

Performance Assessment of Network Virtualization Technologies in Cloud Computing Environments: A Comparative Study of Leading Hypervisors

Irfan Muhammad^{1,*} and Nadeem Iqbal²

¹Coventry University, Coventry, CV1 5FB, United Kingdom.; Email: irfanm908@gmail.com

²Department of Computer Science, University of Central Punjab, Lahore, 54000, Pakistan.;
Email: niqbal@ucp.edu.pk

*Corresponding author: Irfan Muhammad (irfanm908@gmail.com)

Article History

Academic Editor:

Dr. Muhammad Sajid

Submitted: January 27, 2023

Revised: May 1, 2023

Accepted: September 1, 2023

Keywords:

Virtual Machine, Virtual Network, Network Virtualization Environment, Cloud Computing

Abstract

Cloud computing represents a transformation paradigm in contemporary computer science applications. The foundation of cloud infrastructure relies heavily on virtualization techniques, whereby singular physical resources are transformed into multiple virtual counterparts. The virtualization process encompasses various components, including storage systems, computational processors, volatile memory units, networking infrastructure, and additional hardware resources. Both processor and network virtualization technologies deliver substantial benefits, including reduced hardware expenditures, decreased energy consumption, and minimized administrative overhead for resource management. This research concentrates on evaluating and analyzing network virtualization implementations across various cloud computing hypervisors while comparing their performance metrics. The comparative analysis is grounded in experimental studies conducted on three prominent platforms. Our findings demonstrate that virtual network architectures provide enhanced scalability and security compared to traditional physical networks. The capability to assign multiple virtual interfaces to individual virtual machines presents significant operational advantages. Through comprehensive experimentation comparing virtualized against non-virtualized environments across all evaluated hypervisors, we determined that virtual networks function effectively both with and without VLAN tagging mechanisms. Hyper-V exhibits superior flexibility and configuration options regarding VLAN tagging compared to XenServer or vSphere implementations.

1 Introduction

This section establishes the fundamental knowledge required for comprehensive understanding of cloud computing principles. We structure this foundation into essential subsections to support subsequent discussions throughout this paper. The primary divisions include: first, an exploration of the "Cloud" concept itself; second, an examination of "virtualization" technologies; and third, a summary of the remaining article content. In the current technological era, most individuals are familiar with the latest hardware innovations and improvements developed by technology companies from various manufacturers including Intel, Apple, AMD, and other industry leaders. Contemporary machines incorporate

advanced and updated technologies, demonstrating significantly superior performance and capabilities compared to systems utilized fifteen to twenty years ago.

Detailed research into the improvement and enhancement processes underlying these modern hardware technologies reveals that hardware development has approached its theoretical limits. Continuous technological evolution has rendered hardware components so dense and complex that further advancement presents significant challenges, and when possible, remains economically prohibitive. Consequently, manufacturers and technologists have shifted their research focus toward maximizing output from existing hardware resources rather than increasing density and complexity. Cloud computing has emerged as a direct result of these research initiatives.

Performance analysis reveals that XenServer and vSphere deliver superior network throughput for both TCP and UDP protocols compared to Hyper-V, while Round Trip Time (RTT) measurements in virtual networks show increased latency across all hypervisors relative to non-virtualized environments. CPU utilization monitoring indicates that vSphere and XenServer consume more processing resources during network operations compared to Hyper-V. Additionally, processor virtualization demonstrates minimal impact on average scheduling time compared to non-virtualized systems. We conclude that the marginal performance reduction associated with virtualization is negligible when weighed against the substantial advantages gained from implementing virtualized cloud hypervisors. This research provides valuable insights for organizations considering virtualized data center deployments.

2 Literature Review

This section examines previous research conducted in cloud computing and virtualization domains and differentiates our work from existing studies. We conclude this section by explaining our contribution to network virtualization research. Previous research primarily focuses on implementation mechanisms, algorithms, virtualization technologies, and related areas within cloud computing environments. Our primary focus centers on describing previous research work in network virtualization. models are inadequate for properly managing and monitoring interconnections among virtual resources, necessitating the implementation of robust techniques, namely network virtualization. Network virtualization offers enhanced advantages and functionalities compared to physical networks, including software-based management and control capabilities.

