Henna-Riikka Maukonen

Virtual Desktop Infrastructure

Bachelor of Engineering
Information Technology

2021

South-Eastern Finland
University of Applied Sciences

South-Eastern Finland University of Applied Sciences

Author (authors)	Degree title	Time
Henna-Riikka Maukonen	Bachelor of Engineering	December 2021
Thesis title Virtual Desktop Infrastructure		34 pages
Commissioned by -		
Supervisor Matti Juutilainen		

Abstract

This thesis is about investigating Virtual Desktop Infrastructure, the technologies involved with it, such as virtualization, and the different parts the infrastructure is built from including virtual machines, connection brokers, end point devices and so on. Cloud computing is discussed as well, as it is the technology the Virtual Desktop Infrastructure was compared to.

The investigation into these technologies involves the history of virtualization and Virtual Desktop Infrastructure. Cloud computing is introduced in a way that uses mostly the definition of it since the technology is not that old. There was also an investigation of a Virtual Desktop Infrastructure solution called VMware Horizon and a cloud computing equivalent called Citrix Virtual Apps and Desktop. These two were compared at the end of the thesis.

The outcome of the thesis was satisfactory as an introduction to Virtual Desktop Infrastructure and Cloud Computing. The result of the comparison of the VDI solution VMware “Horizon” and the Cloud Computing equivalent Citrix "Virtual Apps and Desktops” was not as expected because the technologies ended up being rather similar, but their use cases were very different.

Keywords

Virtual Desktop Infrastructure, cloud computing, VMware, Citrix

CONTENTS

1 INTRODUCTION

This thesis is about the Virtual Desktop Infrastructure (VDI), how it has evolved and what different technological components must come together to make a working VDI environment. VDI is a virtualization platform in which a desktop instance is hosted on a centralized server in a data center. I want to discuss virtualization, where it came from and the different technologies it brought with it. I will investigate the history of VDI as it pertains to VMware, the company that coined the term in the first place and is at the forefront of creating the technology and continuing to improve on it. I will be investigating how VDI is different from the cloud computing alternative Desktop as a Service (DaaS) model as well. There will be a chapter on what cloud computing is, how it is defined and some important considerations when it comes to it as a technology. At the end of the thesis, I will be comparing VMware Horizon as the VDI solution and Citrix Virtual Apps and Desktops as the Cloud Computing solution to find in which kinds of environments both solutions could be valid as well as how they differ.

2 VIRTUALIZATION

This chapter is based on the descriptions in Virtualization Essentials (Portnoy 2012). Virtualization as a concept is not new in computing. It means the - abstraction of some physical component into a logical object. By virtualizing an object, you can obtain some greater measure of utility from the resource the object provides. For example, Virtual LANs (local area networks), or VLANs, provide greater network performance and improve manageability by being separated from the physical hardware. Likewise, storage area networks (SANs) provide greater flexibility, improved availability, and more efficient use of storage resources by abstracting the physical devices into logical objects that can be quickly and easily manipulated. (Portnoy 2012.)

The first time that virtualization was used in the mainstream was in the 1960s in IBM mainframes. Gerald J. Popek and Robert P. Goldberg codified the framework that describes the requirements for a computer system to support virtualization. The article they wrote in 1974 called “Formal Requirements for

Virtualizable Third Generation Architectures” describes the properties and roles of virtual machines (VMs) as well as virtual machine monitors that are still in use today. In their definition, a VM can virtualize all the hardware resources such as network connectivity, storage, processors, and memory. A virtual machine monitor (VMM) is called a hypervisor nowadays, and it provides the environment where the VMs operate.

The need for virtualization came from needing to utilize servers better and to make data centers smaller. The first commercially available virtualization for x86 computers came from VMware in 2001. Two years after that, an open-source solution Xen arrived on the market. The solutions took the form of being directly installed onto the hardware also referred to as “bare-metal” or it was a layer of software between the operating system and the VMs. What virtualization brought to the table was the ability to consolidate different physical servers into one, and through this, the server was utilized better because it ran different kinds of VMs on it instead of just the one type it used to run. A measure for consolidation is called consolidation ratio, and it is calculated by counting the number of virtual machines on a server. For example, a consolidation ratio of 4:1 means there are four VMs running on a server, and even this number of consolidations could remove three-quarters of the servers in a data center.

2.1 Hypervisor

As I mentioned earlier, the hypervisor used to be called a virtual machine monitor. It was created to solve a specific problem, but VMMs have evolved into something quite different. The first virtual machine monitors were used for development and debugging of operating systems because they provided a sandbox that could be used to run tests rapidly and repeatedly without using all the resources of the hardware. Soon the ability to run multiple environments concurrently was added, carving the hardware into virtual servers that could each run its own OS. This model has evolved into today’s hypervisor.

Nowadays, hypervisors allow us to make better use of processors that are always evolving and becoming faster. Hypervisors allow for the more efficient use of the

larger and denser memory offerings that come along with these new processors as well. Therefore, the hypervisor is a thin layer of software that sits below the virtual machines and above the hardware. Without the hypervisor, an operating system communicates directly with the hardware beneath it and as such. If there were multiple virtual machines without a hypervisor, they would want simultaneous control of the hardware, which would end very badly. The hypervisor manages the interactions between each virtual machine and the hardware.

There are two classes of hypervisors Type 1 and Type 2. The Type 1 hypervisor also known as “bare-metal” is called this because it runs directly on the underlying computer’s physical hardware, interacting with its memory, CPU, and physical storage. A Type 1 hypervisor takes the place of the host operating system. There are some advantages with the Type 1 hypervisor, such as it being very efficient because it has direct access to the physical hardware. This increases their security, because there is nothing in between the hypervisor and the CPU that could be compromised by an attacker. A big disadvantage that the Type 1 hypervisors have is needing a separate management machine to administer different VMs and to control the host hardware. Figure 1 depicts a Type 1 hypervisor.

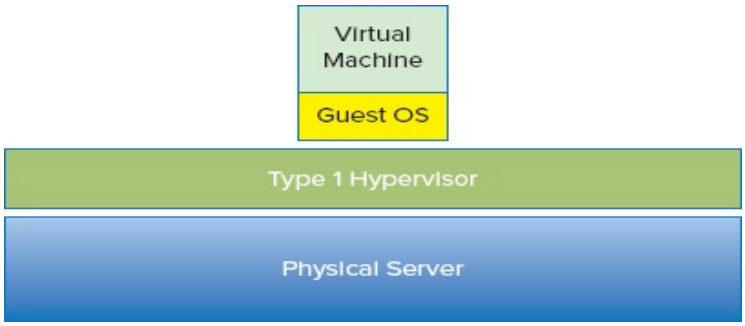

Figure 1 - Type 1 hypervisor (Portnoy 2012)

The Type 2 hypervisor runs as an application in an operating system. This type of hypervisor is rarely used in a server-based environment. Instead, they are suitable for individual PC users needing to run multiple operating systems, for example, a security professional analyzing malware. The Type 2 hypervisors often come with additional toolkits that can be installed to enhance the user experience, such as being able to cut and paste between the guest and host OS or accessing the host OS files and folders from the VM. As with the Type 1 hypervisor, the Type 2 also has its advantages such as being able to quickly access an alternative guest operating system alongside the primary one running on the host system. This is great for end-user productivity. There are also some downsides to the Type 2 hypervisor such as having to access resources via the host OS, which has primary access to the physical machine. This causes latency issues, affecting performance. It also brings the possibility of a security risk if an attacker compromises the host OS because they could then consecutively manipulate any guest OS running in the Type 2 hypervisor. Figure 2 depicts Type 2 hypervisor.

