Enhancement of Performance and Economy of Data Centers by Virtualization

Surya Teja Marella *, K Karthikeya, V Sai Kushwanth, Akhila Bezawada

Department of Computer Science & Engineering,
Koneru Lakshmaiah Education Foundation, Vaddeswaram, Guntur District, A P, India-522502

* Corresponding author: suryatejamarella@gmail.com

Abstract **- Virtualization is defined as creating a virtual (i.e. not actual physical) version of a server, a storage device, an operating system or a network device. The traditional data center has been transformed significantly in recent years due to virtualization. Virtualization has reduced various hardware and energy costs to the benefit of companies that demand more computing power, and has improved performance without increasing infrastructure. At the point when numerous virtual machines keep running on one server, a proficiency of the framework is enhanced and control utilization by the framework is diminished. In effect virtualization does more with less. Though virtualization is profitable to business, many businesses are still unsure whether they should implement it or no. Here, we will study various aspects of virtualization in which the cost of the data center is reduced and the performance of the system is improved due to an implementation of virtualization.**

Keywords - Virtualization Techniques, Data Center, Cloud Computing

I. INTRODUCTION

Virtualization makes helpful IT administrations utilizing assets that are bound to equipment. Virtualization utilizes full capacity of machine by distributing its capacities among connected environments. Let us consider we have 3 physical servers. Each of the physical servers has a dedicated purpose [1]. One of the servers is mail server, other is web server and the third server runs legacy applications. About 30% of server's capacity is utilized. We need to keep all of these servers because the legacy applications are important to our internal operations [2].

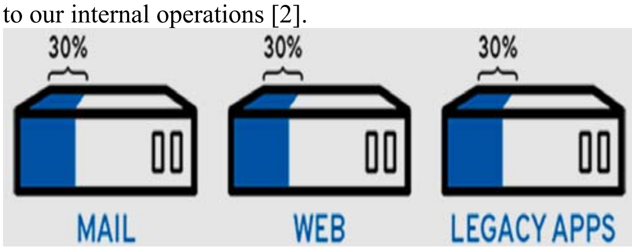

Figure 1. Without Virtualization

Few years ago it was more reliable and easier to run individual tasks on a separate server which means one server, one operating system, one task.

Today with virtualization we can split a mail server into two distinct mail servers. These unique mail servers can handle independent tasks, so that the legacy apps can be migrated. This system uses similar hardware with more efficiency [3].

Considering security we can split the first server to handle various tasks. This strategy will increase the use of the server from 30% to 60% to 90%. Once this strategy is implemented, the empty servers can then be used for other tasks or can be stopped completely to reduce cooling and maintenance tasks [4].

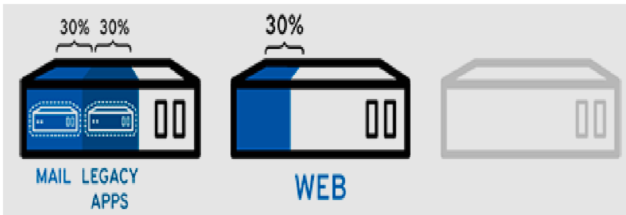

Figure 2. With Virtualization

II. HISTORY AND WORKING OF VIRTUALIZATION

A. History

The technology of virtualization can be traced back to 1960s but it was not adopted till early 2000s. The technologies like hypervisors enable virtualization. The hypervisors were developed decades ago for giving multiple users access to computer simultaneously. The computer performs batch processing [5]. Batch processing used to be a famous computing style in business sector that ran routine tasks faster like payroll. As the time passed, other solutions to the above mentioned problem grew in popularity and virtualization did not One of those elective arrangements was time-sharing. The time-sharing detached clients inside working frameworks. These lead to utilization of working framework like UNIX and it in the long run offered approach to LINUX. In this process virtualization remained an unutilized technology [66].

In 1990s a large portion of the undertakings had physical servers and single-seller IT stacks. They didn't permit inheritance applications to keep running on other merchant's

DOI 10.5013/IJSSST.a.19.06.09 ISSN: 1473-804x online, 1473-8031 print

equipment. At the point when the organizations refreshed their IT surroundings utilizing more affordable servers, working frameworks and applications from assortment of sellers, the companies faced problem of underused physical hardware. This problem was caused because one server could run one task only [7].

This is when people realized importance of virtualization because virtualization solves two problems at a time. First the companies could partition there servers and second the companies could run legacy apps on multiple operating systems and versions. Due to virtualization, companies started to use servers more efficiently and thus costs associated with purchase, set up, cooling and maintenance were reduced. Implementation of virtualization helped reduce vendor lock-in and thus, virtualization became foundation of cloud computing. Today influence of virtualization has reached to a level that companies have specific virtualization the board programming to monitor everything [8].

