DE GRUYTER
OPEN

JACSM 2015, Vol. 7, No. 2, pp. 117 - 132
10.1515/jacsm-2015-0013

PERFORMANCE ANALYSIS OF VIRTUAL COMPUTER NETWORK BASED ON CISCO CLOUD SERVICES ROUTER 1000V IN A PRIVATE CLOUD ENVIRONMENT

Stanislaw Lota[1], Marcin Markowski[2]

[1] Maria Sklodowska-Curie Upper-Secondary School Complex in Olesnica
Wojska Polskiego 67/69, 56-400 Olesnica, Poland
stanislaw.lota@gmail.com

[2] Department of Systems and Computer Networks,
Wroclaw University of Science and Technology,
Wybrzeze Wyspianskiego 27, 50-370 Wroclaw, Poland
marcin.markowski@pwr.edu.pl

Abstract

Virtualization of physical network devices is a relatively new technology, that allows to improve the network organization and gives new possibilities for Software Defined Networking (SDN). Network virtualization is also commonly used for testing and debugging environments, before implementing new designs in production networks. Important aspect of network virtualization is selecting virtual platform and technology, that offer maximal performance with minimal physical resource utilization. This article presents a comparative analysis of performance of the virtual network created by the virtual CSR1000v and virtual machines running Windows 8.1 on two different virtual private cloud platforms: VMware vSphere 5.5 and Microsoft Hyper-V Server 2012 R2. In such prepared testbed we study the response time (delay) and throughput of virtual network devices.

Key words: virtual network, virtualization platforms, network performance, Cisco Cloud Service Router 1000v, Hyper-V, VMware ESXi,

1 Introduction

Currently, the physical network devices such as routers, switches, and servers are the base of computer networks. Popularity of cloud computing forces the research for new network solutions, integrating physical network environments with virtual environments. For a few years works on transferring the functionality of physical network devices to their virtual counterparts and Software Defined Networking (SDN) are being conducted.

Network virtualization technology enables reproduction of the physical network using the virtual network devices running on a single physical server, offering the same functionality as the physical devices. Multiple virtual network devices, such as routers, switches, firewalls, intrusion detection systems, may be used in a virtualized environment. It gives the opportunity to organize network in better, more efficient manner [9]. Other advantages are the optimization of costs, short implementation time, faster network integration, reduction of cabling, independence from expensive physical devices, simpler and faster disaster or failure recovery. Virtualization ensures also high energy-efficiency, since many virtual systems (network devices or servers) utilize single hardware platform. Currently, the virtual computer network can be associated with existing network hardware providing full communication.

One of the important challenge, faced by network virtualization solutions is the performance deficiency. The operational delay, usually very small for hardware network devices, in virtualized devices is being often multiplied as the reason of underlying hypervisor software and hardware platform not optimized for network traffic processing. Also the maximal throughput, offered by virtual network solutions often becomes a bottleneck.

In this context, very important decision in virtual network designing and implementing process in the choice of optimal, efficient virtualization platform, that will not significantly decrease the network performance (Quality of Service - i.e. delay in the network; throughput, etc.). Then, we provide short study on the efficiency of two of the main virtualization solutions.

The main purpose of this article is to present and compare the performance of two virtualization platforms, VMware vSphere 5.5 and Microsoft Hyper-V Server 2012 R2, used as a base for virtual network solution. Experimental virtual network is built of Cisco Cloud Service 1000V Routers and virtual hosts. Performance has been studied in the area of Quality of Service of the virtual network i.e. bandwidth, delay and jitter. Chosen criteria will accurately determine whether network bandwidth corresponds to the capabilities offered by the devices, whether the traffic is not blocked on specific ports, or whether large packet loss and unstable operation of the network may be observed.

In section 2 we present short review of virtual network devices, section 3 gives the look at actual state of art in network virtualization studies. Then, in section 4 we formulate the research problem and present the testbed. Results of experiment are reported in section 5. Finally, we conclude the research in section 6.

