Performance measurement using dpdk in software-defined local and wide area networks

Ki-hyeon Kim [1] *, Dongkyun Kim [1], Yong-hwan Kim [1], Joobum Kim [1]

[1] Korea Institute of Science and Technology Information, 245 DaehangnoYuseongDeajeon Korea 34141, ASIKRKS015, Daejeon
*Corresponding author E-mail: parkroyal@kisti.re.kr

Abstract

Background/Objectives: For data center-based enterprises, the use of data traffic continues to increase. As a result, the increasingly efficient processing of increasing data traffic is becoming a major issue. Effective processing of traffic can be performed through the addition of physical network equipment, but this requires considerable costs.

Methods/Statistical analysis: To solve this problem, researches on server virtualization and network virtualization technology are being actively conducted. However, while network virtualization is a cost-effective solution, open-source virtual switches used in network virtualization technologies have the performance penalty because they handle packets at the kernel of server's operating system. As a result, the Intel has developed the Data Plane Development Kit (DPDK) technology to address the performance degradation of the operating system.

Findings: In this paper, we propose the DPDK performance test environment based on Open Virtual Switch (OVS) to verify the stability and performance of the DPDK-oriented virtualization technology, in association with a new KREONET-S project which aims tothe softwarization of the Korea Research Environment Open Network (KREONET) infrastructure.

Improvements/Applications: We perform conducted experiments on Software-Defined Local Area Network (LAN) and Wide Area Network (WAN) environment using KREONET-S show the result.

Keywords: *Open Virtual Switch, OVS; Data Plane Development KIT; DPDK; SDN; NFV; Kreonet-S*

1. Introduction

There is a growing demand for Web-based services in companies such as Google, Naver, Amazon, and YouTube, which use national and international data centers. And as more and more attention is focused on big data technology and cloud technology, data-driven network traffic is growing. According to Cisco reports, global mobile data traffic grew to 63 percent in 2016. Global mobile data traffic reached to 7.2 exabytes per month in the end of 2016, up from 4.4 exabytes per month at the end of 2015. Monthly global mobile data traffic will be 49 exabytes by 2021, and annual traffic will exceed half a zettabyte[1]. As the amount of data traffic continues to increase, an efficient transmission scheme for handling a large amount of traffic is needed. In order to efficiently process large-scale traffic in the data center, it is possible to solve the problem by expanding the physical network equipment or expanding the network bandwidth, but the problem of capital expenditure (CAPEX) and operating expense (OPEX) cannot be solved. Therefore, network virtualization technology can be used to efficiently configure and operate the network.However, open-source virtual switches used in network virtualization technology have a problem that the bandwidth of the network cannot be guaranteed.

The reason for not guaranteeing the bandwidth of the network is because that buffer memory allocation / unassignment and interruption problems occur during the processing of large packets in the operating system kernel unit [2], [3]. The Data Plane Development Kit (DPDK) [4] is a technology developed to solve these problems. The DPDK supports frames optimized for high-performance packet processing based on Intel architecture, and supports fast packet processing technology can be used. In this paper, we propose a highperformance packet processing system based onOpen Virtual Switch (OVS) with DPDKto ensure high performance when network virtualization technology is applied.

Korea Research Environment Open Network (KREONET) is a principal national Research & Education (R&E) network running by Korea Institute of Science and Technology Information (KISTI) in Korea. It uniquely provides production research network services for around 200 non-profit research and educational organizations, based on hybrid (IP and non-IP) network infrastructure with 21 network centers in Korea (17), USA (2), China (1), and Netherlands(1)[5].KISTI has constructedthe Software-Defined Network (SDN) based programmable and virtualized core and edge / access network infrastructure as an open control platform as the next service model for advanced science research networks and named KREONET-S.KREONET-Sis a new network project to drive softwarization of KREONET Infrastructure.KREONET-S consists of data plane and control plane for SD-WAN[6,7].In order to construct the control plane, Open Network Operating System (ONOS)[8] which can configure the distributed control environment was adopted and we are running 5-node ONOS cluster in the distributed manner for data plane networks of KREONET-S.In order to construct the data plane, 6 core network devices and 6 edge network devices were installed in Daejeon, Seoul, Busan, Gwangju, Changwon, and Chicago area network centers. Each switch is also connected to the ONOS controller.In this paper, we performed performance measurement using Software-Defined Wide Area Network (SD-WAN) between Daejeon and Seoul network centers for OVS-DPDK experiment.

