Research Issues in Virtualization in Cloud Computing

Adesh Kumar
Department of Computer Science
SLBSRSV, New Delhi India

Abstract- Cloud computing technology proposes a complete online platform composed of a large number of services used while needed. Cloud computing provides its services in cost effective way via internet in a reliable and efficient way. Cloud computing reduces the investment on purchasing the hardware, software and software licenses by providing services on rental basis. It reduces the licensing cost and provides backups to keep multiple copies of data.Cloud computing technology proposes a complete online platform composed of a large number of services used while needed In this paper a study has been carried out to find research issues in virtualization in cloud computing.

Keywords- Virtualization, Cloud computing, Services, characteristics, data centre.

I. INTRODUCTION

Cloud computing technology provides scalable and flexible technical infrastructure capabilities as an on-demand service (Singh and Kumar, 2014). The users can access cloud computing resources from any place, anywhere and atany time via any mobile computing devices such as laptops, mobiles, tablets or smart phones. In cloud computing technology, machines with large data centers can be dynamically provisioned, configured, controlled and reconfigured to deliver services in a scalable manner (Zissis and Lekkas, 2012). Cloud computing allows to efficiently manage upgrades and maintenance, backups, disaster recovery and failover functions (Zaharescu, and . Zaharescu, 2012) . According to Vaquero et al. (Vaquero et al., 2008), cloud guarantees are offered by the infrastructure provider by means of customised Service Level agreements (SLA). Cloud computing solves the potential problems of education, climate change, economics and terrorism (Kop and Carroll, 2011).

Cloud Computing Characteristics

Cloud computing has certain characteristics that are illustrated in figure 1.

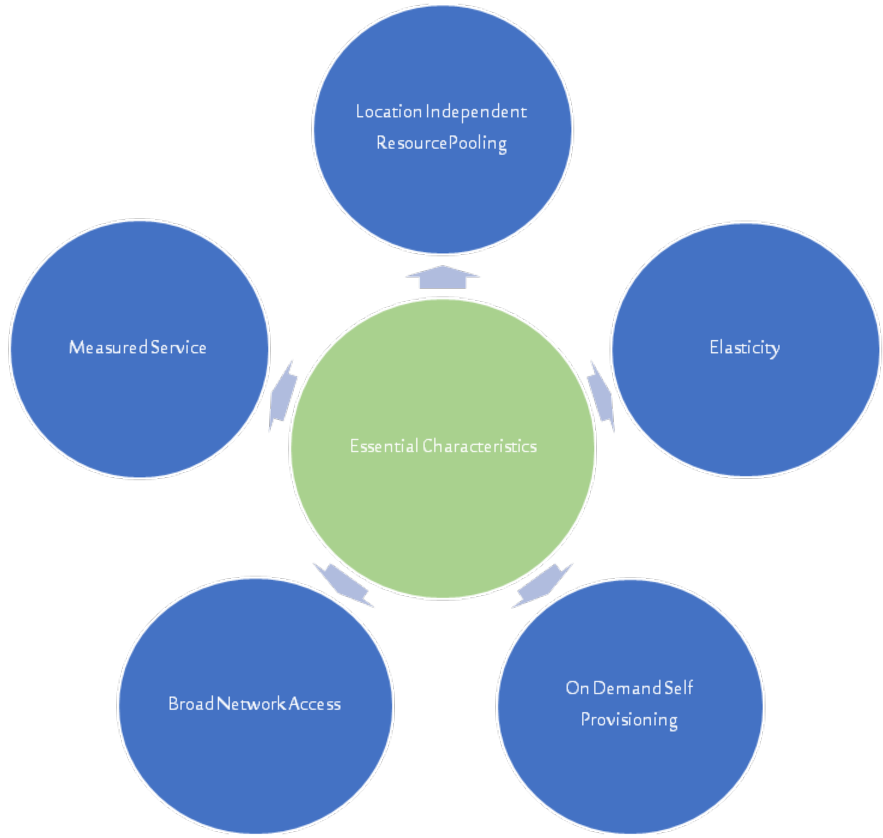

Figure 1.Cloud Computing Characteristics

Location independent resource pooling- Cloud computing provides pool of hardware and software resources to serve request of multiple users (Mell and Grance, 2009). The resources are dynamically allocated and deallocated according to varied user demands. The user has no control or knowledge over the exact location of the given resources.

Elasticity- Cloud computing provides infinite computing resources available on demand (Buyya et al. 2010). Hence cloud provides resources in any quantity and at any time. The cloud can increase the computing resources when load goes high for any particular applications (Mell and Grance, 2009). Similar cloud will release the resources when load is decreased without any human involvement.

