Opinnäytetyö (AMK) | Thesis (UAS)

Degree Program in Information Technology

2012

Joonas Niemi

Empowering IT solutions with server virtualization

Bachelor's Thesis | ABSTRACT

TURKU UNIVERSITY OF APPLIED SCIENCES

Degree Program in Information Technology

5.12.2012 | 56 pages

Instructor: Patric Granholm

Joonas Niemi

EMPOWERING IT SOLUTIONS WITH SERVER VIRTUALIZATION

This thesis explains the migration of traditional server environments to a virtualized server platform. The thesis describes different areas of virtualization, virtualization technologies and explains why virtualization is chosen in corporations. The thesis was carried out as a commission to Erasmus Student Network (ESN), whose server was designed to be migrated to multiple virtual servers with open source platform Xen. Xen is part of the Citrix family and is the creator of paravirtualization technology.

Virtualization is not a new technology. It was first introduced in 1960s by IBM. Virtualization is a technology where logical operations are separated from the physical environment. Virtualization gives the possibility to run several virtual servers, application or complete systems from a single hardware. In virtualization, hardware utilization resources are used more efficiently, compared to one computer-one operating system model. Virtualized operating systems are separated entities, therefore, the data is isolated from each other, making them more secure and easier to administrate.

ESN is a Europeanwide student organization supporting international students in 37 countries and over 400 higher educational institutes. ESN offers several IT services, for example, web pages, online user database and electronic sign up system. All the services are located in one physical server, which was sliced in to five virtual servers. The practical work was carried ou under the supervision of ESN IT Committee.

KEYWORDS: virtualization, Erasmus Student Network, server virtualization, virtualization statistics, XEN

Opinnäytetyö (AMK) | TIIVISTELMÄ

TURUN AMMATTIKORKEAKOULU

Degree Program in Information Technology

5.12.2012 | 56 sivua

Ohjaaja: Patric Granholm

Joonas Niemi

TEHOKUUTTA TIETORATKAISUIHIN PALVELINVIRTUALISOINNILLA

Tämä opinnäytetyö kertoo perinteisten palvelinympäristöjen siirtämisestä virtuaalipalvelin-ympäristöön. Työssä perehdytään virtualisoinnin eri osa-alueisiin, tutustutaan erilaisiin virtualisointitekniikoihin sekä tutustutaan virtuaalisovelluksia toteuttaviin yrityksiin. Työ on tehty Erasmus Student Networkin toimeksiantona, jonka palvelinympäristö virtualisoitiin avoimen lähdekoodin XEN-ohjelmistolla, joka kuuluu osaksi Citrixin virtualisointiperhettä.

Virtualisointi on tekniikkana on ensimmäisen kerran esitelty jo 1960-luvulla. Virtualisoinnilla tarkoitetaan tekniikkaa, jossa loogiset operaatiot eriytetään fyysistä laitteista. Tämä mahdollistaa monen virtuaalisen käyttöjärjestelmän, sovelluksen tai kokonaisen tietokoneen ajamista yhdeltä fyysiseltä palvelimelta. Virtualisointi käyttää tietokoneen resurssit paremmin, mitä yksi käyttöjärjestelmä, yksi tietokone-malli pystyy. Virtualisoidut käyttöjärjestelmät ovat eriytetty toisistaan, jolloin jokainen tietokone toimii itsenäisesti, eivätkä pääse käyttämään toisen koneen tietoja. Tämä parantaa tietoturvallisuutta huomattavasti, sekä mahdollistaa nopeamman viankorjauksen ja parantaa koneiden ylläpitoa.

Erasmus Student Network (ESN) on Euroopan-laajuinen opiskelijaverkosto, joka palvelee kansainvälisiä opiskelijoita 36:ssa maassa ja yli 400:ssa korkeakoulussa. ESN tarjoaa useita tietoteknisiä palveluita, kuten internet-sivuja, web-käyttäjätietokannan sekä sähköisen ilmoittautumisjärjestelmän. Kaikki palvelut sijaitsevat yhdellä palvelimella, joka pilkottiin viideksi virtuaaliseksi palvelimeksi. Työ tehtiin ESN:n IT-komitean alaisuudessa ja toteutettiin veloituksetta.

ASIASANAT: virtualisointi, Erasmus Student Network, palvelinvirtualisointi, virtualisointi-tilastot, XEN

ACKNOWLEDGMENTS

This thesis would not have been possible without the support of my teachers, friends and family.

I would like to thank my supervisor Mr. Patric Granholm for understanding and guidance, my long-time tutoring teacher Mrs. Kalliopi Skarli for the support inside and outside the classroom, and all the lecturers at the ICT-building.

I am grateful for the support from my father who I can always count on and my mother for pushing me forward and motivating me to graduate.

Last, but not the least I would not be here without the support of my friends.
I am in deep gratitude to all of my friends, especially Emma and Jussi, who empowered me to never give up.

TABLE OF CONTENTS

TABLE OF FIGURES

TABLE OF PICTURES

APPENDIX LIST

ABBREVIATIONS

API	Application Programming Interface
CMS	Content Management System
CPU	Central Processing Unit
GB	Gigabyte
DoS	Denial-of-Service attack
IO	Input-Output operation
IP	Internet Protocol
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
LVM	Logical Volume Manager
Mbps	Megabits per second
NGO	Non-Governmental Organization
OS	Operating System
RAID	Redundant Array of Independent Disks
RAM	Random Access Memory
RISC	Reduced Instruction Set Computer
SAN	Storage Area Network
SSH	Secure Shell
SSL	Secure Sockets Layer
SYN	Synchronizing flag operation
TB	Terabyte
TCP	Transmission Control Protocol
VLAN	Virtual Local Area Network
VM	Virtual Machine
VoIP	Voice over Internet Protocol
VPN	Virtual Private Network
VRF	Virtual Routing and Forwarding
VSA	Virtual Security Appliance

1. Introduction

Server virtualization has rapidly grown in the past 8 years (Gartner, 2011). According to research, almost 85% of the IT companies are using or are planning to implement virtual server environment in the near future (Forrester, 2012. Today 52% of the environments are running on virtual platforms and in two years this number is expected to grow up to 75%. (Forrester 2012).

This thesis will explain server virtualization technology and how it benefits IT companies. In addition, server virtualization is argued to be more cost-effective and resource-efficient than normal server environments.

This thesis was carried out for an NGO, Erasmus Student Network. This voluntary-based organization is running multiple internet services, such as multiple websites, email and mailing list server and development environments. All these are located in one physical server where all the data is accessible to anyone who can access it. As ESN is a voluntary organization, the number of people involved is significant. This is causes security and documentation difficulties to the organization. Since last spring, there has been discussion to migrate the server to virtual environment. This would enhance the structure and the security to the administration of the platform. The task was to design and implement the virtual environment in collaboration with the ESN IT committee.

This thesis focuses on virtualization as a strategy for IT corporations. Server virtualization has a long history, but in eight years, it has started to be anincreasingly important factor in IT solutions (Gartner 2011). This thesis concludes from large amount of research where the following questions are answered:

Why corporations choose virtual server environments? How can virtual server environments be more cost effective and efficient? What are the security aspects of virtual environments? What is needed to design and migrate virtual servers to an existing network?

The thesis is divided into three parts. Firstly, the theoretical part will explain what virtualization is, how it works and which technologies are used. This thesis will focus on the Citrix Xen virtualization platform which has been used for the practical work. Citrix Xen was chosen because it is open source software which is the most suitable platform for a low budget implementation.

Secondly, the research part discusses where the cost benefits of virtualization for IT corporations. Virtualization has been rapidly growing since 2005 and will peak in 2018 (Gartner 2011). Several statistics about the growth and usage of virtual environments in IT companies are presented.

There are numerous studies demonstrating how virtualization is a more cost-effective and efficient solution than a normal physical server environment. In addition, the benefits and downsides of server virtualization are described. As the virtualization evolves rapidly, new technologies are listed.

Thirdly, the thesis describes the practical part about server virtualization in a case example for the Erasmus Student Network (ESN). ESN is a Europeanwide student network which operates in 36 countries and over 400 universities. ESN’s main office in Brussels runs ESN web services, for example, the esn.org website. ESN has many services, though only one physical server, which they would like to virtualize in isolated virtual servers. This part describes the detailed virtual environment, together with the technical implementation.

2. Virtualization overview

2.1. What is virtualization?

Virtualization could be defined as separation of logical operations from their physical environment. To give an easy example on virtualization, we can use virtual memory. Virtual memory is an extension to system memory in a computer, which is used from the hard drive, rather than from the RAM memory.

If the system memory runs out, this virtual extension can be used to keep the system running. (von Hagen, 2008, 1).

The advantages of virtualization have been accepted widely, and only having physical servers is decreasing rapidly. (Forrester, 2012)

The key advantages of virtualization are accessibility and isolation. A virtualized environment can be accessed anywhere with any device accessible to the network. This is the base for cloud computing, for example, Google Apps. Virtualized servers can be isolated based on their workload and purpose. One service can be isolated to one virtual environment which gives a clearer structure and increases security by allowing only the designated technicians to access the server configurations.

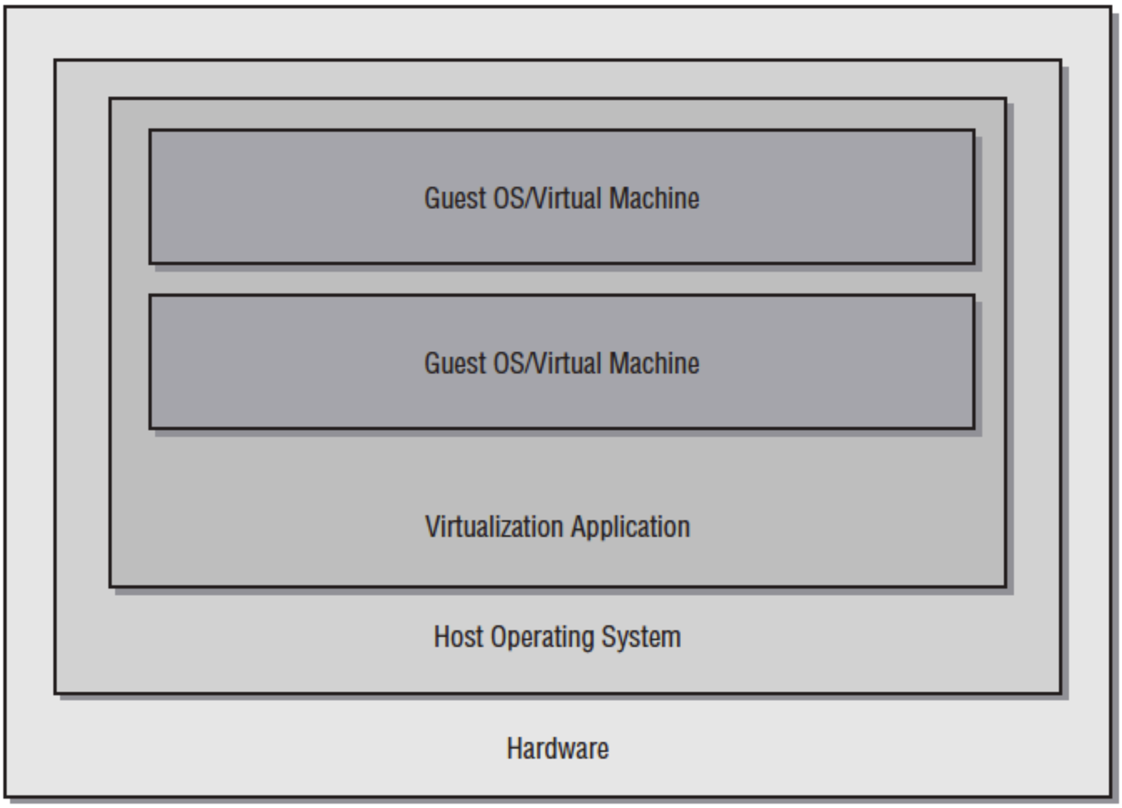

Figure 1. Virtualization overview. (von Hagen, 2008, 10)

There are several other methods of virtualization. This thesis focuses on one of them, server virtualization. Here is a summary of other possible virtualization methods.

