WHITE PAPER

Intel Corporation

Affirmed Networks* virtualized Evolved Packet Core (vEPC) benchmarks with 2nd Generation Intel® Xeon® Scalable Processors

Authors

Hassan El Hariri
Affirmed Networks

David Bastiani
Affirmed Networks

Sarita Maini
Intel Corporation

Shivapriya Hiremath
Intel Corporation

Veera Gurijala
Intel Corporation

Key Contributors

John M Morgan
Intel Corporation

Phil To
Affirmed Networks

1 Introduction

A virtualized Evolved Packet Core (vEPC) enables Communication Service Providers to reduce cost by moving core network components of the Long-Term Evolution (LTE) network from high-cost dedicated hardware to low-cost commercial off-the-shelf (COTS) servers. This document illustrates the high performance of the Affirmed Networks virtualized Evolved Packet Core (vEPC) solution on 2nd Generation Intel® Xeon® Platinum 8280L processors (formerly codenamed Cascade Lake). The test results achieved a **throughput of up to 168 Gbps** and a **packet rate of up to 29 Mpps** on dual-socket servers powered by these Intel® Xeon® Platinum 8280L processors, using Intel® XXV710 25G NICs.

The test environment utilized Red Hat* OpenStack* Platform (RH* OSP*) version10 (Newton) and the Spirent Landslide* physical traffic generator, as a standardized framework for testing and characterizing the performance of the Affirmed vEPC virtual network function (VNF) in an NFVI environment.

Intel® Xeon® Scalable processors incorporate unique features for virtualized network workloads, leading to impressive performance gains compared to systems based on prior Intel processor generations. The increased performance of Intel Xeon Scalable processors can significantly improve the capability for software-centric, carrier-grade virtualization which aids communications service providers in attaining and enforcing service level agreements and increasingly demanding quality of service requirements.

Affirmed Networks* delivers a fully virtualized mobile core solutions supporting both vEPC and Next Generation 5G Core. Affirmed Network's solution enables operators to economically scale networks and deliver differentiated services tailored to specific use cases covering:

- Consumer
- Enhanced Mobile Broadband (eMBB)
- Internet of Things (IoT)
- Mobile Network Operator (MNO)
- Mobile Virtual Network Operator (MVNO)
- Mobile Virtual Network Enabler (MVNE)
- Private Long-Term Evolution (LTE)
- WiFi*
- GiLAN
- Voice over LTE (VoLTE) to Fixed Wireless

Affirmed Networks' vEPC provides high performance, scalable, cost-effective vEPC functions, including:

- Mobility Management Entity/Serving GPRS Support Node (MME/SGSN)
- Serving Gateway (SGW)
- Gateway GPRS Support Node/Packet Data Network Gateway (GGSN/PGW)
- Evolved Packet Data Gateway (ePDG)
- Trusted Wireless Access Gateway (TWAG)
- CIoT Serving Gateway Node (CSGN) – Narrow-Band IoT (NB-IoT)
- Policy and Charging Rules Function (PCRF)
- Authentication, Authorization, and Accounting (AAA)
- Home Subscriber Server (HSS)

Affirmed Networks' solution capabilities include:

- 5G Core
- Control & User Plane Separation (CUPS)
- Network Slicing
- Integrated virtual probes
- Optimized IoT access (NB-IoT/LTE-M/SCEF)
- Virtualized Deep Packet Inspection (DPI)
- Service automation (Affirmed Service Automation Platform)
- WiFi
- GiLAN and analytics services

Table of Contents

1 Introduction **1**
1.1 Intended Audience 3
1.2 Terminology 3
1.3 Reference Document 5

2 Document Overview **5**
2.1 Red Hat* Enterprise Linux* and Red Hat* OpenStack* Platform v10 (Newton*) 5
2.2 Affirmed Networks* 5G vEPC 5
2.3 Spirent Landslide* 6

3 Hardware and Software Components **6**
3.1 Intel® Xeon® Scalable Processors 6
3.2 Hardware Components 7
3.3 Software Components 7

4 OpenStack Nodes and Network Topology **8**
4.1 External/Provisioning Network 9
4.2 Intelligent Platform Management Interface (IPMI) 9
4.3 Management Network 9
4.4 Data Network 9