Measured Services - In cloud computing there is no upfront cost is required for accessing the computing resources (Armbrust et al., 2009). The cloud consumer pays only for what services user is consuming, on rental basis. Metering and billing services is applicable in cloud computing to trace the usage of various computing resources by cloud users (Mell and Grance, 2009)..

On-Demand Resource Provisioning- This feature enables users to directly obtain services from clouds, such as spawning the creation of a server and tailoring its software, configurations, and security policies, without interacting with a human system administrator (Buyya et al. 2010). This feature eliminates the need for more time-consuming, labour- intensive and human driven procurement processes.

Broad Network Access- Cloud capabilities are available over network through which business solutions or applications can be accessed through mobiles, laptops, tablets, and desktops (Mell and Grance, 2009).

Cloud computing provides its services in cost effective way via internet in a reliable and efficient way (Conde et al., 2013). Cloud computing provides its services in three service models to their users based on the layered architecture namely: 1) Software as a service, 2) Platform as a service, and 3) Infrastructure as a service (Mell and Grance, 2009) as shown in figure2.

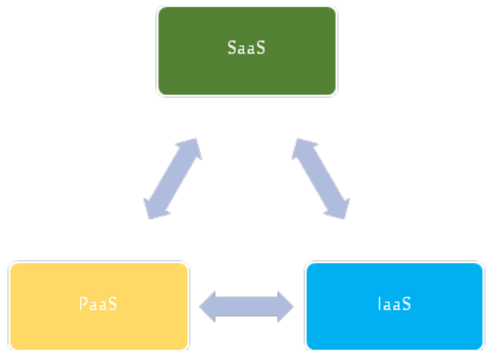

Figure 2. Service models of cloud computing.

Software as a service (Saas) - The Software as a service (Saas) reside on the top of the service models that can be accessed by end users. The users are increasingly shifting from locally installed computer applications to online software services. The use of Saas service model eliminates the burden of software maintenance on local system (Hayes, 2008).

Platform as a service (Paas) -The second abstraction layer, platform as a service (Paas) is the programmable layer of service model used by programmers to develop and deploy codes in the clouds. Developers do not need to know the hardware configuration running in the cloud. The developers can access and program in various types of development environment. Paas platform provides various types of specialized services for data access and authentication (Buyya et al. 2010).

Infrastructure as a service (Iaas) -The Infrastructure as a service (Iaas) is a third layer of service model manages physical resources in the cloud. It manages the virtual machines and other virtualized resources running in the cloud computing environment on demand (Sotomayor,et al., 2009) . The infrastructure as a service runs different types of operating system with customized stack of softwares (Nurmi at al., 2009) . This layer provides on- demand provisioning of server running to satisfy varied user requirements.

II.TYPE OF VIRTUALIZATION

With the increase in applications of virtualization concepts across a wide range of areas in computer science, the girth of the definition has been increasing even more. However, just for the discussions in this paper (Chiueh et al., 2005), we use the following relaxed definition:”Virtualization is a technology that combines or divides computing resources to present one or many operating environments using methodologies like hardware and software partitioning or aggregation, partial or complete machine simulation, emulation, time-sharing,and many others”. Although virtualization can, in general, mean both partitioning as well as aggregation, for the purposes of this paper, we shall concentrate on only partitioning problems (as these are much more prevalent). A virtualization layer, thus, provides infrastructural support using the lower-level resources to create multiple virtual machines that are independent of and isolated from each other. Sometimes, such a virtualization layer is also called Virtual Machine Monitor (VMM). Although traditionally VMM is used to mean a virtualization layer right on top of the hardware and below the operating system, we might use it to represent a generic layer in many cases. There can be innumerous reasons how virtualization can be useful in practical scenarios, a few of which are the following:

- Server Consolidation: To consolidate workloads of multiple under-utilized machines to fewer machinesto save on hardware, management, and administration of the infrastructure

- Application consolidation: A legacy application might require newer hardware and/or operating systems.Fulfillment of the need of such legacy applications could be served well by virtualizing the newer hardware and providing its access to others.

- Sandboxing: Virtual machines are useful to provide secure, isolated environments (sandboxes) for running foreign or less-trusted applications. Virtualization technology can, thus, help build secure computing platforms.

- Multiple execution environments: Virtualization can be used to create mutiple execution environments (in all possible ways) and can increase the QoS by guaranteeing specified amount of resources.

- Virtual hardware: It can provide the hardware one never had, e.g. Virtual SCSI drives, Virtual ethernet adapters, virtual ethernet switches and hubs, and so on.

- Multiple simultaneous OS: It can provide the facility of having multiple simultaneous operating systems that can run many different kind of applications.

- Debugging: It can help debug complicated software such as an operating system or a device driver by letting the user execute them on an emulated PC with full software controls.

- Software Migration: Eases the migration of software and thus helps mobility.

- Appliances: Lets one package an application with the related operating environment as an appliance.