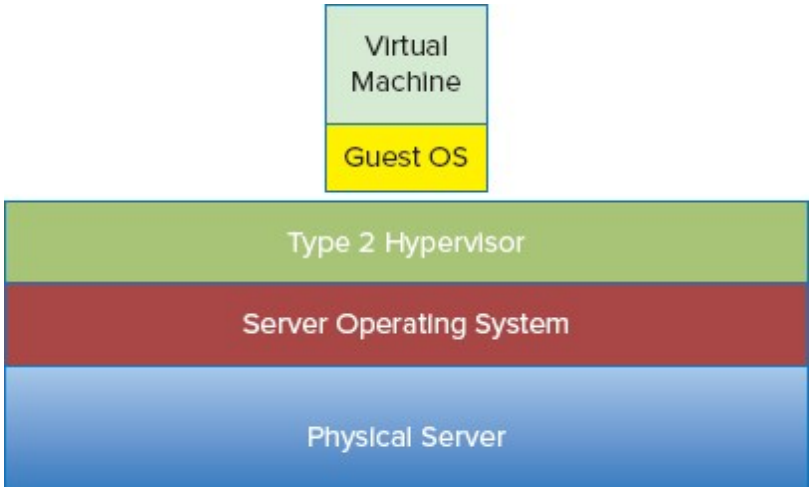

Figure 2 - Type 2 hypervisor (Portnoy 2012)

The role of the hypervisor is simple but it must have the three characteristics that Popek and Goldberg defined: being able to provide an environment identical to the physical environment, providing that environment with minimal performance cost and retaining complete control of the system resources. For the many guest systems to share the physical resources of the host, there are two things that must happen. The first is that, from the perspective of the guest, it must see and

have access to all the hardware resources it needs to function properly. The hypervisor comes in at this point to fool the guest system into believing that it can see and directly interact with the physical devices of the host when each of the guests is presented with a fraction of the resources available on the physical host. The second thing that needs to happen is that the hypervisor must abstract the hardware from each of the guests and it must balance the workload as well. Because the guests are making demands constantly from the different resource subsystems, the hypervisor must service the demands by acting as an intermediary between each guest and the physical device. This must also happen in a way that provides timely and adequate resources to all the guests.

In some ways, a hypervisor has become an operating system for the hardware, but instead of dealing with application and program requests, the hypervisor services entire servers. The hypervisor handles all the storage I/O requests, network I/O, memory processing and CPU work. Since the hypervisor does this for all guests that are being hosted on the server, it has a resource scheduling process to ensure all requested resources are serviced in a reasonable manner. Some hypervisors also have the option of prioritizing guests so that important applications can receive preferential treatment as not to suffer performance issues due to contention. Figure 3 shows how an I/O operation is processed by the hypervisor.

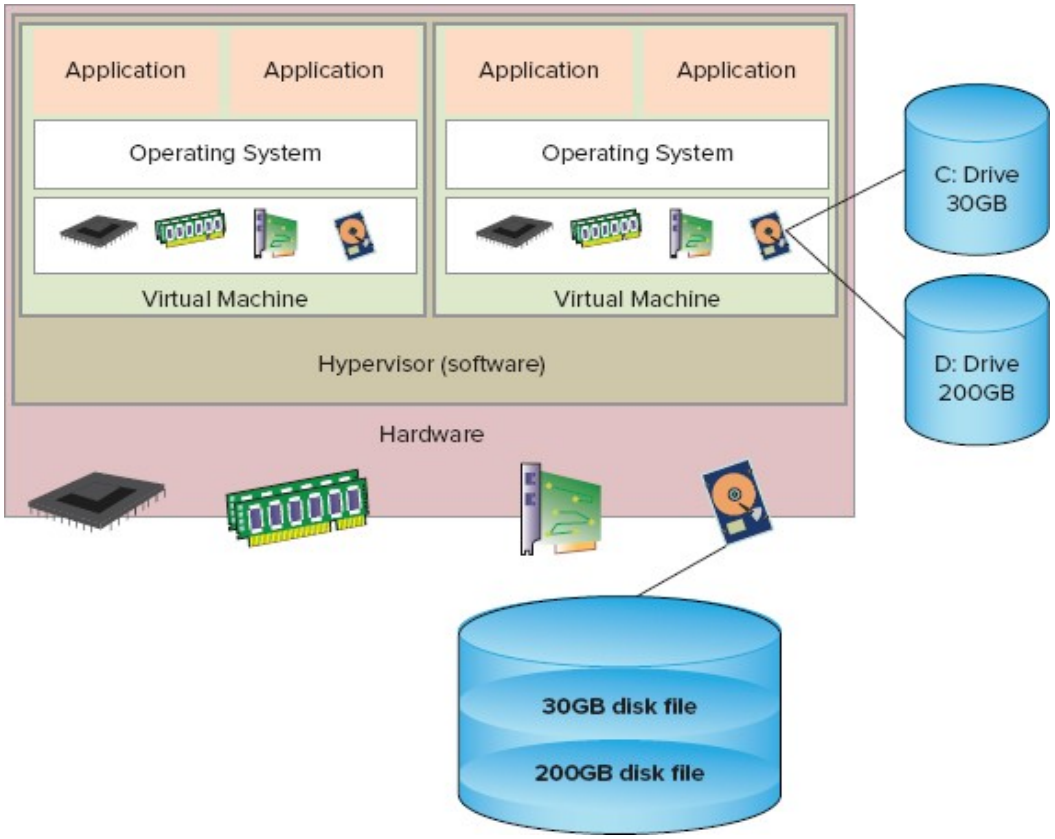

Figure 3 - I/O operation process (Portnoy 2012)

A guest application calls for a disk read and passes that request to the guest operating system. The guest operating system makes a read to the disk that it sees. Here, the hypervisor steps in and traps that call and translates it into a real-world physical equivalent and passes it to the storage subsystem. When the response returns, the hypervisor passes the data back to the guest operating system, which receives it as if it came directly from the physical device. (Portnoy 2012.)

2.2 Virtual machines

A virtual machine (VM) is the virtualization of a computer system. VMs have many characteristics that are similar to a physical server. Like a server, a VM supports an operating system and is configured with a set of resources to which the applications running on the VM can request access, but unlike a physical server, many VMs can run concurrently on a single physical server and the VMs

can also be running different operating systems supporting different applications. A VM is just a set of files that describes and comprises the virtual server, unlike a physical server. These files are the configuration files and the virtual desk files. The configuration files describe the resources that the VM can use. Virtual machines have access to various hardware resources, but from the VMs point of view, it does not know these devices are virtual. The virtual devices the VMs deal with are standard devices, which means they are the same within each VM. This makes them portable across different platforms and virtualization solutions. When looking inside a virtual machine the view is identical to being inside a physical machine. From the point of view of the operating systems or an application, storage, memory, network, and processing are all available.

