JETIR.ORG ISSN: 2349-5162 | ESTD Year : 2014 | Monthly Issue
JOURNAL OF EMERGING TECHNOLOGIES AND INNOVATIVE RESEARCH (JETIR)
An International Scholarly Open Access, Peer-reviewed, Refereed Journal

A STUDY OF VIRTUAL MACHINERIES TECHNOLOGY IN CLOUD COMPUTING ENVIRONMENT

- **Dr. Devendra Singh, Associate Professor, Haryana Institute of Public Administration**

ABSTRACT

One of the issues of cloud computing environments is the problem of virtual machine allocation, particularly in the context of private cloud design. There are a variety of ways in which virtual machines can be assigned to hosts in this context. Calculating how well different Cloud application and service models perform under various performance metrics and system requirements while using a scheduling and allocation policy is a particularly complex and time-consuming challenge to tackle. Using the open source IaaS technology Eucalyptus, the authors provide a framework model for a Virtual Computing Laboratory. A rule-based mapping approach for Virtual Machines (VMs) is also described, based on set theory ideas. As a result of the algorithmic architecture, virtual machines and physical hosts' resources can be dynamically mapped. Furthermore, the paper presents a theoretical study and derivations of some performance evaluation metrics for the chosen mapping policies, such as determining the time it takes for the algorithm to switch between different contexts (waiting, turnaround, and response time), as well as the time it takes for the algorithm to respond.

KEYWORDS: Manger-Worker Process, Private Cloud, Rule-Based Mapping, Virtual Machine (VM) Allocation, Virtualization

INTRODUCTION

In the context of cloud computing, end-users are given the option to outsource on-site services, computational facilities or data storage to an off-site, location-independent centralised facility or "Cloud" (Ioannis & Karatza, 2010). The term "Cloud" refers to a collection of computers and web services that work together to provide cloud computing. Ideally, these machines will have a pool of geographically dispersed physical computational resources, such as processors, memory, network bandwidth, and storage, all of which can be shared across multiple servers. A dynamically logical entity of resources related with cloud computing is frequently organised and leased out on demand as an outsourced service. Elasticity, a key attribute of cloud computing, refers to the ability for cloud resources to grow and shrink on demand (Sarathy et al., 2010). The concept of virtualization technology has made this change possible in cloud computing today. Since a few years ago, the term "virtualization" has become increasingly widely used in the IT industry. Abstraction or decoupling of application payload from the underlying distributed physical host resource is key to this technology's design (Buyyaa et al., 2009; Popek & Goldberg, 1974). In other words, actual resources can be displayed in logical or virtual form, depending on user preferences. In addition, virtualization technology can help cloud service providers save money by increasing machine usage, reducing administrative time, and reducing infrastructure costs. The dynamic provisioning of the logical resources may be achieved by providing a suitable management mechanism on top of this virtualization functionality (as we have proposed in this work) (elastic property of the cloud). According to the notion of dynamic provisioning as it relates to cloud computing, it is essential that each computer resource be able to be dynamically furnished and controlled at any given time. Open source software frameworks for cloud computing, also known as "Infrastructure as a Service," are frequently used by cloud developers to implement the virtualization concept (IaaS). Hypervisor is the name given to this software framework (Nurmi et al., 2009; Chisnall, 2009). Using a hypervisor, also known as a virtual machine management (VMM), a host computer can run multiple operating systems at the same time, called guests. Virtualization may be implemented on a variety of infrastructures, and we offer a variety of virtual infrastructure management tools for that (Sotomayor et al., 2009). In this study, we present a new cloud simulation framework for developers to use. A Manager process and a Worker process are two high-level components of the suggested architecture. In the new system, the manager serves as the cluster controller, while the worker serves as the node controller. However, the primary goal of this project is to provide a framework model that allows VMs to be dynamically assigned to physical hosts based on their resource requirements and the physical hosts' availability. Cloud resources include machines, networks, storage, operating systems, application development environments, and applications.

ROLE OF VIRTUALIZATION IN CLOUD COMPUTING

Cloud computing relies on virtualization to function efficiently and conveniently, and it also presents solutions to major difficulties in data security and privacy protection, all of which are made possible by virtualization. The term "virtualization" refers to the practice of simulating hardware characteristics in software. It is permissible for a single computer to take over the functions of numerous other computers. To achieve all of our goals, such as improved security, maximum utilisation of hardware resources and lower costs for a virtual web or file server, we need to acquire, maintain, depreciate and energy and floor space twice as much for a physical server. There are several advantages to virtualization, such as better utilisation of resources, higher levels of security, more portability, easier management, enhanced flexibility, fault isolation, and quicker implementation.

