An Analysis of System Performance using VMware ESXi Server virtual machines

Thesis submitted in partial fulfillment of the requirements for the award of degree of

Master of Engineering
in
Software Engineering

Submitted By
Hiteshi
(801031011)

Under the supervision of:
Dr. V.P.Singh
Assistant Professor (CSED)

COMPUTER SCIENCE AND ENGINEERING DEPARTMENT
THAPAR UNIVERSITY
PATIALA – 147004

June 2012

Certificate

I hereby certify that the work which is being presented in the thesis entitled, "*An analysis of System Performance using VMware ESXi Server virtual machines*", in partial fulfilment of the requirements for the award of degree of Master of Engineering in *Software Engineering* submitted in Computer Science and Engineering Department of Thapar University, Patiala, is an authentic record of my own work carried out under the supervision of ***Dr. V.P.Singh*** and refers other researcher's work which are duly listed in the reference section.

The matter presented in the thesis has not been submitted for award of any other degree of this or any other University.

Signature
Hiteshi

This is to certify that the above statement made by the candidate is correct and true to the best of my knowledge.

(Dr.V.P. Singh)
Assistant Professor
CSED

Countersigned by

(Dr. Maninder Singh)
Head
Computer Science and Engineering Department
Thapar University
Patiala.

(Dr.S.K.Mohapatra)
Dean(AcademicAffairs)
Thapar University
Patiala.

Acknowledgement

I would like to express sincerest thanks to my thesis supervisor Dr. V.P. Singh, Assistant Professor, Computer Science and Engineering Department for his inspiration, guidance, stimulating suggestions, immense help and support throughout the period of this research work. He has provided me with all the necessary resources including motivation and research environment without which it would not have been possible to complete this work. It was a great opportunity for me to do this work under his supervision.

I would like to thank Dr. Maninder Singh (Head), Computer Science and Engineering Department for his moral support and the research he had facilitated for this work.

I would also like to thank all my teachers for their stimulating discussions and invaluable support I received during this period of research. I am thankful to the authors whose work i have consulted and quoted in this work.

Finally, I wish to thank my dearest family and friends for all their immense love, enthusiastic encouragement and support throughout my life without which it would not have been possible to complete this work. Last but not the least I would like to thank God who has always been with me in my good and bad times.

Hiteshi

Hiteshi

(801031011)

Abstract

Virtualization refers to the abstraction of computer resources. The main purpose of virtual environment is to improve resource utilization by consolidating the operating platforms for users and applications. Recently, virtualization at all levels i.e. system, storage, and network; becomes important as a way to improve system security, reliability, reduce costs and provide greater flexibility.

System-level virtualization has been regained popularity during the past few years because of the availability of efficient solution such as Xen, VMware, Hyper-V and the implementation of hardware support in processors (e.g. Intel-VT, AMD-V). By interpreting the machine, one is able to run a variety of operating systems and environments as needed by the applications. Virtualization allows users to isolate workloads, improving security and reliability and also possible to balance workloads, use migration techniques to relocate applications from failing machines, and isolate fault systems for repair. This thesis work presents an efficient implementation of a system-level performance using virtual machine hypervisor VMware ESXi server.

This technique allows collection of essential workload characteristics for arbitrary and unmodified operating system instances running in virtual machines. For capturing the result efficiently, a command line interface is used which encompasses essential performance metrics including memory usage, CPU usage, disk usage and overall system usage, which further can be represent graphically using esxplot tool. The demonstration of technique can be done by using a benchmark tool to generate workload.

Table of Contents

List of Figures

List of Tables

1. Introduction

Virtualization typically involves using special software to safely run multiple operating systems and applications simultaneously with a single computer. The technology initially allows company to consolidate an array of servers to improve operating efficiency and reduce costs. It has since been applied to dealing with data storage as well as desktop systems. Owning to the success of tools developed by VMware Inc., the technology has become one of the most talk-about technology, and has draw attention from both IS professionals and non-IS executives in virtually all industries. Despite the potentially significant impact on company's operations, this technology has virtually been ignored by the academic researchers.

