IMPLEMENTATION SUMMARY

Gi-LAN Solution Implementation Summary

Intel Corporation
Datacenter Network
Solutions Group

Authors

Eduardo Castro
Solution Software Engineer,
Intel Corporation

Tarek Radi
Lead Technical Program Manager,
Intel Corporation

Table of Contents

1.0 Introduction

As Communications Service Providers (Comms SPs) move to a Software-Defined Networking (SDN) and Network Functions Virtualization (NFV) world, they find themselves needing to enable several use cases. This solution implementation document focuses on the Gi-LAN use case and presents exemplary single host virtualized Gi-LAN architecture. It can serve mobile customer traffic using both local content as well as Internet accessibility while enabling key traffic roaming policies.

The solution is implemented with several virtual network functions (VNFs) from various third-party suppliers on top of a Red Hat Enterprise Linux*-based network function virtualization infrastructure (NFVI), running Red Hat OpenStack Platform* 7. The VNFs used in this solution are given in Table 1.

Table 1. Virtual functions of the solution.

REQUESTED VIRTUAL FUNCTION	SUPPLIER	PRODUCT/VERSION
Switching	Open vSwitch Community	Open vSwitch* 2.3.2
Routing	Brocade	5600 vRouter* build: 3.2.1R6
DNS		
PCRF and PCEF/TDF	Sandvine	Policy Traffic Switch Virtual Series* 7.00.01 Service Delivery Engine Virtual Series* 7.10 Subscriber Policy Broker Virtual Series* 6.50
Management UI for PCRF and PCEF/TDF		Control Center* 6.90.02
Firewall	F5 Networks	BIG-IP* Advanced Firewall Manager 11.6.0
DDoS		
Carrier-Grade NAT		BIG-IP Carrier-Grade NAT 11.6.0
Transparent Proxy and Caching	The Apache Software Foundation	Apache Traffic Server* 6.0.0
Content Delivery Network		
Origin Web Server		Apache HTTP Server* 2.4.10
Outbound Anti-Spam	Cisco Systems	Snort* 2.9.6.0

The intent of this document is to help customers who are interested in implementing this specific use case in an SDN/NFV world. Intel does not aim to promote or recommend any specific hardware, software, or supplier mentioned in this document. In addition, Intel does not aim to tie customers to any specific software and hardware stack.

The primary audiences for this document are architects and engineers planning to implement their own virtualized Gi-LAN architectures. Intel neither aims to promote or recommend any specific hardware, software, or supplier nor tie customers to any specific hardware and software stack mentioned in this document.

2.0 Solution Overview

This section provides a high-level overview on the solution's architecture and deployment. The scope of the work was to deploy and integrate sets of features that would help realize the use cases for this solution as agreed upon with the customer. The uses cases include:

- **Internet web surfing**. In this scenario, the traffic passes through the entire path of the VNFs from the router to the carrier-grade NAT (CGNAT). Assuming the test input machine is trying to access a webpage from the Internet, the request/response should go through the path of VNFs. The Brocade input router routes the user's traffic to the firewall through network policy control (Sandvine VNFs). The F5 BIG-IP* Advanced Firewall Manager applies its default rule by accepting the user's traffic. The domain name system (DNS) function of the Brocade 5600 vRouter* does the address name translation of the website. The F5 BIG-IP CGNAT transforms the private IP address of the user (test input) to a public one.
- **DNS caching** reduces the numbers of DNS queries into the Internet. Basically, instead of sending the query to the Internet (for example, Google DNS, 8.8.8.8) for the IP of a specific website name, the local DNS server will be queried (in this case Brocade VNF).
- **Security** is enforced with a variety of VNFs. The **firewall** scenario demonstrates blocking a website from the Internet. With the use of the F5 BIG-IP AFM before the DNS server, a request can be blocked before saturating the DNS server and the CGNAT VNF. The firewall allows for enabling/disabling rules from the dashboard and viewing the statistics.

