IT IN THE OIL &GAS INDUSTIRES "DOING MORE WITH LESS "

BY:
SENTHIL PALANIAPPAN

GUIDED BY
MR. ASHISH BHARADWAJ
PROGRAM DIRECTOR
UNIVERSITY OF PETROLEUM & ENERGY STUDIES
NEW DELHI, INDIA

FINAL DISSERTATION REPORT SUBMITTED IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR MASTER OF TECHNOLOGY (PETRO INFORMATICS) OF UNIVERSITY OF PETROLEUM AND ENERGY STUDIES, INDIA

COLLEGE OF ENGINEERING
UNIVERSITY OF PETROLEUM & ENERGY STUDIES, INDIA
A1, KAILASH COLONY, NEW DELHI 110048

UNIVERSITY OF PETROLEUM & ENERGY STUDIES

Certificate of Originality

This is to certify that the dissertation report on" **IT in the oil and gas industry doing more with less**" submitted to the University of Petroleum and Energy Studies by **Senthil Palaniappan (R060104017)**, in partial fulfillment of the requirement for the award of the degree of **Masters of Technology (Petro Informatics)**, is a bonafide work carried out by him under my supervision and guidance.

Prof. Ashish Bharadwaj
Head - Academic Computing-cum-Assoc. Professor
University of Petroleum & Energy Studies

ACKNOWLEDGEMENT

This is to acknowledge with thanks the help, guidance and support that I have received during the dissertation.

I have no words to express a deep sense of gratitude to the management of **BPCL & UNIVERSITY of PETROLEUM & ENERGY STUDIES** for giving me an opportunity to pursue my Dissertation, and particularly **Mr. Ashish Bharadwaj** (Project Guide) for his able guidance and support. Also I want to express my gratitude to **Mr.C.P.Miller** Program Manager-OBC-WAN Network for his help and guidance for this dissertation.

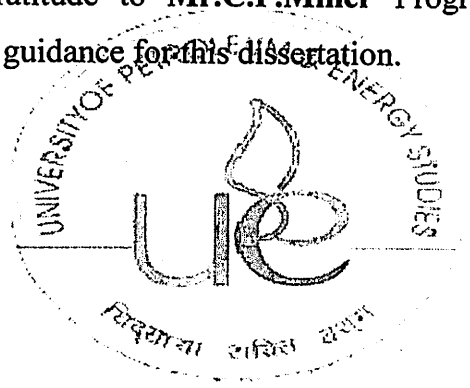

Senthil Palaniappan
U.P.E.S

Table of Contents

1. Introduction

The oil and gas industry, like other mature, commodity industries, is experiencing pressure from today's economic environment. Competing demands are made on companies:

- Increase productivity
- Reduce cost and cycle time
- Improve decision making
- Minimize capital investment
- Find more reserves
- Use fewer resources
- Cut labor expense

These imperatives are most difficult. Doing more with less.

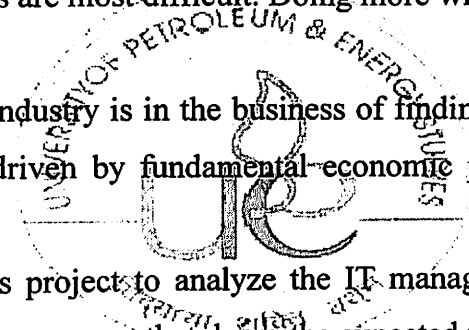

The oil and gas industry is in the business of finding and producing hydrocarbons. It is a "business" driven by fundamental economic principles, with accountability to stakeholders.

The scope of this project to analyze the IT management concepts for Oil and Gas industries in such a manner that doing the expected tasks with fewer resources. For an organization different IT applications will be running to perform a different task to achieve the business goal of the industry.

1.1 Return to Fundamentals in Purchase Decisions

Oil companies will be focused on short-term financial justification of IT investments, and investments that can demonstrate quick payback will be chosen over those that cannot. At the same time, previous investments in major systems will be scrutinized, and implementers will be increasingly asked to prove that the systems have delivered on their promise.

Enterprises will delegate purchasing decisions to their business units rather than the IT departments with financial and business principles playing a strong part in the ROI measurement. Focus will remain on short-term financial justification of technology investments --ROI, cost, value, risk, and flexibility. Cost, though just one factor, will be the leading consideration.

1.2 Offshore Outsourcing for Cost Containment

Thanks to a tough economic environment, there will be significant growth in the use of low-cost, high quality services based out of India, for not only IT-related services but also for outsourcing business processes, including highly technical tasks. For example, an Indian firm, **Scicom**, already provides high-end, technical software maintenance and development services to major players in the industry, at substantially lower costs. Another firm, **Scandent**, has developed Project Administrative Support Services, an IT-enabled service, that has the ability to improve schedule and cost performance in capital projects.

The low-cost advantages of global IT services firms will put increasing pressure on traditional suppliers to the industry, thereby significantly reducing industry IT expenditures.

1.3 Focus on Maximizing Current Infrastructure and Solutions

After few years of spending wildly on new, exciting technologies, oil companies are searching for a way to make all the equipment, software and infrastructure work together- squeeze more out of existing infrastructure, with no need to buy new equipment or solutions.

An example of waste in hardware expenditures is highlighted by IBM data that shows that server utilization is a mere 40 percent of the total installed capacity. Up to 60

percent of an email server's total power can potentially be exploited to perform other tasks, like intranet hosting or even printing.

In 2003, different storage systems, servers, and network devices will begin to be managed through a central console based on virtualization. Virtualization reduces capital expenditures (no need to buy new equipment) and the cost of managing IT equipment.

1.4 Oilfield Service Companies Take on Greater Services Role

Oil companies are increasingly looking to better leverage their asset base to optimize Performance in a market characterized by tight capital budgets. This not only includes managing field assets, but also a more comprehensive view of all assets from facilities to IT to equipment to employees to knowledge. Oil companies will look for real-time information flow of production metrics, rapid, pro-active responses to assets that may fail, and more strategic decisions on how assets will be serviced.

Oilfield Service Companies looking for additional sources of revenue are expanding product and service offerings and accepting more responsibility for their customer's success by supporting those assets. Earlier they occupied huge place at different locations. So the cost involved in establishing the IT infrastructure at different places, to maintain them and to expand the IT infrastructure for future need is more.
So, the intention is to minimize these costs by implementing these IT management concepts like Consolidation of It resources including Consolidation (date storage, network, shared devices, applications, and servers), virtualization (data storage, network, shared devices, operating systems).
This document addresses the advantages of the every concept and issues and impacts.

2. Consolidating Servers

Servers are still the primary focal point for consolidation because they are so obvious. Whether there are 100[2] servers or 5000[2] servers, there will be too many to manage effectively. Today's distributed computing environment lends itself to a proliferation of servers. Reducing and controlling the number of devices to manage and simplifying ways to manage them is the goal of most IT groups. there are two types of scalability.

2.1 Applying Vertical and Horizontal Scalability

When there is a talk about consolidating servers, Organizations generally refer to scaling them vertically or horizontally.

- Vertical scalability enables organization to reduce the number of servers by consolidating multiple applications onto a single server.
- Horizontal scalability enables organizations to deal with increased workloads through the replication of servers and the distribution of workloads across those servers.

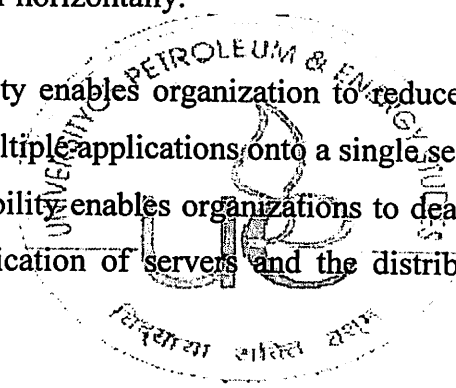

By thinking of consolidation in these terms, the beginning is to define the approach required by that particular consolidation project. Once it has been decided whether this consolidation project requires horizontal scaling, vertical scaling, or a combination of the two, the process can be refined further by the approach of identifying patterns in the server population. Examples of vertical and horizontal scaling are presented in the Sun Blueprints book *Consolidation in the Data Center*

2.2 Identifying Patterns in an End-to-End Architecture

In the end-to-end architectures that are prevalent today, tiers of servers are specialized for particular tasks. When consolidating servers, looking for patterns in server population is very important. After the identification of these patterns within the tiers, start devising a consolidation strategy. Scalability is the key, here. Because it is

expected to deliver predictable service levels in response to unpredictable workloads, it is important to use the right type of scalability for each part of a consolidated architecture. The following sections describe common patterns in an end-to-end architecture.

