RESEARCH ARTICLE OPEN ACCESS

Comparative Analysis Of Virtualization And Containerization: Performance Perspective

Arnold Mashud Abukari [1], Seidu Nimatu [1], Jhansi Bharathi Madavarapu [2]

Department of Computer Science, Tamale Technical University, Ghana [1]

Department of Information Technology, University of the Cumberland's, Williamsburg, Kentucky, USA,40769 [2]

ABSTRACT

In this study, the technical performance of virtualization and containerization is examined. Analyze the starting times of virtual machines (VMs) in cloud environments, a critical performance characteristic. Three researchers have thoroughly researched the context of virtualization and containerization technologies using various suggested methods, each revealing the specific benefits and drawbacks of these tactics. These studies emphasize the importance of carefully weighing the advantages and disadvantages of virtualization and containerization technologies in light of the particular application requirements and larger organizational goals. This study has laid a solid foundation for adopting technology by thoroughly investigating the complex interactions between various technologies and their related effects. The research contributions aim to provide academics, practitioners, and organizations with the knowledge necessary to make wise technical decisions that will enable them to accomplish their goals as the technological landscape continues to change.

***Keywords**:* Cloud Computing, power consumption, Performance, virtualization, hypervisor, KVM, Xen, containerization, Docker, Kubernetes, LXC

I. INTRODUCTION

Although virtualization and containerization are technologies used to build and manage virtual environments, they have different functions.

Using software known as a hypervisor[1], virtualization generates a virtual replica of a physical device, such as a desktop or server. Multiple virtual machines, each having its operating system and resources, can run simultaneously on a single physical machine. Virtual machines can run various operating systems and applications and are segregated from one another.

Contrarily, containerization uses a container engine like Docker[2] or Kubernetes[3] to integrate application code and its dependencies into a single container. Containers share the host operating system and are smaller than virtual computers. Despite being separated from one another, containers are more intimately coupled to the host operating system, which might improve resource efficiency.

In conclusion, containerization embeds an application and all its dependencies into a container that shares the host operating system, whereas virtualization produces a virtual representation of a physical device. Containerization is more resource-efficient since it shares the host operating system, but virtualization consumes more resources because it runs several operating systems.

In this study, virtualization and containerization are compared regarding technical performance. The study will compare resource utilization, scalability, and overall performance of virtual machines and containers. The study will assess the constraints and applications of each technology. In order to conduct a thorough analysis of the performance traits of virtualization and containerization, data will be gathered through benchmarking and actual testing. Based on the performance analysis, the study will make suggestions for organizations on the technology most suited to their requirements.

Even though virtualization and containerization are related technologies that are commonly combined,

---

[1] a type of computer software, firmware or hardware that creates and runs virtual machines.

[2] a set of platform-as-a-service products that use OS-level virtualization to deliver software in packages called containers.

[3] an open-source container orchestration platform for managing, scaling, and automating software deployment.

they have several key differences. The operating system is used in two different ways by these two technologies.

This study compared and contrasted the performance characteristics of virtualization and containerization technologies. The study evaluated and compared the two methodologies' resource utilization, scalability, and overall system efficiency to determine which technology offers more excellent performance.

II. VIRTUALIZATION AND CONTAINERIZATION

Virtualization and containerization are two essential technologies becoming increasingly popular in modern IT design. Virtualization technology makes the underlying hardware resources accessible as virtual resources so that many operating systems can be run on a single physical computer. Thanks to the lightweight virtualization method known as containerization, various isolated applications can run on the same operating system.

Recent research has shown that virtualization and containerization technologies can significantly benefit resource utilization, scalability, and adaptability. Particularly in terms of performance, the trade-offs of the two technologies are still poorly known.

Numerous scholars are actively undertaking many studies to evaluate the performance of virtual machines. Numerous publications in the literature discuss improving the performance of virtual machines. To calculate the overhead of VMs, most of the literature uses the VMM abstraction layer in general.

III. VIRTUALIZATION-INDUCED PERFORMANCE DEGRADATION IN NUMA ARCHITECTURES

Ibrahim et al. (2011) carried out a comprehensive investigation to quantify the performance degradation of virtual machines within Non-Uniform Memory Access (NUMA) systems, particularly regarding the impact of virtualization. Two well-known hypervisors, KVM and Xen, were employed for the experimental evaluation.