Application virtualization is compiling machine independent byte code in to applications which can be executed from any systemn where the needed environment is set up. The best example of this is the Java programming language. (von Hagen, 2008, 2)

Desktop virtualization is a method where a graphical desktop can be displayed from another system. The most known application used is Virtual Network Computing (VNC). Many times computer hobbyists and testing engineers need to access their desktop remotely from another location. Desktop virtualization gives the full access to the desktop. This method is often used in computer help desks where the technician can remotely investigate a problem in a client computer. Although desktop virtualization belongs to this category, it can be argued that this is not real virtualization. When connecting remotely to another computer, an user is looking at the same display which is still running in another location. The changes made remotely, affect the physical computer. (von Hagen, 2008, 2-3)

In network virtualization, isolated separate networks are built on top of physical layers. These networks function independently and are treated as separate networks although they work under the same physical environment. Examples of virtual networks include VLAN which isolates ports of switches or routers to operate in one certain network (Cisco 2012). The virtualized network used in Linux systems is called bridged networking (Debian 2012), shown in Figure 2. One physical NIC is bridging IP-addresses to the virtual servers which all have their unique MAC-address and virtual NIC. This method is used when configuring the servers in this thesis.

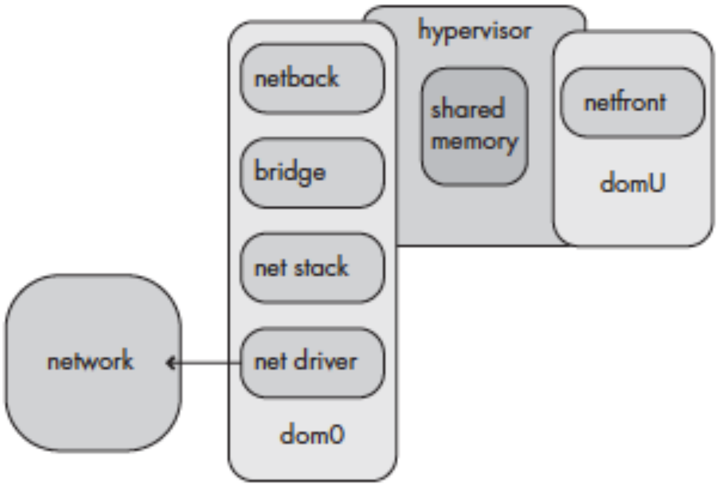

Figure 2. Network virtualization (Takemura & Crawford 2008, 11)

Server or machine virtualization is the most common method of virtualization (von Hagen 2008, 3). In server virtualization, the whole machine runs independently in a virtual system. It has a separate operating system as well as hardware and networking capabilities. The virtual server is displayed to users as a separate physical server although there can be several similar systems in one physical environment. Server virtualization is the focus of this thesis and the practical work is done with server virtualization.

2.2. Virtualization techniques

Server virtualization techniques can be categorized in four types: Full, para, OS and hardware virtualization.

Operating system virtualization

In OS virtualization, the whole operating system is virtualized over another. The OS providing the virtualization is referred to as a host OS which has a virtualization platform. The virtualized OS operates on top of the host OS with the offered hardware, emulated to fit the needs of the guest OS. Specifically, hardware is not virtualized, but only offers the needed services (for example, using the CD-ROM and USB ports) to the guest OS. Common usages of this approach are, for example when a needed software runs in only one OS. For example, many Macintosh users virtualize Windows OS for games or small software that are only made for Windows (Takemura & Crawford 2008, 4). Known products for OS virtualization are Parallels Workstation, VMWare Workstation & Fusion and VirtualBox.

Hypervisor based virtualization

Other virtualization methods, excluding OS virtualization, use hypervisor-technology. A hypervisor is a virtual platform manager that operates above the physical layer and below an operating system (Microsoft, 2012). Hypervisors can be classified into three types: level-1, level-2 and hybrid. Level-2 has the hypervisor located under the host OS. This type is referred to as process virtualization.

Level-1 and hybrid types deal with hardware virtualization; in level-1 the hypervisor is located above the hardware and runs directly on that hardware. The hybrid level emulates the whole hardware, where the virtual systems then lay above the host OS and hypervisor. (Cerling *et al*., 2010, 8-9) Xen uses level-1 hypervisor with paravirtualization.

Full virtualization uses level-1 hypervisor where the unmodified host OS lays on top of the hypervisor. In this approach, the guest OS does not need to be changed and the hardware architecture can differ from the physical one (Takemura and Crawford, 2008, 3, VMware 2012). Full virtualization uses a significant amount of memory and processor which backfires in lack of speed (Takemura & Crawford 2008, 3). VMWare is the most known vendor of full virtualization with their ESX series. (Lowe 2011, 3-4)

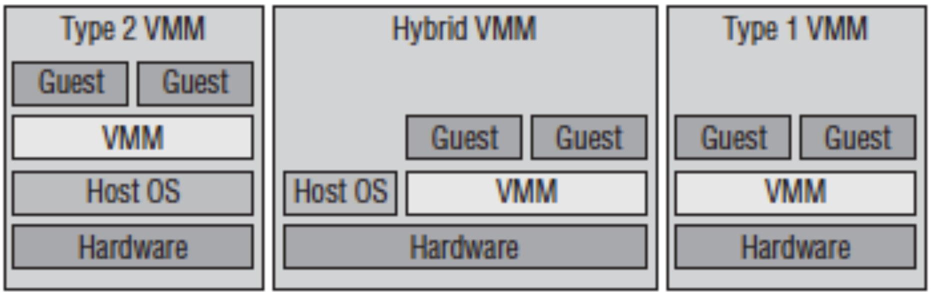

Figure 3. Virtualization levels. (Cerling *et al.*, 2010, 8)

Paravirtualization

Paravirtualization is an approach introduced by Xen in 2005 (Takemura and Crawford 2008, 4), which is now adapted by other vendors as well (VMware, 2012, B). Para comes from Greek origin meaning “alongside” or “besides” (Thayer and Smith, 1999.) In paravirtualization, the hypervisor runs directly from the hardware being very small and compact by size. Paravirtualization does not have a clear division between guest and host OS, instead it uses a trusted guest host referred to as Domain 0 (Dom0) (Takemura and Crawford 2008, 4). Dom0 is the management domain, the virtual machine manager (VMM), which builds the additional virtual machines and controls them. This is described in Figure 4.

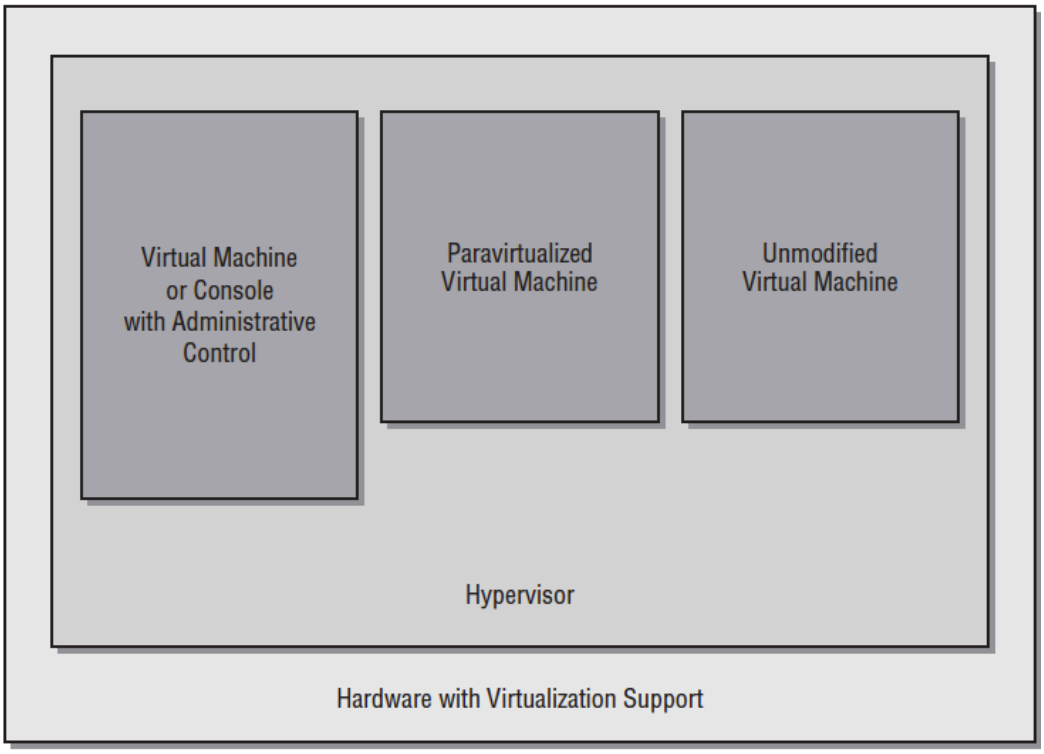

Figure 4. Paravirtualization technique. (von Hagen, 2008, 14)

2.3. Citrix Xen

Xen Hypervisor was originally invented by two Cambrige University students, Keir Fraser and Ian Pratt for their Xenoserver project in the late 1990s. Xen 1.0 was released in 2004 when also XenSource Inc. was founded. The corporate strategy has always been to develop open source products. In 2005, Xen Hypervisors started to spread when corporations like Red Hat, Novell and Sun added Xen as an option of virtualization systems. The same year Xen 3.0 was released. Next year, 2006, Microsoft and VMware adopted the paravirtualization technology, originally introduced in Xen. The milestone in Xen history occurred in 2007 when Citrix System acquired XenServer as their server virtualization solution. When cloud computing started to raise heads, xen.org released Xen Cloud Platform. In the last two years, Xen has been developing Linux kernel. In 2011 the Linux 2.6.37 kernel version was the first kernel to work as Dom0 with “out-of-the-box” principle (Citrix, 2012).

3. Technical overview

3.1. Xen virtualization

Xen uses the technology paravirtualization which was originally invented by the creators of Xen. The basis of this approach is the Domain 0, the privileged kernel and the linux system which operates the Xen daemon (xend), which communicates with the hypervisor. As Dom0 is the management OS, it is responsible for creating the virtual machines, referred as Domain U (DomU), controlling them (start, stop, suspend, save) and deleting them when necessary. DomUs are the actual virtual machines which are all independent of each other. In a normal paravirtualized environment, all the domains use the same hardware setup which cannot differ on various domains. Some of the hardware support virtualization which makes it possible to use unmodified guest system. Unmodified guest system means using OS systems that are not modified for virtual machine usage, for example, Microsoft Windows or standard Linux distributions as a host system. Figure 5 describes the Xen approach.

