



LUND UNIVERSITY
Department of Informatics

Exploring Views on Data Centre Power Consumption and Server Virtualization

Master thesis, 15 ECTS credits in Informatics (INFM02).

Presented: May 2009

Authors: Johannes Radicke, Björn Rodén, and Liu Yunke

Supervisors: Lars Fernebro, Yuanji Cheng

Examiners: Agneta Olerup, Eric Wallin

Abstract

- Title:** Exploring Views on Data Centre Power Consumption and Server Virtualization
- Authors:** Johannes Radicke
Björn Rodén
Liu Yunke
- Publisher:** Department of Informatics, School of Economics and Management, Lund University, Sweden
- Supervisor:** Lars Fernebro, Yuanji Cheng
- Examiners:** Agneta Olerup, Eric Wallin
- Presented:** May 2009
- Thesis type:** Master thesis, 15 ECTS credits
- Language:** English
- Keywords:** Server Virtualization, Server Consolidation, Power Consumption, Green Computing, Green-IT, Interviews, Delphi Survey

Abstract

The primary purpose of this Thesis is to explore views on Green IT/Computing and how it relates to Server Virtualization, in particular for Data Centre IT environments. Our secondary purpose is to explore other important aspects of Server Virtualization, in the same context.

The primary research question was to determine if Data Centre (DC) power consumption reduction is related to, or perceived as, a success factor for implementing and deploying server virtualization for consolidation purposes, and if not, what other decision areas affect Server Virtualization and power consumption reduction, respectively.

The conclusions from our research are that there is a difference of opinion regarding how to factor power consumption reduction from server equipment, both from promoters and deployers. However, it was a common view that power consumption reduction was usually achieved, but not necessarily considered, and thus not evaluated, as a success factor, nor that actual power consumption was measured or monitored after server virtualization deployment. We found that other factors seemed more important, such as lower cost through higher physical machine utilization, simplified high availability and disaster recovery capabilities.

Acknowledgement

We wish to thank our supervisors Lars Fernebro and Yuanji Cheng for their support and assistance, and for agreeing to assist or participate in our research study we also wish to convey our gratitude and thanks to (in no particular order):

- Mats Johansson from RTS AB
- Stephan Åkeborg from IKEA IT
- Thommy Malmström from Logica
- Gabriella Teruggi from HMS Sipac SpA
- Peter Lindholm from Ängelholms Municipality
- Jonas Grubb from SonyEricsson Mobile Communication
- Johan Johansson from Lund University Computing Center
- Claus Colmeus from Lund University Biomedical Center (BMC)
- Bo Freij, Joakim Nygren, Alf Holmberg, Jan Nihlén, and Eric Saxin from Pulsen Systems AB
- Lars Bergfors, Per Fredriksson, Fredrik Lundholm, Mattias Lundström, and Henrik Grönholm from IBM

Table of Contents

1	Introduction & Problem Area	1
2	Green IT/Computing, Data Centre Power Consumption, and Server Virtualization	3
	2.1 Literature Review Conclusions.....	8
3	Research Methodology & Methods	10
	3.1 Research question and Hypotheses	12
	3.2 Integrity and Ethics	13
	3.3 Strengths and limitations	14
	3.4 Validity and Reliability	15
	3.5 Triangulation and Member Validation.....	17
	3.6 Bias and Error.....	17
	3.7 Delphi Method.....	18
	3.8 Interview & Survey.....	21
	3.9 Qualitative analysis.....	25
	3.9.1 Data managing	25
	3.9.2 Reading, memoing	26
	3.9.3 Describing, classifying, interpreting	26
	3.9.4 Representing, visualizing	26
	3.10 Interviews	27
	3.11 Quantitative analysis	27
	3.11.1 Hypothesis testing.....	28
	3.11.2 Sampling.....	28
	3.11.3 Validity and Reliability	29
4	Results & Analysis	30
	4.1 Delphi	30
	4.1.1 Review of Delphi results.....	31
	4.1.2 Statistical analysis of Delphi 2 nd round results	33
	4.2 Interviews	44
	4.2.1 Review of interview results	44
5	Conclusions	59
6	Further Study	63
A	Delphi Statistical Tables 2 nd round.....	64
B	Delphi Questions 2 nd round	70
C	Interview Guides.....	71
D	Interview Transcriptions	73
	Transcription from interview with Promoter A.....	73
	Transcription from interview with Promoter B.....	76
	Transcription from interview with Promoter C	80
	Transcription from interview with Promoter D.....	85
	Transcription from interview with Promoter E	88
	Transcription from interview with Deployers F.....	91
	Transcription from interview with Deployers G	94
	Transcription from interview with Deployers H	97
	Transcription from interview with Deployers I.....	99
	Transcription from interview with Deployers J.....	102
	Bibliography.....	104

Figures & Tables

Figure 1: Rising Energy Consumption in the Data centre (VMware white paper, 2008).....	4
Figure 2: Question formulation process.....	11
Figure 3: Hypothesized relationship between Green IT and server virtualization vs. power consumption reduction.....	13
Figure 4: Types of Surveys (Fink & Kosecoff, 2005)	24
Figure 5: The data analysis spiral (Creswell, 2007).....	25
Figure 6 Model for quantitative analysis	34
Table 1 1 st Delphi Hypotheses for both question group one and two	35
Table 2 Result five-number summary with standard deviation and mode, for question group one.....	36
Table 3 Result five-number summary with standard deviation and mode, for question group two	37
Table 4 2 nd Delphi factor analyses for question group one.....	38
Table 5 2 nd Delphi factor analysis for question group two	40
Table 6 2 nd Delphi Hypotheses question group one and two	40
Table 7 One-Sample Test question group one	41
Table 8 Reduce software licensing costs	41
Table 9 One-Sample Test question group two.....	43
Table 10 2 nd Delphi Hypotheses test for question group one	43
Table 11 Communalities question group one	64
Table 12 Component Matrix question group one	65
Table 13 Total Variance Explained question group one.....	65
Table 14 Communalities question group two	66
Table 15 Component Matrix question group two	67
Table 16 Total Variance Explained question group two.....	68
Table 17 Frequency table question group two	69

1 Introduction & Problem Area

During the last few years, the terms Green IT or Green Computing appears more on the agenda for companies and organizations, and in media, such as from various sources obtainable through Internet.

Green IT/Computing comprises a large area of aspects centred around the main idea of using computing resources more efficiently, reducing the green house gas emission of data centres, developing better power management strategies, handling the increasing demand of data storage, reducing the use of toxin for producing IT equipment, or changing the organizational culture towards a more ecological way of thinking and enabling and improving inter-personal communication with less travel.

According to the Gartner report by (Mingay, 2007) the following definition of green IT seems appropriate to define the term in the context of an enterprise.

“Green IT is the optimal use of information and communication technology (ICT) for managing the environmental sustainability of enterprise operations and the supply chain, as well as that of its products, services and resources, throughout their life cycles.”

The overall power consumption of a Data Centre (DC) is influenced by several factors, e.g. cooling, air-conditioning, floor layout, lighting, types of servers, storage, and networking, and the overall utilization of IT equipment.

Although many companies have deployed virtualized server infrastructure, such as IBM PowerVM, VMware ESX, Citrix XenServer or Virtual Iron, and often motivate their decisions based on simplifying operations management, enhancing usage of computer resources and thus reducing the amount of computers, but also power and cooling requirements, there have been scarce information on actual research studying important (success) factors for virtual server infrastructure deployment, especially in the context of real world power consumption reduction. (VMware white paper, 2008)(IBM, 2008)(IBM Global Services, 2007)

Assuming that “*Virtualization will be the most-important trend for servers through 2012*” (Gartner Press Release, 2008), and that there are uncertainty what the key factors are for successful deployment of virtualized server infrastructure, and it seems that even research, although performed for profit by the research industry, have not yet thoroughly identified what counts as “success” in relation to virtualized server infrastructure deployment, seemingly related to the uncertainty of the proper reasons for companies and organizations choosing to deploy server virtualization.

Since the virtualized server infrastructure seems to be a big part of the Green IT/Computing idea, this thesis will cover some technical descriptions as well as its effects concerning resource savings. However, we limit the scope within the Data Centre IT environments to the server context, excluding other aspects such as networking and storage.

The primary purpose of this Thesis is to explore views on Green IT/Computing and how it relates to Server Virtualization, in particular for Data Centre IT environments. Our secondary purpose is to explore other important aspects of Server Virtualization, in the same context.

Our primary research question, following from the purpose of this Thesis, has been to determine if Data Centre (DC) power consumption reduction is related to, or perceived as, an important (success) factor for implementing and deploying Server Virtualization for consolidation purposes, and if not, what other decision areas affect Server Virtualization and power consumption reduction, respectively, with emphasis on Server Virtualization.

For our fact finding and information gathering, we decided to follow an indirect exploratory approach for knowledge gathering for our research, based on both a modified Delphi survey, with two iterations, and a series of subject matter expert interviews. Indirect in such a way that we relied on the subject matter experts experience and knowledge regarding server virtualization, and data centre power consumption cases pertaining to our subject area. And in order for us to verify our different hypotheses, which consisted of both how (server virtualization) and why (power consumption reduction), we decided to use both qualitative and quantitative methods to gather information.

To not rely on only one source or method, and to perform the research as rigorous as feasible, during our allotted and limited time frame, we decided to focus and rely on triangulation as the core principle for our research. Hence we used researcher triangulation by dividing execution of our selected research methods, but we also employed method and source triangulation by using:

- *Literary review* concerning Green-IT/Green Computing, energy efficiency, and server virtualization, see 2 Green IT/Computing, Data Centre Power Consumption, and Server Virtualization.
- *Delphi method* (modified) survey of Subject Matter Experts (SMEs), who are individuals at companies and organizations that are skilled, knowledgeable, and who are promoting and participating in deployment of server virtualization (classified in our research as promoters), see 4.1 Delphi.
- *Interviews* of both deployers and promoters from companies and organizations, both face-to-face and by phone see 4.2 Interviews.

In the following chapters we will give the background to Green IT/Computing and Server Virtualization following from our literature review. Thereafter we describe the research methodology and methods we have employed to perform our research, followed by the disclosure of the results, and analysis from the applied research methods. Finally we conclude by summarizing our conclusions and suggestions for further studies into the subject area.

2 Green IT/Computing, Data Centre Power Consumption, and Server Virtualization

There is no doubt that global warming and its related issues has become a hot spot in almost every area of social life, human behaviour, research and of course within industries. More and more companies today have "green" on their agenda in one form or another, especially those companies whose business models and strategies rely on a significant amount of IT usage and in particular centred around carbon footprint and power reduction or cost saving.

Referring to chapter 1 "Introduction & Problem Area", the term "Green Computing", which is regarded equivalent to "Green IT" in the reviewed literature during this research, seem to be first mentioned in connection with the EPA in the autumn of 1992. The Washington Post published an article stating that the EPA was trying to promote energy-efficient PCs that were equipped with the new technology of sleep modes that will automatically reduce the power usage in case the machine is not used since a certain time. (Reid, 1992)

Since then the topic of green technology has continuously increased in popularity and regarding the newest discussions about global warming and raising power supply costs organizations should treat it as an additional aspects of their IT/ICT policies and strategies.

Thus Green IT/Computing was voted as the number one topic in a survey about IT trends for the year 2009 concerning the question about the suggested biggest technology development in the coming 5-10 years. The participants of this survey were CIOs, IT managers and middle managers. (CIO Insight, 2009)

The linkage between global warming and technology usage is mainly based on technology power consumption, both for producing and using, which is in turn affected by seemingly continuously increasing data storages and higher network bandwidth demands. Organizations who own large data centres or server parks are interested in handling the balance between satisfying their technology demands and on the other hand decreasing their power consumption. Nevertheless managers do not only focus on saving the environment, moreover green IT technology can save lots of money once deployed.

In order to support business growth and ensure economic sustainability, organization of all size need to establish a strategy to address the impact of their power, cooling, floor space, and environment health and safety (PCFE) which addresses the many different facets of being green for IT data centre. In addition to leverage new and enhanced information services, enable business agility and improve on cost effectiveness to remain competitive while reducing effect on the environment. (Schulz, 2009)

Resources, such as IT equipment, belong to the enterprise and are connected to a certain environmental cost (e.g. power consumption, paper waste, non-recyclable hardware). Accordingly Green IT/Computing is dealing with the management and optimization of enterprises resources following the “green” way. Moreover life cycles of products or IT equipment within an organization can also be influenced and adapted to environmental standards. Consequently the definition is stressing an increasing material and energy efficiency regarding the enterprise’s IT infrastructure and business activities as well as the reduction of environmental impact, for instance pollution or waste of resources. (Mingay, 2007)

A survey of 280 IT executives illustrates that two primary factors lead the companies to considering about this new concept: cost-cutting efforts related to energy efficiency; and efforts to be more socially responsible corporate citizens. This survey also shows the current situation of the implementation of Green IT by those participant enterprises. As can be seen from the survey, over half of enterprises claim that their organization has set up at least one corporate social responsibility program that includes "green" initiatives dedicated to environmental sustainability. (CIO Staff, 2008)

Furthermore, activities organizations are currently undertaking or planning to implement to make their IT operations more environmentally friendly include reducing server power consumption, educating users to turn off equipment at night, configuring desktops not in use to enter sleep mode, and upgrading or reconfiguring data centre cooling infrastructure for improved efficiency. (Schulz, 2009)

Due to the rising consumption of energy, exhaustion of natural resources and increasing concern for the environment, plenty of companies who has the sensitivity to and interest in the issues of Green Computing are high have started to develop a new strategy which is focusing on reducing power consuming and pollution effects of using IT. No doubt building a “Green” data centre is a hot topic in this area. A report by the Environmental Protection Agency (EPA) claims data centres in the U.S. consume 4.5 billion kWh annually, 1.5 percent of the country’s total. Perhaps more importantly, as shown in Figure 1 the consumption has doubled from 2000 to 2006, and is likely to double again until 2011. (Koomey, 2007)




Figure 1: Rising Energy Consumption in the Data centre (VMware white paper, 2008)

The Green Grid is a non-profit trade organization, which focuses on power and cooling in datacenters. From the Green Grid perspective, the cost is the most important reason concerning power reduction (The Green Grid, 2007), and the electrical power usage should

be considered as a key factor when organizations implement their datacenter because most datacenter manager is unaware of what their monthly energy bill is. However, it may cost more on electrical energy in datacenter than the costs of the electrical power system or the cost of the IT equipment itself (The Green Grid, 2007).

According to “Guide lines for Energy-Efficient Datacenter” published by the Green Grid there are two perspectives in the Energy improvements can be made which are equipment-planning perspective and an operational-practices perspective for both IT and physical infrastructure (power, cooling, rack, security, fire suppression, and monitoring) devices. (The Green Grid, 2007)

There are three main ways to deal with the problem of inefficient IT system. First of all, the companies could implement the right-sizing physical infrastructure systems. From the research, using the right size system will reduce 50% of the electrical power. Secondly, among the whole system, the organization could use efficient physical infrastructure devices, which can help power reduction fundamentally. For instance, the closely couple cooling solution make the cool air distribution and hot air removal system more efficiency. With new server and intelligent cooling options available today, a 20,000-square-foot data centre should be able to achieve energy savings in the 40 to 45 percent range over past years, according to SandOaks (2009), as described in a news brief by IT Professional (2008). Thirdly is to design the whole system from the energy-efficient concept (ibid). Especially in data centre, a good well-designed system will contribute more on energy-efficient than only use the efficient physical infrastructure devices. (The Green Grid, 2007)

The amount of data and information managed through information technology usually increases over time within an organization and so does consequently the number of applications and functions. Consequently a higher demand on hardware, in this case mainly servers, occurs.

IBM promote energy efficiency of data centres and virtualization solutions, for which there is no doubt that the cost is one of the main reasons to help their customers to reduce power consumption, both in data centres and at the workplace. Because in the last few year, the energy and power cost has become the majority of IT costs, and many organizations has started their new strategy in power management. In fact, some companies have found that their date centre have been outgrown, and that their available local power supply cannot meet the demand anymore, which force them to relocate or build new data centres, so the properly planed data centre will bring companies financial returns. (Lechner, 2009)

Server virtualization therefore can perform different individual system environments on a single physical machine. Some benefits of virtualization compared to a solution where each system is based on its own physical machine are that it saves hardware costs, reduces the amount of space needed and consequently need less energy consumption in terms of power supply and cooling (Baschab, Piot, & Carr, 2007).

Server virtualization, also referred to as system-level virtualization (Engelmann, Naughton, Ong, Scott, & Vallée, 2008), has several, although similar, definitions, all evolving around the concept of abstracting the physical computer hardware interaction of operating system software, and enabling several instances of operating systems to execute simultaneously within the same physical system, or simplified as described by Popek et al (1974) as “A virtual machine is taken to be an efficient, isolated duplicate of the real machine”, or simply that a “virtual machine, is now generally accepted as a software replica of a complete computer system” (Meyer & Seawright, 1970).

IBM mainframe computers (Encyclopædia Britannica, 2009) was the first commercial system which employed “virtualization” (referred to as logical partitioning) as an abstraction of hardware resources as early as the 1960s (Karger, 2007), and later to IBM POWER computers and others, but its availability for the Intel/AMD based x86 and x64 processor architecture based computers, from the beginning of the 21st century, has contributed to a more widespread use, since the x86/x64 architecture are the most commonly used one for smaller server systems of today (Wikipedia, 2009).

Popek et al (1974) describes a virtualization model, and for their model they identify three essential characteristics: “First, the VMM [virtual machine monitor, which manage access to the physical computer hardware] provides an environment for programs which is essentially identical with the original machine; second, programs run in this environment show at worst only minor decreases in speed; and last, the VMM is in complete control of system resources” (ibid).

The virtualization enabler of a physical computer is commonly referred to as the Virtual Machine Monitor (VMM), or Hardware Virtualizer (Goldberg, 1973), or Virtual Machine Control Program (Meyer & Seawright, 1970), or simply Hypervisor (Young, 1973), and is the software that provides the virtual abstraction, or virtual machine, of the physical computer to operating systems, either with or without computer hardware architecture support (Popek & Goldberg, 1974). Popek et al also describes a hybrid virtual machine (HVM) as a virtual machine system where more instructions are interpreted rather than directly executed, and also the possible property of recursiveness for hypervisors (ibid). Within the virtual machine there are usually a “guest” operating system, which can either not know about the virtualized environment, or it can be modified to be aware, paravirtualized (Barham, o.a., 2003), that it is executing within a virtualized environment (Chisnall, 2007).

It should be noted that the third generation computers Popek et al refers to are computers from the 1970s with “a processor and linear, uniformly addressable memory” (ibid), and their model does not account for I/O instructions or interrupts. However, the three properties of a hypervisor are still relevant today: efficiency, resource control, and equivalence.

The purpose of this thesis is not to in-depth analyzing different hypervisor technologies, so we will not distinguish between monolithic and disaggregated hypervisors (Murray, Milos, & Hand, 2008), or between different physical processor privileged contexts and binary translation by hypervisors (Adams & Agesen, 2006), or different computer hardware architectures providing various resource access control mechanism (Wikipedia contributors, 2009), or the mechanism for controlling the hypervisor.

However, the following will illustrate a common model of hypervisor structure, which can be found in the IBM POWER hypervisor (Armstrong, o.a., 2005), and others, as of this writing. The POWER hypervisor (PHYP) handles virtual machine access to physical memory by using a page table for each virtual machine, mapped to the physical machine page table for the real memory pages within logical memory blocks, it also use separate physical memory area for direct memory transfer of I/O to and from virtual adapters. The virtual network adapter memory transfer is handled by copying the communication frame to and from the virtual machine allocated memory, and the hypervisor only managed memory. Physical adapters can be allocated to virtual machines using hypervisor allocated and restrictive hardware adapter bus bridges, connected to bus adapter slots. The allocation of virtual machine processing is managed by the hypervisor in a pre-emptive, time-shared multi-tasking way, with 1/10 and 1/100 ms timeslots per processor core. Thus the hypervisor controls access to processors, physical memory pages, and bus located I/O adapters. For data transfer by local area or storage area network, either a separate virtual machine can handle the I/O bridging and be a virtual I/O server, or specifically designed hardware can be utilized, such as the Integrated Virtual Ethernet (IVE) for virtualizing a physical Ethernet adapter port, as well as the N_Port ID virtualization (NPIV) for virtualizing a physical fibre channel adapter port (Vetter, o.a., 2008).

According to Alan Dayley, at the time research director at the information technology research and advisory company Gartner “Virtualisation helps organisations to cut costs, better utilise assets and reduce implementation and management time and complexity, all of which are crucial in this economic environment” (Dayley, 2009).

Server virtualization can also assist IT operations organizations with collocation of servers and workload, and disable unused computers from consuming power (Green-IT/Green Computing), or enhance high availability or resilience of server services by reallocating active virtual machines from one physical computer to another, or enhance Disaster Recovery by reactivating virtual machines on another physical computer at possibly another geographic location (Bradford, Kotsovinos, Feldmann, & Schlöberg, 2007).

Not only does the server virtualization affect the processes and tasks within IT operations organizations, but it also affects software design and architecture, and service provisioning. Server virtualization also affects the organizations communication network architecture and design, both from a service perspective but also from a security (IT/IS-security) perspective (Aurelius). Server virtualization is a computer infrastructure provisioning which can be used to deploy operating system servers for Grid computing networks (Ruda, Denmark, & Matyska, 2007), or as infrastructure base for Software as a Service (SaaS) environments (Yu, Guo, Nanda, Lam, & Chiueh, 2006), or Cloud computing environments (Sotomayor, Keahey, & Foster, 2008), or in some instances also High-Performance Computing (HPC) environments (Mergen, Uhlig, Krieger, & Xenidis, 2006).

Most interesting of these are Cloud computing (Wikipedia contributors, 2009), which is a development from Grid computing, SaaS, Infrastructure as a Service (IaaS), and Utility computing, and describe computing environments which dynamically provisions, configures, reconfigures and de-provisions virtual servers, storage, applications and services, as needed. Cloud environments can also include other computing resources, such as storage area networks (SANs), network equipment, firewalls or other security devices, and can be either private/exclusive or external/shared, in relation to the company or organization utilizing

Cloud resources. However in this context, Cloud computing could be viewed as a possible extension to a current private infrastructure, and resources in the Cloud environment being used, and billed for, when additional capacity is needed, thus reducing the requirement of local deployment, reconfiguration, or additional procurement.

Server virtualization is used to consolidate servers from multiple physical servers to fewer physical servers by converting and migrating the original physical servers to become virtual servers, encapsulated as virtual machines and collocated with other virtual servers using the same physical server.

The purpose of server consolidation, from physical to virtual servers, are considered by companies and organizations, to reduce server sprawl, increase physical server resource utilization rates (Goldman Sachs, 2007), decrease floor space, power, and heating, ventilation, and air conditioning and cooling (HVAC), reduce new capital investments and simplifies infrastructure management (IBM ITSO, 2008).

It is claimed, by computer industry vendors (IBM, HP, VMware, Microsoft, and others), that server consolidation with virtualization can reduce costs, not explicitly taking into account server virtualization, that hardware costs by up to 70 percent, hardware maintenance costs by up to 50 percent, support costs by up to 33 percent, energy costs by up to 40 percent, and floor space and facility costs by up to 50 percent (IBM Global Services, 2007), and different vendors give their own opinion on the level of such cost savings. Server consolidation can be performed by using different techniques, such as through geographical centralization and physical consolidation of servers, and through server virtualization (Badaloo, 2008).

Additional considerations for implementing server virtualization, not limited to server consolidation per se, are to significantly reduce time to deployment of new servers and services, and thus providing a more rapid and flexible response to constantly changing business demands (IBM ITSO, 2008), but also to reduce downtime (planned or unplanned), and thus increase business resilience issues, otherwise associated with environmental inflexibility (IBM ITSO, 2008).

2.1 Literature Review Conclusions

The concept of Green IT/Computing is a necessity for all companies and organizations to embrace, to maintain a sustainable growth in delivering increasingly higher computing resources. Both from a company ethical perspective, but also from a cost savings perspective, where the ethical can be viewed as a part of a company or organizations “Goodwill”, whereas the cost savings mostly will be a factor for larger national or international companies and organizations where the bill for power consumption or data centre infrastructure are noticed and accounted for.

The core savings regarding Green IT/Computing, come from the infrastructure and utilization of computing, storage and network resources, within a Data Centre, and from infrastructure and utilization of employee workstations, and from infrastructure and processes for inter personal communication and meetings, and last but not least from material recycling when discarding used, discarded or replaced IT equipment.

Benefits for deploying server virtualization can be summarized into the following groups (IBM, 2008), in a broader context than only for server consolidation purposes to reduce costs, are:

- Infrastructure consolidation and simplification or hardware efficiency
- Software licensing efficiency
- Energy usage efficiency
- Server or Business resiliency
- IT operations responsiveness and systems administration and support efficiency
- IT management and organization alignment

To summarize, in the context of Data Centre environments, it is essential to design the Data Centre and company/organization IT infrastructure for energy efficiency, to right size the physical infrastructure, and to ensure efficient usage of the physical infrastructure, such as through, but not limited to, deploying server virtualization. All of which are easier said than done, especially in a corporate climate with quarterly planning goals.

3 Research Methodology & Methods

In this section we outline the basis for our approach to the research at hand, and how the theoretical scientific framework has influenced our decisions and thus work.

We decided to follow an indirect exploratory approach, based on a series of subject matter expert interviews, for knowledge gathering for our research. Indirect in such a way that we relied on the subject matter experts experience and knowledge regarding server virtualization, and data centre power consumption cases pertaining to our subject area. And in order for us to verify our different hypotheses, which consisted of both how (server virtualization) and why (power consumption reduction), we decided to use both qualitative and quantitative methods to gather information.

Our initial research approach follows Creswell's (2007) iterative approach, and the description by Seale (2007), regarding analytical deduction, following a five-step process:

1. Roughly define the problem.
2. Construct a hypothetical explanation for the problem.
3. Examine a case to see whether it supports the hypotheses.
4. If the case does not fit, either reformulate the hypotheses or redefine the problem to exclude the negative case.
5. Continue this search through several cases until negative instances are no longer found.

