WHITEPAPER

SANDVINE

# Enhancing the Subscriber Experience From the Handset to the Internet

## Network Function Virtualization at Line Rate: Bare Metal Performance with NFV and PacketLogic/V

### VNF PERFORMANCE OPTIMIZATION WITH OVS-DPDK: MARKET OVERVIEW

Network Function Virtualization (NFV) has been championed by network operators globally as a technology that will enable them to transform their networks into an agile infrastructure that is able to adapt to the changing landscape of broadband consumer usage today. In January 2013, the European Telecommunications Standards Institute (ETSI) launched an initiative sponsored by seven of the largest operators in the world: AT&T, BT, Deutsche Telekom, Orange, Telecom Italia, Telefonica and Verizon. To establish requirements and an architecture for the virtualization of network functions. Since that time it has grown to over 220 individual companies, including 37 of the world's major service providers as well as representatives from both telecoms and IT vendors. The NFV member companies see tremendous potential in NFV for telecommunications deployments. The stated goals from the NFV ISG are focused around business benefits, but also include some challenges that need to be overcome.

ETSI NFV GOALS, BENEFITS, AND CHALLENGES

| ETSI NFV ISG Goal | Acheivement Method | Benefit | Challenge |
|---|---|---|---|
| Reduce CAPEX | Leverage COTS Hardware | Cost savings on custom hardware | Maintaining performance on COTS hardware |
| Reduce OPEX | Reduce Installation and configuration times and support costs | Cost savings on hardware support, multiple services active on hardware platforms | Orchestration and Configuration in multivendor networks |
| Reduce TIme-to-Market for new services | Software-only installation and VNF instantiation dramatcially faster than hardware installation | New service launches in days rather than months | Orchestration and Configuration in multivendor networks |
| Improved Network ROI | Less hardware required for new services | More profitable services with lower cost to subscriber | Maintaining performance on COTS hardware, Orchestration and Configuration in multivendor networks |
| Greater network flexibility | Software defined usage of available resources | Use available CPU resources for in-demand services, more efficiency | Orchestration and Configuration in multivendor networks |
| Greater use of software-only solutions | COTS hardware running VNFs | Less "custom hardware" lock-in", Multiple services active on hardware platforms | Maintaining performance on COTS hardware |
| Trial and Deploy innovative services at lower risk | VNF instatiation for service on preinstalled hardware | Lower cost for trial, lower revenue required for service ROI | Orchestration and Configuration in multivendor networks |

Transform your networks into an agile infrastructure that is able to adapt to the changing landscape of broadband consumer usage today.

## THE NFV PERFORMANCE DILEMMA

As shown above, performance is one of the major challenges facing NFV implementations. The performance challenge has proven to be significant in the migration from custom hardware to Commercial Off the Shelf (COTS) hardware for some vendors. NFV has been touted for it's ability to reduce costs, but in order to achieve this, the performance gap between COTS and custom hardware must narrow from what was initally experienced by vendors implementing software-only implementations of their Virtual Network Functions (VNF).

The NFV performance dilemma can be broken into two distinct camps: Control Plane and Data Plane. Control Plane performance can be CPU intensive due to the transactional nature of control plane functions, i.e. millions of small transactions coming to servers on the network. Most control plane applications are implemented on COTS technology today, and the challenge to virtualize those applications was straightforward. Data plane implentations have traditionally relied on specific hardware configurations or hardware-assist for achieving performance (whether through ASICs, FPGAs, network processors, or specialized processors) in order meeting network operator's growing performance needs. The goal of NFV is to remove hardware dependencies, so data plane VNF implementations must optimize around the lowest common denominator for a standard COTS hardware configuration. Data plane applications also require predictable performance despite the CPU and I/O intensive nature of their implementations, and the shift to NFV does not remove that requirement. Network operators require predictable performance from data plane systems so that they can dimension their networks to ensure a high quality of experience in delivering content and services to their subscribers, especially during peak usage times.

In the NFV architecture, the concept is for multiple VNFs to run on a COTS platform, as shown in the **Figure 1**.

