Declaration

This management Research Project is my original work and has not been submitted for another degree qualification of this or any other university or Institution of higher learning.

Signed: ______________________________Date:__________________________

NKONGE ZACHEUS M.

D61/P/8572/2005

This Management Research Project has been submitted for examination with my approval as University Supervisor.

Signed: ______________________________Date:__________________________

KATE LITONDO

LECTURER, DEPARTMENT OF MANAGEMENT SCIENCE

SCHOOL OF BUSINESS

UNIVERSITY OF NAIROBI

Dedication

I dedicate this research to all my family members: Alice, Wakoo, Margeret, John, Rachel, Diana, Hellen and Kang'ara.

Special thanks to my wife, Lauryn Nkonge and mum Alice for their selfless sacrifice that enabled me come this far.

I wish to thank God almighty without whom nothing could be possible.

Acknowledgement

I would like to extend my appreciation to my supervisor; Kate Litondo, Moderator; Lelei and Management Science Department Chairman; Dr. Njihia for their relentless assistance, support, supervision and guidance throughout the entire research undertaking.

My special appreciation also goes to my wonderful wife, friends, colleagues and family for their love, support and encouragement. Thank you for your roles and I will always pray and hold you close to my heart in life.

Further, my cordial appreciation goes to Aaron Mbowa, an expert virtualization consultant at Dataposit Limited for his invaluable advice and sharing of experiences which instilled confidence as regards virtualization technology

I do also acknowledge Mr. Anthony Mukomah the Computer Information Systems Officer at UNON for allowing me to conduct this study in their organization to enable me meet partial requirements towards completion of my MBA. This also goes to all staff at UNON who assisted in making this study a reality.

My acknowledgement also goes to all ICT staff at CIC Insurance Ltd for their contribution during the period I was working on my project.

Abstract

Server virtualization is touted as the panacea to shrinking information technology budgets in today’s economy. The continued need to accomplish more with fewer assets has made the increased availability of server virtualization technologies an attractive option for many organizations. The explosive growth of server virtualization, the advantages associated with virtualization and the seldom advertised disadvantages associated with virtualization must be analyzed and distilled to determine if benefits associated with server virtualization supersede the associated costs.

This study evaluated the implementation virtualization technology at the UNON offices with the objective of establishing level of its usage, its benefits and challenges. Whereas the system users might not know they are operating on a virtual environment, the custodians in ICT department need to be well versed with the technology. The study concentrated on ICT and Finance department to collect analysis data and the findings showed that virtualization technology helped them reduce their physical server from over 20 to only four in the primary data center.

Further, based on findings, the study recommends that besides the many benefits associated with virtualization, organization should conduct its preliminary proof of concept before making decision of its implementation to reduce the effect of its challenges.

Table of Contents

List of Tables vii

Table of Figures viii

Abbreviations ix

CHAPTER 1: INTRODUCTION 1

1.0 Background of the study 1

1.1 How virtualization Works 2

1.1.1 Evolution of virtualization 3

1.1.2 Virtual Machine Monitor Architectures 5

1.1.3 Benefits of virtualization 7

1.2 Challenges of virtualization 8

1.3 Problem Statement 10

1.4 Research Objectives 11

1.5 Significance of the study 11

CHAPTER 2: LITERATURE REVIEW 13

2.0 Introduction to Literature Review 13

2.1 Server Based Computing (SBC) 13

2.2 Comparison between Virtualization and related technologies 14

2.2.1 Cloud computing vs. Virtualization 14

2.2.2 Distributed/Grid computing vs. Virtualization 15

2.3 Server Consolidation (virtualization) 15

2.3.1 High Availability Services 16

2.3.2 Virtualization advantages 17

2.3.3 Virtualization disadvantages 19

2.3.4 Virtual Machine Monitor Architectures 21

2.3.5 Full Virtualization of x86 Hardware 25

2.4 The Challenges of x86 Hardware Virtualization 27

2.5 Taxonomy of virtualization techniques 28

2.6 Virtualization of ICICI Bank in India 30

2.7 Managerial implications from the research 31

2.8 Business case for virtualization 32
CHAPTER 3: RESEARCH METHODOLOGY 34
3.0 Research Design 34
3.1 Population 34
3.2 Data Collection 35
3.3 Data analysis 35
CHAPTER 4: DATA ANALYSIS AND DISCUSSIONS 37
4.1 Introduction 37
4.2 Demographics 37
4.2.1 Gender 38
4.2.2 Age 39
4.2.3 Level of education 40
4.3 Virtualization usage at UNON 41
4.4 Factor Analysis 42
4.4.1 Benefits of virtualization 42
4.4.2 Challenges of virtualization 46
4.5 Descriptive Statistics 50
4.5.1 Advantages / benefits of Virtualization 50
4.5.2 Disadvantages / challenges of virtualization 52
CHAPTER 5: SUMMARY OF FINDINGS AND CONCLUSIONS, RECOMMENDATIONS AND SUGGESTIONS FOR FURTHER RESEARCH 54
5.0 Summary of findings and Conclusions 54
5.1 Limitations of the study 55
5.2 Recommendations 55
5.3 Suggestions for further research 55
REFERENCES 57
APPENDICES 63
GLOSSARY OF TERMS 71

List of Tables

Table 1: Taxonomy of virtual machines monitors architectures. 29

Table 2 : Gender Distribution 38

Table 3 : Education Level 40

Table 4 : Virtualization Benefits Communalities 43

Table 5 : Benefits Total Variance Explained 44

Table 6 : Virtualization Benefits Component Matrix (a) 45

Table 7 : Virtualization Challenges Component Matrix (a) 46

Table 8 : Virtualization challenges Pattern Matrix (a) 47

Table 9 : Virtualization Challenges Structure Matrix 48

Table 10 : Component Correlation Matrix 49

Table 11 : Benefits Descriptive Statistics 50

Table 12 : Challenges Descriptive Statistics 52

Table of Figures

Figure 1 : Virtual Machine Monitor - Virtual Machine Relationship. 26

Figure 2: Virtual Machine Monitor – Hardware simulation. 27

Figure 3: Privilege level architecture without virtualization. 27

Figure 4 : Gender representation 38

Figure 5 : Age Distribution. 39

Figure 6 : Age distribution 39

Figure 7 : Education Level 40

Abbreviations

n.d	Not defined
I/O	Input / Output
IT	Information Technology
ICT	Information and Communication Technology
CPU	Central Processing Unit
VMM	Virtual Machine Monitor
VM	Virtual Machine

CHAPTER 1: INTRODUCTION

1.0 Background of the study

Each year, organizations worldwide spend billions of dollars on their information technology (IT) infrastructure (Ramanathan & Bruening, 2004). Business's today are under ever increasing pressure to reduce operational costs while still ensuring that flexibility, service delivery levels, and business efficiency continue to improve. Organizations are investing considerable time, effort, and hard cash to achieve this (Conoops, 2007). Data center environments are becoming more complex and heterogeneous, with correspondingly higher management costs (VMware, 2006). As part of an effort to rein in operations cost increases, Golden and Scheffy (2007) suggest that a platform-virtualization model offers the opportunity to reduce overall system administration costs in comparison to non-virtual environment, by reducing the overall number of machines that need to be utilized. Gartner (2008) stated that virtualization was the top priority project for IT organizations of all sizes in year 2008 bringing companies real dollar savings. They predicted that virtualization projects will still be a top priority for 2009, despite the tough economic climate, largely because virtualization has been successfully reducing costs (the total cost of ownership, or TCO) and delivering a strong return on investment (ROI). As long as virtualization and related management, automation and optimization tools contribute to the bottom line, projects will continue to be approved

Virtualization can be defined as the creation of a virtual (rather than actual) version of something, such as an operating system, a server, a storage device or network resources. This research intends to study both storage and server virtualization case of the United Nations office at Nairobi which according to a senior virtualization consultant at Dataposit Limited - Kenya, is one of the first organizations to implement the virtualization technology. The UNON offers information technology services to all UN affiliated organization in the region with an estimated staff of 2,000. As befits the fourth largest UN headquarters, Gigiri is home to one of the most sophisticated office automation systems and digital telecommunications networks on the African continent. (UNON website)

1.1 How virtualization Works

Virtualization is the use of software such as VMware Infrastructure (software suite designed to centralize the management of virtualized IT environments) and VMware ESXi to transform or virtualize the hardware resources of an x86-based computer including the CPU, RAM, hard disk and network controller to create a fully functional virtual machine that can run its own operating system and applications just like a “real” computer. Each virtual machine contains a complete system, eliminating potential conflicts. Virtualization works by inserting a thin layer of software directly on the computer hardware or on a host operating system that contains a virtual machine monitor or “hypervisor” that allocates hardware resources dynamically and transparently. Multiple operating systems run concurrently on a single physical computer and share hardware resources with each other. By encapsulating an entire machine, including CPU, memory, operating system, and network devices, a virtual machine is completely compatible with all standard x86 operating systems, applications, and device drivers. You can safely run several operating systems and applications at the same time on a single computer, with each having access to the resources it needs when it needs them.

Techworld (2008) reported that many Australian businesses are using virtualization than the rest of the world, according to analysts. Gartner (2008) says between 8 and 9 % of medium enterprise businesses are virtualizing their x86 servers, while global uptake remains at about 4 to 5 %. Gartner’s servers and storage Vice President Phil Sargeant added that the figures, derived from research into organizational IT maturity, show Australians are spearheading virtualization uptake. According to Unisys (2008), the figure is even higher, since the 8 to 9 % adoption includes only virtualized systems in production.

This study examines platform-virtualization at United Nations Office at Nairobi (UNON), which encompasses the virtualization of the physical hardware and the operating systems that reside on that hardware. Ramanathan and Bruening (2004) state, “Platform virtualization can be defined as the creation of a logically partitioned computing system that runs on top of an actual platform. While virtualization has been applied to storage and servers, the concept of platform-virtualization goes further to include all layers of the platform from applications and operating software to platform components, processors and interconnects. Virtual platforms are perceived by users and function in all respects as if they were physical computers”. Yang (2008) defines platform-virtualization as a technique used to replicate the functionality of real hardware platform so that one physical machine can host more than one reducing cost and increasing efficiency system level software without conflict. Hamm (2005) states, “virtualization has the potential to dramatically change the way corporate computing is done. With servers, it lets one physical machine behave like a half-dozen virtual machines, and it also makes it possible to move an application from one machine to another on the fly”. The Strategic Counsel (2007) states “server virtualization is becoming a mainstream technology, having been adopted by at least 39% of worldwide organizations with more than 500 employees”.

1.1.1 Evolution of virtualization

Virtualization was first developed in the 1960s to partition large, mainframe hardware for better hardware utilization. Today, computers based on x86 architecture are faced with the same problems of rigidity and underutilization that mainframes faced in the 1960s. VMware invented virtualization for the x86 platform in the 1990s to address underutilization and other issues, overcoming many challenges in the process. Today, VMware is the global leader in x86 virtualization, with over 130,000 customers, including 100% of the Fortune 100. Virtualization was first implemented more than 30 years ago by IBM as a way to logically partition mainframe computers into separate virtual machines. These partitions allowed mainframes to “multitask”: run multiple applications and processes at the same time. Since mainframes were expensive resources at the time, they were designed for partitioning as a way to fully leverage the investment.

AIO (2008) say that virtualization was effectively abandoned during the 1980s and 1990s when client-server applications and inexpensive x86 servers and desktops led to distributed computing. The broad adoption of Windows and the emergence of Linux as server operating systems in the 1990s established x86 servers as the industry standard. The growth in x86 server and desktop deployments led to new IT infrastructure and operational challenges (VMware, n.d). These challenges acted as the motivating reasons towards moving to virtualized environments; The first challenge is non-virtualized environments have Low Infrastructure Utilization as typical x86 deployment achieve an average utilization of only 10% to 15% of potential capacity. Secondly, there has been an increase in the operational costs to support growing physical infrastructure due to the requirements uptime, power consumption and cooling. Thirdly, computing environments have become complex requiring specialized and experienced personnel leading to an increase in ICT management costs. Fourthly, non-virtualized environments exhibit Insufficient Failover and Disaster Protection exposing organization to possible downtimes for critical services. Finally, there is a challenge of maintaining numerous end-user desktops like enforcement of security policies.

To address the above challenges, Highleyman (2008) puts it that a fortunate trend is evolving that is making servers become more powerful. Moore’s Law states that server capacities will double every eighteen months and this trend not only has held for decades but is projected to hold well into the future. The result is that data-center servers are carrying less and less of their rated capacity. In fact, recent studies have shown that typical servers in a data-center environment that is governed by a one-application, one-server policy are running at only 10% to 15% of capacity. If only we could harness this excess capacity, we could significantly reduce the number of servers in a data center by a factor of two, three, or even more. This would result in less hardware, less maintenance, less administration, less space, and less energy – in short, less cost by a large factor. This is the promise of virtualization.

