

Aki Tuhkanen

Server Consolidation with VMware ESXi 3.5

Bachelor's Thesis
Degree programme in Information Technology

April 2010



MIKKELIN AMMATTIKORKEAKOULU

Mikkeli University of Applied Sciences

DESCRIPTION


 MIKKELIN AMMATTIKORKEAKOULU Mikkeli University of Applied Sciences		Date of the bachelor's thesis 14.04.2010
Author(s) Aki Tuhkanen	Degree programme and option Information Technology	
Name of the bachelor's thesis Server Consolidation with VMware ESXi 3.5		
Abstract <p>The powerful hardware of modern day computers are designed to run only a single operating system and this typically leaves the hardware highly underutilized. This can lead to a problem called a server sprawl where several underutilized computers are hard to manage, use floor space inefficiently and consume excessive amounts of power. Virtualization provides a solution to server sprawls through server consolidation where several physical servers are migrated to operate on top of single server hardware.</p> <p>This thesis work discusses virtualization in general along with benefits, drawbacks, uses and providers of virtualization technologies. When discussing about different virtualization approaches the main emphasis is on server virtualization and its four main categories. Physical to virtual server consolidation and VMware Infrastructure 3 products are discussed in more comprehensive manner.</p> <p>The practical part of this thesis consists of physical to virtual consolidation of a single physical server that is running Windows Server 2003 operating system. The destination platform is a Windows Server 2008 virtual machine that is running on top of VMware ESXi 3.5 virtualization platform. After the consolidation process the virtual machine provides exactly the same services as the physical server did before the consolidation. The practical part also briefly monitors and evaluates the performance of the virtualization host and the virtual machine.</p> <p>The results of this thesis work show that VMware ESXi 3.5 is a powerful virtualization platform with lots of scalability. The free version of VMware ESXi 3.5 provides a good foundation for a virtual infrastructure but should mainly be implemented for testing environments. Larger production environments should provide host redundancy and powerful backups; therefore at least VMware Infrastructure 3 Standard should be purchased.</p> <p>The main objective of the thesis was to experiment how to eliminate the problem of a server sprawl using virtualization technologies and server consolidation. The results show that the hardware used can host approximately 20 virtual machines and therefore eliminate the server sprawl problem.</p>		
Subject headings, (keywords) VIRTUALIZATION, VMWARE, ESX, ESXI, SERVER CONSOLIDATION, CONSOLIDATION		
Pages 76	Language English	URN http://www.urn.fi/URN:NBN:fi:amk-201004146457
Remarks, notes on appendices		
Tutor Matti Juutilainen	Employer of the bachelor's thesis Mipro Oy	

TABLE OF CONTENTS

TERMS AND ABBREVIATIONS	1
1 INTRODUCTION	3
2 MIPRO OY – SYSTEM INTEGRATOR.....	3
2.1 Network Infrastructure and System Description	4
3 VIRTUALIZATION	5
3.1 What Is Virtualization?	6
3.2 History	8
3.3 Virtualization Approaches.....	9
3.3.1 Server Virtualization	10
3.3.2 Storage Virtualization	15
3.3.3 Network Virtualization.....	16
3.3.4 Application Virtualization.....	18
3.4 Why Virtualize?	18
3.5 Uses of Virtualization.....	19
3.6 x86 Virtualization Providers	20
4 SERVER CONSOLIDATION	22
4.1 Benefits.....	23
4.2 Risks and Drawbacks	24
4.3 When to Consolidate?	25
4.4 Server Consolidation Using Intel Quad-Core Processors	26
5 VMWARE INFRASTRUCTURE 3.....	27
5.1 Hypervisor Architecture	28
5.1.1 Physical Host Server	29
5.1.2 VMkernel	31
5.1.3 Service Console.....	32
5.2 Features	33
5.3 Virtual Networking.....	36
5.4 VMware ESXi 3.5	39
5.4.1 ESXi System Image Design	40
5.4.2 ESXi Startup.....	40
5.4.3 ESXi Management	41
5.5 Summary of Differences Between ESXi and ESX.....	42

6 PRACTICAL IMPLEMENTATION.....	44
6.1 Deciding the Destination Operating System	44
6.2 Hardware, Operating System and Roles.....	44
6.2.1 Source.....	45
6.2.2 Destination	46
6.3 Installing the ESXi Host.....	47
6.4 Defining Virtual Networks	52
6.5 Creating the Virtual Machine	54
6.6 Installing and Configuring the Guest Operating System.....	57
6.7 Migrating Roles.....	60
6.8 Monitoring and Performance of MIPROVRT1.....	65
6.8.1 Performance Data.....	65
6.8.2 Analysis of Performance Data.....	70
7 CLOSING WORDS.....	71
REFERENCES	73

TERMS AND ABBREVIATIONS

BIOS	Basic Input Output System
CIM	Common Information Model
CLI	Command Line Interface
CPU	Central Processing Unit
DAS	Directly Attached Storage
DCUI	Direct Console User Interface
DDR	Double Data Rate
DHCP	Dynamic Host Control Protocol
DIMM	Dual In-line Memory Module
DNS	Domain Name System
DPM	Distributed Power Management
DRAC	Dell Remote Access Controller
DRS	Distributed Resource Scheduling
ECC	Error Correcting Code
EDM	Electronic Document Management
GB	Gigabyte
HCL	Hardware Compatibility List
I/O	Input/Output
IDE	Integrated Drive Electronics
IP	Internet Protocol
iSCSI	Internet Small Computer System Interface
IT	Information Technology
JBOD	Just a Bunch Of Disks
LAN	Local Area Network
MB	Megabyte
NAS	Network Attached Storage
NFS	Network File System
NIC	Network Interface Card
OEM	Original Equipment Manufacturer
OS	Operating System
RAID	Redundant Array of Independent Disks

RAM	Random Access Memory
RPM	Revolutions Per Minute
SAN	Storage Area Network
SAS	Serial Attached SCSI
SCSI	Small Computer System Interface
SDRAM	Synchronous Dynamic Random Access Memory
SMP	Symmetric Multi Processing
TCP	Transmission Control Protocol
VCB	VMware Consolidated Backup
VI Client	VMware Infrastructure Client
VLAN	Virtual Local Area Network
VM	Virtual Machine
VMFS	Virtual Machine Files System
VMM	Virtual Machine Monitor
vNIC	Virtual Network Interface Card
VPN	Virtual Private Network
vSwitch	Virtual Switch

1 INTRODUCTION

The idea to implement this thesis work derives from the fact that virtualization is a vastly growing trend in today's IT world. 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 thesis work will concentrate on server virtualization and server consolidation using a bare-metal hypervisor platform VMware ESXi 3.5. 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 3.5 virtual environment.

2 MIPRO OY – SYSTEM INTEGRATOR

Mipro Oy is an independent and private limited company that specializes in customized automation. Business areas are safety related systems for railways and industry, automation for water and wastewater works, and maintenance. (Mipro Homepage 2009).

Mipro Oy was founded in 1980 and currently has 67 employees. The annual turnover is around 10 million euros. Mipro Oy has a head office in Mikkeli and a branch office in Oulu. Mipro Oy's best expertise can be seen in professional engineering, innovative product devel-

opment and experience in project works. The term “System Integrator” effectively describes Mipro Oy’s operation to safely combine and integrate separate systems together. (Mipro Homepage 2009).

Since 1980 Mipro Oy has extended its operations to a comprehensive supplier of automation and information systems. A good example of a comprehensive project is the renewal of interlocking systems at Ilmala depot in Pasila. Ilmala depot has 55 kilometers of railway tracks and it serves 85 percent of Finland’s all long distance cars. Mipro Oy delivers the whole new interlocking system which replaces three old separate systems. The project started at 2006 and is estimated to continue until 2011. (Mipro Homepage 2009).

2.1 Network Infrastructure and System Description

Windows Server Domain is a group of Microsoft Windows workstations and servers that share a directory database called Active Directory. One or more Active Directory Domains will form a Forest which is the highest level in the Active Directory hierarchy. With the features of Active Directory administrators can manage several Forest and Domain functions, such as handling trusts, user accounts, security policies and software installations for all the machines within the domain. Active Directory database is replicated between several servers called Domain Controllers. Each Domain Controller contains a complete copy of Active Directory and if modifications to any database objects are made, these changes are replicated to all Domain Controllers within the domain. One of the Active Directory Domain Controllers must hold the Forest wide “Schema Master” role and one Domain Controller must hold the Domain wide “Infrastructure Master” role. Domain servers that are not Domain Controllers are called Member Servers and other computers are called Domain Computers. All the domain workstations and servers must be able to locate the Active Directory database. The most common locator method is to use well known Internet standard called Domain Name System (DNS) which is used to resolve computer IP-addresses to domain names. (Cross et al. 2003, 2, 5-6, 11, 28).

Mipro Oy's internal data network is a local area network (LAN) based on TCP/IP protocols. The Active Directory consists of only a single Windows Server Domain. The Domain has 4 Domain Controllers, 3 Member Servers and 120 Domain Computers. All the servers in the Windows domain have Windows 2003 Standard Edition as an operating system and domain workstations are mainly installed with Windows XP Professional. Company LAN also has 2 Linux servers used as file- and database servers. The roles of all the servers in the company LAN are illustrated in Table 1. DC stands for Domain Controller, MS for Member Server and LS for Linux server.

Table 1. Server roles

Role	DC1	DC2	DC3	DC4	MS1	MS2	MS3	LS1	LS2
Domain Controller	X	X	X	X					
DHCP	X								
DNS	X	X	X	X					
File server	X	X	X	X	X				
Print server					X				
EDM						X			
Backup to tape server				X		X			
Access control server							X		
Database server								X	X

Domain Controller and DNS roles are replicated between several Domain Controllers to provide fault tolerance for the domain services. DC1 holds the Schema Master role and the Infrastructure Master role. Domain Controllers have the File server role installed only to replicate user logon scripts and Group Policy Objects between each other, not to provide any file server operations to Domain workstations. At the moment all the servers have their own physical hardware and server virtualization is not utilized.

3 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. This is why the term virtualization needs to be put into better context. The following chapter will define virtualization and its history while keeping the main emphasis on server virtualization.

3.1 What Is 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 thesis 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.

Depending on the virtualization approach, virtualization is used to either separate or combine computing resources. Figure 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. (Waters 2007, 1). (Morgan 2006, 1-2).




Figure 1. Server virtualization
(Rule & Dittner 2007, 19)

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 2 shows an example of storage virtualization where multiple physical disks are combined into a one large logical disk.




Figure 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 Chapter 3.3.2.

3.2 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 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 awoken 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 Microsoft 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 August 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).

3.3 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).

3.3.1 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. With server virtualization the physical characteristics of the server platform is hidden from the user. Instead, the user sees virtual hardware that emulates physical hardware. (Muller & Wilson 2005, 5).

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. VMM is software that functions as an interpreter between the host system and the virtual machines. 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 3. A good example about this type of VMM is VMware ESX Server which will be discussed later in more detail. (Rule & Dittner 2007, 18).

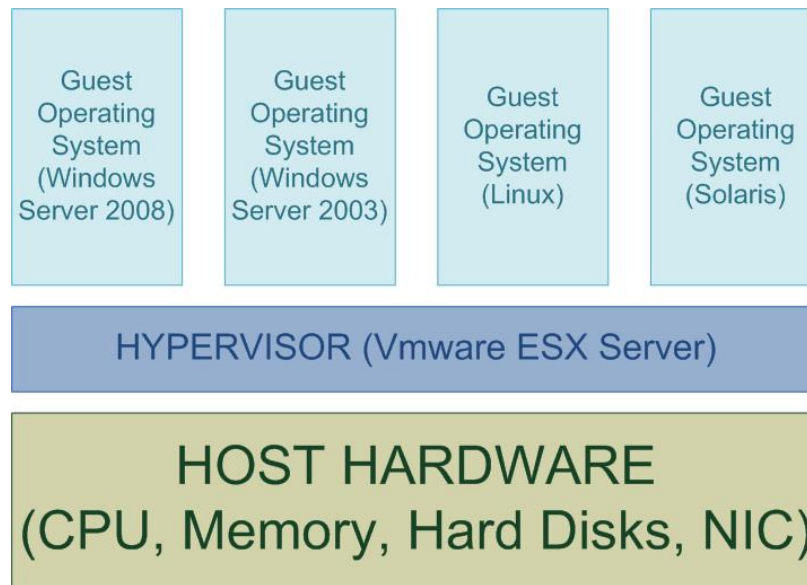


Figure 3. Bare-metal hypervisor
(Marshall et al. 2009, 10)

Type 2 VMM, also known as hosted hypervisor, software runs on top of a host operating system adding one more layer to the model, thus the guest operating systems operate at the third layer above the hardware as Figure 4 shows. Examples of software that utilizes the hosted architecture are VMware Server and Microsoft Virtual Server (Rule & Dittner 2007, 18-19).