Thus the analytical deduction can follow hypothesizing based on a cycle of theorizing from evidence, until the relationship between the problem and hypotheses are found to be understood and relevant, and the researcher is reasonable confident about their general applicability (if that is the purpose). Seale (2007) argues that deviant cases can give additional support to the researchers conclusions, lead to modification of the researchers ideas, but also that they are explainable (in the sense of being understood).

We decided to triangulate our sources and information acquisition methods, and by dividing the scope of work during the information gathering process, also researcher triangulation, referring to the benefits of triangulation approach in scientific research, see 3.5 Triangulation and Member Validation.

Source triangulation was achieved through the use of literary review, Delphi method to solicit important factors from subject matter experts, and interviews.

The schematic, in Figure 2, illustrates the question formulation process we applied for our information gathering, given the time constraints on our research, and the scope of work, we based our initial hypotheses and 1st level of questions on our initial knowledge. We first determined our initial research question, and initial hypotheses.

To enlist open and unbiased responses from all participants and respondents in our research, we decided to ensure their anonymity, see 3.2 Integrity and Ethics, and be clear and open with our purposes, and separation of researcher roles.

For our fact finding research we determined that the most direct way to obtain proper, current and accurate information, would require us to apply different research methods and participants.




Figure 2: Question formulation process

We therefore classified the type of respondents into Promoters, and Deployers, where the promoters are representatives from vendors who promote the use of server virtualization to accomplish power usage reduction, and deployers are those who use server virtualization to accomplish power usage reduction.

Due to the technical complexity of the area of server virtualization, we determined that the best source regarding the advantages and drawbacks, from a technical point of view, would be from promoter subject matter experts, with thorough insight into the technical aspects of server virtualization, but also data centre operations, power consumption, and economical models such as Total Cost of Ownership (TCO) and Return On Investment (ROI), which deployer companies and organizations (we assumed) would require before deciding on deploying server virtualization, and the reasons for doing so.

Thus we achieved source triangulation from using literature review, questioning both those who promote server virtualization, and those who deploy server virtualization, giving several angles or views on the same subject matter.

To ensure that the research method we would use would limit our fact finding, we decided to also use different methods to acquire the information from our sources. Apart from literature review of other researchers accomplishments, we decided to employ a method specifically for subject matter experts, and in our case from the class of promoters of server virtualization, and we selected a modified Delphi Method, see 3.7 Delphi Method for more details, to enable us to build proper and reality founded questions, taking into account our limited time frame. We used the literature review and Delphi results to create more proper questions for our interviews, see 3.8 Interview & Survey.

Thus we decided to gather knowledge through our research, to verify our different hypotheses, which consists of both how (server virtualization) and why (power consumption reduction).

- Create introduction letter for participation in the research, with appropriate statement of anonymity and beneficence.
- Determine classification of respondents into Promoters, and Deployers - The promoters are representatives from vendors who promote the use of server virtualization to accomplish power usage reduction, and deployers are those who use server virtualization to accomplish power usage reduction.
- Determine whom to personally interview - Subject matter experts at companies (deployers) who are or have been responsible for deploying server virtualization within data centre management.
- Determine whom to participate in the Delphi survey - Subject Matter Experts (promoters) who are promoting and participating in deployment of server virtualization.
- Determine whom to online survey - Companies and organizations (deployers) that have deployed server virtualization.

3.1 Research question and Hypotheses

Our primary research question was to determine if Data Centre power consumption reduction is related to, or perceived as, a success factor for implementing and deploying server virtualization for consolidation purposes, and if not, what other decision areas affect server virtualization and power consumption reduction, respectively, illustrated in Figure 3.

We therefore formulated the following initial hypotheses to prove or disprove during our research, and to improve or change during the course of our research according to our iterative approach:

H1₀: Companies do not want to achieve power consumption reduction for IT equipment, especially in a DC environment.

H1₁: Companies want to achieve power consumption reduction for IT equipment, especially in a DC environment.

H2₀: Companies do not use server virtualization to achieve power consumption reduction.

H2₁: Companies use server virtualization to achieve power consumption reduction.

- H3₀: Companies do not want to achieve power consumption reduction primarily for cost reduction.
H3₁: Companies want to achieve power consumption reduction primarily for cost reduction.




Figure 3: Hypothesized relationship between Green IT and server virtualization vs. power consumption reduction

3.2 Integrity and Ethics

By providing confidentiality, with respect to relating published information, based on answers to our questions, to our interview and survey participants, we ensure anonymity and non-maleficence to the individual participants, to ensure that the subject matter experts would not be hindered to give truthful and uninhibited answers to our questions.

Integrity of research is promoted by caring about ethics and acting ethically, which is essential to scientific research since other scientists build incrementally on previous research. Not only does poor practice affect the individual and professional reputation, but also the veracity and reliability of the individual and collective work (in the research community). Topics of ethical focus are informed consent, confidentiality, beneficence and non-maleficence, and that ethics is “a generic term for various ways of understanding and examining the moral life” (Israel & Hay, 2008, s. 12), which is a principlism based ethical approach, based on the four prima-fascia principles of respect for autonomy, beneficence, non-maleficence and justice, where the principles are binding unless conflicting with other principles, in which case it is necessary to choose between the conflicting ones. (Israel & Hay, 2008).

By acknowledging that the “truth is the only value that constitute the goal of research” (Hammersley & Gomm, 1997), and that research immediate goal is the production of

knowledge” (ibid), it can be stated, from a teleological view, that a researcher should not perform misconduct such as *fabrication* (making up data or results), *falsification* (manipulating data or results) and *plagiarism* (appropriation of another person’s ideas, processes, results, or words without giving appropriate credit), since this would will negate the goal of production of knowledge. It could also be argued that it is the self-discipline and integrity of the individual researcher, to behave right and virtuous from a deontological perspective, which prevents fabrication of quotations or data, especially since it can be part of the social research methodology to anonymify the collected research data which would then be obscured from review by others.

We hope that it is evident in this Thesis that we have strived to uphold both demands for rigour by own work, and avoided the short cuts of fabrication, falsification, and plagiarism.

Noting that the main ethical principlism rules, codes, and principles of ethical research in a social and medical context, regarding individual research subjects, emanate from the Nuremberg code (Wikipedia contributors, 2009), and subsequent international and further national legislation, with focus on:

- Respect for persons such as participation in the research based on voluntary and informed consent, with proper confidentiality.
- Beneficence and responsibility of the researcher to avoiding harm and doing good, and that the research should yield beneficial outcome with minimal risk to individuals and communities.

3.3 Strengths and limitations

We believe that our audience would be both academic, due to the interest in research relating to the relation between server virtualization and power consumption reduction, but also commercial, due to interest in knowing whether there is a relationship between deploying server virtualization and reduction in power consumption and thus cost.

We also believe that our selection of multiple research techniques and potential for both qualitative and quantitative analysis, of a contemporary topic, can strengthen the validity of our results, and thus contribute knowledge to both academic and commercial interests.

Regarding research methods it is argued by Yin (2009) that it is important to understand and acknowledge both strengths and limitations of the selected research method. Seale (2007) argue that theoretical speculation can help researchers with ideas to look at problems in different perspectives. Grounding theory in data is an important element in supporting claims with credible evidence, and depends on a separation between data and theory, which is considered necessary. Thus observation alone is not enough, and Seale argue that researchers should not depend on “unthinking adherence to rules of method”, and should engage in methodological debates to “loosen thoughts that are stuck”. (Ibid). Thus one of the reasons for why we have chosen to triangulate research methods, and in particular to employ the Delphi method in an uncommon way, see 3.7 Delphi Method.

Creswell (2007) argues that one methodological research assumption is that researchers can use an iterative approach to penetrate the research topic, with a ground up (bottom up) direction. Starting from empirics and develop theory, thus implying that it is possible to perform qualitative research starting with theories and then evaluate these based on empirical findings, but that the researcher must make an active choice about assumptions when performing qualitative research, and that good research requires making the researcher's assumptions, paradigms, and framework explicit in the writing of a qualitative study. We adopted this methodological assumption, and based our initial research approach on Creswell's (2007) iterative approach, and the description by Seale (2007), regarding analytical deduction, following a five-step process (note that "case" in our context are each subject matter expert):

1. Roughly define the problem.
2. Construct a hypothetical explanation for the problem.
3. Examine a case to see whether it supports the hypotheses.
4. If the case does not fit, either reformulate the hypotheses or redefine the problem to exclude the negative case.
5. Continue this search through several cases until negative instances are no longer found.

Thus the analytical deduction can follow hypothesizing based on a cycle of theorizing from evidence, until the relationship between the problem and hypotheses are found to be understood and relevant, and the researcher is reasonable confident about their general applicability (if that is the purpose). Seale (2007) argues that deviant cases can give additional support to the researchers conclusions, lead to modification of the researchers ideas, but also that they are explainable (in the sense of being understood), whereas Yin (2009) argues that every research method can be used for three purposes: exploratory, descriptive, and explanatory.

3.4 Validity and Reliability

The main approach to validity, which we have applied to our research, focus on the use of triangulation and member validation, and to account for and evaluate, the occurrence of inevitable bias. Referring to section 1 Introduction & Problem Area on page 1, we decided to triangulate using method, source and researcher triangulation, to cover more angles to our research subject matter, and thus increase the accuracy and truthfulness of our research findings.

Norris (1997) argue that "validity enhancing practices does not ensure that research is accurate, correct, certain, trustworthy, objective..." (ibid), and that a practical way to think about validity is to focus on error and bias, to acknowledge that researchers are fallible, and accept the fact that there is no paradigm solution to the elimination of error and bias, but that research demands scepticism, commitment, detachment, and a capacity to accept and use criticism, but also to be self-critical in a constructive manner. He notes however, that researchers have to take some things for granted, since the research has to start somewhere.

Hammersley et al (1997) propose an approach focusing on the distinction between accounts and the studied phenomena, the actual detachment between researchers and the studied phenomena, the description in linguistic terms of a phenomena represents a the phenomena from a point of view and does not reproduce the phenomena itself, and the crucial role of the research community in assessing knowledge claims on criteria of plausibility and credibility. And arguing that in a realism oriented approach, it becomes more difficult and complicated to find errors, since there are no rigid patterns, procedures, protocols, or methodological rules to follow from evidence to conclusion (induction) and thus “truth”.

Regarding quality of research, Yin (2009) argues that it is possible to judge the founding research design by evaluating the concepts of trustworthiness, credibility, confirmability, and data dependability (Yin, 2009, s. 40), and propose the following four tests:

- Construct validity, which is to identify correct operational measures for the concepts being studied, such as triangulation, chain-of-evidence, and review.
- Internal validity, which is seeking to establish a casual relationship (for explanatory or casual studies only), such as pattern matching, explanation building, addressing rival explanations, or using logical models.
- External validity, which is to define the domain to which a study’s findings can be generalized, beyond the immediate case study, such as using theory in sing-case studies, or replication logic in multiple-case studies.
- Reliability, which is demonstrating that the operations of a study can be repeated with the same results and minimizes errors and bias in a study, such as using case study protocol, and developing a case study database.

Seale (2007) argues that researchers should not be too encouraged to use tools and methods for quality validation, reliability etc, and should not be discouraged by criticism of techniques since it will allow the researcher to better understand its strengths and limitations. Seale (2007) also argue that both triangulation and member validation offer methods for testing researchers’ claims by gathering new evidence, though the validation process.

Seale (2007) further argue that concerns about reliability and replicability are relevant to qualitative research. With reliability he states that external reliability involving replication of whole studies has been difficult to achieve in practice. External reliability is concerned with the concept that other researcher studying the same or similar setting should be able to achieve the same findings and results. Whereas internal reliability applies to how other researchers while applying similar constructs would match these to data in the same way as the original researcher. For internal reliability Seale advocate using peer auditing and reflexive methodological accounting, such as how field notes are handled and transcribed, to create low inference, in order to enhance internal reliability. (Seale, 2007)

Hammersley et al (1997) imply that validation techniques found in foundationalism can be used as guidelines, if considerable judgment is involved in applying them, arguing that there is an “absence of any prospect of absolute proof”, and that outcome error is “not necessarily the product of culpable procedural error” (ibid). In this realism oriented approach, Hammersley et al (1997) also argue that that there need not be any differentiation between research procedural errors and erroneous findings, since “the first almost inevitably leads to the second; and the second...is an absolutely reliable indicator of the first” (ibid). And that

indeed the “truth is the only value that constitute the goal of research”, and that research immediate goal is the production of knowledge” (ibid).

3.5 Triangulation and Member Validation

Each application of our research has different concerns regarding triangulation and member validation. By applying the same line of questioning to different categories of research participants, we believe this strengthens the correlation possibility from responses, as do utilizing different inquiry methods such as interviews, online survey, and email survey.

By using the Delphi method, we achieve a form of member validation already early on in the research process, thus strengthening our line of questioning, but also allowing for adjustments.

Both Seale (2007) and Yin (2009, s. 116) argue, and Yin credits Patton (2002), that there are four types of triangulation: *data triangulation*, by using separate data sources, *investigator triangulation*, by using more than one researcher as a team, *theory triangulation*, by approaching the data with different hypotheses or perspectives, and *methodological triangulation*, by applying different research paradigms or methods, such as combining interviews and observations. However, Seale also point out that triangulation can bring knowledge also by not achieving convergence and confirmation, but it is not a guarantee for validity.

Member validation takes it root from ways to ensure that the research subjects accounts are indeed factual and true, and not altered to appease the researcher in some way or capacity. Seale also mention this regarding triangulation, to adjudicate the truth status of the interviewees’ accounts. Member validation can be viewed as the result of the researcher seeking communication (and reassurance); with members of the community of which the research is perform within. (Seale, 2007)

3.6 Bias and Error

Similarly to triangulation and member validation, by our application of the same line of questioning to different categories of research participants, we believe this also reduce the inevitable bias from correlated responses, as do utilizing different inquiry methods such as interviews, online survey, and email survey.

However, by using the Delphi method, we introduce a potential for increased bias through the peer oriented focus of this method, but by limiting the usage to two iterations and accounting for outliers, bias should be identifiable. But since the Delphi method is only used, in our approach, on one category of research participants, there can be built in bias which will not be immediately evident, but should be possible to identify by correlating responses from other categories of respondents, and vice versa.

Bias can be seen as a systemic error, which is based on the viewpoint that it is a deviation from a “true score”, which in turn is based on “the valid measurement of some phenomenon or to accurate estimation of population parameters” (Hammersley & Gomm, 1997), and that bias depends on concepts such as “truth” and “objectivity” (ibid).

Examples of potential sources of bias in research, such as: researcher reactivity with information providers and consumers, researcher affinity or value preferences or commitments with objects or subjects related to the research, researcher ability and personal qualities to conduct the research, selection bias, but also availability and reliability of sources. However, error and bias should be accepted to exist, and that there are no ways to ensure truth, but that validity concepts should still be applied to research, focusing on the cause of error and bias affecting and influencing the researcher. (Norris, 1997)

We have believed it to be important, due to the nature of our inquiries, that we understand that individuals perceive, view, and pursue knowledge about bias in them selves and in others, concluding that in oneself bias identification is performed introspectively, and bias in others by relying on abstract theories to determine if others are biased or not. Thus individuals are also more readily inferring bias when they see others making judgments that serve or coincide with those individuals’ self-interest, and underestimating bias in one self. (Ehrlinger, Gilovich, & Ross, 2005)

3.7 Delphi Method

In our study we used the Delphi method as a means to acquire information and thereafter use additional inquiry from other groups, to get real world performance measure, thereby enabling validity. In addition our selection of Subject Matter Experts is based on assurance from individual organizations and practical accomplishments of the panellists, both of which can reduce the issue of dropouts in subsequent rounds.

The results from the first round with the Delphi method, should be the independent judgment from each panellist, whereas subsequent rounds, which strives to reach some level of consensus among the panellists, can be viewed as having intentionally introduced bias through the presentation of a compilation of the collected views by the group of responding panellists.

However this bias is intentional and forms the basis for reaching consensus, but can also be viewed as possibly exposing the panellists to other viewpoints, which can support, contradict, or influence, the knowledge of the panellists, thus being similar in concept to peer-review in this respect. Riesman (Encyclopædia Britannica, 2009) points out that within a peer-group, individuals can adjust their values to conform to those of the group in a “constant process of change”, and by being a panellist in a Delphi method, each panellist will know that they, and the other panellists, have been selected for their knowledge.

Thus it is a possibility that the perceived results of a consensus based correlation, is not consensus but rather a compromised correlation, where the compromising panellists give in to others point of view, especially if there are several subsequent rounds of questioning where also the motivation to answer can be influential on the panellists responses.

Viewing the criticism of the Delphi method (Sackman, 1974), such as issues with panellist selection as experts on the subject matter, and the truthfulness in perceived consensus, it could be argued that the level of congruence in the first round of a Delphi method, would show who really belongs to the peer-group, and who does not, in so that the panellists who adjust their opinion is an indicator of those not being at the same level of expertise as the majority of the panellists, and those who maintain their non-consensus opinion might either be ignorant or more knowledgeable than the majority of the panellists. Thus following this thought of reasoning, the most valuable outcome from the Delphi method would then be the consensus in the first round, and the outliers from subsequent rounds. Hence we limited our usage of the Delphi method to two iterations or rounds.

Our sample selections for applying the Delphi method to questioning as Subject Matter Experts (SMEs) are individuals at companies and organizations, who are promoting and participating in deployment of server virtualization (classified as promoters), and would be limited to two rounds.

The procedure we have used for the Delphi process has followed the following guidelines, which is a simplification compared to the processes collated by Yousuf (2007):

1. Perform sample selection for participation in the Delphi survey.
2. Promote participation in Delphi questionnaire, email purpose of the thesis and method, and primary questions.
3. SMEs reply to the questionnaire by email, by replying they accept participation in the survey.
4. Follow up phone interview with the respondents who have outlier responses to the questions, to ensure understanding of the questions posed, and to determine the cause for the outlier response.
5. Compile the results from all replied questionnaires, identify common denominators, and prepare second distribution. The second iteration is comprised of an online survey with Likert scale answer, based on the results from the first iteration.
6. SMEs reply to the second iterations, and can leave additional comments or objections, and if possible with ranking of topics as key (success) factors for power consumption reduction through server virtualization.
7. Compile the results from all replied questionnaires, and evaluate the statistical results.
8. Analyze the results.

To solicit selected subject matter experts for their views regarding important (success) factors regarding Server Virtualization and Data Centre power consumption reduction, the following open-ended questions were asked in the initial request to participate in the Delphi survey. Selection of respondents was made using LinkedIn connections of deployers who are known to work with server virtualization. For the first round we also decided to show all solicited participants the group of who was invited to the study.

1. In your view, what are the main success factors for server virtualization?
2. How can these success factors be measured?
3. In your view, is power consumption reduction an important factor for companies and organizations when they decide to deploy server virtualization?

The online survey used fixed answers to questions allowing for Likert scale answers, and additional free format feedback. The questions posed can be found in the Appendix.

Sackman (1974) describes the "conventional Delphi" method, as developed by Rand in the mid 1900s, as a chronological framework following a problem-solving sequence: determine objectives; formulate problem, perform iterative polling (of subject area expert panellists) with feedback and scoring, finalized by interpreting and analyzing the results. Linstone et al (1975) defines the Delphi method as "a method for structuring a group communication process so that the process is effective in allowing a group of individuals, as a whole, to deal with a complex problem". A similar definition of the Delphi method is that it is "an iterative process to collect and distill the anonymous judgments of experts using a series of data collection and analysis techniques interspersed with feedback." (Hartman, Skulmoski, & Krahn, 2007).

It should be noted that Sackman based his criticism of the Delphi method on the 1966 draft standard "Standards for Educational and Psychological Tests and Manuals" by the APA Office of Scientific Affairs, but even so the criticism is informative, although most of the criticism seems to concern the usage of the Delphi method for forecasting, lack of verifiability of the results, and issues with statistical properties and analysis of the results. The main points of Sackmans criticism of the Delphi method is elaborated on and organized into an eight-point checklist by Linstone et al (1975, ss. 560-571).

There are similar methods to sequentially question and review interacting teams of experts, but we will focus on the basic principles of Delphi in this section.

The "conventional Delphi" method as described by Sackman, consists of the following characteristics (Sackman, 1974, s. 7): The format is a structured and formal questionnaire, which is administered by mail, in a personal interview, or online, for two or more rounds. With each subsequent round, the questionnaire consisting of questions and answers options with similar or different scales, are accompanied with instructions, guidelines, and ground rules, and summarized results from the previous round (mostly statistical and anonymous).

Outliers can be requested to provide justification for their responses, and the rounds can continue until it is deemed appropriate due to convergence of opinion, or some level of consensus is reached, at the discretion of the issuer. Thus the essence of the Delphi method consist of experts, anonymity, iteration with controlled feedback, and aggregation of group responses (Hartman, Skulmoski, & Krahn, 2007), or simply comprising of a series of questionnaires answered by a pre-selected group of experts (Arditi).

Sackman (1974) questions the usefulness of the Delphi method, and points out potential issues with determining the level of expertise of subject area expert panellists, and the sampling method used to select the panellists, and the level of formalization of questions. And since the Delphi method relies on iterative questioning, how will survey technical issues affect the outcome, such as how to handle dropouts, and how will dropouts effect the items answered, how will the amount and detail level of the questions effect the results, and how long should the interval be between successive rounds, and how to ensure prompt response from the panellists to expedite the turnaround time.

Other potential issues are anonymity of panellists, or how does lack thereof effect other panellists answers, and will the level of motivation from panellists prevail over several successive rounds of repeated questioning, but also how the results will be used, and how to ensure validity and replicability of the iterative process. Sackman also points out that the Delphi method might not be applicable for rigorous social science experiments, but could be considered a more informal exploratory technique.

3.8 Interview & Survey

Throughout the research process we used interview techniques to collect information from different viewpoints regarding virtualization and energy consumption. Therefore in-person and telephone interviews have been performed in order to collect reliable information from several promoters and deployers of virtualization. Potential interviewees have been contacted by phone to ask about their intention of participating within our studies and schedule interview appointments. Depending on their location and general availability we preferred to meet in-person at their working places or, as an alternative, set up conference calls with them.

Based on the literature reviews several key aspects of virtualization and power consumption have been identified, see chapter 2 Green IT/Computing, Data Centre Power Consumption, and Server Virtualization. These key aspects have been transformed and compiled in an interview guide, stating 16 open-ended question and several connected headwords, see Appendix B. The headwords were mentioned in the first place to make sure and check if the interviewee's answer consists of the details we wanted him or her to speak about. On the other hand they were used to specify our questions in case the interviewee didn't exactly know what to talk about.

According to Kvale (2007) the interview guide provides the topics and their sequences in the interview. It is up to the researcher to use a guide that consists of detailed worded questions or only rough topics and keywords. The guide helps the interviewer to control the process of the conversation, since she knows whether or not a topic has been spoken about. Nevertheless it is not necessary to totally stick to the guide. The researcher should be able to adapt her prepared questions to the answer that have been given so far during a dialog.

Two different versions of the guide, containing slightly different expressions of each question but still aiming on the same topics, were used during the interviews, depending on the fact whether the interviewee was categorized as a promoter or deployer. The interview guide was generally followed during the conversations. In case the interview partner gave explanations or was stating issues of any following-up questions we attempted to encourage them an go into more detail about it right away. This ensured us to keep down the number of redundant information and maybe more important, express our interests in the given answers in order to motivate the interviewee to collaborate with us instead of getting an impression of simply following a strict order of questions without creating a real conversation. That might be particularly true for the phone interviews we conducted, since a positive personal impression is harder to achieve when working with the lack of gesture and mimic. All interviews were to be recorded and transcribed.

The different survey types that have been used throughout our thesis project enabled the collection of valuable and reliable information about our investigated problem areas. Moreover different points of view on the same areas have been discovered.

The procedure we have used for the interview process has followed the following process:

1. Promote participation in interview, email purpose of the thesis and the interview.
2. Create interview guidelines and questions.
3. Schedule appointment for interview.
4. Start with 1st level questions, promote uninhibited answering, and follow up with subsequent level questions.
5. Transcribe the interview.
6. Review the transcription.
7. Compile the transcription.
8. Finalize compilation with comments.
9. Analyze finalized compilation.

We also decided to perform the first one or two interviews, as a pilot test interview in order to check our interview technique and questions. Therefore we chose participants who were willing to give feedback on both the set of questions and general impression about the interview. Moreover they should have a good view of the technology and energy usage aspects, ideally from a promoter point of view, but also knowing about the needs and considerations from a Deployers point of view

When creating our interview questions and interview guidelines, we started with the following core and sublevel questions. The interview questions were then adapted for use during separate interview of deployers and promoters. To develop appropriate questions, based on currently available information, we organized the questions in four level hierarchical structures. The first level were to be the primary opening questions for our interview participants (what questions), the second level questions were to focus on how the first level questions are perceived to be achieved (how questions), whereas the third level questions were to address viewpoints of importance or ranking of important (success) factors for the first level questions, related to the second level questions, whereas the fourth level questions were to focus on exemplification relating to the first, second, and third level questions combined. The following sample questions illustrate the hierarchical structure levels.

1st level questions

1. What is your view regarding that companies, or organizations, want to achieve power consumption reduction for IT equipment, especially in a DC environment?
2. What is your view regarding that companies, or organizations, use server virtualization to achieve power consumption reduction?
3. What is your view regarding that companies, or organizations, use server virtualization to achieve power consumption reduction?

2nd level questions

4. How is it possible to, practically and cost effectively, measure successful accomplishment of power consumption reduction, or other benefits, due to server virtualization?

3rd level question

5. Are there any key factors for successful accomplishment of server virtualization?

4th level question

6. Could you give some examples of how key factors have contributed to successful accomplishment of power consumption reduction due to server virtualization?

The detailed interview questions, which were formulated from the above structuring, can be found in the Appendix.

Our sample selections for online survey are individuals at companies and organizations, that are or have been responsible for deploying server virtualization within data centre management (classified as Deployers).

