

Virtualization in Linux a Key Component for Cloud Computing

Harrison Carranza^a and Aparicio Carranza^a

^aComputer Engineering Technology
New York City College of Technology of The City University of New York
186 Jay Street (V633) Brooklyn, NY 11201
Phone: (718) 260 – 5897 Email: acarranza@citytech.cuny.edu

ABSTRACT

Cloud Computing and Virtualization are becoming very popular and many businesses are embracing this paradigm. Cloud computing enables on-demand access to network services as well as configuration of Computer Systems. Virtualization allows physical machines to be partitioned into several Virtual Machines (VMs); and run applications in each of these VMs as if it were independent physical machines all running at the same time. VMware, Xen and KVM are well known virtualization solutions that provide means to create virtual machines; there after allows us to install different operating systems. VMware is a commercial product that carries a license fee; Xen and KVM are freely available and are included with most of the Linux distributions. In this paper we mainly report our work carried out by using KVM (Kernel Virtual Modules) and briefly mention the high performance solution that VMware server system provides when used; we revise the requirements to implement this approach: Memory (RAM), Storage (hard disk), CPU selection (Intel-VT or AMD-V), etc. Lastly, we create several Virtual Machines and install various Operating Systems to observe their performance.

Keywords: Cloud Computing, Virtualization, Physical Machines, Virtual Machines, VMware, Xen, Linux, KVM, SAN.

I. INTRODUCTION

During the last decade, virtualization and cloud computing have become two driving forces in the engineering and information technology industry. Although both have been in play for more than 40 years by IBM (International Business Machines), these technologies have been promoted to the public as a result of the growing storage and server utilization [7].

Virtualization refers to the abstraction of a physical machine [1], [8]. This means that the information stored inside the system will be compressed as if it were a single device. Cloud Computing is the provision of dynamically scalable and often virtualized resources as a service over the Internet on a utility basis [10], [11]. Cloud computing and virtualization will allow access to the many servers existing in a data center. Since data centers have machines running at only 10 to 15 percent of their computing power, companies are trying to make the other 85 to 90 percent of that power to be put in use rather than have all those resources wasted. [8]

Cloud computing and virtualization are interrelated because they both involve the abstraction of a physical machine. When somebody refers to the “cloud,” they are talking about the abstract network that exists within the Internet and allows access to any information that has been created in a certain data center. Even though we may get access to this data, most of the time we do not know from where we are getting this information, which is not important as long as we are able to access it. [9]

In this paper, we depict in great detail about virtualization which is the driving force to implement either a Private or Commercial Cloud Computing. Particularly we present detailed steps to create virtual machines using the KVM (Kernel Virtual Module). In section II we discuss “How Virtualization Works”, section III details “Virtualization Requirements”, section IV shows “Steps to Create a Virtual Machine in Linux”, section V presents “Comparison of Some Virtualization Technologies”. In the last section (VI) our Conclusion is presented and an appendix follows this last section which shows some sample screen captures.

II. HOW VIRTUALIZATION WORKS

Virtualization and cloud computing are two paradigms that work together in the field of Engineering and Information Technology. Virtualization has been defined in many ways, from virtual people to virtual worlds, and to actually making full use of any machine. This term had been confused many times by companies but in simple terms, “virtualization stands for abstraction” [1]. This means all information inside of a system(s) would be compressed to look as if it were a single device. Although this topic has been in existence for many years, it has become more popular within the last decade because of the many improvements being made by large companies.

Cloud computing and virtualization, specifically storage virtualization functions as Redundant Array of Independent Disks (RAIDs), caching, instant copies, and remote mirroring. This means the objectives of virtualization are for improvement of availability, performance, scalability, and maintainability. The use of virtualization as a part of Storage Area Networks (SANs) is to the user’s advantage, in which the capacity of all the available resources are carefully inspected and controlled from the central part of the network. This means if any malfunction occurred within the network, the administrator can go back into the system to fix the problem. [6]

Storage virtualization has been designed to convert storage data from its physical deployment and abstractly configure it [5]. This shows that there would be more space within a physical machine to actually store important information and create multiple machines, referred to as virtual machines, or VMs, which can provide different running applications all at the same time without having to use other computers that will consume power and space. In order to perform virtualization to provide us with everything that we want it to do, we must get a hypervisor, which is an application that must be installed on the hard drive of the computer so we can later create our virtual machines. It is similar to installing the operating system on the computer and then putting all the software programs that we want to use on top of it.

