



Research Article

A Survey on Performance Metrics in Server Virtualization with Cloud Environment

Vedula Venkateswara Rao and Mandapati Venkateswara Rao

Gitam Institute of Technology, Gitam University, Visakhapatnam

Correspondence should be addressed to: Vedula Venkateswara Rao; venkatvedula2012@gmail.com

Received date: 5 October 2014; Accepted date: 19 February 2015; Published date: 1 June 2015

Academic Editor: KP Lam

Copyright © 2015. Vedula Venkateswara Rao and Mandapati Venkateswara Rao. Distributed under Creative Commons CC-BY 4.0

Abstract

Cloud computing has been the subject of many researches. Researches show that cloud computing permits to reduce hardware cost, reduce the energy consumption and allow a more efficient use of servers. Nowadays lot of servers is used inefficiently because they are underutilized. The uses of cloud computing associate to virtualization have been a solution to the under utilization of those servers and also Data Centers of today are rapidly moving towards the use of server virtualization as a preferred way of sharing pool of server hardware resources between multiple guest domains that hosts different applications. However the virtualization performances with cloud computing cannot offer performances equal to the native performances. The Hypervisors of the virtualized servers, such as Xen use fair schedulers to schedule the guest domains, according to pointers or weights assigned to the domain by administrators. The Hosted application performance is sensitive to the scheduling parameters of the domain on which the application runs. However exact relationship between the parameters of the domain and application performance measures such as response time or through put is not obvious and not static as well, further more due to dynamics present in the system there is a need for tuning of the scheduling parameters. Recently virtual machines have emerged as the basis for allocating resources in enterprise settings and hosting centers. This paper studies various issues in analyzing the performance of Server Virtualization. To be able to meet this aim it has been reviewing at first the previous researches on this area. The aim of this paper is to identify various parameters and issues for evaluating performance of virtualization in cloud computing environment in terms of CPU, memory and I/O devices performances. The paper also studies various issues to understand effectiveness of existing Quality of Service (QoS) controls on resource usage and thereby application performance.

Keywords: Cloud computing, Virtualization, Hypervisor, Virtual Machine, CPU performance, Memory Performance, I/O Performance, Quality of service.

Introduction

Within Information Technologies (IT)

everything is moving very fast and technologies progress in an exponential manner according to Moore's law: "The

complexity for minimum component costs has increased at a rate of roughly a factor of

two per year... Certainly over the short term this rate can be expected to continue, if not to increase" (Moore, 1965). It has been reported that a large majority of server are underutilized. In a research by Anton Beloglazov et al. (2012), Abirami et al., (2012), Bao Rong Chang et al., (2013) and Andrew J. Younge (2013) identified Virtualization and also cloud computing has become the solution to increase the efficiency of servers utilization. Dell has conducted some researches and has found that three quarters of the servers that they were managing globally were not exceeded 20 percent of processor utilization. The processor utilization of those servers low because only one or two applications were running on those servers (Brooks, 2008). To overcome the problem of under utilization of server's solutions such as virtualization and cloud computing have been deployed.

Background

Cloud computing regroups the use of different technologies such as virtualization, clustering, and so on. This technology begins to become really popular because it offers to the customers the choice to pay only for what they use. If a company needs an important computing power for few days they can buy it for few days on the cloud rather than to by machine especially for that purpose. Cloud computing has permitted to optimize server utilization due to the use of virtualization. And with the apparition of processor with multiple cores and virtualization support few years ago virtualization has gain popularity.

Cloud computing is an on demand service able to share resources easily over the internet. It exists different models of cloud that can be used for different usages.

- Public cloud – Manages by a third party service provider to offer services accessible by clients through the internet.

- Private cloud – It is basically the same as public cloud but inside the company network. It loses some

benefits of public cloud but it has the advantages to be more secure.

- Hybrid cloud – It is a combination of public and private cloud.

- Community cloud – It is a cloud shared by multiple organizations.

