Expertise
and insight
for the future

Antti Savola

Server Virtualization with VMware

Metropolia University of Applied Sciences

Master of Engineering

Information Technology

Master's Thesis

22 August 2021

PREFACE

This thesis was a long hard excruciating journey for me but finally it is done now!
I learned a lot about myself during this Master of Engineering program. This whole journey began in 2016 with much enthusiasm towards learning new things. At the same time got a new job and combining stressful demanding job with studying was a difficult task. Time went on and studies progressed and started working with thesis. Then somethings happened at work and was forced to change the subject of thesis. Eventually found the new subject and was thrilled about it.
The practical side of this thesis was fun to do and I learned a lot by doing it. Then began the excruciating part of this thesis, the writing. Oh boy, it was a though one! Somewhere along the way lost my interest towards this thesis and time went on.
I would like to thank Ville and Metropolia for the patience and pushing me towards completing my this and graduating. Finally, it is done now and it's time to end this journey and enjoy the Master's degree.

Nurmijärvi, 22.8.2021
Antti Savola

Author Title Number of Pages Date	Antti Savola Server virtualization with VMware 32 pages 22 August 2021
Degree	Master of Engineering
Degree Programme	Information Technology
Instructor(s)	Ville Jääskeläinen, Principal Lecturer

Today's IT megatrends such as cloud computing, green IT and datacenters rely heavily and are based on virtualization technology. With the help of virtualization technology server's utilization can be used to the maximum by running many virtual servers on one single physical server. This is done by the virtualization software also known as hypervisor. Hypervisor runs and controls the virtual machines resources from the host server in a separated virtual layer.

VMware Inc. began its journey in virtualization technology in 1998 in California and soon released its first product Workstation 1.0. With Workstation one could run multiple operating systems as virtual machines on a single PC. VMware continued to develop its virtualization products and in 2009 VMware released VMware vSphere virtualization platform which become one of the mostly used virtualization platforms today. Many technology advances, product developments and patents later VMware has become the market leader in virtualization technology today.

This thesis was done for a Finnish ISP company and one of its clients. The aim of this thesis was to migrate client's virtual servers from client's dedicated virtualization platform into ISP's shared platform.

This thesis describes the basics of virtualization and server virtualization using VMware's virtualization products.

In the practical side of this thesis virtual servers were migrated into a new virtualization platform. The practical side began by going through the current platform and planning the migration of the virtual servers.

The result of this study shows how the virtual servers were migrated from one platform to another using VMware's virtualization products. The object of this thesis was successfully accomplished. The virtual servers were migrated into the new platform and all services are running as planned. In addition, the old ESXi hosts were wiped and were given to reuse.

Keywords	Virtualization, VMware, Virtualized platform, virtual servers

Tekijä Otsikko Sivumäärä Aika	Antti Savola Server virtualization with VMware 32 sivua 22 elo 2021
Tutkinto	Master of Engineering
Koulutusohjelma	Information Technology
Ohjaaja	Ville Jääskeläinen, Principal Lecturer

Tämpän päivän IT megatrendit, kuten pilvipalvelut, ympäristö- ja energiaystävällinen IT ja konesalit perustuvat virtualisointi tekniikkaan. Virtualisoinnin avulla palvelimien käyttöaste saadaan täysin hyödynnettyä ajamalla monia virtualisoituja palvelimia yhdessä fyysisessä palvelimessa. Tämän mahdollistaa virtualisointi ohjelma nimeltä hypervisor. Hypervisor suorittaa ja hallinnoi virtuaalisten koneiden resursseja isäntä palvelimesta erillisessä virtuaalisessa kerroksessa.

VMware Inc. on virtualisointi tekniikan yritys mikä perustettiin vuonna 1998 Kaliforniassa. Piakkoin VMare julkaisi ensimmäisen tuotteensa Workstation 1.0 minkä avulla pystyy suorittamaan useita käyttöjärjestelmiä yhdessä pc:ssä. VMware jatkoi tuotekehitystään ja vuonna 2009 julkaisi VMware vSphere virtualisointi alustansa mistä tuli sitten yksi käytetyimmistä virtualisointi alustoista tänä päivänä. Monia teknologisia innovaatioita, tuotekehityksiä ja patentteja myöhemmin VMwaresta on tullut tämän päivän virtualisoinnin markkina johtaja.

Tämän YAMK-insinööri työn tavoitteena oli selvittää virtualisoinnin perusteet sekä palvelin virtualisointi käyttäen VMwaren virtualisointi tuotteita ja ratkaisuja.

Työn kokeellisessa osuudessa migratoidaan virtuaalipalvelimia vmware alustasta toiselle.

Lopputuloksena virtuaaliset palvelimet migratoitiin uudelle VMware alustalle ja kaikki palvelut ovat suunnitellusti käynnissä. Lisäksi myös vanhat ESXi host palvelimet tyhjennettiin tiedoista ja laitettiin uudelleen käyttöön.

Keywords	Virtualization, VMware, Virtualized platform, virtual servers

Contents

List of Abbreviations

Blade chassis	Rack for blade server
Blade server	Type of server
Capex	Capital expenditures
COTS	Commercial off-the-shelf harfdware
DRS	Distributed resource scheduler
Enhanced vMotion	Enhanced vMotion Compatibility
ESXi	Hypervisor from VMware
ETSI	European Telecommunications Standards Institute
iLO	Integrated lights out port on HPE's servers, remote management port
Opex	Operating expense
Port group	grouping of switch ports
SAN	Storage area network
Storage array	Data storage system
Storage vMotion	Migrating virtual machines HDDs into a another datastore
vCenter	Centralized management of ESXi hosts
Virtual Connect	Blade chassis' network configuration
VM	Virtual machine
VMkernel port	VMkernel port is a networking interface used by the ESXi host to provide layer 2 or layer 3 connectivity to the host
vMotion	Migrating virtual machine to another ESXi host server
Vmware vSphere	Management software to administer the ESXi platform

1 Introduction

In the early day's computers were bought for one purpose or one application only and the utilization of the computer resources wasn't fully used. As time passed on and IT-technology evolved the utilization of computing resources improved. Today we can utilize the computer's resources to the maximum with the help of virtualization technology. Instead of buying one server for one application as was done before we can have one server to run multiple virtualized servers and therefore utilize the server's resources to the maximum. This also reduces Capital expenditures (Capex) and Operational expenses (Opex) costs and drives the computing environment toward todays green IT trend. Today's IT megatrends such as datacenters and cloud computing relay and are based on virtualization technologies.

This thesis was done for a Finnish Internet Service Provider (ISP) company and also for one of its clients. Due to data security and NDA reasons no more information will be given about the ISP or the client.
It explains the basics of virtualization and how servers are virtualized using VMware's products.

This thesis was done to help both the client and the ISP company. It was beneficial for the client to do this thesis because they got more powerful and newer ending hardware for their virtual servers. This also reduced the clients Opex costs by not having to pay for a dedicated server hardware to run their dedicated environment. The ISP benefited from this thesis by decreasing the amount of dedicated virtual environments that they have to maintain and also by decreasing the number of servers to maintain. So, it was a win-win situation for both of them. Due to data security reasons all of the screenshots in the practical side of this thesis are taken from lab environment.

In the practical side this thesis goes through the phases of migrating virtual servers from one virtualization platform to another. These phases include migration planning, migration process of the VMs and and cleanup process. In migration planning phase the required capacity of the VMs and some ESXi host configuration is done. In the migration process the VMs are of course migrated to the new platform and the migration process is explained. In the last phase the old platform and it's servers are removed and wiped for data security reasons.

The thesis has been divided into four sections. The first section introduces the problem. The second section introduces virtualization technology and the history of it briefly. Also, different types of virtualization technologies are explained. The third section goes through the history of VMware Inc. and different VMWare virtualization technologies and products which were used in this thesis are thoroughly explained.

The last section, chapters 4-7, of this thesis goes through the practical side of this thesis. Section 4 begins by going through the client's environment and planning the migration steps of the VMs and the ESXi host servers. After which the migration process is explained and executed. In the last chapter the results of this thesis are explained and whether the objective of this thesis was successful or not.

2 Virtualization

This section explains what is virtualization and briefly explains how virtualization technology was invented. Different virtualization technologies are also explained related to this thesis and also the pros and cons of virtualization.

2.1 Virtualization

The word virtualization means changing something that exists in a real form into a virtual version. In the IT field virtualization is a process of running a virtual instance of a computer system in a layer separated from the actual hardware as seen in Figure 1.
[1.]

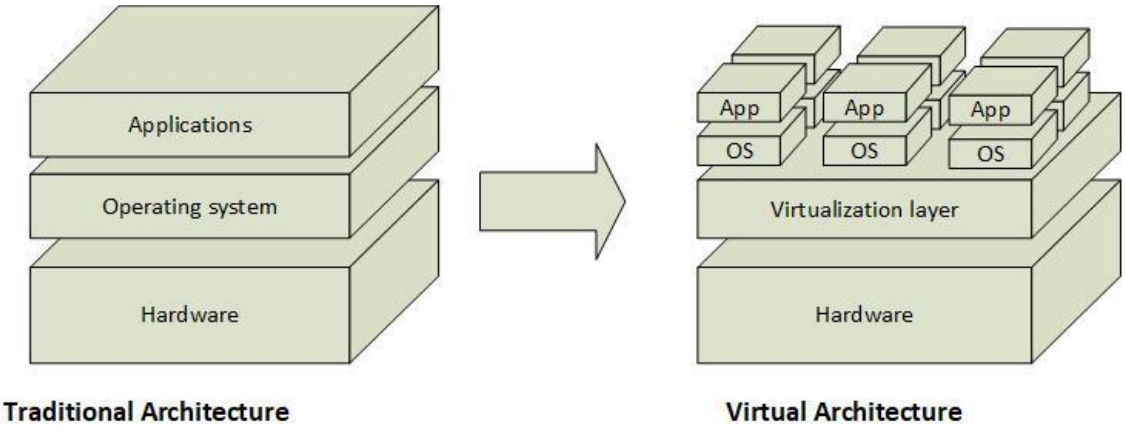

Figure 1. Traditional architecture vs. virtualization architecture [1]

History of virtualization

Virtualization concept was invented in the late 1960's and early 1970's by IBM. IBM began using virtualized memory for first time in the late 1960's. The first fully virtualized computer saw the light of day in 1972 when IBM released its virtualization OS VM/370.
[2.]

During the 1980's and 1990's there was no need for virtualization because the prices of computers came down to consumer level and the number of computers and servers began to rise. This also made big single mainframes obsolete.
[2.]

In the late 1990's virtualization technology started to raise its head again when VMware released its first product VMware Workstation. VMware Workstation was a type-2 hypervisor that ran x86 versions of Windows and Linux operating systems. In the early 2000s VMware released its x86 server virtualization product ESX Server which was a type 1 hypervisor (see section 2.2). Technology kept evolving as years passed and new solutions emerged such as Xen server and Hyper-V.

Nowadays in the 2021 virtualization is the foundation of the cloud driven IT world such as Azure, AWS and Google.