Using any of the three hypervisors: VMware, Xen, or KVM - the command line Interface (CLI) in Linux is preferably used by system administrators to access the servers instead of running them with a GUI (Graphic User Interface) because that would take up resources and risk any security issues for the network [3]. During virtualization, the system administrator will implement several methods to keep the network running efficiently. Although there are many of them that exist and are being used, some of the most fundamental and beneficial ones are the following: building VM hosts, following security rules for physical machine, storage for SANs, duplicate production network with virtual network, installation of virtualization tools, private networks for server-to-server communications, and guest-to-host time synchronization. [4]

III. VIRTUALIZATION REQUIREMENTS

In order to perform virtualization, there are several requirements needed to run efficiently. Although many distributions of Linux allow virtualization in the operating system, there are two specific Linux versions that can make life easier: Fedora and CentOS. It should be considered the extra demands that are put on the system by running guest virtual machines - Particularly, it should be considered if the computer has enough resources to run the virtual guests, check that all the packages needed are installed, and make sure that the necessary services are running. The Minimum requirements for making virtual machines is at least 512 MB of RAM, create virtual guests with 8GB disk space, Virtualization enabled CPUs (*i.e. Intel-VT or AMD-V*). [2]

IV. STEPS TO CREATE A VIRTUAL MACHINE

The System used for these steps was an HP Notebook with 4GB RAM, 500GB hard disk space and Intel CPU (i5-450M: 2.4GHz). For creating a virtual machine on Fedora or CentOS, you must select the Applications bar and then enter the Virtual Machine Manager. The window, as shown in **Figure 1**, should come up on the screen.

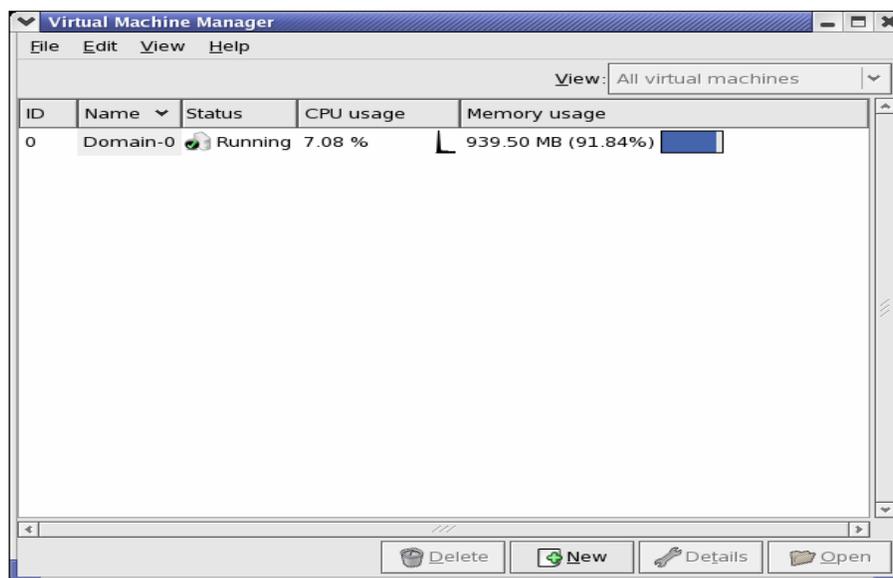


Figure 1: Virtual Machine Manager

When the Virtual Machine Manager window is open on the screen, you select **File**, then **New Machine** to create a virtual machine on the physical computer, as it is seen on **Figure 2**.

