Juha Ahlgren

Reasons to Think About Cloud Computing

Metropolia University of Applied Sciences

Bachelor of Engineering

Media Technology

Thesis

28.2.2013

Metropolia

Authors Title Number of Pages Date	Juha Ahlgren Reasons to Think About Cloud Computing 46 pages + 1 appendix 16 August 2012
Degree	Bachelor of Engineering
Degree Programme	Media Technology
Specialisation option	Digital Media
Instructors	Aleksi Tuovinen, Middleware Software Services Country Manager Jouni Huttunen, IT-Specialist Olli Alm, Senior Lecturer

This thesis presents the essential characteristics of cloud computing such as on-demand self-service, broadband access, resource pooling, rapid elasticity and the measured service. The study also discusses service delivery models and deployment models, which define how services are delivered to the customer. Cloud computing has security issues just like every computing technology. Because these security problems lie in data transfer and software interfaces, as well as data which has been stored and user access control, they are analyzed in this study. This thesis presents technologies which enable cloud computing such as virtualization and hypervisors that are used to virtualize the servers and which can be used to deploy the services.

This thesis was made for a large international IT-company as a part of a customer project. The study presents company products that were used in the project to deliver cost-efficient customer solutions. Cloud computing is a relatively new trend in modern information technology, especially in the constantly evolving Finnish market. This study compares global benefits acquired with cloud computing as well as the effects of this technology to the quality of service and costs. Technologies compared in this study are hypervisors used in server virtualization such as ESXi, KVM, Hyper-V, Xen and PowerVM.

Although ESXi hypervisor is the dominant one, it is also more expensive. Open-source-based KVM hypervisor can manage almost the same features but at a significantly lower cost.

Keywords	cloud computing, virtualization

Tekijä Otsikko Sivumäärä Aika	Juha Ahlgren Pilvipalveluiden mahdollisuudet 46 sivua + 1 liite 1.3.2013
Tutkinto	insinööri (AMK)
Koulutusohjelma	mediatekniikka
Suuntautumisvaihtoehto	digitaalinen media
Ohjaajat	Aleksi Tuovinen, Middleware Software Services Country Manager Jouni Huttunen, IT-Specialist lehtori Olli Alm

Insinöörityössä perehdyttiin pilvipalvelun olennaisiin ominaisuuksiin, joita ovat itsepalvelumahdollisuus, palveluiden hyödyntäminen ajasta, paikasta ja laitteesta riippumatta, resurssien yhteiskäyttö, käytön mittaaminen ja palvelun skaalautuvuus. Näiden lisäksi käytiin läpi jakelumallit ja käyttöönottomallit, joilla määritellään, miten palvelu asiakkaalle toimitetaan. Insinöörityössä perehdyttiin myös virtualisointiin, joka on olennainen teknologia pilvipalveluiden mahdollistamiseksi. Pilvipalveluihin liittyvät turvallisuusriskit piilevät datan siirtämisessä, tallentamisessa, käyttöliittymän turvallisuudessa ja käyttäjien seurannassa, joita tulisi miettiä tarkasti, kun valitaan palveluntarjoajaa sopivalle pilvipalvelulle.

Insinöörityö tehtiin suurelle kansainväliselle IT-yritykselle, ja osana työtä oli asiakasprojekti. Projektissa asiakas otti käyttöön yrityksen tarjoamia tuotteita, joilla asiakalle saatiin luotua kustannustehokas ratkaisu.

Pilvipalvelut ovat suhteellisen uusi suuntaus nykyaikaisessa tietotekniikassa, erityisesti Suomen markkinoilla, ja ne kehittyvät jatkuvasti. Työssä vertailtiin globaalisti pilvipalveluilla saavutettuja etuja ja teknologian vaikutusta yrityksien teknologiakustannuksiin ja palveluntarjoajien palvelunlaatuun. Vertailussa olevat teknologiat olivat virtualisoinnissa käytetyt hypervisorit ESXi, KVM, Hyper-V, Xen ja PowerVM.

Vertailu osoitti, että ESXi on tällä hetkellä markkinoiden valtias alalla, mutta KVM ja Hyper-V koettavat pysyä tiukasti mukana kehityksessä. Vaikka ESXi on vallitseva hypervisori, se on myös selvästi muita kalliimpi. Avoimeen lähdekoodiin perustuvalla KVM-hypervisorilla onnistuttiin saavuttamaan lähes samat hyödyt, mutta selvästi pienemmillä kustannuksilla.

Avainsanat	pilvipalvelut, virtualisointi

Acknowledgments

This study was conducted in IBM Finland to help me and readers of this thesis to get a uniform picture about the characteristics, opportunities, benefits, and market situation of cloud computing. I am very grateful for the opportunity IBM gave me and also the inspiriting atmosphere at the workplace.

At the beginning of this study my understanding of cloud computing was quite blurry, and I had to start from a clean desk to familiarize myself with the topic. Although, my studies at Metropolia did not include in-depth knowledge of cloud computing, I believe I made the right decisions by choosing this topic for my thesis.

I would like to thank my instructors Aleksi Tuovinen and Jouni Huttunen for their guidance and insightful suggestions for this thesis and also Cedric Foo, David Ainsworth, and Donal O'Connell for providing me valuable information. Also, I am grateful for all the colleagues who helped, guided and encouraged me in this study.

Finally, I want to thank my family for standing next to me during the whole writing process, and of course, my friends for their support.

Helsinki, July 18, 2012

Juha Ahlgren

Contents

1 Introduction

Cloud computing is such a new and vague term that even Larry Ellison, the Chief Executive Officer of Oracle Corporation did not have a clue what it means back in 2008: “I have no idea what anybody is talking about. I mean it [cloud computing] is really just complete gibberish. It’s insane.” [1] Even today cloud computing is sometimes compared to similar solutions such as grid computing and almost every supplier or client seems to have their own vision of what cloud computing technically is.

Everybody is interested in cloud computing, but there is sometimes confusion where it can be used or what kind of workloads it can manage. This is one of the topics that I am going to focus in my thesis. What exactly does cloud computing mean? That depends on who you ask. Someone might say that Facebook is an example of cloud computing whereas someone else might provide a very technical and precise answer. It seems almost any form of information technology has been rebranded as cloud computing – if it has anything to do with the internet. In simple terms cloud computing is a business model where a customer rents computing resources over the internet without the actual need of purchasing the physical infrastructure.

You are probably already using cloud-based services even without realizing it. Facebook, Gmail and Google Docs are good examples of popular and free cloud services. Cloud computing means having every piece of data you need for every aspect of your life at your fingertips and ready for use. And today data must be mobile, transferable, and instantly accessible.

This thesis is made to large international IT-company that provides many different computer services from pure infrastructure to comprehensive set of various software. Cloud computing is a relatively new topic, especially in the Finnish market that is also constantly evolving. This study compares benefits that are acquired with cloud computing globally, and also technologies effects to quality of service and costs. Structure of this study is presented below.

Chapter 2 focuses on cloud computing more specifically including definitive characteristics, differences between the delivery models that are used to deliver cloud computing

[1] Larry Ellison, Chief executive officer, Oracle Corporation, September 2008

services to customers. Also, the differences of private, public and hybrid service deployment models are discussed. Finally, security issues concerning cloud computing are dealt with.

Virtualization is the key technology to enable cloud computing and this technology are dealt with in chapter 3. Server virtualization technology and hypervisors such as ESX, KVM, Hyper-v, Xen and PowerVM are also briefly presented. Finally, the current market situation globally and in Finland is analyzed.

Chapter 4 presents a project focusing on customer needs conducted by IBM. The chapter also discusses technologies that were used in the project. Chapter 5 provides information about the Finnish cloud computing markets and what are the factors slowing down these markets. This chapter also covers the business benefits of the cloud computing.

2 What is cloud computing?

One source of confusion around the cloud computing might be that these services can appear in many different forms, and the vast majority of people experience cloud computing as a better service quality. Yet, two organizations, Gartner and National Institutes of Standards and Technology (NIST) have their own definitions for cloud computing. Gartner has defined cloud computing as following:

> "Cloud computing is a style of computing in which scalable and elastic IT-enabled capabilities are delivered as a service using Internet technologies." [1]

A popular definition of cloud computing that is provided by NIST is presented below:

> "Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction." [2, 2]

Cloud computing is an extensive model for real-time IT resource delivery over the internet. It is based on large distributed resource pools and the idea behind it is that the customers' demands vary over the time. Based on these needs the services can be scaled down to cut costs or up to meet the demands.

Cloud computing is ideally built on already established trends to cut costs of the delivery of services while increasing the speed and agility with service deployment. So from one perspective, cloud computing is nothing new because it uses approaches, concepts, and best practices that have already been established. All the capabilities of cloud computing have been made possible due to technologies such as virtualization which is automated, a network that has been improved to provide greater speed and storage capabilities that scales up as time goes by. However, from another perspective, everything is new because cloud computing changes the way we invent, develop, deploy, scale, update, maintain, measure, and pay for applications and the infrastructure on which they run. [3, 4]

The NIST has provided relatively comprehensive set of cloud definitions, including characteristics, delivery models and deployment models in a two-page document [8], which is very concise and uses industry-standard terms.

2.1 Essential characteristics

The NIST definition includes five essential characteristics to define cloud computing. Below characteristics are defined more specifically:

> On-demand self-service: “A consumer can provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.” [2, 2]

Cloud computing services enable the provisioning of cloud resources from self-service portal on demand. On-demand self service is a prime feature of most cloud offerings. The user can scale the required infrastructure up to a substantial level without disrupting the host operations. Service’s monitoring feature keeps track of usage of the users and with this information they can be charged exactly as much as they have used the service.

> Broad network access: “Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, tablets, laptops, and workstations).” [2, 2]

The most vital characteristic of cloud computing is the fact that it is network based and uses web-standards to establish connection from virtual environment to end-user and and establish connection. Resources should be accessible from anywhere from any standardized platform.

> Resource pooling: “The provider’s computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, and network bandwidth.” [2, 2]

The computing resources in the cloud are shared within all the users, in other words, the system does not care who is using which processor or memory. This means that numerous clients may be using the same set of resources and at the same time within the resource pool as shown above.

> Rapid elasticity: “Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time.” [2, 2]

The computing resources can be dynamically re-configured to adjust scale allowing for optimum resource utilization. Users can quickly acquire more resources from the cloud by scaling out. Accordingly, resources can be scaled back to original when they are not needed anymore.

> Measured Service: “Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.” [2, 2]

This means that recourse usage is metered in cloud computing by using different metering solutions. The data can then be used to charge customer by pay-per-use model.

Fundamentally, cloud computing services are just services, however, they appear in a form of infrastructure, platform or software that is delivered to user over the internet. These services can be started up by the consumer on demand without human intervention or lengthy implementations. Cloud computing services are accessed universally, and they scale up or down when required, offering a true utility approach to the consumption.

2.2 Service delivery models

In addition to the five characteristics discussed above, cloud community has used the following models to categorize their services: Infrastructure as a Service (IaaS), Platform as a Service (PaaS) and Software as a Service (SaaS).