5 Affirmed Networks* vEPC Deployment on OpenStack* **10**

6 Performance Benchmarks **10**

7 Summary **11**

Appendix A BIOS Settings **11**

Figures

Figure 1. Spirent Physical Landslide GUI 6
Figure 2. OpenStack Components 8
Figure 3. Network Topology Details 9

Tables

Table 1. Terminology 3
Table 2. Reference Document 5
Table 3. Hardware Components 7
Table 4. Software Components 7
Table 5. Intel® Xeon® Platinum 8280L – Core Allocation for vEPC VMs 10
Table 6. BIOS Settings 11

1.1 Intended Audience

This white paper is intended for communication service providers who are planning and deploying virtualized mobile core infrastructure running on the latest Intel® Xeon® Scalable Processors.

1.2 Terminology

Table 1. Terminology

ABBREVIATION	DESCRIPTION
AAA	Authentication, Authorization, and Accounting
BIOS	Basic Input / Output System
BMC	Baseboard Management Controller
CIDR	Classless Inter-Domain Routing
CIoT	Consumer Internet of Things
CLI	Command Line Interface
CoSP	Communications Service Provider
CSGN	CIoT Serving Gateway Node
CSM	Content Services Module
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name System
DPI	Deep Packet Inspection
ePDG	Evolved Packet Data Gateway
GGSN	Gateway GPRS Support Node
GPRS	General Packet Radio Services
GUI	Graphic User Interface
GW	Gateway
HSS	Home Subscriber Server
IoT	Internet of Things
IP	Internet Protocol
IPMI	Intelligent Platform Management Interface
LTE	Long-Term Evolution
MAC	Media Access Control
MCC*	Mobile Content Cloud*
MCM	Management Control Module
MME	Mobility Management Entity
MNO	Mobile Network Operator
MVNE	Mobile Virtual Network Enabler
MVNO	Mobile Virtual Network Operator

ABBREVIATION	DESCRIPTION
NB-IoT	Narrow-Band IoT
NFV	Network Functions Virtualization
NFVI	Network Functions Virtualization Infrastructure
NG Core	Next-Generation Core
NIC	Network Interface Card
NR	New Radio
NUMA	Non-Uniform Memory Access
OS	Operating System
OSP	Open Stack Platform
OVS	Open vSwitch
PCI	Peripheral Component Interconnect
PCRF	Policy and Charging Rules Function
PGW	Packet Data Network Gateway
PXE	Preboot eXecution Environment
RHEL	Red Hat Enterprise Linux
SAEGW	System Architecture Evolution Gateway
SCEF	Service Capability Exposure Functions
SGSN	Serving GPRS Support Node
SGW	Serving Gateway
SR-IOV	Single-Root Input / Output Virtualization
SSM	Subscriber Services Module
TWAG	Trusted Wireless Access Gateway
UE	User Equipment
UUID	Universally Unique Identifier
vEPC	Virtual Evolved Packet Core
VLAN	Virtual Local Area Network
VM	Virtual Machine
VMM	Virtual Machine Manager
VNF	Virtual Network Function
VoLTE	Voice over LTE
VxLAN	Virtual extensible LAN

1.3 Reference Document

Table 2. Reference Document

REFERENCE	SOURCE
Affirmed Mobile Core* Tests Demonstrate Scalability, Performance for 5G Networks	https://www.affirmednetworks.com/wp-content/uploads/2018/06/Affirmed_Intel_vEPC_performance_report.pdf

2 Document Overview

This document showcases the benchmarks of Affirmed Networks vEPC using Red Hat OpenStack on 2nd Generation Intel® Xeon® Scalable processors (codename Cascade Lake - SP). It demonstrates an enhanced IO throughput of up to 168 Gbps and a packet rate of up to 29 Mpps with 28-core Intel® Xeon® Platinum 8280L @ 2.7GHz, in the control plane as well as user plane servers, using Intel® XXV710 25G NICs.

The test environment consisted of four servers powered with Intel Xeon Platinum 8280L processors, with the Red Hat* Enterprise Linux* (RHEL*) operating system v7.6 and the Red Hat* OpenStack* Platform 10 (Newton*) installed on three bare-metal servers, and a virtual machine (VM) on the fourth server. These servers were configured as an OpenStack* Director VM, an OpenStack Controller Node host, and two OpenStack Compute Node host servers. A physical Spirent Landslide* was used as the traffic generator, and for emulating LTE network components needed for testing vEPC.