For consolidation discussions, it is assumed that there are three server types, or tiers:

- The presentation tier is the closest tier to the end user.
- The business, or middleware, tier is where applications or middleware run in conjunction with the other tiers.
- The resource tier is where large, scalable servers run mission-critical applications and databases.

Although architectures with these characteristics have been around for awhile, most corporations still have many servers running monolithic applications. In many cases, these are older servers running mature applications. These servers are generally excellent candidates for server and application consolidation.

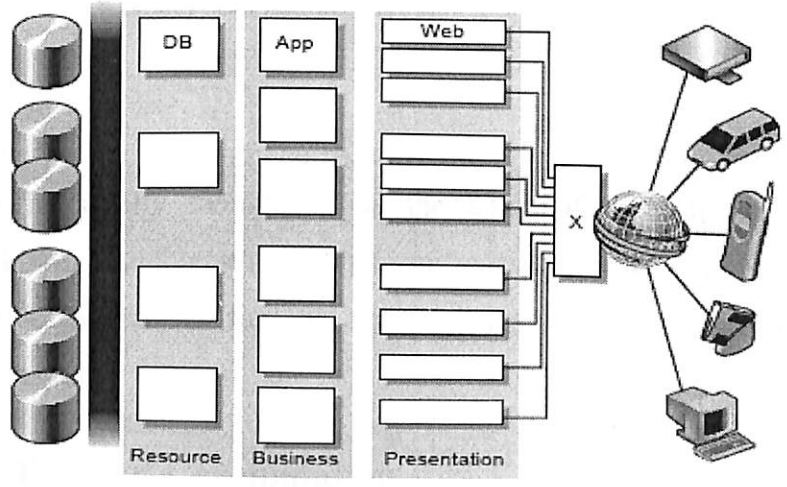

Fig-01 End to End Architect

2.2.1 Presentation Tier

As user gets closer to the Internet and the delivery devices for user applications, the most common techniques for scalability are replication and load balancing across an array of generally small servers. This is usually referred to as horizontal scalability.

On the presentation tier, there will be several small, one or two processor servers running UNIX™, Linux, or Microsoft Windows. Because current architectures call for many small servers in this tier, the opportunities for consolidating to fewer servers are limited. Despite these limitations, there is still ample opportunity to implement standards and best practices to simplify the management of these servers and to reduce total cost of ownership (TCO).

2.2.2 Business Tier

After the presentation tier, there will be business tier servers. These generally run applications or middleware in conjunction with the other tiers. While it is used to see primarily horizontal scaling in this tier, there is growing interest in reducing the number of applications using vertical scalability, as well.

2.2.3 Resource Tier

The resource tier is usually thought of as the heart of the data center. It is characterized by large, scalable servers that run mission-critical applications and databases. Usually organizations find the most interest in vertical scalability and server consolidation on this tier.

2.3 Identifying Types of Consolidation

Some major consulting groups identify different categories of server consolidation. For example, Gartner refers to physical, logical, and rational consolidation groups—categories defined in the following sections. The techniques used to achieve these specific types of consolidations are discussed in later chapters.

2.3.1 Physical Consolidation

Physical consolidation involves consolidating data centers and moving servers to fewer physical locations. The theory behind physical consolidation is that by having servers in fewer physical locations, management consistencies and economies of scale can be achieved more easily when the servers are dispersed. Physical consolidations may also enable organizations to reduce data center real estate costs. It is generally felt that physical consolidation has the lowest risk, but that it also has the lowest payback.

2.3.2 Logical Consolidation

Logical consolidation involves implementing standards and best practices across server population. By doing this, organization can realize substantial benefits in the productivity of their IT staff. They can manage the environment more efficiently and more effectively. This can often result in lower systems management costs and in lower TCO. Logical consolidation is often implemented with physical consolidation and rationalization.

2.3.3 Rationalization

Rationalization involves the deployment of multiple applications on fewer, larger servers and in fewer instances of the operating system (OS). Because TCO reduction is closely tied to the number of instances of an OS that is managed by that organization, reducing the number is the best way to reduce TCO. While rationalization is the riskiest form of server consolidation, it offers the biggest TCO reduction and return on investment (ROI).

Because of its potential to reduce TCO and increase ROI, it is found that most of our customers find rationalization to be the most attractive method for consolidation, despite its higher risks.

2.4 Consolidating Storage

Today, there is as much interest in storage consolidation as there is in server consolidation. As mentioned earlier, every new server is deployed, it results in more storage. In many cases, the cost of storage for a server exceeds the cost of the server, and although the server may not grow very much, the amount of storage required for an application will grow indefinitely.

When it is for application sets, there will be a lot of data replication because multiple applications look at the same data. For example, if an organization has an OLTP server, a batch server, a data warehouse server, and four data mart servers, they may have seven copies of a particular dataset such as a customer master file. While, in theory, all seven copies are identical, because different developers probably created their applications, it is likely that there will be differences in each copy of the data. This situation has the potential to create a situation where reports run on the "same" data, yet yield different results.

As with servers, the goals in storage consolidation are to reduce complexity, increase utilization, and reduce TCO. These goals can be achieved through a variety of consolidation techniques. The ultimate goal is data sharing among applications. Unfortunately, this often requires redesigning and redeveloping applications, so it is a long-term goal, at least for backward consolidation. In a forward consolidation, data sharing should absolutely be a goal.

Other benefits of storage consolidation include:

- Easier backup and recovery
- Increased availability
- Improved scalability
- Storage resource pooling

When undertaking a storage consolidation effort, it is critical that they understand how disk space is utilized. While most of the customers can't tell exactly what their

current disk space utilization rates are, many companies they have surveyed estimate a rate of 40 percent, or less. For companies that have tried to accurately assess disk utilization, complexity often hinders their efforts. For example, one company started counting storage, but stopped at 110 terabytes[8] after realizing that they just couldn't count it all. When it has been evaluated their disk utilization, it has been found that they were utilizing only 20 percent of their available storage.

There are several types of storage consolidation available today. The following sections describe the three most common types.

2.5 Consolidating Servers and Their Associated Storage

With every server consolidation, there is an accompanying storage consolidation. As to move multiple applications to a single instance of the OS, there is a need to move their storage as well, as shown in the following graphic. In theory, once the storage is moved, data will be available to any of the applications. This is the most primitive form of storage consolidation.

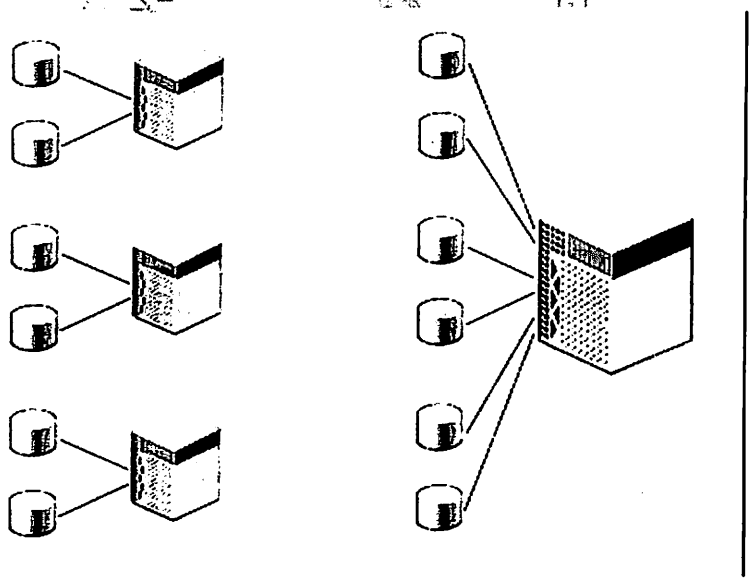

FIGURE 0-2 Consolidating Direct-Attached Storage

2.5.1 Connecting Heterogeneous Environments to a Single Storage Component

Many of our customers have heterogeneous IT environments. They run servers and OSes from many vendors, and they want to access data from a variety of different servers. With direct-attached storage, this was difficult to do. Now, with products like the Sun StorEdge™ 9900 storage array, it is possible to connect the Solaris™ Operating Environment servers, other UNIX servers, Microsoft Windows NT servers, and mainframe servers to the same storage array. The following graphic demonstrates this capability.

