The Role of Virtualization towards Green Computing and Environmental Sustainability

Vincent Motochi, Samuel Barasa, Patrick Owoche, Franklin Wabwoba

Abstract— **The environment has become a key concern by the entire world and global warming is increasingly attracting attention in many conferences. Moreover, energy usage in data centers has become a concern bearing in mind the fact that the more energy is used in data centers, the more it affects the environment with emissions, which eventually cause global warming. This paper studies the role of virtualization towards green computing and environmental sustainability. Server virtualization is emerging as the prominent approach to consolidate applications from multiple applications to one server, with an objective to save energy usage. This research identified the virtualization environments, identified green computing environments, and then established how virtualization could be used to attain environmental sustainability. This paper was developed on an experiment design. The researcher reviewed an empirical experiment to investigate how server virtualization affects the energy usage in physical servers. Through this analysis, the researcher identified a fundamental trade-off between the energy saving from server consolidation and the detrimental effects (e.g., energy overhead and throughput reduction) from server virtualization. This paper found out that a server consumes a substantial amount of energy when idle thus the importance of consolidation. Secondly the energy overhead depends on the type of hypervisor used and the application architecture. Thirdly for a given traffic load, the energy cost can be minimized by launching an optimal number of virtual machines. The analysis and review results show that virtualization brings substantial energy savings, promotes green computing and would be a clear methodology to conserve the environment in the technology world today. Therefore, green computing is a well balanced and sustainable approach towards the achievement of a greener, healthier and safer environment without compromising technological needs of the current and future generations.**

Index Terms— **green computing, virtualization, energy usage, environment**

I. Introduction

Green computing has come out and made ICT to reshape in terms of technologies implementations. According to [1], green computing is (i) the systematic application of ecological-sustainability criteria to the design, production, sourcing, use and disposal of ICT technical infrastructure (ii) as well as within the human and managerial components of ICT infrastructure in order to reduce ICT, business process and supply chain related emissions, waste and water use; improving efficiency and generate green economic rent.

Over the years, information technology (IT) has fundamentally altered our work and life and improved our productivity, economy and social well-being. Today, Enterprises, governments and societies at large have a new important agenda: tackling environmental issues and adopting environmentally sound practices. IT has been contributing to environmental problems, which most people do not realize. Computers and other IT infrastructure consume significant amounts of electricity, which is increasing day by day, placing a heavy burden on our electric grids and contributing to greenhouse gas (GHG) emissions. It has been demonstrated that, by optimizing the data center operations via virtualization, up to 20% of energy consumption can be saved in data centers [2].

We can exploit the power of IT in innovative ways to address mounting environmental issues [3], [4] and make our IT systems – and their use – greener. Therefore, there is need to redesign our data centers by using technologies such as virtualization with a view of leveraging on the reduction of energy usage and subsequently reducing emissions of gases into the air that affect the environment globally. To address the various issues such as resource management in both software and hardware levels to reduce energy consumption. Overload and under load of system is also one of the phenomenon for energy loss [5]. According to [6], the recent studies reveal that the ideal system consumes 70% of energy. The virtualization technology is to resolve those problems [7]. The performance optimization, energy saving, load balancing and on demand resource allocation can be achieved by using good resource management scheme [8]. Virtualization is seen today as a solution to some environmental issues being raised and therefore contributing towards environmental sustainability.

Data centers are plagued with thousands of servers performing the processes for businesses and end users to facilitate and accomplish large business goals. In so doing these servers remain almost 90% idle most of the time performing nothing but utilizing huge energy and generating enormous amount of CO_2 that is hazardous for environmental sustainability. The carbon dioxide so emitted contribute towards global warming. There was need to develop strategies, policies or frameworks for data centers so that they can cope up with ever growing demands from businesses sustainably.

Manuscript received May, 2017.

Vincent Motochi,, *Department of ICT, County Government of Kakamega, Kakamega, Kenya, +254722910293*

Samuel Barasa, *Department of Computer Science, Kibabii University; Bungoma, Bungoma, Kenya, +254723124465*

Patrick Owoche, *Department of Information Technology, Kibabii University; Bungoma, Bungoma, Kenya, +254722843180*

Franklin Wabwoba, *Department of Information Technology, Kibabii University; Bungoma, Bungoma, Kenya, +254722843180*

This paper therefore evaluated the impact of virtualization on environmental sustainability.

II. Literature Review

The practice energy efficiency and green data center is gathering momentum as Kenyan organizations have started realizing its importance in energy conservation and sustainable development. The practice is applied to new technologies that can help in cutting down data center energy costs and in saving energy, which is synonymous to saving money. In fact, it has a big role to play in reducing power consumption in the data centre [9].

Virtualization is a technology that has been there for some time. This concept was firstly introduced by IBM in the 1960s to provide concurrent, interactive access to a mainframe computer—IBM 360, which supports many instances of OSs running on the same hardware platform [10].