2 Virtual Network Devices

The first implementation of virtual network device on PC platform was a computer equipped with two network adapters and specialized software. Such solutions were usually built on FreeBSD distributions and could act as a router, firewall, DHCP server, DNS, or VPN. Nowadays many project, such as pfSense, m0n0wall or VyOS, develop software packages that provides features similar to commercial hardware boxes, and often gaining additional functionalities and greater control of security [2]. These systems can be successfully installed in a VMware environment, as well as Hyper-V virtual machines. In more complex virtual environments, e.g. when we need to connect few virtual servers working on single physical server, we may use physical network device (router or switch) connected to the server to provide connection (and routing) between virtual machines [2]. Another, often chipper and less complex solution is to use virtual network device, located on the same physical server. We can now virtualize routers, switches, firewalls, intrusion detection and intrusion prevention systems (IDS/IPS), as well as load balancers, NetFlow collectors and less common network devices on one server with hypervisor system. These devices can work independently using the same hardware platform. Among virtual device solutions we may distinguish layer 2 and layer 3 switches, routers, firewall and others. An example of virtual Layer 2 switch are the Cisco Nexus 1000V (VMware Distributed Virtual Switch) and open source Open vSwitch. Features of Layer 3 switches are implemented in virtual switch Nexus 1000V, developed by Cisco. The Cisco Virtual Nexus1000V working under the control of the NX-OS and allows to create PVLAN, virtualized DMZ zones and implementation of the policies for advanced network security e.g. ACL together with QoS [4]. ASA 1000V Cloud Firewall is a virtualized version of a hardware firewall ASA 5500 series and provides protection coastline and Virtual Security Gateway (VSG) responsible for the protection of the network using VMware vShield APIs for internal security. For application of routing between virtual machines in a private, public or hybrid cloud environment, Cisco created a virtual router enabling selected functions of the operating system IOS-EX - Cisco Cloud Service Router 1000V. The virtual device is designed for deployment in data centers in the cloud and run as a virtual machine on servers that use virtualization platforms VMware ESXi, Citrix XenServer, Kernel Virtual Machine (KVM) and Microsoft Hyper-V [5]. Noteworthy is also the possibility of implementing Cisco Cloud Service Router 1000V on public cloud Amazon AWS and Microsoft Azure [6]. The main use of the virtual router is acting as a gateway to the WAN for multitenant and secure connection between the provider of public cloud and enterprise. These functions can be implemented using IPSec VPN (DMVPN, EasyVPN and FlexVPN), or MPLS (Multi-Protocol Label Switching). Cisco Cloud Service

Router 1000V is licensed base on a combination of performance and feature set. Virtual machine with CSR1000V deployed on servers requires from 1 to 4 virtual CPUs, from 2.5 to 4GB of RAM, depending on the performance and feature set, 8GB of disk space and three or more virtual NICs vNICs. The device can provide routing functions between virtual machines using protocols such as OSPF, EIGRP, and BGP, Multicast, LISP, GRE [6]. Another solution is to implement a cloud computing environment with virtual router Brocade Vyatta vRouter and community version VyOS. Brocade provides routing based on BGP Multipath, PBR, OSPF, Multicast technologies, IPsec VPN environments for physical, virtual and cloud-based environments. Vyatta vRouter can also function as a firewall but does not support MPLS [1].

3 Related Works

Virtual network operation seems to be important practical problem, but literature connected with the topic is very limited. The most interesting publication is [8], where authors present optimization problem where the objective is to minimize the network resource consumption with virtualization support (NFV-capable nodes), such that the service requirement (order of service chain traversal) for all the traffic flows is satisfied. In [10] authors propose a virtual network architecture for cloud computing and present research about virtual network which can provide communications for virtual resources in cloud computing. It can potentially reduce the global CO2 emission. Furthermore, without purchasing, operating, maintaining, and periodically upgrading local computing infrastructures, cloud computing can lower the cost of IT services for an enterprise. Multiple virtual networks can run simultaneously over a single physical infrastructure without interfering with each other. In this research virtualized network components such as links, bridges and routers were considered. This virtual network can provide the communication between virtual hosts with flexibility. Furthermore, the virtual network can run the customized routing protocol [10]. Furthermore, Ka Ching Chan and Mary Martin in [3] present an infrastructure enabling lecturers to design and set up experiments in not only traditional networking topics such as RIP, OSPF, BGP, and VLAN using a combination of physical and virtual networking devices, but also in the latest technologies such as server virtualization and network virtualization. They present the development of an integrated virtual and physical network infrastructure for the Internetworking Laboratory at La Trobe University’s Bendigo campus. The infrastructure was setup with physical equipment including Cisco routers, Cisco switches, a number of Ubuntu Linux workstations, and a VMware ESXi server hosting a number of virtual machines including Vyatta routers, and

Ubuntu virtual desktops [3]. In article [11] authors evaluated performance of virtual router platforms based on Linux namespaces and show that hardware assisted virtual routers can achieve better aggregate throughput than a non-virtualized router on a multi-core platform [11]. It is noteworthy that in 2014 years Cisco in a white paper [7] describes technology Virtual Extensible LAN (VXLAN) and how to use CSR 1000V to route between VXLAN segments (VXLAN Layer 3 routing) in addition to switch Cisco Nexus 1000V support for VXLAN.

