DTU Informatics
Department of Informatics and Mathematical Modelling

Integration of Virtual Machine Technologies for the support of Remote Development Units

Master's Thesis

Szilárd Csibi

18.08.2013

DTU Department of Informatics and Mathematical Modelling
Technical University of Denmark
Building 303B
Kgs. Lyngby, 2800
Denmark

Integration of Virtual Machine Technologies for the support of Remote Development Units,
M. Sc Thesis

Author:

Szilárd Csibi ,s111336

Supervisors:

Stig Høgh : Associate Professor

Software Engineering

Department of Informatics and Mathematical Modelling Communications

DTU

Allan Møller: Infrastructure Expert Citrix Technologies and Active Directory

Nordea

Project Period: 15.03.2013 – 16.09.2013

Comments: This thesis corresponds to 30 ECTS points. The thesis is submitted in partial fulfillment of the requirements of the Master of Science program in Department of Photonics Engineering at the Technical University of Denmark

Language: English

Acknowledgements

This master project puts an end to one very important period of my life – the pursuit of my master degree. When I turn myself back through time, I can say that I could not make it so far without the precious help of so many people. With these lines here I would like to express my gratitude towards them. My first thoughts go to my family who always supported me in every decision that I have made so far.

Of equal importance were the people that I have met throughout my journey in Denmark, my school colleagues and friend who have offered me support and guidance not only through the education process but also the through the long Danish winters.

I would also like to thank Nordea, for giving me the opportunity to work in the leading edge of IT technologies, my managers and of course my colleagues who have provided me with much needed support, and understanding. I would like to offer my gratitude to a dear colleague and friend Allan Møller, who besides being a good supervisor is also a good teacher and has introduced me in the complicated world of enterprise infrastructure management

Abstract

The exploration of virtual technologies, especially the Virtual Desktop Infrastructure (VDI) based on VMware incorporated and Citrix technologies as a solution for providing a redundant and safe working environment for remotely located development units, including the presentation of a production level implemented solution will be the main focus of this paper.

The advantages presented by virtualization technologies will also be presented together with the integration of these technologies in providing a modern, flexible application delivery solution that complies with the requirements of the Bring Your Own Device concept.

Table of Contents

Chapter 1

Introduction

In the ever changing world of IT, virtualization has been and still can be considered one of the hottest topics. It has become a hype to have virtualized elements in your infrastructure long before desktop virtualization was even considered. In the late 90's virtualization was mostly confined to virtualization of certain applications and lab environments, so that users could conduct tests and simulations without being confined by the onsite limitation of a working lab. The last decade has seen a surge of development in the virtualization area, especially in server virtualization which managed to usher in a new era in IT. After server virtualization is already considered not only a must, but a well-accepted production, solution desktop virtualization is starting to have its zenith. Architectural concepts including migration to the datacenter, and Bring Your Own Device (BYOD)are pushing for a new, quite interesting form of virtualization, that could provide users access to their work environment from any execution platform that has the basic system resources like CPU, memory disk and network.

In the last decades the working tool of many employees has become the individual desktop; because of this large organizations are facing the daunting task managing thousands of individual workstations. The biggest issue is the lack of uniformity in these individual devices, every employee has needs for different applications, has a different style in setting up their own workspace. Before the use of automated application delivery tools, most of the execution platform maintenance was done manually, not only slowing down the production process to which these desktops were essential, but also demanding a large and highly specialized IT workforce. The lack of an easy centralized backup system also meant that in case of a hardware malfunction of an individual desktop, irrecoverable data losses were a reality, a situation which is inadmissible for a production environment. Another encountered problem was the lack of mobility. Initially due to bulky desktops, but ultimately chocked by security, mobility is still a big issue when considering the access to big backend infrastructures, mainly because numerous security requirements have to be met before any access is granted to inner networks. This is usually achieved by a personalized laptop that is configured in a way that can guarantee basic security features, like a valid antivirus, or a valid operating system. These together with a preregistered user profile and credentials enable a user to get access to resources places behind a multitude of firewalls. During this paper the way virtualization solves most of these issues will be presented.

This work is mainly focused on issues regarding Desktop virtualization, and solutions for virtual desktop infrastructure (VDI). The basis for this project is the detailed presentation of a

working VDI solution as implemented at Nordea for providing a secure working environment and better mobility for foreign development units. The initial part contains a description of the Virtual Desktop Infrastructure and a comparison between the implemented solution and other existing solutions, the reasoning for why the implemented solution was chosen. The main part is the actual description of the VDI solution . The final part is a presentation of possible improvements on the current configuration as well as the presentation of new trends and new concepts that are still in the pre-production niche but could be the next big breakthrough in the fast developing world of virtual desktop infrastructures and the description of recovery procedures that are key for production environments.

1.1 Application and Desktop delivery

The main goal of the application and desktop delivery process is to offer users the possibility to work onsite, offsite, offline basically anywhere while using their own device. Bring Your Own Device (BYOD) has become a strategy followed by many leading organizations, which are trying not only to improve employee satisfaction and productivity by granting them the opportunity to work from anywhere on any device, but also optimize application management, ease the procedures through which a user can have access to his/hers needed applications and of course improve redundancy and security.

When choosing an application and desktop delivery solution a few key issues need to be considered. First of all and the most important decision is the choice of an execution platform. All applications need resources like CPU, memory, disk and a network in order to be able to run an operating system (windows), web applications, and mobile applications. The most frequently used execution platforms are: Desktop, Laptop, Tablet, Smartphone, VDI and Server Based Computing (Spruijt,2013)After the execution platform is chosen the way applications will be delivered and managed has to be decided. As presented in the introduction large number of individual desktops can cause difficulties in application deployment and maintenance. Last but not least is the question of accessibility, mobility. With the development of powerful mobile execution platforms like smartphones and tablet the application and desktop delivery is not confined anymore to traditional workstations. Even if these mobile execution platforms can deliver certain functions, in most of the cases they do not possess the hardware to deliver applications with elevated processing requirements. This is where virtualization comes in.

Chapter 2

Desktop Virtualization

2.1 Introduction

In essence virtualization is nothing more than the decoupling of IT resources (Spruijt,2013), it is a smart software layer on top of an existing hardware configuration that is capable of emulating and reproducing the behavior of multiple standard physical units, without the need for dedicated hardware components for every unit. The software that performs this is called a Hypervisor, and the scope of an ideal hypervisor would be to provide an environment for the software that is exactly like a host system, but without dependencies to individual physical components.

Virtual machines are the logical equivalent of physical ones, and the reason for the waste spread of virtualization is that multiple virtual machines can be hosted by one physical machine, this way virtualization not only provides a clean cut alternative for physical machines, but also contributes significantly to the optimization of datacenters, by reducing the number of physical servers, and improving on the percentage of their use. The latter is achieved by enabling multiple virtual machines on one hardware unit, this way a better utilization of expensive equipment can be achieved.

2.2 Main functions provided by a virtual desktop

The main objective of any IT infrastructure is to provide end users access to windows, -web and mobile- applications. The virtual Desktop (vDesktop) is an essential component in any modern Desktop delivery solution, due to its capability of providing the following functions as presented in (VDI whitepaper,2013)

1. Bring your own device (BYOD) : enables the delivery of applications and desktops for BYOD scenarios,
2. Access: vDesktop works independently of locations, endpoint and network
3. Security: It is server hosted, everything is in the data center, hence there is an increased security. Being centrally stored data can be easily backed up, and data thefts can be better avoided.
4. Freedom: every user can be assigned their own desktop with administrative privileges when needed.

5. Management: vDesktop is centrally managed and hardware independent
6. Sustainability: Power Management, handling the necessary resources in an efficient manner

The Bring your own device concept has been and still is a strong motor in the development of virtualization solutions mainly because it is based on user centric computing. Since every user wants applications to be available from a multitude of devices (phone, tablet, desktop, laptop etc.) the functionalities of a vDesktop is needed

2.3 Choosing an application delivery platform

The sheer variety of different platforms the applications have to run on, demands the design of hybrid style and flexible applications. Delivery of applications and data to the user needs to be transparent, and in order to achieve this transparency a set of elements have to be known:

1. Who is the user, what is his role and what is he allowed to do?

   This is the first step in assigning a vDesktop. Based on login credentials and optionally other factors like geo-location the access session will be automatically subjected to a set AD (Administrative Directory) rules that govern every aspect of the connection instance.

2. What applications are being used?

   Since every user has the possibility to be assigned a personal vDesktop, a set of rules have to be defined regarding the availability of certain applications. Starting from an initial uniform application list, based on requirements, internal rules, and subjected to line manager approvals any number of application can be added.

3. What device is being used?

   It is a vital step, because different devices have different operating systems, different limitations and capabilities so the delivery of the applications has to be tweaked or in some cases totally modified in order to be accessible from that particular device.

2.4 Desktop Virtualization types

Desktop Virtualization is the detachment of the desktop, the operating system and the end-user applications from the underlying endpoint or device. This kind of virtualization can be subdivided into two main categories: (Spruijt,2013)

Server Hosted - is where end-user applications are executed remotely and presented at the endpoint via a remote display protocol. Within this there are 3 types:

- Shared desktop (RDSH) – session virtualization, which is also commonly used for publishing single applications

- Personal virtual desktop (VDI) – Virtual Desktop Infrastructure
    - Non-Persistent : . Non-persistent VDI's as the name suggests only exist if they are being used. Every time a user tries to log on a new virtual desktop is created from a master image and is deleted immediately after the user logs off.
    - Layered – desktop components are separated in layers, with both persistent and non-persistent components
    - Persistent: Logical equivalent of physical desktop
- Personal physical desktop - (BladePC)

Client Side - is where applications are executed at the endpoint and presented locally on this workstation.
Within this category there are 2 types:

- Bare-metal –citrix XenClient is running on the machine without an underlying operating system
- Client-hosted –an operating system is used over which additional applications like VMware workstation is installed

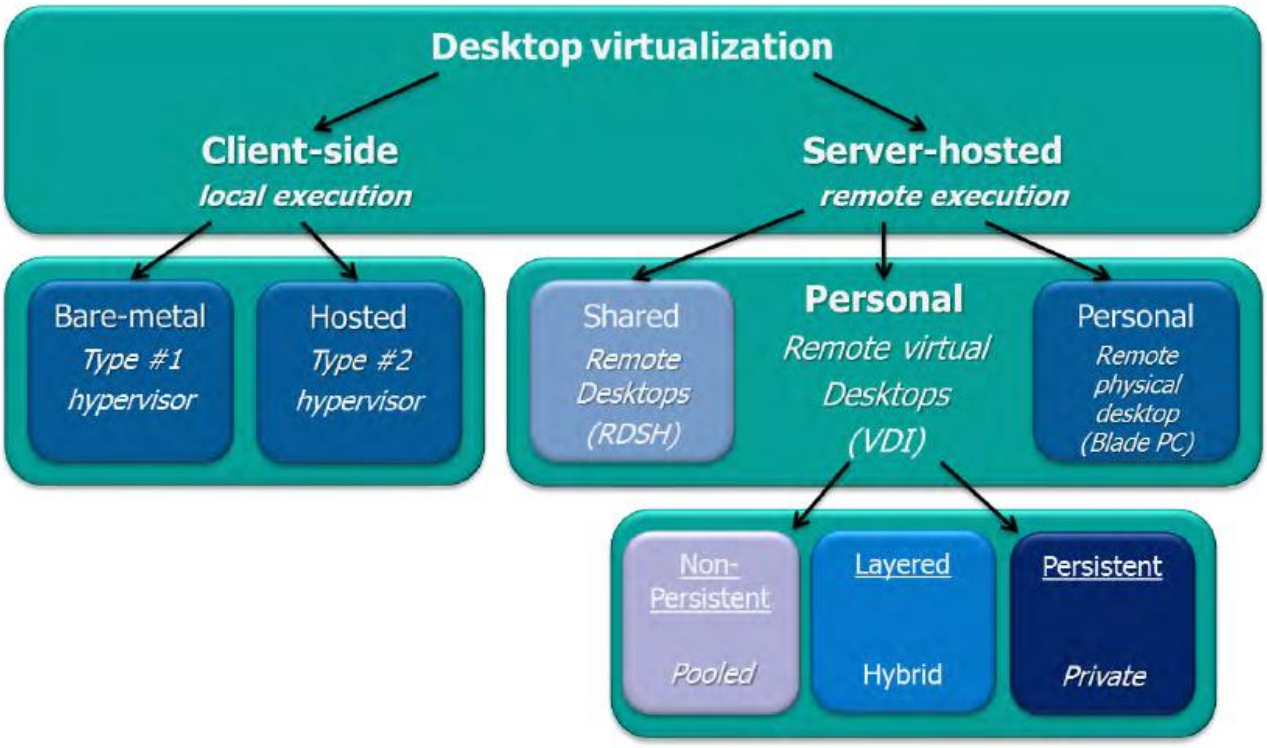

Figure 1. Desktop Virtualization Solutions (VDI smackdown)

The desktop virtualization this project focuses on is the Server Hosted, personal persistent virtual desktop. Server hosted desktop virtualization is a solution for accessing Windows7/8 or legacy Windows desktops that are executed remotely on a virtual machine in a datacenter

(Spruijt,2013).Due to the Enterprise nature of this project the server hosted virtual machines are an easy pick, mainly because the datacenter platform on which this solution can be implemented already exists, but also because the sheer number of the users.

The type of VDI chosen was the persistent, private virtual desktop. This was chosen with the scope of providing users with a similar experience and similar freedoms as those encountered in physical desktops. These include personal settings, possibility to personalize certain applications, freedom to install software within the desktop and most importantly to maintain all these changes in between the reboots of the operating system. Even if the stateless virtual desktops have an advantage in simplicity of management and an ease of rollout due to standardization they are not adequate for a developer-user oriented system.

Another major reason for choosing the persistent VDI solution is to minimize the changes the support elements have to go through. Since the VM (virtual machines) can be considered just as another PC, existing creation and maintenance procedures and processes can be reused.

The main disadvantage of persistent desktops is the high cost of datacenter equipment like SAN (Storage Area Network) storage, but as it will be presented there are solutions to reduce costs by using thin provisioned virtual machines and efficient management systems.

2.5 The lifecycle of virtualization technologies

As previously presented the Virtual desktop infrastructure (VDI) can still be considered a rather new technology, compared to server virtualization it is still in the niche category of virtualization. Currently only 2-3% of the overall desktops in use are VDI's. (Brian Madden 2010). In virtualization, similarly to any emerging technology there is a pattern of expectation/reality fluctuation in time, starting from the deployment phase. This fluctuation is observable in the Gartner Hypercycle (Figure 2,3) where the variation of expectations in accordance with time in relation to virtualization products and solutions can be observed. The initial stage of any virtualization product brings with it an increase in expectations, mostly due to two factors. First of all the developers of the solutions have wild promises on how good their product is and what problems it could solve without any difficulties. This initial expectation high is seriously influenced by the start of the actual deployment of the product in production and test environments that is usually met with the birth of a multitude of unknown issues that have to be troubleshooted, or in some cases spell the end of the use of the virtualization solution for a number of cases. In the case that the issues can be fixed there will still be phases where the customers have to get used to the new product, and have to actually see the improvements introduced. Only after the utility and the improvements are confirmed can a product be considered good enough to be introduced fully into production.