Figure 1

HIGH LEVEL NFV ARCHITECTURE

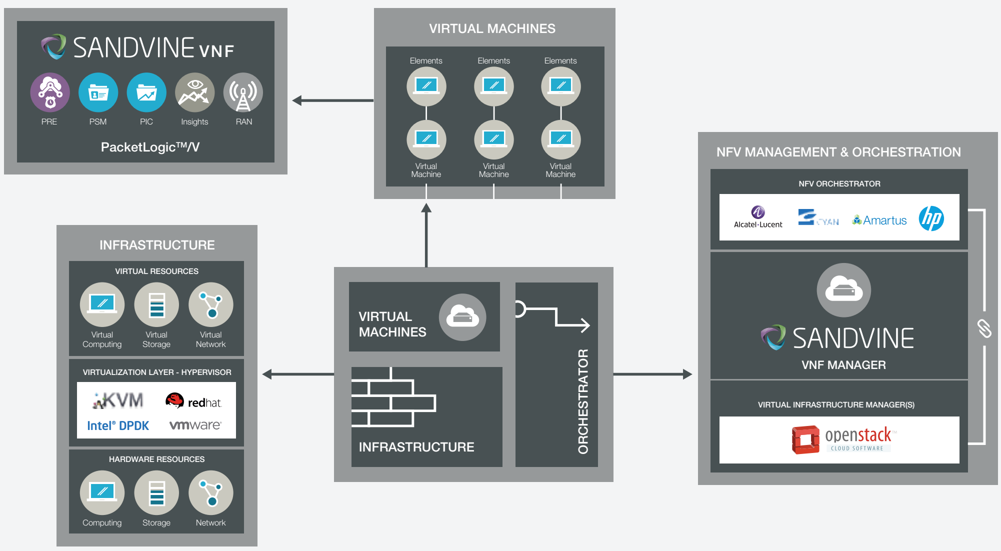

Overcoming the performance challenges requires a close collaboration of all of the above components to achieve the cost savings that are desired by the network operators.

The diagram above shows several connection points that are potential bottlenecks in meeting any performance goals for NFV. First, the COTS hardware should offer the maximum amount of performance available. The ability of the Virtual Switch to offer line rate performance to the Virtual Machines and the COTS hardware is a critical factor, and the Virtual Machine should consume minimal resources in order not to take performance away from the VNF software. Although the concepts are simple, achieving maximum performance from NFV solutions has proven to be more difficult than expected for many VNF suppliers. Most network operators are expecting anywhere from a 10-25% performance penalty when running NFV-based solutions versus hardware-based solutions, which often eliminates any potential cost savings in hardware (although still leaving other benefits to the operator). Overcoming the performance challenges requires a close collaboration of all of the above components to achieve the cost savings that are desired by the network operators.

Meeting the performance requirements of network operators and VNF suppliers requires an extremely stable reference platform that can easily be replicated at the software and virtualization layer, regardless of hardware platform. Network operators need to ensure that the COTS hardware systems that they purchase can run the broadest range of VNF solutions with high performance, and VNF providers need to optimize their software around a common set of interfaces and platforms that they can rely on for their solutions. Although there are multiple virtualization platforms available on the market, the critical consideration for data plane applications is the data plane acceleration tools available in the virtualization platforms. Commercial options are offered from various companies that can offer higher performance to many applications as add-ons to the base virtualization platforms for an additional cost. There are several non-commercial (i.e. bundled) options that are also being offered to solve this problem. Intel has opened up their Data Plane Development Kit (DPDK) and the Open Virtual Switch software to enable performance on open systems.

## INTEL'S ACCELERATION STRATEGY AND SOLUTIONS

Intel's enablement work for SDN and NFV within this emerging model of network infrastructure is embodied in Intel Open Network Platform (ONP). The over-arching goal of Intel ONP is to reduce the cost and effort required for service providers, data-center operators, TEMs, and OEMs to adopt and deploy SDN and NFV architectures. The foundations of this effort are the Intel ONP Server Reference Design, introduced in the second half of 2014, and the Intel ONP Switch Reference Design, introduced in 2013.

Intel delivers a number of platform-level features and capabilities enhance the performance, reliability, and security of the Intel ONP for Servers reference design, which is architected on Intel® Xeon® or Intel® Core™ processor-based systems. The technologies described in this section are open source and optimized to work on open standards and are therefore applicable to SDN and NFV implementations outside the scope of the reference design.

