Comparison of type-2 hypervisor performance on the example of VirtualBox, VMware Workstation player and MS Hyper-V

Borislav Đorđević, *Member IEEE*, Iva Jovičić, Nenad Kraljević and Valentina Timčenko, *Member IEEE*

***Abstract* – This paper presents a comparison of the performances of type-2 hypervisors, on the example of desktop virtualization applications, which include VirtualBox, VMware Workstation Player, and MS Hyper-V. The qualities of all three tested hypervisors, from many aspects of performance, were tested through the performance of the files system. Tests were performed under the same conditions and the same testing methods, using the Filebench program. CentOS 7 was used as the guest operating system. The hypervisor's performances were compared taking into consideration the tests performed for the system with one, two, and three virtual machines in operation. Hypotheses about expected behavior were set, and then they were validated through the obtained results using the Filebench program.**

***Index Terms* – VirtualBox; VMware Workstation; MS Hyper-V; CentOS; hypervisor; virtual machines.**

I. Introduction

Virtualization as a concept is increasingly used and conquers new spaces. It has become a part of everyday life for the simple reason that information technology are all around us. As these technologies are increasingly present in modern life and are constantly advancing, virtualization has taken its place in this development. The main advantages obtained by applying virtualization can be seen in the reduction of costs of IT equipment, electricity and storage space for this equipment. The concept itself provides high security and resistance to failures, and makes administration easier. The choice of virtualization methods and techniques depends on the specific situation and the needs of the end user. This is due to the fact that some virtualization techniques achieve greater flexibility in operation, while others achieve better performance or security. The most commonly used virtualization techniques are: virtualization of hardware, software, data, memory, storage space, virtualization of network infrastructure and virtualization of desktop computers. [1]

Borislav Đorđević – Institute Mihailo Pupin, Volgina 15 Belgrade, Serbia, (borislav.djordjevic@pupin.rs)
Iva Jovičić – VISER, School of Electrical and Computer Engineering of Applied Studies, Belgrade, Serbia, (iva.jovicic@yahoo.com)
Nenad Kraljević – VISER, School of Electrical and Computer Engineering of Applied Studies, Belgrade, Serbia, (nenadk@gs.viser.edu.rs)
Valentina Timčenko - Institute Mihailo Pupin, Volgina 15 Belgrade, Serbia, (valentina.timcenko@pupin.rs)

Virtualization involves the encapsulation and abstraction of computer components, so these components can be used in a way that suits a particular application. Hardware virtualization involves usage of hypervisor, a software layer, which is an intermediary between the hardware and the guest operating system in a virtual machine. This is a simulated environment that could have characteristics equal to the physical environment. The hypervisor can be native, or type-1, which runs directly on the hardware, or hosted, type-2, which runs on the operating system. Examples of this type of hypervisor are VirtualBox 6.1 and VMware Workstation 16 player which have been tested for the purpose of this paper. MS Hyper-V is a bare metal, type-1 hypervisor. However, when activated as a roll, in this case in Windows 10 Pro, it behaves as a type-2 hypervisor. This version of MS Hyper-V was used in the testing process for this paper. [2]

Other than above described classification, hardware virtualization also depends on whether full, partial, or paravirtualization is selected. Full virtualization (Figure 1), topic of this research, represents a simulation of complete hardware, so guest operating systems can be installed and ran without problems. The guest operating system is separated from the physical layer of the host, the hypervisor layer. The advantage of this method of virtualization is increased security and scalability, as well as system flexibility. This solution is the easiest to use, but the performance is slightly lower. [3,4]