All of those models can implement virtualization. For implementing those different models multiple solutions exist. Eucalyptus and OpenNebula are open source solutions available for implementing cloud computing. The different solutions implementing cloud computing don't have the same structures. It can be though that the management of the virtualization is different between the different solutions. Cloud computing doesn't necessarily use virtualization. Different virtualization solution can be implemented within a cloud. The most popular solutions are Xen, Vmware and KVM. Xen and KVM are two open source solutions for virtualization contrary to Vmware which is a proprietary solution. However the utilization of virtualization within cloud computing provides the virtualization and sharing of physical resources within the cloud allow higher utilization rates and reduce in the same time the hardware investment, save space and reduce power consumption. The usage of virtualization introduces a degradation of performances because it introduces additional overhead. Virtualization affects CPU usage, network, memory and storage performances as well as applications performances (Armstrong, 2007). Within virtualization great performances depend essentially of the tasks scheduling and the workload on the system. Bao Rong Chang et al., (2013) in an experiment discovered how to improve I/O performances by the scheduling of virtual machine tasks. It can be found at the fifth place "performance Unpredictability". It is explained in the article

that shared CPU and memory are pretty well managed by virtual machines rather than network and disk I/O sharing which are not. Virtualization in the cloud can meet some I/O interference between virtual machines. The scheduling tasks is an important part of the management of virtualization in cloud computing.

Literature Review

Cloud computing is a recent technology and lot of researches are made in that domain to improve it. Also due to the relation between cloud and virtualization there are as well many researches on virtualization to enhance virtualization performances. Cloud computing is more and more popular and most of the enterprise begin to adopt it. However there are still some obstacles which can restrained the adoption of cloud services by enterprise such as the lack of standardization, reliability associate to the cloud, the security and so on. Qingling Wang et al (2011), Indrani Paul et al., (2010), Sajib et al., (2010) and Xiao Zhang (2013) in their research mentioned that the reason of the adoption of cloud computing by enterprise is principally for economical reasons because cloud computing allows customers to reduce their hardware cost as well as energy consumption and so on. Also there is no waste because customers only pay for what they are using. Because of the popularity of cloud computing a large number of cloud toolkits begin to appear such as Open Nebula, Eucalyptus and so on. Also those cloud toolkits are able to associate with different hypervisors such as Xen, KVM and VMware ESX. As seen previously there are many different types of virtualization. To be able to provide the best performances cloud computing is using para- virtualization as well as hardware-assisted virtualization. Full virtualization is not used in cloud computing due to poor performances cause by its considerable overhead. Virtualization technology is not a new technology however it has regained popularity in 2005 with the apparition of AMD and Intel processors which had support for virtualization.

Virtualization brings many advantages such as the improvement of security, the enhancement of the efficiency of server utilization and so on.

Also during the past few years due to the popularity of virtualization and its utilization in the cloud computing many researchers have been made. From those researches lot of improvements have been made to try to obtain performances near to native performances. In the following it will be outline in a first part the research that have been made on cloud computing. It will be explained the problems due to the lack of standardization in the cloud as well as the economic impact that can have the usage of the cloud and the major obstacles that can be encounter to the growth of cloud computing. Finally it will be compared different solutions of cloud toolkits. In a second part it will be compared the two different types of Virtual Machine Monitor (VMM). Then it will also be compared the native performances and virtual performances. It will end with a comparison of the different hypervisor for both commercial and open source product.

Cloud Computing

Cloud computing is a new technology and evolve rapidly also it is difficult to match a good definition of cloud computing. Because cloud computing is an evolving technology the definition is changes over the time. The U.S. Government's National Institute of Standards and Technology (NIST) tries to give an up to date definition of the cloud computing. According to the NIST cloud computing is on demand service which shares a pool of computer resources over a network. Cloud computing matches five essential characteristics which define the main functionalities provided by the cloud, three service models which give the level of service provided and four deployment models which indicate where the cloud is deployed and who can access to it.