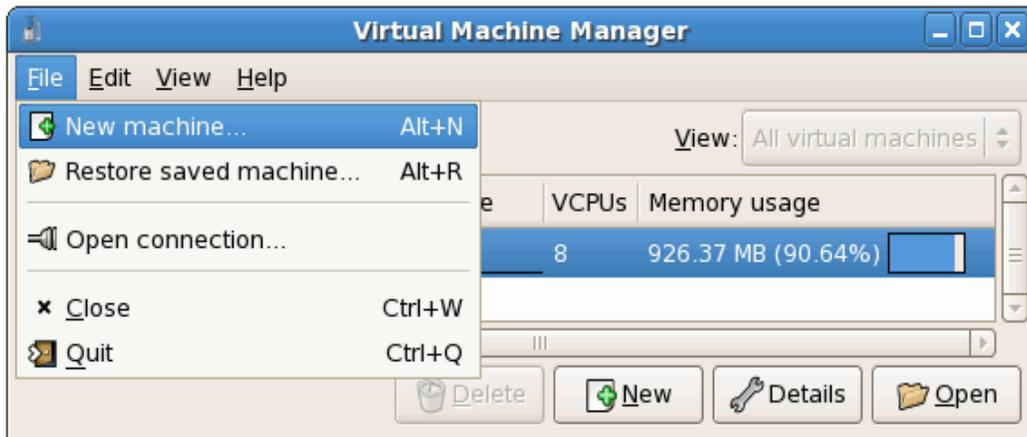


Figure 2: *Creating a Virtual Machine*

In the next window the option to create a new virtual machine is displayed with either a fully-virtualized or para-virtualized system. The location media, storage, memory, and CPU also have to be taken into consideration for efficiency and appropriate performance of the virtual machine. These options are shown in **Figure 3**.



Figure 3: *Creating a new Virtual System*

In our case we chose the fully-virtualized option as it is shown in **Figure 4**.

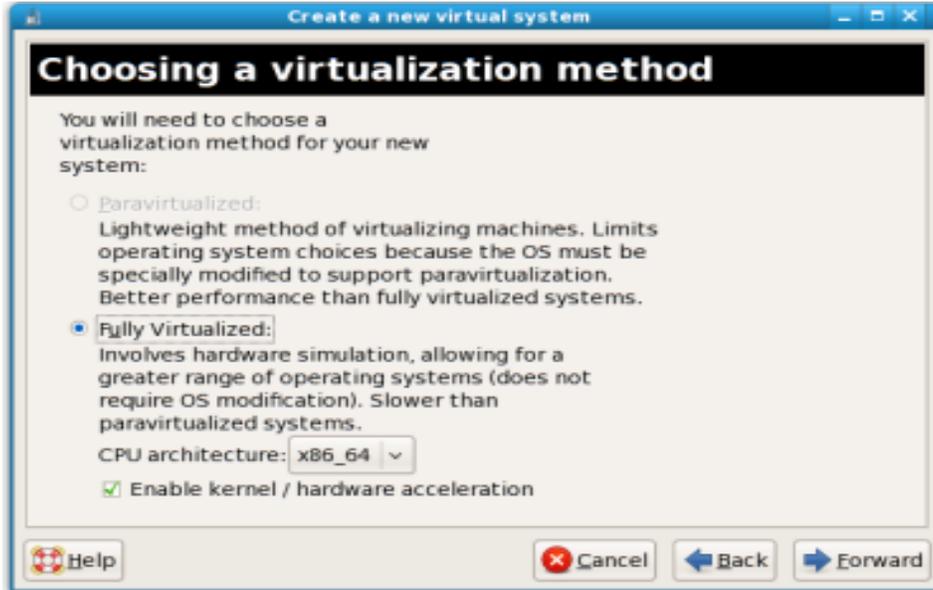


Figure 4: Naming a Virtual System

Figure 5, shows the media location from where the guest operating systems will be installed into the virtual machine.



Figure 5: Locating Installation Media

Another step is to assign storage space for the virtual machine. The user has the option of selecting either normal disk partition or simple file. The simple file option will allow selecting a file location with a certain size for the virtual machine to occupy. This step is shown in **Figure 6**.



Figure 6: Storage Allocation for Virtual Machines in the Physical System

Figure 7, shows the maximum memory on the host machine, the maximum memory and the maximum startup Memory of the Virtual Machine. It also shows the number of virtual CPUs this Virtual Machine will startup with.