Network Function Virtualization (NFV) refers to a technology that allows software services to virtualize network functions, allowing them to be relocated flexibly. On the basis of this, there is a need

for service chaining in order to selectively connect and execute network functions as needed. However, there is a problem of performance degradation when service chaining is configured. In order to solve this problem, some experiments have been performed using the DPDKto construct service chaining.Based on this, we will explain about the research trends using DPDK.

Michail-Alexandros[9] presents the problem that network performance degradation problems exist when VNF is deployed on a VM in a cloud environment in order to composeNFV.In order to solve the problem of performance degradation, this paper performs performance verification by comparing two methods, such as 1) using DPDK and Single Root I/O Virtualization (SR-IOV) 2) not using DPDK and SR-IOV.It shows higher performance when using DPDK and SR-IOV.

Kawashima et al. [10] measured throughput, latency and jitter using virtual switches (Linux Bridge, OVS, DPDK L2FWD, OVS-DPDK, VALE, Lagopus). In terms of hardware, tested the effect on NIC type (Intel X710, Mellanox ConnectX-3 EN). In addition, this paper shows the superior performance of DPDK in the packet processing architecture divided into Linux New API (NAPI), netmap and DPDK. Emmerichet al. [11] performed experiments on performance bottlenecks using virtual switches (Linux Bridge, OVS, DPDK L2FWD, and OVS-DPDK). In terms of hardware, we tested the effect on NIC type (Intel X520, Intel X540). From the software side, we analyze the packet processing overhead of the kernel using the Linux data collection tool (perf).

Wei Zhang et al. [12] is a paper that constructs NFV using DPDK and a system called OpeNETVM shows its performance. In terms of components of OpenNETVM, Docker was used to lighten network service functions, DPDK was used to guarantee high performance and to generate effective service chaining. Service chaining in this paper consists of NF manager and flow table. The flow table is configured using shared memory of DPDK, while NF manager delivers packets to the network interface card (NIC) using the flow table and checks the status of network functions.

In the above papers, DPDK is used to solve the performance problem in constructNFV. Experiments using DPDK have demonstrated the superiority of DPDK. However, in the case of the above papers, the experimental environment is limited because the experiments were performed in the LAN environment.There is a limitation that cannot show all phenomena when users are used in LAN environment. Therefore, in this paper, we perform experiments both in SD-LAN and SD-WANenvironments and show data transmission performance. Also, if OVS with DPDK is connected to the ONOS controller, we want to check whether service chaining works normally. The composition of this paper is as follows. In Section 2, we compare the data transmission performance of each virtual switch in the SD-LAN and SD-WAN, and describe the system configuration and experimental methods for the service chaining performance experiment based on using OVS with DPDK. Section 3 shows the performance measurement results for each experiment. Section 4 describes the conclusion of this paper. Finally, Section 5 explains the future work.

2. System applying environment and experiment method

In this paper, we compared the data transmission performance of each virtual switch in SD-LAN and SD-WAN and conducted experiment of service chaining technology using the DPDK. The reason why this experiment is divided into two experiments is as follows. The background of the DPDK can be divided into two. The first is the need for technology to ensure the performance of the network in a virtualized environment, and the second is the ability to configure service chaining technology as a technology widely used in NFV technology.So the first scenario is to identify good vSwitches among the open-sourcevSwitches that are currently available as open source.The second scenario is to construct a service chaining network environment using the DPDK and to verify

whether the performance of the network can be guaranteed based on the configured environment.

Three physical servers were used to construct the experiment environment, where two servers were used for network traffic generation, and the other server was used for configuring open-source virtual switches. The experimental environment of the server is shown in Table 1. OVS, DPDK, and Quick Emulator (QEMU) were used to construct the service chaining network. The open source version was version 2.6.1 for OVS, version 16.11 for DPDK, and version 2.10.0 for QEMU.