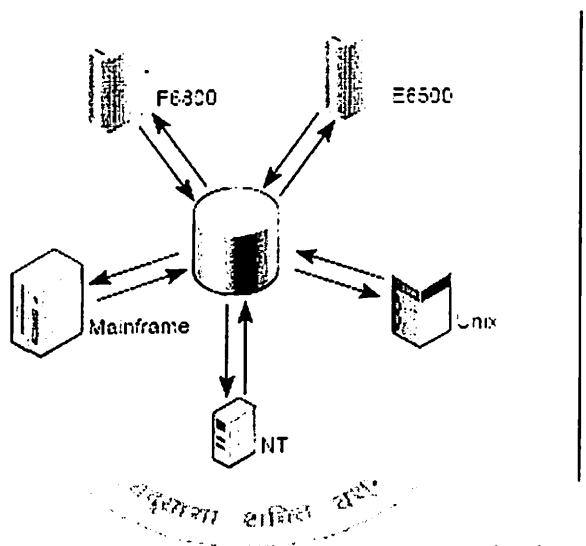

Fig-03 Heterogeneous Environments to Single Storage

This is a very popular type of storage consolidation, especially where there are multiple mission-critical applications running on servers from a variety of vendors.

2.6 Consolidating With Storage Area Networks

Storage area networks (SAN) have been the hottest trend in storage architectures for the last few years. As a technology, the SAN is now mature enough that it can be implemented using standards and standard configurations. As shown in the following

graphic, SAN technology inserts a network, or fabric, of switching devices between servers and storage that enable any server or application to access any storage connected to it. The fabric can then be configured to allow various servers to access various storage.

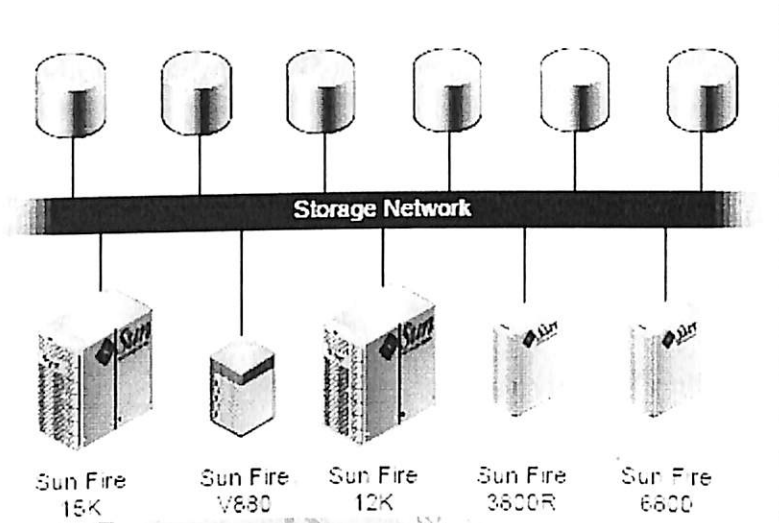

FIGURE 0-4 Storage Area Network (SAN) Configuration

Another hot storage technology is network attached storage (NAS). This technology allows servers and clients to utilize storage directly over the network using common networking protocols like network file system (NFS) and Server Message Block (SMB). Although not used greatly in many server and application consolidations within the data center, it is used extensively in file and print consolidations at the department and workgroup level. As the technology matures, expect it to work itself into data center consolidations.

2.7 Consolidating Shared Services

Another type of consolidation that is rapidly gaining popularity is middleware or shared services consolidation. Over the last few years, companies have implemented

shared services (such as file, print, authentication, and e-mail) in a variety of ways. Business units typically implement their own versions of these services using software from a variety of vendors; however, when they want to share some of this information, they find that the inconsistencies and incompatibilities among these various architectures increase complexity, thereby increasing cost and making it difficult to share. As a result, its often found that companies want to rebuild their shared services as web services using a standards-based, unified architecture such as Sun™ Open Net Environment (Sun ONE).

A common example of this is the use of directory services. Directory services have been implemented over the years using a variety of architectures. Now that standards based architectures are available, products like the Sun ONE Directory Server, which is based on the Lightweight Directory Access Protocol (LDAP), are being used to design and implement corporate-wide directory services. Condensing multiple, disparate directory architectures into a single, corporate architecture allows corporations to simplify directory maintenance and easily replicate directory services, as needed.

2.8 Consolidating Networks

To consolidate servers and applications, network consolidation is usually either a big issue or not an issue at all. When organizations collapse data centers and concentrate servers and applications into fewer physical locations, there may be a severe impact on the network. This impact must be evaluated as part of the consolidation project. There will be certainly a need to ensure that there is sufficient network bandwidth to handle network traffic to and from the consolidated location.

Conversely, when they do a rationalization within a single data center, organizations often find that there are no major network changes that need to take place. The overall traffic does not increase, since new servers and applications are not being introduced. Further, rationalization may actually decrease network traffic, because applications in a single instance of the OS don't need the network to communicate with each other.

2.9 Consolidating Data Centers

Many organizations are looking to consolidate multiple data centers into one site. These consolidations range from simple city-wide consolidations to complex region-wide consolidations. Most companies are being driven toward data center consolidation because of the dramatic drop in the cost of telecommunication wide area network lines, the huge disparity of IT wages between certain regions of the world, and the high real estate costs of large cities (especially in New York, London, and Tokyo). For those considering simple local-site consolidations, consolidation offers cost savings and enables disaster recovery initiatives.

If the organization is seriously looking to consolidate a data center, carefully consider the goals for consolidation. Shutting a data center is a huge task, and before the start down the path, it is vital to articulate and defend the reasons for doing it. Further, once a data center is shut down, the costs of reopening it can be enormous. From there, data center consolidations are similar to other types of consolidation, except that assessment (especially application, networking, and physical planning) and implementation become much more complex. Next one is High Performance Computing in achieving scalability of storage.

2.10 High-Performance Computing (HPC)

Most oil and gas companies have data requirements that are growing exponentially. the IT department is under enormous pressure to meet quality of storage service requirements that dictate a much smaller backup window than the 24 hours it currently takes. Sun StorEdge Performance Suite and StorEdge Utilization Suite software provide maximum scalability, performance, and throughput for data-intensive applications. This allows to automatically stream data to tape devices at fully rated device speeds, eliminating the need for traditional "batch mode" backup, thus increasing the quality of storage service. Costs are dramatically reduced, because fewer hardware, software, and manual system administrators are required. Moreover, it is able to capitalize on the file system, file sharing, and volume management

aspects of the software for rapidly growing capacity requirements, as well as to reduce cost and increase efficiencies.

Some systems require organization to compromise throughput in order to achieve ideal storage utilization. This is not the case with the use of Sun StorEdge Performance Suite software and Sun StorEdge Utilization Suite. In fact, the San Diego Supercomputer Center recently achieved a data transfer rate of 828 megabytes per second to tapes by leveraging its SAN solution based on Sun's HPC SAN architecture..

2.11 Consolidating People Resources and Processes

In any consolidation project, one must not neglect the people and processes they use to manage the environment. Time after time, when there is a work with clients who have both mainframe and distributed-computing environments, it is found that the mainframe side of the house runs smoothly and problem free, while the distributed-computing side of the house is often chaotic, with few developed standards and procedures. The problems that result from this situation demonstrate the importance of resolving these people and process issues. While mainframes consistently run with high availability and high service levels, distributed computing systems often suffer from low service levels and low availability.

Some consulting groups estimate that only 20 percent of data center availability is hardware or technology related; the other 80 percent is estimated to be directly related to people and process issues. The lesson is that successful consolidations must address standards, people, and processes. Without these, availability and service levels may be compromised.

Another benefit of implementing standards and best practices is that 10–20 percent reduction in TCO. Another concept is sever sizing

2.12 Server Sizing Concepts and Considerations

Sizing servers correctly for specific applications is increasingly important. Pc servers are now used to deliver a wide range of enterprise applications: messaging platforms such as Microsoft exchange Server and Lotus Notes, databases like Oracle and Microsoft SQL Server, and ERP software including SAP R/3 and JD Edwards One World.

Any large-scale implementation of applications like these is clearly going to entail a major project, and an important aspect is sizing the server hardware correctly. Organizations need to be able to predict how many users can be supported, what the optimal system configuration will be, how to tune the application and what user response times can be expected. Doing this accurately is a major challenge-infact, it's well possible.