  Distributed denial of service (DDoS) feature is provided with the same F5 BIG-IP AFM VNF and will block any traffic resembling an attack (for example, sending a large amount of packages to an external service). The statistics are in the DDoS dashboard of F5 BIG-IP AFM.

  In the **Anti-Spam** scenario, specific IP addresses and spam messages are identified and blocked using an open source version of Snort* VNF, located just after the firewall.
- **Applying traffic analysis**. This use case utilizes Sandvine VNFs for analyzing traffic for a specific IP using network policy control. Sandvine's policy and charging enforcement and traffic detection functions (PCEF/TDF) and policy and charging rules function (PCRF) are used to throttle traffic speeds. Sandvine VNFs are located just after the input router.
- **CGNAT**. The F5 BIG-IP CGNAT VNF transforms the private IPs of users (in this case the test input) to public ones. From the CGNAT dashboard, it is possible to browse statistics of the NATted requests.
- **Caching (websites)**. Transparent Proxy VNF is located just after the local DNS server and acts as an intermediary for the most recent websites visited by the users from the input network.
- **Content delivery network (CDN)**. This function enables you to deliver content located on various origin servers to end users with high availability and high performance using various mechanisms like switching or load balancing. These origin servers and the CDN controller are located just before the firewall.

Figure 1 depicts how the VNFs are connected to the various networks in the solution. On the diagram, networks are represented as colored circles and slices, while VNF connections are shown as colored lines with end points. The traffic, starting from an Input Network (via the Brocade input router) crosses the path of connected VNFs (blue line) and reaches the Output Network (behind the CGNAT). Tables 2 and 3 provide additional information on the solution's hardware bill of materials.