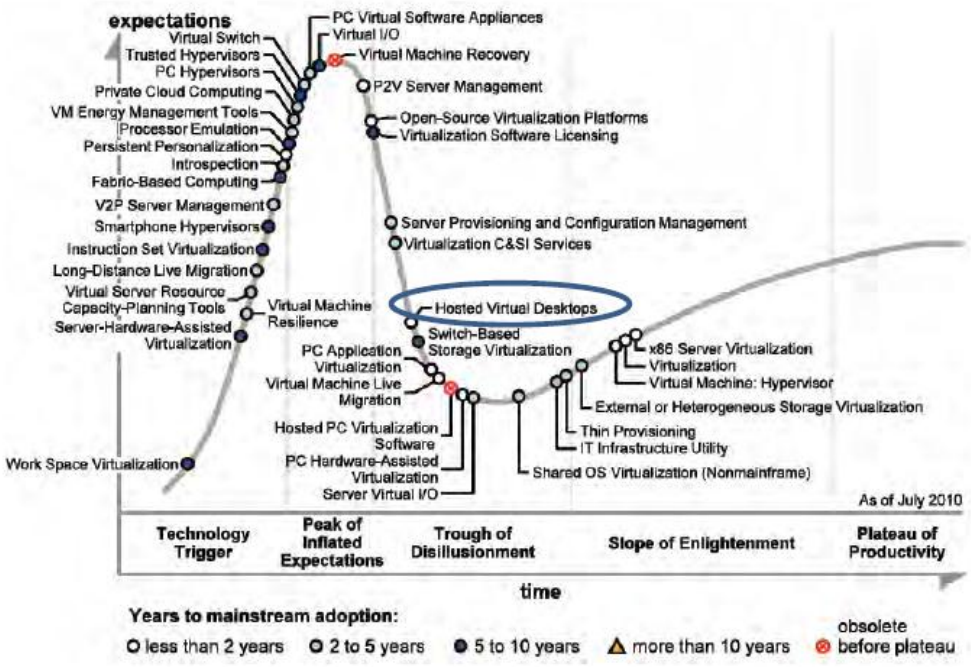

Figure 2 Gartner Hypecycle of virtualization 2010

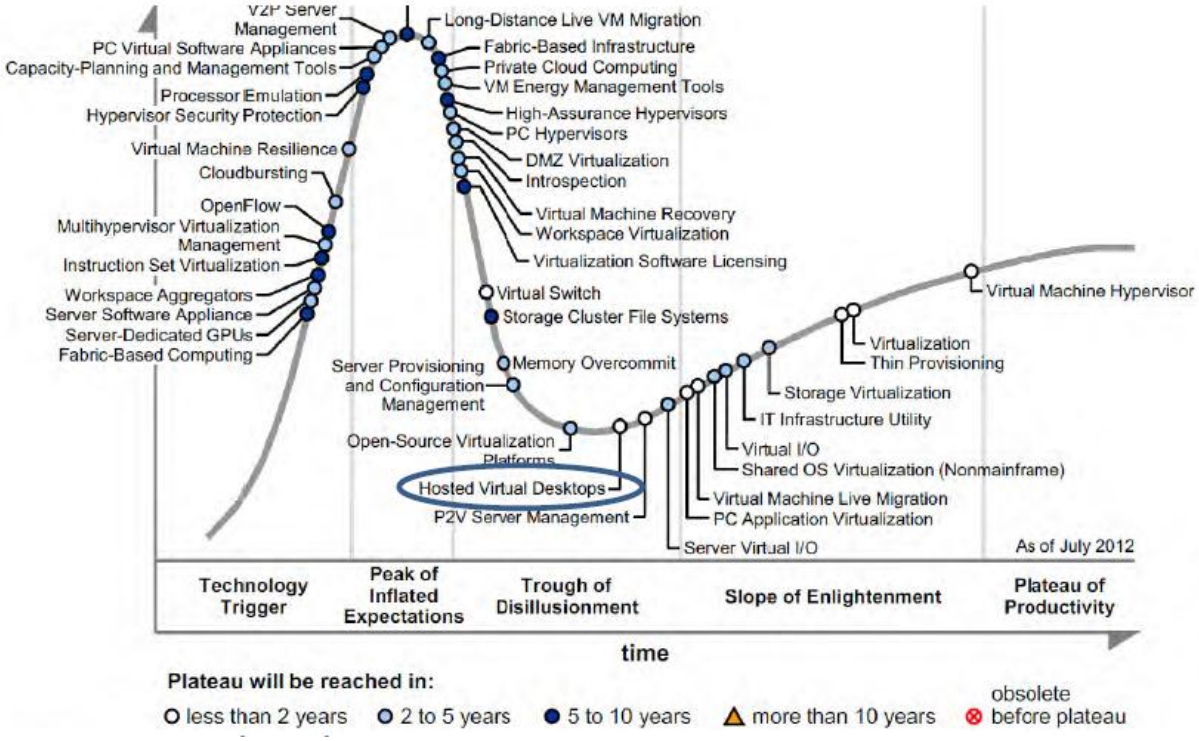

Figure 3. Gartner Hypecycle of virtualization 2012

These observations are applicable for the VDI technology as well. Even if, as will present further on, there are a lot of advantages in using VDI-s, but it is still not the silver bullet that will eliminate all the challenges in the IT world of desktop delivery. There always has to be an objective analysis on the applicability of a new solution in one's infrastructure.

The first Hype cycle representation (Figure 2) presents the location of the VDI solution in 2010. After the initial hype about the advantages of the VDI, in 2010 the lifecycle of the VDI reached the disillusionment stage, where it became obvious that it is not the best solution for everything, and there could be better answers for some technological hurdles than the VDI. On the other hand in figure 4 from 2012 we can already observe the advance of the product towards the plateau of productivity, this being explained by the fact that it has been proven useful and truly a production solution, but still lags behind other virtualization solutions like thin provisioning, storage virtualization, virtual I+O, solutions that are already implemented and used on a daily basis.

True to this graph the predecessor of this project on which this thesis is based on was developed and piloted in 2012, and currently has been in production without many issues for more than 9 months. The pilot project provided the virtual desktop solution for almost a 50+users, mostly developers in India and also 3rd party consultants, but the success of the solution has prompted further development, and the introduction of this technology for more developers in India and the Scandinavian countries. Future prospects are quite good with plans to provide easily accessible virtual workstations for all the users in Nordea, having the scope of eliminating the Virtual Private Network (VPN), and the need to carry a company sanctioned PC in order to be able to work from home.

The presented project is also in a continued morphing stage with currently having more than 300 VDI users compared to the initial 120. This proves the versatility and utility of this particular solution, and puts it on the forefront of future IT development strategies.

2.6 An objective analysis on advantages and disadvantages of the VDI

In accordance with its barely production worthy status the VDI has still to find its customer basis, more precise parameters that define when this virtualizations solution is viable for production. In the following stage I will present a couple of the advantages that this solution can bring, together with some of the drawbacks of particular characteristics, by performing an analysis on opinions presented by two of the main developers on the virtualization stage : Brian Madden and Ruben Spruijt.

2.6.1 Reduction of costs

As true in any product, the positive financial aspect of VDI is a strong motor if trying to assure wide deployment. As presented in the implementation chapter, elements like thin clients, efficient centralized management consoles, lack of physical desktops, lower maintenance costs etc. would ensure a gain in the financial side. This all sounds good, but it is not applicable in every case. The fact that it is a datacenter based solution, meaning that it is inaccessible (in the form which is presented in this project) for smaller companies that do not have the infrastructure to support the virtualization. On the other hand the cost factor of the VDI highly depends on the way we construct and choose the right VDI solution. In this case if the improvements presented in the Improvement chapter will be implemented including provision of the OS and virtual applications the cost of an individual VM will fall due to the reduced costs in operations and maintenance.

2.6.2 Better security

Better security is one of the main requirements of this project. With all the data being stored in data centers and not at the endpoint the safety of the information is assured by the high level security infrastructure already implemented for a large datacenter. The use of private tunnels and the use of SSL (secure socket layer) encryption ensure an enhanced security. The disadvantage is that the security elements do not come with the VDI itself, but are inherited from the infrastructure of the datacenter.

2.6.3 Mobility

This can be considered one of the main advantages and main reasons for the development of the VDI. The BYOD concept can be accomplished to a certain degree thanks to virtual desktops. Since the hardware demanding computation is all done in the datacenter, the device on which the VDI runs can be a basic unit. Most importantly it can be any device with an operating system, CPU and memory. This gives users a lot of mobility, the possibility to access their work desktop from their phones, from hotel computers etc.

The problems with this is that it is still not applicable for every type of device and since the user needs an internet connection to access the VDI it is not usable in offline cases. The offline usage is also limited by the fact the most of the applications used require online access to data and services that are located in the datacenter.

2.6.4 Reduced downtime due to hardware failure and better disaster recovery

These advantages are mostly also inherited from the datacenter structure (VDI in a box also offers a high level of redundancy). By using the datacenter architecture as a building base for VDI, the redundancy, recovery, availability requirements that are demanded from any datacenter automatically apply for the VDI's as well. Elements like high availability, segmentation based on location (existence of multiple connected datacenters), automated recovery procedures that were previously only applied in the server environment become a standard feature for the virtual desktop. The disadvantage of course is that nothing comes for free, all these features are expensive and the degree of the actual application of these measures depends on the requirements of the project, and the financial calculations.

2.6.5 Easier image management

Having a large network of workstations can increase the strain on the local networking when management tools are used. For example when a new update has to be distributed it would take considerable time and effort to apply it to every workstation, not to mention the capacity reduction in the internal network. With VDI when a workstation has to be rebuilt it is not necessary to have physical access to the endpoint hardware(laptop/PC) ,because everything is done through the data centre.
It also has to be considered that there are tools capable of offering similar capabilities when considering image management, like SCCM which does not require a VDI infrastructure, and can still provide the same benefits, and these tools are much cheaper than implementing a new solution like the virtual desktop. Cost issues have to always be considered.

2.6.6 Better user isolation

One of the biggest issues with terminal server based solutions is that multiple users use the same copy of the operating system and the same resources that are distributed to that particular terminal server. Terminal server solutions are quite useful if all users require the same applications, but if every user needs a particular application, publishing all these applications is quite difficult. Besides this some of the windows applications simply do not work under multiple user access situations. VDI solves these issues by providing the user with an isolated machine. Every user has a different copy of Windows, and thus is capable of having their own applications and their own personalized desktop and application palette.

The issues is that in case of a simple VDI solution where provisioning and application layering is not used a better isolation will be achieved but at high costs and only in the case of improved VDI solutions that include provisioning and application layering a solution with good user isolation and reduced costs can be achieved.

2.6.7 Reusable knowledge in applying virtualization

Since server virtualization has been around for more than 10 years, it is already a known technology and the expertise required for its implementation is easy to find, so there should be no difficulties in applying a VDI solution from a manpower perspective.

As a conclusion to this subchapter it is clear that the VDI has quite a few advantages, but cost issues have to be considered and a value over gain variable has to be calculated based on requirements, available resources and already existing infrastructure.

Chapter 3

Virtual Infrastructure

3.1 Introduction

Virtualization of computer hardware dates back to the 1960s when the IBM System 370 Mainframe (Creasy,2011) first introduced this concept, and has matured to a stage where today every fortune 100 company utilizes this technology (VMware, 2011). Virtualization is the technology that is used to create virtual machines from standard physical resources. One of the main purposes of virtualization is to enable a higher utilization of resources, better and easier maintenance, and also a better utilization of space in offices, datacenters.

The virtualization process is enabled by a software layer called a hypervisor, a product that is being produced by all the leading IT companies including Microsoft, Citrix, VMware, Red Hat, Oracle and others. Even so, the leader in this area is VMware. Virtualization technology takes advantage of the resources in Intel and AMD based systems by creating logical machines that do not exist physically, but have the exact same characteristics and performance as a physical counterpart. Each virtual machine is configured with an operating system and software, which makes the virtual desktop undistinguishable from a usual desktop.

The main advantage of a virtual machine is that it can be reached from any location on a multitude of devices, so the user does not have to be in the vicinity of the actual hardware in order to utilize its capabilities. This can be especially useful in high security configurations (the internal network of a bank) or in cases where a lot of processing power is required that would be difficult and inefficient to transport. The only thing needed is either a connection to the local network on which the virtual machine is accessible or an internet connection, in which case certain security requirements have to be met.

Hardware virtualization has seen a huge success in server virtualization, mainly because before virtualization technologies the physical server setup was highly inflexible and in most of the cases inefficient and costly. The mentioned problems were caused by the lack of scalability that led to the underutilization of expensive hardware.

3.2 Virtual Machine Architecture

The logical view of the virtual machine architecture can be seen in figure 4 It is a layered structure, which has a physical base of storage units and physical servers, which are coordinated by virtualization software.

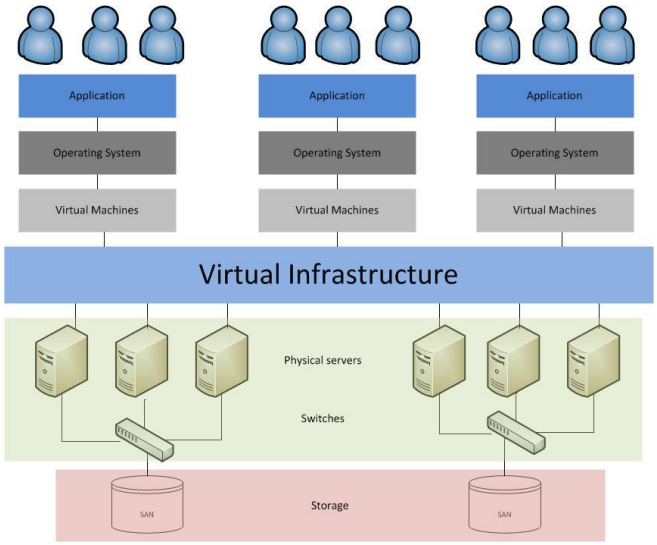

Figure 4 Logical view of architecture

In figure 5 we can see additional details to the virtualization infrastructure layer, the two most important additions being the Symmetrical Multi Processor (SMP) and also the Virtual Machine File System (VMFS). The role of these two elements is to fulfill the virtual desktops processing and storage needs, both components being controlled by the hypervisor.

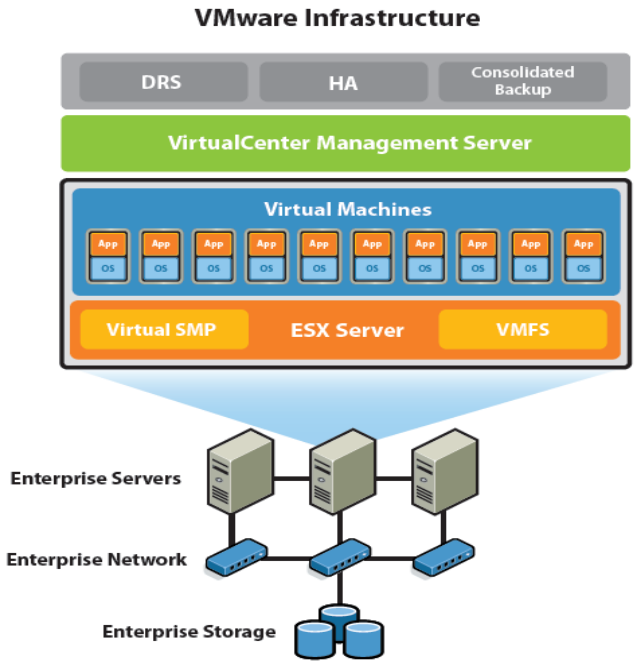

Figure 5 VMware Infrastructure (From VMworld, 2008)

The VirtualCenter Management Server has the role of unifying all the resources from the individual computing servers and spreading these resources to all the virtual units in the datacenter. It also has the role of providing access control, performance monitoring and configuration.

3.3 Virtual Datacenter architecture

The success of virtualization is given by the fact that the entire IT infrastructure (servers, storage and network) is unified into a heterogeneous resource in the virtualized environment. This means that all the resources can be dynamically provisioned with ease and can be assigned to where they are needed without major complications.