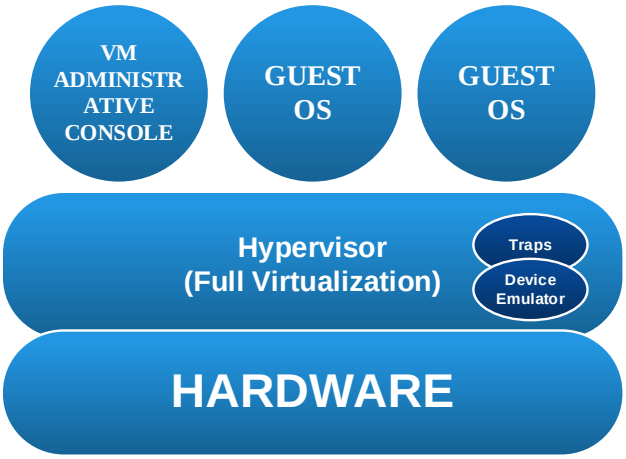

Figure 1. Full virtualization

 ISBN 978-86-7466-930-3

II. Related work, objective and motivation

The focus of this research are the performance characteristics of hypervisors as one of the basic factors in achieving service quality. The problem itself, which can be found in the literature as well as by using these systems in practice, can be viewed from several angles. There are many papers and discussions that use different methodology and approach to evaluating the performance of virtual platforms. The most common is a comparative performance analysis of VMware, Xen, MS Hyper-V and other hypervisors by using various benchmark tools such as HD Tune Pro, ATTO, Filebench, Bonnie ++, etc. Virtualization is a great solution in both desktop and server versions. The needs for virtual platforms for personal use are growing. The main contribution is the mathematical modeling of the file system performance in hypervisor-based virtualization. The modeling of complex virtual environment includes many factors, and modeling expects there is no single winner hypervisor. Similar mathematical model, we used in this paper and most our references [3], [5], [10], [11]. Our model is open for enhancing. We think we are different from related work by our methodology. Its essence is a mathematical model, apply it on a particular case study, and then provide the interpretation of practical results as a validation of model. Using by large number of case studies, we recommend the creation of Knowledge Data Base (KDB) related to the file system performance in virtual environment. Case study in this paper include the performance comparison of three hypervisors in the desktop version, namely VirtualBox, Vmware Workstation player and Hyper-V, in fair-play conditions. This implies identical hardware, the same virtual machines, and an identical version of the guest operating system, which in this case was CentOS 7, an operating system from the Linux distribution family. As VirtualBox and Vmware Workstation player use full virtualization, and MS Hyper-V and paravirtualization as well, the effects of full virtualization for three different hypervisors were examined, using the Filebench benchmark program with four different workloads. Hypotheses about expected behavior are set, followed by a mathematical model for workloads and a hypervisor environment. Performance was measured and the obtained results were interpreted on the basis of models and hypotheses. This paper has some similarities with reference at the serial number five in our literature. The results have the similarities and differences, because the hardware and many other factors are quite different, but we think that both papers are interesting and useful cases of study.

III. VirtualBox, VMware Workstation player and MS Hyper-V

Oracle's VirtualBox is a very powerful program for virtualizing 32-bit and 64-bit operating systems, on computers with Intel or AMD processors. VirtualBox is the only professional solution that is available for free as open source software under the terms of *"GNU"* version 2. This software runs on Windows, Linux, Mac and Solaris operating systems. The technical requirements for running this software are:

- 32-bit or 64-bit operating system with AMD or Intel processor,
- 512MB or more RAM (depending on the number and type of operating systems being virtualized. The RAM memory space allocated to the virtual machine environment can go up to half capacity, as the software itself will not allow more than half of the base system's RAM.
- available hard disk space for the virtual machine environment (recommended size is a minimum of 8GB).

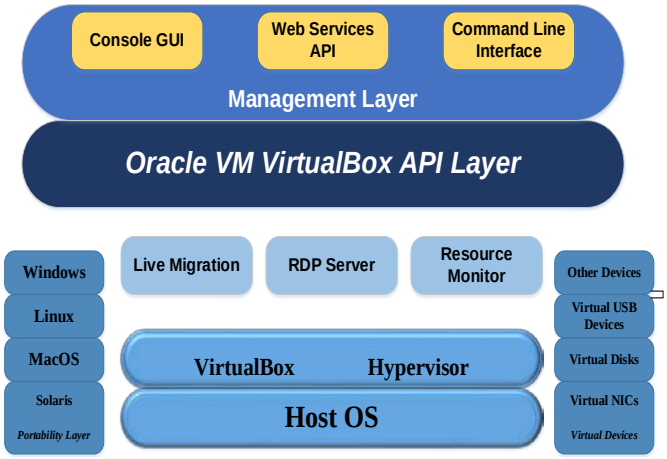

Figure 2. VirtualBox architecture

VirtualBox allows quick and easy data sharing between virtual machines (Figure 2). As VirtualBox is open source, or rather it's software, it tends to solve problems quickly and can be upgraded with new features. Since version 6.1, which is used in this paper, the ability to export and import virtual machines to Oracle Cloud has been added. Software virtualization has been deprecated since this version, and VirtualBox 6.1 uses only hardware-assisted virtualization. [5,6]

VMware® is considered one of the largest manufacturers of software virtualization. The solutions of this company occupy over 70% of the market share in this area, primarily due to the quality of products and the availability of technical support. VMware® has been acquired since 2004 and became part of the EMC Corporation. VMware Workstation 16 player is a software package that runs on standard x86-based hardware with 64-bit Intel and AMD processors and on 64-bit Linux and Windows operating systems. [7,8]

VMware Workstation 16 player (Figure 3) can run existing virtual machines and create their own virtual machines. It uses the same virtualization core as VMware Workstation Pro, a similar multi-featured, non-free program.