- Testing/QA: Helps produce arbitrary test scenarios that are hard to produce in reality and thus eases the testing ofsoftware.

Virtualization is very important for cloud computing and as a result brings another benefit that cloud computing is famous for, scalability. Because each virtual server is allocated only enough computing power and storage capacity that the client needs, more virtual servers can be created. But if the needs grow, more power and capacity can be allocated to that server, or lowered if needed. Clients only pay for how much computing power and capacity they are using, this can be very affordable for most clients, without virtualization, cloud computing as there know it would not exist or would be in a different form. But such is now only in the realm of speculation as virtualization is really here to make Information Technology more affordable for the world (Padhy, 2012). There are many categories of virtualization and researchers have presented a lot of work in these different categories of virtualization. Companies of all sizes are embracing virtualization as a way to cut IT expenses, enhance security, and increase operational efficiency. While the benefits of virtualization are self-evident, many people are still in the dark when it comes to the many different types of virtualization. Virtualization is the ability of a server to run multiple instances of an operating system simultaneously. Each OS is unaware of the existence of the other guest operating systems and runs as if it had sole access to the hardware. There are different types of virtualization depending on requirements (Ercan, 2010).

Data centers filled with servers generate a tremendous amount of heat and consume massive amounts of power. Virtualization is the running of several operating systems on a single host server. The host server runs special software that sits between the server and the guest operating systems and parcels out resources as needed. The guest operating systems are unaware that they are virtualized. Virtualization provides multiple benefits. First it reduces the amount of hardware running in the data center. A properly sized host server can run dozens if not hundreds of virtual servers. Every server that is virtualized is one less physical server consuming power and generating heat. In addition, virtualization reduces the overall footprint of the data center. Where hundreds of servers may have sprawled across the floor, only a few host machines may sit. There are different types of virtualization from several different vendors (Lee, 2014). Basically there are 2 types of virtualization in Server hosting:

1. Hardware

2. Software

However, with the technology advancements, there are many new concepts of virtualization, which have emerged. Presently, the following categories of virtualizations exist:

- Hardware Virtualization

- Software Virtualization
- Desktop virtualization
- Server Virtualization
- Storage Virtualization
- Memory virtualization
- Network Virtualization
- Data Virtualization

III. RESEARCH ISSUES IN VIRTUALIZATION IN CLOUD COMPUTING

Virtualization plays a very important role in the cloud computing technology, normally in the cloud computing, users share the data present in the clouds like application etc, but actually with the help of virtualization users shares the Infrastructure. The main usage of Virtualization Technology is to provide the applications with the standard versions to their cloud users, suppose if the next version of that application is released, then cloud provider has to provide the latest version to their cloud users and practically it is possible because it is more expensive. To overcome this problem we use basically virtualization technology, By using virtualization, all severs and the software application which are required by other cloud providers are maintained by the third party people, and the cloud providers has to pay the money on monthly or annual basis. There are a number of components used to build the cloud infrastructure. At the lowest layer there are actual hardware components like servers, network attached storage and network components. In order to limit the possibility of spreading an infection, networks need be properly separated into multiple DMZs with limiting rules of connectivity between two networks (Khajehei, 2014). The very core of cloud computing is virtualization, which is used to separate a single physical machine into multiple virtual machines in a cost-effective way. Don't get me wrong, running and operating a cloud is certainly possible without virtualization, but requires more work and time to actually pull it off; by using virtualization, there're basically getting a lot of the work done for free. With virtualization, a number of virtual machines can run on the same physical computer, which makes it cost-effective, since part of the physical server resources can also be leased to other tenants. Such virtual machines are also highly portable, since they can be moved from one physical server to the other in a manner of seconds and without downtime; new virtual machines can also be easily created. Another benefit of using virtualization is the location of virtual machines in a data center it doesn't matter where the data center is located and the virtual machine can also be copied between the data centers with ease.