**Table 1:** Experiment Environment of Server Used For Experiment

CPU	Intel Xeon E5-2695 2.3GHz, 14Core
Memory	DDR4 128GB
NIC	Intel X520 10GB DA/SFP+
HDD	400GB SSD SAS Write/12Gbps
Kernel ver.	linux 3.10.0-514.26.2.el7.x86_64

2.1. Comparison of data transfer performance by each virtual switch

Experiments on the performance of data transmission by each virtual switch were performed inSD-LAN and SD-WAN environment. Experiments were conducted on the physical server environment and the virtual server environment in the SD-LAN. The reason for conducting the experiment in the virtual server environment do haveconsists of similarly experiment environment using a large amount of data such as a data center or a cloud service.

2.1.1. Performance experiment of SD-LAN

Experiments in the SD-LAN were conducted by dividing into physical server environment and virtual server environment. Three open source virtual switches were tested using OVS, Linux Bridge and OVS with DPDK. We are Install a virtual switches on one server and connect the two servers to the other NIC for performance measurements.The system configuration of the physical machine (PM) environment is shown in Figure 1(a). The system configuration of the virtual machine (VM) environment is shown in Figure 1(b).VM environment was performed by the same experimental method as PM environment.In the other experiment environment, one VM is respectively installed in two servers for performance measurement and a virtual switch performance measurement experiment is performed between VMs. and VMs are installed on each server to perform performance measurement.
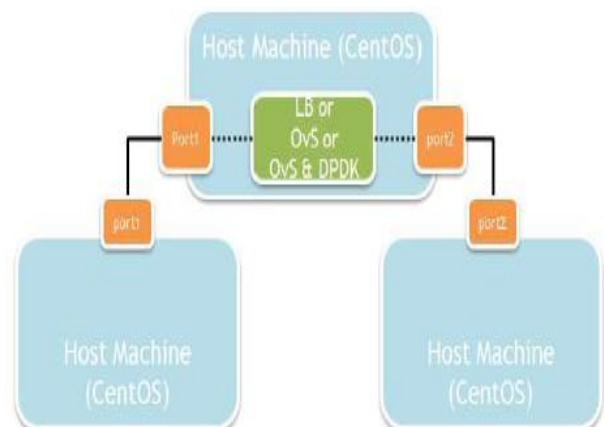
Fig. 1: Physical and Virtual System Diagram for Each Vswitch Performance Measurement in SD-LAN.

2.1.2. Performance experiment of SD-WAN

Experiments in the SD-WAN were performed by installing a performance measurement server and a virtual switch server in Daejeon and Seoul network centers ofKREONET-S. Figure 2 shows the SD-WAN performance measurement system environment. The virtual switches used for performance measurement were OVS, OVS with DPDK. Based on the performance measured results on the SD-LAN, it is difficult to compare the performance in the PM environment in SD-LAN. Therefore, this experiment was performed from VM environment.
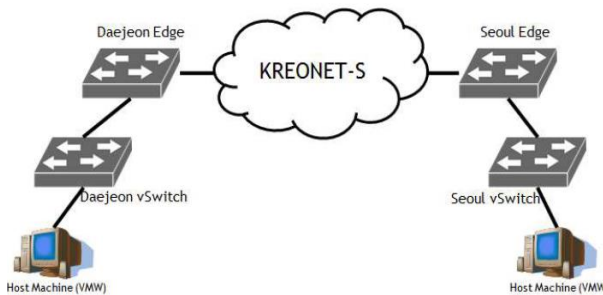

Fig. 2: Network System Diagram for Each Vswitch Performance Measurement in SD-WAN.

2.2. Service chaining technology applying experiment-using DPDK

2.2.1. Performance experiment of SD-LAN

For the performance experiment of SD-LAN, install OVS, DPDK, and QEMU on the server to be used as vSwitch.QEMU is used as a hypervisor to install VMs on the server and to connect the DPDK network between the VM and the host machine. The resources allocated to each created VM are allocated 4core CPU, 4096MB memory, and 30GB HDD. In the experiment, performance measurement is conducted while increasing the number of VMs installed from one to four. The flow rules in OVS are configured to go through all VMs whenever a VM is created. Figure 3 shows the system configuration applying toservice chaining technology using OVS with DPDK. The performance change of OVS with DPDK is measured according to the number of VMs in the SD-LANthrough the system configuration as shown in the figure 3.
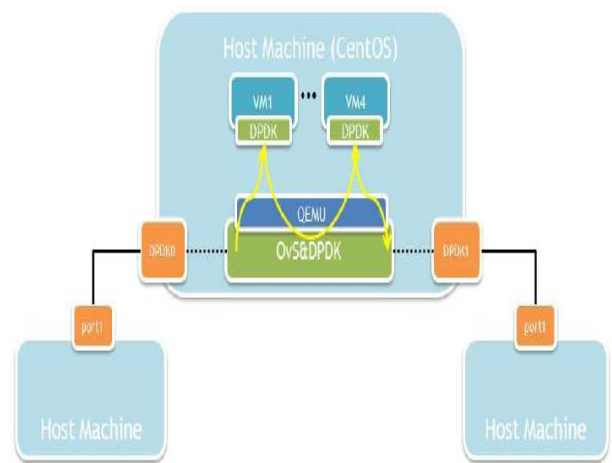