In today's business environment, it is clear that technologies such as virtualization and multicore are particularly important enablers for the consolidated IT infrastructure IT organizations are increasingly seeking to deploy. Each of these technologies is impactful to the market in their own right. However, the use of multi-core technology in conjunction with server virtualization tools has a compounding impact on server configurations, and accelerates the ability of IT organizations to exploit the benefits of multi-core technology. Unlike other previous multicore introductions that took time to become main-stream as customers changed their application code; virtualization allows customers to fully exploit the improvements in x86 processors immediately, accelerating business benefits and thereby increasing adoption rates. Looking forward, International Data Corporation (IDC) [1] believes the server and component vendors will optimize around quad-core technology before moving ahead to octi-core technology.

The need for increased efficiency in managing the server environment, controlling hardware costs, and repurposing hardware for other needs cited by IT respondents makes it clear that IT organizations today need to continue their investment in virtualization technology in order to improve the performance and efficiency of the IT infrastructure. Virtualized infrastructures help simplify IT operations in many ways, from helping to

shield software from hardware, to enabling secure resource sharing, to facilitating software deployment and relocation.

They also increase business agility by:

- Enabling IT staff to dynamically reallocate resources as needed to avoid planned downtime
- Enhancing the efficiency of application testing and development
- Facilitating rapid, cost-effective disaster recovery

This thesis is organized into 6 chapters which include Introduction; Background Information; Literature Review; Problem Statement; setup and implementation details; experimental results and finally Conclusion and Future Scope.

Chapter 1 describes virtualization in general in terms, its background information, basic terminologies and technologies, its architecture and its types. In Chapter 2, discusses survey analysis and current available products in the market. Chapter 3 discusses the problem statement and tasks. Chapter 4 discusses the installation of tools and the simulation environment. Chapter 5 describes the results, evaluates the performance, and analysis and finally Chapter 6 summarizes the conclusions drawn in the thesis along with future research direction.

1.1 Background

Virtualization was first developed in 1960's by IBM Corporation [2], originally to partition large mainframe computer into several logical instances and to run on single physical mainframe hardware as the host. This feature was invented because maintaining the larger mainframe computers became cumbersome. The scientist realized that this capability of partitioning allows multiple processes and applications to run at the same time, thus increasing the efficiency of the environment and decreasing the maintenance overhead [3].

IBM's first professor and the programmers were deeply involved in the first mainframe virtualization effort. According to a Network World [4], IBM's CP-67 software used partitioning technology to allow many applications to be

run at once on a mainframe computer. While its impact was substantial for mainframe users, it took years before a direct descendant of IBM's work came back to life on X86 platforms. In fact the existence of virtualization as a concept went largely unremarked for the nearly two decades of the rise of client/server on x86 platforms. Still, it served as a powerful inspiration for VMware to reviving the concept and applies it to x86 machines.

In the late 1990s,[5] the issues that made caused IBM to application virtualization on its mainframes had begun to impact IT administrators substantially enough for VMware to step in and begin to apply its own virtualization model. These included low x86-platform server utilization, where perhaps 10-15% of server capacity was used, and rising costs associated with electrical power use, cooling and a fast-growing server and storage footprint.

With increased complexity came expanding administrative costs driven by the need to hire more experienced IT professionals, and the need to carry out a wide variety of tasks, including Windows Server backup and recovery. These actions required more manual intervention in processes than many IT budgets could support. Server maintenance costs, especially those tied to Windows Server backup, were climbing and more personnel were required to work through an increasing number of day-to-day tasks. Just as important were the issues of how to limit the impact of server outages, improve business continuity and create more robust disaster recovery plans.

Two primary benefits offered by any virtualization technology are:

- **Resource sharing** - Unlike in non-virtualized environment where all the resources are dedicated to the running programs, in virtualized environment the VMs shares the physical resources such as memory, disk and network devices of the underlying host.

- **Isolation** - One of the key issues in virtualization provides isolation between virtual machines that are running on the same physical hardware. Programs running in one virtual machine cannot see programs running in another virtual machine.

1.2 Basic Terms of virtualization

Virtualization is commonly defined as a technology that introduces a software abstraction layer between the hardware and the operating system and applications running on top of it. This abstraction layer is called virtual machine monitor (VMM) or hypervisor and basically hides the physical resources of the computing system from the operating system (OS). Since the hardware resources are directly controlled by the VMM and not by the OS, it is possible to run multiple (possibly different) OSs in parallel on the same hardware. As a result, the hardware platform is partitioned into one or more logical units called virtual machines (VMs).
The following requirements for a VMM have been defined [6]:

- **Equivalence:** Running an application inside a virtual machine must be equivalent to running the same application on the underlying hardware.