The tested Affirmed Networks vEPC solution implemented a virtualized Serving Gateway (SGW) and Packet Data Network Gateway (PGW). The serving gateway is a critical network function for the LTE mobile core network. The SGW acts as a mobility anchor point and routes data between the eNodeB base station and the PDN gateway. It receives instruction from the Mobility Management Entity (MME) to setup/teardown sessions for UEs (User Equipment). One or more SGWs serve a group of eNodeB's for user plane data. The SGW handles user IP packets between the PGW and eNodeB. The PGW provides external access to Packet Data Networks (PDNs). If the UE has multiple data sessions it can be connected to multiple PDNs. The PGW is responsible for allocating IP-address for the UE as well as the quality-of-service (QoS) and bandwidth parameters for the subscriber session, based on the carrier policy. A combined SGW & PGW is referred to as a System Architecture Evolution Gateway (SAE-GW). Figure 1 shows all the components that are a part of this test setup.

2.1 Red Hat* Enterprise Linux* and Red Hat* OpenStack* Platform v10 (Newton*)

Red Hat* Enterprise Linux* (RHEL*) is a commercial distribution of the Linux* operating system (OS). It provides support for platforms to run varied workloads in physical, virtualized, and cloud environments.

Red Hat* OpenStack* Platform (RH OSP*) is a commercial distribution of the open source project of OpenStack* to deploy easily manageable and configurable pools of compute, storage, and network resources in public and private clouds. RH OSP was installed on RHEL.

Some important OpenStack services that are available within RH OSP are:

- Nova service – compute
- Neutron – network
- Swift – object storage
- Cinder – block storage
- Ironic – bare metal provisioning
- Heat – orchestration
- Keystone – identity management.

RH OSP uses the OSP director to install, upgrade, and manage the cloud. It uses the concept of Overcloud and Undercloud based on TripleO*. The Undercloud is the main director node from where the user can provision and control the Overcloud nodes. The Overcloud is the cluster of nodes performing the roles of controller, compute, and storage that a user creates through the OSP director's Undercloud.

2.2 Affirmed Networks* 5G vEPC

Affirmed Networks* virtualized 5G Evolved Packet Core (5G vEPC) provides the software running on multiple virtual machines (VMs) that are spawned on commercial off-the-shelf high-volume servers based on Intel® Architecture.

In this document, the main services tested in the vEPC are the Serving Gateway (SGW) and Packet Data Network Gateway (PGW), which are both based on the 3rd Generation Partnership Project (3GPP) standards.

The SGW is the point of interconnect between the radio-side and the vEPC. This gateway serves the User Equipment (UE) by routing incoming and outgoing packets. The SGW is logically connected to the PGW. The PGW is the point of interconnect between the vEPC and the external IP networks. The PGW routes the packets to and from the packet data network. The SAEGW's user-plane and control-plane VMs were hosted on distinct servers.

For more information, refer to Affirmed Mobile Core* Tests Demonstrate Scalability, Performance for 5G Networks, a white paper test report of an Affirmed Networks® vEPC solution that shows the line rate performance with linear performance scaling in a given test case on servers powered by the previous generation of Intel® Xeon® Scalable processors.

2.3 Spirent Landslide*

Spirent Landslide* emulates the control and data traffic of mobile subscribers moving through the network while using carrier services. Spirent Landslide's unique testing methods combine sets of node emulators and end-to-end test cases to offer service providers and equipment vendors a fully controlled test environment to validate system scalability and identify capacity limit. It can be used for measuring control plane capacity and stressing data plane performance, as well as for characterizing system performance and identifying performance ceilings. A graphical interface can be used to set up and execute test sessions, monitor progress, and test results.

Spirent Landslide was used to perform end-to-end network validation of vEPC. It emulated UEs, eNodeB, MME and the Network Host which are the nodes in the LTE network as shown in Figure 1 in blue. The nodes in green were a part of Affirmed Networks vEPC implementation of PGW and SGW, running on OpenStack compute nodes with Intel® Xeon® processors.