The main characteristics of the clouds are the following (Mell et al., 2009):

- On demand self-service: Users of the cloud can manage the resources in on demand basis and they only paid for what they consume.
- Broad network access: The resources provided by the cloud can be access by as any normal services through thin or thick clients such as laptop, PDA, mobile phones and so on.
- Resource pooling: The cloud provider serves pool of resources over multiple customers according to the demand. Client which access the service have no knowledge of the exact location of the cloud but may be able to provide a location at higher abstraction level such as country, state, datacenter and so on.
- Rapid elasticity: The resources provided by the cloud are highly scalable. Customer can rapidly scale up the resources that they need and then scale them down if there is no need to use it anymore. The scalability of the cloud gives a real modularity to the cloud. Also resources appear as infinite and customers have no need to make plan for provisioning.
- Measured service: The resources provided by the cloud are controlled and optimized according to the resources capabilities. Also resources usage can be monitored control and reported to be able to provide transparency for both
-

provider and consumer of the resources.

Because of the diversity of the cloud toolkits available and also the diversity of hypervisor there is an important problem of interoperability. The open cloud manifesto tries to establish a set of core principles for cloud providers. There are six principles that are defined by the cloud manifesto which are:

- Cloud providers must work together
- Cloud providers must use and adopt the existing standards
- New standards need to promote innovation
- Community effort should be driven by customer needs
- Cloud computing standard organizations should work together
- Cloud provider must not use their popularity to lock the customer to their platform. Another group called Cloud Computing Use Case group also work to brought together cloud providers and cloud consumer.

Cloud Architecture

The Most significant components of cloud are the Front end and Back end. The Front end is the client side, the User of the applications hosted in cloud. This covers the application used to access the cloud through interface like web browser and client's network. Back end is actually cloud itself with various servers, storage devices, network etc. Theoretically a cloud has all types of dedicated servers present in it to host practically any program. A central Server administers the system, monitoring the traffic and client demands to ensure anything runs smoothly. There are set of rules and middleware which help in networked computers communicated to each other.

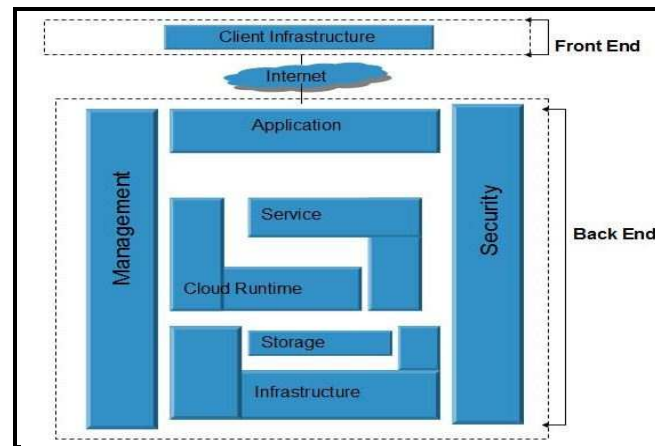


Figure 1:- Cloud Architecture

System Virtualization

Server virtualization is the masking of server resources, including the number and identity of individual physical servers, processors, and operating systems, from server users. Server virtualization is the masking of server resources, including the number and identity of individual physical servers, processors, and operating systems, from server users. The server administrator uses a software application to divide one physical server into multiple isolated virtual environments. The virtual environments are sometimes called virtual private servers, but they are also known as guests, instances, containers or emulations. Virtualization is a method of running multiple independent virtual operating systems on a single physical computer. It is a way of maximizing physical

resources to maximize the investment in hardware. This is described using fig 4.1.

There are three popular approaches to server virtualization: the virtual machine model, the para virtual machine model, and virtualization at the operating system (OS) layer.

Server virtualization offers a way to consolidate servers by allowing you to run multiple different workloads on one physical host server. A "virtual server" is a software implementation that executes programs like a real server. Multiple virtual servers can work simultaneously on one physical host server. Therefore, instead of operating many servers at low utilization, virtualization combines the processing power onto fewer servers that operate at higher total utilization as shown in following figure.