Pros and cons of virtualization

There are many benefits in virtualization and it can improve companies' IT infrastructure compared to the traditional physical IT infrastructure. With virtualization companies can transform their IT infrastructure into modern efficient, agile, scalable and fault tolerant environment and at the same time reduce Capex and Opex costs.

The down side of virtualization is that it is quite expensive at first due to expensive hardware (physical host serves) and licensing investments to make the virtualization even possible.

In some cases, it is also possible that not all hardware or software can be virtualized nor should they. For example, virtualizing Oracle database servers can be very expensive due to the Oracle's licensing costs which are based on server's core account.

2.2 Hypervisor

To be able to run a virtual machine a virtualization software is needed. Virtualization software is also known as a hypervisor. A hypervisor is a software in the virtualization layer that gives virtual machines their computing resources (cpu, memory and storage) from the host server.

There are two types of hypervisors, type 1 native which is also known as bare metal and type 2 hosted hypervisor as descripted in Figure 2.
[3.]

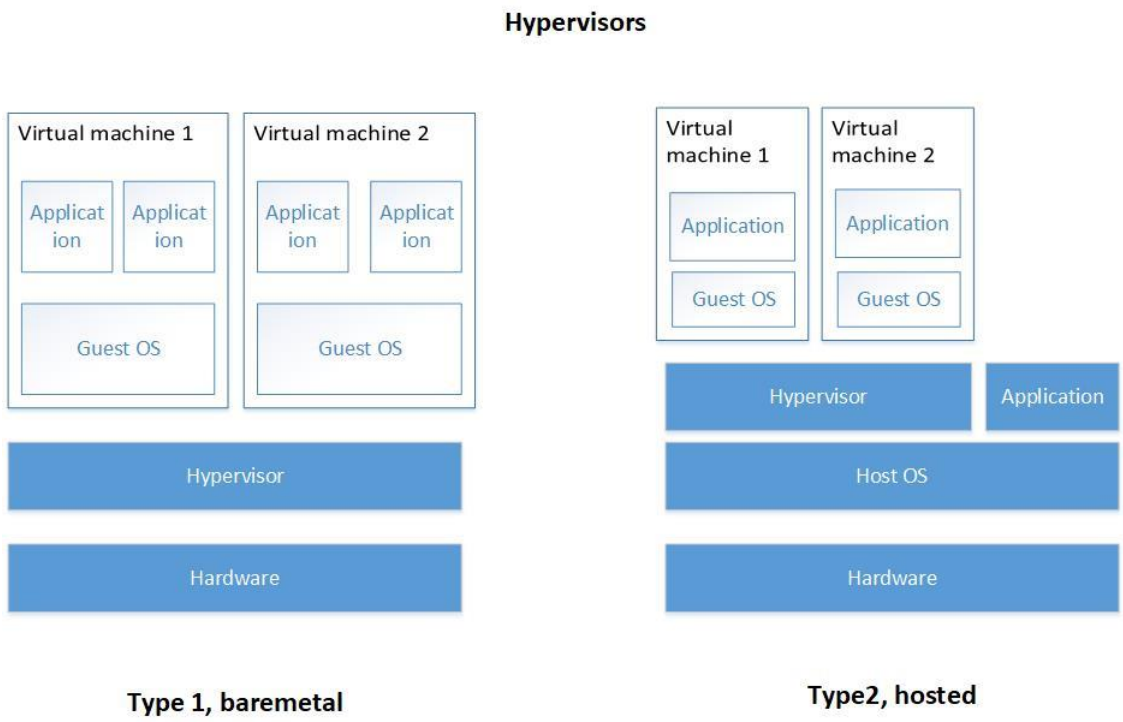

Figure 2. Different hypervisors and their architectures [3]

Type 1 hypervisors are VMware vSphere with ESX/ESXi, Microsoft Hyper-V, Citrix Hypervisor (Xen Server).

Type 2 hypervisors are for example VMware Workstation Pro/VMware Fusion, Oracle VirtualBox.

The differences between bare-metal and hosted virtualization are that in bare metal there is no software or any operating system between the hardware and the hypervisor. Where as in hosted there is a software layer, operating system, running between the hardware and the hypervisor layer.

These different types of hypervisors exist for different use cases. For personal and small development or testing purposes type2 hypervisor is the way to go. Developers can easily set up a development VMs with the help of Oracle's VirtualBox or with VMware's Workstation on their own laptop.

For the enterprise or datacentre use case type 1 is the best choice. Although it comes with licensing costs which need to be taken into account when choosing the hypervisor.

2.3 Different types of virtualizations

This section describes the main types of virtualizations and their use cases.

Desktop virtualization

The two mostly used desktop virtualization types are local desktop virtualization and remote desktop virtualization.

In local desktop virtualization the guest operating system runs on a client device using hardware virtualization. All of the processing and workloads use the client devices resources, such as cpu, ram, storage and connectivity as is described in Figure 3. Local desktop virtualization can be done for example using Oracle's VM VirtualBox, VMware's Workstation, Microsoft's Hyper-V. [4.]

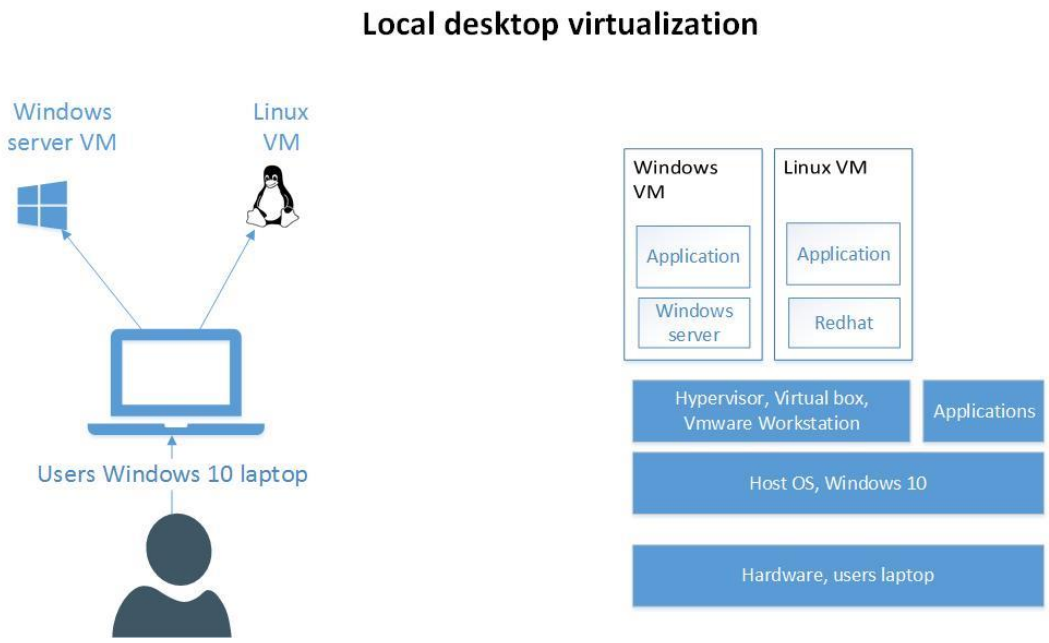

Figure 3. Local desktop virtualization architecture

In remote desktop virtualization the operating systems and applications are running in a server environment while all of the user's interactions with the OS and the applications are coming from the client device. The client device can be a laptop, thin client device or a tablet/smartphone. One popular type of remote desktop virtualization is virtual desktop infrastructure also known as Virtual Desktop Infrastructure (VDI). VDIs consist of client-server environments which use host-based VMs to deliver virtual desktops to the end user as described in Figure 4. Popular VDI solutions are for example Citrix VDI and Azure VDI.
[4.]

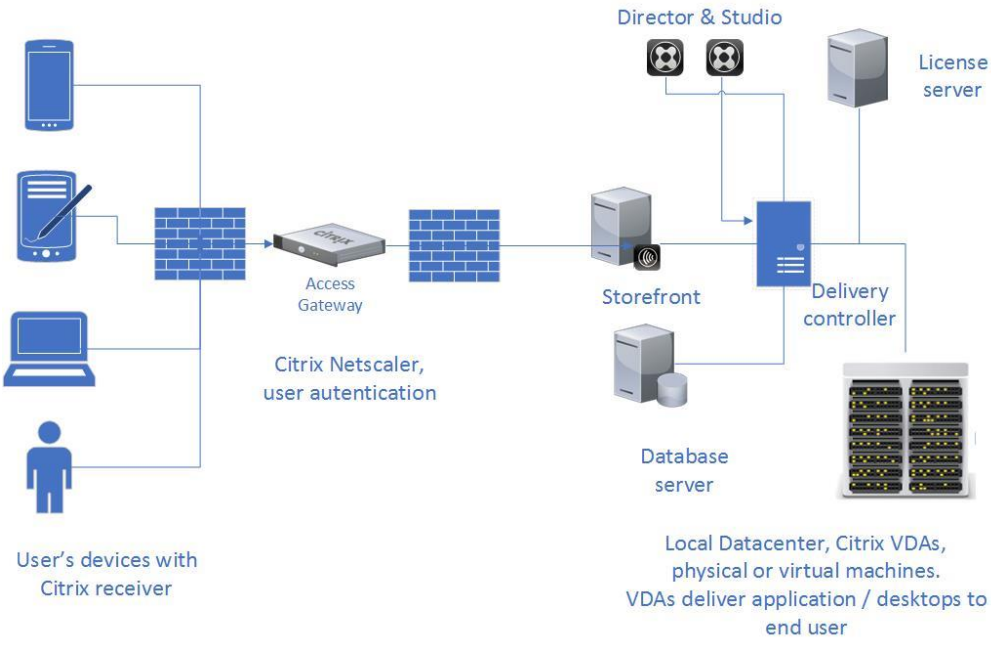

Figure 4. Citrix VDI architecture [5]

Server virtualization

There are three different types of server virtualization technologies available. The differences between these virtualization types are mainly based on the level of isolation they provide which also relates to how much hardware resources these types emulate. The types are

- Full virtualization or virtual machine model
- Paravirtual machine model
- Virtualization at the OS level.

In the full virtualization model, the virtual machine is complete isolated from the host server's hardware by the virtualization layer/software called hypervisor. The virtual machine's guest OS isn't aware that is being virtualized and therefore requires no modification in order to be run.

Hypervisor communicates with the physical server's resources and relays resources from the physical server to the virtual server. Hypervisor also monitors the host servers resource utilization and keeps each virtual machine independent and unaware of other virtual machines.

The down side of the full virtualization is that the hypervisor also requires some resources from the host server for processing. This can impact to the performance of the VMs and their applications.

Paravirtual

Vmware has been using paravirtualization in its products for example vSphere for years. VMware has developed VMware tools and optimized virtual device drivers which allow the VM's guest OS to be able to communicate with the hypervisor. In other words, one could install and run a virtual machine without VMware tools but the VM wouldn't be able to communicate with the hypervisor or the rest of the physical network. These device drivers are for example network interface drivers which allow VM to be able to communicate with the physical network outside of the virtualization layer.
[6.]