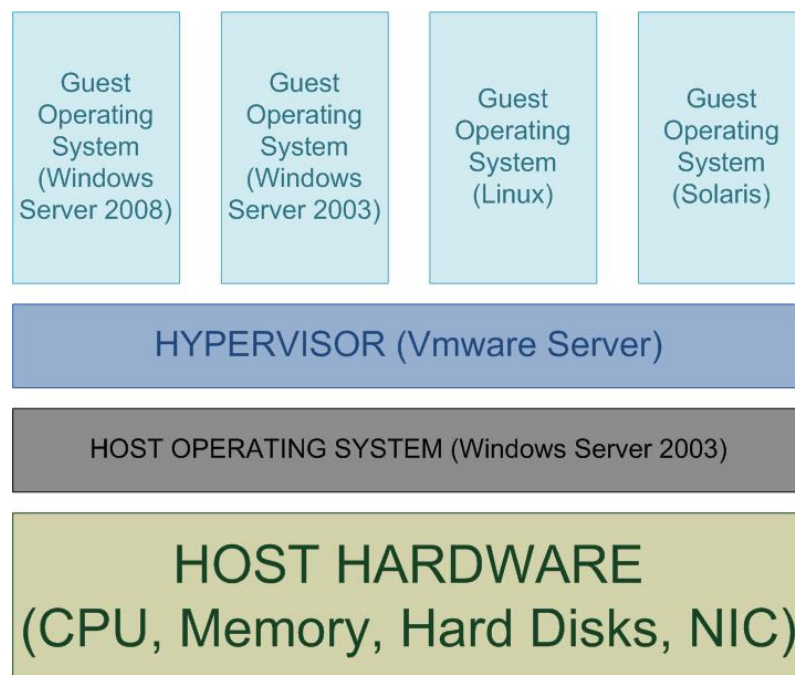


Figure 4. Hosted hypervisor
(Marshall et al. 2009, 9)

Server Virtualization is divided into four categories: Full virtualization, paravirtualization, operating system virtualization and hardware-assisted virtualization.

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 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 operate just 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). (VMware 2007a, 4).

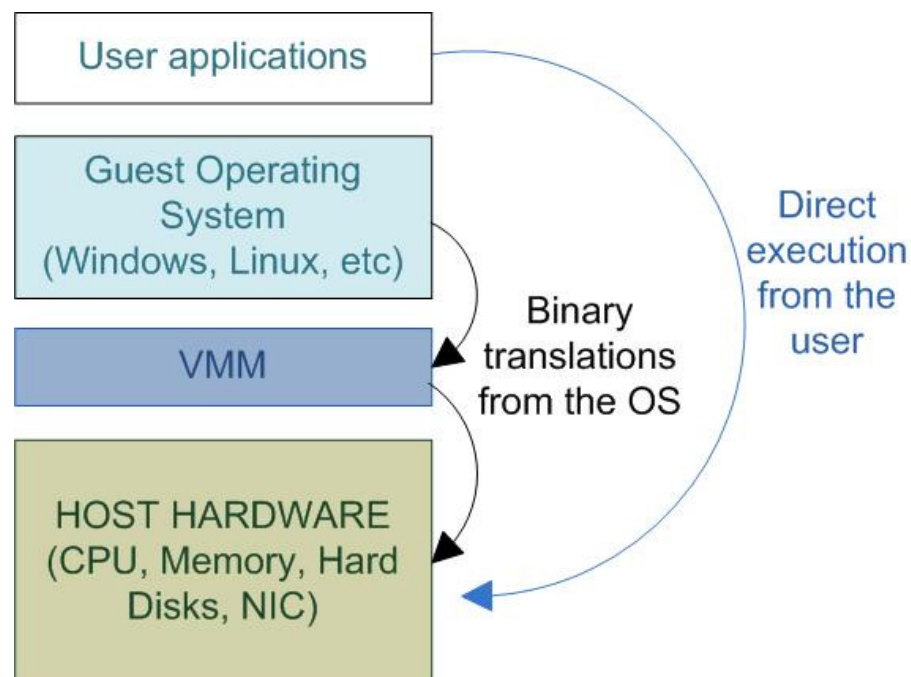


Figure 5. Full virtualization
(VMware 2007a, 4)

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 6. Paravirtualization provides low virtualization overhead and easy implementation but the lack of support 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).




Figure 6. Paravirtualization
(VMware 2007a, 5)

Operating system virtualization, also known as container-based virtualization, runs on top of existing operating system and is used to provide multiple instances of the host operating system. These virtual instances are called containers which are isolated from each other so that the users are not able to see or interact with other virtual instances running on the same hardware. Operating system virtualization imposes little overhead as virtual instances use the operating systems normal system call interface instead of emulated one like in the case of full- or paravirtualization. This makes operating system virtualization very efficient as most of the

hardware resources are available to the applications running in the containers and a VMM is not needed. The downside of operating system virtualization is the fact it is not flexible, the host and the guest operating systems will have the same kernel. For example in Figure 7 the host operating system is Linux, therefore all the hosted virtual instances will also have a Linux kernel. Other guest operating systems such as Windows cannot be implemented. Operating system virtualization is used for web hosting and to some extent for server consolidation. Examples of operating system virtualization are OpenVZ and Virtuozzo. (Golden & Scheffy 2008, 7-8). (Rule & Dittner 2007, 22). (VMware 2006, 10).

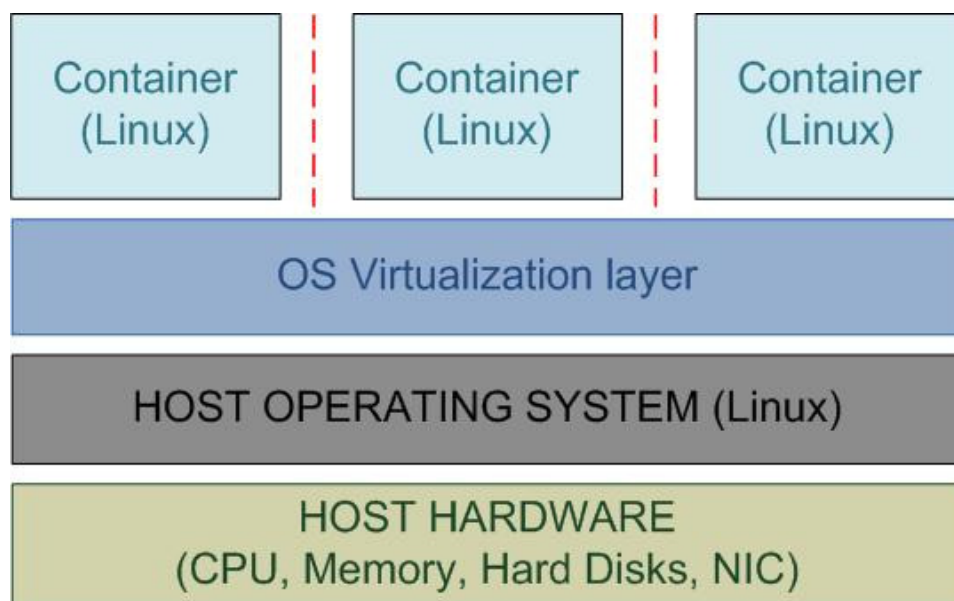


Figure 7. Operating system virtualization
(Parallels Homepage 2009)

Hardware-assisted virtualization is also referred as native virtualization. This approach provides a way to combine full virtualization or paravirtualization with hardware's I/O acceleration techniques. Hardware-assisted virtualization takes advantage of Intel VT and AMD-V techniques which are supported by most of the latest processor models from Intel and AMD. Figure 8 illustrates Hardware-assisted virtualization. (VMware 2007a, 6).

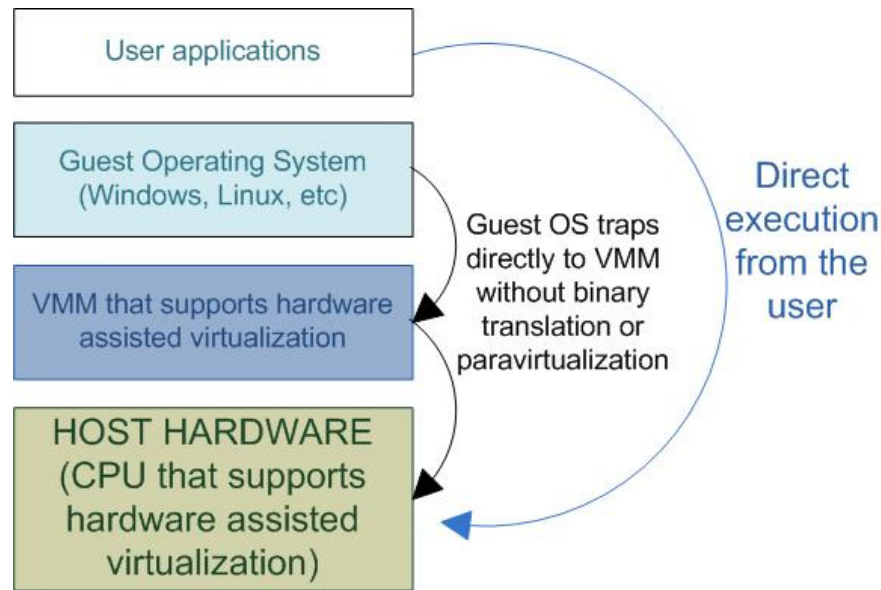


Figure 8. Hardware-assisted virtualization
(VMware 2007a, 6)

When paravirtualization is combined with a hardware-assisted virtualization technique, the guest operating system kernel does not require much modification as the guest operating system has direct access to the hardware resources. This also results in improved I/O performance opposed to any other server virtualization approach. As shown in Figure 8, guest operating system instructions are set to automatically trap to the hypervisor, removing the need for either binary translation or paravirtualization. The downside is that this approach requires modern hardware support and when used with paravirtualization, the guest operating system still needs some modifications so closed source operating systems cannot be installed. Examples of virtualization solutions which take advantage of hardware-assisted virtualization are VMware Workstation, Xen 3 and Linux KVM. (Rule & Dittner 2007, 23-24). (VMware 2007a, 6).

3.3.2 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 networks which appear to be a single physical instance, but actually are a group of several physi-

cal 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. (Rule & Dittner 2007, 24-25).

3.3.3 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. (Rule & Dittner 2007, 25).

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. (Rule & Dittner 2007, 25).

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. (Rule & Dittner 2007, 25).

Virtual networking provides networks between virtual machines and a bridge between virtual and physical networks. A single virtualization host can provide its guest operating systems with several independent virtual network segments as well as connections to the physical NICs. Virtual switches (vSwitch) allow virtual machines on the same host to communicate with each other using the same protocols that would be used over physical switches, without the need for additional networking hardware. Virtual networks can be connected to physical networks in a very flexible manner and for example VLAN tagging is supported. With virtual networking a guest operating system can be provided with multiple virtual NICs (vNIC) which are connected to different networks, both virtual and physical. Figure 9 shows an example of virtual networking inside a single virtualization host. (VMware 2007b, 3-4).




Figure 9. Virtual networking
(VMware 2007b, 6)

In the example of Figure 9, 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. Virtual

networking with VMware Infrastructure 3 will be discussed in more detail in chapter 5.3. (VMware 2007b, 3).

3.3.4 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. (Rule & Dittner 2007, 25-26).

3.4 Why Virtualize?

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. (Sun Microsystems & VMware 2008, 1).

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. Benefits of server consolidation will be discussed in more detail in chapter 4.1. (IBM 2005, 3-6).

Data center reliability and security are key concepts in any business. Traditionally large investments 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. (Rule & Dittner 2007, 7-16).

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 increased. (Rule & Dittner 2007, 7-16).

3.5 Uses of Virtualization

Common uses for virtualization in a business environment are elimination of underutilized or legacy hardware, rapid and flexible test environments, virtual desktops and high availability systems. Altogether these virtualization implementations will help to improve business continuity and reduce total cost of ownership. (Rule & Dittner 2007, 26).

Physical to virtual server consolidation is the most popular use of virtualization in today's IT world. Among other things, consolidation is used to eliminate underutilized and legacy server hardware helping companies to use their resources more effectively. Physical to virtual server consolidation will be discussed in chapter 4. (Marshall et al. 2009, 11). (Rule & Dittner 2007, 26-27).

With virtualization isolated test and development networks can be built very effectively within minutes. Test networks are populated with virtual machines that will form a flexible and effective test and development environment that for example can contain several software versions of the same subject being developed. Virtualization will increase productivity of the developers and accelerate release time. (Rule & Dittner 2007, 29-30).

As in the case of server hardware utilization, desktop and laptop computers are often underutilized as well. As most employees normally use only light applications such as web browsers, e-mail clients and word processing software, there is rarely an actual need for effective workstations. Virtual desktops provide a way for users to have a full desktop experience with minimal set of hardware. The actual desktops and computing resources are virtualized in the data center as users only have a thin client device with keyboard, mouse and monitor on their desks. (Rule & Dittner 2007, 29).

Some IT services of a business can be very critical and the possible downtime of these services can prove to be costly. Virtualization technologies can provide high availability to these systems without the high price of replicated hardware. High availability virtualization systems replicate the data between two or more virtual instances. Availability of business-critical services is assured and business continuity is improved. (Rule & Dittner 2007, 28).

3.6 x86 Virtualization Providers

The market provides several organizations and companies that each offer their own approach to x86 server virtualization. The following chapter will shortly introduce the most popular hosted and bare-metal virtualization solutions in the market.

VMware ESX Server is a bare-metal hypervisor that provides a wide range of supported host hardware and guest operating systems. VMware ESX Server products have a very rich set of features and guest operating systems can run at near native speed. Many different versions of ESX Server are available, from free of charge ESXi to enterprise level Virtual Infrastructure 3 Enterprise. Basic instructions are the same in all versions but features like high availability and live migration of guest operating systems are only available in the Enterprise version. ESX Server has the widest range of supported guest operating systems including most of the 64-bit operating systems available. ESX Server can be used in enterprise production environments for testing, development, enterprise server consolidation and high availability systems. Different versions of VMware ESX Server products and additional features will be discussed in more detail in chapter 5. (Rule & Dittner 2007, 40-42). (Marshall et al. 2009, 9).

Citrix XenServer is bare-metal hypervisor that is based on an open-source Xen hypervisor. XenServer can host Microsoft Windows and Linux operating systems, both 32-bit and 64-bit, but will require Intel VT or AMD-V on the underlying hardware in order to host any Microsoft Windows operating system. Like VMware ESX Server, Citrix XenServer offers many different editions for different situations. Basic bare-metal virtualization technology is the same in every edition but Enterprise and Platinum editions offer features like high availability and live migration of guest operating systems. Performance of guests hosted by XenServer is near native. Similarly to its competitors, XenServer can be implemented on enterprise production as well as testing and development environments. (Citrix Homepage 2009). (Marshall et al. 2009, 9).