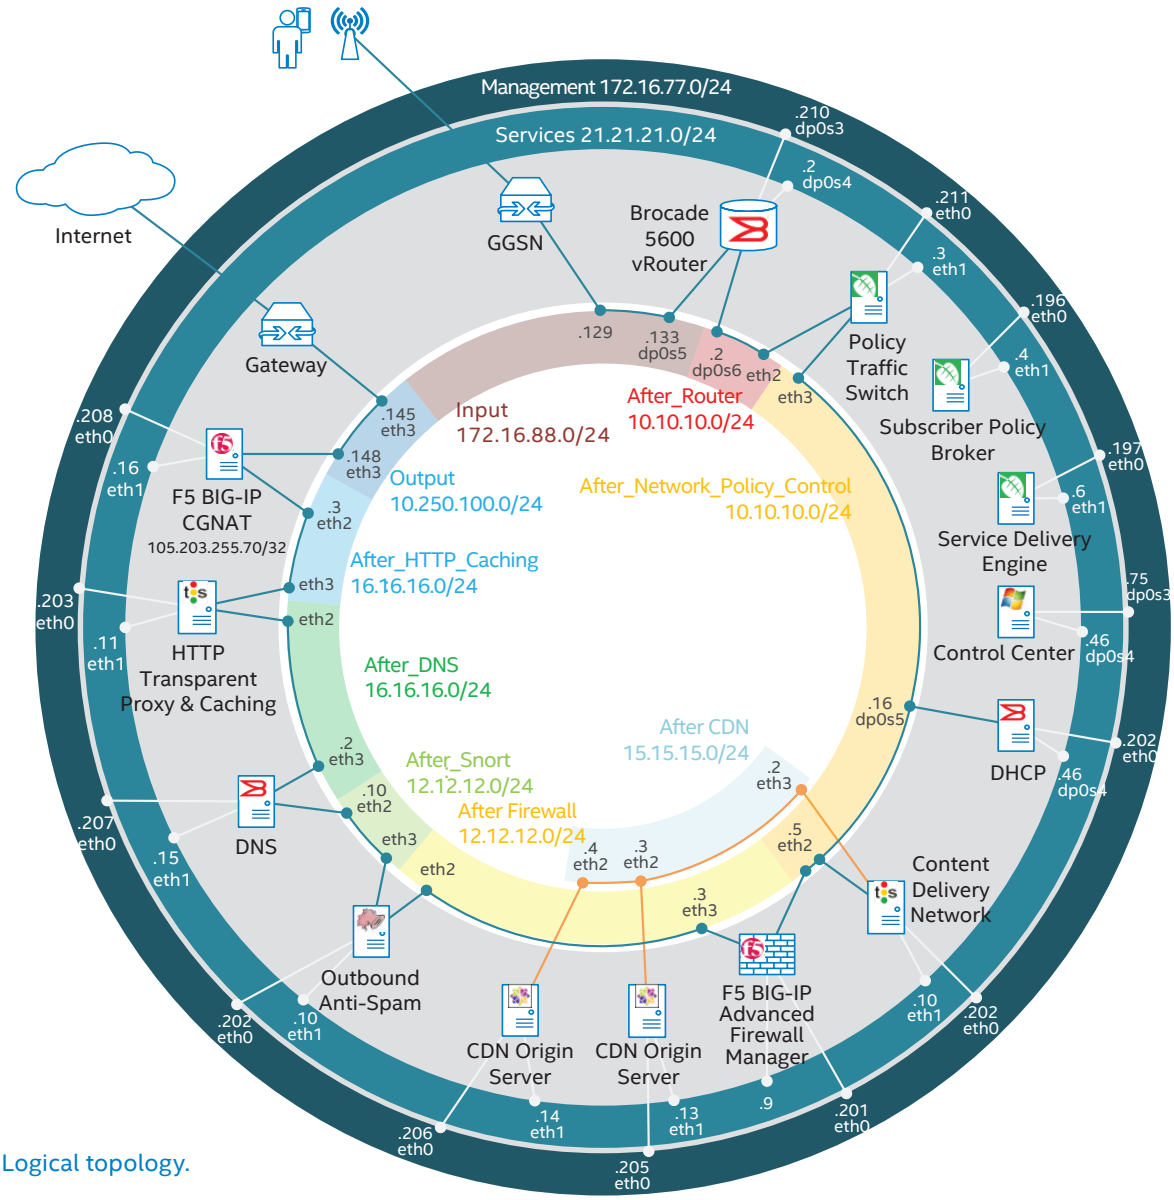

Figure 1. Logical topology.

Table 2. Hardware bill of materials—setup used in integration.

HARDWARE ITEM	VERSION/COMPONENTS	QUANTITY
Dell PowerEdge* R730	• 2× Intel® Xeon® processor E5-2680 v2 2.8 GHz, 25 MB cache, total 20 cores • Memory: 1866 MHz; total 64 GB DDR3 • Hard Disk: 3× 2 TB 7.200 RPM NL-SAS 6 Gbps 3.5" hot-plug, 13G, PERC H730 RAID Controller 1 GB NV cache • NIC: Broadcom* 5720 QP 1 Gb Network Daughter Card • ReadyRails* sliding rails with cable management arm • iDRAC8 Enterprise • DVD+/-RW • SATA internal dual hot plug • Redundant power supply (1+1), 750W	1 (acting both as a compute and a controller)

Table 3. Hardware bill of materials—setup used in performance tests.

HARDWARE ITEM	VERSION/COMPONENTS	QUANTITY
Dell PowerEdge* R730	• 2× Intel® Xeon® processor E5-2690 v3 2.6 GHz, 30 MB cache, total 24 cores • Memory: 2133MHz; total 128GB DDR4 • Hard Disk: 3× 2 TB 7.200 RPM NL-SAS 6 Gbps 3.5" hot-plug, 13G, PERC H730 RAID Controller 1 GB NV cache • NIC: Intel® Ethernet Converged Network Adapter X710 DA2 with Intel® Ethernet SFP+ SR Optics • ReadyRails* sliding rails with cable management arm • iDRAC8 Enterprise • DVD+/-RW • SATA internal dual hot plug • Redundant power supply (1+1), 750W	1 (acting both as a compute and a controller)

OpenStack* is used to spawn and orchestrate the VNFs. The qemu-kvm* is used as virtualization technology run on top of Red Hat Enterprise Linux* 7.1.