For Internet sites, one unknown parameter is the number of users-new e-commerce sites universally seem to totally underestimate the number of visits they will receive after launch. Internal rollouts don't have quite the same challenge-usually; organizations have an accurate idea of how many users of an ERP application there are going to be. However, even within the enterprise, various factors can make the task more difficult than it already is. Applications within the PC arena change very frequently-no sooner have organizations got to grips with sizing exchange when a new version comes out.

Thus the initial guess, even with the benefit of the advice herein, might be overtaken by events and turn out to be wrong. It's therefore a very good idea to adopt hardware that can be readily upgraded, by adding processors, memory, more hard drives and even additional network interfaces if required. Microsoft's IP load-balancing software for Windows NT and various clustering solutions mean that in some circumstances they can even add additional servers to take up unexpected high loads.

2.11.1 The Market

This sizing issue is unique to the PC world. it simply isn't a problem within the high-end system market-midrange and mainframe-where operating systems and applications change far less frequently.Furthermore,because in many cases the entire system comes from a single vendor-hardware, Operating system and application software-the manufacturer will have a full understanding of server sizing and be able to publish accurate and relevant sizing date. It has to be right first time. Customers at this end of the market don't want to find they have spent what could be millions of dollars unnecessarily on a higher specification than they need, and they equally can't afford for their users to suffer inadequate performance.

As PC servers enter these arenas, users are going to expect similar standars.However; organizations are never going to have enough information to be able to accurately size a server. Network server performance is a function of client usage, client hardware and software configuration, network setup and the configuration of the server hardware and software. A detailed knowledge of all the above is very difficult to accomplish outside a lab environment.

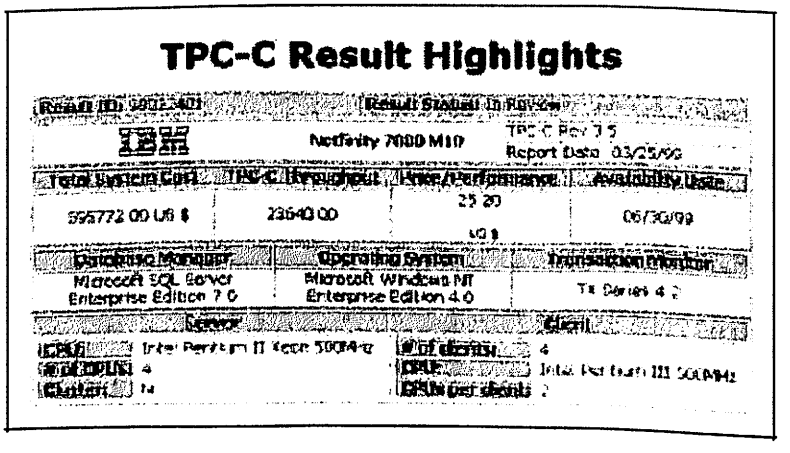

Fig-04 Certified benchmark results are readily obtainable from various websites

2.11.2 Guess

Fortunately, there are ways to arrive at a reasonably good estimate of the server configuration required to meet given load requirements. One is to build and test a prototype. While this would be an accurate technique, cost and time restrictions mean that this is rarely likely to be a serious contender.

The other method is to guess. clearly, a completely wild guess is unlikely to be useful, but guessing is clearly the least expensive and fastest method of server sizing. Thus, if existing measurement data can be used to refine the guess to become something more accurate, then guessing becomes a very worthwhile and useful technique.

The 100 IBM-certified Netfinity Professional Server Experts (PSEs) who qualified to attend IMPACT 99, IBM's TechConnect event abroad the Disney Magic, benefited from far more than a four-day cruise of the Caribbean. Gregg McKnight, who has spent years working on server configurations, outlined to them how they can use published benchmark figures to achieve better results when estimating server sizes.

The first step is to gather as much information as possible about the intended usage. One needs to know the number of concurrent users to be expected –not just the number of potential users. And an accurate picture of what operations these users will be carrying out is also important. Are they light, moderate or heavy users? A task oriented user in a call center will be making much greater use of a customer database than a salesman who pops into the office once or twice a week to look up addresses for the next few visits.

2.11.3 Benchmark Results

Once as much information has been gathered the next is to try to identify a benchmark that as closely as possible mirrors the intended usage –this is extremely important. If something that isn't particularly similar, good results cannot be obtained. Look carefully at the different workloads, of which mail is the one usually quoted by PC manufacturers. The links section at the end of this article contains details of many web sites where benchmark results can be found.

The techniques described in this article can be applied to a variety of different benchmarks such as Notesbench, TPC-C, TPC-D/H, SAP R/3 and SPEC web. Obviously where a benchmark is very specific in the type of use-such as the mail work load for Notesbench-it can be directly related to a target environment. Some of the more general benchmarks are also quite relevant to realistic workloads-TPC-C uses random small transactions of the order entry and inventory management type typical of many sales operations.

Where there is a problem is if there is no standard published benchmark available for the application planned to run, or if the only one available is completely unrealistic. Microsoft Exchange Server Load Simulator, for example, has been used by Microsoft to demonstrate that exchange could achieve a higher user count than LotusNotes, but doesn't in any way mirror real-life operations.

Although Microsoft has renamed it to be MAPI Message Benchmark (MMB), it's still the same product. According to Microsoft it represents the workload of a medium user of email. Infact, it emulates a user who sends 14 messages and reads 70 messages a day probably fairly light by enterprise standards today.

2.11.4 Adjusting results

Published benchmark results are extremely optimistic, and a number of factors contribute to this. Even where benchmarks do accurately reflect a real-life load, the servers submitted for benchmark certification will generally have been tweaked and tuned by a team of experienced experts for a period of some weeks. Furthermore, benchmarks will run a server flat out at 100% CPU utilization. Organization probably wants to allow a bit of headroom and will prefer not to push the server quite so hard. Furthermore, a production server will be running other tasks apart from the basic application. When setting up a server for a benchmark, without having the chances of breaking a record by having backup processes and system management tools running on the same hardware.

Taking all these factors into account, benchmark results cannot be translated directly to production environment. Instead, work has to be done on the basis of just 33% to 50% of the number of users shown in the benchmark results. Which figure has been chosen will depend to some extent on the benchmark itself. **Notesbench** is a fairly realistic reflection of a real-life loading, probably the work based on the 50% figure. For other benchmarks, 33% is a safer bet. So a benchmark result and it shows 20,000 users for a particular server, work instead on a figure of between 6666 and 10,000, as appropriate.

Organizations must read the benchmark results carefully to look for hidden secrets. Software like this allows the maximum number of users to be increased by a factor of between five and 10, by reducing the number of concurrent connections to the server application. Example, that would cut the users to between 666 and 2000. (Alternatively, buy Tuxedo.)

2.11.5 Hardware Specification

After getting a reasonable figure for the number of users the benchmarked hardware could support, so the next step is to adapt the hardware specification to the number of users actually wanted. If organization wants an IBM Netfinity server but can only find Compaq results, the percentage difference will not significantly upset the results.

Conveniently, the benchmark results will show the detailed hardware configuration used to achieve the result - the particular model of server, the number and speed of CPUs, amount of cache, amount of memory, disk subsystem configuration and network adapter information. Thus, it is the starting point for the hardware specification of their desired system.

Organizations learn a great deal about the hardware demands of the application in question from these figures. It should be able to ascertain whether the application is memory-, CPU-, hard disk- or network intensive simply by looking at the figures. Organizations should be prepared for some surprises - 4 GB of RAM and 250 hard disk drives is not an unusual specification for a published benchmark result.

The first step is to scale the number of CPUs to suit for the particular workload - reduce or increase the number of CPUs as necessary. Approximate multipliers it can be used are 1.7 for two way, 3.0 for four-way and 5.0 for eight-way configurations. Hence, if the adjusted benchmark figure is 666 users for a four-way system and to give the support 1,000 users, this could be achieved by increasing the number of CPUs to eight. If there is only need 200 users, a single CPU is enough.

Benchmarks are not usually highly dependent on the network interface. Infact, organizations should be wary of increasing the performance of the network connection. Yes, a faster network interface will allow more user sessions, but will increase the load on the server and mean it gets hit with more requests, and needs to be more powerful to cope with them. Where benchmark figures do suggest that the

load is network-intensive - usually indicated by the presence of multiple network cards – then it's a good idea to put the network adapters on separate PCI buses.