3 VIRTUAL DESKTOP INFRASTRUCTURE

This chapter's description of Virtual Desktop Infrastructure is based on O'Doherty's (2012) book. Virtual Desktop Infrastructure is a virtualization platform in which a desktop instance is hosted on a centralized server in a data center. It is also sometimes referred to as a hosted virtual desktop (HVD). The desktop is accessed over a network with an endpoint device, i.e., laptop, phone, thin client, and so on. The idea of a virtual desktop is not new but the ways it is being applied is constantly evolving.

The idea of a VDI began with VMware's user community around the early to mid-2000s. The users were already using ESX (now ESXi) servers to virtualize servers and deploying a limited number of desktop operating systems. The infrastructure evolved to deal with specific use cases, such as delivering an isolated remote access environments or centralizing developer desktops. In 2006, VMware coined the term VDI when the company created the VDI alliance program in which VMware, Citrix and Microsoft developed VDI products for sale.

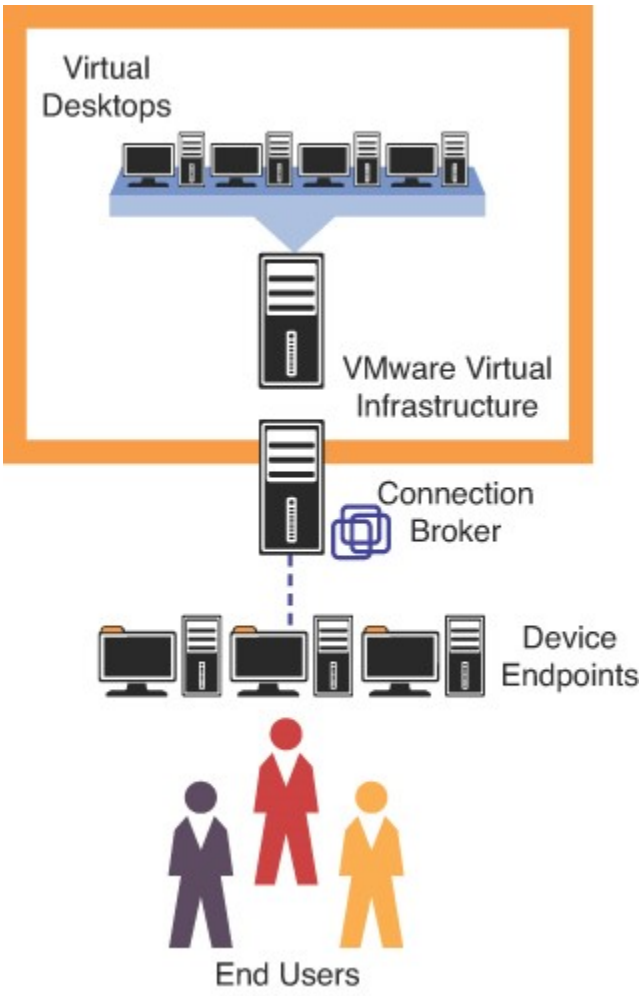

Figure 4 - Early depiction of VDI (O'Doherty 2012)

Figure 4 depicts an early design of VDI with an overview of the technology involved with it. VDI has evolved since then and now involves the hypervisor to segment the servers into Virtual Machines (VMs).

3.1 VDI Challenges

One of the biggest challenges in the early virtual desktop environments was the cost of storage. This was a problem because most of the early VDI environments used the same underlying hardware and storage assets found in their server virtualization environments. Scalable virtual infrastructure consisted of several VMware ESX hosts attached to Storage Area Network (SAN) storage. Virtual

infrastructure is enabled using shared storage, since host clustering, hot migrations and distributed resource leveling are not possible without using some form of shared storage. Even though now there are many different cost-effective and high performance-providing ways of shared storage, that was not the case historically.

Organizations would tier storage to rationalize the cost of VDI, but still it was expensive when contrasted with the price of local storage. For the technology to have practical application, a simpler approach was needed to manage the operating system images and reduce storage costs. The current VDI products such as VMware Horizon allow the incorporation of both local and shared storage to deliver the best cost and performance advantage. The earlier releases addressed the issue by using advanced image management solutions to deliver storage savings on traditional shared storage technology.

The latest-generation Virtual Desktop Infrastructure comes with image management software that allows a single image to be served out to many virtual desktops. This advancement in the technology allows a 15 GB desktop image to be shared between multiple virtual machines while appearing as an independent desktop OS to each user saving costs and simplifying the management of images. As the master could be patched and upgraded once consequently the changes were propagated to all users at the next reboot or login, the operational costs were lowered as well when compared to a distributed desktop environment. VMware has its own proprietary solution called VMware Horizon that is a solution for enterprises to deliver virtual desktops, so they are a competitive alternative to a distributed desktop environment. Even though they are a competitive alternative in some organizations, a percentage of the desktops are still physical although virtual desktops are an integral part of the desktop environment. Other organizations have moved on from physical desktops, because as the technology and the return on investment improved, it might be cheaper to deploy virtual desktops instead of physical ones. This is especially true when looking at the operational cost of managing virtual desktops that incorporate technology such as application virtualization. When a Virtual Desktop Infrastructure is properly

designed, a smaller number of administrators can manage a larger number of desktops.

In the beginning, there was a realization that moving the operating systems to a data center did not mean that there was the ability to take full advantage of the technology as virtualized desktops and decentralized desktop management share some challenges, such as application lifecycle management, testing, deploying, upgrading, and removing applications that are no longer needed. There is also the issue of the more applications you install into the desktop images; the more images need to be maintained. Because of this, the images were segregated based on the applications that they would use. To battle this application, virtualization was created. This allowed a form of packaging that isolates the application in a separate file that has read access to the underlying operating system but only limited or redirected writes. The VMware ThinApp can be configured to write no files to the underlying OS. ThinApp is provided as a part of the VMware Horizon Suite. ThinApp is used to integrate into the virtual desktop environment to handle most application incompatibilities that arose from the way the images were maintained before application virtualization. Application virtualization should be used when there is a different application workload for every business unit to reduce the operational overhead of managing applications. This technology came out in 2008. There were other competitors for the ThinApp besides Ctirix XenApp and Microsoft App-V, but ThinApp had the benefit of not requiring back-end infrastructure to run it.

User data management was also a challenge that needed to be solved. Since there were advanced image management capabilities in a virtual desktop environment, the images had to be homogenized by ensuring that the user data is redirected and not written directly onto the desktop image. VMware solved this by acquiring a company that specialized in profile and user data management RTO Software. The software that was acquired made it so you could use it instead of Windows profiles that were problematic with a large number of users, or it could be used to complement them.

When it comes to virtual desktops, VoIP comes up because the information must travel long routes in traffic from the desktop since it cannot take the shortest possible route between two endpoints. It is possible to integrate VoIP into VDI by using integrated platforms that allow a unified delivery of both virtual desktops and softphone technology. These solutions allow the virtual desktop display and VoIP traffic to be dealt with as separate streams but integrated from the thin client device to provide a seamless experience for the user.

After overcoming these challenges, VDI has now grown to become able to supplant physical desktops in varying environments such as schools and businesses. The VDI desktop can be either persistent or non-persistent depending on the needs of the specific environment it is being deployed in. Although VDI offers advantages such as mobility and ease of access, it is still not necessarily a suitable solution depending on business processes and complexity of applications that will be running on it.

