International Journal of Innovation and Applied Studies
ISSN 2028-9324 Vol. 3 No. 2 June 2013, pp. 528-534

http://www.issr-journals.org/ijias/

Hardware Virtualization towards a Proficient Computing Environment

Shweta Agrawal

Northern Indian Engineering College,
GGSIPU, New Delhi, India

ABSTRACT: In the recent few years Server Virtualization and Green Information Technology have become very popular and are fast becoming the norm in organizations of all disciplines and sizes. Today, different methods of energy savings are in use and in great demand. One of the newest methods in the IT to control the pollution of the environment and the greenhouse effect is Green IT that is directly connected with the Virtualization of Hardware Resources.
Virtualization is the presentation of an environment to one layer in an information technology stack that abstracts or represents a lower layer. It makes it possible for the IT professional to run a number of machines on a single physical machine.
In this study we elicit the concept of Hardware Virtualization. We illustrate the procedure of Hardware Virtualization using a real-world example and then we emulate a virtualized infrastructure to contrast against the physical infrastructure on the basis of CPU utilization. We have used the VMware Workstation 7.1.0 as a software tool for virtualization and AVG PC Tune Up 2011 to present the difference in CPU utilization before and after virtualization.
This paper helps to identify the main reasons for the growing need for data centre virtualization. The results in this paper show that a virtualized infrastructure can potentially increase the CPU utilization by a significant margin, thereby suggesting an efficient and faster way of resource utilization, saving processing time, reducing the cost incurred in building separate physical servers and furthermore reducing the power consumption.

KEYWORDS: Server virtualization, Green information technology, Carbon footprint, CPU utilization, Greenhouse effect, separation, controlled sharing.

1 INTRODUCTION

1.1 HISTORY

Virtualization has been, for a long period of time, in the spotlight of the contemporary information technology. Virtualization software makes it possible for the IT professional to run a number of virtual machines on a single physical machine. VMware, a mainstream virtualization software vendor, gives an overview of virtualization [1] and the history that led up to the popularity of this technology. In the 1960s IBM implemented virtualization as a way to partition mainframe computers. However, in the 1980s and 1990s the popularity of virtualization diminished as more affordable hardware became available. Hardware continued to become more affordable and at the same time more powerful, and the utilization levels fell to less than 10%. The increased focus on better hardware utilization and the drive to reduce the data centre's carbon footprint has revitalized the need for virtualization technology. Various software packages are available on the market and VMware is one of the more popular packages available, with over 170,000 customers at the time of writing.

1.2 ADVANTAGES OF A VIRTUALIZED INFRASTRUCTURE

- Virtualization reduces the amount of hardware required by a data centre. **Less hardware** can be maintained on **less power.** As a result the harmful impact on the environment is reduced thereby lending a hand to the **Green IT.**

- Application of the virtualization technologies vary for use in **simulations** and **testing** of new applications as well as in testing its coexistence.
- A virtual infrastructure is suitable for a few basic but strong security primitives [2] namely, separation and controlled sharing. Hypervisors (term discussed ahead) have traditionally supported strong isolation or separation of Virtual Machines (VMs) and their workloads, including fault isolation limiting an application or operating system fault's effects within a VM.

  A system can achieve separation in several ways:

- Using different hardware facilities for different workloads (physical separation).
- Running workloads at different times (temporal separation).
- Cryptographically protecting workload-specific data (cryptographic separation), and
- Using a reference monitor [3] or security kernel to separate workloads and their resources (logical separation or isolation).
- Moreover, a virtualized infrastructure also brings about better **utilization of processor** and **resources** such as hardware.
- It also reduces the **floor space** requirements of an organization.
- Therefore virtualization proves to be a viable option for any organization which wants to **save funds.**

1.3 THE CONCEPT

Virtualization software decouples the software from the hardware by creating an abstract layer, known as a ***hypervisor***, between the virtual machine and the host operating system. The hypervisor acts as a controller between the hardware and the virtual machines. Its function is to monitor the virtual machines, dynamically assigning the hardware to the virtual machines as and when the hardware is required.

Hypervisors, by their functionalities, can be divided into two groups:

- *Type one hypervisor or native* which is run directly on the server hardware and is called bare – metal hypervisor.
- *Type two hypervisor*, *or hosted*, which is in fact an application run in operating system installation.

