Instituto Superior de Engenharia do Porto

Optimization, High availability and Redundancy in Information Communications and Technology using Networks and Systems Virtualization Architectures

Narciso Artur Caldas Sousa Monteiro

Dissertation for the Master degree in

Engenharia Informática

Area of expertise

Arquitecturas, Sistemas e Redes

Supervisor: António Cardoso Costa

Co-supervisor: António Morim Brandão

Jury:

President:

Maria de Fátima Coutinho Rodrigues, Professora Coordenadora, ISEP – Instituto Superior de Engenharia do Porto

Vowels:

Adriano Carlos Alves Brito Lhamas, Professor Adjunto, ISEP – Instituto Superior de Engenharia do Porto

António Manuel Cardoso da Costa, Professor Coordenador, ISEP – Instituto Superior de Engenharia do Porto

Porto, September 2010

Acknowledgments

This space is dedicated to all of those who, directly or indirectly, gave their support and contribution to the elaboration of this thesis.

To Prof. António Costa, for accepting this work guidance, as well as for his availability and constructive contribution in all development phases.

To Eng. António Brandão, for his support as co-supervisor, for making available means for research and for his patience in having a team member in a mentally exhausted state during several months.

To Metro do Porto, SA, for allowing the use of data from its personal experience to complement this work with a practical application.

To my dear Sonia, for the encouragement and support she has given throughout my academic record and, especially, for her tenderness and patience.

To my parents, Filomena and Artur, for all the love and motivation they always gave me and for the flash of inspiration they had in order to have a child like me.

Thank you all.

Abstract

In an era marked by the so-called technological revolution, where computing power has become a critical production factor, just like manpower was in the industrial revolution, very quickly this dispersal of computing power for numerous facilities and sites became unsustainable, both economically as well in terms of maintenance and management. This constringency led to the need of adopting new strategies for the provision of information systems equally capable but in a more consolidated way, while also increasing the desired levels of the classic set of information assurances: availability, confidentiality, authenticity and integrity. At this point enters a topic not as new as it may seem, virtualization, which after two decades of oblivion presents itself as the best candidate to solve the problem technology faces today. Throughout this thesis, it will be explained what this technology represents and what approach should be taken in the evolution of a systems architecture classical model into a model based in optimization, high availability, redundancy and consolidation of information communications and technology, using networks and systems virtualization. A number of good management principles to adopt will also be taken into account to ensure quality standards in providing services, culminating in a practical application of this model to the infrastructure of the company Metro do Porto, SA. Briefly, this document can almost be classified as a step by step manual for the understanding and implementation of this model, transversal to any infrastructure without extremely particular needs or constraints which derail conversion.

Keywords: virtualization; consolidation; optimization; availability; redundancy

Resumo

Numa era marcada pela chamada revolução tecnológica, em que o poder de computação se tornou um factor crítico de produção, tal como a mão-de-obra o foi na revolução industrial, rapidamente esta dispersão do poder de computação por inúmeros equipamentos e locais tornou-se insustentável, quer ao nível económico como de manutenção e gestão. Este constrangimento levou à necessidade de adopção de novas estratégias para a disponibilização de sistemas de informação igualmente capazes mas de forma mais consolidada, procurando simultaneamente aumentar os níveis desejáveis do clássico paradigma dos factores a garantir na informação: disponibilidade, confidencialidade, autenticidade e integridade. É neste ponto que entra um tema não tão novo como parece, a virtualização, que após duas décadas de esquecimento se apresenta como o maior candidato à resolução do problema que a tecnologia enfrenta. Ao longo desta tese, será explicado o que esta tecnologia representa e qual a abordagem a ter na evolução de um modelo clássico de arquitectura de sistemas para um modelo baseado em optimização, alta disponibilidade, redundância e consolidação de tecnologias de informação e comunicações com recurso à virtualização de redes e sistemas. Serão igualmente tidos em conta um número de bons princípios de gestão a adoptar de forma a garantir parâmetros de qualidade na disponibilização de serviços, culminando numa aplicação prática de todo este modelo na infra-estrutura da empresa Metro do Porto, SA. Resumidamente, este documento pode ser classificado quase como um manual passo a passo de compreensão e implementação deste modelo, transversal a qualquer infra-estrutura sem constrangimentos ou necessidades extremamente particulares e que inviabilizem a conversão.

Palavras-chave: virtualização; consolidação; optimização; disponibilidade; redundância

Index

Figures list

Tables list

Abbreviations list

CPU	Central Processing Unit
DC	Domain Controller
DHCP	Dynamic Host Configuration Protocol
DNS	Domain Name Service
DOS	Denial of Service
DRS	Distributed Resource Scheduler
ERP	Enterprise Resource Planning
GIS	Geographical Information System
GUI	Graphical User Interface
I/O	Input/output
ICMP	Internet Communication Management Protocol
ICT	Information and Communication Technologies
IDE	Integrated Drive Electronics
IDS/IPS	Intrusion Detection System / Intrusion Prevention System
IP	Internet Protocol
iSCSI	Internet Small Computer System Interface
ISO	International Organization for Standardization
ISP	Internet Service Provider
ITIL v3	Information Technology Infrastructure Library version 3
IT	Information Technologies
LAN	Local area network
MMC	Microsoft Management Console
MPIO	Multiple Path Input/output

SAN	Storage Area Network
SAS	Serial Attached SCSI
SATA	Serial Advanced Technology Attachment
SCSI	Small Computer System Interface
SLA	Service Level Agreement
SMS	Simple Messaging Service
SNMP	Simple Network Management Protocol
vCPU	virtual Central Processing Unit
VLAN	Virtual Local Area Network
VMFS	Virtual Machine File System
VT	Virtualization Technology
WMI	Windows Management Instrumentation

1 Introduction

1.1 Problem guidelines

The virtualization concept is commonly seen as an emerging and revolutionary technology whose importance increased in recent years, with a capability never seen before and only made possible by the technological revolution of the 2000s. However, this concept was first launched in the '60s by Christopher Strachey, an Oxford University professor, with the aim of creating a way to take advantage of computing by sharing time, resources, processes and peripherals. Following this vision, two computers considered essential landmarks in the virtualization development history were born: The Atlas and project M44/44X by IBM. If by one hand Atlas pioneered in introducing the concept of virtual memory and paging techniques, the M44/44X project is actually the first physical equipment to run multiple virtual machines. The then current mainframes cost and size were IBM main incentives, which sought the best way to take advantage of the large hardware investment, allowing them to create and define the "virtual machine" concept [Dittner and Rule, 2007]. Interestingly, perhaps by a lack of vision or long-term planning, in the 80s and 90s the computer equipments proliferation and drastic cost reduction, coupled with the birth of more capable and easy to use operating systems, led to the abandon of the sharing resources idea in favor of the massive distribution of computing. What the industry did not realize at the time were the operational costs of managing large size physical infrastructures:

- Maintenance contracts
- Logistics and power consumption
- Complexity / diversity of the architecture and the consequent need for a large number of skilled personnel dedicated solely to maintenance tasks, often manual
- Low usage of the infrastructure total capacities
- Very complex disaster recovery and high availability scenarios
- Lack of scalability, leading to a continuous need for hardware investment, causing a snowball effect in the above factors

Already in the 2000s, when these infrastructures management reached an unbearable situation, very similar to that which led IBM in the 60s (which never abandoned virtualization and continued research and development [Singh, 2004]), virtualization has once again a word to say, most of it due to a company called VMWare.

Based on the work already developed by IBM, VMWare coordinated their efforts for solving a big challenge: the virtualization of x86 systems based on Intel 32bit architecture. The execution scheme of input / output instructions in x86 systems was not designed for this purpose, causing serious problems in the abstraction of resources. This is due to an identified set of instructions that require execution in privileged mode, in other words, direct contact with the hardware, instead of communication with an interface between the virtual operating system and the physical host, thus creating protection exceptions that lead to system blocking. This was a key point because technology would not succeed if the stability of the system was not guaranteed.

As result of a proprietary mechanism to monitor processing, VMware was able to circumvent this constraint and thus successfully implement its conceptual model, becoming a pioneer in x86 virtualization and soon after market leader. This mechanism consists in the dynamic rewriting of operating system kernel key parts in order to capture these sensitive instructions and therefore allow its interpretation to be performed by the virtual machine supervisor [VMWare, 2009].

1.2 Project presentation

This project consists in implementing a virtualization scenario in the Metro do Porto, SA critical services infrastructure. The company is endowed with a classical architecture, which provides good results in the field of availability and, due to its recent existence, is at a high rate of expansion. This is the context in which the company wants this growth and its management made in a sustainable manner and with clear improvement in service quality. For this purpose, virtualization presents itself as the most capable technology and, based on a previous costs comparison study, is also the most economically advantageous. The aim is to provide a development that will achieve higher scalability, lower operation and maintenance costs, as well as an increase in the availability rate, while also improving data/services backup methodologies.

In the end, the results must comply with the following goals:

- 99.99% service availability
- Infrastructure management consolidation
- Operation and maintenance annual costs reduction
- Growth capacity without hardware investment over the medium term
- Ability to recover services in a short space of time and with 24 hours maximum loss

1.3 Used tools

Throughout this paper several methods and tools were used to support the selected technologies assessment and implementation. In laboratory tests, the three major virtualization technologies put to test were: Microsoft Hyper-V 2008 R2, Citrix XenServer 5.5 and VMWare ESX 4. The performance data comparative analysis in these systems was performed by measuring the network and disk operations processing capacity (both measured in MB per second). For the first case, data was collected using the free tool Networx from SoftPerfect. This application counts the ongoing traffic in a given network interface, recording its speed and duration. For the hard disk operations measurement, the free software NBench was the resource used. It is an extremely simple tool that allows the simulation of file reading and copying into a hard drive, according to an operator preset file size and number of threads.

During the implementation, based on the VMWare ESX platform, it was necessary to use a tool for migrating systems into this platform. This tool was required both for the conversion of physical systems as well as virtual systems, due to the existence of a service in VMWare Server format. To serve this purpose, the vCenter VMWare Converter 4.0.1 Standalone was used.

In the outcomes assessment section, two tools were instrumental in obtaining a comparative analysis. For the survey of availability data, Ipswitch's What's Up Gold 14 was used. This is a monitoring system that, through simple protocols such as SNMP and ICMP, checks whether a given service is operating according to the parameters defined as acceptable.

In order to calculate the data center energy usage, and due to the lack of a real measurement process, the APC online calculator was used. This calculator allows the explicit introduction of all the equipment and its characteristics and generates a report of the required power.

1.4 Organization presentation

With around 120 employees, Metro do Porto, SA is a public transportation company, in particular the Porto's light railway system company. Founded in 1993, saw its first line going into operation on December 7 of 2002, in a network that currently has five lines, distributed over 70 stations along 60km [Moura et al., 2007].

Until the middle of 2006, the entire infrastructure of information technologies and communications was contracted and under another entity management. Due to corporate

growth and a need to increase the available services quality, the company decided to create the Office of Organization and Information Systems and consequently integrate the required technology in the organization itself. Currently, this office is responsible for the maintenance and continuous improvement of all the company's services and communications / telecommunications technologies, which are spread over two physically distant data centers, interconnected by optical fiber.

1.5 Work contributions

This paper aims to make a positive contribution in the following aspects:

- Test and implement a technological base for the sustainability of ICT in organizations
- Identify and document key aspects of security, management and control of an infrastructure
- Create a starting point for developing a business continuity management plan, producing a Service Level Agreement based on the Information Technology Infrastructure Library version 3 (ITIL v3) standard
- Demonstrate the actual benefits in moving to a virtual infrastructure

1.6 Thesis organization

This document is divided into four main chapters, guided by a logical sequence of implementation.

Introduction

As the name suggests, this chapter has an introductory note on the subject matter issue. The main motivations for the technological change are listed, with a description of the proposed solution and basic architecture.

Context

In this chapter, a presentation of the problem is made, framed in a broader context, with the description of all aspects related with this work. A survey of virtualization products on the market and their operation patterns is also done. After this identification, a detailed analysis of the individual characteristics of the three most advanced solutions is performed, followed by a direct comparison of the most important functions.

Security is also approached in this chapter where, besides an overview of the most important aspects to consider when implementing a security plan, emphasis is given on the main differences virtualization technologies have at this level.

Technical description

Divided into four subchapters, in this section all the analysis and deployment is performed on the proposed model. Based on the three state of the art products selected, a battery of individual tests is carried out, with a results analysis and potential application in the practical case. Based on these results the entire solution is actually implemented. Besides the whole planning and implementation process description, a comparative analysis of values before and after virtualization is carried out. To operate this model, and with great effect on the ITIL v3 standard, a series of good practices aimed at quality management are detailed. Some of these methodologies are actually applied in the infrastructure used for the practical case.

It is also intended to raise awareness to the importance of developing a continuity management plan, with focus on the most used techniques and the parameters that must be guaranteed.

Conclusion

Finally, a review of all methods and philosophies described in the document is conducted, leaving also guidelines to complement this work in future developments.

2 Context

2.1 Problem statement

As already mentioned, in the first three years of operation the organization was endowed with a classical infrastructure. In this particular case, this option was taken mostly out of necessity because, when it was decided to integrate the IT management in the organization, a working architecture already existed, in this case belonging to the company carrying out the management and services provision. In order to soften and make the departure from the service provider as simple as possible, the choice was to purchase equipment strictly equal to that in operation, requiring only data transfer, which allowed the new infrastructure to enter production as quickly as possible and without major hassles.

This classical structure was large, consisting (as good practices recommend) of a separate server for each service. In total, more than twenty servers were installed, whose storage was shared between the internal disks of each one and a Storage Area Network (SAN), only a couple of years away from becoming obsolete. After three years of operation, the organization's dependence on these services became more evident. This was mainly due to the company’s large number of partners and subcontractors, with which are traded large volumes of information, and the need for electronic communication is fundamental.

This is an organization that, having the public transportation as main business, is constantly growing and therefore has a great workload in the areas of inspection, architecture and engineering, areas that require continued monitoring and generate high volumes of information. The information flows these services provide quickly passed from essential to critical, also increasing their level of care. This is where the department responsible for this management, composed by five elements only, anticipated that the current infrastructure would not allow increasing quality of service as its management was complex and required too much maintenance.

After conducting a market survey of existing solutions for development, and also consulting some business partners, it quickly became apparent that current technology flowed in the direction of virtualization as a way to consolidate and optimize computing resources. In fact, by then the organization already had some legacy services running on a VMware Server platform, in an experimental basis, with good performance and management results. At this stage begins the study of available technologies, including its analysis and suitability to the organization operational requirements, in order to find the product that simultaneously allowed consolidating management, raising service availability levels and providing an

infrastructure sustainable growth. It should also be capable of providing effective redundancy methods and high availability features, seeking to minimize or even eliminate periods of maintenance downtime.

In this analysis, it is also essential to prove that the transition to a virtual architecture has no negative impact on the system's performance. To measure the availability data, the organization has a monitoring tool which records each service observed parameters. These parameters refer to a combination of service status, response time and performance and, if these values do not meet the defined range, in addition to triggering an alert, the service is considered unavailable. All these state changes and their duration are recorded in a database, which then allows report delivering and further analysis.

During the first three operational years, the following values were recorded for services classified as critical and whose continued operation should be ensured:

In 2007, the first full year of the Office of Information Systems operation, an average availability rate of 99.95% was achieved, representing a total of approximately 1836 minutes without service, as shown in Table 1.

Table 1 - Availability data for 2007

Servers - January 01, 2007 00:00:00 - December 31, 2007 00:00:00		
Service	**Unavailability time (minutes)**	**Availability (%)**
DC 1	1253,15	99,73%
DC 2	80,1	99,98%
E-mail	30,02	99,99%
ERP	40,01	99,99%
Maintenance service	20,44	99,98%
File share	80,04	99,98%
Print service	50,63	99,99%
GIS	220,92	99,95%
Webmail	60,36	99,99%
	1835,67	**99,95%**

The following year, and in consequence of some technological changes and various unexpected equipment failures, there was a decline to 99.83%, totaling about 6503 minutes of accumulated staging (Table 2).

Table 2 - Availability data for 2008

Servers - January 01, 2008 00:00:00 - December 31, 2008 00:00:00		
Service	**Unavailability time (minutes)**	**Availability (%)**
DC 1	352,54	99,92
DC 2	810,82	99,81
E-mail	1456,91	99,67
ERP	472,42	99,89
Maintenance service	80,05	99,98
File share	80,06	99,98
Print service	133,89	99,97
GIS	1558,21	99,64
Webmail	1558,07	99,64
	6502,97	**99,83**

In 2009, the year that infrastructure was converted, there was an improvement over the previous year, yet still below desired levels. The rate was 99.90%, which represents about 4420 accumulated minutes without service (Table 3).

Table 3 - Availability data for 2009

Servers - January 01, 2009 00:00:00 - December 31, 2009 23:30:00		
Service	**Unavailability time (minutes)**	**Availability (%)**
DC 1	270,81	99,95
DC 2	250,75	99,95
E-mail	601,19	99,88
ERP	410,89	99,92
Maintenance service	711,74	99,86
File share	790,98	99,84
Print service	570,89	99,89
GIS	180,84	99,96
Webmail	631,71	99,87
	4419,8	**99,90**

Although these values are acceptable, particularly since the downtime is largely related with planned maintenance periods, it is noted that a rate higher than 99.95% was never attained. Moreover, the annual average values discrepancy does not allow, in these circumstances, a

Service Level Agreement to be taken on, as there is no way to ensure compliance. With the expected growth and increasing demand level, the values recorded in this infrastructure will tend to deteriorate with the increasing complexity and available services variety. In addition to these data, this growth would be sustained by the constant addition of equipment, links and resources, potentially making the IT management as the area with the organization's largest operating cost and without a proportional return on investment. This virtualization project aims to solve all these operational difficulties, providing sustainable growth and a greater stability level. Predictably, and counting with the new services addition, it is expected to get a fast return on investment and an availability increase, right from the solution's production entry.

2.2 The virtualization of computing

Virtualization consists in the insertion of a software layer between the operating system and the underlying hardware, which implements an abstraction of that hardware and makes it available in a controlled manner to the upper layer. This layer is called the hypervisor, whose name derives from the original concept of supervisor, introduced primarily by Atlas [Singh, 2004], and represents a component created for managing system processes and provisioning of resources like memory and time sharing, which separates it from the component responsible for the effective execution of applications. This hypervisor layer can be of two types: hypervisor type 1 or type 2. In the first case, the virtualization software acts as an operating system, allowing its execution directly on the hardware, being considered that virtual machines operate on the second layer above hardware (Figure 1). Naturally, this solution has the most resemblance to a real machine and allows better performance of the guest operating systems.

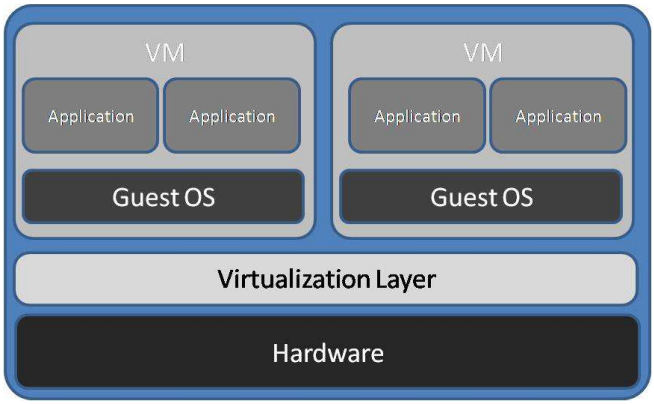

Figure 1 - Hypervisor Type 1

As for hypervisor type 2, it is unable to operate autonomously and requires another operating system to intermediate the communication between the virtualization software and the physical host, placing virtual machines on the third layer above hardware (Figure 2). The introduction of an additional layer naturally reduces the smoothness and performance of the interaction between the host and guest systems, but it becomes useful in testing and investigation scenarios, as it allows the intermediate operating system to capture and/or inject low level virtual machine instructions, typically for diagnosis and simulation purposes.

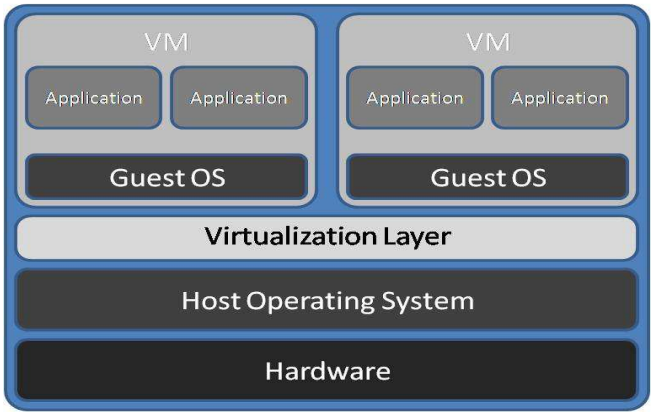

Figure 2 - Hypervisor Type 2

In both cases, the hypervisor function consists in making physical resources available to virtual machines in a transparent and secure manner, a sort of invisible moderator that aims to guarantee the virtual abstraction of hardware to each system and its respective applications. These are totally unaware of the presence of a logical layer and act as if a whole physical machine was under their control. One of the virtues of this abstraction type is the possibility to simulate the presence of a device or peripheral that actually does not exist, as well as the emulation of an interface type that omits the true hardware technology (for instance, the hard drive interface technology). Even so, there are known limitations [Heiser, 2007] to the type of controllers that can be emulated, especially at graphical level and with more specific or less common devices.