Microsoft Hyper-V is a bare-metal hypervisor that competes with VMware ESX Server and Citrix XenServer. Hyper-V comes shipped together with Microsoft Server 2008 and it targets networks from workgroups to enterprises. Hyper-V offers a moderately good range of supported guest operating systems, including 64-bit versions of Windows operating systems and a few Linux platforms. Hyper-V cannot be installed on old systems as it requires hardware that supports Intel VT or AMD-V. It offers a near native performance to the supported guest operating systems. Hyper-V can be used in workgroups as well as in enterprise production environments for example in large scale testing, enterprise server consolidation or high availability systems. As Hyper-V comes free with the new Microsoft Server 2008, it is a notable alternative in the bare-metal hypervisor market although it does not provide the flexibility or performance of its more expensive competitors. (Chappell 2008, 10-11). (Marshall et al. 2009, 10).

VMware Server is a hosted hypervisor that is available free of charge. Installation of the hypervisor must be done on top of either Windows or Linux host operating system. Range of supported operating systems is quite good and VMware Server also supports 64-bit guests. VMware Server is a lightweight solution that cannot compete with bare-metal hypervisor solutions in performance but is useful for example as a testing environment or legacy hardware elimination. VMware Server is a great product to get started with virtualization as it is free and does not set any hardware requirements. (Marshall et al. 2009, 10). (Rule & Dittner 2007, 40-42).

Microsoft also offers a hosted hypervisor called Microsoft Virtual Server. Virtual Server is intended for testing and development environments and can be installed only on a host with Windows Server 2003 or Windows XP. Range of supported guest operating systems is quite narrow and Virtual Server does not support 64-bit guest operating systems. Like its competitor VMware Server, Microsoft Virtual Server offers moderate performance and does not set any requirements on hardware. Microsoft Server can be used for testing and development environments with Windows or Linux operating systems. (Rule & Dittner 2007, 53-56). (Chappell 2008, 10-11).

4 SERVER CONSOLIDATION

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 10. (IBM 2005, 2-3).

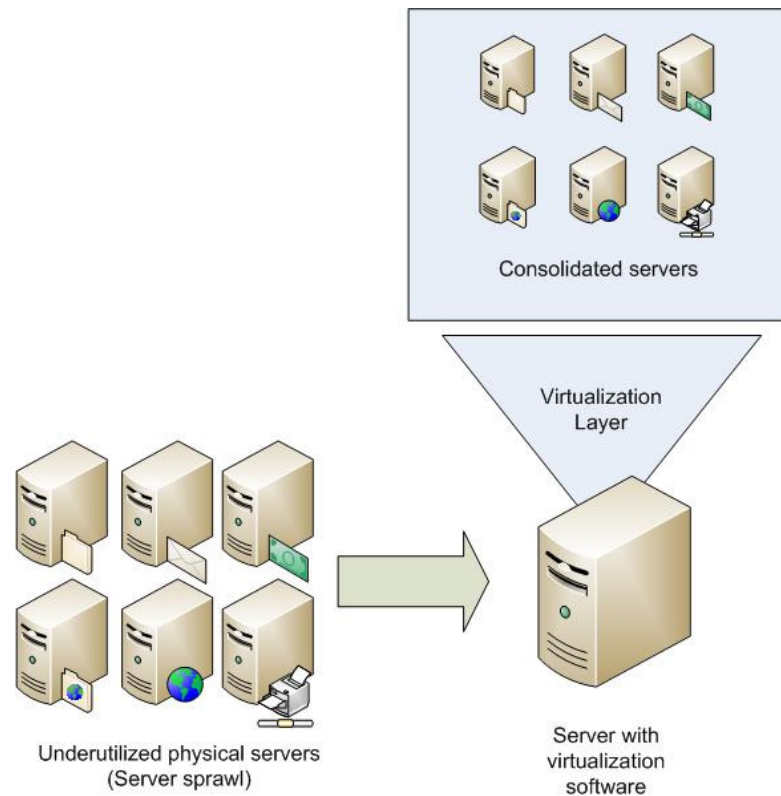


Figure 10. Physical to virtual server consolidation
(IBM 2005, 9).

When server consolidation is discussed later in this thesis, it will refer to physical to virtual server consolidation. Other types of server consolidation will not be discussed. 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 infrastructure is a physical server with a VMware ESX Server hosting several independent virtual machines that were physical servers before the consolidation process. (IBM 2005, 28).

4.1 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%. (IBM 2005, 4-5). (VMware Homepage 2009a).

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. In addition to higher utilization, consolidation provides more efficient system administration than traditional computing model as several server instances run on top of a single host and single administration tools. With consolidation efficiency of the whole IT infrastructure is improved. (IBM 2005, 5).

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. In addition to these business critical benefits, single point of control also further reduces both hardware and operational costs. With the help of consolidation, capabilities of server hardware can be increased more flexibly so the whole IT infrastructure is better prepared for future growth and power requirements. (IBM 2005, 5-6).

4.2 Risks and Drawbacks

When making the decision to implement server consolidation there are some risks and drawbacks that must be taken into serious consideration.

Chapter 4.1 stated that the most important benefit of consolidation is the cost savings. This is true when long term total cost of ownership is considered but implementing consolidation will

increase costs in short term. The need for more powerful hardware, virtualization software licenses, large amount of planning time and additional administration skills are evident. These needs are sure to increase the budget of the consolidation project. (IBM 2005, 22-23).

In consolidation several business critical services can be consolidated on top of single server hardware producing a single point of failure. Failure in the host server will lead into a failure in all the services running on top of the host. For example a broken motherboard can result in a long downtime period to all business critical services. Unless a high availability system is implemented, these services are unavailable until the defective motherboard is replaced with a new one. (IBM 2005, 24).

As several servers across the network are consolidated into a single server, the network traffic produced by these servers will accumulate into a single server. This can lead into network downtime and slow response if the bandwidth for the host server is not sufficient. This situation can be avoided with network load balancing for example with installation of multiple network interface cards on the host server. (IBM 2005, 26).

Backups are very important in any kind of environment. With consolidated environment backups can take a considerably longer time as several servers are backed up from inside a single hosts I/O system. Administrators will have to provide the host with high performance I/O and plan the time of backups in a way that the backup operations will not overlap with production environment operations. (IBM 2005, 24).

4.3 When to Consolidate?

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. This chapter will briefly discuss workloads that are suitable for consolidation, as well as workloads that are better suited to running on dedicated server hardware. (Otey 2006, 45).

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. The occasional utilization spikes will not be a problem as these spikes are only active a short period of time. These peaks might have some impact on individual guest servers but overall performance and responsiveness will stay at an acceptable level. (Otey 2006, 45).

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 percent. Backup servers have high I/O utilization for extended periods of time as data backups are transferred across the network to the backup locations. If backup servers were consolidated, the time for backing up would be significantly increased as they would have to share the hosts I/O system with other consolidated guests. Consolidating servers that have high resource utilization is not suitable as they would alone consume the most of the resources on the physical host. (Otey 2006, 46).

4.4 Server Consolidation Using Intel Quad-Core Processors

This paragraph is based entirely on one white paper from Intel (Intel 2006).

Business data centers typically have old server hardware with single or dual-core processors as well as servers with the more recent powerful quad-core processors. Consolidating the older less powerful servers is probably a good idea as quad-core processors can carry out the same operations faster. In November 2006 Intel released a server consolidation white paper where Intel IT used a series 5300 Quad-Core Intel Xeon processor to consolidate workloads of eight physical machines into virtual machines running on single host server. Comparisons in performance and power consumption were made between eight physical servers based on Intel Pentium III processor, one dual-socket Dual-Core Intel Xeon 5148 series-based server and one

dual-socket Quad-Core Intel Xeon 5300 series-based server. Therefore the Dual-Core server represents four CPU cores and the Quad-Core server represents eight CPU cores.

For the test environment Intel IT built a CPU intensive database application that applied CPU utilization of 10 percent to each of the physical machines with Intel Pentium III processor running at 733 MHz. The same database application was then implemented into eight separate virtual machines running on top of both dual-core and quad-core hosts.

Results show that when all eight consolidated workloads were running the quad-core server completed each workload 66 percent faster than the eight physical servers each running one workload and 34 percent faster than the dual-core server. Also when all workloads were running quad-core server consumed 29 percent less power than the dual-core server and 86 percent less power than the eight physical servers. When less than four of the workloads were running the dual-core server was slightly faster and consumed less power than the quad-core server. The quad-core server showed its advantage as more workloads were added. Job completion times remained relatively uniform throughout all workloads. When running all eight workloads the quad-core server took only 18 percent longer to complete each workload than running only a single workload. The dual-core server showed an average of 85 percent longer completion times with all eight workloads. As more workloads were added, the quad-core server also consumed less power per workload than the dual-core server as the job completion time was faster.

The results confirm that consolidating with the Quad-Core Intel Xeon processor 5300 series has the capability to extend server consolidation strategy to 8:1 consolidation ratio when CPU intensive applications are used. The 8:1 consolidation ratio means that 8 physical servers can be effectively consolidated to virtual machines running on single host hardware with a quad-core processor. Intel evaluates that 8:1 server consolidation with quad-core servers will reduce operational costs by more than 6000 US dollars a year.

5 VMWARE INFRASTRUCTURE 3

Currently VMware Infrastructure 3 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 chapter will discuss architecture of the hypervisors as well as features of different editions of VMware Infrastructure 3, from free standalone version of VMware ESXi 3.5 to enterprise-level VMware Infrastructure 3 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 3.5. From this point on, VMware ESX Server 3.5 will be referred to as ESX and VMware ESXi 3.5 will be referred to as ESXi. To keep this thesis work within limits of the thesis topic, this chapter will not discuss detailed information about physical server hardware, storage or networking systems. (Marshall et al. 2009, 15). (VMware 2008b, 9).

5.1 Hypervisor Architecture

The ESX has three main architectural components that provide the virtualization layer for the virtual infrastructure: Physical host server, VMkernel and Service Console as shown in Figure 11. It is important to state that the architecture of the ESXi is different as the Service Console is highly deprecated. The other components are the same as in the architecture of the ESX.




Figure 11. Architecture of the ESX Server

(McCain 2008, 2)

Physical host server is a server computer that provides all the physical hardware components, mainly CPU, memory, storage devices and network interface cards. 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. Service Console or Console Operating System is a modified installation of Red Hat Enterprise Linux that provides administrators a set of management tools for the virtual infrastructure. (Marshall et al. 2009, 31-32). (McCain 2008, 1-2).

5.1.1 Physical Host Server

The purpose of the physical host server is to provide the hypervisor, namely VMkernel, a platform to execute. VMware designed the ESX and the ESXi to be highly efficient virtualization platforms. Therefore the list of supported hardware is also limited. This list is called Hardware Compatibility List (HCL) and it contains a list of supported server hardware from major manufacturers like Dell, Hewlett-Packard and IBM. The HCL can be obtained from vmware.com website. The list is updated with new releases of the ESX and the ESXi as compatibility and performance is re-evaluated. In order to efficiently perform the challenging task of virtualization, the physical host server must meet certain requirements for CPU, memory, storage and networking resources. To meet all these requirements it is important to verify that the hardware that the hypervisor is installed on can be found from the HCL. (Marshall et al. 2009, 32-33).

CPU resources provided by the physical host are responsible for executing the instructions from the VMkernel and the instructions from the virtual CPUs presented to the virtual machines by the VMkernel. In ESX and ESXi the minimal requirement for CPU is at least two single-core CPUs with clock speeds of 1,5GHz. A maximum number of physical cores per server is 32 and each core can represent 8 virtual CPUs to the virtual machines. The maximum number of virtual CPUs in a single host is 128 and therefore the theoretical maximum amount of virtual machines is also 128. (VMware 2009a, 1, 3).

As virtualization is a growing trend in the whole IT-world, major CPU manufacturers Intel and AMD have developed new virtualization technologies, Intel-VT and AMD-V. These hardware virtualization technologies allow virtualization platforms to handle certain unsafe

x86 instructions from the virtual machines more efficiently. Intel-VT and AMD-V will give virtualization a much needed boost in performance. (Marshall et al. 2009, 7, 33).

Physical memory provided by the host server is the amount of RAM the VMkernel presents to the virtual machines as well as the RAM to operate the VMkernel itself. HCL does not set any explicit minimal requirements for physical memory. However 1GB of RAM should be available for VMkernel and obviously some RAM is needed for the virtual machines. The maximum amount RAM for a single host is 256GB. Best practices are that virtualization hosts should have a large quantity of high-speed RAM as usually each guest server will require at least 1GB of RAM. (VMware 2009a, 3). (Marshall et al. 2009, 33-34).

Storage devices that the host server provides are the physical location for VMkernel and the virtual disks of virtual machines. The proprietary file system used by the ESX and the ESXi is a clustered file system called VMware File System (VMFS) that is optimized for storing and managing virtual machine files. VMware has released three versions of VMFS, the latest file system being VMFS-3. Earlier versions VMFS and VMFS-2 were flat file systems while version 3 offers a directory structure which enables concurrent reading and writing by multiple nodes or host servers. For storage deployment the ESX and the ESXi offer several options including directly attached storage (DAS), network attached storage (NAS), fibre channel storage and Internet Small Computer Systems Interface (iSCSI) storage. The storage system used must also be found in the HCL to keep the virtualization platform running efficiently. VMware has optimized and certified VMFS to effectively work with all of the storage options mentioned in the HCL. VMFS is designed for virtualization environments and it offers capabilities such as live migration of powered on virtual machine from one host to another, clustering virtual machines across hosts for load balancing and high availability of virtual machines for fault tolerance. These advanced features will be discussed in more detail later but it's useful to state that they are not available if local storage, or DAS, is used as only one physical server can access this type of storage. (Marshall et al. 2009, 18, 34-35, 145-150). (McCain 2008, 87-88, 91).