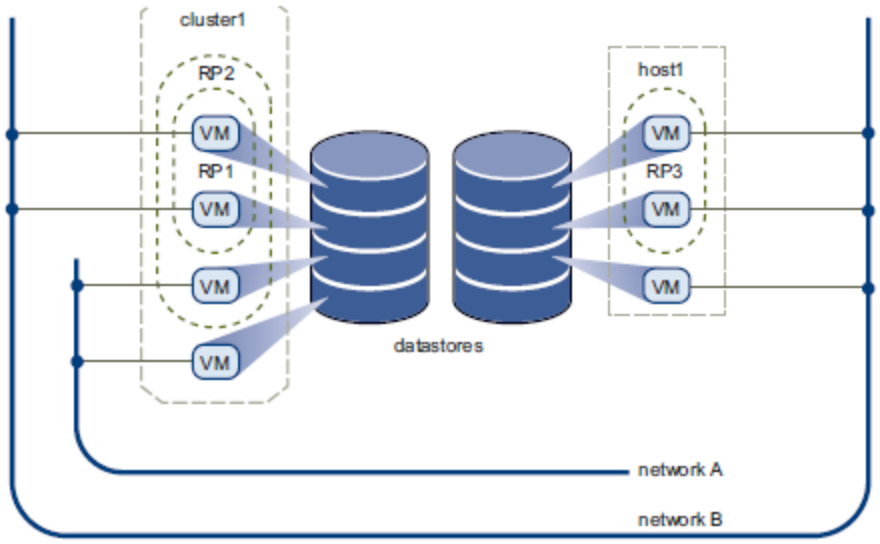

Figure 6. Virtual Data Center architecture (VMware 2012)

The virtual datacenter consists of four virtual elements: computing and memory resources (Hosts, clusters), storage resources (Datastores) ,networking resources and virtual machines. The physical computing units and memory resources are physical machines running ESX Server. These machines are grouped in clusters in order to be managed as one. Physical units can be easily added to the clusters depending on the computational and memory needs. The storage units are represented in the virtual world by the data stores. In the virtual environment the networking layer is augmented by virtual networking, which is essential in providing optimized and more secure connections. Beside the logical connections which mainly are the interconnection of the created virtual machines, the networks also connect the virtual environment to the physical network outside and inside of the datacenter.

3.4 Hosts Clusters and Resource pools

The host and cluster logical configurations are a key element in providing high levels of flexibility for the virtual system. As an example we can have the situation presented in figure 4, 3 physical servers each having 4 GHZ of computing power and 16GB of memory..

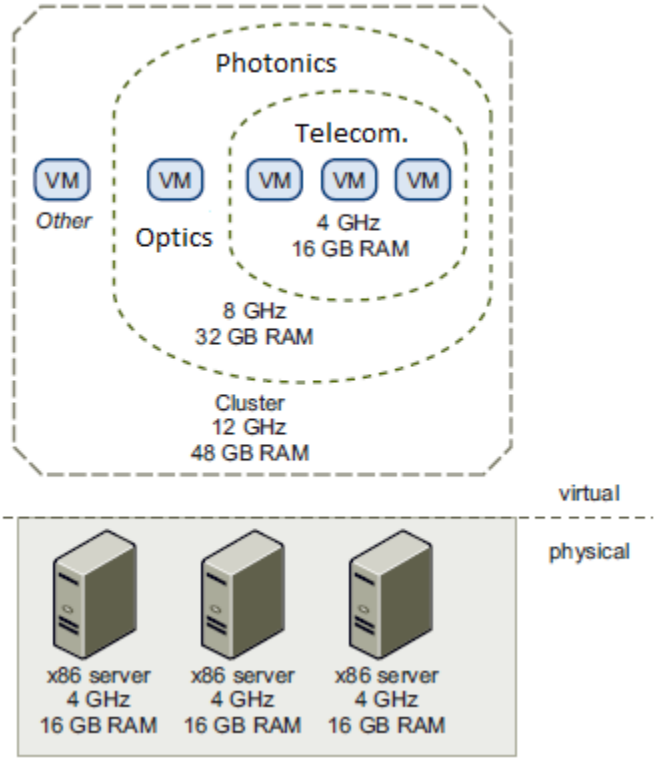

Figure 7. Example of Resource Pools and clusters

This means that the cluster centre has 12 GHZ processing power and 48GB RAM. These resources have to be spread between a few different departments in accordance with their needs. Let's say that the Photonics department at the DTU is a resource pool where the Telecommunications department is assigned a third of the available resources, and Optics department the same amount of resources. This means a third of the resources are still available for other departments inside of Photonics. The main advantage of this configuration is that if the Telecom department is not using all its resources and the optics department is in need of more computational power the system allows the Optics group access to the available resources that are not used by the Telecom. department. Another advantage is that if in the following years the resource demand of any department increases the resources can be dynamically changed as to fit the emerged requirements. These changes can be made without shutting down the virtual machines, which is quite useful in cases where the virtual machines need to be available all the time. As we could see in this example the expensive hardware

resources are optimally used by allowing the flow of resources into departments where the need arises.

3.5 Network architecture

The network architecture of the virtual environment is configured to be similar to the physical environment which include virtual interface cards (vNIC), virtual switches with the addition of certain elements like Port Groups that are absent from physical networking due to fact there is a limited number of network cards that can be added to the system.

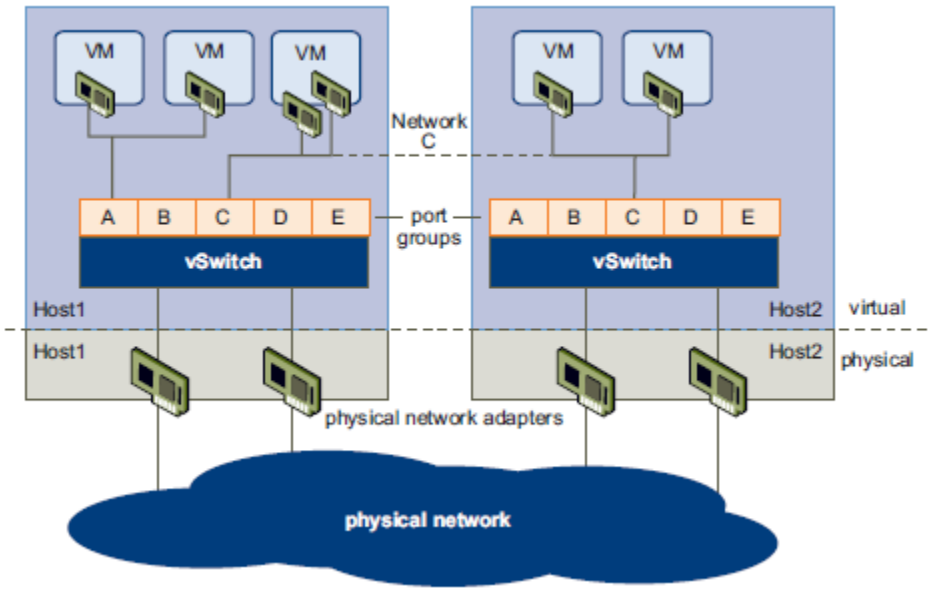

Figure 8. Networking architecture (VMware)

As we can see in figure 8 every virtual machine has its own virtual network card. The operating system talks to the vNICs by using a standard device driver and to the outside world each virtual NIC looks like a normal physical NIC with its own MAC address and its own IP address, and responds to Ethernet protocol as any physical NIC would (VMware) .

Each physical server is assigned its own vSwitch, that has logical connections to the virtual machines through Port Groups and connects to the physical Ethernet adapters that are located in every physical server. For redundancy and load sharing purposes multiple physical NICs can be coupled.

The biggest improvement in networking comes from the utilizations of Ports groups. Through these logical elements virtual machines can be connected to different virtual networks. For

example as we can see in figure 8 if desired separate networks can be created, if a virtual machines vNIC is connected to the C Port Group it means that the mentioned VM is part of the virtual network of all virtual machines that are connected to the C Port Group, even if the virtual machines are positioned on different hosts, and this goes vice versa, or if two virtual machines are located on the same hosts it does not imply that they are located in the same virtual network. This feature is quite useful if a more segmented configuration is needed. It also improves security by ensuring that in the case of a security breach only some of the virtual machines, not all the VM located on a host are affected. By using Port Groups different network policies can be applied for every virtual network and also the traffic management can be improved.

3.6 Storage architecture

The large variety of storage systems such as Fiber Channel SAN, iSCSI SAN, Direct Attached Storage, poses difficulties in management and provisioning. In the virtual VMware environment this problem is solved by using a layer of abstraction that can manage the differences between the different storage systems and present it as a whole.

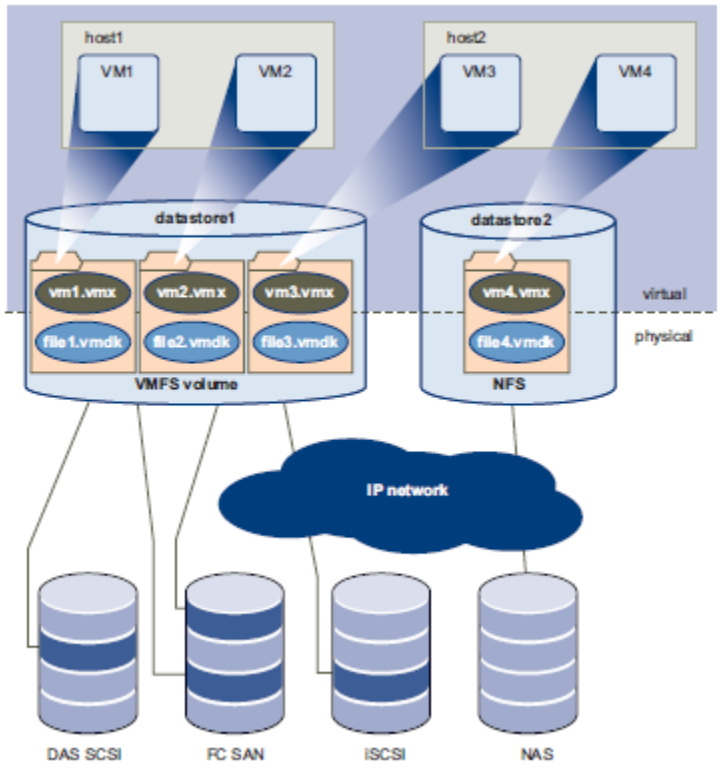

Figure 9. Storage architecture (Vmware 2011)

In figure 9 we can see that the data store is responsible not only for providing storage space for the virtual machines but also for storing the virtual machines. The virtual disks that are provisioned for a certain virtual machine are 'added' by simply adding a set of code lines into the files that represent virtual machines. This feature enables virtual drives to be added easily, it is similar to file manipulation, and they can even be 'hot added' (Vmware 2012) which basically means that they can be added without powering down the virtual machine.

In this project the data store can be regarded as virtual machine file system (VMFS) file system. This file system enables the data store to incorporate multiple storage units. As we can see in figure 10 a data store is capable of accessing many different storage compartments by using their Logical Unit number (LUN) . The LUN is configured when the storage system is set up. The main advantage of using the VMFS file system is that the storage units can be used simultaneously by multiple physical or virtual servers. Another advantage of using the VMFS file system is that, because it can be easily manipulated, a locking mechanism can be implemented.

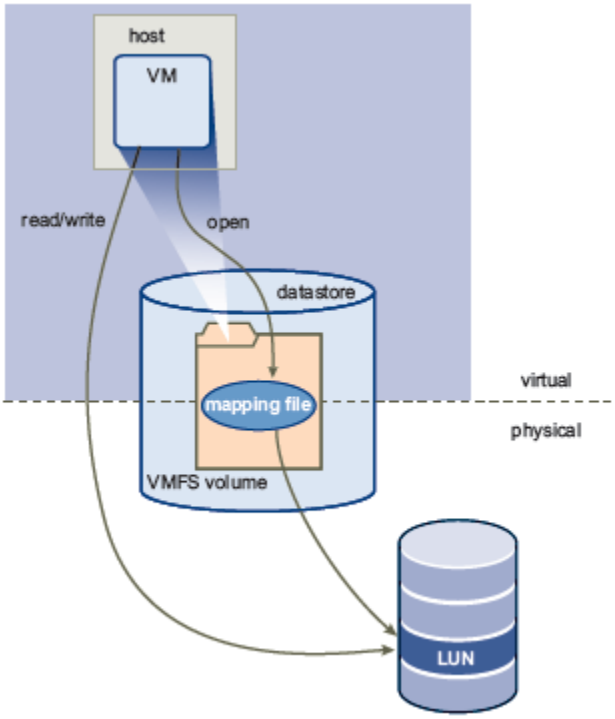

Figure 10 Raw device mapping (RDM)

For a better and faster handling of data stored on the storage units RDM is used. RDM is needed to provide direct access to the LUNs for the virtual machines. As we can see in figure 10 virtual device mapping enables a direct communication between the virtual machine and the LUN on the physical storage. To enable this, a mapping file is created in the data store that instead of storing actual data, maps the files on the storage unit and this 'map' is then presented to the virtual machine. This way the VM knows how the access the data that it needs

directly from the storage units. The mapping system is only used in the incipit of the data transfer, after the location information is transmitted to the VM the data transfer goes directly from the VM to the LUN and vice versa , the information does not pass anymore through the data store.

3.7 VirtualCenter Management Server Architecture

In order to manage all the components of the virtual system which were previously presented a management server is needed. The management server has the role of providing a centralized control point form where all the elements inside the virtual environment can be configured, managed and provisioned.

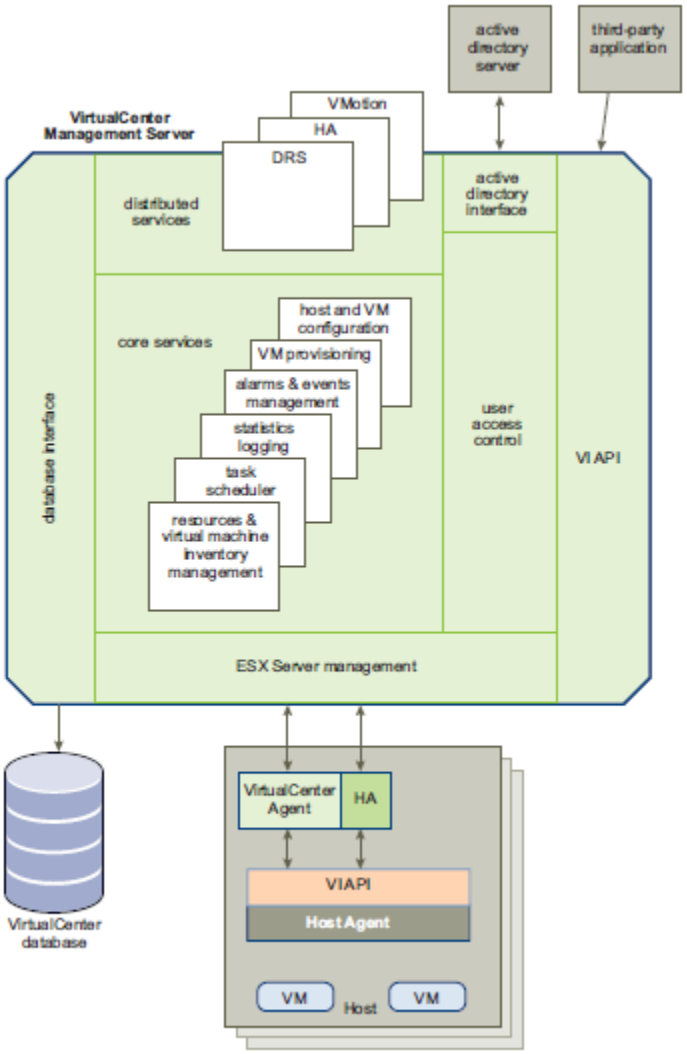

Figure 11 Management Server architecture(VMware)

As visible in figure 11 the management of this virtual cloud is quite complicated and comprises of several elements. The most important ones are the core management services, the user access control, Distributed services and the interfaces to external resources (Vmware).