B. Working

Hypervisor is software that separates physical assets from virtual conditions that require those assets. These hypervisors will sit on the top of OS or can be installed directly in to hardware. Most of the enterprises virtualizes in this way. Hypervisors divide physical resources in such a way that virtual environments can use them.

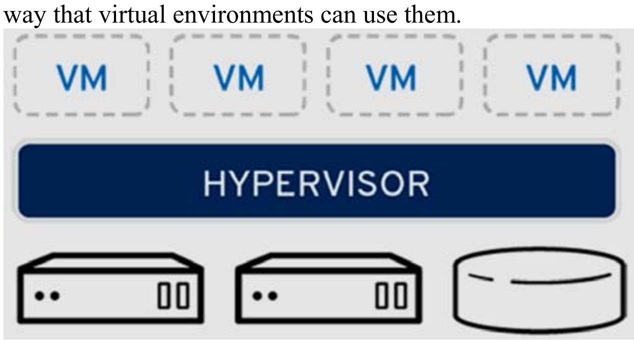

Figure 3. Hypervisor

Resources are divided to many virtual environments from physical environments. Users run computations and interact with the virtual machine [9]. VM can be considered as a single data file. Like any advanced document, it very well may be moved starting with one PC then onto the next. The record can be opened in one PC and is relied upon to deal with the equivalent. In the virtual condition at whatever point a client or a program enrols a guidance requesting an extra asset from physical condition Hypervisor comes into the image. Hypervisor sends the demand to physical framework and recoveries the adjustments in store memory. This procedure occurs at near local speed. For the most part the demand is sent by means of an open source hypervisor dependent on the Kernel-based Virtual Machine for example KVM [10].

III. TYPES OF VIRTUALIZATION

A. Data Virtualization

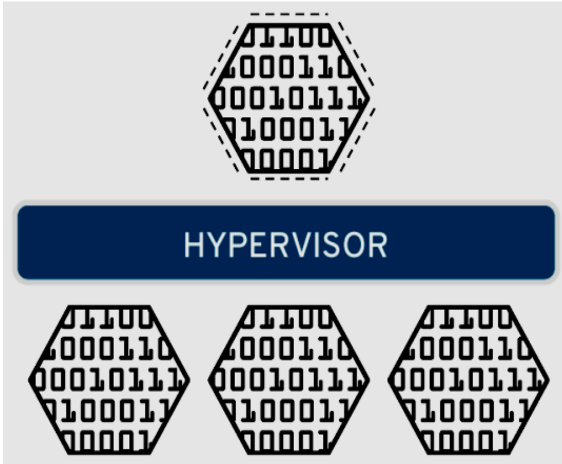

Figure 1. Data Virtualization

In this era of Internet, where data is spread everywhere it is possible to collect all data into a single source. Data virtualization treats data as a dynamic supply [11]. Data virtualization gives handling abilities to unite information from various fluctuated sources, to effectively suit new wellsprings of information and to change information according to client prerequisites. Information virtualization devices like Ret Hat or JBoss sit before different information sources and treat them as a solitary information source. Data virtualization enables enterprises to deliver the needed data in required form at any time to any authorised user[12].

B. Desktop Virtualization

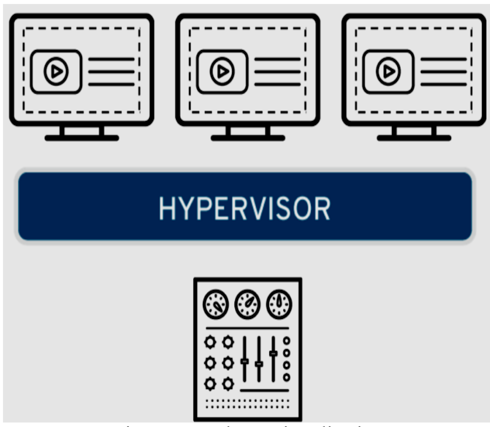

Figure 5. Desktop Virtualization

The Desktop virtualization can be simply confused with an OS virtualization. It enables client to send more than one working framework on a solitary machine. Work area

DOI 10.5013/IJSSST.a.19.06.09

ISSN: 1473-804x online, 1473-8031 print

virtualization empowers a focal director or a computerized organization device to execute reproduced Desktop conditions to numerous physical machines without a moment's delay. Work area virtualization empowers director to play out different number of designs, updates and security minds all accessible virtual Desktop. In traditional Desktop environment it was not possible. It required physical installation, configuration and update on each and every machine [13].