Figure 2:- Server Virtualization

Virtual Machine Model

Virtual machines are based on the host/guest paradigm. Each guest runs on a virtual imitation of the hardware layer. This approach allows the guest operating system to run without modifications. It also allows the administrator to create guests that use different operating systems. The guest has no knowledge of the host's operating system because it is not aware that it's not running on real hardware. It does, however, require real computing resources from the host, so it uses a hypervisor to coordinate instructions to the CPU. The hypervisor is called a virtual machine monitor (VMM). It validates all the guest-issued CPU instructions and manages any executed code that requires additional privileges. VMware and Microsoft Virtual Server both use the virtual machine model.

Para Virtualization

Although virtualization is rapidly becoming mainstream technology, the concept has attracted a huge amount of interest, and enhancements continue to be investigated. One of these is *para-virtualization*, whereby operating system compatibility is traded off against performance for certain CPU-bound applications running on systems without

virtualization hardware assist as described by Fig 3.

The para virtual machine (PVM) model is also based on the host/guest paradigm and it uses a virtual machine monitor too. In the para virtual machine model, however, the VMM actually modifies the guest operating system's code. This modification is called porting. Porting supports the VMM so it can utilize privileged systems calls sparingly. Like virtual machines, para virtual machines are capable of running multiple operating systems. Xen and UML both use the para virtual machine model. The para-virtualized model offers potential performance benefits when a guest operating system or application is 'aware' that it is running within a virtualized environment, and has been modified to exploit this. One potential downside of this approach is that such modified guests cannot ever be migrated back to run on physical hardware. In addition to requiring modified guest operating systems, para virtualization leverages a hypervisor for the underlying technology.

In the case of Linux distributions, this approach requires extensive changes to an operating system kernel so that it can coexist with the hypervisor. Accordingly,

mainstream Linux distributions (such as Red Hat or SUSE) cannot be run in a para virtualized mode without some level of

modification. Likewise, Microsoft has suggested that a future version of the Windows operating system will be developed that can coexist with a new hypervisor offering from Microsoft.

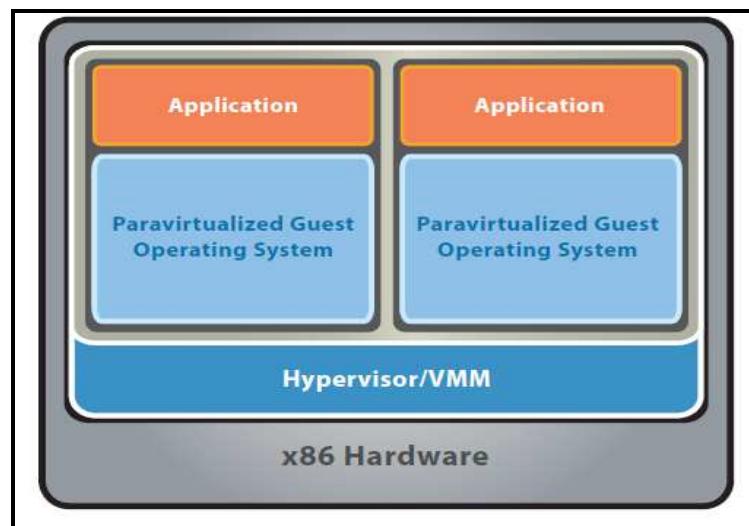


Figure 3:- Para Virtualization

Virtualization at Operating System Level

Virtualization at the OS level works a little differently. It isn't based on the host/guest paradigm. In the OS level model, the host runs a single OS kernels its core and exports operating system functionality to each of the guests. Guests must use the same operating system as the host, although different distributions of the same system are allowed. This distributed architecture eliminates system calls between layers, which reduces CPU usage overhead. It also requires that each partition remain strictly isolated from its neighbors so that a failure or security breach in one partition isn't able to affect any of the other partitions. In this model, common binaries and libraries on the same physical machine can be shared, allowing an OS level virtual server to host thousands of guests at the same time. Virtuozzo and Solaris Zones both use OS-level virtualization.

