

# Performance Evaluation of Private Clouds Eucalyptus versus CloudStack

Mumtaz M.Ali AL-Mukhtar  
Department of Internet Engineering  
AL-Nahrain University  
Baghdad, Iraq

Asraa Abdulrazak Ali Mardan  
Department of Networks Engineering  
AL-Nahrain University  
Baghdad, Iraq

**Abstract**—the number of open source cloud management platforms is increasing day-by-day. The features of these software vary significantly and this creates a difficulty for the cloud consumers to choose the software based on their business and scientific requirements. This paper evaluates *Eucalyptus* and *CloudStack*, the two most popular open source platforms used to build private *Infrastructure as a service* (IaaS) clouds. The performance of virtual machines (VMs) initiated and managed by *Eucalyptus* and *CloudStack* are evaluated in terms of CPU utilization, memory bandwidth, disk I/O access speed, and network performance using suitable benchmarks. Different VM management operations such as add, delete and live migration are also assessed to determine which cloud solution is more suitable than other to be adopted as a private cloud solution. As a further performance testing, a simple web application has been implemented on the both clouds to evaluate their suitability in web application hosting.

**Keywords**—*Cloud Computing; CloudStack; Eucalyptus; IaaS; Virtual Machine; Performance Evaluation*

## I. INTRODUCTION

Cloud computing as a new Internet service concept has become popular to provide a variety of services to users. It is a combination of technologies that have been developed over the last several decades, which includes virtualization, dynamic provisioning, internet delivery of services, grid computing, cluster computing and utility computing [1][2]. According to NIST (National Institute of Standards and Technology), "Cloud Computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction" [3].

There are three deployment models by which Cloud computing services are delivered: public, private, and hybrid. Public Cloud is a cloud that is made available as —pay-as-you-go and accessible to the general public such as Amazon Web Services. Private Cloud refers to a cloud infrastructure that is internal to an organization and is not available to the general public. A private cloud's data centers can be on premise and the physical infrastructure is owned and managed by the organization that owns it [4]. Hybrid cloud is a composition of two or more cloud deployment models that are

bound together by a standardized or proprietary technology [4].

There are certain legal, political, socio-organizational reasons that may discourage an organization from using public cloud infrastructure for certain kinds of activities, for example processing and storing citizens' private data. There is also the issue of privacy, security, location and ownership of data [4]. Many companies hesitate to use public cloud in which computing resource are shared with other companies. These companies do not have any knowledge of where their applications are run and their data are stored or control access to them [1]. Hence, private cloud infrastructure is considered an appropriate alternative.

Another big reason to increase the interest in setting up and managing private cloud is the SLA. The public cloud providers nowadays provide guarantees on their service levels and when service failures occur, they only offer to refund their customers regarding the infrastructure outages. However, service providers are not inclined to pay penalties of low performance level that would refund customers for loss of business revenue. Cloud providers are not only required to supply correct services but, also, to meet their expectations in the context of performance [5]. Also some software systems and applications require different performance levels, quality of services, reliability, and security, which are generally not guaranteed by a public cloud. Private cloud is an alternative to companies or researchers that need more control over that data [1] [6].

There are many commercial and open source cloud management platforms that are used to build Infrastructure as a service (IaaS) private cloud solution such as *Eucalyptus*, *OpenNebula*, and *Vmware* cloud. However the open source solutions are gaining a lot of popularity and momentum with their features, rapid developing with low investment cost which present a viable option for academic and scientific worlds [7], and enterprises who want first to test the cloud computing solution suitability to their business environment before purchase the thousands dollars commercial solution.

The number of cloud platforms related to a private IaaS cloud is increasing day-by-day. The features of cloud management software vary significantly and this creates a difficulty for cloud consumers to choose the software based on their business requirements.

An example for this problem is choosing platform much suitable for hosting web applications or running high performance computing (HPC) applications, or meeting specific user usage way like users that demand a few virtual machines (VMs) but want to run them for a long period of time with guarantees on high-availability, or scientists requiring a large number of resources to conduct actual calculations and analyses of data. The advent of several Open Source Cloud platforms guarantees the performance and uptime. It is not easy for non-expert users to choose from the different platforms without comprehending the characteristics and advantages of each of this platform [6].

As a consequence, performance evaluation of cloud computing platforms has been receiving considerable attention by both the users and service providers as a prominent activity for exploring the limitations of the cloud platforms and improving service quality, infrastructure planning, and making a wiser selection of the platforms. In addition cloud management software vendors can develop and include additional features to their software by fixing the platforms bugs and including the missing features.

The rest of this paper is organized as follows: Section II presents related work. Section III describes the test environment and methodology. Section IV covers the performance evaluation of Eucalyptus and CloudStack VMs. Section V assesses VMs startup and release time. Section VI evaluates live migration of VMs. Section VII presents response time of web application in the both clouds. Finally conclusions are drawn in the last section.

## II. RELATED WORK

Many studies have been conducted to evaluate performance of open source cloud platforms such as Eucalyptus, Opennebula and Nimbus. However these research papers did not perform a complete performance analysis of the cloud platform, and compare only the architectures and features of the cloud management platforms. Nevertheless a little work has been done yet to evaluate CloudStack due to the fact that it is relatively new.

De Sousa et al. [1] evaluated Eucalyptus VMs considered processing and disk I/O performance only while in [6, 7, 8, 9, 10, 11] authors brought out an overview of architectures of open source platforms and comparison of their general features and Characteristics. Mao and Humphery [12] investigated the performance of VM startup and release time of public clouds. However, D. Steinmetz, et al. [13] evaluated performance and studied VM launch time of Eucalyptus and OpenStack but performance benchmarking was not specific and gave a general view of performance. While Folgar, et al. [14] evaluated performance of CloudStack primary storage disk I/O only.

Differently from previous works, this paper evaluates performance of Eucalyptus and CloudStack clouds VMs covering versatile parameters including performance of cloud management platform considering add, delete and live migration of VMs. Performance of VMs in term of CPU utilization ,memory bandwidth, disk I/O speed and networking performance is rated as key point of our evaluation. Also the

performance of VMs are compared with regard to bare-metal or traditional IT infrastructure.

## III. TEST ENVIRONMENT AND METHODOLOGY

CloudStack 4.1 cloud with one zone, pod and cluster has been deployed using 3 identical physical servers. One server is used as a management server including primary and secondary storage and the other two servers are used as host machines. Eucalyptus 3.2 cloud with one cluster has been deployed using 3 identical physical servers each. One server is used as a cloud controller (CLC) including cluster controller (CC) and Walrus storage and the other two servers are used as node controllers (NCs). Our servers are Intel<sup>R</sup> Core<sup>TM</sup> i5-2410M CPU 2.3GHz, 4GB RAM, 500GB SATA Hard Disk and 100MB Ethernet interface. Centos 6.3 (final) is installed on each server as native OS. CloudStack with NFS storage configuration is deployed while Eucalyptus is deployed with local storage configuration. Each host in both clouds is configured with kernel-based virtual machine (KVM) as a hypervisor.

