ISSN 2347-4289

Virtualized Data Center With Server Consolidation In Vmware ESX

Aniruddha P. Tekade, Varsha¬priya J.N.

Department of Computer Technology, VJTI, Mumbai.
tekadeaniruddha@gmail.com, varshapriyajn@vjti.org.in

Abstract: Today Data Centers are critical for corporal operations, monitoring networks and keeping control over all business intelligence activities. Businesses thrive over the operations ofdata center today. Data Center systems are like building or a portion of building whose primary function is to house a computer system and its support areas. This in-turn also saves the cost estimates of the organization plus it makes high-end system easily available virtually which are not just as good as but definitely more and more efficient than the so-called physical machine.Whenever you hear talking about efficient IT-utilization in business environments, you will also hear about virtualization and its advantages over the traditional IT-models. If you want to keep the company's IT-infrastructure up to date, you have to know about utilizing virtualization methods and implementations. Today many organizations find themselves in a situation where they have their data centers filled with underutilized hardware resources. This situation accumulates from the trend of distributed computing and client/server models starting from the late 1980s. As distributed systems and servers were cheap, at the time these solutions provided fast, cost-effective and easy to deploy IT systems. Over time these IT systems became hard to manage, complicated and distributed underutilized systems, resulting in an inefficient use of floor space and excessive power consumption in the data center. Virtualization and more specifically physical to virtual server consolidation provides a solution to these so called "server sprawls". (IBM 2005, 2). This paper goal will concentrate on virtualization usingserver consolidation using a bare-metal hypervisor platform VMware ESXi 4.0/4.1/5.0. Thesis work will introduce virtualization using different server virtualization platforms, both commercial and non-commercial, mainly concentrating on VMware Infrastructure 3 products. In the practical part a single physical production server is consolidated into a VMware ESXi 4.0 virtual environment. Our basic objective is to help all those Poor Educational Institutions, NGOs, Orphanage Schools who actually run a zero-profit "businesses" with a noble intention and who usually have less budget on computer hardware but still desperately in need of a good Computer Center to educate their students with as equal computing experience as the wealthy institutions usually have. This paper is based on the practical implementation of the Virtual Data Center for VJTI Mumbai for educational purposes. The snaps included are taken by snipping tools to demonstrate this live.

Index Terms: Architecture, Consolidation, Bash, Shell, HA, DRS, vMotion,instances, VDC, Shell Script, Bash

1. INTRODUCTION

DATA centerdesign always need to behigh performing, scalable and dependable for any kind of business.A data center has become acritical issue because of the increasing use of the Internet insupporting variousWeb-based services. It has been observed that data centers contributeto approximately 40% of the overall delay, and this delay is likely to grow with the increasing use of dynamic Web contents. Poor response time has significant financial implications for e-commerce applications. User dissatisfaction due to longer response times is becoming a major concern , as reported by a survey on the problems with WWW. Although several hardware and software scale-up techniqueshave been proposed to enhance the performanceof single node servers, theyarenot viable/long- term solutions considering the fact that the number of online usersis growing at about 90% annually. In view of this, cluster-based data centersare envisioned to answer most of these requirements in a cost eﬀective manner is a facility used to housecomputer systems and is associated with the components, such as telecommunications & storage systems. Cloud computing is a method of delivering computingservices from a large, highly virtualizeddata center to many independent end users,using shared applications and pooled resources.While there are many different definitions for cloud computing, it is typically distinguishedby the following attributes: on-demand self-service,broad network access, resource pooling,rapid and elastic resource provisioning, andmetered service at various quality levels. Implementation of these attributes as part ofa large enterprise-class cloud computing servicethat provides continuous availability to a largenumber of users typically requires significantlymore server, networking, and storage resourcesthan conventional data centers. This is onlyachievable through extensive use of virtualization.This is achievable only through extensive use of virtualization.While server virtualization has existed sincethe 1960s, when it was first implemented onIBM mainframes, it has only become widelyavailable on affordable commodity x86 serverswithin the last decade or so.The bottom line is that it all stands only on strong foundation of virtualization. The term virtualization is a broad concept and can have many different meanings in many different contexts. Virtualization technologies are deployed in a wide variety of ways among different manufacturers. The basic reason we are going to consider VMware technology is- VMware provides us a wizard to build the infrastructures and other components and a wizard is always good in developer's perspective plus it also keeps everything a standard.The following concept will define virtualization and its history while keeping the main emphasis on server virtualization.

a) **Definition of Virtualization -** In a traditional computing model software and hardware are coupled tightly together. **Virtualization** provides computing resources as pools, thus decoupling hardware and software. Within the limits of this paper work, **VIRTUALIZATION** can be defined as a technology that adds a layer of abstraction between hardware and the software running on top of the hardware. In other words the layer of abstraction separates the layers of the IT-system from each other.

Copyright © 2014 IJTEEE.

ISSN 2347-4289

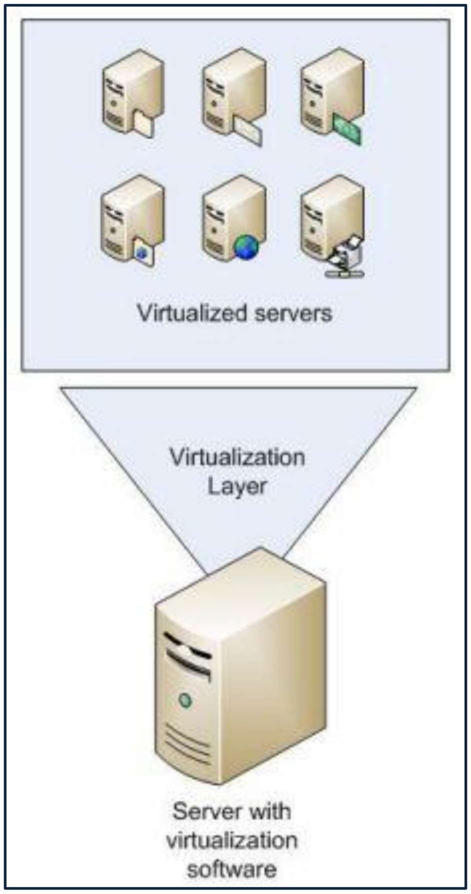

Figure1.1 – Server Virtualization (Rule & Dinner 2007, 19)

Figure 1.1 illustrates server virtualization where one powerful server hosts several independent virtual machines, in this case server operating systems. A virtual machine is a guest operating system that shares the underlying physical hardware with other virtual machines running on top of the same physical host. Each virtual machine has its own operating system and virtual hardware provided by the host, so each virtual machine is independent and completely separated from other virtual machines.In the figure above a single powerful instance represents multiple independent less powerful instances. The users of these servers don't need any knowledge about the virtualization infrastructure below, all they see are the virtualized servers. Figure 1.2 shows an example of storage virtualization where multiple physical disks are combined into a one large logical disk.

