White Paper
Intel, Brocade, Cyan, Red Hat, and Telefónica – NFV Services

End to End Network Function Virtualization Architecture Instantiation

How Intelligent NFV Orchestration based on industry standard information models will enable Carrier Grade SLAs

Executive Summary

For over 2 years now, since the ETSI NFV ISG inception in Darmstadt October 2012[1], NFV has been capturing the Telco industry imagination, its promises and benefits are clear and well understood. Since its inception, the industry has seen huge improvements in base technology layers, including server's (Intel), hypervisor technology (Red Hat), and software libraries (Intel® DPDK), enabling the design of telco-grade Virtual Network Functions (VNFs) (such as Brocade's VRouter 5600*), which have proliferated and evolved extensively during this period. As a direct result, a new and open ecosystem of VNF providers is beginning to flourish.

However, this is only part of the industry issue. These VNFs designed for carrier-class scalability, in order to behave as expected, need to be properly deployed in the underlying infrastructure allowing them to leverage all those new technology advances. NFV Management and Orchestration (MANO) and the associated information models, describing both the infrastructure and VNF requirements, are key to achieve this goal effectively and in a cost efficient manner for the service provider. Hence, legacy cloud management systems (CMS) will simply not suffice for true NFV deployments.

Cyan, Brocade, Intel, and Red Hat have combined with the Telefónica NFV Reference Lab at their GCTO Unit in Madrid to showcase how a realistic network scenario can be designed modelled and deployed via NFV Orchestration (Cyan Blue Planet) onto an NFV-ready infrastructure through the Telefónica design VIM (Virtual Infrastructure Manager). This new and correctly optimized NFV delivery stack is compared to what can be achieved with a typical cloud deployment model as exists today.

The results show the phenomenal benefits achievable through end to end system NFV awareness. The service scenario deployed in Telefónica's NFV Labs in Madrid shows up to a 100x improvement in throughput for a typical routing scenario with respect to the same scenario as deployed in a typical enterprise cloud.

Key to unleashing this performance is the correct modelling of the key attributes required from Virtual Network Functions (VNFs), and exposing this information as the deployment decision criteria in the NFV delivery stack, i.e., the NFV Orchestrator and the Virtual Infrastructure Manager (VIM). The availability of such NFV-ready orchestration components together with appropriate standardized descriptors for VNFs and infrastructure will be the key enablers to large-scale NFV deployments in coming years.

Table of Contents

Introduction

A key component of the NFV vision is one of “The Network becoming a Data Centre,” a radically new network enabled through leveraging the commodity price points and capabilities emerging from the $30-$50 billon per annum global investment in data center technology to enable the delivery of Telco grade virtual network appliances on top as VNFs.

Network functions like Evolved Packet Core (EPC), 3G wireless nodes, Broadband Network Gateways (BNG), Provider Edge (PE), routers, firewalls, etc., have traditionally been delivered on bespoke standalone appliances. NFV aspires to replace this hardware centric approach with a software model which delivers comparable functionality as SW VNFs on standard high volume industry server HW. This transformation is referred to as NFV and the concept is well understood and accepted by the Telco industry as a key lever in the transformation of the network toward a more flexible and mouldable infrastructure.

However this in itself is not sufficient. Deploying a Telco Network service presents additional complexities that typically don’t exist in today’s data center:

- For example, each Telco service delivered to the broadband consumer comes with a service SLA which must be achieved and enforced. This must also take into account how to achieve proper scale as service adoption ensues. This places essential focus on how data plane workloads are handled in terms of throughput, packet loss guarantees, and latency effects. These are the attributes which most affect the virtual application performance and which Telco service providers must ensure as part of a customer’s SLA guarantees.
- Additionally, control on network topology, VNF location, link bandwidths, and QoS guarantees are hugely important in Telco. This is the foundation on which Communication Service Providers must deliver their services. This approach deviates greatly from the centralized data center paradigm, where the topology is mostly static and where visibility and basic placement considerations for the stand-alone VMs are the primary attributes required for service continuity.
- In this new NFV world, the virtual network functions will be delivered by many different vendors. Unless the community embraces a very well understood and open standard based service information model this new flexibility will become difficult to manage and will in itself become a problem. The proposal in this white paper and the associated E2E implementation is to use TOSCA as the service description language. This also enables an easy extension path to including the key NFV attributes required for this deterministic performance.