In order to evaluate and analyze VMs performance of both clouds, we have employed a number of benchmarks each for different evaluation purpose. Table I shows the selected benchmarks.

A customized CloudStack template (image used to establish VM) and Eucalyptus VM image have been created in which all benchmarks are installed and configured to save time and ease of work.

Each benchmark test is repeated five times consequently and the average of results is considered. Different numbers and types of VM are regarded in the performance evaluation. In each cloud the same VM type is used and the same OS is run which it Centos 6.3. Moreover, each cloud is built with similar hardware and uses the same hypervisor (KVM) to achieve a fair comparison between Eucalyptus and CloudStack and eliminate virtualization and hardware differences that may affect evaluation. Table II shows types of VMs that are provided by Eucalyptus and CloudStack.

TABLE I. BENCHMARKS DEPLOYED FOR VM TESTING

No.	Benchmarks	Testing Resource
1	LINPACK	Processor
2	Bonnie++	Disk I/O
3	STREAM	Memory Bandwidth
4	Iperf	Network
5	Lookbusy	Processor
6	UnixBench	Overall System

TABLE II. OFFERING VM TYPES

Type	RAM	CPU core	Disk(GB)
Small	512M	1	10
Medium	1G	1	20
Large	2G	2	40
XLarge	4G	4	60

#### IV. PERFORMANCE EVALUATION OF VMS

VMs of both clouds have been evaluated using selected benchmarks considering different relating metrics. The VM performance has not been compared just between Eucalyptus and CloudStack, but it also has been compared in regard to bare-metal or traditional IT infrastructure.

##### A. Comparison with Traditional IT infrastructure

The first question that comes in the mind of the cloud users or organization that plan to adopt the cloud computing solution is that "does the cloud virtual machine performance is the same as traditional physical machine?". To answer this question, the performance of both machine with same hardware and software is tested using the same benchmark that mimic the real workload. UnixBench benchmark has been run with the traditional hardware stack on the host server of both cloud, then is run on both Eucalyptus and CloudStack Cloud on a single VM utilizing the whole host server resources.

As shown in figures 1 and 2, the performance of Eucalyptus cloud VM is nearly the same as physical one while there is a 7% gain in performance of the CloudStack VM. This result suggests that the cloud computing management system exploits or utilizes the computing resources on the same hardware stack better than the bare-metal or traditional IT system.

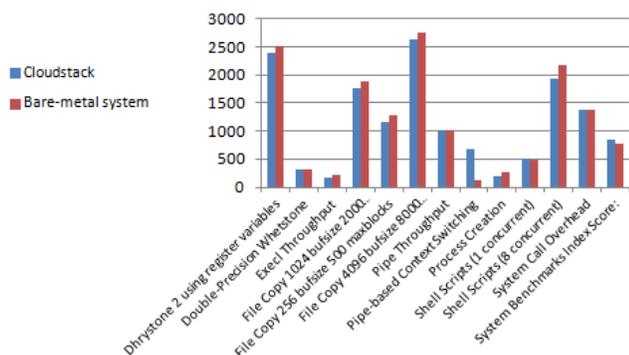


Fig. 1. Performance of CloudStack vs. Bare-metal system

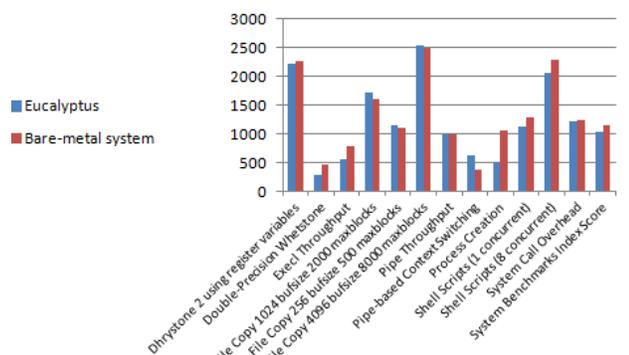


Fig. 2. Eucalyptus VM vs. Bare-metal system

##### B. Processing Performance

The Eucalyptus and CloudStack VM computing power has been assessed to test its ability in running a HPC (High

performance Computing) applications. LINPACK is a benchmark that measures a computer's floating-point rate of execution by solving a dense  $n$  by  $n$  system of linear equations in double precision. Gflop/s is the rate of execution; it refers to billions of floating point operations per second.

In this test three scenarios have been applied. First two types of VMs (small and large) are evaluated as VM computing power varies according to its type. The number of linear equations is set to  $n = 7000$  in small VM and  $n = 10000$  in large one. In the Second scenario, performance of VM is evaluated when there are different numbers of VMs are running the LINPACK simultaneously in order to test CPU isolation of VMs and check if there is any interference among them because of resource sharing. In this scenario a medium type VM with  $n=7000$  has been used.

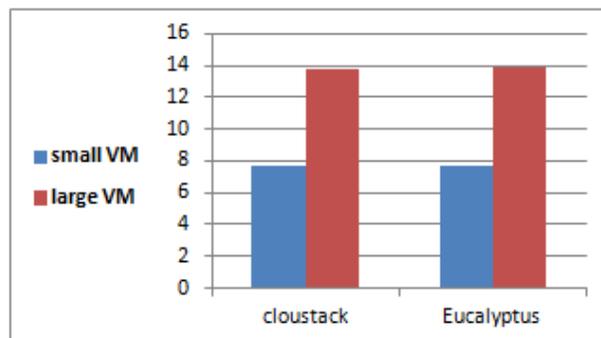


Fig. 3. CPU Performance of VMs

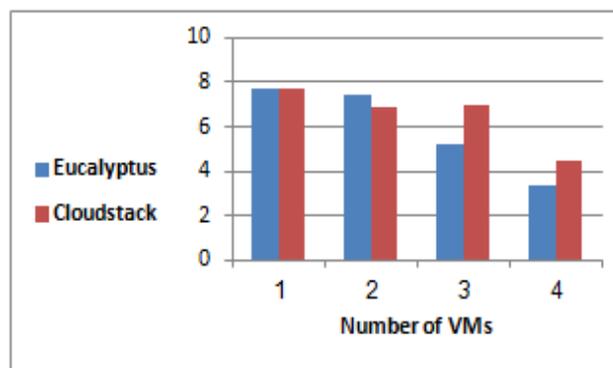


Fig. 4. CPU Isolation

Figure 3 shows the performance of VMs types. Eucalyptus and CloudStack VMs get a similar score. The floating point execution rate is considered very good with 7.7 Gflop/s and 13.8 Gflop/s for small and large VMs respectively in Eucalyptus, and 13.7 Gflop/s and 7.6 Gflop/s in CloudStack as compared to values with performance of physical machines with similar hardware specifications as in [15]. Figure 4 represents the performance when the benchmark is running on a multiple VMs. The figure reveals that CloudStack provides a slightly better VMs CPU isolation than Eucalyptus. In this scenario the VMs have been assigned the entire physical cores of host server.