 ISBN 978-86-7466-930-3

VMware Workstation 16 player is available for personal non-commercial use (free), for distribution or other use by written agreement. The technical requirements for running this software are:

- 64-bit operating system from the Windows or Linux family with AMD or Intel processor,
- 2GB or more RAM (recommended 4GB or more),
- 1.3GHz or higher core speed,
- available hard disk space.

Figure 3. VMware Workstation player architecture

Microsoft is one of the leading companies in information technology. At the server level, virtualization has become the standard, but interest in virtualization has also emerged among users for personal use. Microsoft occupies about 15% of the market with MS Hyper-V virtual platform. With the release of Windows 8 in 2012, MS Hyper-V became an integral part of its Enterprise, Education and Pro editions. MS Hyper-V is a type-1 hypervisor-based system for x86-64 operating system architectures. It is activated in the Windows operating system as a roll, just like any other service in the Microsoft family. There are some MS Hyper-V features that work differently in Windows OS and Windows server. The memory management model is different for MS Hyper-V, where MS Hyper-V manages memory on the server assuming only virtual machines run on the server, and in the Windows operating system, it is managed with the expectation that most client machines run software on the host in addition to virtual machines. [9-11]

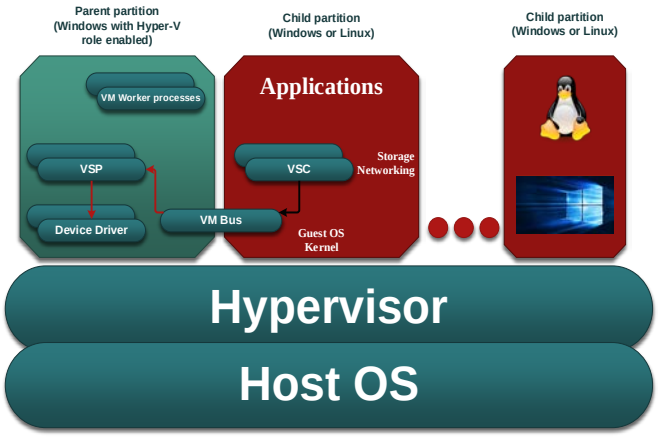

Figure 4. MS Hyper-V architecture

MS Hyper-V (Figure 4) supports virtual machine isolation and uses partitions in which guest operating systems will run.

IV. Hypotheses about expected behavior

As the Type-2 hypervisor running under the guest operating system was used, we can point out that each workload generates typical random and sequential data read times, as well as random and sequential data write times. Each workload is defined by the access time for file systems. Workload represents the total time to complete all operations, the time required to complete all operations related to directories, metadata, file blocks, free lists, house-keeping and journaling operations in the file system. There are five components in a virtual environment that have an impact on workload time (Tw – Time Workload):

$$Tw = f\,(Bn, gFS, VHw\text{-}pr, Hp\text{-}pr, hFS) \qquad (1)$$

The first and second components, Bn (Benchmark) and gFS (guest OS file system) are exactly the same for VMware Workstation player, Oracle VirtualBox and MS Hyper-V. The analysis focuses on the interaction between the reference values and the guest operating system. Because the test environment relies on the use of an identical benchmark, identical virtual machines, and ext4 as the guest file system, these components are expected to have an identical effect on Tw. Processing time for full hardware virtualization is the third component of VHw-pr (virtual hardware processing). Each hypervisor uses its own solution for full hardware virtualization, so the performance will be different too. The fourth component, Hp-pr (hypervisor processing), represents the time it takes for the hypervisor to receive requests from the virtual hardware and forward them to the host drivers. In particular, guest FS requests (guestOS-FS) are forwarded to host FS (hostOS-FS). All of these hypervisors, VMware Workstation player, Oracle VirtualBox and MS Hyper-V, generate different hypervisor processing times. The fifth component, hFS (host OS file system), represents the

 ISBN 978-86-7466-930-3

processing time of the host OS file system. All hypervisors have MS NTFS as hostOS-FS, so this component is expected to have similar processing times for all hypervisors. The dominant influence of the third and fourth components of formula (1) is expected as the tests are focused on the performance of natively virtualized guests (complete hardware virtualization). [12]

V. Test configuration and benchmark application

The prerequisite for quality and adequate testing is the application of one hardware configuration, the same operating system, then the selection of a quality benchmark program and the same measurement methodology for all testing procedures. Testing was done on a personal computer whose characteristics can be seen in Table I, while the characteristics of the disk are given in Table II. CentOS 7 from the Linux distribution family is installed as a guest operating system.