Physical hosts Ethernet adapters and other network equipment provide the network resources for virtual networking and NIC teaming. Networking plays a vital role in the setup. HCL list from VMware contains a list of supported network adapters and if an adapter is not mentioned

in the HCL, the ESX or the ESXi will often not know how access physical network resources. VMware recommends that multiple physical networks adapters are implemented because the host server can be configured to provide redundancy, flexibility and load balancing if more than one network adapter is used in a single host. (Marshall et al. 2009, 35-36).

5.1.2 VMkernel

VMkernel is hypervisor code with device driver modules and it provides the virtualization layer. VMkernel can be referred as the main operating system of the ESX and ESXi as it is responsible for operating the physical host and managing the access of the virtual machines to the underlying hardware. VMkernel is a microkernel that provides executions for the virtual machines, such as CPU scheduling, memory management, disk management and networking functions. (Marshall et al. 2009, 38-39).

VMkernel uses share based priorities to provide CPU, memory, disk and network resources to the Service Console and the virtual machines. If all virtual machines and the Service Console are configured equal shares they all have the same opportunity to access the resources. Resource shares can be configured individually for each virtual machine and the Service Console. This means that if a virtual machine or the Service Console is allocated with more shares than others, it will have a better opportunity to access the resources. In addition VMkernel allows CPU and memory resources to be reserved or limited. If a virtual machine is configured with a 500MHz CPU and 1024 MB memory reservations, this amount of CPU and memory resources are always guaranteed for this virtual machine. If resource limits are configured, the resources this virtual machine can use is limited as defined. (Marshall et al. 2009, 39).

VMkernel offers a unique feature that allows over commitment of memory resources. The technique is called Transparent Page Sharing and it allows VMkernel to keep a single copy of memory instructions that several virtual machines have in common, instead of a copy for each virtual machine. This is very useful when many virtual machines are running the same operating system and have the same applications or services running. For example if five virtual machines have the same information running in their memory space, only a single copy would

be loaded into the physical host's memory, rather than five copies. To the virtual machines themselves the process is transparent. (Marshall et al. 2009, 39).

VMkernel accesses physical network adapters through device modules based on Linux device drivers. VMkernel needs to provide network access to the virtual machines, Service Console and VMkernel itself. Service Console will need connectivity to the physical network for remote administration of the ESX and the ESXi. VMkernel's physical network connection is needed for configuring all of the above as well as accessing possible external storage such as iSCSI system. (Marshall et al. 2009, 39-40).

5.1.3 Service Console

The roles of the Service Console are different in ESX and ESXi. In ESX the Service Console is based on RedHat Enterprise Linux 3 operating system and it works as a boot loader, acts as a text based administration interface and serves as a management point for the VirtualCenter. (Marshall et al. 2009, 36-38).

In ESXi the Service Console is very minimal and the text based administration interface has been removed. The Service Console still works as a boot loader and a management point for the VirtualCenter just like in ESX. (Marshall et al. 2009, 36-38).

As a boot loader the Service Console is responsible for the initial boot process of the ESX and the ESXi. During the boot process the Service Console is loaded and eventually it will hand off the role of operating system to the VMkernel. In an ESX host, after forwarding this role, the Service Console behaves as a virtual machine providing access as a management interface for the host. (Marshall et al. 2009, 36-38).

The Service Console as an interface is a text based administration console to configure operating parameters for the ESX host, viewing system utilization as well as running local and remote scripts that give flexibility to the configuration and usage of the host. The Service Console can be accessed using physical keyboard and monitor connected to the host server or by using a terminal application that supports Secure Shell protocol. As mentioned before, the

ESXi does not have a text based administration console. The management of the ESXi is done only with remote management tools. (Marshall et al. 2009, 37).

As a management point role for VirtualCenter, the Service Console works in co-operation with the VirtualCenter to enable additional features provided by Foundation, Standard and Enterprise editions of VMware Infrastructure 3. These features are VMotion, VMware High Availability, VMware Distributed Resource Scheduling and VMware Consolidated Backup. These features, VirtualCenter and their requirements are covered in chapter 5.2. Communication between the Service Console and the VirtualCenter is handled by VirtualCenter Agent that is installed to the host. The Service Console keeps VirtualCenter up-to-date for management, statistical information and use of advanced features. If the Service Console fails to communicate with VirtualCenter, the ESX or the ESXi host is considered to be offline and features like VMware High Availability are initiated in order to keep the virtual machines available for the VMware Infrastructure 3 environment. (Marshall et al. 2009, 36-38).

5.2 Features

The VMware Infrastructure 3 suite offers several features that make up the full feature set of enterprise virtualization. Four different editions are available: Free edition with ESXi, VI Foundation, VI Standard and VI Enterprise. VMware Infrastructure 3 editions and their features can be seen in Table 2. 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. (Marshall et al. 2009, 101).

Table 2. VMware Infrastructure 3 editions
(VMware Homepage 2009b)

	ESXi – Free License (ESX not available without VI)	VI Foundation (with ESX or ESXi)	VI Standard (with ESX or ESXi)	VI Enterprise (with ESX or ESXi)
Core hypervisor functionality	Yes	Yes	Yes	Yes
Virtual SMP	Yes	Yes	Yes	Yes
VMFS	Yes	Yes	Yes	Yes
VirtualCenter 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

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. (McCain 2008, 2).

Virtual Symmetric Multi Processing (Virtual SMP) enhances virtual machine performance by enabling a single virtual machine to use multiple physical processors simultaneously. Two-way SMP with two virtual processors and four-way SMP with four virtual processors are available. Virtual SMP is a basic feature of the suite and is included in every edition. (VMware 2007c, 1).

VMFS was discussed in chapter 5.1.1. VMFS is also a basic feature of the VMware Infrastructure 3 and is included in every edition of VMware Infrastructure 3 product suite.

VirtualCenter is a centralized management utility for all VMware Infrastructure 3 hosts and virtual machines within the virtual infrastructure. VirtualCenter 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 VirtualCenter Agent. VirtualCenter and its agents are bundled with VI Foundation, VI Standard and VI Enterprise. (McCain 2008, 3).

VMware Update Manager automates patches and update management for VMware Infrastructure 3 hosts and virtual machines. It is designed to track patch levels and manually applying security and bug fixes. Update Manager is available with VI Foundation, VI Standard and VI Enterprise. (VMware 2008a, 3).

VMware Consolidated Backup (VCB) is a Windows application that provides a Fibre 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. After the files are mounted, the entire virtual machine or individual files can be backed up to a chosen directory or backup tape. VCB is bundled with VI Foundation, VI Standard and VI Enterprise. (McCain 2008, 6).

VMware High Availability is a feature that will power on any virtual machines that were previously running on another VMware Infrastructure 3 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. VMware High Availability is available with VI Standard and VI Enterprise. (McCain 2008, 6).

VMware VMotion provides live migration of virtual machines from one VMware Infrastructure 3 host to another. The migration is done with zero down time and continuous virtual machine availability. VMotion is very handy for example when a server needs to be powered off for hardware maintenance. VMotion is bundled only with VI Enterprise. (McCain 2008, 5).

VMware 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. Like VMotion, Storage VMotion is only available with VI Enterprise. (McCain 2008, 5).

VMware Distributed Resource Scheduling (DRS) will continuously monitor VMware Infrastructure hosts to maintain balanced utilization across the load on all servers within the virtual infrastructure by using VMotion technology to migrate virtual machines from one host to another. DRS can be configured to automate the placement of each virtual machine as it is

powered on as well as to manage the virtual machine's location once it is running. DRS is only available with VI Enterprise. (McCain 2008, 5).

Distributed Power Management (DPM) consolidates workloads by monitoring the resource requirements and power consumption of VMware Infrastructure 3 hosts and places unused hosts in standby mode to reduce power consumption. When the work load increases, DPM brings the hosts in standby back online to ensure service availability. DPM is also available with VI Enterprise. (Marshall et al. 2009, 21).

5.3 Virtual Networking

VMware Infrastructure 3 brings a rich set of virtual networking elements. The main elements are virtual networks, vSwitches, vNICs and port groups. The network architecture of the VMware Infrastructure 3 is shown in Figure 12. This chapter discusses the main elements of virtual networking and describes some related advanced features.




Figure 12. VMware Infrastructure 3 network architecture
(VMware 2008b, 21)

A virtual network is a network of virtual machines that are logically 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 VMware Infrastructure 3 host. (VMware 2009b, 22).

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. VMware 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:

- VMXNET – a paravirtualized vNIC that works only with VMware Tools installed 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 3 host. Services clients VMotion, Network File System (NFS) and software iSCSI. Also serves the remote console traffic.

(VMware 2008b, 21). (VMware 2007b, 4).

In a virtual environment a port group is a mechanism for setting policies that controls the network 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. (VMware 2008b, 22).

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.

(VMware 2009b, 22-23). (VMware 2008b, 22).

5.4 VMware ESXi 3.5

VMware ESXi 3.5 is a hypervisor that operates independently from any general-purpose operating system. The architecture is designed for direct integration into a virtualization-optimized hardware, enabling rapid installation, configuration and deployment. Functionally ESXi is equivalent to the ESX offering the same levels of performance and scalability. It supports the entire VMware Infrastructure 3 suite of products and features described in chapter 5.2. (VMware 2007d, 3).

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. The functionality of the Service Console is replaced by new remote command line interfaces and management tools. The small footprint behavior has major benefits over the full Service Console installation, namely better security and faster deployment of the hosts. In ESX, updates have primarily addressed the security vulnerabilities with the Service Console. As the Service Console of ESXi is almost non-existent, the number of security patches is less frequent. Therefore the downtime of the server is reduced significantly. (Marshall et al. 2009, 26). (VMware 2007d, 3).

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.
- VMM process – provides the execution environment for virtual machines.
- Various management agents – used to enable VMware Infrastructure 3 management from remote applications.
- The Common Information Model (CIM) – an interface that enables hardware-level management from remote applications. CIM is an open standard that defines how computing resources can be represented and managed.

(VMware 2007d, 3).

5.4.1 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 VMware.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 VMware Infrastructure Client (VI Client) and VMware Tools images.
- A 110MB core dump partition, normally empty but can hold diagnostic information if problems occur.

The ESXi system has two independent 48 MB boot banks for fail-safe hypervisor updating. While updating the new version is loaded into the inactive bank and the system is set to use the updated one after the reboot. If problem occurs, the system automatically boots from the previously used bank of memory. Also two versions of VI Client and VMware Tools are in the store partition corresponding to the hypervisors in the two banks. (VMware 2007d, 6).

5.4.2 ESXi Startup

When the ESXi boots for the first time, the VMkernel discovers hardware devices and selects appropriate drivers for them. Local hard drives are also discovered and if an empty drive is discovered, it will be formatted to VMFS so it can be used to store virtual machines. During the first boot VMkernel creates configuration files with default values. After the boot up process users can modify the default values using the DCUI or other management tools, as defined in the next chapter. In a normal startup of the ESXi first the hardware drivers are loaded, then the various agents are started and finally the DCUI process is started. (VMware 2007d, 6).

5.4.3 ESXi Management

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.

(VMware 2007d, 5).

Because the ESXi does not include a service console like the ESX, many of the management activities are done using remote management tools. The remote management tools for the ESXi are the VI Client and the remote CLI tools. (VMware 2007d, 6).

VI Client provides a Windows-based graphical user interface with interactive configuration capabilities. VI Client provides sophisticated configuration options that were previously available only in the Service Console. VI Client can be downloaded and installed from the host's website. The VI Client can also manage the VirtualCenter and therefore all the ESXi and ESX hosts in the virtual infrastructure can be managed from a single VI Client. Chapter 6 shows how a virtual infrastructure is managed and configured using a VI Client connected to a single (VMware 2007d, 6).

Remote CLI tools provide an encrypted and authenticated channel to enable scripting and command-line-based configurations of the host and virtual machines. Although remote CLI tools are mainly intended for use with ESXi hosts, they are also fully compatible with ESX hosts. Figure 13 shows the management model of the ESXi and ESX. (VMware 2007d, 6).




Figure 13. ESXi management
(VMware 2007d, 6)

A single PC with remote management tools can manage all hosts through the VirtualCenter Server or just a single host through a direct connection. The same management model also applies for ESX hosts inside the virtual infrastructure.

5.5 Summary of Differences Between ESXi and ESX

ESXi and ESX are both bare-metal hypervisors from VMware and support the same VMware Infrastructure 3 features. The differences between these hypervisors reside in the architecture and the operational management. Although neither relies on an operating system for resource management, the ESX relies on a Linux-based Service Console to perform management functions. In ESXi the Service Console has been deprecated and management is mainly done remotely. Table 3 summarizes the main differences between ESXi and ESX. (VMware Homepage 2009b).

Table 3. ESXi and ESX comparison
(VMware 2007d, 8)

	ESXi	ESX
On-disk footprint	32 MB	2 GB
Bootstrap	Direct from boot loader	Service Console driven
Direct management	DCUI	Service Console shell session
Hardware monitoring agents	CIM plugins	Full applications in Service Console
Other agents	Only with VI SDK	Full applications in Service Console
Scripts, automation and troubleshooting	DCUI, remote CLI	Service Console shell session
Other software	Moved to outside environment	Resident in Service Console

In addition to the architectural differences shown in Table 3, one important difference between the hypervisors is cost. Although the same licensed features are available with both hypervisors, the ESXi is also available as a free standalone version which offers an easily approachable virtualization solution. The fact that ESXi can be obtained for free, makes it a good choice for companies who are planning their first virtualization project keeping in mind that there are also good options for future expandability. (VMware Homepage 2009b).