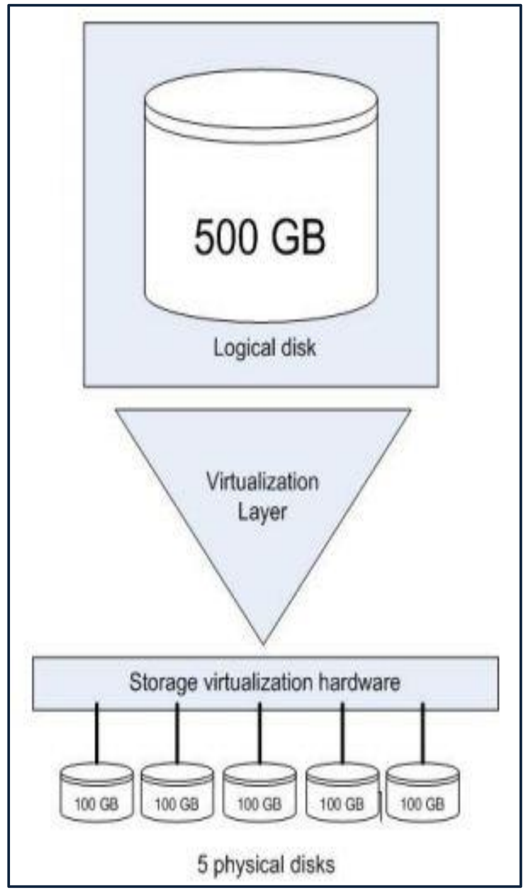

Figure1.2 – Storage Virtualization (Morgan 2006, 3)

In the example above five 100 GB physical disks represent a 500 GB logical disk. This kind of storage virtualization is called "Just a Bunch Of Disks" or JBOD. The basic concepts of storage virtualization are explained in the upcoming concept of Approaches to the Virtualization.

a) History - The roots of present day virtualization go as far as the late 1950s. A term called "Time Sharing", which is considered the driving force behind virtualization, was first introduced in the early 1960s. This technology for dividing one machine's resources to many users was developed because dedicating a machine to a single user was not practical due to extremely expensive technology. During the 1960s several systems which utilized time sharing were introduced, such as System/360, CP-40 and TSS/360. These systems also introduced virtual memory and virtual machines for the first time. (Marshall et al. 2009, 3-4). In 1972 IBM introduced System/370 Advanced Function mainframe computer and the first IBM Virtual Machine Operating System: VM/370. During the 1970s virtual machine technology became popular in the IBM community and in 1973 the MVMUA (Metropolitan Virtual Machine User Association) was founded. The year after this a set of formal requirements for virtualization architectures were created by Gerald J. Popek and Robert P. Goldberg. Following this virtualization had a quiet period which lasted from 1974 to 1987. (Marshall et al. 2009, 4). During the 1980s and 1990s the rise of inexpensive mini-computers and personal desktop

Copyright © 2014 IJTEEE.

ISSN 2347-4289

machines drove virtualization into hibernation. Mainframe model which shares resources centrally was superseded by low cost distributed computing systems. The x86 -architecture was established as the new industry standard. (Marshall et al. 2009, 4). In time the demand and growth of x86 -architecture introduced new challenges to IT-infrastructure and virtualization was awaken from hibernation, introducing the need for x86 server virtualization. In 1997 Connectix Virtual PC 1.0 was introduced. Virtual PC 1.0 was used as a translator of between a virtual Intel x86 processor and a physical Power PC processor used in the Mac. This example of emulation led Connectix into the world of virtualization technologies. (Marshall et al. 2009, 5). In 1999 VMware responded to server virtualization challenge by releasing VMware Virtual Platform, the first commercial x86 virtualization platform. In the late 2000 VMware released the first platform which was designed for x86 server virtualization, VMware GSX Server. In 2001 VMware already took server virtualization to the next level by releasing VMware ESX Server 1.0 which was installed on bare-metal as GSX Server installs on top of existing operating system. Thanks to its native hypervisor VMware ESX Server 1.0 provided great stability and high performance. From that day on to present day VMware has released several updates for both server products. (Marshall et al. 2009, 5). Connectix was acquired by Microsoft in 2003 and thus its x86 server virtualization arena, Connectix Virtual Server, never made it to the market. As Microsoft had acquired property rights for Connectix Virtual PC and Connectix Virtual Server, Microsoft first released Micro-soft Virtual PC 2004 in late-2003 followed by Microsoft Virtual Server 2005 in mid-2004. (Marshall et al. 2009, 5). The focus being on VMware and Microsoft, at the same time an open source virtualization product called Xen was being actively developed. The Xen project included truly notable contributors such as AMD, HP, IBM, Intel, Novell, RedHat and XenSource. The first version of Xen was released in 2003 but a product designed to compete directly with VMware ESX Server was released in 2006 by XenSource. The product was called XenEnterprise 3.0. In Au-gust 2007 XenSource released XenEnterprise v4 which became a closer competitor to the VMware ESX Server. During the same month XenSource was acquired by Citrix. (Marshall et al. 2009, 6). With products of VMware and Citrix being the dominant virtualization platforms in the market, Microsoft had to come up with something to keep up with the pace. With the development of Windows Longhorn (Windows Vista) and Windows Server 2008, Microsoft also started to develop its own hypervisor: Microsoft Hyper-V, which is shipped with Microsoft Server 2008, was released in mid-2008. (Marshall et al. 2009, 6-7).

b) **Virtualization Approaches -**Virtualization can take many forms in present day IT-world. Four main forms of virtualization are server, storage, network and application virtualization. Each form represents an abstraction from physical environments in different manner. Although virtualization is commonly considered as a way of partitioning multiple resources into a single hardware entity, like in server virtualization, it can also do just the opposite: representing multiple hardware instances as a single computing resource. This form of virtualization is often utilized in storage, network and application virtualization approaches. Server virtualization is the most common approach and generally when the term "virtualization" is discussed, indeed people are referring to server virtualization. In the following chapters the primary discussion will be on bare-metal full server virtualization, other main forms of virtualization are discussed briefly to understand some basic concepts. (Rule & Dittner 2007, 21).

c.1. Server Virtualization - Server virtualization means logical partitioning of single server hardware into multiple virtual entities or platforms. Single server hardware with virtualization software can run multiple independent operating systems enabling physical server consolidation.

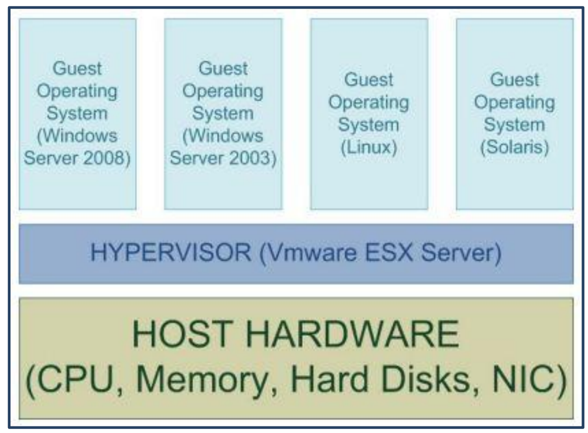

Figure1.3 –Bare Metal Hypervisor
(Martial et al. 2009, 10)

In server virtualization the layer of abstraction between hardware and software is called the Virtual Machine Monitor (VMM), or hypervisor. It is the center of server virtualization as it makes it possible for multiple guest operating systems, which are called Virtual Machines, to run on a single host. A good example about this type of VMM is VMware ESX Server which will be discussed later in more detail. There are two types of VMMs: Type 1 VMM, also known as bare-metal hypervisor, is software that is running directly on the host hardware as illustrated in Figure 1.3.

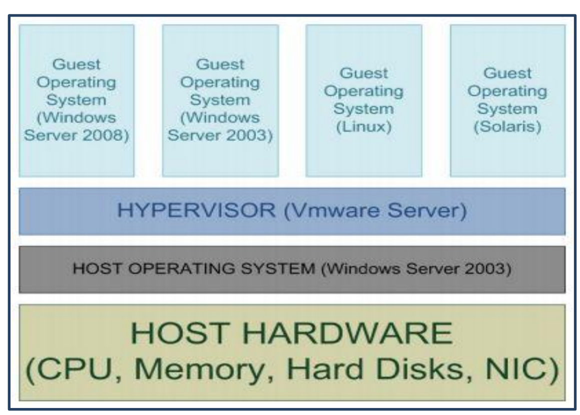

Figure1.4 – Hosted Hypervisor
(Martial et al. 2009, 9)

Copyright © 2014 IJTEEE.

Type 2 VMM, also known as hosted hypervisor, software runs on top of a host operating sys-tem adding one more layer to the model, thus the guest operating systems operate at the third layer above the hardware as Figure 1.4 shows. Examples of software that utilizes the hosted architecture are VMware Server and Microsoft Virtual Server (Rule & Dittner 2007, 18-19).Server Virtualization is divided into four categories: **Full virtualization, paravirtualization,** the operating system virtualization and hardware-assisted virtualization.