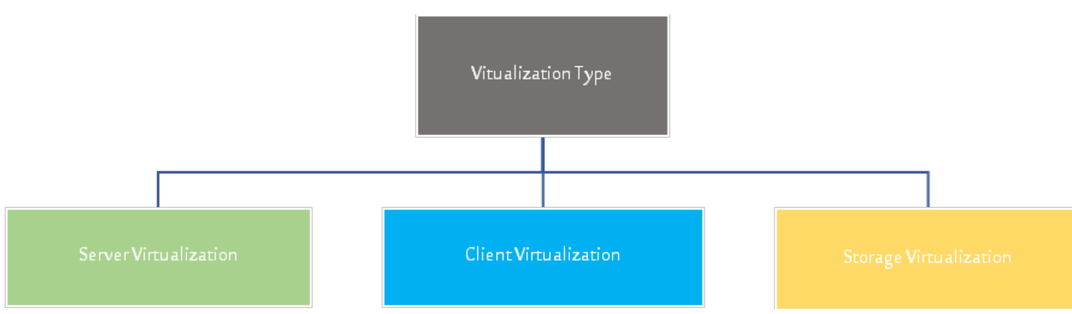

Figure3. Types of Virtualization

Mainly Virtualization means, running multiple operating systems on a single machine but sharing all the hardware resources. And it helps us to provide the pool of IT resources so that we can share these IT resources in order get benefits in the business. Virtualization is a kind of technology that is rapidly transforming the IT landscape and has changed the way people compute. It reduces hardware utilization, saves energy and costs and makes it possible to run multiple applications and various operating systems on the same SERVER at the same time(Lee, 2014). It increases the utilization, efficiency and flexibility of existing computer hardware.

Virtualization provides various benefits including saving time and energy, decreasing costs and minimizing overall risk.

- Provides ability to manage resources effectively.
- Increases productivity, as it provides secure remote access.
- Provides for data loss prevention.

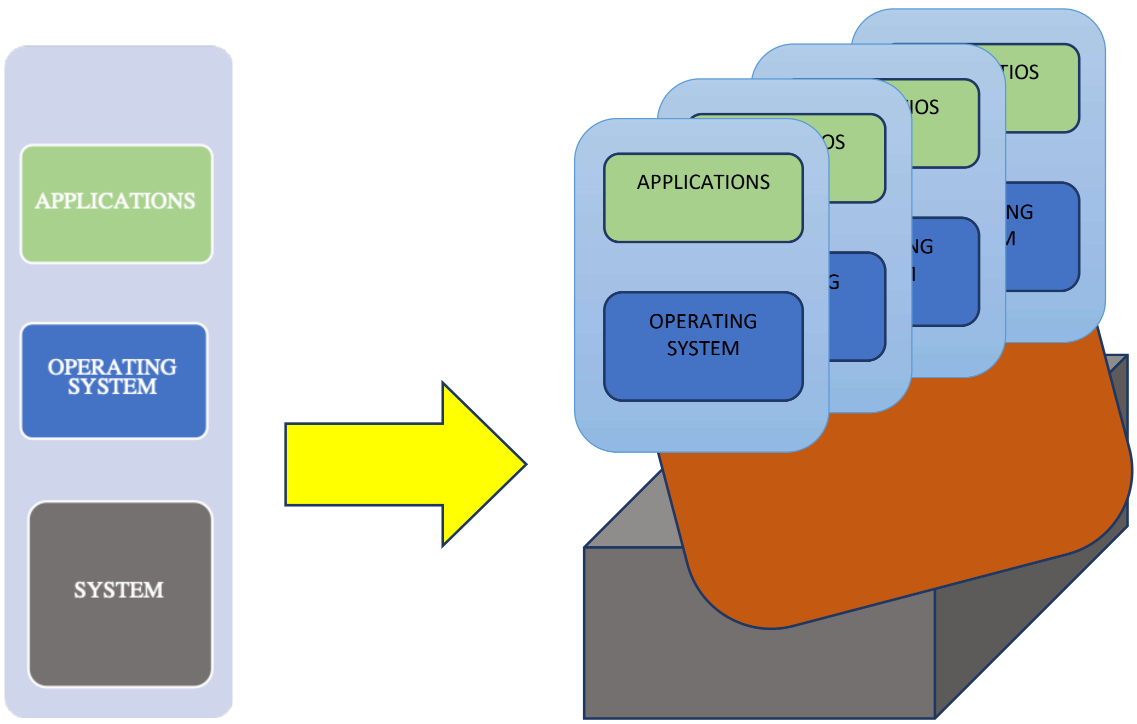

Figure4. Role of Hypervisor in Virtualization

While the main benefit of virtualization is cost-effectiveness, customization increases the cost of such an environment with an advantage of increased security. If there customize the cloud virtualized environments too much, the costs go up considerably. There are three basic cloud computing service models: SaaS, PaaS and Iaas, and three basic cloud computing deployment models: public, community and private. When choosing the right combination there would like to use in our services, there need to be aware that by having greater control over the cloud, there also exists the possibility of increasing the security of the cloud. There can increase security of the IaaS service model as compared to SaaS, because we have greater control over it the same is true for deployment models: we can increase security of a private cloud as compared to a public cloud.

