SINTEZA 2024
INTERNATIONAL SCIENTIFIC CONFERENCE ON INFORMATION TECHNOLOGY, COMPUTER SCIENCE, AND DATA SCIENCE

COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE SESSION

FILE SYSTEM PERFORMANCE COMPARISON WITH THE HYPERVISORS ESXI AND XEN

Borislav Đorđević*,
[0000-0002-6145-4490]
Kristina Janjić,
[0009-0005-0750-6105]
Nenad Kraljević
[0009-0008-7684-5444]

School of Electrical and Computer Engineering,
Belgrade, Serbia

Abstract:

This paper contains a file system performance comparison among type-1 Linux-based hypervisors, with ESXi and Xen chosen as representative examples of such hypervisors. At first glance, both hypervisors may seem to share a similar Linux-based architecture, but upon deeper examination, notable differences emerge. We have used a benchmark called Filebench for experimental measures of file system performance. Filebench is chosen because of its high level of flexibility and adaptability, which enable the emulation of real applications in typical server environments. This paper comprises a mathematical model of the type-1 hypervisor environment, followed by a real file system experiment that serves as a specific case study. The model is employed to interpret the file system performance results obtained from the experimental measures. The guest operating system utilized in our experiments was CentOS 9, selected as a typical representative of Linux distributions. We conducted experimental tests with one, two, and three virtual machines operating simultaneously.

Keywords:

Hypervisor, ESXi, Xen, Centos 9, Filebench.

INTRODUCTION

The virtualization technology allows multiple operating systems to operate concurrently on a single hardware platform, improving system availability and dependability while making better use of resources. In the context of information technology progress, virtualization takes a leading role in IT innovation, transforming the handling and utilization of information and resources. Virtualization simplifies the establishment of virtual entities such as computers, servers, and other resources, resulting in improved resource management, cost-effectiveness, streamlined system administration, and enhanced scalability. Virtualization guarantees the allocation of virtual machines with accurately specified CPU features, RAM memory, and storage space, ensuring optimal utilization of hardware [1].

Despite its benefits, virtualization also brings challenges such as management complexity, vulnerabilities in the security area, licensing costs, and the risk of failure of hypervisors or physical servers. However, despite these challenges, the advantages of virtualization typically surpass its drawbacks, making it a crucial component of modern IT infrastructure.

Correspondence:
Borislav Đorđević

e-mail:
borislav.djordjevic@viser.edu.rs

DOI: 10.15308/Sinteza-2024-10-17

Despite its benefits, virtualization also brings challenges such as management complexity, vulnerabilities in the security area, licensing costs, and the risk of failure of hypervisors or physical servers. However, despite these challenges, the advantages of virtualization typically surpass its drawbacks, making it a crucial component of modern IT infrastructure.

Virtualization can be classified into different types, including hardware, desktop, application, network, memory, and storage virtualization. This paper is based on full hardware virtualization. The primary feature of hardware virtualization is its ability to enable virtual machines to operate as fully independent virtual computers, even though they share the same underlying physical hardware. Hypervisors play a crucial role in achieving hardware virtualization by abstracting hardware from the operating system, enabling multiple operating systems to run concurrently on the same hardware. There are two main classes of hypervisors: type-1 hypervisors, also known as bare-metal hypervisors or native hypervisors, which execute directly on the hardware; and type-2 hypervisors, referred to as hosted hypervisors, which operate within the host operating system.

This paper employs two typical bare-metal hypervisors, ESXi and Xen, with the main objective of a comparative analysis of file system performance when accessing virtual machines through these hypervisors.

2. RESEARCH WORK, MOTIVATION AND GOAL

Many scientific papers in the virtualization area explore various methodologies for evaluating different virtual environments' performance. These papers, as usual, involve well-designed experiments and utilize established benchmarks (like Postmark, Bonnie++, AS SSD, ATTO, FIO, Filebench). However, many of these experiments do not include mathematical modelling in their assessment of different virtual environments [1] [2] [3] [4] [5] [6] [7] [8] [9] [10]. References [11] [12] [13] [14], similar to this paper, incorporate a mathematical model, although they focus on different hypervisors and hardware configurations.

The main contribution of this research paper lies in the comprehensive mathematical model crafted to analyse file system performance within a virtual environment employing Linux-based hypervisors (type-1). The model encompasses a large range of input parameters and is designed for potential future enhancements. In this paper, we used our own methodology, which commences with the creation of a mathematical model and proceeds to experimental testing, as a unique case study. This approach allows for a unique perspective on evaluating virtual environment file system performance by utilizing the model to interpret the experiment's results.