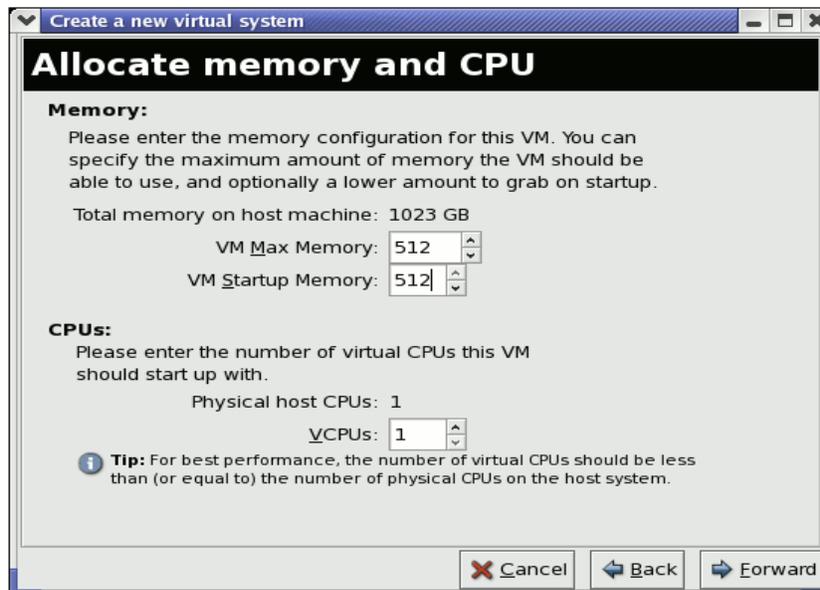


Figure 7: Allocation of Memory and Number of Virtual CPUs to a Virtual Machine

When the creation of the virtual machine is complete, the user will be given the option of installing the operating system of their choice. A window similar to **Figure 8**; should appear. In this instance we chose Fedora 8 to test the functionality of our Virtual Machine.



Figure 8: *Installation of Fedora 8 in the Virtual Machine*

V. COMPARISON OF SOME VIRTUALIZATION TECHNOLOGIES

During our research, we came across several virtualization technologies that we will describe next. At the Computer Engineering Technology Department of the New York City College of Technology, we made use of the VMware ESX Server system installed in 4 Dell Power Edge Servers and 1 HP Proliant Server configured as HA (High Availability) connected via a HBA (Host Bus Adapter) to the Storage Area Network (SAN) with 8TB (Tera-Bytes) of capacity for performing VMotion and Disaster Recovery for Business Continuity solution. The VMware technology rendered a very robust solution with a high performance results. We were able to create over 40 Virtual Machines and installed various versions of the Microsoft Operating Systems (Desktop and Server); and also different Linux distributions - the result was almost no performance degradation. One note to mention here is that VMware is a commercial product that carries an annual license cost little bit high for regular individual user. Next, at home we used the Xen Virtualization System. This system is freely available for desktop and servers. Xen is used by small to medium size companies with a moderate budget. Although it is free; it does not provide the same features as VMware. We played with Xen by installing it into HP laptop with 512MB RAM, 60GB hard disk space; and an AMD 2.3 GHz processor. We encountered some difficulties in trying to work with this system, for this reason we tried the next option. Lastly, we dealt with KVM (Kernel Virtual Module). This solution is also available for free and continuously being in development. Every time that the Linux Kernel gets updated the KVM also gets updated. The procedure for this approach has already been indicated somewhere in this report.

VI. CONCLUSION

Virtualization and cloud computing are integral technologies for the modern data center. Virtualization is the first step for implementing cloud computing; so is the Storage Area Networks (SANs). We have described these technologies and experimented with various Virtualization approaches such as: VMware, Xen and KVM. We presented the detailed steps of virtualization using the KVM approach and also indicated that VMware renders a high performance solution but it is an expensive solution as well. Since Storage Area Networks are the driving technologies of any Data Center and consequently a key component of Cloud Computing, our future plans are to explore SANs and Clustering with Linux.