4 Research Problem and Experimental Environment

The main research problem considered in this paper is the evaluation performance of the network created by the virtual machines and the Cisco Cloud Services Router 1000V. Evaluation criteria are bandwidth, delay, and jitter. Bandwidth is the amount of data that can be sent over the network between two of its points, e.g. router – router or computer - computer in a unit of time. Bandwidth is measured in bits per seconds and is particularly important in the case of transferring large amounts of data over a network. The delay (latency) is the time needed for packet to flow between two designated points. The lower the latency, the better network performance, because the local infrastructure should not exceed a few milliseconds. Jitter is important for real-time transmissions, usually using UDP protocol, e.g. VoIP and multimedia transmissions.

Experimental environment was built on server equipped with 8 core FX-8150 processor clocked at reference period of 3.6 GHz, 16 GB of RAM and two 500GB disks HDD speed 7200rpm and 16MB cache.
Two experimental virtual cloud platforms were implemented on the physical server:

- VMware vSphere 5.5 ESXi 5.5 hypervisor virtualization platform. To create a private cloud environment uses Vmware vCloud Director, vCenter 5.5 vCenter SSO (Single Sign-On) and vSphere Web Client.
- Microsoft Hyper-V Server 2012 R2. To create a private cloud environment uses Microsoft System Center 2012 R2 Virtual Machine Manager 2012 R2.

In order to conduct performance measurement, virtual network based on the Cisco Cloud Service Router 1000V was implemented in both virtual environments. Topology of the virtual network is presented in the Fig. 1. The experimental virtual network consists of three virtualized Cisco Cloud Services Router 1000V, which were connected together by means of virtual interfaces based on virtual standard switches, to define between them three subnetworks. Each of the routers also handle a single network which contains virtual machines with Windows 8.1 Professional. Static routing was implemented in the virtual network. All experiments were performed in a client – server model. The first virtual machine with Windows 8.1 acted as the server, the other virtual machines Windows 8.1 acted as a client.

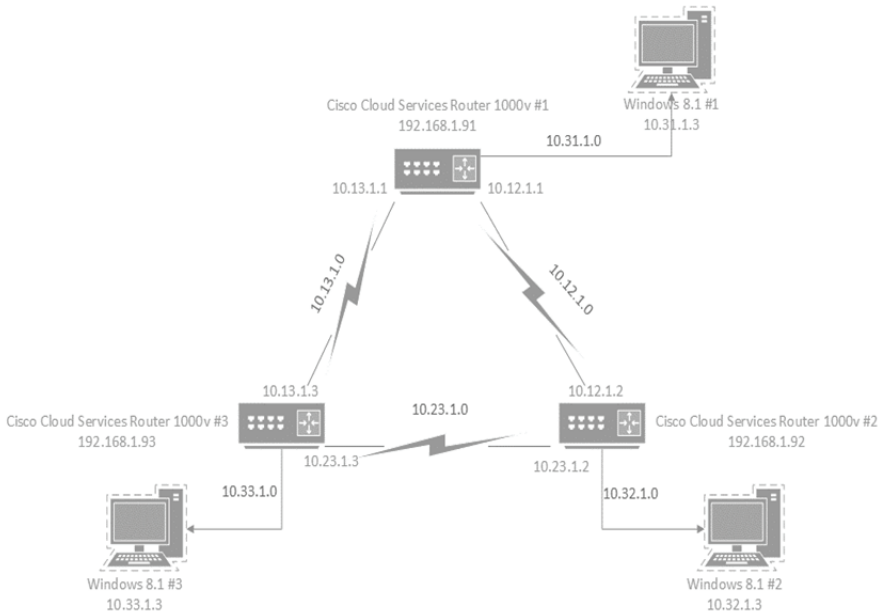

Figure 1. Virtual network topology

A tool Psping included in the PsTools package, which being extended version of the traditional tools Ping, was use for demand generating in the virtual network. The tool allows to check the availability of devices through ICMP, send test packets to any TCP port, perform the measurement of the delay and bandwidth in client - server architecture. Iperf tool was also used to determine the bandwidth between two computers, delay, loss of datagrams, delay variation. Iperf allows testing TCP and UDP on selected ports.