As described earlier, huge industry investment has enabled industry standard high volume servers[1] to deal effectively with I/O-intensive workloads as required in today’s Telco environments. Thus, most recent x86 processors generations working in conjunction with suitably enabled hypervisors and using specialized open source software libraries such as DPDK (Data Plane Development Kit) have enabled standard high volume servers to deal efficiently with edge functions such as BNG, PE router, and EPC workloads. This creates the opportunity of enabling reliable NFV deployments ensuring that true Telco grade SLAs are achieved.

In order to enable these new Telco grade services, it is essential that the appropriate infrastructure resources are properly allocated to the VNF. Thus, practices such as taking into account the internal server memory topology, CPUs and I/O interfaces allocation to virtual machines, the usage of memory in “huge pages” for efficient lookups, or direct assignment of interfaces to the VM, among others becomes essential to assure a given VNF SLA in terms of performance, scalability, and predictability[2]. This type of deterministic resource allocation, including this

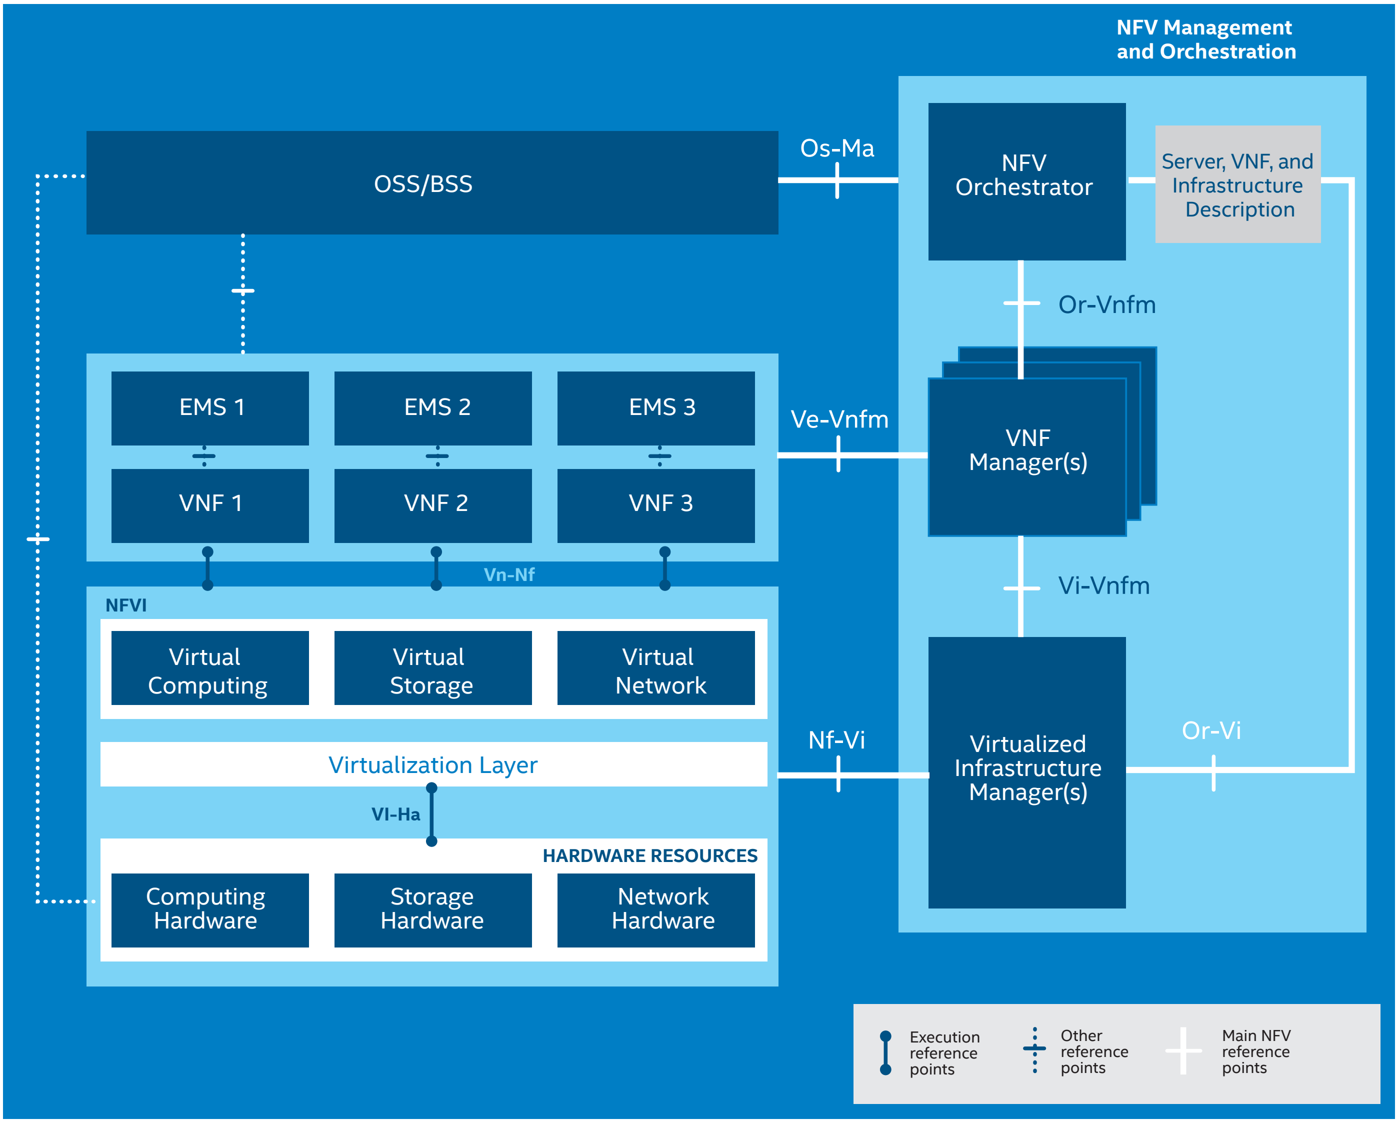

Figure 1. ETSI End to End NFV Architecture.

new enhanced platform awareness capability, not present in cloud computing environments, now becomes a necessity for carrier grade NFV deployments.

The ETSI-NFV reference architecture defines a layered approach to such an NFV deployment (see Figure 1).

Toward ensuring portable and deterministic VNF performance it is paramount to expose the relevant NFVI attributes up through the new NFV delivery stack. This allows the management and orchestration layers to ensure correct allocation of the resources for the end-to-end network service scenario. Likewise, the information models (Network service VNF and Infrastructure description) describing the resource requirements for the Virtual Network Function (VNF) and the global service scenario are key to enable these NFV provisioning layers (NFV-Orchestrator and VIM) to make these intelligent and optimal deployment decisions.

This Enhanced Platform Awareness (EPA) capability allows the NFV orchestration platform to intelligently deploy well-designed VNF workloads onto the appropriate underlying NFVI enabling the optimal SLAs. This also unleashes the favorable total cost of ownership (TCO) NFV promises due to this more efficient use of the underlying infrastructure. This must be achieved through the implementation of a more versatile NFV ready infrastructure, and a more agile and competitive ecosystem of network functions providers enabled through such an open information model.

Toward demonstrating this NFV deployment approach, Intel, Telefónica, Cyan, Brocade, and Red Hat have collaborated to implement and demonstrate a complete ETSI-NFV end to end service deployment solution stack.