## INTEL® DATA PLANE DEVELOPMENT KIT (INTEL® DPDK) USE IN OPEN VSWITCH

The set of software libraries comprising Intel DPDK can be used to dramatically accelerate packet processing by software-based network components, for greater throughput and scalability. Engineering teams from Intel and Wind River collaborated to replace the dataplane switching logic of the open source Open vSwitch project with a new version built on top of Intel DPDK to improve the small packet throughput. The resulting project, which is called the Intel DPDK Accelerated Open vSwitch, offers dramatic improvements to packetswitching throughput and is offered as a reference implementation for use with Intel ONP for Servers and other NFV solutions. The mainstream Open vSwitch project has also started to adopt DPDK connectivity as an optional configuration.

Figure 2

INTEL ONP ARCHITECTURE

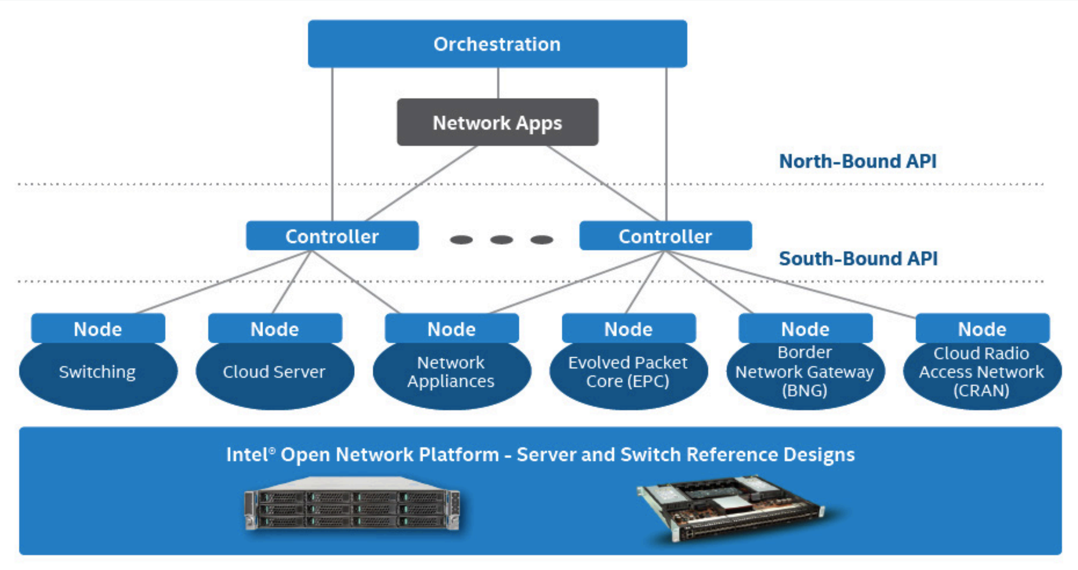


The source code for the base Intel DPDK library as well as the Open vSwitch enhancements that take advantage of DPDK libraries are readily available for development organizations that wish to extend the public projects or to pursue custom development on their own.

## INTEL® QUICKASSIST TECHNOLOGY

Hardware-based acceleration services for workloads such as encryption and compression are supported by Intel QuickAssist Technology, using an open-standards approach that is well suited to use with Intel ONS for Servers and other SDN and NFV implementations. An accelerator abstraction layer provides a uniform means of communication between applications and accelerators, as well as facilitating management of acceleration resources within the control and orchestration layers of the Intel ONP architecture.

## INTEL® VIRTUALIZATION TECHNOLOGY (INTEL® VT)

Capabilities at the virtualization layer itself also play a key role in the robustness of SDN and NFV implemented on Intel architecture. Enablement by Intel for all of the major virtualization environments—including contributions to open-source projects and co-engineering with providers of proprietary offerings—provides robust support for Intel VT. The hardware assists for virtualization offered by Intel VT dramatically reduce overhead, by eliminating the need for software-based emulation of the hardware environment for each VM. As a result, Intel VT enables higher throughput and reduced latency. It also enhances data isolation between virtual machines (VMs), for greater security.

These technologies form a strong reference platform that VNF providers can develop on, and Sandvine is leveraging this architecture fully.