5 Results

The goal of first test was to examine the estimated packet round-trip time [in milliseconds] for 32-byte packet size. We performed 10 experiments, 1000 ping request in each. For each experiment we calculated minimal, maximal and average delay. All attempts were successful and no packets were lost. A large difference between the minimum and maximum delay for both of the test platform was observed during this experiment. Minimal response times for each experiment are presented in the Fig. 2. Lowest delay was noted for

VMware environments, while in the case of Hyper-V delay of each trial decreased in consecutive experiments. Lowest delay obtained in all experiments was 0.6ms for Hyper-V and 0.39 ms for VMware.

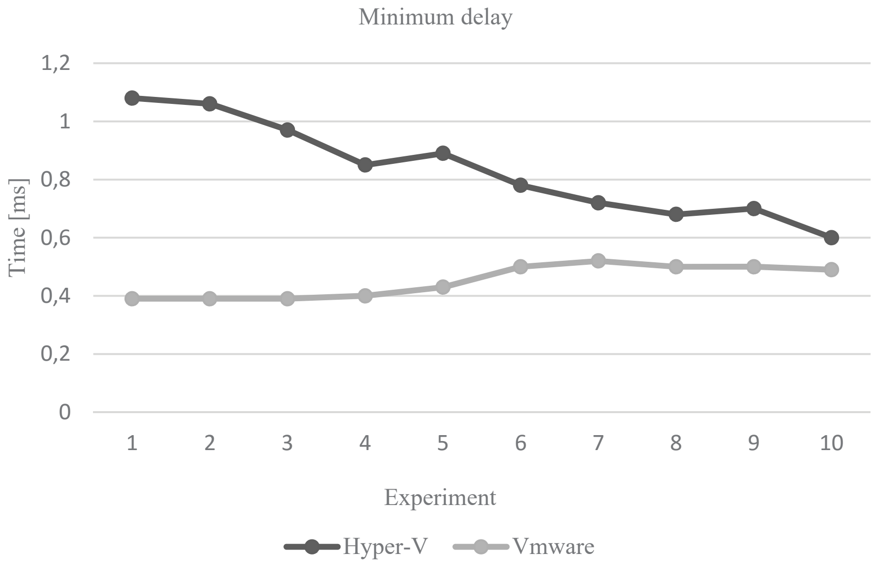

Figure 2. Minimal delay times in each experiment for 32B packet size

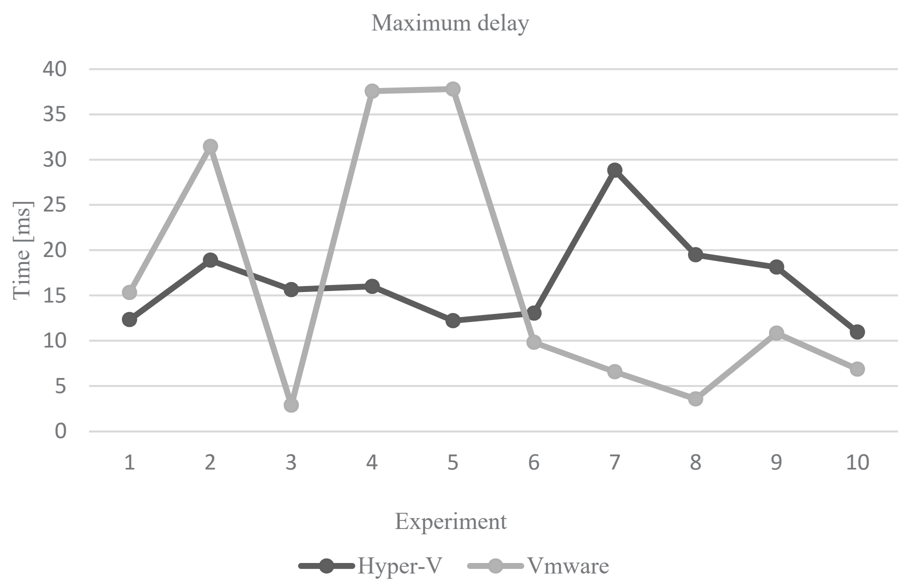

Figure 3. Maximal delay times for 32B packet size

Maximal response times for the same scenarios are presented in the Fig. 3, and average ones in the Fig. 4. VMware several times noted higher delay in comparison to the Hyper-V (in second, fourth and fifth experiment). The biggest delay fluctuated between 2.87 ms and 37.79 ms for VMware and from 11 to 28.91 ms in case of Hyper-V. Average time fluctuated around 2.50ms for Hyper-V, whereas for VMware within 0.60 ms.