TABLE I - Test environment/PC

Components	Characteristics
Processor	Intel Core i5-4590S 3GHz
Memory	8GB DDR3
Cache	6MB L3
Hard drive	Seagate Barracuda 7200.12
Operating system	Windows 10 Pro, 64-bitni

VirtualBox 6.1 virtualization platforms, VMware Workstation 16 and MS Hyper-V, a version for Windows 10 Pro, are installed or activated on the hard drive, where the tests were done. The hard drive was also used to install virtual machines.

TABLE II - Test environment/Hard Disk

	Seagate Barracuda 7200.12
Model Number	Seagate ST3500418AS
Capacity	500GB
Interface	Serial-ATA/300
External Transfer Rate	3.0Gb/s
Max Sustained	300MB/s
Cash	16MB
Average Latency	4.17ms
Spindle Speed	7200rpm
Average Seek Time Read	8.5ms
Average Seek Time Write	8.5ms

All tests were performed using the benchmark program Filebench 1.4.9.-1. This program is designed to measure storage space and performance of file system. Filebench is capable of generating several types of workloads, it can simulate environments when using certain services such as mail server, fileserver, web server, etc. [13, 14]

VI. Testing and results

This paper shows a comparison of the performance of virtual platforms for personal use with their capabilities. Disk performance and data flow were measured. To make testing meaningful, all virtual machines are created exactly the same and have the same characteristics (Table III).

TABLE III - Virtual machine parameters

Components	Characteristics
Virtual processor	1
Memory	2GB
Virtual hard drive	60GB
Operating system	CentOS 7

During the testing, modified base code files were used, such as *varmail.f*, *fileserver.f*, *webserver.f* and *randomfileaccess.f*, which test the web, mail and file server. The appearance of the set parameters of the benchmark program can be seen in Table IV. To achieve the most realistic results, each test lasted 120 seconds. Each test was repeated ten times, and the obtained results were expressed as the average value of these tests.

TABLE IV – Base code parameters of **F* file

	Varmail	Web server	File server	Random file access
nthtreads	16	100	50	5
nfiles	1000	1000	10000	10000
meandir widht	1000000	20	20	20
meanfile size	16k	16k	128k	random

Tests were conducted by first installing VirtualBox and then creating a virtual machine in this program. It was tested, and then this procedure was repeated when two and afterwards three virtual machines were created. The testing system was exactly the same when the virtual machines in the VMware Workstation 16 player application were tested. Of course, before testing on this platform, previous virtual machines and VirtualBox applications were uninstalled. At the end of the testing, the MS Hyper-V roll was activated and tests were performed, as in the previous two cases. In that way, fair-play conditions were created for all three virtual platforms. The results of the Fileserver workload test can be seen in Table V and Figure 5.

TABLE V - Fileserver benchmark results

 ISBN 978-86-7466-930-3

Fileserver	1VM (MB/s)	2VM (MB/s)	3VM (MB/s)
VMware	74,58	64,33	56,41
VirtualBox	42,03	35,16	29,76
MS Hyper-V	33,34	28,12	21,53

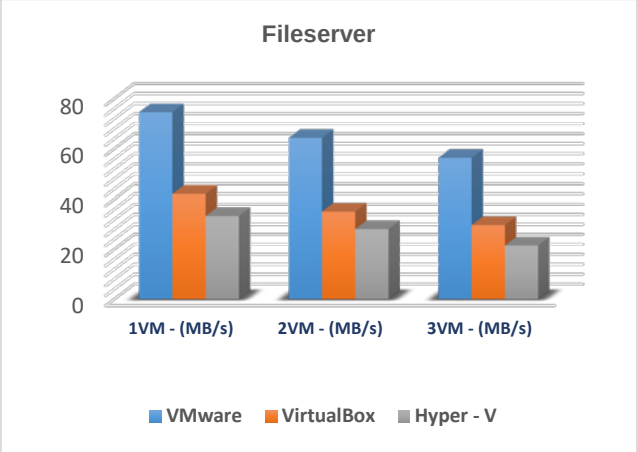

Figure 5. Fileserver test results

The results of testing other workloads are shown in Tables VI, VII and VIII, as well as graphical representations in Figures 6, 7 and 8.