We conducted experiments using ESXi and Xen as Linux-based hypervisors (type-1), both of which are based on full hardware virtualization technology. Hypervisors were tested under identical (fair-play) hardware conditions with CentOS 9 as the guest operating system and using the XFS as guest file system. The experiments utilized the Filebench benchmark tool, covering four different workloads (Fileserver, Webserver, Varmail, and Random-File-Access). Our validation included applying a mathematical model to analyse and interpret the experimental results.

3. ESXI AND XEN

ESXi (Elastic Sky X integrated) is a type 1 hypervisor that is installed directly on the hardware rather than on the operating system, which means it integrates components of the operating system within itself. The ESXi architecture encompasses the underlying operating system, with a kernel called VMkernel and processes running above it. VMkernel serves as a kernel of virtualization created by VMware to oversee and execute all applications, agents, and virtual machines. VMware ESXi employs full virtualization, allowing virtual machines to operate on unmodified operating systems. VMware ESXi includes its own hardware virtualization drivers to provide a communication layer between virtual machines and physical hardware.

In the architecture of the Xen platform [15], the central point is the Xen hypervisor positioned above the physical hardware, and several domains which represent the virtual machines located above the hypervisor. The key components of Xen architecture that collaborate to deliver different virtual solutions are: Xen hypervisor, Domain 0 and Domain U. Xen hypervisor is a core component with the responsibility of managing hardware resources such as CPU cores, RAM memory, and I/O resources for a few concurrent guest operating systems. Dom0 represents a modified Linux operating system with Xen hypervisor as the kernel, with specialized privileges for accessing physical I/O hardware resources and managing the Xen virtual machines (DomU guests). Dom0 is responsible for managing the Xen hypervisor and drivers for physical hardware devices. DomU represents unprivileged guests without the possibility of direct hardware access, which can be launched either as PV guests (modified OS using paravirtualization) or as HVM guests (unmodified hardware-assisted OS).

4. MATHEMATICAL MODEL AND HYPOTHESES ABOUT EXPECTED BEHAVIOUR

The workload time for a hypervisor-based virtual environment, Tw, highlights at least five components that significantly influence it, Equation 1:

$$Tw = f\,(Bn, gFS, VH\text{-}proc, Hyp\text{-}proc, hFS)$$

Equation 1. Time in virtual environment

The first component, *Bn*, denotes the benchmark processing time. The second component, *gFS*, represents the processing time of the guest file system. The guest file system is tightly coupled with two components: the kernel of the guest OS and the FS cache mechanism of the guest OS. *Bn* and *gFS* show similar features for each hypervisor from our experiments. The similarity of effects for the Bn and gFS components is due to the use of the same benchmarks, identical benchmark parameters, the same virtual machines, and the same guest file system, XFS.

The third component, *VH-proc*, denotes the processing time of the virtual hardware, such as the virtual disk drive/drivers. In this third component, *VH-proc*, both hypervisors from our experiment can differ remarkably. ESXi employs only full hardware virtualization, and ESXi has its own solution for full hardware virtualization. Xen employs full hardware virtualization, too. Xen implements QEMU-based open-source software for full hardware virtualization. Xen uses QEMU-based open-source solutions for full hardware virtualization. Due to this, both hypervisors, Xen and ESXi differ significantly in the context of the *VH-proc* component. We highlight that *VH-proc* is coupled with file system caching on the guest and host OS sides.

The fourth component, *Hyp-proc*, represents the time needed for hypervisor processing. This is the time required for the hypervisor to take requests from the virtual disk drivers and then forward them to the host operating system, precisely to the host file system for the virtual machine image file. For our case: XenServer has the original Xen hypervisor, while ESXi has the original VMware hypervisor called *VMkernel*. Both hypervisors are different, so hypervisors must have different performance anyway.

The fifth component, *hFS*, represents the processing time of the host file system. The host file system is coupled with two components: the kernel of the host operating system and the file system caching of the host OS. And here, we expect significant differences between hypervisors. Xen can consume the two most common candidates, namely ext4 and XFS with or without the LVM option, in our case, it was ext4. ESXi uses a totally different cluster-based filesystem, VMFS. So, there are big differences between *hFS*.