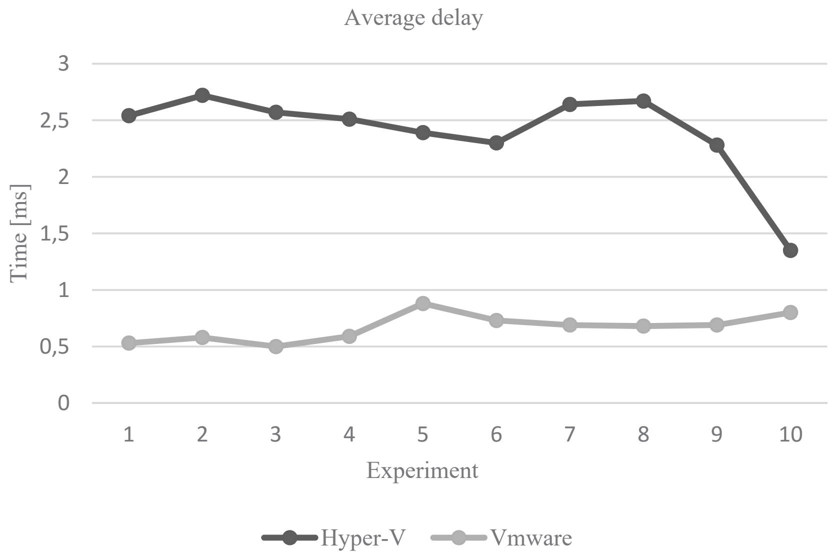

Figure 4. Average delay times for 32B packet size

The second test was to estimate the delay for 8192B packet size. As previous we performed 10 experiments, 1000 ping request each. Results of experiments are presented in the Fig. 5-7.

As we may observe from the results, for VMware environment delay is much more unpredictable and varies in range from 1 to 4000 ms, with the average delay around 300ms. Remarkably better results were obtained for Hyper-V platform, delay fluctuated in range from 7 ms to 38 ms with average around 28 ms.

Although average delay was considerably lower for Hyper-V, the lowest delays were observed for VMware–based virtualization platform. Compiling those results with results for previous experiments for 32B packets size, we may conclude that for some moments VMware is able to ensure better answer time, but generally is very unstable. Taking into account delay experiments, Hyper-V virtualization platform offers better and more stable efficiency.