Service Scenario Overview

VNF Routing Scenario Overview

- The scenario being deployed is a routed VNF forwarding graph using Brocade Vyatta vRouters as VNFs. A three node network forwarding topology achieves a 40 Gbps network throughput between the ingress and egress points at Routers A and C (see Figure 2).
- The exposure of the performance enablers (NFVI attributes) in the VNF Descriptor and importance of a good VNF design (Vyatta vRouter) are crucial toward enabling this service deployment.
- The End to End NFV Service delivery stack with the relevant NFV intelligence built in at each layer, through the information model, the VNF, the NFV Orchestrator, the VIM and finally the NFVI are all required for an optimal VNF service chain deployment.
- The use of Industry standard, open, and extensible information models such as TOSCA and suitable VNF formats are crucial toward enabling an open ecosystem of VNF vendors construct and deliver their services into this new end to end architecture.

The scenario also showcases the importance of a well-designed standard high volume industry server HW based NFVI, which provides the EPA services required for the Deployment of Telco Grade VNFs.

Partners And Contributed Components

The lab environment is located at Telefónica's Global CTO NFV lab in Madrid.

As per Figure 3, the infrastructural components are provided as follow:

Intel components include:

- Intel® Xeon® processor-based servers and Network Interface cards
  - – Intel® Xeon® processor E5-2680 v2 @ 2.80 GHz
  - – Intel® Open Networking Platform (ONP) Ingredients including DPDK R1.6[3]
  - – Intel® X520 10G Network Interface Cards

Brocade components include:

- Brocade Vyatta vRouter 5600 3.2 R2
- OpenFlow switch (Brocade NetIron MLXe)

The **Cyan** components include:

- NFV-Blue Planet Orchestrator release 15.02

The **Telefónica** components include:

- DPDK R1.6 based Traffic generator TIDGEN (Telefónica I+D Generator)
- Telefónica VIM openvim R0.9

The **Red Hat** components include:

- RHEL7.0 (with patches) and QEMU-KVM version 2.0.0 (with patches)

Cloud Versus NFV

As mentioned, the demonstration is hosted at Telefónica's NFV Reference Lab (physically located in Madrid) and provides two separate deployment environments (see Figure 4):

- A NFV-ready NFVI pool, with a Telefónica developed NFV ready VIM implementing the requisite Enhanced Platform Awareness (EPA) and a Cyan NFV-Orchestrator supporting advanced VNF deployment using enhanced NFV information models.
- A standard cloud infrastructure pool ala classic cloud computing, with the same Telefónica VIM connected to the same Cyan NFV-Orchestrator but in this case not using the enhanced information model as the basis for the deployment.