Virtualization technology supports multiple OS's running on a single hardware platform, and provides a convenient means to manage the OSs. The OS and applications running on the virtualization management platform are considered as VMs [11]. Server virtualization has become popular in data centers since it provides an easy mechanism to cleanly partition physical resources, allowing multiple applications to run in isolation on a single server. It categorizes volume servers into different resource pools depending on the workloads they perform, and then server consolidation is applied. This technique decouples softwares from hardware and splits multiprocessor servers into more independent virtual hosts for better utilization of the hardware resources, allowing services to be distributed one per processor. In server consolidation, many small physical servers are replaced by one large physical server to increase the utilization of expensive hardware resources, reducing the consumption of energy and emission of CO_2 [12].

The following virtualization technologies are available:

1) SaaS Level: Since SaaS providers mainly offer software installed on their own data centers or resources from IaaS providers, the SaaS providers need to model and measure energy efficiency of their software design, implementation, and deployment. For serving users, the SaaS provider chooses the data centers, which are not only energy efficient but also near to users. The minimum number of replicas of user's confidential data should be maintained using energy-efficient storage.

2) PaaS level: PaaS providers offer in general the platform services for application development. The platform facilitates the development of applications, which ensures system wide energy efficiency. This can be done by inclusion of various energy profiling tools such as JouleSort. It is a software energy efficiency benchmark that measures the energy required to perform an external sort. In addition, platforms itself can be designed to have various code level optimizations which can cooperate with underlying complier in energy efficient execution of applications. Other than application development, Cloud platforms also allow the deployment of user applications on Hybrid Cloud. In this case, to achieve maximum energy efficiency, the platforms profile the application and decide which portion of application or data should be processed in house and in Cloud.

3) IaaS level: Providers in this layer plays most crucial role in the success of whole Green Architecture since IaaS level not only offer independent infrastructure services but also support other services offered by Clouds. By using virtualization and consolidation, the energy consumption is further reduced by switching-off unutilized server. Various energy meters and sensors are installed to calculate the current energy efficiency of each IaaS providers and their sites. This information is advertised regularly by Cloud providers in Carbon Emission.

Green IT, also known as green computing, is the study and practice of designing, manufacturing and using computers, servers, monitors, printers, storage devices and networking and communications systems efficiently and effectively, with zero or minimal impact on the environment [13], [9]. Green IT is also about using IT to support, assist and leverage other environmental initiatives and to help create green awareness [13]. Thus, green IT encompasses hardware, software, tools, strategies and practices that improve and foster environmental sustainability. Green Computing [14] is defined in various contexts, environmentally, socially and politically with respect to effective and efficient use of energy to achieve competitive advantage in terms of an energy-cost saving strategy, and reduction to carbon emission/footprints, recyclability, biodegradability, and minimal impact to the environment.

Advantages of Green Computing

- Reduced energy usage from green computing techniques translates into lower carbon dioxide emissions, stemming from a reduction in the fossil fuel used in power plants and transportation.
- Conserving resources means less energy is required to produce, use, and dispose of products.
- Saving energy and resources saves money.
- Green computing even includes changing government policy to encourage recycling and lowering energy use by individuals and businesses.
- Reduce the risk existing in the laptops such as chemical known to cause cancer, nerve damage and immune reactions in humans.

Green IT benefits the environment by improving energy efficiency, lowering GHG emissions, using less harmful materials and encouraging reuse and recycling. Thus green IT includes the dimensions of environmental sustainability, the economics of energy efficiency and the total cost of ownership, which includes the cost of disposal and recycling. Increased awareness of the harmful effects of GHG emissions, new stringent environmental legislation, concerns about electronic waste disposal practices and corporate image concerns are driving businesses and individuals to go green.

Green IT is an economic as well as environmental imperative. And, as many green advocates will attest, it is our social responsibility as well [9]. The imminent introduction of more green taxes and regulations will trigger a major increase in demand for green IT products, solutions and services. Hence a growing number of IT vendors and users have begun to develop and offer green IT products and services. As business and governments try to balance growth with environmental risks, we will be legally, ethically and/or socially required to 'green' our IT products, applications, services and practices.

The energy efficient and green data center is gathering momentum as organizations have started realizing its importance in energy conservation and sustainable development. It is applied to new technologies that can help in cutting down data center energy costs and in saving energy, which is synonymous to saving money. It has a big role to play in reducing power consumption in the data centre [9].

To foster green IT, we should understand the following issues: What are the key environmental impacts arising from IT? What are the major environmental IT issues that we must address? How can we make our IT infrastructure, products, services, operations, applications and practices environmentally sound? What are the regulations or standards with which we need to comply? How can IT assist businesses and society at large in their efforts to improve our environmental sustainability?

Lower carbon emission is expected in Cloud computing due to highly energy efficient infrastructure and reduction in the IT infrastructure itself by multi-tenancy. The key driver technology for energy efficient Clouds is "Virtualization," which allows significant improvement in energy efficiency of Cloud providers. Virtualization is the process of presenting a logical grouping or subset of computing resources so that they can be accessed in ways that give benefits over the original configuration. By consolidation of underutilized servers in the form of multiple virtual machines sharing same physical server at higher utilization, companies can gain high savings in the form of space, management, and energy. According to [6], the recent studies reveal that the ideal system consumes 70% of energy. The virtualization technology is to resolve those problems [7]. The performance optimization, energy saving, load balancing and on demand resource allocation can be achieved by using good resource management scheme.