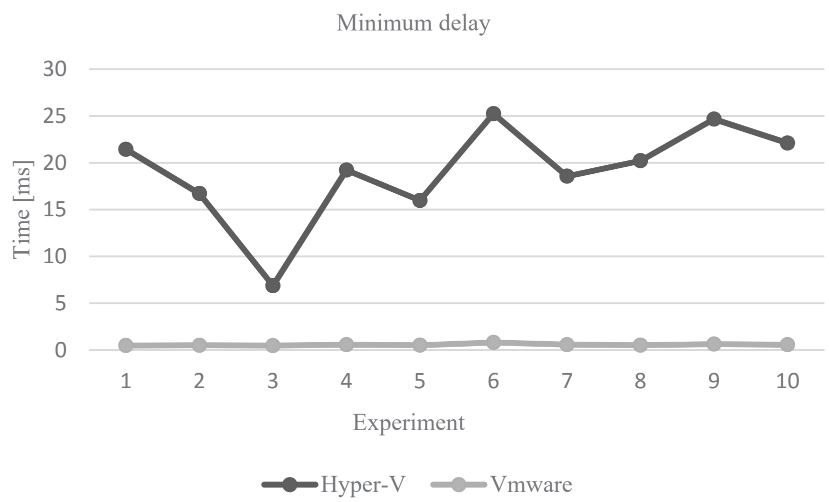

Figure 5. Minimal delay times for 8kB packet size

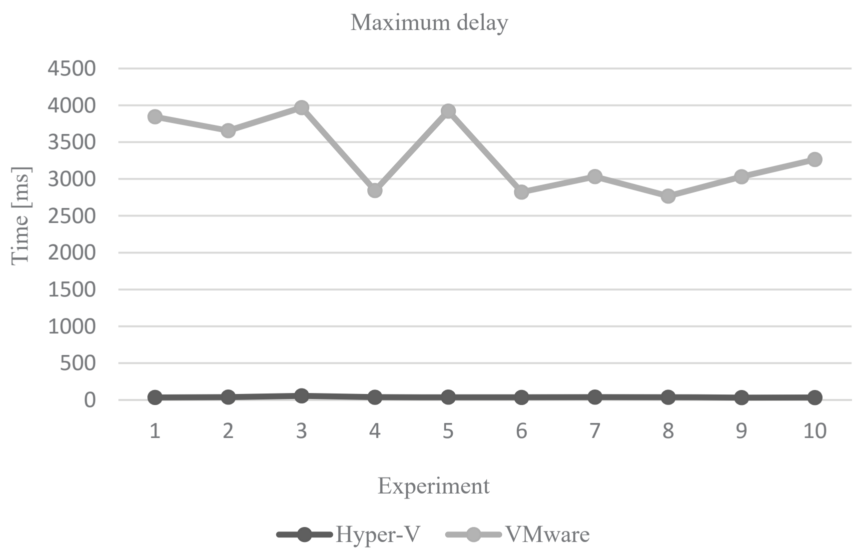

Figure 6. Maximal delay times for 8kB packet size

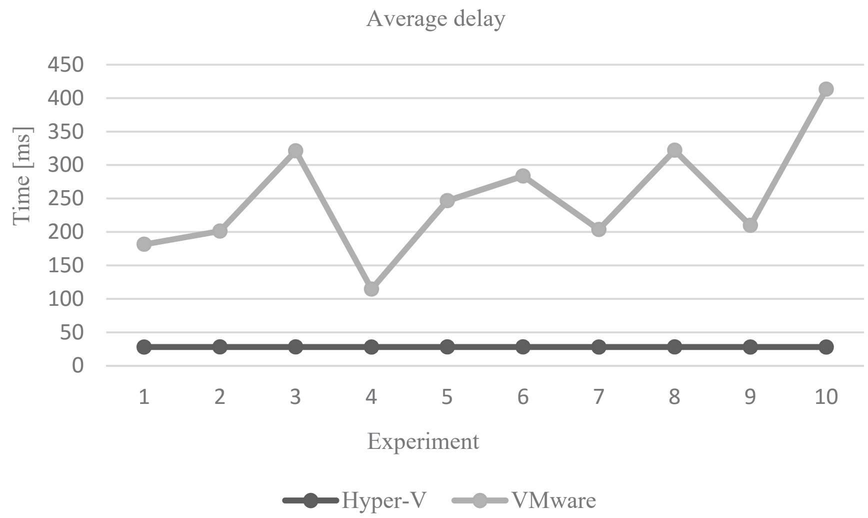

Figure 7. Average delay times for 8kB packet size

The aim of the third experiment was to estimate the capacity of the virtual network. As previously 10 experiments were performed for both evaluated platforms. In each experiment 10,000 packets were sent in the virtual network. Packets were generated with PsPing, and the size of packets was set to 8192 bytes. In the distribution for non-commercial use, Cloud Services Router 1000v offers a maximum bandwidth of 2.5 Mb/s, then we do not expect higher capacity in virtual network. Results of experiment – higher, lower and average throughput are presented in the Fig. 8-10. Unlike during previous experiments, VMware environment was more stable and generally offered better and regular traffic ability performance. Minimal throughput for VMware was newer below 250 Kb/s. In Hyper-V, the smallest throughput was 74.76 Kb / s, and fluctuated up to 236.23 Kb/s.

The greatest maximum throughput was reached in a Hyper-V environment. A little worse efficiency was offered by VMware. In the case of an average throughput results are very similar in both environments with a slight predominance of the VMware environment. Highest average throughput for VMware was 289.33 Kb/s and for Hyper-V 287.42 Kb/s. In the case of 8th attempt for Hyper-V we have noticed the lowest score of at least 331.27 Kb/s and average equal to 260.88 Kb/s.