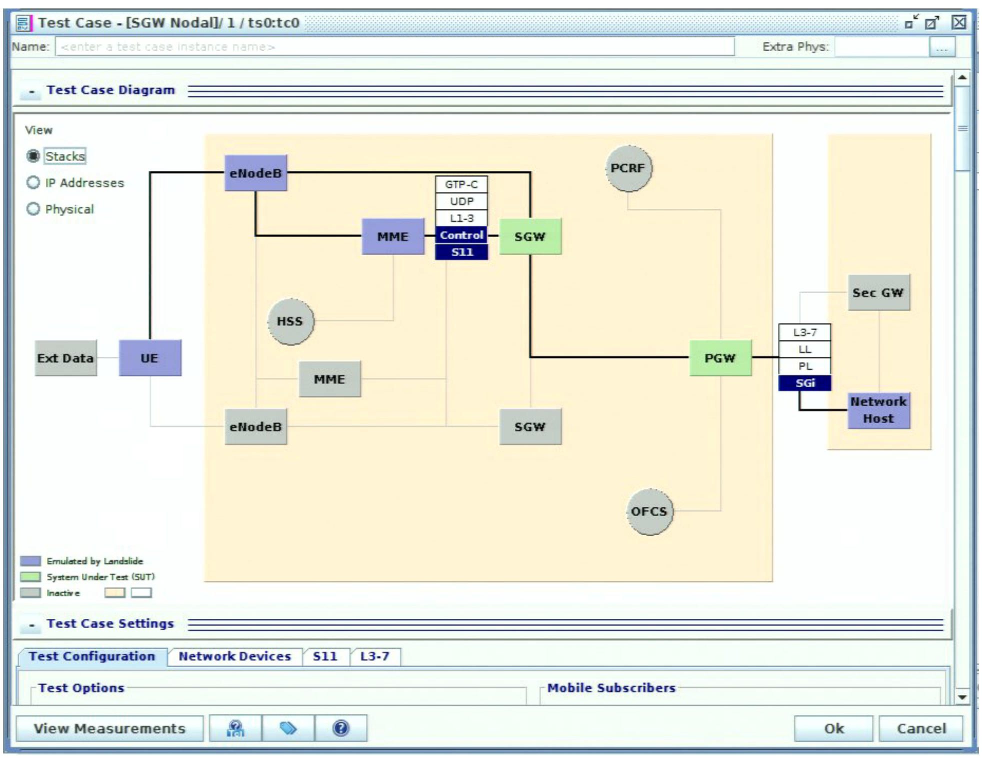

Figure 1. Spirent Physical Landslide GUI

3 Hardware and Software Components

This section details the hardware and software components used for this project.

3.1 Intel® Xeon® Scalable Processors

This document shows the high performance of the Intel® Xeon® Platinum 8280L processors. The Intel® Xeon® Scalable Processor family has architectural enhancements in the processor and CPU cores, higher core counts, larger L2 caches, larger capacity memory, increased numbers of PCIe* lanes, higher memory bandwidth, higher I/O bandwidth, and higher inter-socket bandwidth, compared to previous Intel® Xeon® processors. Hardware and Software components are listed in Table 3 and Table 4.

3.2 Hardware Components

Table 3. Hardware Components

CPU	INTEL® XEON® PLATINUM 8280L PROCESSOR (FORMERLY CODENAMED CASCADE LAKE-SP)
Number of Physical Cores per CPU	28
Number of Logical Processor Cores per CPU	56
Number of CPUs	2
Total Number of Logical Cores with Intel® Hyper-Threading Technology (Intel® HT Technology)	112
Processor Base Frequency	2.7 GHz
Memory Type/ Maximum Memory Speed/Memory Size	RDIMM, DDR4 @ 2933 MHz 384 GB
Number of Memory Channels	6
Number of PCI Express Lanes	48 lanes per CPU
Ethernet Switch	Arista 7280QR C-36
Data Network NICs	8x 25GbE ports in 4 Intel® Ethernet Controller XXV710 for NICs (2 ports per NIC) in each server, NUMA aligned with PCIe IO devices in the same NUMA nodes as the cores using the IO devices
OS Drive	Intel® SSD DC S4600 Series (960GB, 2.5in SATA 6Gb/s, 3D1, TLC) - SSDSC2KG960G7L

3.3 Software Components

Table 4. Software Components

ITEM	SOFTWARE VERSION
Host OS	Red Hat Enterprise Linux v7.6
OpenStack	Red Hat OpenStack Platform v10 (Newton)
vEPC VNF	Affirmed Networks MCC Rel. 9.1.1.0-145
Traffic Generator	Spirent physical Landslide 17.4.0

4 OpenStack Nodes and Network Topology

The Red Hat* OpenStack* Platform setup has four Intel® Xeon® based servers. They are used as OpenStack* nodes; one OpenStack Platform Director Host system, one OpenStack Controller node, and two OpenStack Compute nodes.