The third scenario tests the performance when CPUs overcommitting is implemented. CPU overcommit is the process of allocating more virtualized CPUs (vCPU) to VM

than actual physical CPUs of system. This requires underlying hardware and hypervisor support, and this is one of reason why KVM has been chosen in the clouds deployment. It allows resource utilization and running fewer CPU cores which saves power and money. After testing and customizing the overcommit ratio in our clouds, it has been set to two times the number of physical CPUs in the system.

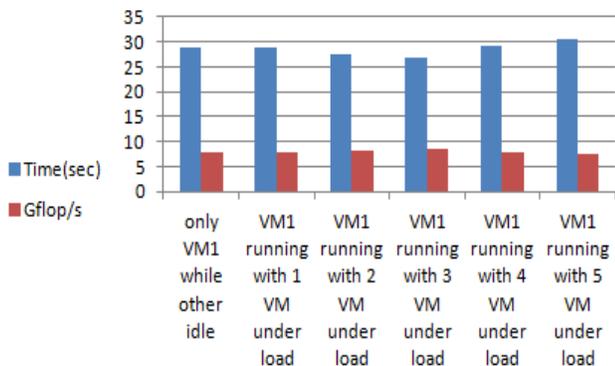


Fig. 5. CloudStack CPU Overcommitted Performance

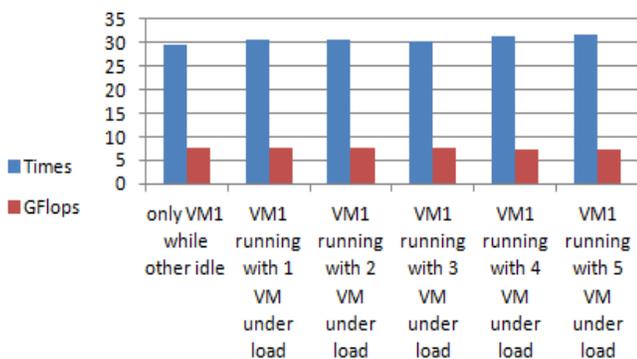


Fig. 6. Eucalyptus CPU Overcommitted Performance

Figures 5 and 6 represent performance of VMs with vCPU. LINPACK is run on medium VM with N= 7000. Then the test was repeated when there are other VMs running with 90% CPU utilization to test the effects of processor interference due to overcommitting. Lookbusy has been used to generate a high CPU utilization in VMs; it is an application for generating synthetic load on a system by generating fixed, predictable loads on CPUs, keeping chosen amounts of memory active, and generating disk traffic.

Figures reveal that assigning VM a vCPUs is appropriate and works as expected, as there is no effects from other VMs on the tested VM that run Linpack. Floating-point rate and time of execution are nearly the same as number of VMs with high utilization increased in each case on the both cloud platforms. This scenario revealed that the cloud vCPU solution is better that using normal CPU core in performance and isolation; this is due to CPU job scheduling and fair sharing techniques implementations of CPU overcommit.

C. Disk I/O Performance

As previously mentioned, Eucalyptus uses host local disk for VM, while CloudStack uses primary storage that access via NFS for VMs disks.

To evaluate and compare the performance of both clouds VM disk I/O, the Bonnie++ benchmark is adopted in this test.

Bonnie++ is a well-known Disk I/O performance benchmark suite that uses a series of tests including data read and write speeds, maximum number of seeks per second, maximum number of file creations, and deletion or gathering of file information per second.

Two scenarios are implemented on both clouds. First, Disk I/O of two types of VMs, small and large are evaluated. Bonnie++ documentation recommend that file size should be double RAM size, therefore files with 1GB and 4GB sizes for small and large VM respectively were considered. Second, performance of VM when there is another VM performing intensive disk I/O operation is inspected.

This is done to test isolation between VMs and check if there is any interference.

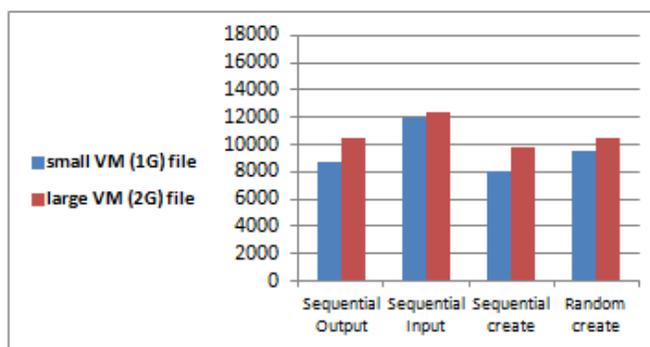


Fig. 7. Disk Access Speed in CloudStack

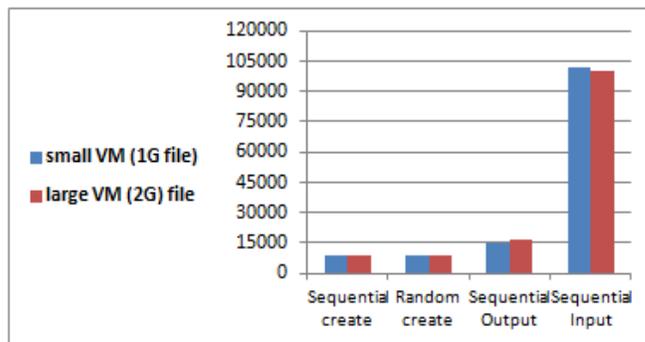


Fig. 8. Disk Access Speed in Eucalyptus

Figures 7 and 8 show the performance of VM types of both clouds. Sequential Output shows the speed in KB/s in which the data has been written. Sequential Input is the speed the data has been read, Sequential and Random create refer to the number of files created per sec.

Eucalyptus has a better overall performance than CloudStack; this is due to using of local disk configuration for VMs in Eucalyptus so VM access the host disk locally, while in CloudStack it accesses shared disk of primary storage over the network via NFS which declines disk I/O speed and performance.