1.1.2 Virtual Machine Monitor Architectures

The Virtual Machine Monitor (VMM) is a software layer that manages one or more virtual machines. A virtual machine is a software abstraction of the real machine that is isolated from the VMM and is also referred to as a domain. Domains can contain any type of software application, but in the context of this paper the application will be an Operating System. The operating system running in a domain is referred to as the guest operating system. The applications running on top of the guest operating system are called the guest applications. The VMM itself can be run on top of hardware or software that are named host architectures. The hardware hosts the VMM while the guest OS and applications make up a domain. Some major VMM architectures being used today as stated by Ferstay, (2001) are:

Nano-kernel systems: The basic idea behind Nano-kernels is that they are a separate, small bit of core functionality that resides within the OS to do some specialized processing (e.g.Real-time interrupt processing and scheduling). Nano-kernels can be considered a lightweight version of a VMM because typically they do not virtualize very many resources. Examples of nano-kernel systems are Jaluna and RTLinux. For the most part, nano-kernel architectures have been dubbed cooperative virtualization environments. This is because the

VMM is not isolated from the domains that it hosts, and thus makes exchange of information between the VMM and any domain very efficient. However, it is the same cooperation that leaves nano-kernel based systems vulnerable to attack. Therefore a nano-kernel based monitor system may not be the best choice for a trusted VMM based security system.

Full Virtualization Systems: Full virtualization systems virtualize every resource on a computer system. Guest operating systems do not have access to physical devices, physical memory, or physical CPUs. Instead, the guest operating systems are presented with virtual devices, virtual memory, and virtual CPUs. In a full virtualization system, the virtualized interfaces presented to the guest operating system look and feel exactly like the interfaces of the real machine and therefore the guest OS and applications may run on the virtual hardware exactly as they would on the original hardware.

Paravirtualization Systems: The architecture of paravirtualization systems has very much the same look and feel as that of a full system virtualization system but with major differences in design decisions. Paravirtualization was born out of the observation that full system virtualization is too slow and complex on today's commodity hardware. For the most part, this is because the VMM must intervene whenever a domain attempts to execute privileged operation. Paravirtualization aims to retain the protection and isolation found in the full system virtualization approach but without the implementation complexity and associated performance penalty. The main idea behind paravirtualization is to make the VMM simpler and faster by relaxing the constraint that guest operating systems must run on the VMM without being modified.

Software Hosted Virtualization Systems: In this architecture, the VMM runs on top of a general purpose operating system. Software hosted VMMs leverage the services of the general purpose operating system (the host) to simplify the process of providing virtual hardware abstractions to guest operating systems.

The hosted VMM can use any virtualization technique to host guest operating systems (e.g. full virtualization or paravirtualization). UMLinux is an example of a software hosted VMM.

1.1.3 Benefits of virtualization

“Virtualization is the highest-impact issue changing infrastructure and operations through 2012.” “It will change how you manage, how and what you buy, how you deploy, how you plan and how you charge. It will also shake up licensing, pricing and component management. Infrastructure is on an inevitable shift from components that are physically integrated by vendors (for example, monolithic servers) or manually integrated by users to logically composed “fabrics” of computing, I/O and storage components.” (Gartner, 2008). Today’s IT intensive enterprise must always be on the lookout for the latest technologies that allow businesses to run with fewer resources while providing the infrastructure to meet today and future customer needs (Intel, 2008). Burger (2008) has listed the following as some of the advantages of an enterprise embracing virtualization technology:

First, Servers consolidation enables multiple operating systems to run in the same server, eliminating the need to dedicate a single machine to one application. Old and new applications can run simultaneously with their respective operating systems in multicore servers with many threads of execution, saving space and power consumption in the datacenter. New versions of an OS and new applications can be deployed without purchasing new hardware.

Secondly, virtualization introduces stability and security eliminating conflicts that can arise between supposedly stable applications, and troubleshooting can be daunting. As a result, cautious system administrators often host each type of application in a separate server even if the server is grossly underutilized. Multiple virtual machines running bread and butter applications are kept safely separated from each other. In addition, since each VM is isolated from the rest, a security breach in one does not affect the others. The fault tolerance and security brought about by the isolation of each virtual machine is a major benefit of virtualization.

Thirdly, there is development flexibility as a virtualized machine can host numerous versions of an operating system, allowing developers to test their programs in different OS environments on the same machine. In addition, with each application running in its own virtual partition, crashing in one virtual machine will not bring down the system.

Fourth, migration and cloning is very easy as virtual machines, each with their own OS and applications, function like self-contained packages that are said to be "decoupled from the hardware." It is relatively easy to move a VM from one server to another to balance the workload, to migrate to faster hardware, as well as to recover from hardware failure. VMs can also be quickly cloned and deployed.

Lastly, desktop virtualization has increase the trend to store a user's desktop (OS and applications) in a separate virtual machine in the server and use the PC or a dedicated terminal as a "thin client" to the server. Each user is isolated from all other users, due to the virtual machine technology, and the maintenance of the applications is shifted from each user's office to the datacenter.

1.2 Challenges of virtualization

“Virtualization is both an opportunity and a threat,” says Patrick Lin, senior director of product management for VMWare. Thomas Ptacek, a security researcher with Matasano Security, says the move to virtualization is the biggest thing happening in IT today. "And every application running in a modern IT organization is on a path to being moved to one big iron machine running multiple VMs,” he says. “And its impact on security touches everything.” Skooliki, (2008) argues that virtual environments will become pretty much standard practice in all data centers and cost advantages, efficiencies, and productivity gains are just so hard to debate. Scooliki (2008) noted the following as some of the disadvantages or issues that should be put into consideration when considering virtualization.

First, there is possible server sprawl since deploying servers is so easy in a virtual environment you may get carried away with the amount of servers you actually deploy. Secondly, some applications may still require dedicated physical resources. There are cases where really memory, CPU, or I/O hungry applications are best on their own physical asset. You have to remember that all VMs are sharing the same pool of resources.

Third, there is the licensing challenge as one needs to make sure the license model works well for what they are doing. E.g. Microsoft defines a server as a physical hardware system that is capable of running server software. Further, Microsoft considers hardware partition or blade to be a separate physical hardware system, and, therefore, a separate server. This therefore requires one to consider about the licensing requirement of each operating system separately

Fourth, there is need to have good tools to manage large number of servers. With virtualization it is easy to start creating many servers and each needs to be administered the same way they were in the physical world hence need for methods or tools to address server expansion. Lastly, one needs to have a robust backup system as when you have lot of VMs on one physical server you may run into challenges associated with backing up all the VMs on one server.

1.3 Problem Statement

Literature tells us that virtualization has both advantages (Gartner, 2008 and Burger, 2008) and disadvantages (Skooliki, 2008). Strategic Counsel (2007) argues that migrating from a non-virtual environment to virtualization reduces IT costs and increases IT efficiency. Samoilenko (2007) cautions that there are some factors associated with migration which could be a barrier to virtualization. Factors such as compatibility and support, licensing, planning deployment, return on investment and training. A previous study of the ICICI Bank of India (ITdrilldown, 2004) shows that they consolidated 230 physical servers to just five virtual servers hosting over 650 applications. This shows that the total cost of ownership went down and return on investment went high in the long-term. The implementers of virtualization at UNON, Dataposit Limited noted that the technology uptake is low besides its immense advantages but predict the trend in Kenya is changing positively through sensitization.

The UN currently has more than 2,000 users at the Gigiri complex who depend on the ICT services of United Nations Office at Nairobi. According to Mr. Anthony Mukomah the Computer Information Systems Officer at UNON during a personal visit to their office in July 2009, they migrated their ICT infrastructure to a virtual environment that they have been running for a couple of years now. The main reason for their migration to a virtual environment was to reduce on cost associated with hardware acquisition, improve security and ensure business continuity through high systems availability. The UNON case study was selected in this research based on the assumption that they have a lot of exposure, experience and knowledge worth researching on. As mentioned earlier, virtualization technology is relatively new in Kenya as evidenced by the low uptake by both large and medium size organizations. Therefore, the findings of this research will be worthwhile and invaluable to those intending to embrace the technology. This research was to investigate the following questions;

1. How much of their processes have migrated to virtual environment?

- Previous research shows that not all operating systems can be virtualized with current technologies

2. What are the benefits of working in a virtual environment?
    - What were the motivating factors of migrating from a non-virtualized to virtualized environments
3. What problems are encountered on the virtual environment?
    - research shows that different virtualization technologies and implementations tend to have some specific challenges manifesting more due to business driving factors into virtualization

1.4 Research Objectives

The objective of this research was to evaluate virtualization at the UN to specifically;

1. Establish the extent to which UN is utilizing virtualization in its ICT infrastructure
2. Investigate the benefits of having a virtualized environment
3. Analyze the challenges being encountered in the virtualized environment

1.5 Significance of the study

According to The Strategic Counsel (2007), there is a strong link between server virtualization and top business drivers for IT. As identified by The Strategic Counsel survey respondents, top business drivers for IT are: reducing IT costs, increasing IT efficiency; improving organizational or business performance; and improving IT service levels. The Strategic Counsel also finds that the benefits of server virtualization are clear and extremely well aligned with organizational expectations and actual server virtualization use cases. Primary benefits include improved server / system utilization rates; improved server reliability and uptime; and improved business continuity and disaster recovery. The study will have benefits to various parties;

To other researchers, the study will provide important reference information about virtualization. To practitioners, the study will enable them understand better the importance of migrating from a non-virtualized to a virtualized environment as well as know expected challenges. To policy makers, the study will help them set long term objectives and remain focused on their business vision and mission. The policy makers will be able to know the critical ICT areas that they should put more emphasis to ensure business continuity.

CHAPTER 2: LITERATURE REVIEW

2.0 Introduction to Literature Review

There are numerous facets to virtualization and a number of authors have presented work in this field. Hardware Level Virtualization presents an interesting method for the full utilization of modern computer systems and for increasing the performance of such systems. In this review, literature in the field of hardware / server virtualization, its performance and relevant aspects of benchmarking is presented, and offers a comparison of some of this literature. The literature also seeks to demystify and explain other relevant terms and technologies that compliment or challenge the virtualization technologies in the market. Further a business case example is presented for one of the banks in India that used the technology successfully hence backing its associated advantages.

2.1 Server Based Computing (SBC)

Local area networks and wide area networks have made it possible to centralize services (Kommeri, 2007). Instead of having all files and programs on local computers, as in the era of single user personal computers (PC), now these can be located on central servers and made reachable by others. Centralization is not only limited to storage services but it can also be used for execution of programs. Due to the development of the networks and the exponential growth of processing power, the servers can now serve increasing amount of client applications. The benefits of this transition from local execution to centralized is the decreased amount of administration work and increased system security.

Server based systems (SBC) resembles the old mainframe systems that were used in the early times of computers. The difference being now that the servers are a lot smaller and offer users more usable user interfaces.

The idea behind Server Based Computing is to move computation from PCs to centralized computing resources because the only things the PC would do is to maintain connection to the server, show graphics and transmit information from its peripherals such as keyboard, mouse or USB device.

This kind of PC is called a thin client that is not a special device. Almost any PC can be used as a thin client. The requirements for thin clients are low since most or all of the real computation is done elsewhere. Thin client does not need hard drives, powerful processors or powerful graphics cards. Having less powerful components means less energy consumption and less heat production. The life cycle of the old and obsolete PCs can be extended when used as thin clients. The requirements for a thin client vary on how much of the software is executed on remote servers

2.2 Comparison between Virtualization and related technologies

2.2.1 Cloud computing vs. Virtualization

Cloud computing is technically defined as a computing capability that provides an abstraction between the computing resource and its underlying technical architecture (e.g., servers, storage, networks), enabling convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interaction.

Clouds have five essential characteristics: on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service. (Armbrust et…al, 2009) define cloud computing as the applications delivered as services over the Internet and the hardware and systems software in the datacenters that provide those services.The services themselves have long been referred to as Software as a Service (SaaS). The datacenter hardware and software is what we will call a Cloud. On the other hand, virtualization is defined as the consolidation of server resources to share amongst applications.

Bluelock, 2009 argue that Cloud computing and virtualization are many times used interchangeably. While they intersect in fascinating ways, they aren't exactly the same thing. Server virtualization provides flexibility that is a great match for cloud computing, and is actually one of the greatest enablers of the technology. Moreover, cloud computing can be defined based on the virtual machine containers that is created with virtualization technology. But that doesn't make virtualization cloud computing. It can therefore be argued that the two technologies can complement one another.

2.2.2 Distributed/Grid computing vs. Virtualization

The convergence of virtualization technologies and distributed computing is an exciting development and the subject of much research in both academia and industry (VTDC, 2006). Wikipedia defines Grid computing as the combination of computer resources from multiple administrative domains applied to a common task, usually to a scientific, technical or business problem that requires a great number of computer processing cycles or the need to process large amounts of data. One of the main strategies of grid computing is using software to divide and apportion pieces of a program among several computers, sometimes up to many thousands. While as in grid computing a process is executed by many different distributed servers, in virtualization, resources are consolidated together and allocated to various processes.

2.3 Server Consolidation (virtualization)

Virtualization is a technique for running more than one operating system instance on a single computer simultaneously, making one computer appear as two, three or several more computers. The two main advantages of virtualization that are interesting in this respect are it makes it possible to run several virtual machines on one physical computer and that the virtual machines can be migrated to other physical computers seamlessly using a method called live migration.

Many applications in the enterprise require as little as 5-10 percent of a server's capacity (IDC, 2008). Business enterprises are opting to consolidate a number of applications on a single physical server, while planning capacity for growth in the total workload. According to a 2007 report by Yankee Group, 85% of virtualization software acquired by enterprises goes to server consolidation projects. The report also states that businesses have been able to better their server utilization from 10% to 80% in certain instances. Consolidating servers by running multiple instances of the Operating Systems (OS) and multiple applications on a single server has a large impact on the total cost of ownership (TCO) through the life of the server. Not only does consolidation reduce hardware expense, but also reductions in power, cooling, and floor space requirements by up to 50%-70% can be obtained, all while increasing service levels to the business.