- **Control:** The VMM must control and synchronize the access of VMs to hardware resources.

- **Isolation:** VMs must be isolated from each other with the purpose of ensuring stability (the crash of a VM should not affect another VMs), security (a possibly compromised VM shouldn't grant access to other VMs) and data consistency.

- **Performance:** The performance overhead caused by virtualization should be minimal, close to "bare metal" performance.

- **Encapsulation:** VMs must exist in form of a file or a directory of files which allows easy migration or cloning of the VM.

1.3 Virtualization Technologies

Virtualization can be done in several ways. There are various virtualization technologies available in the market that helps to virtualize the environment. Depending on the needs and the goals of the organization, one virtualization

technology is better is than the other. The following section gives an overview of some of the existing virtualization technologies. Before going into the details of the different virtualization technologies, basic idea of virtualization is given in Fig.1.1

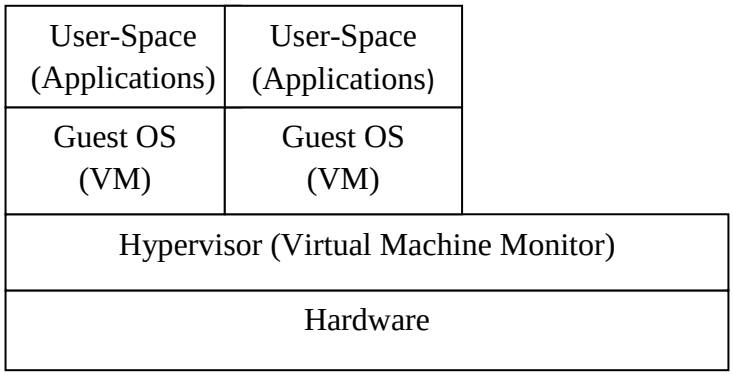

Fig.1.1 The layered abstraction of virtualization.

In the above Fig.1.1 [7] there are two virtual machines running on top of a physical computer processing their own operating system and application. Every guest machine appears to be an independent computer for their running processes. Hypervisor layer is the host software layer that provides the ability to run multiple operating systems on a physical hardware.

Every virtualization technology abstracts a computing resource in some way to make it more useful. Whether the thing being abstracted is a computer, an application's user interface, or the environment that application runs in, virtualization boils down to this core idea. Since, all of these technologies are important, it's fair to say that hardware virtualization gets the most attention today.

1.3.1 Full Virtualization: In this approach, the VMM is called the virtual machine manager that runs on the top of the operating system, commonly as an application in user space. The result is that, in the VMs, the applications and the guest operating systems run on top of a virtual hardware provided by the VMM. However, the virtual machine environment that provides enough representation of the underlying hardware to allow guest operating systems to run without modification can be considered to provide "Full Virtualization"

[8]. In full virtualization setup, I/O devices are allotted to the guest machines by imitating the physical devices in the virtual machine monitor: interacting with these devices in the virtual environment is then directed to the real physical devices either by the host operating system driver or by the "hypervisor driver". This architecture can be observed as below in Fig. 1.2