Figure 1 presents these three cloud computing delivery models:

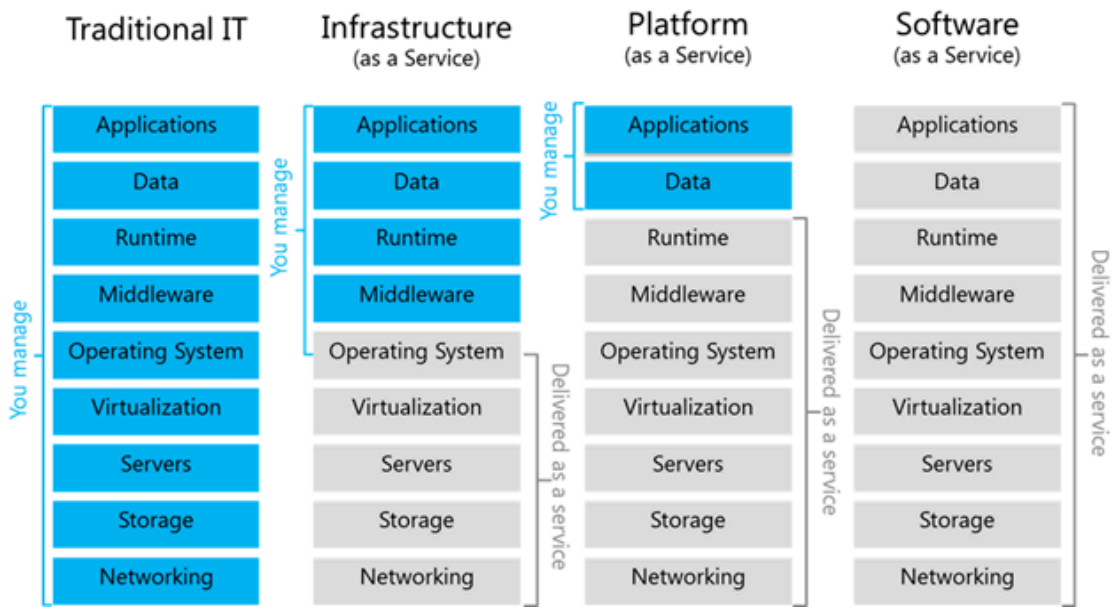

Figure 1. Cloud computing service delivery models. [4]

In figure 1 cloud computing services are assigned into different categories; these categories are divided into layers that describe what subscriber manages personally and what service provider includes in their service.

A commonly agreed upon framework for describing cloud computing services goes by the acronym Software-Platform-Infrastructure (SPI) model. This acronym represents the three major services provided through the cloud: SaaS, Software as a Service; PaaS, Platform as a Service; and IaaS, Infrastructure as a Service. These three categories are often shown in a stack or a pyramid model. In this model Infrastructure is at the bottom, Platform is in the middle and Software on the top. Other “soft” layers can be added on top of these layers, with elements such as cost and security extending the size and flexibility of the cloud.

NIST’s definitions of the three delivery models are presented below:

> Infrastructure-as-a-Service (IaaS): “The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls).” [2, 3]

Infrastructure as a Service provides computer infrastructure, such as virtualization, being delivered as a service. IaaS is popular in the data center where software and servers are purchased as a fully outsourced service. Services are billed based on how much of the resource is used - compared to the traditional method of buying software and servers outright. As can be seen from figure 1, IaaS service provider only the services where to run the operating system and a subscriber has to manage everything else.

The second NIST definition explains the model Platform as a Service

> Platform-as-a-Service (PaaS): “The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.” [2, 2]

Platform as a Service can provide the software platform where systems run on. This frees up developers from having to worry about the underlying infrastructure-related items, such as hardware, memory, network, storage and operating systems to enable rapid delivery of software applications. PaaS-based applications can automatically expand to meet required demand. A customer is only charged for the excess capacity for the period of time required. Well-known examples are the Google Apps Engine, Windows Azure and Engine Yard.

The last NIST definition explains the model Software as a Service.

> Software-as-a-Service (SaaS): “The capability provided to the consumer is to use the provider’s applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.” [2, 2]

Software as a Service is a software delivery model that provides access to software and its functions remotely as a Web-based service. Because the software is hosted remotely, users do not need to invest in additional hardware. SaaS removes the need for organizations to handle the installation, set-up and often daily upkeep and maintenance of the application or service. However, price of the service covers everything behind the service including networking, storage, platform, application and etc. [see

figure 1] Typical and well-known examples of SaaS are Facebook, Apple iCloud, Drop-Box and Google Docs.

2.3 Cloud deployment models

There are many decisions for cloud computing architects to make when moving from a standard enterprise application deployment model to one based on cloud computing. Selection of these models depends on clients' data sensitivity and management requirements. [2]

The differences between cloud deployments models depend on who has access to the different clouds and who owns the infrastructure behind the service. A private cloud (internal cloud) infrastructure is dedicated to a particular organization or group. A public cloud (external cloud) infrastructure is offered as a web application or as a web service and it is owned by an organization selling cloud services. A hybrid cloud deployment model exists due to mixed needs of an organization and it is a combination of two cloud service deployment models.

Organizations may have to consider a number of things when choosing the cloud computing model they wish to use. In order to solve a certain problem, the organization does not have to own a cloud infrastructure. The application they need for a short time might be best suited to be deployed into public cloud because it helps organization to avoid the need of purchasing additional infrastructure. Likewise, a permanent application, or one that has specific requirements regarding the quality of service or the location of data, might best be deployed in a private cloud.

2.3.1 Private cloud

There are two different types of private clouds: On-premise clouds and externally hosted clouds. In on-premise cloud the cloud infrastructure is dedicated to and operated exclusively within a single organization, and managed by the organization or a third party regardless whether it is located on-premise or off-premise. Externally hosted private clouds are also exclusively used by one organization, but they are hosted by a third party organization. In most cases private clouds are more expensive, yet they are more secure when compared to public or hybrid clouds, because private clouds are

isolated from incoming external network traffic and this way they are less vulnerable. [2, 3; 5]

The motivation to setup a private cloud within an organization has several aspects. The first thing to consider is how to maximize and optimize the utilization of the on-premise recourses. Second, security concerns including data privacy and trust also make private cloud an option for many firms. Third, organizations must think about the data transfer cost from a local IT infrastructure to a public cloud. Fourth, organizations always require full control over mission-critical activities that reside behind their firewalls. [6]

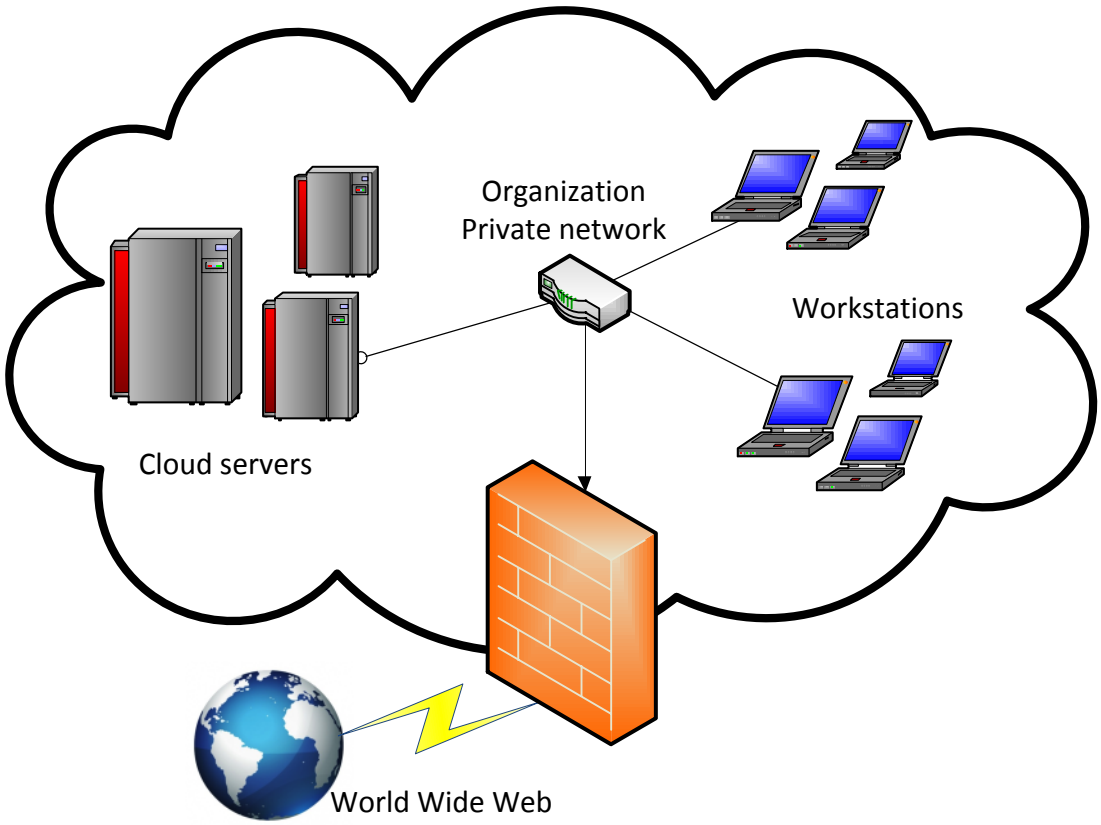

Figure 2. On-premise private cloud environment.

In figure 2 the cloud surrounding the private network presents a company's cloud environment and the local area network, when the whole network is secured by the firewall. An externally hosted private cloud environment is physically located in most of the cases at the vendor's premises, and it is only accessible from the subscriber's premises. Some of these networks might have a possibility to be accessed from outside of the

subscriber's premises, but this remote access requires a Virtual Private Network connection.

2.3.2 Public cloud

Private cloud is the most popular form of current cloud computing deployment models. In public cloud the computing infrastructure is hosted by the cloud vendor at the vendor's premises. Cloud infrastructures maintenance, upkeep and modifications are done by cloud vendor and the client has no visibility or control of the hosted computing infrastructure. However, some of the public cloud service providers give the client an opportunity to choose geographical location for their cloud infrastructure. The biggest difference between public and private is that the computing infrastructure is shared between different organizations. [2, 3; 5]

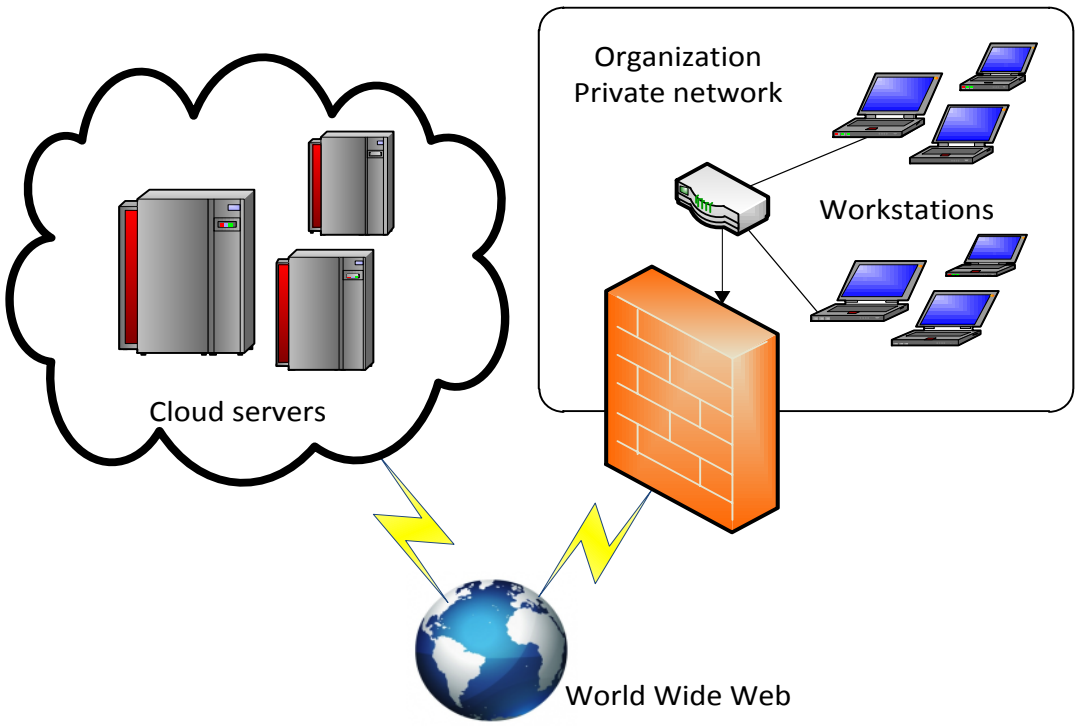

Figure 3. Public cloud environment.