APPENDIX

The following images are some screenshots captured from two different systems showing the final installation of the operating systems on the virtual machines.

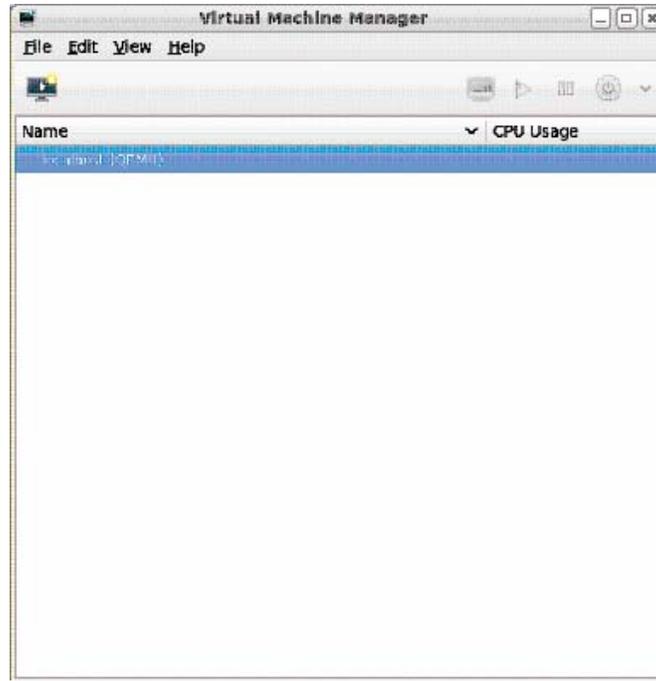


Figure A1: Window showing Virtual Machine Manager

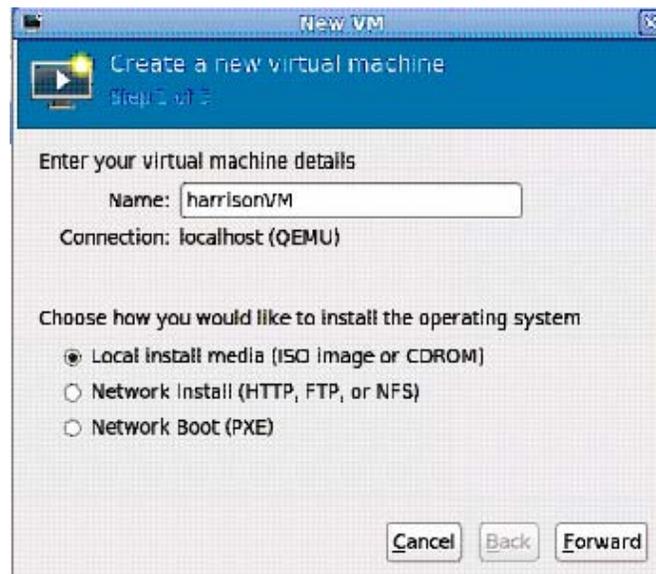


Figure A2: Creating a Virtual Machine, assigning a name, option for installation media

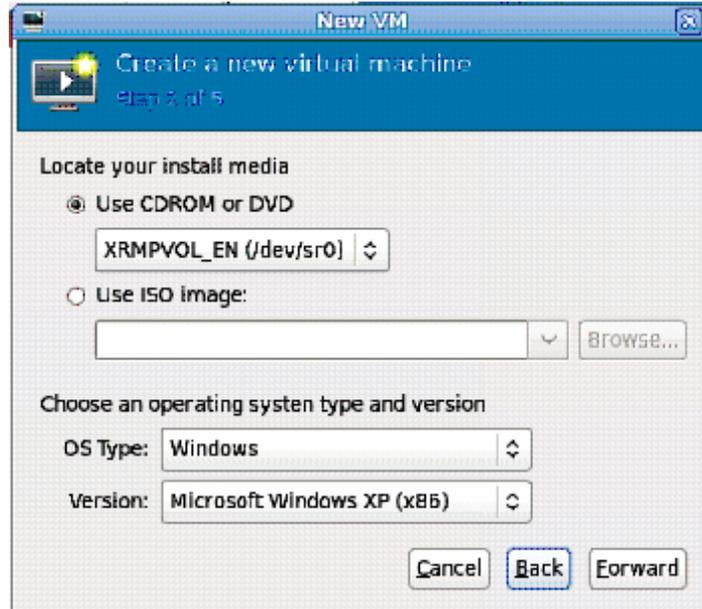


Figure A3: Locating installation media, (Windows XP – Architecture x86)