2.3.1 High Availability Services

High availability computing is, as the term implies, a highly available computer system. How available the system should be depends on what the system should provide (Braastad, 2006). For Internet Service Providers (ISPs), high availability means 24/7 availability, while for other businesses it may stand for availability between for example 8am and 8pm each day. Depending on the service, an outage of one second might be insignificant or disastrous. Therefore, the degree of availability should match the purpose and suit the business needs of the company. The service level is the degree of service that shall be provided by the system according to the service level agreement. It is important that planned and unplanned outages do not exceed this service level degree.

Services that are considered as business critical are often categorized as high availability services. Since computers, software's and networks are not highly reliable in their normal nature, making it difficult to achieve 100% availability for critical services. However, systems running business critical services should be planned and designed from the bottom with the goal of achieving the lowest possible amount of planned and unplanned downtime. One example is redundant power supplies on the computers connected different power sources. If the electricity on one power source is lost, the servers will be unaffected and still available. In case of hard-drive failure, a backup disc is connected in mirror RAID that will do the job and leave the users unaffected from the failure. The availability of a system is given by the elements that the system depends on. These dependencies have to be analyzed in terms of backup and redundancy possibilities. A critical dependency that acts as a single point of failure should not occur in a high availability system. By making the dependencies redundant, we end up with one single physical computer that is completely redundant with no single points of failure resulting to a fault tolerant computer. Another approach is to have duplicates of entire physical computers. The result is that we end up with two or more physical computers which can act as backup computers which can failover services from each other.

2.3.2 Virtualization advantages

Braastad, 2006 noted that the various advantages of virtualization that can be utilized in high availability clusters as elaborated below;

Flexibility: One can run more than one instance of an operating system on a single computer simultaneously without affecting each other. It is possible to migrate a virtualized instance to another physical computer and the virtual instances can be graceful from the host operating system with features like ’pause’, ’resume’, ’shutdown’ and ’boot’. It is also possible to change the specifications of virtual computers while they are running, for example the amount of ram, hard disc size and more.

Availability: One can keep the virtualized instances running even though the physical node has to be shut down, i.e. for hardware upgrade or maintenance. This is done by temporarily migrating the virtual instances to another computer, and migrate them back when the maintenance is finished and the primary computer is ready to serve. Hardware can be changed, upgraded, maintained and repaired without downtime in the services

Scalability is added because is very easy to add or remove nodes. If the demand for capacity increases over time, it is very easy to insert a physical node with the basic cluster installation, and it will contribute in running the existing virtual machines that run services. This way, the cluster will scale with the company as it expands.

Hardware utilization is most likely increased if more than one operating system is hosted simultaneously. This is because virtual machines utilize hardware resources that are left idle by the host operating system. A typical x86 server deployments achieve an average utilization of only 10% to 15% of total capacity, according to International Data Corporation (IDC), a market research firm. Infotech research group, 2009 pointed that organizations can achieve a 40-75% onetime saving on implementation of virtualization through high hardware utilization.

Security is added because greater separation of services is introduced. Using multiple virtual machines, it is possible to separate services by running one service on each virtual machine. This approach is also called jailing of services. If one service is compromised, the other services are unaffected. Because break ins are usually done through software bugs in running software, it is important to only run software that are required for the company. Using virtualization, the server would contain a minimal install that could host several virtual machines. Each virtual machine consists of a minimal operating system install and one service, for example the web server. Let us say that the web server is being compromised. The web pages hosted will be unreliable, but the break in does not affect the remaining running services - the database server, mail server and the file server.

Greener computing / Green IT: Green computing can be defined as the responsible use of computers and their resources efficiently and effectively with minimal or no impact on the environment. The environmental impact and costs of excessive energy consumption are among the leading concerns today. According to the U.S. Environmental Protection Agency, datacenters used 61 billion kilowatt-hours of electricity in 2006 at a cost of $4.5 billion. This accounts for 1.5 percent of all power consumed in the U.S. unless enterprises take action, datacenter usage is predicted to double by 2011. Virtualization is so green that some utility companies, such as Pacific Gas & Electric, provide credits to businesses that incorporate virtualization into their datacenters. “Virtualization technology is helping our customers realize significant energy and cost savings, while addressing critical datacenter capacity issues,” explains Helen Burt, senior vice president and chief customer officer for Pacific Gas & Electric. “By providing financial support, we hope to increase industry adoption of this technology.”

High Infrastructure Utilization: Moore’s Law states that server capacities will double every eighteen months, and not only has this trend held for decades but is projected to hold well into the future. However, typical x86 server deployments achieve an average utilization of only 10% to 15% of total capacity, according to International Data Corporation (IDC), a market research firm. Organizations typically run one application per server to avoid the risk of vulnerabilities in one application affecting the availability of another application on the same server. Infotech research group, 2009 pointed that organizations can achieve a 40-75% onetime saving on implementation of virtualization

2.3.3 Virtualization disadvantages

Brastaad (2006) noted that besides the many obvious virtualization advantages, it also has disadvantages as explained below:

Overheads causing decreased performance has been the biggest con with virtualization. Performance is often being compromised due to flexibility, or contradictory. As the developers have worked hard to decrease the overhead, the amount of overhead has been reduced. Xen, which is para virtualized, has maximum a few percent overhead, bringing it very close to the performance of the standalone physical computer.

Single point of failure (SPOF) in the hardware is still an issue even though the virtual machine is decoupled from the hardware as it is still dependent on the hardware working. Failure in the hardware will most likely lead to failure in the virtual machine, which will force a reboot.

The **management interface** is closely linked to the virtualization platform. This can be a problem as it encumbers consolidation of several platforms into the same environment. However, there is an ongoing effort to create an abstract management interface that will separate virtual machines from the management console.

Complex Deployment and Maintenance is another common issue with virtualization technologies. When a large number of logical resources are consolidated onto a small number of physical resources, the consequences of downtime increases. To compensate, redundant physical hardware configurations are often used, increasing the complexity of deployment. This does not necessarily mean the physical hardware deployment is any more complex than a non-virtualized environment, but it does add a layer of complexity onto typical physical deployments. With the added features that virtualization provides, the possibilities increase, thus complicating the deployment options and management of the environment.

Server Sprawl is a common issue with virtualization technologies. This occurs when administrators feel that the separation of applications by complete operating system environments is more secure and reliable than separation of applications by the security within an operating system. By installing an application in its own dedicated operating system, upgrades can be isolated to either that single application or operating system without affecting other applications or operating systems. Availability is often improved by this technique, however, the management and cost of the operating systems produces a burden not predicted. With the ability to copy virtual machines so easily, this technique is made effortless, without thinking of the consequences of the management and cost burdens.

2.3.4 Virtual Machine Monitor Architectures

This section presents a few different Virtual Machine Monitor (VMM) architectures. However, before they are discussed in detail some basic terminology should be understood by the reader. The VMM is a software layer that manages one or more virtual machines. A virtual machine is a software abstraction of the real machine that is isolated from the VMM and is also referred to as a domain. Domains can contain any type of software application, but in the context of this paper the application will be an Operating System. The operating system running in a domain is referred to as the guest operating system. The applications running on top of the guest operating system are called the guest applications. The VMM itself can be run on top of hardware or software; we name both the host architecture. Figure 2.1 shows the basic organization of a system that uses a VMM; the hardware hosts the VMM while the guest OS and applications make up a domain. (Ferstay, 2001) researched on some of the present different VMM architectures being used today.

Nano-kernel systems

The basic idea behind Nano-kernels is that they are a separate, small bit of core functionality that resides within the OS to do some specialized processing (e.g. real-time interrupt processing and scheduling). Nano-kernels can be considered a lightweight version of a VMM because typically they do not virtualize very many resources. Examples of nano-kernel systems are Jaluna (jaluna, 2005) and RTLinux (Rtlinux website, 2005). In both of these systems, the virtualized resource is interrupt handling and scheduling. In RTLinux, there is a small real-time core that is loaded as a Linux kernel module and runs alongside of the regular Linux kernel. When the RTLinux core is loaded, it takes over all of the interrupt processing and scheduling for the machine by interposing itself at the interrupt handler routine entry points. The RTLinux core then handles all interrupt processing and schedules the actions of the system -including the Linux kernel and its scheduler- according to hard, real-time requirements dictated by the system administrator.

Full Virtualization Systems

Full virtualization systems virtualize every resource on a computer system. Guest operating systems do not have access to physical devices, physical memory, or physical CPUs. Instead, the guest operating systems are presented with virtual devices, virtual memory, and virtual CPUs. In a full virtualization system, the virtualized interfaces presented to the guest operating system look and feel exactly like the interfaces of the real machine and therefore the guest OS and applications may run on the virtual hardware exactly as they would on the original hardware. One such system is VMware [sugarman et al…]. Full virtualization systems have another benefit besides allowing guest operating systems to run unmodified on top of the VMM. The virtualized resources provide a layer of isolation that protects the VMM from the actions of guest operating systems. The virtual memory implementation in the VMM protects it from the guest OS in the same way that the virtual memory implementation in the guest OS protects it from the guest applications. In a similar way, the virtual CPU implementation in the VMM uses the processor's operating modes in the same way that the guest OS uses the CPU operating modes to protect itself from guest applications (Schroeder, Saltzer, 1972).

VMMs built in the full virtualization style are considered the most secure of all VMM architectures because of the strong isolation they provide. The Livewire project is the first attempt to build a security service into the full virtualization VMM; they used VMware as a base. The major negative characteristic of full virtualization systems is degraded performance. This is obvious as every privileged instruction, every interrupt, every device access, and every memory access (etc.) has to be virtualized by the VMM. The performance measurements carried out during the evaluation of the Xen VMM confirmed that a guest OS running on top of VMware can be as much as a one hundred times slower than the same OS running directly on hardware.

Paravirtualization Systems

The architecture of paravirtualization systems has very much the same look and feel as that of a full system virtualization system but with major differences in design decisions. Paravirtualization was born out of the observation that full system virtualization is too slow and complex on today's commodity hardware. Paravirtualization aims to retain the protection and isolation found in the full system virtualization approach but without the implementation complexity and associated performance penalty. The main idea behind paravirtualization is to make the VMM simpler and faster by relaxing the constraint that guest operating systems must run on the VMM without being modified. In a paravirtualization system, the guest operating system code is modified to access the VMM directly for privileged access, instead of going to the virtual resources first and having the VMM intervene. One VMM implementation that uses a paravirtualization approach is the Xen hypervisor (Dragovic et al.., 2004). Xen runs on x86 computer systems and supports commodity operating systems (Linux, NetBSD, WinXP (in development)) once they have been ported to the Xen-x86 hardware architecture. The Xen-x86 architecture is very similar to x86, with extra function calls into the VMM needed for virtual memory page table management and other features of the x86 that are difficult or slow to fully virtualize. Operating systems ported to Xenx86 must also port their device drivers to use Xen's lightweight event notification system instead of using interrupts for communications. In addition, Xen is very lightweight at roughly 42000 lines of code (Barham, et al.., 2003).

Another VMM that uses a paravirtualization approach is Denali (Whitaker et al…, 2002). Denali has the same basic goals as Xen, although Denali makes no attempt to support commodity operating systems. For this reason, Denali gives up certain features that are difficult to fully virtualize or paravirtualize such as virtual memory. In Denali, a domain is meant to support a single application (or OS). If you need two applications to be isolated from one another then you must run them in separate domains (virtual machines). Denali forfeits features needed for operating system support in order to gain simplicity and security.

The main drawback to paravirtualization systems is that they cannot host commodity operating systems without first porting them to run on the VMM. To port an OS to a new architecture takes time, but work on the Xen VMM shows that if the target architecture is very similar to the original architecture and most of the major changes to the OS are architecture independent, then the cost of porting the OS is outweighed by the rewards of paravirtualization. Another worry for VMMs built using a paravirtualization design is that the calls directly to the VMM must remain secure. By providing direct calls into the VMM for guest operating systems to use instead of transferring control indirectly through hardware traps it is possible to expose a security hole. This also adds complexity to the VMMs thin interface to guest operating systems.

It should be noted that some computer architectures have hardware support for virtualization which makes it possible for a VMM that utilizes paravirtualization to run hosted operating systems without modification. Processors that support Intel Virtualization Technology (Intel, 2005) or the AMD Pacica Technology (AMD, 2005) are two examples of such architectures. XenSource is currently working on migrating Xen to Intel VT. Hardware support for virtualization has also been injected into the IBM PowerPC architecture in the form of IBMs Enterprise

Hypervisor Binary Interface. IBM is working on migrating Xen to the PowerPC 970 processor which supports the hypervisor binary interface (Blanchard & Xenidis, 2006).

Software Hosted Virtualization Systems

While the Denali and Xen VMMs run directly on top of the hardware, there is another architecture which has the VMM run on top of a general purpose operating system. software hosted VMMs leverage the services of the general purpose operating system (the host) to simplify the process of providing virtual hardware abstractions to guest operating systems. The hosted VMM can use any virtualization technique to host guest operating systems (e.g. full virtualization or paravirtualization). UMLinux (King, et al, 2003) is an example of a software hosted VMM.