Figure 3 shows how a cloud service can be separated from company's local area network. The idea of a public cloud model is that services running in the server are made available for everyone within the internet.

2.3.3 Hybrid cloud

The cloud service is a combination of two clouds, private and public, which remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds). In cloud bursting an organization uses their own computing infrastructure for normal usage, but has access to the public cloud for peak load requirements as well. This ensures that a sudden increase in computing requirement is handled gracefully. Organizations use the hybrid cloud model in order to optimize their resources to increase their core competencies by margining out peripheral business functions onto the cloud while controlling core activities on-premise through private cloud. [2, 3; 5]

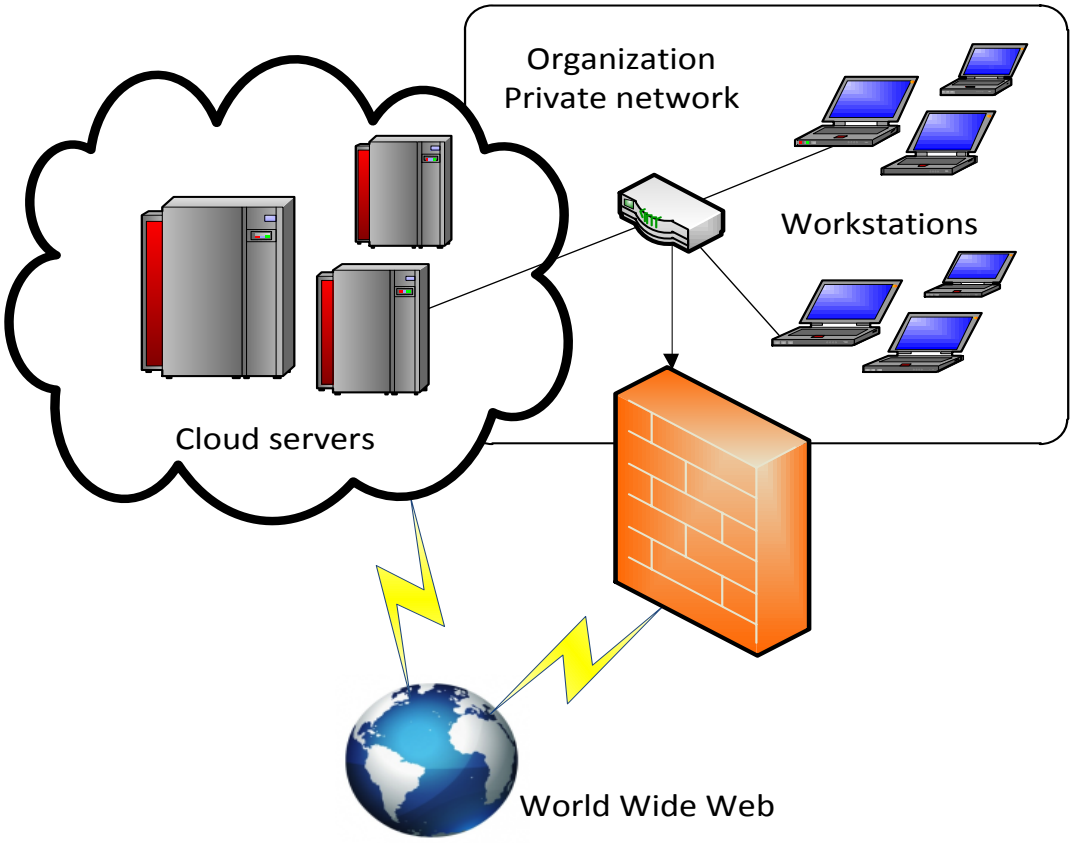

Figure 4. Hybrid cloud environment

In figure 4 the cloud depicts a cloud service which is partially separated from a company's local area network. The benefits of the hybrid cloud model are that one can decide what data applications should reside within, which should be run in the internal private cloud and which should be moved to a public cloud. This will minimize resource overcapacity and balance critical applications and data usage within the private cloud.

Peak-loads and less critical apps/data should be located into the public cloud instead of private.

2.4 Security issues and concerns

Most people are using the internet daily. Certainly many people also store personal data there, yet, some of the data may be very sensitive for the owner. That is why the need for secure storage when managing and sharing large amounts of data is critical. The main security challenge with clouds is that owner of the data may not have control of where the data is physically located. This is because if one wants to exploit the benefits of using cloud computing, one must also utilize the resource distributing provided for clouds. However, information must be safeguarded in the midst of untrusted processes. [7]

Cloud computing is fraught with security risks according to the analyst firm Gartner [7]. Even though cloud computing can offer significant cost savings for small- and medium sized businesses; the service does come with certain security risks. Below there is a list of some security issues that should kept in mind when choosing the potential providers for the cloud-based solution or service [8].

- Secure data transfer. All traffic between local network and the cloud service must travel over the internet. Also one should only use secured channels for data transfer. Data should always be encrypted and authenticated using industry standard protocols, which have been developed specifically for protecting internet traffic, such as IPsec (Internet Protocol Security).

- Secure software interfaces. The Cloud Security Alliance [8, 9] recommends being aware of the software interfaces, or Application Programming Interfaces (API) that are used to interact with cloud services. CSA recommends learning how any organization which can be considered for a cloud provider integrates security throughout its service, from authentication and access control techniques to activity monitoring policies.

- Secure stored data. Data should always be encrypted in the provider's servers. Make sure how the cloud provider secures the data. Before signing an agree-

ment, remember to make sure how the provider disposes the data and the key for encryption.

- User access control. Stored data on the cloud provider's server can potentially be accessed by an employee of the provider, and you do not have personnel controls over those people. First, consider carefully the sensitivity of the data allowed into the cloud. Second, follow the suggestion of the research firm Gartner [7, 2] to ask providers for detailed information about the people who manage your data and the level of access they have to it.

- Data separation. A cloud-based service always shares resources, namely the space on the provider's servers and other parts of the provider's infrastructure. Hypervisor software is used to create virtual containers on the provider's hardware for each of its customers. Cloud Security Alliance notes that "Attacks have surfaced in recent years that target the shared technology inside Cloud Computing environments" [8, 11]. One should investigate the isolation techniques, such as data encryption, which the provider uses to prevent access into your virtual container by other customers.

All cloud-based services have their weaknesses and threats, but security obstacles around cloud computing are being proactively addressed in many different ways. IT-industry creates partnerships to drive cloud computing standards and increase interoperability. Cloud computing alliances are forming to introduce innovative technologies designed to capitalize on the insights provided through cloud computing and produce enhanced cyber security awareness at all layers of the IT stack. Combining interoperability standards with improved cyber tools will give the IT workforce the capabilities needed to safeguard information and add value to mission. [8]

2.5 Cloud computing reference cases

Since cloud computing is here to stay more and more industries are joining the vast virtual world. Cloud computing is used in many industries and for different reasons. The following chapters give a few examples of different industries that have gained from adopting the cloud computing.

2.5.1 Education

New York University (NYU) announced in 2010 that it is moving to Google Apps for Education, providing cloud-based email and collaboration tools to the entire university community. [9]

This move to Google Apps provides more than 60,000 students, staff and faculty at NYU with Google tools like Gmail, Calendar, Docs and Sites and will bring all 18 schools of NYU—including medical, dental and law schools—together on one uniform system. This project is estimated to save the University around $400,000 annually by eliminating the need to purchase, upgrade, and maintain on-premise mail servers and software licenses. [9]

In 2011 University of Jyväskylä in Finland introduced their new solution based on Google Apps for Education cloud service. With this solution students were given access to Gmail-based e-mail service, as well as electronic calendars, instant messaging, Docs, and Sites. These tools provided new opportunities for students to work in groups, and form networks in a way that was not able in the University's earlier systems. [10]

2.5.2 Media and entertainment

- Netflix, an internet subscription service for movies and TV shows is a good example of media industry. Because Netflix streams many movies and shows on demand, the company faces large surges of capacity at peak times. As Netflix began to outgrow its data center capabilities, the company made a decision to migrate its Web site and streaming service from a traditional data center implementation to a cloud environment. This move allowed the company to grow and expand its customer base without having to build and support a data center footprint to meet its growth requirements. [11]

- Multimedia software company Adobe Systems wanted to take advantage of the explosion in digital media and publishing platforms. The company's existing IT infrastructure with many physical servers was relatively slow to respond to new business demands, limiting the company's ability to scale cost-effectively and causing delays in responding to new business requirements. [12]

Working with IBM, Adobe Systems moved from a traditional server environment to a fully virtualized environment, with private cloud capabilities based on IBM Power servers and IBM PowerVM. Hundreds of standalone servers and business-critical SAP systems and Oracle databases were consolidated to just five IBM Power 770 servers running IBM PowerVM software. [12]

By moving into IBM cloud technology Adobe Systems estimates total project savings and cost avoidance in physical systems, license fees, administrative and maintenance costs, cooling and energy from their Infrastructure as a Service project will come to approximately $60 million over five years, and the IBM/SAP portion contributed to that number. They also optimized their infrastructure and simplified the landscape which makes the solution easier and more cost-effective to manage. Also the servers that previously took weeks to procure and set up can now be enabled in a matter of hours, greatly speeding the time to market for new offerings. [12]

2.5.3 Transportation

- An anonymous airline company was awarded a 20-year contract to design, construct and operate a new common-use air-cargo terminal at an international airport. The terminal will provide around-the-clock services for an annual throughput of approximately 2.6 tons while boosting the competitiveness of the airport. To meet these requirements, the company sought to implement robust IT systems under a tight schedule.

  Together with IBM they were able to create a concept that delivers a feature-rich solution for IT support and services, allowing the client to focus on its core business. It also helps the client to meet a tight deadline by preparing a new cargo terminal to open on schedule. The solution builds on years of experience and expertise as a reliable IT services provider offering design and data center services.

2.5.4 Banking

- A leading Chinese commercial bank wanted to differentiate itself from its competitors by launching new business services to meet changing customer re-

quirements. To achieve this objective, the bank needed to increase the responsiveness of its IT infrastructure—facilitating the rapid provisioning of new virtual servers for innovative services. [13]

Seeking to improve its agility in quickly launching new services, the bank decided that a significant upgrade was required for its computing environment. The organization wanted to improve flexibility and simplify workload deployment efforts, and looked for a solution that combined performance, efficiency, scalability and reliability. [13]

With IBM Cloud computing technology they accelerated workload deployment by 99 percent. By consolidating six legacy servers into two IBM Power servers, they reduced the cost and complexity of managing hardware resources. They also improved flexibility of their infrastructure, helping the business to react more quickly and effectively to changing market conditions and customer demands. [13]

3 Virtualization

Complexity can work its way into any IT infrastructure, driven by the rollout of new applications and unanticipated change. However, adding servers in response to each demand for new workloads drives the need for more datacenter space, power, cooling, network cabling, data storage and administrative resources. Such complexity leads to inefficiency. The answer is virtualization, which allows organizations to consolidate multiple operating systems and software stacks on a single platform.