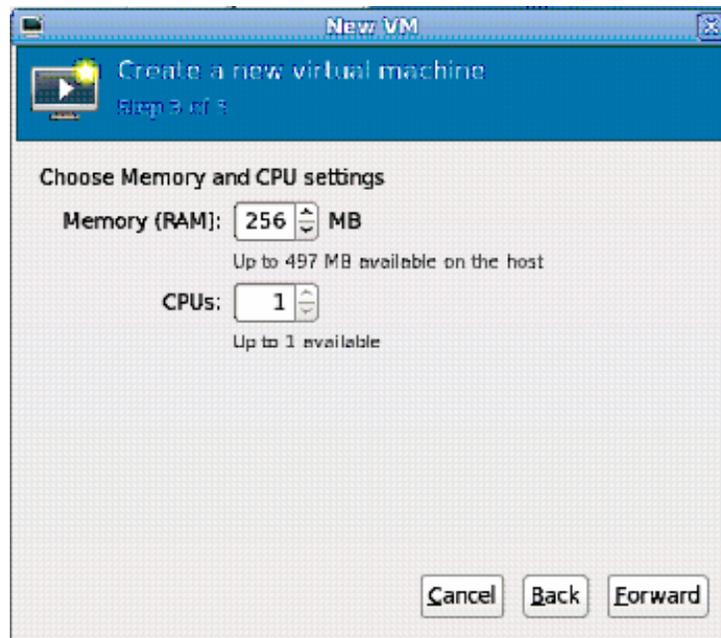


Figure A4: Option of Choosing RAM and number of CPUs

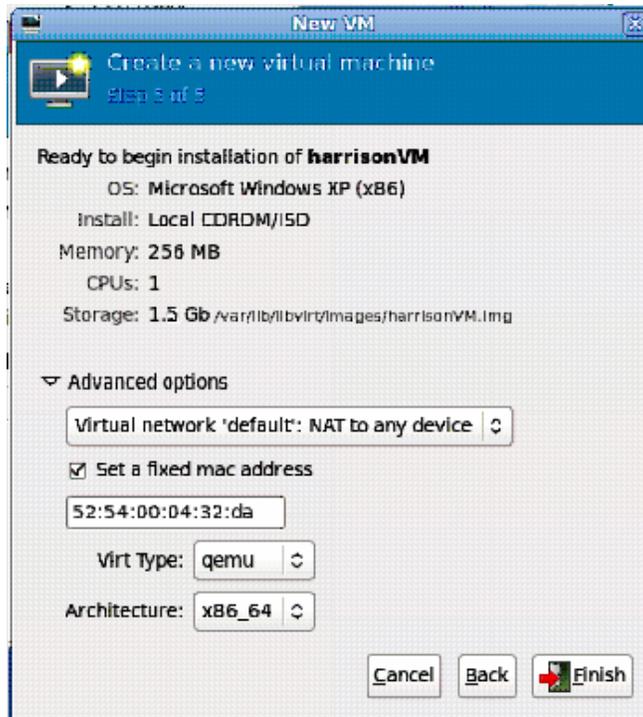


Figure A5: Final Step, Recap of Virtual Machine Details

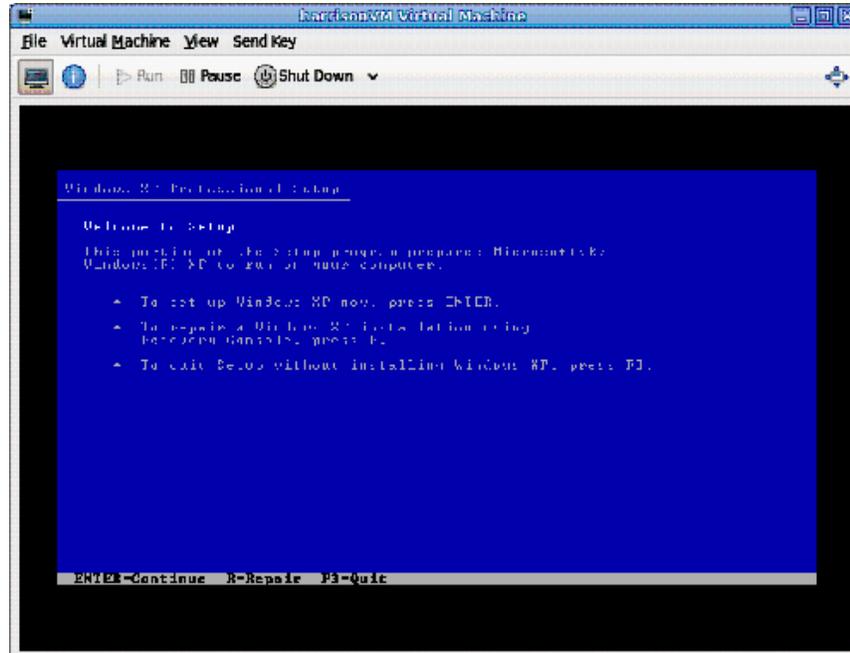


Figure A6: Installation of Windows XP inside the Virtual Machine

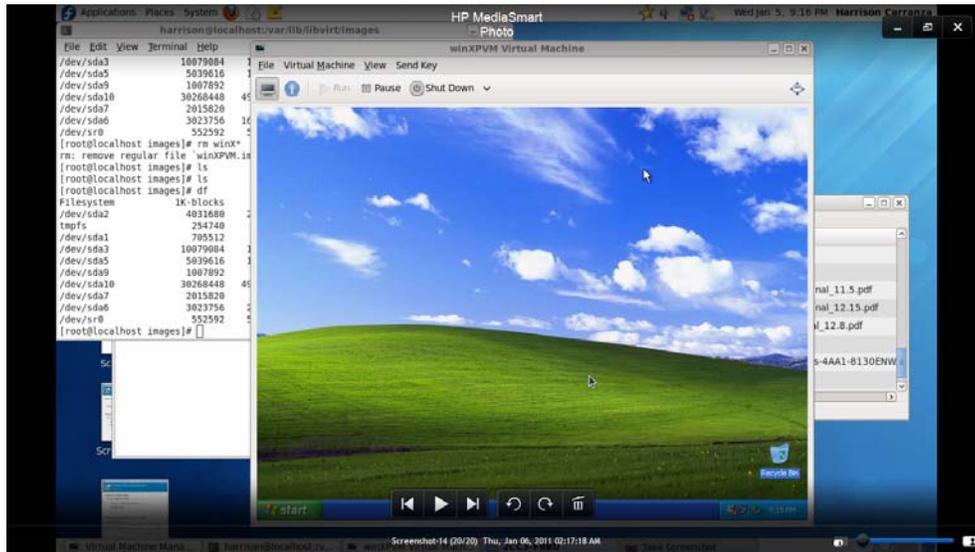


Figure A7: Using HP System, Virtual Machine with Windows XP on the Desktop

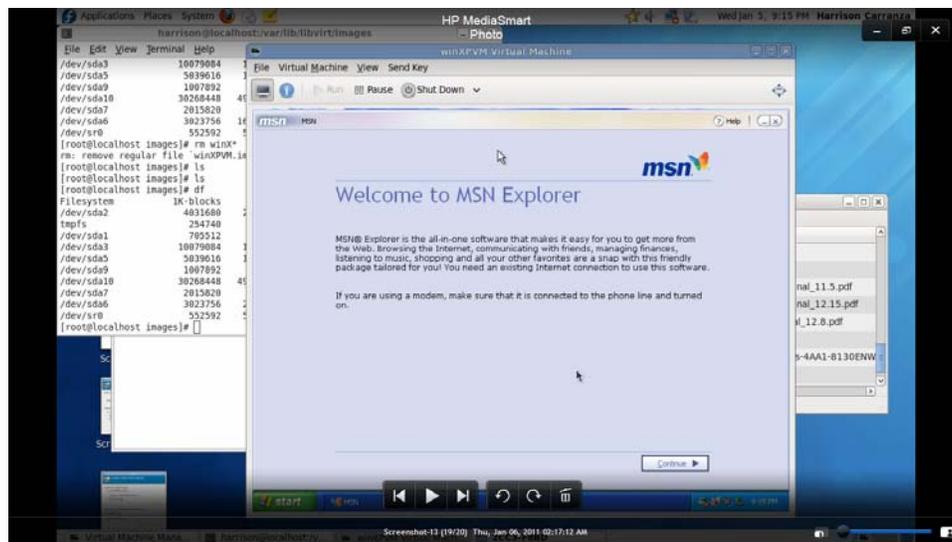


Figure A8: Using HP System, Virtual Machine with MSN Explorer Initiating