Operating system virtualization

Operating system virtualization also known as container technology is a technology where the kernel of an operating system allows multiple isolated user-space instances. These isolated user-spaces are called containers. These containers allow one to package application together with libraries and other dependencies, providing isolated environments for running software services

One popular container format is Docker. Docker is a containerization platform that packages the application and its dependencies together inside a container so that the application works seamlessly in any environment. Containers are excellent in DevOps- model of software development where developers can develop software regardless the environment.

Containers are run on top of the OS kernel directly instead of virtualizing the hardware stack as it is done with the virtual machine approach in type 2 hosted hypervisor model. Due to this the containers are swift and light weighted compared to VMs.
[7,8.]

Network functions virtualization (NFV)

NFV is an architectural frame work developed by ETSI NFV Industry Specification Group. The aim of this frame work is to virtualize networking functions. NFV is a technology where the traditional single service proprietary and purpose-built networking hardware such as routers, firewalls and load balancer are replaced with virtualized VMs/ virtualized appliances to do the same task. The purpose of NFV is to move away from the single purpose hardware into a virtualized network functions that are delivered through software virtualization as virtual network functions (VNFs) with commercial off-the-shelf (COTS) hardware such as HPE ProLiant or Dell's PowerEdge servers. The aim is to make the network more scalable, cost efficient and agile.
[9,10.]

3 VMware

The virtualization platforms used in this thesis were running with VMware's virtualization products so this chapter focuses on VMware and its virtualization products such as vSphere / ESXi.

3.1 VMware Inc.

VMware Inc. is a virtualization technology-based IT-company founded in 1998 in California. VMware's first product Workstation 1.0 brought freedom to IT-professionals by allowing them to run multiple operating systems as virtual machines on a single PC. In 2002 VMware changed the virtualization technology by releasing ESX server hypervisor. ESX server made possible to consolidate many physical servers into virtual servers and thus lowering the number of physical servers in companies.

In 2009 VMware released VMware vSphere virtualization platform which become one of the mostly used virtualization platforms today. Many technology advances, product developments and patents later VMware has become the market leader in virtualization technology today.
[11.]

3.2 VMware's server virtualization products

ESXi is a virtualization hypervisor developed by VMware. ESXi is a type-1 hypervisor also known as a bare-metal hypervisor.

VMware vCenter
Central management for the virtualization environment.
VCenter provides the following management functions:

- Resource management for ESXi hosts and VMs
- Template management
- VM deployment
- VM management
- Scheduled tasks
- Statistics and logging
- Alarms and event management
- ESXi host management

DRS

Distributed resource scheduler is a VMware proprietary algorithm that balances the usage of ESXi hosts' resources by automatically migrating virtual machines between ESXi hosts.

vMotion

vMotion also known as live migration, is used to manually or automatically balance ESXi hosts computing resource utilization between different hosts and clusters. In vMotion virtual servers are migrated from one ESXi host to another without any down time. This means that during migration VMs network connections are not dropped and the applications running inside the virtual machine continues to run without any interruptions as is shown in Figure 5.
[3.] [12.]

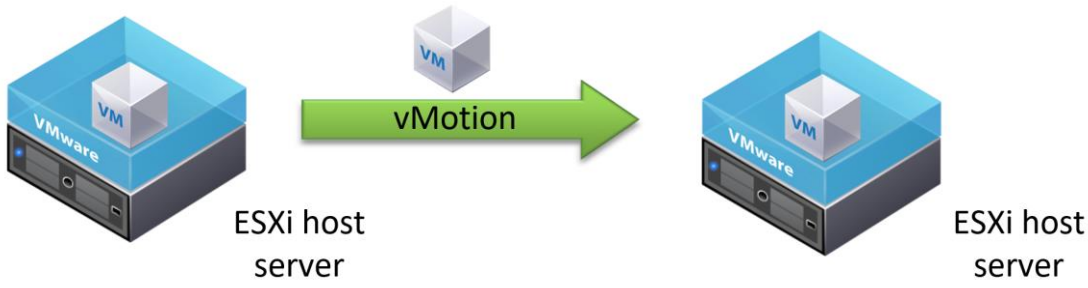

Figure 5. vMotion [3,12]

Storage vMotion

Storage vMotion is used to manually or automatically balance datastore utilization in a vCenter cluster or in ESXi hosts. During the storage vMotion virtual machine's vmdk files are moved from one datastore to another as show in Figure 6.

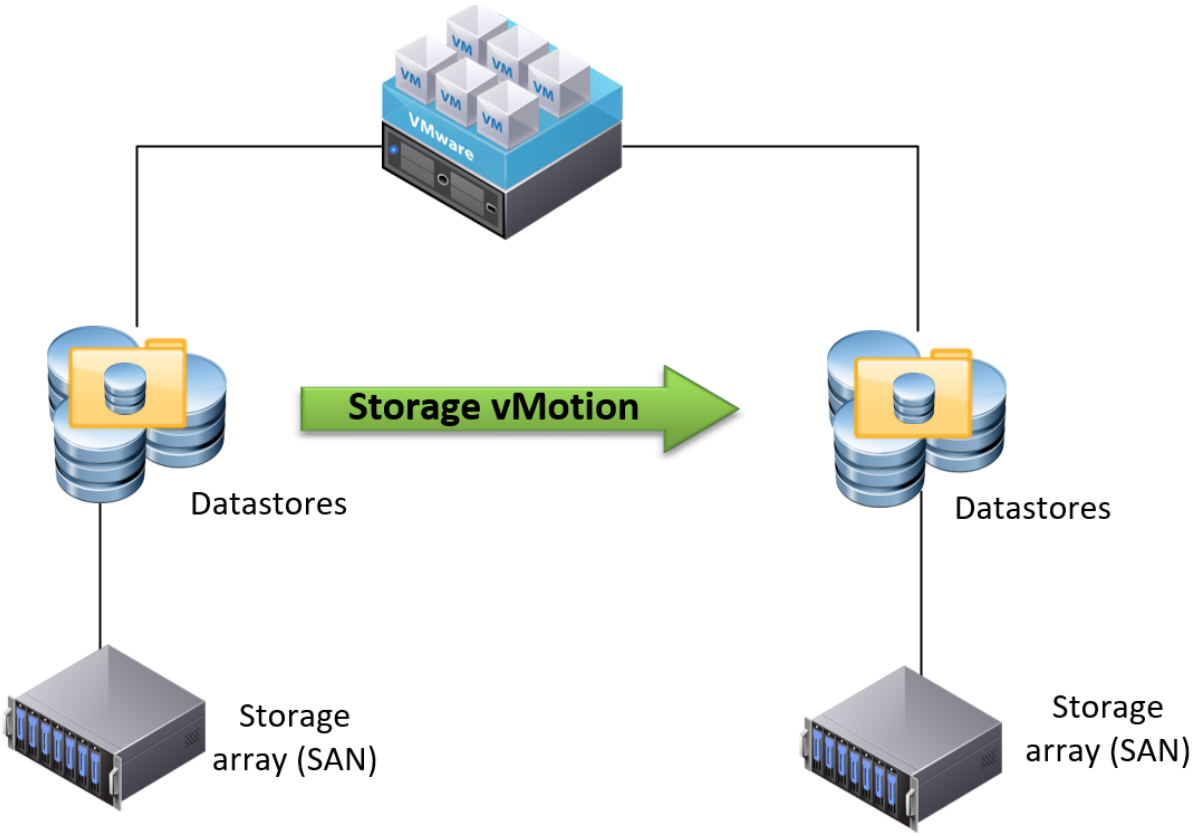

Figure 6. Storage vMotion [3]

HA, high availability

VMware technology that provides high availability at the virtualization layer. It protects against the following failures:

- ESXi host failure – If the host server crashes virtual machines that were running on it will be restarted on other hosts automatically as is shown in Figure 7.

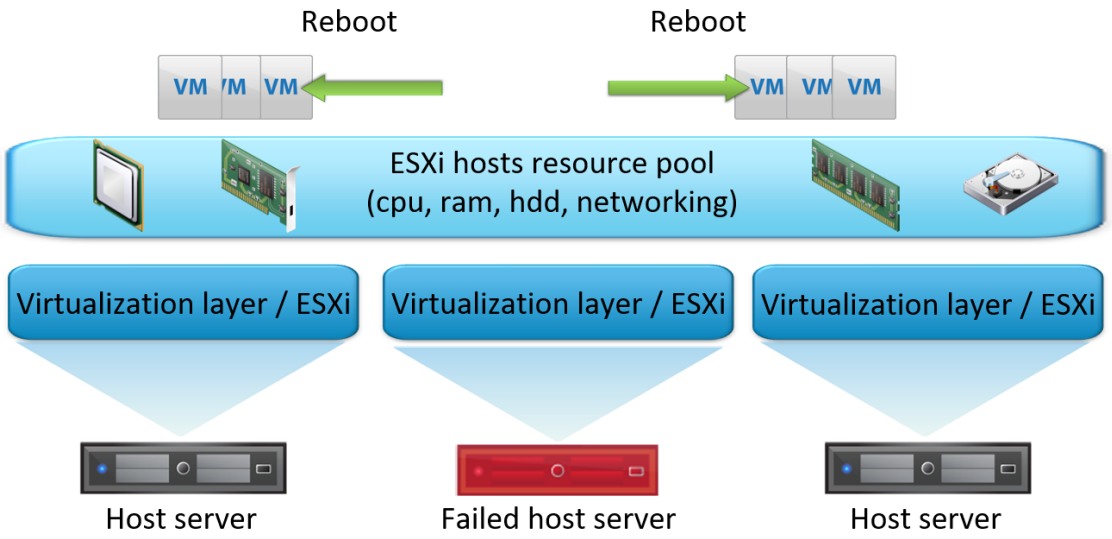

Figure 7. HA ESXi host failure [3]

- Guest OS failure – if the VM's guest OS crashes and the VM stops sending heart-beats then the guest OS is reset

- Application failure – ESXi host server can monitor applications running inside of a VM. If the application fails the VM is restarted.

Virtualized Networking

Networking in the virtualization layer differs quite a bit from the physical networking although there are still some similarities between the virtual and physical networking. Both networking architectures need physical cabling in order to be able to communicate between virtual and physical servers. The architecture of virtualized networking in VMware vSphere is explained in Figure 8.