The study highlighted a notable phenomenon known as "leakage" of page locality that occurs in virtualized systems. High-Performance Computing (HPC) programs had an average performance decrease of roughly 55% when compared to peers with native performance as a result of this phenomenon. Surprisingly, the study revealed that hypervisor-based techniques could only partially resolve memory locality problems in virtualized systems.

Notably, the study found a viable strategy for addressing these performance problems. System partitioning developed as a strategy that could also result in a noticeable gain of up to 60% over VMs with superior NUMA support when used in conjunction with sensible VM selection, runtime support, and replication of native performance.

Ibrahim et al. (2011) narrowed their investigation to the behavior of scientific computing programs in virtualized environments, focusing on memory locality management in particular. The authors thoroughly analyzed the performance of virtual machines (VMs) running on the Xen and KVM hypervisor systems. Notably, the study showed how VMs distributed across numerous sockets in NUMA architectures significantly degraded performance.

A comparative analysis of various programming models was also conducted as part of the research. Their investigation demonstrated that cluster programming models—exemplified by MPI—exhibited superior scalability and performance features when compared to shared memory models like OpenMP.

A critical component of the work was the investigation of strategies for improving memory localization. This entails assessing hypervisor-only tactics and putting system partitioning into practice. The outcomes demonstrated how well optimal system partitioning performed when coupled with suitable VM setup and runtime assistance to simulate native performance. Interestingly, the cumulative impact of the current implementations in the instance of KVM resulted in an average performance reduction of 55%. Partitioning, however, lowered this deterioration to just 11%.

In summary, the study by Ibrahim et al. (2011), emphasizing memory locality control, sheds significant light on the intricate interconnections between virtualized systems and scientific computing applications. The study emphasizes the ability of intelligent system partitioning to lessen performance degradation and, ultimately, bring virtualized systems closer to parity with their native versions. It also emphasizes the challenges brought on by NUMA designs.

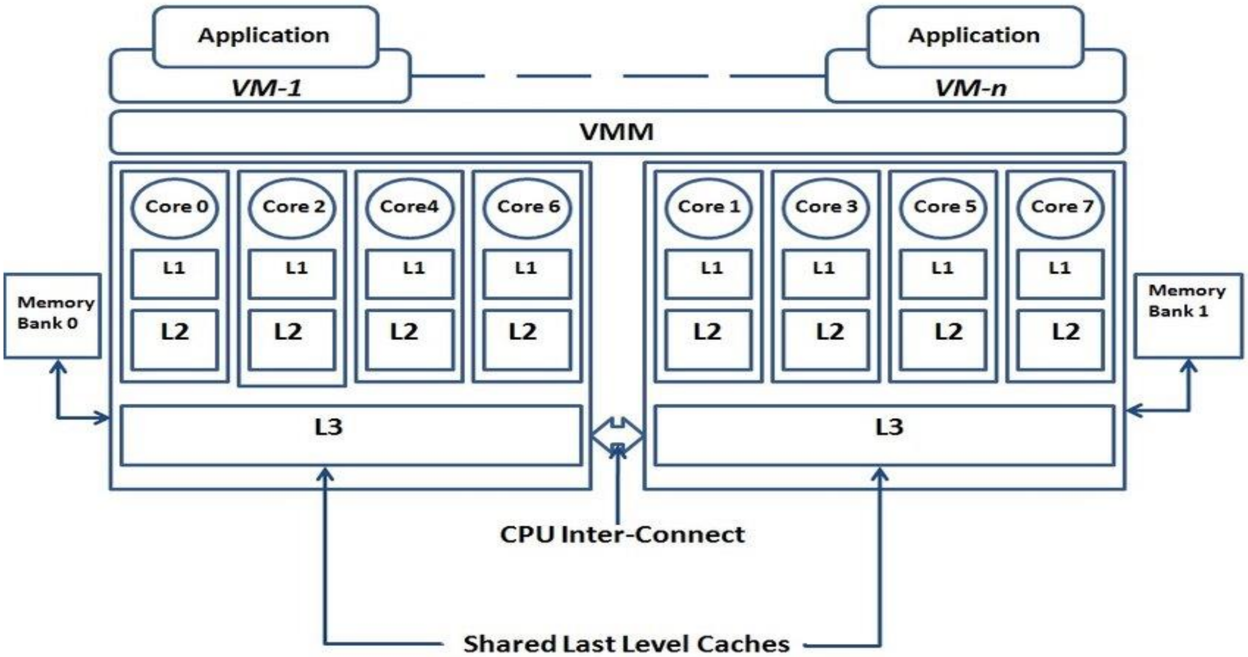

Figure 1: Diagram of NUMA System (Mukherjee et al., 2015)

IV. OPTIMIZING VIRTUAL MACHINE STARTUP TIME IN CLOUD ENVIRONMENTS

The main objective of a noteworthy analysis conducted by Mao and Humphrey (2012) was to evaluate a crucial performance parameter—the startup time of virtual machines (VMs) in cloud settings. This experiment utilized three well-known cloud service providers, Amazon EC2, Rackspace, and Microsoft Azure, for accurate time measurements and subsequent analysis.