SPECweb99	Measures Web server performance.
SAP	Publishes benchmarks for about 12 different SAP R/3 components.
TPC-C	Measures new order transactions while payment, order status, Delivery and stock level transactions are also taking place.
TPC-H	Represents decision support environments where ad hoc queries are made against a database.
TPC-R	Represents decision support environments where users run a Standard set of queries against a database system.
SPECmail2000	Will, when finished, measure a system's ability to act as a mail server Servicing email requests.
Notesbench	Includes Mail, Mail DB, Groupware, Idle Sessions, Replication Hub and Mail Routing Hub workload tests.

Fig -05.Some Benchmarks and their applicability

2.11.6 Disk Subsystem

Getting the right hard drive configuration is difficult, but critical to achieving good performance. A good estimate of the required file or database size, and a sensible estimate of the expected growth is needed. But it's not just a question of size – the number of disk drives is also relevant. For example, the database size may be 18 GB, which today can be accommodated within a single drive. However, to achieve a particular benchmark figure one might need to provide 250 drives.

Clearly, for most organizations, 250 hard drives on a production system will be somewhat impractical, and fortunately one can get away with using fewer drives than

the benchmarked system. However, organizations should not go too far down the road of cutting down on the number of drives - can probably get away with 50% of the drives used in the benchmark tests, but they definitely shouldn't reduce it further than that for the same number of users.

For SQL Server 7, 20-40 disk drives will be needed per CPU. For **Notes bench**, Organizations can make do with between five and eight RAID-1E drives per CPU.

Pay careful attention too to the disk configuration. For many benchmarks, it is found that there are several RAID drive arrays for the main data storage, but a separate controller and drives for the log file - usually configured in a different way to better suit the sequential nature of log file operations. this setup should be certainly replicated.

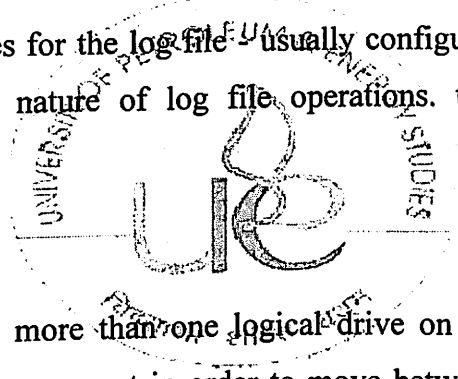

Ideally, never have more than one logical drive on each physical drive -This will require more head movement in order to move between the two logical drives. The first logical drive won't be completely full, and the heads will keep having to move over the empty space - not the case if there is only one logical drive on each physical drive, since all the spare space will be at the end (assuming that the drive has been defragmented).

2.11.7 Software Configuration

The hardware is only the starting point. Full disclosure reports published along with benchmark results will also reveal specifics of the software tuning - both of the operating system and the application. There isn't much to tune on WindowsNT, but on benchmarks that use Windows NT ,there will usually be a number of services have been stopped.

Within the System control panel, Foreground Application Boost will usually be set to a minimum under the Performance tab, and the Server service under the Network

control panel will be set to Maximize Throughput for Network Applications. Some benchmark results will declare these two settings in terms of the registry keys instead, as shown in Figure 2.

Another adjustment sometimes used is the 3gb switch within the BOOT.INI file to allow NT Server to use more user memory and less kernel memory. Take note also of any service pack that might have been applied.

Software application tuning is also important, so do read the published results carefully in this respect. Notesbench results will show the NOTES.INI file settings used, and SQL Server 7 disclosures will reveal any special startup parameters, the stack size (which may well have been tweaked using EDITBIN) and any tuning parameters. Even though most of the complex tuning parameters within SQL Server were removed with the launch of version 7, several still remain.

3. Virtualization -Introduction

The network infrastructure has profoundly impacted the business world. Barriers that existed between companies for most of the 20th century are giving way to accommodate the need for partnership-based opportunities afforded to them through e-business. As the need for interoperability and flexibility increases, IT administrators are scrambling to clean up the debris from the patchwork client-server systems that were pieced together hastily in the storm of the 1980s and 90s. The push for "universal computing" has begun. Much work has been done in the last few years to promote this new initiative. There have been breakthroughs in technologies such as distributed processing (a.k.a. grid computing), "plug-and-play" or "blade" server farms, networked attached storage devices, and many others that support the need for more adaptable technology. In this same mold, the rise of virtualization into this new IT model will drastically change the way information systems are run, by allowing flexibility in both hardware and software systems of which the industry has never seen the like. Through the adoption of virtualization, the CIO will be able to empower the enterprise with systems that adapt quickly to business demands and new

opportunities, a capability that will prove invaluable in today's volatile commercial landscape.

The roots of virtualization are best seen in the computer "time sharing" practices of the late 1950s and early 1960s. Time-sharing was necessary in these distributed computing environments because the technology was extremely expensive. It was not practical to dedicate a computer system to a single user, thus a scheme for dividing the resources among many users was developed. These schemes often used "executive programming" which employed a combination of software and hardware in order to delegate (based on a specified time interval) which user would receive attention from the central processing unit at a particular time. As we will see, this process is similar to that of what we know today as virtualization, in that a layer of abstraction is created in order to logically assign the use of a computer asset. A related concept, multiprogramming, also garnered much attention in this time frame.

In 1967 IBM announced the IBM\360 version 67, which was the first computer to contain "virtual memory," a method in which disk space is used to expand the RAM size of a machine. Later, in the 1970s, the emergence of the "virtual machine" came about. With the virtual machine, an entire system (software and hardware) could be emulated in a contained environment. The virtual machine is perhaps the first form of virtualization, as we know it today. Obviously, these last two instances are an influence on the nomenclature of what we now refer to as virtualization. With the introduction of Intel's recent release of its "Vanderpool" technology (and AMD's subsequent "Pacifica" technology), which provides hardware native server virtualization functionality, the concept of virtualization has taken one step further toward becoming a high-impact, common practice in the enterprise.

3.1 Defining Virtualization

Because of the wide variety of ways in which the technology is deployed, it is difficult to define virtualization in a general manner. Mr. Singh, in his article An Introduction to Virtualization attempts to do so, stating that virtualization is "a

framework or methodology of dividing the resources of a computer into multiple execution environments, by applying one or more concepts or technologies such as hardware and software partitioning, time-sharing, partial or complete machine simulation, emulation, quality of service, and many others" (Singh, 2005). However, the author admittedly leaves out cases in which computer resources are combined rather than separated, such as in storage networks, and unsuccessfully tries to pinpoint the sub-technologies used. In order to find the common thread that ties all virtualization concepts together, it is first necessary to look at the various practices in use today.

3.2 Server Virtualization

Server virtualization enables multiple virtual operating systems to run on a single physical machine, yet remain logically distinct with consistent hardware profiles (Burry & Nelson, 2004). To the contrary, server virtualization can often take the place of the costly practice of manual server consolidation, by combining many physical servers into one logical server. "The idea is to present the illusion of one huge machine that's infinitely powerful, reliable, robust and manageable - whether it's one machine that looks like many, or multiple machines tied together to look like a single system" (Brandel, 2004). The focus of server virtualization is on maximizing the efficiency of server hardware in order to increase the return on investment for the hardware.

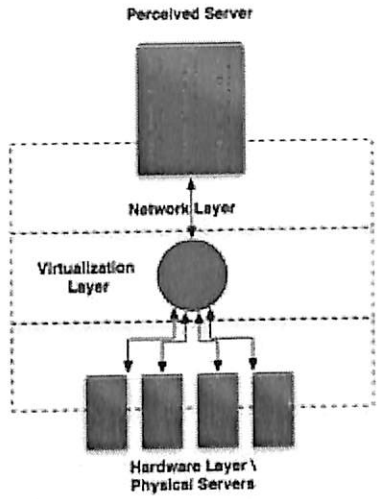

Fig-06: Server Virtualization

3.3 Operating System Virtualization

Through virtualization of operating systems, a single computer can accommodate multiple platforms and facilitate their operation simultaneously. This description is similar to the aforementioned server virtualization, but server virtualization alone does not necessarily provide the ability to run multiple platforms on a single server. Also, in contrast, the goal of OS virtualization is focused more on flexibility than efficiency. OS virtualization can be used to facilitate what is known as universal computing, where software and hardware work together seamlessly regardless of the architecture or language for which they are designed.