Productivity: Virtualization can enhance productivity in development and testing of HPC applications and systems, in several ways. With authorization, the hypervisor can allow one VM to monitor the state (including memory), interrupts and communications (including calls to the hypervisor) of another VM, for debugging or performance analysis. Because this introspection code runs outside the VM being monitored, it can look at any path in the OS or application, without the limitations that may occur when a debugger is inside, and shares hardware state with, the OS or application being debugged. This may be especially beneficial when developing or using a specialized minimal library OS for a specific HPC application, mentioned above (King et al., 2005). The hypervisor(s) can provide a virtual cluster of VMs, one for each node in a specific configuration of an HPC application that uses a cluster programming model like MPI. The VMs can be allocated across fewer real nodes, even a single node, to test at realistic scale but with fewer resources. Virtualization allows testing on the machine to be used for production, using only modest resources, simultaneously with production applications that are using most of the resources. Moving an application from testing to production is only a matter of allocating adequate real resources to the virtual nodes (VMs) used to execute the application. Productivity may also be enhanced by using a virtual cluster, running multiple copies of the OS and application, to achieve scaling in an application originally written for a non-scalable OS, avoiding the rewrite for another OS. This was a primary motivation of the Disco system (Bugnion et al., 1997). The decoupling of hardware and operating system can significantly reduce the restart cycle of a system. The “virtual reboot” avoids the latencies of hardware re-initialization by the BIOS (Choi et al., 2002). A pre-booted and frozen VM image can be shipped to all nodes in a cluster and significantly reduce the startup time of the system. Similarly, changing a VM image to a different OS and application does not reimpose a hardware initialization latency.

Performance: Virtualization raises two performance issues: the cost of virtualization itself and the performance benefits it offers. Recent hardware and software developments have been at least partially motivated to reduce the cost of virtualization. The most popular microprocessors, including Intel, AMD , and IBM Power, all have hardware features to support virtualization and reduce its performance cost (Mergen et al., 2006). Software paravirtualization, used by Xen (Barham et al., 2003) and IBM hypervisors (http://www.research.ibm.com/hypervisor/.), seeks to present hypervisor interfaces, for example to specify memory address translations, that are functionally equivalent to but less performance-costly to implement than exactly virtualizing the analogous hardware, for example page tables. Software pre-virtualization is a technique of semi-automatically annotating OS code so that it can be adapted for a specific hypervisor at load time but remain compatible to real hardware (LeVasseur et al., 2005). The performance of HPC applications may benefit from virtualization in several ways. Virtualization facilitates specialized OSes that are performance optimized for classes of HPC applications, as discussed above. A hypervisor can guarantee resource allocations to a VM, such as a fixed partition of real memory, a fixed percentage of CPU cycles, or a maximum latency to interrupt handling code for a real time VM. The virtual nodes (VMs) of a virtual cluster can be scheduled to run concurrently, in different real nodes or using different processor cores of one or a few real nodes. This gang scheduling allows a cluster HPC application, while running, to communicate between nodes in real time, as it would without virtualization, which may be crucial to efficient forward progress.

Reliability and Availability: Because of VM isolation, hardware or software failures, as in a processor core, memory, OS or application, directly affect only one VM, unless the hypervisor itself suffers a failure.

If the affected VM cannot recover itself and truly fails, its non-failing hardware resources can be unambiguously reclaimed by the hypervisor and used in restarting the failed VM or by other VMs. Other VMs are not aware of the failure unless they have a communication dependency on the affected VM, but they may run slower if they shared a failed processor core, for example. This fault isolation enhances system reliability and increases the probability of completion of long-running HPC applications, without any special effort by the OSes that run in the VMs. When authorized, introspection allows one VM to capture the complete OS and application state of another VM, either on request or periodically. Because this is the state of resources virtualized at a low level, this state can be used to create another VM, with equivalent but different real resources in the same or another real node, and continue processing. This check point/restart capability enables preemption by high- priority work, inter-node migration of work in a cluster for load balancing, and restart from a previous checkpoint after transient hardware or software failures, if the application is idempotent with respect to such a restart. Recovery from hardware failures is particularly important for very long-running computations. Preemption allows new uses, such as what may be called real-time HPC, where a large number of nodes are preempted for a brief time to compute a result needed immediately. All these scenarios enhance system availability, require little or no effort in the OS or application and are important to HPC applications, because they prevent loss of progress in long-running HPC applications (Mergen et al., 2006).

Security Concerns: Virtualization provides many benefits when used in a cloud computing platform, such as cost-efficiency, increased uptime, improved disaster recovery, application isolations (Sherry et al., 2012), etc. Despite all the advantages, when virtualization is used there are also a number of security concerns outlined below.