Fig. 3: Physical System Diagram of Service Chaining In SD-LAN.

2.1.2. Performance experiment of SD-WAN

Performance experiment in the SD-WAN were performed by installing two performance measurement servers and a virtual switch server in Daejeon and Seoul network centers of KREONET-S. Figure 4 shows the service chaining network performance measurement system in the SD-WAN. The server to be used as a vSwitch in Seoul was installed in the same way as the system configuration installed in SD-LAN.Performance measurement was conducted while increasing the number of VMs from one to three.Figure 6 shows the performance of OVS with DPDK depending on the number of VMs in SD-WAN.
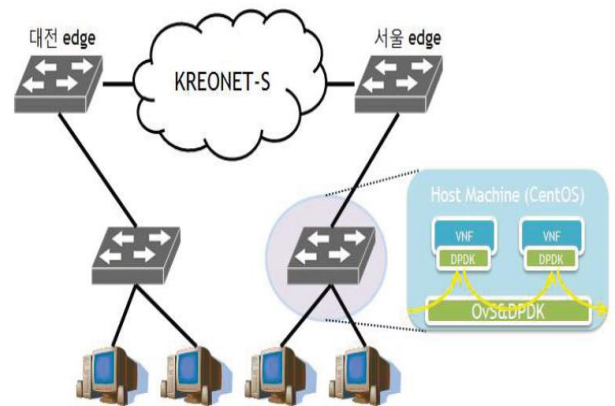

Fig. 4: Network System Diagram of Service Chaining In SD-WAN.

3. Experiment and analysis results

3.1. Experiment data collection method

The performance measurement server was directly connected to the server using the vSwitch in SD-LAN, the server used as vSwitch was installed in the Seoul network center and the performance measurement server was connected to the Seoul and Daejeon network centers in SD-WAN. In order to collect experimental data, iperf3 was used to measure network performance of the server. Data processing and transmission time were measured in a single stream at the intervals of1 secondand 60 seconds total, and the performance resultswere analyzed by performing 10 times.

3.2. Results of virtual switch data transfer performance comparison

Figure 5 shows performance comparison graphs for virtual switches as Linux Bridge, OVS, OVS with DPDK in a physical server environment and a virtual server environment. In the physical server environment shown in Figure 5(a), the Linux Bridge suddenly drops to 5.5Gbits / sec. On the other hand, OVS, OVS with DPDK shows

a stable graph of 9Gbits/sec. In the virtual server environment shown in Figure5 (b), the performance of the Linux Bridge and OVS can be seen to vary greatly. However, in the case of OVS with DPDK, it is seen that performance fluctuation does not appear. This experiment shows that OVS with DPDK is superior to Linux Bridge and OVS in terms of stability and performance.