Virtualization is being used by a growing number of organizations to reduce power consumption and air conditioning needs and trim the building space and land requirements that have always been associated with server farm growth. Virtualization also provides high availability for critical applications, and streamlines application deployment and migrations. Virtualization can simplify IT operations and allow IT organizations to respond faster to changing business demands.

Virtualization is a technology that enables not only the creation of a virtualized hardware platform but also the storage device or network resources. The virtualized hardware platform is called a Virtual Machine (VM) that can be used as a platform for various operating systems. Virtualization is bound to cloud computing, partly because it is a main technology to enable cloud computing, for example in IaaS model, virtualization is the main technology to enable the service.

Virtual machines can provide numerous advantages by installing operating systems and software directly into physical hardware. Isolation of the virtualization ensures that applications and services that run in a VM cannot interfere with the host operating system or other VMs. This will then ensure business continuity which is critical in production environments.

Virtualized service can bring benefits over the physical servers such as:

Cost savings:

- Virtualization decreases the need of the physical servers, hardware devices and data centers thus decreasing the fixed assets investments and mainte-

nance cost but also reducing the power consumption costs up to 40%. [3, 13]

Increase in administrator productivity:

- Administrators have comparatively fewer emergencies to take care of as virtualization takes care of many such situations (e.g. hardware failures, etc.) itself. [14,9]

Flexibility and agility:

- Reducing time-to-market of applications, for example by reducing development and testing, enabling rapid deployment of resources to new locations / offices (for market expansion)

Business continuity and disaster recovery:

- Duplication of critical servers, data, etc. is easier and less expensive, physical duplicates are not required to be maintained for disaster recovery purposes.

- As VM are portable they can run from anywhere, especially in case of a disaster recovery situation.

3.1 Virtualization in the market

The rise of virtualization in IT environments started in the 1960s and in the present day the use of virtualization has gotten more common in both small and large organizations in the way of test environments. Virtualization products have also become easily attainable because of the market growth. Virtualization can be used to decrease the number of physical machines and so increasing the cost-effectiveness. [15]

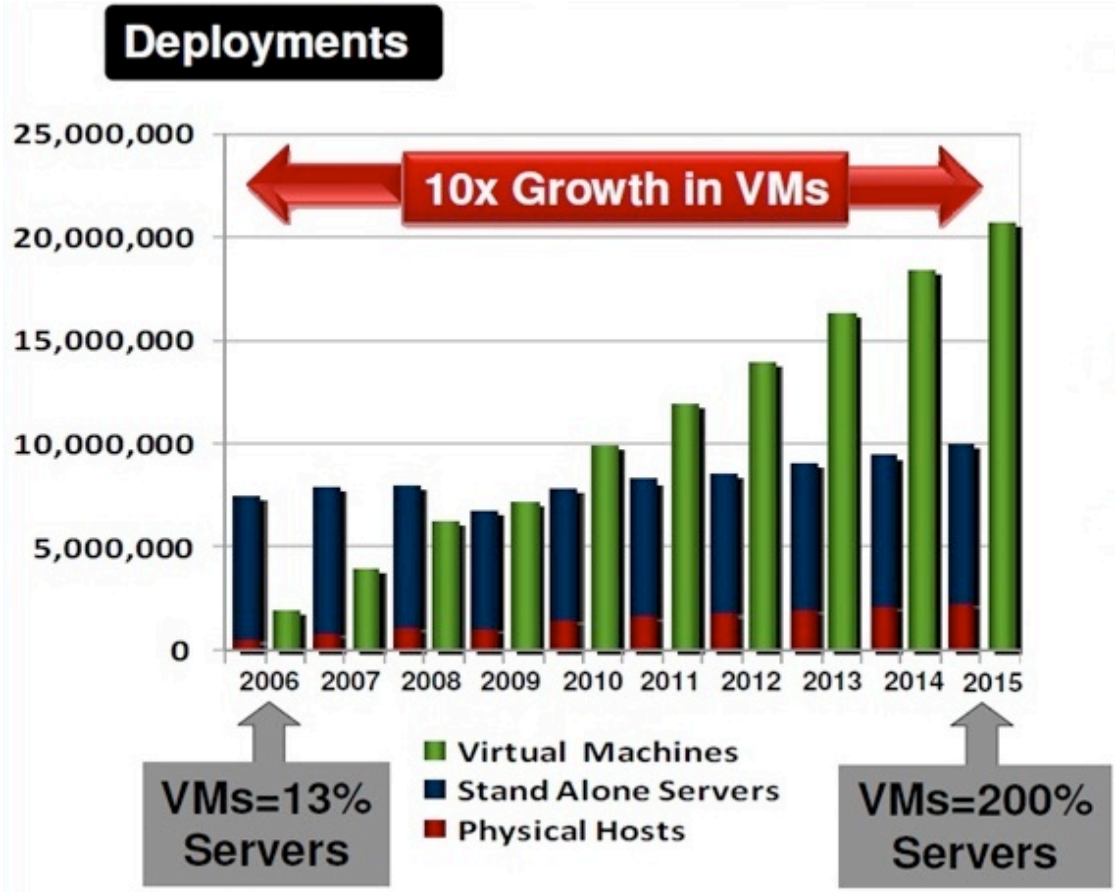

Figure 5. Server virtualization growth from 2006 to 2015. [16]

According to figure 5, the new server shipment will grow over the years as VM deployment will stay stable. IDC also predicts that VM densities will increase 25% from 2010 to 2015 and the rapid increase in VMs deployed versus physical machines will drive customers to rethink systems management tools and automation strategies.

Worldwide virtual machine shipment share from 2010 is presented below:

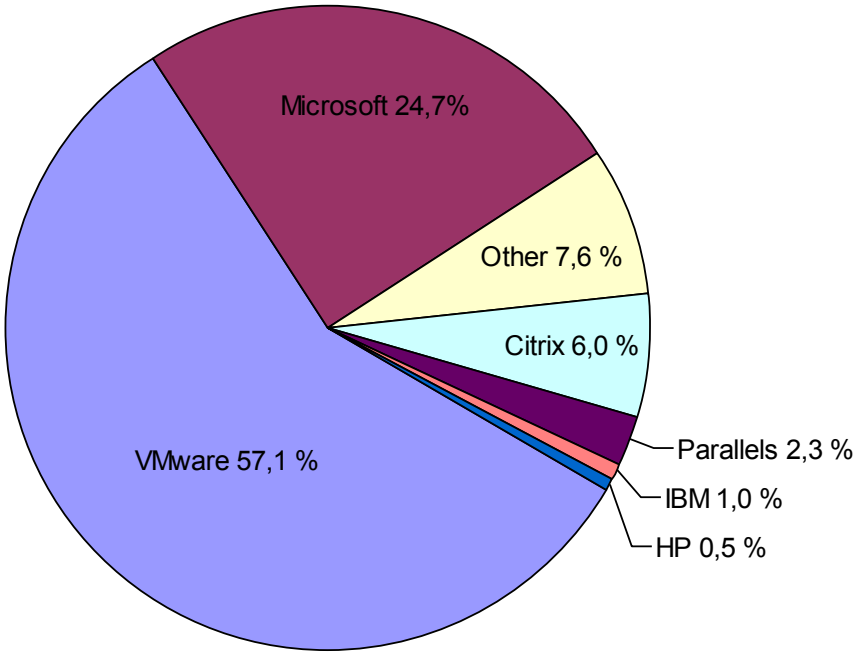

Figure 6. Worldwide virtual machine shipment share by vendor, 2010 [17]

According to figure 6, VMware dominates the virtualization market and they still continue to show strong growth. Although, the other virtual machine software vendors are solid competitors against VMware, head start of the VMware was devastating for others. Today VMware exploits their market position by commanding premium pricing.

3.2 Cloud hardware virtualization technologies

Cloud hardware virtualization technologies also known as hypervisors are the key enabling technology for virtualization. Hypervisor is a program that allows multiple VMs to share a single hardware host. Each operating system has its own portion of the host's processor, memory, storage, and other resources. However, the hypervisor is actually controlling the host processor and resources, allocating what are needed to each operating system in turn and making sure that the VMs cannot disrupt each other.

Hypervisors can be separated into two different types. Below two hypervisor types are presented:

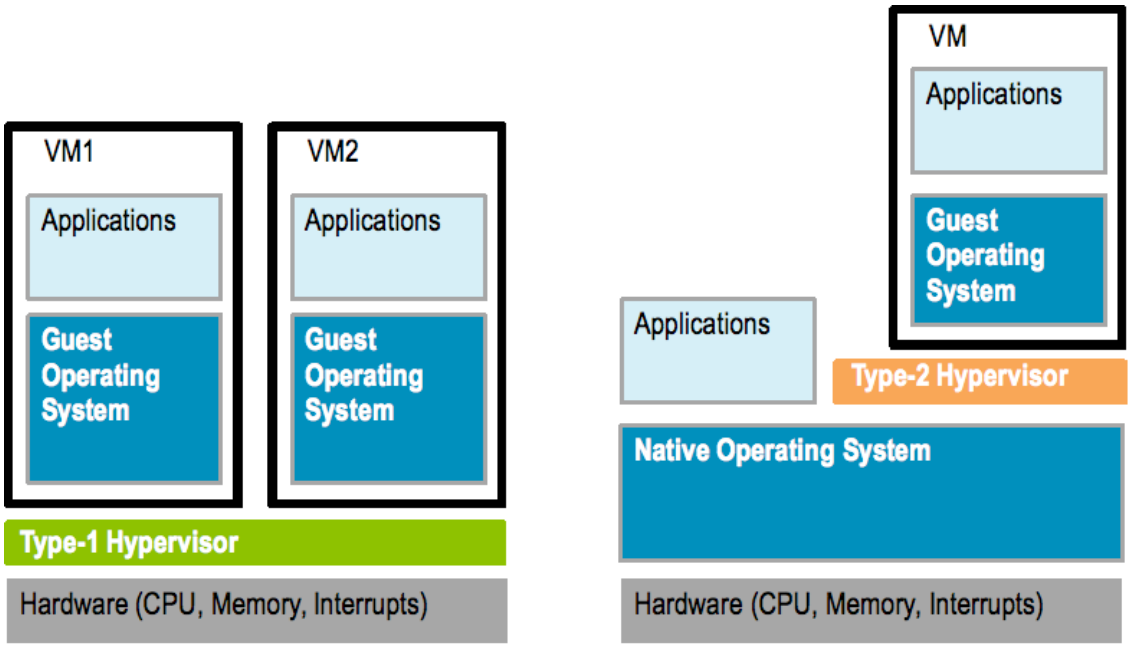

Figure 7. Two different hypervisor classes. [18]

As shown in figure 7, type 1 hypervisors run directly on the host's hardware as the operating system. It controls the hardware and manages the guest operating systems. A guest operating system thus runs on another level above the hypervisor. Modern equivalents of this are the Citrix XenServer, VMware ESX/ESXi, and Microsoft Hyper-V hypervisor. [18]

Type 2 hypervisors run within a native operating system environment. With the hypervisor layer as a distinct second software level, guest operating systems run at the third level above the hardware. Kernel-based Virtual Machine (KVM), Berkeley Software Distribution (BSD) Hypervisor, and VirtualBox are good examples of Type 2 hypervisors. [18]

3.2.1 VMware ESX/ESXi

VMware [19] is global company providing numerous virtualization and cloud computing solutions. ESX (Elastic Sky X) [20] and ESXi (Elastic Sky X integrated) [20] are both an enterprise-level computer virtualization products offered by VMware. However, ESX is replaced with ESXi. ESX and ESXi are both type 1, bare-metal hypervisors [see figure

6] that install directly on the server hardware. Both provide industry-leading performance and scalability; the difference resides in the architecture and the operational management of VMware ESXi.