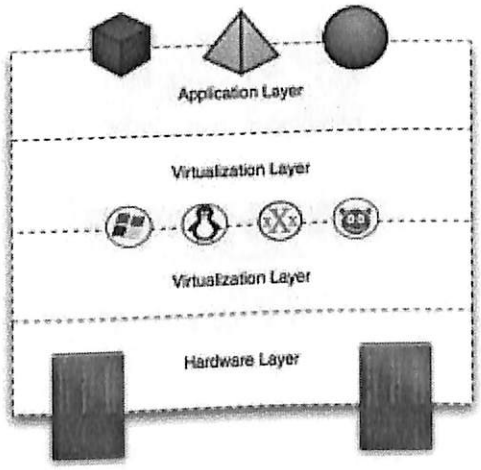

Fig-07: Operating System Virtualization and Application Virtualization

3.4 Application Virtualization

While most of the prevalent virtualization strategies focus on hardware infrastructure, an important and often overlooked method is application virtualization. With application virtualization (also commonly referred to as service virtualization) end-user software is "packaged," stored, and distributed in an on-demand fashion across a network. This virtualization strategy goes hand in hand with the standardized web services initiative that is making waves in the IT industry today. Virtualized applications use a common abstraction layer, which defines a protocol, allowing them to communicate with one another in a standard messaging format. Thus, applications can invoke one another in order to perform requested functions. A virtualized application is not only capable of remotely invoking requests and returning results, but also ensuring that the application's state and other data are available and consistent on all resource nodes executing the application across a grid (Data Synapse 2005).

3.5 Storage Virtualization

Perhaps the most widely deployed and highly regarded virtualization practice, storage virtualization allows separate storage devices to be combined into a perceived single unit. Storage virtualization attempts to maximize the efficiency of storage devices in information architecture.

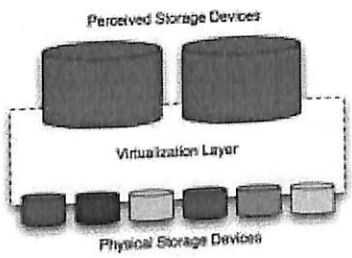

Fig-08: Storage Virtualization

3.6 Data \ Database Virtualization

Data virtualization allows users to access various sources of disparately located data without knowing or caring where the data actually resides (Broughton). Database virtualization allows the use of multiple instances of a DBMS, or different DBMS platforms, simultaneously and in a transparent fashion regardless of their physical location. These practices are often employed in data mining and data warehousing systems.

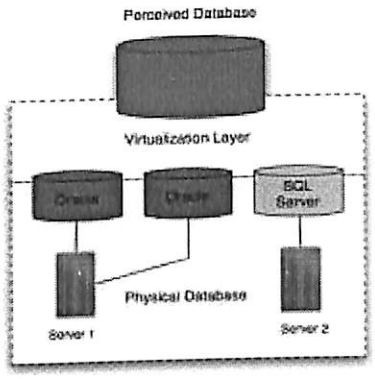

Fig-09: Database \ Data Virtualization

3.7 Network Virtualization

By virtualizing a network, multiple networks can be combined into a single network, or a single network can be separated logically into multiple parts. A common practice it to create virtual LANs, or VLANs, in order to more effectively manage a network.

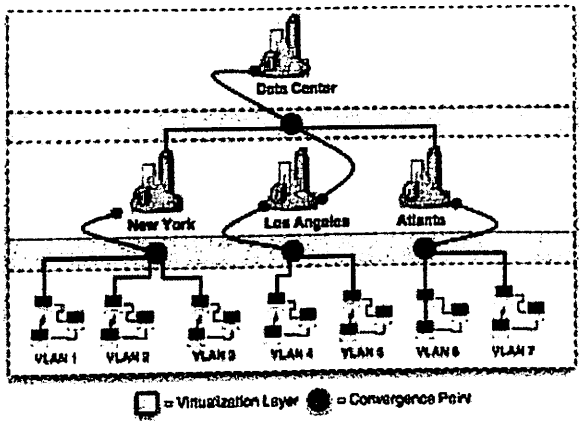

Fig-10: Network Virtualization

3.8 Virtualization Defined

The above descriptions all contain some allusion to either "combining" or "multiplying" a computer asset. By inserting configurative layers in between hardware and\or software components, more control can be exercised on one or both of the separated assets. This middle layer is usually some sort of software that allows for the manipulation of the assets through logical separation or combination. Therefore, this paper will define virtualization as the process of applying a software layer of abstraction in between various computing resources in order to logically separate or combine them.

3.9 The Impact of Virtualization

Adaptability is becoming an increasing focus for the management of modern business. With new opportunities and threats always lurking on the horizon, businesses must be able to quickly, efficiently, and effectively react to their dynamic environment. With regard to IT infrastructure, virtualization is perhaps the most effective tool for facilitating this adaptability. In virtualized systems, the expansion and reduction of technical resources can be performed seamlessly. Because physical devices and applications are logically represented in a virtual environment,

administrators can manipulate them with more flexibility and reduced detrimental effects than in the physical environment. Through the use of virtualization "tools," server workloads can be dynamically provisioned to servers, storage device usage can be manipulated, and should a problem occur, administrators can easily perform a rollback to working configurations. Generally, the addition (or removal) of hardware can be easily managed with virtualization tools.

The deployment of new applications for use across the enterprise is easily performed through varied combinations of application, operating system, and server virtualization. Through virtualization induced "containers," applications can be isolated from both the hardware and one another, preventing configuration conflicts that often complicate their introduction into IT systems.

Increased demand for data or database capabilities can be easily met with data and database virtualization through the management of new DBMSs or physical infrastructure with virtualization tools. All of these examples illustrate the adaptable nature of the virtual enterprise.

In addition to adaptability, the CIO can lower operating costs through implementing virtualization within his or her infrastructure. There is much inherent efficiency that comes with implementing this type of system, because much of its focus is on optimizing the use of resources, thus reducing overhead for maintenance. Any element of current infrastructure can be leveraged more fully with virtualization. Switching costs for new operating systems or applications are lowered with the ability to more flexibly install and implement them. The consolidation of servers and storage space obviously increases the return on investment for this hardware by maximizing efficiency.

Lowering costs will enable the CIO to reallocate the IT budget towards initiatives that are not related to the maintenance of current systems, such as research and development, partnerships, and the alignment of IT with business strategy. A case in which this is well illustrated is that of Welch Foods. Through virtualization, Welch's

IT management was able to increase server usage in their infrastructure from five to ten percent, to a range of 50 to 60 percent, allowing them to increase the number of servers per manager to 70-to-1[6], thus reducing expensive labor costs for day-to-day operations (Connor, 2005).

IT managers will be able to increase the productivity of employees across the entire organization through a properly implemented virtualization system. For businesses that rely on in-house application development, an increase in productivity and increased ease of implementation can be seen. Developers within a platform-virtualized environment can program in languages they are most proficient with. Debugging and testing applications becomes second nature with the ability to create contained virtual environments. In this instance, application and systems testing can be performed on a single workstation that employs a variety of virtual machines without the need to transfer and debug code to external computers. Enterprise-wide testing can be performed in isolated virtual machines, which do not interact with or compromise the resources actually being used on the network. Users in a virtualized environment do not know or care how their use of IT resources is being optimized. They are able to access needed information and perform work effectively and simply, without regard to the complexities that exist behind the scenes.

3.9 Security Impact

In virtualized environments in which resource segmentation takes place, an increase in security can be seen due to the residual complexities for hackers who are not familiar with the configuration of the system they wish to compromise. For example, in application virtualization, virtual applications can run on multiple servers, causing confusion for attackers who are prevented from determining the physical resource that has been compromised (Lindstrom, 2004). In the case of virtual machines, which emulate hardware systems, there can be added confusion for would-be attackers. "It's hard to accomplish much by cracking a system that doesn't exist" (Yager, 2004). Another security benefit brought about by virtualization is related to disaster

recovery. In virtualized server systems, it is not necessary to create identical configurations on backup servers as it is with non-virtualized systems. Because the virtualization layer separates the hardware from the operating system environment, restoring a lost server can be done on a machine with unlike hardware configurations. It is also possible to perform backups from several servers to one secondary server, creating a less expensive method for high availability and disaster recovery.

3.10 Risks

With any benefit, there is always associated risk. This is not an exception for the practice of virtualization. The first problem that IT managers must be aware of occurs in the planning and implementation of virtualization. CIOs and their staff must decide if, in fact, virtualization is right for their organization. The short-term costs of an ambitious virtualization project can be expensive, with the need for new infrastructure and configuration of current hardware. In businesses where cost reduction and flexibility of IT are not currently in alignment with the businesses strategy, other initiatives will be better suited. That is not to say that virtualization is not right for every environment, because most any organization can reap the benefits of a properly planned virtualization initiative. It is the timing and scope of such initiatives that must be scrutinized.