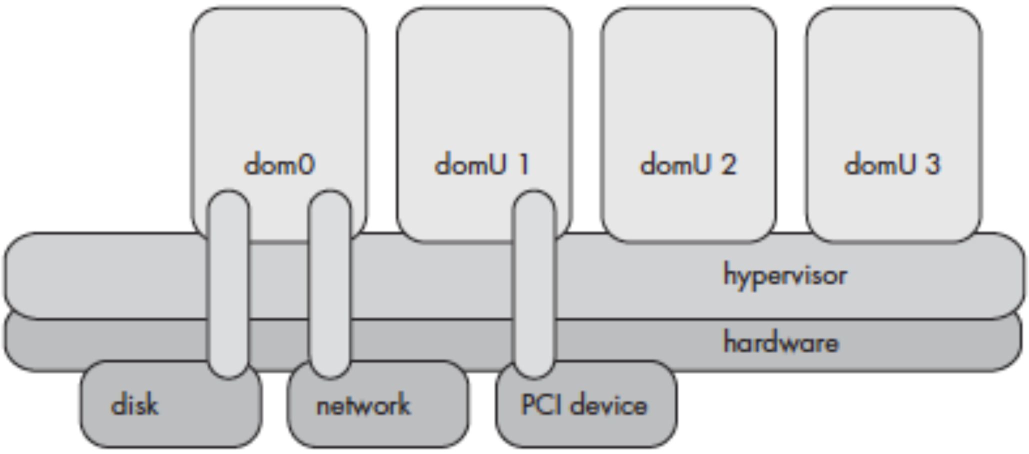

Figure 5. Xen virtualization technique (Takamura & Crawford 2008, 5)

The hypervisors is the heart and soul of the Xen system. It controls that every virtual machine receives the necessary portion of the hardware resources to operate. The hypervisor does not emulate the actual physical hardware, it rather exports a simplified version. For example, the hardware is not any specific model, but it is just generic storage devices or network cards. This makes it easier to have support for those devices, when the guest system does not need specific drivers. (von Hagen, 2008, 31-32, 36)

3.2. Protection rings

In any computing system, there is a need for different authorization levels depending on which actors are used. A virtualization environment has guest, host and user applications. They are all required to operate in different authorization levels. In the used x386 architecture, there are four authorization levels, called protection rings, ranging from 0 which is the highest to 3 which is the lowest. In a virtualized system, the hypervisor operates in the highest level 0, the guest system in level 1 and user applications in level 3. The hypervisor in level 0 has the full control of the hardware, giving the access to guest system according to the system's needs.

The request for higher privileges is made through hypercalls which are similar to system calls made by the user applications to the operating systems (Linux, 2012). If the hardware support is used, the protection rings become less significant because hardware support adds root and non-root modes. (Matthews *et al.*, 2008, 68)

3.3. Scheduling

Many people think that computers are capable of performing multiple tasks at the same time, in other words, multitasking. However, computers do not operate in such a manner. Every sequence is done step by step, furthermore, this is done extremely fast. Scheduling is the method used for prioritizing the tasks and handling the queue (Takemura & Crawford 2008, 6-7). Scheduling is an important part of virtualization as the hypervisor has to run multiple virtual machines in parallel. In virtualization the schedulers determine which VM receives CPU and other system resources and in which order (von Hagen, 2008, 33).

Xen offers two main schedulers: the credit scheduler and the earliest deadline first (sEDF) scheduler. The latter one is planned to be removed in the next update of Xen Hypervisor (Matthews *et al.*, 2010, 489). The credit scheduler'stask is to make sure that every virtual CPU receives enough resources (credits) which are calculated through two parameters: weight and cap. (Lee *et al.*, 2010, 1.)

Weight is the priority of the domain. Its value is anything between 1 and 65535, while the default value is 256. Cap is the maximum percentage of CPU usage, which is displayed in percentages. (von Hagen, 2008, 33-34) Along with the hypervisor scheduler, the guest domains use operating system schedulers, called IO schedulers.

3.4. Virtualization vendors

Besides XEN there is few other vendors who control the virtualization markets:

VMware

VMWare is the leading virtualization corporation in the world (VMware, 2012). They have been specialized in x86 architecture virtualization since 1998. VMware offers both hosted and hypervisor-based architectures. VMWare Workstation, VMware Server and VMware Fusion for Macintosh are their key products in hosted architectures, while VMware ESX is the hypervisor of VMware family. The flagship product is VMware Sphere which gathers all the needed tools for overall virtualization and cloud computing (Lowe, 2011, 1-2).

VMware uses the full virtualization technique in their products and is dominating the virtualization markets. VMware Sphere is used in over 70% of the IT corporations (InformationWeek, 2011).

Microsoft

Unsurprisingly, Microsoft has a foothold in virtualization markets. The 1988-founded Connectix was the early predecessor of Microsoft virtualization. Connectix developed Virtual PC which gave the possibility to run Macintosh based products in Microsoft environments. Microsoft bought Connectix in 2003 and started to enhance Virtual PC products and released Virtual Server 2005. Recently, Microsoft released their own hypervisor Hyper-V. The main Microsoft virtualization products are Virtual Server 2005, Hyper-V server 2008, and Windows Server 2008 Hyper-V (Cerling *et al*., 2010, 5-8). Microsoft comes as clear runner up as VMware is leading the pack in virtualization markets. (InformationWeek, 2011).

Others

There are many other vendors in virtualization, though VMware, Microsoft and Citrix dominates the field. Here is few other known manufacturers:

VirtualBox is Oracle's open source virtualization software mainly used in Windows environment to run Linux distributions. VirtualBox has easy to use graphical interface and it can be seen used for teaching purposes, for example learning another operating system while having the host system the familiar one. (Romero, 2010, 8-14) VirtualBox is mostly used by computer enthusiastic's, though Oracle has it slice from the virtualization markets. (InformationWeek, 2011)

Parallels is a known name in the virtualization field, however, its main focus is desktop and OS-level virtualization. Founded in 1999, Parallels' main products are Parallels Desktop and Virtuozzo. The Parallels desktop is the number one in desktop virtualization, mainly for their Macintosh users who often need to have a virtual Windows desktop. Parallels has server virtualization services even though they are very small vendors in that field. (Parallels, 2012)

There are many options for different types of virtualization. VMware is dominating the world with it enterprise products, while Citrix and Oracle are used by hobbyists, small companies and very field-specified corporations, for example, in the telecommunications field.

4. Virtualization in IT corporations

4.1. Introduction to virtualization in IT

The markets are evolving now that cloud computing is increasing and mobile solutions are becoming more popular.

VMware can be identified as the synonym for virtualization. However, virtualization has a longer history than 13 years of VMware's existence. The real inventor of virtualization is the computing pioneer IBM. The IBM CP-40 operating system with virtualization capabilities was invented in 1967. The CP7 CM systems enable separated virtual machines which could share memory and could be isolated from each other. (IBM, 2012)

After the late 1960s, it was very quiet in the virtualization field. It took almost 20 years before the development of virtualization took the next huge step after being invented.

In the early 1990s, the Disco Project in Standford University started to research virtualization architectures (Bugnion, Devine and Rosenblum, 1997). The project used the IRIX operating system with MIPS processors instead of IBM hardware which was not designed for virtualization. The tests were successful and the team faced an even greater challenge, modifying x86 architecture. This was the key action which founded the VMware corporation. In 2005 processor manufacturers AMD and Intel added additional hardware support to enable higher performance in full virtualization(Matthews *et al.,* 2008, 29-30). After AMD and Intel started to support virtualization in their products, virtualization was here to stay. Virtualization was rapid increased in IT environments.

This chapter explains with the help of statistics how much virtualization is used and in which fields. Several pieces of research were used to evaluate the development of virtualization. Additionally, the key reasons for IT corporations to migrate to virtualization are explained. The hot debate is the costs, followed by arguments why this approach is more cost effective than normal physical hardware solutions. It is important to list the benefits, as well as the downsides.

4.2. Usage of virtualization in corporations

When the commercialization of virtualization started in 2000, the use of virtualized environments was very slow (Gartner, 2011, A). It took until 2005 for virtualization to be more widely used in corporations. There are several statistics to show different aspects of virtualization usage. Gartner, Forrester and IDC are the main research companies in this field. However, their research is commercial as well and significantly costly. During the research, statistics from 2007 until 2012 were gathered. Several pieces of research were carried out and clear images how much virtualization is used can now be introduced.

Based on the statistics, four pieces of research were conducted: the InformationWeek Virtualization management report (InformationWeek, 2011), the Governmental Virtualization report 2010 (CDW-G, 2010), the Virtualization industry quarterly survey (Veeam Software, 2011) and the CISO's guide to virtualization security (Forrester, 2012). In addition, in Finland two surveys has been carried out (meXt, 2009), one in 2007 and again 2009.

The state of virtualization has been measured in three main areas: first, in the virtualization rate which describes the percentage of corporations using virtualization; Second, in the expansion rate where corporations estimate in what time and how much they are planning to expand virtualization; and Last and most importantly, in the penetration rate which describes the overall percentage of virtualization in all environments. When analyzing statistics from the IT field, it is necessary to remember the rapid evolvement of the industry as one year in the IT field is relatively long time. The statistics used here are from

2009, 2010 and 2011. Surprisingly the numbers are at the same level although the data is from three consecutive years.

When it comes to the state of server virtualization, it is highly used and increasing even more. According to this four reports, the virtualization rate is over 78%, while in Finland 2009 it was 68% (InformationWeek, 2011, meXt 2009, Veeam, 2011, CDW-G, 2010). This shows that in 2010-2011 almost four-fifths have been using virtualization solutions.

The expansion rate tells an even better story about which way virtualization is developing. The expansion rate varies from 60 to 80% depending on company sizes (InformationWeek, 2011,Veeam, 2011, CDW-G, 2010). In Finland, such data was not available. These numbers seems to be enormous, if already 78% of the corporations are using virtualization. That is true, but the answer can be found from the penetration rate of virtualization.

Based on the V-index and Gartner penetration rate statistics, the overall virtualization is only 38% (Veeam, 2011, Gartner, 2011). In Finland, in 2009 this number was 31% (meXt, 2009). At first look, the number seems relatively low, on the other hand, there is a couple of easy explanations: virtualization is highly used in corporations whereas physical servers are used, for example, as Active Directory domains, storage servers and desktop usage. Virtualized servers are the most virtualized environment, though alongside it still has many physical servers.

These statistics show the dominance of VMware. Overall, VMware provides almost 70% of the server virtualization, while Microsoft has 15-20% and Citrix 10-15% slice of the market. (InformationWeek, 2011, Veeam, 2011, meXt, 2009)

VMware was the first commercial product in virtualization markets which could explain the situation. It is possible that Citrix Xen and other open source vendors will increase their market share the same way Mozilla did in the web browser markets.

