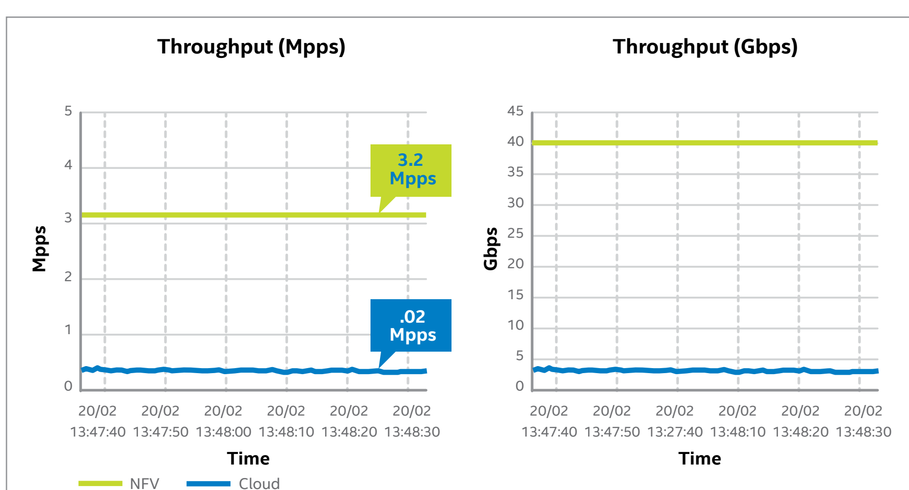

Figure 6. Performance Comparison for 1,518 byte frame size.

Testimonials

Brocade

"Brocade welcomes the advancements in intelligent orchestration, continued partnership within open initiatives and execution toward key NFV standards. The flexibility and openness of Intel's Network Builders Community has brought together committed partners dedicated to accelerating the industry's transition to the New IP. The combined efforts of partners such as Telefónica, Intel, and Cyan highlight key architecture benefits of Brocade's VNF platforms, the Vyatta 5600 vRouter, and its inherent open information data model for facilitating a migration to intelligent architectures with high performance. This also highlights the value of NFV orchestrators and their importance to effective and optimal network deployments, with Telefónica leading the charge to demonstrate NFV without sacrifice."

– Robert Bays, VP of Engineering, Brocade Software Networking

Cyan

"The intelligent NFV orchestration and placement PoC with Telefónica at Mobile World Congress is a clear example of the power of collaboration as it relates to driving real-world NFV use cases," said Mike Hatfield, president, Cyan. "The multi-vendor platform provides a unique framework for showcasing how Brocade's VNF and Telefónica's VIM can expose performance requirements and characteristics to Cyan's enhanced infrastructure aware NFV orchestrator. The orchestrator will intelligently place the VNFs on Intel servers to meet the VNF's specific performance needs and efficiently use compute resources to deliver end-to-end services. This is an important issue that needs to be solved by the industry for deployment of NFV-enhanced services at massive scale."

– Mike Hatfield, President, Cyan

Intel

"Intel believes SDN-NDV is an industry inflection point and is committed to ensuring the new network architecture transformation is built on an open architecture, using open standards enabling an open eco system. Intel is committed to delivering NFV and is actively working through the relevant standards and open source initiatives toward making this a reality. Intel will makes all its ingredients open source[6] though its Open Networking Platform program and is working closely with its Netbuilders SDN-NFV ecosystem community[7] partners such as Cyan, Brocade, and Telefónica to make this a reality."

– Rene Torres, Intel SDN-NFV Marketing Director

Red Hat

"Building the foundation for an open NFV infrastructure requires expertise in Linux, KVM, and OpenStack—all areas of open source where Red Hat is a leading contributor," said Radhesh Balakrishnan, general manager, OpenStack, Red Hat. "By collaborating on the NFV Reference Lab, we're not only bringing features and expertise back to the upstream OpenStack community and our carrier-grade Red Hat Enterprise Linux OpenStack platform, but also enabling CSPs to successfully implement their modernization plans through NFV."

– Radhesh Balakrishnan, General Manager, OpenStack, Red Hat

Telefónica

"Telefónica's vision about Virtualized Network is an E2E virtualization approach, from customer premises to the inner network infrastructure, as a way to improve capacity and flexibility and to obtain better TCO. Telefónica NFV Reference Lab aims to help the ecosystem of partners and network equipment vendors to test and develop virtualized network functions leveraging on an advanced NFV orchestration framework and proper capabilities for deterministic resource allocation in the pool. NFV Reference Lab drives this adoption through the release of open source code, thus encouraging software developers to explore new NFV possibilities and all this from a well-designed and tiered architecture proposal. Its aim is to promote interoperability and provide a more open ecosystem so that telecommunications providers adapt and expand their network services more easily."

– Enrique Algaba, Network Innovation and Virtualisation Director, Telefónica I+D-Global CTO

Acronyms

BNG Broadband Network Gateway

BSS Business Support System

CMS Cloud Management System

CPU Central Processing Unit

vCPU Virtual Central Processing Unit

DPDK Data plane Development Kit

EPC Evolved Packet Core

EMS Element Management System

EPA Enhanced Platform Awareness

GCTO Global Chief Technical Office

IOTLB I/O Translation Look Aside Buffer – Virtualization Technology

NiC Network Interface Card

NFV Network Function Virtualization

NFVI Network Function Virtualized Infrastructure

NFV – O Network Function Virtualization Orchestrator

NUMA Non Uniform Memory Access

OSS Operations Support System

PE Provider Edge Router

PCIe Extensible Peripheral Connect Interface Bus

QoS Quality of Service

SLA Service Level Agreement

TCO Total Cost of Ownership

VIM Virtual Infrastructure Manager

VNF Virtual Network Function

VT-d Intel® Virtualization Technology for Direct I/O

1. ETSI GS NFV-PER 001 V1.1.2 - "Network Functions Virtualisation (NFV); NFV Performance & Portability Best Practises" http://docbox.etsi.org/ISG/NFV/Open/Published/gs_NFV-PER001v010102p%20-%20Perf_and_Portab_Best_Practices.pdf
2. 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.
3. For more information go to http://www.intel.com/performance
4. http://www.intel.com/content/dam/www/public/us/en/documents/white-papers/open-network-platform-server-paper.pdf
5. https://software.intel.com/sites/default/files/managed/72/a6/OpenStack_EPA.pdf
6. http://www.etsi.org/deliver/etsi_gs/NFV/001_099/002/01.02.01_60/gs_NFV002v010201p.pdf
7. https://networkbuilders.intel.com/docs/Intel_Network_Builders_Directory_Sept2014.pdf