3.2 Virtual machine in VDI

A VM has its own storage, memory, CPU, and network interface. They can be run locally on a computer, from the cloud or from a dedicated server. Virtual machines run their own operating systems and function separately from other VMs that may be on the server because of the hypervisor that segments them into their own sandbox. Because of this, the VMs can be deployed to accommodate different processing power needs, run software that require different operating systems or to make a safe sandbox to test software. The Virtual Desktop Infrastructure is used to leverage the VMs to provision and manage applications and virtual desktops.

3.3 Hypervisor in VDI

The hypervisor is a thin layer of software that enables virtualization and as such supports the whole infrastructure. It allows the use of multiple operating systems (VMs) to run alongside each other while sharing the same physical computing resources. The hypervisor separates the VMs from each other and allocates to

them their own memory, storage, and computing power. Because of this, even if one of the operating systems were to crash it would not interfere with the others. Figure 5 shows where the hypervisor sits in the architecture and how the virtual machines are segmented inside it in a server environment.

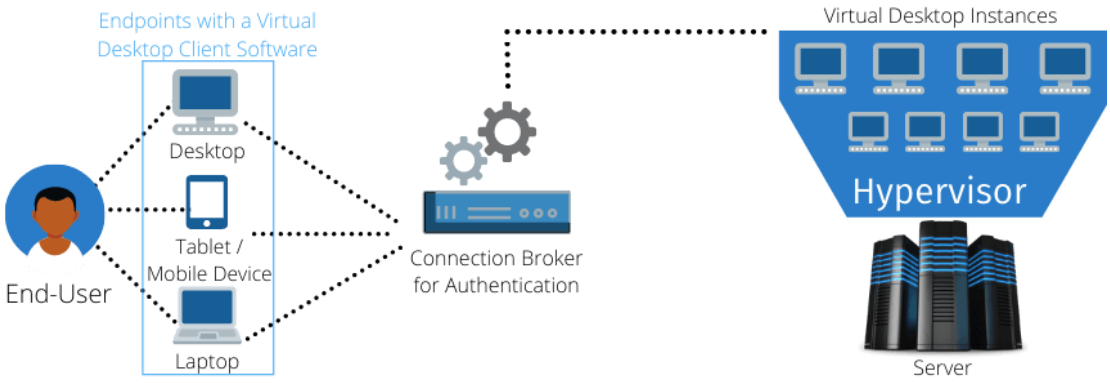

Figure 5 - Hypervisor (Casserly 2020)

When it comes to hypervisors, there are different categories and brands of hypervisors. I will be focusing on the one that VMware provides: EsXi 5 (Elastic Sky X Integrated). This hypervisor is a bare-metal embedded hypervisor also known as Type-1 hypervisor, and it runs directly on physical hardware without any underlying OS. It is its own micro operating system. ESXi runs the vmkernel. This kernel handles access to the underlying CPU and memory resources on the physical server.

3.4 Connection broker

In the early days of VDI, a connection broker provided a new management point by brokering user requests for end-user connections and virtual desktops. This allowed administrators to see the status of the connections. It allowed some administrator functionality as well, such as the ability to reset the connection or reboot the virtual desktop. Deploying virtual desktops without the connection broker did not scale well because the more users were deployed, the less manageable and visible they were. It is very important to see the number of users, because without knowing how many there are, it is impossible to manage the connections. The connection broker also gave the administrator the ability to

enable users' access to a virtual desktop. The early connection brokers did little to enable customers to replace a physical desktop experience although they did enable the deployment of virtual desktops.

After some development from the early days of VDI, the connection broker is now used to associate users to virtual desktops. It is also called a session manager. The connection broker can tunnel connections between users and the virtual desktops to provide additional security such as SSL encryption. In VMware Horizon, the connection broker is a proxy to the user and the Horizon desktop; therefore, it is a critical component that always needs to be available. It is of utmost importance to understand what happens if a connection broker is unavailable. Often the testing of a broker failure uncovers another single point of failure such as where the virtual desktop environment stores its configuration information. It may also uncover the requirement for additional front-end services, such as hardware-based load balancing, to ensure failover occurs without administrator intervention. As with any technology, failures can occur; it is important to fully understand what happens when each component fails so that you can either reduce or understand the associated risks. (O'Doherty 2012.)

3.5 Endpoint device

These endpoint device descriptions are from 10ZIG (2018). By the nature of the name "endpoint device", it is the point at which communication begins and ends over a network. Endpoint devices are the actual hardware pieces used by the end users. They can be PCs, laptops, tablets, mobile phones, Thin Clients and Zero Clients. They are any hardware device that is internet-capable. In the Virtual Desktop Infrastructure, the endpoint devices play a crucial role. Thin Clients and Zero Clients are becoming more the norm with VDI, since all the data is stored and managed on a central server. Endpoint devices are what users use to login from their remote workstation and connect to the server via network. Thin Clients usually replace a PC. It has no hard drive but will have some local storage for the customization of end-user applications. With it, you can use different connection brokers. Zero Clients are also used to replace a PC. They have no local OS or storage. All applications for them are provisioned from the server.

3.6 Virtual desktop

Virtual desktops are preconfigured images of operating systems and applications in which the desktop environment is separated from the physical device used to access it. Users can access the virtual desktops remotely over a network. There are two types of virtual desktops: persistent and non-persistent desktops. Persistent desktops as an implementation means that the end user maintains their personalized settings, stored data and configured instances that can be retrieved each time they log in to the desktop. The persistent desktop is the most like a physical PC. This implementation is popular in environments where end users are always the same. In the non-persistent desktop implementation, the users cannot retain data or configure the desktop instance because the data is destroyed after each session. This type of desktop is an ideal solution when there is a need for one-time access scenarios. Shift, task, and kiosk employees who do not need to save their desktop instances could use them.

3.7 Security

This subchapter's information about the security of VDI is mostly based on VMware (n.d). There are some security risks related to VDI since it is a mission critical technology that stores sensitive data and applications. There are four primary attack surfaces: the hypervisor, the virtual machines, the network, and employees. The hypervisor is vulnerable to an attack called hyperjacking. This means that an attacker has managed to infiltrate the hypervisor using malware, and as a result, they have access to everything connected to the server such as all storage resources and virtual machines. Although this type of attack is rather elusive it still presents a single point of failure in the virtual desktop environment. The second attack surface is the virtual machines themselves. Since they have their own operating systems and configuration, if they are not patched and automatically updated, the delays in deployment of security patches and updates put the entire environment at risk. The third attack surface is the network since all networks are vulnerable to attacks. Because virtual environments share the same physical resources, they are especially vulnerable to network attacks. The

surface can be lessened by implementing network segmentation. The last primary attack surface is the employees. As it is, a growing cause of data breaches are the insiders this is especially true with VDI, because the employees connect to virtual desktops running as part of the VDI system. There is the possibility of someone attempting to breach other employees' desktops or the VDI servers.