In UMLinux, the guest OS and guest applications run in a single process, the guest machine process. The guest machine process communicates with the VMM process via shared memory and IPC. As mentioned earlier, the key benefit to using a hosted VMM is that you get to use the abstractions and services of the host OS to provide virtualization to the guest OS. For example, in UMLinux, the guest-machine process serves as the virtual CPU; host files serve as virtual I/O devices; host signals serve as virtual interrupts; etc. The one area where hosted VMMs fall short is performance. There is extra overhead associated with using the host operating systems services and abstractions instead of working directly with the hardware (Clark, et al…, 2004). Also, hosted VMMs do not make very much sense in a security setting as they rely on the services of the host OS to maintain their integrity in order to provide correct virtualization.

2.3.5 Full Virtualization of x86 Hardware

In 1999, VMware introduced virtualization to x86 systems to address many of these challenges and transform x86 systems into a general purpose, shared hardware infrastructure that offers full isolation, mobility and operating system choice for application environments. Unlike mainframes, x86 machines were not designed to support full virtualization and VMware had to overcome formidable challenges and obstacles to create virtual machines out of x86 computers.

The basic function of most CPUs, both in mainframes and in PCs, is to execute a sequence of stored instructions (i.e. a software program). In x86 processors, there are 17 specific instructions that create problems when virtualized, causing the operating system to display a warning, terminate the application, or simply crash altogether. As a result, these 17 instructions were a significant obstacle to the initial implementation of virtualization on x86 computers.

To handle the problematic instructions in the x86 architecture, VMware developed an adaptive virtualization technique that “traps” these instructions as they are generated and converts them into safe instructions that can be virtualized, while allowing all other instructions to be executed without intervention. The result is a high-performance virtual machine that matches the host hardware and maintains total software compatibility. VMware pioneered this technique and is today the undisputed leader in virtualization technology.

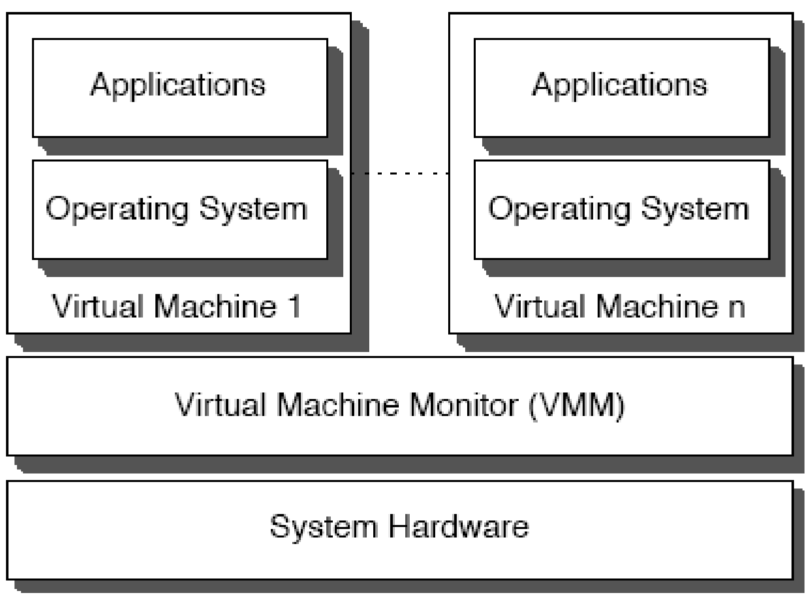

Figure 1 : Virtual Machine Monitor - Virtual Machine Relationship.

Source: Fast Secure Virtualization for the ARM Platform

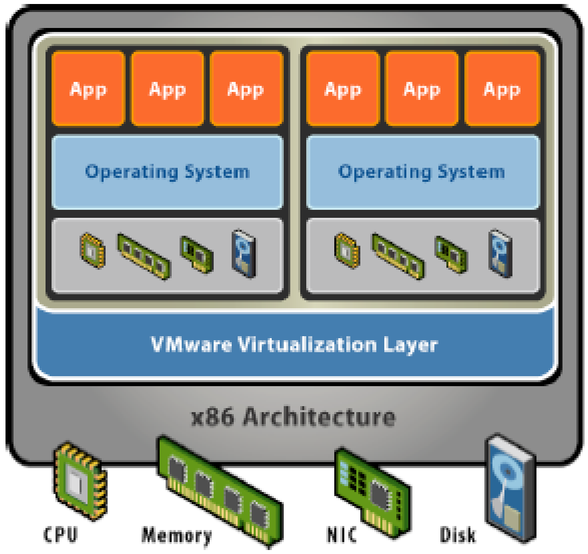

Figure 2: Virtual Machine Monitor – Hardware simulation.
Source: VMware

2.4 The Challenges of x86 Hardware Virtualization

X86 operating systems are designed to run directly on the bare-metal hardware, so they naturally assume they fully 'own' the computer hardware. As shown in Figure 3, the x86 architecture offers four levels of privilege known as Ring 0, 1, 2 and 3 to operating systems and applications to manage access to the computer hardware.

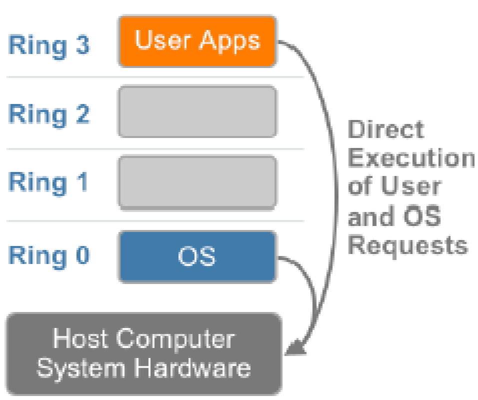

Figure 3: Privilege level architecture without virtualization.
Source: VMware Website

While user level applications typically run in Ring 3, the operating system needs to have direct access to the memory and hardware and must execute its privileged instructions in Ring 0. Virtualizing the x86 architecture requires placing a virtualization layer under the operating system (which expects to be in the most privileged Ring 0) to create and manage the virtual machines that deliver shared resources.

Further complicating the situation, some sensitive instructions can't effectively be virtualized as they have different semantics when they are not executed in Ring 0. The difficulty in trapping and translating these sensitive and privileged instruction requests at runtime was the challenge that originally made x86 architecture virtualization look impossible. VMware resolved the challenge in 1998, developing binary translation techniques that allow the VMM to run in Ring 0 for isolation and performance, while moving the operating system to a user level ring with greater privilege than applications in Ring 3 but less privilege than the virtual machine monitor in Ring 0. While VMware's full virtualization approach using binary translation is the de facto standard today based on VMware's 20,000 customers installed base and large partner ecosystem, the computing industry as a whole has not yet agreed on open standards to define and manage virtualization. Each company developing virtualization solutions is free to interpret the technical challenges and develop solutions with varying strengths and weaknesses. Three alternative techniques now exist for handling sensitive and privileged instructions to virtualize the CPU on the x86 architecture; first there is the Full virtualization using binary translation, secondly, OS assisted virtualization or paravirtualization and thirdly, Hardware assisted virtualization (1st generation)

2.5 Taxonomy of virtualization techniques

Table 2.1 shows taxonomy based on features and requirements where the type is one of NK (Nanokernel), FV (Full virtualization), or PV (Paravirtualization).

Important aspects of VMM systems include the type of virtualization employed, how the VMM is hosted, whether it isolates the CPU and memory from hosted operating systems, whether virtualized devices are present, and how the VMM performs. Finally, it is important that we identify which VMMs are Open Source implementations and which are not. Choosing a VMM implementation to use as the basis for VMM-level research requires access to existing source code.

VMM	**Type**	**Host**	**CPU & Memory Isolation**	**Device Virtualization**	**Performance**	**Open Source**
TRLinux	NK	Linux Kernel	No	NO	Very Good	Yes
VMware	FV	General purpose OS	Yes	Yes	Poor	No
UMLinux	PV	General Purpose OS	Yes	Yes	Mediocre	Yes
Xen	PV	Hardware	Yes	Yes	Good	Yes
Denali	PV	Hardware	Yes	Yes	Good	No

Table 1: Taxonomy of virtual machines monitors architectures.
Source: Fast Secure Virtualization for the ARM Platform, page 18

The performance penalty imposed by certain VMM architectures cannot be ignored. Even as processor speeds increase, the performance penalties will still be noticeable. This is because the performance of a system running a VMM is greatly dependent on the performance of the memory subsystem. Even in a high performance paravirtualization environment such as Xen, memory performance becomes an issue because of the instruction cache, data cache, and translation lookaside buffer (TLB) flushes that occur as a result of switching context from one domain to another. Additional overhead is even more important for handheld devices where every operation is paid for with battery power. A paravirtualization environment allows us to make more efficient use of the processor and memory subsystem on a handheld device, saving its most important resource: its battery. Therefore, an open source paravirtualization architecture would be better than full virtualization alternatives on a mobile device. It would provide all of the isolation that is needed for a secure VMM while incurring minimal performance penalties and more efficient use of battery power. Currently, there are no hardware hosted paravirtualization-style VMMs available for the ARM architecture. A port of Xen from x86 to the ARM architecture would fill this void. Such a port would also be a good starting point for creating secure VMM-level security services.

2.6 Virtualization of ICICI Bank in India

ITdrilldown (2004) presented the case of ICICI Bank as the first, at the scale it operated, to successfully leverage enterprise-wide data warehousing and business intelligence. The second largest bank in the India was also the first bank to virtualize their servers. ICICI Bank's IT team, led by Vohra, used virtualization to arrest an electronic infrastructure spill-over at its datacenters. They consolidated 230 physical servers to just five, running a little under 650 applications on a virtualized environment. It required them to develop the unparalleled technology ability to run 60 virtual machines on a single server but it saved the bank over a crore annually in power, cooling and space.

The result of virtualization was high reduction of server count while that of its closest competitors runs into four or five digits. That was incredibly low for a bank of its size with assets amounting to Rs 384,970 crore (US$7,899), and with 1,400 branches and 4,530 ATMs across the country. The business problem ICICI Bank forever grapples with lies at the core of its standardized Windows NT architecture. Any application typically requires a Web tier, an application tier and a database tier that they say is a necessary evil.

Running that many applications had a domino effect and demanded an ongoing investment in servers, power consumption, rack space, switching gear because as all these servers need to be interconnected to storage and networking sub-systems for management, availability and recoverability. Vohra the ICICI bank general manager says that they were worried that they were ending in a server or an electronic sprawl because they were not utilizing their servers properly but had to keep them because some development or some testing could happen.

2.7 Managerial implications from the research

Today's IT organizations are faced with the daunting task of optimizing all aspects of their departments, including people, processes and technology. Virtualization is one particularly exciting example of this trend and has been defined as the pooling of various IT resources in a way that masks the physical nature and boundaries of those resources from the users. It promises to unlock underutilized server capacity through partitioning, enabling the ability to run multiple operating systems on one server. The most obvious benefits of virtualization are not just the ability to achieve higher server and overall computing infrastructure utilization, but also improved responsiveness and flexibility since virtual resources can be moved or modified dynamically to reflect changing business needs. VMware (2009), one of the leading virtualization technology developers classify the benefits accruing from virtualization as both financial and operational cost saving as described below.

Capital cost savings come in the form of reduced expenses for hardware acquisition and data center real estate. First, virtualization helps you cut capital costs by reducing the number of physical servers necessary to support your infrastructure. With virtualization, server consolidation ratios are typically in the range of 8:1 to 15:1. This leads to a reduction in expenditures that is significantly greater than 50 percent. Secondly, with a virtualized environment, you can reduce your capital expenditures for storage devices by consolidating data storage. The third cost saving is derived from the fact that with fewer physical servers in a virtualized environment, you need less network switching infrastructure estimated at greater than 70 percent reduction in capital expenditures for network hardware. Lastly, there is data center space savings as virtualization helps you make better use of valuable data center real estate. It allows you to remove servers from your environment to open up data center floor space estimated at over 60 percent in cost saving

The **operational cost savings** in a virtualized environment are also substantial stemming from reductions in power and cooling costs, server provisioning and time saving costs, Disaster and recovery costs, management costs and the costs associated with server downtime. With *Power and cooling*, virtualization can help you hold the line on rapidly rising power and cooling costs through reduction in the number of physical servers in your environment. It's estimated that removing one server from your environment can save $835 in power and cooling. Virtualization also reduce *server provisioning costs and time savings* as virtualizing allows you to provision servers in less time, which in turn leads to reduced infrastructure management costs. Virtualization help by reducing *disaster recovery costs* and can also help you keep your business up and running during disasters and other disruptive events. Reducing the number of servers through consolidation reduces the number of servers that must be restored in case of disaster.

Further, virtualization helps reduce *unplanned downtime costs* by enabling faster recovery from unplanned downtime and allowing for better management of planned downtime, virtualization helps you further reduce indirect costs. Virtualization reduces unplanned downtime by up to75 percent.

2.8 Business case for virtualization

Information is the second most valuable firm asset after its human resources and organizations need to protect it zealously. Information is considered as the greatest business enabler of the 21st century and lack or inadequacy of it can lead to making the wrong decisions. This therefore calls for measures and technologies that ensure the safety of data. The following are some of the major strategic business benefits for embracing virtualization.

Business Continuity and Disaster Recovery. Server consolidation helps to increase application availability because economies of scale make it easier to use larger servers with improved reliability, availability, and serviceability over smaller ones. But failures ranging from a single server to an entire regional datacenter can bring a business to its knees. Fortunately, most virtual infrastructure provides straightforward solutions for both disaster recovery and business continuity. Secondly, Virtualization also helps in ICT governance of an organization in that it provides a structure of process and relationships to provide strategic direction, ensure objectives are achieved, risk managed appropriately, and resources used responsibly.