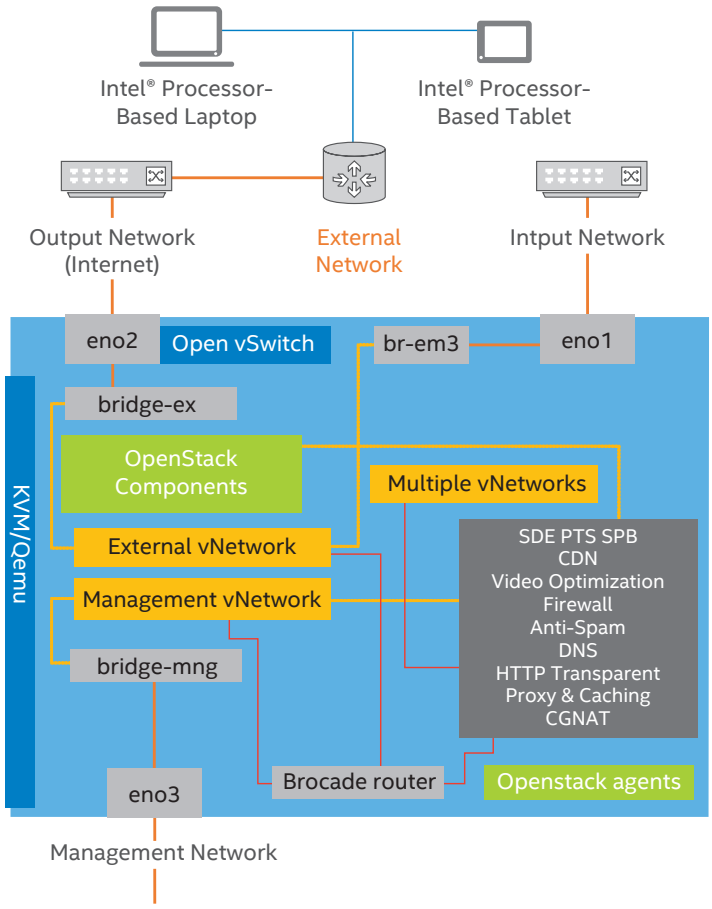

Figure 2. Physical topology diagram

The details on the software stack and the software components versions are presented in Table 4. For the purposes of installing various software components, and for testing or inspection needs, some additional software tools were used. For details, please refer to Table 5.

Table 4. Software stack.

FUNCTION	COMPONENT	VERSION
Operating System	Red Hat Enterprise Linux*	7.1
Hypervisor	qemu-kvm*	2.4.0
Switching	Open vSwitch* plug-in/agent for OpenStack*	2.3.2
Orchestrator	OpenStack	Red Hat OpenStack Platform 7
Routing	Brocade 5600 vRouter*	Build: 3.2.1R6
PCRF and PCEF/TDF	Sandvine PTS* Virtual Series	7.00.01
	Sandvine SDE* Virtual Series	7.10
	Sandvine SPB* Virtual Series	6.50
Management UI for PCRF and PCEF/TDF	Sandvine Control Center*	OS: Microsoft Windows* 7 SP1 6.90.02
CDN Origin Servers	Apache HTTP Server*	OS: CentOS* 7 without GUI Apache HTTP Server 2.4.10
CDN	Apache Traffic Server*	OS: CentOS 7 with GUI Apache Traffic Server* 6.0.0
Firewall	F5 BIG-IP* Advanced Firewall Manager	11.6.0
Outbound Anti-Spam	Snort* image	OS: CentOS 7 without GUI Snort* 2.9.6.0
DNS	Brocade 5600 vRouter	Build: 3.2.1R6
DDoS	F5 BIG-IP Advanced Firewall Manager	11.6.0
Transparent Proxy and Caching	Apache Traffic Server	OS: CentOS 7 with GUI Apache Traffic Server 6.0.0
CGNAT	F5 BIG-IP Carrier-Grade NAT	11.6.0 Kernel 2.6.32-358.23.2.el6.f5.x86_64 Licenses: CGN/AFM/PEM/LTM/ASM