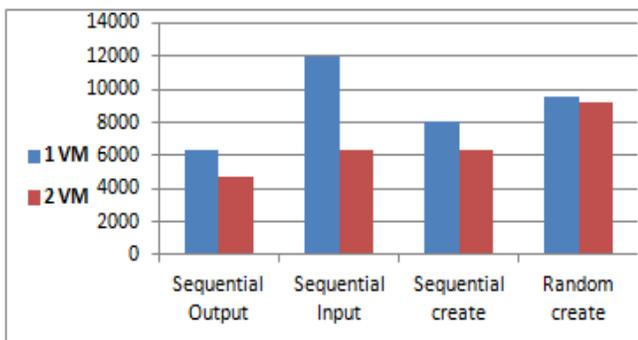


Fig. 9. Disk Isolation in CloudStack

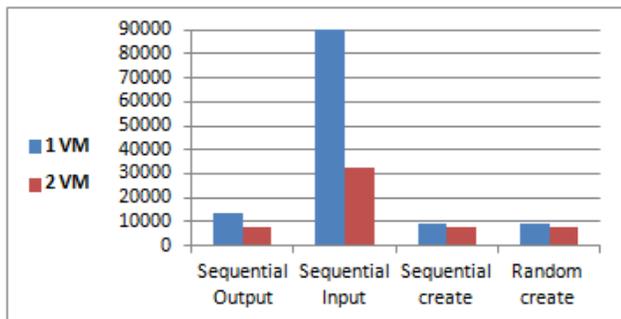


Fig. 10. Disk Isolation in Eucalyptus

Figures 9 and 10 illustrate the performance when two VMs are carrying intensive read write file operations concurrently. In this scenario medium type VM with 1GB file size is dealt with. It reveals that disk I/O performance of VM disk is impacted by the other VM as its performance drops in the both clouds. In CloudStack NFS configuring, this is expected due to primary storage disk sharing and available network bandwidth of VM.

In Eucalyptus, the NC's disk capacity and bandwidth is typically shared between VMs. The capacity is shared in a straightforward way: each virtual machine has a virtual disk image of a determined size that is allocated at the VM starting time. It does not change until the termination of the VM execution. On the other hand, the bandwidth of the disk is shared between all the resident VMs and there is currently no method of dividing this bandwidth or imposing limitations on its consumption by VMs. Therefore, the disk I/O performance of one user would be interfered by another user's VM with intensive disk I/O behavior.

Despite that the interference problem is existed in both cloud platforms; Eucalyptus has a better disk performance than CloudStack. This is due to local storage configuration where the VM disk is accessed locally (within host server) not over the network via NFS.

#### D. Memory Performance

The memory performance stress test is based upon a bandwidth test, as this is what distinguishes between types of memories. To measure the memory bandwidth the STREAM benchmark is used. It is a synthetic benchmark tool that measures memory bandwidth (in MB/s). It is specifically designed to work with datasets much larger than the available

cache on any given system, so that the results are more indicative of the performance of very large, vector style applications.

Figure 11 indicates the results of memory performance of small and large VMs in MB/s of both clouds. The array size applied in the benchmarking is 10,000,000 elements for small VM and 70,000,000 elements for large VM.

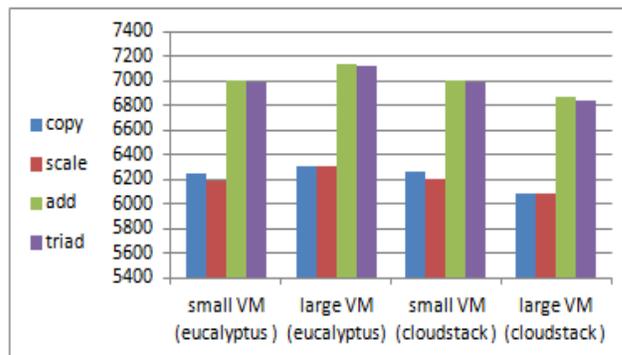


Fig. 11. Memory Performance

Figures 12 and 13 show the memory isolation between VMs, residing on the same host server. In this scenario STREAM benchmark is run on multiple VMs simultaneously. The tests demonstrates that with only one VM provisioned, there are plenty of rooms for further utilization of memory but as the number of VM increase the bandwidth available to each drops. Hence it requires a scheduler to avoid such effects.

Despite that the memory isolation problem is existed in the both cloud platforms; CloudStack shows better memory performance than Eucalyptus.

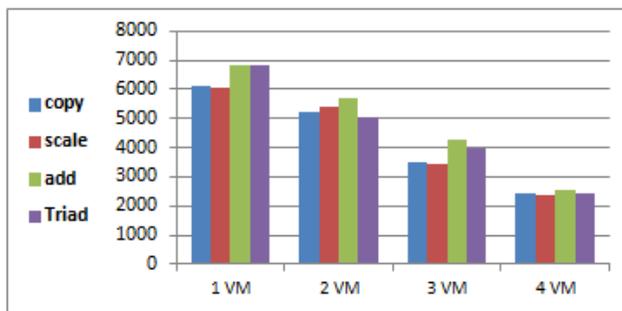


Fig. 12. Memory Isolation in CloudStack

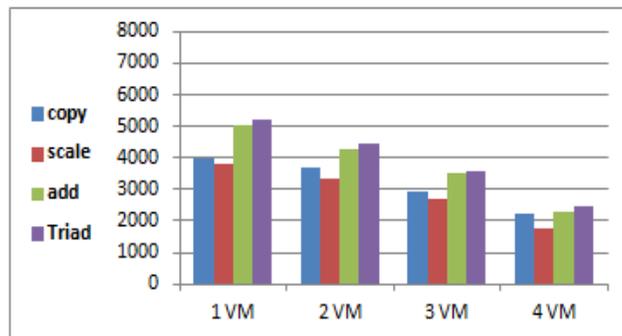


Fig. 13. Memory Isolation in Eucalyptus

E. Network Performance

The network performance tests are performed using Iperf. It is a network testing tool that allows the user to set various parameters that can be used for testing a network. It implements a client and server scheme to measure network performance between two ends, by creating a TCP and UDP data streams and measuring the throughput of network that is carrying them.

Three scenarios have been employed. First, bandwidth of VMs inside the cloud is measured by running two VMs, one as client and other as a server and TCP bandwidth between them is measured. Thereafter, the test is repeated when there are others VMs using the network. Second, packet loss is calculated at different bandwidths using UDP mode with a different number of VMs. Third, jitter is determined using UDP mode when there are more than one VM sending or receiving data over the network.

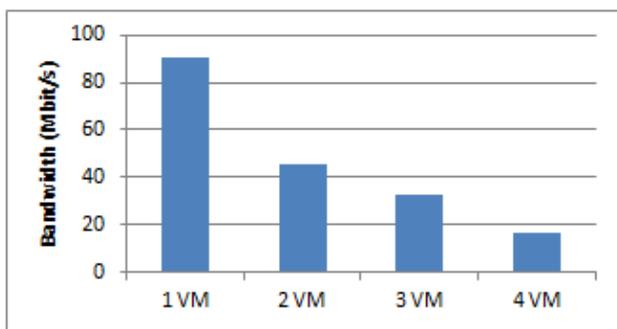


Fig. 14. Network Bandwidths Inside Eucalyptus Cloud.