To accomplish the best efficiency hypervisors of type one are used the most because the operating system is not needed. Type two hypervisors are applicable to simulation and testing environments because it can be run both on client and server versions of the operating systems and in the environments with insufficient hardware resources.

Here, it is also important to understand what a *Host Machine* is? In context of virtualization, Host machine refers to the computer which has one or more virtual machine installed on it, running as the host for the virtual machines [4]. **Fig.1** shows the infrastructure of a Virtualized Host Machine.

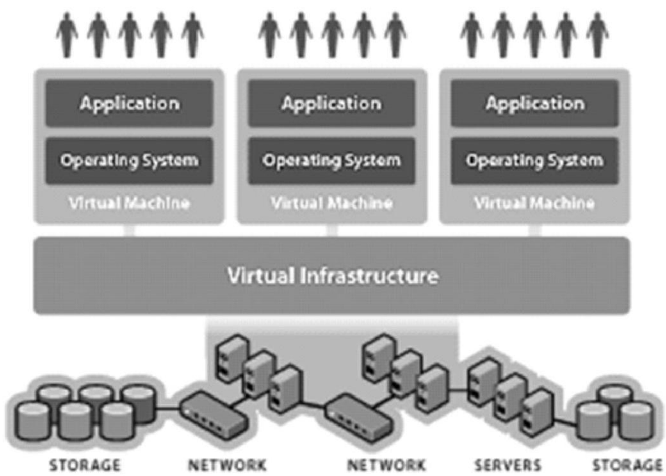

Fig. 1. Virtualized Host Machine

Virtualization software is based on one of three fundamental technologies, namely **full virtualization, paravirtualization** and **hardware assisted virtualization**. VMware [5] explains the key differences between the three virtualization technologies as follows:

Full virtualization is believed to be the most popular form of virtualization. With this ***technology*** all operating systems can be installed on a virtual machine without the need for modification. The hypervisor manages CPU instructions and controls the virtual machine's access to the hardware. The operating system is not aware that it is running in a Virtual environment and that its access is being controlled by the hypervisor.

OS assisted virtualization or paravirtualization is used to overcome the high processor overhead that is created in the full virtualization technique. In this technique the guest operating system is modified to make it aware that it is running on a virtual machine. The virtual machine is then aware of the hypervisor and collaborates with the hypervisor to reduce the overhead on the processor.

Hardware assisted virtualization can be achieved by using hardware that have been developed with features that assist virtualization techniques. This technique requires he hardware vendor to incorporate an inflexible programming model in the product and as a result this virtualization technique is not commonly used.

The rest of the paper is organized as follows: Section II presents related works and case studies where virtualization has proved to reduce the demand for hardware and power consumption and the consequent reduction of carbon emissions. Section III elaborates our virtualized infrastructure setup. In section IV we describe the procedure. In section V we present the data analysis. We conclude in section VI and discuss scope for future work in section VII. In Section VIII we list the references.

2 RELATED WORKS AND MOTIVATION

Extensive investigations to establish the energy requirements of IT data centres has shown that the hardware in data centres are generally underutilized, causing electricity to be wasted. A report compiled by the Environmental Protection Agency [6] found that hardware is not only underutilized, but the energy requirements of data centres in the United States have doubled in the period from 2000 to 2006.

Numerous case studies have proven that server virtualization has the potential to reduce energy consumption, subsequently saving the company money. In 2008 the Royal Borough of Windsor and Maidenhead virtualized 184 of their physical servers . This resulted in a 44% reduction on the council's energy bill; a saving of £1.2 million. The pension company, Standard Life, virtualized 65% of its physical servers. This freed up nearly two thirds of the occupied floor space in the server room. The virtualization also resulted in a saving of £300K per annum on the company's electricity bill. The UK retailer, Tesco, implemented server virtualization technology to reduce 1,500 physical servers to only 120 servers. The IT Director, Nick Folkes explained that the virtualized infrastructure requires less power and is simpler to manage. Frangiskatos et al. states that server virtualization could result in the need for additional security measures to be implemented, however, reiterates the fact that virtualization simplifies the management of data centres.