4.3. Reasons

Rather than starting the sales speech for VMware: why one should choose virtualization, another approach is taken. Research, study papers and surveys were researched (Avande, 2011, Gabriel Consulting Group, 2011, Forrester, 2011, Veeam 2011, Gartner, 2011) to figure out why companies have chosen virtualization. The reasons are quite the same as what the IT professional would say.

Before stating the coolness of virtualization, the problems in IT solutions need to be understood. Most of the companies have servers, desktops and networks. Not all of these companies are professionals in IT solutions. Easily these companies combine many solutions; moreover, they have duplicate and too complex system to administrate efficiently (Avanade, 2011). Another key problem is the usage of the server resources. According to the same study, approximately 20% of the server resources are utilized (Gartner 2011b). This increases energy costs and the overall IT budget.

This leads us to the most important factor, costs. IT hardware costs a significantly large sum of money (Gartner, 2011b). These costs consist of many parts, for example, data centers, networking and client computing. As stated before, the server utilization level is very low, therefore, funds are wasted. According to the Avande study, over $5000 can be saved per virtualized OS. Gartner recommends virtualizing 50-66% of the workload with a noticeable cost savings. (Gartner, 2011b)

From these problems we can already identify why corporations are looking for the alternative solution. According to the surveys, companies see virtualization as saving money, ease to manage and giving higher utilization level (GCC, 2011). The survey made in Finland clearly indicates that cost-cutting and the easier management as the key factors. (meXt, 2009.)

Surprisingly, security was rarely mentioned even though the IT professionals usually state how virtualization enhances security. On the other hand, after corporations have been migrating to virtualization, the worry for the security has increased (Kaspersky, 2012).

4.4. Security

Although security can be seen as important factor to migrate to virtualization, it might often worries the companies. Virtualized environments do not differ from the physical ones in terms of security. Companies cannot rely on having only authorized access as a security method. A recent example of a horrifying security breach happened in February 2011, when a former IT technician of Shionogi pharmaceutical company accessed their servers and deleted 88 virtual servers (Department of Justice, 2011). Another case was researched in Simon Fraser University about the effect of Denial-of-Service attack (Carnegie Mellon University, 1999) in virtualized servers. The research proves that TCP SYN flood (Internet Engineering Task Force, 2007) DoS attack consumes high CPU and memory space in virtualized servers compared to non-virtualized servers. This makes them more vulnerable to the Denial-of-Service attacks (Shea and Liu, 2012). This makes them more vulnerable to the Denial of Service attacks. (Shea and Liu, 2012).

Social engineering or “inside attacks” are common in virtual environments. A former administrator can access the Virtual Machine Manager and, for example, delete the machines. Often intra-virtual machine traffic is not monitored enough. The security professionals do not monitor the intra-virtual machine traffic, as it does not travel through the physical network device. End-point security is another issue security administrators forget. They rely on the host security, while jeopardizing securing the virtual machines (Forrester, 2011).

There are two approaches to the virtual machine security: agent and agentless security. In agent endpoint security, each VM has its own security module or agent attached to the machine. This is the more traditional method. However, it consumes resources as it needs to be installed in every VM separately. This can cause overwhelming workload for companies who have, for example, 20-40

VMs. Although having separate security modules in each VM consumes significant amount of CPU resources in each, this approach is cost effective in a small environment, but is not recommended for large IT solutions.

The agentless approach brings another module, a virtual security appliance to the environment. Here the security to all VMs is handled with one module, which is installed directly to the hypervisor (Forrester, 2011). This appliance can automatically detect the existing and newly installed virtual machines (Kaspersky, 2012). Popular vendors of VSA are McAfee, Kaspersky, Stonesoft and TrendMicro.

Virtual network security

Virtual machine security is vital; however, network layer security plays an important role with the virtual machines. As the virtual machines have virtual NICs or the traffic goes through bridged interfaces, it is important to isolate the traffic from each NIC. To do this, the network can be divided in three areas: access control, path isolation and service edges. Access control in network virtualization is similar to that in any networking system. The goal is to authenticate the users, authorize the access based on access lists and virtual LANs and deny unauthorized access.

Path isolation is the key factor in network virtualization. In this technique independent traffic paths are created for each networking devices by creating VPN tunnels. Each logical traffic travels through their own VRF path.

The service edge area provides the access from the edge of the network inside and outside. It is responsible for controlling the routing of the VPNs to the correct place and restrict unwanted traffic. (Cisco 2012)

4.5. Benefits

Firstly as mentioned in Section 4.3, physical server environments use only up to 20% of the hardware resource. The same hardware consuming one physical server and all the services inside it, could be divided to multiple virtual servers which could employ the resources more efficiently. Today, when manufacturers are introducing multicore processors, virtual servers can consume individual cores to run the virtual machine. (von Hagen 2008, 15)

Secondly, system administration is simplified. When having virtualized servers, the same or increased amount of systems are in use, however, they can be centrally managed. Hardware monitoring is the most important as reaction to failures must occur rapidly. In a virtual system, it is possible to have system monitoring tools, in addition to monitoring all physical and virtual hardware. These monitors are able to schedule backups, cloning, starting and stopping the machines in a case of failure. These centralized systems would not be possible in a physical environment. (von Hagen 2008, 16)

Thirdly, uptime and failure recovery are increased. This is due to the isolation of the virtual servers. When every service is isolated in its own logical unit, they can be recovered or migrated to another virtual machine instantly. The clones or virtual machine images have a capacity to be stored to a network drive where they are adapted to another physical server in a case of hardware failure. (von Hagen 2008, 18)

Fourthly, reduction in hardware and IT infrastructure costs. When existing hardware is effectively in use, there is no need to invest in new hardware at the time of expansion. It is possible to have multiple servers inside the physical hardware. The resources are fully used, external memory and storage can update the system, rather than purchasing a completely new physical server (von Hagen, 2008, 16). Hardware costs are combined with energy savings, as less energy is required. In addition to the servers, fewer resources for cooling, data center space, lightning and power sources are necessary. These reductions incur virtualized OS cost savings of up to 3% of total IT costs compared to a non-virtualized server. If a company has five virtual servers instead of five physical ones, the saving is already 15% (Avanade, 2012). Investment is required to start the virtualized system; yet, virtualization returns the investment averagely in 14 months. This is due to using less hardware and labor requirement. It is estimated that virtualization can save up to 35% of the overall IT costs. (Cassidy and Cassidy, 2009, 133-137.)

4.6. Downsides

The most known and often stated problem is Single Point-of-Failure. Now, the physical server resources are completely in use, by having multiple virtual servers in one piece of hardware. This creates problems if this single physical machine faces system failure, in particular, all the virtual systems will go down. (von Hagen, 2008, 21.)

Second storage availability, server functionality, and performance issues cause great concern for administrators. When local storage is used with a virtualized system, it is highly inflexible for universal use. (von Hagen, 2008, 22). The storage performance problem is a cause of slow hard disk hardware. In a virtual system, the local hard disk is working with multiple workloads, while creating significant bottlenecks to the system performance. (Collect, 2011, 20)

Thirdly, although the back-up system in virtualization was mentioned as a benefit, some might consider it as a weakness. Traditional back-up methods are not effective in a virtualized system, as the servers themselves are individual disk images. Therefore, the storage inside is not real storage, which makes it complicated to back up. The whole virtual machine image and the data inside the machine have to be backed up separately. These are usually backed up in a network storage although it can be challenging to find a proper one. (Collect, 2011, 20)

4.7. Solutions

Single point of failure is the most vital problem to solve. System administrators should keep spare hardware available all times, therefore, no extra time will be used to receiving spare parts from the manufacturers. Virtual machines should have duplicates in a different physical system, if a system failure occurs. The virtual system monitor tools have the capability to achieve this efficiently. Logging and alert software are needed to react immediately to possible failures. (von Hagen, 2008, 22) It has to be stated that the hardware and single point of failure are equally disadvantages of a physical environment, too. On the other hand, virtualized systems have better solutions to recover from these incidents with the solutions mentioned above.

When it comes to storage availability and server functionality, the design element is important. The system needs to be designed in way where virtual machines can work independently without worrying about local storages. It is essential to carefully plan which kind of actions the virtual machine carries and how many resources it would take. If one server has an overly amount of different services, multiple users are created. This can create problems to the system performance (von Hagen, 2008, 22-23). Transparency of the storage is important for the system administrators to see what is happening inside the media. Many storage vendors are providing solutions to monitor the storage, for example, NetApp and Virtual Instruments (Collect, 2011, 21)

There are powerful solutions for data back-ups, nevertheless, system administrators are not familiar with them. Traditional tape back-ups do not work in virtualization. SAN and RAID systems are more redundant for virtualized systems. This approach avoids making overnight back-ups which would slow the server down. (Collect, 2011, 20-22)

When setting up the virtualized environment, the system administrator needs to be aware of the bottlenecks in addition to designing a system that tackles these problems (Collect 2011, 20-22). In the environment designed for ESN, not all of these considerations were possible to carry out due to the low budget and lack of resources.

5. Physical server environment migration to virtual platform

5.1. Erasmus Student Network

Erasmus Student Network (ESN) is a European-wide disciplinary student organization. It is one of the largest European youth organizations by having local organizations in 414 higher educational institutions in 36 countries. ESN covers over 150 000 students by offering services and support for international students.

ESN works on three different levels: local, national and international. In the local level, sections, clubs or registered organizations operate. They offer hands-on support for incoming international students, as well as outgoing and returning students. Mainly, ESN sections organize events, activities and trips, for example, trips to other countries and cities, language and movie evenings or sport and social activities. The main principle is “from students to students”, therefore, the organization is easily approachable. The local level is the basis of the work of ESN.

The national level consists of at least two representatives, National Representative (NR) and vice-National Representative (VNR). In many countries, the national level is operated by the National Board (NB), usually registered national organizations. In Finland this is Erasmus Student Network Finland ry. National Boards are executive bodies which work for the interests of their sections. They can, for example, negotiate national contracts with companies to receive discounts or sponsorships, organize national and international events, most importantly, supervise the interest of their sections in national and international contexts. In the national context, National Boards meet with the country officials, such as National Agencies, governmental offices or higher educational officers. In the international level, the NR liaise with the colleagues of other ESN countries and European Parliamentarians and other interest groups. The highest decision organ of national level is called the National Platform (NP) which is composed of all ESN sections within the country. The NP can be described as the annual or biannual meeting. In Finland the sections are members of ESN Finland and they use the decision-making power of the organization.

The international level is the largest and most complicated one in ESN. The international level consists of many organs, most importantly, the Annual General Meeting (AGM). AGM is the highest decision-making body of the organization, where all 414 sections have a vote. The delegates of the AGM decides on budget, action plans, International Board members, statues and standing orders.

The Council of National Representatives (CNR), which conducted from the National Representatives of all the countries operates under the AGM. The CNR, is a working body of ESN, represents the opinions of the countries, makes proposals for projects, and updates the network about current issues. ESN has working committees supporting the executive work. For example, this thesis is made in cooperation with the IT committee of ESN.