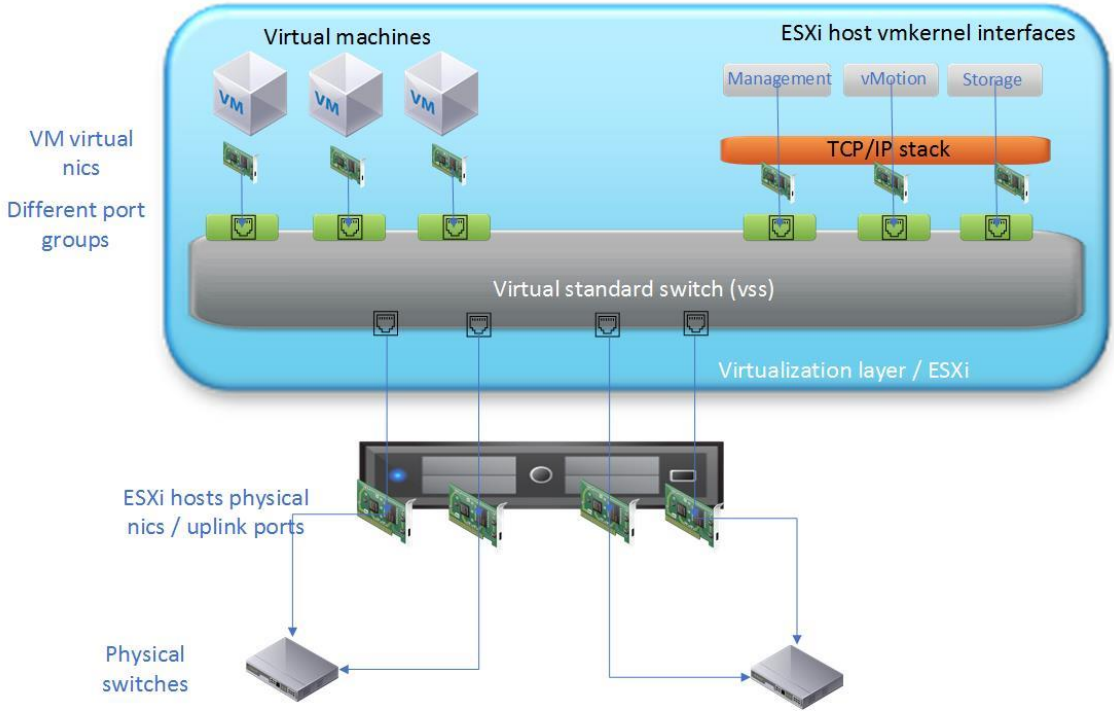

Figure 8. ESXi networking architecture [13.]

- Port groups
  A grouping of ports in a vswitch

- VMkernel,
  VMkernel port is a networking interface used by the ESXi host to provide layer 2 or layer 3 connectivity to the host

- Uplinks,
  are the physical network adapters of the ESXi host.

- vSwitch
  Virtual switch works the same way as a real switch does by forwarding frames to other switches ports at layer 2. vSwitch is connected to ESXi hosts physical nics as uplinks to be able to communicate with the rest of the network.

3.3 VMware platforms

Clients current VMware platform consists of 3 ESXi host servers, a virtual vCenter server, a storage array, networking (SAN, switches, routers, firewalls, etc…) and virtual servers.

The new VMware platform to which the client's VMs will be migrated consists of 15 ESXi host servers, a virtual vCenter server, a storage array and networking (SAN, switches, routers, firewalls, etc…)

4 Migration planning

This section describes virtual server capacity calculations and ip address planning for the case company. The ISP is providing hosting services for the client so all of the hardware and networks are owned by the ISP.

4.1 Capacity calculations

A proper migration plan was needed to make sure that maintenance break in operational services would be as short as possible during the migration of the VMs. Planning began by going through both virtualization environments and their configurations and settings. The capacity (vCPUs, RAM, storage) required to run the VMs was calculated and compared against the free capacity in the new environment to see if the new platform would be able to run the VMs. The VMs would need xxx vCPUs, xxx Gb of RAM and xxx Tb of storage. After calculations and comparisons, it was concluded that the new platform could host the VMs.

4.2 New IP addresses

In order to be able to move the current ESXi hosts to the new platform new management and vMotion vmkernel ip addresses were needed. New IP addresses were reserved from the same subnets that were used in the new platform for ESXi management and vMotion.

4.3 Removing current host servers

Before ESXi hosts can be moved to the new vCenter they have to be removed from the old vCenter. Before disconnecting the hosts Integrated Lights Out (iLO) connections were opened on all hosts. This was required to be able to change the management IP address of the host on later.

After the iLO connections were opened the removal of ESXi hosts from vCenter was done by disconnecting the hosts from the vCenter and finally removing them from the vCenter's database.

4.4 Adding current ESXi hosts and creating new cluster

The management IPs of the ESXi hosts were changed in the management interface of the host. IP address, subnet mask, gateway and DNS servers were inserted into the management interface as shown in Figure 9.

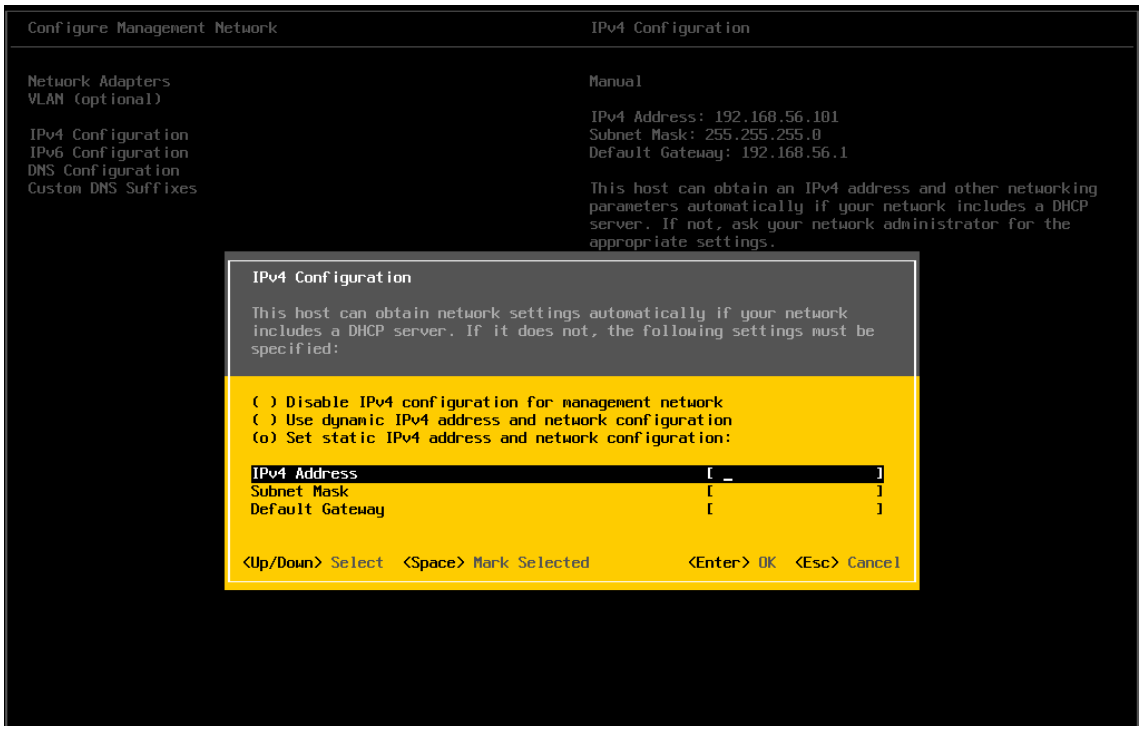

Figure 9. Changing the management IP address of the ESXi host

After the IP change a new temporary cluster was created into the new vCenter which would “house” the current hosts. Hosts were added to the new vCenter and to the new cluster. After this the cluster was configured and the DRS and HA were turned on so that the cluster would be fault tolerant in case a host or a vm crashed. As can be seen in Figures 10 and 11.

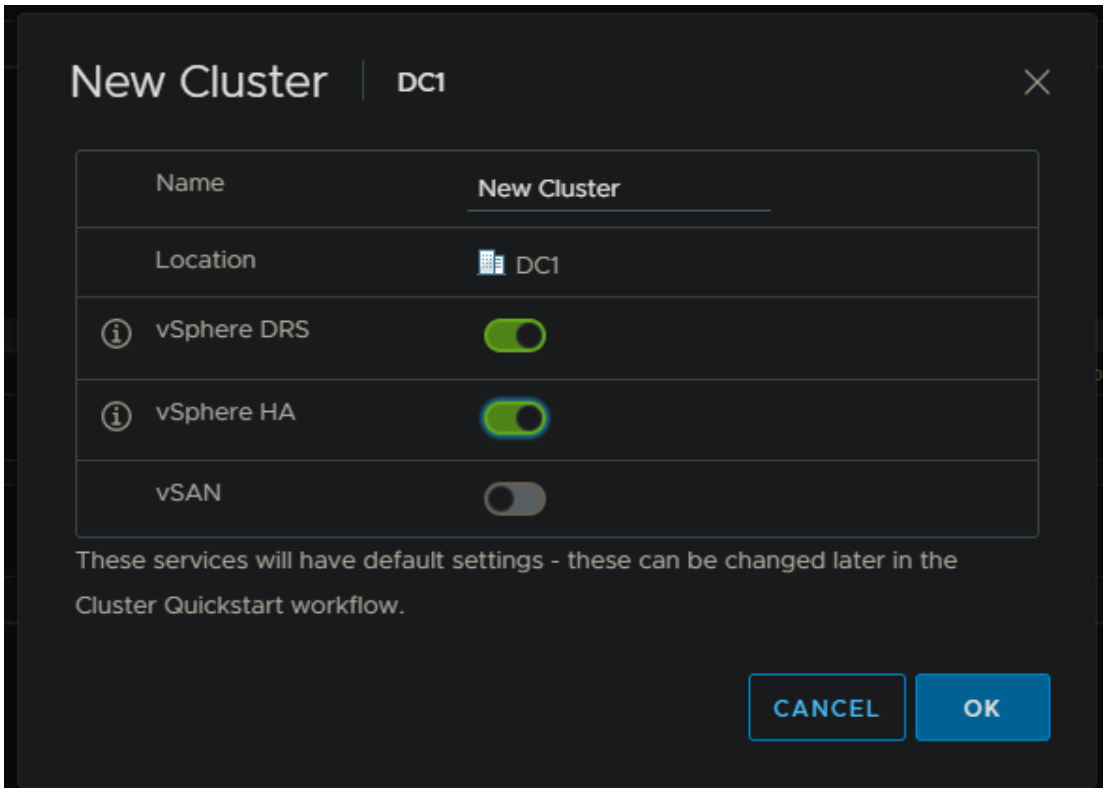

Figure 10. Creating new cluster for the current ESXi hosts with DRS and HA settings.