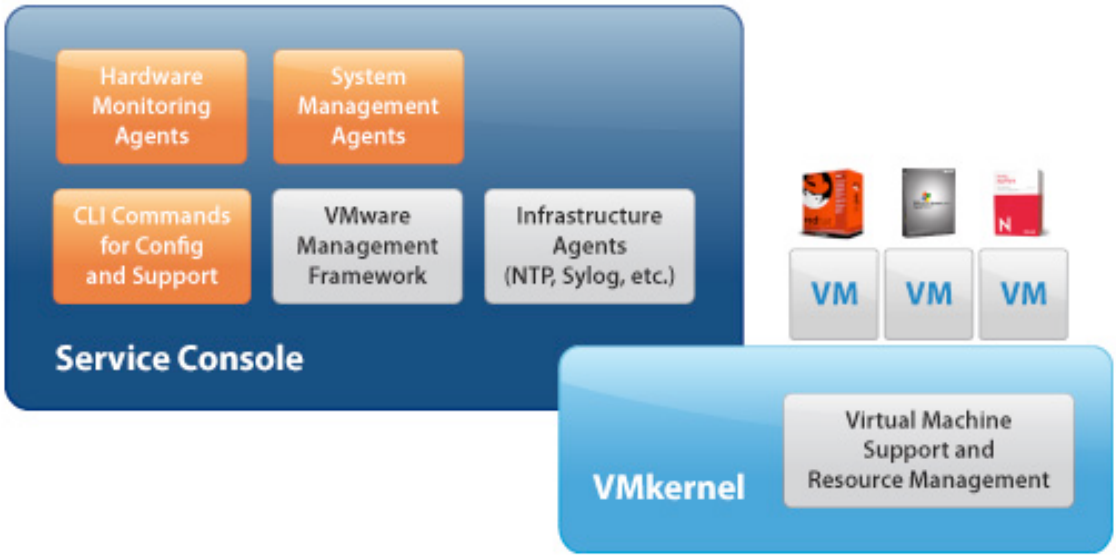

Figure 8. VMware ESX architecture. [21]

As seen in the figure above, although neither hypervisor relies on an OS for resource management, ESX relies on a Linux operating system, called the service console, which is a stripped version of Linux OS (Red Hat Enterprise Linux) that is used for serviceability and partner integration as shown in figure 8. In other words, service console is needed to perform two management functions such as executing scripts and installing third party agents for hardware monitoring, backup or systems management. [21]

The kernel is launched first and it is used to load a variety of specialized virtualization components, including VMware's vmkernel component. This previously-booted Linux kernel then becomes the first running VM and is called the service console. [22]

The architecture of ESXi is shown below:

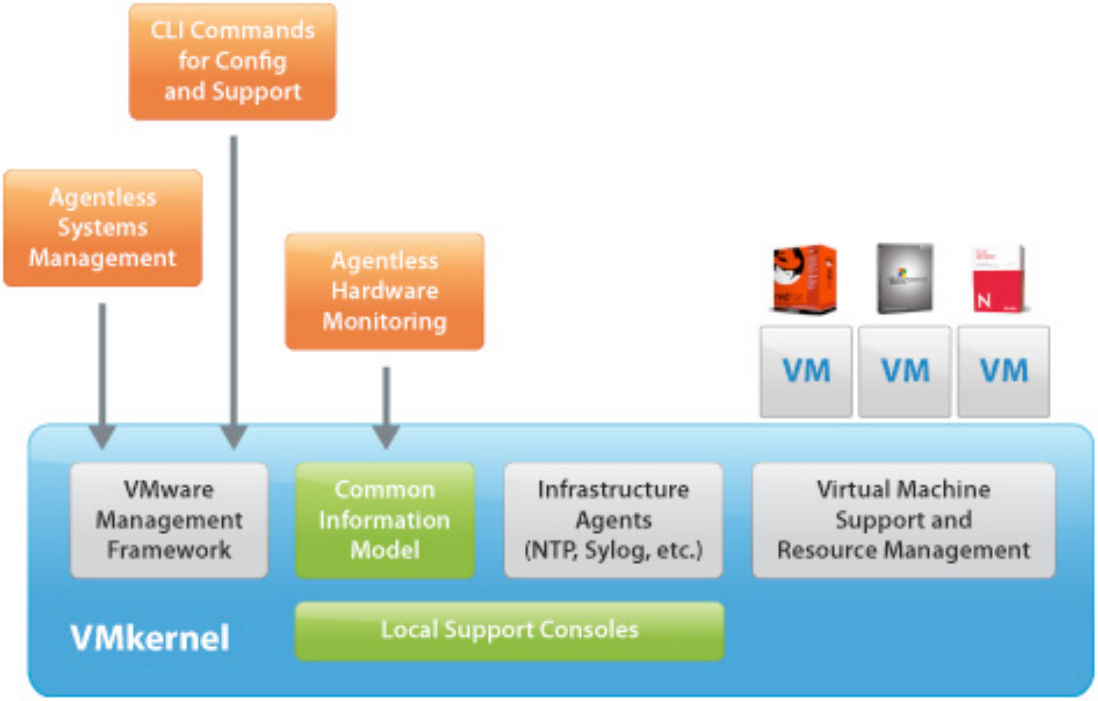

Figure 9. VMware ESXi architecture. [21]

Figure 9 shows how ESXi is very similar to ESX – in fact they have the same underlying code base. The smaller code base of ESXi represents a smaller “attack surface” and less code to patch, improving reliability and security. The functionality of the service console is replaced by remote command line interfaces and adherence to system management standards. The service console has been removed from ESXi, drastically reducing the hypervisor footprint and completing the ongoing trend of migrating management functionality from the local command line interface to remote management tools [22]

Partner integration is API-based in ESXi instead of agent-based as was the case in ESX. This allows us to have a small and stateless architecture that is very easy to deploy and manage. Below, some benefits of the ESXi architecture are presented. [22]

Based on personal experiences, ESXi is the thinnest hypervisor in the market with less than 100MB of code base disk footprint, even if compared with the ESX which is roughly 2GB. A few important benefits have been achieved by modifying the code base, such as: [21]

- Improved security. The old ESX architecture is relied on a Linux-based Console Operating system (COS) for serviceability and agent-based partner integration. In the new, operating-system independent ESXi the COS has been removed and the necessary management functionality has been implemented directly in the core VMkernel. By removing COS reduces the install footprint of the ESXi hypervisor to approximately 150 MB improving security and reliability by removing the security vulnerabilities associated with a general purpose operating system.

- Less patching. Due to its small size and limited components, the ESXi architecture requires fewer patches than early version. Over its lifetime, the ESXi architecture requires approximately 10 times less patches than the ESX hypervisor running with the COS.

- Simpler configuration. The new ESXi architecture has less configuration items simplifying deployment and configuration and making it easier to maintain consistency.

The second main result of removing the service console is that everything that used to be in there mainly for the purpose of management, monitoring and scripting are removed. Communication between software and hypervisor is done by built-in APIs to connect these features to the hypervisor.

3.2.2 KVM

KVM (kernel-based virtual machine) is a young hypervisor to enter the virtualization market, and it joins a market that already has no fewer than four other major solutions in place. Unlike the Type 1 hypervisors, it is embedded into an operating system (OS), namely Linux. However, the implementation of KVM into the OS is unique in several other ways. [23]

KVM is a loadable kernel module within the Linux kernel that allows the Linux operating system to function as a Type 1 bare metal hypervisor. KVM have its advantages compared to other hypervisors, there are some certain things why KVM can be praised, such as: [23]

- Given its open source nature, KVM has a lower total cost of ownership [see figure 10].

- Rapid progress to maturity: A community of experts continuously enhances KVM.

- Exploitation of advances in Linux: KVM is built into Linux and benefits from the entire Linux community.

- Efficiency: KVM takes advantage of modern hardware design to securely execute directly on the host CPU, and is engineered to perform well even in memory- and CPU-constrained environments.

- Truly open source: The code and its repository data are available, continuously inspected, and transparent in modification rationale throughout the product life cycle.

KVM is originally supported by x86 and x86-64 processor architectures, but it has been ported to S/390, PowerPC, and IA-64, and also ARM port is in progress.

IBM has chosen KVM as a strategic hypervisor for x86 based systems. Choosing a primary hypervisor for cloud platforms was not an easy task for IBM. Undoubtedly, several proprietary and open source hypervisors provide sufficient virtualization capabilities. However, a few have the qualities required for to become the stand-out choice for IBM's cloud offerings. Once hardware acceleration was introduced into the Linux kernel in 2007, KVM quickly became the best choice. A number of factors solidify the preference for KVM. [24]

By embracing KVM, IBM is continuing its tradition of hypervisor excellence. IBM helps to lead KVM development, while implementing the features that the clients demand. In addition, IBM uses KVM for workload consolidation and in their cloud offerings. [24]

3.2.3 Hyper-V

Microsoft Hyper-V is type-1 server virtualization hypervisor product that enables one to consolidate workloads, helping organizations improve server utilization and reduce costs.

Hyper-V Server is a dedicated stand-alone product that contains the hypervisor, Windows Server driver model, virtualization capabilities, and supporting components such as failover clustering, but does not contain the robust set of features compared to other hypervisors. [25]

Hyper-V is generally preferred over any other hypervisor, mostly because it does not require special hardware to operate properly. It can basically operate with any hardware that Windows software is installed on. It is also laptop friendly, which means the Hyper-V VPS (virtual private server), is portable and can be transported and used essentially anywhere. [25]

3.2.4 Xen

Xen hypervisor is infrastructure virtualization solution that supports a wide range of guest operating systems including Windows, Linux, Solaris, and various versions of the BSD operating systems. Xen also powers many hosting services and most public cloud services such as Amazon Web Services, Rackspace Hosting and Linode. [26]

Xen hypervisor is open source technology like KVM, but it is type-1 hypervisor unlike KVM. Xen is developed collaboratively by the Xen community and engineers of the data center solution vendors, including AMD, Cisco, Dell, Fujitsu, HP, IBM, Intel, Red Hat, Samsung, and many more. [26]

3.2.5 PowerVM

PowerVM is the virtualization solution for AIX, IBM i, and Linux environments on IBM POWER technology. PowerVM offers a secure virtualization environment, built on the advanced Reliability, Availability, and Serviceability (RAS) features and performance of the Power Systems platform.

3.3 Choosing between VMware or KVM

As the VMware is dominating the virtualization industry and is using its advantage for premium pricing, many of the clients would gladly migrate away from VMware's products because of cost saving purposes. After VMware changed their licensing procedure, the IBM has tried to prefer KVM rather than VMware. After all, the final decision for which hypervisor technology is used in cloud environment is always based on the client needs.