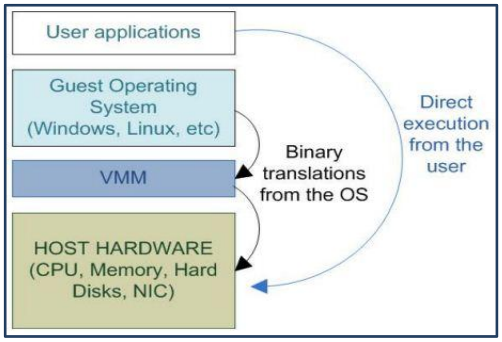

Figure1.5 – Full Virtualization (VMware 2007a,4)

In the case of full virtualization the underlying hardware is fully abstracted into a virtual guest operating system.Full virtualization is provided by the combination of binary translation and direct execution as shown in Figure 1.5. Modification to the guest operating system kernel is not required and the guest is not even aware that it is running in a virtual environment. Unlike other server virtualization technologies, full virtualization does not require any hardware or operating system assistance, the virtual machines operatejust as they were physical systems and they can execute all the same operations as they could on raw hardware.Full virtualization provides complete isolation of each virtual machine with no modifications to the kernel, thus having the widest range of support of guest operating systems. Full virtualization is typically used for server consolidation and testing environments. Examples of full virtualization systems are VMware ESX Server and Microsoft Hyper-V. (Rule & Dittner 2007, 21-22). (VM-ware 2007a, 4).

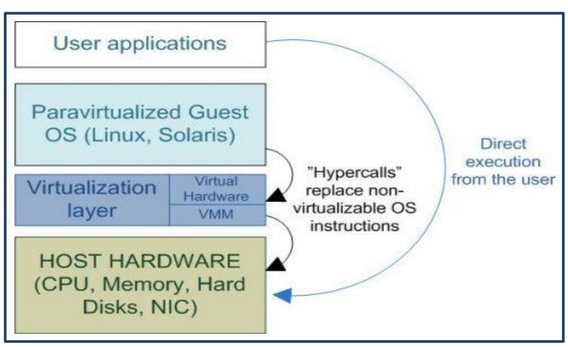

Figure1.6 – Paravirtualization (VMware 2007a, 5)

Paravirtualization provides guest operating systems with software interface that resembles the underlying hardware. As the hardware is only partially simulated, paravirtualization involves modifying the kernel of the operating system to replace instructions that were not virtualized. These replaced instructions are called hypercalls and they provide communication between critical host kernel operations and the guest operating systems as shown in Figure 1.6. Paravirtualization provides low virtualization overhead and easy implementation but the lack of sup-port for closed source operating systems, such as Windows Server 2003 or Vista, makes paravirtualization incompatible and poorly portable. Typically paravirtualization can be used for developing environments and server consolidation. Examples of paravirtualization products are XenServer and IBM z/VM. (Rule & Dittner 2007, 21-22). (VMware 2007a, 5).

c.2. Storage Virtualization

In storage virtualization several physical storages are multiplexed into a single logical storage. Storage virtualization in its most basic form is used with RAID implementations, where two or more physical hard drives are combined to a single logical hard disk to provide data redundancy. Bigger scale implementations of storage virtualization are Storage Area Network (SAN) and Network Attached Storage (NAS) technologies. They are distributed storage net-works which appear to be a single physical instance, but actually are a groupof several physical interconnected disks. With the help of storage virtualization workload of a single hardware can be uniformly distributed to multiple ones. Storage virtualization also provides fault tolerance as data can be easily replicated to multiple physical locations throughout the network.

c.3. Network Virtualization - Network virtualization is used either to separate multiple virtual networks from a single physical network, or to provide networks within virtual environments without using any physical networking devices. Most common approaches to network virtualization are Virtual Local Area Network (VLAN), Virtual IP, Virtual Private Network (VPN) and virtual networking inside a virtualization host. (Rule & Dittner 2007, 25). Virtual Local Area Networking divides a single physical network into several independent logical networks called VLANs. Each VLAN has its own identification tag and is located on its own network segment. This way the VLANs can be securely separated from each other even if they use the same physical network switch. VLAN technology provides good traffic flow management, security and easy network administration. VLANs can also be used in virtual networks within VMware Infrastructure 3 hosts.Virtual IP is an IP-address that is not connected to a specific network interface. It is assigned to several physical instances on the network with a single IP- and MAC-address. Traffic to the Virtual IP address is redirected to the physical instances. Virtual IP technology provides fault tolerance and load-balancing for example in virtualized storage environments. VPN is used to securely connect to a specific network, such as the office local area network, over a public network. VPN provides a secure channel from site to another with the help of tunneling protocols such as IP Security. When VPN connection is established, resources of the remote network are available just as if they would be available when connected to the physical network. This way a user can experience the internal resources of the office network anywhere in the world, as long as there is an Internet connection available.

Copyright © 2014 IJTEEE.

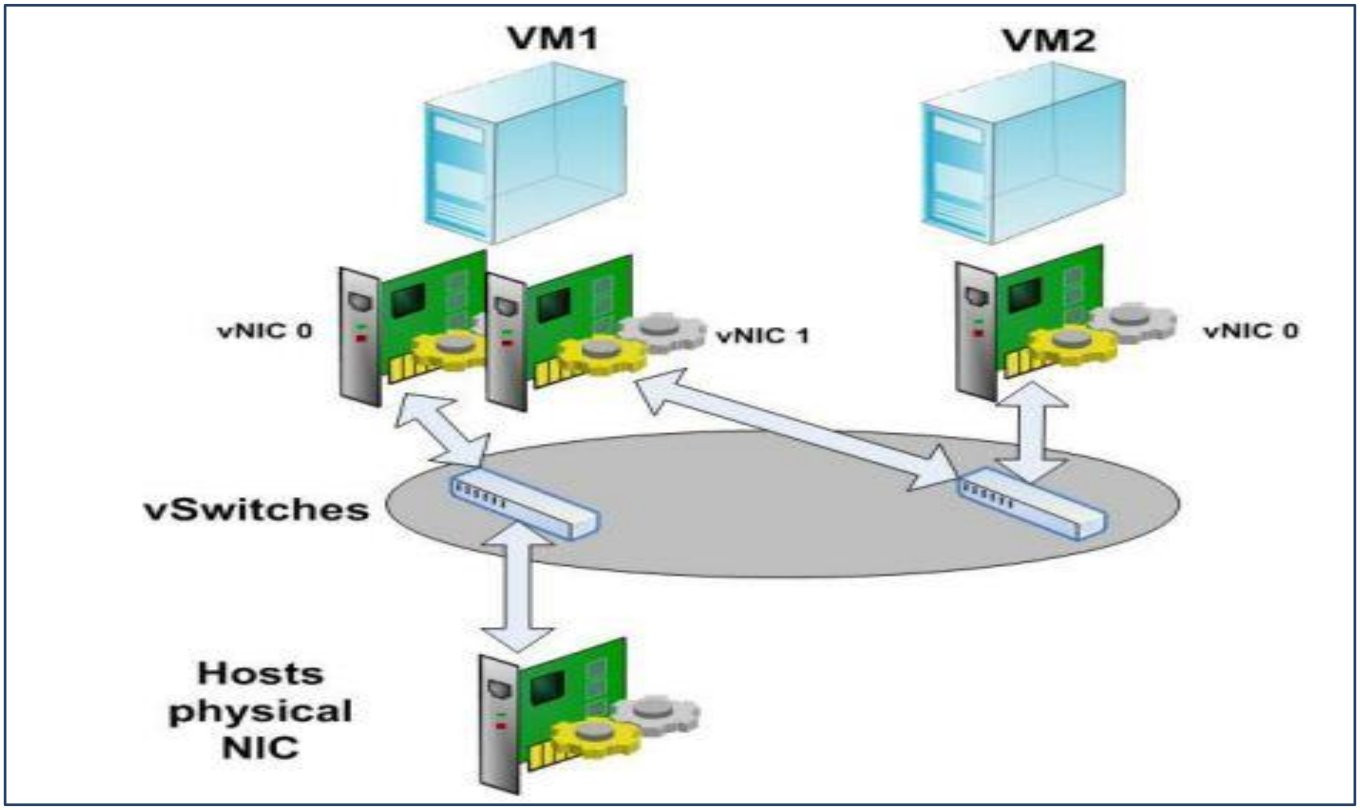

Figure1.7- Virtual Networking

In the example of Figure 1.7, VM1 has two NICs, vNIC 0 is connected to the hosts physical NIC and vNIC 1 is connected to a virtual network inside the host. VM2 has only one vNIC and it is connected only to the virtual network. This way VM1 has connection to both physical and virtual networks, while VM2 only has connectivity inside the virtual network.