6 PRACTICAL IMPLEMENTATION

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. This chapter will discuss hardware, operating system and roles of the server under consolidation and the hardware of the destination server. 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. Chapter will show how to install and configure the ESXi installable edition to provide a platform for the virtual machines and the guest operating systems. A single virtual machine is created and Windows Server 2008 Standard Edition is installed as a guest operating system. Finally this practical part shows how the roles of the source server are migrated to the newly created guest operating system. The management of the ESXi host and virtual infrastructure is done with VI Client software installed on a Windows XP Professional computer connected to the company LAN.

6.1 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 be 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. The domain preparation procedure will be briefly explained in Chapter 6.7.

6.2 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.

6.2.1 Source

The physical server under consolidation is a Dell PowerEdge 600SC based on Intel Pentium 4 Processor. It is a tower server as shown in Figure 14 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
- Single 250 W power supply unit
- Single embedded Gigabit NIC
- IDE CD-ROM drive

**Figure 14. Dell PowerEdge 600SC
(Dell 2003)**

The only redundant components in this server are the RAID 1 mirrored hard drives so this server's hardware is quite outdated. If for example the power supply unit of this server would fail, the server would be unavailable until a new power supply unit was bought and installed. As the warranty of this server has already expired and older replacement parts are hard to get, this server could be out of production for weeks if a hardware failure would occur.

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

Services provided by the roles of this Domain Controller are replicated between other Domain Controllers in the domain as stated in the system description in Chapter 2.1.

6.2.2 Destination

The company has made a decision to buy server and workstation hardware directly from Dell. ESXi HCL offers several options from Dell server family, including Dell PowerEdge 2950 III (Figure 15). 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 was bought with 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



Figure 15. Dell PowerEdge 2950 III
(Dell 2007)

This PowerEdge 2950 III provides hardware fault tolerance with redundant power supply units, hard drive configurations and NICs. Six hot-pluggable hard drives are configured so that two 73GB drives are configured with RAID 1 and four 450GB drives are configured with RAID 5. This way the virtualization platform is provided with one logical drive with the size of 73GB and one with the size of 1350GB. DRAC is a controller card with its own Ethernet port and it provides comprehensive remote management tools that can be used via web browsers. All of the things described above are preinstalled and preconfigured during the manufacturing process by Dell.

6.3 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.

As mentioned in the previous chapter, there are two logical hard drives available for the virtual infrastructure. The 73GB drive will be called Datastore1 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. After downloading the image is burned to a CD and inserted into the server's CD/DVD-drive. Installing the Host is a fairly straightforward process: After booting up from the installation CD the only thing that has to be decided is to where the ESXi will be installed, in this case Datastore1. 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 thesis work the hostname will be "miprovr1". 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 16.



Figure 16. DCUI configurations

Figure 17 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.




Figure 17. DCUI

After the initial installations steps the management tools for the virtual infrastructure are installed 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 company LAN. After installation the management client is ready to be connected to the ESXi. The connection window of the VI Client is shown in Figure 18.




Figure 18. VI Client connection window

Connection to the host is established and the client is ready to start managing the virtual infrastructure. Figure 19 shows the main view of the management client. The main window provides general information about the host, such as information about the processors, memory, hard drives and networks as well as the number of physical NICs and virtual machines. The view also shows the current CPU and memory utilization. Before any virtual machines are created, the host's CPU usage is 217 MHz and memory usage is 553 MB. To ensure host's functionality the host is configured with 600 MHz CPU reservation and 600 MB memory reservation.

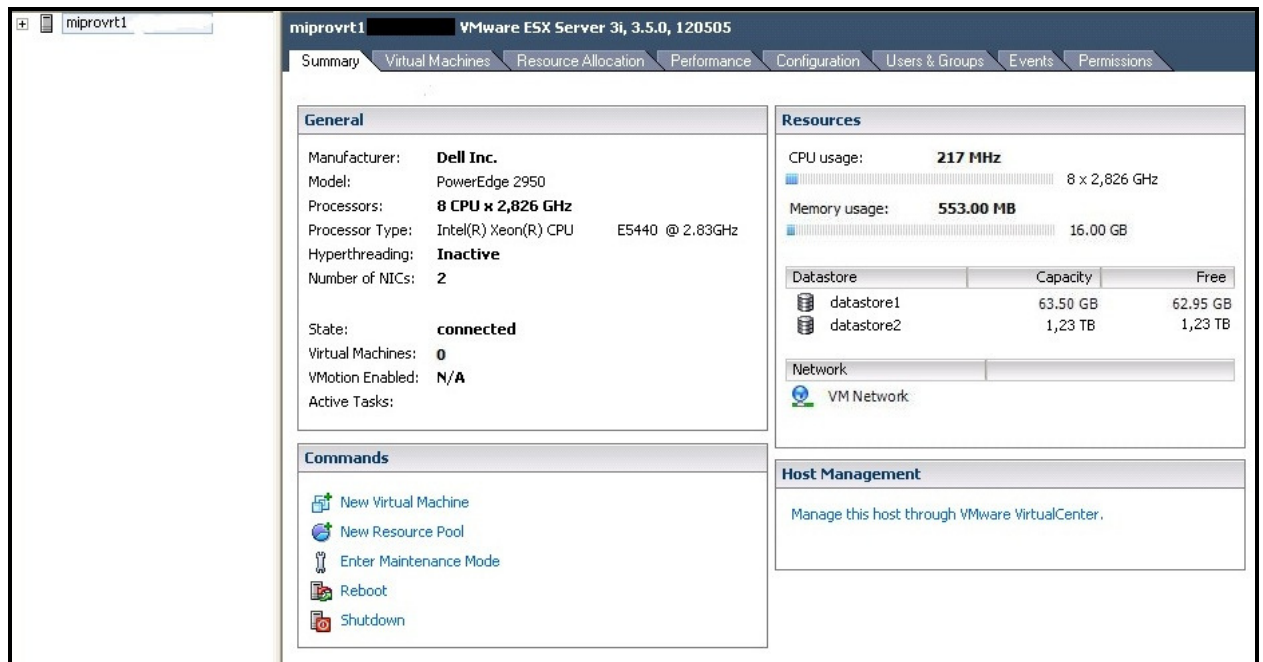


Figure 19. VI Client main view

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. Figure 20 shows the “Configuration” –tab of the VI Client. 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.

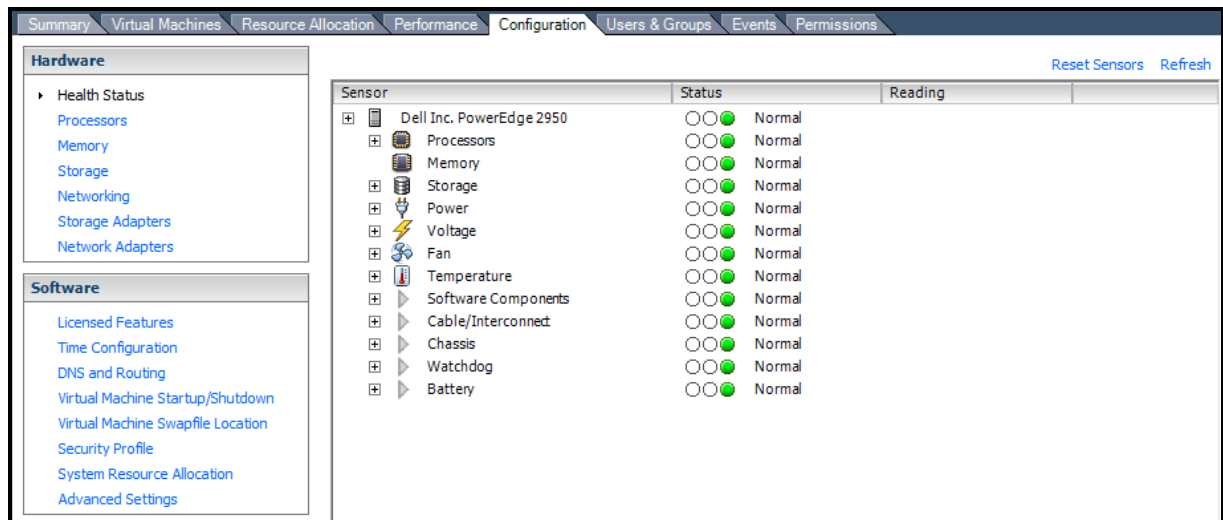


Figure 20. Configuration and Health Status –view

Before configuring the virtual networking, some initial configurations are made under “Licensed Features”, “Time Configuration” and “DNS and Routing” -tabs: Defining the license serial number, configuring the host’s access to a time server and configuring the DNS settings for the host. The free license serial number is obtained when the ESXi image is downloaded from vmware.com website. Configuring time and DNS -servers are done according to the existing company policies. The configurations are done so that two physical servers inside the company LAN provides the ESXi host with time and DNS information. After these configurations the host is ready to start defining virtual networks.

6.4 Defining Virtual Networks

In this thesis work the network for the virtual machines and the network for management use the same vSwitch connected to the physical network, this way the whole virtual infrastructure can be accessed from any computer connected to the company 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. This view shows the current configurations of the host as well as provides a way to create new virtual networks and modify existing ones. 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 company LAN. The Virtual Machine Port Group is named “Mipro LAN” and the VMkernel Port is named “Management Network”. Figure 21 shows the configured properties for vSwitch0.

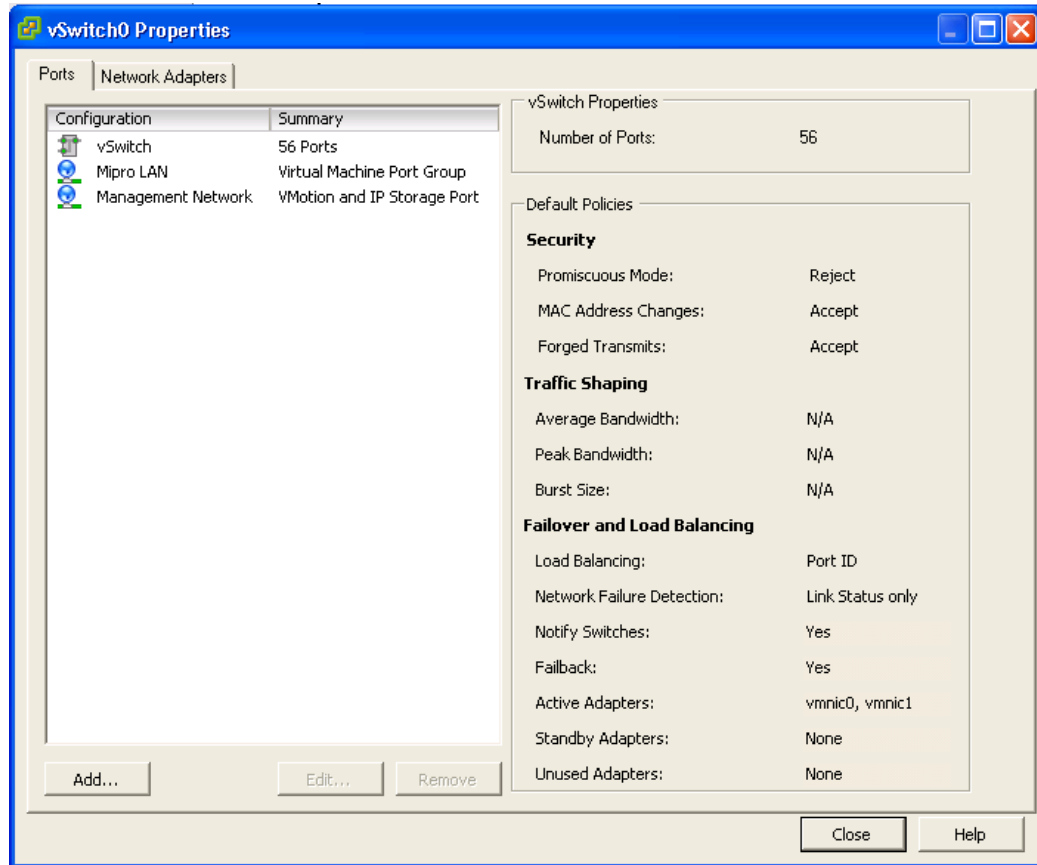


Figure 21. vSwitch0 Properties

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 22 displays the “Configuration → Hardware → Networking” view of the VI Client after the virtual networking configurations described above have been done.

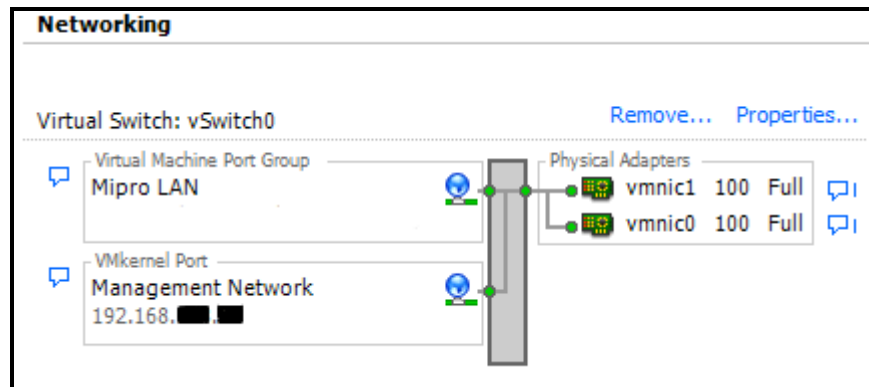


Figure 22. VI Client networking view

The virtual network for the virtual machines is now defined and connected to the physical network. All the network services available for computers connected to the physical switches of the company LAN are now also available for virtual machines that are configured to use the vSwitch0. Virtual infrastructure is now ready to start creating virtual machines.