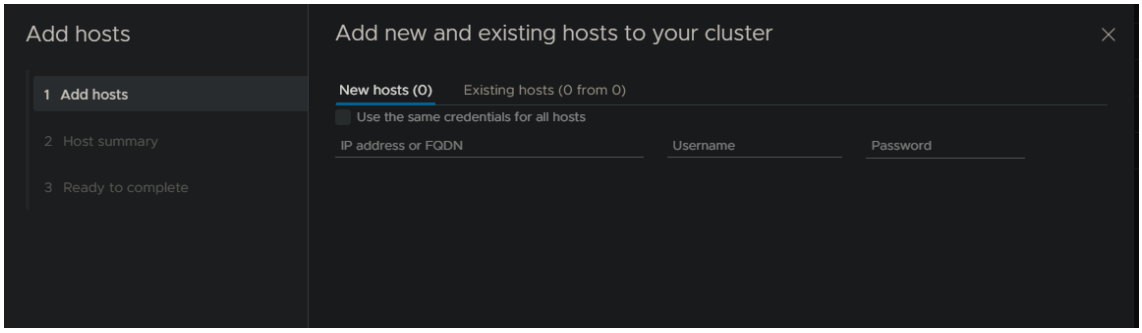

Figure 11. Adding current ESXi hosts into the temporary cluster

4.5 Adding client's production networks

Client's production networks' and their corresponding VLAN-IDs were distributed to the new platform's switches and blade chassis' uplink ports through virtual connect. In order to get the client's networks to the virtualization layer new port groups had to be created into the new hosts. The creation of a new port group is shown in Figure 12. Port groups were named identical as in the old platform so that the migration would go smoothly without any errors from different network names. VLAN-IDs were also added to the port groups so that the networks would be separated from other networks.

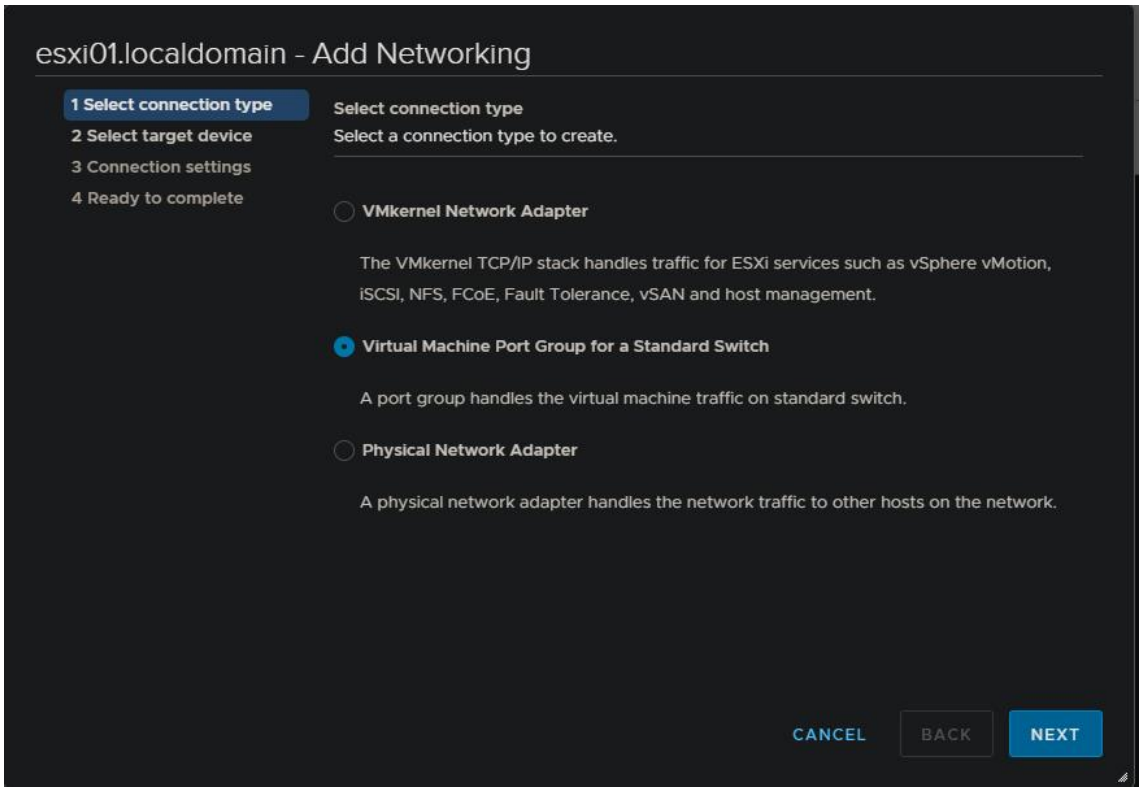

Figure 12. Creating client's networks into the new ESXi hosts

4.6 Creating migration datastore and adding it to both platforms host servers

A temporary migration volume was created into the storage array and it was “presented” through Storage Area Network (SAN) to both virtualization platforms. A new datastore was then added to the vmware. In order for the ESXi hosts to see the new datastore hosts’ HBAs had to be rescanned as can be seen in Figure 13.

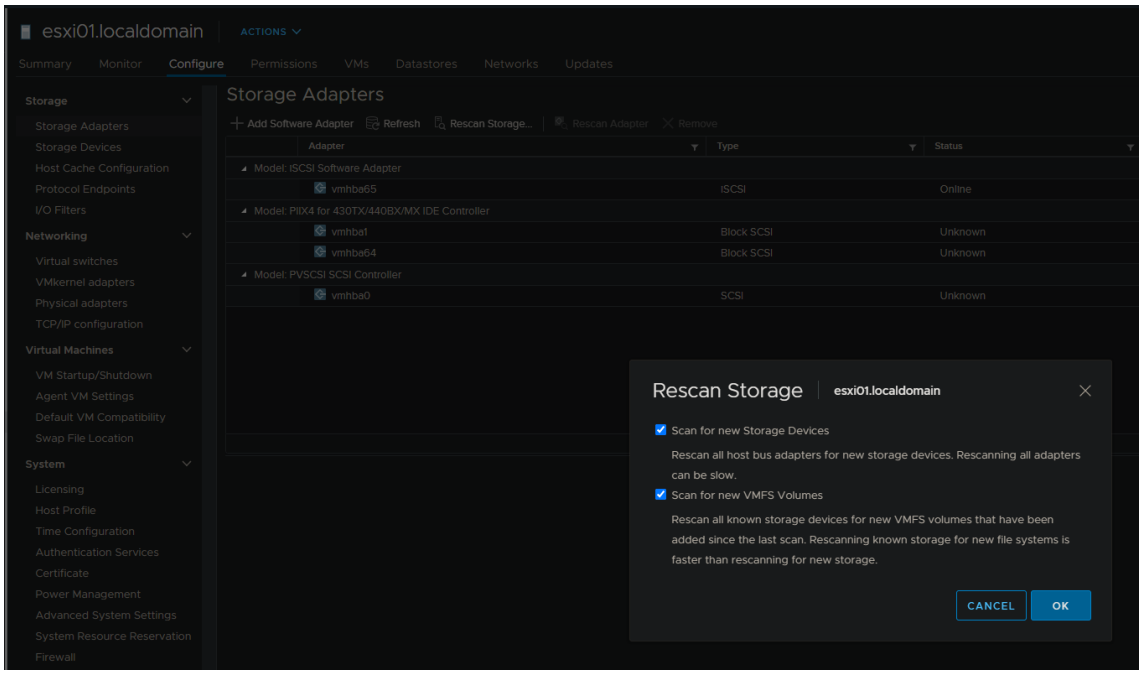

Figure 13. Scanning the ESXi’s HBAs

After the HBAs were scanned the migration volume was visible to the hosts and the new vmfs datastore was created as seen in Figures 14,15,16 and 17.

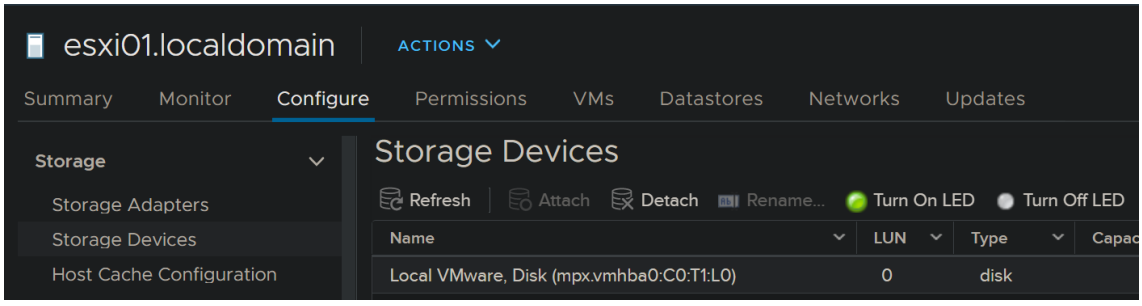

Figure 14. Migration volume appeared into the storage devices section of the hosts after the HBA scan