The initial description of Ostermann et al. (2008) is a thorough and in-depth investigation of efficiency evaluations of cloud computing applications for scientific workloads. The study first highlights the problems of dependability and cloud efficiency before underlining the crucial function of Infrastructure-as-a-Service (IaaS) to fulfill the demands of scientific workloads within a cloud paradigm.

The study discovered that VM initiation times inside cloud infrastructures considerably impact both application performance and end-user satisfaction. The size of the operating system (OS) image, the type of instance, the location of the data center, and the quantity of concurrently purchased instances are a few variables that affect how soon VMs start up.

The study warns cloud users that spot instances have more unpredictable startup times and lengthier wait times than on-demand instances.

The study summarizes the critical topic of VM starting time in cloud computing systems. The authors carefully examine the launch times of virtual machines on three key real-world cloud providers: Amazon EC2, Windows Azure, and Rackspace. The relationships between starting time and numerous factors are then thoroughly investigated, considering temporal considerations, the size of the OS image, instance classification, the location of data centers, and the number of concurrent instance acquisitions.

Additionally, the EC2 architecture's spot instance launch time dynamics are explicitly examined and contrasted with those of on-demand instances. This study's primary objective is to give crucial information to cloud clients so they can make informed and strategic decisions.

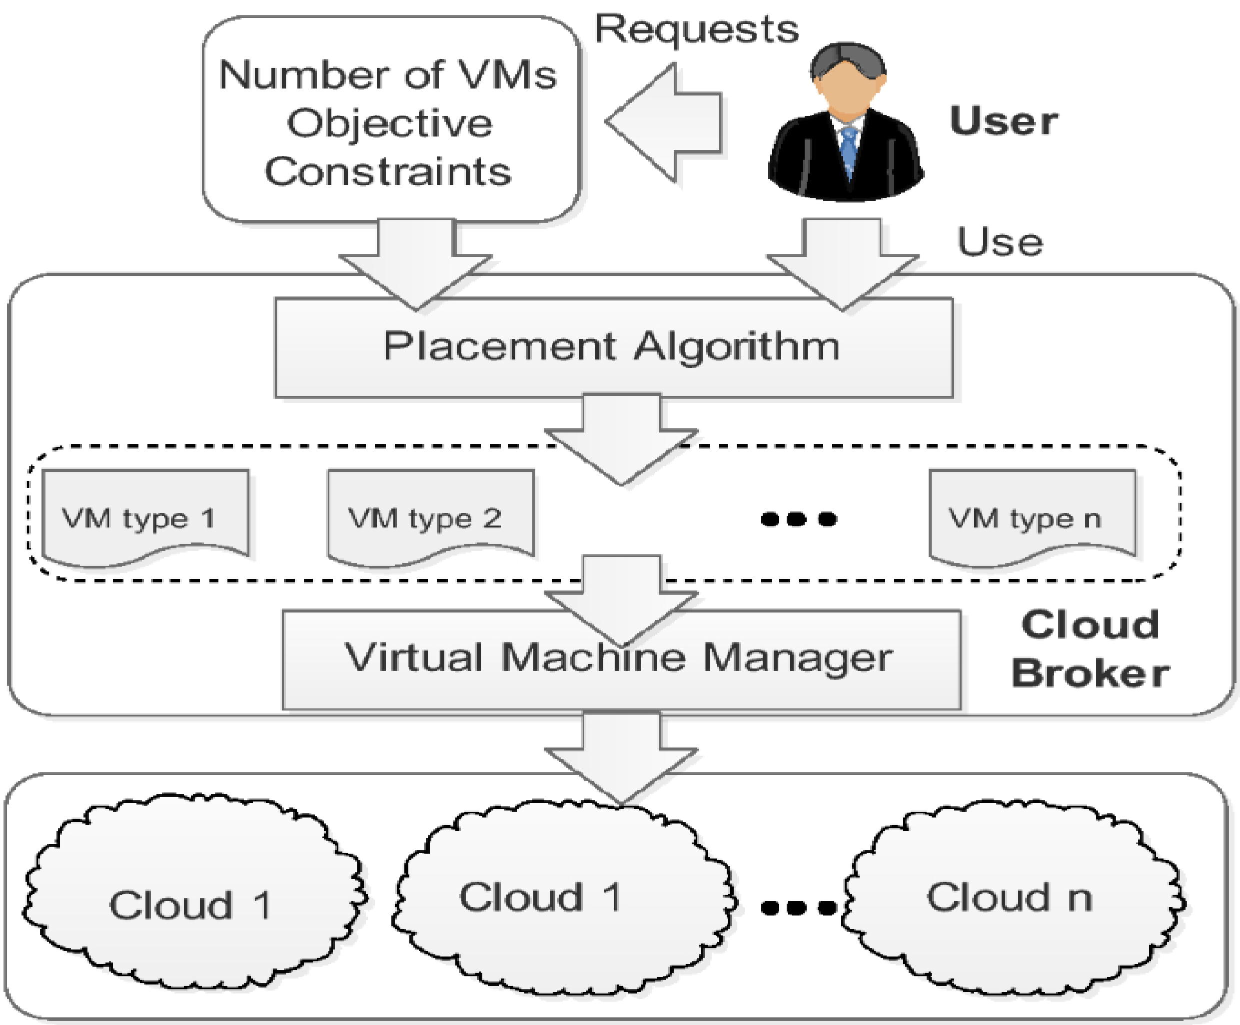

Figure 2: Optimizing IaaS Data centers (Talebian et al., 2020)

EXPLORING VIRTUAL MACHINE PERFORMANCE OVERHEAD IN INFRASTRUCTURE-AS-A-SERVICE CLOUD ENVIRONMENT