Table 5. Software and tools.

FUNCTION	COMPONENT
OpenStack* Installer	Packstack*
Sandvine VNF installer	OpenStack Heat Script
Traffic Inspection	Wireshark*
Bandwidth management	iPerf* netperf*

The details of the solution, including installation and configuration can be found in the Gi-LAN Solution Implementation Installation Guide.

Link to guide: **https://networkbuilders.intel.com/network-technologies/solution-blueprints**

3.0 Test Results

This section presents example performance test scenarios defined to benchmark the setup, and the associated latency and throughput results. The setup was benchmarked with iPerf*/iPerf3* tool. For the purpose of benchmarking, several additional VM were created.

3.1 Scenarios

- **Scenario 1: Performance at the input router.** The test input is located in the Subscriber (Input) network, and the results are collected at the egress of the input router.
- Sc**enario 2: Performance at the firewall.** The test input is located in the Subscriber (Input) network, and the results are collected at the egress of the firewall.
- **Scenario 3: Performance after the firewall.** For this test, CGNAT VM was skipped. The test input is located after the firewall, and the results are collected at the extra router (output router) created in place of CGNAT.
- **Scenario 4: Performance of the network chain.** This test measures the performance of the entire network chain—from Subscriber (Input) network to the performance router (Output network).

3.2 Results

Table 6 presents the throughput and latency results measured for each scenario. Note that the latency was measured with the ping tool.

Table 6. Test Results.

	AVERAGE THROUGHPUT [GBPS]	AVERAGE AVERAGE LATENCY [MS] LATENCY [MS]
Performance at the input router	Packstack*	1.03
Performance at the firewall	OpenStack Heat Script	1.93
Performance after the firewall	Wireshark*	2.22
Performance of the network chain	iPerf* netperf*	7.43

The throughput performance measured for the duration of the tests (60 seconds) for each of the scenarios presented is shown in Figure 3.