In August 2005 the energy company E.ON [8] migrated their existing physical IT infrastructure to a virtualized infrastructure. The company required a flexible environment with the ability to accommodate all the individual IT requirements of the separate business units. A virtual environment met these requirements. Following the move to virtualization, E.ON reported a 56% reduction in power consumption. The heat in the data centre was also reduced by more than half [9].

The related works establish a strong argument to utilize the full potential of the virtualization technology in order to reduce the ever increasing demand for computer hardware and electrical energy supply, increase the CPU utilization and provide every advantage that virtualization brings with itself.

3 VIRTUALISED INFRASTRUCTURE SETUP & CPU UTILIZATION

In our setup for estimating the CPU utilization we have used Full Virtualization, as discussed in section I.

In this setup we have used the following hardware and software:

- Lenovo R500 laptop (Processor: 2.26GHz, RAM: 2 GB, OS: Windows 7Ultimate).
- Samsung PC (Processor: 2.20GHz, RAM: 1 GB, OS: Windows 7 Ultimate).
- CPU usage indicator software (AVG PC Tune Up 2011)

Virtualizations software: VMware Workstation Version 7.1.0 was used to create virtual machines.

Virtual machines created for the analysis: Red Hat Enterprise Linux 5 and Windows 7 Ultimate.

Fig. 2. The complete setup

For our purpose we define **CPU Utilization** as the time spent by the computer not in idle state. We assume that:

CPU Utilization = CPU Usage

Therefore, we can use CPU Usage as a measure of CPU Utilization.

4 PROCEDURE

Non Virtualized V/S Virtualized

For recording the results from a **Non-Virtualized Infrastructure**, we installed AVG PC Tune Up 2011 on one of the machines. Then we connected both the machines (the machines specifications have been mentioned in the previous section) on LAN and recorded the CPU usage readings by running the usage indicator software (AVG PC Tune Up 2011) in **Table 1** in section V.

For recording the results from a **Virtualized Infrastructure** we installed VMware workstation on the same machine on which we installed AVG PC Tune Up 2011. We refer to this machine as the Host machine. Then we created two Virtual Machines (VM) by using VMware Workstation, having Linux Operating systems on one VM and Windows 7 Ultimate on the other. First we ran Linux - VM1 (refer **Fig.3**) and recorded the CPU Usage readings from AVG Tune Up software (refer **Fig.5**). Then we shut down Linux (VM1) and ran Windows 7 Ultimate - VM2 (refer Fig.3) and recorded the readings for it in Table 1. Finally we took the average readings for all the three cases mentioned above and recorded them.