Another work by Xu et al. (2011) thoroughly analyzes the performance effects of virtual machines (VMs). The authors begin a study that analyses the performance overhead caused by VMs and charts how this overhead increased as virtualization spread from a single server to numerous geographically scattered data centers.

The study makes a crucial finding: Infrastructure-as-a-Service (IaaS) cloud computing incurs noticeable performance overhead despite providing a compelling means of delivering scalable and economical virtual machines. The competition for resources in data centers is the leading cause of this overhead. However, the research shows that there have been coordinated attempts to reduce VM performance overhead in various IaaS cloud scenarios. These scenarios range from single-server virtualization to enormous mega-datacentres, and they go even further to include the difficulties of having numerous geo-distributed data centers.

The discovery of potential future research concerns about controlling VM performance overhead in the IaaS cloud is a significant takeaway from this work. The accuracy and efficacy of performance modeling techniques must be improved, which is a significant problem. The study also emphasizes how important it is to choose and put into practice effective performance overhead reduction techniques.

The study effectively provides a thorough grasp of the performance dynamics interacting with VMs in cloud computing's Infrastructure-as-a-Service (IaaS) model. The critical balancing act between the IaaS's capacity for growth and cost-effectiveness and any potential performance obstacles brought on by the resource-sharing model in data centers is highlighted. The paper examines various techniques for evaluating and minimizing this performance overhead and explains its sources. The research comes to a close by outlining the challenges that still need to be overcome to manage VM performance overhead effectively in the evolving IaaS cloud environment.