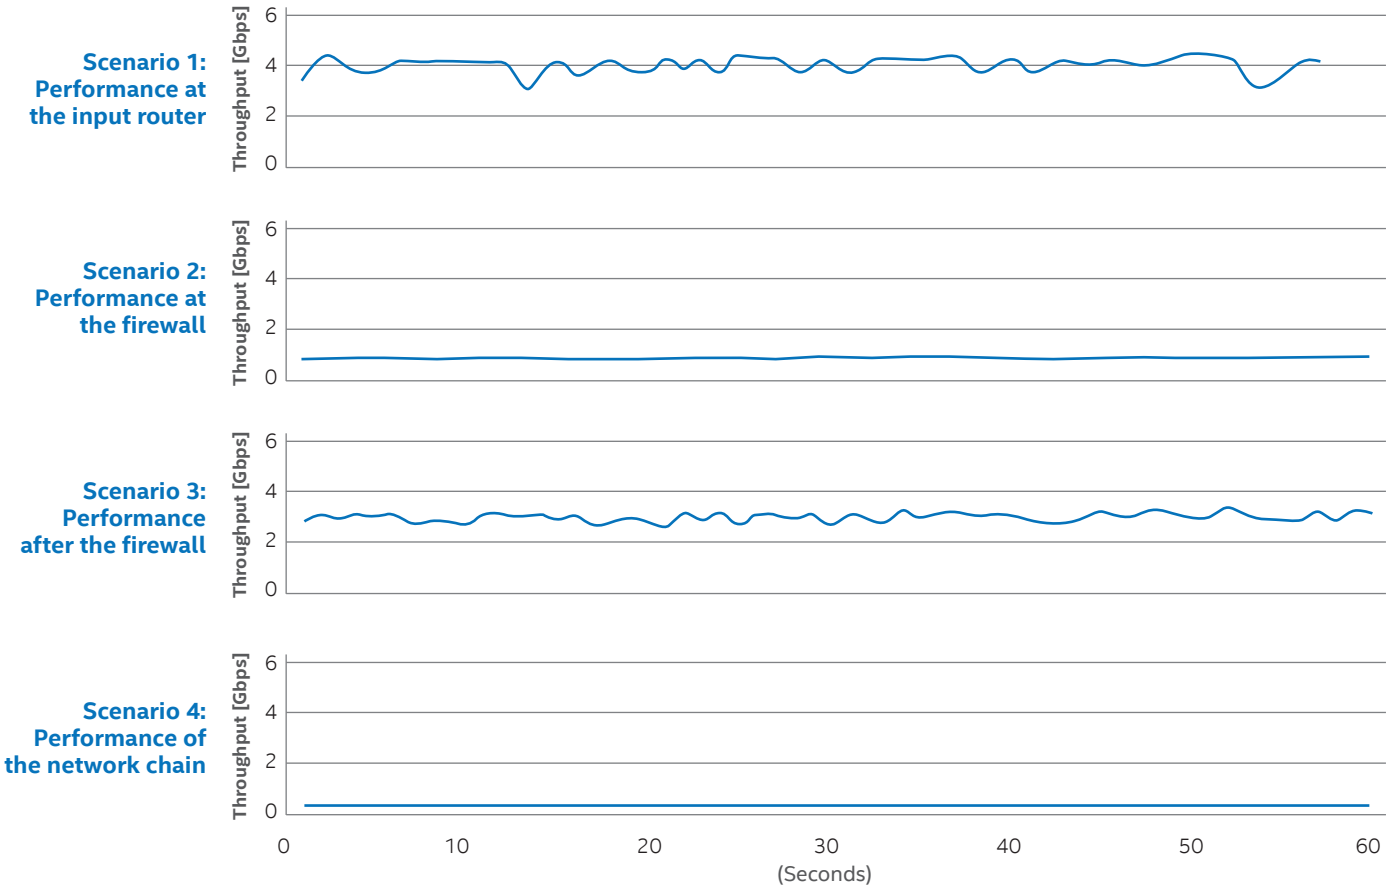

Figure 3. Throughput performance measured for each scenario over 60 seconds.

Figure 4 and Figure 5 present the general packet processing performance on the data plane interfaces of the input router and the output router respectively.