Figure A9: Using Dell System, Virtual Machine with Windows XP on the Desktop

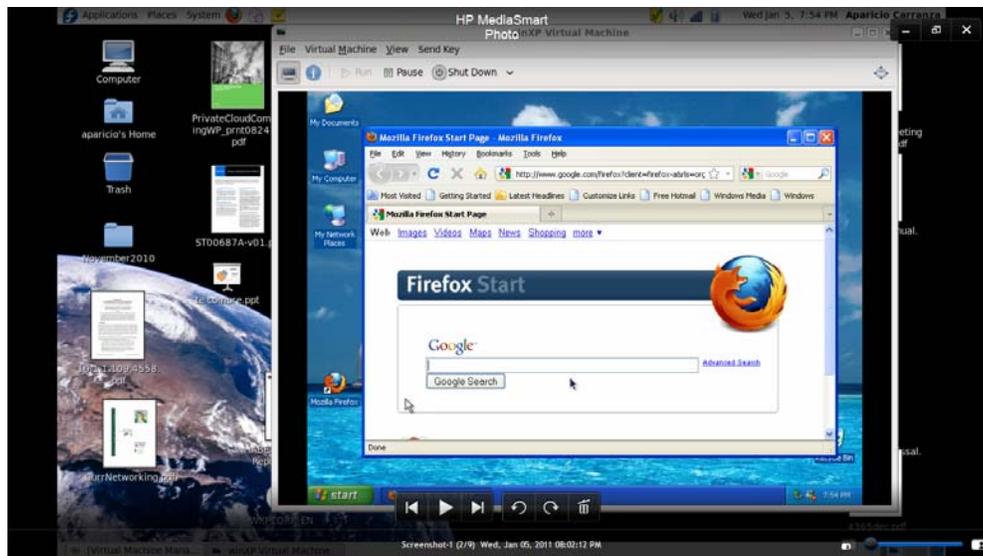


Figure A10: Using Dell System, Virtual Machine with Mozilla Firefox Initiating

REFERENCES

- [1] C. Poelker and A. Nikitin, Storage Area Networks for Dummies, Wiley: Hoboken, New Jersey, 2009.
- [2] C. Negus and E. Foster Johnson, Fedora 2010 Edition, Wiley: Indianapolis, Indiana, 2010.
- [3] G. Bledsoe, Linux Journal, October 2010. www.linuxjournal.com
- [4] R. Barnes, Linux User and Developer, 2010. www.linuxuser.co.uk
- [5] T. Clark, Designing Storage Area Networks, 2nd Edition, Addison Wesley: Boston, Massachusetts: 2003.
- [6] W. Muller, U. Troppens, R. Erkens. Storage Networks Explained, Wiley: West Sussex, England: 2004.
- [7] “White Paper: Ten Things You Need to Know About Virtualization.” www.datacore.com
- [8] B. Golden. Virtualization for Dummies, Wiley: Hoboken, New Jersey: 2008.
- [9] J. Engates. “The 101 Guide to the Cloud,” 2009.
- [10] “Gartner says Cloud Computing will be as Influential as E-business.” www.gartner.com
- [11] G. Gruman. “What Cloud Computing Really Means.” Infoworld. www.infoworld.com