C. Server Virtualization

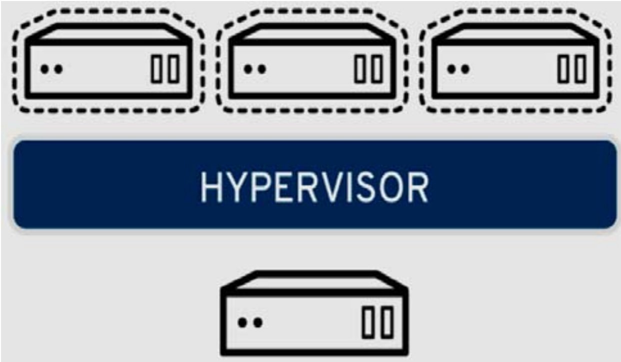

Figure 6. Server Virtualization

Servers can be characterized as PCs that are utilized to run an extensive volume of explicit undertakings great with the goal that alternate PCs like workstations or work areas can do assortment of different assignments, when we virtualizes a server it lets the other computer perform other important tasks. When we virtualizes a server we also partition it so that specific components can perform multiple functions[14].

D. Operating System Virtualization

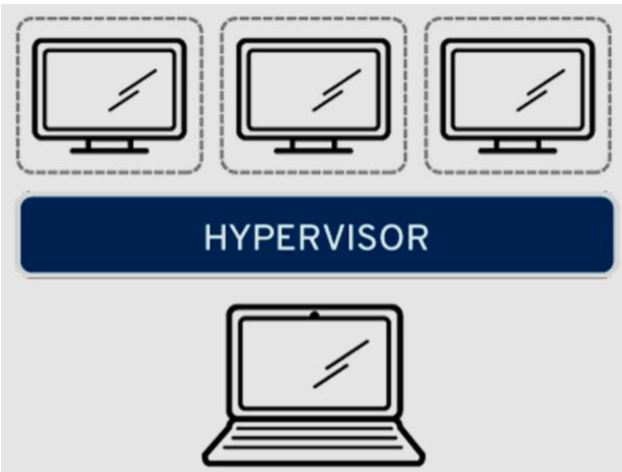

Figure 7. OS Virtualization

Operating system virtualization generally takes place in the kernel which is the central task manager of all available operating systems. It is an effective method to run Linux and Windows environments in same machine along with each other. Enterprises push virtual operating systems to computers. This has following effects [15].

- Since the requirement of out of box capabilities is not there, the bulk hardware costs are reduced.
- Increased security, due to monitoring and isolation of all virtual instances.
- Time spent on IT benefits like programming refreshes is limited [16].

E. Network Functions Virtualization

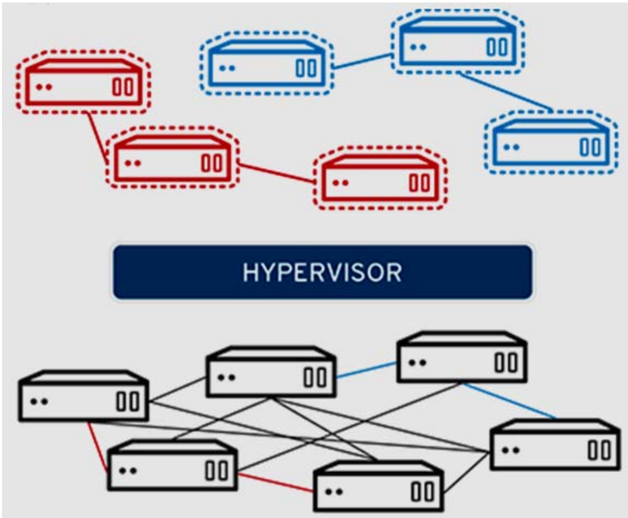

Figure 8. Network Function Virtualization

In Network functions virtualization (NFV) system's key capacities (like index administrations, document sharing, and IP design) are isolated. These key capacities can be circulated among situations [17].

Presently the product capacities are autonomous of the physical machines they once lived on. It is possible that the specific functions can be packed together into a new network and can be assigned to an existing environment. Virtualization of system capacities diminishes the quantity of physical segments of the system like switches, switches, servers, links, and hub. These components of network are needed to create multiple, independent networks. Due to the above mentioned benefits the network function virtualization is very popular in the telecommunications industry [18].

F. Application Virtualization

Application Virtualization encapsulates computer programs from underlying operating system. A completely virtualized application is executed as if it is installed in a traditional sense. The virtual application behaves like a real application by interfacing with original operating system at

DOI 10.5013/IJSSST.a.19.06.09 ISSN: 1473-804x online, 1473-8031 print

the runtime. Virtualization layer is necessary to implement full virtualization. The function of virtualization layer is to intercept all disks operations of virtualized application and redirects them in a single file to a virtualized location. Now the application is working with one single file instead of many files. This makes it easier to run single application on different computers [19].