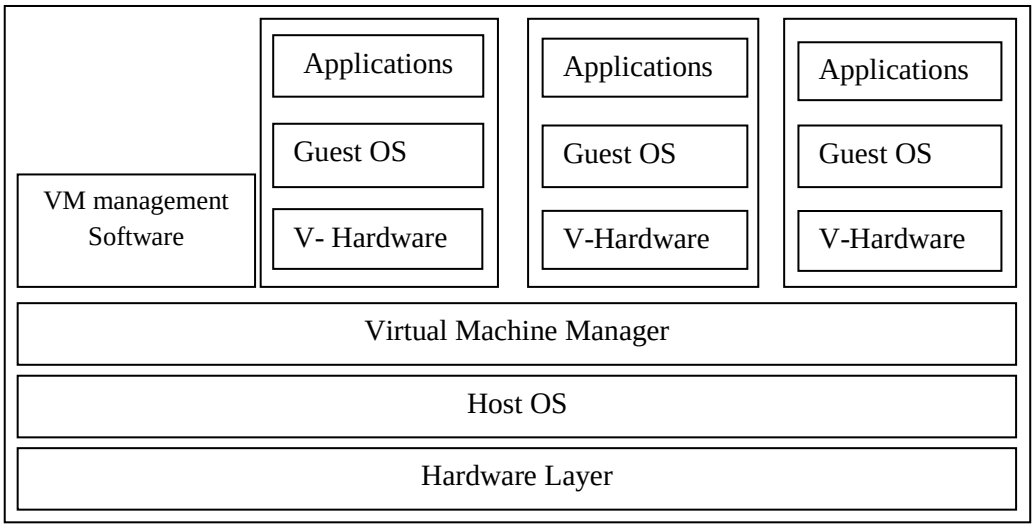

Fig. 1.2 Full virtualization

Main **advantage** of this approach is that, it is very easy to use. A common user can install a software product like VMware workstation just like any other software product on its own OS. Inside VMware workstation, a guest OS can be installed and used just like it would be running on directly hardware.

The main **disadvantage** of this approach is the poor performance, which can be less than when running directly on hardware.

1.3.2 Para virtualization:

Para virtualization is the technique used by Xen which provides a virtual machine interface representing a slightly modified copy of the underlying hardware, where the non-virtualizable portions of the x86 original instruction set are replaced with their easily virtualized equivalents. Unlike full virtualization, in para virtualization the running guest OS should be modified in order to be operated in virtual environment. Para virtualization is the subset of server virtualization, which provides a thin software interface between the host hardware and the modified guest OS. An interesting fact in this

technology is that the guest machines are aware of the fact that they are running in virtualized environment. One of the main characteristics of the para virtualization technology is, the virtual machine monitor is which simply allows para virtualization to achieve performance closer to non virtualized hardware. Device interaction in para virtualized environment is very similar to the device interaction of full virtualized environment; the virtual devices in para virtualized environment also rely on physical device drivers of the underlying hardware [9].

1.3.3 Hardware layer Virtualization: This approach is commonly used on the server market to due to its high virtual machine isolation and performance. Here, the VMM runs directly on hardware, controlling and synchronizing the access of the guest OSs to the hardware resources. Fig.1.3 depicts the architecture.

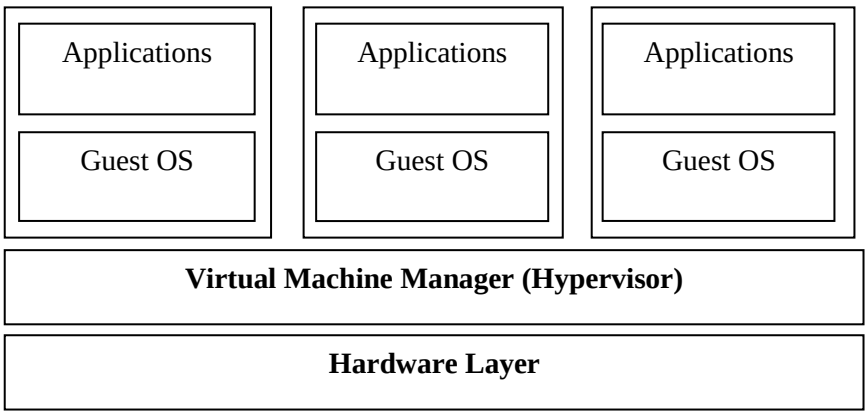

Fig. 1.3 Hardware Virtualization

VMware ESX Server and Xen are the main two competing VMMs that use this approach. Since x86 architecture was not developed with virtualization in mind, and this is not what Goldberg would refer to as a new virtualizable architecture, new techniques were developed to implement CPU virtualization. Microsoft also launches some of its products for hardware virtualization.

1.3.4 Application Virtualization: In application virtualization, the user is able to run a server application locally using the local resources without needing the complexity on completely installing this application on his/her

computer. Such virtualized applications are designed to run in a small virtual environment containing the only resources needed for the application to execute. Thus, in application virtualization each user has an isolated application environment virtually. This small isolated virtual environment acts as a layer between the application and the host and the host operating system. Application virtualization in shown in Fig.1.4