The procedure we had intended to use for the online survey process has followed the following guidelines:

1. Analyze the results from interviews, and subject matter expert questionnaires, and literature review, identify denominators and outlier responses.
2. Create online survey.
3. Request assistance from selected companies and organizations, to identify staff at companies and organizations that have or are responsible for decision to deploy server virtualization.
4. Promote participation in the online questionnaire, email purpose of the thesis and method, to at least 30 staff from companies and organizations who have or are responsible for decision to deploy server virtualization.
5. Compile the results from all replied questionnaires.
6. Perform statistical validation and reliability verification on the data.
7. Perform statistical analysis and inference on the data.
8. Analyze the statistical results.

The online survey used fixed answers to questions allowing for Likert scale answer, and additional free format feedback. The questions posed can be found in the Appendix.

Surveys can be conducted in a self-administered way or through interviews. Both types distinguish between different sub-types. Therefore the self-administered survey can either be performed via mailing, on-site, or online. Interviews on the other hand can be held in person or through the use of telephone communication.




Figure 4: Types of Surveys (Fink & Kosecoff, 2005)

According to Fink et al (Fink & Kosecoff, 2005) several advantages and disadvantages exist for each type of survey, as shown in Figure 4.

Mailed self-administrated surveys can contact participants in a large geographical area. Moreover the receivers do not have any technical difficulties to answer the paper-based survey by simply using a pen. In addition they are able to complete the survey wherever they want. Nevertheless the contacted people need to be motivated enough to answer the survey and must be able to read, see and write. A correct address list is essential to achieve a good rate of respondents. Additionally follow up mailings and incentives might be considered. Associated costs will be based on the amount of used paper, printings, envelopes, stamps and incentives if used.

Doing a self-administrated on-site survey has the advantage that questions can be asked in-person and therefore enable both sides to ask comprehensive respectively follow up questions about the topic. Moreover information is obtained immediately since answers are given right away to the researcher. The on-spot setting does also enable group questioners. On the other hand received information is limited to the attendant people and space as well as privacy issues for completing the survey should be considered in advance. The costs are comparable to the mailing survey ones, besides possible charges for the provided working space.

Performing an online self-administrated survey in contrast does mainly rely on technical knowhow and equipment. It can be reached world wide as long as the participants have Internet access and know how to use a browser. Online surveys do usually only accept correct answers regarding syntax and format, which is a big advantage and reduces the amount of error-affected data. Moreover researcher can control the answering order of the questions, which might be important in certain situations. Also hyperlinks can be provided to explain unfamiliar words or instructions to the participants. Another advantage is that collected information is already available digitally and can be analyzed easily. Issues can occur concerning server down times and ensuring privacy and confidentiality to the contributors.

Telephone or in-person interviews both need trained interviewers who know how to communicate with the interviewees. What is more can either of them can be performed by using a written or computer assisted guide that leads through the investigated topic and makes sure every important information will be covered.

Through the flow of the in-person or telephone based conversation answers can be explored and the researcher is able to assist the respondent with unacquainted phrases. A schedule for

reaching the contacts might be needed regarding telephone interviews since they can only be conducted if the respondents are available.

Performing in-person interviews alternatively requires a suitable place to meet that possesses of enough space and appropriate privacy protection. Accordingly, visiting the interviewee at home can be dangerous or difficult for the researcher. The costs consist of expenses trough training, space, travel and incentives.

3.9 Qualitative analysis

Throughout the conducted surveys we collected a lot of information. A crucial step towards a useful analysis was to identify and compile the mentioned facts that met our interests regarding the research questions. The process of analyzing qualitative data can be visualized in a data analysis spiral (Creswell, 2007), see Figure 5. It starts with certain data and generates through several analytical circles a final account or narrative. The spiral emphasizes that researcher do usually not follow a linear approach in order to analyze their data (ibid).




Figure 5: The data analysis spiral (Creswell, 2007)

The following paragraphs are related to the different procedures of the spiral to the assumed conducted activities of this study:

- Data managing
- Reading, memoing
- Describing, classifying, interpreting
- Representing, visualizing

3.9.1 Data managing

A collaboration tool (Microsoft Groove) is used throughout the whole time of the project, which offers several features regarding data managing and storing and teamwork element such as instant messages and calendar functionalities. Organizing data becomes important to get an overview of all available information and make further working activities easier.

The majority of data is already digital, e.g. digital records of the interviews, returned survey results and literature sources or first have to be transformed in a digital file, e.g. through transcribing the interviews. Afterwards the data needs to be clustered in different units, such as folders. This will help to keep track of all the available and collected information and the possibility of losing the overview is much less compared to using non-digital file systems.

3.9.2 Reading, memoing

Reviewing the in-person interviews and several topics related articles was done to identify main aspects, motivation and concerns of virtualization and power consumption managing in organizations. Taking notes and remembering important details at this stage is the basis of the later analyzing part, where our findings have to be connected to common explanations.

Therefore the part of reading and memoing consists of reviewing the available information, understanding the overall meanings, identifying important details and breaking them into different parts. In addition, taking notes while reviewing the data will support this process. The notes will consist of short phrases, ideas, critical remarks or key concepts that are worth to focus on later

3.9.3 Describing, classifying, interpreting

Having dealt with divers interview techniques and partners as well as focusing our literature reviews on both the SV and power consumption opportunities regarding green IT, categories could be established accordingly. Once classified, the available information could be rephrased and inserted into the thesis document, which consists of different sections that are also based on divers clusters of information. Moreover this part covered the interpretations of the essential information that has been collected throughout the whole research process. This phase is quite important since the reviewed data will be split up in categories, which consist of the same characteristics in order to establish different themes or patterns of data.

Consequently these categories consist of all available information linked to the specific issues and will then be investigated by the researcher. Categories will be established based on the different identified factors influencing power consumption reduction and virtualization. This is followed by the interpretations of the essential information that has been collected in throughout the research process.

General assumptions or rules therefore need to be identified and also exceptions should be considered and discussed. The term *interpretations* in research is linked to the act of making sense of the data. It can be based on hunches, insights and intuition and personal views.

3.9.4 Representing, visualizing

Our findings have been visualized through the creation of this thesis document. It mainly consists of text, but also presents diagrams and figures to visualize complex relationships in a comprehensible way.

Furthermore cross tables have been used to visualize relationships and frequency distributions of certain variables. Also hypothesizes or propositions are used that point out certain relationships among categories of information. Based on the existing relationships a model could be created in order to illustrate them.

3.10 Interviews

Getting the right information out of the in person interviews was crucial in order to come up with valid statements and reliable final conclusions in this paper. Therefore the characteristics of the interviewee needed to be taken into consideration. Facts, such as job position, main responsibilities and the situation or size of the organization she or he is working for need to be clarified before the actual interview was performed. This gave us essential information that helped us to interpret the interviewee's answers in the right way.

Kvale (2007) explains that during the interview the researcher should compare his or her understandings of the interviewee's answers. Consequently the interview partner is able express her answers more precisely until only one possible interpretation is left. After transcribing the interview an additional interpretation step follows conducted by the researcher. Therefore the material gathered through the interview needs first be structured, e.g. by computer programs, to make the huge amount of information manageable. Moreover essential parts must be identified and separated from the non-essential ones. This action relies on the purpose of the actual study. Finally, the understandings and meaning of the interviewees concerning the investigated topic should be presented.

Regarding our research, the recorded interviews have first been encoded into documents (see section 3.9), which only consists of essential information. This was followed by the clustering of all answers expressing similar characteristics, which led us to our final aspects of SV and energy consumption issues (see section 4.2.1). Accordingly our conclusions are built on the general opinions that have been stated in these clustered aspects.

3.11 Quantitative analysis

Collected quantitative data needs to be electronically organized and managed, as well. Different datasets needs to be labeled in order to keep track of all of them. Before using analyzing tools like Excel or SPSS the data needs to be checked for completeness and errors. Afterwards the statistical tools can be used to generate data representations such as frequency tables, correlations, regressions, distribution graphics or independence tests to describe the data and make patterns among variables visible.

For our purposes, we intend to use the five-number summary and standard deviation, to analyze the basic statistical properties from our quantitative data, which we will collect during Delphi round two and online surveying. The five-number summary consists of the max, min, mean, quarter percentiles, and standard deviation values for each question or sub-question.

3.11.1 Hypothesis testing

Hypothesis testing is based on a tentative assumption about a population parameter, which can be based on the normal distribution, or not, and is used to test validity of a claim or research hypothesis (Anderson, Sweeney, Williams, Freeman, & Shoemith, 2009, s. 291), where the tentative assumption is referred to as the null hypothesis, denoted by H_0 , and the opposite hypothesis to the null hypothesis is called the alternative hypothesis, denoted by H_1 (ibid p288). The H_1 hypothesis is the one we want to reach, and to be proved, by rejecting the H_0 hypothesis (ibid p290). The probability of making a Type I error when the null hypothesis is true as an equality is referred to as the level of significance, which is denoted by α , and if the cost of making a Type I error is high, then a small value of α is preferred (Anderson, Sweeney, Williams, Freeman, & Shoemith, 2009, s. 294), however, a smaller value will increase the Type II errors (ibid p328), thus selecting a proper level of significance determines the power of the test, and thus the probability of correctly rejecting the H_0 when it is false.

3.11.2 Sampling

For our usage of the Delphi method we intend to use a non-probabilistic judgement sampling, which would never the less limit potential non-sampling errors due to lack of knowledge (Anderson, Sweeney, Williams, Freeman, & Shoemith, 2009), adhering to the requirement of the Delphi method to ensure that the respondents are SMEs, with knowledge both about server virtualization, but also knowledgeable about company and organizations reasoning for choosing and deploying server virtualization.

Due to the nature and requirements of applying a Delphi method, such as questioning subject matter experts in their field of expertise, we anticipate that for this method, due to the allotted timeframe, we might not receive a samples size, which will allow us to statistically analyze the data and reach our desired certainty of 95%. However, nevertheless we will apply the methods, allowing for the possibility of uncertainty, and correlate the statistical findings we do achieve, with those of Delphi first round, and Interview series.

For our Online Survey method we also intend to use a non-probabilistic judgement sampling, which would also in this case limit potential non-sampling errors due to lack of knowledge (Anderson, Sweeney, Williams, Freeman, & Shoemith, 2009), since we would require selection assistance from participating Promoter companies to provide us with potential respondents, who all should have knowledge about server virtualization, but also knowledgeable about their company and organizations reasoning for choosing and deploying server virtualization.

The responding sample size for our Online Survey, should preferably superseded 30, thus given a response rate of 25% we would have to distribute the survey to at least 120 potential respondents. The general form of the interval estimate of the population mean μ being $x \pm \text{Margin of Error}$ (Anderson, Sweeney, Williams, Freeman, & Shoemith, 2009, s. 260), denoted by Δ , which can be calculated using $z_{\alpha/2}(\sigma/\sqrt{n})$, thus $x \pm z_{\alpha/2}(\sigma/\sqrt{n})$, depending on the confidence level, where 99% gives $z_{\alpha/2}=2.576$ and 95% (our choice) gives $z_{\alpha/2}=1.960$, depending on that the population follows a normal distribution in which case the calculated confidence interval will be exact, otherwise it will be approximate (ibid).

3.11.3 Validity and Reliability

Our approach to achieve statistical reliability and validity, are based in our attempt to embrace the scientific concepts of rigor, thoroughness, openness to alternative or rival conclusions and tests, and adherence to measuring protocols, and especially for quantitative methodology also thorough sampling, and determining appropriate statistical tests and methods to correlate data, would also be appropriate, if the measured data and the test methods, are appropriate to be able to evaluate and draw correlated conclusions from.

The purpose of quantitative validity is to reduce, or eliminate, error and bias in the variables, and for the test of tests, to ensure high correlation between the data, tests and criteria. Validity evolves around the concept of ensuring that the measurements, tests, and criterion are indeed measuring, testing for which the criterion is set. Validity in this respect can be viewed in two ways for our purposes. The first involve validity of the test itself and is the degree to which evidence and theory support the interpretation of the intended tests to test for what was intended to test for.

Reliability on the other hand is concerned with the consistency of the measurement and that measures are not taken in error or with bias, thus allowing for erroneous or biased data to be evaluated (whether for sample or population). Thus the reliability determines how repeated measures, by the same collection methods or techniques, render the same measurements and thus evaluation. So with proper sampling methods, the sample results will provide “good” estimates of the population parameters (Anderson, Sweeney, Williams, Freeman, & Shoesmith, 2009, s. 224), and there are certain methods that can be used to evaluate statistical reliability, such as the Cronbach’s alfa (Wikipedia contributors, 2009).

4 Results & Analysis

As mentioned in chapter 1 Introduction & Problem Area, our primary research question was to determine if Data Centre power consumption reduction is related to, or perceived as, an important (success) factor for implementing and deploying server virtualization for consolidation purposes, and if not, what other decision areas affect server virtualization and power consumption reduction, respectively.

In this chapter we describe the results and analysis for each method we employed, Delphi (qualitatively and quantitatively), and Interviews (qualitatively).

4.1 Delphi

From our simplified Delphi surveying, as described in 3.7 Delphi Method, we selected 35 Subject Matter Experts (SMEs) from different server virtualization promoting companies. The choice was made using LinkedIn profiles and through personal knowledge of the SMEs, thus a non-probabilistic judgement sampling, which would never the less limit potential non-sampling errors due to lack of knowledge (Anderson, Sweeney, Williams, Freeman, & Shoesmith, 2009), adhering to the requirement of the Delphi method to ensure that the respondents are SMEs, with knowledge both about server virtualization, but also knowledgeable about company and organizations reasoning for choosing and deploying server virtualization.

The first round of Delphi consisted of three (3) questions with open-ended answers, to allow the SMEs to answer in their own words, and outlining that we would conduct a Delphi survey, and would keep individual responses anonymous and confidential. Thus answering the survey would imply willingness to participate in the survey. Over a period of two weeks for the first round, and an additional week for the second round, we received eight (8) responses resulting in a 23% response rate. Unfortunately, the number of known subject matter experts was limited, given the limitation that SMEs should be knowledgeable both about server virtualization, but also knowledgeable about company and organizations reasoning for choosing and deploying server virtualization. The respondents were from IBM, Pulsen Systems AB, RTSAB, Logica, and HMS Sipac SpA (Italy).

From the responses from the first round, we built a direct email addressing online survey, using SurveyGizmo (<http://www.surveygizmo.com/>), with Likert scale answers to our posed questions, see Appendix for the detailed questions, and submitted it to our first round respondents, including some from the first round who did not complete the first round on time.

The staff at SurveyGizmo (<http://www.surveygizmo.com/>) also provided us with SPSS structured data from the resulting responses.

During our analysis of the responses, we found some outliers, mostly concerning the view regarding if power consumption reduction is an important factor for companies and organizations when they decide to deploy server virtualization. The outliers were two fold, both that power consumption reduction was not relevant as a decision factor, and that it was indeed the main cost reduction incentive companies and organizations to perform server virtualization.

To allow those respondents with outlier views to further clarify their reasons, we also scheduled more in-depth interviews with them, but the results will be accounted for in the Interview section, see 4.2 Interviews.

4.1.1 Review of Delphi results

To summarize the views during the first Delphi round from the three questions, we got the following results.

From the first open answer question, “What are the main success factors for server virtualization”:

- Cost reduction and more efficient utilization of made investments and lower Total Cost of Ownership (TCO), not only servers but also the whole infrastructure and data centre facilities, such as reduction in floor space, cooling, air-conditioning, power consumption, and staffing.
- Increase resilience and stability, and reduction of downtime, both planned (such as maintenance windows) and unplanned (such as hardware or power failure) and simplified disaster recovery.
- More efficient usage of assets, such as through resource pooling of machines allowing for higher utilization and handling of peak workloads.

From the second open answer question “How can the main success factors for server virtualization be measured”:

- Reduction in license fees required to server applications (depending on vendor licensing methods), calculate before and after deploying server virtualization.
- Reduction in service downtime, measure uptime before and after deploying server virtualization, or calculating and evaluating tasks required for disaster and recovery or business resilience plans with and without server virtualization (NOTE: in the contest that both environments will exist in parallel during the transition period which can last for years).
- Reduction in overall costs for IT operations, such as overtime and consultant fees, calculate before and after deploying server virtualization.
- Reduction in investment and maintenance costs for IT operations, such as server and storage hardware, calculate spending before and after deploying server virtualization.
- Increase in company “Goodwill”, if part of a company strategy to reduce environmental impact of IT operations, measure energy consumption before and after

deploying server virtualization, such as from the data centre UPS (Uninterrupted Power Supply), or calculate approximated utilization before and after deploying server virtualization.

From the third open answer question “Is power consumption reduction an important factor for companies and organizations when they decide to deploy server virtualization”:

- No – it is mainly an additional benefit, it is not the primary objective, the cost for power is not high enough in Sweden, organizations with fewer servers it will not have any major economical value.
- Yes – one of the most important, if organizations deploy large number of servers it can have substantial economic value, especially when current data centres needs more power than is available, it is also politically correct and an important corporate management issue (“Goodwill” related).

From the responses from the first round, we built a direct email online survey, with Likert scale answers to our posed questions, see Appendix for the detailed questions, and submitted it to our first round respondents, including a few from the first round who did not complete the first round on time.

The following division of questions into two groups, was made to determine the view of the subject matter experts on whether companies and organizations are choosing server virtualization to reduce energy costs for powering servers, cooling and air conditioning of data centre, and if companies and organizations who have deployed server virtualization monitor and evaluate the power consumption of physical servers and other data center equipment, such as storage, or if they measure and calculate possible power savings from their deployed virtualized server environment, or if other reasons are more motivating, based on the results from the first Delphi round.

We decided to arrange the questions in no particular order, the questions in bold below are related to our initial hypotheses regarding Server Virtualization and power consumption reduction.

When companies and organizations *are CHOOSING server virtualization*, do they do so to:

- Reduce infrastructure server sprawl by consolidating to fewer physical servers
- Reduce time for server provisioning and deployment of new servers
- Reduce time for systems management of infrastructure and servers
- **Reduce energy costs for powering servers, cooling and air conditioning of data centre**
- Reduce data center floor space for fewer physical servers
- Reduce software licensing costs
- Reduce server maintenance costs
- Increase server resiliency and minimizing planned and unplanned downtime
- Increase the speed for deploying new servers

When companies and organizations *have DEPLOYED server virtualization*, do they then:

- Monitor and evaluate the resource usage of virtual servers (CPU, memory, disk, network)
- **Monitor and evaluate the power consumption of physical servers and other data center equipment, such as storage**
- Test and evaluate the virtualized server resilience, such as high availability or disaster recovery functionality
- Implemented structured work processes and procedures for system management of their virtual server environment
- **Measure and calculate possible power savings from their deployed virtualized server environment**
- Measure and calculate possible time savings from systems management for the deployed virtualized server environment
- Calculate the Return on Investment for their virtual server environment
- Consider that their virtual server environment was worth the investment
- Consider that their virtual server environment has increased server resilience against planned or unplanned downtime
- Have a valid and up to date disaster recovery plan for the their IT environment
- Have a Green-IT agenda or plan for the their IT environment

4.1.2 Statistical analysis of Delphi 2nd round results

This section contains the presentation of the findings from the 2nd round of Delphi method; therefore SPSS has been used as an analysis tool. Each question, see Appendix B, of the Delphi survey was linked to a fixed answers option, which was represented by the following Likert scale values:

- Strongly disagree = 1
- Disagree = 2
- Neither agree or disagree = 3
- Agree = 4
- Strongly Agree = 5
- No opinion/No answer = 0

It should be noted that our collected sample size (8), based on response from Subject Matter Experts (SME), was inadequate for drawing statistically certain conclusions. Our target response samples size was intended to be above 15, but would have been desirable if it had become 30 or more. Nevertheless, we perform our planned statistical analysis on the data available.

The response value for “no opinion/no answer”, which is equal to zero (0), has been changed throughout the dataset into NULL values. The “no opinion/no answer” represents when the respondent did not wish to, or could not answer. Consequently the regarding value are not able to affect further analysis in a misleading way.

There are 20 sub-questions in primary question one and two, so at first we create 20 hypotheses, as showed in Table 1, and then we use factor analysis to reduce the hypotheses into six,

new hypotheses. After which we use t-test and frequency of the data to analyze, see Figure 6 for the approach we used for the qualitative analysis.




Figure 6 Model for quantitative analysis

The following table (in Table 1) consist of the 20 hypotheses we initially created for the 20 sub-questions in our primary question group one and two.

Hypothesis
H _{1a} : Companies and organizations are choosing server virtualization so to reduce infrastructure server sprawl by consolidating to fewer physical servers
H _{1b} : Companies and organizations are choosing server virtualization so to reduce time for server provisioning and deployment of new servers
H _{1c} : Companies and organizations are choosing server virtualization so to reduce time for systems management of infrastructure and servers
H _{1d} : Companies and organizations are choosing server virtualization so to reduce energy costs for powering servers, cooling and air conditioning of data centre
H _{1e} : Companies and organizations are choosing server virtualization so to reduce data centre floor space for fewer physical servers
H _{1f} : Companies and organizations are choosing server virtualization so to reduce software licensing costs
H _{1g} : Companies and organizations are choosing server virtualization so to reduce server maintenance costs
H _{1h} : Companies and organizations are choosing server virtualization so to increase server resiliency and minimizing planned and unplanned downtime
H _{1i} : Companies and organizations are choosing server virtualization so to increase the speed for deploying new servers
H _{2a} : After companies and organizations having deployed server virtualization, they monitor and evaluate the

resource usage of virtual servers (CPU, memory, disk, network)
H _{2b} : After companies and organizations having deployed server virtualization, they monitor and evaluate the power consumption of physical servers and other data centre equipment, such as storage
H _{2c} : After companies and organizations having deployed server virtualization, they test and evaluate the virtualized server resilience, such as high availability or disaster recovery functionality
H _{2d} : After companies and organizations having deployed server virtualization, they implemented structured work processes and procedures for system management of their virtual server environment
H _{2e} : After companies and organizations having deployed server virtualization, they measure and calculate possible power savings from their deployed virtualized server environment
H _{2f} : After companies and organizations having deployed server virtualization, they measure and calculate possible time savings from systems management for the deployed virtualized server environment
H _{2g} : After companies and organizations having deployed server virtualization, they calculate the Return on Investment for their virtual server environment
H _{2h} : After companies and organizations having deployed server virtualization, they consider that their virtual server environment was worth the investment
H _{2i} : After companies and organizations having deployed server virtualization, they consider that their virtual server environment has increased server resilience against planned or unplanned downtime
H _{2j} : After companies and organizations having deployed server virtualization, they have a valid and up to date disaster recovery plan for the their IT environment
H _{2k} : After companies and organizations having deployed server virtualization, they have a Green-IT agenda or plan for their IT environment.

Table 1 1st Delphi Hypotheses for both question group one and two

The following table (Table 2) shows the results summarized with five-number summary, with standard deviation and mode, for question group one.

	N Valid	Mean	Mode	StDev	Min	Max	Percentiles		
							25	50 Median	75
Reduce infrastructure server sprawl by consolidating to fewer physical servers	8	4.50	5	.756	3	5	4.00	5.00	5.00
Reduce time for systems management of infrastructure and servers	8	4.25	4	.463	4	5	4.00	4.00	4.75
Reduce energy costs for powering servers cooling and air conditioning of data centre	8	4.00	4	.535	3	5	4.00	4.00	4.00
Reduce data center floor space for	8	4.25	4	.707	3	5	4.00	4.00	5.00

fewer physical servers									
Reduce software licensing costs	7	4.00	5	1.414	2	5	2.00	5.00	5.00
Reduce server maintenance costs	8	4.38	4	.518	4	5	4.00	4.00	5.00
Increase server resiliency and minimizing planned and unplanned downtime	7	4.14	4	.690	3	5	4.00	4.00	5.00
Increase the speed for deploying new servers	8	4.50	4 ^a	.535	4	5	4.00	4.50	5.00

Table 2 Result five-number summary with standard deviation and mode, for question group one

The next table (Table 3) shows the results summarized with five-number summary, with standard deviation and mode, for question group two.

	N Valid	Mean	Mode	StDev	Min	Max	Percentiles		
							25	50 Median	75
Monitor and evaluate the resource usage of virtual servers CPU memory disk network	8	3.9	4	0.64	3	5	3.3	4.0	4.0
Monitor and evaluate the power consumption of physical servers and other data center equipment such as storage	8	3.0	3	0.76	2	4	2.3	3.0	3.8
Test and evaluate the virtualized server resilience such as high availability or disaster recovery functionality	8	3.5	4	1.31	1	5	2.5	4.0	4.0
Implemented structured work processes and procedures for system management of their virtual server environment	8	3.6	4	0.92	2	5	3.0	4.0	4.0
Measure and calculate possible power savings from their deployed virtualized server environment	8	2.9	3	0.35	2	3	3.0	3.0	3.0
Measure and calculate possible time savings from systems management for the deployed virtualized server environment	8	3.1	3	0.83	2	4	2.3	3.0	4.0

Calculate the Return on Investment for their virtual server environment	8	3.6	3	1.06	2	5	3.0	3.5	4.8
Consider that their virtual server environment was worth the investment	7	4.0	4	0.82	3	5	3.0	4.0	5.0
Consider that their virtual server environment has increased server resilience against planned or unplanned downtime	8	3.8	4	0.71	3	5	3.0	4.0	4.0
Have a valid and up to date disaster recovery plan for the their IT environment	8	3.4	4	1.30	1	5	2.3	4.0	4.0
Have a Green IT agenda or plan for the their IT environment	8	3.4	4	0.74	2	4	3.0	3.5	4.0

Table 3 Result five-number summary with standard deviation and mode, for question group two

4.1.2.1 Factor Analysis

Factor analysis is based on the purpose of reducing a certain amount of variables by extracting a lower number of factors out of them (Eckstein, 2006). Therefore variables with high correlations among each other will be clustered into a common factor. Accordingly, each factor consists of information of several variables. Created factors are independent among each other (Martens, 2003).