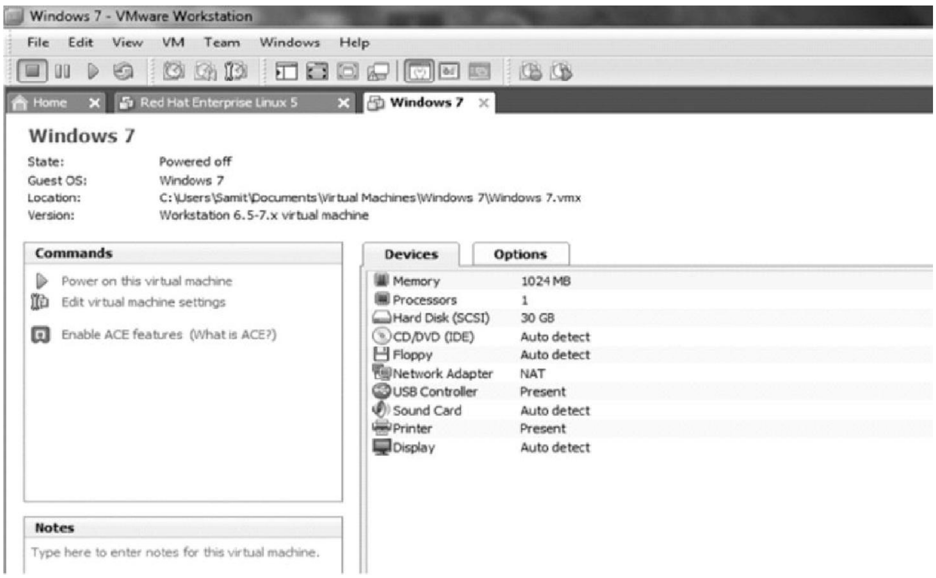

Fig. 3. Virtual Machines on VMware

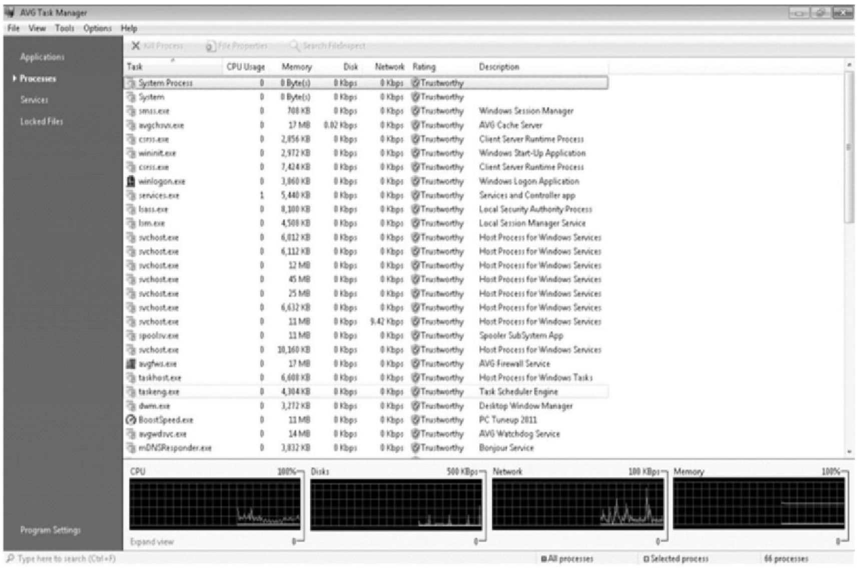

Fig. 4. Applications running on the host machine

Fig. 5. CPU Usage

5 DATA ANALYSIS

Table 1, shown below gives the CPU Usage readings obtained from AVG PC Tune Up 2011 for all the three cases:

- No virtualization
- Virtualization-Linux Virtual Machine
- Virtualization-Windows 7 Virtual Machine

Table 1. CPU Usage readings obtained from AVG PC Tune Up 2011

Readings	No virtualization (%)	Linux VM (%)	Windows7 VM (%)
1.	2	9	24
2.	4	20	16
3.	3	12	9
4.	5	10	12
5.	51	13	24
6.	2	18	16
7.	5	11	17
8.	2	12	13
9.	3	8	21
10.	2	6	16
11.	3	12	17
12.	5	19	19
13.	2	27	6
14.	3	18	25
15.	2	11	14
16.	4	12	11
17.	5	8	14
18	2	10	26
19.	48	11	18
20.	2	14	27
Average	**7.75**	**12.55**	**16.5**

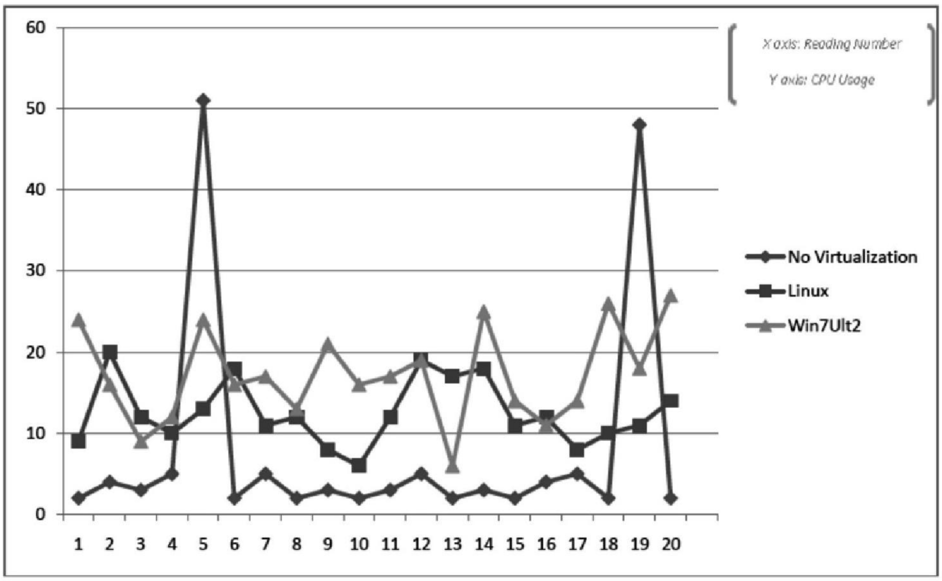

Fig. 6. Graphical representation of the recorded readings

Increment in CPU Usage:

- **Linux VM**

IncLinux = Usage (VM1) - Usage (No Virtualization)

= 12.55% - 7.75% = 4.8%

- **Windows 7 Ultimate VM**

IncWindows = Usage (VM2) - Usage (No Virtualization)

= 16.5% - 7.75% = 8.75%

From the above analysis we observed that the CPU usage increases by 4.8% in case of **Linux Virtual Machine** and **8.75%** in case of **Windows 7 Virtual Machine.**

As we assumed:

CPU Utilization = CPU Usage

We can also say that CPU utilization increases by **4.8%** in case of **Linux Virtual Machine** and **8.75%** in case of **Windows 7 Virtual Machine**.

6 CONCLUSIONS

This case based study compared physical and virtualized infrastructure on the basis of CPU utilization. Two virtual machines were created for this purpose using **VMware Workstation 7.1.0**. The readings were taken using **AVG PC TUNE UP 2011**. We observed an increment of **4.8%** in case of **Linux Virtual Machine** and an increment **8.75%** in case of **Windows 7 Virtual Machine.** in the CPU Utilization. Therefore we can say that virtualization is a better alternative to running separate physical machines. If we apply Virtualization technology in data centres the amount of hardware required would reduce thereby decreasing the floor space requirements significantly. Furthermore, it would also save power, reducing the carbon footprints. In other words Virtualization saves the company's funds and also paves way for a greener tomorrow.

7 FUTURE WORK

There is a lot of scope for future work. Firstly we have performed our experiment on a small scale. The same method applied on a large scale would be significantly beneficial. Secondly, a difference was observed in the CPU utilization of Linux and Windows operating systems. Thorough evaluation is needed to understand the reason behind this. Third, the power savings from a virtualized environment need further exploration. It is believed that a company with more than five physical servers should invest in virtualization technology. Also further the carbon emissions could be calculated by using the nameplate ratings provided by the manufacturer of the machine so that we get an idea about the harmful impact on the environment. For future work we can also consider using paravirtualization or hardware assisted virtualization and compare results with that obtained from full virtualization. Another avenue that could be further investigated is the sustainability of the virtual environment.

REFERENCES

[1] VMware, "Virtualisation Basics," VMware Inc, Palo. Alto, California, Technical Report 2010.

[2] Ronald Perez, Leendert van Doorn, and Reiner Sailer, “Virtualization and Hardware-Based Security,” *IEEE SECURITY & Privacy*, Volume 6, Number 5, pp. 24-31, September / October, 2008.

[3] P. Anderson et al., *Computer Security Technology Planning Study*, tech. report ESD-TR-73-51, vols. I & II .

[4] Wikipedia, “Hypervisor”, 2013. [Online] Available: http://en.wikipedia.org/wiki/Host_machine (22 May 2013)

[5] VMware, “Understanding full virtualization, paravirtualisation and hardware assist,” White Paper WP-028-PRD-01-01, 2007.

[6] Energy Star, “Report to Congress on Server and Data Centre Energy Efficiency, Public Law 109 - 732,” Energy Star Program, 2007.

[7] Dimitrios Frangiskatos, Mona Ghassemian, and Diane Gan, "Technology perspective: Is Green IT a threat to IT security?" School of Computing and Mathematic Sciences, University of Greenwich, London, Whitepaper, 2009.

[8] Gemma Simpson, "Virtualisation powers up energy group," Silicon, Case Study, 2007.

[9] Pretorius, M., Ghassemian, M. and Ierotheou, C. (2010) An investigation into energy efficiency of data centre virtualisation. In: Proceedings International Conference on P2P, Parallel, Grid, Cloud and Internet Computing 3PGCIC 2010. IEEE Conference Publications. The Institute of Electrical and Electronics Engineers, Inc., CA and NJ, USA. Tokyo, Japan, pp. 157-163. ISBN 978-0-7695-4237-9.