Virtualization is a major strategy to reduce data center power consumption. In virtualization, one physical server hosts multiple virtual servers. Virtualization enables data centers to strengthen their physical server infrastructure by hosting multiple virtual servers on a smaller number of more powerful servers, using less electricity and simplifying the data center. Besides getting much better hardware usage, virtualization reduces data center floor space, makes better use of computing power, and greatly reduces the data center's energy demands. Many enterprises are using virtualization to curtail the huge energy consumption of data centers. In order to tackle the issue of data centers' huge power consumption, leading IT enterprises formed a non-profit group called the Green Grid in February 2007. This group has the responsibility to define and propagate the best energy-efficient practices in data center operation, construction, and design, and drive new user-centric metrics and technology standards.

Rather than a single hardware machine supporting a single OS, that in turn hosts a single server application (Web server, mail server, etc.), a virtualized system enables a single hardware machine running a single OS to host multiple virtual machines, which may or may not be running the same OS. This leads to a single host machine running multiple virtual machines. Virtualization presents the opportunity to scale the service provision within a data centre, make far more efficient use of hardware resources, reduce running costs and reduce energy consumption. Hence, it proved to be a breakthrough technology in data centre IT infrastructure design, enabling a far more efficient use of resources with a much reduced impact in terms of energy consumption.

Virtualized cloud data centre environments introduce a number of operational degrees of freedom not possible in traditional data centres. Firstly, applications hosted in virtualized cloud data centres run on virtualized OSs, that is, the OS does not run on the bare metal but is mediated through a virtualization hypervisor. The practical effect is that applications are no longer bound to a physical host and can be moved around within a pool of servers to optimize the overall power and thermal performance of the pool. Secondly, the loose binding between applications and hosts allows treating a group of hosts as pooled resources, allowing optimizations as a group that were not possible with individual machines, such as powering down some equipment during low demand.

III. METHODOLOGY

The research study used qualitative approach to conduct this research. This included reviewing a published paper based on an empirical finding from an experiment and other papers relating on virtualization in relation to energy usage and green computing.

A physical setup experiment was designed and implemented, which consisted of three identical servers as shown in Figure 1 below. The machines under test were Inspur 3060 servers, each of them contained a quad-core Intel 2.13 GHz Xeon processor, 2 GB RAM, 500 GB hard disk and 1 Gigabit Ethernet card. All of them were connected to a test intranet over a D-link GDS-1024T 1000 Base-T switch. Kill-A-Watt power meters, with a standard accuracy of 0.2%, measured the energy usage of each server. CentOS 5.6-final- x86 64 with Linux kernel 2.6.18 was used as our OS platform for both host and guest systems. Xen 3.0.3 and KVM 83 were installed on server B and C respectively. The 3 guest virtual machines were allocated with 4 VCPUs, 512 MB RAM and 50 GB image. All the software parameters were left intact.

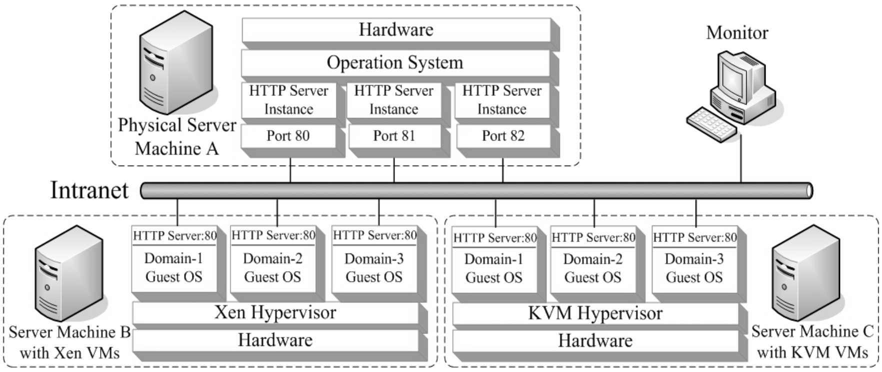

Experimental Setup: 3 configured systems, a non-virtualized server, a Xen-virtualized server, a KVM-virtualized server

Figure 1: Physical experiment setup for three identical servers

The experiment was controlled by another computer, which was also connected to the intranet as a monitor to obtain the benchmarking time, the energy and power consumption. And each server was responsible for gathering its average CPU usage.

The test begun by collecting background energy consumption when all the servers were idle. Following that, a set of local and network traffics were launched to stress all three servers. Detailed test cases were explained below.

IV. Findings

This section presents findings on the objectives as per reviewed desktop research and the empirical findings from the experiments reviewed, which further generalize a few fundamental insights.

A. Impact of Virtualization on Environmental Sustainability: Empirical Experiment Findings

Servers were configured for a simulation experiment as per the first setup when idle and when they were executing requests and findings recorded.

Finding (a): All the servers consumed the same power (about 2.8W) when turned off, but plugged into the power supplies.