Various examples of this technology for the Windows platform include:

- Cameyo
- Citrix XenApp
- Numecent Cloudpageing
- Sandboxie
- Symantec Workspace Virtualization

Application virtualization allows applications to run in environment even if the environment does not suit native application. Benefits of application virtualization include reduced system integration and reduced administrative costs. Application virtualization maintains software baseline across multiple computers in an organization. Operating System and other applications are protected from poorly written or buggy code by lesser integration. Application virtualization provides memory protection in some cases. This feature helps to implement principle of least privilege. It does so by removing requirement of administrative privileges for end users. Application virtualization enables incompatible applications to run simultaneously with minimum regression testing against each other. It also has security benefits and allows simple operating system migrations. Application virtualization has also reduced resource utilization [20].

Though the application virtualization has many benefits, it has some limitations also. It is difficult to virtualizes applications that require heavy OS integration. It is mandatory for application virtualization software and the virtualized applications to have correct license. It causes many problems. If the application virtualization is not able to handle Registry-level compatibility issues between applications and operating systems then applications will still allocate memory in the traditional way, even if they are virtualized. It creates requirement of compatibility fixes for a virtualized application.

G. Storage Virtualization

Previously employees used to save files to storage devices. Nowadays storage devices are abstracted from server. IT companies need to choose storage for particular task. Important decisions need to be taken during the process. Differences between block level storage and file level storage are discussed here [21].

IV. BLOCK AND FILE LEVEL STORAGE

Most data centres who have used Storage Area Network, have used block level storage. Block level storage works with servers using iSCSI connectivity mechanisms and industry standard Fibre Channel. Difference between Block level storage and hard drive is that hard drive can be accessible using Fibre Channel or iSCSI. Block level storage is flexible and versatile. Block level storage device creates raw storage volumes and then these storage volumes are used as hard drives. Thus, the block level storage can be used for any kind of application including database storage, file storage and virtual machine file system. In windows, volumes are formatted with NTFS used by VMware servers. Files are shared with users using a file storage device. When a block based volume is created and then an operating system is installed user can share files using native operating system. It is important to remember that when user is using a block based volume, user is actually using a blank hard drive. Replication capabilities of storage devices are used for backup. The organizations usually prefer using third-party backup tools or operating system native backup tools like data protection manager for taking back up of files. As the storage looks like normal hard drive, special back up steps are not necessary. Block level storage is more complicated than file based storage.

A. Applications of Block-Level Storage

- Databases: Block level storage is very efficient for clustering databases because clustered databases require shared storage.
- Exchange: Despite of improvements made by Microsoft for Exchange, the company only supports block level storage.
- VMware: Though VMware can use file level storage through Network File system, it is a common practice to deploy VMware servers that utilize shared volumes on block level storage.
- Server Boot: If the storage device used is efficient enough, then it is possible to configure servers to boot from block level storage [22].

B. File Level Storage

Though block level storage is very flexible, file level storage is preferred for its simplicity. It is implemented where requirement is limited to storing files. Files should be stored in a centralized, available and accessible place. The file level devices are network attached storage devices that provide more space for a lower cost as compared to block level storage. File level storage is accessed using common file level protocols like SMB/CIFS protocols for Windows and NFS protocols for Linux, VMware. In the file level storage, the storage device manages the files and folders of the device. The file level storage device handles user access control and permissions assignment for this purpose. Special handling is required for backup in case of file level storage devices. File level devices are easier to set up as compared to block level devices. File level storage option is not recommended if high level of performance is required. File level devices are preferred where capacity is sole requirement.

B1. Cases for File Storage:

- Mass File Storage: If user's requirement is limited to storing files, in that case file storage should be used.
- mVMware: Storage presented via NFS can be connected with VMware hosts. In this case block level storage is used.

C. Converging Block and File Storage

The block and file storage can be converged. Modern storage devices include both the block and file storage. If user is confused whether to opt for a block or file storage, user can go for a converged or hybrid device which will suit all needs of user.

V. DATA CENTRE OPTIMIZATION BY VIRTUALIZATION

Virtualization makes it possible for experts and skilled IT managers to deploy creative solutions to cost effective utilization of infrastructure. Virtualization separates resources for a service from physical delivery of service. Virtualization runs several application environments on the same machine. These machines in the same environment are isolated from each other.