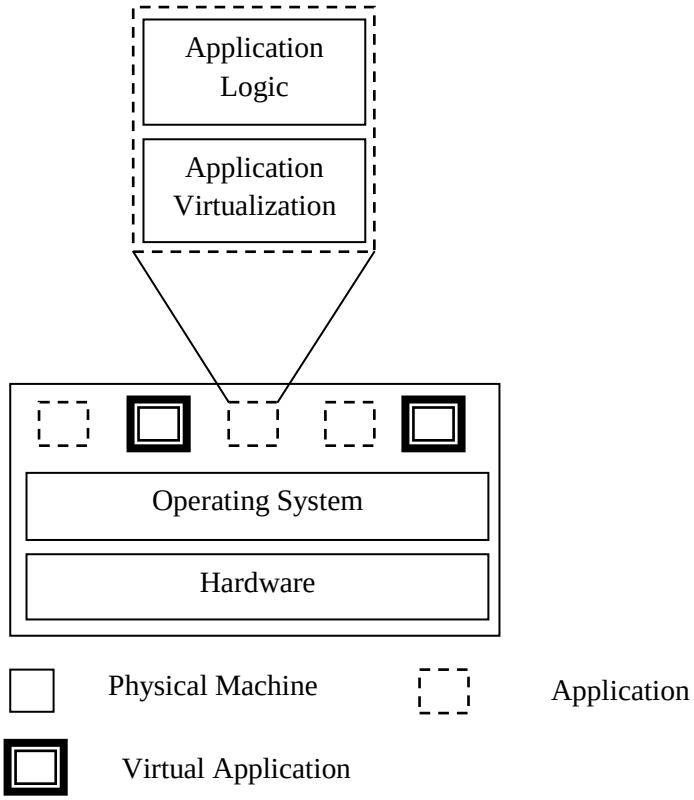

Fig.1.4 Application Virtualization

As an **organization-level** [10], there are three main types of virtualization for all important business tasks are:

1.3.5 Storage Virtualization: Storage virtualization is a form of Resource Virtualization, where a logical storage is created by abstracting all the physical storage resources that are scattered over the network. First the physical storage resources are aggregated to form a storage pool which then forms the logical storage. This logical storage which is the aggregation of scattered physical resources appears to be a single monolithic storage device to the user. In storage virtualization, several physical storages are multiplexed into a single

logical storage. Storage virtualization in its most basic form is used with RAID implementations, where two or more physical hard drives are combined to a single logical hard disk to provide data redundancy. Bigger scale implementations of storage virtualization are Storage Area Network (SAN) and Network Attached Storage (NAS) technologies. With the help of storage virtualization workload of a single hardware can be uniformly distributed to multiple ones. Storage virtualization also provides fault tolerance as data can be easily replicated to multiple physical locations throughout the network. [11]

1.3.6 Network virtualization: It is the creation of a virtualized network addressing space within or across network subnets i.e. it combines the computing resources in a network by splitting the available bandwidth into independent channels that can be assigned to a particular server or device in real-time. Network virtualization is used either to separate multiple virtual networks from a single physical network, or to provide networks within virtual environments without using any physical networking devices. Most common approaches to network virtualization are Virtual Local Area Network (VLAN), Virtual IP, Virtual Private Network (VPN) and virtual networking inside a virtualization host [11]. Virtual Local Area Networking divides a single physical network into several independent logical networks called VLANs. Each VLAN has its own identification tag and is located on its own network segment. This way the VLANs can be securely separated from each other even if they use the same physical network switch. VLAN technology provides good traffic flow management, security and easy network administration. VLANs can also be used in virtual networks within VMware Infrastructure [11].

1.3.7 Server virtualization: The main area of virtualization which hides the physical nature of the server resources, including the number and identity of individual servers, processors and operating systems, from the software running on them.

This last category i.e. server virtualization is far and away from the most common application of today's technology, and is widely considered the primary driver of the market. When most people say about the term "virtualization", they are likely talking about the server virtualization.

1.4 Hypervisor: A hypervisor, also called a virtual machine manager, is a program that allows multiple operating systems to share a single hardware host. Each operating system appears to have the host's processor, memory, and other resources all to itself. However, the hypervisor is actually controlling the host processor and resources allocating, what is needed to each operating system in turn and making sure that the guest operating systems (called virtual machines) cannot disrupt each other or we can say, It's the software program or part of the code in firmware that manages either multiple operating systems or multiple instances of the same operating system on a single computer system. The hypervisor's job is to manage the system's processor, memory and other resources to allocate what each operating system requires. Hypervisors provide the means to logically divide a single, physical server or blade, allowing multiple operating systems to run securely on the same CPU and increase the CPU utilization. Where hardware partitioning allows for hardware consolidation, hypervisors allow for flexibility in how the virtual resources are defined and managed, making it a more-often used system consolidation solution.