The factor analysis was conducted by using the Principal Components Method in SPSS. Due to the communalities of the variables linked to question one (Table 11 in Appendix) it can be assumed that the variable *Reduce time for server provisioning and deployment of new servers* is not a crucial factor for companies or organization when they choose Server Virtualization (the extraction loading is below 0.6). Thus this variable has been removed from the further factor analysis study.

The final result of the FA leads to three different factors, which are able to explain 80.6% of the total variance related to the former variables (Table 13 in Appendix). Actual loadings of each variable on a certain factor due to their highest correlations are mentioned in Table 12 in Appendix.

Consequently the three following variables have a high correlation with factor (or component) one:

- *Reduce time for systems management of infrastructure and servers.*
- *Increase server resiliency and minimizing planned and unplanned downtime.*
- *Increase the speed for deploying new servers.*

Thus it can be concluded that time reduction is a main reason for companies and organization for choosing Server Virtualization.

In addition, the following three variables have a high correlation with component two:

- *Reduce energy costs for powering servers cooling and air conditioning of data centre.*
- *Reduce software-licensing costs.*
- *Reduce server maintenance costs which are about the companies and organization consider cost saving as a crucial factors when they implement the new technology.*

The third factor consists of the variables of:

- *Reduce infrastructure server sprawl by consolidating to fewer physical servers.*
- *Reduce data centre floor space for fewer physical servers.*

Thus this component is apparently related to the purpose of achieving consolidation through Server Virtualization. Table 4 summaries the assumed characteristics of the factors, related to their included variables.

Factors	Variables
Time reduction	Reduce time for systems management of infrastructure and servers
	Increase server resiliency and minimizing planned and unplanned downtime
	Increase the speed for deploying new servers
Achieve consolidation	Reduce infrastructure server sprawl by consolidating to fewer physical servers
	Reduce data center floor space for fewer physical servers
Cost reduction	Reduce software licensing costs
	Reduce server maintenance costs
	Reduce energy costs for powering servers cooling and air conditioning of data centre

Table 4 2nd Delphi factor analyses for question group one

Regarding the second question group of the Delphi survey, eleven questions had to be answered concerning Deployment, or after the fact

Another factor analysis with the same settings as before has been performed in order to shrink the overall hypothesizes of this part as well.

According to the analysis results (Table 14, Table 15, and Table 16 in Appendix) the variance of all variables could be explained up to 85.2% by three factors.

The *first factor* is based on the following five variables:

- *Measure and calculate possible power savings from their deployed virtualized server environment.*
- *Measure and calculate possible timesaving from systems management for the deployed virtualized server environment.*
- *Consider that their virtual server environment has increased server resilience against*

planned or unplanned downtime.

- *Calculate the Return on Investment for their virtual server environment.*
- *Consider that their virtual server environment was worth the investment.*

Accordingly the characteristic of this first component is linked to fact that organizations do care about cost savings when SV has been deployed.

The *second factor* includes the information of the following variables:

- *Monitor and evaluate the power consumption of physical servers and other data centre equipment such as storage.*
- *Test and evaluate the virtualized server resilience such as high availability or disaster recovery functionality.*
- *Implemented structured work processes and procedures for system management of their virtual server environment.*
- *Have a valid and up to date disaster recovery plan for the their IT environment.*
- *Have a Green IT agenda or plan for the their IT environment.*

The characteristic of factor two can thus be described as related to improved performance of IT equipment.

The *third factor* is described only by the variable of:

- *Monitor and evaluate the resource usage of virtual servers CPU memory disk network.*

Therefore it is linked to the general awareness of actual consolidation within an organization. In Table 5 the results are summarize for the second factor analysis part.

Factors	Variables
Cost reduction	Consider that their virtual server environment has increased server resilience against planned or unplanned downtime
	Calculate the Return on Investment for their virtual server environment
	Consider that their virtual server environment was worth the investment
	Measure and calculate possible power savings from their deployed virtualized server environment
	Measure and calculate possible time savings from systems management for the deployed virtualized server environment
Improve performance	Monitor and evaluate the power consumption of physical servers and other data center equipment such as storage
	Test and evaluate the virtualized server resilience such as high availability or disaster recovery functionality
	Implemented structured work processes and procedures for system management

	of their virtual server environment
	Have a valid and up to date disaster recovery plan for the their IT environment
	Have a Green IT agenda or plan for the their IT environment
Awareness of actual consolidation	Monitor and evaluate the resource usage of virtual servers CPU memory disk network

Table 5 2nd Delphi factor analysis for question group two

According to the results of the FA the original hypothesizes could be narrowed down into six more general ones, stated in Table 6.

Hypothesis
H_{1a} : Companies and organizations are choosing server virtualization so to reduce time
H_{1b} : Companies and organizations are choosing server virtualization so to achieve consolidation
H_{1c} : Companies and organizations are choosing server virtualization so to reduce cost
H_{2a} : After companies and organizations having deployed server virtualization, they still consider the cost reduction of IT equipment.
H_{2b} : After companies and organizations having deployed server virtualization, they still consider the improve performance of IT equipment.
H_{2c} : After companies and organizations having deployed server virtualization, they are awareness of actual consolidation of IT equipment.

Table 6 2nd Delphi Hypotheses question group one and two

The next step was to perform t-tests in order to check if the population means are different from a test value equal to three and therefore test if a certain variable is generally treated as important regarding SV.

Throughout the performed t-tests, the calculated *p*-value has been compared to the level of significance in order to recognize if either a hypothesis should be support or not. Regarding a significance level of $\alpha = 0.05$, almost all the *p*-values of the first Delphi part are smaller than the confidence value of $\alpha = 0.05$ expect for the hypothesis of “*Companies and organizations are choosing server virtualization so to reduce software licensing costs*” (see Table 7). Therefore the majority of hypotheses are supported by the data.

Hence most of the facts that are expressed by the hypotheses are involved in the decision process of organizations when considering deployment of Server Virtualization, except considering the reduction of licensing costs.

Part of Hypothesis		Test Value = 3	95% Confidence Interval of the Difference	
		Sig. (2-tailed)	Lower	Upper
1a	Reduce infrastructure server sprawl by consolidating to fewer physical servers	.001	.87	2.13
1a	Reduce time for systems management of infrastructure and servers	.000	.86	1.64
1a	Reduce energy costs for powering servers cooling and air conditioning of data centre	.001	.55	1.45
1b	Reduce data center floor space for fewer physical servers	.002	.66	1.84
1b	Reduce software licensing costs	.111	-.31	2.31
1b	Reduce server maintenance costs	.000	.94	1.81
1c	Increase server resiliency and minimizing planned and unplanned downtime	.005	.50	1.78
1c	Increase the speed for deploying new servers	.000	1.05	1.95

Table 7 One-Sample Test question group one

Considering the low sample size and taking into account the frequency distributions of the variable (Table 8) about cost reduction due to lower software licensing it should be noticed that 71 percent of the respondents did agree/strongly agree about its importance linked to SV.

Consequently our own interpretation of this variable, based on the frequency table and the experience made throughout the first round of Delphi and interviews, differs from the result of the t-test. Thus the variable “*Reduce software licensing costs*” should also be considered as important.

		Frequency	Percent	Valid Percent	Cumulative Percent
Valid	Strongly Agree	4	57	57	57
	Agree	1	14	14	71
	Disagree	2	29	29	100
	Strongly Disagree	0	0	0	100
	Total	7	100	100	

Table 8 Reduce software licensing costs

Moreover a second t-test, with the same settings as before, has been performed to check the new generated hypotheses of the second Delphi part. The results show that only the following two hypotheses are supported by the test:

- *Monitor and evaluate the resource usage of virtual servers CPU memory disk network;*
- *Consider that their virtual server environment was worth the investment and Consider that their virtual server environment has increased server resilience against planned or unplanned downtime*

Once more it should be noted that the test results needs to be handled with caution due to the small number of our respondents. Accordingly each scale unit accounts a lot in this t-test, which affects the accuracy of our data analysis.

Part of Hypothesis		Test Value = 3	95% Confidence Interval of the Difference	
		Sig. (2-tailed)	Lower	Upper
2a	Monitor and evaluate the resource usage of virtual servers CPU memory disk network	.006	.34	1.41
2a	Monitor and evaluate the power consumption of physical servers and other data center equipment such as storage	1.000	-.63	.63
2a	Test and evaluate the virtualized server resilience such as high availability or disaster recovery functionality	.316	-.59	1.59
2a	Implemented structured work processes and procedures for system management of their virtual server environment	.095	-.14	1.39
2a	Measure and calculate possible power savings from their deployed virtualized server environment	.351	-.42	.17
2b	Measure and calculate possible time savings from systems management for the deployed virtualized server environment	.685	-.57	.82
2b	Calculate the Return on Investment for their virtual server environment	.140	-.26	1.51
2b	Consider that their virtual server environment was worth the investment	.018	.24	1.76
2b	Consider that their virtual server environment has increased server resilience against planned or unplanned downtime	.020	.16	1.34
2b	Have a valid and up to date disaster recovery plan for	.442	-.71	1.46

	the their IT environment			
2c	Have a Green IT agenda or plan for the their IT environment	.197	-.25	1.00

Table 9 One-Sample Test question group two

Nevertheless the related frequencies mentioned in Table 17 (see appendix) showed that only the following two variables are not a indispensable considered by organization after deploying the SV technology:

- *Monitor and evaluate the power consumption of physical servers and other data centre equipment, such as storage*
- *Measure and calculate possible power savings from their deployed virtualized server environment*

The results of testing the new hypotheses are shown in Table 10 Table 14below.

Hypothesis	Supported
H_{1a} : Companies and organizations are choosing server virtualization so to reduce time	Supported
H_{1b} : Companies and organizations are choosing server virtualization so to achieve consolidation	Supported
H_{1c} : Companies and organizations are choosing server virtualization so to reduce cost	Supported
H_{2a} : After companies and organizations having deployed server virtualization, they still consider the cost reduction of IT equipment.	Partially Supported
H_{2b} : After companies and organizations having deployed server virtualization, they still consider the improve performance of IT equipment.	Partially Supported
H_{2c} : After companies and organizations having deployed server virtualization, they are awareness of actual consolidation of IT equipment.	Supported

Table 10 2nd Delphi Hypotheses test for question group one

According to Table 10 and the collected information during the Delphi method it can be assumed that reduction of time and costs, as well as consolidation of servers are important factors that lead to the deployment of SV.

After successfully having deployed the technology organizations do have a higher awareness of their actual IT equipment consolidation. Benefits linked to costs reduction and improved IT performance on the other hand are not considered significantly.

4.2 Interviews

For our personal interview series, we selected five (5) Subject Matter Experts (SMEs) from different server virtualization promoting companies, where the selection of SMEs was made using LinkedIn profiles and through personal knowledge, thus a non-probabilistic judgement sampling, in order to limit potential non-sampling errors due to lack of knowledge (Anderson, Sweeney, Williams, Freeman, & Shoemith, 2009). We also selected five (5) individuals responsible for server technology, management or deployment, or data centre operations, who were chosen based on personal knowledge and with the assistance of Pulsen Systems AB, a company in Sweden, who work with implementing server infrastructure solutions, and promoting server virtualization, to their clients.

The companies and organizations were coded based on our classification into Deployer and Promoter, in no particular order, as Deployers “DEP# ORG”, and Promoters “PRO# ORG”, and the staff from respective classification as “DEP#” and “PRO#” respectively, in the compilation and analysis below, from interview transcripts. Transcripts can be found (obfuscated) in Appendix.

4.2.1 Review of interview results

Many things have so far been stated in this thesis that point out advantages of virtual technologies, based on current multifarious literature sources. To find out about common thoughts, issues and constraints in practice, several interviews with virtualization vendors and users (see above) have been conducted during the research process. This resulted in discovery of the following common topic areas:

- Important decision points
- Other types of virtualization
- Importance of power reduction
- Concerns of Server Virtualization
- High Availability and Disaster Recovery
- Backup/Restore and Disaster Recovery
- Storage for Server Virtualization
- Key success factors and pitfalls
- Cooperation and managing expectations
- Measuring benefits of Server Virtualization

4.2.1.1 Important decision points

According to PRO3 (PROB ORG), who has been in this business for several years, the main attitude towards Server Virtualization has changed over the past years. Server Virtualization started out as an easy to manage test and development environment for the companies who implemented it during the time of '01 till '04. Since the technology developed continuously the most important reason today is regarding the cost savings that can be achieved through this technology. PRO3 stated that the costs of a physical server compared with having the same functionality in a virtual environment is round about 40% less in the latter one. Money can be saved through new licensing models, faster implementation and deployment, less hardware costs, lower operating costs through easier management, less cooling and space

demands as well as less power consumption in general. PRO2 (PROA ORG) stated that 5–10 servers could easily be replaced by only one virtual server environment.

Changing hardware demands based on capacity growing or the adjustment of system requirements of a running application are also arguments towards Server Virtualization. PRO4 (PROC ORG) offered that in the past Server Virtualization has been done due to well-known benefits of system flexibility, rapid deployment and easier management of the server environment. Today, IT departments are forced to do more with less workforce and within less time and the importance of lowering the overall costs of the IT has become much more important. Hardware costs, as traditional capital expenses and operating costs are lower in virtualized environments. Moreover, the ability to reschedule workloads during operations instead of having a service team working at weekends and nights on it can save costs.

Concerning the more energy efficiency, it needs to be pointed out that stand-alone servers are usually not running on full capacity. PRO2 (PROA ORG) mentioned a normal server is usually only utilized by 5%-10% of its full capacity, which consequently leads to a waste of a huge amount of the supplied energy of a machine since power can either be on or off. Putting several virtual environments into one physical machine allows up to 50%-60% of utilization and therefore uses around 85% of the maximum power consumption.

However, in the first place hardware costs can be realized by the consolidating of physical machines. Depending on their actual hardware demands several servers can be put into one virtual environment. Therefore fewer machines have to be bought, maintained and repaired if necessary.

Another fact that has been mentioned by the promoters is the feature of high availability. Especially for smaller companies, the amount of saved costs is not always worth the manpower and investment that is needed for a virtualization project. Therefore their main motivation lies in the high availability of business processes and functions, which are crucial today in terms of business activities that rely on data and applications. Having set up 2-3 identically virtual machines inside one physical machine, enables one to move applications or functions from one machine to another one. This also affects the costs regarding the total amount of hardware that is needed in a Data Centre and of course has impacts on the power-, cooling- and space- demands as well.

There are even examples of customers who did not invest in the technology because of cost reduction at all. Instead they decided to go with a more cost intense solution due to disaster recovery features in the virtual technologies that make your systems more robust against failure. Moreover, small companies are dealing with the Server Virtualization because it recently has become more interesting in terms of pollution policies. The demand of having a green image is increasing nowadays.

PRO3 (PROB ORG) summarized that about 60%-70% of the organization which invest into Server Virtualization do it because of cost reduction, around 20% are driven by better disaster recovery, safety aspects or high availability and about 10% still only uses it for more efficient test and development purposes. Usually, less energy consumption is not a main reason of deployment in PRO3 opinion.

Regarding the deployers' points of view DEP4 (DEPD ORG) mentioned that he wanted to minimize the physical hardware resources. Therefore, the goal was to use the current hardware resources as good as possible to lower costs of hardware maintenance, licensing and also minimizing the demand of space in the data center. Moreover, minimizing the overall power consumption and cooling demand within the Data Centre should be achieved through virtualization.

As another aspect, the increased ease of management of a fewer hardware, especially in terms of upgrading VMware was an important argument towards Server Virtualization. However, a complete 100% virtualized server environment is not DEP4 goal. Instead DEP4 is trying to virtualize as much as possible and if feasible putting all their high I/O machines into virtual environments. But main emphasis is still to minimize hardware, lower overall costs and get a easier platform to run.

When DEP1 (DEPA ORG) was facing the situation of purchasing new physical hardware for their purposes, he also thought about how to use it more efficiently in the future regarding the high energy consumption they usually have to deal with. Therefore Server Virtualization came into DEP1 focus as well as the benefit of easier system management for DEP1 technical workforce. Moreover the upper management of the DEPA ORG he works for wanted to centralize the data to open up the opportunity of outsourcing their systems in the future.

Regarding the DEPB ORG, they also an investment in new servers leads to the idea of putting their money on Server Virtualization because near future investments regarding more physical machines should not be an issue due to the expansion possibilities of virtual environments. Besides their old physical servers needed too much cooling and software licensing costs could be saved through applying this technology.

DEP3 (DEPC ORG) argued that the increased speed of server deployment within a virtualized environment was the main argument, which leads to their Server Virtualization project. Faster deployment is also related to less demands on hardware, space and any other connected consumptions and therefore result in lower overall costs. Moreover using the high availability feature of Server Virtualization was important, too. Half of their server park has been virtualized by now and even more virtual environments will substitute current physical servers. Accordingly new stand-alone servers are implemented very rarely these days.

According to DEP5 (DEPE ORG) a good pre-study has been performed in order to identify the actual benefits and features of this technology. The study covered divers topics like consolidation, power savings, better utilization, and reduction of hardware costs and reduction of staff resources in terms of less needed system administrators.

4.2.1.2 Other types of virtualization

According to the statements of the promoters of virtualization technology, Server Virtualization is by far the most common deployed solution. The main reason of this fact is the existence of its widespread knowledge and positive reputation. Other options such as storage and desktop virtualization are upcoming technologies, which still need some development. PRO3 (PROB ORG) said that 90% of PROB ORG customers want a Server Virtualization solution. Moreover the whole area of desktop virtualization might be more spotlighted during the next year. The average size of a Server Virtualization project is 100

servers or more and typical customers are big banks, insurance companies, and government institutions and large companies in general.

Accordingly the investigated deployers have a strong focus on Server Virtualization within their system architecture. The common opinion is to virtualize as many servers as possible in the future. Desktop virtualization might be an option for standard working machines to reduce the overall amount of used PCs due to cost reduction. DEPE ORG is struggling with low footprint capacities, regarding their external contractors who work from outside the building but all need an accessible (via VPN) workstation inside as well. Several hundred of them are in use, resulting in capacity issues linked to space and cooling. Therefore virtualizing those machines could be one of their future projects.

However, desktop virtualization will not be applied to people who work in positions where the amount of used applications is broad and new software is tried frequently, such as administrator, technician, scientific worker.

4.2.1.3 Importance of power reduction

The possibilities of saving power through Server Virtualization might not, as stated before, be a main reason of deployment regarding the opinions that have been collected of the interviewed promoters, but however it is still used as a common argument during discussions with potential buyers. The representatives of PROA ORG said that calculations can be made which show the difference between power consumption of normal physical servers and servers consisting of several virtual systems. An environment of 50 virtualized servers will save several thousand Swedish crowns a year. Another raw calculation is that an investment of SEK 500.000–600.000 has a payback for power only within five years of deployment. Taking all the cost benefits of the Server Virtualization technology into account will enable a payback time often after one or one and half years. Moreover they argued that power reduction less important in the private sector since only few companies do already split their IT costs from the rest of the whole power consumption costs.

A major issue about energy saving is that in most cases IT organization do not measure power consumption by themselves. It belongs to another responsibility and budget. Typically there is only one overall power invoice available covering a whole office building. Therefore, potential costs savings within the IT infrastructure would not directly benefit the IT department or IT manager since the saving are added to somebody else's budget. Another fact that describes the incomplete awareness of IT and power consumption is the situation that occurs for many providers of server outsourcing. They themselves buy facilities for their servers, which are often rented only by floor space that includes a fixed amount of power consumption costs. Consequently no cost benefits are achieved in this field through virtualization.

Within the public sector there is already a stronger emphasis on deploying greener technology. Also PRO4 (PROC ORG) confirmed that the general attitude towards Green IT/Computing has recently changed and customers want to show they have an environmental conscious due to the fact that power and cooling is becoming a pretty big part of IT operations. For instance, PRO4 (PROC ORG) has worked with a customer who couldn't install enough cooling in a server room, simply because the maximum power availability was already used in that area.

For DEP1, a representative of the public sector, the reduction of power consumption is one of the bigger benefits they are expecting from their Server Virtualization-project. He stressed that the local environmental office announced a new strategy, stating that local government should be more environmental focused. Moreover he is planning to improve the policy systems regarding the workstations used at some locations. They should shut down themselves automatically in a certain timeframe during the night or evening, when they are not in use.

DEP4 (DEPD ORG) mentioned that DEPD ORG platform is an expensive one and he is responsible for its cost reduction. Therefore he needs to focus on different aspects such as cooling, floor space, and cabling, licensing, maintaining and of course power usage. In addition a lot of new technologies is influencing the new build data centers of DEPD ORG. For instance cooling on different spots related to actual heat development and also re-using the heat that is created through the cooling process.

DEP5 (DEPE ORG) emphasized that lower energy consumption was their highest motivation for virtualization. Over the last years DEPE ORG has grown very much as a company. This has an impact on the IT demands as well, e.g. the number of hosted servers, storage and backup solutions. He explained that about two years ago they reached the limit of energy consumption within their Data Centre, regarding power supply and cooling demands. Therefore consolidating hardware through virtualization was one big alternative to rebuild or renew the whole Data Centre.

When asked about alternative ways to reduce the Data Centre power consumption DEP5 pointed out that due to the general low server utilization (10-15%) the large vendors of processor technology have released a new technology that allows setting the proportional power consumption of the CPU. For example it can run only on 50% power supply. If the application reaches a peak it can automatically request more power from the CPU until the operation is finished. Following new technology faster, through shorter lifecycles of one's current hardware is a general way to deploy the most energy efficient hardware.

Moreover building a new Data Centre in correct way, e.g. by placing the servers back to back and consequently making it easier to identify hot spots that need to be cooled down. Furthermore, studies have proven that a correct airflow within the DC allows the cooling temperature to be only at 25 degrees Celsius instead of a much colder temperature.

DEP2 (DEPB ORG) didn't consider the power benefits when discussing about the plans of a virtual environment. However, he is aware of this side effect and argues that another department is dealing with the issues of reducing the power consumption of the whole building, for instance through improvements of heating, ventilation and lighting.

DEP3 (DEPC ORG) offered that benefits of Server Virtualization regarding energy consumption were not a key aspect that led them towards its deployment. More important were cost savings concerning hardware consolidation, the faster implementation of environments and easier administration. Nevertheless, the heat produced in their server park due to cooling processes is used to heat their office- and a nearby school building.

PRO4 (PROC ORG) went deeper into the server technology when thinking about alternative ways to save energy: Higher utilization through Server Virtualization is the common way.

Another one is to analyze the whole infrastructure you are working with and identify individual processes that might not need to run at full speed and therefore not under full power supply. Tools exist which enable the restriction to a specific amount of Watts used by a single machine. In addition a complete shutdown of servers during non-peak hours is an option, especially when having a redundant version of the application, which is still available and can handle the minor usage by itself. Using (PROC ORG) blade servers can also reduce the power demand since several applications are sharing one set of power supply for all blade computers, within a blade chassis.

Using better cooling options, e.g. water instead of air-cooling, makes the cooling process more energy efficient. As an example of high scale heat re-usage, IBM Stockholm uses heat exchangers in one of their biggest cooling rooms (2000-2500 servers) which enable them to heat their office building during winter time (1200 employees) plus selling back some of the heat to the local utility company.

PRO5 (PROA ORG) also emphasized the over-dimension of most components in a DC. Powering up or down servers according to their usage peaks can save a lot of energy. The same applies to storage, since having discs running 24/7 without the possibility of spinning them down e.g. overnight causes gratuitous energy demands.

Low-power CPUs used in DCs are a further possibility according to PRO3 (PROB ORG). PRO1 (PROA ORG) argued that the investment motivation towards more efficient energy solutions would stay low as long as the energy consumption of the IT department is not accounted separately.

4.2.1.4 Concerns of Server Virtualization

All interviewees who gave insights from a promoter perspective could name several drawbacks and crucial topics that need to be thought about wisely in order to be successful with Server Virtualization technology. The PROA ORG representatives mentioned that hardware costs are almost the same compared to standard machines, but licensing costs can depend on the actual agreement with the third-party software vendors. Therefore customers need to have a good consultation before finally deciding whether to deploy Server Virtualization or not. Moreover there are applications, which should be managed isolated from the virtual environment due to easier security management, for instance firewalls and back-up systems.

In addition a proper overall system management is required. IT departments are responsible for observing the actual capacities of the virtual environment, e.g. dealing with capacity growth forecasting that more often results in the re-adjustments of dedicated system resources than buying simply new servers like in the old days. This is often a total new aspect for the responsible department when server consolidation through virtualization is performed. In cases where networks are also virtualized and managed via a server the issue of responsibility is raised once more, since it needs to be clarified if either the networking team or the server staff has to deal with its configuration.

PRO4 (PROC ORG) explained that the success of Server Virtualization implementation strongly depends on the existence of a structured service management process that is dealing with the deployment and commissioning of the servers. PRO4 stated that the phase where the old machine is substituted by the new virtual ones is crucial. Concerns regarding the handling

of the existing data, users, future licenses and warranties need to be figured out. Furthermore if there is a lack of control regarding server management and deployment processes it is more likely that Server Virtualization will even worsen the situation, since it enables the easy creation of new (and more) server environments in one physical box.

PRO5 (PROA ORG) mentioned that consolidation of servers into fewer physical machines leads to higher susceptibility regarding the overall IT robustness against failure. Vendors have to face this fact and provide the customers with satisfying solutions to protect them against any critical impacts of a significant hardware breakdown.

PRO3 (PROB ORG) argue that customers need to realize the linked costs to this new technology. Since it is generally new, staff needs to be educated as well as new hard- and software must be purchased. Acquiring knowledge does often lead to hiring different consultants than worked with before. Consequently new contract conditions will be faced. Moreover the costs for storage solutions are quite high, because advanced storage is needed when deploying Server Virtualization.

DEP4 (DEPD ORG) stated that depending on the kind of virtualization, new hardware is needed. Also the right amount of required software and therefore licensing costs should be considered. At the DEPB ORG the licensing problem is not present since unlimited Linux licenses are available and used as the main operating systems. However, monitoring the system has to be done from time to time to figure out if the memory dedication for individual virtual servers is still sufficient. The user satisfaction is assumed to be high, since they don't even know about the recent switch from physical to virtual environments and no complaints have been recognized so far.