The core services provide an automated platform for provisioning and facilitate the control of all the virtual elements to the administrator by incorporating elements like logging, alarms, virtual machine inventory.

One of the most important elements in the Managements server is the User Access and control unit. This is connected to the active directory, and thus can provide a controlled and regularized environment where the administrators can manage through different policies the access rights of every user to the virtual infrastructure.

Chapter 4

Implementation and description of the Virtual Infrastructure components used

4.1 Introduction

The design of the VDI solution follows on the baseline of the standard VMware virtual infrastructure architecture. This is essential when implementing, because the existing virtual infrastructure rack design can be used for this solution as well, this way reducing implementation and configuration risks that can occur when using a hardware configuration.

The novelty introduced in the design is the strict definition of the capabilities of every virtual machine. Due to contractual obligations every VM created during this implementation has the same performance characteristics, and is part of a standard developed for the VDI project. As a result every VM has the characteristics visible in Figure 9.

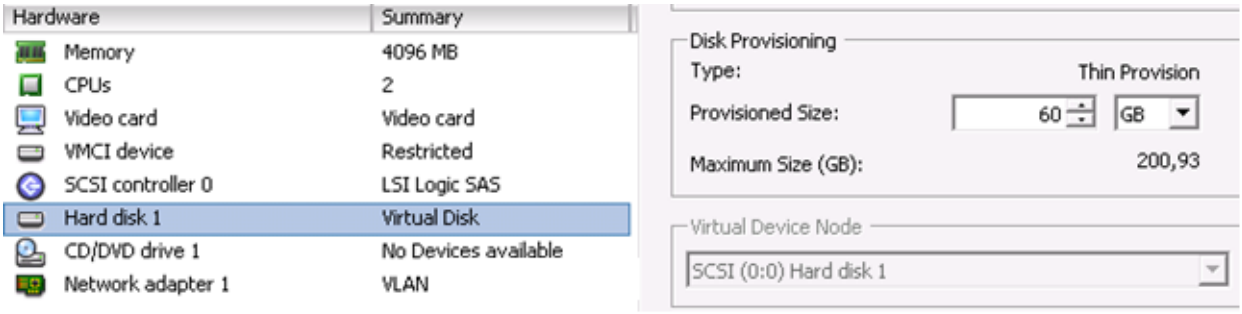

Figure 12 Virtual Machine characteristics

The values have been decided based on requirements presented by the project leaders. This way every VM (Virtual desktop workstation) is allowed to use a maximum of 4 GB of memory and 60GB of storage. For an easier deployment of the virtual machines a template has been defined through which a new VM can be created more efficiently and with similar capabilities without having to configure them one by one, this way improving both speed of deployment but also reducing the risk of human error when creating the machines.

Even if the values presented are fixed, if required by the developer a virtual machine can be enhanced by adding more memory and more storage, however there are some limitations due to the physical component of the virtual infrastructure that have limited capacity. This flexibility

is important to ensure future dynamicity and to facilitate the support for newer and newer applications.

4.2 Design architecture

For security reasons the VDI architecture overview can only be found in the classified section of this paper (Appendix A).

As it is observable in Annex1 the different levels of the VMware structure are clearly divisible. The first layers of this architecture, the storage which is represented by the SAN-s and the switches that connect the VDI system to the storage are located in a network called the management network. The third layer is the layer represented by zone 1 that is basically a LAN specially created for the VDI solution. On top of virtual infrastructure we have the access external connection layer that provides the user with the actual desktop, the mechanism of which will be further presented in the following chapter.

4.3 Resource pools and hosts

One of the main design characteristics for the VDI project is its modularity. Every component has a modular structure, a characteristic that enables more flexibility and room for growth.

The basic modular element of the resource pool for this project is the Fujitsu Primergy BX922. This is a powerful server that has specialized hardware-based virtualization support (Figure 13).

Due to the design of the VDI project the presented servers have an extra physical component. That is an extra network card that enables the creation of the local VDI network in zone 1.

The initial design was made to host 150 users, more than half positioned in Pune India and some developers situated in Finland. Based on this requirement and the initial 6GB/virtual machine it was decided to implement the project on the 192GB Ram version of the Fujitsu blade. The total resource pool was designed to be 8 BX922 blades each having 192GB of Ram, with an expected 32 users per blade. This was later modified to the 4GB/user configuration, partly because based on initial testing it was seen that that amount of resource would be satisfactory for the developers and the used applications but mainly because the expansion of the size of the project. Overall the exact amount of ram allocated is mostly for management and calculation purposes, since over commitment of rams is permitted ,which means that even if all the ram resources have been logically distributed , new VM can be added.

Main Features	Benefits
Performance due to processor technology	
■ Two Dual-Core, Quad-Core or Six-Core CPUs with Intel® Xeon® processor 5500 and 5600 series with Turbo Boost technology, Demand Based Switching, QuickPath Interconnect (QPI) and internal Memory Management Unit. The Intel® QuickPath architecture memory controllers provide the BX922 S2 with a high-speed bandwidth of up to 25 Gigabytes/second (GB/s) between the individual processors, the processors and the memory, as well as between the processors and the I/O hub.	■ Performance that can be tuned with constant energy consumption and heat dissapation. ■ The optional use of two processor generations offers a choice for both the price-conscious user as well as the demanding high-performance user.
Top virtualization support	
■ Two integrated dual-channel Intel® 82576 Gb Ethernet controllers are standard. The integrated Intel® virtualization technology for connectivity contains - in addition to I/O acceleration technology and the Virtual Machine Device Queues - Single Root IO virtualization SR-IOV as well. ■ Two PCI Express 2.0 Mezzanine slots with a combination of quad-channel 1 Gb Ethernet, dual-channel 10Gb Ethernet, dual-channel 8 Gb Fibre Channel, dual-channel 10 Gb CNA (FCoE), and dual-channel 40 Gb Infiniband (QDR) offer excellent connection features via a high-performance midplane. The high server blade I/O capacity allows optimal use of various I/O protocols, ensuring smooth operations for demanding applications.	■ Best-in-class I/O connectivity. ■ High flexibility regarding the type of I/O connection. ■ As a result of bypassing the internal hypervisor virtual switch SR-IOV enables virtual machines to reach a performance level which is almost the same as pure physical machines.
Flexible boot options	
■ Various server boot options, e.g. from the network, from Hard Disk or Solid State Drives or from a USB Flash module (for VMware ESXi) make the server ideal for every application. It is an excellent platform for both virtualized and physical environments.	■ Multiple use, optimized for virtualization and extremely energy efficient regarding usable local boot media.
Worry-free administration	
■ Management via the integrated Remote Management Controller (iRMC S2) enables access to each server and extensive control, even at remote locations. The integrated Pre-failure Detection and Analysis function provides reliable operations in all circumstances.	■ Easy and reliable management and control.

Figure 13 Fujitsu Primergy BX922

The eight host were distributed four by four(Figure 12) to the existing datacenters in order to improve the redundancy of the system and also to comply with Nordea internal policies.

The ESX, which is present on every blade provides the virtualization layer needed over the physical infrastructure, enabling the provisioning of the resources like CPU, memory network resources . This provisioning ensures that the available resources can be used by multiple virtual machines in parallel.

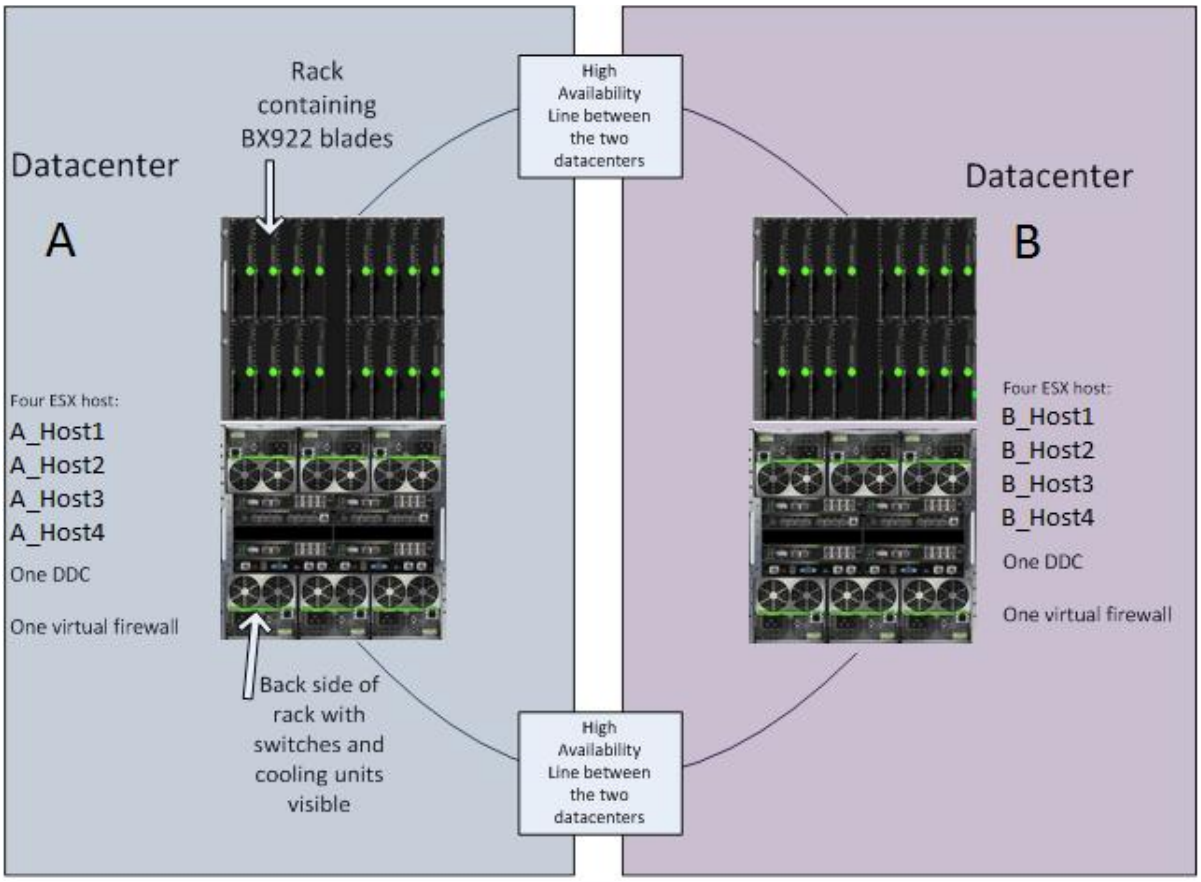

Figure 14 Hardware setup and interconnection of data center's

As presented in figure 14 the design consists of two identical physical configurations in two separate locations. These two locations are connected by two distinct lines. A broader picture can be seen in figure 10. Each rack has two management switches which connect the management blades to the management network, two SAN switches that provide the access to the storage by optical connections and two access switches , through which all the traffic from the virtual machines to zone 1 is directed. Each Fujitsu blade is connected in both access and SAN switches.

4.4 Networking components

From a networking perspective the Port Group System has been utilized to provide a logical separation between two virtual machine groups. For redundancy purposes it has been decided that all even numbered VM will be connected to an even numbered virtual local area network and all odd ones to another virtual network. This procedure provides more flexibility and a reduced risk of total system failure, because the virtual machines not only have a physical isolation, depending on which database they are located but also a logical separation based on which virtual LAN they reside in. As previously mentioned the virtual machines share the

resources of individual servers that are running through VMware ESX server. For increased efficiency in the use of the available resources DRS and VMware vMotion (Figure 15) is used.

The main function of vMotion is the migration of virtual machines from one physical server to another if one high load situations arise. This not only increases the efficiency of the system as a whole but also ensures a higher probability that even in extreme cases the predetermined VM requirements are met. Without this procedure in the case of over provisioning the performance of virtual machines would be affected.

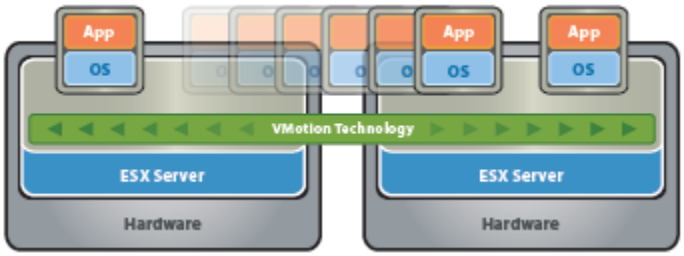

Figure 15. VMware Vmotion (VMware 2012)

The storage system is based on IBM SAN Volume Control Manager (VCM). This storage solution was developed for the support of virtualization architectures. Its main advantage is the enabling of thin provisioning which ensures a better utilization of the available storage resources. What IBM VCM does is to create a virtualization layer over the physical storage units using the SAN Volume Controller hardware unit. By using a volume controller the storage section becomes highly modular (new capacity can be easily added without major modifications to the existing architecture) and also the redundancy of the system is increased thanks to symmetric disk mirroring.

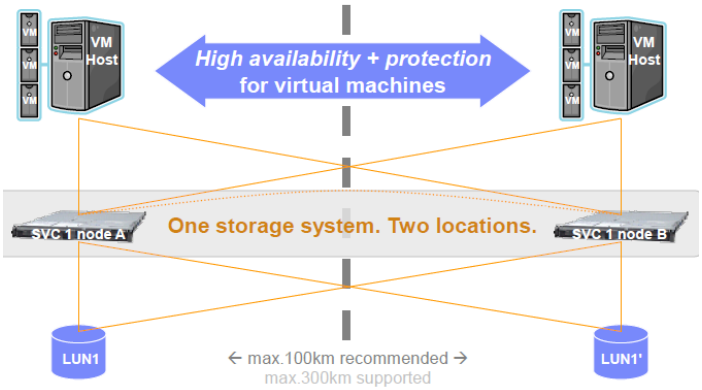

Figure 16 SVC split cluster Symmetric Disk Mirroring (IBM 2012)

As we can see in figure 16 the storage system even if placed in two separate locations is regarded as one unit. As previously presented (Figure 14) the design of the VDI solution is based on the interconnection of two separate physical locations so the use of symmetric disk

mirroring was an obvious solution. This configuration ensures a high availability high redundancy storage system that thanks to a central management console can be easily manipulated and efficiently used for virtual machines.

4.5 Storage technologies

One of the main enablers of virtual machine technology is the use of thin provisioning. Thin provisioning eliminates the problems found in classical 'fully allocated' solutions where disk capacity is consumed even when not in use, thus making storage a scarce resource. The basic principle of this method is to allow over committing of the existing physical resources, meaning that in the virtual environment more storage space can be provisioned than it actually exists in the datacenter. This is possible because thin provisioning operates on the virtual machine disk (VMDK) level. The storage capacity can be assigned to virtual machines in two ways: 'thin' or 'thick'. The thick disk can be considered as a standard storage disk, that no matter what circumstances will always take away the set amount from the existing physical resources (In figure 17 it is 20GB)

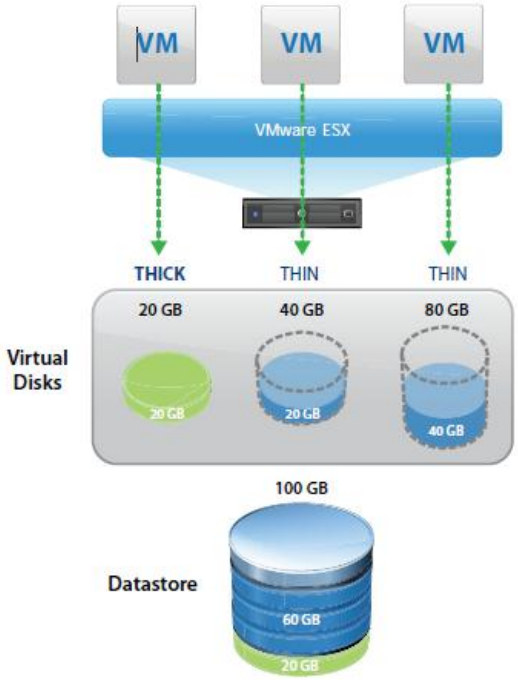

Figure 17. Thin provisioning example (VMware 2011)

If the disk is assigned as thin the blocks that represent the data in the VMDK will not be backed by actual physical storage until a writing process is completed. This means that the no matter what is the capacity of the extra virtual disk in a virtual machine the amount of physical storage used will be equal to the amount of data stored on the assigned drive. As an example if on the extra 80GB virtual hard drive assigned to the virtual machine only 10 GB are used the extra

70GB are free to be used by a different virtual machine. This procedure can increase the visible storage capacity to a great extent mostly because in many cases the assigned virtual disks are underutilized. Of course there is a limitation, and during this project it was decided that over 100% provisioning will not be allowed. By using two 10TB storage units and thin provisioning the storage system assigned to the VDI can have up to 40TB of visible storage capacity, meaning that a 40GB thin layered disk is assigned to every virtual machine the maximum capacity could be up to 1000 VM just by using this configuration, which leaves plenty of room for development and growth.