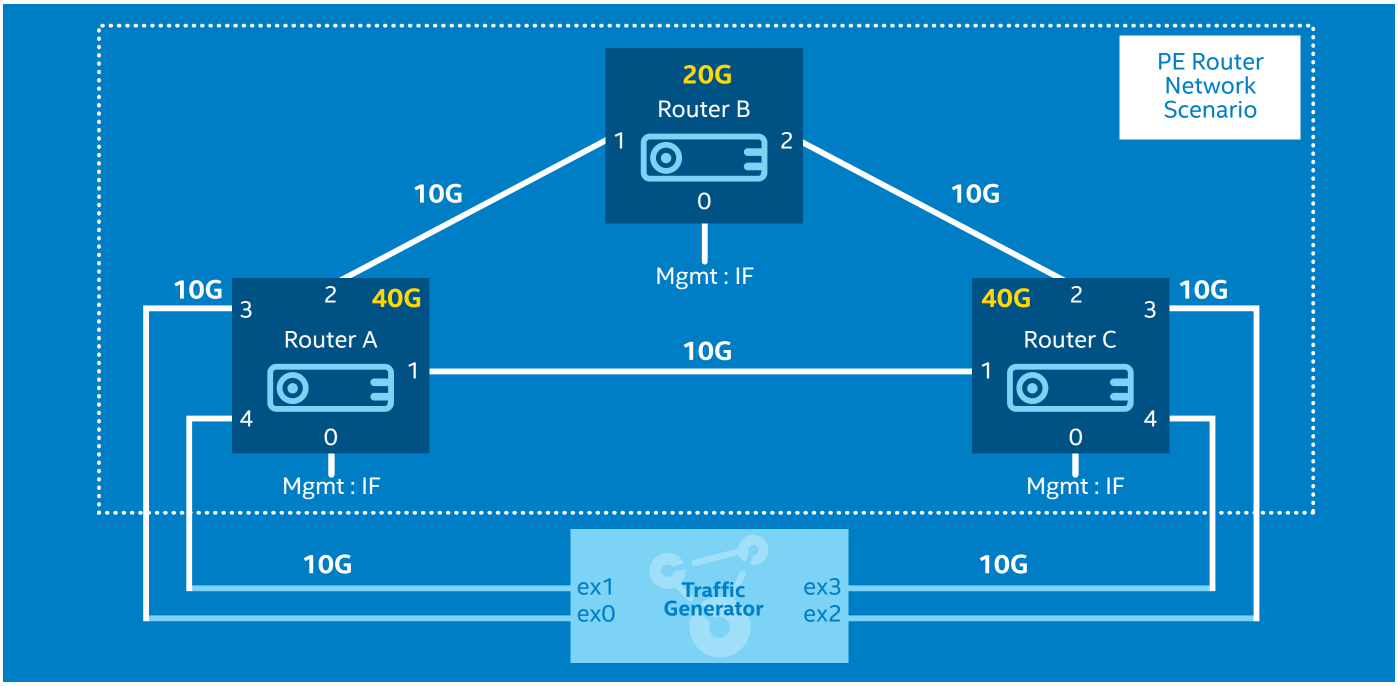

Figure 2. PE VNF Routing Service Chain.

Starting with both server pools empty (no VNFs deployed), the demo scenario is deployed onto each platform through the Orchestrator. With both setups running, performance measurements are displayed in real time, showing much higher and stable throughput in the NFV-ready scenario.

Information models for both scenarios are compared, showcasing the key additional attributes and end to end EPA awareness required for the optimized NFV deployment.

Scenario Execution Results

Initial Sub Optimal Cloud Deployment

The initial deployment demonstrates the typical issues with doing a "blind" enterprise-cloud like deployment of a typical VNF onto an underlying non NFV optimized infrastructure.

The Brocade routing scenario is deployed. Since a suboptimal NFV information model is used, the Brocade vRouter is incorrectly deployed through the non-aware MANO stack and is unable to fully achieve the 23 Mpps (40 Gbps @ 192 byte packet size) it is designed to achieve, instead reaches a mere 270 Kpps largely because of the following (see Figure 5):

- No PCIe* pass through: The NIC is not directly connected to vRouter, which now receives and transmits packets via the vSwitch, this is a suboptimal networking path to the VNF as compared to PCIe pass through mode and limits the throughput at acceptable packet loss.
- No NUMA affinity: vCPUs are arbitrarily allocated from CPU socket that may not be directly attached to the NICs and may also use a non-local memory bus.
- No CPU pinning: vCPUs allocated to vRouter may be shared or dynamically rescheduled limiting determinism.
- No 1G Huge Page setup. This greatly limits the performance achievable in DPDK (Vyatta) performance and doesn't correctly leverage the recent advances in server IOTBL and VTd architecture especially for small packet sizes.