For the *"Fileserver"* workload, we note that VMware is by far the best, while VirtualBox is better than Hyper-V. In a complex workload such as Fileserver with sequential and random write components, the FS cache effect on guest and hostOS is significant, so VMware wins convincingly primarily because of the 3rd component of formula (1) and the best cooperation with FS caching.

TABLE VI - VARMAIL BENCHMARK RESULTS

Varmail	1VM (MB/s)	2VM (MB/s)	3VM (MB/s)
VMware	44,58	42,62	39,87
VirtualBox	18,46	17,83	16,62
MS Hyper-V	14,06	12,77	10,11

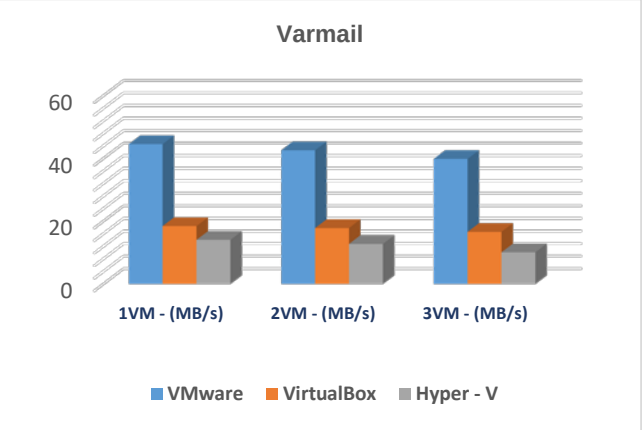

Figure 6. Varmail test results

For the *"Varmail"* workload, we note that VMware is by far the best, while VirtualBox is again slightly better than MS Hyper-V. In the Varmail workload, in addition to random reading, there are also synchronous components of random write, the impact of FS caching is small, so VMware and then VirtualBox obtain wins, primarily because of the 3rd and 4th components of formula (1).

TABLE VII - WEBSERVER BENCHMARK RESULTS

Webserver	1VM (MB/s)	2VM (MB/s)	3VM (MB/s)
VMware	84,68	81,04	77,68
VirtualBox	47,73	42,26	37,86
MS Hyper-V	80,92	76,61	71,92

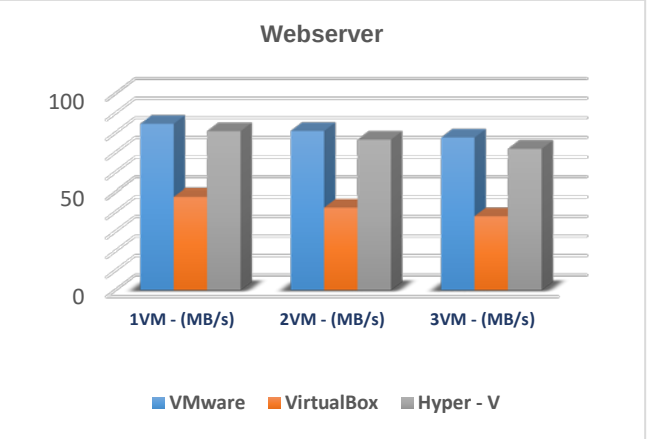

Figure 7. Webserver test results

For the *"Webserver"* workload, we note that VMware is slightly better than MS Hyper-V, and VirtualBox is significantly weaker. In the Webserver workload, which has both random read components and very few random write components, there is less influence of FS caching, so VMware and MS Hyper-V did better than VirtualBox, and win primarily because of the 3rd and 4th and the components of formula (1) and the cache effect in random reading, the dominant workload in Webserver environment.