Another risk associated with virtualization can occur in businesses that do not have an efficient element of redundancy in their systems. Because the convergence of resources often takes place in virtualization environments, especially in that of server virtualization, the physical failure of one piece of hardware will impact all virtual elements that it manages. It is therefore necessary to ensure that backup systems are in place to deal with such problems. Fortunately, because of the isolation inherent in virtualized systems, backup processes can be greatly simplified.

A final problem that can occur in virtualized systems is increased overhead. The software layers inserted in between resources can chew up processor cycles,

sometimes up to double-digit percentages. Users and vendors say overhead can range from 2% or 3% to as high as 20%, depending on the product and application (Mitchell 2005). However, the efficiency that virtualization provides for hardware infrastructure should negate any problems associated with overhead when properly implemented.

With the positives far outweighing the negatives, virtualization is a technology that will soon be a universal practice. "Ultimately, virtualization will become just a standard layer of the infrastructure stack" (Mitchell, 2005). As costs for virtualization technology begin to decline, and more hardware manufacturers such as Intel and AMD begin to include built-in virtualization functionality in their products, it will become increasingly difficult to justify not using virtualization in an IT system. The unmatched effectiveness of virtualization to provide adaptability and reduce costs for the enterprise will empower IT managers and position their organizations for growth. Because of the inevitable induction of virtualization technology into the standard architecture stack, CIOs from all types of businesses should begin sketching the path to their future in virtualization

4. Case Study

4.1 Server Infrastructure Needed Updating

Bharat Petroleum Corporation Limited (BPCL), a leading petroleum refining and distribution company is India's second largest company by revenue. This Fortune 500 listed company markets engine oils, gasoline, liquefied petroleum gas (LPG), kerosene and industrial fuels to the domestic and industrial sector. It manages a network of more than 5,701 petrol stations and 1,980 LPG distributors across the country.

BPCL is known to be a technology-savvy company with an IT department that was formed in 1964. Today, BPCL has an IT team of more than 200 professionals comprised of ERP, infrastructure, developers and services IT professionals who support over 5,000 desktop users and 3,100 SAP users across 350 locations.

The company was using industry-standard tower servers more than five years old that occupied significant floor space in the data center. It became a challenge to manage the server sprawl. Deputy General Manager of BPCL, IIS Infrastructure explains, "Each of our applications was residing on its own dedicated server. There were 17 servers running 17 different applications. We wanted to consolidate all the applications onto a single server."

The company was also rapidly running out of server resources to support its testing and production requirements. "Each of our in-house developers needed his own

server for testing and development, and the number kept growing," says Deputy General Manager of BPCL. "We needed additional storage, RAM and processor speed to meet their requirements."

In August 2002, BPCL started to look for a solution. After an exhaustive search, Deputy General Manager of BPCL found that VMware ESX Server was the only robust and versatile virtual machine technology that would meet his needs, enabling server consolidation, centralized management and rapid application development through virtual machines. Soon after, BPCL looked for a local partner with the expertise and capability to deliver VMware solutions.

4.2 VMware ESX Server Proves Its Mettle in Simulation

As part of the selection criteria, Deputy General Manager of BPCL wanted to ensure that VMware ESX Server could successfully support Bharat Petroleum's multiple applications and work well in test and production environments. Deputy General Manager of BPCL emphasizes, "Our department plays a critical role in the company; we are responsible for providing the IT support to the various business units to enable them to carry out their respective activities efficiently and effectively at optimal costs. It is critical to maintain high service levels." their team wanted assurance that the VMware software would support the company's IT infrastructure 24 hours a day, seven days a week, with no extended downtime.

With the help of IBM India, which provided a 4-CPU server for the proof of concept exercise, their infrastructure services team attached 100 GB of external storage to VMware ESX Server for testing purposes in January 2003. They simulated the interoperability of migrating various applications by creating seven virtual machines with Windows 2000 server as the guest operating system, and tested applications such as Domain Controller, Exchange 2000 Server, SQL 2000, Web Server and IIS based application servers.

A group of system administrators and applications developers took one month to test the system rigorously for interoperability between virtual machines as well as

physical servers. The proof of concept was successful; VMware ESX Server dramatically reduced the set-up time for OS configuration, easily supported multiple applications and worked seamlessly with other software applications and hardware devices.

4.3 VMware ESX Server Wins over New Fans

Over a two month period, BPCL configured and set up the ESX Server and migrated the applications in phases to the new server. Explaining the cautious approach, Deputy General Manager of BPCL says, "Virtualization technology is still a very new concept. To give the team confidence that this really works, we led by example and started the migration of applications residing in servers for the infrastructure group. Once the applications were up and running smoothly, the application and services group began to use it extensively as they saw that ESX Server was working well."

Using ESX Server, BPCL now has over 17 virtual machines running on one physical machine, and plans to adopt a scale-up approach by adding more CPUs to further boost the processing power of its IBM X440 server. BPCL's in-house application team runs tests of new B2B and B2C services on the VMware-enabled server. The team also configured a two-node Microsoft Cluster server on VMware software and used it for testing new applications before porting on production cluster server.

More than 50 people in the IT team use ESX server locally and also access it remotely from BPCL's Chennai and Mumbai offices. According to Deputy General Manager of BPCL, the scalability and robustness of ESX Server is apparent from the number of users accessing virtual machines at any one point in time, which ranges from **five to 2,000** users, depending on the application.

VMware Software Enables BPCL to Scale with Ease

4.3.1 VMware ESX Server helps BPCL to meet both its business and technical goals, providing benefits, that include:

- **Server Consolidation**. Currently, BPCL is able to put 17 virtual machines on one server. With additional CPUs and RAM, the server can support up to 30 virtual machines, allowing BPCL to increase server and CPU utilization dramatically. "Overall, server consolidation was a major gain for us," Deputy General Manager of BPCL says. "It has reduced the population of standalone test and development servers. We are now able to virtualize our infrastructure to meet changing demands without additional investment."
- **Cost Savings**. "We've reduced the recurring cost for additional servers for items such as annual maintenance, power consumption, administration and manpower," Deputy General Manager of BPCL says. He estimates total cost savings to be approximately US$30,000.
- **Ease of Migration and Faster Server Deployment**. The migration process was smooth. By creating a master copy of its Windows 2000 server operating system, Deputy General Manager of BPCL says the IT team was able to dramatically reduce server set-up time from many hours to few minutes. "We now enjoy great flexibility in creating any number of virtual machines and assigning resources as per the individual requirement without delay," he says. "We use VMware software to reallocate resources dynamically as to meet ad-hoc requirements."
- **Faster Application Testing and Development**. VMware software gives BPCL the ability to quickly react to changing business needs. "The application developers and testers like the software because we can create a virtual machine immediately for them to work on," Deputy General Manager of BPCL says. "Before we had ESX Server, it would take up to three months to procure a server. Now, we don't have to wait for the hardware before starting our application development. Once the testing is completed, we migrate the virtual machine application to the actual production server environment." Another benefit is the

reusability of the virtual machine resources. Once a testing job is completed, the virtual machine is available for other purposes, thereby optimizing server resources.

4.4 Future Plans to Optimize Server Utilization with VMware Software

Pleased with the performance of ESX Server, Deputy General Manager of BPCL has plans to add more CPU-processing power to his current server. "Today, we are working in a single CPU environment," says Deputy General Manager of BPCL. "We are going to test some of the applications in a dual CPU environment. If that works out well, we intend to scale up to 8 CPUs, which gives us greater processing power and performance within the same server.

"This means further savings for us without the need to invest in another standalone server," he adds. "We also have plans to test Linux on ESX Server and other VMware solutions to discover new ways to maximize our resources. VMware software is definitely the solution of choice for maximum cost savings from server consolidation and for greater flexibility and scalability in rapid application testing and development."