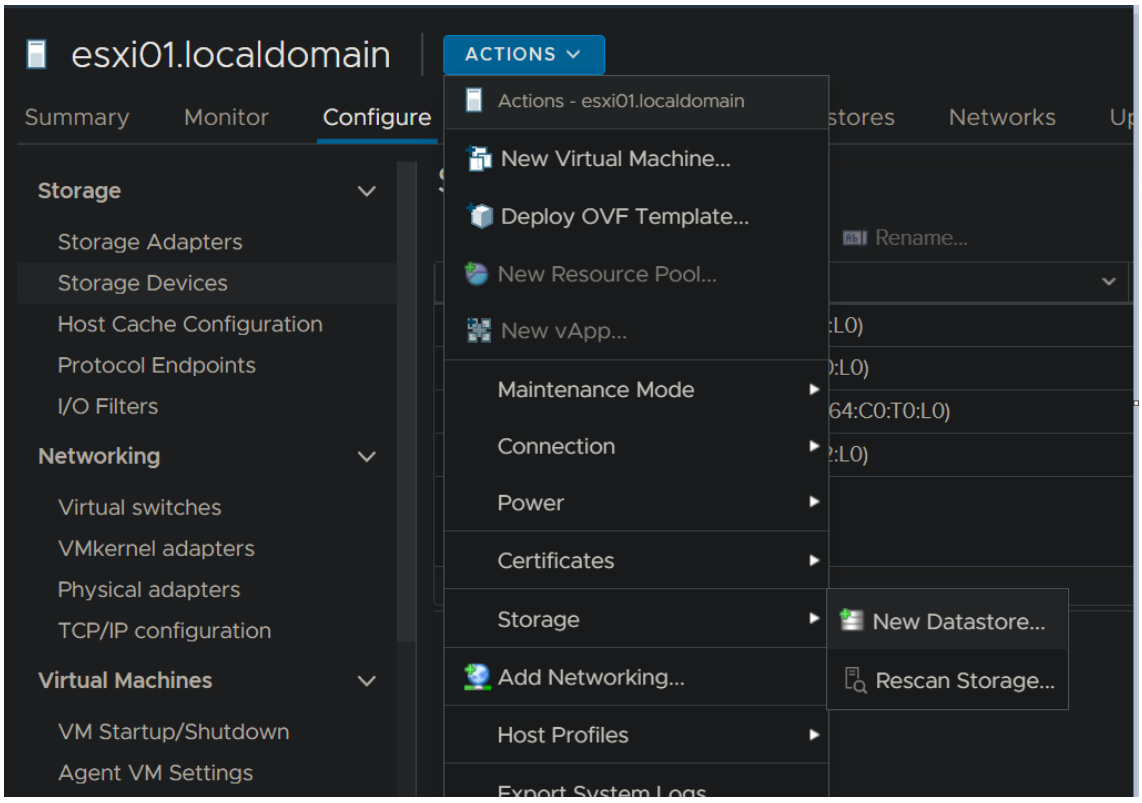

Figure 15. Creating a new VMFS datastore

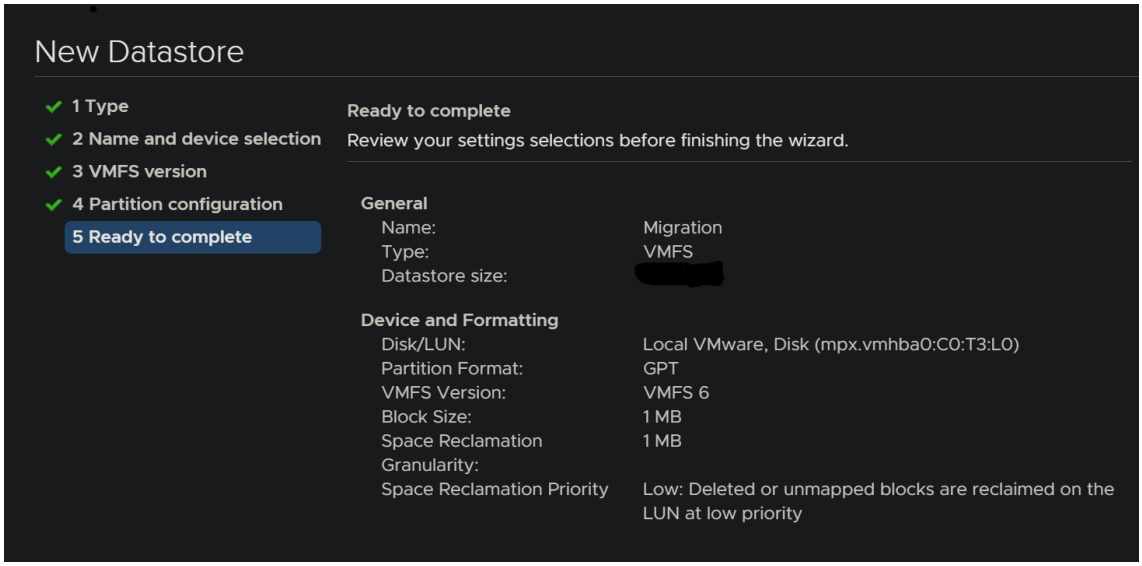

Figure 16. Choosing the correct datastore settings for the migration volume

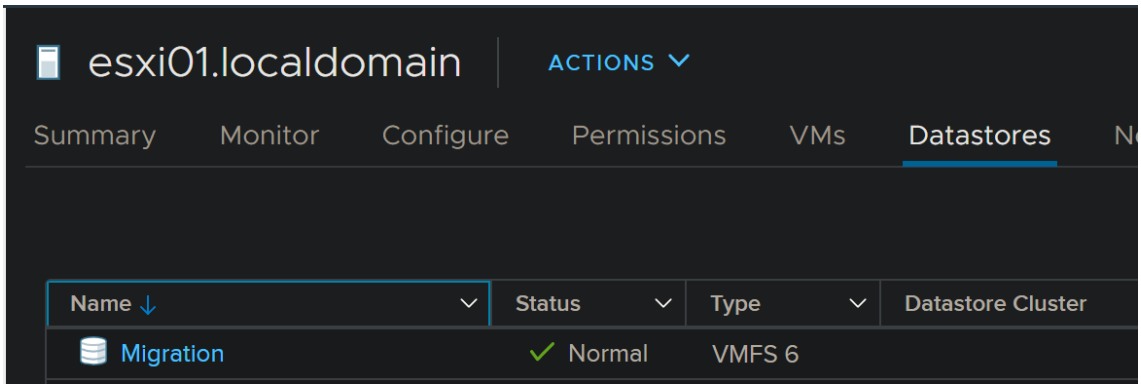

Figure 17. Migration datastore created

5 Migrating virtual servers

This chapter describes the virtual server migration process from one vCenter to another and how it was done.

5.1 Planning maintenance breaks with customer

After the underlaying infrastructure was configured for the migration, it was time to schedule the maintenance break with the customer.

It was agreed with the customer that the migrations would be done only on weekends starting on Friday evenings and ending at Sunday midnight.

5.2 Migrating VMs

The migration began by moving test server's disks to the migration datastore with Storage vMotion as described in Figures 18 and 19.

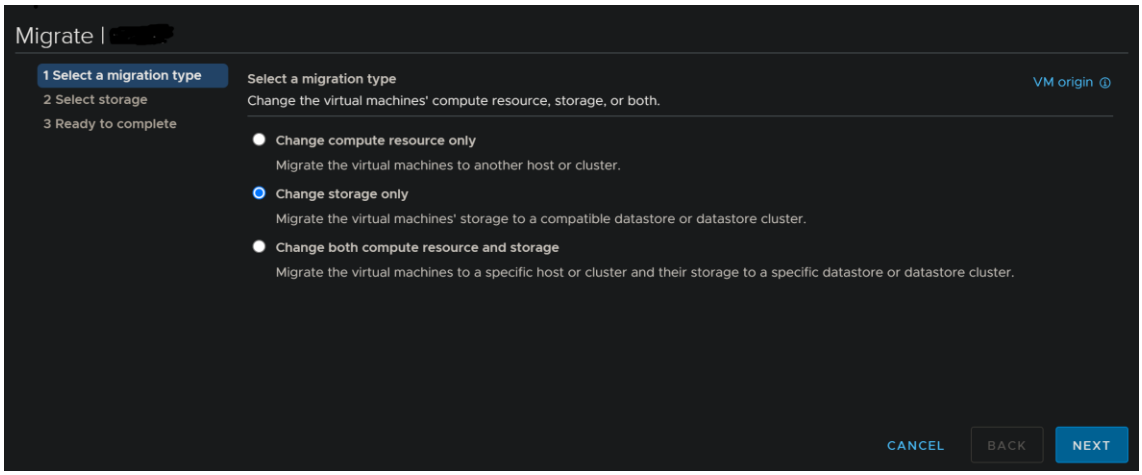

Figure 18. VM's storage vMotion

Figure 19 shows how the VMs were migrated from the old platform to the new platform via storage vMotion. This was possible by showing the migration datastore to both platform's ESXi hosts through the SAN network. By sharing the migration datastore the VMs could be migrated from the old platform into the new platform.

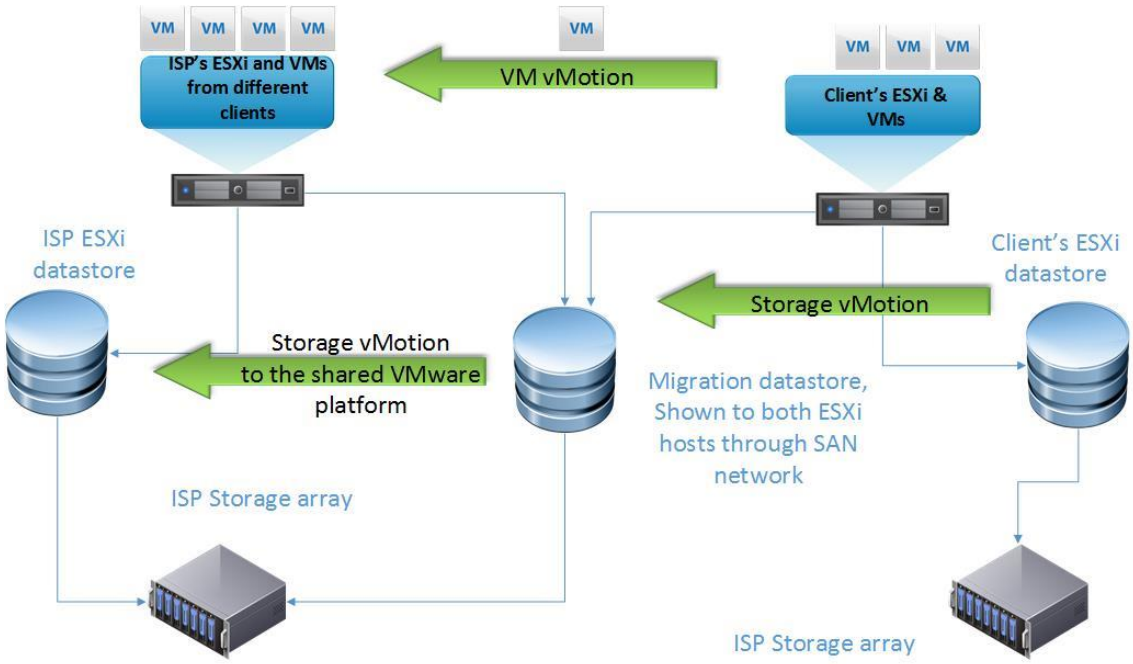

Figure 19. VM storage vMotion to the shared datastore

After that the test VMs were switched off and the VM's were migrated to the new ESXi hosts using vMotion and they were powered on as seen in Figures 20 and 21.

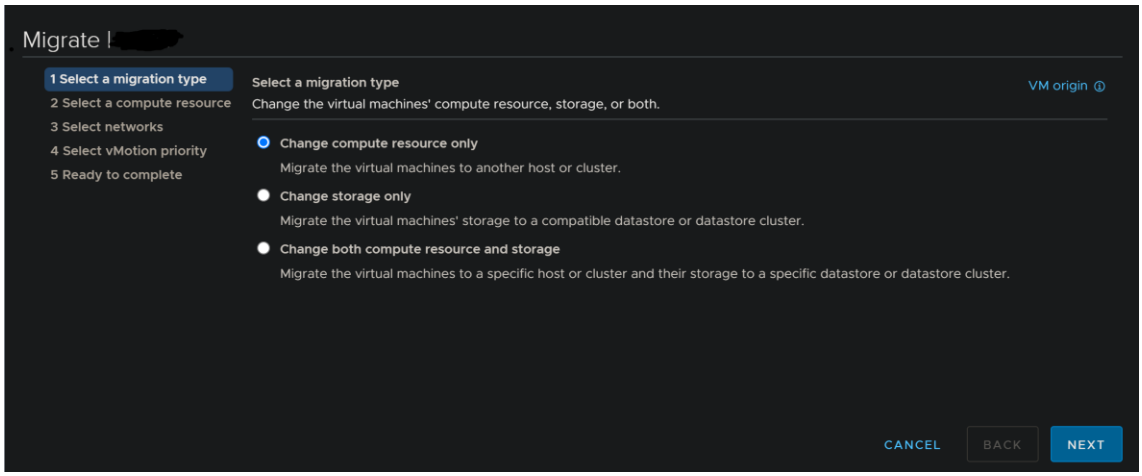

Figure 20. VM vMotion

The CPUs on the old and the new ESXi hosts were too different to be able to use enhanced vMotion. After the VMs were up and running the functionality and the network connectivity of the VMs were tested. Once the VMs were concluded to be functional and connected it was time to move the VM's disks to the new platforms datastores using storage vMotion, as described in Figure 21. The test servers were successfully migrated to the new platform.