VDI's security architecture is critical in minimizing desktop security vulnerabilities that are common to virtual environments. The key components to it are unified management platform, real-time compliance monitoring, vulnerability-scanning and data loss prevention. The unified management platform means that a single virtualization platform keeps track of the virtual and remote desktops, furthermore it also simplifies and expedites the provisioning of virtual desktops. It protects the data center and workloads better as well. The real-time compliance-monitoring technology is used to monitor the virtual infrastructure for anomalies and sudden changes. If such a change or anomaly is found the monitor automatically gives an alert and the problem can be swiftly dealt with therefore it ensures the preservation of integrity of the virtual desktop data and resources. Vulnerability-scanning is the process of identifying security weaknesses and flaws in the system automatically. It has a built-in response to suspicious activity that can be blocking network traffic, quarantining a virtual machine or other such actions. Data loss prevention is the encryption of virtual machine files, virtual disk files and core dump files. Encrypting the virtual machines helps better protect the organization's sensitive data and meet compliance standards.

There are some best practices like with most technology. Here is a list of some from VMware.

- Setting controls to disable a device in local mode if it is not synchronized within a predetermined time interval. This allows companies to stay ahead of hackers on the lookout for new ways to outsmart security protocols.

- Preventing unauthorized access by establishing stringent policy-driven access controls for desktops and apps across corporate and employee-owned devices.

- Quarantining intrusions with micro-segmentation so that they don’t spread across the network.

- Protecting data by leveraging built-in encryption capabilities, data at rest encryption and distributed firewall support.

- Investing in employee training to minimize data leakage from lost or stolen devices.

- Ensuring endpoint protection by applying the latest security patches to operating systems, constantly updating antimalware software, and taking advantage of an endpoint device’s built-in hardware security capabilities. (VMware n.d.)

Following these best practices, the VDI environment is better protected.

3.8 Benefits and disadvantages

Some of the benefits of VDI are device flexibility, security, user experience, scalability, and mobility. One of the biggest benefits of VDI is device flexibility, since it can be used on pretty much any device. Older devices can be used as end devices for VDI as well so their life span can be extended. If, for example, a company would want to switch from normal PCs to VDI there would not be a need to purchase new devices such as thin clients just for using VDI. It should also be noted that when the time comes to purchase new end devices the company could just purchase less powerful end devices that would in turn be less expensive.

Security is one of the benefits as well since the information is not on the physical device but in the data center. Because of this, if the end device would be stolen, no data could be gained from it. The user experience is very good when it comes to VDI since the desktop is standardized and the user has access to a consistent workspace wherever they are working.

The VDI environment is easily scalable as well. If there was a need to expand it for seasonal employees or other such occasions it can be done quickly within the limits of the data center. VDI includes the ability to support remote and mobile employees. As currently, remote working is becoming more popular due to the

COVID-19 pandemic, and companies have realized that remote working is just as effective as working at an office. The mobility that VDI brings is very important in this instance. The mobility of VDI is also useful in the case of a company having consultants and such since they do not need a physical device from the company and can work from their own devices or devices that were provided by their own organization. This can be seen as a security risk in some compliance requirements such as Katakri by Ministry of Foreign Affairs of Finland.

There are some disadvantages to VDI as well, such as additional costs, complex infrastructure, licensing issues and reliance on network connectivity. There are some additional costs in the form of IT infrastructure expenses, personnel, licensing, and other things that might be more expensive than expected. When it comes to VDI there is a huge consideration and that is licensing, it is very easy to make a licensing breach when using VDI because some licensing and support agreements do not allow software to be shared among multiple devices/users. The initial procurement for VDI licensing, ongoing maintenance and support agreements is very time consuming and the costs can be high as well. To add to this there are the licenses for Microsoft Windows workstations.

There is also the complex infrastructure of the environment to take into consideration. VDI requires several components to work flawlessly together to provide the virtual desktops. If there is a problem with any of the back-end components, they may stop users from making virtual desktop connections. Another thing that will prevent users from making connections is if there is no internet connectivity. A poor connection will most likely cause a poor user experience. The reliance on network connectivity is probably the biggest disadvantage of VDI. To add to these disadvantages there is also the potential for poor user experience if there is no sufficient training provided to users. Having access to the local desktop and virtualized desktop may be confusing for the users.

4 CLOUD COMPUTING

Cloud computing is an on-demand network access to computing resources which most of the time are provided by an outside entity. These resources require little management. The term “cloud” started as a convenient way to refer to an abstracted network used by engineers. It was convenient from a marketing perspective as well since it could describe different types of solutions that had one thing in common and that was the use of the Internet. The key aspects of cloud computing are scalability, elasticity and delivery as a service over the internet. Scalability in cloud computing aims to bring up the contrast between on-premises computing and the fact that it is easier to scale computing and storage in a cloud computing model. Elasticity aims at the fact that scalability goes both ways: it will scale down when capacity is not needed. According to The US National Institute of Standards and Technology (NIST) the definition of cloud computing is “*Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction (Mell and Grance 2011).”*

According to this definition, the cloud model is composed of five essential characteristics, three service models and four deployment models. The five essential characteristics are:

1. On-demand self-service.
2. Broad network access.
3. Resource pooling.
4. Rapid elasticity.
5. Measured service.

The three service models are:

1. Software as a Service (SaaS). It is a type of software that is deployed over the Internet. In this model, the vendor manages everything, and the user can program nothing. Only thing that the user can do is configuring and using the software through a web browser. Examples of these types of

software are Microsoft Office 365, Google docs and enterprise applications such as SAP SuccessFactors and ServiceNow.

2. Platform as a Service (PaaS). It is a complete cloud platform containing hardware, software, and infrastructure for developing, running, and managing applications without the cost, complexity and inflexibility that often comes with building and maintaining the platform on-premises. The PaaS provider hosts everything at their data center—servers, networks, storage, operating system software, databases, development tools— and users typically pay a fixed fee for specified amount of resources, or they can choose to pay only for the resources they use. Some examples of these platforms are Google Cloud, IBM Cloud and Amazon Web Services.
3. Infrastructure as a Service (IaaS). It delivers fundamental computing, network, and storage resources to users' on-demand, over the internet and on a pay-as-you-go basis. IaaS allows users to scale and shrink resources on an as-needed basis, reducing the need for "owned" infrastructure. IaaS gives the most control but also requires the most work out of these three service models.

The four deployment models are:

1. Private Cloud. It is a Cloud computing environment dedicated to a single customer. It can be owned, managed, and operated by the organization; a third party or some combination of the two and can exist on or off premises. Private cloud is the only way to meet regulatory compliance requirements. It is often chosen because a company works with confidential documents, intellectual property, medical records, or other sensitive data. Private cloud is more expensive than the other cloud models, but it provides more security than they do.
2. Community Cloud. It is a type of cloud model that is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns, be it industrial groups, research groups, standard groups and so on. It is also a hybrid of the private cloud, so it provides a level of privacy, security, and policy compliance.
3. Public Cloud. It is a type of cloud infrastructure provisioned for open use by the public. In this model, users do not need to purchase hardware, software or supporting infrastructure because it is owned and managed by the provider. Some examples of public cloud are Amazon Elastic Compute Cloud (EC2), Microsoft Azure, IBM's Blue Cloud, Sun Cloud, and Google Cloud.
4. Hybrid Cloud. It is a type of cloud model that combines two or more distinct cloud infrastructures, which remain unique entities. These entities are still bound together by standardized or proprietary technology that allows data and application portability.