1.4.1 Types of Hypervisors

IBM breaks hypervisors down into two different types [12]:

a. Type 1 hypervisor
b. Type 2 hypervisor

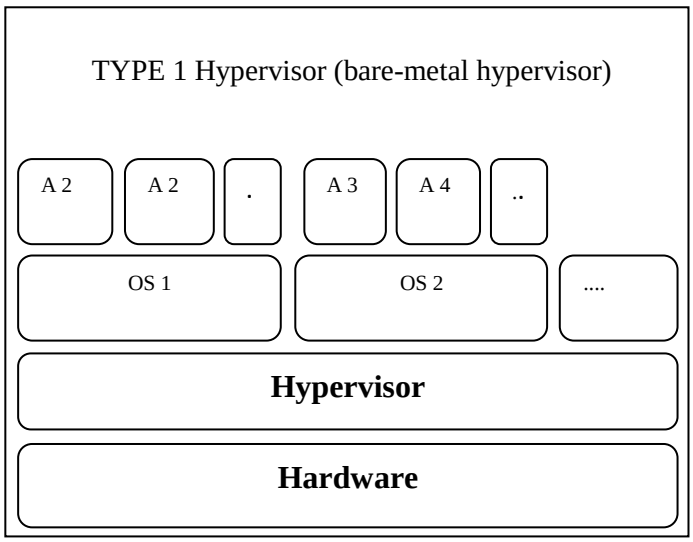

Fig. 1.5 Type 1 hypervisor

Type 1 hypervisors are also called Bare-metal hypervisor and are those that run directly on the system hardware and offer a higher level of virtualization and hardware security as shown in Fig. 1.5

Type 2 hypervisors are those that run on a host operating system that provides virtualization services, such as I/O device support and memory management. These are used mainly on client systems where efficiency is less critical, and are also commonly used for systems where support for a broad range of I/O devices is needed and can be provided by the host operating system. Fig. 1.6 shows type 2 hypervisors

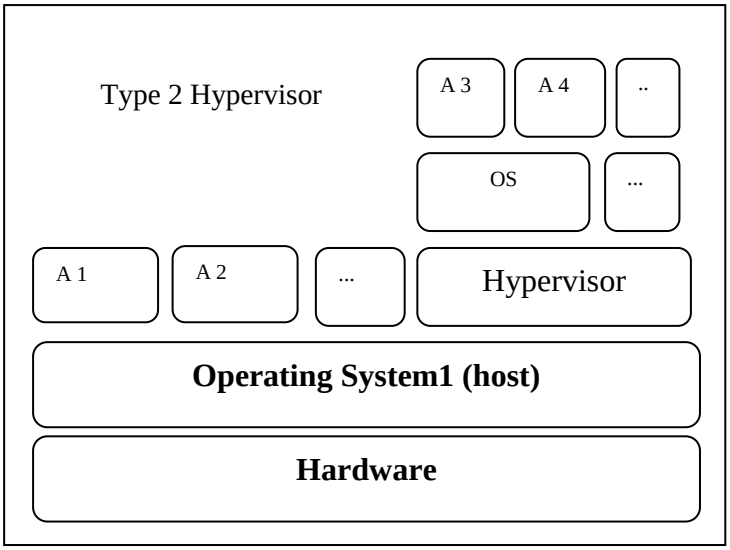

Fig.1.6 Type 2 Hypervisor

1.4.1 Hypervisor Architecture

ESX Server is a type 1 hypervisor that creates logical pools of system resources so that many virtual machines can share the same physical resources. ESX Server is an operating system that functions like a hypervisor and runs directly on the system hardware. ESX Server inserts a virtualization layer between the system hardware and the virtual machines, turning the system hardware into a pool of logical computing resources that ESX Server can dynamically allocate to any operating system or application. The guest operating systems running in virtual machines interact with the virtual resources as if they were physical resources. The hypervisor architecture of VMware vSphere 5.0 plays a critical role in the management of the virtual

infrastructure. ESXi runs independently of a host operating system (OS) and improves hypervisor management in the areas of security, deployment and configuration, and ongoing administration.