6.5 Creating the Virtual Machine

Before creating a new virtual machine the resource requirements, reservations and limitations should be considered. The source server under consolidation is a standard Domain Controller without any heavy applications or services. The server's utilization is very low and therefore the destination virtual machine could be configured with the minimum requirements. As the available virtual hardware resources can be modified later and this will be the only virtual machine in this host for the time being, recommended system requirements will be used. According to Microsoft the minimum system requirements and recommendations for Microsoft Windows Server 2008 are:

Minimum:

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

Recommended:

- Single 2 GHz CPU (x86 or x64)
- 2048 MB of RAM
- 40 GB of available disk space

(Microsoft 2007).

As stated in Chapter 5.1.2, the VMkernel is also capable of resource reservation and limitation individually for each virtual machine. Especially reservation is a useful feature as necessary amount of resources is guaranteed for the virtual machine. Considering the resource recommendations, the virtual machine will have the following resources available for the guest operating system:

- Single virtual CPU with 400MHz resource reservation
- 2048 MB of RAM with 512MB resource reservation
- Single 40 GB hard drive
- CD/DVD drive
- Single NIC

Creating a new virtual machine is done with New Virtual Machine Wizard by selecting “File → New → Virtual Machine” from the VI Client main view. The New Virtual Machine Wizard offers several options to configure the virtual machine platform. There are two Wizard types available: “Typical” and “Custom”. The “Custom” –type is used as it provides more advanced configuration options. The following list will show the configuration steps and options selected for the virtual machine:

- Virtual machine name: **VRTSRV1**
- Host: **MIPROVRT1**
- Datastore where the virtual machine will be located: **DATASTORE2**
- Guest operating system: **Microsoft Windows Server 2008 (32-bit)**
- Number of virtual CPUs: **1**
- Amount of RAM: **2048 MB**
- Number of vNICs: **1**

- Virtual network for the vNIC: **Mipro LAN**
- Type of vNIC: **E1000**
- Storage adapter type: **LSI Logic**
- Create a new virtual disk with size of: **40 GB**
- Location of the virtual disk: **Store with the virtual machine**

Figure 23 shows the configuration summary for the virtual machine VRTSRV1.

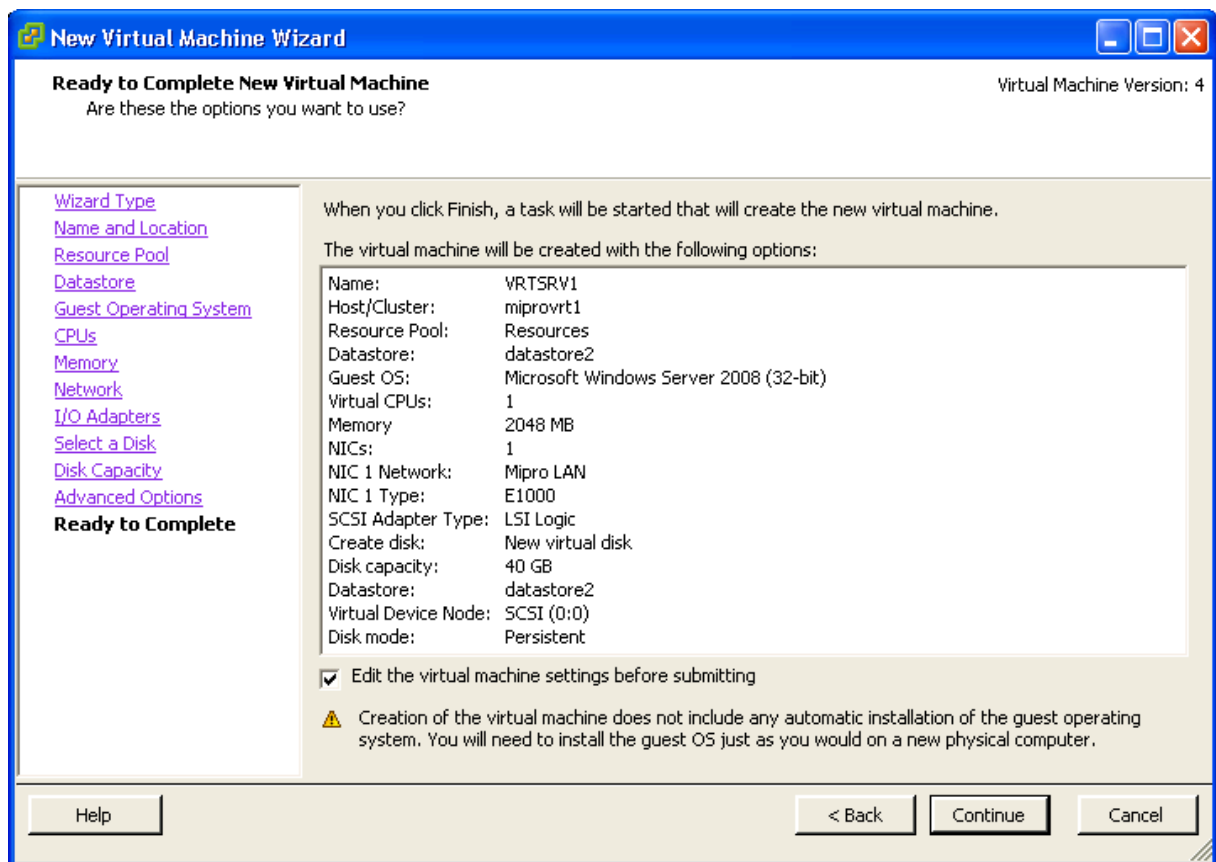


Figure 23. Virtual machine configuration summary

By checking the box “Edit the virtual machine settings before submitting” the settings are configured before creating the virtual machine in a way that virtual floppy drive is removed and resource reservations are defined. Floppy drive can be removed in the “Hardware” –tab and resource reservations can be defined in the “Resources” –tab. The virtual machine is configured with CPU reservation of 400MHz and memory reservation of 512 MB as shown in Figure 24.

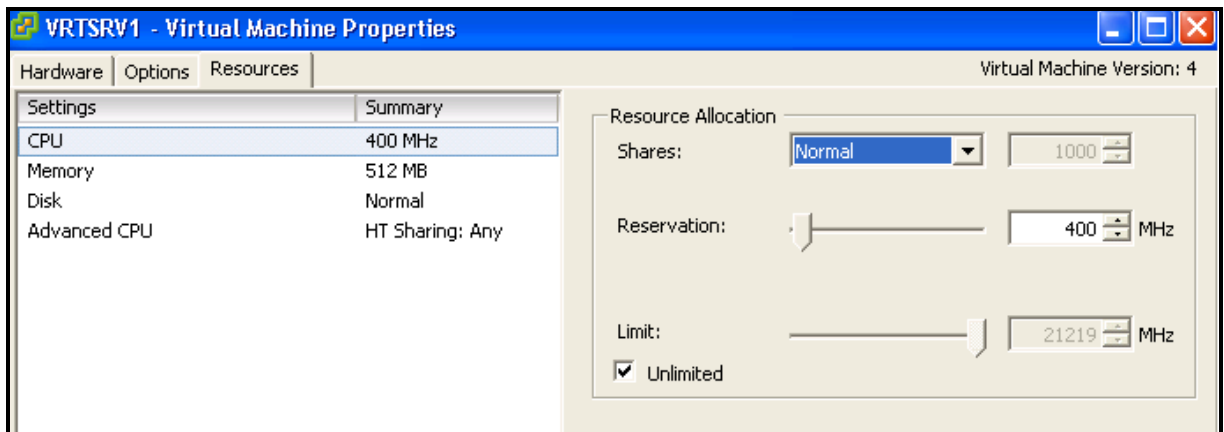


Figure 24. VRTSRV1 resource reservations

The virtual machine VRTSRV1 can now be powered on and the guest operating system installation can begin.

6.6 Installing and Configuring the Guest Operating System

A virtual machine operating system installation is almost identical with an installation into a physical machine. Operating system installation DVD is inserted into the computer using the VI Client and the DVD-drive is connected to the virtual machine. Virtual machine is powered on and configured to boot from the DVD-drive. The Windows Server 2008 installation procedure will begin. After the regional settings and the product key have been defined, the installation proceeds to choosing the hard drive where the operating system will be installed. The installation automatically identifies the 40 GB virtual drive so there is no need to load any additional drivers as seen in Figure 25.




Figure 25. Windows Server 2008 installation to the virtual hard drive

After the destination hard drive is chosen, the installation will begin. During the installation phase the virtual machine is rebooted several times. When the installation is complete and the operating system starts up for the first time, a new local administrator password must be defined. As the ESXi host provides its virtual machines with generic well known virtual hardware, there is no need for driver installations. The Windows Server 2008 identifies the hardware and installs all the required drivers automatically. The guest operating system is now installed and ready for initial configurations.

Configuring the operating system starts with VMware Tools installation for better performance and use convenience. VMware Tools can be installed choosing “VM → Install/Upgrade VMware Tools” from the VI Client user interface. To complete the installation of the VMware Tools the virtual machine must be rebooted.

By default the Windows Server 2008 operating system obtains a dynamic IP address from a DHCP server and the default hostname in this case is “WIN-EMGSXPRU57H”. The server is configured with a static IP address and a hostname “VRTSRV1”. At this time the server will not join the company domain. Next the most recent updates for the Windows Server 2008 operating system are downloaded and installed. Figure 26 shows “Initial Configuration Tasks” –window of the Windows Server 2008 operating system after the computer information have been provided and the updates have been installed.

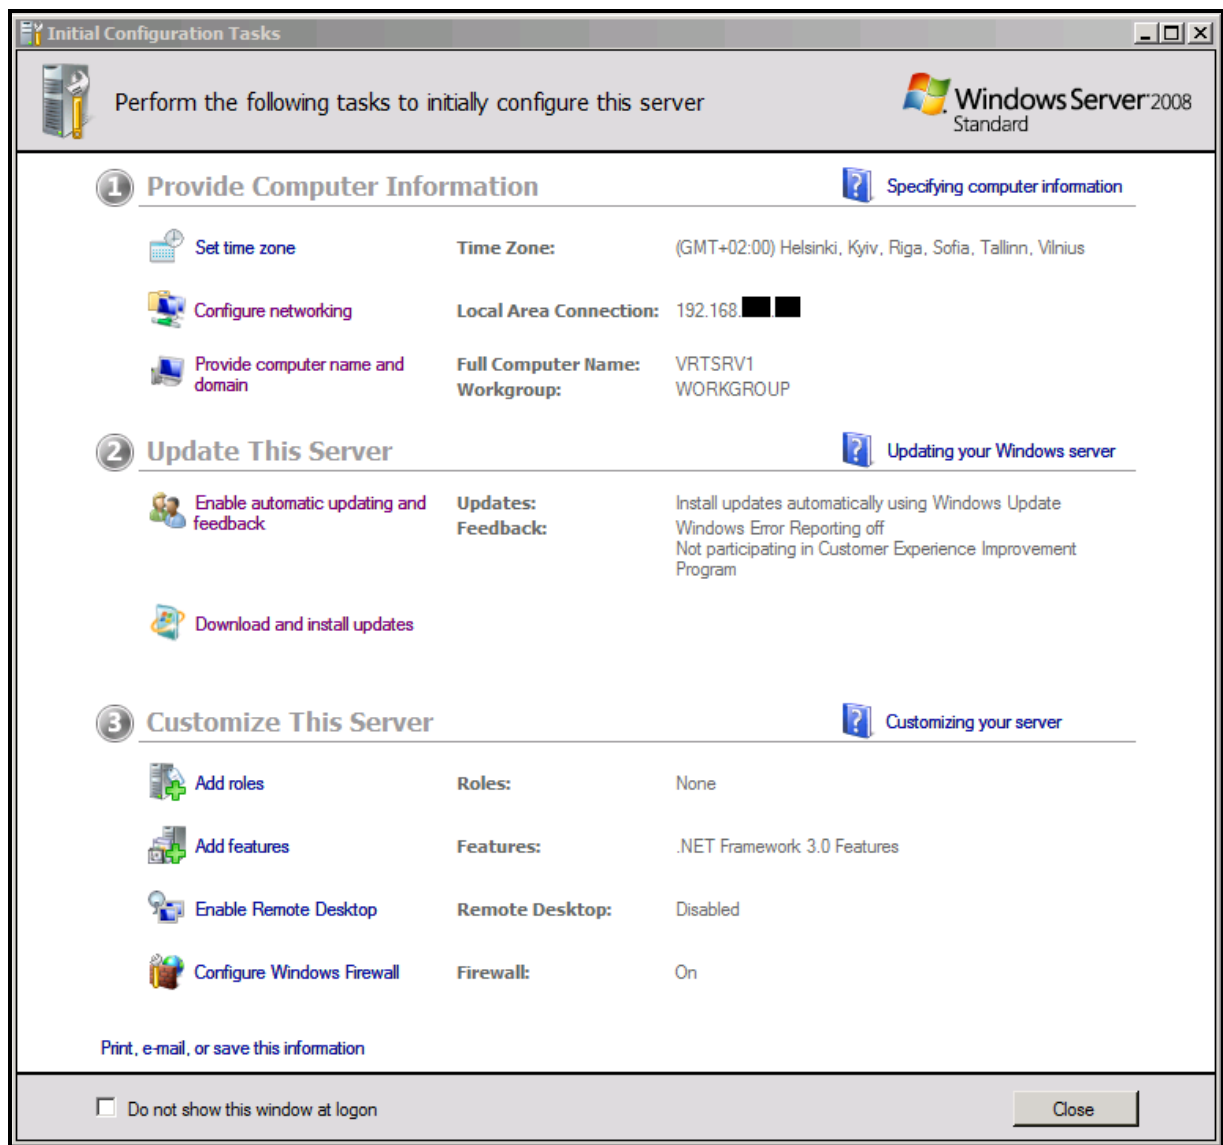


Figure 26. Initial Configuration Tasks –window

After these basic steps VRTSRV1 can be customized by adding roles and features.