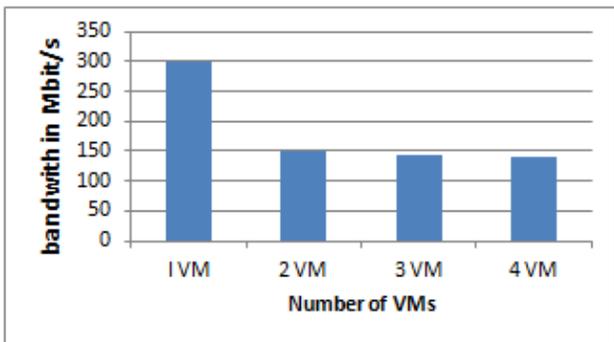


Fig. 15. Network Bandwidths Inside CloudStack Cloud.

Having seen the disk I/O interference problems, it is expected to find similar issues in the process of sharing another resource that is the network adapter. Figure 14 shows that performance of VM degrades as number of VMs increase. It proves that Eucalyptus has no built-in system of bandwidth fair-sharing between VMs. Every time concurrent TCP connections in the network are started from the VMs, each of them gets a different share of the link bandwidth and has the ability to starve the other depending on which connection begins first.

Figure 15 shows network performance of CloudStack. It reveals that when one VM is communicating, it utilizes all available network bandwidth but when there are others VMs

using the network, the bandwidth is fairly divided among them.

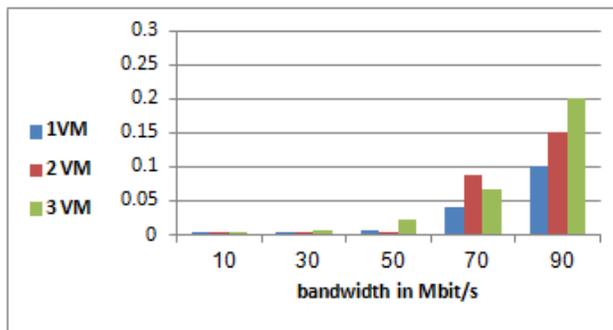


Fig. 16. Packets Loss in CloudStack

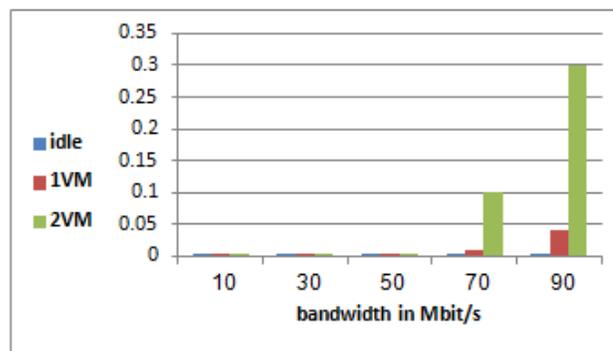


Fig. 17. Packets Loss in Eucalyptus

As depicted in figures 16 and 17, the packet loss is persisting around zero when each VM is communicating at a small bandwidth but as the bandwidth increases the packet loss increases considerably. However it does not arrive to a critical loss value in the both clouds.

Figure 18 expresses jitter when the VM is using 100Mbit/s bandwidth. As the number of VMs concurrently using network increases, the jitter is slightly increased in Eucalyptus while the jitter value is nearly the same in CloudStack. This is due to that bandwidth is fairly divided among VMs.

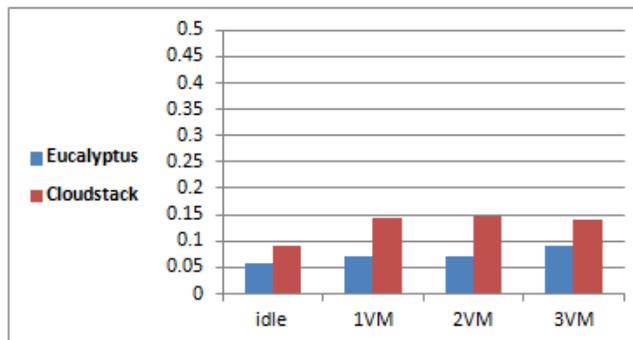


Fig. 18. Jitters in Eucalyptus and CloudStack

CloudStack has better network connection performance than Eucalyptus, due to better internal design and using of vRouter system virtual machine in cloudstack. Also the network internal traffic does not have to go through the master node (the CLC in Eucalyptus that act as internal router).

Therefore, the network connection between internal VMs will be solely determined by the physical network card which is around 1Gbps. CloudStack also provides a good bandwidth sharing among VMs. The network performances for all cloud solutions are restricted by the physical network environment.

### V. VM PROVISIONING AND RELEASE TIME

One of many advantages of the cloud is the elasticity that is the ability to dynamically acquire or release computing resources in response to demand. However, this elasticity is only meaningful to the cloud users when the acquired VMs can be provisioned in time and be ready to use within the user expectation. The long unexpected VM startup time could result in resource under-provisioning, which will inevitably hurt the application performance, hence it is required to evaluate the VM startup and release time to help cloud users to plan ahead and make in-time resource provisioning and releasing decisions [12]. A systematic study of VM provisioning and releasing time has been done for the Eucalyptus and CloudStack considering different related factors.

#### A. Number of VMs

The average provisioning time of VMs in CloudStack cloud is 16 seconds while in Eucalyptus, it is 127 seconds. This difference is due to CloudStack NFS storage configuration, in CloudStack the VM uses primary storage as its disk access via NFS while other resources (CPU, memory ...) are provided by host server so there is no need to copy VM image file from image repository in primary storage to host machine disk. However in Eucalyptus VMs use host local disk. Therefore when a new VM is provisioning, the image file (size in Mbytes) is copied from Walrus storage to host machine (node controller) which is time consuming.

Figure 19 reveals that when the number of VMs requested increases, the launch time increases accordingly in both clouds. This is due to that both cloud platforms handle each VM requested as if it is launched individually (one requested after other). The provisioning time of 2 VMs request in CloudStack is 31 seconds which equals the sum of two VMs startup time requested alone, and the same applies for VMs request. In Eucalyptus, the launch time of multiple VMs shows a time difference; for example the time for 3 VM provisioning is 134 second, which it not a 3 times of provisioning one VM. This is due to Eucalyptus is not resending the image file for multiple VM. So when a new VM is creating, the Eucalyptus checks if the image file exists in the images cache on host server. Therefore there is no need to copy it again from the walrus. The little difference in multiple VMs provisioning is the time consumed in each VM resources allocation.