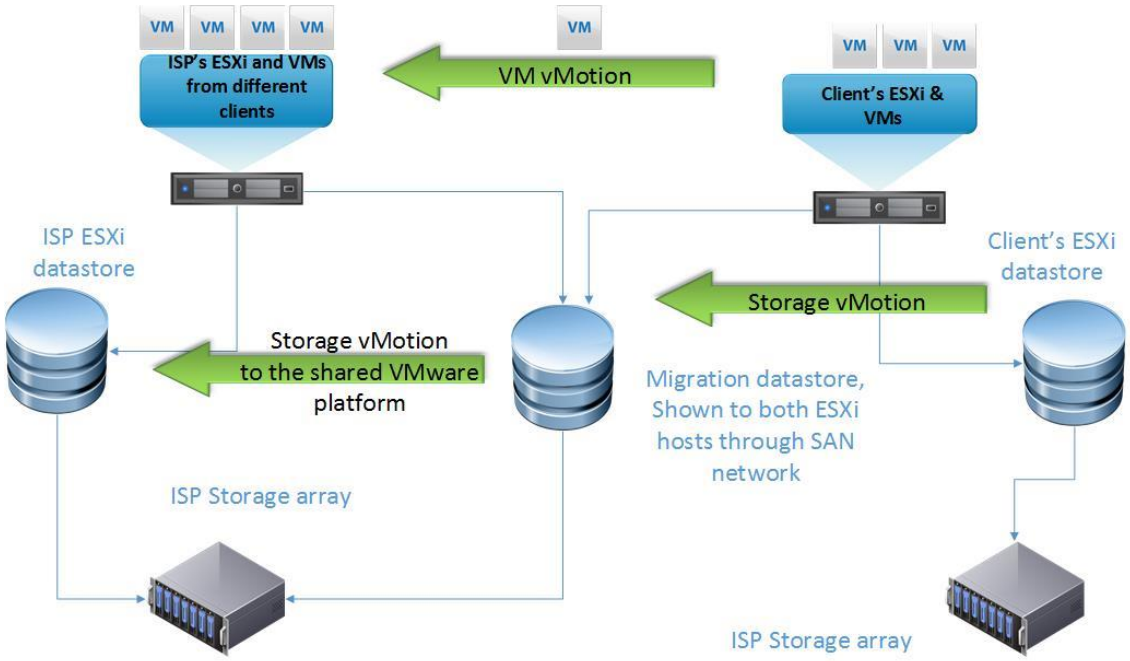

Figure 21. Storage vMotion from migration datastore to the new platform's datastore

6 Securing old host servers

This chapter explains the cleanup tasks done after the migration.

6.1 Removing client's networks from old host servers

After the VMs were moved to the new platform it was time to remove the client's networks from the old host servers. Starting from networking switches and blade chassis virtual connect the client's VLANs were removed from the switch's ports and from the blade chassis uplink ports and also from the ESXi's port groups. Removing of port groups in ESXi can be seen in Figure 22.

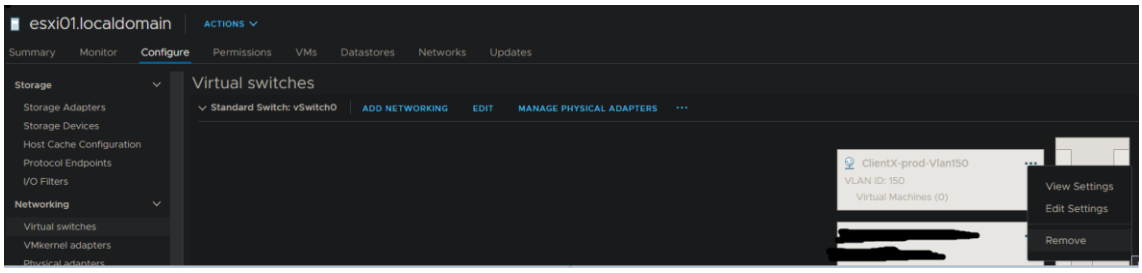

Figure 22. Removing port groups in ESXi

6.2 Deleting old datastores

The old datastores were deleted from the VMware and from the storage array. The Zonings done in the SAN network between the ESXi servers and the storage array were also deleted.

6.3 Removing the old ESXi hosts from the monitoring system

In today's IT world with thousands and thousands of servers it is impossible to monitor them by hand. That's why many corporations and ISPs use some sort of monitoring system to monitor the usability and performance of the servers.

Such's is the case also in this thesis. The old host's datastore, CPU, memory usage and network connectivity were monitored. All of the performance statistics of the old hosts were deleted and any old alarms were cleared from the monitoring system.

6.4 Securing data on the old servers

It was decided to reuse the old ESXi hosts in another virtualization environment because the hosts were still pretty usable and powerful although being HP's generation 8 blade servers. Servers were planned to be used in a testing / R&D purpose.

Before old hosts could be reused local HDDs had to be wiped for data security purposes. The local HDDs were wiped with DBAN application. DBAN overwrites HDD's data blocks full of zeros with different algorithms many times so that the data can't be recovered.

Erasing the server's disks began by booting up the server from the DBAN installation media as shown in Figure 23.

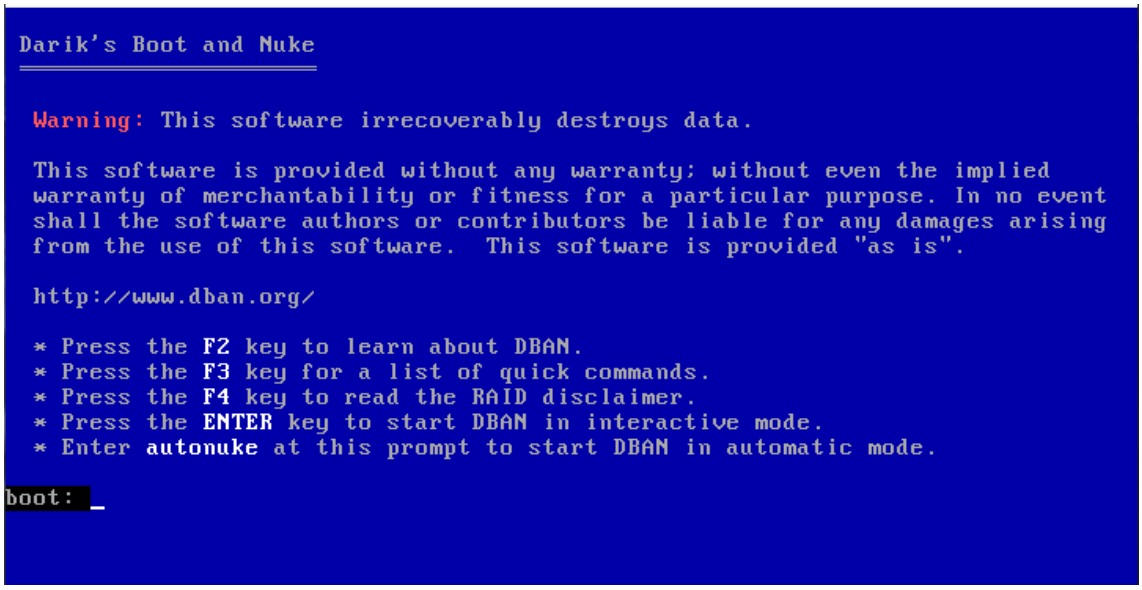

Figure 23. DBAN startup

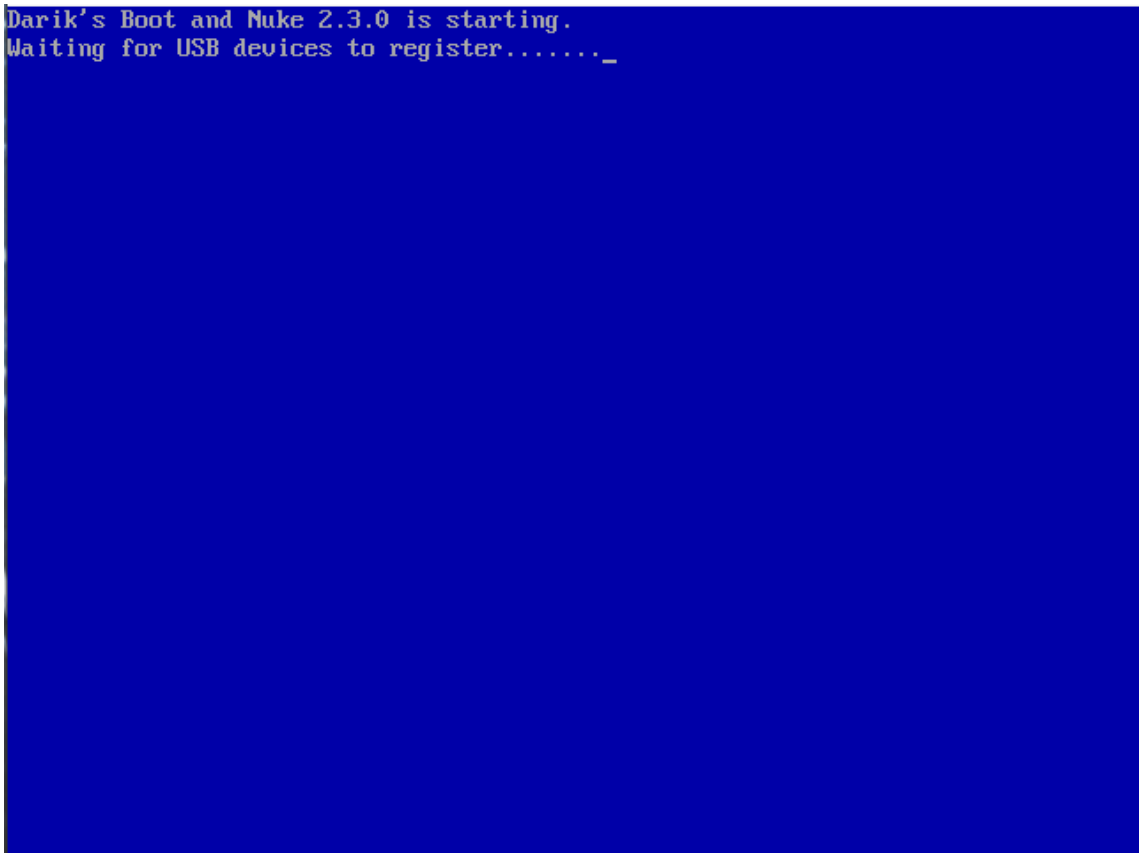

Figure 24. Booting the server from DBAN media and starting up DBAN

After DBAN was up and running overwriting algorithm was chosen to overwrite the disks data blocks, USA's Department of Defence 5220.22-M algorithm was chosen to overwrite the disks for three rounds to completely erase the disk.