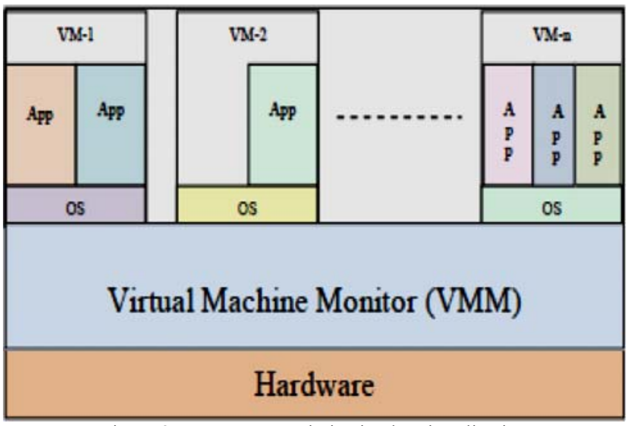

Figure 8. Data centre optimization by virtualization

The system is divided into three components:

1. Hardware: The Hardware provides basic resources like CPU, memory, I/O devices, etc.
2. Virtual Machine Monitor (VMM): The VMM is hardware abstraction layer. It provides interface between hardware and virtual machines. Access to shared resources is controlled by VMM. The VMM schedules processes and allocates processor cycle to them. The VMM layer maps and remaps virtual machines to hardware resources. The VMM scheduling is quite similar to system scheduling. When a computer goes offline or it fails and a new machine comes online, VMM remaps virtual machines according to it. The virtual machines can be easily replicated. Thus, the administrators are able to bring new services online as needed. [23]
3. Virtual Machine: The virtual machine represents a real machine using software. The software provides environment, which can run or host a guest operating system. It also supports various application environments [24].

A. Virtual Machine Migration

This defined as the task of moving a virtual machine from one physical hardware environment to another. It is performed to manage hardware virtualization systems. It is also called teleportation. VM migration is done from one physical machine to another physical machine. It is used for load balancing. For this, the physical machine should be fault tolerant. It also reduces power consumption in cloud data centres [25].

The virtual Machine Migration is divided in to two types, which are:

1. Hot (live) Migration: Virtual machine keeps on running while migrating. The virtual machine does not lose its status. In live migration virtual user does not feel any interruption. In live migration, the state of virtual machine is transferred. The state consists of memory and local file system. Local file system is not transferred [26].
2. Cold (non-live migration): In this type of migration, status of virtual machine loses and user can notice the service interruption. In cold migration, initially, virtual machine is suspended and then its state is transferred. Lastly the virtual machine is resumed at destination host [27].

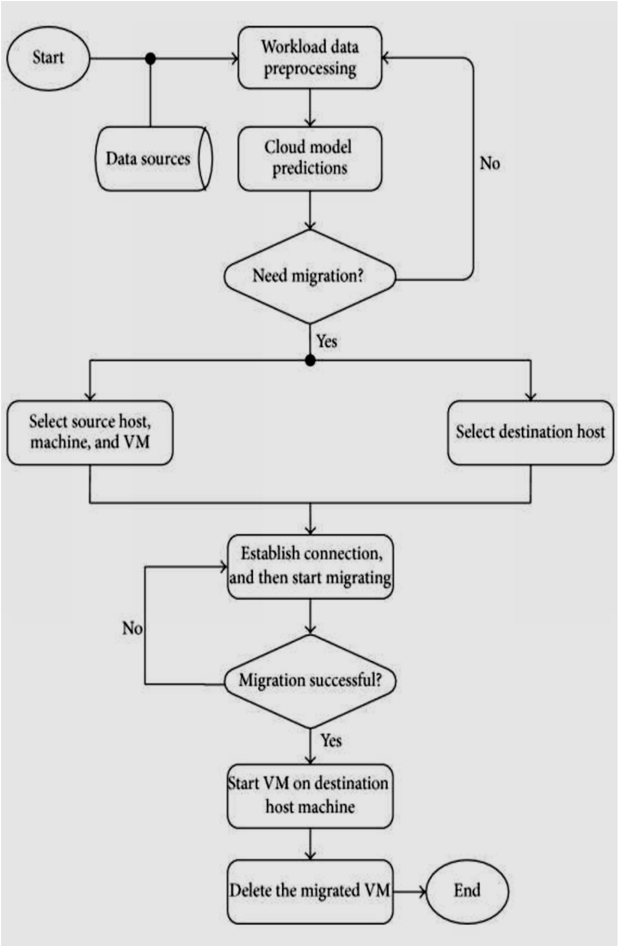

Figure 10. Flow Chart of VM Migration Process

B. Importance of Virtual Machine Migration

- *Load balancing:* Load balancing reduces the inequality of resource usage levels across all the PMs in the cluster. It prevents some machines from getting overloaded in the presence of lightly loaded machines which have sufficient spare capacity. Live migration is used to balance the system. The system load is balanced by migrating virtual machines from overloaded physical machines to under-loaded physical machines [28].