2.2.1 Security

With the pace of changes and developments that businesses are subjected at the IT level, security is not always seen as one of the key guidelines to consider. However, this is clearly a critical success factor and systems stability, and its proper definition, implementation,

testing and maintenance of high importance [International Organization for Standardization, 2005]. With the emergence of new systems paradigms, including virtualization (which due to the context of this work will be the most emphasized point), there are new concerns to add to the already complex traditional security plan.

According to data collected by Dataloss DB [Dataloss DB, 2010], only in the first quarter of 2010 were reported 102 security incidents in organizations, totaling 6,383,771 affected records, which clearly shows this problem's extent and the insufficient guidance for the need of defense mechanisms. There is still a long way to go in raising awareness, especially with the introduction of new technologies which, as already mentioned, conform primarily to timing issues, placing security in the background. The introduction of virtual systems brings new dimensions to this scenario, starting by the need to distinguish between physical and virtual events. Although the three major security areas whose planning must be ensured are the same (physical, logical and personnel security [Silva et al., 2003]), with systems virtualization much of the physical area is transformed into logical. If each server so far represented a physical host that could be protected and segmented by equally physical devices, its transition into a logical component and the hosting of several logical systems in a single place[1] brings a new dimension of difficult control and for which the protection technology is still too "embryonic" to use. This resource pooling brings an unwanted promiscuity among shared systems and makes it harder to audit inter-virtual machine traffic because, as they reside on the same host, communication is directly engaged by the hypervisor, eliminating any type of inspection, typically IP.

Even more, it is necessary to realize that the protection of the virtual operating systems is no longer sufficient because there is now another layer - the hypervisor with the management console - which is also an operating system, consequently having vulnerabilities and requiring protection. It is precisely here that the current technology is still limited, although there are some approaches that will minimize the risks [EMA 2010]. There are many discussions about how best to protect this layer but they are all still full of negative issues [VMWare, 2009], such as a strong impact on performance or the loss of virtualization functions. For instance, if a strategy of segmenting each server data network is adopted, though logical, features such as VMotion[2] may no longer be used, limiting the solution's potential and one of the major reasons for change [VMWare, 2009]. For this reason, the desired and recommended paradigm in depth defense becomes harder to apply due to the

[1] It is estimated that each company with virtualization has an average of 5 to 10 virtual servers per physical.
[2] VMware's technology that performs automatic resource management and high availability by migrating virtual systems between physical nodes.

elimination of fringes between systems [Oltsik, 2009]. In addition to all these factors, virtualization brings another problem, which is the diversity and ease of growth in the services and systems range. This exponential growth thus increases the management complexity and event recording, forcing a greater capacity in the organization for the correlation of these events. It is then essential to have a solution that provides the IT team with a quick, objective and easy view of status reports in order to properly filter what is actually important and requires intervention and monitoring [Oltsik, 2009].

Moreover, protection tools (IDS / IPS, firewall, etc.) set for each system can result in very high maintenance and licensing costs. In a survey [Ritter, 2009] conducted among several companies that offer their approach to security solutions optimized for virtualization, Trend Micro's Virtualization Security stood out. This tool, which results from the acquisition of Third Brigade by Trend Micro, brings a mixed concept. Initially the product is installed on the hypervisor itself in order to monitor the inter-VM traffic. However, as it is recognized that this approach has a strong impact on performance, the product has agents to install on each virtual machine to perform such work, freeing the hypervisor from that task. When a virtual machine is started, the application basis verifies if it has an agent. If so, it makes sure that it has a conformant security policy and then passes the defensive function into the agent. If not, control is then undertaken by the hypervisor (Figure 3). In short, this solution claims to have a coordinated application of defense mechanisms between the core system and the virtual machines themselves, enhancing protection and minimizing the impact on performance [Third Brigade, 2008].

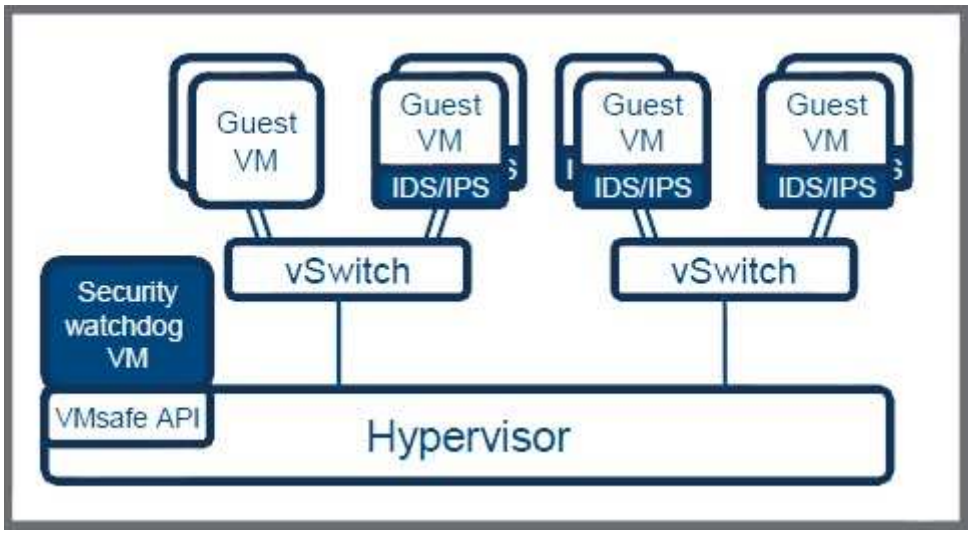

Figure 3 - Virtualization Security architecture

VMWare itself released a few white papers [VMWare, 2009] on different security architectures that can be implemented, which are recommended readings for anybody that seeks to adequately protect its virtual infrastructure. One of them is oriented to network segmentation, providing solutions for physical, logical or mixed distribution, with a comprehensive table of pros and cons for each one [VMWare, 2007]. In fact, many of these protections are even available in ESX standard version and it is only necessary to configure them properly. One of the recommendations consists in the use of resource constraints functionality, available in any VMWare version. If we consider that one or more virtual machines are targeted by an attack or infected with malware, this situation may lead to an unexpected increase in the host physical resources usage. Therefore, if resource constraints are not set, this disproportionate increase will impact other systems and, instead of isolating the problem in a single service, will ultimately lead to extraneous resource consumption, potentially causing a Denial of Service (DOS) throughout the whole structure [VMWare, 2009].

All these are new and unique concepts in virtualization and require a new approach to the security plan. Note that virtualization does not require a break with traditional security models, as all existing concerns must prevail. It is also important to predict typical situations, to be concerned with internal and external protection[3], with the backups, event logging, access control, basically with everything that is applicable in the physical world [Silva et al., 2003]. The important thing to retain is that all the traditional security paradigms are relevant but insufficient, for new virtualization oriented methodologies are required to maintain the optimum level of protection. It is necessary to take into account that virtualization is software and, just like any application, its strength lies also in the proper configuration. The most common security flaws are configuration errors and mistakes and, in this particular case, simple (and typical) errors like the assignment of a physical network adapter to an undue virtual machine expose the system and seriously undermine its security. A classic case is the development of templates for new systems, as virtualization solutions allow the creation of a base template with a particular operating system and configuration. This means that the addition of any new system is at a few minutes distance, simply requiring just two or three operations for creating and placing into production a new server based on that template. If that template is misconfigured or has unpatched vulnerabilities, these problems are sprawled into all new systems based on it [VMWare, 2009] [Hau and Araujo, 2007]. In consequence, it

[3] It is well known that peripheral protection is not enough, statistics show that 37% of attacks, accidental or otherwise, are of internal origin.

is of great importance to be aware of these differences and consider all risks right from the first interaction with the architecture [EMA 2010].

There is also new data to consider in the backups chapter, and this is clearly an important asset in an infrastructure. The fact that servers are turned into mere files, hardware independent, gives a never seen before portability which is ideal for creating recovery, development and testing environments, as well as for the creation of points in time (snapshots) that make change management much safer and more convenient. In this area, good practices also recommend a mixed approach attitude [Christensen and Howitson, 2008]. On one hand, it is recommended that each service is seen as physical, therefore classical backup techniques should be used. On the other hand, since there are new possibilities, these copies can go further and protect not only data but the service as a whole. A backup that includes data, together with an image of the system, will allow a total restore in minutes, in the same or another physical host. This rapid replacement, in many businesses, may actually mean millions of cost reduction due to a forced stop and, in parallel with a good critical services warranty policy (to be discussed in the chapter on business continuity), eliminates the potential loss of business opportunities due to service unavailability.

In conclusion, the key point is balance and common sense. Each organization should mitigate an appropriate identification of the risks it is exposed to and sort them by criticality [Yourdon, 2002] [Silva et al., 2003]. This will give an insight into what must be addressed in the security plan, which aspects should be focused (and which may be released) and what cost is associated with the required protection. It is often made the mistake of carrying out a security plan based on an audit checklist [Silva et al., 2003], but these examples are too generic and could result in unnecessary investments for the organization's reality and, on the other hand, leave in oblivion too much particular situations that general plans do not foreseen.

In the special case of virtualization, and because one can never insist too much on safety awareness, it should be clear that these considerations are paramount:

- Isolation and protection of the hypervisor and service console
- Periodical auditing of the configuration
- Segmentation and resource constraining native features usage
- Combining classical defense and safeguard methods with new virtual infrastructures oriented technology

2.2.2 State of the art

Following the general growing demand and interest in virtualization, many manufacturers and solutions arise in the market, with many flavors and goals. At least 73 different developments have been identified, ranging from proprietary to the most common systems, including the world of workstations. These developments are listed in the table below (Table 4), with a ratio between product, manufacturer, required operating system and licensing type.

Table 4 - Virtualization solutions in the market

Name	Manufacturer	Host operating system	License
Bochs	Kevin Lawton	Windows, Windows Mobile, Linux, IRIX, AIX, FreeBSD, OpenBSD, BeOS, Mac OS X	LGPL
CHARON-AXP	Stromasys	Windows 2003/2008 x64	Proprietary
CHARON-VAX	Stromasys	Windows 2003/2008, OpenVMS	Proprietary
Containers (also 'Zones')	Sun Microsystems	Solaris 10, OpenSolaris 2009.06	CDDL
Cooperative Linux	Dan Aloni and other programmers	Windows 2000, XP, 2003, Vista[1]	GPL version 2
Denali	University of Washington	Denali	?
DOSBox	Peter Veenstra and Sjoerd with community help	Linux, Windows, Mac OS Classic, Mac OS X, BeOS, FreeBSD, OpenBSD, Solaris, QNX, IRIX, MorphOS, AmigaOS	GPL
DOSEMU	Community project	Linux	GPL version 2
FreeVPS	PSoft	Linux	GPL version 2
GXemul	Anders Gavare	Unix-like	BSD
Hercules	Written by Roger Bowler, held by Jay Maynard	Unix-like	QPL
Hyper-V	Microsoft	Windows 2008 w/Hyper-V Role, Windows Hyper-V Server	Proprietary
Imperas OVP Tools	Imperas	Microsoft Windows, Linux	Tools under proprietary license, development under Apache 2
iCore Virtual Accounts	iCore Software	Windows XP	Proprietary
Integrity Virtual Machines	Hewlett-Packard	HP-UX	Proprietary
FreeBSD Jail	FreeBSD	FreeBSD	BSD License

Name	Manufacturer	Host operating system	License
JPC (Virtual Machine)	Oxford University	Java Virtual Machine	GPL version 2
KVM	Qumranet [4]	Linux	GPL version 2
LinuxOnLinux	Gelato@UNSW	Linux	GPL
Linux- VServer	Community project	Linux	GPL version 2
Logical Domains	Sun Microsystems	Solaris 10	?
LynxSecure	LynuxWorks	No host operating system	Proprietary
Mac-on-Linux	Mac On Linux	Linux	GPL
Mac-on-Mac	Sebastian Gregorzyk	Mac OS X	GPL
OKL4	Open Kernel Labs	No host operating system	BSD
OpenVZ	Community project, supported by SWsoft	Linux	GPL
Oracle VM	Oracle Corporation	No host operating system	Proprietary
OVPsim	OVP [5]	Microsoft Windows, Linux	Apache 2.0
Padded Cell for x86	Green Hills Software	INTEGRITY Real-time OS	Proprietary
Padded Cell for PowerPC	Green Hills Software	INTEGRITY Real-time OS	Proprietária
Palacios VMM	The V3Vee Project	Operating system independent	BSD
Parallels Desktop for Mac	Parallels, Inc.	Mac OS X (Intel)	Proprietary
Parallels Workstation	Parallels, Inc.	Windows, Linux	Proprietary
PearPC	Sebastian Biallas	Windows, Linux, Mac OS X, FreeBSD, NetBSD	GPL
PowerVM	IBM	No host operating system	Proprietary
Proxmox Virtual Environment	ProxMox	Debian Lenny com ProxMox Role	GPL v2
QEMU	Fabrice Bellard and other programmers	Windows, Linux, Mac OS X, Solaris, FreeBSD, OpenBSD, BeOS	GPL/LGPL
QEMU w/ kqemu module	Fabrice Bellard	Linux, FreeBSD, OpenBSD, Solaris, Windows	GPL/LGPL
QEMU w/ qvm86 module	Paul Brook	Linux, NetBSD, Windows	GPL
QuickTransit	Transitive Corp.	Linux, Mac OS X, Solaris	Proprietary
RTS Hypervisor	Real-Time Systems	No host operating system	Proprietary
SimNow	AMD	Linux (64bit), Windows (64bit)	AMD Proprietary
SIMH	Bob Supnik / The Computer History Simulation Project	Windows, BSD, Linux, Solaris, VMS	Unique, BSD type license
Simics	Virtutech	Windows, Linux, Solaris	Proprietary
Sun xVM Server	Sun Microsystems	No host operating system	GPL versão 3

Name	Manufacturer	Host operating system	License
SVISTA 2004	Serenity Systems International	Windows, OS/2, Linux	Proprietary
TRANGO	TRANGO Virtual Processors, Grenoble, France	No host operating system	Proprietary
User Mode Linux	Jeff Dike e outros programadores	Linux	GPL version 2
View-OS	Renzo Davoli and other programmers	Linux 2.6+	GPL version 2
VDSmanager	ISPsystem LLC	FreeBSD	Proprietary
Sun xVM VirtualBox	Sun Microsystems	Windows, Linux, Mac OS X (Intel), Solaris, FreeBSD, eComStation	GPL version 2; versão completa com funcionalidades empresariais é proprietária
Virtual Iron Virtual Iron 3.1	Virtual Iron Software, Inc.	No host operating system	All the product has proprietary license; some componentes are GPL version 2
Virtual PC 2007	Microsoft	Windows Vista (Business, Enterprise, Ultimate), XP Pro, XP Tablet PC Edition	Proprietary
Windows Virtual PC	Microsoft	Windows 7	Proprietary
Virtual PC 7 for Mac	Microsoft	Mac OS X	Proprietary
VirtualLogix VLX	VirtualLogix	No host operating system	Proprietary
Virtual Server 2005 R2	Microsoft	Windows 2003, XP	Proprietary
Hyper-V Server 2008	Microsoft	No host operating system	Proprietary
Hyper-V Server 2008 R2	Microsoft	No host operating system	Proprietary
CoWare Virtual Platform	CoWare	Windows, Linux, Solaris	Proprietary
Virtuozzo	SWsoft, agora Parallels, Inc.	Linux, Windows	Proprietary
VMware ESX Server	VMware	No host operating system	Proprietary
VMware ESXi	VMware	No host operating system	Proprietary
VMware Fusion	VMware	Mac OS X (Intel)	Proprietary
VMware Server	VMware	Windows, Linux	Proprietary
VMware Workstation 6.0	VMware	Windows, Linux	Proprietary
VMware Player 2.0	VMware	Windows, Linux	Proprietary

Name	Manufacturer	Host operating system	License
Wind River hypervisor	Wind River	No host operating system	Proprietary
Wind River VxWorks MILS Platform	Wind River	No host operating system	Proprietary
Xen	Citrix Systems	NetBSD, Linux, Solaris	GPL
XtratuM	Universidad Politecnica de Valencia	No host operating system	GPL
z/VM	IBM	No host operating system	Proprietary
z LPARs	IBM	Z Mainframes tool	Proprietary

To carry out a detailed comparative analysis focusing on the key features and technological orientation of each product, the three most reputable manufacturers and with the largest implementation base in the market were chosen, which prima facie appear as prime candidates for choosing a solution. Since we are dealing with server virtualization and with high performance and availability requirements, within each manufacturer the analysis will be restricted to products targeted at more featured servers and with professional support. We are referring to the solutions ESX from VMWare, Hyper-V Server 2008 R2 from Microsoft and Xen from Citrix Systems.

2.2.2.1 VMWare

Headquartered in California, VMware appears in 1998, launching its first product in 1999, VMWare Workstation. In 2001 enters the servers segment, now with hypervisor 1 and 2 solutions, and very quickly conquers the market, which hails the work done by the company. Albeit with some speculation, the company was a major driver for the smooth and well received way virtualization had when it returned as a hot topic (of course that the fact the product has free versions also had a tremendous influence on its success). After the introduction follows a detailed analysis of the ESX product, the hypervisor 1 solution for servers. Equipped with the latest technology, ESX version 4 has the following technical details [IT 2.0, 2009]:

Host characteristics

Hypervisor type: 1

Hypervisor nature*:* ESX (VMWare proprietary)

Licensing (without support): Free for the base version

Includes licensing for guest operating systems: No

Maximum number of logical CPUs: 64

Maximum memory: 1TB

Scalability: 320 virtual machines with a maximum of 512 vCPUs

SCSI disk support: Yes

SAS disk support: Yes

IDE/SATA disk support: Yes

iSCSI disk support: Yes

Fibre channel disk support: Yes

SAN MPIO support: Yes

Clustered file system: Yes

Live snapshots: Yes

Built-in thin-provisioning: Yes

VLAN support: Yes

Guest characteristics

Maximum number of vCPUS for each Linux system: 8

Maximum number of vCPUs for each Windows system: 8

Maximum memory: 255GB

Hot Add CPU: Yes

Hot Add memory: Yes

Hot Add disks: Yes

Hot Add network adapters: No

Supported operating systems: DOS, Windows 3.1, Windows 95/98, Windows NT4, Windows 2000, Windows 2003 x32, Windows 2003 x64, Windows 2008 x32, Windows 2008 x64, Windows XP x32, Windows XP x64, RH Linux Enterprise x32, RH Linux Enterprise x64, SUSE Linux Enterprise x32, SUSE Linux Enterprise x64, Netware, Ubuntu Linux, Debian, FreeBSD, CentOS, Sun Solaris, SCO Unixware, SCO Openserver, IBM OS/2 Warp

Management characteristics
Product name: vCenter

Centralized management console: Yes

Web management console: Yes (with limited functionality)

Centralized network management: Yes

Host live migration: Yes

Storage live migration: Yes

Fault tolerant: Yes

High availability module: Yes

Disaster recovery automation: Yes

2.2.2.2 Microsoft

Although exempt from introduction, Microsoft exists since 1975 and is headquartered in Washington. Microsoft could not be out of this race and in 2003 acquired Connectix, absorbing the technology of the company that developed the first Virtual PC and Virtual Server versions. Since then, Microsoft has been investing in improving its virtualization solution and in 2008 finally launches its first hypervisor 1 version, conveniently called Microsoft Hyper-V 2008 (currently in version R2) ,which is presented with the following characteristics [IT 2.0, 2009]:

Host characteristics
Hypervisor type: 1

Hypervisor nature*:* Hyper-V (Microsoft proprietary)

Licensing (without support): Free for base version

Includes licensing for guest operating systems: No

Maximum number of logical CPUs: 64

Maximum memory: 1TB

Scalability: 384 virtual machines

SCSI disk support: Yes

SAS disk support: Yes

IDE/SATA disk support: Yes

iSCSI disk support: Yes

Fibre Channel disk support: Yes

SAN MPIO support: Yes

Clustered file system: Yes

Live snapshots: Yes

Built-in thin-provisioning: Yes

VLAN support: Yes

Guest characteristics

Maximum number of vCPUS for each Linux system: 4

Maximum number of vCPUs for each Windows system: 4

Maximum memory: 64GB

Hot Add CPU: No

Hot Add Memory: No

Hot Add disks: Yes

Hot add network adapters: No

Supported operating systems: Windows NT 4.0, Windows 2000, Windows 2003 x32, Windows 2003 x64, Windows 2008 x32, Windows 2008 x64, Windows XP x32, Windows XP x64, Windows Vista x32, Windows Vista x64, RH Linux Enterprise x32, RH Linux Enterprise x64, SUSE Linux Enterprise x32, SUSE Linux Enterprise x64

Management characteristics
Product name: System Center Family

Centralized management console: Yes

Web management console: No

Centralized network management: No

Host live migration: Yes

Storage live migration: No

Fault tolerant: No

High availability module: Yes (through Microsoft Cluster Server)

Disaster recovery automation: No

2.2.2.3 Citrix Systems

Specialized in applications virtualization and remote access, Citrix Systems was founded in Texas in 1989 by a former IBM programmer, having rapidly moved to Florida, the birthplace of the founder. Despite the expertise from IBM, the company had a difficult beginning with no profits, being forced to resort to a partnership with Microsoft and Intel to survive. As a result of this relationship with Microsoft, Citrix specialized in remote access applications and would be the maker of Microsoft Windows Terminal Server Edition, one of the most popular remote access servers. They are currently the market leader in applications and workstations virtualization. Only in 2007 the company acquired the Xen open-source technology and launched itself in the hypervisor 1 virtualization for servers market. XenServer (currently in version 5.5) is the result of this release and is endowed with these characteristics [IT 2.0, 2009]:

Host characteristics
Hypervisor type: 1

Hypervisor nature*:* Xen derived

Licensing (without support): Free for base version

Includes licensing for guest operating systems: No

Maximum number of logical CPUs: 32

Maximum memory: 128GB

Scalability: 50 virtual machines

SCSI disk support: Yes

SAS disk support: Yes

IDE/SATA disk support: Yes

iSCSI disk support: Yes

Fibre Channel disk support: Yes

SAN MPIO support: Yes

Clustered file system: No

Live snapshots: Yes

Built-in thin-provisioning: No

VLAN support: Yes

Guest characteristics
Maximum number of vCPUS for each Linux system: 8

Maximum number of vCPUs for each Windows system: 8

Maximum memory: 32GB

Hot Add CPU: No

Hot Add memory: No

Hot Add disks: No

Hot add network adapters: No

Supported operating systems: Windows 2000, Windows 2003 x32, Windows 2003 x64, Windows 2008 x32, Windows 2008 x64, Windows XP x32, Windows Vista x32, RH Linux Enterprise x32, RH Linux Enterprise x64, SUSE Linux Enterprise x32, SUSE Linux Enterprise x64, Debian, CentOS, Oracle Enterprise Linux

Management characteristics
Product name: XenCenter

Centralized management console: Yes

Web management console: Yes (with limited functionality)

Centralized network management: No

Host live migration: Yes

Storage live migration: No

Fault tolerant: No

High availability module: Yes

Disaster recovery automation No

Conducting a comparative analysis of characteristics, and consulting some performance evaluations available, the reason for VMWare‘s leadership becomes very obvious, also derived from their larger experience in this area. However, looking at data from Microsoft’s previous versions, it is possible to identify a clear evolution in the product capabilities, which should approach VMWare in the short term, at least in technical characteristics since the performance difference is much more pronounced. For a more specific look at the technological differences, the values that are most relevant in a real implementation of a typical information systems scenario were chosen.