The Spirent Landslide physical traffic generator was used to send traffic to the vEPC VNF on the Compute nodes in an OpenStack environment connected on the same data network via the Intel® Ethernet Network adapter XXV710 25 GbE NICs. These are low power, standard volume NICs.

Figure 2 and Figure 3 present the Red Hat OSP platform and network setup.

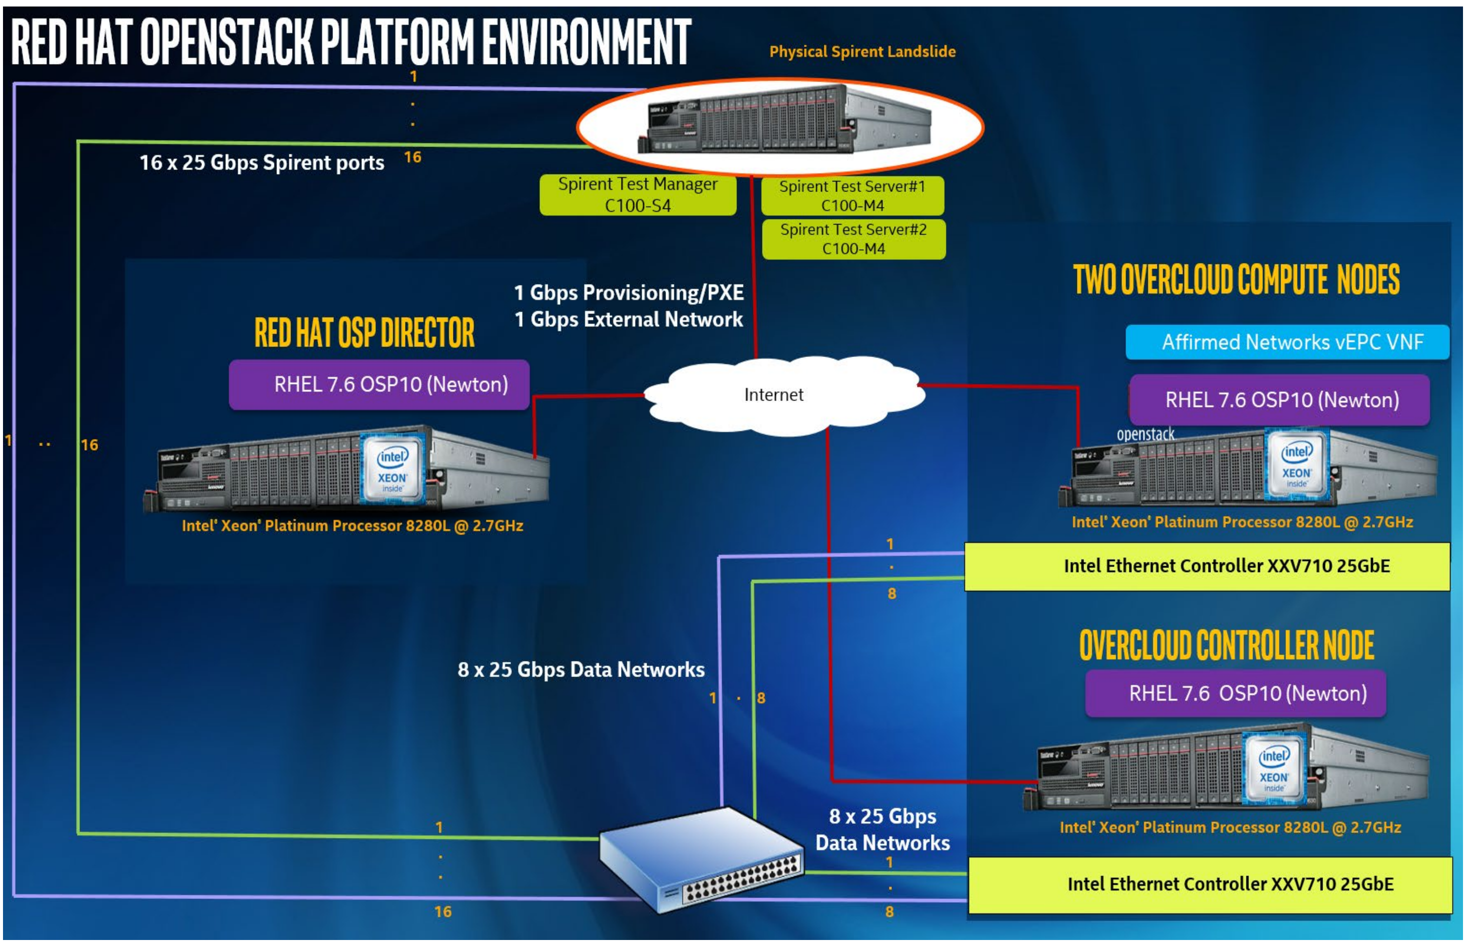

Figure 2. OpenStack Components

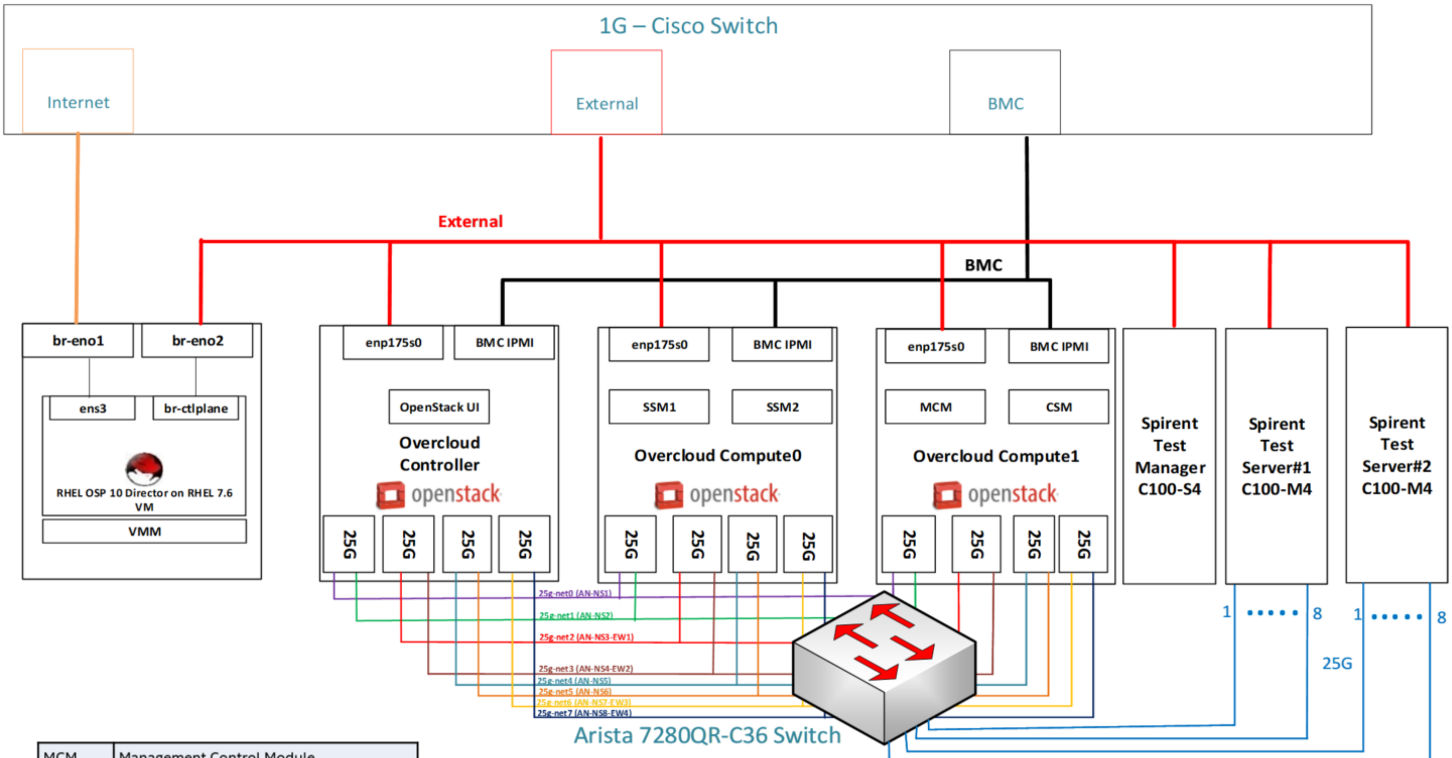

MCM	Management Control Module
CSM	Content Services Module
SSM	Subscriber Services Module