Server Consolidation

Virtualization may be used for server consolidation. Figure 4.3(a) shows the typical architecture of an e-commerce site with separate machines for Web servers, applications servers, and database servers. Figure 4.3(b) shows these servers being consolidated on a Single, albeit more powerful, server. Virtual infrastructure initiatives often spring from data center *server consolidation* projects, which focus on reducing existing infrastructure "box count", retiring older hardware or life-extending legacy applications. Server consolidation benefits result from a reduction in the overall number of systems and related recurring costs (power, cooling, rack space, etc.) While server consolidation addresses the reduction of existing infrastructure, *server containment* takes a more strategic view, leading to a goal of infrastructure unification. Server containment uses an incremental approach

to workload virtualization, whereby new projects are provisioned with virtual

machines rather than physical servers, thus deferring hardware purchases.

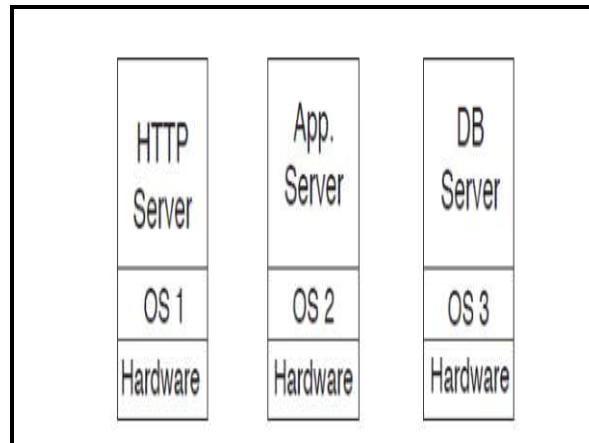


Figure 4:-(a)

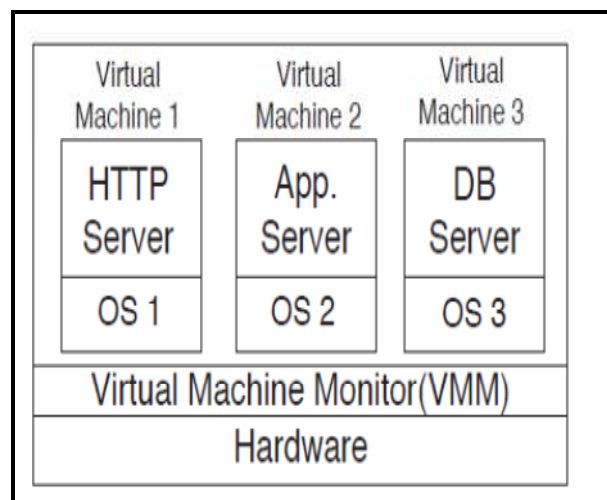


Figure 4:-(b) Server Consolidation Example

CPU Utilization and Performance

Virtual Machines will increase the number of context switches and also consume considerable CPU for address translation. Monitoring CPU utilization shows whether all the CPUs on the host are completely occupied. The following are key metrics in evaluating CPU performance in Virtualization environment. CPU utilization is a key

performance metric. It can be used to track CPU performance improvements, and is a useful data point for performance problem investigations. CPU utilization has critical implications on other system performance characteristics, namely power consumption. We may think the magnitude of CPU utilization is only important if you're bottlenecked on CPU at 100%, but that is not at all the case. Each additional % of CPU

Utilization consumes a bit more. In the performance evaluation the virtual machine, threads that are scheduled by a virtual machine instead of natively by the underlying operating system. Threads emulate multithreaded environments without relying on any native OS capabilities, and they are manageable in user space instead of kernel space, enabling them to work in environments that do not have native thread support. Performance On a multi-core processor, native thread implementations can automatically assign work to multiple processors, threads can be started much faster on some VM. Also, a thread may block all other threads if performing a blocking I/O operation. To prevent the problem, threads must use asynchronous I/O operations, although the increased complexity can be hidden by implementing separate native I/O processes which cooperate with threads.

- CPU Usage
- Total-all CPUs
- Per CPU
- Delta between CPUs

The CPU performance can be measured as a number of virtual machines running on the node.