Figure 6 is a graphs of the experimental results in the virtual server environment in the SD-WAN. In the case of OVS, performance in 1, 2, 8, and 9 experiments drops below 8Gbits/sec. This is a problem when processing packets using OVS at the Linux kernel level. However, in the case of OVS with DPDK, it is seen that performance fluctuation does not appear. OVS with DPDK are superior to OVS in terms of stability and performance.
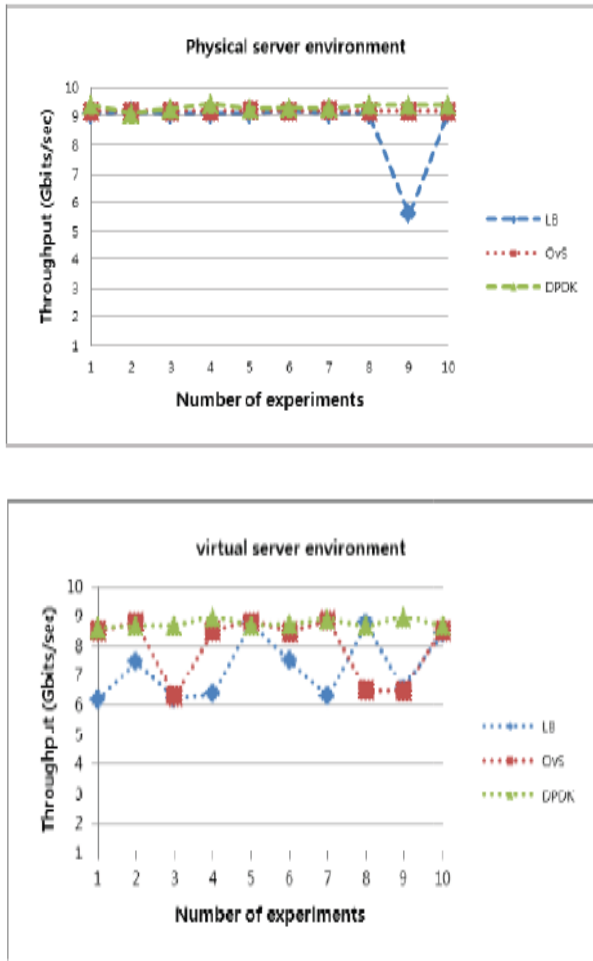

Fig. 6: Bandwidths Comparison per Each Experiment Time in Physical and Virtual Server Environment on SD-WAN.

3.3. Experiment results applying service chaining technology-using DPDK

Experiments on the SD-LAN have been carried out about performance changes as increasing the number of VMs whencomposed of service chaining technology. Figure 7(a) is a graph showing the performance change according to the experiment number of data when the data is transmitted as a single stream as the number of VMs increases. The graph shows that the performance of 1VM and 2VM is not much fluctuating, and it is close to 9Gbits/sec. However, performance starts to drop in 3VM and 4VM. Therefore, in this experimental environment, the optimal number of VMs to perform Service Function Chaining (SFC) is two. It can be seen that the high performance is not guaranteed when the number is three or more in this experiment environment.

When composed of the service chaining environment in the SD-WAN, the performance change is experimented as the number of VMs increases. Figure 7(b) is a graph showing the performance change according to the experiment number of data when the data is transmitted as a single stream as the number of VMs increases. The graph shows that the performance of 1VM is not fluctuating and it is close to 9Gbits/sec. However, as the number of VMs increases from 2VM, performance begins to decrease. Therefore, in this experimental environment, the optimal number of VMs to perform SFC is one. It can be seen that the high performance is not guaranteed when the number is one or more.

Fig. 5: Bandwidths Comparison per Each Experiment Time in Physical and Virtual Server Environment on SD-LAN.

Experimental results in the SD-LAN and SD-WAN are due to the use of PMD (Poll Mode Drivers), which are responsible for packet processing in the DPDK, rather than processing packets at the Linux kernel. PMD consists of the APIs provided by BSD drivers running in user space to configure devices and their queues. In addition, the PMD directly accesses the RX and TX descriptors without interrupts (except link state change interrupts) to allow the user application to quickly receive, process, and forward packets[4].

Fig. 7: Performance Change According To the Number of Vms in Service Chaining Environment.

4. Conclusion

In this paper, we have dealt with DPDK, a high-performance packet processing technology based on the Intel architecture that overcomes network problems in the operating system. We also performed experiments comparing data transmission performance per virtual switch, and service chaining technology using DPDK. In the experiments of comparison for the virtual switch-specific data transmission performance, three types of virtual switches as Linux Bridge, OVS, OVSwith DPDK were installed in SD-LAN and SD-WAN environment to compare network performance between the virtual switch. When performance measurement is conducted using a single stream in a SD-LAN environment, the virtual server environment shows a bigger change than a physical server environment. In the virtual server environment, the Linux Bridge has a variation of about 3Gbits/sec from 6Gbits/sec to 9Gbits/sec. In the case of OVS, like Linux Bridge, the performance varies from 6Gbits/sec to 9Gbits/sec, which is about 3Gbits/sec. Therefore, it is difficult to guarantee network performance. In the case of OVS with DPDK, there is almost no variation in terms of performance, and the performance of 8.6Gbits/sec is maintained. As a result, OVS with DPDK proved the superiority and safety in terms of data transmission performance in the VM environment as well as physical machine environment.