DEP5 (DEPE ORG) said, that a big part of the cost benefits of Server Virtualization is reduced through high licensing costs of VMware (x86 virtualization), and new needs of external and internal knowledge about that technology. In addition not all applications are suitable for Server Virtualization, e.g. the ones that demand high memory and CPU capabilities. Virtualizing these applications doesn't make sense (currently) since they usually need their own physical machine anyway.

DEP1 (DEPA ORG) had doubts about the user satisfaction within DEPA ORG, mainly because of speed issues, which might appear in their virtual environment. In the past, local servers were set up in each facility. Nowadays the Server Virtualization projects include plans about a central point where the data is stored. Consequently bigger bandwidth is needed but only a 100MB connection between the facilities and the central point is available. DEP1 also mentioned future education costs for being able to manage the new technology.

All deployer participants mentioned a strong cooperation with their virtualization vendor during the time of designing, implementing and testing the new servers. This relationship loosens up during the later stages and more personal decisions of the actual users are taken into account. DEP4 (DEPD ORG) states that delivered servers are pretty much pre-tuned when they arrive, meaning that images of operations systems are already there and only have to install. Moreover during the process of implementation and testing, unused standard applications will be removed and additional ones will be added in order to customize each virtual machine regarding its special needs.

At DEPE ORG a review of service characteristics have been performed for physical machines, which had been planned to virtualize. Test and development environments were the first ones that have been virtualized, which has saved their application development team a lot of time and money. The actual virtualization process that followed, was managed step by step since new virtual servers needed to set up first, before any movement of old servers towards them was possible, also with respect to the low energy capacity within the Data Centers.

4.2.1.5 High Availability and Disaster Recovery

According to PRO4 (PROC ORG) the field of high availability is strongly related to the demand of business availability. Questions about the need of supported business processes, their availability level and service agreement with customers and their individual ranking compared to other processes need to be discussed in order to supply them with the right amount of hardware according to their priority.

PRO1 (PROA ORG) and PRO2 (PROA ORG) mentioned that the possibility of moving applications in case of any problem within one physical machine would result in no extra hardware and software licensing costs. Still, monitoring and continuously tests are needed for sustainable availability. PRO5 (PROA ORG) pointed out high availability is one of the main features of Server Virtualization and therefore it's not an option to set it aside. Having several functionalities in only one physical machine requires high availability concerns in case of breakdowns.

PRO3 (PROB ORG) explained there are two different versions of high availability. The first one comes as a standard feature of Server Virtualization. High availability in this case is much more cost effective than any other comparable solution. On the other hand does a real solution of high availability cause additional costs that would shrink the overall cost advantage of Server Virtualization, because available, but unused, redundant hard- and software is required.

From a user perspective, DEP1 (DEPA ORG) related high availability to better server performance and the fact troubleshooting could be done faster through virtualization. For instance no more on-site service is required to restart a server. DEP2 (DEPB ORG) offered their goal was to achieve 100% of server availability and so far they almost got it. Downtime, due to upgrades or other maintenance work are usually around 10 or 15 minutes and therefore much shorter compared to the old physical server strategy. At DEPD ORG, cluster software is used within the virtualization environment to achieve a high availability.

According to DEP4 (DEPD ORG) two nodes (Virtual Machines/Logical Partitions) are allocated resources (physical memory, virtual/physical processors, virtual/physical adapters). Moving these resources from one node to another is possible. Since the second node features parallel running resources it can take over workload from node one (fast restart). Moreover less downtime is achieved through the redundancy of virtual engines. Whole partitions therefore can be moved from one physical machine to another.

At DEPE ORG the feature of high availability is used for maintenance purposes, when some servers have downtime due to it. They have also created resource pools by connecting several (10) servers with each other and installing 50 virtual machines on them. In a scheduled downtime for maintenance actions, they can run firmware updates on 50% of that resource

pool, bringing down five of the physical servers, upgrade- and then take them up again. Having all services online again the other half will be upgraded, leaving no impact of the maintenance on the production environment.

DEP3 (DEPC ORG) said that one big advantage of Server Virtualization is the ability to set up- or making copy of existing environments in a fast and easy way. However, no life depending services are running on their servers. They usually deal with communication services for which high availability is not necessary but a nice-to-have feature (convenient).

4.2.1.6 Backup/Restore and Disaster Recovery

PRO2 (PROA ORG) explained that a common backup solution is IBM's TSM technology. It was used before with physical machines only but can also be adapted to virtual servers. The servers have back-up clients that send information over the network to the TSM physical machine, which has a database. Therefore the TSM knows which data is on the server and can create back-up files on tape or other media. In the first place, a full back-up will be performed, followed by incremental back-ups updating changes of the data. In case of restoring a certain file or directory can be recreated based on the backup information. The back-up server (TSM) is still kept as a physical machine to enable easier restore procedures.

PRO5 (PROA ORG) thought that one big advantage of Server Virtualization regarding backup or restore is the ease of handling the servers as virtual images or files. This makes it a lot easier to perform backup procedures. Once available, they can simply be restored on the servers. No hardware dependencies such as drivers are needed because virtualization (current x86 virtualization) allows a separation of the server's hardware demand by providing a virtualization layer within the machine that covers the device driver communication (front and back device drivers). This saves a lot time compared to a stand-alone solution where an OS restore is a complicated process since more sources for error are available.

PRO3 (PROB ORG) stated no good solution so far exists that covers both, virtual and physical server environments, at the same time. Therefore a good solution has to be chosen for each of them resulting in a mix of backup tools. About the size of the actual backup files PRO3 continued that full backup images are useful in case of disaster recovery, but since PROB ORG rely on large amount of available storage, tools that can restore single files are also popular. Server Virtualization in general offers possibilities of disaster recovery for a larger volume of servers. Before this technology was available, deployers could only equipped the most business critical servers with disaster recovery plans in terms of mirroring.

The majority of the promoters pointed out that Server Virtualization have also advantages when speaking about disaster recovery, but it is a rather complex solution that requires proper planning and handling. Disaster recovery, according to PRO5 (PROA ORG), is inherently connected with Server Virtualization. Having a one-server-one-function architecture, a failure will stop that individual service and affect a set of users. In case of having 40-60 services in one physical box which might affect a lot more users, a sophisticated recovery plan is needed that covers break downs and repair services of the box. These plans must be built in into the virtual environment from the very beginning. Otherwise no guaranty regarding a high quality of services is given, because eventually something will break and affect each running system of the server.

The DEPD ORG IT department is dealing with bootable system images, which are saved, in a repository server. Disaster recovery installation will be done through these images. Since less hardware is used, due to Server Virtualization, the restore time is positively affected (fewer disk images need to be restored). On the other hand it could take quite long to restore if a loss of a whole physical server (or storage system) with multiple virtual environments on it (without high availability or disaster recovery enablement active). Therefore proper disaster recovery planning is required that covers the worst-case scenarios of complete hardware failure. In the case of DEPD ORG, external storage is used. Consequently a complete crash of one of the physical machines doesn't have any major impact. Moreover mirroring is used between their storage subsystems. Concerning backup and restore in DEPD ORG virtualization environment, every virtual machine features a backup system (IBM TSM), enabling them to run incremental backups based on an elementary complete one, which has been created as described above.

The DEPB ORG is using mirror disks and regular backup tools on their servers. Backups rely on IBM technology and are tape based and store up to eight generations (versions) of individual files. To overcome a disaster situation, the backups are stored at two different locations.

DEP5 (DEPE ORG) stated that due to their diversity of systems no default set up for disaster recovery is available, since it is highly system dependant. About 80% of the systems only need a common tape backup of their servers and databases. Other systems require more investment by having the backup storage on two different locations. Not much virtualization technology is used for backup processes in their case.

DEP3 (DEPC ORG) argued that the creation of new environments or restoring of them in failure scenarios has become easier in the virtualized environment.

4.2.1.7 Storage for Server Virtualization

The common reaction of the promoters regarding the question of storage handling in virtual environments was that it is cost intensive due to the additional demand of storage hardware. Correct planning is the basis of keeping the storage costs down and achieving good utilization, and also to assure that storage is available on different systems within the virtual environment. Storage area network (SAN) is quite common and enables central storage of data and is commonly used, especially by larger companies and organizations.

DEP1 (DEPA ORG) agreed to that fact, since data centralization was one of the goals during the virtualization process.

PRO5 (PROA ORG) explained that shared storage is mandatory. A dual storage environment makes sure that there are available online copies on disk in a location that is separated from the normal production environment. In other words, two separated physical servers placed at different locations, with separate storage equipment, are needed. A common issue about the right storage treatment is the inconsideration of the I/O volume. Most of the time the storage volume is discussed in terms of gigabytes instead of I/Os per second when selecting a storage environment. Stand-alone servers have the ability to perform high I/O since they are underutilized. When consolidating many of these servers into one central storage system, it seems to be the view that it is often overlooked that the connected reduction of disks leads also to reduction of the I/O capability of the storage system. Therefore proper dimensioning

of the storage system has to be figured out in order to handle the I/O bandwidth, which all those servers did un-virtualized before, because the overall I/O will still be the same, and possibly increase over time.

DEPD ORG-IT uses virtualization for storage whereas the DEPB ORG is still looking for a good solution. They are struggling with forecasted storage shortness in the near future but no investments have been made so far to overcome it. The lack of money and expertise has been stated as the two main reasons.

4.2.1.8 Key success factors and pitfalls

According to the promoter fraction a good project preparation is significant for a successful outcome. Identifying needs and demands of the customer consists of knowing which kinds of servers are interesting for a specific virtual environment and analyzing the original loadings of the machines. Furthermore characteristics and behaviors of services that will be put into the virtual environment in terms of memory consumption, workload and I/O need to be explored. Projects will fail if the machine's workload hasn't been checked and too many (e.g. 15) of the heaviest loaded physical machines, having peaks at equal times, will be put in the same virtualized physical machine. In addition, trying to keep it simple in the first stages of Server Virtualization is important. Primarily the easiest and most efficient application servers should be virtualized.

PRO2 (PROA ORG) advised to virtualize approximately half of the systems in the first place, evaluate it and then decide about further steps of virtualization. Systems that are easy to virtualize are usually test- and development servers, administrative services, file- and email-services. More complicated ones are systems with sophisticated hardware dependency as well as high performance systems, e.g. huge databases or other unique environments with high computing and memory demands.

PRO4 (PROC ORG) also emphasized the preparation phase with a focus on the handling of processes. According to PRO4 opinion it is quite easy to achieve physical consolidation but there is always the threat that the virtual environment will be growing without having control over it.

PRO3 (PROB ORG) saw the time aspect as the most important success factor. Spending more time than estimated to migrate the planned amount of servers into the new environments will impact the ROI analysis negatively. Consequently the payback calculations will turn out to be wrong and cost benefits will be achieved later since having an investment of several million SEK on the one hand and still running the majority of their cost intensive physical servers on the other.

Also high availability was mentioned by PRO3, and the fact to plan and build a redundant solution right away in order to have satisfied customers, to enable them to keep their services online in case of failures. This also covers the opinion DEP2, who suggested more uptime and lower downtime of the systems. DEP2 also stressed the economical aspect of saving costs through hardware consolidation. DEP4 (DEPD ORG) agreed with the scrupulous project planning factors, for instance avoiding potential problems as much as possible through the analysis of the current environment in terms of utilization.

Server Virtualization projects might fail in case of a bad maintenance management, which is not able to keep track of all existing disks, partitions, memory and processor usage. Regarding any problems during the implementation phase of virtual systems DEP4 (DEPD ORG) argued to keep it quite simple and skip most of the sophisticated features. Furthermore the service personnel who are responsible for the system maintenance and management need the right level of expertise. Another main success factor is the customer satisfaction and the decision to invest in a proven technology.

DEP5 (DEPE ORG) pointed out the need of a correct pre-study regarding the demands and possibilities of Server Virtualization technology, covering different hard- and software solutions. Problems during the deployment process of virtualization might appear when switching from a physical to a virtualized environment and being in a consolidation mode. Some applications might also not be supported to run in a virtual environment. Often this is due to their licensing models and -specifications, which are sometimes formulated unclear in licensing agreements, and as a result can cause administrative problems or even lawsuits against the deploying organization.

DEP3 (DEPC ORG) stated that is important to convince the organizations staff regarding the usage of Server Virtualization technology. There are still some doubts about it and therefore key aspect is to make everyone feel confident with its deployment. A big issue of Server Virtualization is the creation of quite complex dependency chains within the virtual environment. It was easier to keep track of the stand-alone servers. Consequently it is important to have an up to date environment-map of all currently used systems.

DEP1 (DEPA ORG) noticed that low energy consumption and easier system management, in terms of centralized data, user handling and technical services are DEP1 personal success factors. However, environmental benefits have never been considered, they are just considered a positive side effect.

According to PRO1 (PROA ORG) Server Virtualization project can fail respectively don't match the expectations of the customer, if high availability or lower power consumption does not appear to be important to them since prices for standard server hardware are dropping continuously. Furthermore the business owner or upper management needs to support the whole project. Sometimes there are doubts about the technology because it is hard to visualize for non-experts compared to dedicated physical servers. Moreover a lack of motivation can be present if the old hardware setup is working well and without major problems. Consequently the estimated timeframe of the project will be exceeded and result in much lower cost benefits and longer payback times.

The team around PROA ORG has experience with supporting and rescuing almost failed Server Virtualization projects where companies tried to deploy the technology by themselves without having the required knowhow. However, problems due to technical difficulties do usually not occur. Nowadays, the Server Virtualization technology is well developed and getting a delivered machine into working mode is generally done by click and install. However, there are essential things to estimate beforehand during planning, as already mentioned.

4.2.1.9 Cooperation and managing expectations

Concerning different considerations of promoters and deployers regarding the Server Virtualization technology PRO3 (PROB ORG) stated that some of the customer's software doesn't support Server Virtualization. Either it cannot run in the new environment due to compatibility or licensing issues from third party software vendors. Furthermore PRO3 often gets positive feedback about the ease of use of the technology. In contrast, some users are not fully satisfied about the actual ROI since they didn't realized the full amount of costs and the high demand of a fast deployment in the first place.

Sometimes customers have causeless doubts about security aspects when being aware that systems of different companies could run virtualized on the same physical machine.

PRO4 (PROC ORG) pointed out that some customers might not have expected the full range of related costs. They build the cost case on what they can see and touch, regarding how much money they have to invest in the first place on new servers, new virtualization software, energy consumption, etc. but they miss to assume some of the big soft values that will also make investment worth while. It is hard to put figures on these values in terms of cost efficiency. They are related to common benefits such as speed of deployment, flexibility and the possibility to rapidly shift workloads or resources from one virtual environment to another one.

According to PRO2 (PROA ORG) customers are often surprised about the overall increased system speed, since a 100% workload is usually not reached in any sub-system.

However, high customer satisfaction regarding Server Virtualization is the common case based on the collected feedback of the promoters and deployers.

PRO5 (PROA ORG) explained that some customers don't know their current system well enough which makes correct capacity planning more difficult and time intense. Furthermore he stated that customers with the right level of expertise usually don't have much trouble with the technology. On the other hand, there have been cases where all systems are doing fine in the beginning but after putting in too much services into one physical machine, the environment did slow- or eventually break down.

All of the interviewed deployers pointed out to have a good collaboration with their vendors, having more or less regular meetings and giving and receiving feedback continuously and being informed about future technologies. Efficient project management methods and expertise were the main reasons of their satisfaction.

Nevertheless DEP5 (DEPE ORG) argued that large IT vendors are creating new technology "hypes" every year. Green IT/Computing is the current one and virtualization was so to, 1-2 two years ago, which is interesting given that this technology has been around since the 1960s. It is important to not become totally dependent on one of these vendors and trust only them and their advices blindly. Organizations should always try to get a second opinion by building up long term relationships with different vendors, especially smaller local ones who can consult (assumedly) less biased.

4.2.1.10 Measuring benefits of Server Virtualization

Regarding methods that enable the actual measurement of deploying Server Virtualization technology within an organization PRO3 (PROB ORG) mentioned that ROI calculations for the customer is developed based on their existing environment and current costs. Accordingly PRO5 (PROA ORG) stated that doing the paper exercise is one way, but the really interesting part is to measure the energy benefits in the real computer room and be able to conclude after a year of deployment if goals regarding power consumption have been reached. Tools from virtualization soft- and hardware vendors can be used for that, as well as observing the actual UPS rate or Watts that is supplied to the Data Centre.

PRO4 (PROC ORG) explained hardware costs and their related savings are easy to calculate. The finance department is able to compare the investments linked to Server Virtualization and stand-alone solutions. Available monitoring tools facilitate the ascertainment of certain energy usages (specific servers or services), which can then be referred to actual costs. A final reduction of management costs is harder to estimate since comparable figures must be known beforehand which is normally not the case. On the other hand, customers are able to disclose numbers regarding licensing costs, power and cooling, as well as maintenance costs. However it is also possible to put figures on the project afterwards, for instance through facts like how many people could be moved away or laid off from the IT department due to lower workload.

PRO1 (PROA ORG) pointed out the issue of correct measurement by stating the problem of comparison, which lacks through the unavailability of correct figures. For instance savings based on software licensing costs are easy to calculate since detailed invoices and price information are available. Power consumption is rather difficult because most organization don't have specific invoices for their servers. Manpower could be compared in terms of needed staff capacities during the deploying and maintaining phases, but on the other hand the excess staff can easily work on other projects or have new duties assigned. The investment in new hardware will pay off in the long run, as newer machines are more efficient and lead to lower service and maintenance costs.

PRO2 (PROA ORG) added that if cost saving calculations is correct, a payback should be reached after 1-1½ years of deployment. If this is not the case the previous ROI calculations is wrong.

DEP1 (DEPA ORG) has to deal with the fact that there are no available tools within their server environment yet, that can monitor the power consumption and power savings (to their knowledge). Nevertheless DEP1 argued that Server Virtualization results in a logical reduction of hardware and therefore should lead to less power and cooling demands. As a consequence less power is consumed within the server park. DEP1 hoped to gain the main energy saving through the future power management tool that is able to shut down individual workstations, when not in use, automatically at facilities. In DEP1 case the real amount of costs savings cannot be figured out or calculated due to the lack of access respectively responsibility connected to the current energy costs.

Furthermore IT costs are included into the total rent a specific facility has to pay. Therefore it's not the economical factor in the first place, which motivated DEP1 towards the Server Virtualization project. Instead DEP1 is trying to increase the ease of server management and do this in part towards improved local energy efficiency. A possible way to measure the cost

savings is to talk to the technical office, which is responsible for paying these costs, and make them compare the costs before and after the Server Virtualization project. However there are other factors that need to be considered, for instance if building heating must be increased if the individual workstations in a facility are shut off.

Also DEP2 (DEPB ORG) is not able to measure the actual energy consumption related to the DEPB ORG server park. DEP2 noticed that there is more unused space in the server room, as a result of the server consolidation. Still DEPB ORG cannot see if the power consumption is going down since once again the payment of energy belongs to another department.

DEP3 (DEPC ORG) mentioned that no costs saving measurements have been made so far. Related figures do exist, but a comparison doesn't make much sense since their server park is continuously growing at the moment.

At DEPD ORG things are different, DEP4 (DEPD ORG) and team have insights into monthly and even daily power consumption of their machines. DEP4 also agreed that hardware purchases and linked cost benefits are easy to calculate. Licensing costs and possible costs savings are more difficult to figure out, since different licensing models need be considered and kept track of. In DEP4 case there is insight into the invoices of the Data Centers floor space costs, power consumption, hardware, and maintenance costs.

Moreover capacity forecasting tools are used to react quickly regarding new hardware demands and utilization balance issues. He pointed out that knowledge about the actual utilization is important in order to identify the situations when one has to stop putting in environments into one box and instead create another new one. Forecasting costs is also crucial because an investment request regarding more hardware demands takes some time to get through different management authorities.

DEPE ORG is not able to measure the actual benefits of virtualization technology at the moment. To overcome this lack DEPE ORG started a project, making figures of power consumption-, cooling-, hardware-, staff- and licensing costs more transparent and enable more correct ROI calculations.

5 Conclusions

To summarize and conclude the research results, and our separate analysis from the applied methods, Literature review, Delphi survey and Interviews, we offer the following conglomerated conclusions, referring to chapter 2 Green IT/Computing, Data Centre Power Consumption, and Server Virtualization, chapter 4.1 Delphi, and chapter 4.2 Interviews.

The concept of Green IT/Computing is perceived as a necessity for all companies and organizations to embrace, to maintain a sustainable growth in delivering increasingly higher computing resources. Both from a company ethical perspective, but also from a cost savings perspective, where the ethical can be viewed as a part of a company or organizations “Goodwill”, whereas the cost savings mostly will be a factor for larger national or international companies and organizations where the bill for power consumption or data centre infrastructure are noticed and accounted for.

The core savings regarding Green IT/Computing, come from the infrastructure and utilization of computing, storage and network resources, within a Data Centre, and from infrastructure and utilization of employee workstations, and from infrastructure and processes for inter personal communication and meetings, and last but not least from material recycling when discarding used, discarded or replaced IT equipment.

In the context of Data Centre environments, it is essential to design the Data Centre and company/organization IT infrastructure for energy efficiency, to right size the physical infrastructure, and to ensure efficient usage of physical infrastructure. All of which are easier said than done, especially in a corporate climate with quarterly planning goals.

Important benefits that usually lead towards a server virtualization deployment have, as mentioned, different reasons. A general increase of IT equipment in today’s organizations forces IT-managers to save costs wherever possible. Therefore costs savings are considered to be a main factor mostly due to hardware consolidating, which can lead to better server utilization and less demands for additional hardware, cooling, space and system management.

Moreover certain software licensing models do result in lower software usage costs within a virtualized environment. Furthermore the high availability feature of server virtualization can be considered to be very important to some groups of companies and organizations, because long downtimes of certain services will cause deficits e.g. in payment, regarding the business processes that rely on IT infrastructure. Doubts regarding server virtualization do usually exist about licensing models. They highly depend on the software product and therefore a lot of research has to be done by the IT department in order to avoid additional costs or even indictments due to misuse. Moreover there are applications that shouldn’t be put into a

virtualized environment at all owing to their inconsistencies with hardware or licensing support.

Maintaining business availability can be realized through providing the right amount of hardware resources to a certain service, which in some cases can end up in a more cost intense solution. In addition, server virtualization is also utilized because test- and development environments can be deployed fast and easily.

Besides server virtualization, the virtualization of desktop PCs or workstation machines might grow in importance in the near future. Large organizations need to run a lot of these machines and could save a significant amount of money through virtualization technology due to the same reasons mentioned above regarding Server Virtualization.

Consequently new capacity growths can only be realized if more efficient technology is used, as an alternative to build additional data centres. A minority of organizations is deploying server virtualization to gain a green image, so far. This small trend might become more important in the public sector first, before affecting the private sector.

Other solutions to save energy within a data centre have been identified during this research process. For instance re-using the heat that has been created through common cooling procedures in order to heat up nearby office buildings. Other options are to follow new technology approaches faster, e.g. using newest CPU products to maximize server utilization or using tools that allow the up- and down powering of CPUs, or entire servers, due to their general workload peaks. A lot of energy can be saved by sophisticated cooling, for instance through using the more efficient water cooling, instead of air cooling, and by placing the server racks in such a way that it facilitates air cooling, or optimizing the air flow within the server room.

In general, very good internal and external expertise is demanded to analyze IT infrastructures beforehand to identify opportunities as well as problems that might appear through server virtualization project. In particular the transition from a stand-alone to a virtualized services is linked to many issues that need to be clarified. Companies and organizations are not always aware of the overall range of costs that will affect them when deploying a server virtualization solution. Therefore external knowledge is one of the main impacts that need to be taken into account of such an investment, along with new hard- and software demands.

Accordingly, backup-, restore-, as well as disaster recovery- procedures are in general faster and easier to manage in a virtualized environment but demand elaborated planning. Especially in case of whole server machine failures, affecting several virtualized services at once, need to be considered when creating such a plan. A common tool that covers back-up activities of stand-alone and virtualized environments seems not to be available currently. Moreover a default solution is very hard to achieve because of the diversity of systems in the Data Centre.

High availability and decent backup systems result in additional demands of storage hardware. Therefore correct planning is needed to keep the storage costs down and gain a high utilization of storage resources. The Storage Area Network (SAN) solution is a

commonly used approach, sometimes accompanied with Storage Virtualization (not covered in this research).

According to both the Promoters and Deployers, a good project preparation is significant for a successful outcome of a server virtualization project. Identifying the current IT infrastructure is essential to deduce demands and constraints of a virtual server solution. Most important is analysis of the original workloads of the machines and the characteristics and behaviours of all services that are to be run in a virtualized server environment. The easiest and most efficient applications should be virtualized in the first place, to get familiar with the technology. In the long perspective the goal should be to virtualize as much as possible, but always having in mind that 100 percent will most likely never be reached since some applications are simply not suitable as mentioned before.

The overall time that is needed to accomplish the server virtualization project is also crucial since it has a big impact of the ROI calculations.

Moreover the existence of a good maintenance management is needed. Having consolidated hardware and the possibility to create new server environments fast and easy leads to the demand of keeping track of the whole server infrastructure and the related hard- and software resources (especially software licenses).

Other important (success) factors that have been identified throughout this study are Deployer (customer) satisfaction, and high availability of business processes. High availability combined with the benefits of less power consumption shouldn't be underestimated since a decision towards server virtualization, which is only based on less hardware costs, might end up as a failure, since prices for stand-alone servers are falling continuously (price/performance ratio) and the additional cost factors that comes with server virtualization are significant. In particular the full range of costs is often not apparent for potential deployers, as previously mentioned.

Therefore a good collaboration with external knowledge owners (Promoters) could be beneficial, which was always the case for organizations that have taken part in this study. Acquiring professional advices from multiple sources decreases the dependency of any one promoter, and can also help to open the deployers eyes in terms of distinguishing between their real needs and "hyped" solutions.