Memory Performance

The Virtual Machines themselves will consume a small portion of RAM to maintain shadow page tables required to map virtual addresses to physical addresses. RAM plays crucial part in performance optimization by reducing the number of physical I/O requests. With multiple virtual machines managing their own foot print, the amount of memory available to the applications will reduce and will adversely impact the performance. Memory is also a limiting factor in scaling out the virtual machines and generally decides how many virtual machines can be hosted on a physical server simultaneously.

Memory is typically the most constricted host resource, and the most likely to begin suffering bottlenecks if proactive monitoring is not in place. There are many factors that go

into memory management making this a complicated resource area to monitor for optimal performance allocation. At the time of writing, there is a lack of metrics which reveal sufficient detail to derive optimal memory allocation for a VM. Yet, proactive monitoring of existing metrics can help easily spot memory problems, or can serve as warning signals for impending problems.

The following are different metrics in evaluating memory performance in virtual machine environment.

- Memory Balloon
- Memory Heap
- Memory Heap Free
- Memory Reserved Capacity
- Memory Shared
- Memory Consumed
- Memory Active
- Memory Granted
- Memory Usage
- Memory Swap Used
- Memory Sway In
- Memory Swap Out
- Memory Used by VMkernel
- Memory Unreserved
- Memory State

Storage utilization and I/O Performance

Since virtual machines host their own complete guest operating system, the size of typical virtual machine, applications running on it and data might run into GBs and requires large amount of disk storage. Significant amount of I/O activity will be involved to keep multiple virtual machines live on a single server. Applications like Data base servers require the high amount of I/O amount might face significant bottlenecks if I/O subsystems are not able to cope with the increased load.

Storage is a key resource for a virtualized environment. After all, a VM is technically a VMDK file, which can become a large data file as an operating system and other supporting software is included. Virtualization has made provisioning servers much easier and cheaper leading to a proliferation of VMs, and thus data creation in the data center. Also with additional data objects related to VMs

such as snapshots and VM templates, each VM is often the cause of several times more storage usage than just its own VMDK file. Also important to note, different VMs will grow their storage needs at different rates based on usage levels and user activities. Free storage availability is particularly important to track as without enough storage, a VM will crash.

The following are key metrics in evaluating storage performance in virtual machine environment.

- Disk Usage
- Disk I/O Latency Levels
- Total free
- Total used

Research Highlights

- To monitor and study various performances issues in Server Virtualization and Network Virtualization.
- To Study the optimization process and various issues in analyzing the performance of Server Virtualization.
- To identify various parameters and issues for evaluating performance of virtualization in cloud computing environment in terms of CPU, memory and I/O devices performances.
- To study various issues to understand effectiveness of existing Quality of Service (QoS) controls on resource usage and thereby application performance.
- The goal is to consolidate the Data Center and increase its performance.

Conclusion

Server virtualization is always necessary for optimizing the IT Infrastructure. In this paper we have described about the architecture and several metrics (parameters) in evaluating the performance of Server Virtualization. This technology will become increasingly important in the future and be introduced for widespread as a

common technology for building the IT infrastructure.

References

1. Abirami S.P. and Shalini Ramanathan (2012), 'Linear Scheduling Strategy for Resource Allocation in Cloud Environment', International Journal on Cloud Computing: Services and Architecture (IJCCSA), Vol.2, No.1, February 2012.
2. Andrew J. Younge, Robert Henschel, James T. Brown, Gregor von Laszewski, Judy Qiu, Geoffrey C. Fox(2011), 'Analysis of Virtualization Technologies for High Performance Computing Environments', Pervasive Technology Institute, Indiana University 2729 E 10th St., Bloomington, IN 47408, U.S.A. {ajyounge,henschel,jatbrown,gvonlasz,xqiu,gcf}@indiana.edu, Proceedings of the IEEE 4th International Conference on Cloud Computing (CLOUD), 9-16.
3. Andrea Arcangeli, Izik Eidus and Chris Wright, 'Increasing memory density by using KSM', Red Hat, Inc. aarcange@redhat.com, ieidus@redhat.com, chrisw@redhat.com.
4. Anton Beloglazov, Jemal Abawajy, Rajkumar Buyya (2012), 'Energy-aware resource allocation heuristics for efficient management of Data Centers for Cloud Computing', Future Generation Computer Systems, Elsevier, also available at Science Direct, Sponsored by Cloud Computing and Distributed Systems(CLOUDS)'.
5. 'Application Performance Management in a Virtualized Environment Growing', WHITE PAPER: APPLICATION PERFORMANCE MANAGEMENT.
6. Aravind Menon, Jose Renato, Yoshio Turner, G. (John) Janakiraman, Palo Alto john, Willy and Zwaenepoel(2005), 'Diagnosing Performance Overheads in the Xen Virtual Machine Environment', VEE'05, June 11-12, 2005, Chicago, Illinois, USA. Copyright2005ACM1-59593-047-7/05/0006...\$5.00.