2.2.2.4 Maximum number of logical CPUs

The fact that Xen only supports 32 processors, compared to their competitors 64, may not be a major factor at the moment but with the growing number of processors installed on the hardware by the manufacturers and in the presence of large clusters, this is a handicap for the Citrix solution.

2.2.2.5 Host maximum memory

Xen gets once again behind the competition and this time in a critical factor. The system scalability depends on the available memory and a 128GB limit puts the solution out of large projects where there are a high number of systems to virtualize.

2.2.2.6 Thin-provisioning capability

Thin-provisioning represents the ability to partially allocate resources to a system, i.e., although certain resources are made available to a system it will only use those that are actually needed, allowing the surplus to be used for functions in other systems. Xen's inability to perform this process adds to it a further loss in value.

2.2.2.7 Maximum number of vCPUs and memory for each host system

Just like in the first point, the growth in the number of CPUs per machine will allow a larger number of processors to be made available to guest systems. The drawback here goes for Hyper-V, which allows only half of the eight supported by the other products. Memory is another very important factor and if the 64GB from Hyper-V seem yet sufficient for the current reality, Xen's 32GB are likely to be scarce in highly demanding systems. Anyway, both are well below the 255GB from ESX.

2.2.2.8 Hot add of CPU, disk and memory

One of the virtualization main goals is to increase the systems availability rate along with the resource management versatility. Therefore, the possibility to adjust memory, CPU or disk to a guest system without interrupting its production is essential, especially imagining a scenario where a particular system reaches an unexpected production peak and, in order to smooth and expedite its operations, a larger number of resources are assigned during this period of time. All this is possible in ESX, only the addition of disks is possible in Hyper-V and none of these operations is supported in Xen.

2.2.2.9 Storage live migration

This point represents the ability to migrate data between storages without stopping the corresponding system, useful to provide an infrastructure with a superior high availability capacity. Only ESX supports this feature.

2.2.2.10 Centralized network management

If one of the goals of virtualization is systems consolidation, it is equally true that management tools consolidation has a significant role in the system's management ease and convenience, causing a direct impact on service intervention times. Once again, only ESX is able to perform network management in the same module of the remaining system management.

2.2.2.11 Fault tolerance

Represents the capacity that a server has to continue operation even in the presence of a fault, either physical or logical (and not critical to the system integrity). Again, it is a feature only available on ESX.

2.2.3 Conclusion

Unsurprisingly, ESX comes out clearly as a winner in this comparison and presents itself as the most suitable solution for implementing a virtualization project in an infrastructure with heterogeneity, a large dimension and without a very specific need for virtualizing a system only supported in another application (like Oracle Enterprise Linux, available only on Xen), even though a mixed technology scenario could be considered, although this invalidates the desirable scenario of consolidation and solution centralized management.

3 Technical description

3.1 Analysis

In the services listed as critical, the current situation allows a classical infrastructure where each service is provided by dedicated equipment and, in some cases, with data divided between local storage and network storage (Figure 4).

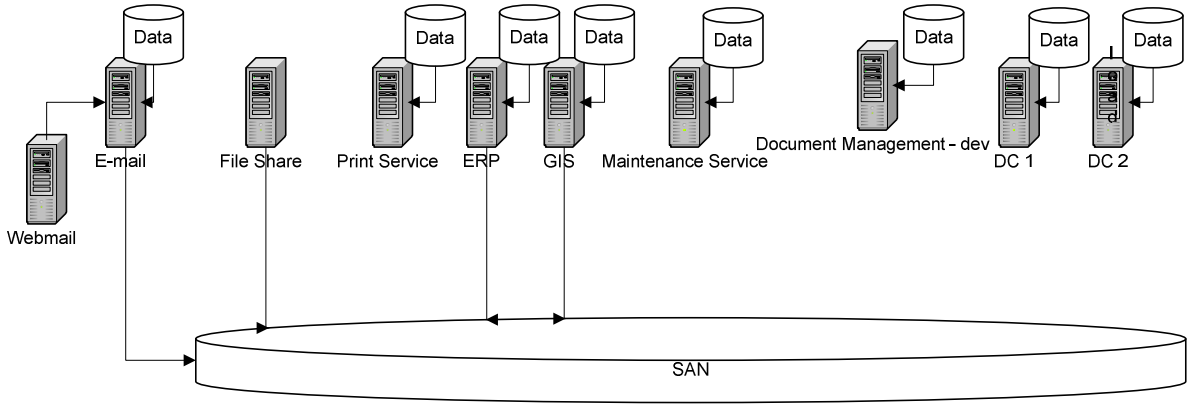

Figure 4 - Initial architecture diagram

This scenario, in addition to multiple failure points, has several points of decentralization that make its management harder, particularly:

- One management interface for each service
- The data storage dispersion increases management and backup methods complexity
- Different versions and generations equipment
- Each service depends on the physical platform where it is housed
- Lack of scalability, both in terms of current and future services

Besides these factors, it must also be taken into account that each one of these devices requires a maintenance contract. Bearing in mind that these are distinct generations of technological platforms, the contracts management is administratively complex. In addition, equipments with high environmental and energy supply needs are required. The desired solution must therefore address these problems in order to create an abstraction of these services relative to the physical platform where they are housed. This approach should concentrate all these services in a single location, where the underlying hardware management is transparent and seen as unified. Thus, not only will the management be

centralized as the available resources may be distributed according to each service needs (Figure 5).

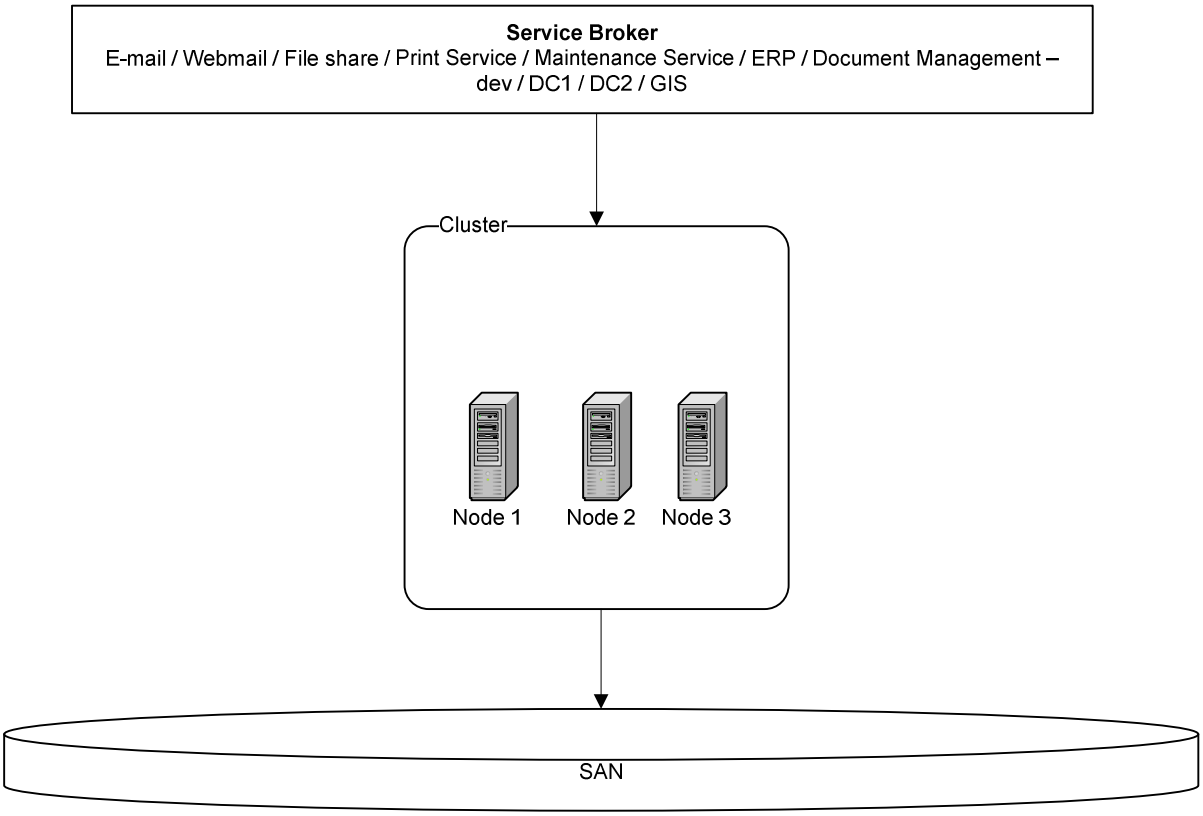

Figure 5 - Desired architecture diagram

As the previous figure illustrates, the goal is to have services a level above the hardware, leaving the resource management for the technological platform, enabling the transparent change of services location according to its needs. With this solution, it is also intended that all data, including each service operating system, is placed in the SAN, likewise allowing it to run regardless of the physical platform. In this scenario, the capacity of the ten baseline equipments should be assured by only three, with the possibility of adding nodes to this cluster if future services growth requires increasing the available resources.

3.1.1 Requirements / Features

The system to be implemented should ensure transparent access to services, regardless of their physical location (Figure 6). This access requires continuous availability and be under

the same performance conditions, which means it must manage a resource pool in order to properly balance each service allocated resources, as well as compensating for possible equipment failures by automatically and immediately moving the affected service to the most available node.

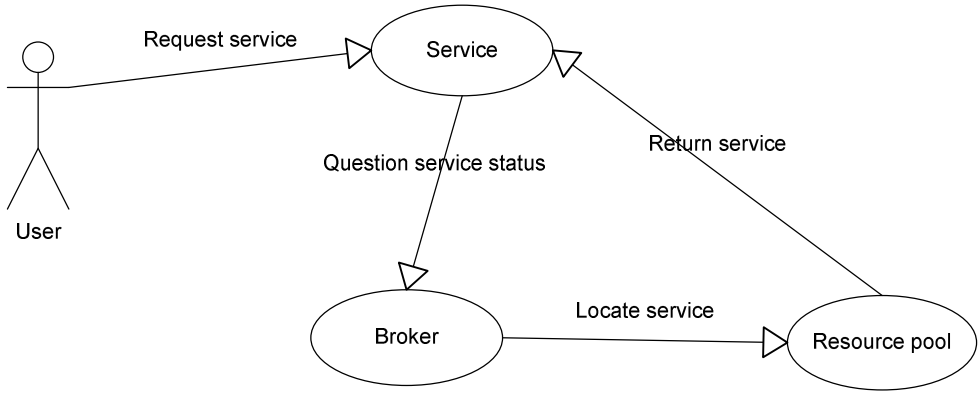

Figure 6 - Interaction set for a service request

Similarly, by placing services at a higher level, it should bring them enough portability to allow its temporary placement in another independent physical platform, providing "immediate" recovery facilities in extreme failure situations.

3.1.2 Processes

Access to these services can be classified according to the processes:

- Authentication
- E-mail management
- Webmail
- Printing
- File sharing
- Perform maintenance
- Development
- Management of geographical information
- Resource planning

The following diagram (Figure 7) shows the actors and their interaction with these processes:

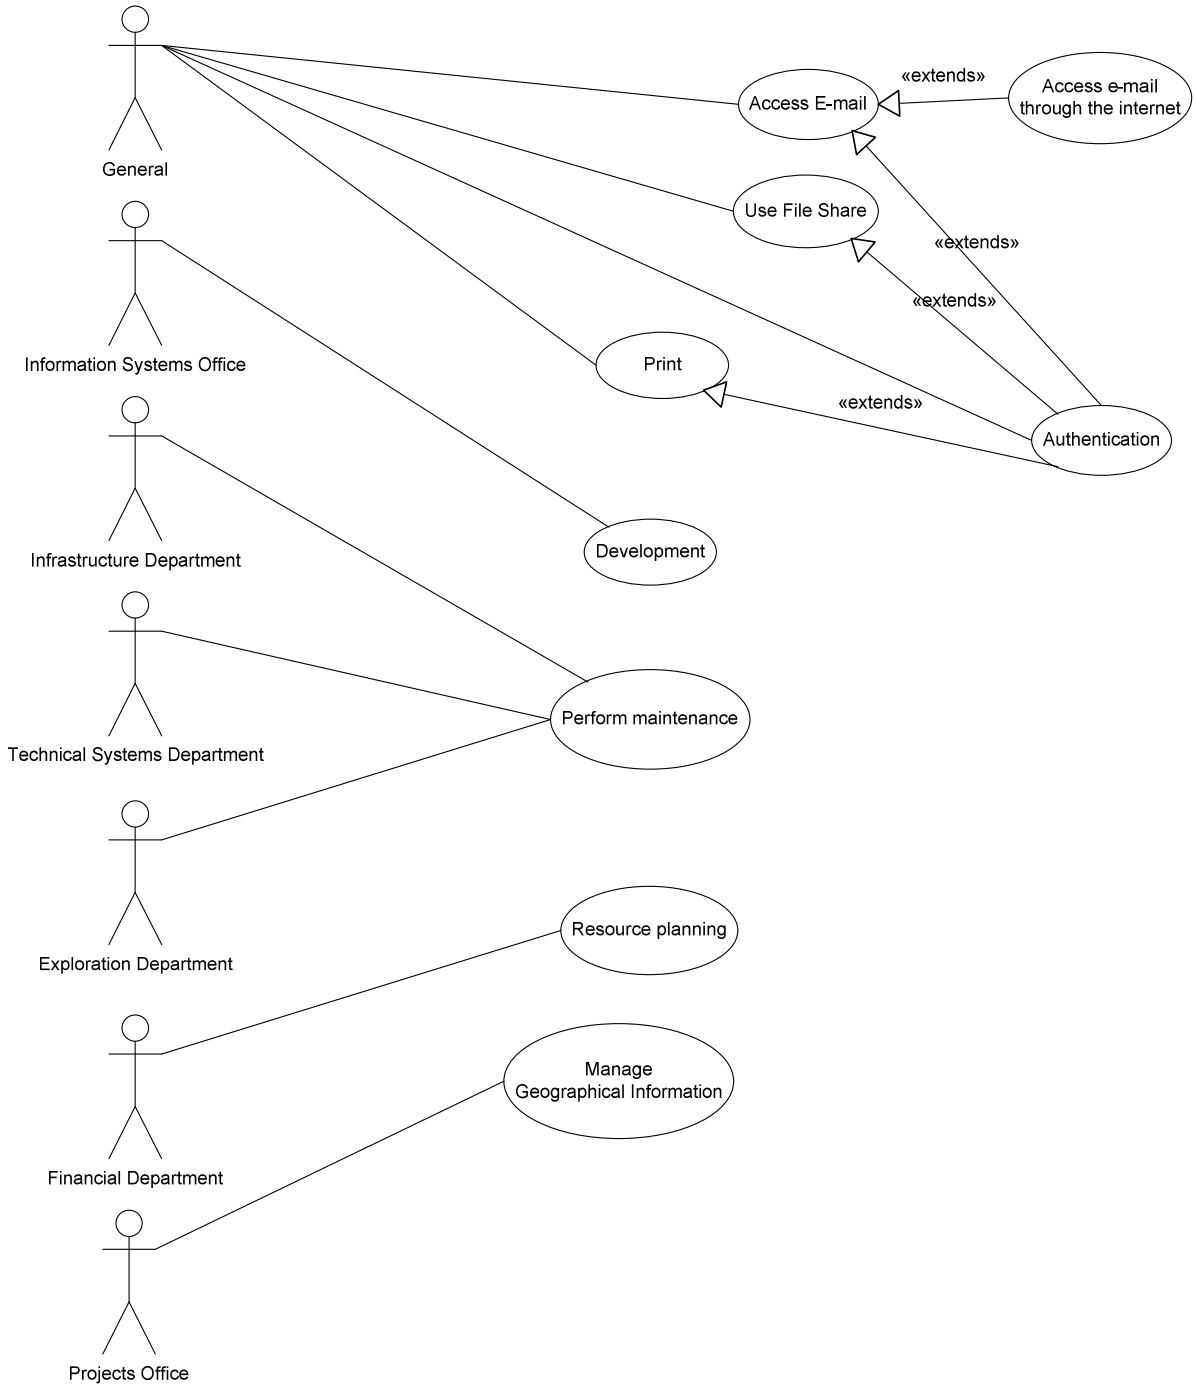

Figure 7 - Interaction diagram

3.1.3 Logical structure

As identified in the previous section, there are seven distinct entities interacting with nine system processes.

General Entity

Fits all the elements of the organization, without exception. Their interaction is accomplished through the processes:

- Authentication – in this process, the services represent domain controllers that allow credentials validation for accessing the system.
- E-mail management – dependent on the authentication service, this process provides access to the e-mail system. The external access, which is just a portal for interconnecting with this system, depends on it.
- Printing – printing services are made from this process. It is also performed using the authentication service.
- File sharing – also dependent on the authentication service, this process is associated with the file-sharing service.

Information Systems Office Entity

This entity includes solely the elements of the information systems office, responsible for the whole IS / IT area in the organization. Only a unique process of interaction is necessary:

- Development – development of the document management service is accomplished in this process. The development team uses it to prepare the future document management system.

Technical Systems Department, Exploration Department and Infrastructures Department Entities

These entities have a more direct relationship with the effective operation of the light metro system. In the Technical Systems lies the whole electrical and signaling area, in the Exploration all the operation planning and supervision is done and the Infrastructures is responsible for the civil, geological, environmental and archaeological areas.

- Perform maintenance – the maintenance service provides an application that connects to all equipments placed at stations (escalators, vending machines, HVAC, etc.) and controls its condition.

Projects Office Entity

This is the area responsible for preparing and reviewing architectural projects.

- Management of geographical information – this process includes the geographical information service which, as the name suggests, manages the detailed geographical data of the Porto metropolitan area.

Financial Department Entity

This department includes all the accounting and human resources area.

- Resource planning – the access to the ERP service allows the execution of all accounting, commercial and staff management.

3.1.4 Physical Structure

The system's physical elements identification, with corresponding interfaces, relates to the following components diagram (Figure 8).

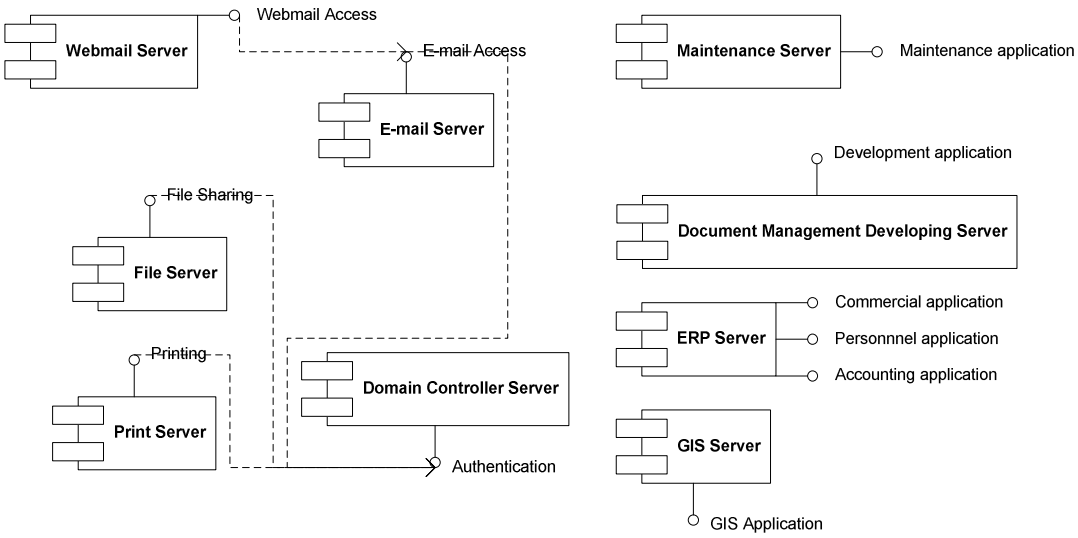

Figure 8 - Components diagram

In order to support these components, and according to the desired architecture features, the required hardware consists of three servers, fiber optical storage and its corresponding connection modules between servers and storage, as shown in Figure 9.