If we examine the host operating system (a mandatory part of a virtual environment), we can conclude the following: Both hypervisors have the same architecture (Linux-based), but in detail, they can be very different. Each hypervisor has its own Linux distribution: XenServer employs Xen-adopted Linux distributions and, ESXi consumes VMware adopted Linux distributions. That way, host operating systems have different versions of the following: host kernels, physical disk drivers, host file systems, and different OS systems and graphical environments. Although both hypervisors are Linux-based, they vary significantly in the context of host operating systems and host file systems.

In this paper, we interpret the performance of different hypervisors by using our own mathematical model, considering the identified differences and similarities among the hypervisors.

5. TEST CONFIGURATION AND BENCHMARK APPLICATION

We highlight the fair-play performance examination by using identical hardware, virtual machines, operating systems, measurement methodologies, and a benchmark program. We ensured this fair-play by configuring adequate hardware, selecting consistent operating systems (on the guest and host sides), and employing a single benchmark program for all testing phases. The virtual platforms used were VMware ESXi 8.0 and Xen Citrix Hypervisor 8.2.1, while the experiment was carried out on an HP server with CentOS Stream 9 as the guest operating system. HP server has the following configuration:

- **CPU:** Intel® Xeon® Silver 4116 CPU @ 2.10GHz
- **RAM:** 32GB DDR4 2400 MHz
- **Hard disk:** 2x HPE 480GB SATA 6G RI SSF SSD RAID1, SATA 3, Sequential read up to 535 MB/s, Sequential write up to 495 MB/s

- **Host Operating Systems:** Xen Citrix Hypervisor 8.2.1 and VMware ESXi 8.0

All experimental tests were performed using the benchmark tool Filebench 1.4.9.1-3. Filebench enables the simulation of various real server environments through different workload definitions. Filebench provides detailed information on performance, including file read/write throughputs for different workloads [16]. For storage, we employ two identical hard drives as RAID-1, mounted on the server (HPE ProLiant BL460 Gen10). Testing was performed with both virtual environments while virtual machines were stored on the same RAID-1 physical disks. The virtual machine parameters are shown as follows:

- **Number of virtual CPU per VM:** 4
- **Virtual memory per VM:** 8GB
- **Virtual hard disk per VM:** 64GB (/dev/sda), 32GB /dev/sda1 root FS, 32GB /dev/sda2 testing FS (XFS)
- **Guest OS:** CentOS Stream 9

6. TESTING AND RESULTS

In this paper, the main objective is to measure the file system performance of two different type-1 hypervisors, which would assist in choosing the most efficient hypervisor for specific requirements. We conducted assessments using various workloads like Fileserver, Webserver, Mailserver, and Random file access workloads. Initially, we evaluated performance using a single virtual machine and then repeated the assessment with two, three, and finally four virtual machines simultaneously running.

In general, for performance explanation, all features from Chapter 4 are very important. These are *VH-proc*, guest file systems, FS-pair, FS-cache-pair, *Hyp-proc*, virtual and physical disk drivers, and most components from the host OS, such as the kernel, host file system, and OS/graphical environments. We mention that most components are very different for ESXi and Xen.

The obtained results of the Fileserver workload test are presented in Figure 1.

The Fileserver workload contains all types of transfers: random read, random write, sequential read, and sequential write transfers. This workload includes a large number of I/O operations and a large data flow. For repeated read and asynchronous write transfers, file system caches may have a remarkable impact.

In the case of Fileserver workload, ESXi is solidly better than Xen. ESXi is better than Xen by 22-76%, it can be considered a big difference. By analysing the obtained Fileserver workload throughputs and the maximum disk speeds (the disk interface about 600MB/s and the maximum sequential speeds of SSD disks are about 500MB/s), on 1VM both hypervisors expose a higher throughput than the maximum disk speeds. On 2VMs Xen throughputs drop below the maximum disk speeds, while ESXi drops below the maximum disk speeds on 3VMs.