7. Bao Rong Chang, Hsiu-Fen Tsai, Chi-Ming Chen(2013), 'Evaluation of Virtual Machine Performance and Virtualized Consolidation Ratio in Cloud Computing System', Journal of Information Hiding and Multimedia Signal Processing c 2013 ISSN 2073-4212 Ubiquitous International Volume 4, Number 3, July 2013.
8. ' Better virtualization of XenApp and XenDesktop with XenServer", XenApp and XenDesktop with XenServer White Paper.
9. Bhukya, D.P. ; Ramachandram, S. ; Reeta Sony, A.L, 'IO Performance Prediction in Consolidated Virtualized Environments".
10. B.Thirumala Rao, N.V.Sridevi, V.Krishna Reddy and L.S.S.Reddy (2011), 'Performance Issues of Heterogeneous Hadoop Clusters in Cloud Computing', Global Journal of Computer Science and Technology Volume XI Issue VIII May 2011 .
11. Diego Ongaro Alan L. Cox Scott Rixner(2008), ' Scheduling I/O in Virtual Machine Monitors', VEE'08, March 5-7, 2008, Seattle, Washington, USA. Copyright c 2008 ACM 978-1-59593-796-4/08/03...\$5.00.
12. Diego Ongaro Alan L. Cox Scott Rixner(2008), ' Scheduling I/O in Virtual Machine Monitors", VEE'08, March 5-7, 2008, Seattle, Washington, USA. Copyright c 2008 ACM 978-1-59593-796-4/08/03...\$5.00.
13. Diogo M. F. Mattos and Lyno Henrique G. Ferraz(2012), ' Virtual Network Performance Evaluation for Future Internet Architectures, JOURNAL OF EMERGING TECHNOLOGIES IN WEB INTELLIGENCE, VOL. 4, NO. 4, NOVEMBER 2012.
14. Diwaker Gupta1, Ludmila Cherkasova, Rob Gardner, Amin Vahdat(2006), ' Enforcing Performance Isolation Across Virtual Machines in Xen', Enterprise Software and Systems Laboratory HP Laboratories Palo Alto HPL-2006-77 May 4, 2006*.
15. 'Experimental Evaluation of the Performance-Influencing Factors of Virtualized Storage Systems', Computational Intelligence and Computing Research (ICCIC), 2010 IEEE International Conference on Digital Object Identifier: 10.1109/ICCIC.2010.5705753, Publication Year: 2010 , Page(s): 1 - 4.
16. Himanshu Raj, Ripal Nathuji(2009), ' Resource Management for Isolation Enhanced Cloud Services", CCSW'09, November 13, 2009, Chicago, Illinois, USA. Copyright 2009 ACM 978-1-60558-784-4/09/11 ...\$10.00.
17. Horacio GonAlez Velez, Maryam Kontagora (2011), 'Performance evaluation of mapreduce using full virtualization on a departmental cloud', International Journal of Applied Mathematics and Computer Science, Vol. 21, No. 2, 275-284 DOI: 10.2478/v10006-011-0020-3(AMCS).
18. Indrani Paul, Sudhakar Yalamanchili, Lizy K. John(2012), 'Performance Impact of Virtual Machine Placement in a Datacenter", 2012 IEEE 31st International Performance Computing and Communications Conference (IPCCC), 1-3 Dec. 2012, Austin, TX, Page(s):424 - 431 , ISSN :1097-2641,Print ISBN:978-1-4673-4881-2, DOI:10.1109/PCCC.2012.6407650 .
19. Nikolaus Huber, Marcel von quast, Michael Hauck, Samuel Kounev(2011), 'Evaluating and modeling virtualization performance overhead for cloud environments', CLOSER 2011 - Proceedings of the 1st International Conference on Cloud Computing and Services Science, Noordwijkerhout, Netherlands, 7-9 May, 2011. SciTe Press 2011 ISBN 978-989-8425-52-2.
20. Nikolaus Huber, Marcel von Quast, Michael Hauck, Samuel Kounev(2013) ; 'Evaluating and modeling virtualization Performance overhead for cloud environments', *Journal of Information Hiding and Multimedia Signal Processing* 2013 ISSN