c.4. Application Virtualization -- Application virtualization provides a layer of abstraction between specific applications and the host operating system. Virtual software packages allow applications to execute normally without being installed in a traditional sense. Each virtual software package have their own application-specific copy of computers resources and each application can be packaged, stored and activated on demand independently from each other. Application virtualization eliminates conflicts between applications and enables running several versions of the same application on the same host. Performance will be improved as there is no need for separate virtual machine for each application. Application virtualization also enables the applications to run in a non-native environment for example Windows applications can be run in Linux environments. Examples of application virtualization software are Windows Application Virtualization and VMware ThinApp.

c.5. Virtualize the Data Center – Business data centers provide critical IT-services to the company. A common data center of a company is equipped with server, storage and network hardware. Usually companies rely on their data centers heavily, so it is very important that they provide business continuity, fast disaster recovery and high security of data. Without virtualization technologies the data centers are typically inefficient, underutilized, unnecessary expensive and complex to administer. With the help of virtualization all of the problems above can be solved cost-effectively. Server consolidation is an effective way to dramatically increase efficiency and server utilization. In consolidation a physical server instance and its workload is moved into a virtualized environment where several server instances can run on a single hardware. Consolidation offers several benefits such as lower total cost of ownership, increased hardware utilization, easier administration of resources and removal of legacy hardware. Data center reliability and security are key concepts in any business. Traditionally large in-vestments are made into IT-systems in order to keep critical applications running, improving business continuity. Virtualization technologies provide reliability, fault tolerance and fast disaster recovery by virtually partitioning and replicating the IT services, thus reducing total cost of ownership and risk of data loss. In traditional computing model when a partition of an IT-system is compromised, the physical host is also compromised. Isolating the compromised physical instance is often considered very difficult if not impossible. With virtualization a compromised partition can securely be isolated from the rest of the IT-system, preventing the compromise of other partitions. In a virtualized environment security configurations can be specific to a single partition rather than to the whole IT-system or a physical servers. As each partition can be configured with different credentials, system access can be effectively isolated and overall system security is in-creased.

Copyright © 2014 IJTEEE.

ISSN 2347-4289

I. SERVER CONSOLIDATION AND DATA CENTERS

A.Server Consolidation is a process where two or more servers are merged together into a single server. Server consolidation is implemented to solve the problem of a server sprawl. There are four types server consolidation; centralization, physical consolidation, data integration and application integration. In context with virtualization, server consolidation is described as a process of physical consolidation where several physical servers are being transformed into a virtual environment running on top of single server hardware, as illustrated in Figure 2.1.

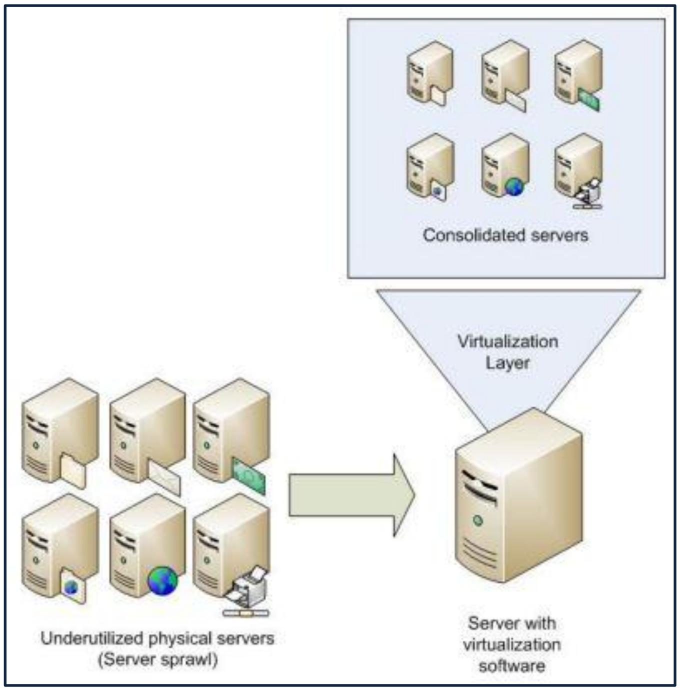

Figure2.1 – Physical to Virtual Server Consolidation (IBM 2005, 9)

After the server consolidation process the virtual infrastructure provides exactly the same services to the users as the physical servers did before the consolidation. A typical example of a virtual infra-structure is a physical server with a VMware ESX Server hosting several independent virtual machines that were physical servers before the consolidation process.

B.Benefits - Server consolidation provides an efficient way to eliminate the problem of a server sprawl. When physical servers are consolidated into the virtual environment, amount of physical hardware is reduced and the remaining hardware will host more than one servers. This offers several significant benefits over the traditional computing model with single purpose server hardware.The most important benefit of consolidation is money. Implementing server consolidation a business can reduce total cost of ownership. Hardware costs are reduced as the need for physical servers has dropped along with the need of cabling for those devices. Server management costs are decreased as consolidation offers efficient and consistent system administration centrally, thus IT staff can be reduced or assigned to other tasks. Also operational costs are reduced as after the implementation there will be less physical servers consuming floor space and power. According to VMware.com, with the help of consolidation, a business can reduce hardware and operating costs by 50% and energy costs by 80%.Traditionally server utilization is only around 10 to 30 percent. This means that most of the high performance provided by expensive server hardware is not taken advantage of. With consolidation server usage is brought up to a higher level and IT resource optimization can be achieved.Consolidation of servers provides a single point of control that has several advantages. Single point of control allows improved service levels, better availability of services, improved systems management, advanced version control management and increased security.

C.Decision of Consolidation - The key to successfully consolidate server workloads is to know which services are to be consolidated and which are not. The most important issue when determining the workloads suitability for consolidation is the CPU and I/O utilization rates. A server with low utilization rates is a better candidate for consolidation than a server with utilization rates close to 100 percent. In a typical file and print server, average CPU and I/O utilization rates are below 15 percent while occasionally rates may spike at 100 percent for a short period of time. File and print server workloads are excellent candidates for server consolidation as their resources are unused most of the time. Combining several standalone server workloads of this type into a single server workload effectively increases server utilization while overall performance is not degraded.High CPU and I/O resource utilization servers, such as database or backup servers, are not very suitable candidates for server consolidation. Database servers tend to have very high CPU, disk and network utilization, therefore there is also no reason or need to consolidate database servers as the overall utilization of the server is typically already close to 100%.

II. VMWARE ARCHITECTURE

The very basic reason we are considering a proprietary virtualization product is because a proprietary product works in a wizard – the set of rules which is always good from a designer’s perspective. Additionally it also resolves the problem of standardization which is very much needful to have a uniqueness throughout the connected network. Currently VMware Infrastructure is a widely deployed virtualization solution on the market. This virtualization product designed for companies and enterprises was built and architected to provide comprehensive virtualization, management, resource optimization, application availability, and operational automation capabilities in an integrated virtualization platform. VMware Infrastructure 3 customers can choose between two hypervisors; VMware ESX Server 3.5 or VMware ESXi 3.5. The bare-metal architecture of both hypervisors provide CPU virtualization, advanced memory management, operating system virtualization, application virtualization, storage virtualization and network virtualization. This section will discuss architecture of the hypervisors as well as features of different editions of VMware Infrastructure, from free standalone version of VMware ESXi 4.0 to enterprise-level VMware Infrastructure Enterprise. The basic architecture and features of all editions are more or less the same, the only architectural difference being the almost non-existent Service Console of the VMware ESXi 4.0. From this

Copyright © 2014 IJTEEE.