The executive body of ESN International is the International Board (IB). The five member IB is working full time in Brussels for a year. The IB is responsible for representing the interests of the whole network to the stakeholders, for example, the European Union and the Council of Europe. The IB is responsible for setting up international projects, such as study surveys or higher education policies. The most known ESN projects today are PRIME, which focuses on problems on recognition of studies, SocialErasmus, where local students are encouraged to integrate with the local community and ESNSurvey, an annual survey to international students about a certain topic. (Erasmus Student Network 2012)

ESN operates completely by volunteers and by the support of grants. Without the all the volunteers, ESN would not have the resources to operate at such a high level.

5.2. Current IT solutions

ESN offers multiple services to their members and interest groups. The most visible and important one is the esn.org website. The website is the key information channel of the network. Projects, publications, partners and information packages for sections can be found through this platform. ESN uses its own webpage template called the Satellite. This is a Drupal-based CMS for ESN websites. All sections or national ESN organization can download the template and install it to their servers.

Picture 1. esn.org, Satellite website. (ESN, 2012)

ESN has a membership card, the ESNcard. With the card, students can receive discounts and offers in negotiated places, such as restaurants, student cafeterias and ESN organized events. In addition, ESNcard-holders can receive international benefits from New Yorker, Brussels Airlines and Hostelworld, to name a few. The card benefits are promoted through the esncard.org website. The stability of this particular website is highly important, because it is presented to the new possible partners.

The two websites are the most visible IT services that ESN offers. However, ESN has multiple supporting services, such as mailing lists, event registration, project websites and testing platforms. The user database of ESN is called Galaxy. In Galaxy, the user can be identified based on their position in ESN and assigned privileges can be given. All ESN members can sign up to the Galaxy, where the country representative, for instance, the NB president can give access to certain areas. As an illustration, the event registration system is based on Galaxy accounts. A user has to have at least a section member role to be able to register to Annual General Meeting or regional platforms.

Picture 3. ESN Galaxy (ESN, 2012)

In addition, ESN has multiple project websites, for PRIME, SocialErasmus, ExchangeAbility and ESN meetings. Mailing lists are an important part of ESN

IT services. They are used for communication with the CNR, National Boards and local sections. The description of the current systems is described in detail in Appendix 1, 3 and 4.

5.3. Hardware and system

Currently ESN has two servers which are provided by Hetzner Online (www.hetzner.de), a hosting company based in Germany. The servers are EQ6 and virtual server VQ19. The EQ6 is the main server of ESN where everything else than esncard.org is located. The server is rather old and cannot be purchased from Hetzner anymore. The current server has 12GB RAM, 1,5TB hard disk and 10TB for network traffic.

The virtual server VQ19 was purchased for the esncard.org. It was decided to isolate to separate server because it has to be very stable as it is presented to prospective international partners of ESN. The VQ19 is running with 2 GB of RAM and 80 GB Hard disk. It has a total traffic limit of 4 TB per month. Both systems are running Debian Linux, version Lenny (kernel 2.6.26-2-amd64).

5.4. Structure and services

Appendix 3 lists all web services and their purposes. All together ESN has 51 different domains. Many of them are redirections to different sites, usually esn.org. According to the ESN IT committee, the amount of redirections to the main website is due to the removal of old sites. For example, there are three redirections which were used for Galaxy 3 website. The 25th anniversary of Erasmus intended to have a separate website although it was never used.

Appendix 4 displays the file structure of the current server. The two main directories in use are home and data. The home-directory contains a directory for all users who have access to the ESN server. At the time of writing, there are fourteen user directories, two for testing purposes and one for internal backup system. The data-directory contains all data used in ESN servers. This includes production websites, testing and development platforms.

There are several problems in the current server:

All users who need to access a service need full access to the ESN server. This creates security and maintenance risks. Any user can modify anything on the server, can manage the services and install additional software. The changes made will affect the physical server, if configured wrong, causing server failures.

In addition, the server structure is disorderly arranged. It is not documented which service can be found in which directory. The space usage is not efficient. The server uses 307 GB of hard disk space which is not a large amount. However, internal backup space uses 130 GB from it and the whole home-directory takes up 200 GB in total. The data directory uses only 21 GB of space although it contains all the IT services. Appendix 1 explains more thoroughly the structure of the current system.

5.5. Designed solution

Based on the research on the current IT solution, a design document (Appendix 1) was made in order to describe the possible solutions. The design document describes the current situation, recommendations for the new physical hardware and design schemes for the virtual environment.

The new server will be purchased from Hetzner Online. ESN is willing to invest more funds to receive sufficient hardware which would have more resources than the current one. Four options were researched: Dedicated servers EX5, EX6, EX6S and EX10. The specifications of each server were better than the current EQ6 and VQ19. The total cost of current solutions is 924 euros.

After consultation with the ESN IT representatives, server options were narrowed to EX6S and EX10. The proper description of these two is found in Appendix 1. The most significant differences between the options were the processor and the RAM memory. Both have capability over the designed resources and it gives room for expansion. The EX6S has 4-core and EX10 6-core processor. More cores are welcomed specially in virtual environment where individual cores can be used separately in each virtual machine.

The decisive factor for choosing the server was price-quality comparison. The ESX6S would cost 960 euros in the first year and 840 euros in the second year and after. The ESX10 would cost 1248 euros in the first year and 1127 euros in the second year and after. By choosing the ESX6S, the cost would lower and the hardware is significantly upgraded, therefore, ESX6S was chosen.

5.5.1. Virtualization plan

Currently the ESN server has 8 core services in addition to multiple testing and production websites. In the virtualization plan, the ESN officials were interviewed to prioritize the services. It was decided to have six virtual servers, in addition to the domain host, esn.org and esncard.org would operate alone in the first two servers. These two websites are the most vital IT services of ESN, therefore, stable operation is required. The Galaxy together with event registration system is designed to operate in the same virtual server. The Galaxy has the user information which uses the LDAP database. This server will be the hardest to migrate to the new system. The fourth server would host the rest of the websites, for example project and event websites, including Mailman, the mailing list service. Virtual machines five and six are for testing and development purposes.

For each virtual machine partition schemes were made. The new server has two times 3TB of storage in comparison the old server which used altogether 307GB.

The Xen host will use LVM partition for the virtual disks. LVM is a logical volume manager for Linux OS. LVM establishes the possibility for more flexible disk usage. The storage volumes can be resized inside the LVM. (Lewis, 2012)

In the original plan, the internal backup space was designed for all the machines. However, after consultation with ESN, internal backup is not used. There is high risk relying on it, because if the disk has failures, the back-ups are unusable. A remote back-up system will be used, together with software RAID level-1.

The virtual machines are partitioned similarly to the Xen host. Space is used for base system, home directories and the data files for the current service. Each partition contains SWAP space, which is extended software memory if the actual RAM-memory runs out. SWAP memory is a valuable back up, especially in virtual servers.

The new server has 32GB of RAM memory and the current system is equipped with 12GB. According to the server statistics of the ESN server, around 8GB of memory is in use. When isolating the servers, RAM usage per server is very low although it is important to have enough free space, for the servers to operate flawlessly. The memory is designed to be allocated based on the usage of the server. Servers 1 and 2 would have 8 GB, 3 and 4 with 6 GB and the testing servers 4 GB. The domain host Dom0 was allocated with 8 GB of memory. There were different options inside the committee about the memory usage. Discussion is still ongoing, whet ever flexible or adjusted RAM memory is used for the servers.

Networking in virtual environment operates differently than in physical environment. Virtual machines do not have physical hardware. With the network card, it is more complicated compared to other devices. Moreover, several IP addresses travel through one network card. The physical device does not by default know where the addresses should be assigned to. For this reason, a network bridge is needed. Network bridging creates an additional network device which forwards the IP addresses to the virtual machines. In the ESN networking scheme, seven network addresses are needed. A subnet for fourteen addresses is seen as the wisest solution (/28 subnet) which can be configured as DHCP server from the domain host.

5.6. Implementation

Implementation was done in separate stages. First, the testing environment was set to familiarize with Xen platform. The initial testing was the most time-consuming part of the project. Testing was conducted in a desktop computer equipped with the Debian Squeeze operating system together with the Xen system. The domain host was installed to the server with three virtual servers. Each of the virtual servers included one service: one satellite website, web interface administration platform and mail service. These servers were operating in LAN, receiving IP addresses through the network bridge from the domain host. All servers were accessed with SSH from an external computer. The initial testing gave data to establish the technical implementation document (Appendix 2) where the implementation is described in detail for the IT technicians.

At time of writing, the project is in the next testing phase, where the planned environment is set up in testing platform at the ESN headquarters. After this phase, the new server will be installed with the base server. The actual server will be installed with a network installation. All installation packages required for installation are identified in the implementation document. The most important packages are web server, Apache, remote control service, SSH server, mySQL database, OpenSSL, encryption for the web sites, PHP5, web-programming support and Webmin, graphical web interface administration.

The security aspects are considered to be important. The usual problem with the old environment was the authorized access. Too many members of personnels have access to the complete system. Limited access to systems is created. The Web Project Administrator of ESN will have access to all systems. The virtual environment administrator has access to the Domain host. Only selected IT committee members can access the production websites, where they implement IT projects of ESN. All the committee members have access to the testing environment. These are described in detail in the implementation plan. (Appendix 2)

IPtables will be used as the firewall of the system which control the network traffic towards the virtual servers. Hardened authorization will be used to access the Domain host. Normal SSH authorization is used for the virtual environments.

5.7. Outcomes

As mentioned before, during the time of writing, the project is still in the design and testing phase. While ESN is a voluntary organization the project members are involved with variety of others positions and tasks. It has been challenging to implementing the project on a rapid schedule. However, the need for the new server environment has been identified for long period of time. The tendency to keep the old environment can be described as the “easy way out” solution as the new system would need resources, documentation and adjustment using it. However, ESN is determined to migrate to the new system in order to make it more efficient and stable.

After the testing phases, the implementation will start. The work will be carried out by the IT committee. Scheduled migration will be conducted during January and the new system is planned to run a live environment in early February. Documentation is done throughout this whole period in addition to constant documentation of changes to the new system.

6. Conclusion

6.1. Virtualization is today

Based on the arguments presented in this thesis, virtualization is an important factor in today’s IT infrastructure. Virtual environments consume more efficiently the hardware resources, reduce IT costs, enhance security aspects and eases the management. According to surveys made on the corporations, cost reductions and hardware utilization are the most important reasons for migrate to virtual environments. (InformationWeek, 2011,Veeam, 2011, CDW-G, 2010). ESN International choose migration to virtualization based on these arguments. Low budget organizations, such as ESN benefit greatly from the possibility to have multiple servers with the cost of one.

However, virtualization has a long road to travel, while the penetration rate is under 40% (Veeam 2011). Virtualization has many usage methods, therefore, offering several options for companies, organizations and IT hobbyists to enhance their systems. One option is to make several virtual servers with one physical hardware, another is to use Microsoft Windows with Mac computers, and a third option could use specialized software through application virtualization.

Competition in virtualization markets is dominated by VMware. Virtualization has been commercialized only for 14 years, therefore, VMware having the advantage is no surprise. While virtualization is becoming more popular with small companies and hobbyists, it could be expected to see open source solutions, for example Xen and VirtualBox to grow their market share in the near future. Virtualization solutions are significant, yet an huge financial investment. However, small and specialized companies might see it as unnecessary cost while competitive open source-solution are available.