The third benefit is that virtualization helps organizations lower Total Cost of Ownership and increase Return On Investment. It should however be noted that the initial outlay cost for the technology is high but with little subsequent costs. With climate change being a challenge worldwide, virtualization introduces green computing that is environment friendly

CHAPTER 3: RESEARCH METHODOLOGY

3.0 Research Design

The chapter contains the methodology used to perform the research. It describes the research design, process used to collect data and the process used to analyze the results detailed from the data collected. The research design used was a case study because the researcher could learn more effectively when actively involved in the learning process (Bonwell and Eison, 1991; Sivan et al, 2001). The case study approach is one way in which such active learning strategies can be implemented in our institutions.

3.1 Population

The study was modeled on a census of all the ten staff in both infrastructure and software section of ICT department and two Finance staff at the United Nation Office at Nairobi, Gilgil. The population was classified into two groups with each group having a different questionnaire. The first group population consisted of both the heads ICT and Finance departments who were interviewed face to face as key informants. They were selected as they were to provide both strategic and business case for implementing virtualization. The second group population consisted of all other staff in ICT department. The ICT staff in infrastructure and software section were selected as they had good understanding of virtualization technology while as the finance staff gave facets like user experience and Return On Investment and total cost of ownership over the period of usage.

3.2 Data Collection

Primary data was collected through a self-administered questionnaire (See Appendices) that were constructed using open-ended, closed-end and Likert-type questions. The questionnaire was divided into two major parts consisting of key informants and general questionnaire. The key informant questionnaire intended to get more candid and in-depth answers regarding virtualization at UNON. For the key informants, a face to face interview was conducted to enable collect as much information as possible as well as information that cannot be collected using the general questionnaire. Due to the design of the general questionnaire, it was administered using both interview and 'drop and pick' method. The interview was conducted to avoid subjectivity resulting from limiting the respondents' answers to the questions asked as well as enabling the researcher get information that otherwise cannot be gotten through 'drop and pick' questionnaires. This facilitated more in-depth interaction with the respondents will comprising of selected UNON employees in the ICT department and finance department. The guiding questionnaire was divided into four parts. The first part was to capture the general information regarding UNON as a service lending organization. The second part covered the extent to which UNON has virtualized their data center and the various offices they support. The third part covered the various benefits they get from virtualizing their servers while the fourth part focused on the various challenges they encountered during and after virtualizing their servers.

3.3 Data analysis

Content analysis was used to analyze the data collected through the questionnaires and interviews. This was described as a systematic qualitative description of the composition of objects or material of study. It involved observation and detailed description of objects, items or things that comprise the study (Mugenda, 1999). The data collected for the three objectives was presented as follows; for general information questions, frequencies and percentages were used. Data in part B was presented using descriptive analysis. Data in part C and D was analyzed using descriptive and statistical methods to establish their relevance to the organization.

Statistical Package for the Social Sciences (SPSS) for windows version 12.0; a statistical software application was used to analyze quantitative data collected for the study.

CHAPTER 4: DATA ANALYSIS AND DISCUSSIONS

4.1 Introduction

This section presents the data analysis, presentation, and discussion on data gathered for the case study of evaluation of virtualization technology at UNON offices at Gilgil. Factor Analysis and percentages have been used extensively to make any necessary conclusions. The presentation of the analysis was captured in two parts the first part captures analysis made using percentages and the second part cover factor analysis. The major findings of the study as they relate to each of the research objectives are presented.

The following are the research objectives that guided the study:

1. To establish the extent to which UN is utilizing virtualization in its ICT infrastructure
2. Investigate the benefits of having a virtualized environment
3. Analyze the challenges being encountered in the virtualized environment

A population of 12 respondents that includes a census of all ICT ten staff (in both Infrastructure and Software section of the department) and two Finance department staff was selected. All the ICT staff were served with a questionnaire as it was believed they have virtualization knowledge while as a face to face interview was administered to the two Finance department staff. The collected data was cleaned, coded and analyzed using SPSS software Version 12.0.

4.2 Demographics

The study recorded a 91% response rate with 11 respondents out of the targeted 12 being able to respond successfully. The study interviewed both the ICT Manager and the Deputy Finance Manager to get in depth information. Further the study administered a questionnaire to all the ICT staff through the drop and pick method.

4.2.1 Gender

The study did not have control of the gender of the respondents as its population was largely targeting the staff in ICT department. The study had 27% of women and 73% of men responding from both the ICT and finance department.

Table 2 : Gender Distribution

Gender	Frequency	Percentage
Men	8	27%
Women	3	73%
Totals	11	100%

Source: research data

Figure 4 : Gender representation

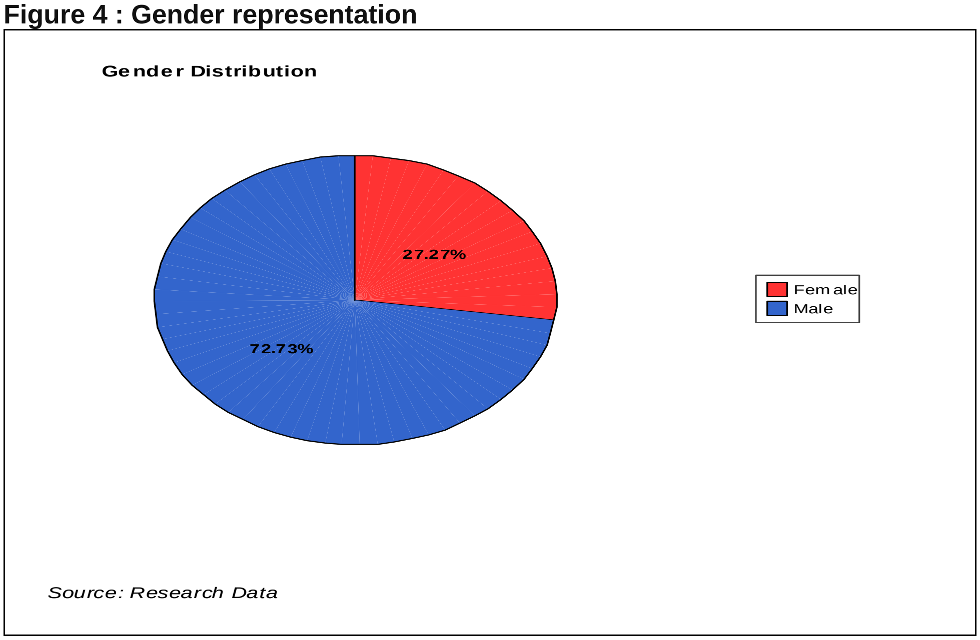

4.2.2 Age

Figure 5 : Age Distribution

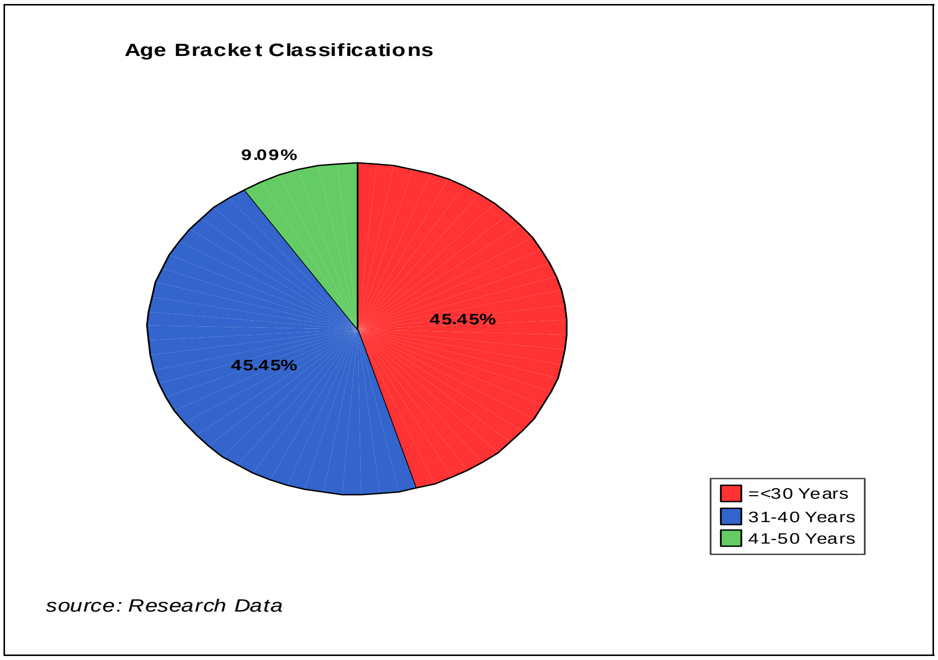

Source: research data

Figure 6 : Age distribution

Age Bracket	Frequency	Percentage
Below 30 years	5	47%
31 – 40 years	5	46%
41 – 50 years	1	9%
Total	11	100%

Source: research data

The age of the respondents interviewed were less 50 years old. Of the 11 responses received, 47% fall in less than 30 years age bracket, with another 46% belong to the 31-40 age bracket. Only one respondent represented by 9%

was over 40 years old. The various age brackets representations clearly shows that the majority of workers fall under 40 years of age.

4.2.3 Level of education

Table 3 : Education Level

Respondents	PHD	Postgraduate degrees	Undergraduate degree	Diploma
Number	1	1	8	1
Percentage	9%	9%	73%	9%

Source: Research Data

Figure 7 : Education Level

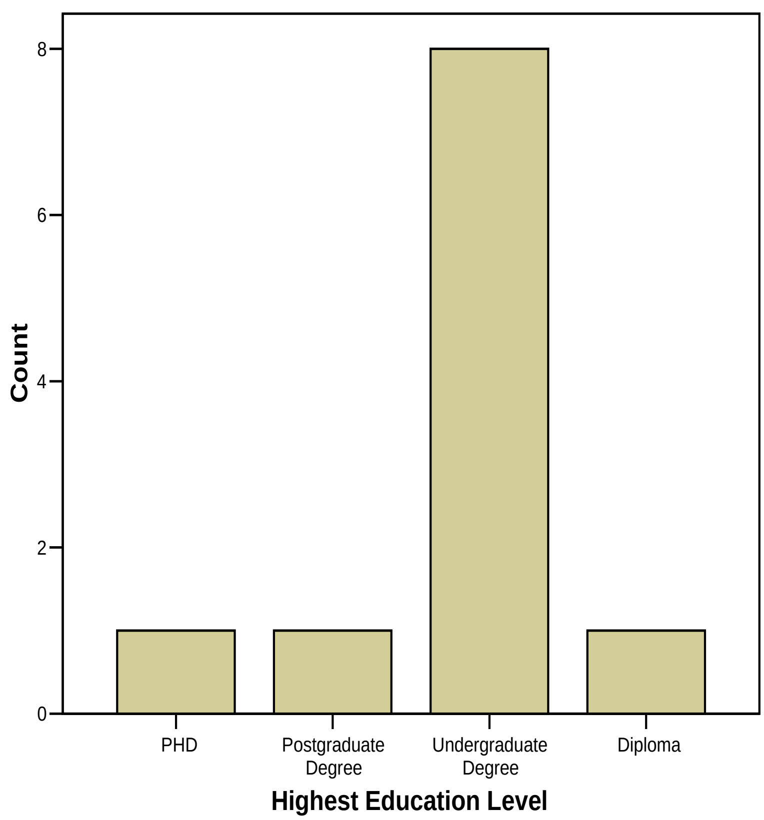

Source: Research Data

The highest level of education is a PhD degree with most staff having an undergraduate degree representing 73% of all the respondents. There is a representation of 9% each for both master’s degree and diploma level.

4.3 Virtualization usage at UNON

The UN offices in Nairobi support over 10 agencies within Kenya and the same number outside Kenya. In particular, the ICT services at UNON serve the whole region where various agencies access different types of systems as per their needs. The number of UNON staff in Kenya is estimated to be over 3000 with over 40 of them in IT but different stations. The virtualization technology used at UNON is VMware ESX3. They preferred this as it’s the most developed, robust virtualization software in the international market. There are four Virtualized Certified Professionals at UNON who support the virtual infrastructure both at the primary site and DR site.

Before virtualization in the year 2007, UNON had over 20 physical servers in their data center using the concept of one server one application. The servers were hosted different operating systems including Solaris Unix, Windows operating systems running in the servers. The initial objective at UNON was to virtualize windows servers to counter frequent server failures through data replication to their Disaster and Recovery Center.

Upon virtualization, UNON replaced over 20 servers running windows server operating systems with four servers at the primary datacenter and four servers at the DR site. The four servers currently host over 30 virtual machines each running an application that was previously consigned to a dedicated physical server. However, during virtualization, the Solaris Unix operating systems could not be virtualized as its architecture is complex and not supported by VMware software. The four servers running such systems were left intact during the implementation but the systems have since been migrated to Linux Enterprise server operating systems that can be virtualized.

4.4 Factor Analysis

Principal Component Analysis (PCA)

The study used PCA to test the variability between the research variables to establish the kind of relationship they have. Further the factor analysis reduces the variables by comparing the variability of the components tested in the study.
From the study the Principal Component Analysis tool showed that there is low variability between the various virtualization benefits variables being tested with ease in server provisioning being the least affected in the test.