## SANDVINE'S NFV STRATEGY AND SOLUTIONS

Sandvine's strategy for NFV is to deliver software-based VNF solutions that meet the performance expectations of network operators by matching the performance capabilities of our hardware-based solutions. Sandvine has always been a software company, and our customers recognize that we deliver more performance per CPU cycle than anyone else in our market. We have historically utilized COTS hardware in specific configurations to run our high performance software solutions, and this leveraged our knowledge of how to tightly integrate with hardware to get maximum performance out of that hardware configuration.

The current PacketLogic solutions run on a mixture of COTS appliances and servers, and the main software components will remain the same, as shown in **Figure 3.**

Figure 3

PACKETLOGIC/V ARCHITECTURE

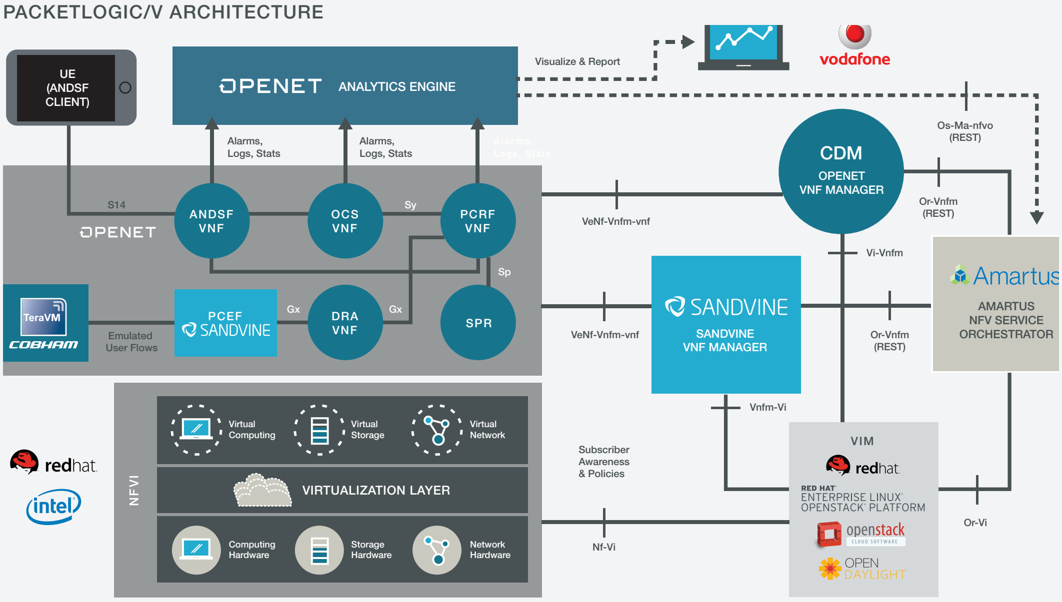


There is one major new component introduced into PacketLogic with PacketLogic/V, and that is the VNF-Manager solution shown in the diagram, and that is designed to solve the second major issue that we identified in the overview, but outside the scope of this performance whitepaper. Outside of the VNF Manager, the goal is to deliver carrier-scale performance across all of our VNFs. The strategy for PacketLogic/V was to leverage the Intel toolkit above to deliver our solution as a pure software VNF offering, and match the performance that we achieve in our hardware-based offerings. Sandvine has worked to maximize our performance on SandyBridge and IvyBridge on our existing Intel-based hardware solutions with Intel in the past, and when we began our investigation into achieving high performance on NFV systems, we once again opened that channel.

## 2014 PERFORMANCE: PACKETLOGIC/V

We initially began testing on a Virtual CPE solution with HP and Intel and were able to get ~700Mbps per instance on a stock HP system running VMware, which maxed out the I/O capacity of the VMware virtual interfaces at the time, as our CPU utilization was only 12%. This was a good start for the low end of the market (sub-1Gbps), but not attractive for most broadband operators that have multiple 10Gbps links to the network. We began working with Intel and optimizing our interfaces to the NFV reference architecture, and eventually achieved the performance numbers in the graph in October of 2014:

As **Figure 4** shows, with Intel® Xeon® Processor E5-2600 v3 Haswell CPUs we are achieving the maximum performance possible with the I/O that can be placed in a standard COTS server (4x10GE).