- *Server Consolidation:* Server consolidation algorithms are required in order to reduce server sprawl in data centres. The algorithms perform as VM packing heuristics and try to pack as many virtual machines possible [29]. This improves resource usage by turning off unused or under-utilized machines [30]. Server consolidation results in reduced power consumption and also reduced operational costs for data centre administrators [31].

- *Hotspot & Coldspot Migration:* The discovery of hotspots and coldspot depends on edges which are set by the server farm proprietor. It is likewise founded on the Service Level Agreements indicated by the customers. A higher asset use an incentive close to greatest is set as an upper limit. A least asset utilization esteem is set as the lower limit. Physical machines having asset use esteems past the upper edge have hotspots. The physical machines having use esteems beneath the lower edge have coldspot. The hotspot infers over-use. The coldspot infers under-usage [32].

C. Virtual Machine Scheduling

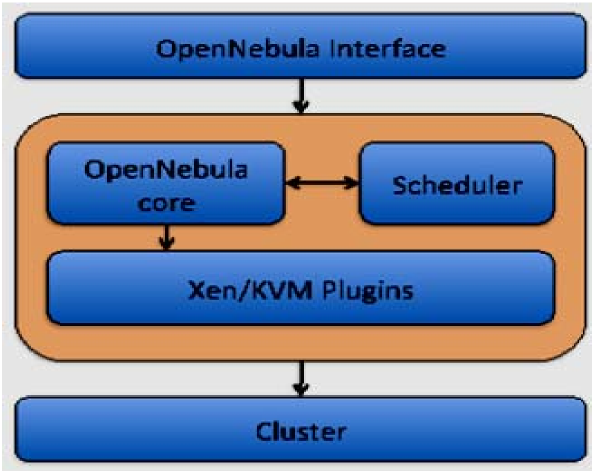

Figure 11. Open Nebula: Default Scheduler

Virtual Machine scheduling is a balancing scenario in which tasks are scheduled according to the given requirements and algorithm used. The virtual machine scheduling algorithms are used to schedule the VM requests to the Physical Machines of the Data Center [33]

Today, different data centres and Physical Machines are available. Leading cloud providers of 2018 are Sales Force, Amazon, Microsoft office 365 and Windows Azure, Oracle Cloud, and Google Apps. In general a scheduling algorithm three levels of working are:

1. Find the appropriate Physical Machine for all virtual machines.
2. Identify the proper provisioning scheme for the virtual machines.
3. Schedule the tasks on the VMs [34].

Scheduling algorithms are primarily named static or dynamic calculations. First-Come-First-Severed (FCFS) is a case of static virtual machine booking calculation while Genetic Algorithm is a case of dynamic virtual machine planning calculation. The OpenNebula default scheduler gives a rank booking approach that places virtual machines on physical machines as indicated by the position of physical machines [35].

Algorithm	Time Efficient	Power Aware	Cost Effective	Security Aware	Memory/ Bandwidth
Round Robin	Yes	No	No	No	No
Genetic Algorithm	No	Yes	Yes	No	No
Match Making	No	No	Yes	No	No
TADCS	Yes	Yes	No	Yes	No
Memory-Aware	No	No	No	Yes	Yes

In this table we have depicted the VM planning and provisioning calculations utilized for streamlining of various variables like Time, Cost, Energy, Security and Bandwidth. Less calculations are accessible to arrangement the VMs for the security of neighboring VMs or hubs [36].

VI. CASE STUDIES

Case Study 1: Virtualization for Financial Services Company

Due to a significant increase in computing, financial services company are running out of space for physical servers. The existing infrastructure is unable to keep up with increasing workload [37]. More time is spent configuring systems and balancing power distribution. The company then implemented server virtualization technology. The technology solved the space problem as well as it improved the efficiency of the system [38]. Today, 75% of the company works on virtualization. It runs 200 virtual machines on 10 physical servers. The virtualization saved power consumption by 33% [39].

Case Study 2: School District Adopts Virtual Exchange 2010 Server

Even little school areas with low metropolitan spending plans can do the change to virtualization. The idea and confirmation are clarified for this situation ponders, in which a school region swung to virtualization to broaden its equipment revive cycle and enhance accessibility [40]. Virtualization enabled the region to decrease the buy of new servers. Awed by the achievement of this task, the region is attempting to virtualize its three area controllers alongside its requesting applications [41].