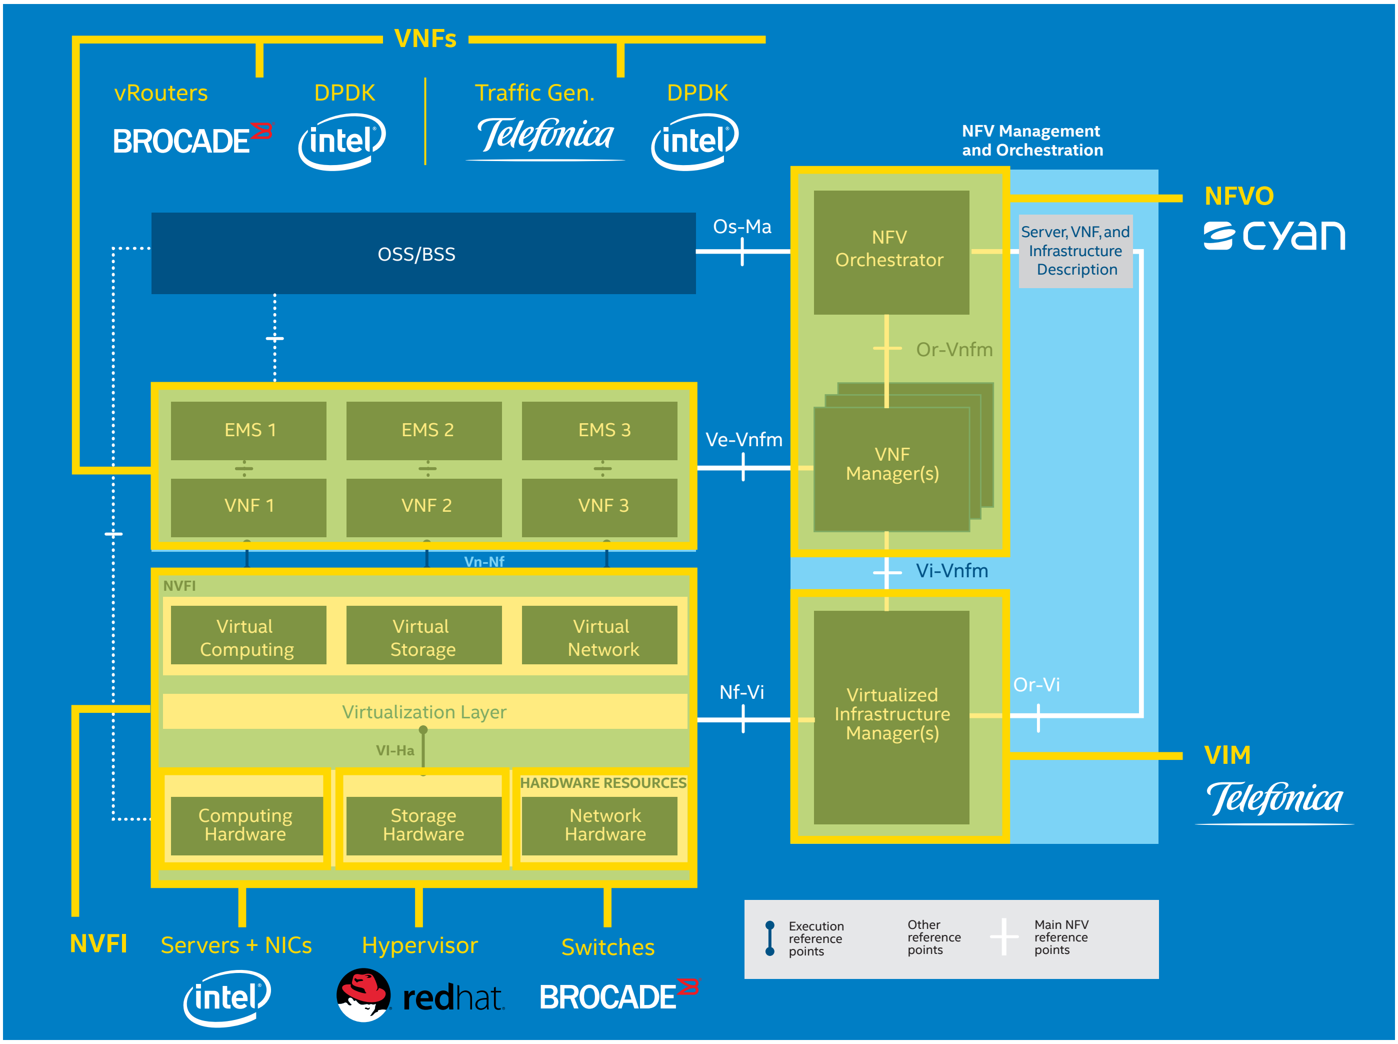

Figure 3. Partners and System Component Contributions.

Figure 5 highlights the sub optimal performance achieved.