Figure 4

2014 PACKETLOGIC/V PERFORMANCE

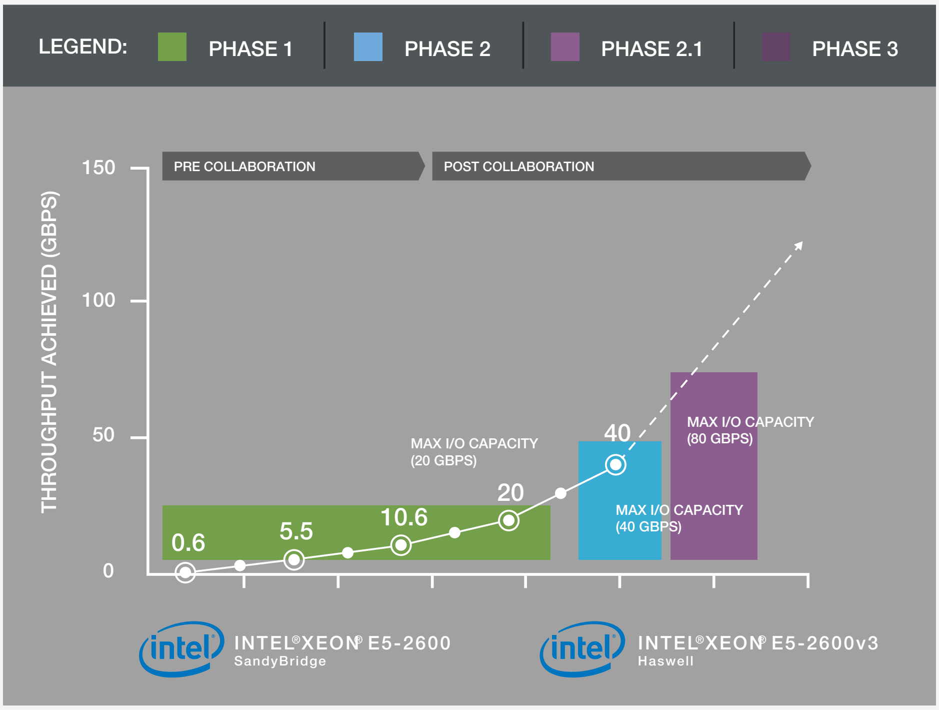


This is also roughly equivalent to the performance achieved in the COTS appliances that Sandvine delivers to our network operator customers today, which delivers on the strategy of allowing the customer to choose their own platform with no performance penalty for NFV or hardware-based solutions.

In this type of deployment, the limiting factor will be I/O ports, and not the CPU in the server. There are several server vendors that are tacking the I/O density challenge, and if that solution is solved, it will enable NFV to solve the dense 10GE network deployments that are commonplace in the core network today. It also highlights the need for 100GE interfaces for servers, which the performance numbers that are being achieved on Haswell show can be supported with a VNF solution.

As NFV and the data plane server market matures, NFV has unlimited potential to meet even the most demanding network performance required by network operators today. Using blade server solutions, scalability into the hundreds of Gbps is easily achievable, and as 100GE matures in the data center, network operators will be able to take advantage of those interfaces in their deployments.

## 2015 PERFORMANCE: PACKETLOGIC/V

In 2015, Sandvine continued to push the performance limits with the Intel® Xeon® Processor E5-2600 v3 Haswell CPUs by moving from 4x10GE ports to 4x40GE ports for the I/O on the server. In this scenario, the system is now able to achieve up to 160Gbps of I/O performance rather than the 40Gbps that was tested in 2014. This is the same CPU configuration that is supported in the PacketLogic PL9420 system, and provides a good guideline for comparing a purpose-packaged appliance to an NFV equivalent COTS server. In this configuration, the performance graph is updated to the results below:

As shown in **Figure 5**, the same system used in the initial 40Gbps test is now delivering over 150Gbps with the enhanced port density, raising the performance to parity with appliance configurations. This delivers the PacketLogic software as a true software solution, which has no performance or feature penalties for virtualization.