Virtualization in Cloud Computing:

- For combining local and network resources data storage virtualization.
- For grouping physical storage devices into the single unit
- For reaching the high level of availability or improving availability using virtualization
- Improving performance using virtualization
- Using virtualization using stripping and caching
- Capacity improvement

In Cloud Computing, virtualization refers to the use of a central server or computer to host an application for numerous users, rather than installing software on each device individually. One location for all of the data from many databases, hard drives, and USB drives makes it more accessible and secure. Creating virtual hardware, software, an operating system, a storage or network device is referred to as virtualization in the cloud computing context. In an IT setting, virtual changes occur far more quickly than physical ones. Because of the scalability and agility provided by cloud virtualization, it is necessary to keep up with the always changing environment.

Importance of virtualization:

Virtualization is a prerequisite for cloud computing environments since it simplifies resource management. By securing both the integrity of the cloud components and guest virtual machines, virtualization in Cloud Computing allows for an increase in security. Additionally, Cloud Component virtualized computers are scalable and reliable. The Managed Service Provider VA also allows high use of pooled resources, resource sharing, and speedy provisioning. Reasons why you should use Managed Service Provider VA:

- Simplified Management
- Reduced system administrative work
- Resource Optimization
- It saves Money
- Easier software installation
- Data center consolidation and decreased power consumption
- Testing of CD’s live without even burning them

- Better use from the hardware
- Increased CPU utilization
- Virtual machine can run on any x86 servers

NETWORK VIRTUALIZATION

It refers to the management and monitoring of a computer network's available resources from a single software-based administrator's interface that combines all of the network's resources. Scalability, dependability, adaptability, and security are just a few of the benefits that users may expect from this new technology. Many network administration duties can also be automated by this software. If you're dealing with a surge in network traffic that's sudden and unexpected, network virtualization can help. As a result, IT operations become more efficient, network productivity increases, and backups become easier. Two categories of Network Virtualization:

- **Internal:** Provide network like functionality by combining multiple networks to a single system.
- **External:** Combine many networks as software containers, or parts of networks into a virtual unit.

MEMORY VIRTUALIZATION

In order to provide a distributed, shared, or networked function, it introduces a means to detach volatile random access memory (RAM) from the server this method enhances performance by supplying additional memory capacity without increasing the primary memory's available capacity. Why does a section of the disc drive function as a virtual memory extension?

Implementations:

Application level integration – Applications and programs running on connected computers directly connect to the memory pool through an API or the file system.

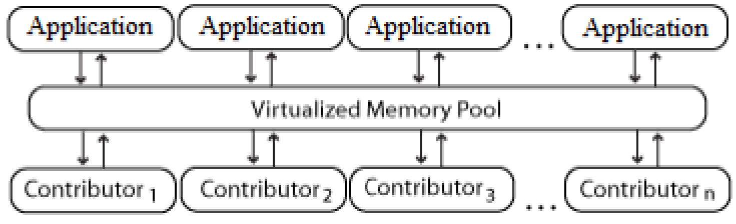

Figure 1. Application level integration

Operating System Level Integration – The operating system firstly connects to the memory pool, and makes that pooled memory available to applications.

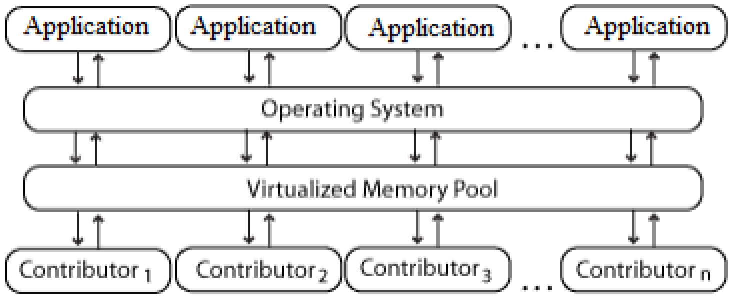

Figure 2. Operating system level integration

CLOUD PLATFORM ENVIRONMENT

Based on service-level agreements worked out between the service provider and customers and defined by Buyya, et al. (2008), cloud computing is a distributed parallel computing system made up of a network of interconnected, virtualized computers that are dynamically provisioned and presented as a single or more combined computing resources. Amazon, Google App Engine, Apple Mobile Me, and Microsoft Clouds are some of the best cloud environments now on the market that offer effective, efficient, and reliable services to cloud customers (Amazon, 2009; Chu et al., 2007).