ISSN 2347-4289

point on, VMware ESX Server 4.0 will be referred to as ESX and VMware ESXi 4.0 will be referred to as ESXi.

Component **(Feature)**	ESXi – free License **(ESX not available without VI)**	VI Foundation **(with ESX or ESXi)**	VI Standard **(with ESX or ESXi)**	VI Enterprise **(with ESX orESXi)**
Core hypervisor functionality	Yes	Yes	Yes	Yes
Virtual SMP	Yes	Yes	Yes	Yes
VMFS	Yes	Yes	Yes	Yes
Virtual Center Agent		Yes	Yes	Yes
Update Manager		Yes	Yes	Yes
Consolidated Backup		Yes	Yes	Yes
High Availability			Yes	Yes
vMotion				Yes
Storage VMotion				Yes
DRS				Yes
DPM				Yes

Table 3.1 – VMware Infrastructure Editions (VMware Recent Website& kb articles)

1. Hypervisor Architecture - The ESX has three main architectural components that provide the virtualization layer for the virtual infrastructure: Physical host server, VMkernel and Service Consoleas shown in Figure 3.1. It is important to state that the architecture of the ESXi is different as the Service Consoleis highly deprecated. The othercomponents are the same as in the architecture of the ESX.

a. **Physical Host server -** Physical host server is a server computer that provides all the physical hardware components, mainly CPU, memory, storage devices and network interface cards.
b. **VMkernel -** VMkernel, often called as the operating system of ESX and ESXi, is the hypervisor that operates the host hardware and provides the virtualization layer between the host hardware and the virtual machines.
c. **Service Console -** Service Console or Console Operating System is a modified installation of Red Hat Enterprisethat provides administrators a set of management tools for the virtual infrastructure.

2. Features - Four different editions are available: Free edition with ESXi, VI Foundation, VI Standard and VI Enterprise.In addition to these features VMware Tools are a very important set of features that boost virtual machine performance. VMware Tools are a set of applications, services and drivers that allow smooth communications between the host and the guest operating system. VMware Tools should be installed to every virtual machine to provide the best available guest operating system performance.**ESX** and **ESXi** are the core of the product suite providing the core hypervisor functionality for the virtual infrastructure. Either of the available **core hypervisors** can be chosen for any edition with the exception that with a free license the only available hypervisor is the ESXi. The ESX hypervisor requires at least the VI Foundation edition.**Virtual Symmetric Multi-Processing** (Virtual SMP) enhances virtual machine performance by enabling a single virtual machine to use multiple physical processors simultaneously.

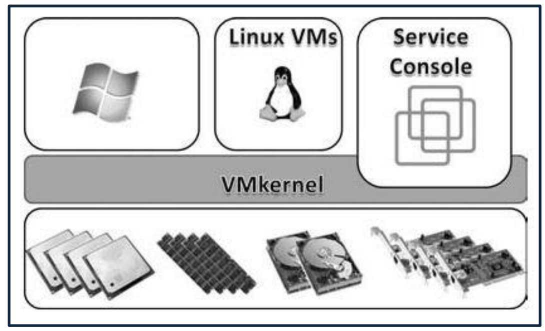

Figure3.1-Architecture of ESX Server

VMFS is file systems for VMware virtualization. VMFS performs I/O via VMDK to the actual disk of the system. **VirtualCenter** is a centralized management utility. It database stores all the data about the hosts and virtual machines and allows administrators to deploy, manage, monitor, automate and secure the virtual infrastructure. Connection to hosts is provided by the **Virtual-Center Agent**. VirtualCenter and its agents are bundled with VI Foundation, VI Standard and VI Enterprise.**VMware Consolidated Backup** (VCB) is a Windows application that provides a Fiber Channel or iSCSI-based backup solution that offloads the backup processing to a dedicated physical server. VCB can take a snapshot of a running virtual machine and mount it into the file system of the VCB server.**High Availability** is a feature that will power on any virtual machines that were previously running on another VMware Infrastructure host which has experienced a failure and is unavailable. This process is fully automated and provides a cost-effective failover solution for the virtual machines.**VMotion** provides live migration of virtual machines from one VI host to another. The migration is done with zero down time and continuous virtual ma-chine availability.**Storage VMotion** allows running virtual machines to be moved between datastores. This feature ensures that outgrowing datastores or moving to a new SAN does not force an outage for the virtual machines. Distributed Resource Scheduling (DRS) will continuously monitor VI hosts to maintain balanced utilization across the load on all servers within the VI by using VMotion technology to migrate virtual machines from one host to another.When the work load increases, **Distributed power management**(DPM) brings the hosts in standby back online to ensure service availability.

3. Virtual Networking - VMware Infrastructure brings a rich set of virtual networking elements. The main elements are virtual networks, vSwitches, vNICs and port groups.A virtual network is a network of virtual machines that are logically

Copyright © 2014 IJTEEE.

ISSN 2347-4289

connected to together so that they can send and receive data from each other. Each virtual network is serviced by a single vSwitch. A vSwitch works much like a physical Ethernet switch with physical ports. It detects which virtual machine is connected to which of its virtual ports and forwards traffic to the correct virtual machine. A vSwitch can connect virtual networks to physical networks using one or more of hosts physical NICs. If no physical NICs are associated with a vSwitch, all the traffic in that virtual network is restricted to as host only traffic.A maximum number of **127 vSwitches** can be created on a single VI host.Each virtual machine can have one or multiple vNICs. The guest operating system and applications talk to the vNIC through a standard device driver or a VMware-optimized device driver just as though the vNIC was a physical NIC. When building up a virtual network vNICs are logically connected to vSwitches and assigned one or multiple IP addresses. A vNIC also responds to the standard Ethernet protocol exactly as a physical NIC would, therefore an outside agent does not know whether it's communicating with a virtual or a physical machine. VM-ware Infrastructure 3 has six different vNICs. Four of the vNICs are used by virtual machines, one is used by the Service Console and one is for the VMkernel. The four virtual machine vNICs are:

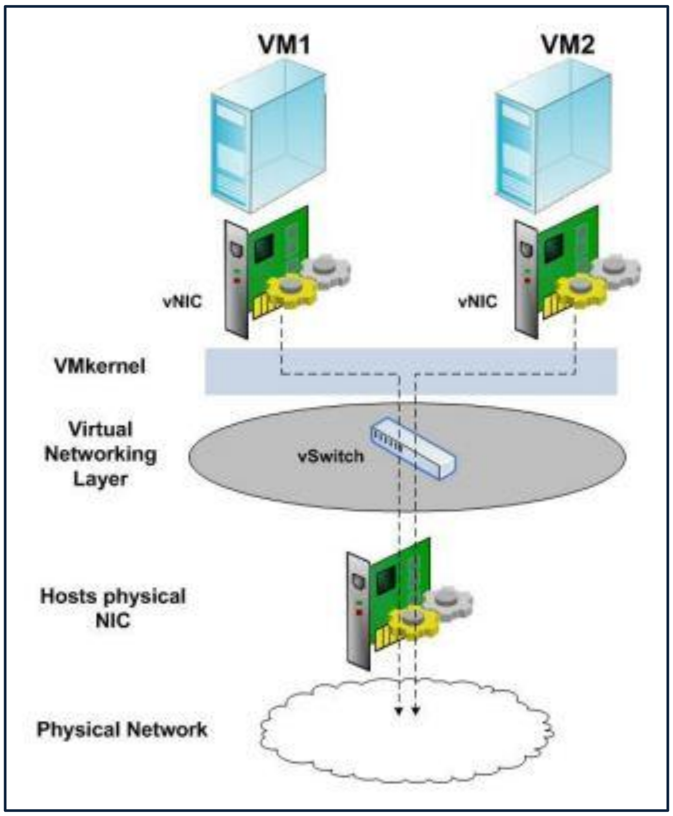

Figure3.2-Virtual Networking

- **VMXNET** – a paravirtualized vNIC that works only with VMware Tools in-stalled in the guest operating system. This vNIC is designed for high performance and is recommended if available for use.
- **Vlance** – provides strict emulation of the AMD Lance PCNet32 Ethernet adapter, which is an older 10 Mbps NIC with drivers available for most 32bit guest operating systems. This vNIC is usually used if the VMware Tools are not installed in the guest operating system.
- **Flexible** – Identifies itself as a Vlance vNIC if VMware Tools are not installed. With VMware Tools installed, the vNIC is changed to higher performance VMXNET adapter.
- **E1000** – provides strict emulation of the Intel E1000 Ethernet adapter. This vNIC is used by default in 64-bit guest operating systems, but is also available for 32-bit guests. Typically also used with Windows Vista and newer Windows guest operating systems.

The other two vNICs are:

- **vswif** – a paravirtualized vNIC similar to vmxnet that is used only by the Service Console of the ESX.
- **vmknic** – vNIC of the VMkernel, the software layer that manages most of the physical resources on the VMware Infrastructure host. Services clients VMotion, Network Files System (NFS) and software iSCSI. Also serves the remote console traffic.

In a virtual environment a port group is a mechanism for setting policies that controls the net-work connected to it. A single vSwitch can have multiple port groups and the maximum number of port groups on a single Host is **512**. Instead of connecting to a specific port on a vSwitch, a virtual machine connects to **port group** assigned to its vNIC. All virtual machines that have their vNICs assigned to the same port group belong to the same network inside the virtual environment even if they are running on top of different physical servers.Port groups can be configured to enforce a number of policies that provide networking security, network segmentation, better performance, higher availability and traffic management. The most common uses of port groups within a Host are:

- VLAN – provides logical groupings of switch ports with the help of VLAN tagging, allowing traffic as if all ports were on the same physical LAN segment.
- NIC Teaming – provides a way to associate multiple physical NICs to a single vSwitch to form a team. A team can either share the traffic load between the physical NICs or provide passive failover in the event of a hardware failure.
- VMkernel TCP/IP stack – the VMkernel TCP/IP networking stack provides network connectivity for the vmknic and supports iSCSI, NFS and VMotion. VMkernel TCP/IP stack has its own port group called VMkernel port group.

4. VMware ESXi 4.0 is a hypervisor that operates independently from any general-purpose operating system. Functionally ESXi is equivalent to the ESX offering the same levels of performance and scalability. The main difference between ESX and ESXi is that the Linux-based Service Console has been deprecated, reducing the footprint of the ESXi to around 32MB. An ESXi host consists of two main components: the VMkernel and the processes that run top of it. As mentioned before VMkernel has control of all hardware devices on the physical server and resources for the applications on the host. The main processes that run on top of the VMkernel are:

- Direct Console User Interface (DCUI) – a low-level configuration and management interface used primarily for basic initial configurations of the host.
- DCUI can be accessed only through the console of the server.

Copyright © 2014 IJTEEE.

- VMM process – provides the execution environment for virtual machines. Various management agents – used to enable VI management from remote applications.
- The Common Information Model (CIM) – an interface that enables hard-ware-level management from remote applications. CIM is an open standard that defines how computing resources can be represented and managed.

4.1. ESXi System Image Design - ESXi is distributed in two different formats, embedded and installable. Embedded version is directly embedded into the firmware of the server. The embedded version can be ordered from several major server manufactures and it offers faster deployment of the virtual infrastructure as no installation is needed. Installable edition of ESXi can be downloaded from the VM-ware.com website and is installed into the server's boot disk. Regardless of the edition used, the same system image design is present:

- A 4 MB boot loader partition.
- A 48 MB boot bank, contains the 32 MB core hypervisor code.
- An alternative 48 MB boot bank, initially empty.
- A 540 MB store partition, contains various utilities including the VI Client and VMware Tools images.
- A 110MB core dump partition, normally empty but can hold diagnostic in-formation if problems occur.

4.2. ESXi startup and management – When the ESXi boots for the first time, the VMkernel discovers hardware devices and selects appropriate drivers for them.The DCUI has several configuration tasks that are mainly used during initial server configurations and troubleshooting. These tasks include setting administrative password, configuring IP-addresses, performing network tests, viewing logs and restarting agents. All other configurations are done using remote management tools such as the VI Client.

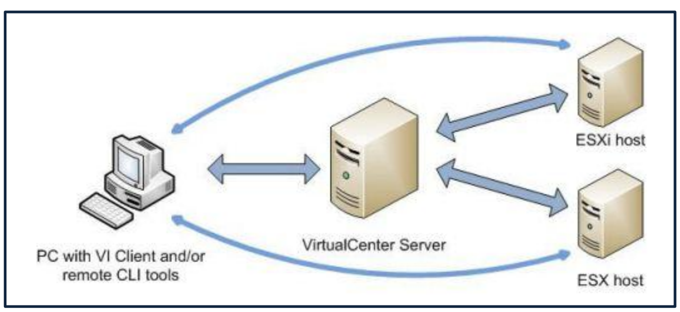

Figure3.3-ESXi Management

A single PC with remote management tools can manage all hosts through the VirtualCenter Server or just a single host through a direct connection.

Parameter	ESXi	ESX
On-disk footprint	32MB	2GB
Bootstrap	Direct from boot loader	Service Console Driven
Direct Management	DCUI	Service Console Shell Session
Hardware Monitoring Agents	CIM plugins	Full application in service console
Other Agents	Only with VI SDK	Full application in service console
Scripts, Automation & Troubleshooting	DCUI, remote CUI	Service Console Shell session
Other Software	Moved to outside environment	Resident in service console

Table 3.2-ESXi and ESX comparison

III. DATA CENTER IMPLEMENTATION & CONFIGURATION

The practical part of this thesis work shows how to consolidate a single physical server with Windows Server 2003 R2 Standard Edition into a virtual machine running Windows Server 2008 Standard edition. A single physical server and a single free license of the ESXi is used, therefore features like VirtualCenter, VMotion and DRS are not available during the practical work.

1. Deciding the Destination Operating System

According to the licensing model of the source operating system, the same license cannot be used with any other hardware than it was originally shipped with. This kind of licensing is called Original Equipment Manufacturer (OEM) licensing. For this reason a new license had to be purchased for the operating system of the destination virtual machine. A decision was made that the destination operating system will be Microsoft Windows Server 2008 Standard Edition as all the roles and services of the Microsoft Windows 2003 Standard Edition source server can me migrated to the newer operating system. It is important to notice that the existing Microsoft Windows Server 2003 domain needs to be prepared before a Microsoft Windows Server 2008 can be promoted to function as a Domain Controller.

2. Hardware, Operating System and Roles

Modern day servers should provide good fault tolerance for both hardware and software in order to keep the company's production running. Services provided by the servers must be available to the network at all times. Redundancy of these services is provided by fault tolerant hardware setups and by replication of server roles.

- **Source -** The physical server under consolidation is a Dell PowerEdge 600SC based on Intel Pentium 4 Processor. It is a tower server and was bought in 2003 with the following hardware:
    - Intel Pentium 4 processor, 2.4GHz with 512KB cache
    - 400MHz Front Side Bus
    - 2 x 256MB DDR ECC SDRAM memory
    - 2 x 80GB 7200 RPM IDE drives, RAID 1 configured
    - CERC IDE RAID controller

Copyright © 2014 IJTEEE.