Figure 5

2015 PACKETLOGIC/V PERFORMANCE

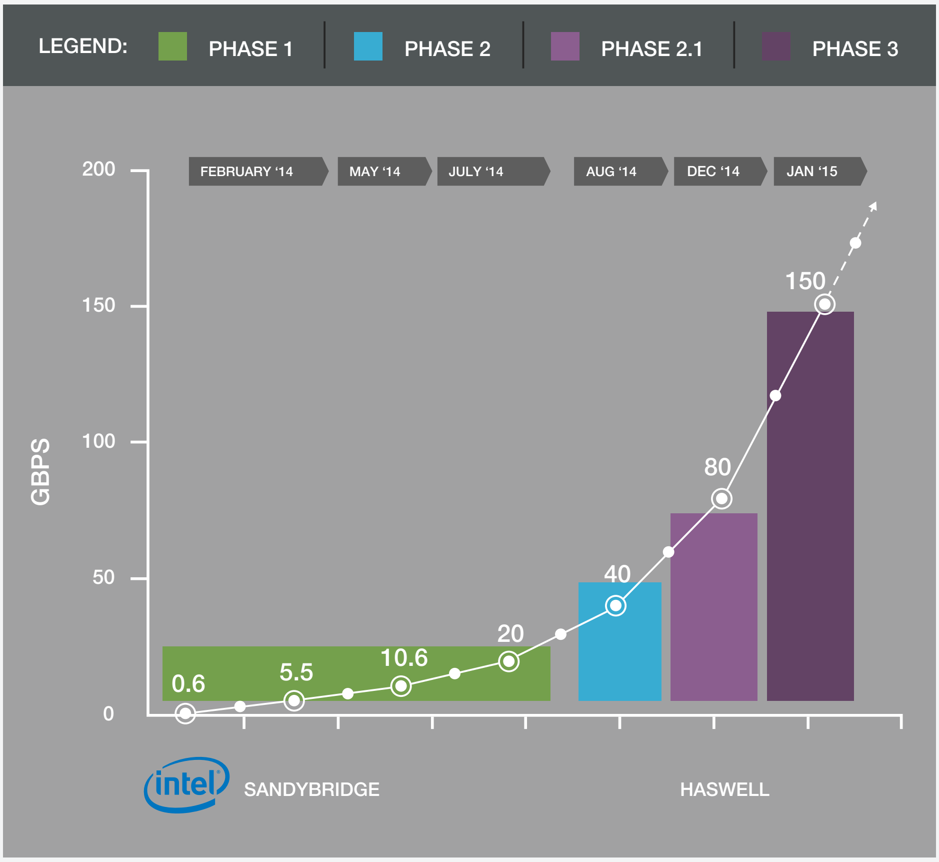

To test the performance of the PacketLogic/V as shown on the previous pages, the following environment was used:

| Component | Details |
|---|---|
| **Hardware** | Intel Haswell Server with 2x E5-2699 v3 @ 2.30GHz (18 cores each)<br>64GB 2133MHz DDR4 RAM<br>4x Intel 82599 10Gbps NICs<br>1x 1TB SATA hard drive |
| **Software** | Debian 7<br>DPDK Version 1.7.0<br>OVS/DPDK version 1.1<br>QEMU 1.6.2 + OVS/DPDK patches |
| **Guest VNF Software** | PacketLogic PRE version 15.0.5<br>VirtIO Driver model, 2 channels |
| **Traffic Generation** | Ixia Breaking Point |
| **Test Setup** | 4 x 10Gbps ports<br>iMix traffic<br>10,000 flows |

v20171221

## ABOUT SANDVINE

Sandvine helps organizations run world-class networks with Active Network Intelligence, leveraging machine learning analytics and closed-loop automation to identify and adapt to network behavior in real-time. With Sandvine, organizations have the power of a highly automated platform from a single vendor that delivers a deep understanding of their network data to drive faster, better decisions. For more information, visit **sandvine.com** or follow Sandvine on Twitter at **@Sandvine.**

SANDVINE

**USA**
47448 Fremont Blvd,
Fremont,
CA 94538,
USA
T. +1 510.230.2777

**EUROPE**
Birger Svenssons Väg
28D
432 40 Varberg,
Sweden
T. +46 (0)340.48 38 00

**CANADA**
408 Albert Street,
Waterloo,
Ontario N2L 3V3,
Canada
T. +1 (0)519.880.2600

**ASIA**
Ardash Palm Retreat,
Bellandur, Bangalore,
Karnataka 560103,
India
T. +91 80677.43333