In recent years, network virtualization has garnered considerable interest from researchers. Chowdhury et al. [10] examined network virtualization and four extant technologies that closely align with its principles. They identified areas and initiatives for the deployment of network virtualization and examined the architectural characteristics of network virtualization in cloud settings. According to these authors, rather than redesigning entire systems from scratch, network virtualization with support for coexisting multiple current architectures represents the solution for overcoming challenges and issues in existing systems and internet infrastructure. According to their research [8], Open vSwitch acts as a bridge between Physical Interfaces (PIFs) and VMs. While functioning as a Layer 2 switch, it replicates Ethernet switching behavior using VLAN, RSPAN, and ACL technologies. Open vSwitch exposes interfaces for state management and forwarding, but its connectivity management adheres to a model akin to OpenFlow [11]. Pfaff et al. [8] elucidated the implementation of Open vSwitch for centralized management, virtual private networks, mobility, distributed configuration, network host visibility, and IP subnet segregation. Our investigation concentrates on evaluating Open vSwitch based on various parameters using XenServer hypervisor.

3 Methodology

For some practitioners, cloud computing signifies a fascinating concept; however, many remain unclear regarding its exact definition [1]. Cloud computing constitutes a distributed computer system design. Before examining the applications and implementations of cloud systems, we must establish a clear definition of cloud computing. Numerous definitions exist for cloud computing, yet none provide a completely concrete and comprehensive characterization. Each definition emphasizes different aspects

of cloud computing, with primary focus areas including scalability, resource monitoring, automation and optimization, collaboration capabilities, cost-effectiveness, utility computing, user accessibility, service delivery, and hardware/software virtualization [1]. Vaquero et al. [1] proposed a comprehensive definition that considers all features and factors inherent in cloud systems: Clouds are extensive reservoirs of readily available and accessible virtualized resources, including hardware, development platforms, and services.




Figure 1: Service Provider User Relationship Model

Cloud systems seek to delegate these services to cloud Infrastructure Providers (IP) as a service model. This configuration reallocates computer resources from service providers to infrastructure providers, leading to improved flexibility and cost savings for service providers. Figure 1 depicts this relationship [1, 2]. According to cloud service categories, these services can be categorized into three separate models:

Infrastructure as a Service (IaaS): Infrastructure providers generally manage data centers housing numerous physical machines using various hypervisors that facilitate multiple virtual machines. These virtual machines encompass virtualized resources. Users generally access one or many virtual computers via internet connectivity. For users, these virtual machines operate as real computers that may be accessed remotely [1, 3]. **Platform as a Service (PaaS):** PaaS offers users a computing platform, unlike IaaS. Users access the system via many software interfaces and employ the system's resources. The requirements and management of hardware resources remain apparent to users.

4 Understanding Virtualization Technologies

In earlier parts, we have often mentioned the term "Virtualization," but what does "Virtualization" truly denote? Mike Adams, director of product marketing at VMware, a leader in virtualization and cloud software services, states, "Virtualization software enables the simultaneous operation of multiple operating systems and applications on a single server." It allows enterprises to decrease IT expenditures while enhancing the efficiency, use, and flexibility of current computing gear. According to Adams' explanation, virtualization denotes an environment where resources function within a regulated virtual framework instead of on tangible physical resources. Virtualization enables a single physical machine's

resources, such as main memory, CPU, and storage, to be employed as several virtual processors, memory units, and storage systems.

4.1 Distinguishing Virtualization from Cloud Computing

A prevalent misperception is that virtualization and cloud computing are identical; however, this is inaccurate [6]. This misconception generally stems from the profound interrelationship between virtualization and cloud computing technologies. Virtualization and cloud computing jointly offer a range of services, including private cloud implementations [4]. Virtualization, specifically known as virtual machines (VMs), is founded on physical computer resources. These virtual machines possess their own virtual storage devices and network connectivity, while the cloud dictates the allocation, delivery, and presentation of these virtualized resources. A cloud environment can be created without employing virtualization technology. Nonetheless, virtualized resources render the environment more scalable in comparison to non-virtualized implementations [6, 7].

