Correction to "File System Performance Comparison in Full Hardware Virtualization with ESXi, KVM, Hyper-V and Xen Hypervisors"

Borislav DJORDJEVIC[1,2], Valentina TIMCENKO[1], Nenad KRALJEVIC[2], Nemanja MACEK[2]
[1]*Mihajlo Pupin Institute & School of Electrical Engineering, University of Belgrade, Belgrade, Serbia*
[2]*VISER, School of Electrical and Computer Engineering of Applied Studies, Belgrade, Serbia*
borislav.djordjevic@pupin.rs

***Abstract*—This paper encompasses the corrections of the paper "File System Performance Comparison in Full Hardware Virtualization with ESXi, KVM, Hyper-V, and Xen Hypervisors." The corrections have been made in section II. RELATED WORK, OBJECTIVE AND MOTIVATION, and the References section. The paper "File System Performance Comparison in Full Hardware Virtualization with ESXi, KVM, Hyper-V, and Xen Hypervisors," authored by B. Djordjevic, V. Timcenko, N. Kraljevic, N. Macek, represents a significant extension of the IEEE conference paper "File system performance comparison in full hardware virtualization with ESXi and Xen hypervisors," authored by B. Dordevic, V. Timcenko, N. Kraljevic, and N. Davidovic. This paper explains the corrections of the "File System Performance Comparison in Full Hardware Virtualization with ESXi, KVM, Hyper-V, and Xen Hypervisors" paper and compares the mentioned papers in detail.**

***Index Terms*—file systems, operating systems, performance evaluation, platform virtualization, virtual machine monitors.**

I. Corrections in the Paper

In our paper [1], section II. RELATED WORK, OBJECTIVE, AND MOTIVATION was incomplete. A new passage has been added at the end of the third paragraph of this section, as follows:

"This paper represents an extended version of our IEEE conference paper [27], which focuses on the measurement and interpretation of filesystem performance for the case of only two hypervisors, ESXi and Xen." Although both papers use the same experiment as a practical methodology, this one provides several extensions compared to the conference paper. The first significant difference is the covering of four hypervisors included in the experiments and the generation of the results for two different case studies that rely on the implementation of two different hardware configurations. The second significant difference is that the focus of this paper was shifted entirely to the extended modeling of the file system performance in a virtual environment and the generation of knowledge database, which includes different case studies related to the filesystem and hypervisor performances in various virtual environments.

In our paper [1], section References was incomplete. We have added a new reference, the 27^{th,} as follows:

[27] B. Dordevic, V. Timcenko, N. Kraljevic, and N. Davidovic, "File system performance comparison in full hardware virtualization with ESXi and Xen hypervisors," 2019 18th International Symposium INFOTEH-JAHORINA (INFOTEH), 2019, pp. 1-5. doi:10.1109/INFOTEH.2019.8717664.

II. Explanations for Corrections

The authors state that their AECE journal paper entitled "File System Performance Comparison in Full Hardware Virtualization with ESXi, KVM, Hyper-V, and Xen Hypervisors," authored by B. Djordjevic, V. Timcenko, N. Kraljevic, N. Macek, here referenced as [1], represents a significant extension of the IEEE conference paper "File System Performance Comparison in Full Hardware Virtualization with ESXi and Xen Hypervisors," authored by B. Dordevic, V. Timcenko, N. Kraljevic, and N. Davidovic, here referenced as [2].

Although the mentioned papers deal with similar topics, we state that the conference paper [2] was just an initial basis for the research and a part of case studies presented in the AECE journal paper [1]. Authors responsibly claim that paper [1] is far more comprehensive in both scientific and experimental contexts.

The main extensions, differences, and improvements of paper [1] compared to paper [2] are as follows:

- The focus of the manuscripts is entirely different. In the conference paper, the focus is on the hypervisor performance comparison, whereas in the journal paper, the focus is on the mathematical modeling;
- One of the major theoretical contributions of paper [1] compared to paper [2] is the extended and improved mathematical model presented in paper [1];
- Paper [1] provides a significantly expanded experimental part and results that additionally cover four hypervisors evaluated in a different environment compared to only two hypervisors and the environment analyzed in paper [2];
- Authors have provided a significantly better interpretation of experimental results in paper [1] when compared to the methodology and scope of the analysis in paper [2];
- As the second major theoretical and practical contribution, in paper [1], the authors have introduced the knowledge database that has further found a practical usage by some of their colleagues and clients.

Digital Object Identifier 10.4316/AECE.2021.03013

III. Detailed Comparisons

When comparing paper [1] and [2] chapter by chapter, authors can highlight the following differences.