The study conducted by Clabby Analytics "VMware vs. KVM: A Functionality/Price Discussion" [27] introduces the differences in pricing between VMware and KVM. [See table 1].

Both VMware and KVM can offer similar functionalities when it comes to performing the functions required by most IT managers (the ability to build, deploy, and manage virtual machines). However, KVM can cost up-to 40% less to implement and operate as compared with VMware. Service and support fees are generally similar when comparing VMware to KVM (when using the Red Hat implementation), but VMware's first year license fees are significantly higher and are primarily responsible for the large difference in cost. [27]

The chosen server configuration for the research to illustrate VMware vs. Red Hat and Red Hat/IBM consisted of a 20 server, 2-socket server Linux environment. The license count/price data for this particular configuration were gathered directly from the vendor websites or from the vendors themselves. [27]

Document notes that:

- Prices in figure 10 are list prices — not the actual "street" prices.

- License counts fluctuate between the products being compared. This is because the licensing terms are different between each product – some are charged on a per socket basis with memory limits, some are charged on a per server basis.

- Quantity required license counts doubled between second and the third year. This is because buyers were paying for two years worth of license/support fees (as compared to a single year).

Year 1	List Price	Quantity reguired	License &/or SnS cost
Configuration 1: A VMware Environment on Red Hat Linux			
VMware vCenter Server 5 Standard (licence)	6 244	1	6 244
VMware vShpere 5 Enterprise Edition (License)	3 594	40	143 760
Red Hat Ent. Linux 2-sockets (unlimited VMs 1 yr. Premium)	3 249	20	64 980
			214 984
Configuration 2: A Red Hat Virtualization/KVM Environment			
Red Hat Ent. Linux 2-sockets (unlimited VMs 1 yr. Premium)	3 249	20	64 980
Red Hat Enterprise Vitualization for Servers (1 yr. premium)	749	40	29 960
			94 940
Configuration 3: A Red Hat Linux Environment w/IBM Sys Dir and Vmcontrol			
Red Hat Ent. Linux 2-sockets (unlimited VMs 1 yr. Premium)	3 249	20	64 980
IBM Sysmtes Director Std Edition (x86 Support & virt. Media key)	450	20	9 000
IBM Vmcontrol Enterprise Edition x86 (w/1 year SW maintenance)	1 400	20	28 000
IBM SmartCloud Provisioning Resource Value Unit License + SW subsciption	125	20	2 500
			104 480
SnS Costs in Year 2 and 3			
Configuration 1: A VMware Environment on Red Hat Linux			
VMware vCenter Server 5 Standard (licence)	1 249	2	2 498
VMware vShpere 5 Enterprise Edition (License)	719	80	57 520
Red Hat Ent. Linux 2-sockets (unlimited VMs 1 yr. Premium)	3 006	40	120 240
			180 258
Configuration 2: A Red Hat Virtualization/KVM Environment			
Red Hat Ent. Linux 2-sockets (unlimited VMs 1 yr. Premium)	3 006	40	120 240
Red Hat Enterprise Vitualization for Servers (1 yr. premium)	693	80	55 440
			175 680
Configuration 3: A Red Hat Linux Environment w/IBM Sys Dir and Vmcontrol			
Red Hat Ent. Linux 2-sockets (unlimited VMs 1 yr. Premium)	3 006	40	120 240
IBM Sysmtes Director Std Edition (x86 Support & virt. Media key)	90	40	3 600
IBM Vmcontrol Enterprise Edition x86 (w/1 year SW maintenance)	280	40	11 200
IBM SmartCloud Provisioning Resource Value Unit License + SW subsciption	115	40	4 600
			139 640
Total 3 Year Cost of Ownership			
VMware with vSphere and vCenter			**$ 395 242**
Red Hat			**$ 270 620**
Red Hat w/IBM System Director and VMControl			**$ 244 120**

Figure 10. Comparative Acquisition and Service-and-Support Costs Over 3 Years [27]

The most significant finding in figure 10 is that first year start-up costs for VMware were comparatively high ($214,984) — but dropped $90,213 annually during the second and the third year. By comparison, Red Hat subscription costs dropped to $89,630 per year during the second and the third year. Meanwhile, the cost for the IBM Systems Director/IBM Systems Director VMControl and SmartCloud Provisioning solution drops to $69,810 per year during the second and the third year. [27]

This data illustrates that a 20 server VMware solution costs approximately 40% more than the lowest cost KVM solution (the IBM Systems Director/VMControl/SmartCloud Provisioning solution). This is a large price premium for products that are functionally equivalent. VMware users might claim that Vmware's products are deeper, better integrated and simpler to use than products from Red Hat and the KVM ecosystem, this might be true. However, KVM users can virtualize, provision, and manage their environments just as effectively using KVM ecosystem products. Even VMware offers a lot of product extensions and nice-to-have things, but are these worth a 40% price premium?

Although, VMware products have good functionality and they are generally very well integrated with one another. As research shows, KVM and its ecosystem can manage most if not all of the same tasks as VMware such as: virtualization, provisioning, workload management, and etc. but at a far lesser cost.

4 IBM technology to build cloud environment

IBM aims to bring customers infrastructure, software and platform together in one solution to help customer to save time and money. All offerings are designed for business use and in order to solve some particular problem IBM has lots of different solutions from infrastructure to accounting and managing cloud computing systems.

In addition to software solutions IBM has several hardware solutions that are designed to work as a platform for cloud solution. These hardware solutions are usually powered with x86-64-bit processor architecture, but also with more unique processor architectures such as Power Systems and System Z, both of these architectures are designed by IBM.

4.1 Practical cloud computing project

I was given a chance to participate in the largest cloud computing project in IBM carried out in Finland so far. In this project IBM’s customer wanted to simplify their complex infrastructure and enable resource sharing by using cloud-computing solution based on IBM Power Systems servers, IBM PowerVM technology and IBM Systems Director and IBM Service Delivery Manager software to optimize their resource use, decrease time to market and increase flexibility of the virtual image deployment.

This solution was first made available for customer as a proof of concept. A proof of concept (POC) is a demonstration if the product or technology is viable and capable of solving an organization's particular problem. POC is, therefore, a prototype that is designed to determine feasibility, but does not represent deliverables. To enable POC customer provided two existing Power servers from their data centers as the platform for the service.

Power servers were located in different of the customer’s two data centers, whereof the other is mirrored data center. Mirroring is a function that replicates all new data from primary data center into the mirrored one, therefore mirrored data center should be the exact copy of the primary one. In fail state Power servers’ automatic-deployment features will shut the failed server after the mirrored server has taken the place upon the failed one, ensuring business continuity.

We installed IBM PowerVM Enterprise Edition technology to virtualize the IBM Power Systems servers and IBM Systems Director VMControl software to manage the virtualized environment. Customized IBM Service Delivery Manager-software were chosen for a cloud management platform. The software delivers a preconfigured, automated service-management solution to oversee cloud deployment and provision resources in virtual environments.

By implementing a cloud-computing environment with these technologies, the client successfully simplified their infrastructure and enabled resource sharing. The Power Systems servers helped improve system stability and ensure business continuity by providing multiple redundant server resources for use in case a business system fails.

In the end the client lowered total cost of ownership by improving capacity use by 75 percent, reduced IT labor costs by 50 percent, Enabled deployment of a self-service cloud, increased flexibility by gaining support for multiple platforms and hypervisor technologies and increased scalability and agility for responding to business demands.

4.2 Technology and software behind the project

I have listed here commercial software and hardware solutions that are used to build cloud computing environment for the previous project. All the software and hardware solutions in the project were based on the Power Systems.

Briefly Power Systems is the name of the IBM's Power architecture-based server portfolio. Power Systems processor architecture is designed by IBM just like x86 was designed by Intel. The starting point is that operating systems used in x86 environment such as Windows cannot be used on Power Systems and operating systems used for Power Systems such as AIX cannot be used in x86 systems. However, based on my own experiences Power Systems has many unique features compared to x86 for example the features of the PowerVM.

4.2.1 PowerVM

In the world of Power systems the term VM or Virtual Machine has been replaced with a different name LPAR (Logical Partition). Fundamentally LPAR is based on the same

idea as VM but with the difference that it works on the different infrastructure architectures.

PowerVM's dynamic logical partitioning (LPAR) allows a single partition to act as a completely separate AIX, IBM i, or Linux operating environment. Logical partitions can have dedicated or shared processor resources. With shared resources, PowerVM can automatically adjust pooled processor resources across multiple operating systems, borrowing processing power from idle partitions to handle high transaction volumes in other partitions.

PowerVM's feature Micro-Partitioning supports up to 20 dynamic logical partitions per processor core. Depending on the Power server, up to 1024 independent virtualized servers can be run on a single physical Power server — each with its own processor, memory, and I/O resources. These partitions or workloads inside the LPAR can be assigned at a granularity of 1/100th of a core. Consolidating systems with PowerVM can reduce operational costs, improve availability, ease management and improve service levels, while allowing businesses to quickly deploy applications.

When Power Systems have been designed and manufactured, the PowerVM feature has been partly built into the Power Infrastructure. To enable a PowerVM feature, the Power server must have Virtual I/O Server installed into a separate LPAR. VIOS owns the resources that are shared with clients. A physical adapter assigned to the VIOS partition can be shared by one or more other partitions. VIOS is designed to reduce costs by eliminating the need for dedicated network adapters, disk adapters and disk drives, in each client partition.

In most cases software licenses are sold and assigned per core and the same goes for Power Systems. If you purchase 100€ license for one Power processor core and then "virtually" divide it into maximum of 20 pieces, with fast mathematics you have paid 5€ per one virtual computer or LPAR in Power language. As the result of the study, in the end the customer was able to reduce the total cost of ownership and increase agility and scalability of the environment with PowerVM technology.

4.2.2 IBM System Director and VMcontrol

IBM Systems Director supports the PowerVM environment and is the IBM management tool for multiple, heterogeneous servers. IBM Systems Director supports advanced management functions such as workload lifecycle management, health check, and topology mappings, as well as the ability to take action on monitored events.

IBM Systems Director VMControl is a plug-in option for IBM Systems Director that represents a transformation from managing virtualization to using virtualization to better manage an entire IT infrastructure. IBM Systems Director and VMControl are designed to help reduce the total cost of ownership in a virtual environment by decreasing management costs, increasing asset utilization, and linking infrastructure performance to business goals.

VMControl is available in three editions, to suit the varying levels of virtualization deployment at client sites. These three editions are Express, Standard and Enterprise and the Express was the one used in the project. VMControl Express Edition provides only basic virtual machine lifecycle management.

Because of the automation of the Systems Director and VMcontrol the customer was able to save IT-labor costs by 50%.

4.2.3 IBM Service Delivery Manager

IBM Service Delivery Manager is software for enterprise customers who want to get started with a private cloud computing model. The software enables the user to rapidly implement a complete software solution for service management automation in a virtual data center environment, which in turn can help the organization move towards a more dynamic infrastructure. It is delivered at a pre-integrated software stack, deployed as a set of virtual images that automate IT service deployment and provide resource monitoring, cost management, and provisioning of services in a cloud. [28]

IBM Service Delivery Manager is ideal for organizations that want to get started with a private cloud computing model. The product enables rapid implementation of a complete software solution for service management automation in a virtual data center environment, which in turn can help organizations move towards a more dynamic infrastructure. [28]

IBM Service Delivery Manager includes several functionalities such as: [28]

- A self-service portal interface for in-advance computing reservations of virtualized environments, including storage and networking resources
- Automated provisioning and de-provisioning of resources
- Real-time monitoring of physical and virtual cloud resources
- Integrated usage and accounting chargeback capabilities that can help system administrators to help, track, and optimize system usage
- Built-in high availability of the cloud management platform
- Prepackaged automation templates and workflows for most common resource types

In the project described in this thesis, IBM Service Delivery Manager (ISDM) were in a key position. With ISDM, the customer was able to customize unique self-service portal where users were able to request any of the software from the catalog or optionally just an empty logical partition. The result of this study indicates that with ISDM the customer was able to increase flexibility of the service.