Chapter 5

Accessing the VDI infrastructure

5.1 User logon process and communication flow for VDI access

For security reasons the overview of the communication and access process can only be found in the classified section of this paper (Appendix B). In the following section a detailed description of the communication flow will be presented in accordance with the numbering on the figure presented in Annex 2.

1. The user points the browser at https://<NORDEA_VDI_URL>

There are two main addresses that can be used to access the VDI logon interface. The nordea.external.VDI -1 is the general URL for outside access and the nordea.internal.VDI-1 is for internal access through which the VDI platform can be accessed from the inside network. There are also backup addresses that were used as residual access bearers (nordea.external.VDI-2 and nordea.internal.VDI-2), this was needed due to domain differences. The original project was designed for users in a different domain, but once the virtual desktop platform was integrated into the new general domain the users that still used the old domain could continue accessing their VDI through the additional secondary URL-s.

2. Netscaler prompts endpoint scan

After the Netscaler receives the connection request the endpoint scan process is initiated on the PC on which the user wants to access the VDI. The Endpoint scan checks for a valid Windows license and for a valid, up to date antivirus on the connection requesting machine. Only if the endpoint scan is successful the following steps are accessible. The user is prompted to accept the scan; in case of refusal the connection will be halted.

3. NetScaler 1 queries for user name and password.

This is the first step of the authentication. The netscaler requests the internal user name, comprised of original registration location information about the user and a series of numbers.

4. NetScaler 1 sends user name and password to Entrust.

The previously mentioned user name and password is sent to the verification system and further analyzed as presented in the Entrust Authentication chapter 5.3.

5. NetScaler 1 queries user for second factor code.

After the first level authentication is completed the Entrust verification system requests the second level authentication credentials, which can be provided by a Gridcard, SMS, e-token and will be further discussed.

6. NetScaler 1 sends second factor to Entrust, which verifies access.

This is the last step of the authentication process. After the information provided in the second level authentication is verified the role of the Entrust system is completed and the user information is forwarded to the inner system, into the zone 2.

7. NetScaler 1 forwards to Web Interface on Netscaler 2.

8. NetScaler 2 verifies credentials by contacting NetScaler 1.

This is an automatic step, and it is caused by the fact that the internal virtual netscaler VPX has the same configuration and functions as the physical unit. This means that it also has to do a verification, similar to the one done by the Entrust system, but in this case the only process is a verification request to the first Netscaler that confirms that the user credentials have been verified, and are valid.

9. Web Interface on Netscaler 2 passes user credentials to the Citrix Desktop Delivery Controller (DDC).

This step is essential in making sure that the user only gets access to what has been assigned to him by the active directory system. The domain controller DDC has the role of coordinating the process of assigning the appropriate VDI to the user.

10. DDC verifies user authorization by performing a Microsoft Active Directory query with the end user's credentials.

11. DDC queries the site database for the end user's assigned desktop groups, by using port 1433. Using the desktop group obtained from the database, controller queries the hypervisor about the status of desktops within that group.

12. DDC identifies to the Web Interface running at NetScaler 2, the desktop it assigned for this particular session.

13. Web Interface sends an ICA file to the Citrix Receiver through NetScaler 1. The ICA file points to the virtual desktop identified by the hypervisor.

From the user perspective this is the first time the inner system is visible. The user can see the VDI icon in the Netscaler web interface.

14. Citrix Receiver establishes an ICA/HDX connection to the specific virtual desktop that was allocated by the DDC for this session through NetScaler 1 which sends the request to NetScaler 2 (Next Hop).

15. NetScaler 2 proxies the ICA/HDX request to the VDI.

16. The VDI contacts the DDC's Virtual Desktop Agent for verification of a valid license file.

This is the step through which the VDI license is verified. The licenses for VDI are acquired in bulk(for 500 or 1000 VDI) and before use the validity of the VDI license is verified by a dedicated license server. Licenses can be 'per device' or 'per user'.

17. DDC queries Citrix license server to verify that the end user has a valid license.

18. DDC passes session policies supplied by the active directory (AD) to the Virtual Desktop Agent (VMA), which then applies those policies to the virtual desktop.

This is one of the most important steps in ensuring the security of our VDI system. The group policies through which the VDI is controlled, have been previously defined, and are usually used to disable certain features in the VDI. In our case the policies disable voice and video communication, prohibit admin access for the standard VDI user and many more. The are several levels of policies that ensure that all users that have access to the inner layer stay within the clearance margins they have been provided by their leading managers.

19. Citrix Receiver displays the virtual desktop to the end user.

This is the last step in the process; the user is now capable of seeing the VDI window provided by the citrix receiver that is installed locally. After this process the VDI can be handled as any standard physical desktop.

5.2 Components

5.2.1 Netscaler

The Netscaler is a versatile hardware device that is mainly used as a transport layer load balancer and a security component. Its basic function is to make the decision where to route traffic fast and efficiently, to accomplish this it uses different techniques than network routers, to ensure a much higher speed of routing. Beside its Level 4 and Level 7 load balancing functions it also provides content switching, data compression, content catching, SSL acceleration, network optimization and security features.

The main reason a Netscaler hardware device is needed is because real enterprise applications are complex ,and conventional solutions that could provide all the features that one hardware Netscaler can provide like SSL(secure socket layer) acceleration, compression protection are too complicated and quite slow. By using smart routing techniques the netscaler is capable of achieving speeds up to 5-10 times faster than conventional configurations. One of main characteristics through which the nescaler obtains the increase in speed is the request switching technique that incorporates the use of persistent connections, multiplexing over persistent connections. Because the HTTP traffic is usually many short lived connections and servers perform much better with persistent connections the netscaler uses multiplexing over a few persistent connections, basically the segmented connection requests of client are gathered into one continuous connection to the server.

Another speed increasing technique is the compression used by the netscaler. Basically it uses Gzip to compress data, and by maximizing the packet payloads it increases application performance and speed. The versatility of the netscaler is ensured by the fact that is able to do multi-protocol compression, and this way any type of data existing in the cloud gateway can be processed, compressed and sent along with a high efficiency.

Because of the high level of control on all protocols the security level can be improved dramatically. The Netscaler includes a built in hardware component used for encrypting data. This way it can be considered as a highly efficient firewall device adding an extra security layer to the system. This procedure also speeds up the communication between the clients and the servers because the servers do not have to spent time on encrypting and unencrypting data all this is done inside the netscaler on a hardware level .

As presented in the network drawing the netscalers used in this project are the second layer of protection situated just behind the main external firewall. The outside traffic after passing the initial firewall is directed to the physical boxes. The netscaler is visible from outside as a web page (nordea.internal.VDI). Basically the netscaler is represented by one assigned IP address on the external side and one IP address on the internal side. The routing between the two IP addresses is done inside the netscaler, and in this process all the others steps like encription/decryption ,data compression/decompression and so on are done. This way the inside network is totally separated from the outside network, this ensures high level of security.

In the VDI architecture there is an extra virtual netscaler. The virtual ‘box’ is necessary for spanning the additional security level. The virtual netscalers are basically virtual servers that are assigned the same task as the actual physical boxes, and also add an extra protection layer but all the encryption is done by software and not hardware.

An add-on security feature of the Netscaler is the Extentrix software that has the task of checking the security configuration of computers through which a connection request is made. It checks existing antiviruses, their update status and also the firewall and operating system settings If any of the requirements are not met the connection requester receives a denial of connection message, and it is asked to update/install an antivirus and to configure the correct firewall settings.The Extentrix software requires that the operating system from which a connection is requested has the Extentrix client installed, in the case this is missing it will prompt for an install.

As previously stated one of the main functions of the netscaler is layer 4 load balancing. The load balancing provided by the netscaler also includes health monitoring, session persistence and network integration. The health monitoring is not only responsible for the basic ping, TCP checks but also performs scriptable health checks, dynamically checks the servers response times. The load balancer implemented in the netscaler uses the health checks to ensure that only optimally functioning servers are included in the load balancing process.

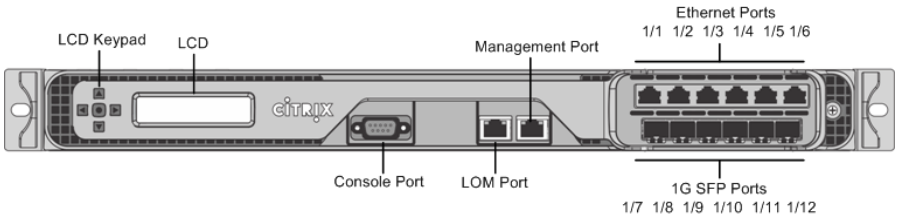

Figure 18. Netscaler 8200

The NetScaler choice for this project is the NetScaler 8200, which is the lower end product from citrix for the nCore series. The main advantage from the previous versions as the series name states is the multiple core attribute, through which the 8200 version has a better performance when dealing with multiple tasks. The 8200 Netscaler as presented in figure 18 incorporates an LCD screen of small dimensions and an LCD keypad that are mainly used for the initial manual configuration if the Ip , subnet mask and gateway addresses. This feature enables a fast and redundant configuration not only in the first configuration stage but also in case of failure. The Netscaler also includes a serial (console) port which is a classic redundant feature of any large scale hardware. In this system this is connected to a serial management console than can be used as a backend connection in case the classical network connection fails. The management port is connected to a central management unit . Out of the 12 ports available 3 Ethernet ports are used 1/1-management 1/3 inside network and 1/5 outside network . The other Ethernet ports can be enabled if further expansion is needed. The 1G optical ports are not used because the current configuration of the networking does not support that feature.

After the initialization of the Netscaler (after it receives an IP address) the configuration console can be accessed either by PUTTY or through a management web interface. This management interface is accessible only from the colored zone in which the netscaler is located for security reasons. The configuration of this hardware can be done both from a command line and the GUI on the management web interface.

The most important configurations include enabling the full duplex speeds on the used ports, configuring the VLAN, adding the routing table, specifying the user authentication requirements (AAA groups),specifying encryption list, configuring which servers the netscaler will be in contact with (by specifying the servers IP addresses) etc.

For this VDI solution for redundancy purposes two 8200 Netscalers are used. As it can be seen in the main architecture diagram there are two sites , two datacenters located in different places. The two Netsclares also include a high availability feature that ensures that even if one of the units fails the operations are not affected. After enabling the high availability features on both 8200 hardware, the one in site 1 is set as the primary unit and the on in site 2 as the secondary.

In the proof of concept and previous VDI solutions the Netscaler of choice was the 7000 which beside performance lacks was also missing a user web interface. This was ensured by the addition of the virtual servers on which the web address was placed. The 8200 edition supports a user web interface as well, this way also reducing the additional resources needed for implementing the VDI solution.

The encoding of information is done by using a cipher. A cipher is an encryption method that can have different strengths; in this case strong 256-bit ciphers are used. The used cipher group is selected from a list of available cipher groups and are changed after a certain time of being used to avoid security threats that could arise from using the same ciphers for an extended period of time. As mentioned SSL is also used. An SSL certificate consists of a private key, a public key, optional intermediate certificates and a root signing certificate.The SSL certificate can be used with several different ciphers, but has to be changed (renewed) periodically every 12 months.

5.2.2 Netscaler VPX

The Netscaler VPX is a virtual netscaler that is located in the zone 2 and has the exact same role as the hardware unit previously presented. It includes all Netscaler functions like load balancing, traffic management, application acceleration, application security (Firewalls in Access Gateway and Citrix application). There are two main reasons that the virtual netscaler is needed. First it offers an added layer of security for our system because by having firewall and cyphering functions it acts as a second barrier for possible intrusions. As well as the physical counterpart it has an outside and an inside IP address, the traffic between the two being encoded inside the virtual netscaler. The second reason for this virtual element is its architecture function. Due to the different security zones it is not advisable to have a direct connection between zones that are not logically adjacent. This means that if it is not possible to have a direct jump from zone 3 which is regarded as a less secure zone to zone 1 which is the inner network of the security system. Because of this there needs to be an element in each zone, and thus the virtual netscaler plays the role of the unit in the zone 2, which can be considered as a jump area from the inside to the outside of the secure region.

The VPX is often used as a cheaper alternative to the Netscaler Hardware, because it can perform the same tasks at a fraction of the cost. Its main advantage is that it does not require a specialized hardware unit, that is not only expensive but also requires shelf space in the data store; which for a lot of companies is quite costly and could be used for more important elements. The cost efficiency of the VPX is quite good due to the fact that it requires only a basic server configurations on which it can be implemented and the licensing fees from Citrix. It

has to be mentioned that in order to save costs not all elements available in the VPX and hardware Netscaler have been purchased.

The virtual Netscaler even though can perform every function of the hardware unit has lower performance margins. The VPX cannot be configured on any traditional server that meets the minimum requirements, it can only run on a hypervisor, and then when installed it needs the exact configuration as set in the files provided by Citrix. In this case the server used has 4GB RAM and 50 GB hard drive. This means that the additional performance that is provided by specialized hardware like the dedicated SSL acceleration in the physical Netscaler cannot be matched with the VPX. There is a fixed limit of 500 new SSL transactions per second with the VPX and this cannot be increased by the addition of more performing hardware to the hosting server. The performance of the virtual Netsclaler is also limited by licensing issues. The platinum license offers the best performance but is quite costly and it would still not offer the performance of the physical Netscaler dedicated hardware unit.

Because the VPX has the same configuration and the same functions as the hardware version they communicate well, using port 443.In this project as presented in the logon process and communication flow chapter the VPX is mainly used as a transitional unit, that does not perform checks on its own but ask for information from Netscaler 1 and send users credentials forward to the citrix desktop delivery controller. It is also the platform on which the secondary web interface is running. The VPX will eventually proxy the ICA/HDX file to the request to the VDI.

The VPX has been specially designed to offer increased security and performance if combined with the physical netscaler. Even if it does not possess the hardware elements for increased performance it has a valuable contribution to the security level of the system and also helps with the integration and connection of the performance wise superior hardware unit into the security architecture.

5.2.3 Citrix controller

The DDC is the component that ensures the delivery of the virtual desktops to the users. As presented in the previous chapter it receives the user credentials from the virtual Netscaler, after which it is responsible to provide the appropriate ICA file based on the setting and groups policies which the user is subjected to. It has quite a wide array of tasks to perform. After receiving information from the VPX it does two main verifications. It gets the user rights from the active directory to ensure that it does not provide more authorization than the user is allowed to have. After that it checks the user has the valid license for receiving connection to the VDI. Based on the gathered information it enables the visibility of the VDI icon on the second netscalers web interface. This ICA file contains the access and the correct information

on the VDI that is then used by the user's citrix receiver to connect to the appropriate virtual desktop.