This concludes the brief overview of the general idea of Cloud Computing.

4.1 Security

There are some security concerns when it comes to cloud computing. Most of these have been brought up in the paper by Morsy et al. (n.d). Some of these are that organizations outsource their security management to a third party that hosts their IT assets which results in loss of control. Different organizations sharing the same infrastructure without knowing of the strength of security controls is an issue that needs to be addressed. The Service-level agreements do not have very strong security guarantees. Depending on the cloud model the infrastructure is publicly available and thus more susceptible to attacks.

There are some security advantages to cloud computing as well. The following are some of them according to Walker (2019). Protection against Distributed Denial-of-Service (DDoS) attacks, because there are measures that stop huge amounts of traffic aimed at a company's cloud servers. Data security, because the top cloud computing solutions have protocols in place to protect sensitive information and transactions. The top cloud solutions help companies to achieve data compliance in regulated industries by managing and maintaining enhanced infrastructure for compliance and to protect personal data. One of the necessities of a cloud computing solution is high availability, which is achieved by real-time support that includes live monitoring. There are redundancies in place as well that ensure the availability of the service.

Here is a list of cloud security best practices:

- Have a working cloud security program in place.
- Leverage data classification tools to identify sensitive or regulated data, then evaluate how that data is being accessed, used, and shared. Evaluate access permissions for files and folders containing sensitive data, as well as specific roles, locations, and the types of devices used to access that data.
- Ensure that your most sensitive data is segregated from the cloud provider's resources and those of other customers. A private cloud is often the best choice for highly sensitive data or data subject to strict regulatory requirements.
- Leverage user and entity behavior analytics (UEBA) tools to monitor for suspicious activity, so you can act quickly to protect your sensitive data in the event of unauthorized access.

- Implement cloud security controls. In addition to data classification and UEBA, implement tools for managing permissions and access privileges, password control, disaster recovery, malware prevention, and encryption.
- Perform regular monitoring, vulnerability scans, system audits, and patches to detect and fix cloud security risks.
- Ensure that cloud data security practices are in line with industry regulations and compliance requirements. (Brook 2020.)

These practices are a good guideline for a company thinking of moving forward with Cloud Computing models.

4.2 Benefits and disadvantages

The information in this subchapter is mostly based on the paper by Avram (2014). One advantage that cloud computing brings especially to smaller companies is the lower cost of entry to benefit from compute-intensive business analytics that historically have only been available to larger corporations. Cloud computing is a huge opportunity to third-world countries that do not have their own IT infrastructure since some computing providers are using the cloud platform to enable IT services in these countries. It is a way to gain immediate access to hardware resources without the need for upfront capital investments by users, which leads to faster time to market in many businesses. When treating IT as operational expense this enables the reduction of upfront costs in corporate computing. Cloud computing can help lower IT barriers for innovation. It also makes it easier to scale services according to client demand. Cloud computing has enabled new types of applications and services that were not possible before.

Some of the disadvantages that come up with cloud computing are security and privacy concerns as was mentioned in the previous chapter. The full potential of cloud computing is tied to the availability and high-speed access to everything involved with it. The reliability of the resources that cloud computing provides must be high and contingency plans must be in place in case of failures or outages. These must be negotiated when going over the SLA and tested in failover drills. The costs may be surprisingly high for applications because of pricing plans that cloud computing providers offer. The cloud applications may

not have the same features as an application that would be run in-house and these need to be noted when choosing an application to use. There is also the incompatibility between a cloud application and other applications. For example, it is not possible to insert a spreadsheet made in another application into Google Sheets.

5 COMPARING VMWARE HORIZON AND CITRIX VIRTUAL APPS AND DESKTOPS

In this chapter, I will compare the VDI solution by VMware, “VMware Horizon”, and the Cloud computing equivalent from Citrix, “Citrix Virtual Apps and Desktops”. I want to compare the use cases for these two and the differences they may have. First, I will be introducing the two solutions on their own and then making the comparison between the two.

5.1 VMware Horizon

The information in this subchapter is based on the book by Von Oven and Coombs (2019). There are different product editions within the Horizon solution portfolio for different platforms and use environments, each having different functionality and features. As it is, the Horizon solution is rather secure because no data leaves the data center unless there is a policy made to allow it. End users cannot introduce malware or other malicious content to the desktop as well, because all that is being transmitted to the client device is screenshots of the virtual desktop, with mouse interactions being sent back to the virtual desktop. Since Horizon is a VDI solution, it has centralized management that makes updating and patching easy. It makes it easier to troubleshoot issues as well. The architecture of Horizon is rather simple as shown in Figure 6.

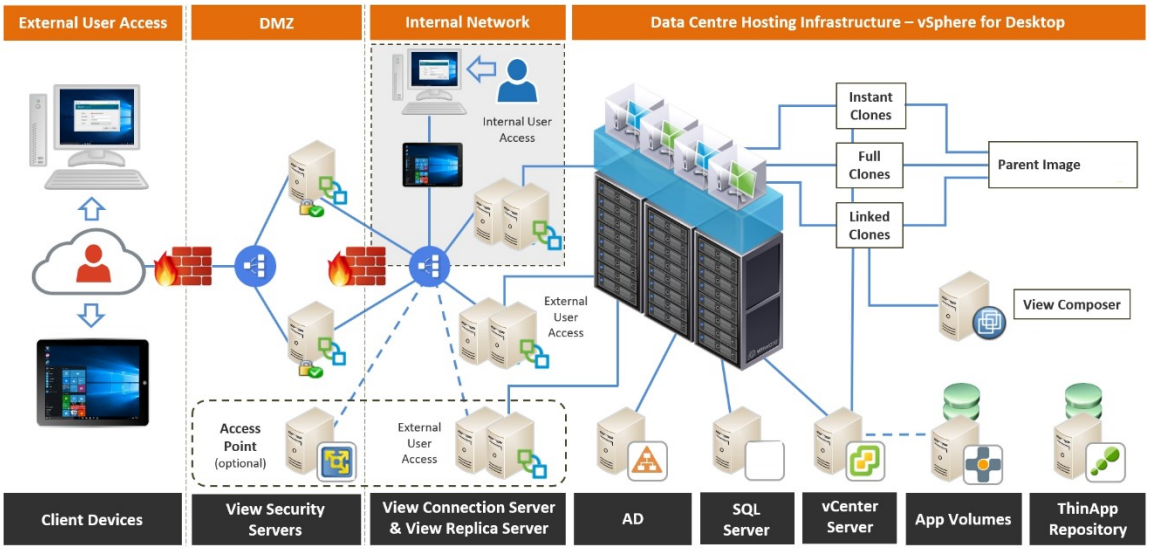

Figure 6 - Horizon View architecture (Von Oven & Coombs 2019)

The infrastructure components run as applications on the Microsoft Windows Server, except for the Unified Access Gateway, which is a hardened Linux-based appliance.

The Horizon View Connection Server is the connection broker of this architecture. The user is authenticated with Active Directory. Depending on what resources are available to the user, they will see a few different virtual desktop machine icons on the launch screen. These icons represent the desktop pools the user is entitled to use. Once the user is authenticated, the connection server makes a call to the vCenter server to create a virtual desktop machine. Either the vCenter makes a call to the View Composer or it will create an instant clone using the VM fork feature of vSphere to start the build process of the virtual desktop if there is not one already available for the user to log in.