Finding (b): When servers were turned on but active VMs stayed idle, the power consumption on different servers varied from each other. The power consumed by the Xen-based server and the physical server were almost the same. In particular, 63.1W (3 active VMs) and 63.0W (2 active VMs) of power was consumed by the Xen-based server, which was 0.47% and 0.32% more than 62.8W consumed by the physical server. While the KVM-based server incurred a much higher overhead, consuming 70.1W (11.6% overhead) for 3 active VMs and 68.8W (9.55% overhead) for 2 active VMs. Moreover, the power usage of the KVM-based server fluctuated within a wider range.

The Finding (b) could be explained by the different impact on CPU and RAM usage by Xen and KVM. The CPU utilization of the idle physical server was generally less than 0.3 %, compared to 0.7-0.8 % for the Xen-based server, and 0.8%-2.6% for the KVM-based server. The extra CPU usage of virtualized servers accounted for a portion of the energy overhead. The rest energy overhead for the KVM-based server could also be attributed to the large memory footprint in KVM, as indicated by the results of memory test in [19].

B. Local Calculation Benchmark

Observations on them are given as follows.

Finding (c): The virtualized server could consume less energy than the physical server does. Specifically, when 5 instances are executed (the instance number is one more than the number of CPU-cores), the energy overhead was negative for Xen-based servers, as the valley point.

Such an observation could be understood as the inter-play between the concurrent processes and the CPU cores in a multi-core server. For the physical server, we observe that 4 instances were finished first and the last instance was completed much later. This was further verified by the observation that the CPU usage maintained nearly 100% until the first 4 instances were completed, and afterwards it dropped to around 25%. In comparison, in the virtualized servers, all the instances were completed almost at the same time. The CPU usage on the Xen-based server maintained at a high level of 99.8%, compared to 87.7% for the physical server. In this case, the Xen-based server, either running 2 or 3 VMs, took around 10% less time and consumed 11% less energy than that of the physical server. For the KVM-based server, the advantage of the CPU scheduler was reversed by the extra penalty of hypervisor in most cases, except for the case when 2 active VMs were configured, resulting in a saving of 2% energy than the physical server. This finding suggested that, if there was no binding between running processes and CPU-cores, native operation system cannot truly take advantage of multi-core architecture; in contrast, virtualized systems, based on either Xen or KVM, was able to partition computing resources into smaller pieces to achieve better resource allocation across active VMs to save energy.

Finding (d): The KVM-based server consumed more energy than that of the Xen-base server. For example, when processing 7 parallel tasks, 2 KVM VMs consumes 5.4% energy more than that based on 2 Xen VMs, and the gap reaches 23% between 3 KVM VMs and 3 Xen VMs. It was because the KVM hypervisor consumed more CPU cycles and occupied higher memory footprint, compared to the Xen hypervisor. The additional requirement translates into higher energy consumption.

Finding (e): The number of active VMs affected the energy usage for the KVM-based server. Particularly, when configuring 3 active VMs, the KVM-based server consumed more energy than that consumed by 2 active VMs configured on the same server. This could be attributed to the frequent Lock Holder Preemption (LHP) mechanism, investigated by [14]. A guest VCPU in the KVM-based server may be preempted when the host de-schedules the VCPU threads. If the preempted VCPU was running in a critical section, the lock will be held a certain time from the perspective of the guest VM. The probability of LHP was higher with more active VMs. Once LHP occurs, CPU resources were simply wasted in the lock holding period, which in turn increased the task completion time. So the average power consumption for the KVM-based server with 3 active VMs was the lowest, but the task completion time was the longest.

C. HTTP Request Benchmark

Results as recorded from the server application setup and energy consumption based on the architecture.

Finding (f): The virtualization overhead for network-intensive traffic was much larger than that for computing-intensive traffic. For the Xen-based server, the energy overhead for computing-intensive traffic was less than 5%, while the overhead for network-intensive traffic could rise up to 70%.

The same situation happened to the KVM-based server. The cause of this finding is at least two-fold. First, for networking traffic, the CPU usage of virtualized server was much higher than that of the native server; while for local computing tasks, the CPU usages of all the servers were almost equal. That difference suggests that dramatic CPU cycles were budgeted for VFR/VIF in Xen or TUN/TAP in KVM. Second, according to (Uhlig, et.al 2004) the probability of Lock Hold Preemption (LHP) for I/O- intensive workloads was 39% for virtualized server. The high frequency of LHP translated into high energy cost.

Finding (g): The energy overhead for the virtualized server was correlated with the number of active VMs. For 3 active KVM VMs, the energy overhead was around 1.5 times higher than that for 2 active VMs; similarly, 3 active Xen VMs consumed almost twice overhead of that for 2 active VMs. Moreover, the gap for the KVM-based server is higher. For example, in the case of 15,000 req/s, the overhead gap between 3 active VMs and 2 active VMs for KVM was more than 80%; while it was around 20% for Xen.