Security, design, and testing are equally important in virtual and physical environments. Security aspects differ from physical solutions where the system is located in one centralized position. Virtual environments have the base system together with multiple virtual systems to make secure. Careful planning with suitable solutions, based on the size and the purpose of the environment is recommended. In the project made for this thesis the security aspects were carefully considered. Linux hardening and strong authorization limits were implemented to conclude this task.

Proper documentation and designing could save the administrators resources in addition to giving structure to the environment and easier recovery from system failures.

Xen is a highly considerable option when searching for low-cost, lightweight and simple solution for virtualization. The open source community opens the possibility for several additional tools, for example, graphical user interfaces. With Xen virtualization, technology becomes familiar to all types of users. The documentation is broad, additionally easy to understand. Xen can be installed to older hardware, therefore, the threshold for testing is low. Xen was perfect solutions for ESN, because the lightweight and management capabilities. Xen was already familiar software for the main technician and it was easy to learn for other as well.

The ESN migration project is still ongoing and will be finished in February 2013. Consultation and documentation for the project is continued after the this thesis.

6.2. Future

Virtualization is a relatively young technology in commercialized markets. The hardware used in computer architecture were not designed to run virtualized environments. Therefore, the growth of virtualization is mild (Gartner, 2012a). However, the two largest CPU manufacturers, AMD and Intel, are developing CPU architecture to meet the demands of virtualization. Hardware assistance is growing, therefore, full virtualization becomes a more approachable solution (VMware, 2012b)

Virtualization covers a wide area of different technologies. In addition to server virtualization, new approaches, for example, wireless, home appliance and phone virtualization are researched.

Wireless technology, especially Wireless LAN, has significant potential for virtualization. WLAN is often used with wireless API through DSL-line. DSL connection speed is relatively low differing from 1 to 24 Mbps (ITU, 2012), compared to WLAN 54 to 600 Mbps (IEEE, 2012). In this sense, WLAN card capacity is unused. Research describes the possibility to use multiple channels, isolated from each other, using different wireless APIs. This would increase the utilization level of WLAN cards and, therefore, increase the connection speed. The computer could use several DSL lines with a single network card. (Bianchi *et al*., 2012, 2-7)

The energy savings of virtualization could benefit home environments greatly. In Future Home Environment design (FHE), virtualization is used in end user computers as well as with networking. A home which is having multiple computers could save energy and costs with OS virtualization. The architecture enables load balancing, where the resources of unused computers, are shifted to the ones in use. The unused or passive computers can be hibernated which then saves energy. (Berl *et al*., 2009)

Smart phones are today more popular compared to laptop computers (Canalys, 2012). Therefore, smart phone virtualization is a hot topic. The majority of smart phones are running either Google Android or iPhone iOS operating system, which both bases on Linux systems, with a Unix kernel (Apple, 2012, Google, 2012). Smart phones can be described as small computers, because they have a CPU, RAM memory and a storage space. With smart phones virtualization could mean application and OS virtualization. CPU phone virtualization is already possible with XEN on ARM (Chen, 2011). ARM is a manufacturer of RISC microprocessors, mainly used in embedded systems (ARM, 2012). In addition, Android OS already has a virtual machine as a core service, called Dalvik Virtual Machine. (Google 2012.)

Another approach to phone virtualization is syncing smart phones together with desktop phones. Enterprises usually have desktop phones working with VoIP in their offices. In addition, employees have smart phones. Syncing the data together is an efficient way to keep up-to-date, wherever the work takes place. This can be achieved with microkernel-based phone virtualization. (Acharya, Buford and Krishnaswamy, 2009)

As we can see from the examples above, virtualization has many forms. As a concept virtualization is much more, than running Windows XP on Mac OSX. Hypervisor-based virtualization can be used in network, smart phone and computer virtualization. The virtualization technologies are evolving and in the future the technology will be embedded to all computing. The future will tell.

REFERENCES

Acharya, A., Buford, J., Krishnaswamy, V. (2009). 'Phone Virtualization Using a Microkernel Hypervisor'. Available from: IEEE Electronic Databases [Accessed 25.11.2012].

Apple Inc. (2012). *Mac Technology overview*. Available at: https://developer.apple.com/library/mac/#documentation/MacOSX/Conceptual/OSX_Technology_Overview/SystemTechnology/SystemTechnology.html (Accessed 25.11.2012)

ARM Ltd.(2012). *Company Profile*. Available at: http://www.arm.com/about/company-profile/index.php (Accessed 25.11.2012)

Avanade Perspective (2009). *Server Virtualization: A Step Toward Cost Efficiency and Business Agility*. Available at: http://www.avanade.com/en-us/approach/research/Pages/Server-Virtualization-A-Step-Toward-Cost-Efficiency-and-Business-Agility.aspx (Accessed 28.10.2012)

Bianchi, G., Gallo, P., Goma, E., Grunenberger, Y., Tinnirello, I. (2012). 'Wireless Card Virtualization: From Virtual NICs to Virtual MAC Machines'. Available from: www.FutureNetworkSummit.eu/2012 [Accessed 25.11.2012].

Berl, A., de Meer, H., Hlavacs, H., Treutner, T. (2009). 'Virtualization in Energy-Efficient Future Home Environments'. *IEEE Communications Magazine* [Internet]. December 2009, 62-67. Available from: IEEE Electronic Databases [Accessed 25.11.2012].

Bugnion, E., Devine, S., Rosenblum, M. (1997). *Disco: Running Commodity Operating Systems on Scalable Multiprocessors*. Available at: http://www-flash.stanford.edu/Disco (Accessed 28.10.2012)

Canalys (2012). *Smart phones overtake client PCs in 2011*. Available at: http://www.canalys.com/newsroom/smart-phones-overtake-client-pcs-2011 (Accessed 25.11.2012)

Carnegie Mellon University (1999). *CERT/CC Denial of Service*. Available at: http://www.cert.org/tech_tips/denial_of_service.html#1 (Accessed 5.11.2012)

Cassidy, A. and Cassidy, D. (2009) *Practical Guide to Reducing IT Costs.* Fort Lauderdale, Florida: J. Ross Publishing Inc.

CDG-W (2010), *Governmental Virtualization Report*. Available at: http://www.cdwnewsroom.com/2010-government-virtualization-report-agencies-report-reduced-costs-improved-continuity-of-operations-and-security-benefits/. (Accessed 28.10.2012)

Cerling, T., Buller, J., Enstall, C., Ruiz, R. (2010) Mastering Microsoft Virtualization. Indianapolis, Indiana: Wiley Publishing Inc.

Chen, X. (2011). 'Smartphone Virtualization: Status and Challenges'. IEEE Electronic Databases.

Cisco (2012). *Network Virtualization Solutions*. Available at: http://www.cisco.com/en/US/netsol/ns658/index.html. (Accessed 23.10.2012)

Citrix Systems Inc (2012). *Xen History*. Available at: http://www.xen.org/community/xenhistory.html (Accessed 24.10.2012)

Collect, S. 2011. Virtual Vexations. Computerworld October 2011, 16-22

Debian (2012). *Bridging Network Connections*. Available at: http://wiki.debian.org/BridgeNetworkConnections (Accessed 23.10.2012)

Gabriel Consulting Group (2011). *Survey: x86 Virtualization Picking Up Steam, Paying Off*. Available at: http://gabrielconsultinggroup.com/gcg-news-and-views/20-general-blog/241-survey-x86-virtualization-picking-up-steam-paying-off.html (Accessed 28.10.2012)

Gabriel Consulting Group (2011). *Unix Survey: Is Virtualization Hitting the Wall?* Available at: http://www.gabrielconsultinggroup.com/gcg-news-and-views/20-general-blog/288-unix-survey-is-virtualization-hitting-the-wall.html (Accessed 28.10.2012)

Gartner (2011). *Virtual Machines Will Slow in the Enterprise, Grow in the Cloud*. http://www.gartner.com/DisplayDocument?id=1573914 (Accessed 23.10.2012)

Gartner (2011). *Ten Key Actions to Reduce IT Infrastructure and Operations Costs*. Available at: www.gartner.com -> Research. (Accessed 29.10.2012)

Google (2012). *About the Android Open Source Project*. Available at: http://source.android.com/index.html (Accessed 25.11.2012)

Granneman, J., Ottenheimer, D., Savage, M. (2012). Information Security Essential Guide. Newton, Massachusetts: TechTarget Inc.

Department of Justice, The United States 2011. 'Former Shionogi employee arrested, charged with hack attack on company servers'**.** Available from: http://www.justice.gov/usao/nj/Press/files/Cornish,%20Jason%20News%20Release.html. [5.11.2012]

Erasmus Student Network AISBL (2012). *About ESN*. Available at: http://www.esn.org (Accessed 10.11.2012)

Forrester (2012). *The CISO's Guide To Virtualization Security*. Available at http://www.forrester.com/The+CISOs+Guide+To+Virtualization+Security/fulltext/-/E-RES61230?al=0. (Accessed 23.10.2012)

Hagen von, W. (2008) *Professional XEN Virtualization*. Indianapolis, Indiana: Wiley Publishing Inc.

IBM (2012). *IBM and HP Virtualization*. Available at: http://www.ibm.com/developerworks/aix/library/au-aixhpvirtualization/index.html. (Accessed 26.10.2012)

IEEE Standards Association (2012). *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications.* Available at: http://standards.ieee.org/about/get/802/802.11.html. (Accessed 25.11.2012)

InformationWeek Analytics (2011). *State of Virtualization*. Available at: http:/reports.informationweek.com. (Accessed 25.10.2012)

International Telecommunications Union (2012). *G.992.3 : Asymmetric digital subscriber line transceivers 2 (ADSL2)*. Available at: http://www.itu.int/rec/T-REC-G.992.3/en. (Accessed 25.11.2012)

Internet Engineering Task Force (IETF) (2007). *TCP SYN Flooding Attacks and Common Mitigations*. Available at: http://tools.ietf.org/html/rfc4987 (Accessed 5.11.2012)

Kaspersky Lab (2012). *The Kaspersky Security for Virtualization eBook*. Available at: http://www.kaspersky.com/berady. (Accessed 29.10.2012)

Kushetzky, D. (2011) *Virtualization: A Manager's Guide*. Sepastobol, California: O'Reilly Media Inc.

Lee, M., Krishnakumar, A., Krishnan, P., Singh, N., Yajnik, S. (2010). XenTune: Detecting Xen Scheduling Bottlenecks for Media Applications. Georgia Institute of Technology, Atlanta, Georgia: IEEE Electronic Databases.

Lewis, A. (2012). *LVM HOWTO.* Available at: http://tldp.org/HOWTO/LVM-HOWTO/ (Accessed 28.11.2012)

Linux (2012). *Linux manual pages. Syscalls(2).* Available at: http://linux.die.net/man/2/syscalls (Accessed 30.11.2012)

Lowe, S. (2011) *Mastering VMware vSphere 5*. Indianapolis, Indiana: John Wiley & Sons, Inc.

Matthews, J., Dow, E., Deshane, T., Hu, W., Bongio, J., Wilbur, P., Johnson, B. (2008) Running XEN: A Hands-On guide to the Art of Virtualization. Boston, Massachusetts: Pearson Education Inc.