ISSN 2347-4289

- Single 250 W power supply unit
- Single embedded Gigabit NIC
- IDE CD-ROM drive

The only redundant components in this server are the RAID 1 mirrored hard drives so this server's hardware is quite outdated. As an operating system the source server has Windows Server 2003 Standard Edition. The primary role of this server is to work as a Domain Controller in the company's Windows Server 2003 Domain. In order to function as a Domain Controller and a DNS server, these three roles have to be installed:

- Domain Controller role for Active Directory
- DNS Server role
- File Server role

▪ **Destination -** The VJTIfinally made a decision to deploy available dell servers for first step. ESXi offers several options from Dell server family, including Dell PowerEdge 2950 III. From previous experience it was known this affordable 2U rack server provides solid and reliable performance. Dell PowerEdge 2950 III offers highly customizable hardware configuration for multiple different purposes, for example virtualization. The destination server had the following configurations:

- 2 x Quad-Core Xeon E5440 CPU, 2,8GHz with 2 x 6MB cache
- 1333 MHz Front Side Bus
- 4 x 4GB Fully Buffered DIMM (667MHz)
- 2 x 73GB 15000 RPM SAS hard drives, RAID 1 configured
- 4 x 450GB 15000 RPM SAS hard drives, RAID 5 configured
- PERC 6/i SAS RAID controller
- 2 x 750 W power supply unit
- Dual embedded Broadcom NetXtreme II 5708 Gigabit2 Ethernet NIC
- IDE DVD-ROM drive
- Dell Remote Access Controller (DRAC)
- 3 year Dell ProSupport, 4 hour on-site response time

3. Installing the ESXi Host

After installing the server to the rack and plugging in the power cables, a few BIOS modifications are needed. Enabling Intel VT provides hardware-assisted virtualization support and configuring the DRAC with an IP address and root password provides remote management from any computer within the LAN. These steps are not required but they give a boost to virtual machine performance and remote server management capabilities. The server is connected to the company LAN with three Ethernet cables: Two cables are plugged into the Dual embedded Broadcom NIC and one cable is plugged into the DRAC. There are two logical hard drives available for the virtual infrastructure. The 73GB drive will be called Datastrore1 and the 1350GB drive will be called Datastore2. The ESXi is installed on Datastore1 while Datastore2 is intended for virtual machines. Later on Datastore1 will also contain installation images for guest operating system installations. Installation image for the ESXi is obtained from vmware.com website.Installation takes less than five minutes and after the server reboots the virtual infrastructure is up and running with default settings. When the ESXi host is running with default settings root password is not set, Management network IP address is assigned by a DHCP server and the server has a default hostname "vmvisor.vmware.com". The default settings are not the best practices for virtual infrastructures, so a static IP address and a hostname that describes the server's role in the LAN are assigned. For this paper work the hostname will be "vjtivdc". IP and DNS configurations are done in the pane "Configure Management Network". A single IP address is bound to both physical NICs of the server to provide fault tolerance. Root password is also enabled for security reasons. All of the settings described above can be modified from the DCUI seen in Figure 4.1. Figure 4.2 shows the DCUI view when the root user is logged out. This view shows the version of the ESXi, manufacturer and model of the server hardware, model of the processors, the amount of installed physical memory and the IP address of the server. After the initial installations steps the management tools for the virtual infrastructure are in-stalled to a remote computer. The VI Client can be downloaded from a web address shown in DCUI, in this case from http://192.168.X.X/. The website also contains a link to download remote CLI tools, but for this thesis work only the VI Client is used for management. VI Client software is downloaded and installed to a Windows XP Professional computer within the VJTI LAN. After installation the management client is ready to be connected to the ESXi.The connection window of the VI Client is shown in Figure4.3.

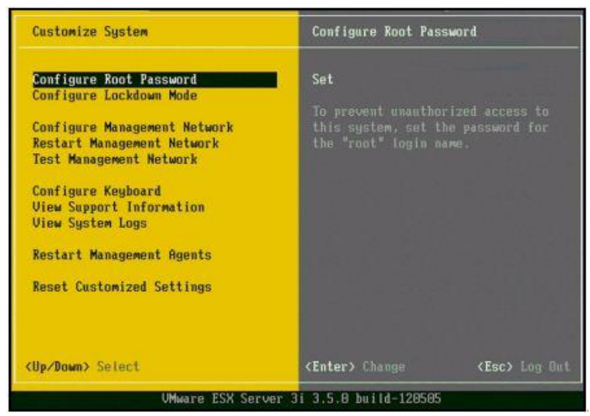

Figure 4.1 – DCUI Configuration

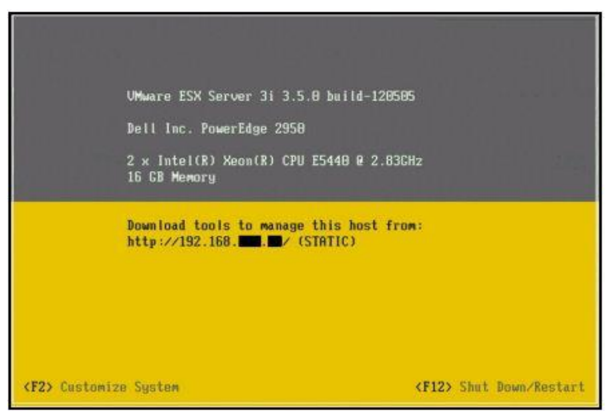

Figure 4.2 – DCUI

Connection to the host is established and the client is ready to start managing the virtual infra-structure.

Copyright © 2014 IJTEEE.

ISSN 2347-4289

Figure 4.3 – VI Client Connection Window

After this you are able to see the main management window where from we can actually manage the whole network and all the virtual machines. The ESXi host and the VI Client are ready to start configuring the virtual infrastructure and virtual machines. The "Configuration" –tab shows information about the host's physical health as well as contains all the important configurations for the virtual infrastructure. The "Health Status" -view shows comprehensive and real time health information about the physical components of the host server. If a physical component would fail or break down, the status of that particular component will change to either warning or alert, depending on the severity of the failure. According to the health status the host is now ready for initial VI Client configurations, as all physical components are in the "Normal" state.

4. Creating Virtual Network

We use a vSwitch here for connecting multiple VMs and management network with different ports. So we would be able to access the complete network from anywhere in LAN. There will be **two port groups** inside a single vSwitch: Virtual Machine Port Group is for the virtual machine network and VMkernel Port is for the management network. In "Configuration→Hardware→Networking" –tab, administrator can access the network configurations of the ESXi host.A single vSwitch called "vSwitch0" is configured to provide physical network connectivity for the Virtual Machine Port Group and the VMkernel Port. The two physical NICs, vmnic0 and vmnic1, are configured to form a NIC team that provides vSwitch0 with connectivity to the LAN. The Virtual Machine Port Group is named "VJTI LAN" and the VMkernel Port is named "Management Network". After these configurations, the vSwitch0 connects the two port groups to the physical network using the two physical NICs, vmnic0 and vmnic1, configured with failover and load balancing to provide redundancy.

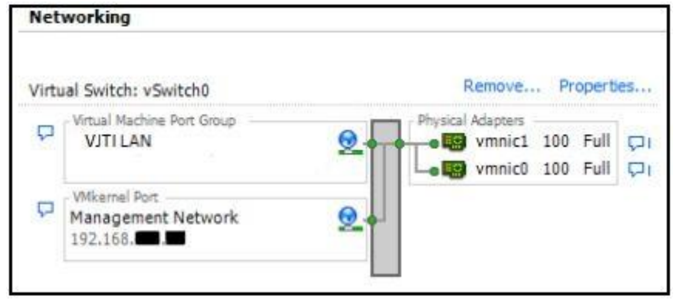

Figure 4.4 – vSwitch and network connectivity

Now we can proceed with the creation of the virtual machines in different clusters. Minimum requirements are :

- Single 1GHz CPU(x86) or single 1.4GHz CPU (x64).
- 512 MB of RAM
- 10 GB of available disk space
- DVD-ROM drive

Now the very final step is that – Once you are logged in with Virtual Center Sever, you need to add the a Datacenter→ Add Clusters into it→ Add the already created VMs or we can also add our own machine. Interesting thing – You can also add your own physical system virtually into the network as a VM which also belongs to one of this Datacenter Machines. Other than this we also need to customize the migrating roles, monitor the performance of the datacenter peripherals such as storage, network, CPU, Running status, memory usage etc. At the time of writing this, the ESXi host has been online for about 2-3 months and is currently hosting 11 running virtual machines. Guest machines have Linux or Windows operating systems and the memory allocated varies between 128 MB and 1536 MB. Virtual machines have different purposes for both testing and product environments. The following figures are theperformance check of the overall systems.