Finding (h): The network throughput for the KVM-based server reached its maximum between 10,000 reqs/s and 15,000 reqs/s. When the request rate was 15,000 reqs/s, the KVM-based server took longer time to complete the task and thus consumed more energy, compared to the case of 10,000 reqs/s. As a comparison, the task completion time and the energy cost for the physical machine and the Xen-based server monotonically decreased as the request rate increased up to 15,000 req/s.

This observation was largely due to the extra memory footprint for KVM. In Apache server, each serving request took certain memory. The maximum number of requests that could be served simultaneously was thus proportional to the amount of available resources. For the case of KVM, the extra memory footprint shrinks the amount of available memory for request serving.

Finding (i): The marginal power consumed by the server under different load conditions was limited, compared to the power consumed when the server was idle. Specifically, the additional power consumed by the server under different level of networking requests was at most 37.3% against the idle state, and the maximum additional power consumption for the local computation benchmark was 57.6%. Moreover, the marginal power consumption was highly correlated with the CPU usage as observed. As a result, the experiment verified a previously power consumption model for the server, in which the power consumption of the server could be viewed almost as an affine function of CPU usage with the idle power consumption as the y-intercept [15].

Finding (j): The energy overhead for virtualized servers was expansive, where when plotted the lines are curved by one degree polynomial fitting based on the power consumption of different configurations, the power gap between the baseline and the virtualized servers for both Xen and KVM increases as the throughput increases, before the maximum throughput of the KVM-based server is reached. When no network traffic occurs, the gap between Xen and baseline is around 1% (0.8 Watt), and the gap between the KVM-based server and the baseline server is approximately 10% (6.9 Watt). When the throughput grows to 10000 reqs/s, the gap becomes 15.2% (10.8 Watt) for Xen and 11.2% (7.9 Watt) for KVM.

V. Discussion

In this section, various researches have been reviewed, compared and analyzed including the empirical findings described by the experiment setup.

Objective: To evaluate the impact of virtualization on environmental sustainability. Based on the objective under study the researcher deduced the following:

The server is still far from energy-proportional. The idle server even consumes approximately two thirds of the energy when its computing resource is fully occupied. As a result, it will be advantageous to consolidate applications from multiple servers to one and turn off those idle servers to save energy.

The virtualized server in general consumes more energy than the physical server does. The energy overhead for virtualized servers increases as the resource utilization increases. When the virtualized servers are idle, Xen incurs nearly 1% energy overhead, and KVM contributes around 10% extra energy cost. For networking benchmark, Xen’s virtual firewall router, virtual network interface, and virtual event mechanism add an energy overhead ranging from 2.8% to 70.2% for different workloads and VM configurations; and KVM’s Linux kernel based virtual network bridge and x86 VTx invoking results in an energy overhead between 59.6% and 273.1% for various combinations of configuration and workload.

Energy saving could be achieved by launching an optimum number of virtual machines. During measurement, the virtualized server with 2 active VMs consumes less energy than the one with 3 active VMs. Specifically, about 20% energy for KVM and 15% energy for Xen on average could be conserved for all cases under networking benchmark, by migrating tasks from one VM to another and turning off the idle VM.

When a multi-core server is running multi-process applications, the physical machine could consume more energy than virtualized servers. It is due to a lack of the multi-core optimization in the physical machine. While both Xen and KVM are able to distribute physical CPU cores into virtual CPU cores, avoiding starvation. This demonstrates an essential advantage of virtualization in improving resource utilization.

With regard to the above findings from the empirical experiment, the researcher would summarize that:

i. The energy usage in data centers can be reduced by virtualization for example by consolidating applications from multiple servers to one server and shutting down the rest. This is based on the observation of the power consumption model for native server.
ii. On the other hand, for the virtualized servers, there are two detrimental effects would hamper the energy efficiency. First, the hypervisor introduces a potential energy overhead over the physical machine, by occupying system resources for its execution. This overhead is expansive as a function of the 'goodput', which denotes the portion of computation capabilities used for support applications. Second, the maxi- mum supportable goodput for virtualized server is reduced, compared to its native server. The combination of these two detrimental effects would offset the energy benefit of server consolidation. Moreover, the impact of these detrimental effects also depends on the type of hypervisor chosen.
iii. There are many other upcoming technologies based on virtualization (application development, migration, metrics and algorithms, CPU design and many other designs) are key to reducing energy use and green computing.
iv. This fundamental trade-off dictates how server consolidation should be designed for green data centers. Specifically, the decision of server consolidation should balance those two competing forces, to reduce the energy usage by data centers.

VI. Conclusion

In this paper the researcher reported an empirical study on the impact of virtualization and other published papers on energy efficiency and towards promoting green computing in attaining ICT development and second phase of the Millennium Development goals.

Objective: To evaluate the impact of virtualization on environmental sustainability. Based on the objective under study the researcher deduced the following: That virtualization is a key technology towards reducing energy usage and making data centers green. In the long run it conserves the environment by reducing the amount of gases remitted into the atmosphere because the gases cause global warming effect.