EUCALYPTUS CLOUD ENVIRONMENT

Java-based cloud management solution Eucalyptus consists of five high-level modules. It includes cloud controllers for the clusters and nodes as well as node controllers for storage and the walrus. An independent Web service is used to implement each high-level system component (Nurmi et al., 2009; Yoshihisa and Garth, 2010). Users can access and manage the virtualized resources (machines, networks, and storage) through cloud controller APIs. Amazon EC2 API and a web-based user interface are now available for this system component (AEC, 2011). Block-level network storage is provided by the storage controller, which can be dynamically attached by virtual machines. Amazon Elastic Block Storage (EBS) semantics are supported in the current storage controller implementation (Kleineweber et al., 2011). There are third-party interfaces supported by Walrus, which makes it possible to store and access virtual machine (VM) images and other user data. In addition, the cluster controller gathers the necessary information and schedules VM execution on certain node controllers, as well as maintains the virtual instance network that runs in the cloud environment. The node controller is responsible for the execution, inspection, and termination of VM instances on the host where it is installed.. The VM instance is under the control of the node controller. When a cloud controller receives a request for a VM image, the node controller will start and operate that image and make it available to the network for usage by the cloud end user. Eucalyptus' cloud model is depicted in Figure 3.

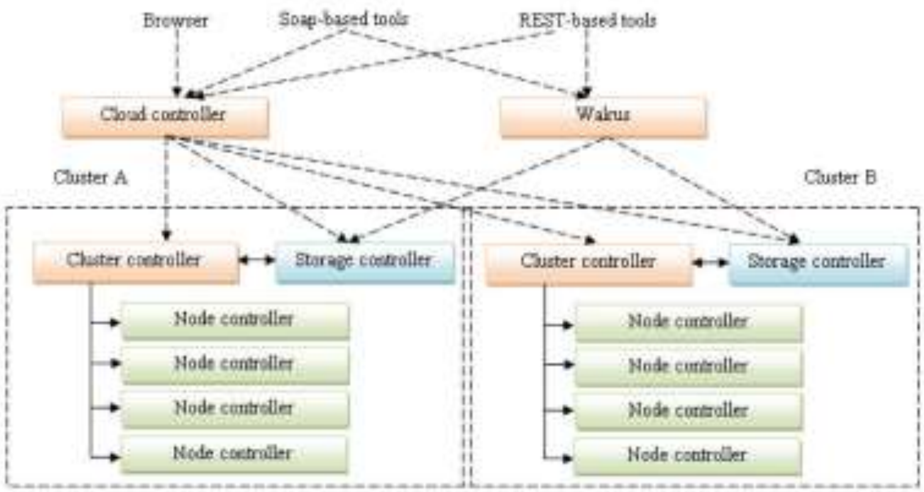

Figure 3. Eucalyptus architecture

Users and administrators enter the cloud computing platform using the cloud controller. It makes high-level scheduling decisions and implements them by sending requests to the cluster controller, which it gets from the node managers. When the cluster manager is asked to implement a cloud controller's request, it uses node controllers to do so. In order to communicate with the cloud controller, users employ a user interface that makes use of an access protocol called Simple Object Access Protocol (SOAP) or REST messages. Requests are routed to cluster controllers via a cloud controller. Numerous cluster controllers and multiple node controllers can be assigned to a single cloud controller. To improve the efficiency of handling virtual machine images in eucalyptus-based clouds, Upatissa and Atukorale (2012) provided an abstract architecture for an improved private cloud model. It is assumed in this study that in Cloud computing, any hardware or software entity such as high-performance systems, storage systems, laboratory device servers, or apps shared amongst Cloud users is considered a resource. The term "resource" refers to computer hardware, such as nodes in a network or storage systems, for the remainder of this paper. The term "resources" is used interchangeably with "laboratory devices" and "laboratory device servers." A service is a resource's network-enabled functionality that may be called upon by other resources, applications, or users.