Optimal NFV Deployment

The secondary deployment uses the correct NFV TOSCA and VNFD Models and the information in the model allows the Planet Blue orchestrator to optimally deploy the Brocade configuration through the Telefónica VIM and achieved full line rate performance of 23 Mpps (40 Gbps @ 192 Bytes). See Figure 5 for performance.

This deployment scenario demonstrates the benefits in doing an “intelligent” NFV deployment through the EPA aware delivery stack onto underlying NFVI using the correct extended information model containing the attributes required for deterministic VNF performance.

The Brocade vRouter is deployed with the correct EPA parameters correctly exposed via VNF Descriptor and enforced by the Cyan NFVO and the VIM. The Brocade vRouter is able to achieve the high performance as expected by design correctly implementing the PCIe pass through, the NUMA awareness, CPU pining, and huge page requirement as required by the Brocade VNF.

Figure 6 demonstrates similar line rate performance but for larger packet sizes.

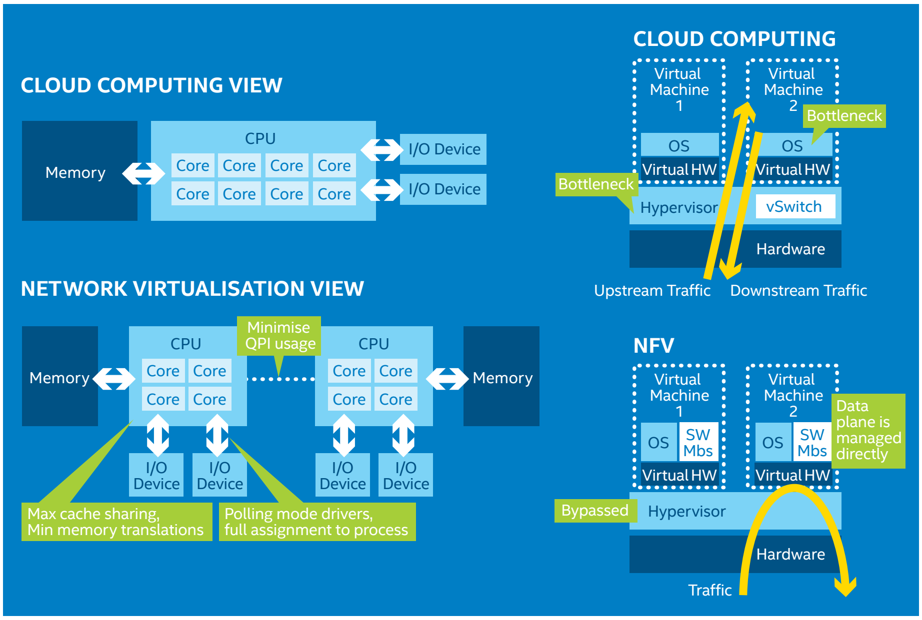

Figure 4. Cloud vs. NFV.

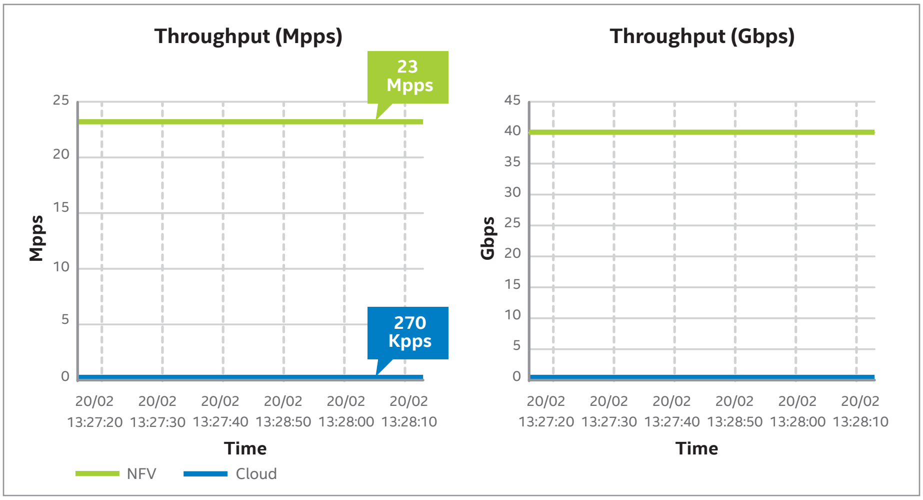

Figure 5. Performance Comparison for 192 byte frame size.