 ISBN 978-86-7466-930-3

TABLE VIII – RANDOMFILEACCESS BENCHMARK RESULTS

Randomfile access	1VM (MB/s)	2VM (MB/s)	3VM (MB/s)
VMware	5141,56	5034,02	5008,76
VirtualBox	2595,11	2546,56	2517,36
MS Hyper-V	4888,85	4816,64	4790,44

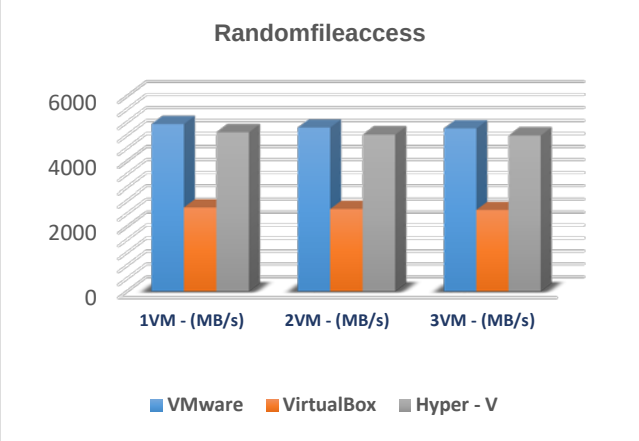

Figure 8. Randomfileaccess test results

For the *"Randomfileaccess"* workload, we can see that Hyper-V is slightly weaker than VMware, and VirtualBox is significantly weaker again. In the Randomfileaccess workload, which has both random read components and a lot of asynchronous random write components, there is a solid impact of FS caching primarily for random writing, so VMware and MS Hyper-V fared better than VirtualBox, primarily because 3rd and 4th components of formula (1) and solid cache effect in random entry.

VII. CONCLUSION

When it comes to virtualization, it should be noted that it brings major changes in the information technology and computer industry, primarily in reducing investment in infrastructure and saving electricity consumption. In area of personal computers, virtualization has made great strides by bringing a large number of software for this purpose. Some of them are also presented in this paper. The advantage of using these applications is the ease of installation and usage, and the fact that the platforms tested in this paper are completely free. For this case study, VMware is the absolute winner in all workloads. We believe that the differences is made by the 3rd and 4th components of formula (1), as well as by the powerful hypervisor usage of the FS cache effect. For workloads with many sequential features and weak cache effect (Fileserver and Mailserver) VirtualBox is better than MS Hyper-V, and for workloads with random dominance and solid cache effect (Webserver and Randomfileaccess) MS Hyper-V is better than VirtualBox. Future work on this topic will focus on testing other virtual platforms in the field of desktop computers, as well as the use and testing of various operating systems that are applied in practice.

ACKNOWLEDGEMENT

The work presented in this paper has partially been funded by the Ministry of Education, Science and Technological Development of the Republic of Serbia.

LITERATURE:

[1] Matthew Portnoy: "Virtualization Essentials", 2016.
[2] Hypervisor Type-2. Online: https://www.ibm.com/developerworks /cloud/library/cl-hypervisor compare/, 2021.
[3] B.Đorđević, V.Timčenko, N.Kraljević, N.Davidović, "*File system performance in full hardware virtualization with ESXi and Xen hypervisors*", 18th International Symposium INFOTEH-JAHORINA, 2019.
[4] Full Virtualization. Online: https://www.sciencedirect.com/topics /computer-science/full-virtualization, 2021.
[5] Dejana T. Vojnak, Borislav S. Đorđević, Valentina V. Timčenko, Svetlana M. Štrbac: "*Performance Comparison of the type-2 hypervisor VirtualBox and VMware Workstation player*", 27th Telecommunication Forum, 2019.
[6] VirtualBox, Online: https://www.virtualbox.org/, 2021.
[7] Perfomance Evaluation of VMware and VirtualBox Online: https://pdfs.semanticscholar.org, pdf-2021.
[8] VMware, Online: https://www.vmware.com/product/workstation-player.html, 2021.
[9] Microsoft, Online: https://www.microsoft.com/, 2021.
[10] Jadran Torbić, Ivan Stanković, Borislav S. Đorđević, Valentina Timčenko: "*Hyper-V and ESXi hypervisors comparison in Windows Server 12 virtual environment*", 17th International Symposium INFOTEH-JAHORINA, 2018.
[11] B. Đorđević, V. Timčenko, N. Kraljević, N. Maček: "*File System Performance Comparison in Full Hardware Virtualization with ESXi, KVM, Hyper-V and Xen Hypervisors*", Advances in Electrical and Computer Engineering, vol.21, iss.1, 2021.
[12] M. Polenov, V. Guzik, V. Lukyanov: "*Hypervisors comparison and their performance*", Computer Science Online Conference, 2018.
[13] Filebench, Online: https://github.com/filebench/, 2021.
[14] Christopher Stracheyje, "*Time Sharing in Large Fast Computers*". Online:https://archive.org/details/large/fast/computers/page/n5/mode /2up, 2021.

 ISBN 978-86-7466-930-3