4.4.1 Benefits of virtualization

The table 4.below shows the communalities between the various virtualization benefits factors. The study noted that the proportional variance between the various factors is low with. The initial communality variance for all the items is one and the study recorded Ease of server provision to have the lowest variability at 0.07and Reduced power and cooling to have the highest the highest variability at 0.17. In general the study noted that the factors are closely correlated showing that they all support server virtualization as a technology.

Table 4 : Virtualization Benefits Communalities

	Initial	Extraction
Reduced Capital costs, resource & space utilization	1.000	.850
Stability & Security	1.000	.940
Flexibility in OS deployment	1.000	.924
Ease in server provision	1.000	.993
Reduced OS licensing costs	1.000	.975
Reduced power and cooling	1.000	.830
Ease of OS and application testing	1.000	.893
Run multiple Operating systems on one single server	1.000	.953
Improved flexibility, High Availability of application and performance	1.000	.921
Build up business continuity through improved D&R	1.000	.969
Improved enterprise desktop management & control	1.000	.977
Green Computing that is environment friendly	1.000	.970
Improves system compliance	1.000	.932
Dynamic load balancing	1.000	.962
Ease of backup and restore	1.000	.978
Reducing operational risks	1.000	.917
Support for IT governance	1.000	.991
Real time response to changing business needs	1.000	.928
Reduced Total Cost of Ownership	1.000	.851

Extraction Method: Principal Component Analysis.
Source: Research Data

Table 5 : Benefits Total Variance Explained

Component	Initial Eigenvalues			Extraction Sums of Squared Loadings		
	Total	% of Variance	Cumulative %	Total	% of Variance	Cumulative %
1	5.288	27.834	27.834	5.288	27.834	27.834
2	4.727	24.882	52.716	4.727	24.882	52.716
3	3.318	17.465	70.181	3.318	17.465	70.181
4	2.069	10.888	81.069	2.069	10.888	81.069
5	1.340	7.054	88.123	1.340	7.054	88.123
6	1.010	5.317	93.440	1.010	5.317	93.440
7	.494	2.599	96.039			
8	.396	2.087	98.126			
9	.234	1.230	99.356			
10	.122	.644	100.000			
11	8.396E-16	4.419E-15	100.000			
12	4.854E-16	2.555E-15	100.000			
13	3.041E-16	1.600E-15	100.000			
14	8.016E-17	4.219E-16	100.000			
15	2.771E-17	1.458E-16	100.000			
16	-1.730E-16	-9.105E-16	100.000			
17	-2.712E-16	-1.427E-15	100.000			
18	-4.018E-16	-2.115E-15	100.000			
19	-6.061E-16	-3.190E-15	100.000			

Extraction Method: Principal Component Analysis.
Source: Research Data

The study investigated 19 generic benefits of virtualization as shown in table 5 above. Variances extracted by the factors in the second column (Eigenvalue) above, we find the variance on the new factors that were successively extracted to be 6 with a value greater than 1. In the third column, these values are expressed as a percent of the total variance. As we can see, factor 1 accounts for 27.8 percent of the variance, factor 2 for 24.8% percent, and so on. As expected, the sum of the eigenvalues is equal to the number of variables of 19 being tested. The fourth column contains the cumulative variance extracted and reaches by the sixth component, the total variance is at a significant 93.4%. The contributions of the six extracted or reduced factors are mapped in the component matrix. (See table 6 below)

Table 6 : Virtualization Benefits Component Matrix (a)

	Component					
	1	2	3	4	5	6
Reduced Capital costs, resource & space utilization	-.390	-.110	-.401	-.287	.626	.223
Stability & Security	.589	-.488	-.402	-.336	.165	-.230
Flexibility in OS deployment	.705	.140	-.008	.085	-.139	-.617
Ease in server provision	.585	-.531	.441	-.398	.130	.000
Reduced OS licensing costs	.123	.597	.770	.052	-.079	-.052
Reduced power and cooling	.224	.024	.399	.721	.147	-.281
Ease of OS and application testing	.436	-.351	-.632	.096	-.388	.143
Run multiple Operating systems on one single server	.866	.216	.156	-.351	-.095	-.003
Improved flexibility, High Availability of application and performance	.191	.672	.040	.266	.600	-.001
Build up business continuity through improved D&R	.590	-.536	.468	-.044	.194	.272
Improved enterprise desktop management & control	.814	-.497	.085	-.222	.086	-.063
Green Computing that is environment friendly	.516	.357	-.622	.423	-.022	.104
Improves system compliance	.476	-.685	.072	.415	.208	.125
Dynamic load balancing	.424	-.350	.202	.693	-.082	.362
Ease of backup and restore	.659	.561	.370	-.216	.152	.152
Reducing operational risks	-.331	-.414	.642	-.126	-.405	.211
Support for IT governance	.495	.805	-.173	-.091	-.168	.176
Real time response to changing business needs	.194	.868	.246	-.164	-.058	.217
Reduced Total Cost of Ownership	.632	.241	-.571	-.037	-.149	.208

Extraction Method: Principal Component Analysis.
a 6 components extracted.
Source: Research Data

The 6 Principle components above were extracted using The Kaiser criterion (see table 5) that takes only Eigenvalues greater than 1. It maps the benefits of virtualization against the six reduced factors from the initial 19. It was noted that the two first extracted factors carried a total weight of 52.7% in the cumulative percentage eigenvalues of which ten of the tested variables lie (represented by positive values in variable 1 & 2). Further, the statistics show that there are only four benefits represented in both reduced factor 5 & 6 which account for only 12.3% of the weight. The statistics affirm to the benefits of virtualizations in literature review (Braastad, 2006).

4.4.2 Challenges of virtualization

Table 7 : Virtualization Challenges Component Matrix (a)

	Component		
	1	2	3
Server Sprawl	.648	.702	.295
Some application need dedicated resources	.848	-.443	.291
Licensing of operating systems and applications	.505	-.065	.861
Need of a robust backup and restore system	.667	-.712	-.221
Performance degradation of some applications	.962	.087	.260
Unanticipated running costs	.973	.154	-.172
Unused virtualization features	.962	.087	.260
Increased storage requirements	.935	-.279	-.219
Management complexity	.971	-.119	-.206
Risk of total failure	.825	.557	-.092
IP address mapping	.935	-.279	-.219
Incompatible hardware with virtual operating system	.457	.889	.023
Slow applications	.972	-.211	.107
Maintenance of the virtual infrastructure	.528	.661	-.534
Need of trained specialists	.935	-.279	-.219

Extraction Method: Principal Component Analysis. a 3 components extracted.
Source: Research Data

The study had initial 15 variables representing the generic challenges of virtualization as shown in table 7 above. The factor analysis reduced the factors to three principle components hence showing their low variability. This is confirmed by the fact that the statistics did not extract a variance table (see table 5) as in the benefits of virtualization. First, all the study variables have been represented in the first extracted component with positive values showing their central tendency confirming that disadvantages of virtualization are seldom mentioned during virtualization.

Secondly, the statistics represented four variables within the three reduced components and six variables the first two components. Also it was noted that only four variables were not represented in at least two extracted components showing their relevance in virtualization is almost the same.

Table 8 : Virtualization challenges Pattern Matrix (a)

	Component		
	1	2	3
Server Sprawl	-.225	.895	.404
Some application need dedicated resources	.781	-.275	.489
Licensing of operating systems and applications	-.095	-.082	1.065
Need of a robust backup and restore system	1.154	-.521	-.110
Performance degradation of some applications	.486	.335	.445
Unanticipated running costs	.702	.494	-.048
Unused virtualization features	.486	.335	.445
Increased storage requirements	1.026	.026	-.086
Management complexity	.925	.207	-.074
Risk of total failure	.244	.868	-.001
IP address mapping	1.026	.026	-.086
Incompatible hardware with virtual operating system	-.344	1.091	.054
Slow applications	.808	.047	.288
Maintenance of the virtual infrastructure	.212	.974	-.557
Need of trained specialists	1.026	.026	-.086

Extraction Method: Principal Component Analysis. Rotation Method: Promax with Kaiser Normalization.
a Rotation converged in 4 iterations.
Source : Research Data

The initial matrices were rotated using promax with Kaiser Normalization to give the above results after four iterations. Kaiser normalization was preferred due to the assumption of correlation among variables. The total Variance explained tables shows the final output of the principle axis factoring which expresses the sum of the squared loading of the variables, however the sum of the squared loadings are irrelevant because the variables are assumed to be correlated.

From the above table, there are four variables that scored a rating over one in the first component emphasizing their importance as regards challenges of virtualization. These included; need for robust backup and restore systems at 1.154 followed by increased storage requirements, IP address mapping and need for trained personnel all at 1.026 rating.

Table 9 : Virtualization Challenges Structure Matrix

	Component		
	1	2	3
Server Sprawl	.365	.935	.608
Some application need dedicated resources	.889	.242	.762
Licensing of operating systems and applications	.371	.244	.991
Need of a robust backup and restore system	.871	-.044	.255
Performance degradation of some applications	.846	.706	.791
Unanticipated running costs	.900	.790	.454
Unused virtualization features	.846	.706	.791
Increased storage requirements	.997	.454	.407
Management complexity	.983	.594	.435
Risk of total failure	.630	.976	.415
IP address mapping	.997	.454	.407
Incompatible hardware with virtual operating system	.168	.956	.270
Slow applications	.965	.507	.685
Maintenance of the virtual infrastructure	.383	.875	-.120
Need of trained specialists	.997	.454	.407

Extraction Method: Principal Component Analysis. Rotation Method: Promax with Kaiser Normalization.
Source: Research Data

The statistics from the challenges structure matrix above explain the variability of the variables after Promax rotation with Kaiser Normalization. The statistics represents a high correction in reduced factors of 0.74, 0.58 and 0.50 respectively for the reduced virtualization challenges components asserting the findings of table 7 that shows low variability between amongst challenges of virtualization.

Table 10 : Component Correlation Matrix

Component	1	2	3
1	1.000	.446	.472
2	.446	1.000	.347
3	.472	.347	1.000

Extraction Method: Principal Component Analysis. Rotation Method: Promax with Kaiser Normalization.

Source: Research Data

Table 10 above shows the summary of the extracted 3 components of virtualization challenges after Promax rotation with Kaiser Normalization. The statistics show that the components are almost evenly correlated;

Component 1 & 2 at 44.6%

Component 1 & 3 at 47.2 %

Component 2 & 3 at 34.7 %

The three components have a correlation average of 42.2% showing that the challenges of virtualization are highly correlated hence none of them needed special attention for this study.

4.5 Descriptive Statistics

4.5.1 Advantages / benefits of Virtualization

Table 11 : Benefits Descriptive Statistics

	N	Kurtosis							
	Statistic	Statistic	Statistic	Statistic	Statistic	Statistic	Std. Error	Statistic	Std. Error
Reduced Capital costs, resource & space utilization	11	3	5	4.73	.647	-2.420	.661	5.510	1.279
Build up business continuity through improved D&R	11	3	5	4.73	.647	-2.420	.661	5.510	1.279
Real time response to changing business needs	9	3	5	4.67	.707	-2.121	.717	4.000	1.400
Ease of backup and restore	9	3	5	4.67	.707	-2.121	.717	4.000	1.400
Ease in server provision	11	1	5	4.64	1.206	-3.317	.661	11.000	1.279
Dynamic load balancing	10	3	5	4.50	.707	-1.179	.687	.571	1.334
Improved flexibility, High Availability of application and performance	11	3	5	4.45	.688	-.932	.661	.081	1.279
Run multiple Operating systems on one single server	11	3	5	4.45	.820	-1.153	.661	-.254	1.279
Flexibility in OS deployment	11	4	5	4.45	.522	.213	.661	-2.444	1.279
Stability & Security	11	3	5	4.45	.820	-1.153	.661	-.254	1.279
Reducing operational risks	9	4	5	4.44	.527	.271	.717	-2.571	1.400
Ease of OS and application testing	11	2	5	4.18	.982	-1.204	.661	1.136	1.279
Reduced power and cooling	10	3	5	4.10	.876	-.223	.687	-1.734	1.334
Improves system compliance	11	3	5	4.09	.701	-.123	.661	-.453	1.279
Reduced OS licensing costs	11	3	5	4.09	.539	.155	.661	1.862	1.279
Reduced Total Cost of Ownership	9	2	5	4.00	1.000	-.964	.717	.786	1.400
Green Computing that is environment friendly	11	2	5	4.00	1.000	-.733	.661	-.133	1.279
Improved enterprise desktop management & control	11	1	5	3.82	1.079	-1.907	.661	4.964	1.279
Support for IT governance	8	2	4	3.63	.744	-1.951	.752	3.205	1.481
Valid N (listwise)	6								

Source: Research Data

The descriptive statistics of the study in table 11 above shows the statistical ranking for the contributions of the various benefits of virtualization. The ratings for the variables that scored over one include Reduced Capital costs, increased resource & space utilization, Building up business continuity through improved Disaster & Recovery, Real time response to changing business needs, ease of backup and restore and Ease in server provision at 1.206. The study shows that all the Kurtosis values are positive within the range of 0.522 (Flexibility of OS deployment) and 1.206 (Ease in server provision) meaning that the distribution is highly peaked or with leptokurtic distribution.