1. The differences related to the Abstract

In the case of paper [1], the main focus is on the mathematical modeling of the file system performance in a virtual environment when using type-1 hypervisors and taking four different hypervisors as the experimental basis whereas, in paper [2], the focus is on the hypervisor type-1 performance comparison, taking into consideration only two virtualization platforms, VMWare ESXi and Citrix XenServer.

2. The differences related to Section II – "Related Work, Objective and Motivation"

There are significant differences between paper [1] and paper [2] in the section describing the main contributions. More precisely, in paper [1], the main contributions are provided through all-encompassing, comprehensive mathematical modeling of the file system performance in a virtual environment when using type-1 hypervisors, whereas in paper [2], the main contribution is to examine the performances of two hypervisors, ESXi and Xen, in fair play conditions. Paper [2] deals with only one performance measurement (single case study). Paper [1] introduces two new ideas. The first one is a complex mathematical model that can be used to predict and interpret filesystem performance in a virtual environment. Another one is to collect a large number of case studies (similar as in the [2]) analyzed for different hardware/software environments, and stored into the knowledge database, which can be further useful for the system administrators.

3. The differences related to section III – "Hypervisor Description"

The research presented in paper [1] describes four different hypervisors: ESXi, Xen, KVM, and Hyper-V. Paper [2] describes only two hypervisors: ESXi and Xen.

4. The differences related to the sections related to mathematical modeling and expected behavior

In [1], mathematical modeling and expected behavior are included in chapters: "IV. Mathematical Modeling of FS Performance for Virtual Environment" and "V. Case Study and Expected Behavior" In [2], they are included in chapter "IV. Hypotheses of Expected Behavior"

The main extension and improvement presented in paper [1] compared to paper [2] are exactly in the mathematical modeling and details related to the expected behavior. Paper [1] covers a much more comprehensive model expressed through seven equations, whereas paper [2] contains only three basic equations. Both papers start from three general equations that apply to all filesystem performances in a virtual environment. In eq. 3, paper [1] includes all types of virtualization, whereas paper [2] contains only full hardware virtualization. In paper [1], the mathematical model is extended remarkably by adding the equations 4-7, which defines the filesystem pairs interactions, combining a large number of parameters such as filesystem types, caches, and related hypervisors parameters. The proposed model includes the disk technologies, dependency on the filesystem version, as well as some specific hypervisor parameters such as CPU scheduling. In addition, the paper [1] is discussing in detail the expected behavior for the case of full hardware virtualization provided in the experimental part and through the obtained results. With a significantly extended model, authors believe that paper [1] has a considerably higher scientific contribution compared to paper [2].

5. The differences related to the experimental part sections

The experimental part has significant improvements regarding results and scope in the paper [1]. Paper [2] only contains results for two hypervisors, ESXi and Xen, whereas paper [1] contains results for four hypervisors: ESXi, Xen, KVM, and Hyper-V. The authors have also added the experimental results for the native guest operating system performance, which represents the system without the virtualization. It is vital to highlight that, for the needs of the research presented in paper [1], the authors have repeated the set of the experiments in a different environment relying on a server HPE ProLiant DL 180 G9, which has much better characteristics, and therefore the authors have provided a different case study. Both case studies have resulted in relatively similar filesystem performances, which is actually in line with the authors' theoretical assumptions, thus allowing to make the conclusions related to the used hypervisors and this way gaining greater validity for the further practical application. Due to the all-encompassing mathematical model, the interpretation of the obtained results is more exhaustive and gives more quality to paper [1] when compared to paper [2].

6. The differences related to the conclusion section

The Conclusion Section is significantly expanded in paper [1]. The results from both papers indicate that there is no best hypervisor, the solution that would be the best option for any possible use case, but the behavior of a hypervisor depends on the workload characteristics. It is clear that paper [1] has significantly richer future work, but the main difference between the two is the conclusion part in the journal paper, where the authors highlight again the contribution related to the introduction of knowledge database, which corresponds to the second theoretical and practical contribution of paper [1]. There, the authors have formed two case studies for this database. Based on the model, the authors can indicate that there are numerous factors affecting filesystem performance in different virtual environments, and only a rich knowledge database, with a large number of case studies, can be helpful to administrators to create their optimal virtual environments.

References

[1] B. Djordjevic, V. Timcenko, N. Kraljevic, N. Macek, "File System Performance Comparison in Full Hardware Virtualization with ESXi, KVM, Hyper-V and Xen Hypervisors," Advances in Electrical and Computer Engineering, vol. 21, no. 1, pp. 11-20, 2021. doi:10.4316/AECE.2021.01002

[2] B. Dordevic, V. Timcenko, N. Kraljevic, and N. Davidovic, "File system performance comparison in full hardware virtualization with ESXi and Xen hypervisors," 2019 18th International Symposium INFOTEH-JAHORINA (INFOTEH), pp. 1-5. 2019. doi:10.1109/INFOTEH.2019.8717664