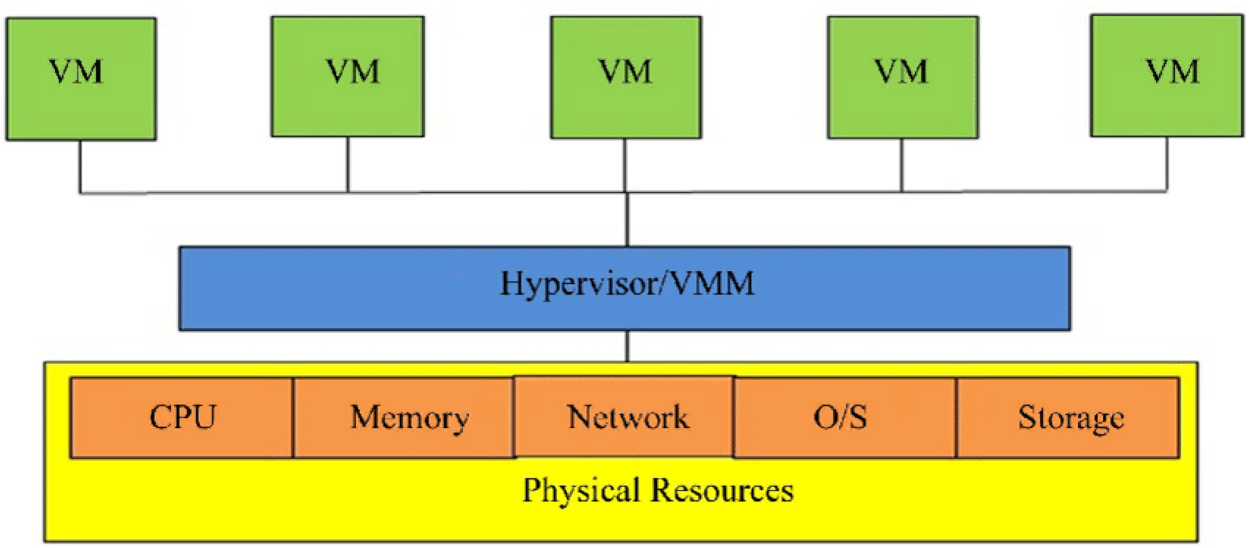

Figure 3: Virtual Computing Platform Environment (Riskhan and Muhammed, 2015)

DRAWBACKS OF REVIEWED SYSTEMS

A comparison of the reviewed proposed systems is tabulated in Table 1.

Table 1: Drawbacks

Proposed Systems	Drawback and Limitation
Ibrahim et al. (2011)	- Limited Hypervisor Scope: This does not consider other hypervisors that might affect performance differently, only looking at KVM and Xen.
	- Limited Workload Applicability: Because the findings concentrate on High-Performance Computing (HPC) software, they might not apply to all workloads.
	- Lack of Architectural Diversity: Only considers NUMA architectures, ignoring performance difficulties in other system types.
	- Narrow Performance Metrics: Focuses only on performance decline without considering gains that could be made in other areas.
Mao and Humphrey (2012)	- Limited Investigation of Cloud Providers: The study only looks into Amazon EC2, Rackspace, and Microsoft Azure, leaving out information from other providers.
	- Temporal Limitation: Because the study was done in 2012, its findings might not be consistent with recent developments in cloud computing.
	- Limited Overhead Consideration: Ignores other performance overheads in cloud settings and concentrates only on VM starting time.
	- Lack of Security Consideration: The study does not explore the potential security ramifications of VM startup optimization.
Xu et al. (2011)	- Timeframe Restrictions: The study's results may not accurately reflect how Infrastructure-as-a-Service (IaaS) changes over time.
	- Definition of overhead: The study focuses exclusively on performance overhead, ignoring other elements like security or compliance.
	- Insufficient Cloud Provider The study does not consider a broader selection of IaaS cloud providers that can influence overhead.
	- Inadequate Mitigation Detail: Although the study discusses mitigation, it doesn't thoroughly examine the efficacy of those strategies.

COMPARISON OF REVIEWED CASE STUDIES

Comparisons of the reviewed proposed systems are tabulated in Table 2.

Authors	Ibrahim et al. (2011)	Mao and Humphrey (2012)	Xu et al. (2011)
Focus	Virtualization-induced performance degradation in NUMA	Optimizing VM startup time in cloud environments	VM performance overhead in IaaS cloud

Methodology	Comprehensive investigation, experimental evaluation	Accurate time measurements and analysis of cloud providers	A thorough analysis of performance implications, IaaS evaluation
Hypervisor Analysis	KVM, Xen	N/A	N/A
Cloud Providers Analyzed	N/A	Amazon EC2, Rackspace, Microsoft Azure	N/A
Virtualization Scenarios	NUMA architectures	N/A	IaaS cloud environments
Performance Phenomenon	"Leakage" of page locality in virtualized systems	VM startup time impact on application performance	Performance overhead in IaaS cloud
Findings	Significant performance degradation in VMs on NUMA	VM startup time can significantly impact user experience	Noticeable performance overhead in IaaS, efforts to mitigate
Research Challenges	Limited potential of hypervisor-only approaches	Identifying factors affecting startup time, awareness of spot instances	Improving performance modeling accuracy, effective mitigation
General Applicability	Specifically focuses on scientific computing applications	Broadly applicable to cloud environments, regardless of workload	Broadly applicable to IaaS cloud environments
Scope of Analysis	Limited NUMA architectures	Startup time optimization in cloud settings	Examines performance overhead in IaaS cloud environments
Temporal Consideration	No temporal limitation is considered	Conducted in 2012, it might not align with current cloud trends	No temporal limitation is considered
Diversity of Factors	Focuses on performance degradation and memory locality	Explores startup time and its factors	Examines performance overhead and efficiency in an IaaS environment