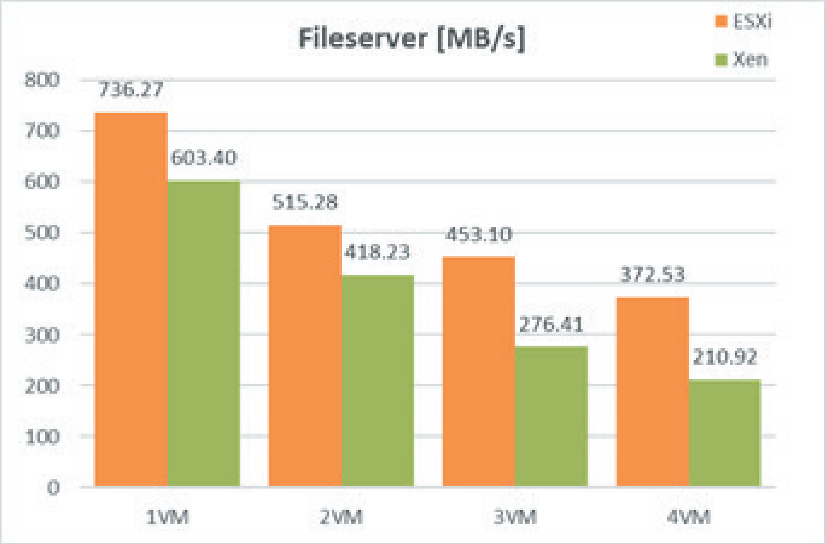

Figure 1. Fileserver test result chart.

High throughputs show that pairs of file system caches (guest/host) have a big impact, but also a lot of IO random/sequential requests processed through the virtual-physical drivers. High throughputs highlight that FS cache in pairs has a great impact, especially with ESXi, and is somewhat weaker with Xen.

This means that the FS-caches absorbed a large number of disk requests, but also that a solid number of IO operations were processed by virtual and physical disk drivers.

For the Fileserver workload, all features from Chapter 4 are important. We believe that the most impactful features are VH-proc with file systems in pairs and with the cache effects of this pair of file systems. Because a lot of cache misses, the virtual disk drivers of the guest operating systems and the physical disk drivers of the host operating systems are also important. In the context of fileserver workload features for random/sequential requests, we assume that ESXi shows the better combination of features: *VH-proc* with file system cache effects and a better combination of virtual and physical disk drivers than Xen.

The obtained results of the Mailserver workload test are presented in Figure 2.

The mailserver workload contains random read and synchronous random write components, involving a moderate number of I/O operations and data flow. Because of the dominance of such components, the efficiency of file system caches (both the guest and host operating systems) is minimal.

For the Mailserver workload, ESXi is better than Xen by about 1-5%, it can be considered a small difference. By analysing the obtained Mailserver workload throughputs and the maximum disk speeds (disk interface of about 600MB/s and maximum sequential speeds of SSD disks of about 500MB/s), on all virtual machines, both ESXi and Xen have solidly lower throughputs than the maximum disk speeds.

Low mail throughputs demonstrate that the influence of file system caches in pair is very small, meaning that most IO operations are processed to virtual and physical disk drivers.

For Mailserver workload, most features from Chapter 4 are very important. We believe that the most impactful features are *Hyp-proc* and *VH-proc* with file system pair, but for Mailserver workload with the minimal cache effects of file system cache-pair. Due to the weak file system cache influence, virtual drivers of the guest operating system and physical disk drivers of the host operating system dominate, for random read/random write performance. For the Mailserver workload, we assume that ESXi offers a better combination of *Hyp-proc* and *VH-proc,* with zero file cache effects, and a better combination of virtual and physical disk drivers than Xen.

The outcomes of the Webserver workload test are presented in Figure 3.

Webserver workload involves random read and small random write components, with a moderate number of I/O operations and data flow. The impact of file system caches on both the guest and host operating systems can be limited for random read components, except for repeated reading scenarios.

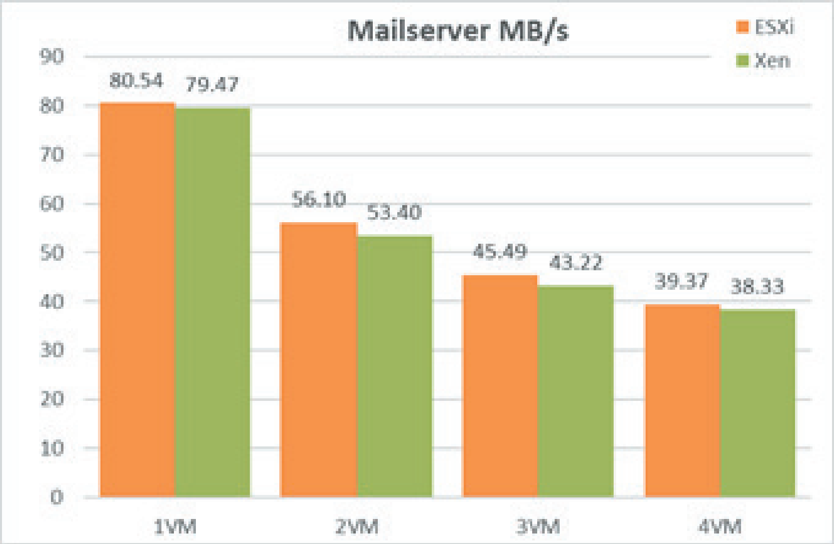

Figure 2. Mailserver test result chart.