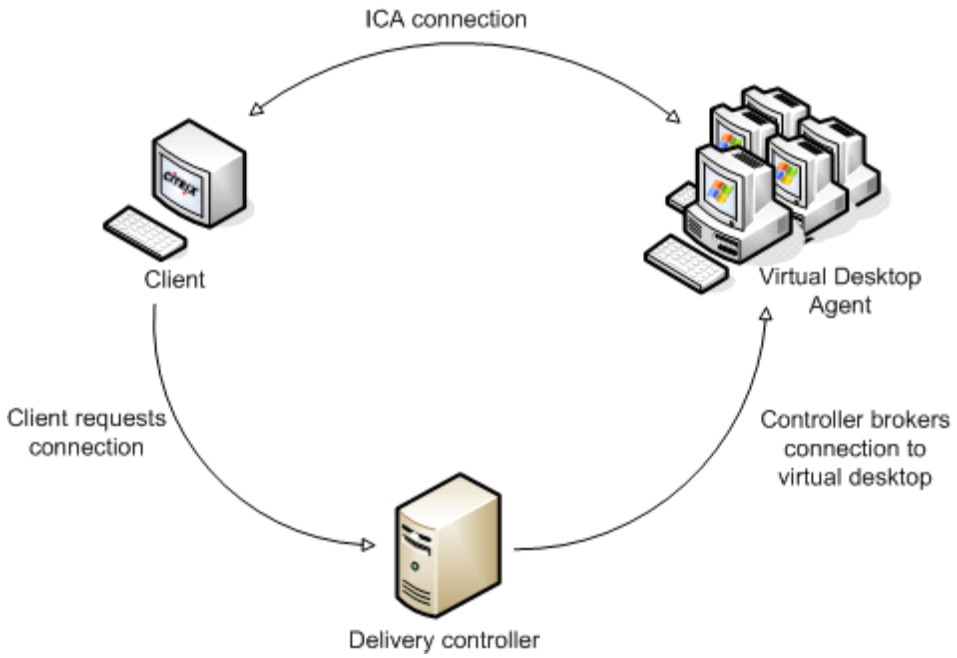

Figure 19. Citrix delivery controller

The process that si performed by the desktop delivery controller is quite easily understandable in figure 19. The virtual desktop agent is responsible with managing the virtual desktop sessions. The main function of the DDC is to broker the ICA connection for the user based on the user credentials provided, or in easier terms to tell the virtual desktop agent what is the exact ICA/HDX connection is has to provide to the client.

The Virtual Desktop Agent (VMA) installed on each VM is responsible for using the capabilities if HDX to optimize graphics ,sound ,printing etc., and send this with the ICA protocol to the Citrix Receiver installed on the end user's computer. (Citrix website) Every desktop logic executive like mouse movement, screen updates, keystrokes, audio etc. are transmitted though the ICA protocol. ICA/HDX is used to provide a better performance by incorporating Image and Browser acceleration, multi-monitor support, printing using the universal printer driver and many more.

The main reason in using the ICA/HDX based architecture is to optimize connection speeds and to reduce latencies. This is achieved by sending only screen shots, keystrokes and mouse clicks through the network leaving all the heavy computing to be done on the hosting servers and not the clients PC, in this case the VMware hosts that have enough resources and hardware to be able to provide a smooth working environment even when using demanding software. This method not only speeds up the connections but ensures a high level of security because the

used applications are not copied to the PC from which the user connects and thus the applications can be managed centrally and safely. Providing a central management platform for the applications also reduces costs and increases management efficiency.

5.3 Entrust Authentication

Entrust authentication is one of the main security features that ensures a safe connection from the VM to the protected zone 1. There are three main steps on the way to providing access to the inner Netscaler . In the first instance when accessing the VDI URL, the users are prompted to provide internal login credentials. After the Netscaler receives this data it forwards it to the Entrust system in the Server Network. The coordinating Entrust server (Entrust GI) forwards the login information to the Entrust repository server. The repository holds a copy of the internal Active Directory, and thus a copy of the user information inside the AD. If the user credentials are considered valid the Entrust sends a second factor authentication request to the Netscaler 1, which is forwarded to the user. The last step of the entrust authentication is the verification of the second factor credential submitted by the user. If the verification is successful the web interface is transferred to Netscaler 2 and there is no more authentication requested from the user. It is important to mention that the failure of providing any of the previously mentioned credentials will result in a denial of connection to the VM.

5.3.1 Two factor authentication

Entrust authentication is based on a two-factor authentication solution that not only asks for a user name and a password but also and additional security element. This second authentication usually requires the user to physically possess something (i.e USB Token, Grid Card E-Token) that can provide an additional unique key . This extra security element is requested during the login process.

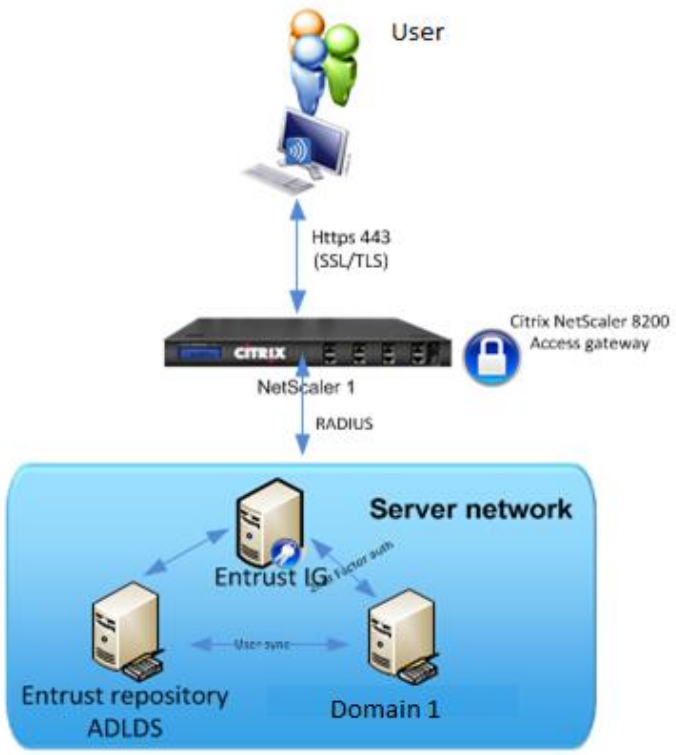

Figure 20. Entrust Authentication

.

During the initial phases of this project the Grid Card option was considered and implemented. The Grid Card (Figure 20) is basically a table that contains a cypher. During authentication the user is provided with 3 random letter number groups (i.e A3 D4 I7) and he has to introduce the corresponding element from the cypher. The main advantage of the Grid Card is its simplicity, because when assigning a remote user a specific grid card a secure mail or Fax can be used to send it. This way there are no difficulties in replacing the security element or providing a new one. The disadvantage of Grid Cards is that it can be easily taken out of certain security zones that workers in remote locations use and thus is more susceptible to malicious interferences. Also anyone can easily take a picture of someone's Grid Card and thus the security element is compromised.

	A	B	C	D	E	F	G	H	I	J	
1	E	Q	X	3	T	5	N	4	M	Q	1
2	E	3	K	6	J	M	9	F	8	6	2
3	C	1	6	M	3	J	H	M	P	Y	3
4	T	W	W	1	4	V	6	0	7	2	4
5	8	6	7	W	6	J	5	M	P	X	5

Serial #

Figure 21. Example of Grid Card

A second solution with the Entrust system is the physical USB token. This solution consists of a USB that has to be inserted in the remote PC. Due to the fact that having USB ports open poses a grave security threat this was not considered as a viable method.

A more widely used approach is the SMS verification as the second phase of the authentication. Basically the users phone number is registered in the Entrust system that sends the login code directly to the phone of the user in an SMS. This solution is quite time sensitive (the user needs to have the login key immediately when it requires it) and is dependent on the phone network of the users location. This made it unusable in the Pune project because the local mobile network has huge delays in processing SMS and sometimes it can take even over an hour for an SMS to be received.

The solution that was considered the most suiting for this project is the use of E-Tokens. An E Token is a small electric device that can provide a seemingly random key that the user uses as the second authentication key. Every VM user is assigned his or her own token. Each token can be identified by a unique serial number that is registered in the Entrust platform before assigning it to anyone. The solution for this project also implies the use of a security area that is a closed location where the use of the tokens is allowed. This means that taking out the token from this secure area is prohibited. Due to the fact that the codes provided by the tokens are time sensitive, they expire meaning a code given at a certain time cannot be used later. This improves security by making sure that the virtual workplace is not utilized outside working hours in environments that could pose a security risk for Nordea.

Figure 22 EToken

The tokens are assigned using a special Entrust platform that is placed on the Entrust Identity Guard (Entrust IG) server as showed in figure 22. With any authentication method a main requirement is to have an easy method of giving and taking away access so that the security of the network is ensured, and also in case the token is lost or malfunctions the user can be provided with a new one fast and securely. In this project this problem was solved by having a large initial pool of Etokens, that are registered in the system but are not assigned to anyone. They will be sent to the 'Red Room' location to Pune, and this way in case any of the

tokens fail or is compromised a new one can be assigned fast and the user can resume the work on the WDW.

5.3.2 Communication protocols used for the Entrust authentication

The communication protocols used during this procedure can be divided into two main areas. Both protocols are specially selected to provide enhanced security, and their use is specified in internal security standards. Also the equipment (the Netscalers and Entrust servers) was designed specifically to utilize these transfer protocols which are rated as having an advanced security.The first one is the communication between the Netscaler and the user. This is done through secure HTTP (HTTPS). This is not a new protocol in itself it is just a layering of normal HTTP on top of SSL, this way the security benefits of an encryption used in SSL is added to the Hypertext transfer Protocol. The main reasoning behind using HTTPS is to ensure that middle-man eaves dropping is not possible , this way every time a connection is established both sides know exactly to what they are connected to .The extra security on the connection is achieved by using an SSL certificate. All the certificates used by Nordea for HTTPS is provided by VeriSign through a subsidiary company. The basic idea of a certificate is using a special key, which in essence is a set of prime numbers that are unique to the certificate issued to the client. The web browsers have a preinstalled knowledge on how to handle https websites, they know what key ‘prime numbers’ to use to decrypt the information sent or received through the secure connection. This way the only person who can access the information is somebody who has the security certificate.The second special protocol is used between the Netscaler and the server network. In this stage Remote Authentication Dial In User Service is used. The main purpose of RADIUS is to provide centralized Authentication, Authorization and Accounting (AAA) management for the computers when connecting to the network service. if some specialists consider RADIUS as outdated and possibly a security issue it is embedded into the Netscaler and it is considered as the standard for AAA management.

Chapter 6

Other VDI solutions

6.1 VDI-in-a-box from Citrix

It is a solution developed for the low-cost end of the virtual desktop market. Due to its design VDI in a box eliminates up to 60% of the infrastructure needed for a working VDI solution, including management servers and expensive SAN (Storage area network). This is achieved by using a shared grid of off-shelf standard servers.

Architecture

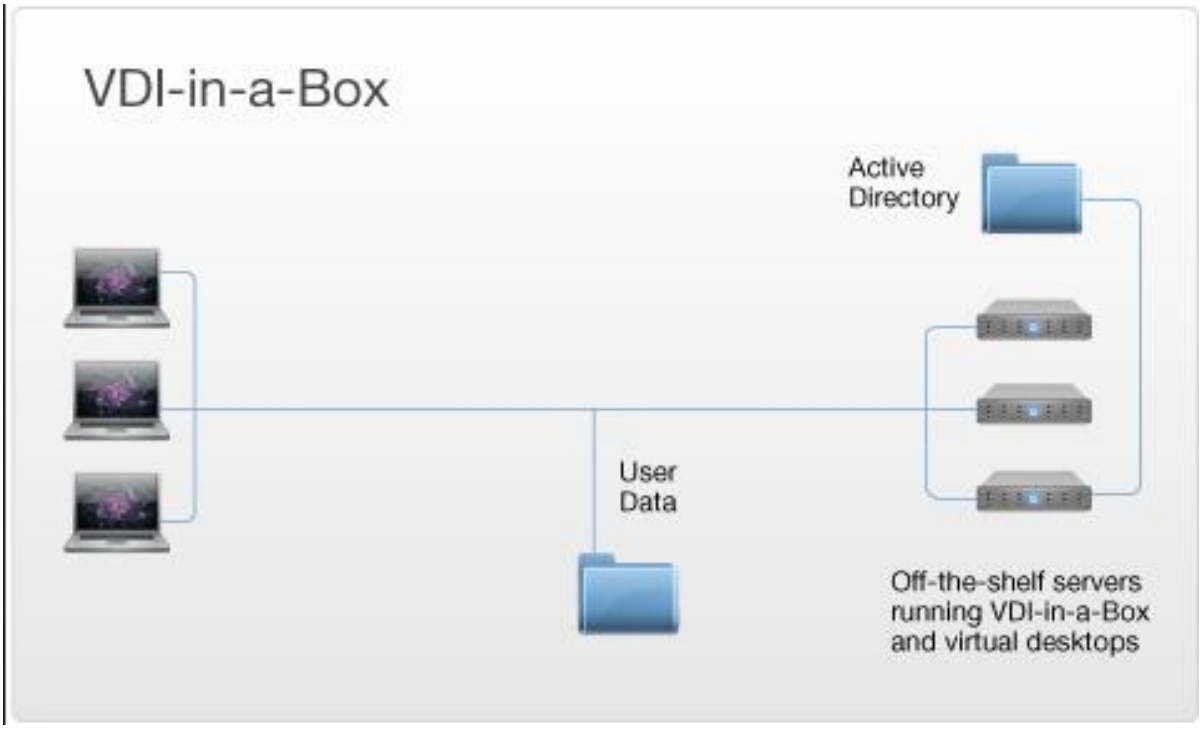

Figure 23 VDI in a box Architecture (citrix 2012)

As shown above the architecture of VDI in a box is a simple configuration that consists of one or more standard physical servers, with attached storage each running a hypervisor and the vdiManager described below. One of the biggest advantages of this solution that contributes greatly to its low cost is its modularity. The vdiManager can be configured to run on one server, or to be part of a cohesive grid of servers. This offers great scalability, if more resources are needed a simple addition of a new server to the already existing grid will suffice.

The vdiManager performs the following functions:

1. Creates virtual machines from predefined templates.

   A template consists of two main parts. One is the image that includes the operating system, a set of applications and a VDI-in-a-box agent that is used to relay information to the VDIManager about users connections and desktop health issues. Provisioning is also implemented in this solution; multiple templates can use the same image of the operating system. The template also contains policies that govern the resources, how much ram is to be allocated, how many desktops to create, and the use of peripheral components (USB).

2. Load balancing is usually achieved through expensive hardware units, but in this solution the vdiManager performs this task. It creates new virtual desktops on one of the servers on the grid based on resource availability, to ensure that the load is evenly spread across the grid. Every time a user logs in, the virtual desktop is created on one of the least loaded server, this way ensuring good performance.
3. By having a good control over the server grid the vdiManager can not only provide load balancing but can also assure high availability. The vdiManager instances communicate between each other, and share key operational and configuration information. The VDI templates and images based on which the virtual desktops are created are stored on each and every server, this way if one of the servers fails, the other units in the grid are capable of restoring, recreating the lost desktops. After the fail is repaired the provisioning and control over the specific virtual machines will be reassigned to the original server.(Spruijt,2013)
4. The management interface is web based.