Regarding the measurement of the server virtualization benefits the Deployers made clear that possibilities exists to calculate the concrete ROI based on the most impacts of the server park. On the other side, Deployers usually lack the needed figures due to nonexistent access or availability of them. Especially power consumption is hard to figure out, since linked costs are usually included in the overall building costs where the data centres facilities are located. Therefore most of the interviewed Deployers, regardless of their company size, don't measure the actual energy savings. Overall no real measurement of the benefits is done generally in practice.

Our conclusions during this Thesis research are that there exist a difference of opinion regarding how to factor power consumption reduction from server equipment, both from Promoters and Deployers.

However, it was a common view that power consumption reduction was usually and possibly achieved, but not necessarily considered, and thus not evaluated, as neither a important (success) factor, nor that actual power consumption was measured or monitored after server virtualization deployment. In most cases the IT department is simply not in charge of analyzing its current power supply invoices, which might be especially true for small and medium sized organizations. On the other hand, organizations with huge server farms can be considered to be highly aware of this benefit, since they often operate at the local limits of available power supplies.

We found, as mentioned above, that other factors seemed more important, such as lower cost through higher physical machine utilization, simplified high availability of application server systems, and overall disaster recovery capabilities through a virtualized server environment.

Referring to the main purpose of our thesis, see chapter 1 Introduction & Problem Area, *the reduction of energy consumption cannot be considered to be a main argument towards the deployment of server virtualization, which is supported both by our Interviews and Delphi qualitative and quantitative analysis.*

Therefore the benefits of server virtualization technology linked to less power consumption is usually noticed as a positive side effect, are also supported both by our Interviews and Delphi analysis.

Although, there are notable examples that deviate, see 4.2.1.3 Importance of power reduction, warranting further study to identify possible differentiators between organizations and companies in this regard, not covered in these studies.

6 Further Study

We found that another topic, which could be beneficial to explore, is to also review the Deployer organizations technical IT environment, to be able to correlate their responses to the size and complexity of their IT environment. This was not part of our research, but it could account for some outlier responses regarding in particular the view on the importance of power consumption reduction in a Data Centre environment.

We would also suggest a larger survey of Deployer companies and organizations, which in our research was limited to interviews.

A Delphi Statistical Tables 2nd round

	Initial	Extraction
Reduce infrastructure server sprawl by consolidating to fewer physical servers	1.000	.774
Reduce time for server provisioning and deployment of new servers	1.000	.549
Reduce time for systems management of infrastructure and servers	1.000	.844
Reduce energy costs for powering servers cooling and air conditioning of data centre	1.000	.950
Reduce data center floor space for fewer physical servers	1.000	.752
Reduce software licensing costs	1.000	.822
Reduce server maintenance costs	1.000	.747
Increase server resiliency and minimizing planned and unplanned downtime	1.000	.991
Increase the speed for deploying new servers	1.000	.825
<i>Extraction Method: Principal Component Analysis.</i>		

Table 11 Communalities question group one

	Component		
	1	2	3
Reduce infrastructure server sprawl by consolidating to fewer physical servers	-.548	-.591	.346
Reduce time for systems management of infrastructure and servers	.893	-.109	-.130
Reduce energy costs for powering servers cooling and air conditioning of data centre	-.831	.400	.322
Reduce data center floor space for fewer physical servers	-.827	.030	.307
Reduce software licensing costs	.233	.705	.544
Reduce server maintenance costs	.453	.751	-.044
Increase server resiliency and minimizing planned and unplanned downtime	.735	-.039	.674
Increase the speed for deploying new servers	.533	-.646	.421

Extraction Method: Principal Component Analysis.

Table 12 Component Matrix question group one

Component	Initial Eigenvalues	Extraction Sums of Squared Loadings		
	Cumulative %	Total	% of Variance	Cumulative %
1	41.471	3.732	41.471	41.471
2	66.566	2.259	25.095	66.566
3	80.597	1.263	14.031	80.597

Extraction Method: Principal Component Analysis.

Table 13 Total Variance Explained question group one

	Initial	Extraction
Monitor and evaluate the resource usage of virtual servers CPU memory disk network	1.000	.979
Monitor and evaluate the power consumption of physical servers and other data center equipment such as storage	1.000	.850
Test and evaluate the virtualized server resilience such as high availability or disaster recovery functionality	1.000	.820
Implemented structured work processes and procedures for system management of their virtual server environment	1.000	.789
Measure and calculate possible power savings from their deployed virtualized server environment	1.000	.656
Measure and calculate possible time savings from systems management for the deployed virtualized server environment	1.000	.813
Calculate the Return on Investment for their virtual server environment	1.000	.832
Consider that their virtual server environment was worth the investment	1.000	.907
Consider that their virtual server environment has increased server resilience against planned or unplanned downtime	1.000	.891
Have a valid and up to date disaster recovery plan for the their IT environment	1.000	.839
Have a Green IT agenda or plan for the their IT environment	1.000	.995

Extraction Method: Principal Component Analysis.

Table 14 Communalities question group two

	Component		
	1	2	3
Monitor and evaluate the resource usage of virtual servers CPU memory disk network	.177	.444	.866
Monitor and evaluate the power consumption of physical servers and other data center equipment such as storage	-.724	.509	.259
Test and evaluate the virtualized server resilience such as high availability or disaster recovery functionality	.303	.842	.139
Implemented structured work processes and procedures for system management of their virtual server environment	.331	.807	-.170
Measure and calculate possible power savings from their deployed virtualized server environment	.699	-.409	.010
Measure and calculate possible time savings from systems management for the deployed virtualized server environment	.900	-.056	.014
Calculate the Return on Investment for their virtual server environment	.876	.027	-.253
Consider that their virtual server environment was worth the investment	.947	.054	.085
Consider that their virtual server environment has increased server resilience against planned or unplanned downtime	.922	-.093	.176
Have a valid and up to date disaster recovery plan for the their IT environment	.136	.905	.014
Have a Green IT agenda or plan for the their IT environment	-.009	.777	-.625

Extraction Method: Principal Component Analysis.

Table 15 Component Matrix question group two

Component	Initial Eigenvalues	Extraction Sums of Squared Loadings		
	Cumulative %	Total	% of Variance	Cumulative %
1	41.693	4.586	41.693	41.693
2	72.824	3.424	31.131	72.824
3	85.181	1.359	12.357	85.181

Extraction Method: Principal Component Analysis.

Table 16 Total Variance Explained question group two

Item	Strongly disagree	Disagree	Neither agree or disagree	Agree	Strongly Agree	No opinion	Total
Monitor and evaluate the resource usage of virtual servers (CPU, memory, disk, network)			25.0%	62.5%	12.5%		8
Monitor and evaluate the power consumption of physical servers and other data center equipment, such as storage		25.0%	50.0%	25.0%			8
Test and evaluate the virtualized server resilience, such as high availability or disaster recovery functionality	12.5%	12.5%		62.5%	12.5%		8
Implemented structured work processes and procedures for system management of their virtual server environment		12.5%	25.0%	50.0%	12.5%		8
Measure and calculate possible power savings		12.5%	87.5%				8

from their deployed virtualized server environment							
Measure and calculate possible time savings from systems management for the deployed virtualized server environment		25.0%	37.5%	37.5%			8
Calculate the Return on Investment for their virtual server environment		12.5%	37.5%	25.0%	25.0%		8
Consider that their virtual server environment was worth the investment			25.0%	37.5%	25.0%	12.5%	8
Consider that their virtual server environment has increased server resilience against planned or unplanned downtime			37.5%	50.0%	12.5%		8
Have a valid and up to date disaster recovery plan for the their IT environment	12.5%	12.5%	12.5%	50.0%	12.5%		8
Have a Green-IT agenda or plan for the their IT environment		12.5%	37.5%	50.0%			8
Average %	2.3%	11.4%	34.1%	40.9%	10.2%	1.1%	

Table 17 Frequency table question group two

B Delphi Questions 2nd round

The online survey used fixed answers to questions allowing for Likert scale answer, and additional free format feedback. The scales for all questions are:

strongly agree – agree – disagree – strongly disagree – no opinion/no answer

1. When companies and organizations are choosing server virtualization, do they do so to:
 - a. Reduce infrastructure server sprawl by consolidating to fewer physical servers
 - b. Reduce time for server provisioning and deployment of new servers
 - c. Reduce time for systems management of infrastructure and servers
 - d. Reduce energy costs for powering servers, cooling and air conditioning of data centre
 - e. Reduce data centre floor space for fewer physical servers
 - f. Reduce software licensing costs
 - g. Reduce server maintenance costs
 - h. Increase server resiliency and minimizing planned and unplanned downtime
 - i. Increase the speed for deploying new servers

2. When companies and organizations have deployed server virtualization, do they then:
 - a. Monitor and evaluate the resource usage of virtual servers (CPU, memory, disk, network)
 - b. Monitor and evaluate the power consumption of physical servers and other data centre equipment, such as storage
 - c. Test and evaluate the virtualized server resilience, such as high availability or disaster recovery functionality
 - d. Implemented structured work processes and procedures for system management of their virtual server environment
 - e. Measure and calculate possible power savings from their deployed virtualized server environment
 - f. Measure and calculate possible time savings from systems management for the deployed virtualized server environment
 - g. Calculate the Return on Investment for their virtual server environment
 - h. Consider that their virtual server environment was worth the investment
 - i. Consider that their virtual server environment has increased server resilience against planned or unplanned downtime
 - j. Have a valid and up to date disaster recovery plan for the their IT environment
 - k. Have a Green-IT agenda or plan for the their IT environment

C Interview Guides

PROLOG

All interviews should start with an introduction of the participants, and the purpose of the Master Thesis, and that the responses will be handled confidentially, transcribed and analyzed before using in the Thesis.

EPILOG

All interviews should end with thanking the interviewed for participation, and asking if there is anything the participant would like to add.

Ask the participant if it is ok that we mention that we have performed the interview with them in the Thesis, and that any transcriptions will be obfuscated before inclusion.

NOTE

Interview questions needs to be adjusted depending on the participant - is it a Deployer (company/organization using servers) or a Promoter (company/organization selling servers or consultancy).

Speak calm and clear, and ensure that the participant understands each question.

INTERVIEW QUESTIONS

Introduction

We are interested in your views regarding success factors for server virtualization, and especially if and how power usage is related to successful server consolidation with virtualization, but also other factors that you think are important for reducing power consumption of IT equipment.

Questions to server virtualization Promoters (who sell to customers)

Please tell us what you believe are important decision points for companies and organizations, when they are deciding if and how to implement server virtualization.

1. Why do you promote server virtualization?
 - Proportion of virtualized servers compared to usual ones
2. What do customers usually want to virtualize?
 - Server (types?)

- Desktop
 - Data Centre
 - Application
 - Storage
 - Other
3. How do you suggest customers go about server virtualization?
 4. Do you propose server virtualization to customers to reduce power usage?
 - Do you believe there are other areas where there are more power usage reduction to be made?
 - Regarding the size of the server park / data centre
 - Cooling (free cooling, hot spots)
 - UPS (uninterruptible PS)
 - Lighting
 5. Do you believe there are any drawbacks to server virtualization for customers?
 - Hidden, unanticipated costs
 - Hardware
 - Staff
 - System management
 - Software licensing issues
 - Security
 6. What is your view regarding High Availability with server virtualization?
 7. What is your view regarding Disaster Recovery with server virtualization?
 - Less hardware
 - Many VM on one piece of HW
 - Downtime, mirroring
 8. What is your view regarding storage for server virtualization?
 9. What is your view regarding backup and restore for server virtualization?
 - Single data files vs. entire servers (images)
 - Less backup licenses
 10. What are the key success factors for server virtualization?
 11. What are the key pitfalls for server virtualization?
 - Business owner/ management buy-in
 12. In your view, do companies and organizations have other considerations regarding server virtualization?
 13. Do you get feedback; do the results meet their expectations in general?
 14. How do you think organization that deployed virtualization can monitor and measure its benefits (ROI through cost savings, lower energy consumption)?
 - Utilization
 - Forecast capacity growth
 - Measures:
 - Hardware costs
 - Server and storage consolidation
 - Reduced management costs
 - Lower disaster recovery costs
 - Higher system availability (less downtime)

D Interview Transcriptions

Transcription from interview with Promoter A

- 1) Please tell us what you believe are important decision points for companies and organizations, when they are deciding if and how to implement server virtualization.
There is a difference between larger and smaller companies. In the earlier days, larger companies mainly focused on consolidation of physical machines in order to save money. Today smaller companies deploy virtualization to achieve high availability. The amount of saved power (and combined costs) is usually not worth the manpower and investment of virtualization. Therefore the main motivation is high availability, to have the same functions (apps) in several places. About mid-sized companies, power consumption reduction and security starts to become factors when thinking about virtualizations. An issues concerning power consumption is its measurement, since companies usually don't consider IT power consumption for itself, instead they just receive one big bill covering all the costs of a building e.g.. Anyhow, they know about the power consumption issue and future importance.
- 2) Why do you promote server virtualization
Two big issues: high availability part (licensing costs can be saved), and control. It's much easier to control virtual environments than a high amount of physical machines. High availability is crucial today in terms of business activities that rely on applications and data (1 hour break down can cost a lot).
- 3) What do customers usually want to virtualize?
Today it's mainly server virtualization, since it fits to most organizations needs and enough common knowledge about it exists. Other types of virtualization, for instance desktop virtualization can be deployed if the customers have special needs.
- 4) How do you suggest customers go about server virtualization?
They usually think about ease of control and ease of use in terms of changing/ adapting the IT environment.
- 5) Do you propose server virtualization to customers to reduce power usage?
In the public sector, yes. In the private sector it's more important to address high availability, ease of use and control. If companies do already split their IT costs from the rest of the whole power consumption costs they do care about power usage and future actions regarding decreasing power consumption. Therefore power usage can be proposed, however it's not the biggest part of the business in the private sector. Moreover the willingness to invest in environmental purposes only is very low in Sweden (low budgets are available).

- 6) Do you believe there are other areas where there are more power usage reduction to be made (DC) apart from virtualization?
It's very hard as long as the energy consumption (of IT) is not accounted separately (from the whole building). It's an issue if you build a new data center. The motivation is quite low in terms of investing in power usage reduction technology within a working DC. Instead investments are made (or needed more) for new servers for example.
- 7) Do you believe there are any drawbacks to server virtualization for customers?
Expertise (consultants). The hardware costs are almost the same compared to standard machines. Licensing costs can be cheaper (depends on the actual contract), e.g. if based on cores / license.
- 8) What is your view regarding High Availability with server virtualization?
Costs can be saved on licensing since moving of applications within one physical machine is possible. Monitoring and tests are needed for sustainable availability. You can use any application for virtualization.
- 9) What is your view regarding Disaster Recovery with server virtualization?
It's a complex solution and you need to know how to handle it. Do your catastrophe plans. Have a image outside the virtual environment.
- 10) What is your view regarding storage for server virtualization?
Not a lot companies are willing to pay the high price for it today.
- 11) What is your view regarding back-up and restore for server virtualization?
The interviewer couldn't catch any good phrases here due to bad quality of the record.
- 12) What are the key success factors for server virtualization?
High availability. (Satisfy the customers, keep your services online.)
- 13) What are the key pitfalls for server virtualization?
Security concerns are common. If you don't need high availability or think about lower power consumption there is no real need to virtualize servers since hardware prices are dropping. Also the fact that business owner (or management) do not really see the advantages respectively cannot visualize virtual machines instead of several "real" ones leads to pitfalls of virtualization projects or intentions.
- 14) In your view, do companies and organizations have other considerations regarding server virtualization?
Depends on their special needs.
- 15) Do you get feedback, do the results meet their expectations in general (satisfy).
They are satisfied in general. There is no reason why you shouldn't virtualize everything today.

16) How do you think organizations who deploy virtualization can monitor and measure its benefits (ROI, costs savings, lower energy consumption)?

In order to measure you need something to measure against. You can easily compare licenses costs. Power consumption is rather difficult. Manpower could be compared in terms of deploying and maintaining, but on the other hand the needed staff can easily work on other projects and won't be laid-off. Hardware costs can be decreased on the long run, since in the first place investments in new hardware have to be done that normally pay off a couple of month later (newer machines are more efficient, lower service costs).

Transcription from interview with Promoter B


- 1) Please tell us what you believe are important decision points for companies and organizations, when they are deciding if and how to implement server virtualization.

A very important part is to consolidate machines for more efficiency usage compared to physical machines (and the connected investment). Normally a server is used for 5-10 percentage, virtualization allows up to 50, 60 percentage of utilization. So you can put together 5-10, maybe more machines into the same part. The affect will be that small ideal machines will take about 80-85% of the max power consumption in a virtual environment.

In a small environment it might be not so important regarding costs but it becomes more interesting in terms of pollution policy. You can show that you want to make the world a better place to live, even though the economical aspects do not play a big part in your decision. But in a very big environment, if you have hundreds or thousands of servers, virtualization and combination of several machines will give you a lot of payback or less (expensive) power (consumption). Also if you have 2, 3 virtual machines inside a physical machine you can increase safeness, you can move application / function from machine to another, for maintenance or for low balancing etc. This way you lower the costs, since you use less energy and on the other hand make the server environment safer and reliable.


- 2) Why do you promote server virtualization

If you set up several applications in one machine you have the ability to have much more safer machine. You also spread the costs of several applications. For single machines (physical servers with only one app) you cannot have a high safeness (not always a redundant interface available to the network for example). Regarding power consumption, not all the application will use its max consumption at the same time. By putting several applications together that don't have their maxes at the same time you get a quicker system by placing them together. Power usage of CPU. Comparing usage of two single machines with a combination of two applications placed in one virtual environment.



The best effect is reached when combining 10 up to 50 apps in one virtual environment.

The delta between max and min load is decreasing, and constant power consumption can be estimated. Therefore machines can be utilized nearly up to their max capacity. Individual peaks do not longer matter. Overall demands of the systems are better predictable (regarding the needed sizes of machines).



The same concept applies to storage virtualization.

3) What do customers usually want to virtualize?

Don't think about data (storage) in the first place. High availability purposes of servers that are in touch with central data. If one server crashes, application and data is still available, and business processes can be kept running. This is important in high performance, banking or security environments. Also there is the possibility to move and start up applications from one to another machine in case of problems. Consolidation of physical machines is one reason for server virtualization (regarding power and space issues). Also test environments are fast to set up in a virtualized way.

4) How do you suggest customers go about server virtualization?

In the earlier days it was most important is to save storage. Now when speaking about server virtualization it's about saving costs regarding hardware. You can replace 5-10 old servers through one virtual server environment. Also the power consumption of the newer server is more efficient. Capacity growing and replacement of broken hardware are arguments for buying new (virtual) servers. You can adjust the hardware needs of an application continuously through using virtualization. More flexibility is also given, e.g. through setting up a test environment and a new version of a running application on the same machine which hosts the original application. Switching between the two versions is easy and setting up or deleting one, as well.

5) Do you propose server virtualization to customers to reduce power usage?

Yes, it's a good argument. Calculations can be made that show the difference between power usage of physical single usage servers and servers consisting of several virtual systems regarding power consumption (because none of them are running under full capacity, you can simply reduce the amount of used hardware). Moreover the cooling aspect are also lower when less power is consumed. Therefore a payback of the investment of virtualization is often reached after 1 – 1,5 years.

- 6) Do you believe there are other areas where there are more power usage reduction to be made (DC) apart from virtualization?

Use produced heat to warm up you buildings during the winter times. Data centers will be smaller through virtualization and consume less power in general.

- 7) Do you believe there are any drawbacks to server virtualization for customers?

Specific application that you want to isolate from virtualization are existing (firewalls, back-up systems). A separated back-up system might be easier to use in case of data loss, disaster recovery. Licensing costs can be an issue depending on the contract (different sorts of licensing) based on the software company.

- 8) What is your view regarding High Availability with server virtualization?

Several sites (instances) are available of individual applications and enable switchovers in case of problems. By switching (off and on) an application on different (active or passive) sites will lead to no extra hardware costs.

- 9) What is your view regarding Disaster Recovery with server virtualization?

It's an extra bonus of SV. Machines in several DCs enable load balancing in case a machine breaks down or is overloaded.

- 10) What is your view regarding storage for server virtualization?

Most system have a central storage. SAN is quite common today.

- 11) What is your view regarding back-up and restore for server virtualization?

IBM TSM is used a lot. The old fashioned way is that you have some servers, one TSM server and tape libraries. The servers consist of little back-up clients that send information over the network to the TSM physical machine that has a database on it. Therefore the TSM knows which data is on the servers and can create back-up files on tape. In the first place a full back-up will be performed, after that incremental back-ups follow continuously where data changes will be recognized. In the case of restoring a certain file or directory can be recreated (related to a timestamp). The same procedure applies if the physical machines will be put in a virtual environment. The back-up server (TSM) is still kept as a physical machine to enable easier restore procedures.

- 12) What are the key success factors for server virtualization?

First of all a study will be made to identify the needs/demands of the customer. What kind of servers are needed in a virtual environment. Primarily the easiest and best (most efficient) applications will be virtualized. There are types of physical servers that are impossible to virtualize (e.g. print servers and adapters?). Take approximately half of the system to virtualize it first, evaluate it and decide about further steps of virtualization. Focus

on the main purpose first before considering other benefits of virtualization, e.g. reducing the amount of (physical) hardware versus high availability.

The easiest systems to virtualize are test- and development servers, more complicated are huge data bases and unique environments.

13) What are the key pitfalls for server virtualization?

Too less expertise if developed in-house. Access to data especially for huge databases can be crucial if wrong implemented. VM ware specific problems; everything has to go through the VM ware environment.

14) In your view, do companies and organizations have other considerations regarding server virtualization?

Sometimes they think everything will be put in one/two machines. Also security issues occurs in case virtual systems of several organizations should be put on the same physical machine..

15) Do you get feedback, do the results meet their expectations in general (satisfy).

Average load might be quite low in general if several systems are combined in one virtual environment. Therefore a 100% workload is usually not reached and the overall speeds of all systems are increased through virtualization.

16) How do you think organizations who deploy virtualization can monitor and measure its benefits (ROI, costs savings, lower energy consumption)?

Power consumption will be less. If costs saving calculation are correct a payback can be reached after 1-1,5 years. If you are not able to save money after this time the previous calculations must be wrong. The ability of offloading systems (if redundant) in virtual environments can also save money, since in case less capacity is needed some systems can be shut down or offloaded. On the other hand more knowledge (expertise) is needed for this (might lead to additional costs). Managing more physical machines can demand more time (monitoring) and effort (in case of hardware reparation).

Transcription from interview with Promoter C

- 1) Please tell us what you believe are important decision points for companies and organizations, when they are deciding if and how to implement server virtualization.
One of the more important decision points I think is that to gather with reducing cost for hardware and also about the energy efficiency. If you do it right some resilience alone with implementation, you get robustness against the failure because you can build a cluster or you are more required to do it if you are running production environment. So that's, I think, one of the drivers, because you get that almost, well, for free, of course normally licenses for the implementation we cost but it's a minors charge and we get extra robustness from the solution.
- 2) Why do you promote server virtualization
I promote it because of the same reason, normally you can get savings in energy, if you do it and for mmm...
- 3) What do customers usually want to virtualize?
Servers mostly. Many have heard about the storage virtualization, but it's not as common as SV. After the SV is desktop virtualization. So it's first SV, and then Desktop virtualization and the last one is Storage virtualization.
- 4) How do you suggest customers go about server virtualization?
Saving money. You get some additional flexibility, but it will end up in cost saving. What will it cost me, how can I make this paid? Those customers we met are maybe in the situation that they have a lot of old hardware which have to renew them that often quite expensive or they have to buy new servers. They know there is a technology with virtualization so they are curious. Mostly we try to get them analysis the current environment and when they do that they realized compare with the old one, the new technology is a good solution. But on the anther hand, there is new issue come up about the robustness against the failure, because all the functions are in one sever, and you have to think about that problem. Or you just do the half of job of SV, if you want to want to have the benefit you have to consider that.
- 5) Do you propose server virtualization to customers to reduce power usage?
It's a consequence of SV and it's always there. It's quite a lot. If we have an environment with approximately 40 or 50 servers, it will save 100 plus thousand Swedish Kroner a year. So you have an investment about half a million, we speak Intel platform mostly Windows or Unix, you can probably virtualized the environment with the solution for 500,000/ 600,000 SEK, you have a payback for power for only in five years. You don't need to buy the servers, you have a monthly charge for your hardware, and you will see you can get better things and still saving costs and energy. The major issue about energy saving is that in most cases IT organization didn't measure on power consumption. It's another budget, the propriety when you run the whole building is a part of building budget. So even the IT manager can pointed on saving, the saving are getting somebody else budget, so he can't get any use out of it. That's the drawback. The other thing is if

you looking on the major maybe outsourcing provider, or IT organization who sell server services to customers, they themselves buy facilities which are computer rooms with cooling and power, and that computer room is often rented by floor space/ footprint, and not based on power consumption. But the power consumption is built into that footprint cost so they can't actually benefit from SV from energy perspective. So this is normal situation that IT is not involved in power consumption.

- 6) Do you believe there are other areas where there are more power usage reduction to be made (DC) apart from virtualization?

There is an over dimension of most components in a DC. Power up or down server according to their peaks. The same applies to the storage. Having discs running 24/7 costs a lot of power and you can save a lot if you are able to spin down the discs.

- 7) Do you believe there are any drawbacks to server virtualization for customers?

The one thing is that it will require the IT server provider to be more conscious about the power you dedicate to give to different servers or virtual servers. In the old structures, when one server one function, the IT provider could say "well, the server is low, you need to buy a new one", then the customers will buy a new one.

If you virtualize you more or less put that responsibility of low balancing on the IT department instead of application owner. They have to be sure that everybody get what they need or pay what they need and they have a projection about how computer power will increase over the time. So that's a new challenge. That will probably affect datacenters because most smaller or medium size companies have enormously oversize servers so it they put that into one server, it's often not a problem. As far as server consolidation is concerned,

If you also have desktop on virtualization, then you have a big computer power consumer introduced into the computer center which is most IT provider don't have yet. It's also a new technology to put desktop power into the server instead of everybody's desktop.

Training fee (Knowledge and expertise) is not a problem because if it's a new thing you have to be educated. But that's go for everything.