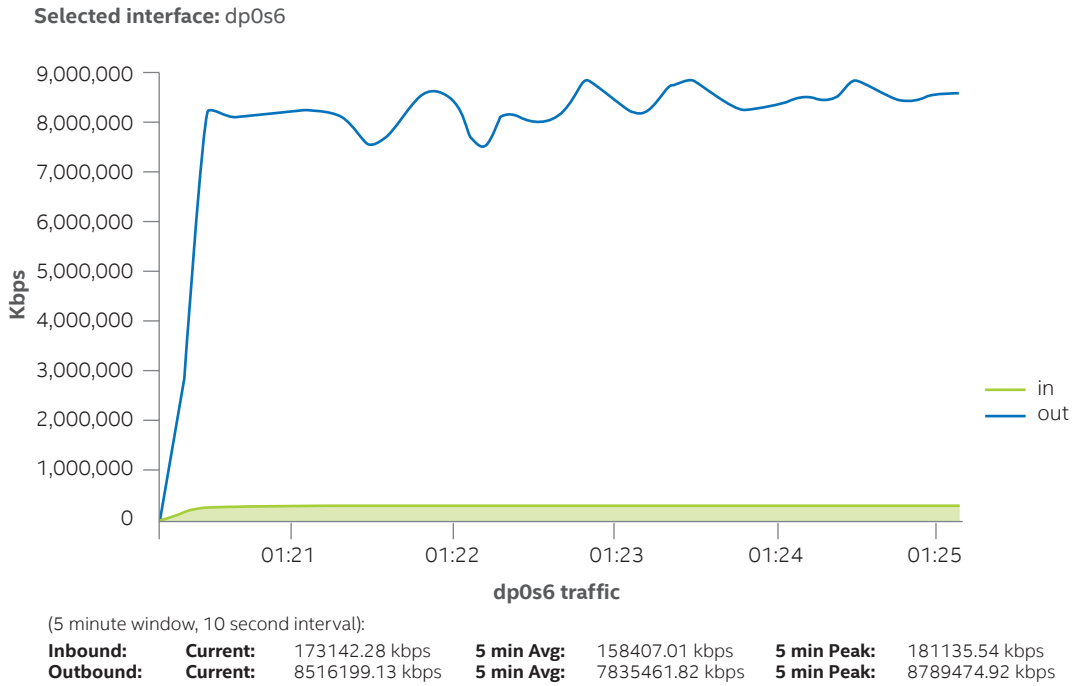

Figure 4. General packet processing efficiency of the input router.

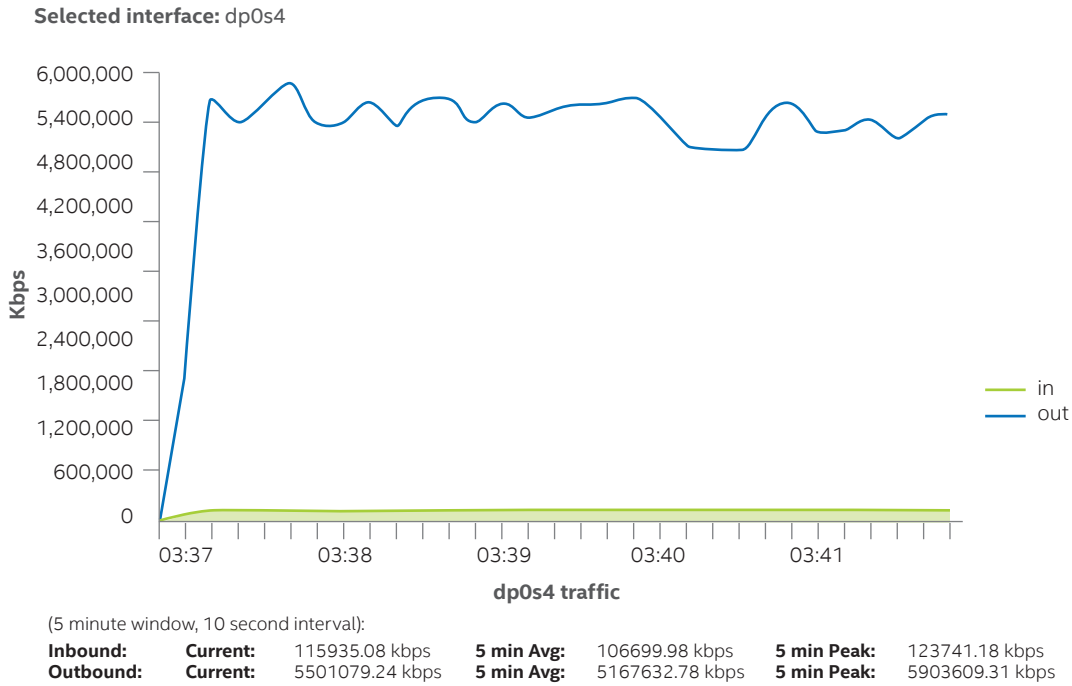

Figure 5. General packet processing efficiency of the output router.