5 Business value through cloud computing

According to a Gartner Research report, "Reimagining IT: The 2011 CIO Agenda." cloud computing ranks No. 1 on the CIO technology priority list for the year 2011. The report is based on survey responses from 2,014 CIOs in 50 countries, as well as interviews conducted by Gartner analysts with leading CIOs. In fact, Gartner predicts an enterprise boom in cloud computing adoption in the coming years. Mobile technologies and virtualization are also top priorities for CIOs. These solutions dovetail with the top three business drivers: Increasing enterprise growth; attracting and retaining new customers; and reducing enterprise costs. [29]

CIO technologies	Ranking of technologies CIOs selected as one of their top 3 priorities in 2011			
Ranking	***2011***	**2010**	**2009**	**2008**
Cloud computing	1	2	16	*
Virtualization	2	1	3	3
Mobile technologies	3	6	12	12
IT management	4	10	*	*
Business intelligence (BI)	5	5	1	1
Networking, voice and data communications	6	4	6	7
Enterprise applications	7	11	2	2
Collaboration technologies	8	10	5	8
Infrastructure	9	14	7	6
Web 2.0	10	3	15	15

*New response category

Figure 11. Cloud becomes top CIO technology priority in 2011 [29]

As can be seen from figure 11, cloud computing has been a new topic among the CIOs in 2009. Because of the hype around cloud computing, this topic ranked as the second most discussed technology in 2010.

There are numerous fields that most of the clients in every industry are trying to achieve with cloud computing. The most obvious one is that every company wants to do better business with less money, without any effects on the quality of service. However most of the customers in various industries are facing also other challenges such as: [30]

- Global operations and competition

- Size of the data and information
- Data center costs
- Customer expectations

Currently, competition is so intense that companies have to stand out from one another and while everyone wants to do more with less, quality of service will separate competitors from one another. As cloud computing is quite recent topic in markets and because of this some of clients are not ready to take risks to achieve these benefits such as cost effectiveness. However the distributed hardware independent datacenters mean fewer outages and the risk of downtime, also newest storage systems ensures that your data will always be retained. Below there is listed some acclaimed benefits of the cloud computing technology for small- and medium-size businesses: [30]

- Low start-up costs
- Low cost for sporadic use
- Ease of management
- Scalability
- Device and location independency

Cloud computing keeps start up expenses low by eliminating the need for expensive servers, software and full-time IT professionals. Almost fully automated management also cuts a part from expenses, making server managing much easier. Cloud computing can also help companies to become more productive and mobile. With increased mobility, employees can access into cloud from anywhere within the internet connection. This improves productivity of employees who need to travel and work outside the office while meeting clients. Also the information is accessible by all authorized employees. [31]

5.1 Cloud computing in the market

As shown in figure 12, in the last quarter of 2011 HP took home 17.4 percent of worldwide cloud hardware revenues. IBM was close behind with 15.9 percent. Overall, IT vendors booked $39.4 billion in cloud equipment revenues last year, a 15 percent increase over 2010. In the fourth quarter alone, they brought in a healthy $10.6 billion haul. [32]

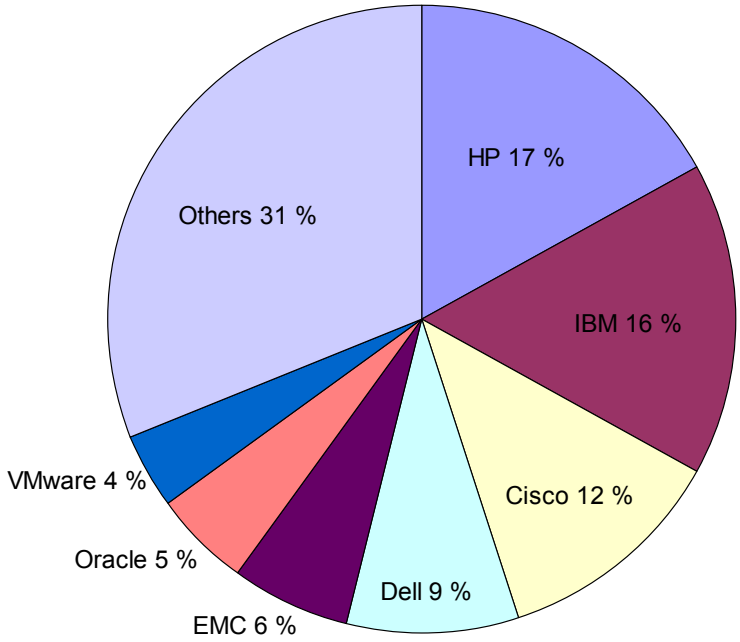

Figure 12. Cloud equipment market share 2011. [32]

Jeremy Duke, Synergy Research Group's founder and chief analyst said (2011): “The vendors that provide infrastructure for cloud services have the opportunity to benefit greatly from this surge -- you don't see many $40 billion markets growing by 15% per year” [32]. The majority of that growth is coming from Europe, the Middle East and Africa (EMEA) and Asia-Pacific. North America still accounts for the largest share, however, with 43 percent.

5.2 Business in Finland

Although cloud computing should be as secure as standard computing and hosting, still many doubt this and sometimes the media reports that it is not. According to a study conducted by Finnish Information Technology Association, many CIOs in Finland have doubts about moving business-critical systems into cloud. Results of the research are shown in figure 13.

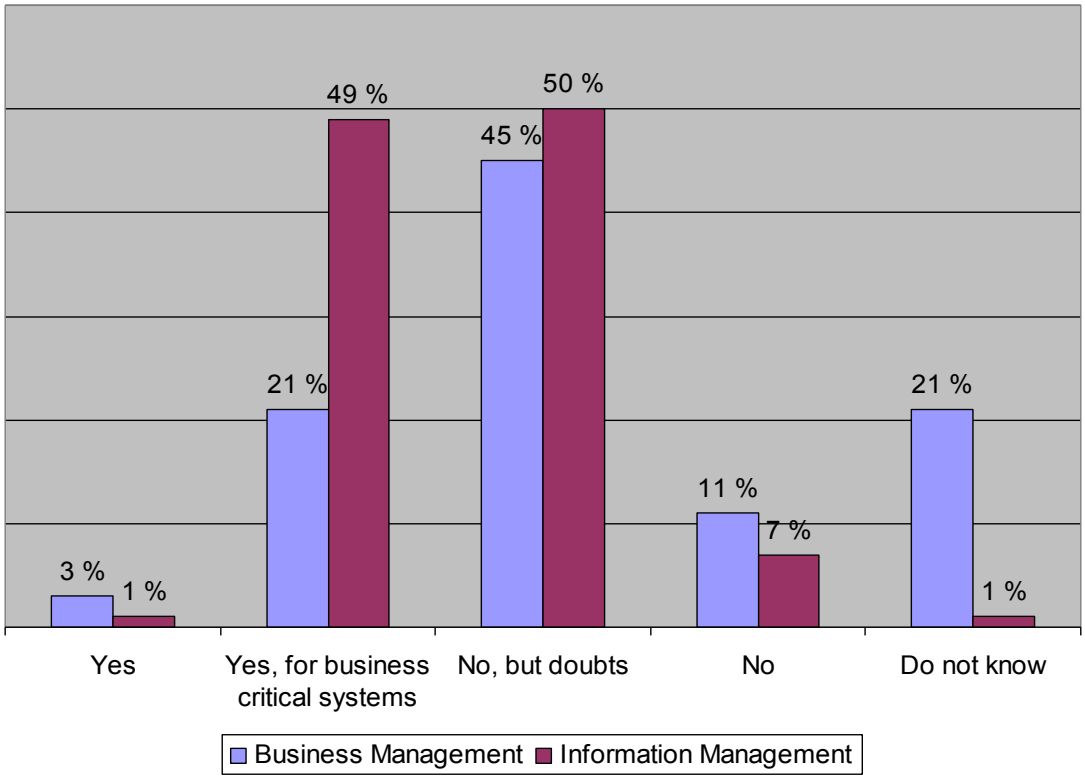

Figure 13. Is security an obstacle for buying cloud computing services? [33]

Every fifth of the senior vice presidents would not export business-critical systems to cloud, and 40 percent of the CIOs say no to cloud computing because of security concerns. However, 50 percent of the CIOs say that security problems are not an obstacle for non business-critical cloud systems. [33]

Below the progression of cloud computing between 2010 and 2011 are presented:

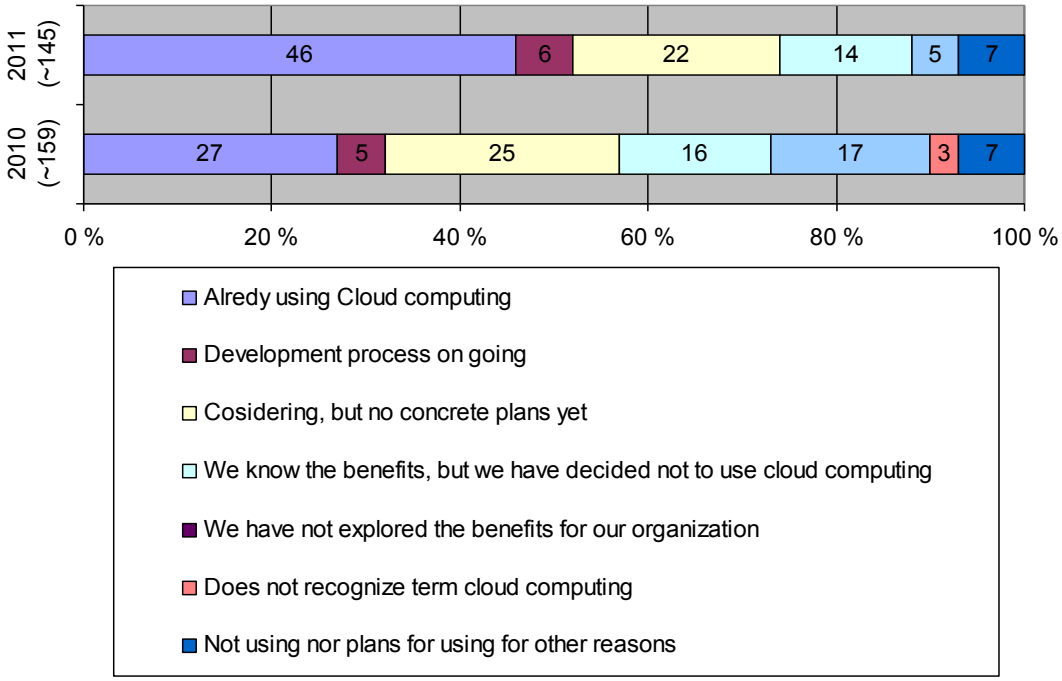

Figure 14. Progression of cloud computing in Finland in years 2010-2011 [34]

Demand of the cloud computing is rising all the time in Finland as shown in figure 14. However, 64 percent of large Finnish enterprises use cloud computing in their business according to Finnish IT-accosiation TTL (tietotekniikan liitto). Executive Director of TTL Robert Serén (2012) predicts cloud computting markets to triple during next three years. Serén and Sofigate's director Jari Raappana (2012) point that cloud computing in small and growing business can bring huge opportunities. [34]

6 Conclusion

Cloud computing is one of the most discussed emerging IT-technology today, it offers a variety of benefits but also creates an opportunity for new market opportunities. However, markets in the field of cloud computing are relatively new and still developing in Finland. They are led by large international enterprises such as IBM, Microsoft, Amazon and HP. These companies offer various services from computing infrastructure to smart software solutions. Today many of the small, medium and large-sized companies are interested in cloud computing because of the significant benefits this technology offers. However, there are a number of security risks and other concerns that should be kept in mind when choosing the potential provider for the cloud solution.