6.7 Migrating Roles

This chapter will show how the roles of the source server are migrated into the newly installed VRTSRV1. The source server has three roles installed: Domain Controller, DNS Server and File Server.

In a Windows Server 2003 Domain environment the following five steps are required when the roles are migrated into the server running Windows Server 2008:

- Preparing the existing Windows Server 2003 Domain for a Windows Server 2008 Domain Controller.
- Adding the Active Directory Domain Services -role to the VRTSRV1
- Promoting the VRTSRV1 to function as a Domain Controller in the existing Windows Server 2003 Domain.
- Demoting the source Windows Server 2003 Domain Controller and removing the DNS Server -role.
- Removing the source Windows Server 2003 from the Domain and the physical network.

Preparing the existing Domain is carried out using a command line tool “ADPREP.EXE” which can be found from the Windows Server 2008 installation DVD. The preparation procedure must be executed with Active Directory Forest wide administrator privileges from a Domain Controller that holds the Schema Master and the Infrastructure Master roles. To prepare the Forest and the Domain, the following commands are executed from the command line of a Domain Controller that holds the appropriate roles:

- `adprep.exe /forestprep`
- `adprep.exe /domainprep`

(Microsoft 2009).

The existing Windows Server 2003 Domain environment is now prepared for the new Domain Controller. The next step is to add the Active Directory Domain Services –role to the VRTSRV1. In Windows Server 2008 all the server roles are added using the “Add Roles Wizard” seen in Figure 27.

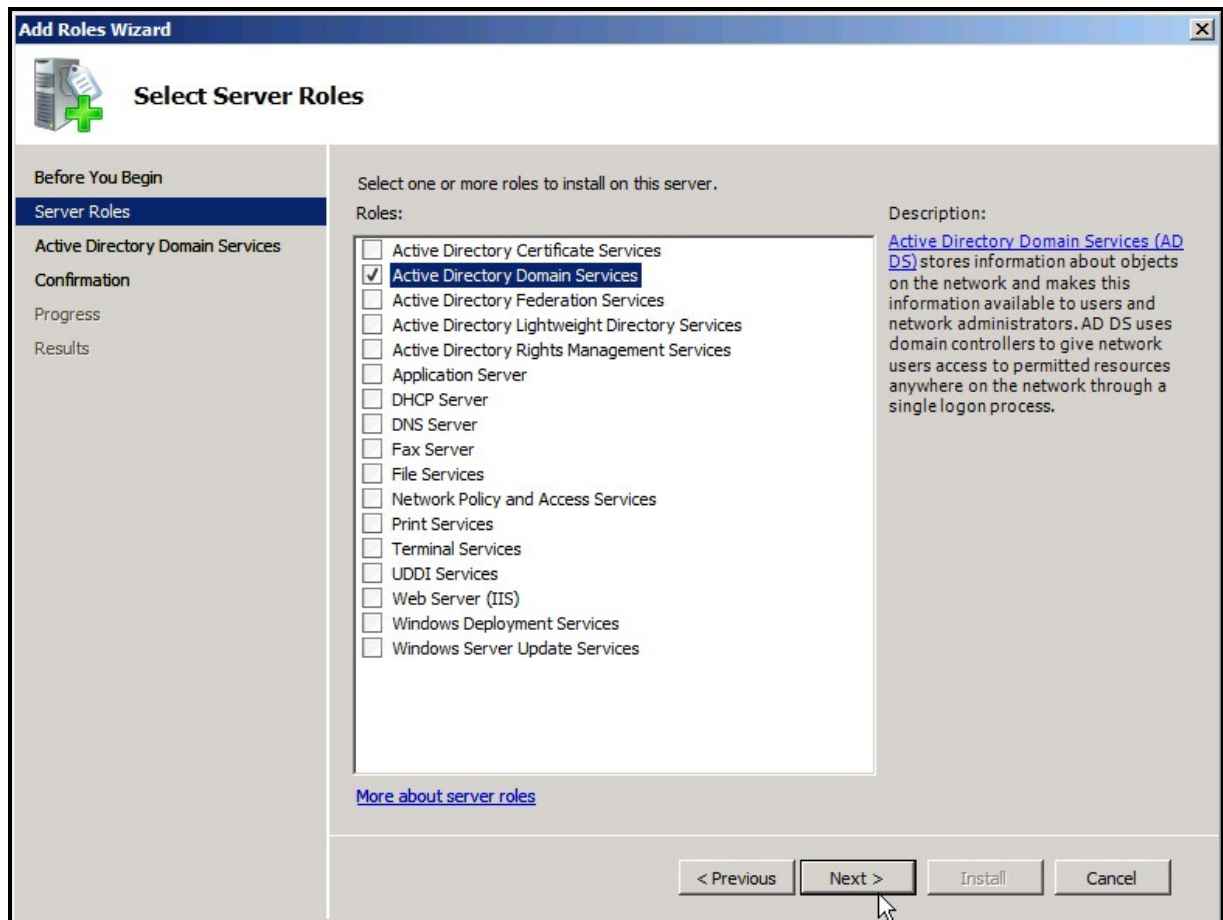


Figure 27. Add Roles Wizard

Installing the Active Directory Domain Services -role is a straightforward process without additional configuration options. All the configurations are defined during the server promotion process in the next step.

After the Domain Controller –role is installed the “Active Directory Domain Services Installation Wizard” is used to promote the server to function as a Domain Controller. The wizard will start by running “DCPROMO.EXE” from the command line. The following list shows the configuration options that are defined during the Active Directory Domain Services Installation Wizard:

- The Domain Controller is added to an existing Domain and Forest.
- Existing Domain name and Domain Administrator credentials are specified.
- Additional options specified for the Domain Controller: **DNS Server and Global Catalog.**
- Active Directory data is replicated over the network from an existing Domain Controller.
- Active Directory Database and Log folder: **C:\Windows\NTDS**
- Active Directory SYSVOL folder: **C:\Windows\SYSVOL**
- Directory Services Restore Mode Administrator account password is defined.

Rebooting the server will complete the installation of Active Directory Domain Services and after the reboot VRTSRV1 is a fully functional Domain Controller. Figure 28 shows the “Server Manager” -window that displays the roles installed on VRTSRV1.

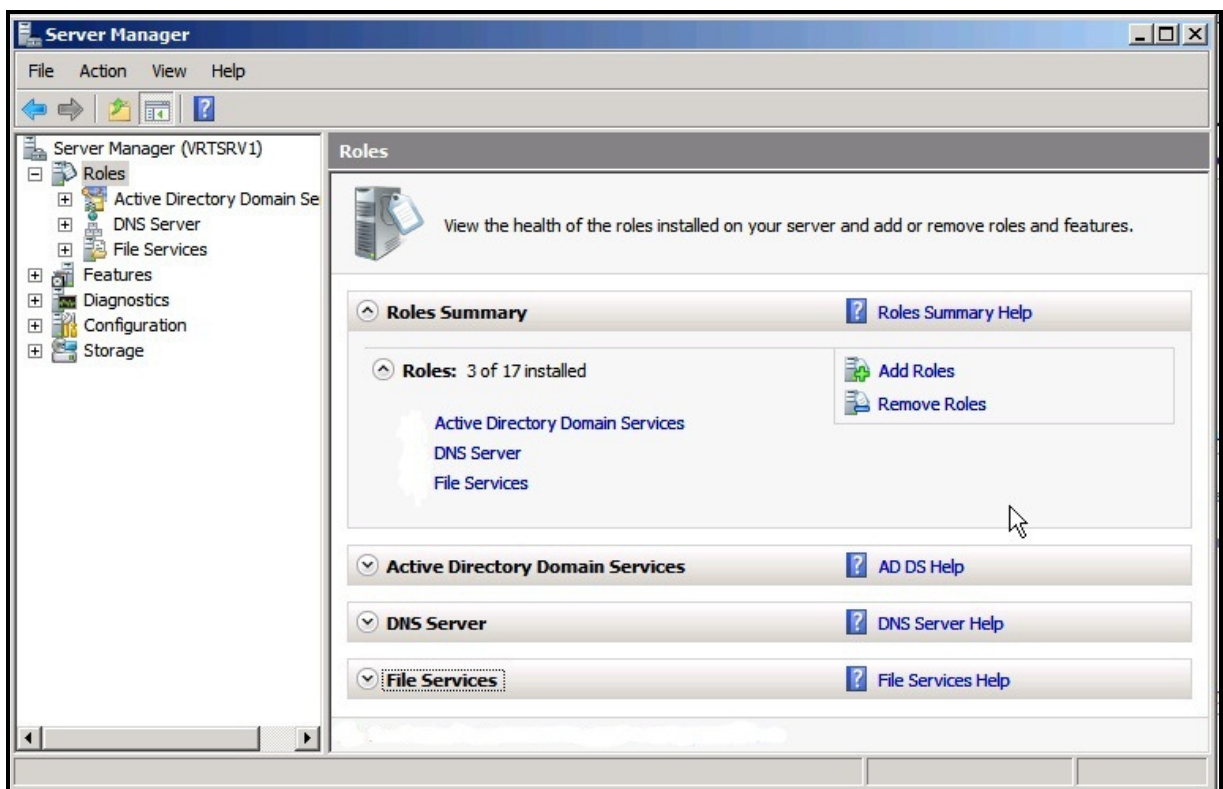


Figure 28. VRTSRV1 roles

The destination server VRTSRV1 now holds the same roles and services as the source server does. The next step is to demote the source server so that it will no longer function as a Domain Controller and then remove it from the Domain and physical network.

Demoting the source Domain Controller is carried out using a command “DCPROMO.EXE” that is executed from the command line with Domain Administrator credentials. Running the command will open the “Active Directory Installation Wizard”. The installation wizard detects that the server already is a Domain Controller and automatically starts the “Remove Active Directory” –procedure, shown in Figure 29. The option “This server is the last Domain Controller in the Domain” is not selected. During the “Remove Active Directory” –process the demoted server has to be defined with a new local administrator password.

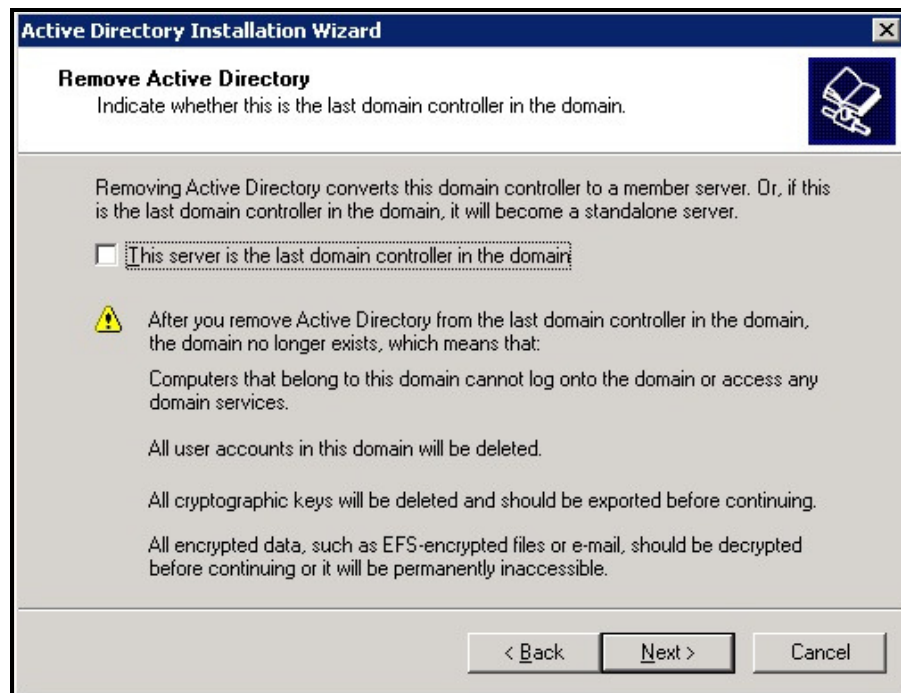


Figure 29. Remove Active Directory

To complete the removal of Active Directory the server has to be rebooted. After the reboot the server functions as a Member Server of the Domain. After the removal of Active Directory the DNS Server –role is removed using the “Add or Remove Roles” –process of Windows Server 2003. The Active Directory and DNS –databases are now up to date as the destination server holds the appropriate roles and the source server is functioning as a Member Server only.

The source server can now be removed from the Domain. Removing the server from the Domain is done by renaming the server and changing the membership to a workgroup instead of a Domain. Figure 30 shows that the name of the server is changed to “OLDSRV” and membership to “WORKGROUP”.