Therefore, to promote green computing and environmental sustainability, virtualization has emerged as one of the most important design requirements for modern computing systems, such as data centers, as they continue to consume enormous amounts of electrical power. Apart from high operating costs incurred by computing resources, this leads to significant emissions of CO_2 into the environment. For example, currently IT infrastructures contribute about 2% of total CO_2 footprints. Unless energy-efficient techniques and algorithms to manage computing resources are developed, ICT's contribution in the world's energy consumption and CO_2 emissions is expected to rapidly grow. This is obviously unacceptable in the age of climate change and global warming. Thus it is clear from the research study that virtualization plays a significant role in ensuring green computing and subsequently environmental conservation and sustainability as envisaged by the second theme in the second millennium goals.

VII. Recommendations And Way Forward: Case Of Kenya

Objective: To evaluate the impact of virtualization on environmental sustainability. Based on the objective under study the researcher made recommendations and way forward with regard to Kenya as both a technological and environmental hub:

In Kenya, there is a shortage of electricity or power and the government is improvising techniques to inject another 5000kwatts. However, KenGen and Kenya Power and Lighting Company are not taking any initiatives to advocate for better means of power usage to save some energy. Moreover, it is not doing enough to educate its citizens of the dangers of environmental changes and global warming which reduces rainfall and eventually increases food security. Our country relies heavily on rain fed agriculture, yet the environmentalists proclaim that today rainfall unlike yester years is so unpredictable. This unpredictability has very adverse effects to the citizens more so those living in Northern, parts of Eastern and other arid and semi-arid areas of Kenya. Similarly, in the technological world, the government should start creating awareness such that as data centers are being developed, they consider designing them using virtualization. This would go a long way in promoting green data centers, which would reduce energy intake and reduce emission of global warming effect gases to the environment. By so doing this would conserve the environment and put it to sustainable levels with other countries. Based on the researcher and researchers cited in this paper the following would be recommended for further research:

- For designing the holistic solutions in the scheduling and resource provisioning of applications within the datacenter, all the factors such as cooling, network, memory, and CPU should be considered. For instance, consolidation of VMs even though effective technique to minimize overall power usage of datacenter, also raises the issue related to necessary redundancy and placement geo-diversity required to be maintained to fulfill SLAs with users.

- Last but not the least, the responsibility also goes to both providers and customers to make sure that emerging technologies do not bring irreversible changes which can bring threat to the health of human society. Deploy the datacenters near renewable energy sources and maximize the Green energy usage in their already established datacenters. Before adding new technologies such as virtualization, proper analysis of overhead should be done real benefit in terms of energy efficiency.
- This paper also proposes the development of a Green IT framework that would use virtualization technology as basics for implementing energy efficient data centers. The framework should provide an imminent solution to the data center managers to improve the performance of their existing data center by implementing the proposed Green IT framework. It would also help them to reduce the emission of GHG so that global warming effects can be eliminated or reduced. It should also provide a set of green metrics to be implemented in data centers to measure the performance in terms on energy efficiency and CO_2 emissions. The proposed Green IT framework would help IT business enterprises specifically data center industry to follow up and implement a virtualized Green IT framework, to save huge amount of energy and at the same time reduce the CO_2 emissions, which ultimately reduces global warming effects.

.

REFERENCES

[1] Molla, A. (2009). The Reach and Richness of Green IT: A Principal Component Analysis, pp. 754 - 764.

[2] Silicon Valley Leadership Group, Accenture, Data Centre Energy Forecast Report-Final Report. July, 2008.

[3] Aronson, J.; Blignaut, J. N.; Milton, S. J.; le Maitre, D.; Esler, K.; Limouzin, A.; Fontaine, C.; de Wit, M. W.; Mugido, W.; Prinsloo, P.; van der Elst, L. and Lederer, N. (in press) Are socio-economic benefits of restoration adequately quantified? A meta analysis of recent papers (2000-2008) in Restoration Ecology and 12 other scientific journals. Restoration Ecology 18(2).

[4] Ruth, S. (2009). Green it - more than a three percent solution? Internet Computing, IEEE, 13(4), 74-78.

[5] Mayank Mishra, Anwesha Das, Purushottam Kulkarni, and Anirudha Sahoo, Dynamic Resource Management Using Virtual Machine Migrations, IEEE Communications Magazine, 0163- 6804/12, September 2012.

[6] Xiaomin Zhu, Laurence T.Yang, Huangke Chen, Ji Wang, Shu Yin and Xiaocheng Liu, Real- Time Tasks Oriented Energy-Aware Scheduling In Virtualized Clouds, IEEE Transactions On Cloud Computing, Vol. 2/April-June 2014.

[7] Zhen Xiao, Weijia Song, and Qi Chen, Dynamic Resource Allocation Using Virtual Machines for Cloud Computing Environment, IEEE Transactions on Parallel and Distributed Systems, vol. 24, no. 6, June 2013.

[8] Elizabeth Sylvester Mkoba, Mokhtar Abdullah Abdo Saif, A Survey On Energy Efficient with Task Consolidation in The Virtualized Cloud Computing Environment, IJRET: International Journal of Research in Engineering and Technology.