#### B. Type of VMs

The VM provisioning and release time in both cloud platforms is not influenced by its type as illustrated in figures 20 and 21. VMs with different types have nearly the same startup and release time around 16 and 28 seconds respectively in CloudStack and 128 and 10 seconds in Eucalyptus. This reveals the satisfactory and quick VM resource allocation schedulers of both Eucalyptus and CloudStack.

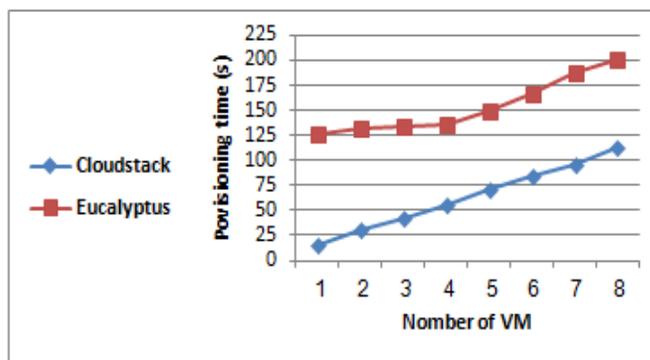


Fig. 19. VMs Launch Time vs. Number of VMs

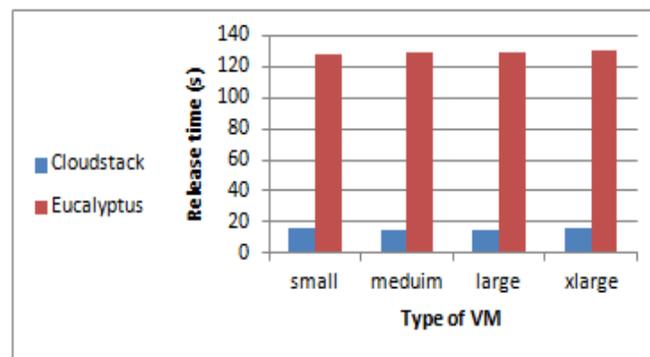


Fig. 20. VMs Startup Time vs. Type of VMs

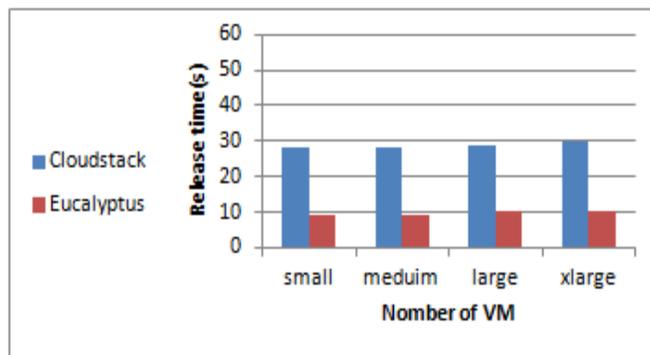


Fig. 21. VMs Release Time vs. Type of VMs

#### C. Image Size

The VM provisioning time is not influenced by size of image or template used to initiate it in CloudStack as depicted in figure 22. VMs with different image sizes have nearly the same startup time around 16 seconds. This is due to using primary storage as shared disk for VMs in CloudStack access via NFS. Therefore there is no need to copy templates (of different sizes) from primary storage to hosts which results in reducing the time of VM startup regardless of template size.

In Eucalyptus, the size of VM image (which depend mainly on OS) can largely impacts the provisioning time as it is shown in figure 23. This is due to local disk configuration of Eucalyptus which requires VM image copying from image repository in walrus to disk of sever that hosts VM. The larger the image file, the longer the VM provisioning time will be.

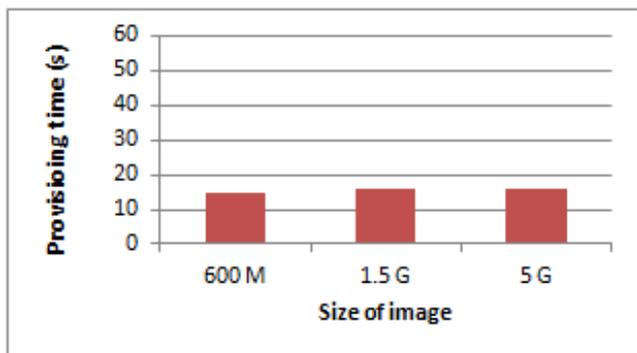


Fig. 22. VMs Startup Time with Different Image Size in CloudStack

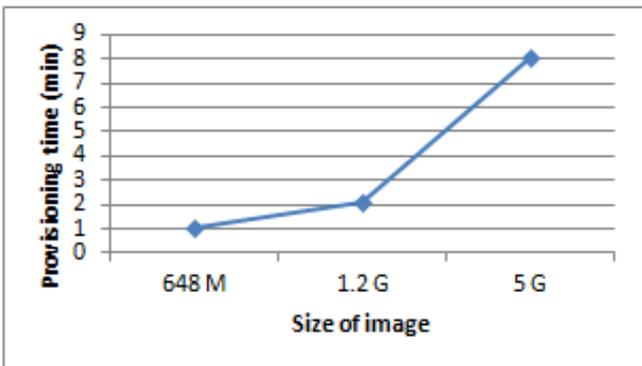


Fig. 23. VMs Startup Time with Different Image Size in Eucalyptus

#### D. Adding Additional Disk Space

CloudStack allows users to attach additional volume to VM disk at time of creation. The VM provisioning and release time is not affected by adding additional disk volumes as being requested by the user. The VM startup and release time is nearly the same when adding different disk size to the VM which is rated 16 and 29 seconds respectively as shown in figure 24. This is probably due to that CloudStack uses the primary storage to provide disks to VMs with a quick resource allocation scheduler. Eucalyptus allows attaching disk volume to running VM only.

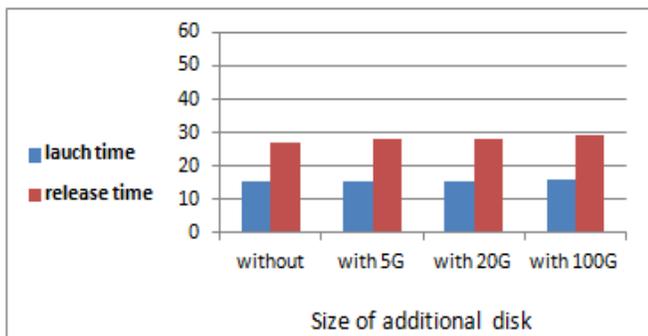


Fig. 24. VMs Startup and Release Time vs. Additional Disk

## VI. PERFORMANCE EVALUATION OF VM LIVE MIGRATION

CloudStack supports live migration of VMs between host servers while Eucalyptus supports only cold migration of VMs due to local disk configuration and lack of central sharing of VMs. Cold migration requires stopping the running VM and then moving it with its data disk to another host machine where it starts and runs again. So the VM will have a downtime which may affect user works [16]. Cold migration has no advantage in disaster recovery since VM disk is located at the host machine. So if the host fails, the VM and its data disk will be lost. This is contrast to the live migration in CloudStack where the VM data disk is at high available primary storage. Therefore in case of host failure, the VM can migrate to another host and resume work and access its disk via NFS.