4.5 VMware ESX Server at Work

- VMware ESX Server on IBM xSeries 440 2CPU server with 6GB RAM, connected to SAN
- IBM Fast T 900 SAN with two 16-port switches connected to 3 TB of total storage
- IBM Tivoli Storage Manager client on servers for virtual machine backup to IBM LTO Tape Library model no. 3584
- Guest operating system: Microsoft Windows 2000 Server
- Production applications include: Domain Controller, Exchange 2000 Server Instant Messaging, IBM Tivoli Gateway Servers, GAAP Application Server with SQL 2000

- Test applications include: IIS and SQL 2000, .Net Server with SQL 2000, IBM WebSphere, Etrust Antivirus 7.0, IBM MQSeries and EAI servers , IBM Tivoli Storage Manager

4.6 Virtual Infrastructure - VMware ESX Server

VMware ESX Server is virtual infrastructure software for partitioning, consolidating and managing servers in mission-critical environments.

Ideally suited for enterprise data centers, ESX Server minimizes the total cost of ownership of computing infrastructure by increasing resource utilization and its hardware-independent virtual machines encapsulated in easy-to-manage files maximize administration flexibility.

4.6.1 VMware ESX Server allows enterprises to:

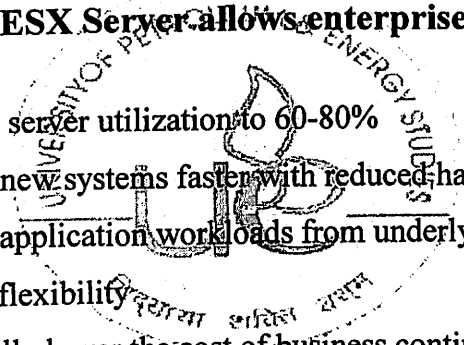

- Boost x86 server utilization to 60-80%
- Provision new systems faster with reduced hardware
- Decouple application workloads from underlying physical hardware for increased flexibility
- Dramatically lower the cost of business continuity

4.6.2 Use of VMware ESX Server

Implement Server Consolidation

Usage Scenarios

Consolidate branch office and data center mission-critical applications and infrastructure services onto fewer highly scalable, highly reliable enterprise-class servers. Virtualizing blade servers is an especially effective approach.

Benefits

- Improve utilization rates up to 60-80% for x86 servers
- Reduced TCO across computing infrastructure
- Built-in headroom for expansion and scaling
- Maximum hardware utilization
- Simplified system management
- Improve Flexibility and Responsiveness

Improve Flexibility and Responsiveness

Usage Scenarios

Use Virtual Center to deploy and manage ESX Server Virtual Infrastructure Nodes and manage the hardware resources as a single pool of compute, storage and networking power. Dynamically load balance workloads across the pool, and avoid planned downtimes.

Benefits

- Provisioning times for new applications measured in tens of seconds, not days
- Response times for change requests measured in minutes
- Zero-downtime hardware maintenance without waiting for maintenance windows

Deliver High Availability and Guarantee Service Levels

Usage Scenarios

Protect critical data in secure virtual machines and isolate multiple servers that run together at near-native performance levels on standard x86-based hardware

Benefits

- Protection against non-hardware errors and single point of failure for higher availability
- Run resource-intensive SMP applications such as Oracle, SQL Server, Microsoft Exchange server, SAP, Siebel, Lotus Notes, BEA WebLogic and Apache More control over IT performance metrics
- Run IT as an enterprise service provider, delivering better service levels to customers

Dramatically Lower the Cost of Disaster Recovery Capability

Usage Scenarios

Create a unified disaster recovery platform that increases availability, reduces recovery time and lowers hardware and operational costs

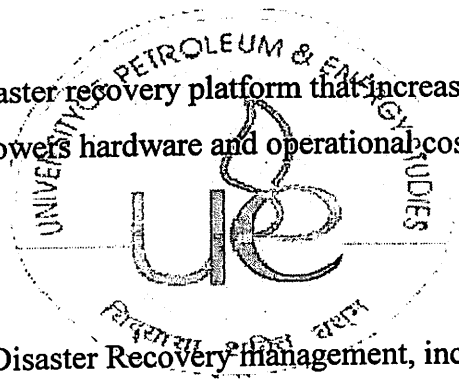

Benefits

- Streamline Disaster Recovery management, increase availability, reduce recovery time and lower hardware and operational costs
- Eliminate the need for costly 1-to-1 mapping of production and Disaster Recovery servers
- Recover virtual machine images on any x86 hardware platform

4.7 Bharat results

- 17:1 server consolidation ratio
- Created efficient test, development, integration and production environment for in-house applications
- Cut server deployment time from hours to minutes
- Improved CPU utilization rate by 400 percent
- Reduced annual hardware expenditures
- Sped up application testing and deployment from months to weeks
- Improved time –to – market of B2B and B2C in-house applications
- Saved physical space and cooling costs
- Created scalable infrastructure to support more virtual machines with additional CPUs

5. Conclusion

Building and connecting every aspect of IT landscape is very important for any organization. Doing more with less i.e., doing the expected tasks with fewer infrastructures.

Today's scenario for IT expenditure for any industry is shrinking. So the systems are expected to deliver more per the amount spent on it. This will in turn reduce the cost involved in maintenance and also reduction in running cost to carry out these things this document suggest some of the concepts like Consolidation(servers, applications, shared devices, network devices), sever sizing, Virtualization.

Virtualization is the technique of managing and presenting storage devices and resources functionally, regardless of their physical layout or location. On virtualization this document analyzes the concept of server, operating system, application, storage, database and networks. And also it covers what are the impacts and risks in carrying out it.

Server sizing is one of the important concept being carried out for scalability. To configure a server for a particular load is very important. For any application availability is very important to achieve the business goal. If the server could not handle the workload, it would cost more .so it is very important to configure the server for particular workload keeping things in mind for future. This document asserts various methods of predicting the future workload like guessing and prototyping.

A case study has been taken to analyze the case. This case deals with how consolidation was carried out for doing more with less by installing VMware. The benefits which BPCL got is 17:1 server consolidation ratio, 17:1 server consolidation ratio, Created efficient test, development, integration and production environment for in-house applications, Cut server deployment time from hours to minutes, Improved CPU utilization rate by 400 percent ,Reduced annual hardware expenditures, Sped up

application testing and deployment from months to weeks ,Improved time –to – market of B2B and B2C in-house applications ,Saved physical space and cooling costs ,Created scalable infrastructure to support more virtual machines with additional CPUs .

6. References

6.1 Articles

1. Clark Tom ,Storage Virtualization , Technologies for Simplifying Data Storage and Management ,[2005]
2. IBM Redbooks, Storage Networking Virtualization: What's It All About?, [2000]
3. IBM Redbooks, IBM Virtualization Engine Version 1 Planning and Installation guide. [2005]
4. Jacobs David, Joel Samuel Yudken ,The Internet, Organizational Change and Labor: The Challenge of Virtualization[2003]
5. Muller Al,Advanced Server Virtualization: VMware and Microsoft Platforms in the Virtual Data Center,[2006]
6. Scott Michael Baker, Server virtualization[2006]
7. W. Luke Timothy, Vectors of Virtualization[2006]

6.2 Books

1. IBM Redbooks, Advanced Power Virtualization on IBM P5 Servers: Architecture And Performance Considerations[2005]
2. IBM Redbooks, Server Consolidation with the IBM eServer xSeries 440 and VMware ESX Server [January 21, 2003]
3. Kumar Reddy, Victor Moreno, Network Virtualization [2006]
4. Marshall David,_Advanced Server Virtualization: VMware and Microsoft Platforms in the Virtual Data Center [2006]
5. Orenstein Gary, IP Storage Networking: Straight to the Core-page-75[Computres-2003]
6. Ulf Troppens, Rainer Erkens, Wolfgang Müeller, Storage Networks Explained: Basics and Application of Fibre Channel SAN, NAS ISCSI and InfiniBand [2004]
7. Venugopal.v Reddy (ed.) Manjulika S. (ed.), Towards Virtualization: Open and Distance Learning, Publisher: Kogan Page India [2002].

8. Wolf Chris , Erick M Halter, Virtualization: From the Desktop to the Enterprise, [2005]

6.3 Web

1. www.itp-journals.com
2. www.sun.com/blueprints/online.html
3. www.teamquest.com/datacenter/server-consolidation
4. www.ibm.com.
5. www.ibm.com/servers/eserver/pseries
6. www.ibm.com/servers/eserver/pseries/linux.
7. www.vmware.com
8. www.dell.com/oracle
9. www.dell.com/servers
10. www.sun.com/servers
11. www.sun.com/storage
12. www.docs.hp.com
13. www.software.hp.com
14. www.polyserve.com/pdf/Petroleum_case_study.pdf
15. www.panasas.com