Figure 4.5 – Datacenter Management Window

Copyright © 2014 IJTEEE.

ISSN 2347-4289

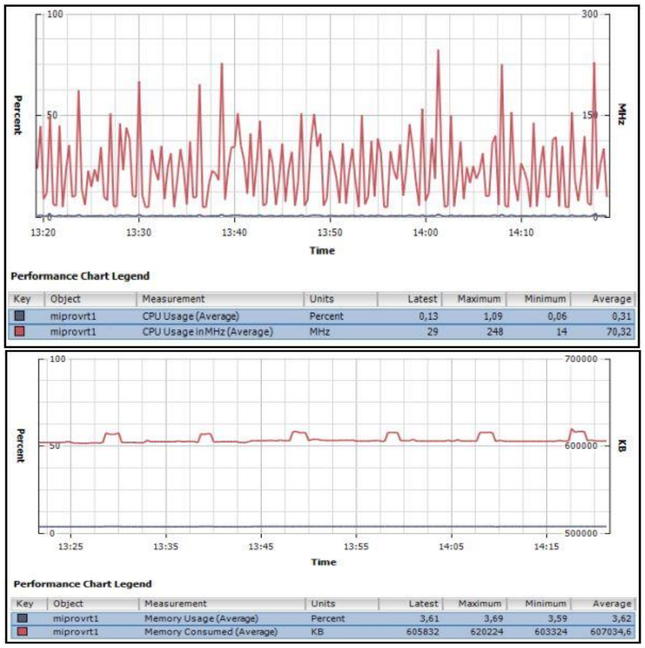

Figure 4.6 – Hypervisor only CPU usage and Memory usage

IV. CONCLUSION

It is evident that the VMware ESXi 3.5/4.0 and above is a powerfulvirtualization platform with versatile features and lots of scalability. VMware ESXi 4.0 Free Edition provides a good foundation forvirtual infrastructures and its licensing model is easily upgradable to meet higher needs. It is important to keep in mind that the free edition of ESXi is best suited for testing environments and the most business critical services should not be consolidated to be hosted on top of a single hardware. The system is suitable for production server consolidation but for high avail-ability of services, there is a need to purchase redundant hardware and at least VI Standard Edition. After implementation the virtual infrastructure has operated as expected and the availability of the services it provides has been very high. Services have been unavailable only when the ESXi has been applied with important security patches that required a server reboot. Creating new virtual machines and the allocation of physical resources has been easy and effortless. The primary objective of this thesis work was to experiment how to eliminate the problem of a server sprawl using server consolidation. The results are more than satisfying as the hard-ware used could host up to 20 virtual machines using VMware ESXi virtualization platform.Virtualization is a fairly new subject in the IT world so the works of reference are still quite limited. There is a lot of material available online, but usually they are written in cooperation with the virtualization providers and therefore objectivity is often controversial. Indeed the biggest challenge in this thesis work was to find good books about virtualization that dis-cusses the topic objectively. Once suitable source material was found the project was carried out without any bigger problems. So, again the bottom line is that all the educational institutions that require the data center for their colleges and institutions and their students can very easily implement this virtual data center. It is all the way much efficient, highly uptime and guaranteed than the hardware data centers for your purposes.It also serves the Cloud Networks.

REFERENCES

[1] Chappell, David 2008.
Virtualization for Windows: A Technology Overview. [on-line][referred 3.3.2009]. Available in www-format:<http://download.microsoft.com/download/E/0/4/E04E1183-E5FC-4C1B-A2AA89E3F4D6E4C/MSVirtualizationOverviewv3dot0.docx>.

[2] Citrix Homepage 2009.
Citrix XenServer – Features and Benefits. [online][referred 1.3.2009]. Available in www-format:<http://www.citrix.com/English/ps2/products/subfeature.asp?contentID=1686942>

[3] Dell 2007. Dell PowerEdge 2950 Server. [online][referred 9.8.2009]. Available in www-format:<http://www.dell.com/downloads/global/products/pedge/en/PE2950_SS_072007.pdf>.

[4] IBM 2005. Server Consolidation with VMware ESX Server. [online][referred 21.2.2009] Available in www-format: <www.redbooks.ibm.com/redpapers/pdfs/redp3939.pdf>.

[5] Intel 2006. Server Consolidation Using Quad-Core Processors. [online][referred 14.4.2009]. Available in www-format: <www.intel.com/it/pdf/consolidate-using-quadcore.pdf>.

[6] Morgan, Bradley 2006. Virtualization. [online][referred 14.1.2009]. Available in www-format: <http://www.windowsecurity.com/whitepapers/Virtualization.html>.

[7] Converged Optical Network and Data Center Virtual Infrastructure Planning by Konstantinos N. Georgakilas, Anna Tzanakaki, Markos Anastasopoulos, and Jens Myrup Pedersen, VOL. 4, NO. 9/SEPTEMBER 2012/J. OPT. COMMUN. NETW.

[8] VMware KB articles

[9] http://www.google.com/about/datacenters/locations/index.htm

[10] ADVA Optical Networking, “Virtualized optical networks forsustainable cloud services,” White Paper.

[11] E. Escalona, S. Peng, R. Nejabati, D. Simeonidou, J. A. Garcia-Espin, J. Ferrer, S. Figuerola, G. Landi, N. Ciulli, J. Jimenez,B. Belter, Y. Demechenko, C. de Laat, X. Chen, A. Yukan,S. Soudan, P. Vicat-Blanc, J. Buysse, M. De Leenheer, C. De-velder, A. Tzanakaki, P. Robinson, M. Brogle, and T. M. Bohn-ert, “GEYSERS: A novel architecture for virtualization and co-provisioning of dynamic optical networks and IT services,” in Future Network & Mobile Summit (FutureNetw), 15–17 June2011, pp. 1–8.

Copyright © 2014 IJTEEE.

ISSN 2347-4289

[12] sNICh: Efficient Last Hop Networking in the Data Center byKaushik Kumar Ram, Rice University, Alan L. Cox, Rice University, Scott Rixner, Rice University, Parthasarathy Ranganathan,Jose Renato Santos,Jayaram Mudigonda, HP Labs.

[13] The Cost of a Cloud: Research Problems in Data Center Networks byAlbert Greenberg, James Hamilton, David A. Maltz, Parveen Patel Microsoft Research, Redmond, WA, USA.

[14] Cisco. Data center ethernet. http://www.cisco.com/en/US/-netsol/ns783/networking solutions package.html

[15] Cisco systems: Data center: Load balancing data center services, 2004

[16] M. Isard. Autopilot: Automatic data center management. Operating Systems Review, 41(2), 2007

[17] E. R. Hinden. Virtual router redundancy protocol (VRRP). RFC 3768, 2004

[18] A Mathematical Programming Approach for Server Consolidation Problems in Virtualized Data Centers by Benjamin Speitkamp and Martin Bichler.IEEE TRANS ACTIONS ON SERVICE S COMPUTING, VOL. 3, NO. 4, OCTOBER -DECEMBER 2010

[19] Google -Wide Profiling : A continuous Profiling Infrastructure For Data Centers.

[20] VMware , Inc. VMware virtual networking concepts http://www.vmware.com/files/pdf/ virtual_networking_concepts.pdf, 2007.

Copyright © 2014 IJTEEE.