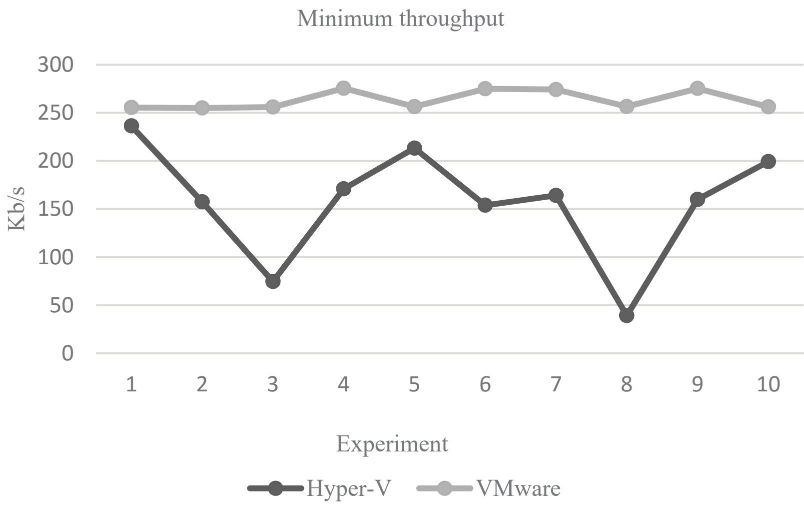

Figure 8. Minimal throughput in virtual networks

In the fourth test we have estimated bandwidth in virtual network using iperf traffic generator. The goal was to estimate the throughput in both directions (client-server and server-client). Results of experiments are presented in the Table 1. As compared to throughput of a server and a client server to a client in both environments have a difference in the throughput of 0.05 Mb/s.

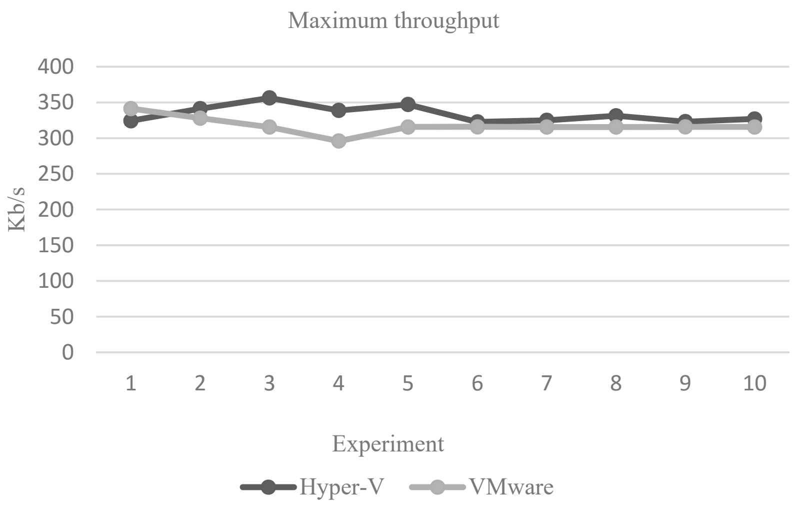

Figure 9. Maximal throughput in virtual networks

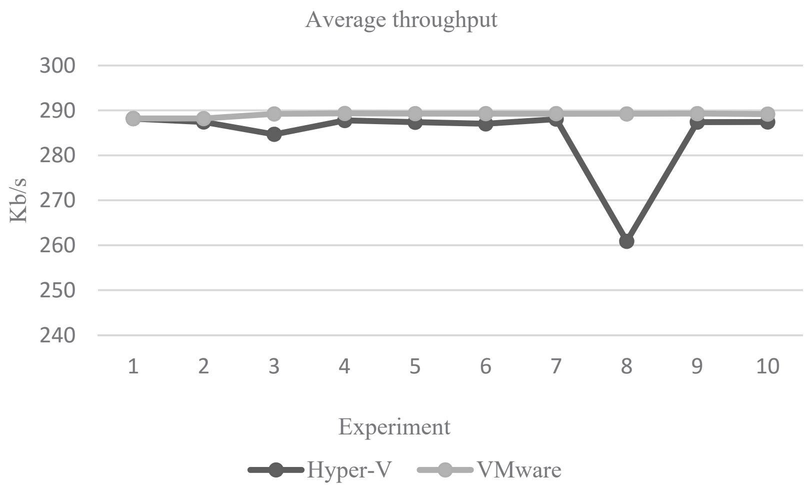

Figure 10. Average throughput in virtual networks

The next experiment consisted of assessing jitter and provide information about the number of packets lost by iperf. For Hyper-V smallest jitter was 0.130 ms, while the largest 0.5 ms. For VMware smallest jitter was 7.429 ms, and the largest 8.802 ms. In all of the experiments packet loss does not exceed 1%. Results of jitter studies are presented in the Table 2.