SYSTEM ARCHITECTURE: VIRTUAL COMPUTING LABORATORY MODEL

A single physical node may host numerous virtual machines thanks to virtualization, which is the primary technology underlying the majority of private cloud solutions. Virtualization is only one component of a private cloud. In order to automate the deployment of virtual machines, storage, networking, and other infrastructure services, a cloud controller uses an API. In this article, our primary goal is to identify an alternate yet appropriate way to manage the provisioning and mapping of virtual machines to physical hosts in a typical private cloud computing environment. Virtual machines will be allocated to real hosts at runtime using an approach similar to the work provided in Malgaonkar et al. (2011) for the planned virtual computing laboratory. Similar to the work done by Upatissa and Atukorale (2012) and Nurmi et al. (2009), we plan to enhance the eucalyptus private cloud architecture with two high-level management components (i.e. the manager process and worker processes). Virtual machines known as managers are responsible for keeping track of assigned and unassigned query information about resources in this paradigm. Responding to questions and control requests from its manager, the worker runs system software on its node and performs various operations on it. As a result of providing one assignment to each employee, the workload is evenly distributed. The manager/worker model is the foundation of our cloud design since it ensures that workloads are evenly distributed.

1. Manager Process

It is the manager procedure that serves as a gateway to the cloud management system. Cloud controllers and worker nodes can query any system with network connectivity for information about available resources via cloud controllers. Subsequently, the manager process is assigned the task of

i) Making scheduling decisions (such as scheduling of incoming instances to run on specific nodes)

ii) Controlling the instances of virtual network overlay, and

iii) Gathering information about a set of nodes from the worker process.

Explain instances, describe resources, and so on are common formats for management requests to worker processes. When the manager process gets a set of operation instances, it first asks the worker processes to provide information about the resources that are available and appropriate through the function "describe Resources." Managers receive a report from workers that includes lists of resource parameters (processor, memory storage and bandwidth) as well as a summary of the results of their search. Manager processes use this information to compute the number of concurrent instances of the given kind that can be executed on the lists of available nodes and submit this value to the cloud controller for delegation and allocation.

2. Worker Process

Each VM instance's data is handled by a separate worker process. The Worker process runs on every node that hosts virtual machine instances. To respond to the management process's requests for queries and control over system software on its node, a worker processes the programme. In order to find out about the physical nodes' resource profiles, the manager process asks the worker processes. These profiles include the number of CPUs, the capacity of memory, the available disc space, and the status of VMs on the node. Once it has been collected, it is sent back to the cloud controller for further processing and delegating to the manager (see Figure 4). It is also possible to include some local storage in the cluster or resource pool, which can be either local storage that is physically tied to the node, or shared storage that can be accessed via a shared pool of storage using storage local area network fibre channel or similar technologies. Multiple clusters of managers and workers can be added to the architecture shown in Figure 5 to boost both capacity and redundancy, which can be utilised to increase total infrastructure availability.

Figure 4. In a manager/worker-style, a manager process send the elements of a set is often denoted by lower case letter (e.g. a b, ,c, . . .)