[9] Murugesan, "Going Green with IT: Your Responsibility Toward Environmental Sustainability," Cutter Business - IT Strategies Executive Report, vol. 10, no. 8, 2007.

[10] Rivoire S, Shah MA, Ranganathan P (2007). A balanced energy efficiency benchmark." Proceedings of the ACM SIGMOD International Conference on Management of Data.

[11] Feng, S., Huang, J., Wang, J. (2011). Loss of the Polycomb group gene polyhomeotic induces non-autonomous cell overproliferation.

[12] Mueen U, Azizah AR (2010). Server consolidation: An Approach to make Data Centers Energy Efficient and Green. Int. J. Sci. Eng. Res. (IJSER), 1(1): 1-7.

[13] Murugesan, San, "Harnessing Green IT: Principles and Practices," IT Professional, vol.10, no.1, pp.24,33, Jan.-Feb. 2008.

[14] Pushtikant Malviya, Shailendra Singh "A Study about Green Computing", International Journal of Advanced Research in Computer Science and Software Engineering, Volume 3, Issue 6, June 2013.

[15] X. Fan, et al., Power Provisioning for a Warehouse-sized Computer. In Proc. of the ACM Int'l Symposium on Computer Architecture, 2007.

[16] Ching-Hsien Hsu, KennD.Slagter, Shih-Chang Chen, Yeh –Chinh Chung, Optimizing Energy Consumption with Task Consolidation in Cloud‖, Information Sciences, No. 3, March 2014.

BIBLIOGRAPHY

R. Buyya, R. Ranjan, and R. N. Calheiros, "Modeling and simulation of scalable cloud computing environments and the CloudSim toolkit: Challenges and opportunities," in Proceedings of the 7th High Performance Computing and Simulation Conference (HPCS'09). IEEE Press, NY, USA, 2009

Carlo Mastroianni, MichelaMeo, and Giuseppe Papuzzo,‖ Probabilistic Consolidation Of Virtual Machines In Self-Organizing Cloud Data Centers‖, IEEE Transactions On Cloud Computing, Vol. 1, No. 2, July-December 2013.

J.Che, et.al., Asynthetical performance evaluation of Open VZ, Xen and KVM. IEEE Services Computing Conf., China, Dec, 2010.

IBM Corp., IBM System Virtualization. http://publib.boulder.ibm.com-/infocenter/eserver/v1r2/topic/eicay/eicay.pdf, Dec, 2005.

P. Barham, et al., Xen and the Art of Virtualization. In Proc. of the 19th ACM symposium on operating systems principles, Dec, 2003.

A. Kivity, et al., kvm: The Linux Virtual Machine Monitor. In Proc. of the Linux Symposium, Ottawa, 2007.

X. Fan, et al., Power Provisioning for a Warehouse-sized Computer. In Proc. of the ACM Int'l Symposium on Computer Architecture, 2007.

L. Cherkasova et al., Comparison of the three CPU schedulers in Xen. Sigmetrics Performance Evaluation Review 35(2): pp. 42-51, 2007.

Xen Credit-Based CPU Scheduler, http://wiki.xensource.com/xenwiki/C-reditScheduler

W. Jiang, et al., CFS Optimizations to KVM Threads on Multi-Core Environment. 15th Int'l Conf. on Parallel and Distributed Systems, 2009.

A. Menon, et al., Diagnosing Performance Overheads in the Xen Virtual Machine Environment. 1st ACM/USENIX Conf. on Virtual Execution Environments (VEE'05), Jun, 2005.

P. Bohrer, et al., The case for power management in web servers. Power Aware Computing, 2002.

W. Jiang, et al., CFS Optimizations to KVM Threads on Multi-Core Environment. 15th Int'l conf. on Parallel and Distributed Systems, 2009.

V. Uhlig, et al., Towards Scalable Multiprocessor Virtual Machines. In Proc. of Virtual Machine Research and Technology Symposium, 2004.

J.Che, et.al., Asynthetical performance evaluation of OpenVZ, Xen and KVM. IEEE Services Computing Conf., China, Dec, 2010.

V. Uhlig, et al., Towards Scalable Multiprocessor Virtual Machines. In Proc. of Virtual Machine Research and Technology Symposium, 2004.

Aruna Malik and Harsh K.Verma, " Performance Analysis of Virtual Private Network for Securing Voice and Video Traffic", International Journal of Computer Application, Volume 46-N0.16,May 2012.

Timothy Wood, "Improving data center resource Management, deployment, and availability with virtualization", PHD thesis June, 2009, (Unpublished)

Jalal Frihati, Florica Moldoveanu, Alin Moldoveanu, General guidelines for the security of a large scale data centre design, U.P.B. Sci. Bull., Series C, Vol. 71, Issue 3, 2009.

Jakub Szefer and Ruby B. Lee. Architectural Support for Hypervisor-Secure Virtualization. In Proceedings of International Conference on Architectural Sup- port for Programming Languages and Operating Systems, ASPLOS, pages 437– 450, March 2012.