The Horizon View security server is a role performed by the connection server, but this portion of the architecture sits within the demilitarized zone (DMZ) and not inside the domain. This allows users to connect securely to their virtual desktops without the need for a virtual private network (VPN). Both the Horizon connection server and Horizon security servers are in fact brokers but the Horizon security server proxies connections to the connection server.

The Horizon View replica server is a copy of a view connection server. It has two key roles. The first of these roles is to enable high availability for the Horizon view environment since the replica server in the case of server failure takes over. The second role is that it enables the scaling of users and virtual desktop connections. An individual instance of a connection server can support 2,000 connections and adding another connection server supports another 2,000 up to 10,000 users per Horizon view pod. There is a lot more to the architecture but for the purposes of this thesis this overview will suffice.

Good use cases for Horizon are listed here:

Static task employees that are typically fixed to a specific location with no need for remote access. These types of employees could include call center employees, administration employees and retail users. This type of employee often uses a small number of Microsoft Windows applications, does not install their own applications nor do they require SaaS application access. Might require location-aware printing.

Mobile knowledge employee that could be hospital clinician, a company employee or have a finance or marketing role. This type of employee typically uses applications from a corporate location but might access applications from mobile locations. They might need access to SaaS applications but do not install their own applications. They require access to USB devices. Might require access to two-factor authentication because of working from remote locations as well as location-aware printing.

External contractors often require access to specific line-of-business applications, usually from a remote or mobile location. Use subset of core or department applications based on the project they are working on. May require SaaS application access. They have restricted access to the clipboard, USB devices

and so on. Require two-factor authentication. Might use different type of operating system but requires access to Windows based applications for the duration of their contract.

5.2 Citrix Virtual Apps and Desktops

The information in this chapter is based on the document by Citrix Staff (2021). Citrix Virtual Apps and Desktops is a Desktop as a Service solution made by Citrix. It is a virtualization solution that gives IT control of virtual machines, applications, and security while providing access for any device. This solution allows the use of hybrid cloud model to manage on-premises data center and public cloud. In this solution, the control over applications, policies, and users stays with the company instead of Citrix. Citrix manages most of the installation, setup, and upgrades. Figure 7 shows the components and services of Citrix Virtual Apps and Desktop.

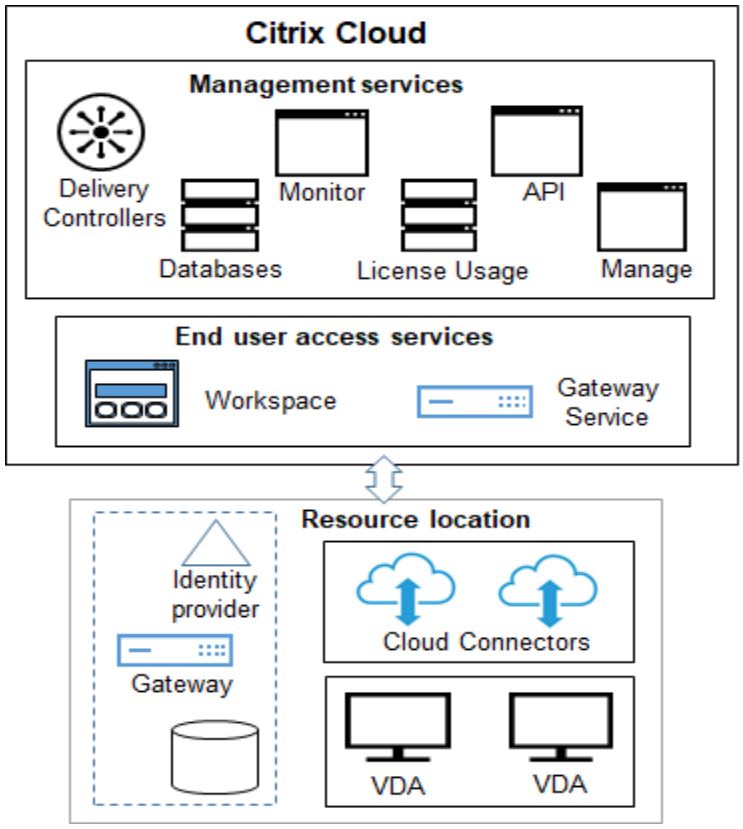

Figure 7 - Citrix Virtual Apps and Desktop site overview (Citrix Staff 2021)

The Citrix Virtual Apps and Desktop service production deployment is called a site. Citrix manages the user access, and management services and components in Citrix Cloud. The applications and desktops reside on machines in one or more resource locations. In this service deployment model, a resource location contains components from the access layers and resource layers. Each of these resource locations is called a zone which are used to keep the applications and desktops closer to the user, which improves performance because the connections do not need to traverse large segments of the Wide Area Network (WAN). The zones can be used for disaster recovery, geographically distant data centers, branch offices, a cloud, or an availability zone in a cloud. The components and services that Citrix manages are:

1. Delivery controls
2. Databases
3. Licensing
4. Management interfaces
5. Monitor interface
6. Cloud Connectors

The delivery controls are the connection broker in this deployment. The monitor interface provides administrators and help desk teams an environment to monitor, troubleshoot problems and perform support tasks for end users. Cloud connectors are the communication channel between the components in the Citrix Cloud and components in the resource location. In the resource location, the Cloud Connector acts as a proxy for the Delivery Controller in Citrix Cloud.

The components and technologies that the customer manages are:

1. Citrix Gateway
2. Active Directory
3. Identity Provider
4. Virtual Delivery Agents (VDA)
5. Hypervisors and cloud services
6. Citrix StoreFront

Citrix Gateway enables connections from remote locations while securing the connections with Transport Layer Security (TLS). Citrix Gateway is a Secure

Socket Layer Virtual Private Network (SSL VPN) appliance deployed in the DMZ. The Identity Provider can be on-premises Active Directory, Active Directory plus token, Azure Active Directory, Citrix Gateway, and Okta. Each physical or virtual machine that delivers resources must have Citrix Virtual Delivery Agent installed on it. VDAs establish and manage the connection between the machine and the user device and applies policies that are configured for the session. The Citrix StoreFront is used as the web interface for access to applications and desktops. When it comes to security the Citrix Virtual Apps and Desktop has been evaluated for compliance requirements and they have solution architectures for General Data Protection Regulation (GDPR). The security is well built into the deployment and when using the cloud best practices for security it is enhanced even more.

Good use cases for Citrix Virtual Apps and Desktop are listed here:

Temporary workforce or expanding workforce. Whether the expansion is seasonal or the need to expand into new territories this solution brings the scalability to instantly provision desktops to employees and if needed shut them down when they are not needed anymore.

In mergers and acquisitions, the transition period can be smoothed over by using Citrix Virtual Apps and Desktop to provide equal access to shared company resources.

5.3 Comparison

I think the main comparisons that can be made between Horizon and Citrix Virtual Apps and Desktop are related to costs, control and management, flexibility and scalability as well as access to resources.