- The Hypervisor: the hypervisor is used to separate operating systems of virtual machines from the physical hardware. When there adds a new virtual machine on top of the same physical machine, there must ensure that the operating system has the latest security updates installed and the software has been properly patched. Once the attacker has gained administrative access to the guest operating system, there can further advance to exploiting the vulnerabilities that exist in a hypervisor by successfully exploiting such a vulnerability, the attacker can gain complete access to the host physical machine. Even more, once the attacker has access to the host hypervisor, there can easily access all the virtual machines running on that physical machine. In order to understand the importance of having a secure hypervisor, there must first say a few words about it. A hypervisor was specifically built for virtualization in mind and has therefore a required limited function specific to the hypervisor (Wang et al., 2010). It doesn't need to be as complete as a whole operating system, because it must be able to perform a limited set of functions; therefore, it also has a smaller number of entry points that can be attacked and exploited. The concerns are great, since hypervisors are rarely updated and therefore new security vulnerability would most likely be unpatched for a long period of time. In 2006 two root kits were developed to demonstrate how a rootkit can be used to take control over the host machine – hypervisor and all of its virtual machines: Blue Pill in Sub Virtual.

- Resource Allocation: when a certain resource, like data storage or physical memory, is allocated to a virtual machine, the virtual machine can use it for storing its data. If the resources are later reallocated to another virtual machine due to this virtual machine not being needed anymore and is therefore removed, the new virtual machine could read the data from hard drive or memory. The data on such resources must be properly deleted when being transitioned from one virtual machine to the next.

- Virtual Machine Attacks: if attackers can successful gain access to one virtual machine by exploiting vulnerability in one of the applications running in that virtual machine, he can attack other applications running in different virtual machines over the network. If the virtual machine is running on the same physical host as the compromised virtual machine, it may be hard to detect such network attacks. There must be able to monitor the traffic coming to and from each virtual machine on the same physical host.

- Migration Attacks: when there migrates a virtual machine, the migration usually consists of transferring the whole virtual machine over the network from one physical machine to the other. For a successful attack, the attacker already needs to have access to the network where the migrations are taking place and reading/writing into memory at the time of migration. When the migration communication is not encrypted, the attacker can perform a MITM attack to copy the whole virtual machine VHD file. Mitigation of Security Concerns: Following can mitigate the security concerns as outlined previously (Pek, 2013).

- The Hypervisor: attacks on the hypervisor can easily be prevented by regularly updating it. Since there is a low number of vulnerabilities in the hypervisor itself, there is also a low number of times there will have to update, but it's important that can do. This will not protect against zero-day vulnerabilities that may still be present in a hypervisor, but at least there're protected from known vulnerabilities that exist in the wild.

- Resource Allocation: prior to allocating resources to new virtual machine, the resources need to be properly protected. When assigning the physical memory to a new virtual machine, the memory needs to be filled with zeros in order to prevent any data used in the previous virtual machine from being leaked. The same is true for hard drive data, which might still be present even after easing and recreating the partition. A new virtual machine can use different forensic tools to read the whole unstructured contents of the hard drive into the file and analyze it for possible leaked data. To prevent this, there must also overwrite the hard drive used by the old partition with zeros when assigning it to the new virtual machine.

- Virtual Machine Attacks: there must ensure that the underlying virtualization software can differentiate against traffic coming from and going to different virtual machines. The traffic also needs to be analyzed for possible known attacks; one way of doing this is by using port mirroring to copy any traffic on a specific port on a switch to another port (mirroring) for data to be analyzed by IDS/IPS.

- Migration Attacks: migration attacks can be handled by using proper mechanisms for detecting and preventing MITM attacks. To enhance security, the migrations should also be done over a secure communication channel like TLS.

IV.CONCLUSION

In this paper a study has been carried out to find research issues in virtualization in cloud computing. Virtualization is very important for cloud computing and as a result brings another benefit that cloud computing is famous for, scalability. Because each virtual server is allocated only enough computing power and storage capacity that the client needs, more virtual servers can be created. But if the needs grow, more power and capacity can be allocated to that server, or lowered if needed. Clients only pay for how much computing power and capacity they are using, this can be very affordable for most clients.Without virtualization, cloud computing as there know it would not exist or would be in a different form. Cloud Computing may be the most overused exhortations in the tech industry, often thrown around as an umbrella term for a wide array of different platforms, systems, and services. It's thus not entirely surprising that there's a great deal of confusion regarding what the term actually entails. here are several different breeds of virtualization, though all of them share one thing in common: the end result is a survey to find main research issues in cloud computing. In most cases, virtualization is generally accomplished by dividing a single piece of hardware into two or more 'segments.' Each segment operates as its own independent environment.

REFERENCES

[1] Buyya, R., Yeo, C. S., Venugopal, S., Broberg, J., &Brandic, 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.

[2] Sbihi, B., Eddine, K., &Aknin, N. (2013). Towards a collaborative learning process based on the hybrid cloud computing and Web 2.0 tools. International Journal of Engineering and Technology, 1(2).

[3] Conde-González, M. Á., García-Peñalvo, F. J., Zangrando, V., García-Holgado, A., Seoane-Pardo, A. M., Alier, M., ... &Brouns, F. (2013). Enhancing informal learning recognition through TRAILER project. In Proceedings of the Workshop on Solutions that Enhance Informal Learning Recognition (WEILER 2013) (Vol. 1039, pp. 21-30). CEUR-WS.