Time duration of VM live migration in CloudStack has been expressed considering different factors as pursued.

#### A. Image Size

Duration of VM migration is influenced by image or template size used to initiate it as shown in figure 25. There is no difference when using 1G and 5G image size for VMs using shared disk. There is no need to move data disk from source host to destination host. That is the size should not affect migration time. However, when 600M image has been deployed, it takes a shorter time than 1G and 5G. This is due to that these are GUI OS images while 600M is non. This means that it is lighter and its applications consume less CPU and memory; the context switch compromising CPU status and memory pages copied from source to the destination host, is of small size thus it migrates faster.

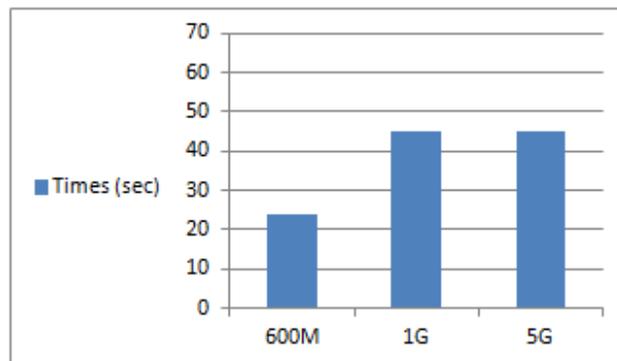


Fig. 25. Live Migration with Different Image Size

#### B. Types of VMs

We have measured migration time of different types of VMs running normal application. The type of VM can largely influence the duration of migration as shown in figure 26. This is due to increasing memory size assigned to VM in each type, so the duration of live migration increases linearly with it.

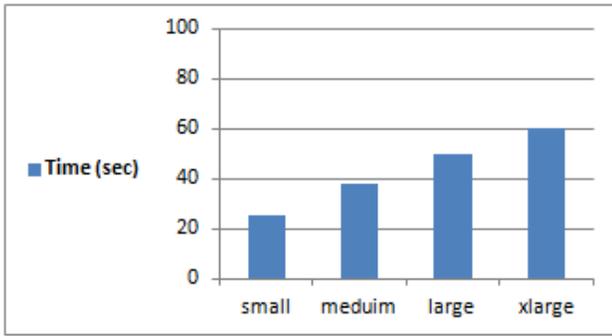


Fig. 26. Live Migration with Different VM types

### C. Number of VMs

The average time of live migration of VM in the CloudStack cloud is 40 seconds. When the number of migrating VM increase, this time increases accordingly, as it is shown in figure 27.

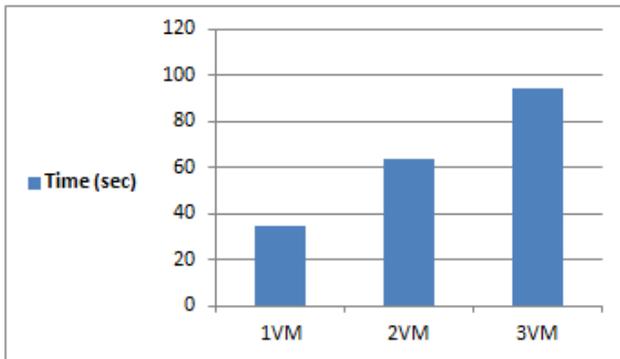


Fig. 27. Live Migration with Different Number of VMs

### D. CPU load

We have measured the migration time of VM when the CPU is running an intensive application to assess its effect on migration as a relating factor. We have tested two types of VMs, medium and large and have used Lookbusy tool to generate a 90% CPU utilization. We have found that the CPU load can have an impact on the duration of migration as shown in figure 28.

We can conclude that live migration depends on CPU utilization and applications running on the VM.

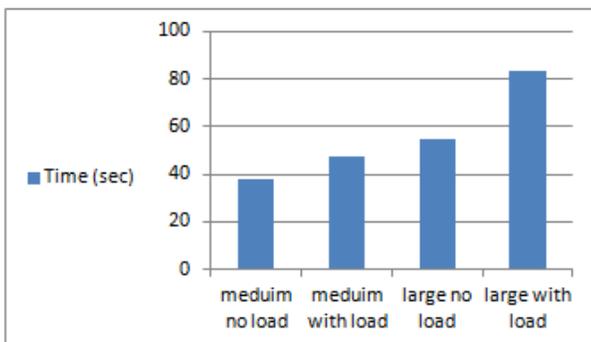


Fig. 28. Live Migration with CPU Load

## VII. WEB APPLICATION OVER CLOUD COMPUTING

One of most popular usage of VM in cloud is web application hosting. Cloud hosting has many advantages over traditional web hosting like cost reduction, scalability, flexibility, backup, security and isolation, and unlimited storage capacity. The main issues of running web application over cloud are performance and stability.

To evaluate Eucalyptus and CloudStack Clouds in hosting a web application, we have implemented a web application on VM in both clouds and measured the response time of web application to test the stability of running the application on VM. Etherpad is an open source online office suite similar to Google Docs. It is a web-based collaborative real-time editor. Etherpad has been implemented on VMs of both Clouds with MySQL as the database and Nginx as the web server.

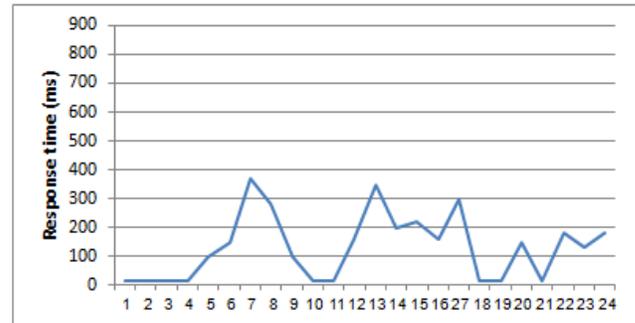


Fig. 29. Response Time in Eucalyptus Cloud

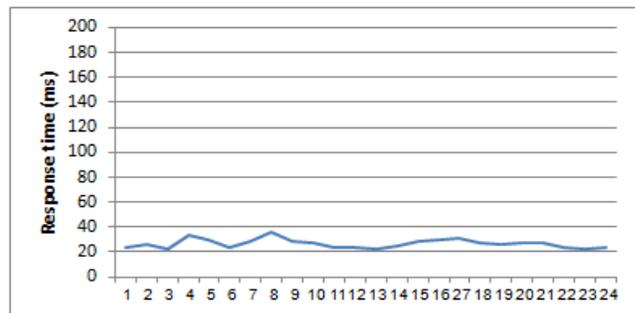


Fig. 30. Response Time in CloudStack Cloud

Figures 29 and 30 showed the response time of Etherpad which has been run on both clouds for 24 hours. The data is collected every one hour. The goal is to test VM stability on running the web application with changing the background load. Different numbers of VMs are let to run and invoke different disk and CPU intensive operations to test whether the cloud resource sharing could affect the running web application.