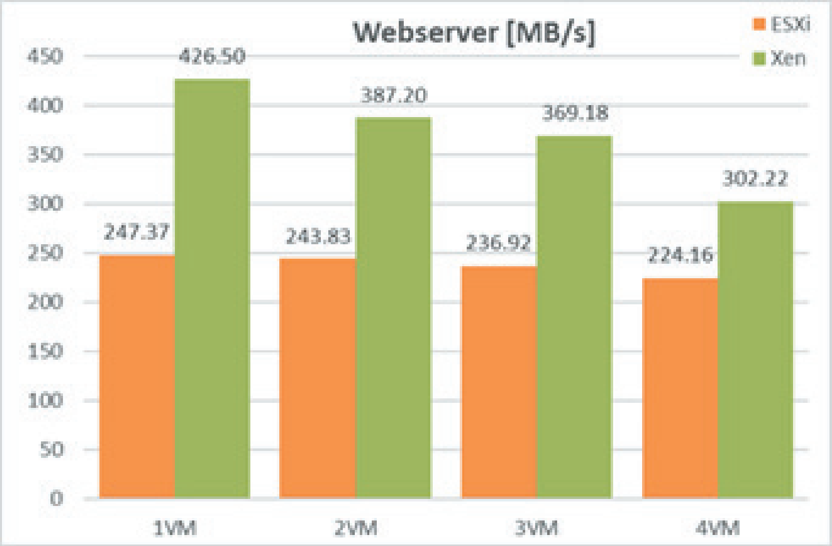

Figure 3. Webserver test result chart.

Figure 4. Randomfileaccess test result chart.

For the Web workload, Xen is remarkably better than ESXi, about 34-72%, which can be considered a remarkable difference. By analysing the obtained Webserver workload throughputs and the maximum disk speeds (disk interface of about 600MB/s and maximum sequential speeds of SSD disks of about 500MB/s), Xen and ESXi show good random-read speeds related to the maximum disk speeds.

High web speeds indicate that there is an impact of file system caches for random read, but a lot of random disk IO requests pass to virtual and physical drivers. High throughputs for random read workload (close to max disk speeds) indicate that the FS cache pair realized a solid success with ESXi and Xen.

For Webserver workload, all features from Chapter 4 are important. But we believe that the most impactful components are *Hyp-proc* and *VH-proc*, with limited file system cache effects. Due to a lot of misses in both caches, the virtual drivers of the guest operating system and the physical disk drivers of the host operating system are important. In the context of Webserver workload with random reads, we assume that Xen shows much better combinations of features: *VH-proc* with limited reading file system cache impact, *Hyp-proc*, and virtual and physical disk drivers related to ESXi.

The obtained results of the RFA workload test are presented in Figure 4.

The RFA workload includes the random read and asynchronous random write components, a moderate number of I/O operations, and data flow. For the RFA workload, especially due to asynchronous random writing, the influence of file system caches can be great.

For the RFA workload, XEN is solidly better than ESXi, from 39% to 2 times, it can be considered a remarkable difference. By analysing the obtained RFA workload throughputs and the maximum disk speeds (the disk interface is about 600MB/s and the maximum sequential speeds of SSD disks are about 500MB/s), both hypervisors show significantly higher speeds than the maximum disk speeds. High RFA speeds indicate that file system caches are absolutely dominant, meaning that file system cache pairs demonstrate big success with XEN and ESXi.

For RFA workload, all features from Chapter 4 are very important. However, we assume that the feature *VH-proc*, with the strong cache effects of FS-cache-pair is the most impactful feature. In the RFA workload, Xen shows a much better combination of *VH-proc* with strong file system cache effects related to ESXi.

7. CONCLUSION

We've investigated the disparities in file system performance between ESXi and Xen, two Linux-based type-1 hypervisors. Our mathematical model indicates that the divergence in Linux-based hypervisors stems from various crucial factors: hypervisor processing, virtual hardware processing, file system caching effects on both sides (guest/host), host file systems, and main components of the operating system (kernel, operating system/graphical environments). Although some features may appear similar at first glance, upon closer examination, these features can exhibit significant differences for different Linux-based hypervisors like Xen and ESXi, causing differences in performance.