V. CONCLUSION

The context of virtualization and containerization technologies has been extensively investigated through multiple proposed systems by three researchers, each revealing distinct advantages and disadvantages of these strategies. These studies underline how crucial it is to carefully consider the benefits and drawbacks of virtualization and containerization technologies in light of the specific requirements of the application and the overarching organizational goals.

Despite providing isolation and administrative benefits, virtualization may cause performance loss in some situations, such as Non-Uniform Memory Access (NUMA) architectures, as seen through the lens of Ibrahim et al. (2011). Mao and Humphrey (2012) add to this subject by highlighting the necessity of optimizing VM starting times in cloud environments and showing that, despite the advantages of the cloud, even little inefficiencies can significantly impact performance. Xu et al. (2011) significantly contribute by addressing the broader context of Infrastructure-as-a-Service (IaaS) cloud systems and highlighting the severe performance overheads connected with VMs.

With the help of the various insights gleaned from these case studies, the proposed study makes an effort to add to the body of literature by offering a complete comparative analysis of virtualization and containerization technologies. The focus has examined crucial performance metrics, including resource utilization, scalability, and response time. This research has provided a holistic perspective that enables decision-makers to compare these indicators across the virtualization and containerization paradigms and choose the optimum option for their specific use cases.

By carefully examining the complicated relationships between various technologies and their accompanying repercussions, this study has established a strong foundation for technology adoption. The research contributions are meant to equip researchers, practitioners, and organizations with the wisdom required to choose strategic technologies to help them achieve their objectives as the technological landscape continues evolving.

REFERENCES

K. Ibrahim, S. Hofmeyr, and CostinIancu, “Characterizing the performance of parallel applications on multi-socket virtual machines,” The 11th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, 2011.

S. Zaghloul, “Virtualization: the key to cloud computing,” MASAUM International Conference on Information Technology (MICIT’12), 2012.

P. Perera and C. Keppitiyagama, “A Performance Comparison of Hypervisors,” The International Conference on Advances in ICT for Emerging Regions (ICTer2011), 2011.

Z. Magsi, M. Y. Koondhar, M. H. Depar, Z. H. Pathan, F.-U.-D. Memon, and S. Solanki, “Conceptual framework transformation of converged infrastructure (CI) into hyper-converged technology for virtualization of server infrastructure,” in 2020 IEEE 7thInternational Conference on Engineering Technologies and Applied Sciences (ICETAS), Dec. 2020, pp. 1–4, doi:10.1109/ICETAS51660.2020.9484233.

Mao, Ming, and Marty Humphrey. "A performance study on the VM startup time in the cloud." Cloud Computing (CLOUD), 2012 IEEE 5th International Conference on. IEEE, 2012.

Xu, C.; Zhao, Z.; Wang, H.; Shea, R.; Liu, J., "Energy Efficiency of Cloud Virtual Machines: From Traffic Pattern and CPU Affinity Perspectives," Systems Journal, IEEE, vol.PP, no.99, pp.1,11

J. Mukherjee, D. Krishnamurthy and J. Rolia, "Resource Contention Detection in Virtualized Environments," in *IEEE Transactions on Network and Service Management*, vol. 12, no. 2, pp. 217-231, June 2015, doi: 10.1109/TNSM.2015.2407273.

Talebian, H., Gani, A., Sookhak, M. *et al.* Optimizing virtual machine placement in IaaS data centers: taxonomy, review, and open issues. *Cluster Comput* **23**, 837–878 (2020). https://doi.org/10.1007/s10586-019-02954-w

Riskhan, B. and Muhammed, R. (2015) Enhancing the Performance of Current Online Education System—A Study of Cloud Computing and Virtualization. *Journal of Computer and Communications*, **3**, 43-51. doi: 10.4236/jcc.2015.310006.

Simon Ostermann, Alexandru Iosup, Nezih Yigitbasi, Radu Prodan, Thomas Fahringer, andDick Epema. An Early Performance Analysis of Cloud Computing Services for ScientificComputing. Technical Report PDS-2008 006, Delft University of Technology, The Netherlands, December 2008.