[4] Pocatilu, P., Alecu, F., &Vetrici, M. (2010). Measuring the efficiency of cloud computing for e-learning systems. Wseas transactions on computers, 9(1), 42-51.

[5] Zaharescu, E., &Zaharescu, G. A. (2012). Enhanced virtual e-learning environments using cloud computing architectures. International Journal of Computer Science Research and Application, 2(1), 31-41.

[6] Zissis, D., &Lekkas, D. (2012). Addressing cloud computing security issues. Future Generation computer systems, 28(3), 583-592.

[7] Vaquero, L. M., Rodero-Merino, L., Caceres, J., & Lindner, M. (2008). A break in the clouds: towards a cloud definition. ACM SIGCOMM Computer Communication Review, 39(1), 50-55.

[8] Buyya, R; Pandey, S; Vecchiola, C (2009), "Cloudbus Toolkit for Market-Oriented Cloud Computing, Cloud Computing", Proceedings, 2009, 5931, No. 21, pp. 24 – 44

[9] Cai C, Wang L, Khan SU, Jie T (2011), "Energy-Aware High Performance

[10] Addis B., Ardagna D., Panicucci B., Squillante M., Zhang L. (2013), "A Hierarchical Approach for the Resource Management of Very Large Cloud Platforms", IEEE Transactions on Dependable and Secure Computing, Vol. 10, No. 5, pp. 253-272, 2013

[11] Ali, Q., I., Alnawar, J., Mohammed (2013), "Optimization of Power Consumption in Cloud Data Centres Using Green Networking Techniques", Received: 2-10-2013 Accepted: 9- 6-2013, Al-Rafidain Engineering, Vol.22, No. 2, March 2014, pp. 13-27.

[12] Andersen, D. G., Franklin, J., Kaminsky, M., Phanishayee, A., Tan, L., &Vasudevan, V. (2009, October),"FAWN: A Fast Array Of Wimpy Nodes", In Proceedings of the ACM SIGOPS 22nd symposium on Operating systems principles, ACM, pp. 1-14.

[13] Ardagna D., Trubian M., Zhang L. (2007), "SLA Based Resource Allocation Policies in Autonomic Environments", Journal of Parallel and Distributed Computing, Vol. 67, No. 3, pp. 259-270, 2007

[14] Ardagna, D., Panicucci, B., Trubian, M., & Zhang, L. (2012), "Energy-Aware Autonomic Resource Allocation in Multi-Tier Virtualized Environments", Services Computing, IEEE Transactions on, Vol. 5, No. 1, pp. 2-19.

[15] Armbrust, M., Fox, A., Griffith, R., Joseph, A., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Stoica, I. &Zaharia, M.(2009), "Above the Clouds: A Berkeley View of Cloud Computing", Electrical Engineering and Computer Sciences University of California at Berkley, USA, Technical Report UCB/EECS-2009-28, February 10, 2009, pp. 1-23,

[16] Auer, G., Giannini, V., Desset, C., Godor, I., Skillermark, P., Olsson, M., ...&Fehske, A. (2011). How Much Energy Is Needed To Run A Wireless Network?" IEEE Wireless Communications, Vol. 18, No. 5, pp.40-49.

[17] Baliga, J., Ayre, R. W., Hinton, K., & Tucker, R. S. (2011), "Green Cloud Computing: Balancing Energy in Processing, Storage, and Transport", Proceedings of the IEEE, Vol. 99, No. 1, pp.149-167.

[18] Baliga, J., Ayre, R., Hinton, K., & Tucker, R. S. (2011), "Energy Consumption in Wired and Wireless Access Networks", IEEE Communications Magazine, Vol. 49, No.6, pp.70-77.

[19] Baliga, J., Ayre, R., Hinton, K., Sorin, W. V., & Tucker, R. S. (2009), "Energy Consumption in Optical IP Networks", Journal of Lightwave Technology, Vol. 27, No. 13, pp. 2391-2403

[20] Barroso, L. A. (2005), "The Price of Performance", Queue, Vol. 3, No. 7, pp. 48–53, 2005.

[21] Barroso, L. A., &Hölzle, U. (2007), "The Case for Energy-Proportional Computing", Computer, pp. 33–37.

[22] Barroso, L. A., Clidaras, J., and Hölzle, U. (2013), "The Datacenter as A Computer: An Introduction to the Design of Warehouse-Scale Machines", Synthesis lectures on computer architecture, Vol. 8, No. 3, pp. 1-154.

[23] Beaty, D., & Schmidt, R. (2004), "Back to the Future: Liquid Cooling: Data Center Considerations", Ashrae Journal, Vol. 46, No. 12, pp 1-42.