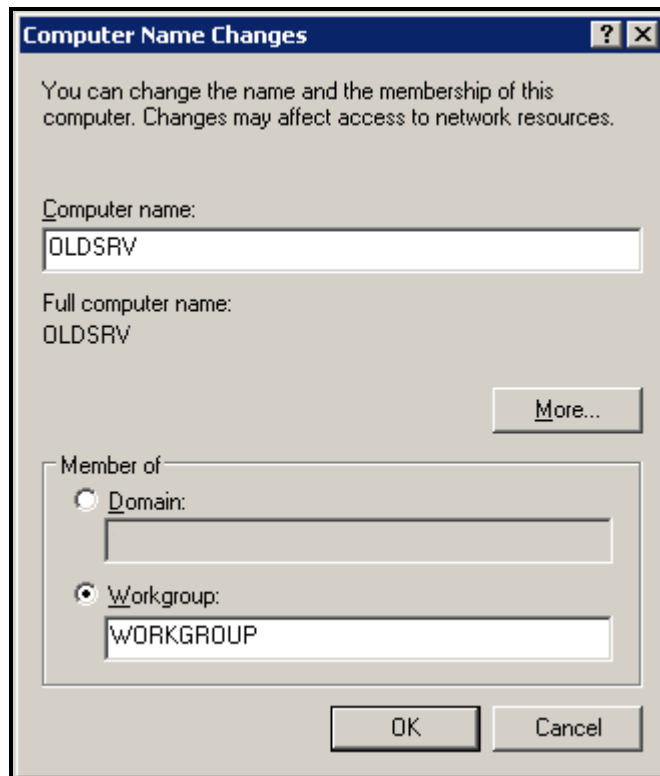


Figure 30. Computer name and membership

To finalize the removal process, the server is removed from the Active Directory using the “Active Directory Users and Computers” –tool with Domain administrator privileges. The server can now be powered off and removed from the physical network.

The role migration is now complete. The destination server has all the same features and services as the source server did before the migration process and the source server is completely removed from the network. The next chapter will include some monitoring and discussion about performance of MIPROVRT1.

6.8 Monitoring and Performance of MIPROVRT1

This chapter consists of performance monitoring of the host MIPROVRT1 and a brief analysis of the collected data. Monitoring is done using the “Performance” –tab of the VI Client. With the free version of ESXi, the performance data can be collected using a one hour interval. VI Client features monitoring for CPU, memory, network, and hard drives as well as information and health status of the host’s hardware. In this chapter the monitored resources are CPU usage (MHz), memory usage (MB) and disk usage (KBps).

6.8.1 Performance Data

The hardware performance of MIPROVRT1 is monitored in two separate one hour periods of time:

- When only the ESXi hypervisor is running
- When the virtual machine VRTSRV1 is running on top of MIPROVRT1

In addition the memory usage and consumption of the virtual machine VRTSRV1 is monitored when it’s functioning as a fully operational Domain Controller.

The results of the collected data during the monitoring periods are represented by screenshots from the VI Client and the results are shown below. Figures 31-33 show the performance data of the host when only the ESXi hypervisor is running and Figures 34-36 show the host’s data when the virtual machine VRTSRV1 is running. Figure 37 shows the memory usage and consumption of VRTSRV1.




Figure 31. Hypervisor only CPU usage

The average CPU usage of the ESXi hypervisor is 70,32 MHz or 0,31 percent. Peaking at 249 MHz, the maximum CPU usage of the hypervisor is 1,09 percent.




Figure 32. Hypervisor only memory usage

On the average the ESXi hypervisor uses 607 MB or 3,62 percent of the host's physical memory. The maximum memory usage is 620 MB or 3,69 percent.




Figure 33. Hypervisor only disk usage

Disk usage of the hypervisor is minimal as during the one hour measurement period the host's average disk usage is only 0,45 KB per second. Figure 33 shows that the disk is idle for almost the whole one hour monitoring period.




Figure 34. CPU usage with VRTSRV1 running

With VRTSRV1 running the host's average CPU usage increases to 104,82 MHz or 0,46 percent. Compared to results seen in Figure 31 the average CPU usage is increased by 34.5 MHz or 0,15 percent and the maximum CPU usage by 105 MHz or 0,47 percent.




Figure 35. Memory usage with VRTSRV1 running

The host's average memory usage with VRTSRV1 running is 2459 MB or 14,66 percent. Compared to the hypervisor only results the increase in average memory usage is 1852 MB or 11,04 percent.




Figure 36. Disk usage with VRTSRV1 running

With VRTSRV1 running the host's average disk usage is 12,34 KB per second and shows an increase of 11,89 KB per second compared to hypervisor only disk usage.




Figure 37. VRTSRV1's memory usage and consumption

On the average VRTSRV1 uses 87,6 MB or 4,17 percent of the memory allocated for its use. The average memory consumption on the host is 1895 MB.

6.8.2 Analysis of Performance Data

The results show that the ESXi hypervisor's maximum CPU usage is 1,09 percent and maximum memory usage is 3,69 percent. This means that MIPROVRT1 can continuously provide at least 98,91 percent of its physical CPU resources and 96,31 percent of its physical memory to the virtual machines running on top of it. The amount of physical memory installed on the host hardware is 16384 MB, therefore the amount of memory available for the virtual machines is 15779 MB. The disk usage of the hypervisor is very minimal as 15000 RPM SAS drives typically offer throughputs of at least 100 MB per second.

Powering on VRTSRV1 will naturally increase the resource usage on the host. CPU and disk usage are not greatly increased as it is typical for a Domain Controller not to significantly use these resources. This is true especially in an environment where Domain Controller services are redundant. Memory usage increases much more significantly, by 11,04 percent, because the amount of memory allocated to VRTSRV1 follows the recommendations set by Microsoft.

Chapter 5.1.1 stated that the maximum number of virtual CPUs for a single host is 128. This means that if every virtual machine uses a single virtual CPU the host could theoretically have 128 virtual machines running on top of it. Considering the amount of CPU and disk usage data this system could even host the theoretical maximum number of virtual machines if all guests were identical to VRTSRV1. However the performance data clearly shows that the bottleneck of this environment is the amount of physical memory. Figure 35 shows that on the average the VRTSRV1 consumes 1895 MB of the host's physical memory when the Microsoft hardware recommendation for the amount of memory is used. Therefore the maximum number of Windows Server 2008 virtual machines with 2048 MB of allocated memory would be limited to eight as they would consume 15160 MB or 96,1 percent of the physical memory available for virtual machines.

As stated in Chapter 6.5 the VRTSRV1 is allocated with 2048 MB of memory because it is currently the only virtual machine on top of the host. The performance data shows that VRTSRV1 could easily be allocated with a significantly lesser amount of physical memory as the maximum memory usage is only 167,8MB or 7,99 percent. The amount of memory allocated to VRTSRV1 can be reduced down to the minimum requirement level of 512 MB with-

out any significant loss in performance. In theory the host could have four times more virtual machines running if they were identical in resource usage. With 16 GB of physical memory MIPROVRT1 could theoretically host 32 virtual machines with 512 MB of allocated memory.

It is harder to say how many virtual machines could be hosted by MIPROVRT1 in practice as each virtual machine has different requirements and hardware utilization levels. Considering the results above a decent amount of hosted virtual machines with 512 MB of memory could be between 20 and 25. These numbers are of course speculative and in practice the amount of hosted machines is determined experimentally.

The performance data shows that a single somewhat powerful server hardware with VMware ESXi hypervisor offers a very good platform for server consolidation. If consolidated servers were highly underutilized this single server could ideally host up to 25 guest machines. Without any doubts this amount of guest operating systems on a single hardware will produce significant savings in hardware, electricity and cooling costs.

7 CLOSING WORDS

At the time of writing this, the ESXi host has been online for about one year 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 host's resource usage levels are still at a fairly low level. As an one-day average the CPU usage is 8,05 percent, memory usage is 50,77 percent, disk usage is 2,28 MB per second and network usage is 0,6 MB per second. This data shows that with the current hardware the system could host around 10 additional virtual machines. Installing additional 16 GB of memory this host would have potential to host at least 15 additional virtual machines with low physical utilization levels. All together this would mean that with 32 GB of memory this single hardware could host around 35 virtual machines.

It is evident that VMware ESXi 3.5 is a powerful virtualization platform with versatile features and lots of scalability. VMware ESXi 3.5 Free Edition provides a good foundation for

virtual 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 availability of services, there is a need to purchase redundant hardware and at least VMware Infrastructure 3 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 hardware used could host up to 20 virtual machines using VMware ESXi 3.5 virtualization platform. As a by-product this project provided Mipro Oy's software developers a very flexible and effective testing environment. Since the implementation of the ESXi no new server hardware has been purchased and three older physical servers have been consolidated successfully.

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 discusses the topic objectively. Once suitable source material was found the project was carried out without any bigger problems.

REFERENCES

Books

Cross, Michael – Martin, Jeffrey A. – Walls, Todd A. – Grasdal, Martin – Littlejohn Shinder, Debra – Shinder, Thomas W. 2003. Planning, Implementing, and Maintaining a Windows Server 2003 Active Directory Infrastructure Study Guide & DVD Training System. Rockland, Massachusetts. Syngress Publishing, Inc.

Golden, Bernard – Scheffy, Clark 2008. Virtualization for Dummies, Sun and AMD Special Edition. Indianapolis, Indiana. Wiley Publishing, Inc.

Haletky, Edward L. 2008. VMware ESX Server in the Enterprise: Designing and Securing Virtualization Servers. Crawfordsville, Indiana. Pearson Education, Inc.

Marshall, David – Beaver, Stephen S. – McCarthy, Jason W. 2009. VMware ESX Essentials in the Virtual Data Center. Boca Raton, Florida. Taylor & Francis Group, LLC.

McCain, Chris 2008. Mastering VMware Infrastructure 3. Indianapolis, Indiana. Wiley Publishing, Inc.

Muller, Al – Wilson, Seburn 2005. Virtualization with VMware ESX Server. Rockland, Massachusetts. Syngress Publishing, Inc.

Rule, David – Dittner, Rogier 2007. The Best Damn Virtualization Book Period. Burlington, Massachusetts. Syngress Publishing, Inc.

Electronic sources

Chappell, David 2008. Virtualization for Windows: A Technology Overview. [online][referred 3.3.2009]. Available in www-format: < <http://download.microsoft.com/download/E/0/4/E04E1183-E5FC-4C1B-A2AA-D89E3F4D6E4C/MSVirtualizationOverviewv3dot0.docx>>.

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>>

Dell 2003. Dell PowerEdge 600SC Server. [online][referred 9.8.2009]. Available in www-format: < <http://www.dell.com/downloads/emea/products/pedge/600sc.pdf>>.

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>.

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>.

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>.

Microsoft 2007. Windows Server 2008 System Requirements. [online][referred 3.9.2009]. Available in www-format: < <http://msdn.microsoft.com/en-us/windowsserver/cc196364.aspx>>

Microsoft 2009. Scenario for Installing AD DS. [online][referred 12.9.2009]. Available in www-format: < <http://technet.microsoft.com/en-us/library/cc771433%28WS.10%29.aspx>>.

Mipro Homepage 2009. MIPRO OY (Ltd.) - System Integrator. [online][referred 24.1.2009]. Available in www-format:< http://www.mipro.fi/mvhome/yritys_e.html >.

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

Otey, Michael 2006. Server Consolidation Essentials. [online][referred 8.4.2009]. Available in www-format:
<<http://windowsitpro.com/ebooks/index.cfm?fuseaction=ebook&ebid=bac1e8b0-640a-460f-8313-1ab25a86a8cc>>.

Parallels Homepage 2009. OS Virtualization. [online][referred 1.2.2009]. Available in www-format: < <http://www.parallels.com/eu/products/virtuozzo/os/>>

Sun Microsystems – VMware 2008. The Business Benefits of Virtualization and Consolidation with Sun and VMware. [online][referred 10.2.2009]. Available in www-format:
<www.sun.com/sunray/whitepapers/sun_vmware_benefit.pdf>.

Waters, John K. 2007. ABC: An Introduction to Virtualization. [online][referred 14.1.2009]. Available in www-format: < <http://www.cio.com/article/print/40701>>.

VMware 2006. Virtualization overview. [online][referred 21.1. 2009]. Available in www-format:
<<http://www.vmware.com/pdf/virtualization.pdf>>.

VMware 2007a. Understanding Full Virtualization, Paravirtualization, and Hardware Assist. [online][referred 20.1.2009]. Available in www-format:
<http://www.vmware.com/files/pdf/VMware_paravirtualization.pdf>.

VMware 2007b. VMware Virtual Networking Concepts. [online][referred 6.2.2009]. Available in www-format:
<http://www.vmware.com/files/pdf/virtual_networking_concepts.pdf>.

VMware 2007c. VMware Virtual SMP. [online][referred 22.4.2009]. Available in www-format: <http://www.vmware.com/pdf/vsmp_datasheet.pdf>.

VMware 2007d. The Architecture of VMware ESXi. [online][referred 20.5.2009]. Available in www-format: < http://www.vmware.com/files/pdf/ESXi_architecture.pdf>.

VMware 2008a. VMware Infrastructure 3 Pricing, Packaging and Licensing Overview. [online][referred 4.5.2009]. Available in www-format: <http://www.vmware.com/files/pdf/vi_pricing3.pdf>.

VMware 2008b. Introduction to VMware Infrastructure Update 2 Release for ESX Server 3.5, ESX Server 3i version 3.5, VirtualCenter 2.5. [online][referred 14.5.2009]. Available in www-format: <http://www.vmware.com/pdf/vi3_35/esx_3/r35u2/vi3_35_25_u2_intro_vi.pdf>.

VMware 2009a. Configuration Maximums, VMware Infrastructure 3: Update 2 and later for ESX Server 3.5, ESX Server 3i version 3.5, VirtualCenter 2.5. [online][referred 20.4.2009]. Available in www-format: <http://www.vmware.com/pdf/vi3_35/esx_3/r35u2/vi3_35_25_u2_config_max.pdf>.

VMware 2009b. ESX Server 3i Configuration Guide. [online][referred 4.5.2009] Available in www-format: <http://www.vmware.com/pdf/vi3_35/esx_3i_e/r35/vi3_35_25_3i_server_config.pdf>.

VMware Homepage 2009a. Server Consolidation. [online][referred 25.3.2009]. Available in www-format: < <http://vmware.com/solutions/consolidation/consolidate.html>>.

VMware Homepage 2009b. VMware ESX and ESXi 3.5 Comparison. [online][referred 22.4.2009]. Available in www-format: <<http://kb.vmware.com/kb/1006543/>>.