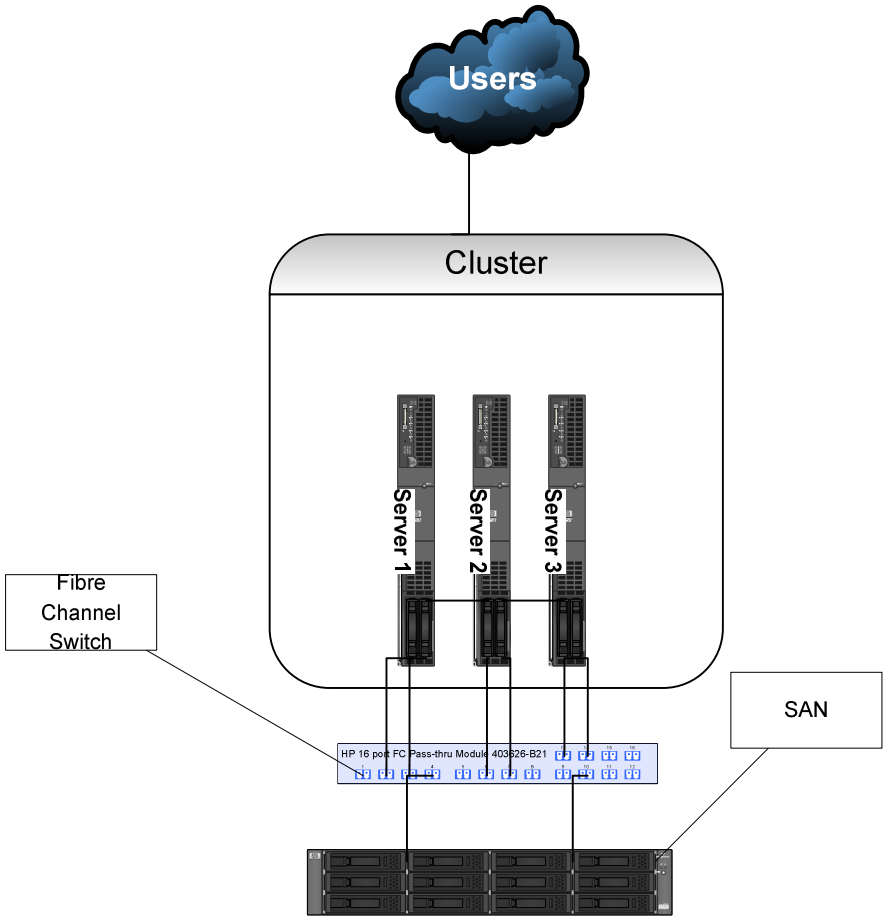

Figure 9 - Deployment diagram

3.2 Development

3.2.1 Testing laboratory

For a better understanding of the characteristics collected for each solution, a small testing lab was created, focused on the validation of its suitability to the infrastructure that will be used as a case study. It essentially seeks to ascertain the ease and feasibility of implementing each technology, simulating its application in a realistic scenario.

To support this laboratory, Metro do Porto, SA installed an HP Proliant DL380 G4 server with the following characteristics:

Processor: 2 x Intel Xeon 3.2Ghz

Virtualization Technology capable: No

Memory: 2GB

Storage: 36GB RAID 1 + 72GB RAID 5

Although this equipment was previously used in the organization and accomplished a good performance record, for this laboratory the memory capacity shortage and the absence of VT[4] may cause some limitations. However, it is also an opportunity for testing the solutions in minimal conditions without the need for hardware investment, which is certainly the case for many companies wishing to use virtualization without having budget to renew existing servers. Still, this limitation will not allow emulating the complete environment that is intended for the practical case, therefore the option was to choose three essential services with different operating characteristics:

- Server 1 with Microsoft Windows 2003 Server and Microsoft Exchange Server 2007
- Server 2 with Red Hat Enterprise Linux and Trend Micro Interscan Messaging Security Suite mail relay virtual appliance
- Server 3 with Microsoft Windows 2003 Server and file sharing service

It should be taken into account that neither the high-availability nor the advanced automatic migration capabilities of each solution can be tested because they are not available in their free versions.

3.2.1.1 Outcome metrics

To reach a conclusion on each technology, relevant evaluation metrics must be identified to clearly differentiate each solution and hence its eligibility for the practical case, in detriment of others. Among dozens of parameters that could be chosen for this purpose, it was decided to stick to three key points which clearly demonstrate if the solution serves the intended purpose.

[4] A technology that enables the sharing of processor resources between several operating systems simultaneously.

- Supported operating systems: The candidate solution must support without limitation all the operating systems required for the replication of the practical case scenario.
- Performance: The data input / output operations and access times to services should be as similar as possible to the values recorded before virtualization.
- Scalability: Each technology has its own approach to the growth factor and the optimal scenario is one that allows a transparent growth, without significant impact on users and easily performed.

3.2.1.2 Conclusions

Put to test, the systems showed quite different results.

For the technical analysis and listing of the collected data, XenServer and Hyper-V cases will be treated solely in terms of installation and management interface as none of them allowed the testing environment completion, therefore evaluation metrics could not be determined.

XenServer

The XenServer installation is simple and quick, allowing immediate configuration of key parameters and leaving the system ready to use. The media, which can be downloaded for free from the Citrix website, also contains the remote management console setup. However, at the start of the product installation a warning (host processor did not have VT technology) was issued and it was not possible to host Microsoft systems, invalidating the solution for this particular scenario. Even so, the test continued, at least to verify the management console potential. It is noteworthy that the system local interface is highly elaborated and allows virtual machines basic management, including the storage location, something that does not exist in other products (Figure 10).

Figure 10 - XenServer local console

The management console was within the expected range for this type of solution: it is well organized, easy to use and practical for most common parameters in VMs creation and maintenance. In the deployment of a virtual machine two limitations were confirmed. One of them, already identified in the product features survey, has to do with the very restricted number of supported operating systems. With the exception of corporate versions of Linux, it is not possible to install another system. Moreover, it was confirmed that the existence of the host VT feature is a mandatory requirement for installing Windows systems. As such, the Xen test was abandoned and deemed an unworkable solution for this scenario.

Hyper-V 2008 R2

Next the Hyper-V product laboratory was tested. One immediate difference is the installation media, which was supplied on a complete DVD, compared to several CDs in the other solutions. This hypervisor has actually a 2GB footprint which, in comparison with XenServer 400MB and ESX 200MB seems excessive for a "mere" operations intermediary. As in Microsoft Server 2008, there are two installation modes: full (with a Graphical User Interface (GUI)) or core (command line only). Besides having a GUI is unnecessary, for a direct comparison the core option was chosen. Almost an hour later, without any required intervention, the product asked for basic data such as credentials and network configuration. After restarting, the system was ready to use. Once the session started, the management console shown was essentially a powershell script with basic configuration options, lacking virtual machines management (or even visualizing). Once this step was finished, and since the network configuration was specified during the installation, the next step was configuring the remote access, which allows managing all the product features.

By default, access via Terminal Server is active and authorized by the advanced firewall, which has been recently introduced in Microsoft server's product range. This was, among the tested products, the only one base equipped with a security feature. However, Terminal Server only allows accessing the local console so, on this same console, three additional remote access types must be turned on: powershell, Server Manager and Microsoft Management Console (MMC). This means that only a Windows client can manage the solution, which was hardly a surprise. Excluding Server Manager, which essentially allows managing the services state, user accounts, storage and checking the event log, the other two alternatives have different contexts. The powershell implies a mastery of the language, with the addition of Hyper-V specific commands, and is not the quickest or the simplest method for recurring maintenance tasks. The MMC allows access to a management console similar to the other products.

On the client workstation, it should be noted that a specific snap-in in the MMC is required to manage Hyper-V (Figure 11).

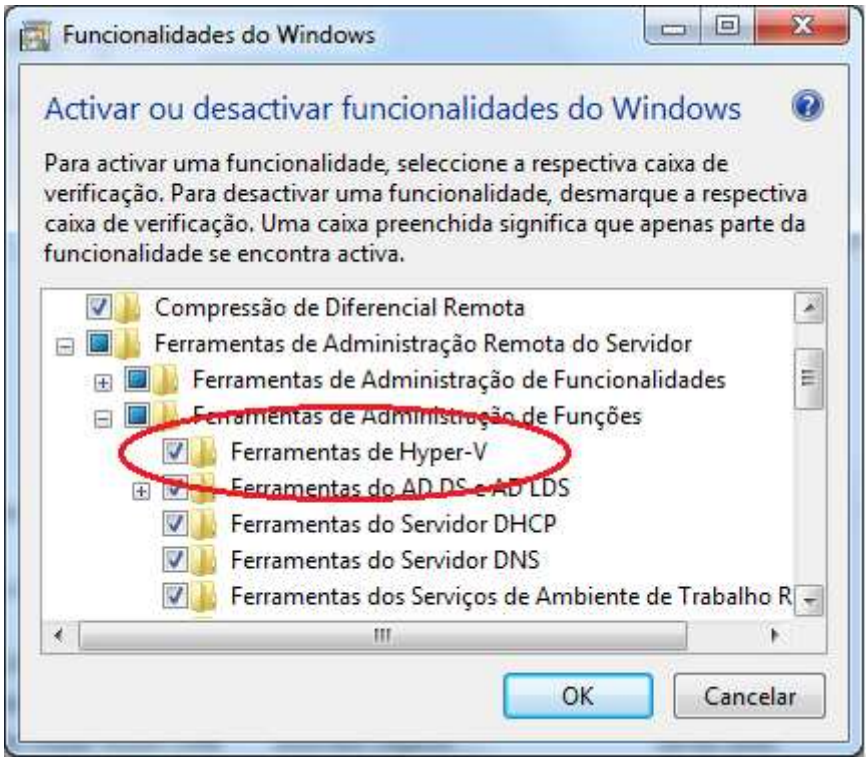

Figure 11 – Hyper-V administration tool

Once this configuration was completed, the "Hyper-V Manager" option was added in the administrative tools. The first attempt to connect to the server failed and it took some time to understand why. The fact that this server is not on the same domain of the client workstation is mandatory and requires some configuration maneuvers. As there was no obstacle, it was more practical to add Hyper-V to the existing domain. This is an interesting option for those who work with Active Directory structures, for the entire permissions hierarchy and group policies are implicitly applied to Hyper-V. However, it seems unnecessary exposing and basing critical hypervisor functions in such high level structure. Already with full console access, the visual elements disposal was easy to understand because the interaction method is similar to the other virtualization products (Figure 12). In fact it seems that all followed a standard line in the creation of management consoles. Once again, for those who deal with Microsoft systems, working with this interface is extremely intuitive.

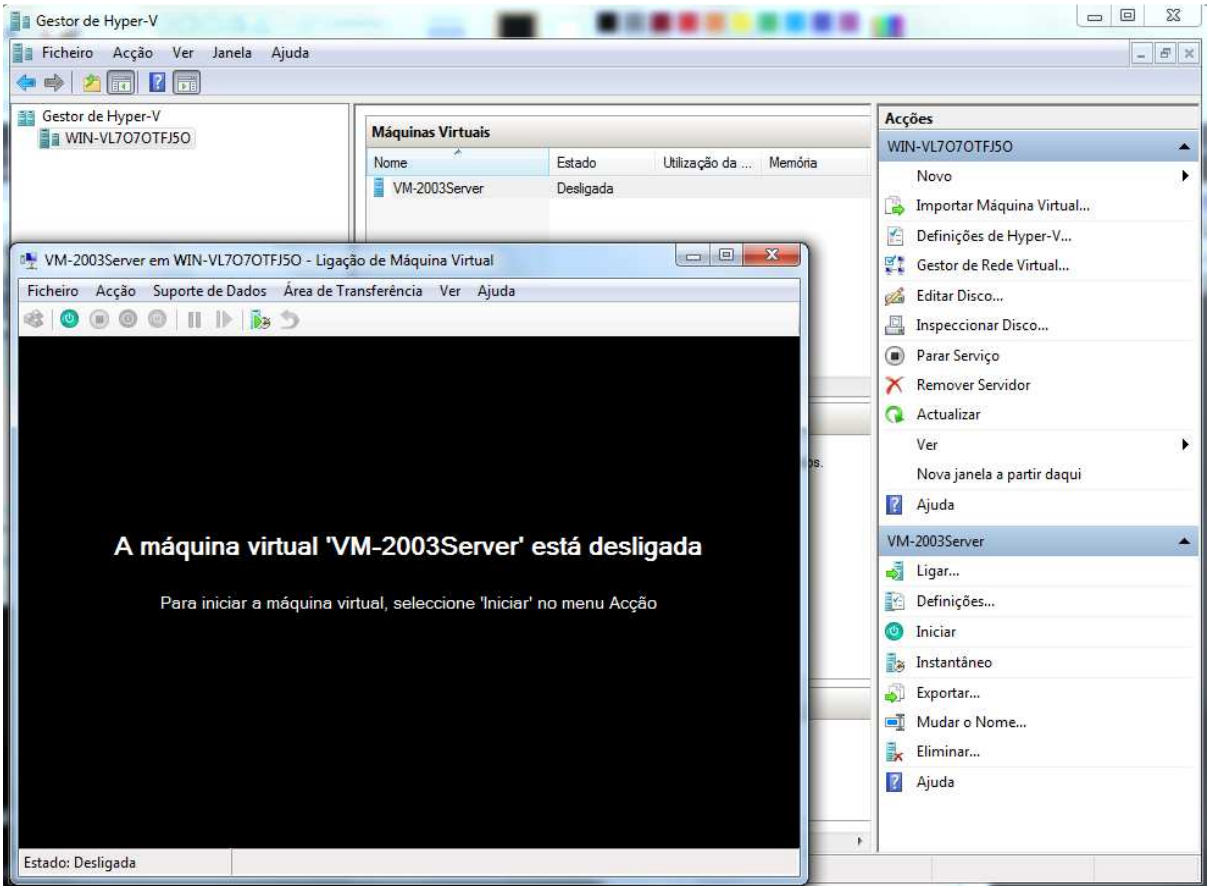

Figure 12 - Hyper-V management console

The steps in creating virtual machines are also identical, as it is necessary to create virtual networks, storage sites and finally the virtual machines themselves. All these steps are helped by wizards which greatly simplify any task and quickly provide a base configuration to

proceed for the creation of new virtual machines. Unlike XenServer, which warns right at the setup about the limitations of not having a VT capable processor, Hyper-V allows the whole creation process to be carried out until a virtual machine startup attempt is made. Whatever operating system was chosen, the absence of VT expressed itself once again and prevented further testing. This limitation is identified in the application requirements but, faced with the impossibility of using the best equipment for this lab, and since this issue did not apply to ESX, we decided to confirm if it was indeed possible to boot a virtual machine, even with performance degradation. Yet again, faced with this limitation, the testing battery could not be carried out.

ESX 4

This solution is, right away, the only one where the issue of lacking hardware-assisted virtualization does not imply any limitation (due to the way VMWare circumvented the issue). It should be noted that VMware offers two installation types: ESX, which is the complete version based on a Linux kernel, and ESXi, which is equipped with a reduced proprietary kernel with a smaller footprint. There are benefits in both cases. ESXi has a local interface very similar to XenServer, providing a simple screen with basic network configuration options and administration privileges management. In ESX, there is only a local service console, completely identical to a Linux one. Since the reduced version is still in its early stages, ESX is still the recommended version for more complex scenarios, especially involving clustering, so it was the selected version for this laboratory. In addition, the Linux service console is very useful for any emergency situation.

Turning to the deployment itself, it is very similar to the XenServer case. It was simple and quick to set up basic network parameters and access credentials. After some other setup procedures, the management console was opened. VWare's experience in this domain was easily noticed on this console and everything seemed well thought and tested in detail. Navigation is done by tabs, which instantly provides a settings summary and an insight into the host hardware characteristics, along with storage and networking management. There is also an option that quickly draws a schema map with connection dependencies, useful for a quick overview of what is deployed and to understand its relationships.

After setting up the network and storage, the first virtual machine was created. VMware has an online repository of Vapps (free and commercial ready to use system images) which can be immediately downloaded and imported into ESX. The images range from full installations of operating systems to electronic mail systems, network management, security tools, etc. In

testing or even production scenarios, a feature like this can save a lot of implementation time. The deployment itself boils down to six sequential procedures:

1. Configuration choice – prompts if a typical template based on the chosen operating system should be used or if the user wants to set specific configuration options.
2. Name and location – the soon to be virtual machine must be identified along with the host where it will relay on (in this case, there is only one).
3. Storage – assuming that storage space has already been created, in this place the destination data store must be specified.
4. Guest operating system – based on this choice ESX applies all the hardware specific recommendations for the chosen system.
5. Create disk – within the defined storage, a virtual disk is created in order to be used by the operating system. This option allows defining its size and choosing whether the whole disk space should be immediately assigned or only according to the system needs (thin-provisioning).
6. Ready to finish – a summary of the chosen settings is presented, which can be edited. Otherwise, the operation is confirmed.

After finishing this process, the three proposed testing services were created (Figure 13).

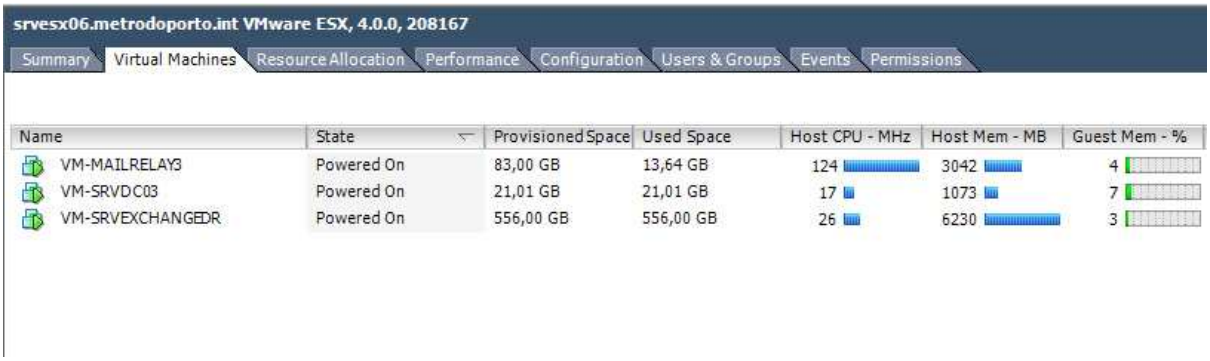

Figure 13 - Virtual machines view in ESX console

All Windows systems were root installed and the Trend Micro, based on an ISO supplied by the manufacturer, optimized for deployment in virtual environments. An important detail refers to the installation, in each guest, of an application called VMWare Tools which installs a set of optimized drivers, a clock synchronization mechanism and enables the possibility of interacting directly with the ESX operating system. This interaction allows performance data

collection, status monitoring and operations such as restarting or shutting down the guest through the operating system itself.

After a few days of operation, performing data copying and messaging usage (though obviously in a single user scenario), values were collected to classify the system's behavior according to the defined evaluation metrics.

Supported operating systems

This goal was completed right at the virtual machines installation. All the systems required for this test were easily deployed.

Performance

With the help of a tool called Networx [SoftPerfect, 2010], the network throughput was measured using a file share with an equivalent physical system by copying a 279MB file, followed by a 1.5 MB one. The result was as follows:

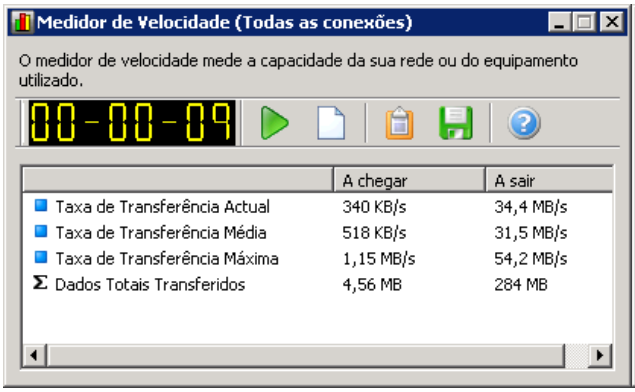

Figure 14 - Throughput for 279MB

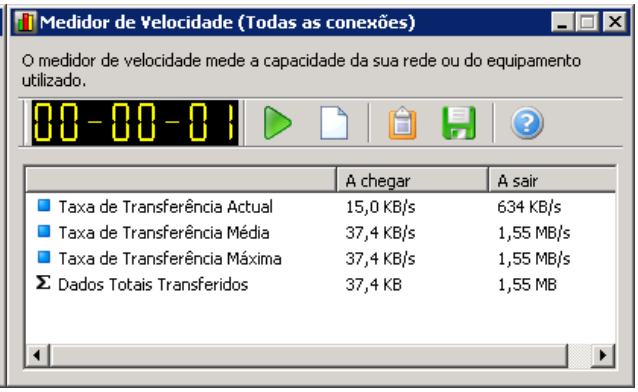

Figure 15 - Throughput for 1,5MB

As shown in Figure 14, in the first case there was an average rate of 31.5 Mbps with a peak at 54.2 Mbps. As for the second copy, the values were 1.55 MBps, both average and maximum (Figure 15). In a similar but virtual system case, the observed values were:

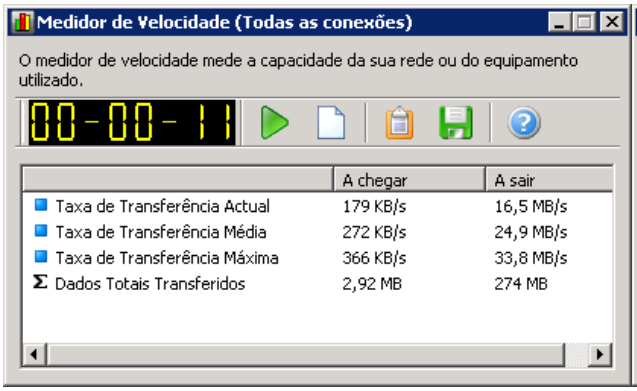

Figure 16 - Throughtput for 279MB

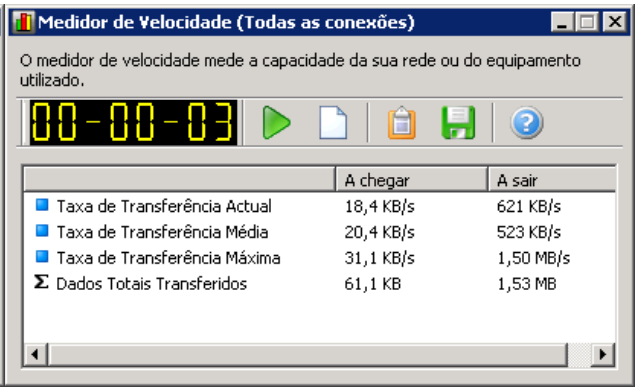

Figure 17 - Throughput for 1,5MB

In this case, it can be seen that the larger file had a 24.9 Mbps average and a peak at 33.8 MBps, which delayed the operation by two seconds (Figure 16). The difference in the average rate is not significant, bearing in mind that the server's network card used in the laboratory is shared with the virtual machines it hosts, plus the management network. In realistic scenarios (like the deployment described in the next chapter), more and better network adapters would be used. Interestingly, the time difference needed to copy the smaller file was exactly the same. The maximum speed is almost equal but in the case of the virtual machine the time required for reaching this peak is higher (Figure 17).