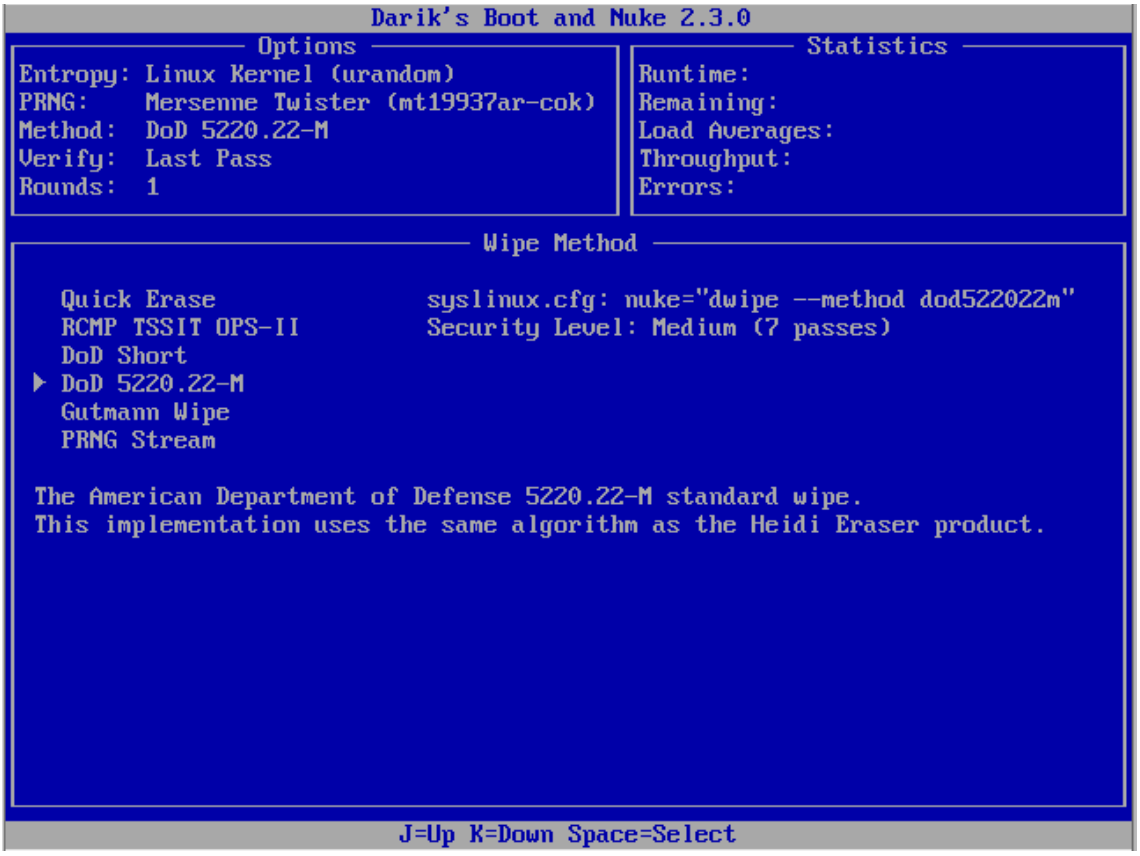

Figure 25. DBAN's algorithms

After the wiping settings were configured to the DBAN it was time to erase the disks. The wiping took about one day per server as can be seen in Figures 26 and 27.

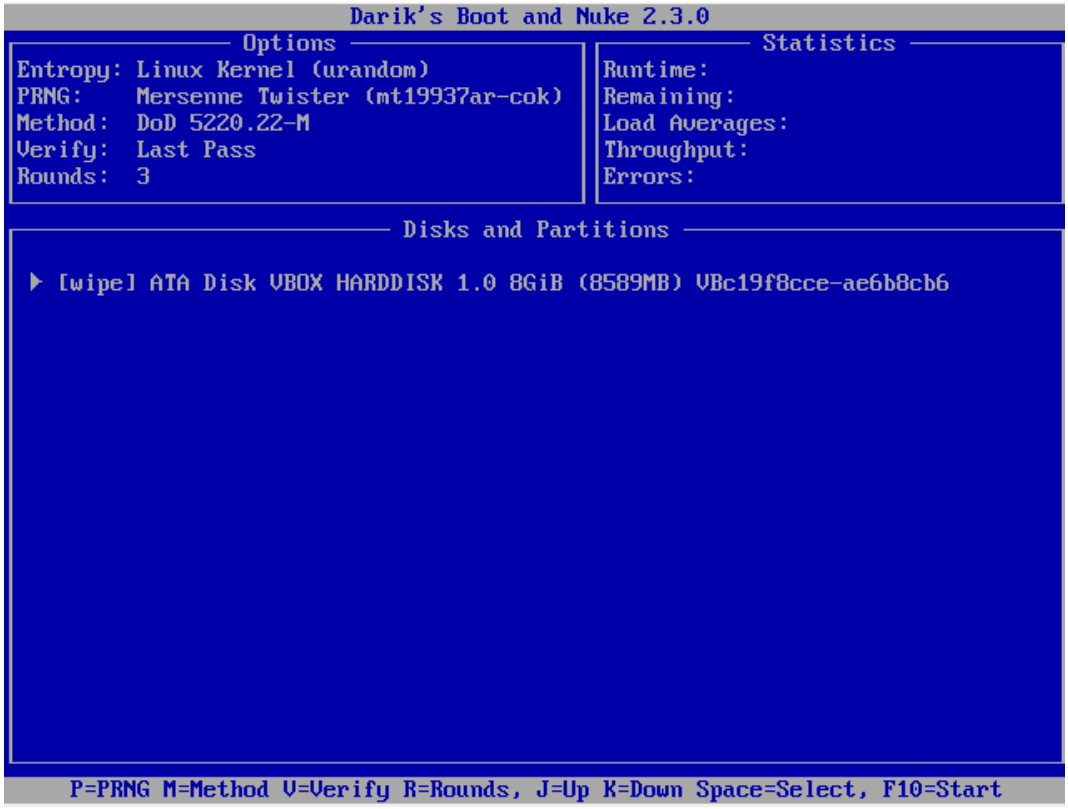

Figure 26. DBAN configured to wipe the disk

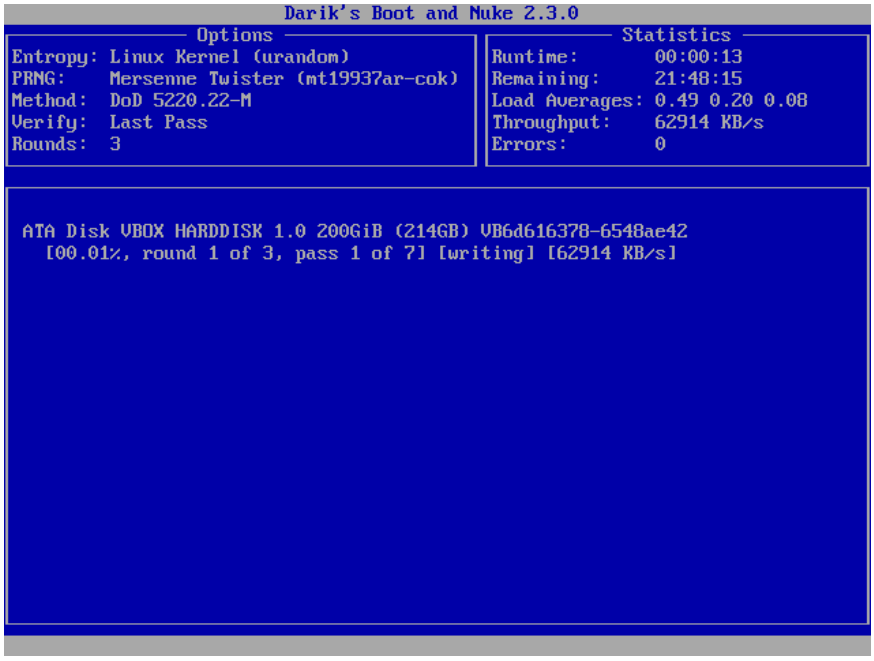

Figure 27. DBAN wiping the disk

7 Discussions and Conclusions

The aim of this thesis was to explain the basics of virtualization technology and the functionality of VMware's virtualization platform vSphere. These related to the practical side of this study where the aim was to migrate virtual servers to another platform.

The study began by thoroughly explaining virtualization technology, its architecture and some use cases. Also, it's history was briefly gone through.

After this the study focused on VMware and its vSphere products. The core components of vSphere, vCenter, ESXi, DRS, vMotion, Storage vMotion, HA and virtualized networking were explained.

The study's practical side began by going through the old and the new platforms and their capacities. After which began the migration planning and preparations for the migration. The preparations included some networking changes and SAN changes. After the changes the migrations began by doing couple of test migrations and testing was the migration successful or not. It was concluded that the test migrations were successful and the production servers could be migrated. The migrations were done in co-operation with the client.

After the migrations were done began the cleanup of the old platform. In the cleanup the old networks, storages, monitoring's were deleted.

This study benefited both the ISP and the client by decreasing the number of physical servers and environments the ISP had to maintain and also gave to the client's servers new and upgraded virtualization platform. Also, the client didn't have to pay anymore for the dedicated server hardware which was used to run the old platform. This lowered the client's Opex costs.

As a result of this study the virtual servers were successfully migrated into the new VMware platform without a long service break. The old host servers were successfully wiped of any data and were given to the R&D team. Both the ISP and the client were happy that the server migration was successful.

References

1 Virtualization Essentials. https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/ebook/gated-vmw-ebook-virtualization-essentials.pdf Accessed 29 Mar. 20

2 Timeline of virtualization development. https://en.wikipedia.org/wiki/Timeline_of_virtualization_development. Accessed 25 Jan. 2020

3 Nick Marshall, Grant Orchard, Josh Atwell Mastering VMware vSphere 6. John Wiley & Sons. 2015

4 What is desktop virtualization. https://www.citrix.com/fi-fi/glossary/what-is-desktop-virtualization.html. Accessed 09 Aug. 2020

5 Citrix Virtual Apps and Desktops techical overview. https://docs.citrix.com/en-us/citrix-virtual-apps-desktops/technical-overview.html accessed 21 Aug. 2021

6 Understanding Full Virtualization, Paravirtualization, and Hardware Assist. https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/techpaper/VMware_paravirtualization.pdf Accessed 15 Aug. 2021

7 Docker – Architecture. https://www.tutorialspoint.com/docker/docker_architecture.htm accessed 15 Aug. 2021

8 Docker overview. https://docs.docker.com/get-started/overview/ accessed 15. Aug. 2021

9 Network Functions Virtualization (NFV). https://www.vmware.com/topics/glossary/content/network-functions-virtualization-nfv. Accessed 16 May 2021

10 Network Functions Virtualization Overview. https://docs.vmware.com/en/VMware-vCloud-NFV/2.0/vmware-vcloud-nfv-openstack-edition-ra20/GUID-FBEA6C6B-54D8-4A37-87B1-D825F9E0DBC7.html. Accessed 16 May 2021

11 VMware Timeline. https://www.vmware.com/fi/timeline.html. Accessed 16 Aug. 2020

12 VMware VMotion. https://www.vmware.com/pdf/vmotion_datasheet.pdf. Accessed 27 Jun. 2020

13 vSphere Standard Switches. https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.networking.doc/GUID-350344DE-483A-42ED-B0E2-C811EE927D59.html. Accessed 27 jun. 2020