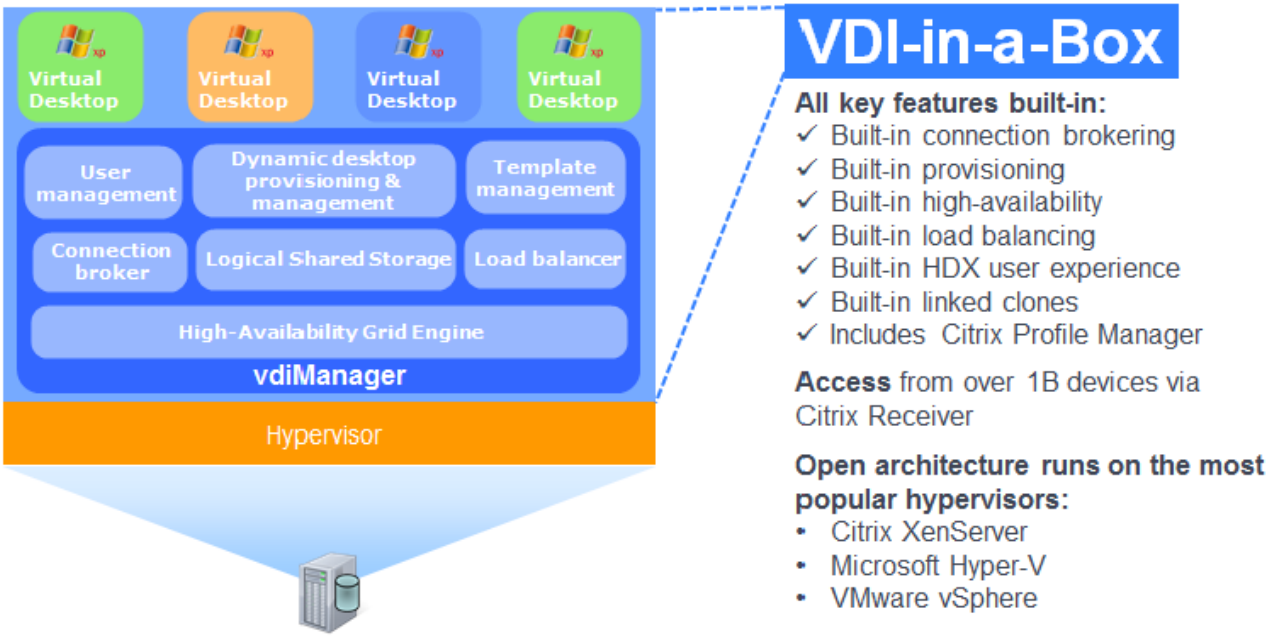

Figure 24 The main features of VDI-in-a-box (Citrix 2012)

The VDI-in a box presents a series of advantages, most importantly reduced costs related to low hardware needs, good redundancy, easier management(less components easier to manage), high level functionalities like load balancing, high availability, provisioning, web based management center. The cost is also lower in comparison with enterprise VDI solutions due to reduced licensing costs. A license is only considered in use when a user's device has established connection to a virtual desktop. This allows multiple users or multiple devices to share licenses.(Spruijt,2013)

6.2 Microsoft RDVH-Virtual Desktop Infrastructure

The original technology that provided VDI-like features from Microsoft was called the Terminal Services, which allowed users to connect to a full desktop or individual programs on the users own device. The term 'Remote Desktop Virtualization Host' (RDVH) was introduced by the Windows Server 2008 R2 edition, and allowed users to have a dedicated virtual desktop running a windows client operating system. With the addition of RemoteFX the quality and capability of the remote session was improved over the standard Remote Desktop Protocol.

The RDVH virtual infrastructure was designed with a special focus on the BYOD concept, with improved features for touch-screen devices, thus reflecting the mobile application pedigree that can be seen in the latest Windows products, namely Windows 8 and Windows server 2012.

The VDI solution provided by Microsoft is based on the already existing RDP and an enhanced RemoteFX technology and consists of the following Windows Server 2012 roles as presented in (VDI whitepaper, 2013)

1. Remote Desktop Gateway (RDG)
   This is an optional role to provide secure access to the Microsoft Virtual Desktop Infrastructure from internet-based clients.

2. Remote Desktop Web Access (RDWA)
   This role provides access to the desktops and/or remote applications available for a specific user. After the user browses to the Web Access URL and authenticates, Web Access provides a webpage displaying the shortcuts to the resources available to this user. If the client device is running Windows 7/8 and is on the corporate LAN the shortcuts can be also integrated in the user's Start Menu.

3. Remote Desktop Connection Broker (RDCB)
   The Connection Broker tells Web Access which resources are available to the user. The RDCB role is the broker which connects the client to the correct resource selected by the user in Web Access. The Connection Broker also contains the Remote Desktop Management Service. The Remote Desktop Management Service maintains a database with the static configuration of the deployed RDG, RDWA, RDCB, RDSH and RDVH roles, and dynamic session information of the managed RDSH and RDVH servers.

4. Remote Desktop Session Host (RDSH)
   Formerly known as a Terminal Server, RDSH provides server hosted desktops or remote applications to the client. The RDSH role is not required for a Windows Server 2012 virtual desktop infrastructure, but could be added to provide a hybrid solution.

5. Remote Desktop Virtualization Host(RDVH)
   A Virtualization Host is a Microsoft Hyper-V host with the Virtualization Host agent service installed. RDVH provides virtual desktops or remote applications to the client. The Virtualization Host agent service manages the starting of the virtual machines or remote applications (in a virtual machine) when a user wants to connect.

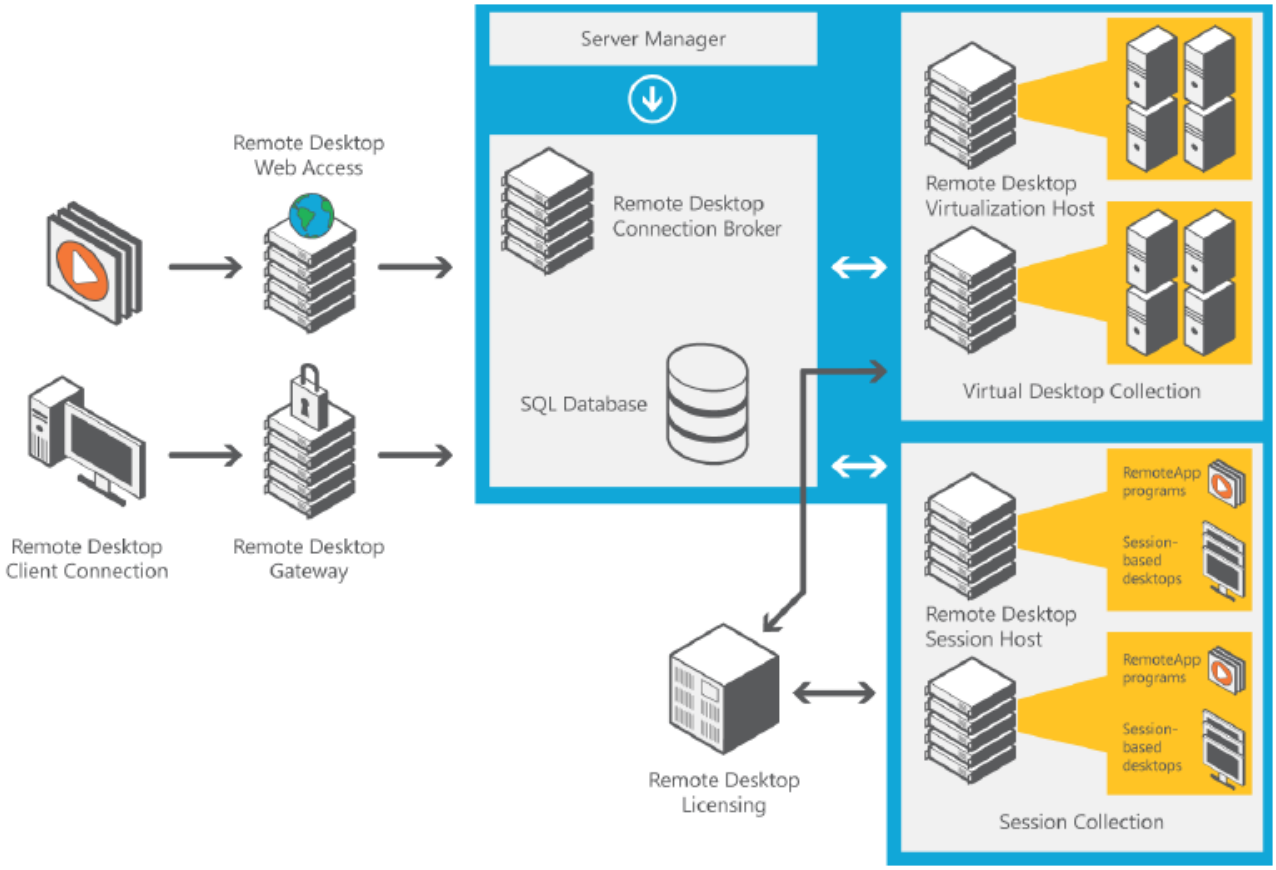

Figure 25 Microsoft VDI architecture(Microsoft 2012)

Being a Microsoft product one would consider that the RDHV VDI solution would be the inherent solution for an enterprise architecture, in which the biggest percentage of elements are using a variation of a Microsoft products, but in this case it is not the optimal solution.

One of the main reasons is that, as it will be presented in the comparison segment it is not the best VDI solution lacking performance parameters (the quality of the remote connection is lower than that of a Citrix solution) mainly because it was developed as an integrated Microsoft solution rather than a dedicated VDI product. Being an integrated solution it strictly requires a Windows operating system; the clients device should run Windows 7 SP1 or Windows 8 with RDP8, which is also a disadvantage when considering the sheer variety of operating systems in the mobile devices sphere.

6.3 VDI feature comparison from the project perspective

The VDI solution chosen for this project was the XenDesktop solution provided by Citrix and base on a VMware infrastructure. As presented in more detail throughout this project,

there are a lot of parameters influencing this decision, ranging from available infrastructure to costs, compatibility with project requirements, but the features provided by this solution are also essential.

Since all the solutions in the desktop visualization domain are quite complex there are a lot of features that are similar between the solutions, features which in essence are the complete description of what a VDI solution is capable of. This feature comparison is presented in more detail in annex 2 in the comparison sheet provided by the (Spruijt ,2013)

As can be noticed in the comparison sheet, the VDI-in-a- box solution is a good choice for simple low costs designs, because it's not only cheaper and simpler than other VDI solution but also is quite similar in the feature comparison segment. Beside some missing monitoring and management tools, and some security elements its feature palette is comparable to the more expensive and complex big brother Citrix XenDesktop solution.

When comparing with the other presented VDI products presented the XenDesktop solution is capable of offering some features that are not available in the other solutions, and are essential to the requirements of this project.

Some key features that make the XenDesktop solution fit the project requirements are the following:

1. The two factor authentication– this is a main requirement in the security descriptions of the project. Due to the fact that the VDI have a direct access to the inner network a heightened access control has to be utilized.
2. Restrict functionality based on time/location/device : this feature is not available in any other VDI solution. Since the main scope of this project is to provide a desktop delivery solution for remotely located users (in this case out of the EU territory) having a tight control on the provided VDI is a must.
3. Image delivery to VDesktop through LAN: this feature is essential for implementing OS streaming
4. Additional instrumentation and monitoring tools : this feature is essential for a large VDI solution, this also ensures the scalability of the solution, by permitting a expansions and the addition of possibly thousands of VDI users while maintaining a high level of monitoring, fast debugging and error correction
5. It uses citrix ICA/HDX protocol, which has been proven usable on connections with up to 500-700ms latency where as Microsoft RDP/Remote FX is only usable with up to 300-500ms latency. This is essential in the case of this project, because the latency encountered by developers situated in India is quite high.

Chapter 7

Improvements

7.1 Introduction

The virtual desktop solution presented so far can be considered a major leap from the traditional desktops, not only in its capability in offering a greater mobility, and permitting the realization of the bring your own device concept, but also in reducing infrastructure costs, reducing management difficulties, providing a more clear overview of large infrastructures thanks to its modular structure and better security and redundancy provided by its data center based architecture.

Even if it can be regarded as a leading edge technology, the VDI solution presented in this paper consisting of persistent virtual desktops is only the first step towards a truly virtualized application delivery platform.

The fact that the individual desktops still take up quite a large chunk of expensive storage, in this case roughly 50GB plus the storage needed for the applications and user profiles means that there is still room for improvement. Also the fact that each VDI is constructed from a OS template that is permanently installed with a set of applications can cause issues in case of template modifications, updates, new application delivery.

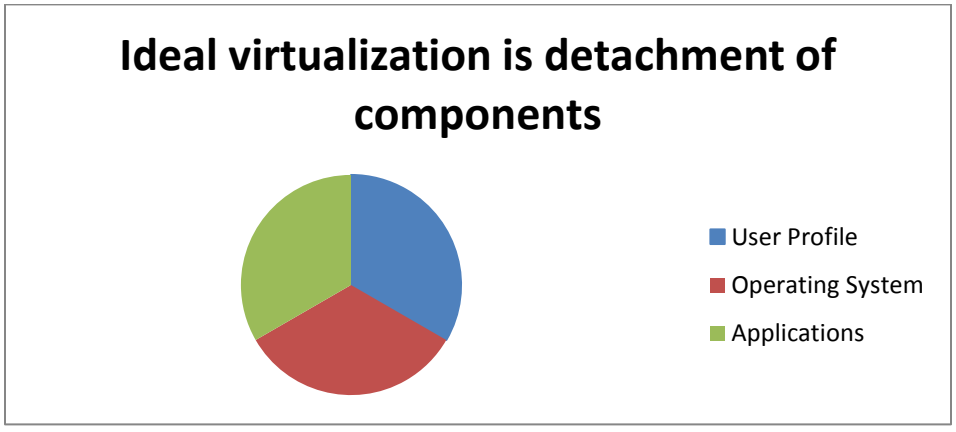

Figure 26 Components of virtualized desktop

To solve these issues and to have a truly virtualized desktop and application delivery platform a complete separation of the following three layers has to be achieved: operating system, user profiles and applications (Figure 26). Traditionally all these are bound together on a physical desktop, which makes mobility quite poor, the application palette rigid and hard to modify and of course updating and hot fixing a tedious adventure. Desktop virtualization solves

the mobility issues, by putting everything in the data center so users can have access to their desktops through a simple internet connection, but application delivery to these virtual desktops is still time consuming and inefficient.

7.2 Workspace virtualization improvement

As we can see in Gartner's Hypercycle of virtualization a complete workspace virtualization is still in the peak of inflated expectations phase, which means that it still far from being implemented in a production environment. Even so if achieved it could make large infrastructure management easier and could lower storage and infrastructure needs making the VDI a truly powerful application delivery solution.

To achieve a complete workspace virtualization the dependencies that exist today between the operating system , applications and user profiles has to be removed. As can be observed in figure 27 the future of virtual desktop will be a concatenation of three separately manageable streams with all of them configured and controlled from inside the datacenter infrastructure.

Figure 27 Complete desktop Virtualization

The first step toward this architecture is the introduction of non-persistent VDI's. Non-persistent VDI's as the name suggests only exist if they are being used. Every time a user tries to log on a new virtual desktop is created from a master image and is deleted immediately after the user logs off. Even if this has already been achieved, as previously presented it does not have the same proprieties as a persistent desktop especially when it comes to user profiles and application personalization. In some cases this might not pose difficulties but in the case developers are the clients for the solution more profile and application management freedom has to be possible.

7.3 User profile management

Traditionally the user profile is created when a user first uses a desktop; this consists of a folder where all the changes made by that specific user are stored, and are loaded during future logins. With a persistent VDI this is done exactly in the same way, meaning that if the user logs into another virtual desktop the profile previously created will not be accessible. This poses issues when implementing non-persistent VM's, because since the VM is deleted if the user logs of the information stored in the user profile folder is also lost.

The way to solve this issue is to have a centralized database that contains all the user profile information and which provides this information to the VDI provision services. This way every time a user logs in, a new virtual machine will be created that has the same user setting as the last virtual desktop used. This tool is also useful in the case of multiple domain architectures, where one user profile can be used for different domains, improving user experience by always providing the same familiar desktop.