Microsoft (2012). *MSDN, Hyper-V Architecture*. Available at: http://msdn.microsoft.com/en-us/library/cc768520%28v=bts.10%29.aspx (Accessed 23.10.2012)

meXt (2009). *Virtualisointi suomalaisissa organisaatioissa 2009*. Available at http://www.mext.fi/tutkimukset/virtualisointi-suomalaisissa-organisaatioissa-2009 (Accessed 26.10.2012)

Parallels (2012). *About Parallels*. Available at: http://www.parallels.com/eu/about/ (Accessed 25.10.2012)

Romero, A. (2010). *VirtualBox 3.1 - Beginner's Guide*. Birmingham, United Kingdom: Packt Publishing.

Ruest, D. and Ruest, N. (2009). *Virtualization, A Beginner's Guide*. New York, New York: The McGraw-Hill Companies.

Shea, R. and Liu, J. (201). *Understanding the Impact of Denial of Service Attacks on Virtual Machines*. Available at: IEEE Electronic Database (Accessed 5.11.2012)

Takemura, C. And Crawford, L. (2009) *Book of XEN*. San Francisco, California: No Starch Press.

Thayer and Smith (1999). *Greek Lexicon entry for Para. "The NAS New Testament Greek Lexicon".* Available at: http://www.biblestudytools.com/lexicons/greek/nas/ (Accessed 24.10.2012)

Veeam Software (2011). *V-index.com Virtualization Industry Quarterly Survey*. Available at: http://www.v-index.com/. (Accessed 28.10.2012)

VMware (2012). *Virtualization overview*. Available at: www.vmware.com/pdf/virtualization.pdf. (Accessed 24.10.2012)

VMware (2012). *Understanding Full Virtualization, Paravirtualization, and Hardware Assist*. Available at http://www.vmware.com. (Accessed 26.11.2012)

APPENDIXES

APPENDIX 1
DESIGN PLAN — Joonas Niemi
SERVER ENVIRONMENT — 22.11.2012
ESN INTERNATIONAL

In this design plan ESN International's web server current functionality, planned and designed the new virtual server environment are explained. The current ESN server is outdated, everybody who has access to the server, has access to almost everything. ESN server needs to be better structured, different web sites will be isolated from each other, moreover working independently.

Current server have been examined and transferrable services are listed to be transferred to the new environment. New server hardware options have been researched and recommendations for the best options given.

1. Current ESN server

Current ESN server has two most important directories

/home/ has directories for users in ESN.org server. There is still many old members files. These should be archived
/data/ has all the website files and sub-site files. There is many unused sites and domains, which can be removed
In total there is 51 domains, which from 12 redirects to ESN.org, 4 to redirects somewhere else, 12 is some kind of testing sites. From this we can already tell, there is too many and useless domains. Appendix you can find the list from the sites.
/var/www has monitoring tools for the server. All "real" websites are under / data/sites

Services in current ESN server:
-esn.org
The main website for the organization

-Galaxy
User and section database about ESN sections and their members who have signed in to galaxy.esn.org. This database is planned to be used for multi-signups (one account, access to all services)

-events.esn.org
Event registration page, where ESN members can sign up for ESN meetings. Uses galaxy accounts

-Satellite
Web site package for sections and projects. Satellite has many development and testing environment in ESN server.

-Projects web sites
All ESN projects like Prime, ExchangeAbility and SocialErasmus has their own web sites under esn.org

-Web shop
Not used for now, but will have possibility to buy gadgets and clothing

-Mailman
ESN mailing lists are operated by mailman. There is plenty of lists, many unneeded as well.

-ESNcard.org
Web site for ESNcard. Currently operated from another server, but will migrated to the virtual environment

Esn.org and esncard.org are the most important services of ESN. Those web sites have the most visibility and they represent ESN to the outside world. It is crucial that they are stable and fast. It is recommended that these two sites are isolated from each other to their own virtual servers. Recommendation is to have both esn.org and esncard.org in another server as well, if the main virtual server causes failures.

All servers should have unique IP addresses, so they are easier to control and can be accessed from outside. Purchasing seven IP-addresses is recommended.

2. New server

There are two very good options for the new server: Dedicated servers EX 6S and EX 10. EX 6S costs 79 euros per month with 149 euros one time setup fee, EX 10 costs 109 euros with 149 euro setup fee.

EX 6S:
Intel® Xeon® E3-1245 Quadcore incl. Hyper-Threading Technology
RAM **32 GB DDR3 RAM ECC**
2 x 3 TB SATA 6 Gb/s 7200 rpm (Software-RAID 1)
NIC **1 Gbit OnBoard connected at 100 Mbit**
Backup Space **100 GB**

EX 10:
Intel® Core™ i7-3930K Hexacore incl. Hyper-Threading Technology
RAM **64 GB DDR3 RAM**
2 x 3 TB SATA 6 Gb/s 7200 rpm (Software-RAID 1)
NIC **1 Gbit OnBoard connected at 100 Mbit**
Backup Space **100 GB**

Traffic **Unlimited**	*Traffic* **Unlimited**

The current server has 12GB of RAM memory which from averagely 1,2GB is free, 3,3GB is inactive and 4,88GB is active. It can be concluded the 12GB being enough currently, however when allocated better when virtualized memory can be used more efficiently.
EX 6S has 32GB of memory which should be considered be the minimum amount. It also uses the Error-Controlling Code method to solve memory errors.
EX 10 has enormous 64GB of RAM. This is more than is required, though it can be said that there is never enough RAM memory.
The current server has Intel i7-920 Quadcore 2,67Mhz processor, which is outdated. The two new options have very efficient processors from Intel according to the cpubenchmark.net's Passmark test:

Intel Core i7-3930K @ 3.20GHz	13,593
Intel Xeon E3-1245 V2 @ 3.40GHz	10,148

I7-3930 is better in this comparison as it has six cores compared to E3-1245 four. In the end E3-1245 would be enough for ESN server.

According to the monitoring stats the current server uses average 30% of CPU performance. There has been some peaks, though in general the CPU usage has been low.

Current server has the following filesystem:

Filesystem	Size	Used	Avail	Use%	Mounted on
/dev/md2	1.4T	307G	1005G	24%	/
tmpfs	5.9G	0	5.9G	0%	/lib/init/rw
udev	10M	784K	9.3M	8%	/dev
tmpfs	5.9G	0	5.9G	0%	/dev/shm
/dev/md1	2.0G	79M	1.9G	5%	/boot

The server has one partition of which it uses 307GB out of 1,5TB. From this 307GB /data uses 21G and /home 201G. The home directory is significantly large. This can be explained by having backups store in home directories which takes 131G. However, this is not the most reasonable solution not to have external backups.

The new servers have both 2x3T of storage. As we can see from the current server around 500G is a safe amount of memory. Generally, when having virtualized system, we need to reserver slightly more storage as all servers have their own operating system. 6TB of memory all together is good number to work with. It is highly recommended to use RAID system, where the disks are mirrored. The offered amount gives the possibility for detailed backup systems.

3. Cost calculation

At the moment ESN pays for two servers: One virtual for esncard.org and one for everything else. The servers are hosted by Heztner (www.hetzner.de). The servers are EQ 6 (not available anymore) for 60e per month and VQ 19 for 17e per month. ESN as an NGO is a free from VAT tax. In total this makes 924 euros per year.

The new servers would cost either 64e per month plus the one time setup fee of 120,70e (EX6S) or 88,30e per month plus setup fee of 120,70e (EX10). Additional costs include extra IP-addresses, in our case the 6 IP’s, which is 5,67e per month.
Yearly costs would be 960e first year, 839,30e after that (EX6S) or 1248,34e first year and 1127,63e after (EX10).
Generally having EX6S will save in yearly costs, EX10 would bring more. However virtualization systems ease the administration, gives more security and they are more efficient. Invest in powerful hardware now, will give savings in the long run.

4. Conclusion

To summarize, there is a need for a new server. Based on the research, ESN would not require the most powerful servers. The possible servers were narrowed in to two options EX 6S and 10. Both servers have same services, programs and data centers. The difference comes in hardware. EX 10 has more powerful processor and has more memory.
EX 6S has enough memory and the processor is fast enough for the server usage. EX10 has hexacore CPU which makes it even more powerful. EX 10 is costly, too large investment for ESN’s needs. The suggestion is to choose EX6S.

5. ESN International Virtualization plan

At the moment ESN server is very messy and needs to be structured and cleaned. Suggestion is to divide the services to virtual servers according to their usage and bindings to other services. Below you can find the scheme:

XEN Host (Dom0) Debian Squeeze Base server 192.168.1.100 (external IP from the ISP) br0 10.0.0.1

VS1 10.0.0.2 esn.org	VS2 10.0.0.3 esncard .org	VS3 10.0.0.4 Galaxy Events	VS4 10.0.0.5 Producti on Mailman	VS5 10.0.0.6 Testing	VS6 10.0.0.7 Aegir

Networking: 10.0.0.0/28 network, mask 255.255.255.240, Gateway 10.0.0.1, usable addresses 10.0.0.2-14, Broadcast 10.0.0.15

Partition scheme, Dom0:

XEN host Dom0 Base system Mount point / 50G XFS	SW 4G	Home directories Mount point / home 50G XFS	Virtual machine disks LVM 1892G

esn.org:

VS1 esn.org Base system / 20G	SW 2G	Users /home 10G	Website files no mount point 28G	non allocated space no mount point 40G

esncard.org

VS2 esncard Base system / 20G	SW 2G	Users /home 10G	Website files no mount point 28G	non allocated space no mount point 40G

Galaxy:

VS3 galaxy Base system / 20G	SW 2G	Users /home 10G	Website files no mount point 28G	non allocated space no mount point 40G

Production sites:

VS4 Prod sites Base system / 20G	SW 2G	Users /home 10G	Website files no mount point 28G	non allocated space no mount point 40G

Testing and Development:

VS5 Testing Base system / 30G	SW 2G	Users /home 30G	Website files no mount point 15G	non allocated space no mount point 23G

Aegir:

VS6 Aegir Base system / 10G	SW 1G	Users /home 10G	Website files no mount point 15G	non allocated space no mount point 14G

In this partition scheme 2TB disk is to be partitioned from the host system. Base system is the Linux installation mounted to root.
50G is fairly enough, as the system will use around 10G with, cause we can choose completely which software we install. SWAP space is designed to be used for all disk partitions, which the current server is missing. Swapping is a process where RAM memory is transferred to a virtual memory space in a hard disk. This could happen, for example when memory runs out, or there is idle services which can be swapped from reserving memory. However, SWAP memory is not as fast as RAM memory. If RAM memory runs out constantly SWAP memory is significantly slower as it comes from a hard drive disk.
In our case this should not happen (running out of memory), though SWAP is vital for server stability. If the server do run out of memory and it does not have any place to retrieve more, it will crash. Another huge benefit is to swap unnecessary services from using the RAM memory.

Virtual machines have been given 1.3TB storage, which gives maximum allocation 100G to each server. Each virtual machine has its own “disk”, which is partitioned inside the virtual server. First four machines uses the same partition scheme as they all have live websites, so fort easier to configure all together. The testing server uses different partition, moreover it has different purpose.