In the experiment using the service chaining technology, the service chaining environment was configured using DPDK, and then we increased the number of VMs from 1 to 3 both in the SD-LAN and SD-WAN environment. In the SD-LAN environment, the variation of performance from 1 VM to 4 VM is not significant, but the data transmission performance becomes lower below 8Gbits/sec as the number of VMs increases to 3 or more. Therefore, the number of ideal VMs in this experimental environment is vitrified as two. However, if the performance of the server is improved or the performance of the network is enhanced, the performance can be guaranteed up to three number of VMs. In case of 1VM in the SD-WAN environment, there is not much variation in performance. However, as the number of VMs increases to 2 or more, the performance becomes lower than 8Gbits/sec. In the SD-WAN environment, it is identified that data transmission performance is lower than the SD-LAN environment.

5. Future work

In this experiment, we verified the network performance using OVS with DPDK. Also, service chaining was configured using OVS with DPDK. However, when the service chaining is activated, the performance decreases as the number of VMs increases. Therefore, our future work will search for a method to guarantee performance as the number of VM increases when service chaining is applied using both OVS and DPDK.

Acknowledgment

References

[1] Cisco, https://www.cisco.com/c/en/us/solutions/collateral/service-provider/ visual-networking-index-vni/mobile-white-paper-c11-520862.html

[2] Intel, http://www.intel.com/content/dam/doc/application-note/ gbe-controllers-interrupt-moderation-appl-note.pdf

[3] Walravens, Cedric, and Benjamin Gaidioz. "Receive descriptor recycling for small packet high speed Ethernet traffic." In Electrotechnical Conference, 2006. MELECON 2006. IEEE Mediterranean, 2006, pp. 1252-1256.

[4] Data Plane Development Tools web site, Retrieved Sep., 25, 2016, from http://dpdk.org//

[5] KREONET web site, Retrieved Sep., 25, 2016, from http://www.kreonet.net/

[6] KREONET-S web site, Retrieved Sep., 25, 2017, from http://www.kreonet-s.net/

[7] Kim, Dongkyun, Hyunhun Cho, Yong-Hwan Kim, Kyu-Il Kim, Ki-Sung Yu, and Joon-Min Gil. "User-Oriented Software-Defined Wide Area Network Adopting Virtual Dedicate Networks." Advanced Science Letters 22, 2016, no. 9, pp. 2262-2267.

[8] Open Network Operating System web site, Retrieved Aug., 25, 2016, from http://onosproject.org/

[9] Kourtis, Michail-Alexandros, Georgios Xilouris, Vincenzo Riccobene, Michael J. McGrath, Giuseppe Petralia, HarilaosKoumaras, Georgios Gardikis, and Fidel Liberal. "Enhancing VNF performance by exploiting SR-IOV and DPDK packet processing acceleration." In Network Function Virtualization and Software Defined Network (NFV-SDN), 2015 IEEE Conference on, 2015, pp. 74-78.

[10] [Kawashima, Ryota, Shin Muramatsu, Hiroki Nakayama, Tsunemasa Hayashi, and Hiroshi Matsuo. "A Host-Based Performance Comparison of 40G NFV Environments Focusing on Packet Processing Architectures and Virtual Switches." In Software-Defined Networks (EWSDN), 2016 Fifth European Workshop on, IEEE, 2016, pp. 19-24.

[11] Emmerich, Paul, Daniel Raumer, Florian Wohlfart, and Georg Carle. "Assessing soft-and hardware bottlenecks in PC-based packet forwarding systems." ICN 2015, 2015, pp.78-83.

[12] Zhang, Wei, Guyue Liu, Wenhui Zhang, Neel Shah, Phillip Lopreiato, GregoireTodeschi, K. K. Ramakrishnan, and Timothy Wood. "OpenNetVM: a platform for high performance network service chains." In Proceedings of the 2016 workshop on Hot topics in Middleboxes and Network Function Virtualization, ACM, 2016, pp. 26-31.