4.2 Virtualization Implementation in Cloud Computing

Numerous research articles discuss these virtual components. In cloud environments, multiple virtual resources function collectively as a cloud system and deliver services to end clients or users. For these virtual resources, or more specifically, virtual machines, to provide services, they must work collaboratively. If VMs collaborate, communication mechanisms must exist, as collaboration would be impossible otherwise. The answer is affirmative: there exists a mechanism called networking, or more specifically, network virtualization [5]. Communication mechanisms require various management and control methodologies. Although virtual environments have multiple advantages, they also pose some obstacles, such as scale, migration, isolation, and security issues [8]. These difficulties necessitate resolution [8]. In physical networks, we may monitor and regulate network traffic between nodes utilizing specialized tools and dedicated hardware. Nonetheless, in virtual contexts, this becomes even more challenging. In virtual networks, communication is directed at the virtual layer, where it can potentially bypass any physical network hardware, making network communication tracking and control significantly more challenging [5, 8, 9]. Network virtualization promises to address all these challenges.

5 Experimental Setup and Implementation

In prior sections, we examined XenServer, Hyper-V, vSphere, virtual machines, virtualization, Open vSwitch, XenCenter, and associated concepts. This section focuses on experimental environment setup and evaluation procedures. Our primary objective involves establishing virtualized cloud environments for all the mentioned hypervisors.

5.1 XenServer Implementation

This section delineates the experimental configuration for XenServer hypervisors. All hypervisor host servers employ machines with identical setups. Initially, we acquired the ISO file from the Citrix website, available for download on their official portal [34]. Upon completion of the download, we generated a bootable USB device using this image. Our trials employed XenServer version 6.2.0, which incorporates Open vSwitch version 1.4.6. XenServer is constructed on the Linux kernel, and its installation procedure is akin to that of other Linux distributions. Nonetheless, hardware specifications impose specific limitations. XenServer requires a maximum of 1TB and a minimum of 1GB of RAM, supports up to 19 and at least 1 100Mb or faster NICs, accommodates a maximum of 160 logical CPUs, and necessitates a minimum of 100MB of disk space. An Intel VT or AMD-V 64-bit x86-based system with one or more CPUs is necessary for operating Windows-based VMs [36].

Upon system boot, the installation process commences. XenServer, in contrast to other operating systems, does not permit the construction of custom-sized partitions. XenServer autonomously administers storage disks and file systems. It is advisable to install XenServer on workstations devoid




Figure 2: Cloud Environment

Figure 2: XenServer Management Interface

of essential data or to back up data prior to installation, as the installation process would obliterate any data on the storage drive. The installation procedure necessitates the acceptance of the end-user licensing agreement and the selection of a clean installation option. Select a primary drive for installation from the available local hard disks. We employed a single hard disk that concurrently stores virtual machines. Subsequently, choose the suitable installation supply; we used local media for this objective.

The system prompts for root password entry. This password is used for all operations requiring administrative privileges and for connecting to the server using XenCenter at later stages. The network configuration stage follows [36].

5.2 VMware vSphere Implementation

This section outlines the experimental framework for the installation and configuration of VMware vSphere ESXi. vSphere is accessible at no cost for non-commercial purposes via VMware's online download portal. Nonetheless, the download process mirrors their purchasing approach. We enrolled on their website and asked permission to download vSphere. If the system lacks supported hardware, installation cannot proceed. A compatibility guide is available on their website for consultation before proceeding.

Installation is relatively straightforward, similar to other Linux installations and resembling XenServer installation procedures. Upon installation completion, We can employ the CLI interface on host machines with vSphere installed or utilize the vSphere client for server contact. We selected the vSphere client method. We obtained the vSphere Client via the same portal utilized for the vSphere download. The vSphere Client is a Windows-based program deployed on a separate workstation with specifications comparable to the host machines. Following installation, we accessed the host utilizing the IP address and root password for subsequent server interaction.

6 Experimental Results and Analysis

Following the actions discussed in previous sections, we completed experimental environment setup, installation, and configurations. This section summarizes the experimental configuration used for



Figure 3: Hyper-V Management Console

experiments and evaluations in subsequent sections.