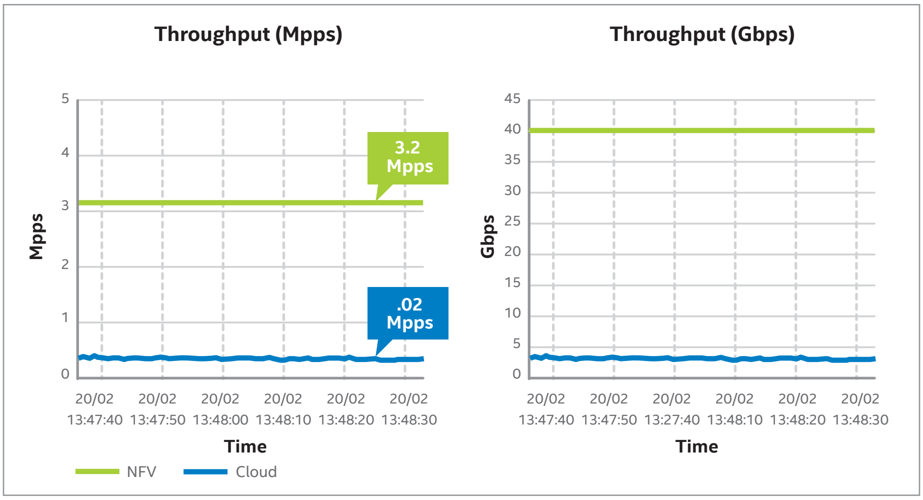

Figure 6. Performance Comparison for 1518 byte frame size.

Conclusions

- True End to End NFV aware system designs will deliver huge VNF performance improvements (23 Mpps v 270 Kpps) necessary for Telco grade performance.
- Properly developed end to end NFV solutions will reduce Network TCO and allow a new ecosystem of VNF providers flourish.
- The community must understand these enhanced performance attribute (EPA) capabilities and ensure proper exposure up through the End to End system. This requires taking and end to end system view toward implementing the appropriate levels of intelligence up through the NFV delivery stack to maximize the application performance and determinism required for Telco grade SLA deployments.
- The VNF community must understand these capabilities and build/model their VNFs accordingly.
- An intelligent EPA aware Orchestration and VIM are the key components toward releasing complete NFV TCO value.
- Intel, Red Hat, Cyan, and Telefónica will continue to work to enable Open Stack (VIM) with these critical NFV EPA enhancements[4].
- Standard and open information models are also crucial to enable the open VNF ecosystem and enable the transition from the world of monolithic, vertically integrated network appliances to SW defined network functions.
- 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 leveraging best of breed cloud open source technologies.

Testimonials

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

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

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

Testimonials (continued)

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[5] though its Open Networking Platform program and is working closely with its Netbuilders SDN-NFV ecosystem community[6] 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

Acronyms

BNG	Broadband Network Gateway
BSS	Business Support System
CMS	Cloud Management System
CPU	Central Processing Unit
vCPU	Virtual Central Processing Unit
DPDK	Dataplane 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

References

1. https://portal.etsi.org/NFV/NFV_White_Paper.pdf
2. 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
3. http://www.intel.com/content/dam/www/public/us/en/documents/white-papers/open-network-platform-server-paper.pdf
4. https://software.intel.com/sites/default/files/managed/72/a6/OpenStack_EPA.pdf
5. http://www.etsi.org/deliver/etsi_gs/NFV/001_099/002/01.02.01_60/gs_NFV002v010201p.pdf
6. https://networkbuilders.intel.com/docs/Intel_Network_Builders_Directory_Sept2014.pdf