Ritu Malik Rupali Syal "Performance Analysis of IP Security VPN "International Journal of Computer Applications (0975 – 8887) Volume 8– No.4, October 2010

Olalekan Adeyinka Analysis of problems associated with IPSec VPN Technology 2008 978-1-4244-1643-1/08

Xiaomin Zhu, Laurence T. Yang, Huangke Chen, Ji Wang, Shu Yin and Xiaocheng Liu, Real- Time Tasks Oriented Energy-Aware Scheduling in Virtualized Clouds, IEEE Transactions On Cloud Computing, Vol. 2/April-June 2014.

Zhen Xiao, Weijia Song, and Qi Chen, Dynamic Resource Allocation Using Virtual Machines for Cloud Computing Environment, IEEE Transactions on Parallel and Distributed Systems, vol. 24, no. 6, June 2013.

Srikantaiah S, Kansal A, Zhao F (2008), Energy aware consolidation for cloud computing in Conference on power aware computer and systems (HotPower '08)

Mehiar Dabbagh, Bechir Hamdaoui, Mohsen Guizani, and Ammar Rayes, Toward Energy- Efficient Cloud Computing Prediction, Consolidation, And Over commitment, IEEE Network, March/April 2015.

Rosenblum M. and T. Garfinkel, "Virtual Machine Monitors: Current Technology and Future Trends," Computer, May 2010, pp. 39-47.

Chonglin Gu, Hejiao Huang, and Xiaohua Jia, Power Metering for Virtual Machine in Cloud Computing - Challenges and Opportunities, IEEE Access, Vol. 2, Sep 2014.

Dong Jiankang, Wang Hongbo, Li Yang yang, Cheng Shiduan, Virtual Machine Scheduling for Improving Energy Efficiency in Iaas Cloud, China Communications, March 2014.

UttamMandal, M.FarhanHabib, Shuqiang Zhang and Biswanath Mukherjee, Davis Massimo Tornatore, Davis and Politecnico Di Milano Greening the Cloud Using Renewable-Energy-Aware Service Migration, IEEE Network, November/December 2013.

Fahimeh Farahnakian, Adnan Ashraf, Tapio Pahikkala, Pasi Liljeberg, Juha Plosila, Ivan Porres, And Hannu Tenhunen, Using Ant Colony System to Consolidate Vms for Green Cloud Computing, IEEE Transactions On Services Computing, Vol. 8, No. 2, March/April 2015.

Li Hongyou, Wang Jiangyong, Peng Jia, Wang Junfeng, Liu Tang, ‖Energy-Aware Scheduling Scheme Using Workload-Aware Consolidation Technique in Cloud Data Centres, China Communications, December 2014.

Molla, A., & Cooper, V., 2009: Green IT Readiness: A Framework and Preliminary Proof of Concept. Australasian Journal of Information Systems, 16(2), 5-23

Malone C, Belady C (2007). Metrics and an infrastructure model to evaluate data center efficiency. In Proceedings of the Pacific Rim/ASME International Electronic Packaging Technical Conference and Exhibition (IPACK).

Authors' Profiles

Vincent Motochi is a PhD student in Information Technology at Kibabii University, which is in Bungoma County (Kenya). He holds a Master of Science in Information Technology and Bachelor of Science in Computer Science both from Masinde Muliro University of Science and Technology (Kenya). Currently he is the Head of ICT for the County Government of Kakamega. Before that he was the ICT Manager of the County Government of Vihiga on Interim terms for one year. Before that he was the Computer Programmer for the former County Council of Kakamega. He has special interests in information security and has published papers on the same topics and he is also a member of the Association of Computing Machinery (ACM).

Samuel Barasa is a PhD student in Information Technology at Kibabii University, (Kenya). He holds a Master of Science in Data Communications from KCA University (Kenya) and Bachelor of Science in Computer Science from Masinde Muliro University of Science and Technology (Kenya). Currently he is a Tutorial Fellow in the Department of Computer Science, Kibabii University, Kenya. Initially served as an ICT Maintenance Engineer at Masinde Muliro University of Science and Technology (Kenya). He also possesses CCNA, CCNA Instructor, and A+ Certifications. He is a member of the Association of Computing Machinery (ACM).

Patrick Owoche is a Tutorial Fellow at Kibabii University Bungoma where he is pursuing a PhD in IT. He is an MSc in IT from Masinde Muliro University of Science and Technology (MMUST), a Certified Information and Communication Technologist (CICT). He also holds BEd Science (Maths & Physics) from Egerton University, Diploma in Office System Technology from KTI and CCNA II.

Dr. Franklin Wabwoba is a senior lecturer in Information Technology and Dean of the School of Computing and Informatics at Kibabii University (Kenya). He holds a PhD (Information Technology) from Masinde Muliro University of Science and Technology, Master of Science (Computer Applications) from Kenyatta University; endorsement (Education Management), from University of South Africa and Bachelor of Education (Science: Mathematics and Computer Science) from Egerton University. He has taught Computer Science and Information Technology courses for many years. He has presented several papers in scientific conferences and has many publications in referred journals as well as university level computing books. He has a strong research interest in green ICT, the impact and integration of ICT into education. He is a professional member of the Association for Computing Machinery (ACM).