Figure 3. Network Topology Details

The OpenStack network topology in Figure 3 includes the following networks:

- One 1 Gbps External network
- One 1 Gbps PXE/Provisioning network
- One 10 Gbps Management network, with one Intel® Ethernet Converged Network Adapter X710 10 GbE NIC
- 4 x 25 Gbps East-West networks, using two ports of two Intel® Ethernet Network Adapter XXV710 25 GbE NIC
- 8 x 25 Gbps North-South networks, using two ports of four Intel® Ethernet Network Adapter XXV710 25 GbE NICs

4.1 External/Provisioning Network

The external network was used for Internet access from the controller and compute nodes. The role of the provisioning network was to allow the OSP director to deploy the Overcloud nodes, including the Intelligent Platform Management Interface (IPMI) and iPXE and communicate with them. Each Overcloud node had one static IP address reserved for IPMI access.

4.2 Intelligent Platform Management Interface (IPMI)

All Overcloud bare metal systems had an Intelligent Platform Management Interface (IPMI), used by the OSP Director to control power management when using its Ironic service.

4.3 Management Network

The management network was used to manage OpenStack services and instances.

4.4 Data Network

The Data network was composed of the four 25 GbE links used for the East-West traffic between the vEPC VMs within the OpenStack environment and eight 25 GbE links for the North-South traffic between the Spirent Landslide physical Test Servers and the two vEPC data plane Virtual Machines called Subscriber Services Module (SSM). Refer to Section 5, Affirmed Networks* vEPC Deployment on OpenStack* for additional information and definition.

5 Affirmed Networks* vEPC Deployment on OpenStack*

For optimal user-plane performance the Affirmed Networks* vEPC required the RH OSP* to be installed and Single-Root Input / Output Virtualization (SR-IOV) to be enabled on all the Data network NICs in the Compute nodes. In addition, NUMA, CPU Pinning, and Huge pages were configured on the Compute nodes.

This vEPC uses 4 VMs that have both internal and external networks. The VMs are the basic building blocks that provide the ability to perform vEPC functionality. The Affirmed Networks vEPC implementation includes adding appropriate flavors and deploying the following four types of virtual machines (VMs) on the compute nodes.

- Management Control Module (MCM**)** which controls Operations, Administration, and Management, CLI, etc.
- Content Service Module (CSM**)** - a VM instance that runs the tasks needed for call control, IP routing and providing advanced services like video optimization, TCP Proxy, HTTP Proxy, etc.
- Subscriber Services Module (SSM**) –** This type of VM is a User Plane VM, responsible for receiving packets into the MCC and sending the packets out and providing workflow services. We spawn two SSMs in the performance data shown in this document.

Table 5 lists the vEPC VM flavors and core allocations used for Intel® Xeon® Platinum 8280L servers.

Table 5. Intel® Xeon® Platinum 8280L – Core Allocation for vEPC VMs

VM	# OF VCPUS	RAM (GB)	ROOT DISK (GB)	COMPUTE NODE (CONTROL PLANE + DATA PLANE)	NUMA NODE
MCM	16	32	112	Compute Node 0	Socket 0
CSM	16	128	50	Compute Node 0	Socket 0
SSM1	48	128	50	Compute Node 1	Socket 0
SSM2	48	128	50	Compute Node 1	Socket 1

Note: Linux kernel optimizations as well as OpenStack Enhanced Platform Awareness (EPA) features in the VMs were used to optimize performance by using Intel hardware-specific platform technologies.

6 Performance Benchmarks

The Affirmed Networks vEPC solution driven by Intel® Xeon® Platinum 8280L servers delivered a network throughput of up to 168 Gbps and 29 Mpps. The test results in Table 6 are using UDP packets with two user plane VMs and default bearer sessions, with a packet size of 650 Bytes, one million sessions established, using a physical Spirent Landslide. A packet size of 650 Bytes is typically used for benchmarking, as it is the average of package sizes used in product deployment.