Using NBench [AC&NC, 2010], the disks I / O performance was measured. This tool allows setting a test file size and number of simultaneous threads for a copy simulation. Obviously it is not totally comparable to a real production scenario, but provides a comparative performance capacity overview. For the physical server, as shown in Figure 18, in a test with a 1MB file, the rate was 9.17MBps for writing and 62.50MBps for reading.

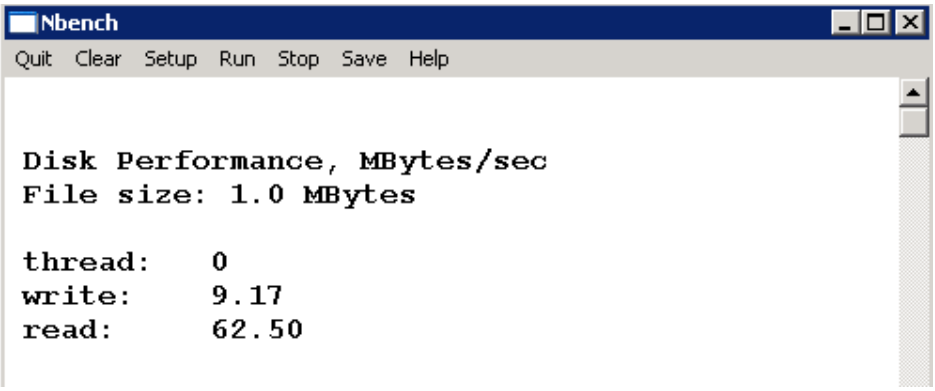

Figure 18 - I/O rate for 1MB in a physical server

The virtual server result was different. The fact that the hypervisor accessed the disk through the Virtual Machine File System (VMFS) greatly improved the write speed. The reading was similar, reaching 66.67MBps. The writing was just as fast as reading from one physical server and reached 62.50MBps (Figure 19).

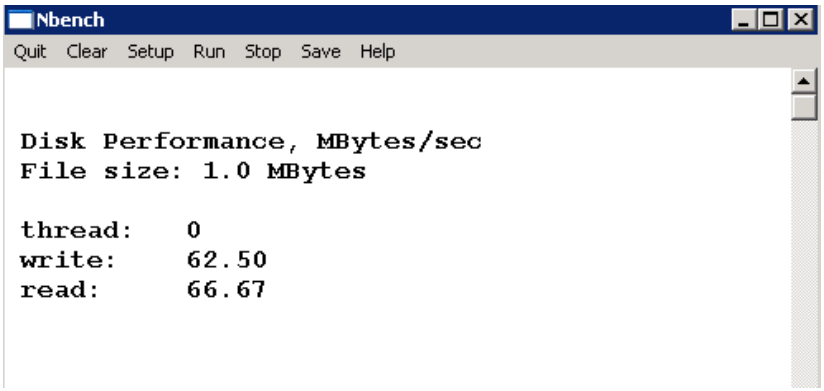

Figure 19 -I/O rate for 1MB in a virtual server

Next, a test was performed for a 200MB file. In the physical server, the performance was similar, accomplishing 12.33MBps at writing and 61.26MBps at reading (Figure 20).

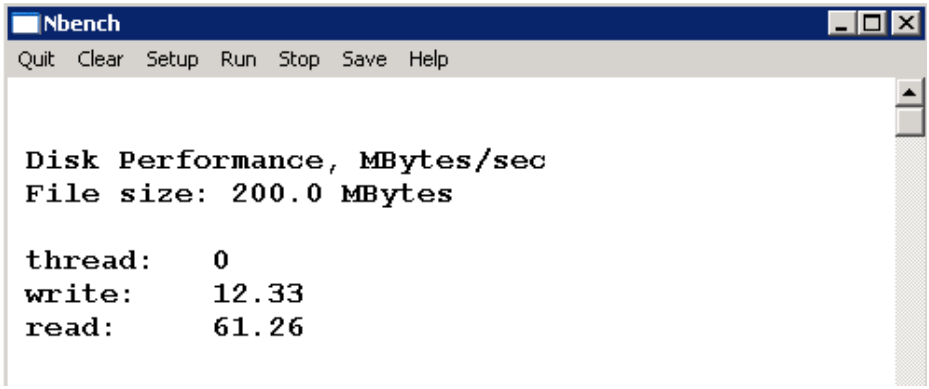

Figure 20 - I/O rate for 200MB in a physical server

The same test on the virtual server was even more surprising. The reading reached an impressive 129.29MBps and the writing peaked 44.75MBps (Figure 21).

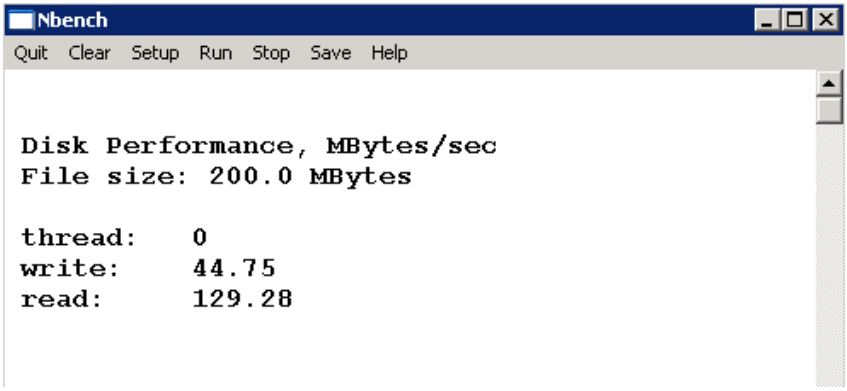

Figure 21 - I/O rate for 200MB in a virtual server

Finally, the test was performed in both using a 50MB file with three simultaneous threads, in order to evaluate the I/O performance in more demanding situations. Besides the speed differences, the virtual server produced fairly similar values in all threads. In the physical server there are significant differences, probably caused by a bottleneck in the disk access subsystem.

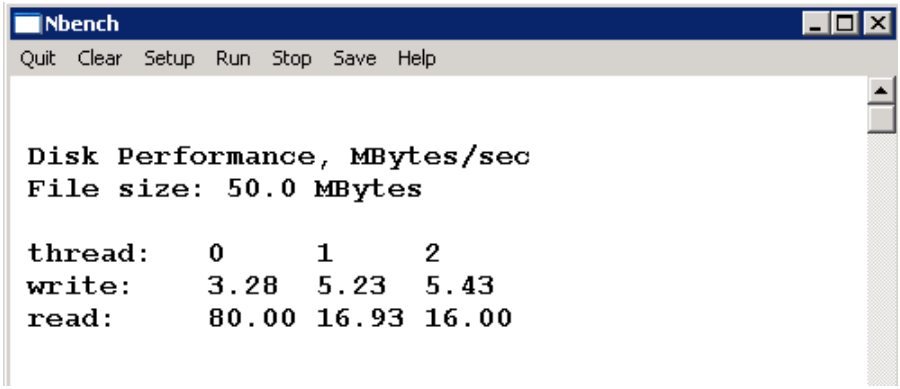

Figure 22 - I/O rate for 50MB with 3 threads in a physical server

The reading speed decline was quite pronounced in the physical server, where the first thread‘s 80MBps lowered to 16Mbps (Figure 22). The writing improved slightly in the following threads, running at an average speed of 4.64MBps. In the virtual scenario, the values were almost constant, with an average of 22.96MBps for writing and 62.74MBps in reading (Figure 23).

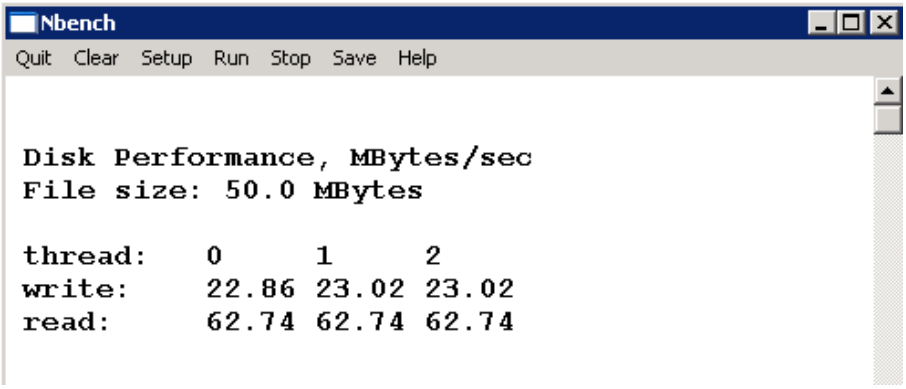

Figure 23 - I/O rate for 50MB with 3 threads in a virtual server

The physical server data is probably influenced by the disks technology and should fairly improve if a high performance SAN is available. Still, it is quite impressive the way ESX handles disk operations.

Scalability

The final tests concerned the alleged growth capacity of ESX. After some initial failed tests, we realized that the hot addition of processors and memory does not work on any operating system. Rather, the list of compatible systems is quite reduced and Windows Server 2003 Standard Edition (used in this lab) is not one of them. However, the disk addition and/or capacity increase is compatible with any system. Both situations were tested, performing smoothly and fast.

3.2.2 Implementation

3.2.2.1 Planning

The migration plan included two phases: the needs assessment and the implementation methodology definition. For the assessment, all services to be migrated into the virtualization cluster were identified, along with their current usage state, in order to properly size the required storage space and allocate processing power. Table 5 lists the result of this process.

Table 5 - Services migrating to the virtual structure

Service	Memory(GB)	Number of CPUs	Space (GB)
File share	2	2	2300
Print service	1	2	41
ERP	2	2	111
GIS	2	2	317
DC 1	1	2	35
DC 2	0,512	2	14,5
Document management – development	4	2	63
E-mail	8	4	553
Maintenance service	2	2	52
Ticketing	4	2	1000
Intranet	1	1	77
Complaints	4	1	114
Operational security	4	2	75
Webmail	2	2	32
Logical security	2	2	82
TOTAL	39,512	30	4866,5

As shown in Table 5, the required resources are much lower than the available ones, fulfilling the commitment to allow infrastructural growth in the mid-term without additional investment. The twenty four available cores may seem insufficient compared to the identified thirty, but this is one of virtualization advantages (optimization). By distributing processing as a whole for independent services, the underlying technology coordinates efforts to provide only the processing resources that each service requires at a certain moment, freeing the surplus to other services.

When defining the deployment methodology, there are several aspects to consider. In the physical services case, it must be taken in account whether the conversion to virtual is feasible[5]. Otherwise it will be necessary to perform a total reinstall. Such conversion usually runs without any problems on Windows systems and in most Linux versions, as long as it is conveniently prepared. For the second domain controller, the task was simplified because this service was already virtualized, using the free VMWare Server version, which actually served as a test platform for the technology and where some legacy services were installed. It would then be sufficient to convert the VMWare Server format into ESX, directly to the

[5] With a technology called Physical to Virtual (P2V), available with the VMWare Converter tool.

destination cluster and also using VMWare Converter. In short, the services / migration processes are:

Table 6 - Services / migration processes

Service	P2V	VMWare Server to ESX	Reinstall	New deployments
File Share	X			
Print Service	X			
ERP	X			
Geographical Information System	X			
Document Management - development	X			
Maintenance service	X			
DC1	X			
DC2		X		
E-mail			X	
Webmail			X	
Ticketing				X
Intranet				X
Operational security				X
Complaints				X
Logical securiy				X

Regarding the reinstalls, although these systems were P2V conversion capable, this option was chosen because they are both Microsoft Exchange 2003 systems and this intervention has been used for updating them to the 2007 version. The new deployments refer to services whose delivery was planned and approved but, due to this project, were put on hold in order to be deployed right into the new infrastructure.

3.2.2.2 Execution

After identifying all the necessary migration details and having their methodology defined, the deployment was based in three sequential steps: installing the base technology, performing the system/storage volumes setup and importing services. These are three logical steps that start with the assumption that all physical architecture is installed, including servers, SAN and network connections, with all redundancies properly implemented, as shown in the following section.

As verified in the laboratory tests, installing ESX is simple and fast. Once completed in the three cluster nodes, the vSphere Client management module must also be installed, allowing

the product to be centrally managed. Although it is possible to connect to each node individually, for managing these physical resources as a whole a fourth system was also configured. This central management console can be installed on any system, including a virtual machine inside one of cluster's nodes, but for failure prevention it was advisable to set it outside the cluster. In this case the console was configured in a physical machine, totally independent from the virtualization central unit. Before setting up the cluster, storage volumes must be configured in the SAN according to the provisioning requirements. VMWare good practices recommend that no more than 16 virtual machines share the same volume (defined as data store), even though this provisioning depends on the system purpose, data type and required I/O capabilities. In this provisioning, it is necessary to consider that one of the virtualization advantages is the possibility to take system snapshots. For this reason, each presented volume should have an exceeding size of 30%, as recommended by the manufacturer. In this deployment the decision was to comply with the recommended margin, leading to the volume attribution matrix specified in Table 7.

Due to each system large dimension, sharing volumes would result in the creation of huge data stores, so it seemed more appropriate creating individual ones. In the file server case, this size issue is even more important and, in order to simplify maintenance, monitoring and eventual migration tasks, it was decided to create separate volumes for each department. This generated a large number of volumes and made data stores visualization in the management console slightly more complex.

Table 7 - Volume allocation per system

Service	Volume (GB)	Margin 30%(GB)	Total (GB)
Ticketing	1000,00	300,00	1300,00
DC 1	33,00	9,90	42,90
DC 2	14,00	4,20	18,20
E-mail	545,00	163,50	708,50
Document management – development	59,00	17,70	76,70
Print service	40,00	12,00	52,00
ERP	109,00	32,70	141,70
Maintenance service	50,00	15,00	65,00
File share – DAF	168,00	50,40	218,40
File share – DEX	145,00	43,50	188,50
File share – DIN	393,00	117,90	510,90
File share – DST	94,00	28,20	122,20
File share – GCI	158,00	47,40	205,40
File share – GJU	90,00	27,00	117,00
File share – GPC	299,00	89,70	388,70
File share – GPR	380,00	114,00	494,00
File share – GSI	102,00	30,60	132,60
File share – MetroGeral	97,00	29,10	126,10
File share – Orgãos Sociais	48,00	14,40	62,40
File share – Outros	244,00	73,20	317,20
File share – Operating system	33,00	9,90	42,90
Intranet	76,00	22,80	98,80
Complaints	110,00	33,00	143,00
Logical security	80,00	24,00	104,00
Operational security	71,00	21,30	92,30
GIS	275,00	82,50	357,50
Webmail	30,00	9,00	39,00
TOTAL	**4743,00**	**1422,90**	**6165,90**

This step finished, the rest of the configuration was performed in the management console. The application licensing was applied, then created the allocated volumes and next added the cluster nodes. This way, the management became centralized and the high availability module was installed. Finally, the networking architecture was defined. The blade center platform used in this deployment was equipped with four network interface cards and, complying with security recommendations, two of them were exclusively dedicated to management network, which included the service console. It is in this network that inter-virtual machine traffic occurs and tasks like VMotion are executed, hence the need for high

isolation. The remaining network cards were assigned to normal traffic, in load balancing mode, directly connected to the physical network equipment. In order to prevent problems, each pair of network cards was configured as an alternative to the other pair, improving global fault tolerance. As this is a multi-service network, communication was divided into Virtual Local Area Networks (VLANs) therefore requiring their identification in the virtual switches so that each service can run on the intended network. The following figure (Figure 24) shows the network configuration described above.

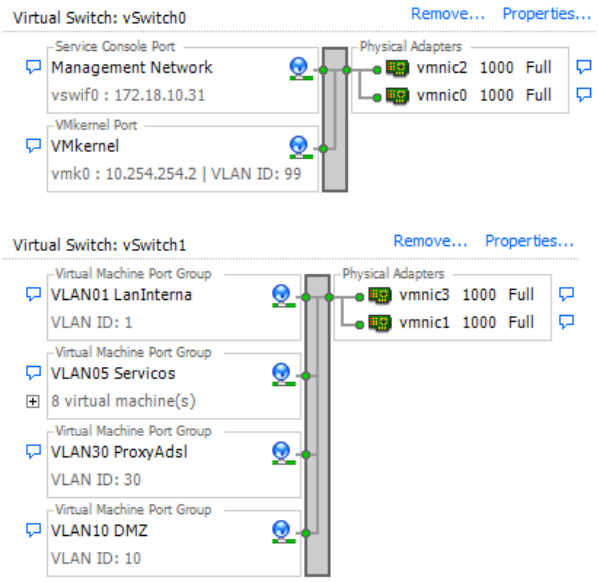

Figure 24 - Network configuration in ESX

As it can be seen, the production network was spread over four separate VLANs. The VLAN01 - LanInterna represents the client's network and is necessary because domain controller 1 hosts the Dynamic Host Configuration Protocol (DHCP) service. The VLAN10 - DMZ is the demilitarized zone network for services, in this case Webmail. The VLAN30 is a network segment specifically created for delivering packets to the proxy server and connecting customers to the Internet via ADSL. Finally, VLAN05 - Servicos is the main network thread, where all remaining services are located. Once this process was complete, the system was ready for receiving virtual machines. The creation process is a rather simple step and, since it was already described in the laboratory tests section, only the P2V conversion procedure will be addressed.

Using the VMWare Converter Standalone tool, two conversion types were required: physical to virtual and VMWare Server to ESX. Right at the application's first screen, the source of this conversion was set (Figure 25).

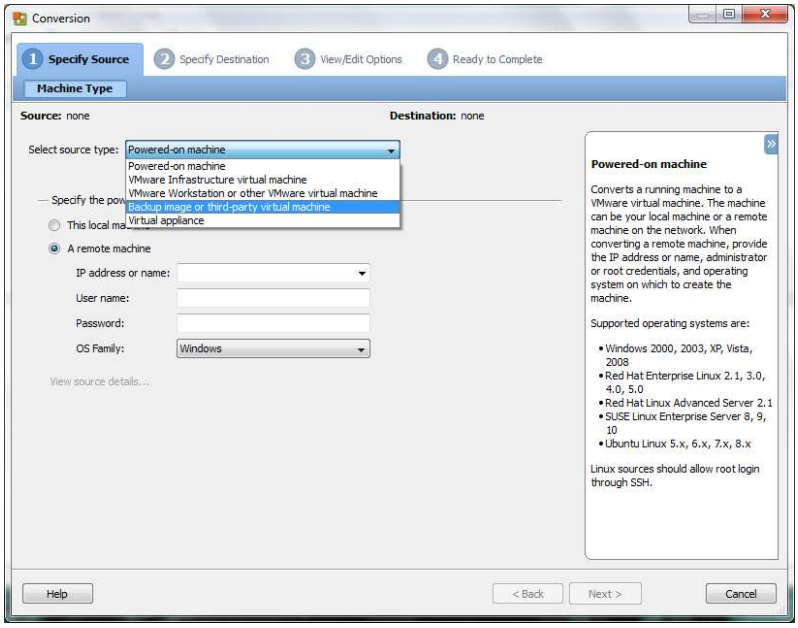

Figure 25 - VMWare Converter, step 1

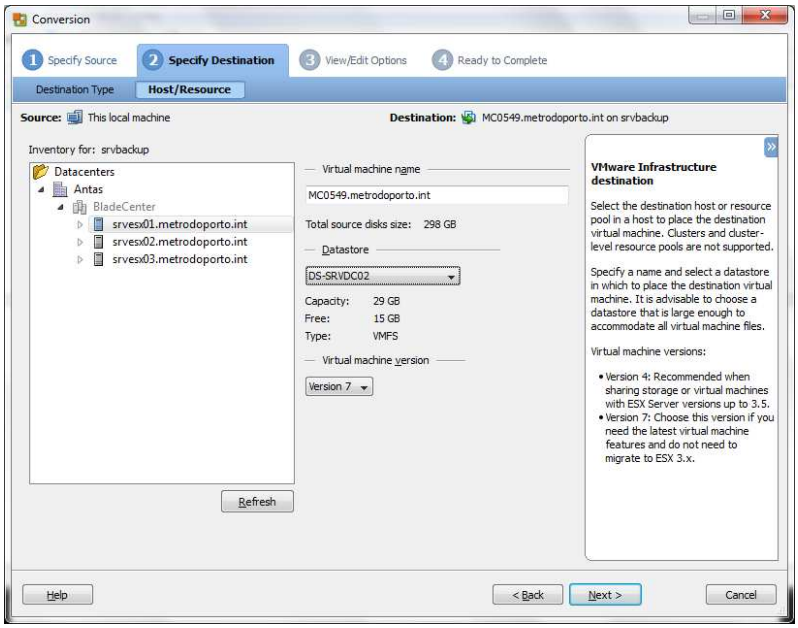

Figure 26 - VMWare Converter, step 2

Before completing the operation, a series of virtual machine configuration options can be setup right at the conversion. In a further step, the soon to be implemented virtual hardware was listed. At this point it was possible to define on which network would the machine reside and, among others, perform memory and processor adjustments. The conversion process itself was quite fast, although this depends on the data size. Upon completion, the service was available on the ESX console and ready to take off. The conversion should be performed without any information update on the source machine, so that data integrity is not compromised. Once completed the new services migration and creation, the solution was essentially ready to go into production.