Table 1. Comparison of throughput for one and two-directions scenarios

Scenario	Hyper-V		Vmware	
throughput by iperf	Min	Max	Min	Max
	2.42 Mb/s	2.43 Mb/s	2.42 Mb/s	2.43 Mb/s
throughput in both directions by iperf [client – server]	Min	Max	Min	Max
	2.37 Mb/s	2.43 Mb/s	2.42 Mb/s	2.43 Mb/s
throughput in both directions by iperf [server – client]	Min	Max	Min	Max
	2.28 Mb/s	2.42 Mb/s	2.37 Mb/s	2.38 Mb/s

Table 2. Results of jitter studies

Experiment	Hyper-V	VMware
1.	0.5 ms	8.29 ms
2.	0.13 ms	8.802 ms
3.	0.397 ms	7.429 ms
4.	0.172 ms	8.248 ms
5.	0.194 ms	8.042 ms

6 Conclusion

Analyzing results of all performed experiments, we may conclude that Hyper-V virtualization platform offers better efficiency and better Quality of Service for virtual computer network. Most important conclusion is, that Hyper-V based virtual network behaves more stable and predictable with much lower delay and throughput fluctuation. We may recommend this platform rather than VMware, for virtualization of small virtual computer networks, built on one physical machine.

Quite important problem observed during experiments was the weak recurrence of the obtained results in both environments. Another problem was physical server performance, which is very well illustrated in performance tests for the delay and throughput. Probably, the source of those problems was too great load on the physical server processor, during which the hypervisor ESXi had to decide which virtual machine first should use the computational power of the processor. This leads to a situation where all traffic is handled by only one logical processor, while the other ones are asleep. VMware and Microsoft also points out, that the results of performance tests may be disrupted by the power plan of virtual machines running Windows 8.1. By default, Windows sets balanced mode of performance and energy consumption. Despite the efforts to define stable and identical test conditions, it turned out that the individual measurements for the delay and bandwidth can be quite significantly different in each of the trials. Moreover, we should observe the directions of development of these technologies, because similar situation may take place in the future, as in the case of cloud computing, which was a novelty two years ago, and today is an essential tool for facilitating daily life.

References

1. Brocade Communication Systems, 2016, *Brocade 5400 vRouter (formerly Vyatta 5400 vRouter) Data Sheet*, Brocade Communication Systems, Inc., San Jose, CA.
2. Ka Ching C., 2015, *Integration of Physical Equipment and Simulators for On-campus and Online Delivery of Practical Networking Labs*, Technical Report CSIT, La Trobe University, Bendigo, Australia.
3. Ka Ching C., Martin M., 2012, *An Integrated Virtual and Physical Network Infrastructure for a Networking Laboratory*, The 7th International Conference on Computer Science & Education.
4. Cisco Systems, 2016, *Cisco Nexus 1000V Switch for VMware vSphere. Cisco Systems*, Inc., San Jose, CA.
5. Cisco Systems, 2016, *Cisco Cloud Services Router 1000V 3.16 Series Data Sheet*, Cisco Systems, Inc., San Jose, CA.
6. Cisco Systems, 2016, *Q&A Cisco Cloud Services Router 1000V with Cisco IOS-XE® Software Release 3.10.*, Cisco Systems, Inc., San Jose, CA.
7. Cisco, 2016, *White Paper Flexible Workload Mobility and Server Placement with VXLAN Routing on Cisco CSR 1000V and Cisco Nexus 1000V.*

8. Gupta A., Farhan M., Chowdhury P., Tornatore M, Mukherjee B., 2015, *Joint Virtual Network Function Placement and Routing of Traffic in Operator Networks*, Department of Computer Science, University of California Davis.
9. Khan A., Zugenmaier A., Jurca D., Kellerer W., 2012, *Network Virtualization: A Hypervisor For the internet?,* IEEE Communications Magazine, January 2012 , pp 136-143.
10. Ruay-Shiung C. Chia-Ming W., 2010, *Green Virtual Networks for Cloud Computing*, Dept. of Computer Science and Information Engineering National Dong Hwa University Hualien, Taiwan, R. O. C.
11. Siraj M. R., Hidell M., Sjödin P., 2012, *PC-based Router Virtualization with Hardware Support*, 26th IEEE International Conference on Advanced Information Networking and Applications.