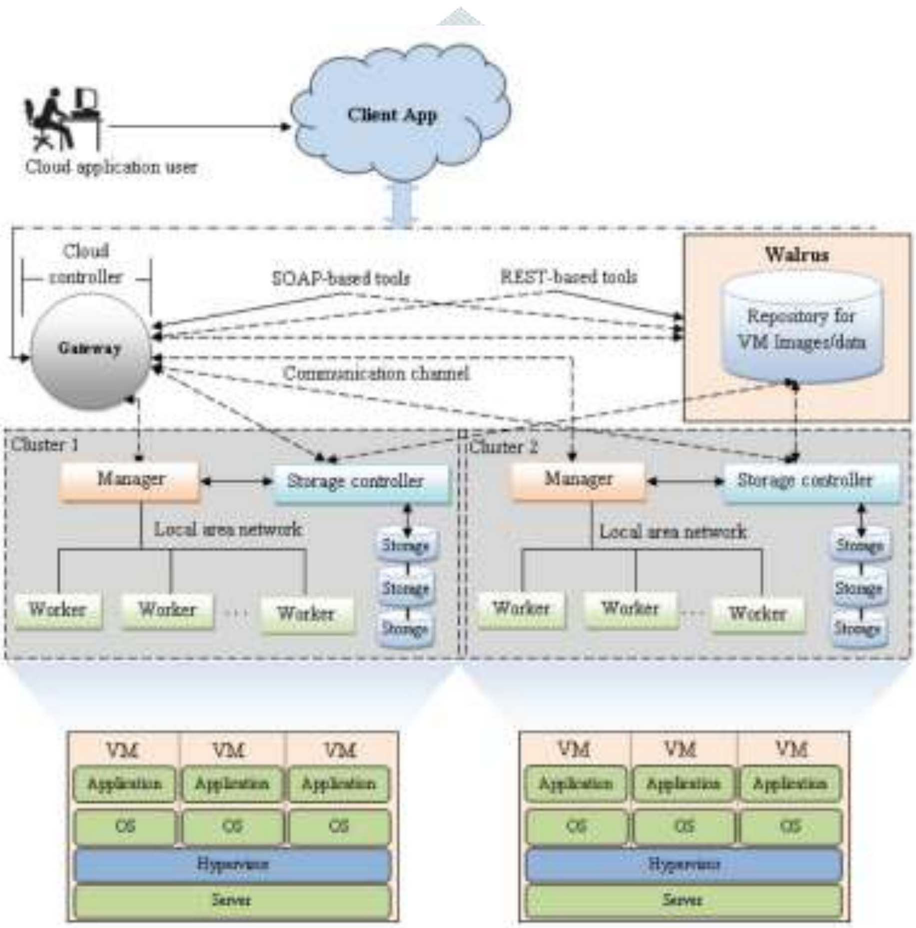

Figure 5. The proposed private cloud architecture

APPLYING VIRTUALIZATION TO CLOUD COMPUTING

Rather of selling products, cloud computing offers computing as a service, making it possible for computers and other electronic devices to access shared resources, software, and data. The following characteristics should be included in a cloud computing platform that has been properly designed: This is known as "scalability in dynamic conditions." High availability, high performance, and equilibration of resources can be achieved through dynamic resource allocation. There are a wide range of research topics related to cloud computing, including energy management and stability as well as virtualization and scalability. The cloud computing industry relies on a number of different technologies, including virtualization. Using virtualization, cloud computing hardware may be pooled together and shared among multiple users. Virtualization can optimise the utilisation of computer resources and the dynamic allocation of those resources. Currently, cloud-based infrastructure is typically built using VMware, vCloud, and Xen Cloud Platform solutions. Cloud computing platform based on the powerful Xen hypervisor and equipped for virtualization, Xen Cloud Platform is an enterprise-ready platform for virtualization and cloud computing. For example, users can access their virtual machines' consoles, view their properties, do power management operations, manage virtual machine "instantaneous events," and even move their virtual machines between servers in the same pool. An online cloud of VMware infrastructure instruments is provided via the vCloud service. A cloud infrastructure can be built utilising the VMware vCloud tools that are available.

Terminal applications are the primary focus of the SAAS layer. Developers can build cloud applications using the PAAS layer, which provides an application-based cloud management platform. There are two layers to the IAAS layer, and they are developed using VMware vCloud. The VMware vCloud is made up of five components. VMware vCloud Director, VMware vSphere, VMware vShield, VMware vCenter Chargeback, and VMware vCenter are the most important components to consider. One hundred percent of the department's infrastructure is made up of them. Customers can create a Virtual Data Center resource pool by using the VMware vCloud Director component. Consumption of means may be requested by pupils. Additionally, it can take advantage of technologies like clones and fixtures that can greatly speed up the availability of infrastructure. Security services provided by VMware's vShield component include firewall, DHCP, and VPN. Additionally, it empowers users to improve application and data security by supporting virtualization protection for cloud-based settings. It is possible to automate daily tasks in a large data centre using VMware vCenter, which provides a single point of control for all aspects of virtual infrastructure. vCenter Charger is back in VMware Resource and cost accounting templates are included in the component. Cloud computing services cost estimates and analyses are the key goals of this tool. It may also help consumers better understand how much cloud computing services cost in terms of resources, as well as how to make better use of those resources and cut down on the overall infrastructure expenses. vSphere-based vCloud services provided by VMware vCloud Server virtualization, storage virtualization, and networking virtualization are just a few of vSphere's many capabilities. Virtual real-time machine migration and load balancing are also supported. Storage migration is also done in a non-disruptive manner, eliminating I/O bottlenecks and freeing up valuable storage capacity. With the VMware vCloud tools, we can create an IAAS layer for the cloud services platform's infrastructure.

CONCLUSION

Cloud Computing Environment and virtual machine allocation were discussed in detail in this study. The purpose of the design is to model an easy-to-deploy private cloud and define the allocation problem of virtual machines in terms of set theory. Cloud customers can run their Virtual Machines efficiently on a limited number of real machines by using a simple but effective rule-based mapping method, which has been proposed and presented together with this algorithm. To put it another way, this method will allow for maximal processing power with the smallest amount of physical data centres infrastructures.