There were two advanced options left to configure, which are crucial to achieve the desired level of service. One of them was the Distributed Resource Scheduler (DRS), which represents the capability to automatically distribute virtual machines within the cluster nodes, aiming to obtain the best performance for each one and allowing the available resources effective management. This option was divided into three working modes, with a sensitivity meter. These three options are: manual, partially automated and fully automated. In manual mode, DRS generates recommendations and informs the system administrator, which must manually move virtual machines in case he agrees. The partially automated mode distributes virtual machines automatically when they are first powered on. After that, only recommendations are generated. As the name suggests, the fully automated mode does all the work, without need for any kind of approval. This sounds like the best option, since this functionality was already strongly tested with good results in all ESX versions. With this mode, no permanent human monitoring of the recommendations is required. Whatever mode is chosen, the sensitivity meter allows controlling, from a conservative to an aggressive level, which priority degree must be met in order to perform a node change decision. These priorities are classified from 1 to 5 to indicate which improvement level can be obtained with a node change. If virtual machines are expected to run in the most available node, even so it does not make sense that node change immediately occurs after a momentary consumption increase. For that reason, this meter was selected at the middle position, where only priorities from 1 to 3 are effectively processed (Figure 27).

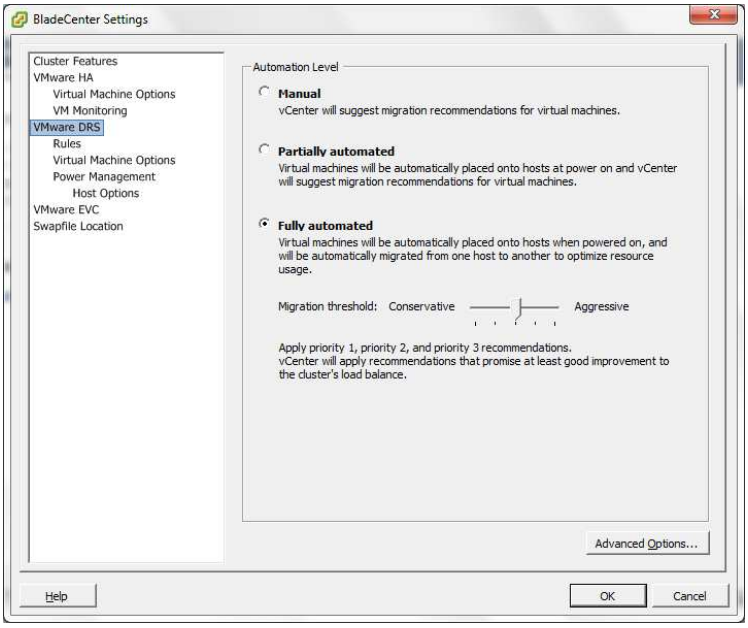

Figure 27 – DRS settings

It is also possible to define rules that must prevail during a change. For instance, it is very important to prevent the junction of machines with a very high processing load. By means of a rule, it was decided that the email and file sharing services could not share the same node. A similar rule was implemented to prevent domain controllers from co-existing on the same host so that, in case of failure and during the seconds in which high availability moves virtual machines between nodes, there is always a fully available domain controller.

Finally, the alarming system was the last to be setup. ESX is equipped with a set of predefined events and alarms for the most common situations so, for a classic services set like this one, the manual creation of additional alarms was unnecessary. As it will be discussed with further detail in the "Management Methodology" section, a complete control and alarming tool already existed, so the ESX alarming module for operating system and service availability events was not used. In this case, only hardware and virtualization warnings have been setup, the ones that may not be checked by the other tool. This context includes events like:

- Host temperature
- Hardware anomalies (fans, memory modules, etc)
- Excessive CPU and/or memory usage by a virtual machine
- Loss of connectivity/redundancy with the SAN
- Connection lost with one or more cluster nodes

These events are classified by the system as “Normal”, “Warning” or “Critical” and their corresponding alarms can be defined in order to perform different actions according to the event severity (Figure 28). In this scenario the cluster becomes the heart of the whole infrastructure so a tight control was desirable. For this reason, every state change, including the return to a normal state, was configured to generate an e-mail.

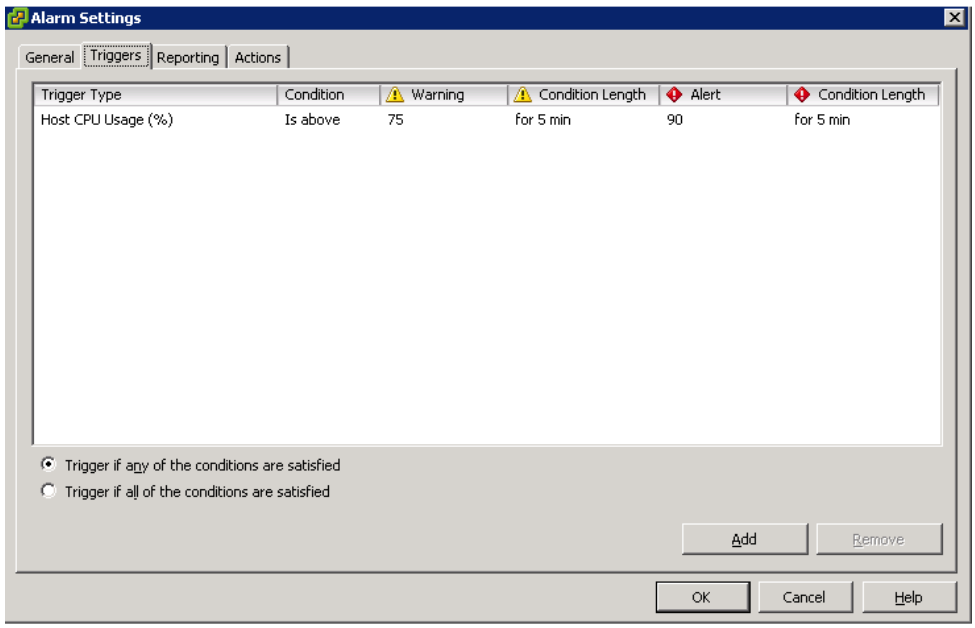

Figure 28 – Alarm triggers setup

According to the base plan, an e-mail alert will be sent to the management team for every state change. Note that it is possible to repeat the action every x minutes until the alarm is cleared (Figure 29). This feature has not been setup because it was not considered important as the management team is aware of all the alarms relevance.

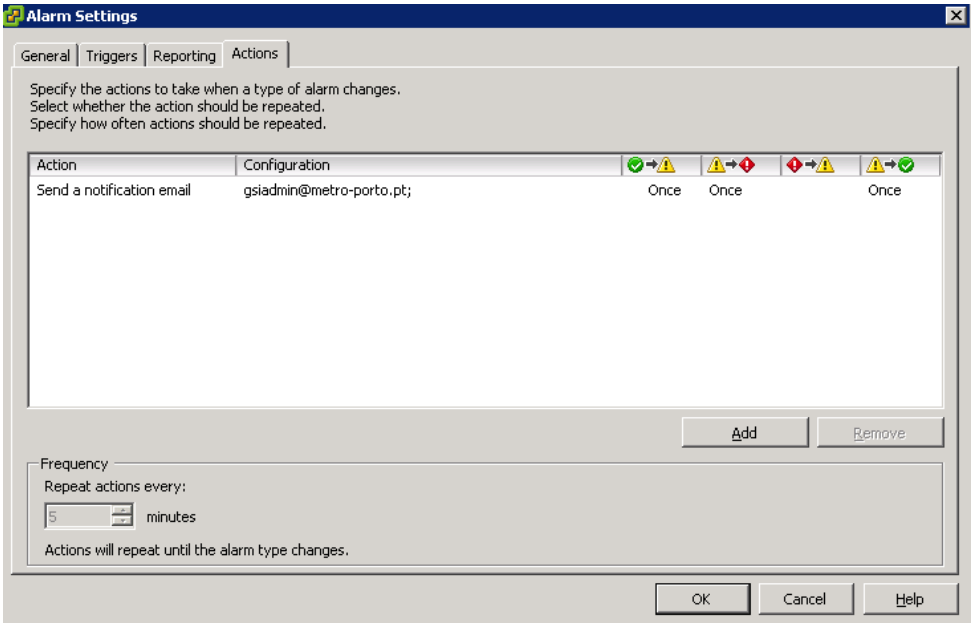

Figure 29 – Alarm actions setup

Before considering the cluster fully implemented, there was another security recommendation to account for. It is highly advisable to use the ESX native functionality that restricts the resources used by each virtual machine, in order to prevent extraordinary situations that can affect other platforms. This process requires some learning time, so the most effective option was to define appropriate limits to each service characteristics. With virtualization running, it will be easier to understand each system needs and thus adjust restrictions for achieving the best performance without compromising the remaining systems.

After these steps, the transition from the classical model to the proposed infrastructure was complete. This kind of migration process, when performed with a comprehensive study of needs and with a properly defined and planned deployment scheme, is expected to run as smoothly like in this case study, where no abnormal situation was reported and the commitment to be done without user impact was fully achieved.

Starting production was the next step, including a close monitoring of the IT team in order to verify each service correctness and configuration suitability. After reaching the “optimal” configuration, the implementation was completed and outcome indicators could be collected, allowing a comparative analysis of results.

3.3 Testing and deployment

To support this technology, it was decided to purchase a cluster of three servers with VMWare ESX and the high availability module, based on a HP Blade Center platform and a SAN HP EVA 4400 storage equipment (Figure 30).

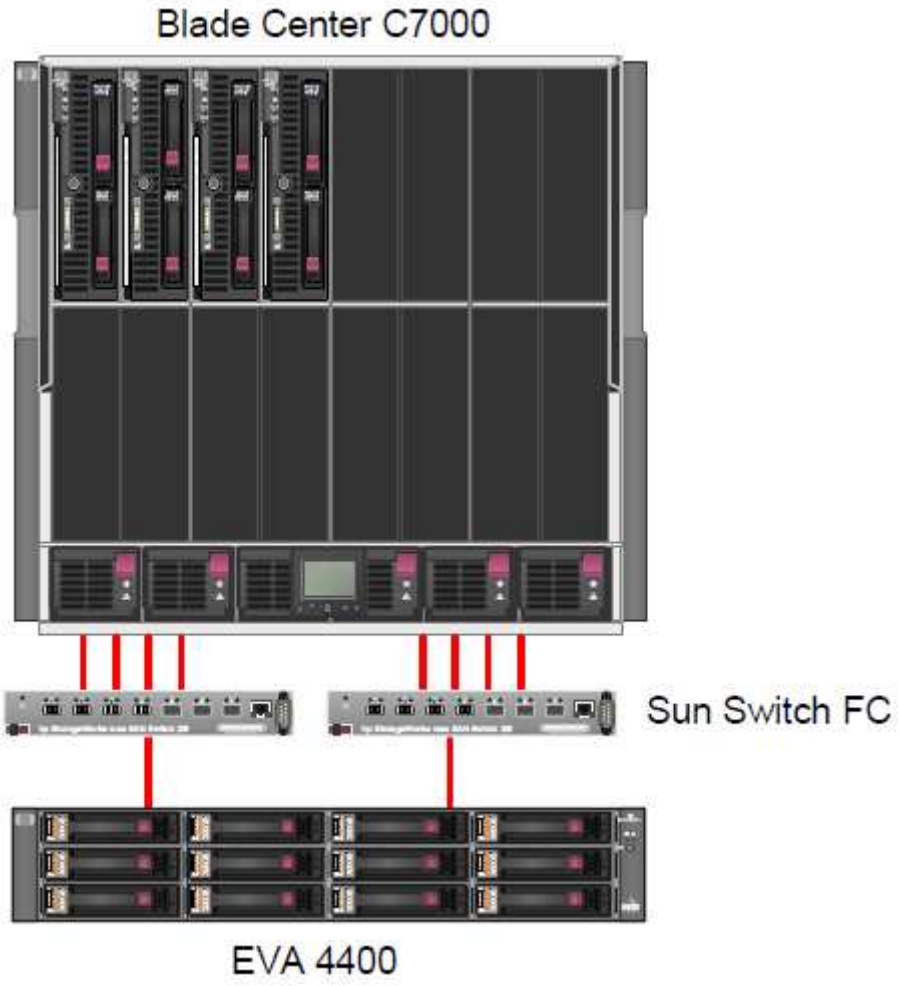

Figure 30 – Physical architecture

This SAN is composed of four drawers with twelve 300GB drives, totaling 12TB of raw space. This capacity was estimated to be able to store all systems and their data and provide growth capability for a year without adding disks. It should be noted that the company, due to their business area (and the absence of content filtering and usage limits policies) has atypical storage needs, with an average quarterly growth of about 11%. At the processing level, the blade servers were equipped with dual quad-core Intel Xeon 2.83 GHz processors and 32GB of memory. In terms of cluster, a total of 128GB of memory and 24 processing cores were made available. Since this central unit will accommodate 15 servers

with minimal standard requirements, these features will satisfy the goal of providing growth capabilities in the medium term without any investment.

3.3.1 Outcome assessment indicators

The process of evaluating outcomes consisted in direct comparison between values verified before and after the infrastructure conversion. The gathering of such data required that the capture had been carried out in an acceptable time period for an effective comparison. In this case, figures were available for the classical model since the first half of 2007 and it was easy to compute an average for a large time period. For the virtual infrastructure the time period under review was just the first quarter of 2010. Although the implementation occurred before, this was the period when actually all the adjustments were considered as being in a stable point. Using the reporting module application the company has ready for monitoring systems and services, data could be exported into spreadsheet format for a given time period. These values, among others, include the response time and service status, which allows evaluating the availability in a period of time. It was based on these sheets that annual average estimates were generated and subsequently converted into reports. This is a current practice in the institution, where a memo with this kind of data is sent semiannually to the top administration. In order to obtain a direct comparison, only the existing services in the previous structure were included in this analysis, because for all new services already placed on the virtual platform there was no comparative data available. The document management development service was left out because, although it existed before the virtual infrastructure, it was only put into production in the previous semester and consequently sample data was scarce. Once this process defined, the data collection returned the following results.

3.3.1.1 Before virtualization

In 2007, the first full year of the Office of Information Systems operation, the average availability rate of 99.95% was achieved, corresponding to a total of 1836 minutes out of service, as shown in Table 8.

Table 8 - Availability data for 2007

Servers - January01, 2007 00:00:00 - December 31, 2007 00:00:00		
Service	**Unavailability time (minute)**	**Availability (%)**
DC 1	1253,15	99,73%
DC 2	80,1	99,98%
E-mail	30,02	99,99%
ERP	40,01	99,99%
Maintenance service	20,44	99,98%
File share	80,04	99,98%
Print service	50,63	99,99%
GIS	220,92	99,95%
Webmail	60,36	99,99%
	1835,67	**99,95%**

In the following year, as a consequence of technological changes and unexpected equipment failures, the average declined to 99.83%, totaling 6503 minutes of accumulated unavailability (Table 9).

Table 9 – Availability data for 2008

Servers - January 01, 2008 00:00:00 - December 31, 2008 00:00:00		
Service	**Unavailability time (minutes)**	**Availability (%)**
DC 1	352,54	99,92
DC 2	810,82	99,81
E-mail	1456,91	99,67
ERP	472,42	99,89
Maintenance service	80,05	99,98
File share	80,06	99,98
Print service	133,89	99,97
GIS	1558,21	99,64
Webmail	1558,07	99,64
	6502,97	**99,83**

In 2009, the year infrastructure was converted, there was an improvement over the previous year, yet still below desired levels. The average availability rate was 99.90%, which corresponds to 4420 accumulated minutes out of service (Table 10).

Table 10 – Availability data for 2009

Servers - January 01, 2009 00:00:00 - December 31, 2009 23:30:00		
Service	**Unavailability time (minutes)**	**Availability (%)**
DC 1	270,81	99,95
DC 2	250,75	99,95
E-mail	601,19	99,88
ERP	410,89	99,92
Maintenance service	711,74	99,86
File share	790,98	99,84
Print service	570,89	99,89
GIS	180,84	99,96
Webmail	631,71	99,87
	4419,8	**99,90**

3.3.1.2 After virtualization

As stated before, the values for this scenario were still premature for a medium term comparison. However, the data was extrapolated to the rest of the year. In the first quarter of 2010, the much sought rate of 99.99% was achieved, resulting in less than 153 total minutes without availability (Table 11).

Table 11 – Availability data for 2010

Servers- January 01, 2010 00:00:00 - March 31, 2010 23:30:00		
Service	**Unavailability time (minutes)**	**Availability (%)**
File share	70,4	99,95
Print service	10,01	99,99
ERP	0	100,00
DC 1	0,33	100,00
DC 2	10,34	99,99
E-mail	20,33	99,98
Maintenance service	10,34	99,99
GIS	20,34	99,98
Webmail	10,33	99,99
	152,42	**99,99**

These results are only possible thanks to the capability to perform maintenance tasks, without any downtime window, on a cluster that can migrate and change the services characteristics without interruption. It should be also noted that no subsequent purchase of equipment was made and five new services were added, with more in line for entering production. Nine physical servers were discontinued, with their corresponding maintenance

contracts, and data center power consumption was reduced by 42%. With this type of solution the IT management also acquired a new consolidated way of dealing with the infrastructure. Most maintenance tasks can now be performed in work time, without the need for hard and stressful overtime interventions, sometimes just for small operations. Every time one of the nodes is placed in maintenance mode (or when it actually fails), ESX transfers that node's virtual machines to other nodes, informing the high availability module that there is an unavailable node, in order to reorganize all the resource provisioning plan. All of this is made without service rupture. It is a solution like this that allows complying with an official Service Level Agreement (SLA) without requiring an incredibly high budget.

3.4 Operation

3.4.1 Management methodology

If all solution's technical and technological components are key factors for the project's success, of equal or greater importance is its management methodology. This is a critical process, either in the paradigmatic approach as well as in responsibilities and human skills performance. It is essential to clearly identify action responsibilities, designating who takes decisions, especially in contingency situations, and who coordinates functions and tasks designed to ensure smooth and continuous system operation [Serrano and Jardim, 2007]. This is mainly an internal organizational process but it must be planned with utmost rigor and meticulous selection of human resources. These resources should not only be endowed with excellent technical ability but also have good administrative and leadership skills, as critical infrastructure management requires above than average attributes (personal, technical, etc). The goal lies primarily in the approach, where the typical approach to systems management is reactive, meaning that problems are addressed only after they happen. This means that the applied resolution may not be the best but the fastest one. In addition, the required methods for such a fast resolution destroy evidences of what happened, leaving no trail for an assessment aimed to prevent future occurrence.

From the users perspective, or customers, which is probably the most appropriate word to describe those whom the services are intended for [Van Haren Publishing, 2007], as much as they can appreciate the IT team for speedy and effective interventions, that does not prevent them from acknowledging a failure, typically causing two behavioral phenomena. If one causes an insecurity feeling regarding the system operation and consequently a concern with data and services availability, others often observe it as a relationship of

unrelated situations. For example, it is usual the extrapolation of a simple local problem, like a text document that refuses to open possibly due to lack of available resources on the client computer, into a much larger problem, "associated" to a service failure in the near past. Conjectures like these are real and contribute negatively to the necessary sense of stability that one must convey to the users.

First, a clear definition of services to be provided is required, literally creating a catalog describing services with their features and usage conditions [newScale, 2005]. This is a useful document for both management and clients, because it allows a "distant" view over the systems, but also identifies quality, availability and usage parameters. The assignment of these values must therefore be realistic and viable, since they represent a commitment to comply with between who provides the service and its customers, a vital factor in building loyalty and stable relationships. This catalog[6] represents the starting point for the correct implementation of a quality-oriented management policy [Van Haren Publishing, 2007]. Also based in the Information Technology Infrastructure Library version 3 (ITIL v3), it is advisable to develop two catalog types: the corporate, customer-oriented, and the technical, suited to the IT management team. The latter should allow a rapid and clear identification of each service technological dependencies, physical and logical. In the limit, this document should fully describe each service with a depth that includes maintenance / operation contracts identification and underlying contract templates [Leopoldi, 2002]. For the practical case study a corporate service catalogue was developed, provided as Appendix I. Due to the size and detail level required by rigorously prepared technical catalog, the developed version only identifies, in a simplistic way, the technical dependencies of critical services that ensure the company's business continuity, without going into detail, which could risk confidential information exposure. This document is provided as Appendix II.