Figure 4 provides a general overview of network virtualization implementation followed in our experimental setup. Through comprehensive experimentation and evaluation, we can conclude several important findings:

- Virtual networks demonstrate superior scalability and security compared to physical networks
- Multiple Virtual Interfaces (VIFs) can be assigned to individual Virtual Machines, with limits depending on underlying hypervisor capabilities
- Network security and traffic isolation can be ensured using VLAN tags or IDs
- VIFs with identical VLAN tags can communicate with each other on virtual networks, ensuring security and isolation
- Hyper-V provides the most advanced VLAN tagging features for network isolation
- Network throughput for TCP and UDP protocols shows minimal differences between VMs with and without VLAN tags
- Round-Trip Time results in non-virtual networks are marginally lower than virtual networks due to hypervisor layer overhead
- CPU allocation shows similar patterns where multiple Virtual Machines share single processors compared to non-virtual networks with no CPU distribution




Figure 4: Experimental Network Architecture

7 Performance Comparison and Discussion

Based on our experimental evaluations, we can summarize performance characteristics across different hypervisors. In TCP throughput measurements using both tagged and untagged VLAN config-

urations, vSphere outperformed other hypervisors, demonstrating optimal performance. XenServer demonstrated performance very close to vSphere with differences of approximately 0.01 MB/s for both configurations.

In terms of UDP throughput with both tagged and untagged VLAN setups, analogous to TCP outcomes, UDP exhibited the poorest performance with Hyper-V hypervisors, whilst XenServer was comparatively near vSphere yet somewhat surpassed it. In TCP CPU utilization with both tagged and untagged VLAN setups, Hyper-V exhibited optimal performance with average CPU usage of 2.49% and 2.14%, respectively, but vSphere displayed the highest resource consumption. XenServer exhibited performance that was halfway between Hyper-V and vSphere for both configurations.

8 Conclusions and Future Work

Our thorough research and experimental findings indicate that network virtualization within cloud computing hypervisors facilitates improved network scalability and security. The capability to assign one or multiple Virtual Interfaces (VIFs) to individual Virtual Machines provides significant operational flexibility, with limitations dependent on underlying hypervisor capabilities. Security and network traffic isolation can be effectively implemented using VLAN tags or identifiers. Tagged VIFs enable communication exclusively among VMs with identical VLAN tags on virtual networks, ensuring both security and isolation while maintaining general communication capabilities. While security advantages gained through VLAN tagging implementation do not negatively impact network performance. Based on our experimental evaluation, VMware vSphere can be considered the optimal hypervisor choice, followed by Citrix XenServer. While vSphere demonstrated superior performance in certain scenarios and XenServer excelled in others, the performance differences between these two hypervisors remain minimal and can be considered negligible. Consequently, both hypervisors can be regarded as nearly equivalent in performance and used interchangeably. For most performance parameters, the impact remains insignificant compared to non-virtualized environments. While virtualized environment setup costs may increase, the benefits obtained significantly outweigh those of non-virtual environments.

References

- [1] Vaquero, L. M., Roderio-Merino, L., Caceres, J., & Lindner, M. (2008). A break in the clouds: Towards a cloud definition. *ACM SIGCOMM Computer Communication Review*, 39(1), 50–55.
- [2] Popa, L., Yu, M., Ko, S. Y., Ratnasamy, S., & Stoica, I. (2010). Cloud police: Taking access control out of the network. In *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks* (p. 7). ACM.
- [3] Cloud Computing. (2014). *Wikipedia*. Retrieved July 4, 2014, from http://en.wikipedia.org/wiki/Cloud_computing
- [4] Angeles, S. (2014). Virtualization vs. Cloud Computing: What's the Difference? *Business News Daily*. Retrieved July 7, 2014, from <http://www.businessnewsdaily.com/5791-virtualization-vs-cloud-computing.html>
- [5] He, Z., & Liang, G. (2012). Research and evaluation of network virtualization in the cloud computing environment. In *Networking and Distributed Computing (ICNDC), 2012 Third International Conference on* (pp. 40–44). IEEE.
- [6] Intel Information Center. (2013). Virtualization and Cloud Computing. Retrieved June 26, 2014, from <http://www.intel.com/content/dam/www/public/us/en/documents/guides/cloud-computing-virtualization-building-private-iaas-guide.pdf>
- [7] Pu, Y., Deng, Y., & Nakao, A. (2012). Cloud rack: Enhanced virtual topology migration approach with open vswitch. In *Information Networking (ICOIN), 2012 International Conference on* (pp. 1–6). IEEE.