This study set out to analyze how cloud computing can help companies to cut costs and provide better service. In addition the following questions were answered: What is cloud computing? What kind of problems can be solved with cloud computing? What can be achieved by using cloud computing?

The answers were obtained by analyzing the client reference database of IBM, by interviewing IT-specialists in IBM and by project results. Also, studying scientific articles as well as evaluating personal experiences of those involved in the project played an important part during the writing process.

The NIST's (National Institute of Standards and Technology) presents a comprehensive set of cloud computing definitions that were discussed in this study. Those definitions include the following cloud deployment models: private, public and hybrid clouds, which are used to provide services to clients over the internet without a need to download or install any additional software. Also it seems that any internet service can be categorized into some of the service models: IaaS, PaaS, or SaaS.

Things concerning to security is always a good idea to take into account and some of the security risks were analyzed. Even though cloud computing is acclaimed to be a safe opinion; this might not always be the truth. However, we should not forget that the same security problems exist in ordinary server implementations as well.

Virtualization has become popular over the last years partly because it enables cloud computing as a technology. Hypervisors have their own role in virtualization and so far

there are a few well-known hypervisors available. This study compares five hypervisors: ESXi, KVM, Hyper-v, Xen and PowerVM. The comparison shows that ESXi is currently the ruler of the market in the field of virtualization, but the KVM and Hyper-V are trying to keep up with development. Although ESXi hypervisor dominates the market, it is most expensive option. An open-source-based KVM hypervisor can manage almost the same features, but at a significantly lower cost. Based on these findings, the clients should consider KVM or other open source hypervisors over the VMware's ESXi hypervisor.

This thesis presents a project where one IBM customer wanted to simplify their complex infrastructure and enable resource sharing by using a cloud-computing solution based on IBM Power Systems servers, IBM PowerVM technology and IBM Systems Director and IBM Service Delivery Manager software to optimize their resource use, decrease time to market and increase flexibility of the virtual image deployment. By implementing a cloud computing environment, the client succeeded in reaching the desired features.

In the end, the client managed to lower the total cost of ownership by improving capacity use by 75 percent and was also able to reduce IT labor costs by 50 percent. These measures enabled deployment of a self-service cloud and increased flexibility by gaining support for multiple platforms and hypervisor technologies. Also, the customer was able to increase scalability and agility in business.

What comes to future of cloud computing it is very difficult to make predictions about the future. In my opinion this interesting and versatile technology still has lots of undiscovered opportunities to be found. However, I truly believe that our today's youth will definitely meet one another more significant cloud computing solution.

References

1. Gartner, [Online]. URL: http://www.gartner.com/it-glossary/cloud-computing/. Accessed 12.2.2013.

2. Peter Mell, Tim Grance, the NIST Definition of Cloud Computing [Online]. National Institute of Standards and Technology, Information Technology Laboratory, Special Publication 800-145; September 2011. URL: http://csrc.nist.gov/publications/nistpubs/800-145/SP800-145.pdf. Accessed 15 July 2012.

3. Michael Armbrust, Armando Fox, Rean Griffith, Anthony D. Joseph, Randy Katz, Andy Konwinski, Gunho Lee, David Patterson, Ariel Rabkin, Ion Stoica, and Matei Zaharia. Above the Clouds: A Berkeley View of Cloud Computing [Online]. RADlab Berkley; February 2010. URL: http://www.x-integrate.de/x-in-cms.nsf/id/DE_Von_Regenmachern_und_Wolkenbruechen_-_Impact_2009_Nachlese/$file/abovetheclouds.pdf. Accessed 19 September 2012.

4. SAAS, PAAS and IAAS – Making Cloud Computing Less Cloudy [Online].CIO Research Center, 1 December 2010. URL: http://cioresearchcenter.com/2010/12/107/. Accessed 14 June 2012.

5. James Baty, Glenn Brunette. What is Cloud Computing [Online]. Sun Microsystems; June 2009. URL: http://www.alejandrobarros.com/media/users/1/50369/files/4363/CloudComputing.pdf. Accessed 10 April 2012.

6. M. Armbrust, A. Fox, R. Griffith, A. Joseph, R. Katz, A. Konwinski, G Lee, D. Patterson, A. Rabkin, I. Stoica. Above the clouds: A Berkeley view of cloud computing. EECS Department, University of California, Berkeley, Tech. Rep. UCB/EECS-2009-28, 2009.

7. Jay Heiser, Mark Nicolett. Assessing the Security Risks Of Cloud Computing [Online]. Gartner; 3 June 2008. URL: http://cloud.ctrls.in/files/assessing-the-security-risks.pdf. Accessed 14 June 2012.

8. Top Threats to Cloud Computing [Online]. Cloud Security Alliance; March 2011. URL: https://cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf. Accessed 14 June 2012.

9. New York University, [Online]. URL: http://www.nyu.edu/about/news-publications/news/2010/11/01/NYU-Google-email.html Accessed 13.2.2013.

10. Jyväskylän Yliopisto, [Online]. URL: https://www.jyu.fi/itp/ohjeet/how-to-nain-tehdaan/opiskelijoiden-saehkoepostin-uudistus. Accessed 13.2.2013.

11. J. Ciancutti. Four Reasons we chose amazon's cloud as our computing platform, [Online]. Netflix, Techblog. December 2010. URL: http://techblog.netflix.com/2010/12/four-reasons-we-choose-amazons-cloud-as.html. Accessed 12.3.2013.

12. IBM Client Reference material, [Online]. IBM, SAP. URL: ftp://submit.boulder.ibm.com/sales/ssi/ecm/28/0crdd-8vtd28/SPC03405-WWEN-00_approved_for_external_use.PDF Accessed 12.3.2013

13. IBM Client Reference material, [Online]. IBM. URL: ftp://submit.boulder.ibm.com/sales/ssi/ecm/we/0crdd-936ewe/Leading_bank_invests_Anonymous_Final_Approved_for_External_Use.PPT. Accessed 12.3.2013.

14. Jouko Ruuskanen, IT Architect, IBM. Discussion 19 September 2012.

15. Jeff Daniels. Server Virtualization Architecture And Implementation, [Online]. Crossroads, Fall 2009.URL: http://delivery.acm.org/10.1145/1620000/1618592/p8-daniels.pdf?ip=195.212.29.173&acc=OPEN&CFID=180088732&CFTOKEN=88679369&__acm__=1360763545_f38996abb1a3f68ad556a6cbb8206fce. Accessed 12.3.2013.

16. Server Virtualization growth [Online]. IDC 2012. URL: http://www.smartertechnology.com/c/a/Smarter-Strategies/Closing-the-Virtualization-Management-Gap/. Accessed 7 march 2013.

17. Worldwide Virtual Machine Software 2010 Vendor shares [Online]. IDC; 3 August 2011. URL: http://www.marketresearch.com/IDC-v2477/Worldwide-Virtual-Machine-Software-Vendor-7352531/. Accessed 13 February 2012.

18. Gernot Heiser. Much Ado About Type-2 [Online]. Open Kernel Labs; 14 October 2010. URL: http://www.ok-labs.com/blog/entry/much-ado-about-a-type-2/. Accessed 15 June 2012.

19. VMware Corporation, [Online]. URL: http://www.vmware.com/. Accessed 17 August 2012.

20. VMware ESX & ESXi hypervisors [Online]. VMware. URL: http://www.vmware.com/products/vsphere/esxi-and-esx/index.html. Accessed 17 August 2012.

21. Transitioning to the ESXi Hypervisor Architecture – What Customers Need to Know [Online]. VMware Inc; February 2011. URL: www.vmgu.ru/ext/ppt/transitioning_to_ESXi.pptx. Accessed 18 June 2012.

22. VMware ESXi and ESX Hypervisor Frequently Asked Questions [Online]. VMware Inc. URL: http://www.vmware.com/products/vsphere/esxi-and-esx/faqs.html. Accessed 18 June 2012.

23. Gary Chen, Al Gillen. KVM for Server Virtualization:An Open Source Solution Comes of Age [Online]. IDC; October 2011. URL: http://www-03.ibm.com/systems/resources/systems_virtualization_idc__kvmforservervirtualization.pdf. Accessed: 27 June 2012.

24. George Wilson, Michael Day, Beth Taylor. KVM: Hypervisor Security You Can Depend On [Online]. IBM Linux Technology Center; November 2011. URL: ftp://public.dhe.ibm.com/linux/pdfs/LXW03004-USEN-00.pdf. Accessed 4 June 2012

25. Comparing HyperVisor Vitualization: Hyper-V vs. Xen [Online]. VPSHosting; 31 August 2011. URL: http://www.vpshosting.com/2011/08/comparing-hypervisor-vitualization-hyperv-xen/. Accessed 28 June 2012.

26. Server Virtualization with the Xen Hypervisor [Online]. Xen. URL: http://www.xen.org/products/xenhyp.html. Accessed 3 July 2012.

27. VMware vs. KVM: A Functionality/Price Discussion [Online]. Clabby Analytics Research Report. May 2012. URL: ftp://public.dhe.ibm.com/linux/pdfs/Clabby_Analytics_-_VMware_v_KVM.pdf. Accessed 10 July 2012.

28. IBM Software Group. IBM Service Delivery Manager [IBM Internal document]. IBM

29. Mark McDonald, Dave Aron. Reimagining IT: The 2011 CIO Agenda [Online]. Gartner EXP. URL: http://www.gartner.com/technology/cio/cioagenda_findings.jsp. Accessed 28 April 2012.

30. Glenda Lyon. Tivoli Cloud Computing [IBM Internal Document]. IBM.

31. Why Your Company Must Consider Cloud Computing [Online]. QuickCert Blog; 19 march 2012. URL: http://www.quickcert.com/blog/index.php/why-your-company-must-consider-cloud-computing/. Accessed 25 June 2012.

32. Patrick Grady. HP and IBM Battle for Cloud Equipment Market Leadership; Cisco Closes the Gap [Online]. Synergy Research Group; 2011. URL: http://www.ispsjb.com/content/hp-and-ibm-battle-cloud-equipment-market-leadership-cisco-closes-gap. Accessed 4 June 2012. Accessed 15 June 2012.

33. Tiina Siltala. Kaikki ei sovi pilveen. [Online]. Tivi CIO, 2012. URL: http://www.tietoviikko.fi/cio/kaikki+ei+sovi+pilveen/a807197. Accessed 7 march 2013.

34. Leena Mäntysaari. Pilvipalvelujen eteneminen Suomessa 2011 [Online]. Marketvisio; 21 December 2011. URL: https://www.marketvisio.fi/fi/tutkimukset/ict-johtaminen/1066-pilvipalvelujen-kaytto-yleistyy-vauhdilla. Accessed 26 June 2012.