4.0 Next Steps

- To learn more about the technologies mentioned in this paper, please follow the links in the document.
- To learn more about Intel's technology for NFV, attend the courses available in the Intel® Network Builders University at https://networkbuilders.intel.com/university.
- To learn more about Intel® Network Builders partners for NFV products, visit https://networkbuilders.intel.com/solutionscatalog.
- To build a test bed using the Intel® Open Network Platform Reference Architecture, download the documentation at https://01.org/packet-processing/intel%C2%AE-onp.
- To get the highest performance from your NFV systems, specify compatibility with the Data Plane Development Kit in your infrastructure and VNF procurements.
- To get the highest return on investment from your NFV systems, specify use of Enhanced Platform Awareness in your orchestration, infrastructure, and VNF procurements.

Appendix A: Abbreviations

ABBREVIATION	DESCRIPTION
AFM	Advanced Firewall Manager
CDN	Content Delivery Network
CGNAT	Carrier-Grade Network Address Translation
CPU	Central Processing Unit
DDoS	Distributed Denial of Service
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
GGSN	Gateway GPRS Support Node
Gi-LAN	Gateway-Internet LAN
GPRS	General Packet Radio Service
GUI	Graphical User Interface
HTTP	Hypertext Transfer Protocol
iDRAC	Integrated Dell Remote Access Controller
LAN	Local Area Network

ABBREVIATION	DESCRIPTION
NFV	Network Functions Virtualization
NL-SAS	Near-Line Serial Attached SCSI
OS	Operating System
PCEF	Policy and Charging Enforcement Function
PCRF	Policy and Charging Rules Function
PTS	Policy Traffic Switch
RAID	Redundant Array of Independent Disks
RPM	Revolutions per Minute
SDE	Service Delivery Engine
SPB	Subscriber Policy Broker
TDF	Traffic Detection Function
UI	User Interface
VM	Virtual Machine
VNF	Virtualized Network Functions

Appendix B: References

REFERENCE	SOURCE
Brocade 5600 vRouter Data sheet	https://www.brocade.com/content/dam/common/documents/content-types/datasheet/brocade-5600-vrouter-ds.pdf
Evaluating Dynamic Service Function Chaining for the Gi-LAN White Paper	http://www.intel.com/content/dam/www/public/us/en/documents/white-papers/evaluating-dynamic-service-function-chaining-for-the-gilan-paper.pdf
F5 BIG-IP Carrier-Grade NAT Data Sheet	http://www.f5.com/pdf/products/big-ip-cgnat-datasheet.pdf
F5 BIG-IP Advanced Firewall Manager Data Sheet	http://www.f5.com/pdf/products/big-ip-advanced-firewall-manager-datasheet.pdf
Open vSwitch	http://openvswitch.org/
Sandvine Policy Traffic Switch Virtual Series	https://www.sandvine.com/platform/policy-traffic-switch/pts-virtual-series.html
Sandvine Service Delivery Engine Virtual Series	https://www.sandvine.com/platform/service-delivery-engine.html
Sandvine Subscriber Policy Broker Virtual Series	https://www.sandvine.com/downloads/general/platform/subscriber-policy-broker/sandvine-subscriber-policy-broker.pdf
Snort	https://www.snort.org/ https://www.snort.org/downloads/archive/snort/snort-2.9.6.1.tar.gz
Apache HTTP Server	https://httpd.apache.org/ https://archive.apache.org/dist/httpd/
Apache Traffic Server	http://trafficserver.apache.org/ https://www.snort.org/downloads/archive/snort/snort-2.9.6.1.tar.gz

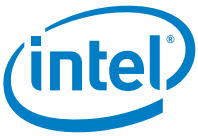

 1116/MH/MESH/PDF 335201-001US