21. 2073-4212, Ubiquitous International Volume 4, Number 3, July 2013.
22. Paul Barham*, Boris Dragovic, Keir Fraser, Steven Hand, Tim Harris, Alex Ho, Rolf Neugebauer†, Ian Pratt and Andrew Warfield(2003), 'Xen and the Art of Virtualization', SOSP'03, October 19–22, 2003, Bolton Landing, New York, USA. Copyright 2003 ACM 1-58113-757-5/03/0010 ...\$5.00.
23. Sajib Kundu, Raju Rangaswami, Kaushik Dutta, Ming Zhao(2010), ' Application Performance Modeling in a Virtualized Environment", School of Computing & Information Sciences, College of Business Administration Florida International University {skund001, raju}@cs.fiu.edu kaushik.dutta@business.fiu.edu,zhaom@cs.fiu.edu.
24. Shicong Meng, Ling Liu(2013), 'Monitoring-as-a-Service in The Cloud', ICPE'13, April 21–24, 2013, Prague, Czech Republic. ACM 978-1-4503-1636-1/13/04.
25. Sriram Govindan, Jeonghwan Choi, Arjun R. Nath, Amitayu Das(2009), ' Xen and Co.: Communication-Aware CPU Management in Consolidated Xen-Based Hosting Platforms', 2009 IEEE Published by the IEEE Computer Society.
26. Waldspurger and Carl A(2002), 'Memory Resource Management in VMware ESX Server', VMware, Inc. Palo Alto, CA 94304 USA carl@vmware.com, Proceedings of the 5th Symposium on Operating Systems Design and Implementation.
27. Wang, Carlos A. Varela (2011), ' Impact of Cloud Computing Virtualization Strategies on Workloads 'Performance" Qingling Department of Computer Science Rensselaer Polytechnic Institute, Troy, NY, USA <http://wcl.cs.rpi.edu/> {wangq9, cvarela}@cs.rpi.edu, *UCC, page 130-137. IEEE Computer Society, (2011).*
28. Xiao Zhang Eric Tune Robert Hagmann Rohit Jnagal Vrigo Gokhale John Wilkes (2013), ' CPI²: CPU performance isolation for shared compute clusters", Google, Inc, 2013 ACM 978-1-4503-1994-2/13/04. \$15.00
29. Yiduo Mei , Ling Liu , Xing Pu and (2010), 'Performance Measurements and Analysis of Network I/O Applications in Virtualized Cloud', 2010 IEEE 3rd International Conference on Cloud Computing (CLOUD), 5-10 July 2010 Page(s): 59 – 66, E-ISBN: 978-0-7695-4130-3, Print ISBN: 978-1-4244-8207-8, INSPEC Accession Number: 11499480.
30. Zongjian He, Guanqing Liang(2012), 'Research and Evaluation of Network Virtualization in Cloud Computing environment', , IEEE Third International Conference on Networking and Distributed Computing (ICNDC), 2012 at Hangzhou,40-45, ISSN :2165-5006,Print ISBN:978-1-4673-2858-6,INSPEC Accession Number:13263829,Digital Object Identifier :10.1109/ICNDC.2012.18.