3.4.1.1 Reactive and pro-active monitoring

In a large deployment, with a lot of care in creating redundant and high availability systems, the key factor to ensure a proper management is monitoring. However, it is necessary to understand that monitoring by itself does not provide any added value. It is essential to clearly define what is to be monitored, at what depth, how often, who should be alerted and which corrective action should be taken. There is a large range of possible failure points that, without additional tools support, is impossible for the IT management team to predict and correct.

[6] The choice of the word catalog for this document stems precisely from the service level management section of the international best practices standard Information Technology Infrastructure Library version 3.

A problems list follows to clarify this issue:

- Equipments and data center temperature above normal
- CPU or hard drive usage above normal
- Unusual increase in bandwidth usage
- Eventual mechanical failures (fans, memory modules, etc.)
- Hard disk in degraded status
- Loss of internet connectivity
- Reduction on available storage capacity

The early detection of problems allows corrective actions to be taken before they actually become real problems, reducing or eliminating unnecessary production stop situations which have a bad impact on the availability metric and may even violate customer's service level agreements [Serrano and Jardim, 2007]. It may be excessive to create a human resources team for this purpose, working 24 hours a day, subjected to human error and to a large diversity of alarm causes, therefore implying specialized hardware and software purchase. The implementation of monitoring and alarming solutions that meets these requirements will be addressed by the service catalogs described earlier. Additionally, low and high level control points are identified, such as disk space and usage rates, to define the state of operating system processes essential for an adequate service functioning. For this type of control, there are various protocols (ICMP, SNMP, WMI, etc.) which provide state data to the monitoring management application. Based on that data, it is up to the monitoring management application to consolidate this information, filter it by severity and then take actions. These actions may be like issuing a warning or even performing an intervention, using the referred protocols. All these combinations should be setup with maximum care and in accordance to services needs, bearing in mind that such a system requires some learning and refining time until it reaches the running state. Besides this component, a professional monitoring tool will also allow the production of availability and intervention time statistics, generating evidences on the fulfillment of service levels, both for clients and management. It is based on evidences that management can review the most error-prone services and define ways of optimizing them, as well as conducting internal performance assessments and validating the services correct deployment. Still, an additional relevant question is left unanswered: how is the monitoring system monitored? Having the whole system under the control of this monitoring component, it is critical that it is in constant operation or a false sense of security may appear. As such, it should not be monitored by itself neither by another automatic system, as this may trigger a monitoring over monitoring snowball effect.

This indeed has to be under human supervision, ensuring that all necessary resources are available, preferably with some sort of services state polling.

3.4.1.2 Alarms and resolutions actions

Based on the previous survey it is possible to define an alarm creation policy. Yet again, it is critical to properly identify the type and means of support for sending alarms, as well as defining the recipients. A clear assignment of responsibilities is needed for each service, as it relies precisely in each recipient the power to make a decision and trigger corrective measures. Until now, the monitoring methodology emphasis has resided primarily in technology but the human factor is a crucial factor. There are two critical issues for reaching success. Primarily, each monitored service responsible person must have knowledge and experience in the organization. A thorough understanding of not only the IT component but also the organization working habits will allow a quick and practical analysis, leading to higher efficiency in solving problems. If an unexpected increase in a file sharing server processing is observed, generating an alarm, an IT well-prepared manager can quickly understand the source and deal with the issue faster and with more appropriate corrective actions. This performance breakdown could come from one organization area that occasionally causes high increase in server requests, due to occasional extraordinary events. This knowledge will allow the service manager to quickly realize that this may be the cause and then release this server from less important processes or temporarily increase memory and processing, looking for the desired performance. Secondarily, a thorough understanding of the quality oriented policy is required. If the responsible manager does not have a proactive management attitude, he can potentially jeopardize the entire system. The proactive attitude requires that any alarm is not ignored and neither its resolution delayed. A common example is the reception of a storage space shortage alarm in a particular service. If an alarm was set for this, it means that a corrective measure must be taken, as further usage may cause an outage situation.

Another key factor in setting alarms concerns predicting timing for administrative resolution processes. In most cases, the best method for defining the values for which an alarm is to be triggered requires a constant study of growth rates and services usage, allowing a fairly accurate forecast of supply requirements and perception of what is considered above normal usage. For applying these methods in the practical case study, the table provided as Appendix III was designed, which lists precisely the critical services identified in the technical catalog, with its monitoring, alarm and action type definitions. The chosen alarm methods

are mainly based in sending e-mail notifications, yet an SMS dispatch is contemplated for systems whose criticality requires an immediate acknowledgement or in cases where it is not feasible using e-mail.

3.4.1.3 The hidden side of transparent management

As in most real situations, there is a downside in adopting a methodology that seeks to be transparent. The common user generally lacks sensitivity and knowledge to understand the operational and management complexity of large systems, along with the corresponding responsibilities it entails. Since it is a goal to smoothly develop all infrastructure phases without clients noticing it, for them there is a general impression of inaction on the part of IT management. Instead, they tend to think that IT management only solves sporadic situations with direct impact on the customer. However, the management (from the moment the services it provides are crucial for doing business, generating income and efficient organization production) has a criticality of equal importance to any other organizational area [Nascimento, 2006].

To acknowledge this is decisive, since it awakens the top management for the need of permanent support and for awareness to the required means to support a correct and tuned operation. There are several ways to overcome this situation, depending on each organization's personality. Once capturing the initial attention, it is essential to deliver periodic status and requirements reports, emphasizing the consequences and production impacts due to the underperformance of essential IT activities. In addition, demonstrations, documentation provisioning and public projects presentations also provide a broad/deep view of developed work and explain the benefits for the organization, reinforcing the point that IT exists to support the organization and not otherwise.

3.4.2 Business continuity

Although this issue is not one of the major goals in this document, it is important to address it and to have in mind the important principles in establishing a business continuity plan. From the moment that IT becomes an active and vital part of the organization's production process, this plan existence also becomes crucial. Besides the typical failure situations, there are far more catastrophic scenarios whose probability of occurrence, low as it may be, poses a risk and, as such, should have its own action and recovery plan. Except for very dramatic or extreme events, in all other disaster related events the organization will continue to exist. Even so, it is impracticable for its survival to contemplate a total recovery scenario

[Serrano and Jardim, 2007]. For addressing this need there are several studies and methodologies that address the most effective ways of creating scenarios (known as Disaster Recovery).

Assuming that each organization has its own personality, there is not a general recipe for implementing disaster recovery, but only a series of guidelines, because recovery complexity may even exceed the one of the main infrastructure. In the creation of disaster recovery scenarios there are essentially two desirable protection aspects: data and services. For the first one, which should be the solution starting point, data safeguard is the main concern so that its recovery is possible for any event. This method is typically implemented in an offsite context, providing backup copies storages in places other than the organization premises. This requires the definition of an acceptable disaster data loss period, which is the periodicity in which offsite copies are performed and stored. This solution will ensure right away that, for an extreme event (even with full installation of the entire technological infrastructure) there is a recent timeframe from which the organization may continue to work after restoration. For services protection, the complexity level increases considerably but benefits typically dominate. In this approach, a replica of the primary data center, besides data, is implemented at another geographically distant location (usually more than 150km [U.S Securities and Exchange Commission, 2002]). This distance may vary according to several factors. In places where a higher likelihood for natural disasters, a much larger separation between primary and secondary data center is required, aiming to ensure that a disaster will not destroy both. A secondary infrastructure is implemented for continuing or replacing essential services within a short time delay, so that, in case of stoppage, production can resume as soon as possible.

The acceptable staging time definition requires a cost analysis that encompasses the organization not just at financial level, but also at other levels, some intangible, like external image and lost business opportunities [Serrano and Jardim, 2007]. It is therefore necessary to define the ratio between the solution's complexity and acceptable stoppage time, in order to find a convergence point for the solution to be built upon. Naturally, this ratio is also valid for the financial analysis, as greater is the urgency of restoring services, also greater is the required budget to meet that goal. This is not just a hardware, software and information replication, it is also everything required for services operation such as facilities, communications and telecommunications. In this issue, distance is also a key factor, as for short distances a copper or fiber optic LAN can be installed, but for long distances communications it will require other means, probably causing a network supplying dependency on external entities. There is also room for balance as these are systems that

will be usually idle or stopped, therefore opting for lower cost equipment is understandable, lowering the acceptable availability level for running the datacenter to allow normal, not optimal, services usage.

It is also important to bear in mind that a disaster recovery solution is not just a plan for the IT but for the entire organization. There is no purpose in restoring services if other company's resources are unavailable and not duly informed of how to proceed in such disaster situations. Their displacement to another location may be required, where should be available a functional workstation, enabling access to services for resuming production. Printers, internet access or even desks and chairs are just some examples of what may be needed for restoring a good close to normal functioning. Disaster recovery requires a detailed and specific plan for each area or function, with clear instructions on what should be their individual and collective procedures and corresponding leadership, emphasizing the role of a communication manager [Serrano and Jardim, 2007]. This will ensure that everyone is always in tune with the latest operating information.

4 Conclusion

4.1 Thesis summary

This work provides a basis for a better understanding of the sustainability problem that the IT world has to deal with. The quick everyday changes which IT professionals are subjected to make adaptation speed and ability to respond to challenges their success key factors [Yourdon, 2002]. With IT evolution showing the need to align business with information systems, large scale supplying with uninterruptible and high quality IT services becomes more critical. The focus will have to shift from technology itself to the results it produces. As Nascimento describes in his remarkable work about the changes that information system professionals have faced along the last years, the informatics department is no longer an area that independently defines and implements working methods, but instead, an area, like every other in an organization, that works in a collaborative way [Nascimento, 2006]. This business alignment is only viable if everyone is aware of the importance of defining processes together and matching technology orientation to organization's purposes. In this scenario, the IT is only a vehicle that provides these services to the upper layer (information systems), which does supply essential data for development and decision making. Supplying these services is unarguably a complex task that requires specialized professionals and high level practices, but emphasis should be put on data and not on technology. Sooner or later there will become a time where accessing a service will be like turning on a light. Everybody expects it to work, without questioning which method or by which means the energy is supplied. It is unquestionable that it exists, what is its goal and its unacceptable unavailability [Nascimento, 2006]. Just like virtualization abstracts resources and makes them available to services in a controlled manner, so does IT for the creation of an abstract cloud for which only its managers and technicians understand the working details, allowing their clients to get benefits in a correct and continuous way with a single concern: productivity.

4.2 Completed objectives

This paper is about research and analysis of virtualization methodologies for creating high availability, redundancy and optimization in information technologies. After an initial survey, a practical application with clear objectives was conducted in a real organization. In terms of results, those objectives have been achieved. Even though, availability data should be monitored for a period of no less than two years (the equipments warranty period) in order to assess the solution effectiveness, once the hardware life cycle reaches the decline phase and actual failures become more likely to occur.

The only setback identified in the test platform for the three biggest manufacturers was the absence of expensive hardware, which prevented carrying out a total comparative analysis and ultimately limited the laboratory experiments, also influencing performance data collection for all solutions.

4.3 Other performed works

In the "Management Methodology" section special emphasis was given to the standard ITIL v3, for which two documents have been produced and will serve as a starting point for creating an SLA in the underlying organization. The development of these documents, corporate services catalogue and technical services catalogue, was not included in the initial project planning. Both documents are useful for better understanding the importance of generating documentation to support the services delivery, thus contributing to objectives and allowing a high level perspective on the solution's working patterns.

4.4 Limitations

Although no specific limitation was found on this work, there is an obvious operation condition: the dependence on VMWare ESX. From the moment services start operating in this platform, there is no known tool to reverse the process or to convert it to another system. Once this path is followed, for the moment a rollback is only possible with total services reinstallation.

4.5 Future work

This is the beginning of a philosophy change, as this work is a starting point for a whole new culture of strategic organizational IT positioning. Like any other starting point, many new plans to be developed were left unfinished. The issue of continuity management, with increasing chances for natural disasters and growing exposure to logical hazards (sabotage, piracy, etc.), is a crucial factor and its development requires planning and forecasting. This is one of the most complex and difficult managerial areas, where continuous improvement does not necessarily bring real improvement, and requires continuous adaptation. It is a plan difficult to improve, as increasing risks along with steady infrastructure growth require a persistent effort just for keeping the plan updated, not leaving much room to study its evolution. It is imperative that all security measures and service levels are applicable to this plan, as the capability to produce alternative scenarios should not negatively impact the standard operational parameters. Operating in exceptional circumstances without ensuring

security and data integrity can be even more catastrophic than stopping. Still, it may not be reasonable for most organizations to create continuity scenarios, never expected to be used, with the same technology available on the production site. This is a constraint that requires high elasticity in the definition of the compromise between usability and suitability.

The issue of continuous adaptation is transversal to every IT area. Any work of this nature only sets the beginning and will not be useful without constant re-evaluation and validation of its adequacy with reality. The contexts change quickly and, if background documentation and procedures reviewing tasks are not performed, the model quickly becomes obsolete, with no practical application. At the other extreme, without this constant monitoring, it will eventually be inevitable to start the entire procedure from scratch. For many IT professionals, the administrative role is considered boring and unnecessary, but it is precisely this that settles the pillars for providing quality services. To be an IT professional today is far more than simply holding and acquiring technical knowledge. Other expertise, coupled with organizational skills, not only grants recognition but allows cultural expansion that will increasingly be a fundamental factor for personal and professional development. In the limit, it may even become a “survival factor” in the immense and competitive market where the management paradigms evolution reinforce intellectual over mechanical property as a higher value.

On the virtualization side, with the return on investment rates presented in the case studies, a general spread over almost all infrastructures is predictable. Technology is still evolving and will very rapidly reach even more abstract operation levels, resulting in a commitment to provide solutions that simply will not allow any downtime period for any intervention. With hardware development, it will possible to add more processor cores and memory on a single device, further reducing the need for physical platforms. Nowadays, where no alternative is presented with even closer consolidation and optimization values, a relevant question remains: VMware is currently the manufacturer of choice in 60% of installations worldwide. If it reaches the targeted 100%, and with it the monopoly on virtualization market, what can be expected for their business model and licensing costs? As stated in the constraints, rolling back will probably be more costly and unattainable...

To finish, once this whole concept is put into practice, why not applying for an international standards certification? The standards adoption is becoming more common and desirable, so there is available common guidance in this area. Proprietary technologies are declining and very soon the same may happen with working methods. The chapter on management methodology introduced some ITIL v3 practices that, once put into operation, allow room for

developing all levels of the standard and get public recognition for an organization. For now that is only a differentiating factor but in the future it will represent the minimum standard required for service provision. Several other standards apply, such as ISO 27001 for security and COBIT for IT management, so there are many progression paths to address in the near future. The mission is shared: delivering on time, safe, reliable and quality services. In a world full of unstoppable technological changes, this may be the only way for organizations to survive and to assure their assets...

5 References

AC&NC. 2010. Disk I/O Benchmarks – Windows Benchmarks http://www.acnc.com/04_02_02.html. (accessed February 2010)

Christensen,Kent and Alan Howitson. 2008. Data protection for virtual server environments. USA: Datalink

DataLoss DB. 2010. Data loss news, statistics and research. http://datalossdb.org/yearly_reports/dataloss-2010.pdf. (accessed April 2010)

Dittner, Rogier and David Rule. 2007. The Best Damn Server Virtualization Book Period. USA: Syngress

EMA. 2010. Securing the administration of virtualization. USA: EMA

Hau, William and Rudolph Araujo. 2007. Virtualization and Risk – Key Security Considerations for your Enterprise Architecture. USA: Foundstone

Heiser,Gernot, 2007. Virtualization for Embedded Systems. Australia: Open Kernel Labs

IT 2.0. 2009. Next generation IT infrastructures. http://www.it20.info/misc/virtualizationscomparison.htm. (accessed December 2009)

International Organization for Standardization. 2005. ISO/IEC 27001:2005 - Information technology -- Security techniques -- Information security management systems – Requirements. Switzerland: International Organization for Standardization

Leopoldi, Rick. 2002. IT Services Management – A Description of a Service Catalogue. UK: RL Consulting

Moura, António, Artur Carvalho, Carlos Magno, Germano Silva, Joana Fillol, Jorge Fiel, Jorge Morgado. 2007. A história do Metro do Porto. Portugal: Metro do Porto e Calendário

Nascimento, José Carlos. 2006. Gestão de Sistemas de Informação e os seus Profissionais. Portugal: FCA

Newscale. 2005. How to Produce an Actionable IT Service Catalogue. UK: Newscale

Oltsik, John. 2009. The new security management model. USA: Enterprise Strategy Group

Ritter, Ted. 2009. Virtualization Security – Achieving Compliance for the Virtual Infrastructure. USA: Nemertes

Serrano, António and Nuno Jardim. 2007. Disaster Recovery. Portugal: FCA

Silva, Pedro Tavares, Hugo Carvalho, Catarina Botelho Torres. 2003. Segurança dos Sistemas de Informação. Portugal: Centro Atlântico

Singh, Amit. 2004. An introduction to virtualization. USA: kernelthread.com

SoftPerfect. 2010. NetWorx – Free Bandwidth Monitoring and Usage Reporting. (accessed February 2010)

Third Brigade. 2008. Virtualization security: A Coordinated Approach for Intrusion Detection and Prevention. Canada: Third Brigade

U.S. Securities and Exchange Commission. Interagency Concept Release: Draft Interagency White Paper on Sound Practices to Strengthen the Resilience of the U.S. Financial System. http://www.sec.gov/rules/concept/34-46432.htm (Accessed August 2010)

Van Haren Publishing. 2007. ITIL v3 A pocket guide. USA: Van Haren Publishing

VMWare. 2007. Virtual networking concepts. USA: VMWare

VMWare. 2009. Network segmentation in virtualized environments. USA: VMWare

VMWare. 2009. Virtualization basics: History of virtualization. http://www.vmware.com/virtualization/history.html. (accessed December 2009)

VMWare. 2009. Technical White Papers http://www.vmware.com/resources/techresources. (accessed December 2009)

Yourdon, Edward. 2002. Byte Wars: the impact of September 11 on Information Technology. USA: Prentice Hall

Appendix I. Corporate Services Catalogue

Corporate Services Catalogue

Service	Components	Users	Availability	Supported by	Usage conditions	Support hours
E-mail						
	E-mail	All departments	24/7	GSI	No storage limits. Max. Message size : 20MB	Mon-Fri 08:30 - 19:00
	Webmail	All departments	24/7	GSI	HTTP,IMAP and POP3 access	Mon-Fri 08:30 - 19:00
	Anti-spam with quarantine management	All departments	24/7	GSI	15 days retention time	Mon-Fri 08:30 - 19:00
File Share						
	File Share - Comum	All departments;Prosegur;Garra	24/7	GSI	No storage limits	Mon-Fri 08:30 - 19:00
	FTP	Internal and external users	24/7	GSI	Physically limited to 2TB	Mon-Fri 08:30 - 19:00
	File Share	GSS	24/7	GSI	Physically limited to 4TB	Mon-Fri 08:30 - 19:00
Print						
	Print server	All departments	24/7	GSI	Default printing in black	Mon-Fri 08:30 - 19:00
	Canon equipments	All departments	24/7	Canon Copicanola		Working days 09:00 - 17:00
	Ricoh equipments	GSS	24/7	Ricoh Portugal		Working days 09:00 - 17:00
	HP equipments	GPR	24/7	GSI		Mon-Fri 08:30 - 19:00
Internet						
	Proxy	All departments	24/7	GSI	Only FTP, HTTP, HTTPS access	Mon-Fri 08:30 - 19:00
	Linha ADSL	All departments	24/7	Vodafone	8Mbps download / 1Mbps upload	24h
	Linha FWA	All departments	24/7	Vodafone	1Mbps upload e download	24h
Document Management						
	Document management - desktop and web access	All departments	Mon-Fri 08:00 - 23:00 Sat 19:00 - 24:00 Sun 24h	GSI		Mon-Fri 08:30 - 19:00
ERP						

Corporate Services Catalogue

Service	Components	Users	Availability	Supported by	Usage conditions	Support hours
	Financial module	DAF;GPC;GSI	24/7	GSI		Mon-Fri 08:30 - 19:00
	Administrative Modulo	DAF;GPC	24/7	GSI		Mon-Fri 08:30 - 19:00
	Staff module	DAF	24/7	GSI		Mon-Fri 08:30 - 19:00
GIS						
	SIG	GSI; GPR	24/7	GSI		Mon-Fri 08:30 - 19:00
	Expropriation management	GPR;GSI;GJU	24/7	GSI		Mon-Fri 08:30 - 19:00
Web services						
	Attendance management	All departments	24/7	GSI		Mon-Fri 08:30 - 19:00
	Tram-train content transfer	GSI:GCM	24/7	GSI		Mon-Fri 08:30 - 19:00
	Corporate website replica	GCM	24/7	GSI		Mon-Fri 08:30 - 19:00
Network access						
	Structured network	All departments	24/7	GSI		Mon-Fri 08:30 - 19:00
	Multi-services network	Internal and external users	24/7	GSI		Mon-Fri 08:30 - 19:00
	Wireless network	All departments	24/7	GSI		Mon-Fri 08:30 - 19:00
	Corporate TV wireless network	GSI;GCM	24/7	GSI		Mon-Fri 08:30 - 19:00
	Tram-train wireless network	GSI;GCM	24/7	GSI		Mon-Fri 08:30 - 19:00
	Guests wireless network	Internal and external users	Mon-Fri 09:00 - 20:00	GSI	No local access, just internet	Mon-Fri 08:30 - 19:00
Corporate TV						
	Content development	GCM	24/7	GSI		Mon-Fri 08:30 - 19:00
	Content transfer	GCM	24/7	GSI		Mon-Fri 08:30 - 19:00
	Contenet management	GCM	24/7	GSI		Mon-Fri 08:30 - 19:00
	Content segmentation	GCM	24/7	GSI		Mon-Fri 08:30 - 19:00
	Content viewer	Internal and external users	24/7	GSI		Mon-Fri 08:30 - 19:00