In our experiment, ESXi vs. Xen and four workloads, the total score was 2:2 per workload. The greatest differences were caused by host file systems (VMFS vs. ext4), virtual hardware processing (*VH-proc*), and hypervisor processing (*Hyp-proc*). In this case study, the ESXi hypervisor is better for Fileserver and Mailserver workloads, while Xen is better for Webserver and RFA workloads. Differences are solid for Fileserver workload (case with strong cache effects and impact of virtual/physical drivers for big data flow), small for Mailserver workload (case with no cache effects and impact of virtual/physical drivers for moderate random read/random write data flow), relatively big for Web (case with limited cache effects and impact of virtual/physical drivers for moderate random read data flow), and strong for RFA (case with solid cache effects and small impact of virtual/physical drivers for random read/random write data flow).

To draw robust conclusions regarding hypervisor performance, it's essential to conduct diverse experiments. All of them should represent different case studies. We suggest potential possibilities for future experimental work. This work should involve the different type-1 Linux-based hypervisors (ESXi, Xen, KVM, and Proxmox) under different hardware configurations and different workload benchmarks. These experiments should include evaluating new releases (versions) of Linux-based hypervisors, different guest operating systems (some versions of Linux and Windows), different file systems (ext4, Btrfs, and XFS), comparing different benchmark tools (Fio, HD Tune Pro, and AS SSD), and variations with factors that may affect performance (such as RAM memory and CPU cores).

8. ACKNOWLEDGEMENTS

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

9. REFERENCES

[1] E. Correia, *Hypervisor based server virtualization*, Encyclopedia of Information Science and Technology, Third Edition, IGI Global, 2015.

[2] A. Kumar, S. Shiwani, "Guest operating system-based performance comparison of VMware & Xen hypervisor," Int. *J. of Sci, Eng. and Tech,* vol. 2, no. 5, pp. 286-297, 2014.

[3] H. Kazan, L. Perneel, M. Timmermann, "Benchmarking the performance of Microsoft Hyper-V server, VMware ESXi and Xen hypervisors," *J. of Emerging Trends in Computing and Information Sciences*, vol. 4, no. 12, pp. 922-933, 2013.

[4] A. Bhatia, G. Bhattal, "A comparative study of various hypervisors performance," *Int. J. of Sci, Eng. and Tech. Research*, vol. 7, no. 12, pp. 65-71, 2016.

[5] S. Pawar, S. Singh, "Performance comparison of VMware and Xen hypervisor on guest OS," *IJICSE*, vol. 2, no. 3, pp. 56-60, 2015.

[6] C. D. Graziano, *A performance analysis of Xen and KVM hypervisors for hosting the Xen worlds project*, Iowa: Iowa State University, 2011.

[7] V. P. Singh, *Analysis of system performance using VMware ESXi server virtual machines*, PhD Thesis, 2012.

[8] S. Ally, "Comparative analysis of Proxmox VE and XenServer as type 1 open source based hypervisors," *Int. J. of Sci. & Tech. Research,* vol. 7, pp. 72-77, 2018.

[9] S. A. Algarni, M. R. Ikbal, R. Alroobaea, A. S. Ghiduk, F. Nadeem, "Performance evaluation of Xen, KVM, and Proxmox hypervisors," *Int. J. of Open Source Software and Processes*, vol. 9, no. 2, 2018.

[10] V. K. Manik, D. Arora, "Performance comparison of commercial VMM: ESXi, XEN, HYPER-V & KVM," in *3rd Int. Conf. on Computing for Sustainable Global Development*, New Delhi, 2016.

[11] B. Đorđević, M. Piljić, N. Kraljević, V. Timčenko, "Comparison of file system performance in full virtualization with MS Hyper-V and KVM hypervisors," in *PROCEEDINGS X ICETRAN*, Novi Pazar, 2022.

[12] 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, 2021.

[13] B. Djordjevic, V. Timčenko, N. Kraljevic, N. Jovicic, "Performance comparison of KVM and Proxmox Type-1 Hypervisors," in *30th TELFOR*, Belgrade, 2022.

[14] B. Đorđević, D. Sakić, N. Kraljević, V. Timčenko, N. Davidović, "File system performance for type-1 hypervisors on the Xen and VMware ESXi," in *21st INFOTEH-JAHORINA,* East Sarajevo, 2022.

[15] "XenServer," [Online]. Available: https://docs.xenserver.com/en-us/citrix-hypervisor. [Accessed 2023].

[16] "GitHub," [Online]. Available: https://github.com/filebench/filebench. [Accessed 2023].