VII. CONCLUSION AND FUTURE WORK

We have seen in the above case studies that virtualization can save up to 65 % of energy without compromising availability or performance of network security. Virtualization requires standard architecture that can accelerate and save costs for multiple business initiatives and projects. Effective implementation of virtualization will reduce data consumption of overall data centre, but each server is expected to draw more power which may lead to costly downtime in the data centre. The number of servers will be reduced, but each server will be more critical than ever [42]. The support infrastructure cannot move around as needed but the applications can move around dynamically as and when required. Efficiency of operation of the data centre increases optimally due to virtualization. If a virtualized data centre is not managed properly than it can be less reliable. This kind of data centre is even more expensive than its non-virtualized counterpart. Further research on these aspects can be carried out for more details.

REFERENCES

[1] Xing, Y. and Zhan, Y., 2012. Virtualization and cloud computing. In *Future Wireless Networks and Information Systems* (pp. 305-312). Springer, Berlin, Heidelberg.

[2] Malhotra, L., Agarwal, D. and Jaiswal, A., 2014. Virtualization in cloud computing. *J. Inform. Tech. Softw. Eng*, *4*(2).

[3] Swathi, T., Srikanth, K. and Reddy, S.R., 2014. Virtualization in cloud computing. *International Journal of Computer Science and Mobile Computing*, *3*(5), pp.540-546.

[4] Lombardi, F. and Di Pietro, R., 2011. Secure virtualization for cloud computing. *Journal of network and computer applications*, *34*(4), pp.1113-1122.

[5] Wilde, N. and Huber, T., Virtualization and Cloud Computing. *University Of West Florida*.

[6] Marella, Surya Teja, and Akhila Bezawada. "Technology That Made Life Easier: Cloud Computing." (2018): 1625-1635.

[7] Sharma, G.P., Singh, S., Singh, A. and Kaur, R., 2016. Virtualization in Cloud Computing.

[8] Thakur, P. and Mahajan, M., 2016. Virtualization in Cloud Computing. *International Journal of Recent Trends in Engineering & Research (IJRTER)*, *2*, pp.308-315.

[9] Krutz, R.L. and Vines, R.D., 2010. *Cloud security: A comprehensive guide to secure cloud computing*. Wiley Publishing.

[10] Thakral, D. and Singh, M., 2014. Virtualization In Cloud Computing. *JCSMC*, *3*, pp.1262-1273.

[11] Vaezi, M. and Zhang, Y., 2017. Virtualization and cloud computing. In *Cloud Mobile Networks* (pp. 11-31). Springer, Cham.

[12] Guide, P., 2013. Virtualization and Cloud Computing. *online]* *http://www.intel.com/content/dam/www/public/us/en/documents/ guides/cloud-computingvirtualization-building-private-iaas-guide. pdf (accessed February 2015)*.

[13] Rachamalla, S., Mishra, D., Kulkarni, P., Kumar, R., Martínez, A., Gonzalez, A., MapReduce, W.S.E.A.U., Houngkaew, C. and Suzumura, T., Virtualization & cloud computing.

[14] Armbrust, M., Fox, A., Griffith, R., Joseph, A.D., Katz, R., Konwinski, A., Lee, G., Patterson, D., Rabkin, A., Stoica, I. and Zaharia, M., 2010. A view of cloud computing. *Communications of the ACM*, *53*(4), pp.50-58.

[15] Jain, R. and Paul, S., 2013. Network virtualization and software defined networking for cloud computing: a survey. *IEEE Communications Magazine*, *51*(11), pp.24-31.

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

[17] Hadar, E., Vax, N., Jerbi, A. and Kletskin, M., CA Technologies Inc, 2013. *System, method, and software for enforcing access control policy rules on utility computing virtualization in cloud computing systems*. U.S. Patent 8,490,150.

[18] Suen, C.H., Kirchberg, M. and Lee, B.S., Hewlett-Packard Development Co LP, 2014. *Block level storage*. U.S. Patent Application 14/371,709.

[19] Suen, C.H., Kirchberg, M. and Lee, B.S., Hewlett-Packard Development Co LP, 2014. *Block level storage*. U.S. Patent Application 14/371,709.

[20] Lee, L.W., Hines, R.C., Gurkowski, M.J. and Blankenbeckler, D.L., DPHI Acquisitions Inc, 2009. *Block-level storage device with content security*. U.S. Patent 7,549,044.

[21] Jewett, D.E., Radford, A.J., Strand, B.D., Chung, J.D., Jacobson, J.D., Haigler, R.B., Thompson, R.S. and Couch, T.L., Applied Micro Circuits Corp, 2008. *Architecture for providing block-level storage access over a computer network*. U.S. Patent 7,392,291.

[22] Berl, A., Gelenbe, E., Di Girolamo, M., Giuliani, G., De Meer, H., Dang, M.Q. and Pentikousis, K., 2010. Energy-efficient cloud computing. *The computer journal*, *53*(7), pp.1045-1051.