Other thing is that because of the server not only virtualizes memory, CPU, it also often virtualizes the network. So you have a virtual switch that's often handled by separate department in the large IT operations. So they have network group deal with the network configuration and now that part is taken away from the network guys and put into server virtualization environment, who's the responsibility is configure the right way. That's an issue.

- 8) What is your view regarding High Availability with server virtualization?

If you do SV, you should consider a High Availability solution. It's not a option not to select it. Because now you put a lot of functions in one box, if you don't think of it, you don't do your job well.

- 9) What is your view regarding Disaster Recovery with server virtualization?

It's more or less the same answer as the previous one. You have to have a disaster recovery plan. Disaster Recovery and robustness or the ability to cope with the hardware and

software problems is inherently connected with SV. You have to deal with it, because with one server one function; If you do a service on that function you stop one thing / one box, one set of users of that function are affected. But if you put 40-60 services in one box, how should you do service of the box itself, because it will break, everything does. So you have to have some way or plan your servers in a virtualized environment and the only way to solve it is to build it in from the beginning. So normally you build it; I guess you worked or looked at different technologies, moving services online or between servers. You have to have that solution from the start, otherwise you can't deliver a quality of service because eventually something will break and then it will affect every systems, everything will be down. And that's normally not acceptable. So you have to think about the many aspect of Disaster Recovery.

10) What is your view regarding storage for server virtualization?

You need to have a shared storage. You should consider on dual storage environment (and that's for the redundancy part) so you can always be sure you have a copy online on disk in some locations separated from the normal production environment. So you should consider having always two computer rooms. They don't need to be big ones because it's a lot less hardware in the virtualized solution so you can shrink them but you nevertheless need to have two separate places so you need two storage systems, as you need two physical servers. Another issue is that you almost never consider the IO volume. You only consider the storage volume; you think in gigabytes instead of in IOs per second when you select the storage environment. But one of the major things that will go wrong is that storage system is under dimensioned compared with server environment. Because we have a lot of individual servers; you have one or several discs in each one of them, the discs are underutilized with respect of storage use but they have the ability to perform high IO. It's a fast and quick storage in the stand-alone server environment. When you then consolidate that, and put all those serves, which we are running in physical machine into one central storage system, you usually don't consider that the reduction of (numbers of) discs leads also to reduction of the IO capability of the storage system. So we have to think about dimensioning the storage system to handle the IO that all those servers did un-virtualized because it will be the same IEO after the virtualization as well.

11) What is your view regarding back-up and restore for server virtualization?

That's one of the positive side effects, because you virtualized, you often could hold together your server as virtual image or file and this is a lot easier to operate on and doing backups with. Because normally if you have failure, you have the need to restore the server and application itself and then you have to have that ability to restore the data. That's two parts in the stand alone server, that's an awkward procedure normally to do an re-install because there are in the Intel environment no good ways to do an image restore because you have a hardware dependency of the OS to the hardware; so in the virtualized environment you have separated the servers demand on the HW by putting a virtualization layer, so you don't have to bother about drivers and problems like that. So doing the image backup or restore is a big advantage of SV.

The data backup themselves they don't differ more from the normal environment; you have to do backup because even you if you built a cluster virtual environment, it will be not protect you against corrupted data you always have a need to do restore of corrupted data bases and stuff like that. So for that you need a backup system as you always did.

12) What are the key success factors for server virtualization?

You have to analysis the load before you do the visualization. I think you should also build a redundant solution from the start. If you don't do it you should be aware of the side effect of not having it. Most virtualization platforms are quite easy to operate. There are some, which are very hard to operate, but you need not select those as your platform, VMware and some of the XEN based platforms are very easy to run and operate. So the operation part of the environment is not that difficult it's more that you should know what you put in them. So the service you put in the virtualized environment you should know the characteristics, their behaviour or what strain do they put on the new virtualized solution, in respect of load characteristics and IOs, memory consumption and stuff like that. They are always machines which are not suitable for virtualized environment; don't need to be a big number of them, but number of services in more or less every environment is not suitable. If you are not aware of that you will probably run into some problems.

(How can you figure out if it suitable or not?)

You should have experiences. You should analyze the behaviour of them before you virtualized them.

(But you cannot generally say the application is not suitable?)

Normally database applications are fine, most administered application are fine, files and email's services are fine. Compute service are not good of obvious reasons. There are some applications, which are very aggressive memory consumer, which are not good for SV, they are servers, which have hardware dependency, and they need special adapters. If you want to attach tape systems, or if you have a switchboard, which have special card you put in it, you cannot virtualize that. Number crunching high performance compute, bad choice. There are some general guidelines. But how do you know in this gray area if the services are good or not good, you can ask the application vendor, and you can get probably a random answer. Most cases they don't know or they think they are uncomfortable with virtualized environment, so they say no. But they should say yes because it will be a benefit even for them. You get a lot of different answers from the vendors about supporting for the specific application in the virtualized environment. So eventually it comes back to you who operate the environment to check and test if that service is working and have a backup plan if it seems to be working and overtime does not. So you have a recovery method of that putting it back in a physical environment. Or run it and you have a spare machine which would be used for those applications where the vendor are not allowed that you do put it into the case. For correction, if that service virtualized, you put it back on hardware, you put it the case and that you put it back in the virtualized environment. Because most often it is not related to the SV, it s something else that is wrong.

13) What are the key pitfalls for server virtualization?

You need to know what components you are able to selected in that environment, especially with the storage part (regarding movement of services).

(The most obvious one is under dimensioned memory. Because if there is something you need for the virtualized platform running properly it is the memory, and the other is the IO capability. If you are not aware of the IO capability or you think it is so easy to put in

new services and you just keep adding them, and don't see how that will affect the attached storage, as you know there is a shared storage in the virtual environment.

If you want to move applications you are dependent on several resources: memory, CPU, disk. So both services must be able to access the same disk, they need to be able to communicate memory with each other (transferring data is related to each of their memory). During the transaction other applications should not be able to get access to two disks, only to one. Stop incoming transaction during that process and start it again if everything is completely copied.

- 14) In your view, do companies and organizations have other considerations regarding server virtualization?

It's a very big different between customers, many of them are very clear about what they want and familiar with technology. Most do not know about the capacity planning, because it's a common problem. Some customer who wants to go with SV, they often need to know more about their own systems but they don't know enough about this. So they don't know the capacity and of course it's a new thing so they are a little afraid of it, having doubts...

- 15) Do you get feedback, do the results meet their expectations in general (satisfy).

The project we handled are mostly successful. Because it's easy you use standard components and you buy VMware or XEN but you don't do the performance calculations. once it starts running and I working and its easy so you add services then eventually it stops running or it's slow and you don't know why. So it depends if you know what are you doing, it's easy.

- 16) How do you think organizations who deploy virtualization can monitor and measure its benefits (ROI, costs savings, lower energy consumption)?

Yes, there are several. Most of them are delivered from a vendor, which are delivering some component of the solution. So there are a lot third party products (VMware power cabs, HP...), which measure the capacity or capacity forecast, or something like that.

About power consumption you can do the paper exercise but the really interesting part is to measure it in the actual computer room and see a year ahead if you reached your target and that's more or less up to you as a customers to do it. We usually point out you should do it, because most UPSs supply the current rate they are outputting into the DC or the Watts. So you can read that of a small screen on the UPS, and if you do that you can start virtualizing and that process often take longer than you think because you have to do the normal operations part of the IT environment as well. So the virtualization process is not as fast as you thought in most cases, as it would be. So it takes time to go through server-by-server, shutting them down and it will take some time after the last shutdown that you will get the benefits. You often underestimate the time it takes, both as a supplier and as a vendor.

Transcription from interview with Promoter D

- 1) Please tell us what you believe are important decision points for companies and organizations, when they are deciding if and how to implement server virtualization.

It has changed over time. First of all, what we are focusing on the last year is Intel based SV. The customer want to have optimized and have more efficient test environment for 2001- 2004, that was the first scenario that made companies implement SV. It is so easy to handle test environments

Then the technology developed the main reason has become the cost reduction, server hardware, licenses etc. In our opinion the virtual machine is 40% less costs involved than in physical machines.

There are companies who are not investing in the technology because of cost reduction; they even decided to go with a more cost intense solution because of the disaster recovery features in the virtualization technologies.

60-70% who invest do it because of cost reduction, 20% are driven by DR/high availability reasons and only 10% are driven by getting more efficient test and develop environments.

- 2) Why do you promote server virtualization

There are three main reasons:

- *70 %Cost reduction*
- *20% Disaster recovery / high availability*
- *10% Get more efficient tested or developmental environment*

- 3) What do customers usually want to virtualize?

The main part is the sever virtualization, and for the desktop there is no one so far have really implement because of the technology is not so developed. There are 90 per cent of companies who want to implement the server virtualization and 10 per cent want to implement the storage virtualization. The Desktop is maybe coming next year. The average size of the sever virtualization is about 100 servers or more. Typical customers are big banks, insurance companies, government institutions and large companies.

- 4) How do you suggest customers go about server virtualization?

Cost reduction is the first reason and the High availability/DR in some cases and the third one is the faster development in the tested development environment. There are three different reasons so there are three different solutions.

- 5) Do you propose server virtualization to customers to reduce power usage?

The power reduction is not a main reason (Very very small part of them). The cost reduction is still the main reason.

- 6) Do you believe there are other areas where there are more power usage reduction to be

made (DC) apart from virtualization?

It is a big area of development, the server virtualization is one part. And so is the lower power CPUs used there is another area which take down the power consumption in the datacenter. But the power consumption is not a problem. There are some special cases, where the companies cannot get more power into their data center the power companies said that we cannot give you more power so the companies have to find some new solution which make their power going down and there is one case like that. People want to be green, but the Green IT is not something you invest in.

- 7) Do you believe there are any drawbacks to server virtualization for customers?
There are a lot of hidden (or obviously) cost:

First of all, for the new technology you need to get people trained and you have to invest in a new platform both on software and knowledge and hardware. Especially knowledge is a factor since new systems require new knowledge from other contacts to other conditions as you maybe long term contract with windows consultants. Then the storage costs much, because if the companies want to implement server virtualization they need advanced storage technology, which has high cost.

- 8) What is your view regarding High Availability with server virtualization?
What you need to do is often to 2 separate issues. You could reduce your cost a lot with server virtualization, and if you also want HA you can get it with virtualization much more cost effective you could before but it is still an investment. and if you go in with the high availability you will be end up with more cost than if you only want to go with cost reduction. Because HA means a lot of unused hardware and software that you have spare somewhere.

- 9) What is your view regarding Disaster Recovery with server virtualization?
SV gives you Disaster Recovery possibilities for large volume of servers. Before you just choose couple of servers that you have DR plans for which is the most business critical servers you have in your organization you mirrored somewhere. Being SV you can get DR with large volume of IT infrastructure, that's the big difference.

- 10) What is your view regarding storage for server virtualization?
You need to use it to get the storage cost down. Don't overbook it.

- 11) What is your view regarding back-up and restore for server virtualization?
That is a tricky technology question today because there is no good solution, which can cover both virtualized and physical environment in the same time. You need to choose a good one for each of them.

(Do you usually to promote to save or backup only single file or the entire servers?)

It depends. If you are going to the DR the image size is quite efficient. But if you want to use your own backup tools or agency to restore single files that is very costly because the image will take up a lot of area on the storage side, so you need to combine with storage virtualization to be more efficient. We have a lot customers said which solution/tools they should use, today they have a mix environment, which combine the physical and virtual-

ized machines. There is no good single one solution. It depends on how much cost they recover and how faster and how easy he want to save the single files.

12) What are the key success factors for server virtualization?

Do it quick. Time is the most important success factor. If they are not fast enough migrating the planned amount of servers into the virtual environment their ROI calculations turn out wrong and pay back is achieved later. Since they invested 5 mio SEK and still run the majority of their physical servers, which costs a lot of money.

13) What are the key pitfalls for server virtualization?

Its click and install solutions. So there are no technology problems, maybe five year ago there are some choosing the right hard- software, but today you just install it.

14) In your view, do companies and organizations have other considerations regarding server virtualization?

One problem is very common is that their software doesn't support SV. There are two perspective, one is the software doesn't virtualization which means we only support software it in the physical environment the other is the support the SV from the license perspective.

Licensing from third party software vendors is an issue.

15) Do you get feedback, do the results meet their expectations in general (satisfy).

Often it is positive according to the technology (easy). It's more negative regarding the ROI because of the cost of storage, HW, SW, knowledge what they didn't realized in the first place. They also haven't realized that they need to deploy very fast to make the ROI be correct.

16) How do you think organizations who deploy virtualization can monitor and measure its benefits (ROI, costs savings, lower energy consumption)?

We look into the customer's costs today and then develop a ROI calculation based on their existing environment.

In the Intel space it's very hard to measure if a solution is cost affective or not, in the mainframe environment you know how much a transaction costs. In Intel you don't, which makes it quite hard.

Transcription from interview with Promoter E

- 1) Please tell us what you believe are important decision points for companies and organizations, when they are deciding if and how to implement server virtualization.
The most important fact is costs. Saving money in DC through new licensing models, how you implement and run it (efficiency). Regarding energy efficiency and virtualization you have to take into account that usual servers run not on full capacity (utilization). Virtualization is driven by less hardware (less costs) and also lowering operating costs, since lots of these costs are connected to power, cooling and space issues.
- 2) Why do you promote server virtualization
See question 1
- 3) What do customers usually want to virtualize?
Majority wants SV. VMware is the leading "system". After that storage virtualization is next. Desktop virtualization is harder to gain interests, it's not as important as server and storage.
- 4) How do you suggest customers go about server virtualization?
If you would ask me a year ago, a lot would focus on flexibility, rapid deployment and easier management (image management, same configs and patches on all servers). Nowadays, IT departments are forced to do more within less people and less time and also costs have become much more important. Hardware costs (traditional capital expense) and operating costs are lower in virtualized environments. Reschedule workloads during operations instead of having a service team working weekends and nights on it can save costs.
- 5) Do you propose server virtualization to customers to reduce power usage?
Yes. That has changed the last year and month regarding the attitude of our customers, because before that green IT and energy savings were not a topic. Today customers want to show that they are green and have an environment conscious due to the fact that power and cooling is becoming a pretty big part of IT operations. And also I work with customers who cannot deploy new applications because they cannot get enough cooling into the server room.
- 6) Do you believe there are other areas where there are more power usage reduction to be made (DC) apart from virtualization?
One thing is to use a server more efficiently. Higher utilization through server virtualization. The other part is looking at your infrastructure and your server and saying I don't mind if this transaction takes 15% longer. You can actually say that this server cannot consume more than x amount of WATTS through different kinds of technologies (tools). You can set power caps within blade centers. Shutting down servers during non-peak hours. The other option is to look into cooling and cooling technologies. If you look at different ways of operations you have your classic server, each with its own power supply. AND if you move that into a blade server the energy consumption per computer goes

down. Because you have fewer power supplies that can be used more efficient. Also you have more efficient cooling instead of lots of small vents, which have inefficient airflows, so you can engineer better cooling basically. You can use other cooling options, e.g. water-cooling in the rack, which is more efficient than the air-conditioned cooling systems. In IBM Stockholm, in one of their biggest cooling rooms (2000-2500 servers) they have heat exchangers that enable the heating of the office building during winter time (1200 employees). And also we are selling heat back to the utility company. Of course not everyone can do that. You need big rooms and the abilities (technologies) to sell heat.

- 7) Do you believe there are any drawbacks to server virtualization for customers?
If you don't have a structured process, you don't have service management in place for deployment and the commissioning for the server. Deploying servers is easier than taking them out. Why? You have to figure out who is/ is not using the server, what happens to the data, licenses what warranties do we need... So if you don't have control over your processes and deployment and then make the step towards server virtualization where it is very easy to create new servers it's gonna be messy. Since you don't have only the physical servers, you can have several virtual ones in one box (several OS...).
- 8) What is your view regarding High Availability with server virtualization?
It starts with business availability. What is it that needs to be done?! What kind of business process are we actually supporting with this piece of technology?! What kind of availability level and service level agreement do we have for that business process. And then you have to rank all your processes and supply them with the right amount of hardware regarding their priority.
- 9) What is your view regarding Disaster Recovery with server virtualization?
It can help you if you set it up right. Once more you have to figure out how much downtime you can afford and linked to that you chose the right technology option.
- 10) What is your view regarding storage for server virtualization?
Planning needs to be done. You need to make sure that storage is available on the different systems in your virtual environment.
- 11) What is your view regarding back-up and restore for server virtualization?
The information that changes within your application when you run it has to be backed up, using same procedures as always basically. Most virtual technologies come with functionalities that help you to restore the OS (image management VMware).
- 12) What are the key success factors for server virtualization?
Planning preparation and processes. Processes in terms of handling things, because otherwise you might end up with more of a mess you started with. It's quite easy to achieve physical consolidation but there is always the thread that your virtual environment is growing without having control over it.
- 13) What are the key pitfalls for server virtualization?
Again planning. If you miss checking the workload of the servers and you deploy 15 of your heaviest loaded machines (or that peak the exact same time) it is not going to work

out in one box. Issues are also about the software licensing. Software vendors of the world (such as IBM) are not 100% comfortable with SW-licensing in virtual environment. You need to discuss what are the regulations, what are the licensing models, how do they work. Do you pay per core or socket or OS?

- 14) In your view, do companies and organizations have other considerations regarding server virtualization?

...

- 15) Do you get feedback, do the results meet their expectations in general (satisfy).
In general yes. If we go back to the cost case: the customers build the cost case on what they can see and touch, regarding how much money they have to invest in the first place (new servers, new virtualization sw, how much energy it will consume...). But they miss to assume the biggest gains are more expressed in soft values such as speed of deployment, flexibility / the possibility of shift workload or resources from one virtual to another rapidly. It's hard to put a value on these ones when you start out.

- 16) How do you think organizations who deploy virtualization can monitor and measure its benefits (ROI, costs savings, lower energy consumption)?
Hardware costs are easy to calculate (less, finance department...). Regarding monitoring we have possibilities to look at the energy usage and can refer them to costs. You can calculate costs regarding different servers or even services and can narrow them down to an individual (user or projects). Management costs reduction are hard (but of course theoretical possible) to figure out since you have to know them beforehand and usually customers don't know them. They usually know about licensing costs, power and cooling, maintenance costs. Also there should be education costs included in the managing costs, which depend on the customer. However it is possible to put figures on it afterwards (e.g. we fired/ moved two guys since the workload is not that high anymore in our department).

Transcription from interview with Deployers F

- 1) Please tell us what you believe are important decision points, when you decide if and how to implement server virtualization & 2) Why do you consider server virtualization?
I do want to minimize the physical hardware resources, so I want to use the current hardware resources as good as possible to lower hardware maintenance costs, to lower licensing costs and also minimizing space in the data center, minimizing POWER consumption, cooling. Also it's easier to manage less hardware especially for upgrading VMware. Also we need some competence for the people who are running the virtual environment. We don't want to virtualize 100% (but as much as possible) of our servers, for example all of our high I/O machines. Also we use mixes of virtualization, so we use virtualized CPU, virtualized network and dedicated fiber adapters if there is a server, which does a lot of I/O. But the big decision is to minimize hardware and lower costs and get a somehow easier platform to run.
- 3) What do you consider are appropriate to virtualize?
We have running thin clients projects and we will have fewer PC in the near future. Desktop virtualization might be an option for the standard users, but not for everyone (administrator, technician). My field is the virtualization of the POWER machines. For windows environments we are using VMware. Linux is almost 100% non-virtualized since the Linux guys don't want it to run on VMware. I hope I will run it virtualized [like] AIX (IBM Unix) in the future.
- 4) How do you do the server virtualization (steps involved, design, implementation, test...)?
We are working with POWER virtualization since 2005, so it's a proven technique. It's well described but first of all you need to get some expertise. Then I decide how to use the virtualization and also who is going to maintain the virtualized environment. You need to make some kind of baseline. Designing is a cooperation with IBM, regarding performance. But the servers are pretty much pre-tuned when they arrive (OS images are already there, you just have to install the engines). Installation and tests are followed up. Then we take the result of it and make a delivery, which consists of virtualization, the OS images for hosts to install. So we got an AIX image, we take away lots of things and add other things, so it becomes an [redacted] image. And after that image we add 7-10 applications, which are pre-packaged. Now you can choose which ones to install and you get your server.
- 5) Did you consider server virtualization do reduce power usage of the servers?
My platform is an expensive one. So I have a focus on reducing costs. It is power, cooling, floor space, cabling, licensing and maintaining. I try to figure out all those things to get the biggest picture possible and try to get the right people to look at different costs from different angles.

- 6) Do you believe there are other areas where there are more power usage reduction to be made (DC)?

The new data centers, which we build last year, they use lots of nice techniques. They can cool on different spots. We put in half the power in each rack so we just populate the racks to the half. Reusing the heat (created through cooling) is also planned.

- 7) Do you believe there are any drawbacks to server virtualization for customers?

Depends on the kind of virtualization. For POWER virtualization you need extra hardware (costs) and you also need the software maintenance for software virtualization. But you need to know how many licenses you need (how many software is needed). Imo there are no security issues, we are counting on the vendors that there is isolation between different partitions.

- 8) What is your view regarding High Availability with server virtualization?

We have been using virtualization on the POWER [architecture] since POWER4. We created partitions with whole CPUs, memory and dedicated network storage. So we created a partition of resources in bigger machines. And those resources we put on AIX and put it on some clustering software for high availability. We are using 2 nodes and running resources on node 1 which could be moved to node 2. We have also in parallel running other resources on node 2 so they can take over for each other. We got also less downtime through redundancy of virtual engines. We can move a partition from a running POWER6 to another POWER6 hardware.

- 9) What is your view regarding Disaster Recovery with server virtualization? & 11) What is your view regarding backup and restore for server virtualization?

If you are running AIX then we are dealing with the bootable system images which are saved in a repository server. Disaster recovery installation will be done though these images. If there is less hardware of course it will minimize restore time since there are fewer drives to restore. One the other hand it could take quite long to restore if you lose a whole server with different virtual environments on it. So you need to make better planning for your disaster recovery in a virtual environment in case the complete hardware dies. We are using external storage, so it doesn't affect us that much if a whole machine crashes. We also use mirroring between our subsystems.

Regarding backup and restore issue: In the POWER virtualization it's not like in the VMware where you can just save the whole disk files. On the POWER every VM every partition installs a backup system (TSM). So we are running incremental backups, so the first one is a complete one and then you just continue to do incremental backups.

- 10) What is your view regarding storage for server virtualization?

Regarding Storage, we are using virtualization for storage. There are some big vendors providing it, e.g. IBM or HP. Most of them are running on some servers and there is lots of storage behind those servers. The storage may vary between vendors and all the hosts are accessing the storage through these 1,2,3,4,5,6..., or whatever front-end or front servers. So behind the front you can use different storage drivers to each storage subsystem. So it's easy to migrate from one disk vendor to another one behind the storage engines.

12) What are the key success factors for server virtualization?

When you start to virtualize things you don't want to get into problems so you must figure out how your environment actually can utilize virtualization. Customers must be as happy as before. Keep it quite simple and skip the fancy stuff, figure out what you actually [are] trying to gain by it and if everything is running nice announce it somewhere (webpage). Main thing is customer satisfaction. Also the people that maintain and run the virtual environment must be aware of the routines of doing it and it must be a proven technology (small overhead).

13) What are the key pitfalls for server virtualization?

Bad maintenance management, know about your disks, partitions, memories, processors. Document it!

14) In your view, do promoters of SV have other considerations regarding server virtualization?

No.

15) Are you staying in touch with the promoters after the deployment process of the virtualization? Do you give them feedback? And are the results meeting your expectations in general?

Yes, still in touch and give feedback (perform meetings) about hardware software issues.

16) How do you monitor and measure its (SV) benefits in your organization? (ROI through cost savings, lower energy consumption)?

Power is very easy to measure. You see each month and each day how much power you consumed. It's also easy to calculate on hardware purchase and also what a data hall will cost. Licensing costs are more difficult to figure out (need to keep track of them, buy new ones). It's getting trickier to keep track of all these licensing costs. And you also have the management / staff costs which are hard to calculate. I don't have a total overview of all costs. But I know the floor space costs, power consumption, hardware, maintenance costs We are using some different software when we need to buy more hw and also to see if things are not in balance and we should move them from one place to another. We are looking for a (one) more complete software solution for monitoring performance capacity between platforms. You really need to know about the utilization, so when to stop putting in things and create another one. Forecasting costs is also crucial since you need to know in advance if you need more hardware (money needs to be requested).

Transcription from interview with Deployers G

- 1) Please tell us what you believe are important decision points, when you decide if and how to implement server virtualization & 2) Why do you consider server virtualization?

A good pre-study is needed about the purpose of the technology. This applies to any new IT technology, not only SV. The study should cover facts such as how, where, what you will need it for. Today, the big vendors are promoting this fashion world. Every year there is something new. Virtualization was the fashion topic one and two years ago, now it is Green IT. But it is nothing new, virtualization has been around since the 60s.

In our case the pre-study covers topics like consolidation, power savings, better utilization, what is your goal to apply it. Is it to reduce hardware costs, to reduce resources in terms of system administrators.

We started out by testing it first. Building a small server farm and adjust it for minor applications, e.g. test servers. Which saved us a lot of money and time to deploy new test servers.

- 3) What do you consider are appropriate to virtualize?

We haven't planned anything more than the servers yet. We are starting to have issues regarding desktop machines, could be the next target for virtualization. We have a lot of outside contractors, doing a lot of development for our applications. They connect through VPN and work from home, but also need a connected workstation here (1to1). Several hundred are in use and we get into space, cooling troubles. Pre-studies are currently on.

- 4) How do you do the server virtualization (steps involved, design, implementation, test...)?

See before and below, first tests... then others, review the characteristics of the server is very important. What applications are running on it?! There are a lot of applications that are not supported to run in a virtual environment.

- 5) Did you consider server virtualization do reduce power usage of the servers?