Back ups are vital, therefore external back up system is used. The back ups will be sent to the external server with special software, which is decided later.
The new server offers RAID-1 technology. This means that the data is written to two disks identically (mirroring). RAID is recommended to use for stability, hence one disk breaks the server is still fully functional and all the data is on the another disk.

RAM memory allocation with 32GB:

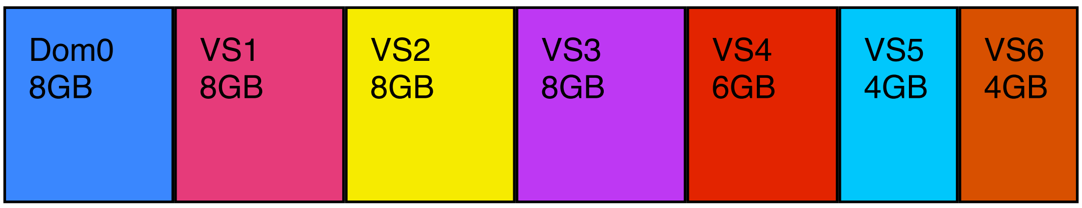

As virtual servers are in use virtual more RAM memory is needed than in normal environments. It is wise to allocate enough memory for all servers, thus keeping in mind the priority of the services, which are esn.org and esncard.org. Memory allocation is based on the current servers statistics on how much memory is used (described in page 3). When the servers are isolated the memory is used better on the corresponding services. Domain host needs 8GB to stably host virtual servers. Testing environments VS5 and VS6 needs less memory as their usage level is lower.

APPENDIX 2
TECHNICAL IMPLEMENTATION DOCUMENT
ESN INTERNATIONAL

Joonas Niemi
13.11.12

Description

This document describes the technical implementation of the new ESN server environment. Design of the system is described in the design plan. This document explains detailed installation procedures for each server.

Defining requirements for virtual environment

Currently ESN server is operating all services from one physical server. The services are explained in the design document. The new environment has 6 virtual servers installed on host server. All ESN websites have different requisites, which need to be identified before installations. In addition, it is needed to identify softwares needed for base system and all virtual installations.

All operating systems will be Debian Squeeze 6.0 x68 AMD installations. Debian was chosen, cause it is used by the former server and most of the technicians are most familiar with it.

The base system needs Xen Cloud Platform as well, to host virtual servers. In Debian Squeeze this can be installed directly from distributed packages via apt. All servers will be installed normal CLI server installation with the following additional softwares:

-Apache2 web server
-MySQL database server and web server module
-PHP5 and web server module
-SSH server
-OpenSSL
-Webmin

The virtual servers are created and partitioned based on the design document. All servers have different work purpose, therefore different administrators as well. We will create user privileges for all servers separately as follows:

Base server:
-Web project administrator
-IT committee chairs
-Virtual system administrator

Esn.org:
-Web project administrator

-Communication manager
-IT intern

Esncard.org:
-Web project administrator
-Nuvole developers

Production websites:
-Web project administrator
-Web site administrators
-Project coordinators (IT committee members)

Aegir:
-Web project administrator

Testing and development:
-Committee members

Directory structures

All servers will have separate partitions for root (/), home (/home) and website files (/data). Each website will have their own directory under data, for example /data/esn.org.

Installation procedures for the virtual environment

The servers are installed separately, although there is a possibility of cloning. However, installing only the core installation, does not take that much time. We can also ensure, that all servers are unique and no network conflicts will happen.

All servers will be installed through command line and no graphical desktops are needed. Network connections will go through a bridged network and interfaces will be configured through DHCP. The Dom0 has static external IP address and the bridge network has the gateway for the virtual servers.
Subnet /28 is configured for the virtual network.
The IP addresses below are for explanation purposes, which are not final.

IP address assignment:
-Dom0: Host IP 192.168.1.100, br0 10.0.0.1/28, Gateway for virtual servers
-Server 1: 10.0.0.2/28
-Server 2: 10.0.0.3/28
-Server 3: 10.0.0.4/28
-Server 4: 10.0.0.5/28
-Server 5: 10.0.0.6/28
-Server 6: 10.0.0.7/28
Virtual server gateway 10.0.0.1, mask 255.255.255.240

Virtual server installation

The base server is installed straight to the hardware. Host system is installed through net installation with all additional packages needed. Most importantly Xen is installed and configured with instruction from http://wiki.debian.org/Xen. After the installation virtual servers are created. We can use either command line tools or GUI for managing the machines. Both options are tested and this is more personal preference.
It is recommended that first one virtual machine with Debian Squeeze is created and tested. If the test is successfully run, we can proceed to creating the virtual machines.
The machines are created one by one, with the specifications from the design plan. It is recommended first install the testing and Aegir servers. They are not vital for ESN services and can be used as verification that the system works. Second servers one, two and four (ESN.org, esncard.org, production websites) are installed. After installation one testing Satellite page should be configured to all servers to test the functionality. Testing should happen inside and outside the network. When the functionality is tested, the data can be transferred to the servers. We have to remember to use alternative names when configuring servers as the old servers still host the live sites. The last server to be installed is the with Galaxy and event registration (Server 3). This is the most challenging server, because of the LDAP-services needed. The project coordinators are more familiar on the migration of this service.

Below you can find the installation plan:

Aegir and Testing	Task
1	Installation of the OS
2	Installing needed packages
3	Creating the web environments
4	Testing the functionality
5	Transferring the needed data

ESN, esncard, production	Task
1	Installation of the OS
2	Installing needed packages
3	Installing testing Satellite pages
4	Testing the functionality

ESN, esncard, production	Task
5	Transferring the needed data
6	Configuring the real websites
7	Testing with alternative names
8	Turning live

Galaxy	
1	Installation of the OS
2	Installing needed packages
3	Installing LDAP
4	Configuring the environment
5	Testing
6	Transferring the data
7	Testing
8	Going live

Security

Security aspects needs to be considered. Authorized access, encryptions and firewall are the most important factors.

Authorized access:
-SSH server should be used in all remote control. The domain host will have hardened authorization based on this tutorial (http://howto.biapy.com/en/debian-gnu-linux/system/security/harden-the-ssh-access-security-on-debian)
-To preventing DDoS and brute-force attacks fail2ban is used. It seeks through logs that have malicious signs, for example too many wrong passwords
-Root login should be disabled in all systems. Root privileges can be used via su and sudo
-For this purpose unprivileged users should be made and password generated with apg
-Private and public key authentication hardens the access
-Changing the SSH port will lower the amount of automated attacks

-Using Port Knocking to hide SSH server from unwanted visitors

To access web-based configurations SSL-based access authorization should be used, together with the normal user authorization. IPtables-firewall will be configure in all servers.

APPENDIX 3
ESN.ORG SITE LIST
TABLE
ESN INTERNATIONAL

Joonas Niemi
13.11.12

#	Site list:		Description
1			AGM registration page
2			AGM registration page
3			Committee website, outdatet
4			Development site, authorized access
5			Event registration
6			ESNcard test site?
7			ESN Galaxy
8			ESN KAPA test website
9			ESN update server
10			ESN Wagenningen test site
11			ESN wiki page, unused
12			ESN.org
13			Events registration
14			Facebook playground testing site
15			Infocenter
16			IT committee help center
17			LDAP testing environment
18			Mailing list subscription page
19			Monitoring tools
20			PRIME website
21			Project website
22			Project website
23	 .org://agm.esn.org		Redirects to current AGM website
24	<a href="http://zabbix.esn.org">		Redirects to ESN.org
25	<a href="http://www.esn.org">		Redirects to ESN.org
26	<a href="http://pg.www.esn.org">		Redirects to ESN.org
27	<a href="http://new.galaxy.esn.org">		Former staging site for Galaxy 3
28	<a href="http://galaxy.pg.esn.org">		Former development site for galaxy 3
29	<a href="http://supernova.esn.org">		former blog and voting tool
30	<a href="http://pg.satgal.esn.org">		Redirects to ESN.org
31	<a href="http://pg.satupdate.esn.org">		playground for satellite upgrade
32	<a href="http://erasmus25.eu">		Redirects to ESN.org
33	<a href="http://25erasmus.eu">		Redirects to ESN.org
34	<a href="http://supernova.esn.org">		Redirects to ESN.org

35	<a href="http://feng.esn.org">		Redirects to ESN.org
36	<a href="http://alumni.esn.org">		Former website of the alumni
37	<a href="http://it.esn.org">		Redirects to esn.org/content/committe
38	<a href="http://cnr.esn.org">		Redirects to events.esn.org
39	<a href="http://rp.esn.org">		Redirects to events.esn.org
40	<a href="http://pg.event.esn.org">		Satellite event playground
41	<a href="http://pg.rc1.esn.org">		Satellite release candidate
42	<a href="http://johnny.pg.esn.org">		Satellite testing site
43	<a href="http://erasmus25.dev.esn.org">		Satellite testing site
44	<a href="http://satellite.esn.org">		Satellite testing site
45	<a href="http://pg.theme.esn.org">		Satellite theme testing site
46	<a href="http://pg.theme2.esn.org">		Satellite with another layout
47	<a href="http://sitelist.esn.org">		Sitelist in a website
48	<a href="http://socialerasmus.esn.org">		SocialErasmus website
49	<a href="http://webshop.esn.org">		Web shop
50	<a href="http://webshop.dev.esn.org">		Web shop development
51	<a href="http://git.esn.org">		

APPENDIX 4
SERVER STORAGE USAGE
TABLE
ESN INTERNATIONAL

Joonas Niemi
12.09.2012

/home/:	**Web sites:**
8K ./josefin	root@esn:/data/sites/esn.c
1.8G ./jurrien	106M ./exchangeability
92M ./git	131M ./feng
32G ./benjamin	103M ./it
12K ./joonas	4.2G ./www
16G ./erikr	274M ./helpcenter
28K ./marco.larosa	4.0K ./questionaire
4.4G ./johnny	8.0K ./cnr
130G ./esnbu	24M ./wiki
36K ./damian	12K ./list
3.2G ./hendrik	91M ./webshop
36K ./jirka	16K ./agm
2.9G ./rania	171M ./sat_depl_files
538M ./marc	203M ./agm-registration-
7.6G ./fabian	34M ./alumni
241M ./gitlab	44K ./maintenance
1.8G ./adamw	275M ./eefie
200G .	260M ./socialerasmus
200G total	232M ./committee
	684M ./galaxy
/data/:	75M ./satellite
24K ./public_keys	2.4G ./playground
18G ./sites	2.8G ./dev
235M ./download_area	475M ./agm-registration
296M ./sites_archive	357M ./supernova
383M ./svn	214M ./prime
2.5G ./playground	60M ./drush_backup
29M ./dev	4.0G ./infocenter
32K ./ssl	472K ./trac
21G .	302M ./rp
21G total	18G .
	18G total

Disk usage:

Filesystem		Size	Used	Avail	Use%	Mounted o
/dev/md2		1.4T	307G	1005G	24%	/
tmpfs		5.9G	5.9G	0	0%	/lib/init/rw
udev		10M	784K	9.3M	8%	/dev
tmpfs		5.9G	5.9G	0	0%	/dev/shm
/dev/md1		2.0G	79M	1.9G	5%	/boot