4.5.2 Disadvantages / challenges of virtualization

Table 12 : Challenges Descriptive Statistics

	N	Std. Deviation		Kurtosis								
	Statistic	Statistic	Statistic	Statistic	Statistic	Std. Error	Statistic	Statistic	Statistic	Std. Error	Statistic	Std. Error
Some application need dedicated resources	6	3	2	5	3.67	.422	1.033	1.067	-.666	.845	.586	1.741
Server Sprawl	5	3	2	5	3.20	.490	1.095	1.200	1.293	.913	2.917	2.000
Licensing of operating systems and applications	6	4	1	5	3.17	.601	1.472	2.167	-.418	.845	-.859	1.741
Need of trained specialists	7	4	1	5	3.14	.634	1.676	2.810	-.309	.794	-1.468	1.587
Maintenance of the virtual infrastructure	7	2	2	4	3.14	.340	.900	.810	-.353	.794	-1.817	1.587
Increased storage requirements	6	4	1	5	3.00	.730	1.789	3.200	.000	.845	-1.875	1.741
Unused virtualization features	6	4	1	5	3.00	.577	1.414	2.000	.000	.845	-.300	1.741
Risk of total failure	6	4	1	5	2.83	.749	1.835	3.367	.362	.845	-2.103	1.741
Unanticipated running costs	6	3	2	5	2.83	.477	1.169	1.367	1.586	.845	2.552	1.741
Need of a robust backup and restore system	5	3	1	4	2.80	.490	1.095	1.200	-1.293	.913	2.917	2.000
Management complexity	6	4	1	5	2.67	.667	1.633	2.667	.383	.845	-1.481	1.741
Slow applications	7	4	1	5	2.57	.571	1.512	2.286	.620	.794	-.809	1.587
Performance degradation of some applications	6	3	1	4	2.33	.494	1.211	1.467	.075	.845	-1.550	1.741
IP address mapping	6	4	1	5	2.33	.615	1.506	2.267	1.270	.845	1.531	1.741
Incompatible hardware with virtual operating system	7	2	1	3	2.29	.286	.756	.571	-.595	.794	-.350	1.587
Valid N (listwise)	4											

Source: research data

The study descriptive statistics rated need for dedicated resources and server sprawl as the main challenges in operating in a virtualized environment. Further, incompatible hardware with virtual operating systems and IP address mapping were among the least fears in their virtual environment. However the statistics still showed that the Kurtosis average of 2.86 with the highest variable value of 3.67 and the lowest of 2.29. This still supports the earlier statistics that showed low variability and high correlation between the challenges of implementing virtualization.

CHAPTER 5: SUMMARY OF FINDINGS AND CONCLUSIONS, RECOMMENDATIONS AND SUGGESTIONS FOR FURTHER RESEARCH

5.0 Summary of findings and Conclusions

The objective of the study was to evaluate the virtualization implementation at UNON to establish the extent to which UN is utilizing it. Further the study was to establish the benefits and challenges of working in a virtualized environment. The study had a sample size of eleven respondents from ICT and Finance departments. The research findings reveal that:

A majority 94% indicated that indicated that the two most important reasons for virtualization are to reduce capital costs, resources, space utilization and to build business continuity through improved disaster and recovery. On the other hand, 40% of respondents indicated that the top three disadvantages of virtualization include that some application need dedicated resources, maintenance of the virtual infrastructure and the need of trained specialists

On average the study recorded an average of 81.4% on all the benefits / advantages of virtualization against an average of 31.8% for the disadvantages of implementation of virtualization technology in an organization. This clearly shows that the benefits associated with virtualization far much out way the challenges or its disadvantages.

Regarding level of virtualization usage at UNON, the study established that they migrated all their applications to a virtual environment. Going it future, they will not invest more in infrastructure but tap on the benefits of the virtualization technology that minimizes on hardware requirements.

As regards the benefits of virtualization, it was resounding from the study that virtualization implementation was a worthwhile investment with benefits being recouped in the long term.

The study also noted the virtual environment comes with its own challenges from its implementation through its running. However, it was apparently clear that the benefits far much outweigh the challenges.

5.1 Limitations of the study

Virtualization is a relatively new technology in the ICT industry and in particular to Kenyan organization that are embracing it now. Further, the technology business concept is yet to be known and appreciated by other departments in the organization. This limited my study largely to ICT department of UNON with only one respondent in finance department.

Secondly, there was limitation of academic material both in text books and previous research regarding virtualization technologies. This reduced the study to rely heavily on online materials that were not always specific to the objectives of the study.

5.2 Recommendations

The study has clearly shown that besides virtualization technology having its own challenges, the benefits of its implementation are immense. It’s therefore recommendable for organizations to analyze their specific requirements and establish if the technology is appropriate for them to embrace.

5.3 Suggestions for further research

One of the reasons that UNON was selected as a research case study is that it has been using the technology for over four years in Kenya. This study wanted to know their experience over time they have been using the technology. This study would recommend the following areas for further research:

Establish the challenges that organizations are encountering while embracing server virtualization technology. Further, studies can be done to establish the driving force in embracing virtualization by both medium and large size organizations in Kenya.

This study also would recommend research to be done in Green Computing or Green IT to establish how organizations in Kenya are tackling environmental issues caused by e-waste.

REFERENCES

AMD (2005) **'AMD64 Virtualization Codenamed "Pacifica" Technology Secure Virtual Machine Architecture Reference Manual'** [online] Available at: http://ci.nii.ac.jp/naid/10020830751/en [Accessed 12 June 2009]

Armbrust M, Armando Fox, Rean Griffith, Anthony D. Joseph, Randy Katz, Andy Konwinski, Gunho Lee, David Patterson, Ariel Rabkin, Ion Stoica, and Matei Zaharia (2009) ***'Above the Clouds: A Berkeley View of Cloud Computing'*** [online] Available at: http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.pdf [Accessed 29 October 2009]

Barham P. (2003) ***'Xen and the art of virtualization'*** [online] Available at http://people.cs.ubc.ca/~norm/508/2007W1/summaries/paper19/paper19.pdf [Accessed 2 May 2009]

Beng-Hong L, Sugerman J, Venkitachalam G. (2001) **'Virtualizing I/O Devices on VMware Workstation's Hosted Virtual Machine Monitor'**, *USENIX Annual Technical Conference*.

Blanchard H. and Xenidis J. (2001) ***'Fast Secure Virtualization for the ARM Platform'*** Available at: http://www.cs.ubc.ca/grads/resources/thesis/May06/Ferstay_Daniel.pdf

Bonwell and Eison, 1991; Sivan et al, 2001 ***'case study theory'*** [online] Available at http://www.cases.bham.ac.uk/casetheory.htm [Accessed 23 October 2009]

Burger and Gartner, (2008) '*Press release'* [online] Available at: http://www.gartner.com/it/page.jsp?id=883312 [Accessed 25 April 2009]

Braastad, (2006) ***'Management of high-availability services using virtualization'*** [online] Available at: http://eternity.iu.hio.no/theses/pdf/master2006/espenbraastad_thesis.pdf [Accessed 7 July 2009]

Bruening F and Ramanathan R. (2004) ***Bringing Flexibility and New Capabilities to Computing Platforms*** [online] Available at: http://download.intel.com/technology/computing/archinnov/teraera/download/Virtualization_0604.pdf [Accessed 3 August 2009]

BSI (2006) '**Business *Continuity Management Code of Practice*'** Available at: http://www.bsigroup.com [Accessed 27 August 2009]

Chen P, Dunlap G, and King S (2003) ***'Operating system support for virtual machines'.*** In USENIX Annual Technical Conference, General Track, pp 71-84

Clark B, Deshane T, Dow E, Evanchik S, Finlayson M, Herne J, and Neefe J Matthews (2004**).'*Xen and the art of repeated research'***. In USENIX Annual Technical Conference, FREENIX Track, pp 135 -144

Datasynapse, (2006) **'*Virtualization increases business continuity'*** Available at: http://www.datasynapse.com/ [Accessed 12 June 2009]

Dragovic B, Fraser K, Hand S, Harris T.,Pratt I, War_eld A, Barham P., and Neugebauer R (2003) **'Xen and the art of virtualization'***. In Proceedings of the ACM Symposium on Operating Systems Principles*

Ferstay D. (2001) **'*Fast Secure Virtualization for the ARM Platform'*** *[*online] Available at: http://citeseerx.ksu.edu.sa/viewdoc/download;jsessionid=C70696475A651C6C65

36D1253E8A8988?doi=10.1.1.88.5596&rep=rep1&type=pdf [Accessed 2 May 2009]

Fritz Nielson (2007) ***Strategies for Delivering Dynamic Security for Virtual Data*** **Centers,** virtualization security playbook, 2007 Available at: http://www.findwhitepapers.com/content1582/ [Accessed 2 September 2009]

Gribble S, Shaw M, and Whitaker W (2002) ***A scalable isolation kernel***, In Proceedings of the Tenth ACM SIGOPS European Workshop, St. Emilion, France

Hamm (2005) '***A Virtual Revolution'*** [online] Available at http://www.businessweek.com/magazine/content/05_25/b3938622.htm [Accessed 2 June 2009]

Highleyman W. (2008) ***'Fault Tolerance for Virtualized Environments'*** [online] available at: http://www.stratus.com/download/index.cfm/whitep/documents/AvailDigest-FTforvirtual3-11-08.pdf [Accessed 23 May 2009]

How Does Virtualization Work? Available at: http://www.vmware.com/technology/virtualization.html [Accessed 2 April 2009]

IDC (2008) ***'benefits of server consolidation'*** [online] Available at: https://partner.microsoft.com/download/global/40079090 [Accessed 7 July 2009]

Info-Tech Research Group (2009) ***"Free Top 5 IT Cost Cutting Strategies"*** [online] Available at: http://whitepapers.itnews.com.au/whitepaper4823/ [Accessed 27 August 2009]

Intel (2005) **“I*ntel Virtualization Technology Specification for the IA-32 Intel Architecture”*** [online] Available at http://andrewl.brainstemprojects.com/library/docs_intel/virtualization_Apr05.pdf

Kommeri J, (2007) ***‘System management in Server Based Computing with virtualization’*** [online] Available at: http://lib.tkk.fi/Dipl/2009/urn012953.pdf [Accessed 25 April 2009]

ITdrilldown (2004) **‘Bank Scores with Server Virtualization’** [online] Available at: http://www.cio.com/article/471401/Bank_Scores_with_Server_Virtualization [Accessed 2 May 2009)

Mugenda M. O. and Mugenda A. (1999), **Research Methods: Qualitative and Quantitative Approaches**, African Centre for Technology Studies, Nairobi, Kenya

Stratus Technologies (2008), ***Best Practices for Server Virtualization in Mission-Critical Healthcare IT*** [online] Available at: http://www.stratus.com/download/index.cfm/pdf/whitepapers/Whitepaper_Healthcare_ServerVirtualization.pdf [Accessed 18 May 2009]

UNON website (2006), ***‘About United Nation office at Nairobi’*** Available at: http://www.unon.org/about.php [Accessed on 5 April 2009]

Rasmussen E (2009) **‘*Reducing IT Costs and Increasing IT Efficiency by Integrating platform-Virtualization in the Enterprise*’** [online] Available at https://scholarsbank.uoregon.edu/xmlui/bitstream/handle/1794/8571/Rasmussen-2009.pdf?sequence=1 [Accessed 25 April 2009]

Saltzer J, Schroeder M. (1972) ***'The Protection of Information in Computer Systems'*** [online] Available at: http://www.cs.virginia.edu/~evans/cs551/saltzer/ [Accessed 1 August 2009]

Samoilenko (2007) ***'Brief Annual Report'*** [online] Available at http://www.nbuv.gov.ua/science/2007_en.pdf [Accessed 23 May 2009]

Second International Workshop on Virtualization Technology in Distributed Computing, **2007** [online] Available at: http://workspace.globus.org/vtdc07/ [Accessed 29 October 2009]

Skooliki (2008) ***Virtualization Industry Challenges*** *- Report 2008* [online] Available at: http://www.virtualization.info/challenges/ [Accessed 12 May 2009]

Strategic Counsel, (2007) ***'Global Server Virtualization Survey - Summary Report'*** [online] Available at: http://ca.com/files/SupportingPieces/consolidatedreport52607v1.pdf [Accessed 25 April 2009]

Techworld (2008) **'*Inside Application Virtualization'*** [online] Available at: http://vmblog.com/archive/2008/07/24/john-sheehan-inside-application-virtualization.aspx [Accessed 4 May 2009]

Thomas Burger (2008) **"*The Advantages of Using Virtualization Technology in the Enterprise*"** [online] Available at: http://software.intel.com/en-us/articles/the-advantages-of-using-virtualization-technology-in-the-enterprise/ [Accessed on 15 April 2009]

'Virtualization Pretty big down under 2008' [online] Available at: http://virtualization.com/guides/2008/07/24/virtualization-australia/ [Accessed 27 May 2009]

‘Virtualization technology’ **(2007)** [online] Available at: http://www.fsmlabs.com/blog/virtualizatio/ [Accessed on 2 April 2009]

Waldspurger C, (2003) “***Memory resource management in vmware esx server*”** [online] Available at: www.waldspurger.org/carl/ [Accessed 13 April 2009]

VMware (2009) *‘**Building an Adaptive Infrastructure with HP and VMware virtualization solutions’*** [online] Available at: http://www.zdnetasia.com/2007/i/pwc/hp_intel/s4/Virtualization_businesscase.pdf [Accessed on 28 October 2009]

Yang (2008) in **‘*Reducing IT Costs and Increasing IT Efficiency by Integrating platform-Virtualization in the Enterprise*’** by Rasmussen E (2009) [online]. Available at https://scholarsbank.uoregon.edu/xmlui/bitstream/handle/1794/8571/Rasmussen-2009.pdf?sequence=1 [Accessed 25 April 2009]

Bruel & Price, (2007) Conoops, (2007) EconomicExperts (n.d) Golden and Scheffy (2007) Goth, (2007) Kirch, (2006) Schwab (2006) TechWeb, (2008) Vanover (2007) Yang (2007) *ALL FOUND IN* **‘*Reducing IT Costs and Increasing IT Efficiency by Integrating platform-Virtualization in the Enterprise*’** by Rasmussen E ([online] Available at : https://scholarsbank.uoregon.edu/xmlui/bitstream/handle/1794/8571/Rasmussen-2009.pdf?sequence=1 [Accessed 25 April 2009]

APPENDICES

Appendix 1: Questionnaire Introduction Letter:

Nkonge Zacheus
School of Business
University of Nairobi
P.O Box 30197
Nairobi.