Yes we did very much. That was actually our highest for moving towards virtualization. Over the last year SE has grown very much and that reflexes on IT as well, e.g. the number of hosted servers, storage and backup solutions. About 2 years ago we reached the limit about energy consumption in our DC, regarding power supply and cooling. Consolidating hardware through virtualization was one big alternative to rebuild/renew the whole DC. The virtualization process was step by step since new virtual servers needed to set up first, before any movement of old servers towards them was possible, with respect to the low available energy in the DC.

- 6) Do you believe there are other areas where there are more power usage reduction to be made (DC)?

Yes, utilization is very low on servers (10-15%, peaks up to 80%). Therefore INTEL and AMD are providing new processor technologies (dynamic power cabling) which features

that you can set your proportional CPU's power consumption. If the application reaches a peak it can request more power from the CPU automatically if needed.

You can refresh your current hardware with new hardware within a shorter life cycle and try to follow the newest solutions, e.g. change from rack to blade servers reducing the footprint within a DC. Also building your DC in a correct way to enable more efficient cooling. Place your racks back to back to get more specific cooling spots. Also the temperature shouldn't be as cool as possible, studies showed that 25 degrees are perfectly all right when having the right air flow.

- 7) Do you believe there are any drawbacks to server virtualization for customers?
There are a lot of different technologies, e.g. VMware or open source ones. Going with VMware is expensive. It is a licensing product and at the same time you save hardware costs through virtualization you need to pay a lot for the VMware solutions. Also external and internal knowhow is demanded. Not all applications are suitable for virtualization, e.g. the ones, which need a lot of memory and CPU, are demanding a own physical machine. So virtualization doesn't make sense there and will cost you more in the end. Pre-studies are needed.
- 8) What is your view regarding High Availability with server virtualization?
We are using it for maintenance purposes, when we have server downtimes due to that. We create resource pools by connecting several (10) servers with each other. And the virtual 50 machines can use that, by still having enough capacity. In a scheduled downtime for maintenance we actually can run firmware Updating on 50% of that resource pool, bring down 5 of them, upgrade, take them up again and switch over to the rest without getting the services offline and having an impact on the production environment.
- 9) What is your view regarding Disaster Recovery with server virtualization? & 11) What is your view regarding backup and restore for server virtualization?
We have a lot of systems, so we don't have one default set up for DR that applies for everything. It's very system dependant. Systems belong to different owners within the company and their owner decides for themselves what kind of DR requirements they need. 80% of systems only need a normal tape backup of the servers and databases and so on. Other systems require more investment, like having the backup files on two different locations. We don't use that much virtualization technology for that, it's more physically. We are also using SAN technology from HP, regarding high available storage.
- 10) What is your view regarding storage for server virtualization?
SAN storage is used.
- 12) What are the key success factors for server virtualization?
Do a proper pre-study and sort out your intention about virtualization. What's the purpose of it and so on. And if you have done that choose the V-technology (soft- and hardware) that suits best your purposes.
- 13) What are the key pitfalls for server virtualization?
Bad pre-studies. When switching from a physical to a virtualized environment and being in a consolidation mode you will run into problems with some applications running on

those physical servers and are not supported in a VE. Licensing models and specifications are sometimes unclear and can also cause problems or even lawsuits.

- 14) In your view, do promoters of SV have other considerations regarding server virtualization?

As I said before you have this fashion world of IT and the big vendors want you to buy it. What you really need to do is to build up a long-term relationship not only one but maybe 2,3 different suppliers. The big vendors want to be the only supplier to your company. If you follow that you end up very dependant to them. It is ok to work with HP or IBM in some parts, but always try to work with some local companies, because they very often have experts as well, that can see things a bit different and unbiased. Always get a second opinion.

- 15) Are you staying in touch with the promoters after the deployment process of the virtualization? Do you give them feedback? And are the results meeting your expectations in general?

Yes, we still talk to vendors. Reviewing is done, future needs, updates on new technologies etc.

- 16) How do you monitor and measure its (SV) benefits in your organization? (ROI through cost savings, lower energy consumption)?

We don't have a good model for that today; we are trying to find a solution. We currently run a project named [REDACTED], so we can get these figures and calculate our actual ROI based on power consumption, cooling, hardware, staff costs, licensing costs... No extra energy invoice is available right now for the DC.

Transcription from interview with Deployers H

- 1) Please tell us what you believe are important decision points, when you decide if and how to implement server virtualization & 2) Why do you consider server virtualization?
The most important factor was that our servers were quite a few years old, so we needed new hardware. So alternatives were to buy 6 new servers and increase the amount step-wise if we need more than 6 OR virtualization. Also we had cooling problems before (in the datacenter). Also licensing costs are able to be lower for us.
- 3) What do you consider are appropriate to virtualize?
Servers only. The way people work here; they want to have their own computers and don't be dependent on someone's server configurations (thin client use). No future plans exist, since the variety of OS and applications is too broad.
- 4) How do you do the server virtualization (steps involved, design, implementation, test...)?
████████ provided us with its expertise. The ██████████ made the decision to hire ██████████ and go for virtualization.
- 5) Did you consider server virtualization do reduce power usage of the servers?
Not really, but it's a cost advantage.
- 6) Do you believe there are other areas where there are more power usage reduction to be made (DC)?
Another department spends time on reducing power consumption of the whole building, heating and ventilation, room lighting, not so much regarding computers (but heating and cooling issues connected to computers).
- 7) Do you believe there are any drawbacks to server virtualization for customers?
Not really since we use mainly Linux and don't have licensing problems (unlimited). The users are also satisfied, they don't even know about the change we made. Memory dedication might be complicated or needs to be checked several times.
- 8) What is your view regarding High Availability with server virtualization?
We want to have 100% of availability and we almost have it. Downtimes are 10 or 15 minutes usually, which is less time than the old strategy with physical servers.
Downtime happened when making the system transfers of the old to the new servers.
- 9) What is your view regarding Disaster Recovery with server virtualization? & 11) What is your view regarding backup and restore for server virtualization?
We do have mirror disks and regular backups on the servers, but mainly for crash recovery. About the servers, we store 6 or 8 generations of old files (at least one year old). Backups rely on tape backup (IBM).
- 10) What is your view regarding storage for server virtualization?

That's one of our main problems. We don't have resources for long time data storage. So far it hasn't affected us but it probably will in the near future since we will work with new technologies that create and need terabytes of data. We don't have the expertise and money yet to focus on that issue.

12) What are the key success factors for server virtualization?

Uptime, no downtime. Costs, instead of having a lot of small servers that don't perform on their max level we have fewer machines that operate nearly at full capacity (less hardware costs). It's a more economical way of doing it.

13) What are the key pitfalls for server virtualization?

Disaster cases when everything is going down (fire in server room). But we have locally divided backup disks and servers in 2 different buildings.

14) In your view, do promoters of SV have other considerations regarding server virtualization?

No, the whole project was a cooperation.

15) Are you staying in touch with the promoters after the deployment process of the virtualization? Do you give them feedback? And are the results meeting your expectations in general?

Yes, still working on the project. We are satisfied with their knowledge.

16) How do you monitor and measure its (SV) benefits in your organization? (ROI through cost savings, lower energy consumption)?

Well, we don't have actual figures or measure it. We only see that there is more room in the server room but we don't see if the power consumption is going down, since we are not able to see it. The bill is paid by someone else, not me. I only see bills about the network, but so far I haven't seen any decreasing there (but project is not finished now). Less hardware costs of course.

Transcription from interview with Deployers I

- 1) Please tell us what you believe are important decision points, when you decide if and how to implement server virtualization & 2) Why do you consider server virtualization?

Three main reasons that lead to the project:

Massive hardware upgrade situation leads to the idea to use the current hardware (save costs), and also use it more efficiently regarding energy consumption. Moreover a virtual environment is easier to manager for the technical staff.

Also the upper management is thinking about completely out- or in-source the IT used in schools and in the organization. An important step towards any of this decision is to centralize the data. (It can easily get outsourced now)

- 3) What do you consider are appropriate to virtualize?

The purpose only lies on server virtualization right now. There should be virtualized as much as possible in the future, e.g. applications and even the desktops. But no concrete plans are existing by now. The proportion between physical machines and machines used for the virtual environment are around 15:10 at the moment. Before the project started we had 24 physical servers. At the end we are planning to only have 4 physical servers left in total. The virtual environment will consist of 12-15 machines.

- 4) How do you do the server virtualization (steps involved, design, implementation, test...)?
- ████████ provided us with a strategy that we just agreed on and followed. Now we are coming to the point where our own personal decisions come more and more into the focus (regarding the specific circumstances we are working with).*

- 5) Did you consider server virtualization do reduce power usage of the servers?

Yes, as said before it was one of the benefits we were looking for. Also the local environmental office announced in its new strategy that local government should be more environmental focused.

- 6) Do you believe there are other areas where there are more power usage reduction to be made (DC)?

We are going to try to improve our policy system regarding the workstations at the schools. They should shut down themselves automatically in a certain time frame during the night or evening. That is a big issue for us.

- 7) Do you believe there are any drawbacks to server virtualization for customers?

There is a doubt about user satisfaction, mainly because of speed issues that might appear in the virtual environment. In the past local servers were set up in each school, now there is a central point where the data is stored and bigger bandwidth is needed. We have only a 100MB connection between the schools and the central point. Also education costs will appear, since we have to learn about managing this new kind of environment.

- 8) What is your view regarding High Availability with server virtualization?

Yes, we thought about better server performance. Also notification about troubles in the system will be received faster and easier than before. No more on-site service is needed to restart the server.

- 9) What is your view regarding Disaster Recovery with server virtualization? & 11) What is your view regarding backup and restore for server virtualization?

Virtual server backup programs are implemented. Every virtual server is backed up in our system through snapshots. This way we are hoping to improve the way of doing backups compared to our old solution.

- 10) What is your view regarding storage for server virtualization?

All files are located in a central point. This step was included in the server virtualization project.

- 12) What are the key success factors for server virtualization?

Lower energy consumption is one of them. Easier environment to handle technically. Easier way to handle the users. Before we had it spread out in many locations and now we have them more centralized and can implement policies in a much easier way. We hadn't any environmental thoughts when we were thinking about it. It was mostly considered how to set up a new and reliable system, which is easy to manage. However it's a side effect.

- 13) What are the key pitfalls for server virtualization?

User satisfaction as mentioned before.

- 14) In your view, do promoters of SV have other considerations regarding server virtualization?

No, we had a very long relation with [REDACTED] before and we trusted them.

- 15) Are you staying in touch with the promoters after the deployment process of the virtualization? Do you give them feedback? And are the results meeting your expectations in general?

We are still working with them and give feedback. We are very satisfied atm. They have a very good project method, since they do a lot of theoretical work before they start setting up the hardware.

- 16) How do you monitor and measure its (SV) benefits in your organization? (ROI through cost savings, lower energy consumption)?

There are no tools existing that monitor our power consumption and power savings. Anyhow it's logical that reducing of hardware in general leads to less cooling demands and therefore less power consumption. The big energy saving will be the power management tool that is going to shut down PCs automatically at schools. That's more important regarding energy savings than the virtualization of servers. The real amount of costs savings cannot figured out or calculated, since I don't have access to the current energy costs (different office is responsible) and its combined with the total rent the school has to pay. Therefore it's not the economical factor that motivates me towards this projects, instead I have the responsibility regarding the environment and to improve the local efficiency. I was aware of some problems, e.g. that computers won't shut off at night usually,

but I didn't had actual figures. A way to measure the cost savings is to talk to the technical office, which is responsible for paying these costs, and make them compare the costs before and after this project. However there are other factors that need to be considered, for instance if building heating must be increased if the PC's are shut off. So measuring is complicated.

Transcription from interview with Deployers J

- 1) Please tell us what you believe are important decision points, when you decide if and how to implement server virtualization & 2) Why do you consider server virtualization?
I don't think the power consumption was the main track why we started virtualization. It was more about the faster deployment of virtualized machine. That's the main reason. The power consumption is a benefit of it. Faster deployment means lower costs (hardware, staff) and also reduction space demand and further consumptions related to servers. High availability was also an issue for us.
- 3) Do you still have the usual machines or only the virtualized machines? And you could give some figures?
Half of our server park has been virtualized and it's increasing all the time. There are very fewer physical servers we implement today. The number of the servers is roughly 250.
- 4) How do you do the server virtualization (steps involved, design, implementation, test...)?
I can't give your good answer on that.
- 5) Did you consider server virtualization do reduce power usage of the servers?
Power reduction is cost saving of power was one of the things that we in the decision when we get what we will do. But I don't think it's the main question. Then main question is costs, hardware and fast implementation, easy administration.
- 6) Do you believe there are other areas where there are more power usage reduction to be made (DC)?
There must be something else that we can reduce power consumption I really hope but we have come quite far. On the heating of this house and the [REDACTED] next the door, we use the heat, which comes as the overflow from our server park. When we cool the computers in the server racks, we get so much heat that we can fire up this house and the one next door.
- 7) Do you believe there are any drawbacks to server virtualization for customers?
I don't see any licensing problems. Of course, the staffs who works with it has to have to be educated in what they are doing, but it doesn't matter what they do, because people always has to be trained on what they do.

I don't have clear view about the software licensing models, but maybe it is a problem.
- 8) What is your view regarding High Availability with server virtualization?
The ease of setting up new environments or making copies of exiting environments is one thing that is very good in SV. The speed of accessing the environment.

(Interviewer: But regarding to server downtime, I guess you save downtime because you implement SV which is the common opinion of high availability. Do you have downtime

before? Maybe you don't. Are there any crucial businesses processes going on which really rely on high availability here?)

I won't say that. That's nothing life depending. But if we ask our customers at the [REDACTED], they probably would say it. But not if I look at it from more board angle, we don't have any like that the hospital has equipment that life dependent. We don't have that things. We have email in many things that is regarded as life crucial, but actually it isn't. But of course, for the system (e.g. email communication) we need to have high availability on.

- 9) What is your view regarding Disaster Recovery with server virtualization? & 11) What is your view regarding backup and restore for server virtualization?

We have store the data into separate place. If there is something happen in the server, it will be easier for us to create a new environment or restore the old one a new with SV.

- 10) What is your view regarding storage for server virtualization?

Today we have dedicated areas on the storage for each system. The way we handle data in the future will be more dynamic.

- 12) What are the key success factors for server virtualization?

You need to have the staff behind you, since there are some doubts about the technology. Make everyone feel confident with it is one key factor.

- 13) What are the key pitfalls for server virtualization?

The business shouldn't have to care if we deliver our service from physical or virtual environment, so I don't see any problems there. The problem we can have is that we build quite complex chains. Everything is depending on each other. Before we had one system here and on there, now you have everything at the same place. So it's very important to update the map of the environment all the time.

- 14) In your view, do promoters of SV have other considerations regarding server virtualization?

I don't think that's a big issue.

- 15) Are you staying in touch with the promoters after the deployment process of the virtualization? Do you give them feedback? And are the results meeting your expectations in general?

Yes, still in touch. We have quite good skill here. Nowadays we only ask for help when we face very complicated problems that we need service on something. And we are satisfied in the high level. On the lower, support chain, area we are not impressed by our vendors. They don't deliver what we think they should support.

- 16) How do you monitor and measure its (SV) benefits in your organization? (ROI through cost savings, lower energy consumption)?

I don't think they did any measuring before. It's hard to say how much we measure since we are busy right now with increasing our services and servers and no real comparison can be made. If we look for it, we'll find how much we paid (hardware costs) and how much power consumption we save. But I can't give you the figure now. We might do it in the future.

Bibliography

- Adams, K., & Agesen, O. (2006). A comparison of software and hardware techniques for x86 virtualization. *Proceedings of the 12th international conference on Architectural support for programming languages and operating systems* (pp. 2-13). San Jose: ACM.
- Anderson, D. R., Sweeney, D. J., Williams, T. A., Freeman, J., & Shoemith, E. (2009). *Statistics for Business and Economics*. Cengage Learning EMEA.
- Arditi, D. (Ed.). (n.d.). *The Delphi Method*. Retrieved 03 23, 2009 from Illinois Institute of Technology: <http://www.iit.edu/%7Eit/delphi.html>
- Armstrong, W. J., Arndt, R. L., Boutcher, D. C., Kovacs, R. G., Larson, D., Lucke, K. A., et al. (2005). Advanced virtualization capabilities of POWER5 systems. *IBM Journal of Research and Development*, 49 (4/5), 523-532.
- Aurelius. (n.d.). *IT Security*. Retrieved March 2, 2009 from Open Security Architecture: <http://www.opensecurityarchitecture.org/cms/definitions/it-security>
- Badaloo, M. (2008, 08). *An examination of server consolidation: trends that can drive efficiencies and help businesses gain a competitive edge*. . From Server Consolidation - Server Optimization and Integration Services: <ftp://ftp.software.ibm.com/common/ssi/sa/wh/n/ssw03004usen/SSW03004USEN.PDF>
- Barham, P., Dragovic, B., Fraser, K., Hand, S., Harris, T., Ho, A., et al. (2003). Xen and the art of virtualization. *Proceedings of the nineteenth ACM symposium on Operating systems principles* (pp. 164-177). Bolton Landing: ACM.
- Baschab, J., Piot, J., & Carr, N. (2007). *The Executive's Guide to Information Technology*. John Wiley & Sons.
- Bradford, R., Kotsovinos, E., Feldmann, A., & Schlöberg, H. (2007). Live wide-area migration of virtual machines including local persistent state. *Proceedings of the 3rd international conference on Virtual execution environments* (pp. 169-179). San Diego: ACM.
- Chisnall, D. (2007). *The definitive guide to the Xen hypervisor*. Boston: Pearson Education.
- CIO Insight. (2009). *Research IT Trends for 2009: Cost Cutting Goals Expand- But don't take over, Social Responsibility Drive Environmental Moves*. From CIO Insight: <http://www.cioinsight.com/>
- CIO Staff. (2008, 03 17). *CIO's Green IT Survey: Cost Cutting, Social Responsibility Drive Environmental Moves*. Retrieved 03 17, 2009 from CIO: http://www.cio.com/article/197952/CIO_s_Green_IT_Survey_Cost_Cutting_Social_Responsibility_Drive_Environmental_Moves

- Creswell, J. W. (2007). *Qualitative inquiry and research design : choosing among five traditions*. Thousand Oaks: Sage Publications.
- Dayley, A. (2009, February 12). *Gartner Press Release*. Retrieved March 2, 2009 from <http://www.gartner.com/it/page.jsp?id=503192>
- Eckstein, P. P. (2006). *Angewandte Statistik mit SPSS: Praktische Einführung für Wirtschaftswissenschaftler*. Gabler.
- Ehrlinger, J., Gilovich, T., & Ross, L. (2005). Peering Into the Bias Blind Spot: People's Assessments of Bias in Themselves and Others. *Personality and Social Psychology Bulletin* , 31 (5), 680-692.
- Encyclopædia Britannica. (2009). *Computer*. Retrieved 03 29, 2009 from Encyclopædia Britannica Online: <http://www.britannica.com/EBchecked/topic/130429/computer/235890/Mainframe-computer#toc235890>
- Encyclopædia Britannica. (2009). *David Riesman*. Retrieved 03 28, 2009 from Encyclopædia Britannica Online: <http://www.britannica.com/EBchecked/topic/503294/David-Riesman>
- Engelmann, C., Naughton, T., Ong, H., Scott, S. L., & Vallée, G. (2008). System-Level Virtualization for High Performance Computing. *16th Euromicro Conference on Parallel, Distributed and Network-Based Processing* (pp. 636-643). IEEE.
- Fink, A., & Kosecoff, J. (2005). *How to Conduct Surveys: A Step-by-step Guide*. Sage.
- Gartner Press Release. (2008, 04 23). *Gartner Press Release*. Retrieved 03 02, 2009 from Gartner: <http://www.gartner.com/it/page.jsp?id=654011>
- Goldberg, R. P. (1973). Architecture of virtual machines. *National Computer Conference* (pp. 309-318). ACM.
- Goldman Sachs . (2007). *United States: Technology: Hardware report*. Goldman Sachs.
- Hammersley, M., & Gomm, R. (1997). Bias in Social Research. *Sociological Research Online* , 2 (1).
- Hartman, F. T., Skulmoski, G. J., & Krahn, J. (2007). The Delphi Method for Graduate Research. (P. Jerry, Ed.) *Journal of Information Technology Education* , 6, 21.
- IBM Global Services. (2007). *IBM Server Optimization and Integration Services – server consolidation efficiency study* . From <http://www-935.ibm.com/services/us/its/pdf/scon-sces-ds-gtd01525-usen-00-080607.pdf>
- IBM ITSO. (2008). *Systems Architect Handbook A Reference Guide for Systems Architects*. Poughkeepsie, NY, USA: IBM ITSO.
- IBM. (2008). *The New Enterprise Data Center* . Armonk,: IBM.
- Israel, M., & Hay, I. (2008). *Research Ethics for Social Scientists*. London: Sage Publications.
- IT Professional. (2008, 01). Virtualization Offers IT Departments the Way to a Greener Path. *IT Professional News Brief* , 10 (1), pp. 7-11.

- Karger, P. A. (2007). Performance and security lessons learned from virtualizing the alpha processor. *Proceedings of the 34th annual international symposium on Computer architecture* (pp. 392-401). San Diego: ACM.
- Koomey, J. G. (2007). *Estimating total power consumption by servers in the U.S. and the world*. Lawrence Berkeley National Laboratory.
- Kvale, S. (2007). *Interviews: an introduction to qualitative research interviewing*. Thousand Oaks,: Sage Publications.
- Lechner, R. (2009). *Building a Smarter Planet – Green and Beyond*. IBM.
- Linstone, H. A., & Turoff, M. (1975). *The Delphi Method Techniques and Applications*. Addison-Wesley Educational Publishers Inc.
- Martens, J. (2003). *Statistische Datenanalyse mit SPSS für Windows*. Oldenbourg.
- Mergen, M. F., Uhlig, V., Krieger, O., & Xenidis, J. (2006). Virtualization for high-performance computing. *ACM SIGOPS Operating Systems Review* , 40 (2), 8-11.
- Meyer, R. A., & Seawright, L. H. (1970). A virtual machine time-sharing system. *IBM Systems Journal* , 9 (3), 199.
- Mingay, S. (2007). *Green IT – The New Industry Shock Wave*. Retrieved 04 04, 2009 from Gartner: www.netdesign.dk/manedens-tema/telepresence/green-it-the-new-industry.pdf
- Murray, D. G., Milos, G., & Hand, S. (2008). Improving Xen Security through Disaggregation. *ACM/Usenix International Conference On Virtual Execution Environments* (pp. 151-160). Seattle: ACM.
- Norris, N. (1997). Error, bias and validity in qualitative research. *Educational Action Research* , 5 (1), 172-176.
- Popek, G. J., & Goldberg, R. P. (1974). Formal Requirements for Virtualizable Third Generation Architectures. *Communications of the ACM* , 17 (7), 412-421.
- Reid, T. R. (1992, 11 16). EPA Nudges firms to take the lead in green computing.
- Ruda, M., Denmark, J., & Matyska, L. (2007). Scheduling Virtual Grids: the Magrathea System. *Proceedings of the 3rd international workshop on Virtualization technology in distributed computing* (p. 7). ACM.
- Sackman, H. (1974). *Delphi Assessment: Expert Opinion, Forecasting, and Group Process*. Rand Corporation, Santa Monica.
- SandOaks. (2009). Retrieved 04 04, 2009 from <http://www.sandoaks.com/>
- Schulz, G. (2009). *The Green and Virtual Data Center*. Auerbach Publications.
- Seale, C. (2007). *The Quality of Qualitative Research*. London: Sage Publications.
- Sotomayor, B., Keahey, K., & Foster, I. (2008). Combining batch execution and leasing using virtual machines. *Proceedings of the 17th international symposium on High performance distributed computing* (pp. 87-96). Boston: ACM.
- The Green Grid. (2007, 02 16). *Guide lines for Energy-Efficient Datacenter*. From Microsoft Global Foundation Services:
http://www.globalfoundationservices.com/environment/documents/Green_Grid_Guidelines_WP.pdf

- Vetter, S., Dimmer, I., Haug, V., Huche, T., Singh, A. K., Venkataraman, A. K., et al. (2008). *IBM PowerVM Virtualization Managing and Monitoring*. IBM.
- VMware white paper. (2008). *How VMware Virtualization Right-sizes IT Infrastructure to Reduce Power Consumption*. VMware.
- Wikipedia contributors. (2009, 04 27). *Cloud computing*. Retrieved 04 28, 2009 from Wikipedia, The Free Encyclopedia: http://en.wikipedia.org/w/index.php?title=Cloud_computing&oldid=286525879
- Wikipedia contributors. (2009, 05 05). *Cronbach's alpha*. Retrieved 05 20, 2009 from Wikipedia, The Free Encyclopedia: http://en.wikipedia.org/wiki/Cronbach's_alpha
- Wikipedia contributors. (2009, 03 29). *IOMMU*. Retrieved 03 29, 2009 from Wikipedia, The Free Encyclopedia: <http://en.wikipedia.org/w/index.php?title=IOMMU&oldid=267266542>
- Wikipedia contributors. (2009, 02 21). *Nuremberg Code*. Retrieved 03 28, 2009 from Wikipedia, The Free Encyclopedia: http://en.wikipedia.org/w/index.php?title=Nuremberg_Code&oldid=272338525
- Wikipedia. (2009, February 11). *X86*. Retrieved March 2, 2009 from Wikipedia, The Free Encyclopedia: <http://en.wikipedia.org/w/index.php?title=X86&oldid=269984456>
- Yin, R. K. (2009). *Case study research: design and methods*. Thousand Oaks: Sage Publications.
- Young, C. J. (1973). Extended architecture and Hypervisor performance. *Proceedings of the workshop on virtual computer systems* (pp. 177-183). Cambridge: ACM.
- Yousuf, M. I. (2007). Using Experts' Opinions through Delphi Technique. *Practical Assessment, Research & Evaluation*, 12 (4), 8.
- Yu, Y., Guo, F., Nanda, S., Lam, L.-c., & Chiueh, T.-c. (2006). A feather-weight virtual machine for windows applications. *Proceedings of the 2nd international conference on Virtual execution environments* (pp. 24-34). Ottawa: ACM.