Corporate Services Catalogue

Service	Components	Users	Availability	Supported by	Usage conditions	Support hours
	On-board TV	Internal and external users	24/7	GSI	Available for all Eurotram and Tram-train vehicles with on-board TV	Mon-Fri 08:30 - 19:00
Security						
	Logical protection	All departments	24/7	GSI	127 seats	Mon-Fri 08:30 - 19:00
	Antivirus	GSI; External users	24/7	GSI	110 seats; only for servers and external users	Mon-Fri 08:30 - 19:00
	Firewall and routing	All departments	24/7	GSI		Mon-Fri 08:30 - 19:00
	Access control	All departments	24/7	GSI		Mon-Fri 08:30 - 19:00
	Video surveilance	All departments	24/7	GSI		Mon-Fri 08:30 - 19:00
	Certificates server	All departments	24/7	GSI	Certificates valid for 1 year	Mon-Fri 08:30 - 19:00
	Operating system updates	All departments	24/7	GSI	Installed every Friday at 13:00	Mon-Fri 08:30 - 19:00
	Card production	All departments; Inspection and Security	Mon-Fri 08:30 - 19:00	GSI	Must be requested by Human Resources Department	Mon-Fri 08:30 - 19:00
Remote access						
	VPN	Internal and external users	24/7	GSI	Requires signing a responsability term for external users	Mon-Fri 08:30 - 19:00
	Webmail	All departments	24/7	GSI	HTTP,IMAP and POP3 access	Mon-Fri 08:30 - 19:00
	Terminal Server	Internal and external users	24/7	GSI	Limited to 10 simultaneous accesses	Mon-Fri 08:30 - 19:00
Attendance management						
	Data collection	All departments	24/7	GSI		Mon-Fri 08:30 - 19:00
	Terminals	All departments	24/7	ELO		Working days 09:00 - 17:00
Databases						
	Security	GSS	24/7	GSI		Mon-Fri 08:30 - 19:00
	Complaints	DEX	24/7	GSI		Mon-Fri 08:30 - 19:00

Corporate Services Catalogue

Service	Components	Users	Availability	Supported by	Usage conditions	Support hours
	Case files management	DEX	24/7	GSI		Mon-Fri 08:30 - 19:00
	Maintenance	DEX;DIN;DST	24/7	GSI		Mon-Fri 08:30 - 19:00
	Document Management	All departments	Mon-Fri 08:00 - 23:00 Sat 19:00 - 24:00 Sun 24h	GSI		Mon-Fri 08:30 - 19:00
	Hardware management	GSI	24/7	GSI		Mon-Fri 08:30 - 19:00
	GIS	GSI;GPR	24/7	GSI		Mon-Fri 08:30 - 19:00
	ERP	DAF;GPC;GSI	24/7	GSI		Mon-Fri 08:30 - 19:00
	Attendance management	DAF	24/7	GSI		Mon-Fri 08:30 - 19:00
	Ticketing	All departments	24/7	GSI		Mon-Fri 08:30 - 19:00
	Taxation	GSI	24/7	GSI		Mon-Fri 08:30 - 19:00
Backup						
	Server images archive	GSI	Weekly	GSI		Mon-Fri 08:30 - 19:00
	ERP server images archive - includes data	GSI	Daily	GSI		Mon-Fri 08:30 - 19:00
	DR - offsite backup	GSI	Daily	GSI	File share, e-mail and document management	Mon-Fri 08:30 - 19:00
	Annual offsite backup in magnetic support	GSI	Annual	GSI	All data	Mon-Fri 08:30 - 19:00
	Document Management	All departments	previous 3 weeks	GSI		Mon-Fri 08:30 - 19:00
	File share - Comum	All departments	previous 3 weeks	GSI		Mon-Fri 08:30 - 19:00
	[Databases]	All departments	previous 3 weeks	GSI		Mon-Fri 08:30 - 19:00
	[GIS]	All departments	previous 3 weeks	GSI		Mon-Fri 08:30 - 19:00
	[E-mail]	All departments	previous 3 weeks	GSI		Mon-Fri 08:30 - 19:00
Users management						
	Users/groups management	All departments	Mon-Fri 08:30 - 19:00	GSI		Mon-Fri 08:30 - 19:00
	Authentication srvice	All departments	24/7	GSI		Mon-Fri 08:30 - 19:00
Productivity						
	Office applications	All departments	24/7	GSI	127 licenses	Mon-Fri 08:30 - 19:00

Corporate Services Catalogue

Service	Components	Users	Availability	Supported by	Usage conditions	Support hours
	Project management	All departments	24/7	GSI	15 licenses	Mon-Fri 08:30 - 19:00
	Document portability	DAF;GSI;GPC;DEX;DST	24/7	GSI	1 license per department	Mon-Fri 08:30 - 19:00
	CAD	All departments	24/7	GSI	10 network licenses and 10 single user licenses	Mon-Fri 08:30 - 19:00
Operating systems						
	Clients : Windows and MacOS	All departments	24/7	GSI		Mon-Fri 08:30 - 19:00
	Servers : Windows and Linux	GSI	24/7	GSI		Mon-Fri 08:30 - 19:00
Multimedia						
	Projectors	All departments	24/7	GSI	Meeting rooms	Mon-Fri 08:30 - 19:00
	Electronic boards	All departments	24/7	GSI	Meeting rooms	Mon-Fri 08:30 - 19:00
Technical support						
	Services/Applications	All departments	Mon-Fri 09:00 - 17:30	GSI		Mon-Fri 08:30 - 19:00
	Hardware	All departments	Mon-Fri 09:00 - 17:30	GSI		Mon-Fri 08:30 - 19:00
	Software	All departments	Mon-Fri 09:00 - 17:30	GSI		Mon-Fri 08:30 - 19:00

Corporate Services Catalogue

Service	Components	Users	Availability	Supported by	Usage conditions	Support hours
Legend						
GPR	Gabinete de Projectos					
GSI	Gabinete de Organização e Sistemas de Informação					
DAF	Departamento Administrativo e Financeiro					
GSS	Gabinete de Segurança					
GPC	Gabinete de Planeamento e Controlo					
DEX	Departamento de Exploração					
DST	Departamento de Sistemas Técnicos					
DIN	Departamento de Infraestruturas					
GCM	Gabinete de Comunicação					

Appendix II. Technical Services Catalogue

Technical Services Catalogue
Metro do Porto,SA

Version 1.1
Revised date 06/01/2010

Content

Introduction

This Service Catalogue documents the services delivered by IT to the business. Such a catalogue is an essential foundation for Service Level Management (SLM). SLM is the process of documenting and agreeing service targets in Service Level Agreements (SLA), then monitoring and reviewing the actual service levels against those targets. The objective is to maintain and gradually improve business-aligned IT service quality.

This catalogue describes and defines the services and the service targets. It is used as a reference for SLAs negotiated with the business. That is, SLAs will in general reference the services and service targets as defined in this document. The SLAs themselves need only specify exceptions and variations from this document.

The improvements in service quality and the reduction in service disruption that can be achieved through effective SLM can ultimately lead to significant financial savings. Less time and effort is spent by IT staff in resolving fewer failures and IT customers are able to perform their Business functions without adverse Impact. Other specific benefits from SLM include:

- IT Services are designed to meet Service Level Requirements
- improved relationships with satisfied customers
- both parties to the agreement have a clearer view of roles and responsibilities - thus avoiding potential misunderstandings or omissions
- there are specific targets to aim for and against which service quality can be measured, monitored and reported - 'if you aim at nothing, that is usually what you hit'
- IT effort is focused on those areas that the business thinks are key
- IT and customers have a clear and consistent expectation of the level of service required (i.e. everyone understands and agrees what constitutes a 'Priority One' Incident, and everyone has a consistent understanding of what response and fix times are associated with something called 'Priority One')
- service monitoring allows weak areas to be identified, so that remedial action can be taken (if there is a justifiable business case), thus improving future service quality
- service monitoring also shows where customer or user actions are causing the fault and so identify where working efficiency and/or training can be improved
- SLM underpins supplier management (and vice versa) - in cases where services are outsourced the SLAs are a key part of managing the relationship with the third-party - in other cases service monitoring allows the performance of suppliers (internal and external) to be evaluated and managed
- SLA can be used as a basis for charging or cost allocation - and helps demonstrate what value customers are receiving for their money.
- The cumulative effect should lead to a gradual improvement in service quality and an overall reduction in the cost of service provision.

Scope

The Service Catalogue lists all of the IT services currently being provided to the organisation.

With ongoing development, the Service Catalogue will describe the service targets, and details of the users and those responsible for ongoing maintenance of each service.

For the purpose of this document, a service will be defined as the following:

One or more IT Systems which enable a business process

Assumptions

- Service level commitments assume “normal” load. “Normal” implies typical average loadings over the previous monthly interval. It does not include extraordinary events such as, for example, surges in external activity triggered by a media announcement or promotion, or acquisition of new business units.
- Cost models assume that IT Services department will continue to be funded by annual budget allocation, and procurement remains under IT Services control.

Customer Responsibilities

The customer should fully respect the DMP/05/03 Ver. 1 directive on the rules for computing resources usage.

Services

Application services

ERP
Document Management
GIS
Expropriation management
Attendance management
Case files management
Complaints management
Maintenance
Hardware management
Card production
Standard desktop

Technical services

Communication
Standard desktop
Files / Printing
Desktop productivity tools
Network acccess
Remote accesss
Internet
Technical support
Threat management
Backups and archiving
Access control

Professional services

Service level management
Projects management
IT consulting
Security architecture
IT architecture
Architectural Reviews of new technology
Application enhancement

Applications maintenace
Vendor relations
Training
Service delivery
Service support
On-call support
Field support

ERP

Description

Enterprise resource management application. Provides accounting, personnel and commercial management.

Scope

Departamento Administrativo e Financeiro
Gabinete de Planeamento e Controlo
Gabinete de Organização e Sistemas de Informação

Service level

Gold

Availability	24/7
Performance	2 vCPU 2.83Ghz; 2GB RAM; LAN 1000Mbps
Capacity	35GB for OS / 75GB for DB
Continuity	Availability to be recovered within 4 hours of an outage
Service desk	Gold level incident response

Control

Service manager: GSI
Access permissions through to the approval of the soliciting department

Training

User manual available

Dependencies

Standard desktop application service

Sub-services and systems

VMWare high availability cluster with three nodes

Owner
GSI

Availability
24/7

Blade Center

Owner
GSI

Availability
24/7

SAN HP EVA 4400

Owner
GSI

Availability
24/7

Document management

Description

Document management and workflow application

Scope

All departments

Service level

Gold

Availability	Mon-Fri 08:00 - 23:00 Sat 19:00 - 24:00 Sun 24h
Performance	Dual Xeon 3.4Ghz; 2GB RAM; LAN 1000Mbps
Capacity	30GB for OS / 20GB for DB / 860GB for user files
Continuity	Availability to be recovered within 24 hours of an outage
Service desk	Gold level incident response

Control

Service manager: GSI
Access permissions through to the approval of the soliciting department

Training

User manual available

Dependencies

Standard desktop application service

Sub-services and systems

Physical server HP DL360 G4

Owner

GSI

Availability

24/7

SAN HP EVA 4400

Owner

GSI

Availability

24/7

GIS

Description

Set of applications for consulting and defining geographical information

Scope

Gabinete de Organização e Sistemas de Informação
Gabinete de Projectos

Service level

Silver

Availability	24/7
Performance	2 vCPU 2.83Ghz; 2GB RAM; LAN 1000Mbps
Capacity	25GB for OS / 40GB for Workspaces / 250 GB for SIG and DB files
Continuity	Availability to be recovered within 8 hours of an outage
Service level	Silver level incident response

Control

Service manager: GSI
Access permissions to GSI and GPR

Training

No supporting documentation or training plans

Dependencies

Standard desktop application service

Sub-services and systems

VMWare

Owner
GSI

Availability
24/7

Blade Center HP

Owner
GSI

Availability
24/7

SAN HP EVA 4400

Owner
GSI

Availability
24/7

Communication

Description

E-mail, public folders and webmail service

Scope

Internal and external users

Service level

Gold

Availability	24/7
Performance	4 vCPU 2.83Ghz; 8GB RAM; LAN 1000Mbps
Capacity	30GB for OS / 600GB for DB
Continuity	Availability to be recovered within 4 hours of an outage
Service desk	Gold level incident response

Control

Service manager: GSI
Access permissions through to the approval of the soliciting department

Training

User manual available

Dependencies

Standard desktop application service

Sub-services and systems

VMWare high availability cluster with three nodes

Owner

GSI

Availability

24/7

Blade Center

Owner

GSI

Availability
24/7

SAN HP EVA 4400

Owner
GSI

Availability
24/7

Files / Printing

Description
Files sharing, storage and printing service

Scope
All departments; properly authorized external users

Service level

Gold

Availability	24/7
Performance	4 vCPU 2.83Ghz; 2GB RAM; LAN 1000Mbps
Capacity	35GB for OS / 2.5TB for user files
Continuity	Availability to be recovered within 4 hours of an outage
Service desk	Gold level incident response

Control
Service manager: GSI
Access permissions through to the approval of the soliciting department

Training

User manual available

Dependencies

Standard desktop application service

Sub-services and systems

VMWare high availability cluster with three nodes

Owner

GSI

Availability

24/7

Blade Center

Owner

GSI

Availability

24/7

SAN HP EVA 4400

Owner

GSI

Availability

24/7

Terminology

Cluster:
A computer cluster is a group of linked computers, working together closely so that in many respects they form a single computer. The components of a cluster are commonly, but not always, connected to each other through fast local area networks. Clusters are usually deployed to improve performance and/or availability over that of a single computer, while typically being much more cost-effective than single computers of comparable speed or availability.

DB:
Short for Database

IT Service:
One or more technical or professional IT capabilities that enable a business process.
An IT service exhibits the following characteristics:

- Fulfils one or more needs of the customer
- Supports the customer's business objectives
- Is perceived by the customer as a coherent whole or consumable product

IT System:
An integrated composite that consists of one or more of the processes, hardware, software, facilities and people, that provides a capability to satisfy a stated need or objective. It is a collection of resources and configuration items or assets that are necessary to deliver an IT service
An IT system is sometimes referred to as a Technology Solution

OS:
Short for Operating System

SAN:
Storage Area Network: A network of storage disks. In large enterprises, a SAN connects multiple servers to a centralized pool of disk storage. Compared to managing hundreds of servers, each with their own disks, SANs improve system administration. By treating all the company's storage as a single resource, disk maintenance and routine backups are easier to schedule and control. In some SANs, the disks themselves can copy data to other disks for backup without any processing overhead at the host computers.

Appendix III. Alarms Map

Serviço	Componente	Monitor	Protocol	Alarm	Action
Comunicação	Correio electrónico	OS free space	SNMP	Less than 4GB	Send alarm e-mail
Comunicação	Correio electrónico	DB free space	SNMP	Less than 8GB	Send alarm e-mail
Comunicação	Correio electrónico	Disk usage rate	WMI	Above 90%	Send alarm e-mail
Comunicação	Correio electrónico	Network availability	ICMP	No reply in the last 30s	Send alarm e-mail
Comunicação	Correio electrónico	Network response time	ICMP	Above 80ms	Send alarm e-mail
Comunicação	Correio electrónico	SMTP service status	WMI	Not "Running"	Send alarm e-mail; Restart service
Comunicação	Correio electrónico	Mailbox access service status	WMI	Not "Running"	Send alarm e-mail; Restart service
Comunicação	Webmail	OS free space	SNMP	Less than 4GB	Send alarm e-mail
Comunicação	Webmail	Network availability	ICMP	No reply in the last minute	Send alarm e-mail
Comunicação	Webmail	Network response time	ICMP	Above 80ms	Send alarm e-mail
Comunicação	Webmail	Web access service status	WMI	Not "Running"	Send alarm e-mail; Restart service
Ficheiros / Impressão	Partilha de ficheiros Comum	OS free space	SNMP	Less than 4GB	Send alarm e-mail
Ficheiros / Impressão	Partilha de ficheiros Comum	Free space for each storage logical volume	SNMP	Less than 6GB	Send alarm e-mail
Ficheiros / Impressão	Partilha de ficheiros Comum	Disk usage rate	WMI	Above 90%	Send alarm e-mail
Ficheiros / Impressão	Partilha de ficheiros Comum	Network availability	ICMP	No reply in the last 30s	Send alarm e-mail
Ficheiros / Impressão	Partilha de ficheiros Comum	Network response time	ICMP	Above 80ms	Send alarm e-mail
Ficheiros / Impressão	Partilha de ficheiros Comum	File share service status	WMI	Not "Running"	Send alarm e-mail; Restart service
Ficheiros / Impressão	Impressão	OS free space	SNMP	Less than 4GB	Send alarm e-mail
Ficheiros / Impressão	Impressão	Network availability	ICMP	No reply in the last 30s	Send alarm e-mail
Ficheiros / Impressão	Impressão	Network response time	ICMP	Above 80ms	Send alarm e-mail
Ficheiros / Impressão	Impressão	Print service status	WMI	Not "Running"	Send alarm e-mail; Restart service
Sistema de Informação Geográfica	SIG	OS free space	SNMP	Less than 4GB	Send alarm e-mail
Sistema de Informação Geográfica	SIG	DB free space	SNMP	Less than 8GB	Send alarm e-mail
Sistema de Informação Geográfica	SIG	Disk usage rate	WMI	Above 90%	Send alarm e-mail
Sistema de Informação Geográfica	SIG	Network availability	ICMP	No reply in the last 30s	Send alarm e-mail
Sistema de Informação Geográfica	SIG	Network response time	ICMP	Above 80ms	Send alarm e-mail
Sistema de Informação Geográfica	SIG	Licensing service status	WMI	Not "Running"	Send alarm e-mail; Restart service
Sistema de Informação Geográfica	SIG	DB engine service status	WMI	Not "Running"	Send alarm e-mail; Restart service
Gestão Documental	Gestão Documental	OS free space	SNMP	Less than 4GB	Send alarm e-mail
Gestão Documental	Gestão Documental	DB free space	SNMP	Less than 8GB	Send alarm e-mail
Gestão Documental	Gestão Documental	Disk usage rate	WMI	Above 90%	Send alarm e-mail
Gestão Documental	Gestão Documental	Network availability	ICMP	No reply in the last 30s	Send alarm e-mail
Gestão Documental	Gestão Documental	Network response time	ICMP	Above 80ms	Send alarm e-mail
Gestão Documental	Gestão Documental	DB engine service status	WMI	Not "Running"	Send alarm e-mail; Restart service
Gestão Documental	Gestão Documental	Client access service status	WMI	Not "Running"	Send alarm e-mail; Restart service
Gestão Documental	Gestão Documental	Web access service status	WMI	Not "Running"	Send alarm e-mail; Restart service
Gestão Documental	Hardware de suporte	System temperaure	SNMP	Above 28 degrees	Send alarm e-mail and SMS
Gestão Documental	Hardware de suporte	General health status	SNMP	Not "Normal"	Send alarm e-mail and SMS
Gestão Documental	Hardware de suporte	Fibre channel connections status	SNMP	No connection	Send alarm e-mail
Gestão Documental	Hardware de suporte	Network interfaces status	SNMP	No connection	Send alarm e-mail
Gestão Documental	Hardware de suporte	Hard drives status	SNMP	In degraded state	Send alarm e-mail
Aplicação de gestão	Todos os módulos	OS free space	SNMP	Less than 4GB	Send alarm e-mail
Aplicação de gestão	Todos os módulos	DB free space	SNMP	Less than 8GB	Send alarm e-mail
Aplicação de gestão	Todos os módulos	Disk usage rate	WMI	Above 90%	Send alarm e-mail
Aplicação de gestão	Todos os módulos	Network availability	ICMP	No reply in the last 30s	Send alarm e-mail
Aplicação de gestão	Todos os módulos	Network response time	ICMP	Above 80ms	Send alarm e-mail
Aplicação de gestão	Todos os módulos	DB engine service status	WMI	Not "Running"	Send alarm e-mail; Restart service

Serviço	Componente	Monitor	Protocol	Alarm	Action
Aplicação de gestão	Todos os módulos	Client access service status	WMI	Not "Running"	Send alarm e-mail; Restart service
Cluster VMWare	Todos os nós	Cluster status	SNMP	Not "Normal"	Send alarm e-mail and SMS
Cluster VMWare	Todos os nós	High availability service status	SNMP	Not "Running"	Send alarm e-mail
Cluster VMWare	Nó 1	Node status	SNMP	Not "Normal"	Send alarm e-mail
Cluster VMWare	Nó 2	Node status	SNMP	Not "Normal"	Send alarm e-mail
Cluster VMWare	Nó 3	Node status	SNMP	Not "Normal"	Send alarm e-mail
Blade Center	Blade Center	System temperaure	SNMP	Above 28 graus	Send alarm e-mail and SMS
Blade Center	Blade Center	General health status	SNMP	Not "Normal"	Send alarm e-mail and SMS
Blade Center	Blade Center	Fibre channel connections status	SNMP	No connection	Send alarm e-mail
Blade Center	Blade Center	Network interfaces status	SNMP	No connection	Send alarm e-mail
SAN	SAN	General health status	SNMP	Not "Normal"	Send alarm e-mail and SMS
SAN	SAN	System temperaure	SNMP	Above 35 degrees	Send alarm e-mail and SMS
SAN	SAN	Fibre channel connections status	SNMP	No connection	Send alarm e-mail
SAN	SAN	Hard drives status	SNMP	In degraded state	Send alarm e-mail