October 2010

Dear Sir / Madam,

<u>RE: COLLECTION OF RESEARCH DATA</u>

I am a postgraduate student at the University of Nairobi, School of Business. As part of my course work assessment, I am required to submit a management research project. In this regard, an evaluation of Virtualization Technology: The case of United Nation Office in Nairobi.

This is to kindly request you to assist me to collect the data b filling our the accompanying questionnaire. The information you provide will be used exclusively for academic purposes

My supervisor and I assure you that the information you give will be treated with strict confidence and at no time will your name appear in my report. A copy of the final paper will be availed to you upon request.

Your co-operation will be highly appreciated.

Thank you in advance.

Yours faithfully,

NKONGE ZACHEUS	KATE LITONDO
MBA Student	Supervisor & Lecturer

APPENDIX 2: Questionnaire

a. **Key Informant Interview Guide**

Respondent

Names________________________________Title_________________________

Gender [ ] Male [ ] Female Section /Department__________________

Date of Contact______________________ Date Completed__________________

1) Do you believe that it was a good move to implement virtualization at UNON? Yes ( ) No ( ) IF YES, Why do you believe so?
2) Has virtualization helped UNON offer better its services? Yes ( ) No ( ) IF YES, in what ways? How can it compare with your services before virtualization?
3) What are the top most benefits and challenges of operating in a virtualized environment?
4) What are your future plans as regards virtualization at UNON?

b. General Research Questionnaire

Section A: Background information

Names (optional) ____

Gender [ ] Male [ ] Female

Age bracket [ ] over 60 [ ] 51- 60 [ ] 41- 50 [ ] 31- 40 [ ] 30 and below

Please tick your highest level of education;

[ ] PHD

[ ] Postgraduate degree

[ ] Undergraduate degree

[ ] Diploma

[ ] Other (please specify) ______________________________

Do you support UN agencies outside Kenya? [ ] YES [ ] NO. If Yes, how many? ____________

How many UN agencies and programmes do you support within Kenya? ____________

What is the total number of UNON staff in Kenya

[ ] over 200 [ ] 151- 200 [ ] 101- 150 [ ] 51- 100 [ ]. 50 and below

What is the total number of staff of UNON in ICT services department?

[ ] over 40 [ ] 31 – 40 [ ] 21 -30 [ ] 11- 20 [ ]. 10 and below

Section B: Extent of Usage of virtualization at UNON

a) Which of the following virtualization technology have you implemented at UNON (please tick where applicable)

Virtualization Technology	**Version(please specify)**
VMware	
Microsoft Virtualization	
Xen	
TRLinux	
UMLinux	
Denali	

Other (if any)

____

____

b) Do you have trained Virtualised Certified Professionals at UNON?
[ ] YES [ ] NO If Yes, how many? _________________

c) For how long has UNON been operating on a virtualized environment? _______ years

d) How many physical servers did you have;
i) Before virtualization ______________ ii) after virtualization_____________

e) How many virtual servers do you have in your virtual environment? _________

f) What is the total number of applications hosted in your virtual environment? ______________________________

g) What is the total number of Operating Systems used in UNON? ___________

h) Are there some Operating System you use that cannot be virtualized? Yes / No (Please list them)

____

____

____

Section C: Benefits of virtualization to UNON

The following have been attributed as some of the benefits of operating in a virtualized environment as opposed to a non-virtualized environment. Kindly rate each of the benefits as it applies to UNON

<table>
<tr><td rowspan="2">Benefits of virtualization</td><td colspan="5">Please tick the weight you attribute to each of the benefits

5 (highest) … 1 (lowest)</td></tr>
<tr><td>5</td><td>4</td><td>3</td><td>2</td><td>1</td></tr>
<tr><td>Reduced capital costs, resource & space utilization</td><td></td><td></td><td></td><td></td><td></td></tr>
<tr><td>Stability and Security</td><td></td><td></td><td></td><td></td><td></td></tr>
<tr><td>Flexibility in OS deployment</td><td></td><td></td><td></td><td></td><td></td></tr>
<tr><td>Ease in server provision</td><td></td><td></td><td></td><td></td><td></td></tr>
<tr><td>Reduced OS licensing costs</td><td></td><td></td><td></td><td></td><td></td></tr>
<tr><td>Reduced power and cooling</td><td></td><td></td><td></td><td></td><td></td></tr>
<tr><td>Ease of OS and application testing</td><td></td><td></td><td></td><td></td><td></td></tr>
</table>

Benefits of virtualization	Please tick the weight you attribute to each of the benefits *5 (highest)* ... *1 (lowest)*				
	5	**4**	**3**	**2**	**1**
Run multiple Operating systems on one a single server					
Improved flexibility, High Availability of application and performance					
Build up business continuity through improved D&R					
Improved enterprise desktop management & control					
Green Computing that is environment friendly					
Improves system compliance					
Dynamic load balancing					
Ease of backup and restore					
Reducing operational risks					
Support for IT governance					
Real time response to changing business needs					
Reduced Total Cost of Ownership					

List other benefits of operating in a virtualized environment.

____

____

____

____

Section D: Virtualization challenges

The following are identified as the general challenges of operating in a virtualized environment. In order of challenge level, kindly indicate on a scale of 1 (lowest) to 5 (highest), the implication of each of the challenges.

Challenges of virtualization	Please tick the weight you attribute to each of the challenges *5 (highest)* … *1 (lowest)*				
	5	**4**	**3**	**2**	**1**
Server Sprawl					
Some application need dedicated resources					
Licensing of operating systems and applications					
Need of a robust backup and restore system					
Performance degradation of some applications					
Unanticipated running costs					
Unused virtualization features					
Increased storage requirements					
Management complexity					
Risk of total failure					
IP address mapping					
Incompatible hardware with virtual operating system					
Slow applications					
Maintenance of the virtual					

Challenges of virtualization	Please tick the weight you attribute to each of the challenges *5 (highest)* *1 (lowest)*				
	5	**4**	**3**	**2**	**1**
infrastructure					
Need of trained specialists					

List other challenges of operating in a virtualized environment.

____

____

____

____

THANK YOU FOR YOUR TIME

GLOSSARY OF TERMS

The glossary contains specialized terms used in the document. Definitions are designed to describe concepts, environments, terms and things that relate to virtualization and virtualization environments, as these concepts are utilized in this study. Terms are mined from literature reviews, initial concept terms, and key word searches (Rasmussen, 2009)

Business drivers: The factors in the industry or the broader business environment that either impact the financial institution or provide opportunity for business expansion (Bruel & Price, 2007).

Service level agreement (SLA) is a contract between the provider and the user that specifies the level of service expected during its term. SLAs are used by vendors and customers as well as internally by IT shops and their end users. (PCMAG, n.d)

Business functions (Business Process): Series of logically related activities or tasks (such as planning, production, sales) performed together to produce a defined set of results (Business Dictionary, n.d.).

Data center: Computer facility designed for continuous use by several users, and well equipped with hardware, software, peripherals, power conditioning and backup, communication equipment, security systems, etc. (Business Dictionary, n.d.).

Emulation: The binary translation of one instruction set into another, whether from x86 instructions to PowerPC instructions, or from x86 instructions to virtualized x86 instructions (Microsoft, 2003).

Enterprise: A business organization. (Collins Essential English Dictionary 2nd Edition, 2006).

Hardware virtualization: Partitioning the computer's memory into separate and isolated "virtual machines" simulates multiple machines within one physical computer. It enables multiple copies of the same or different operating systems to run in the computer and also prevents applications from interfering with each other (Computer Language Company, n.d.).

Host operating system: A virtual machine monitor or host called a hypervisor to enable multiple operating system instances to run on a single physical server (Goth, 2007).

Guest operating system: The guest operating system running on a layer above the hypervisor. It can also run within an operating system, with the guest OS running on the third layer above the hardware (Goth, 2007).

Hypervisor: A thin layer of software that generally provides virtual partitioning capabilities that runs directly on hardware, but underneath higher-level virtualization services (VMware, 2006).

Information Technologies (IT) Infrastructure: Set of tools, processes, and methodologies (such as coding/programming, data communications, data conversion, storage and retrieval, systems analysis and design, systems control) and associated equipment employed to collect, process, and present information (Business Directory, n.d.).

Legacy application: An older application that is often no longer supported (Microsoft, 2003).

Legacy operating system: An older operating system, often incompatible with up-to-date hardware. Virtual machines allow you to run legacy operating systems on new hardware (Microsoft, 2003).

Legacy Systems: Existing computer system that provides a strategic function for a specific part of a business (BNET, n.d.).

Migration: A change from one hardware or software technology to another. Moving data from one storage system to another (data migration). Moving data and applications from one computer to another (TechWeb, 2008).

Non-virtual environment: All of the resources on the physical computer are permanently dedicated to the applications running on that computer (Kirch, 2006).

Paravirtualization: The name for another approach to server virtualization. In this approach, rather than emulate a complete hardware environment, the virtualization software is a thin layer that multiplexes access by guest operating systems to the underlying physical machine resources (Golden & Scheffy, 2007).

P2V: Term for physical-to-virtual when discussing environment migration (Vanover, 2007).

Partitioning: The ability to run multiple operating systems on a single physical system and share the underlying hardware resources (VMware, 2006).

Platform Virtualization: A technique used to replicate the functionality of real hardware platform so that one physical machine can host more than one system level software without conflict (Yang, 2007).

Return on Investment (ROI): The monetary benefits derived from having spent money on developing or revising a system (Computer Language Company, n.d.).

Sandbox: A restricted environment in which certain functions are prohibited (Computer Language Company, n.d.).

Server Virtualization: Server virtualization is where the base hardware of a system is virtualized, allowing multiple guest operating environments to run directly on top of the hardware, without requiring a complete host Operating System (Schwab, 2006).

SMB: Small to Medium-sized Business also called "SME" (small to medium-sized enterprise) typically refers to companies with 25 to 500 employees; however, some SMB/SME ranges use an upper limit of 5,000 employees (Computer Language Company, n.d.).

Streaming: A subset of other virtualization technologies, which provides a way for software components, including applications, desktops, and complete Operating Systems, to be dynamically delivered from a central location to the end-user over the network (Schwab, 2006).

Total Cost of Ownership (TCO): Includes the cost of the hardware, software and upgrades as well as the cost of the in-house staff and/or consultants that provide training and technical support (Computer Language Company, n.d.).

Virtual Machine Monitor (VMM): Software that runs in a layer between a hypervisor or host operating system and one or more virtual machines that provides the virtual machine abstraction to the guest operating systems. With full virtualization, the virtual machine monitor exports a virtual machine abstraction identical to a physical machine, so that standard operating systems (e.g., Windows 2000, Windows Server 2003, Linux, etc.) can run just as they would on physical hardware (VMware, 2008).

X86 virtualization: Refers to the Intel x86 family of CPU chips which are the most commercially successful instruction set architecture in the history of personal computing, which includes the Core, Pentium and preceding models such as the 486 and 386. x86-based computers, including machines with x86-compatible CPUs from AMD, Cyrix, VIA, Intel and others, make up the world's largest hardware platform (EconomicExpert, n.d.). As the x86 term became common after the introduction of the 80386, it usually implies a binary compatibility with the 32-bit instruction set of the 80386. This may sometimes be emphasized as x86-32 to distinguish it either from the original 16-bit x86-16 or from the newer 64-bit x86-64 (also called x64). Although most x86 processors used in new personal computers and servers have 64-bit capabilities, to avoid compatibility problems with older computers or systems, the term x86-64 is often used to denote 64-bit software, with the term x86 implying only 32-bit.

Green computing or **green IT**, refers to environmentally sustainable computing or IT. It is "the study and practice of designing, manufacturing, using, and disposing of computers, servers, and associated subsystems such as monitors, printers, storage devices, and networking and communications systems efficiently and effectively with minimal or no impact on the environment. Green IT also strives to achieve economic viability and improved system performance and use, while abiding by our social and ethical responsibilities. 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.

Operating System (OS) porting refers to design of a device or application to ensure it integrates well with the hosting OS thereby utilizing the resources availed to it.

Principal component analysis (PCA) involves a mathematical procedure that transforms a number of possibly correlated variables into a smaller number of uncorrelated variables called principal components. The first principal component accounts for as much of the variability in the data as possible, and each succeeding component accounts for as much of the remaining variability as possible.