-
- [8] Pfaff, B., Pettit, J., Amidon, K., Casado, M., Koponen, T., & Shenker, S. (2009). Extending networking into the virtualization layer. In *Hotnets*.
 - [9] Pettit, J., Gross, J., Pfaff, B., Casado, M., & Crosby, S. (2010). Virtual switching in an era of advanced edges. In *2nd Workshop on Data Center-Converged and Virtual Ethernet Switching (DC-CAVES)*.
 - [10] Chowdhury, N., & Boutaba, R. (2010). A survey of network virtualization. *Computer Networks: The International Journal of Computer and Telecommunications Networking*, 54(5), 862–876.
 - [11] Casado, M., Koponen, T., Moon, D., & Shenker, S. (2010). Rethinking packet forwarding hardware. In *ACM SIGCOMM Computer Communication Review* (Vol. 40, No. 4, pp. 458–459).
 - [12] Tseng, H. M., Lee, H. L., Hu, J. W., Liu, T. L., Chang, J. G., & Huang, W. C. (2011). Network virtualization with cloud virtual switch. In *Parallel and Distributed Systems (ICPADS), 2011 IEEE 17th International Conference on* (pp. 998–1003). IEEE.
 - [13] Dong, Y., Yang, X., Li, J., Liao, G., Tian, K., & Guan, H. (2012). High performance network virtualization with sr-iov. *Journal of Parallel and Distributed Computing*, 72(11), 1471–1480.
 - [14] Soares, J., Carapinha, J., Melo, M., Monteiro, R., & Sargento, S. (2012). Resource allocation in the network operator’s cloud: A virtualization approach. In *Computers and Communications (ISCC), 2012 IEEE Symposium on* (pp. 000800–000805). IEEE.
 - [15] Lee, J., Tourrilhes, J., Sharma, P., & Banerjee, S. (2010). No more middlebox: Integrate processing into network. *ACM SIGCOMM Computer Communication Review*, 40(4), 459–460.
 - [16] Duan, Q. (2011). Modeling and performance analysis on network virtualization for composite network-cloud service provisioning. In *Services (SERVICES), 2011 IEEE World Congress on* (pp. 548–555). IEEE.
 - [17] Duan, Q., Yan, Y., & Vasilakos, A. V. (2012). A survey on service-oriented network virtualization toward convergence of networking and cloud computing. *IEEE Transactions on Network and Service Management*, 9(4), 373–392.
 - [18] Kakadia, D., & Varma, V. (2013). Network virtualization platform for hybrid cloud. In *Cloud Computing Technology and Science (CloudCom), 2013 IEEE 5th International Conference on* (Vol. 2, pp. 69–74). IEEE.
 - [19] Yu, R., Yang, X., Huang, J., Duan, Q., Ma, Y., & Tanaka, Y. (2012). QoS-aware service selection in virtualization-based cloud computing. In *Network Operations and Management Symposium (APNOMS), 2012 14th Asia-Pacific* (pp. 1–8). IEEE.
 - [20] Wang, G., & Ng, T. E. (2010). The impact of virtualization on network performance of amazon ec2 data center. In *INFOCOM, 2010 Proceedings IEEE* (pp. 1–9). IEEE.
 - [21] Hwang, J., Zeng, S., Wu, F., & Wood, T. (2013). A component-based performance comparison of four hypervisors. In *Integrated Network Management (IM 2013), 2013 IFIP/IEEE International Symposium on* (pp. 269–276). IEEE.
 - [22] Moniruzzaman, A. (2014). Analysis of memory ballooning technique for dynamic memory management of virtual machines (VMs). *arXiv preprint arXiv:1411.7344*.
 - [23] Citrix Systems. (n.d.). Introducing XenServer. Retrieved from <http://docs.vmd.citrix.com/XenServer/4.0.1/installation/ch01s03.html>
 - [24] Xen Project. (n.d.). Xen tools. Retrieved from http://wiki.xen.org/wiki/Xen_tools
-