There are different products through which this can be achieved, but in general the architecture of the majority of the solution is quite similar. This consists of an agent which is installed together with the operating system, and can be part of the golden master image from which the VM is constructed from. The agent is connected to a management centre and a database. Every time changes are made to a user profile, the agent sends this information to the coordinating server; these changes are recorded and stored in the database. During the initiation of a new user session this profile setting are preloaded to the new non persistent VM. This way even if the virtual desktop is deleted every time a user logs off the profile information is stored and is can be reused.

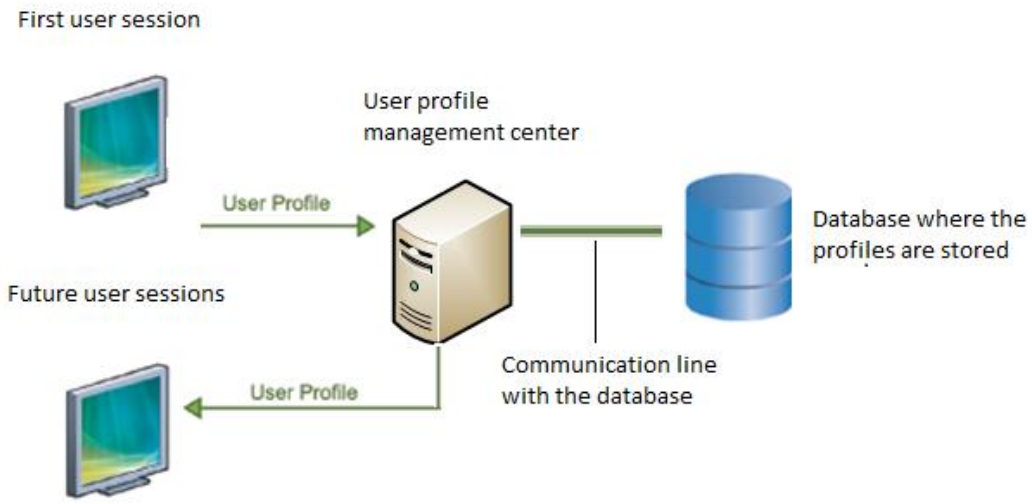

Figure 28 User profile management

To have an up to date profile store it is essential that during every session the modifications made are recorded and the information in the database is synchronized with the user information on the virtual machine.

7.4 Provisioning and OS streaming

This step is fundamental in lowering storage costs and improving the management of virtual machines. In large enterprise architectures, managing every machine as an individual entity is quite difficult and requires a lot of resources. This is also the case in the persistent VDI model

where on each of the virtual machines a separate copy of Windows 7 is installed, not only taking up a lot of storage space, but also making OS upgrades and management more difficult.

This issue can be solved by having a 'master image' of the operating system that is used for the initialization of virtual machines. Using a master image allows for better management since if a change has to be made to the OS, or an upgrade has to be pushed through the modification of the master image will cause all future virtual machines that are created by using this image to be upgraded as well. If some virtual machines are required to be set up differently tailored master images can be created.

The provisioning is accomplished by a provisioning server that creates a master image of a physical disk of the pre-configured machine and virtualizes it into multiple virtual disk images. By using this virtual image the OS can be streamed to the VDI servers.(Citrix provisioning services 2012).

With OS streaming the virtual machine boots and runs from the master image file stored on the network, the actual operating system is streamed to the desktop from a central server. The main advantage of using OS streaming is that compared to a standard instance of an operating system which might consume up to 1 GB of the machines memory resources, with streaming only a fraction of this is used at any given moment, only the files required for the desktop to function are downloaded. For this to be achievable a high level separation between the OS and applications is needed.

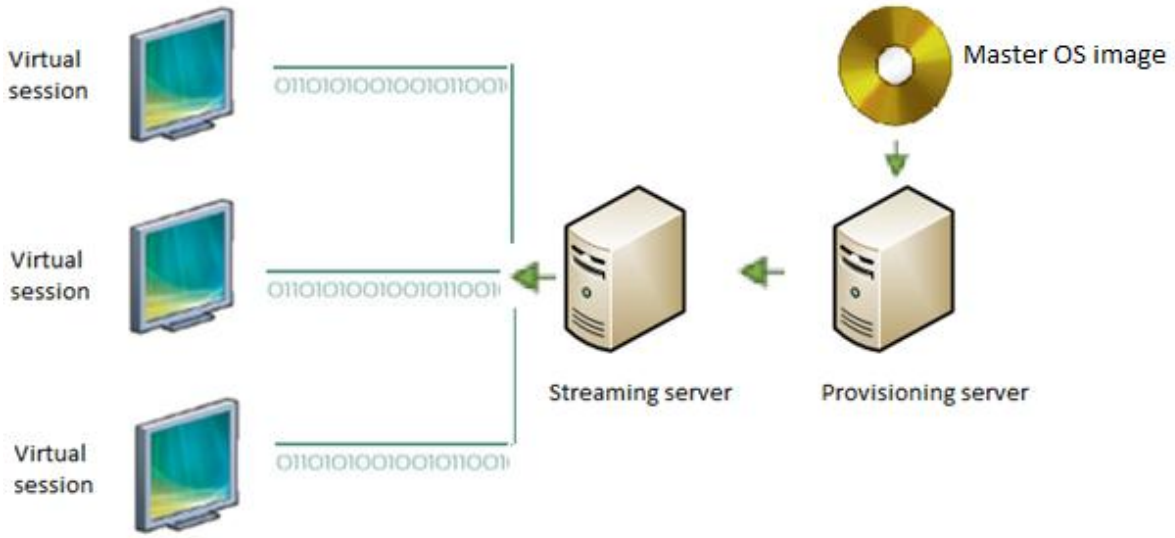

Figure 29 Provisioning and OS streaming

7.5 Application virtualization

The last and one of the hardest steps is application virtualization. This is an essential part, because the core function of any VDI infrastructure is application delivery. Application virtualization can be separated into two main functions: application isolation and application streaming.

Application isolation- is the procedure through which an abstraction layer that encapsulates the application is introduced between the application and the operating system of the client device. Each virtualized application is contained in their own abstract layer.Resulting in the elimination of application conflicts and errors caused by operating system instabilities

Application streaming-it is similar to OS streaming, a virtualized application is delivered to a user's isolated environment from a centralized application repository or application hub when requested. The advantage of streaming instead of having it locally installed is the reduction of storage requirements

By combining these two functions costs derived from certification and application for regulatory compliances, deployment, maintenance and updates can be significantly reduced.The management and support costs are also lowered in the same fashion as the master image contributes to the easing of management in OS streaming. By having a centrally accessible application hub, updates to any application can be made efficiently, it is not necessary to update on each individual machine, and the location of the users is irrelevant to this operation.

Another advantage in application streaming is the selective nature of this procedure, only required applications are streamed to the user's device. There is also the option of catching the virtual application on the users device which can reduce the traffic of the streamed applications, and it is useful for applications that are not used frequently.

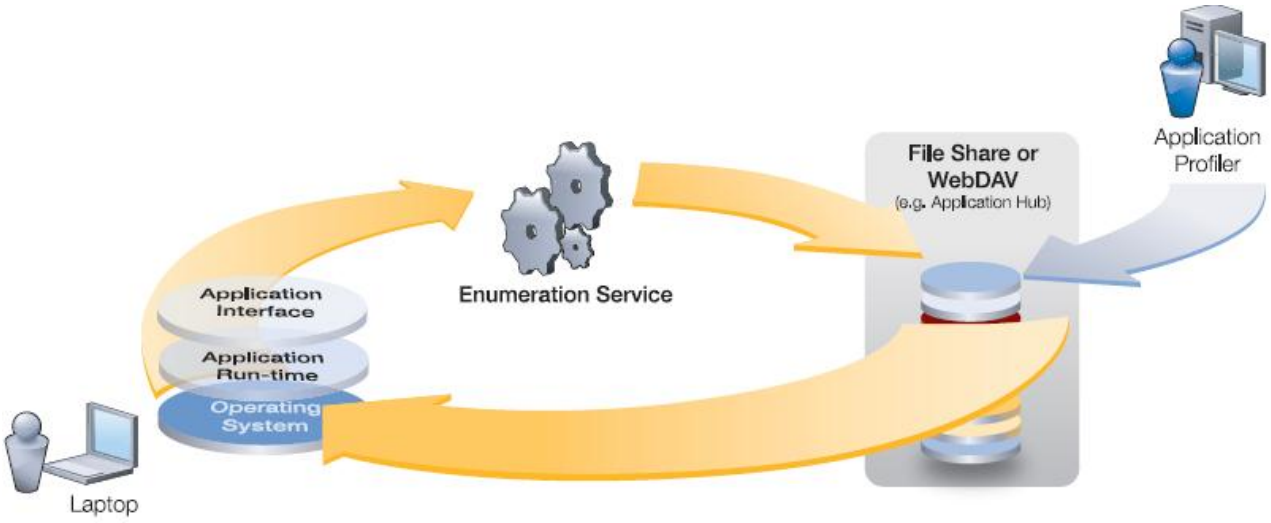

Figure 30 Application streaming, layering (V)

7.6 Summary

The ultimate goal would be to reach a stage where every virtual desktop component (OS, User profile, Applications) are virtualized and layered, this way reducing hardware requirements, easing management, eliminating errors caused by application-OS incompatibility and errors, difficulties caused by upgrades hot fixing.

As we can see in figure 31 this could be achieved by componentizing the desktop into discrete containers, independent of each other. Each layer/container could be independently provisioned, patched, inserted, and replaced without affecting the others. These containers could then be dynamically merged to ensure, that all desktops have the same IT-compliant OS and applications but retain all the user and machine personalization.

Figure 31 Layered 'composite' Virtualization (Unidesk 2012)

Chapter 8

Conclusion

As it has been presented throughout this paper, one of the most difficult decisions is choosing the best, most fitting VDI solution. Since it is quite a new technology, and it has a production status of not more than 1-2 years, the exact design parameters are vague, and are not clearly defined for every particulate requirement. There are still some elements from the VDI design, like the incapability of virtualizing certain applications that can deter infrastructure architects from using this technology.

Even so, this project and the previous Proof of Concept project (PoC) proves that for special requirements, like remotely located developers it a viable solution. Also it can be stated that with the help of this technology the BYOD concept is becoming a reality, with users being able to connect to their workstations from any locations, on a multitude of different devices just by using a simple web connection and a receiver application installed on their local device.

8.1 Status Update from 17/08/2013

This status update was made on the 17/08/2013

The VDI project has been in production for over three months, currently there are over 430 users that have been assigned a virtual machine from the platform that was presented in this project. Beside the occasional problems that I will present further on, there were no significant errors than can be related to a miss-design or faulty architecture. The success and reliability of the VDI project has ushered in requests from different business lines, which are currently requesting their own virtual machines.

The project has been in a continuous overhaul, the 100% resource utilization has been reached(memory vise) currently the number of VM exceed the resource pool that was originally calculated with the 60GB, 4GB RAM configurations and future expansion is planned. This proves that over utilization of memory and a storage resource is a viable solution with VDI technology and it does not affect the performance of the virtual machines in a significant matter.

Future plans comprise of adding 1000+ new VM to be used as remote access points, so that issues arising from the need for the secured work laptop to connect to the inner network would be eliminated. This would increase greatly the mobility and accessibility of the employees, who would be capable of working remotely and responding to emergency incident with more ease. This is a clear sign that the BYOD plan is in full implementation progress.

On the technological front, work has begun towards the implementation of a better user-profile management system, which as was presented during this paper is necessary for a complete virtualized solution. There are different products considered for this purpose, and testing is being done to determine the one that fits current requirements the best. These include products from RES Software, Microsoft and Citrix. The goal is to have a delivery platform, OS and domain independent user profile management system.

A better control over the VM has been implemented, that allows individual users to restart their virtual machines. Having this feature eases the pressure on the support units, if problems are encountered with the VM it can easily be restarted from the access interface.

The introduction of XenDesktop 7 will enable in the integration of application virtualization, and a better VM management. It will also offer a more complete control point and better visibility over the VM farm.

8.1.1 Problems Encountered and their solutions

Most of the issues reported by users were connection related, especially due to the sometimes unreliable long distance lines. There have also been issue regarding inter-domain migrations that caused Active directory mismatches that resulted in connection denials for users.

One of the biggest issues encountered was caused by regulatory differences between EU and India. The tokens, being a chip based electronic devices had issues passing customs, and even currently the grid card authentication has to be used until the legal issues can be solved. This has not affected production.

Even if the current state of the VDI solution is missing the improvements stated in this chapter, provisioning, non-persistent VM and application virtualization, these are currently being considered, and could be integrated in the next 6-9 months. Overall the project can be considered a success, and it has been a great opportunity working in the leading edge of virtualization and IT technology in general.

List of Acronyms

AD	Active Directory
BYOD	Bring Your Own Device
DDC	Desktop Delivery Controller
DRS	Distributed Resource Scheduler
DTU	Danmarsk Tekniske Universitet
IG	Identity Guard
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
ICA	Independent Computing architecture
HDX	High Definition user eXperience
IP	Internet Protocol
LAN	Local Area Network
SSL	Secure Socket Layer
LUN	Logical Unit number
NIC	Network Interface Card
RAM	Random Access Memory
RADIUS	Remote Authentication Dial In User Service
RDM	Raw device mapping
RDP	Remote Desktop Protocol
RDSH	Remote Desktop Session Host
RDVH	Remote Desktop Virtualization Host
RDCB	Remote Desktop Connection Broker
RDG	Remote Desktop Gateway
RDSH	Remote Desktop Session Host
RDVH	Remote Desktop Virtualization Host

RDWA	Remote Desktop Web Access
SAN	Storage Area Network
SCCM	System Center Configuration Manager
SMP	Symmetrical Multi Processor (
URL	Uniform Resource Locator
USB	Universal Serial Bus
VDesktop	Virtual Desktop
VDI	Virtual Desktop Infrastructure
VM	Virtual Desktop Workstation
VMA	Virtual Desktop Agent
VI	Virtual Infrastructure
WMDK	Virtual Machine Disk VMDK
VMFS	Virtual Machine File System
VM	Virtual Machine
vNIC	Virtual Network Interface Card
VCM	Volume Control Manager
VSM	Virtual System Manager

References

1. The VDI Delusion : Brian Madden with Gabe Knuth and Jack Madden
2. VDI Smackdown : Ruben Spruijt 2013
3. VDI whitepaper : Brian Madden 2012
4. http://support.citrix.com/
5. www.Citrix.com
6. Vmware.com
7. Desktop Decisions Computerworld. 3/11/2013, Vol. 47 Issue 4, p28-31. 3p.
8. http://www.citrix.com/products/netscaler-application-delivery-controller/overview.html
9. http://computersight.com/software/what-is-netscaler/#ixzz2MrYgVwOU
10. http://citriland.wordpress.com/2011/08/24/what-is-netscaler/
11. http://support.citrix.com/proddocs/topic/netscaler-ssl-93/ns-ssl-user-defined-cipher-groups-tsk.html
12. vSphere documentation center http://pubs.vmware.com/vsphere-50/
13. A secure sharing and migration approach for live virtual desktop applications in a cloud environment: Jianxin Li, Yu Jia; Lu Liu
14. Virtualizations : Issues , Security Threats, and Solutions : Michael Pearce, Sherali Zeadally, Ray Hunt
15. VDI in hastily formed networks in the support of humanitarian relief and disaster recovery missions : Albert Barreto III (2011)
16. Internal Nordea documents
17. Virtual Desktop Infrastructure vmware 2012
18. www.Intel.com
19. OS Streaming Deployment : David Clerc, Luis Garces-Erice, Sean Rooney
20. www.Microsoft.com