To start with, the costs the infrastructure needed to host and support a data center can be high especially at first. Even though at first there may be a lot of costs in the future if the workforce stays relatively the same there most likely would-be savings. With Citrix Virtual Apps and Desktops, the cost at first is lower

and with a pay-per-user subscription model the costs can be predicted better than when you have your own data center that you must maintain and maybe expand. There are cost from licensing as well that need to be taken into consideration since with Horizon that falls on the company and it is a lot of work to keep up with all the licenses of a company. This is outsourced to Citrix in Citrix Virtual Apps and Desktop and so there is no additional cost on the company to have the workforce keep up with them. From this, a conclusion can be drawn that if there is a lot of fluctuation in a company maybe DaaS might be the better choice.

Control and management with VDI are all encompassing; the company must keep up with everyday maintenance and security, troubleshooting, and all software and hardware updates. The need for data compliance considerations is also on the company. However, control and management may be difficult to hand over to a third party in some instances such as in government organizations. With Citrix Virtual Apps and Desktop, there is less internal control on the company. Citrix has been evaluated for some compliance requirements, and they can help the company to achieve compliance, but if it is not possible to meet the requirements VDI might just be the better choice.

The flexibility and scalability with Horizon are dependent on how the infrastructure was built in the first place. If there is a need to expand and there is no buffer built into it there will be need for upgrades in hardware, this of course means more costs. Citrix Virtual Apps and Desktop is rather flexible since Citrix provides the infrastructure. There will not be a need to expand your own infrastructure, but adding more subscriptions means more costs as well.

Access to resources with VDI is pretty much guaranteed since the servers belong to the company and no one else has access to them meaning that there is no risk of interference or interruption because of a sudden rush of other users. However, if there would be a time that the data center was down all the company resources are unusable. With Citrix Virtual Apps and Desktop, you share resources in the cloud with others and this either could result in disruptions from a lot of users or

just attacks on the cloud such as DDoS. There is also the need to always be connected to the internet, which can affect accessibility.

With these considerations, I think Horizon is a good solution for companies that have a stable workforce that does not have a lot of fluctuation. Of course, with a buffer on resources in the infrastructure, the fluctuation is not a bad thing, but then there might be unused resources just sitting in the data center. In the case of being a governmental organization or something comparable to it this might be the better choice between the two solutions. As for Citrix Virtual Apps and Desktop it is a good solution when there are different branches in different territories or otherwise need for flexibility and scalability in the number of users.

6 CONCLUSION

I think that this thesis is a successful introduction to Virtual Desktop Environment and Cloud Computing. I learned a lot during the process starting from the history to the different technologies that make both these deployment models possible. There are so many considerations when it comes to VDI and cloud computing that it is difficult to make comparisons that are not just dependent on the overall cost of them, because in part cloud computing has been created to make it easier to gain access to high level computing without your own infrastructure with less upfront costs in infrastructure. The comparison could have been a bit more in depth but as an introductory level to these technologies, it felt like it was enough.

REFERENCES

10ZIG. 2018. Endpoint Devices – A Understanding by Definition and Review of the Different Types WWW document. Available at: https://www.10zig.com/resources/vdi-blog/endpoint-devices [Accessed 23 October 2021]

Avram, M.-G. 2014. Advantages and challenges of adopting cloud computing from an enterprise perspective. PDF document. Available at: https://www.sciencedirect.com/science/article/pii/S2212017313007l0X [Accessed 12 November 2021]

Brook, C. 2020. What is Cloud Security?. WWW document. Available at: https://digitalguardian.com/blog/what-cloud-security [Accessed 12 November 2021]

Casserly, P. 2020. What is VDI and how it works? – Boston IT Support. WWW document. Available at: https://www.casserlyconsulting.com/the-edge/what-is-vdi-and-how-it-works-boston-it-support [Accessed 25 September 2021]

Citrix Staff. 2021. Citrix Virtual Apps and Desktops service Overview. WWW document. Available at: https://docs.citrix.com/en-us/citrix-virtual-apps-desktops-service/overview.html [Accessed 12 November 2021]

Citrix. 2021. Privacy & Compliance. WWW document. Available at: https://www.citrix.com/about/trust-center/privacy-compliance/ [Accessed 13 November 2021]

Griesemer, M. 2021. VDI vs. DaaS: What's the difference?. WWW document. Available at: https://www.citrix.com/blogs/2021/08/24/daas-vs-vdi-comparison/ [Accessed 12 November 2021]

IBM Cloud Education. 2019. Hypervisors. WWW document. Available at: https://www.ibm.com/cloud/learn/hypervisors [Accessed 16 October 2021]

IBM Cloud Education. 2019. IaaS (Infrastructure-as-a-Service). WWW document. Available at: https://www.ibm.com/cloud/learn/iaas [Accessed 9 November 2021]

IBM Cloud Education. 2021. PaaS (Platform-as-a-Service). WWW document. Available at: https://www.ibm.com/cloud/learn/paas [Accessed 9 November 2021]

IBM Cloud Education. 2021. Private Cloud. WWW document. Available at: https://www.ibm.com/cloud/learn/introduction-to-private-cloud [Accessed 9 November 2021]

Innocent, A. A. T. Cloud Infrastructure Service Management – A Review. no date. PDF document. Available at: https://arxiv.org/ftp/arxiv/papers/1206/1206.6016.pdf [Accessed 10 November 2021]

Mell, P. Grance, T. 2011. The NIST Definition of Cloud Computing - Recommendations of the National Institute of Standards and Technology. PDF

document. Available at: http://faculty.winthrop.edu/domanm/csci411/Handouts/NIST.pdf [Accessed 10 November 2021]

Morsy, M. A. Grundy, J. Müller, I. no date. An Analysis of the Cloud Computing Security Problem. PDF document. Available at: https://arxiv.org/ftp/arxiv/papers/1609/1609.01107.pdf [Accessed 12 November 2021]

O'Doherty, P. 2012. VMware View 5: Building a Successful Virtual Desktop. Ebook. VMware Press. Available at: https://learning.oreilly.com/ [Accessed 25 September 2021]

Portnoy, M. 2012. Virtualization Essentials. Ebook. Sybex. Available at: https://learning.oreilly.com/ [Accessed 7 November 2021]

Raghavendran, CH. V. Satish, G. N. Varma, P. S. Moses, G. J. 2016. A Study on Cloud Computing Services. PDF document. Available at: https://www.ijert.org/research/a-study-on-cloud-computing-services-IJERTCONV4IS34014.pdf [Accessed 12 November 2021]

VMware. No date. Virtual Desktop Infrastructure Security. WWW document. Available at: https://www.vmware.com/topics/glossary/content/virtual-desktop-infrastructure-security [Accessed 9 November 2021]

VMware. No date. Virtual Desktops. WWW document. Available at: https://www.vmware.com/topics/glossary/content/virtual-desktops [Accessed 8 November 2021]

Vmware. No date. Virtual Machine. WWW document. Available at: https://www.vmware.com/topics/glossary/content/virtual-machine [Accessed 2 October 2021]

Von Oven, P. Coombs, B. 2019. Mastering VMware Horizon 7.8 - Third Edition. Ebook. Packt Publishing. Available at: https://learning.oreilly.com/ [Accessed 8 November 2021]

Walker, S. 2019. 5 Benefits of a Cloud Computing Security Solution. WWW Document. Available at: https://blog.tbconsulting.com/5-benefits-of-a-cloud-computing-security-solution [Accessed 12 November 2021]