REFERENCES

AWS. Amazon Web Services LLC (2009). Amazon elastic compute cloud (EC2). Retrieved from http:// aws.amazon.com/ec2/

Buyya, R., Yeo, S. C., & Venugopal, S. (2008). Marketoriented cloud computing: Vision, hype, and reality for delivering IT services as computing utilities. In Proceedings of the 10th IEEE International Conference on High Performance Computing and Communications.

Buyyaa, R., Yeoa, S. C., Venugopala, S., Broberga, J., & Brandicc, I. (2009). Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility. Future Generation Computer Systems, 25(6), 599–616. doi:10.1016/j.future.2008.12.001.

Calheiros, N. R., Buyya, R., & De Rose, F. A. C. (2009). A heuristic for mapping virtual machines and links in emulation test beds. In Proceedings of the 9th International Conference on Parallel Processing (ICPP) (pp. 518–525). Washington, DC: IEEE Computer Society.

Chisnall, D. (2009). Guide to the xen hypervisor. Prentice Hall Press.

Chu, X., Nadiminti, K., Jin, C., Venugopal, S., & Buyya, R. (2007). Next-generation enterprise grid platform for e-science and e-business applications. In Proceedings of the 3rd IEEE International Conference on e-Science and Grid Computing.

Kleineweber, C., Keller, A., Nieh¨orster, O., & Brinkmann, A. (2011). Rule-based mapping of virtual machines in clouds. In Proceedings of the 2011 19th International Euromicro Conference on Parallel, Distributed and Network-Based Processing (pp. 527-534).

Malgaonkar, P., Koul, R., Thorat, P., & Zawar, M. (2011). Mapping of virtual machines in private cloud. International Journal of Computer Trends and Technology, 2(2).

Mohanty, R., Das, M., & Lakshmi, P. M., Sudhashree. (2011). Design and performance evaluation of a new proposed fittest job first dynamic round robin (FJFDRR) scheduling algorithm. International Journal of Computer Information Systems, 2(2).

Moschakis, I. A., & Karatza, H. D. (2010). Evaluation of gang scheduling performance and cost in a cloud computing system. Journal of Supercomputing, 1-18. doi doi:10.1007/s11227-010-0481-4, Springer, Berlin.

Nurmi, D., Wolski, R., Grzegorczyk, C., Obertelli, G., Soman, S., Youseff, L., & Zagorodnov, D. (2009). The Eucalyptus open-source cloud-computing system. In Proceedings of the 2009 9th IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGRID '09). Washington, DC: IEEE Computer Society.

Nurmi, D., Wolski, R., Grzegorczyk, C., Obertelli, G., Soman, S., Youseff, L., & Zagorodnov, D. (2009). The eucalyptus open-source cloud-computing system. In Proceedings of the 9th IEEE/ACM International Symposium on Cluster Computing and the Grid (CCGRID'09) (pp. 124-131). IEEE.

Popek, G. J., & Goldberg, P. R. (1994). Formal requirements for virtualizable third generation architectures. Communications of the ACM, 17(7), 412–421. doi:10.1145/361011.361073.

Sarathy, V., Narayan, P., & Mikkilineni, R. (2010). Next generation cloud computing architecture: Enabling real-time dynamism for shared distributed physical infrastructure. Retrieved from www.kawaobjects.com/resources/PID1258479.pdf

Smith, E. J., & Nair, R. (2005). Virtual machines: Versatile platforms for systems and processes. Morgan Kauffmann.

Sotomayor, B., Montero, R. S., Llorente, I. M., & Foster, I. (2009). Virtual infrastructure management in private and hybrid clouds. Internet Computing, IEEE, 13(5), 14–22. doi:10.1109/MIC.2009.119. Sun Microsystems, Inc. (2009). Introduction to cloud computing architecture. White Paper (1st ed.).

Upatissa, D. T., & Atukorale, A. (2012). Low cost virtual lab environment to the university by using cloud environment. In Proceedings of the International Conference on Computer Engineering and Technology (ICCET 2012) (IPCSIT vol.40). Singapore: IACSIT Press.

Yoshihisa, A., & Garth, G. (2010). pWalrus: Towards better integration of parallel file systems into cloud storage. In Proceedings of the Workshop on Interfaces and Abstractions for Scientific Data Storage (IASDS10), Co-Located with IEEE International Conference on Cluster Computing 2010 (Cluster10), Heraklion, Greece.