[23] Bichler, M., Setzer, T. and Speitkamp, B., 2006. Capacity planning for virtualized servers.

[24] Nelson, M., VMware Inc, 2009. *Virtual machine migration*. U.S. Patent 7,484,208.

[25] Rangegowda, D. and Fries, R., Microsoft Corp, 2013. *Virtual machine migration*. U.S. Patent 8,479,194.

[26] Nelson, M., VMware Inc, 2010. *Virtual machine migration*. U.S. Patent 7,680,919.

[27] Nelson, M., VMware Inc, 2012. *Virtual machine migration*. U.S. Patent 8,260,904.

[28] Baran, M.E. and Wu, F.F., 1989. Network reconfiguration in distribution systems for loss reduction and load balancing. *IEEE Transactions on Power delivery*, *4*(2), pp.1401-1407.

[29] Verma, A., Dasgupta, G., Nayak, T.K., De, P. and Kothari, R., 2009, June. Server workload analysis for power minimization using consolidation. In *Proceedings of the 2009 conference on USENIX Annual technical conference* (pp. 28-28). USENIX Association.

[30] Speitkamp, B. and Bichler, M., 2010. A mathematical programming approach for server consolidation problems in virtualized data centers. *IEEE Transactions on services computing*, *3*(4), pp.266-278.

[31] Padala, P., Zhu, X., Wang, Z., Singhal, S. and Shin, K.G., 2007. Performance evaluation of virtualization technologies for server consolidation. *HP Labs Tec. Report*, *137*.

[32] Vogwill, T., Fenton, A., Buckling, A., Hochberg, M.E. and Brockhurst, M.A., 2009. Source populations act as coevolutionary pacemakers in experimental selection mosaics containing hotspots and coldspots. *The American Naturalist*, *173*(5), pp.E171-E176.

[33] Kaur, P. and Rani, A., 2015. Virtual machine migration in cloud computing. *International Journal of Grid Distribution Computing*, *8*(5), pp.337-342.

[34] Kim, H., Lim, H., Jeong, J., Jo, H. and Lee, J., 2009, March. Task-aware virtual machine scheduling for I/O performance. In *Proceedings of the 2009 ACM SIGPLAN/SIGOPS international conference on Virtual execution environments*(pp. 101-110). ACM.

[35] Lago, D.G.D., Madeira, E.R. and Bittencourt, L.F., 2011, December. Power-aware virtual machine scheduling on clouds using active cooling control and DVFS. In *Proceedings of the 9th International Workshop on Middleware for Grids, Clouds and e-Science* (p. 2). ACM.

[36] Ongaro, D., Cox, A.L. and Rixner, S., 2008, March. Scheduling I/O in virtual machine monitors. In *Proceedings of the fourth ACM SIGPLAN/SIGOPS international conference on Virtual execution environments* (pp. 1-10). ACM.

[37] Hu, J., Gu, J., Sun, G. and Zhao, T., 2010, December. A scheduling strategy on load balancing of virtual machine resources in cloud computing environment. In *Parallel Architectures, Algorithms and Programming (PAAP), 2010 Third International Symposium on* (pp. 89-96). IEEE.

[38] Xu, C., Gamage, S., Rao, P.N., Kangarlou, A., Kompella, R.R. and Xu, D., 2012, June. vSlicer: latency-aware virtual machine scheduling via differentiated-frequency CPU slicing. In *Proceedings of the 21st international symposium on High-Performance Parallel and Distributed Computing* (pp. 3-14). ACM.

[39] Xi, S., Xu, M., Lu, C., Phan, L.T., Gill, C., Sokolsky, O. and Lee, I., 2014, October. Real-time multi-core virtual machine scheduling in xen. In *Embedded Software (EMSOFT), 2014 International Conference on* (pp. 1-10). IEEE.

[40] Xi, S., Xu, M., Lu, C., Phan, L.T., Gill, C., Sokolsky, O. and Lee, I., 2014, October. Real-time multi-core virtual machine scheduling in xen. In *Embedded Software (EMSOFT), 2014 International Conference on* (pp. 1-10). IEEE.

[41] Lin, C.C., Liu, P. and Wu, J.J., 2011, July. Energy-aware virtual machine dynamic provision and scheduling for cloud computing. In *2011 IEEE 4th International Conference on Cloud Computing* (pp. 736-737). IEEE.

[42] Jang, J.W., Jeon, M., Kim, H.S., Jo, H., Kim, J.S. and Maeng, S., 2011. Energy reduction in consolidated servers through memory-aware virtual machine scheduling. *IEEE Transactions on Computers*, *60*(4), pp.552-564.