-
- [25] XenServer.org. (n.d.). XenCenter Development. Retrieved from <http://www.xenserver.org/partners/developing-products-for-xenserver/21-xencenter-development/88-xc-dev-home.html>
- [26] Citrix Systems. (2014). XenServer 6.2.0 Technical FAQ. Retrieved from <http://support.citrix.com/article/CTX137836>
- [27] Microsoft Corporation. (2014). Overview of Hyper-V. Retrieved from [https://technet.microsoft.com/en-us/library/cc816638\(WS.10\).aspx](https://technet.microsoft.com/en-us/library/cc816638(WS.10).aspx)
- [28] Nishikiori, M. (2011). Server virtualization with VMware vSphere 4. *Fujitsu Scientific and Technical Journal*, 356–361.
- [29] Qi, L., Chen, Y., Yuan, Y., Fu, S., Zhang, X., & Xu, X. (2020). A QoS-aware virtual machine scheduling method for energy conservation in cloud-based cyber-physical systems. *World Wide Web*, 23(2), 1275–1297.
- [30] Alharbe, N., Rakrouki, M. A., & Aljohani, A. (2022). An improved ant colony algorithm for solving a virtual machine placement problem in a cloud computing environment. *IEEE Access*, 10, 44869–44880.
- [31] TechRepublic. (n.d.). More tools for easy network testing: NTttcp and NetCPS. Retrieved from <http://www.techrepublic.com/blog/data-center/more-tools-for-easy-network-testing-nttcp-and-netcps/>
- [32] Microsoft TechNet Gallery. (n.d.). NTttcp Version 5.28. Retrieved from <https://gallery.technet.microsoft.com/NTttcp-Version-528-Now-f8b12769>
- [33] Xing, H., Zhu, J., Qu, R., Dai, P., Luo, S., & Iqbal, M. A. (2022). An ACO for energy-efficient and traffic-aware virtual machine placement in cloud computing. *Swarm and Evolutionary Computation*, 68, 101012.
- [34] Wireshark Foundation. (n.d.). What is Wireshark. Retrieved from http://www.wireshark.org/docs/wsug_html_chunked/ChapterIntroduction.html#ChIntroWhatIs
- [35] Saxena, D., Gupta, I., Kumar, J., Singh, A. K., & Wen, X. (2021). A secure and multiobjective virtual machine placement framework for cloud data center. *IEEE Systems Journal*, 16(2), 3163–3174.
- [36] Li, R., Zheng, Q., Li, X., & Yan, Z. (2020). Multi-objective optimization for rebalancing virtual machine placement. *Future Generation Computer Systems*, 105, 824–842.
- [37] Huang, Y., Xu, H., Gao, H., Ma, X., & Hussain, W. (2021). SSUR: an approach to optimizing virtual machine allocation strategy based on user requirements for cloud data center. *IEEE Transactions on Green Communications and Networking*, 5(2), 670–681.
- [38] WALboaneen, D., Tianfield, H., Zhang, Y., & Pranggono, B. (2021). A metaheuristic method for joint task scheduling and virtual machine placement in cloud data centers. *Future Generation Computer Systems*, 115, 201–212.
- [39] Bermejo, B., & Juiz, C. (2020). Virtual machine consolidation: a systematic review of its overhead influencing factors. *Journal of Supercomputing*, 76(1).
- [40] Masdari, M., & Zangakani, M. (2020). Green cloud computing using proactive virtual machine placement: challenges and issues. *Journal of Grid Computing*, 18(4), 727–759.
- [41] Dias, A. H., Correia, L. H., & Malheiros, N. (2021). A systematic literature review on virtual machine consolidation. *ACM Computing Surveys (CSUR)*, 54(8), 1–38.
- [42] Potdar, A. M., Narayan, D. G., Kengond, S., & Mulla, M. M. (2020). Performance evaluation of docker container and virtual machine. *Procedia Computer Science*, 171, 1419–1428.
-