[24] Beloglazov A., Abawajy J., Buyya R. (2012), "Energy-Aware Resource Allocation Heuristics for Efficient Management of Data Centers for Cloud Computing", Future Generation Computer Systems, Vol. 28, No. 5, pp.755–768.

[25] Beloglazov, A., Buyya, R., Lee, Y. C., &Zomaya, A. (2011), "A Taxonomy and Survey of Energy-Efficient Data Centers and Cloud Computing Systems", Advances in computers, Vol. 82, No. 2, pp. 47-111.

[26] Benini, L., Bogliolo, A., & De Micheli, G. (2000), "A Survey of Design Techniques for System-Level Dynamic Power Management", IEEE transactions on very large scale integration (VLSI) systems, Vol. 8, No. 3, pp. 299-316.
[27] Bernsmed, K., Jaatun, M. G., Meland, P. H., &Undheim, A. (2011, August), "Security SLAs for Federated Cloud Services", In Availability, Reliability and Security (ARES), 2011 Sixth International Conference on, IEEE, pp. 202-209.
[28] Bianchini, R., &Rajamony, R. (2004), "Power and energy management for server systems", Computer, Vol. 37, No. 11, pp. 68-76
[29] Bianco P.,Lewis G. A., Merson P. (2008), "Service Level Agreements in Service- Oriented Architecture Environments", Software Architecture Technology Initiative, Integration of Software-Intensive Systems Initiative, Software Engineering Institute, CMU/SEI-2008-TN-021,September 2008, pp. 1-49, http://www.sei.cmu.edu
[30] Bianzino, A. P., Chaudet, C., Rossi, D., &Rougier, J. L. (2012). A Survey of Green Networking Research", IEEE Communications Surveys & Tutorials, Vol. 14, No. 1, pp. 3-20.
[31] Blough, B., Bean, J., Jones, R., Patterson, M., Jones, R., & Salvatore, R. (2011), "Qualitative Analysis of Cooling Architectures for Data Centers", The Green Grid White Paper, pp. 1-30, www.thegreengrid.org.
[32] Bolla, R., Bruschi, R., Davoli, F., &Cucchietti, F. (2011), "Energy Efficiency in the Future Internet: A Survey of Existing Approaches and Trends in Energy-Aware Fixed Network Infrastructures. IEEE Communications Surveys & Tutorials, Vol. 13, No. 2, pp. 223-244.
[33] Bouley, D., & Brey, T. (2009), "Fundamentals of data centre power and cooling efficiency zones", The Green Grid White Paper, 21, pp. 1–17.
[34] Brooks, D., &Martonosi, M. (2001), "Dynamic Thermal Management for High- Performance Microprocessors", In High-Performance Computer Architecture, 2001, HPCA, the Seventh International Symposium on, IEEE, pp. 171-182.
[35] Brown, D. J., & Reams, C. (2010), "Toward Energy-Efficient Computing", Communications of the ACM, Vol. 53, No.3, pp. 50-58.
[36] Kop, R., & Carroll, F. (2012). Cloud computing and creativity: Learning on a massive open online course. European Journal of Open, Distance and E-learning, 15(2).
[37] Mell, P., &Grance, T. (2011). The NIST definition of cloud computing.
[38] Hayes, B. (2008). Cloud computing. Communications of the ACM, 51(7), 9-11
[39] Buyya, R., Broberg, J., &Goscinski, A. M. (Eds.). (2010). Cloud computing: Principles and paradigms (Vol. 87). John Wiley & Sons.
[40] Sotomayor, B., Montero, R. S., Llorente, I. M., & Foster, I. (2009). Virtual infrastructure management in private and hybrid clouds.IEEE Internet computing, 13(5), 14-22.
[41] Keller, E., Szefer, J., Rexford, J., & Lee, R. B. (2010, June), "NoHype: Virtualized Cloud Infrastructure without the Virtualization. In ACM SIGARCH Computer Architecture News, Vol. 38, No. 3, pp. 350-361, ACM.
[42] Yuan, H., Kuo, C. C. J., & Ahmad, I. (2010, August), "Energy Efficiency in Data Centers and Cloud-Based Multimedia Services: An Overview and Future Directions", In Green Computing Conference, 2010 International, IEEE, pp. 375-382.
[43] Sonnek, J., Greensky, J., Reutiman, R., & Chandra, A. (2010, September), "Starling: Minimizing communication overhead in virtualized computing platforms using decentralized affinity-aware migration", In 2010 39th International Conference on Parallel Processing, IEEE, pp. 228-237.
[44] Lent, R. (2010, December), "Simulating The Power Consumption Of Computer Networks" , In 2010 15th IEEE International Workshop on Computer Aided Modeling, Analysis and Design of Communication Links and Networks (CAMAD), IEEE, pp. 96-100.