KEY PERFORMANCE INDICATOR	INTEL® XEON® PLATINUM 8280L
Packet Size	650 Bytes
Sessions Established	1 Million
IO Throughput (Total Gigabits per second)	167.8 Gbps
IO Throughput (Total Million Packets per second)	28.8 Mpps
Packet Loss	~0% (0.0005%)
SSM Average Core Utilization - %	IO cores – 71% HD cores - 86%

7 Summary

The Affirmed Networks* vEPC Key Performance Indicators, with a million subscribers and a packet size of 650 Bytes with Red Hat* OSP* v10, one Spirent* physical TAS VM, and two physical TS VMs, showed an impressive sustained and stable IO throughout, with low CPU utilization suitable for commercial deployments.

The high core count of the Intel® Xeon® Scalable processors, combined with architectural improvements, feature enhancements, and high memory bandwidth, is a tremendous performance and scalability advantage over previous Intel® Xeon® processor generations, especially in today's NFV environments.

Affirmed Networks and Intel have collaborated to bring together the best of both technologies accelerating time to deployment with deterministic behavior and transforming network technology with value-propositions for Communication Service Providers, including top-of-the-line performance, lower cost per bit, better ROI and as a future proof platform for 5G services.

Appendix A BIOS Settings

Table 6 shows the Basic Input / Output System (BIOS) settings that were used to enable the best achieved performance on the SUT.

Table 6. BIOS Settings

MENU (ADVANCED)	PATH TO BIOS SETTING	BIOS SETTINGS	REQUIRED SETTING FOR DETERMINISTIC PERFORMANCE
Power Configuration	Power and Performance	CPU Power and Performance Policy	Performance
		Workload Configuration	I/O Sensitive
	Power and Performance → CPU P-State Control	Enhanced Intel® SpeedStep Technology	Enabled
	Power and Performance → Hardware P States	Hardware P States	Disabled
	Power and Performance → CPU C State Control	Package C-State	C0/C1 state
		C1E	Disabled
		Processor C6	Disabled
	Uncore Power Management	Uncore Frequency Scaling	Disabled
		Performance P-limit	Disabled
Memory Configuration	Advanced → Memory Configuration	IMC Interleaving	2-Way Interleaving
Virtualization Configuration	Processor Virtualization Feature	Intel® Virtualization Technology (VT)	Enabled
	Integrated IO Virtualization Configuration	Intel® VT for Directed I/O	Enabled
Thermal Configuration	Advanced → System Acoustic and Performance Configuration	Set Fan Profile	Performance

Legal Information

By using this document, in addition to any agreements you have with Intel, you accept the terms set forth below.

You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

Information in this document is provided in connection with Intel products. No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document. Except as provided in Intel's terms and conditions of sale for such products, Intel assumes no liability whatsoever and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right.

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. For more information go to intel.com/benchmarks.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

Performance results are based on testing as of March 25, 2019 and may not reflect all publicly available security updates. See configuration disclosure for details. No component or product can be absolutely secure.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request. Contact your local Intel sales office or your distributor to obtain the latest specifications and before placing your product order.

Intel processors of the same SKU may vary in frequency or power as a result of natural variability in the production process.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at Intel.com.

Intel® Turbo Boost Technology requires a PC with a processor with Intel Turbo Boost Technology capability. Intel Turbo Boost Technology performance varies depending on hardware, software and overall system configuration. Check with your PC manufacturer on whether your system delivers Intel Turbo Boost Technology. For more information, see https://www.intel.com/content/www/us/en/architecture-and-technology/turbo-boost/turbo-boost-technology.html.

All products, computer systems, dates and figures specified are preliminary based on current expectations, and are subject to change without notice. Any differences in your system hardware, software or configuration may affect your actual performance.

Intel® Hyper-Threading Technology requires a computer system with a processor supporting HT Technology and an HT Technology-enabled chipset, BIOS and operating system. Performance will vary depending on the specific hardware and software you use.

Intel does not control or audit third-party web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Intel Corporation may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights that relate to the presented subject matter. The furnishing of documents and other materials and information does not provide any license, express or implied, by estoppel or otherwise, to any such patents, trademarks, copyrights, or other intellectual property rights.

Intel, the Intel logo, Xeon, and others are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

* Other names and brands may be claimed as the property of others.

© 2019 Intel Corporation
0419/SM/PTI/PDF 606430-001US