The figures 29 and 30 reveal that CloudStack VM is more stable than Eucalyptus VM in running the web application as the response time values over the 24 hours are nearly constant between 20 and 38 ms, while in Eucalyptus, this dramatically varies between 14 and 380 ms. Hence CloudStack is more suitable in hosting web application than Eucalyptus cloud.

## VIII. CONCLUSION

In this paper, we analyzed and compared the performance of Eucalyptus and CloudStack cloud with different storage configuration thoroughly to assess its suitability to be adopted as an open source private cloud solution for different business and scientific purposes. We have considered the performance of VM as the key point of evaluation. It has been found that storage configuration of the cloud largely affects VMs performance. CloudStack NFS configuration is 69% faster in VMs provisioning than Eucalyptus local disk configuration, while VM disk I/O performance in Eucalyptus local disk configuration outperformed the VM disk performance in CloudStack NFS configuration.

VMs performance of both clouds was evaluated in regard to CPU utilization, disk I/O speed, Memory bandwidth, Network performance, and VM management operations such as VM provisioning time and live migration. The result showed that there is always a performance decrease due to co-located VMs running resource-intensive tasks. The drop in performance is slight for CPU and memory intensive workload and very significant for disk and network I/O intensive workloads. The major lessons learned related to the performance evaluation of VM management operation are: (1) the duration for the live migration changes with the CPU load; (2) the duration for the live migration increases linearly as the memory assigned to the VM increases; (3) the startup and release time have not been impaired by the VM type; (4) the startup and release time have not been impaired by image size or by adding additional disk volumes in CloudStack, while the startup time is largely affected by image size in Eucalyptus.

Also, Eucalyptus and CloudStack clouds ability in hosting web applications was tested by measuring the response time of web application that was hosted on their VMs. It has been found that CloudStack is more suitable in hosting web applications and as private cloud solution in general due to its stability and fair VMs performance. On the other hand, Eucalyptus is easier in deployment and more modular, it can be used in testing a specific application on the cloud so it's a good choice for developers and researchers in this field.

## IX. FUTURE WORK

As a future work we intend to analyze security aspects of Eucalyptus and CloudStack by evaluating the compliance of them with security standards related to cloud security. Also, the methodology and the benchmarks used for performance evaluation in this paper can be used for different cloud management platforms whatever open source or commercial platforms (OpenNebula OpenStack, HP cloud, VMware, etc.) and compare their performance results with this paper to extend the evaluation of the cloud management platforms.

## REFERENCES

- [1] E. T. de Sousa, P. R. Maciel, E. M. Medeiros, D. S. de Souza, F. A. Lins, and E. A. Tavares, "Evaluating Eucalyptus Virtual Machine Instance Types: A Study Considering Distinct Workload Demand", Third International Conference on Cloud Computing, GRIDs, and Virtualization pp. 130- 135, July 2012.
- [2] H. Tianfield, "Cloud Computing Architectures", Proceeding of the 2011 IEEE International Conference on Systems, Man, and Cybernetics (SMC), pp. 1394-1399, Oct 2011.
- [3] P. Mell and T. Grance, "The NIST Definition of Cloud Computing" National Institute of Standards and Technology Special Publication 800-145, September 2011.
- [4] Z. Pantic and M. A. Babar, "Guidelines for Building a Private Cloud Infrastructure", Technical Report TR-2012-153, IT University of Copenhagen, Denmark, 2012.
- [5] D. N. Chorafas, "Cloud computing strategies", CRC Press, 2011.
- [6] A. S. Pilla, and L.S. Swasthimthi, "A Study on Open Source Cloud Computing Platforms", EXCEL International Journal of Multidisciplinary Management Studies, Vol. 2, Issue 7, July 2012.
- [7] I. Voras, et al., "Evaluating Open-Source Cloud Computing Solutions", Proceedings of the 34th International Convention MIPRO, pp. 209-214. May 2011.
- [8] [8] P. Sempolinski and D. Thain, "A Comparison and Critique of Eucalyptus, OpenNebula and Nimbus", Proceedings of the IEEE Second International Conference on Cloud Computing Technology and Science (CloudCom), pp. 417-426, November 2010.
- [9] [9] C. El Amrani, et al, "A Comparative Study of Cloud Computing Middleware", Proceeding of the 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, pp. 690-693, 2012.
- [10] [10] G. V. Laszewski, J. Diaz, F. Wang, G. C. Fox, "Comparison of Multiple Cloud Frameworks", proceeding of the 5th IEEE International Conference on Cloud Computing (CLOUD 2012), pp.734-741, June 2012.
- [11] [11] O. Sefraoui, M. Aissaoui, and M. Eleuldg "Comparison of multiple IaaS Cloud platform Solutions", Recent Researches in Information Science and Applications, Proceedings of the 7th WSEAS International Conference on Computer Engineering and Applications (CEA '13), pp. 212-217, January 2013.
- [12] M. Mao and M. Humphrey, "A Performance Study on the VM Startup Time in the Cloud", Proceeding of the 5th IEEE International Conference on Cloud Computing (CLOUD 2012), pp.423-430, June 2012.
- [13] D. Steinmetz, B. Perrault, R. Nordeen, J. Wilson, X. Wang, "Cloud Computing Performance Benchmarking and Virtual Machine Launch Time", Proceedings of the 13th annual conference on Information technology education, ACM New York, pp. 89-90, 2012.
- [14] F. Gomez-Folgar, A. Garcia-Loureiro and T. F. Pena, R. Valin, "Performance of the CloudStack KVM Pod Primary Storage under NFS Version 3", Proceeding of 10th IEEE International Symposium on Parallel and Distributed Processing with Applications, pp. 845-846, 2012.
- [15] J. J. Dongarra, "Performance of Various Computers using Standard Linear Equations Software", scientific report, University of Manchester, February 5, 2013.
- [16] A. Agarwal and R. Shangruff, "Live Migration of Virtual Machines in Cloud", International Journal of Scientific and Research Publications, Volume 2, Issue 6, pp.1-5, June 2012.