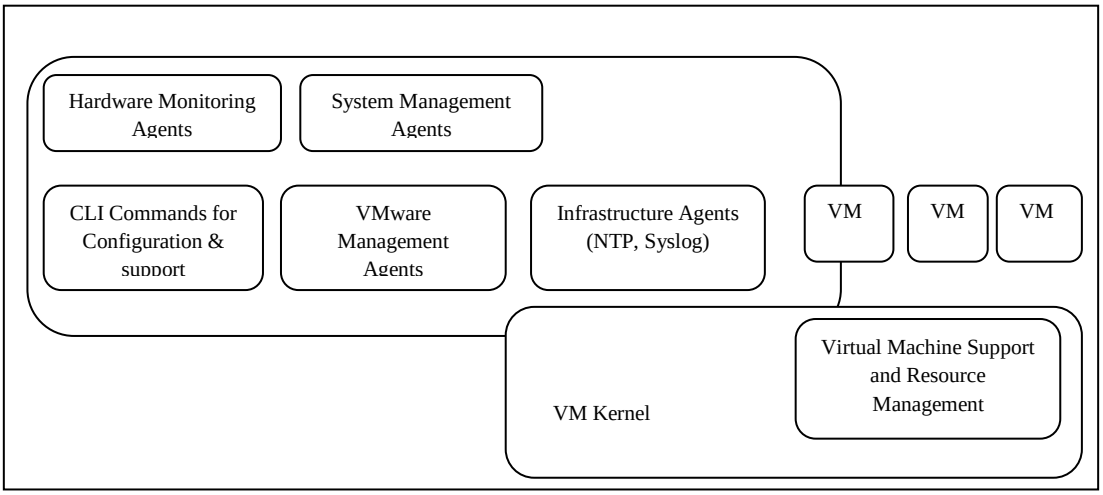

Fig. 1.7 Architecture of ESX

In the original ESX architecture [13] (Fig. 1.7), the virtualization kernel (VMkernel) is augmented by a management partition known as the console operating system (COS) or service console.

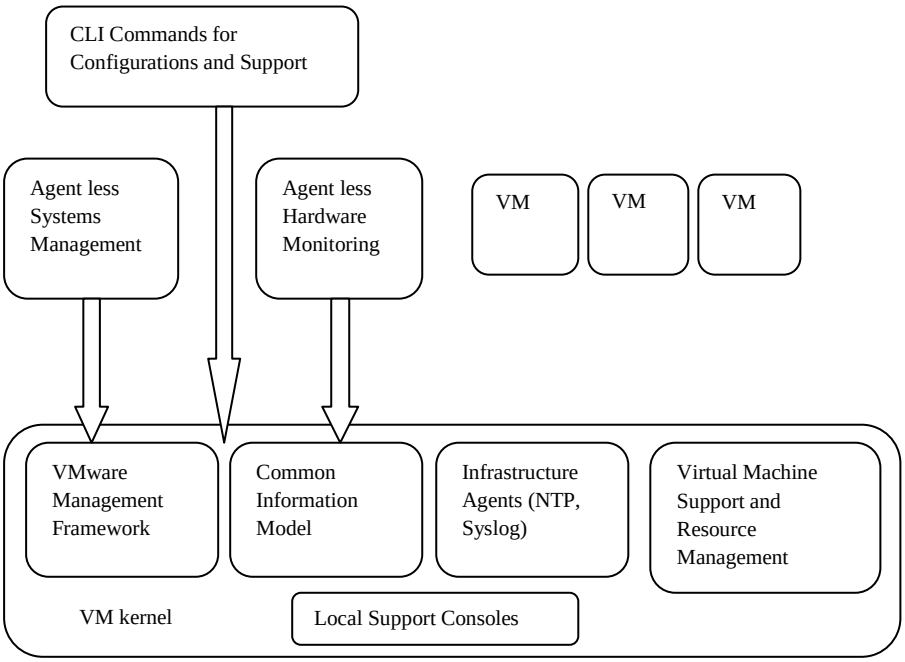

Fig. 1.8 Architecture of ESXi

In the ESXi Architecture, the COS has been removed, and all of the VMware agents run directly on the VMkernel. Infrastructure services are provided

natively through modules included in the VMkernel. Other authorized third-party modules, such as hardware drivers and hardware monitoring components, can run in the VMkernel as well. Only modules that have been digitally signed by VMware are allowed on the system, creating a tightly locked–down architecture. Preventing arbitrary code from running on the ESXi host greatly improves the security and stability of the system, as shown in Fig. 1.8 [14].

The **key components** of the ESX Server architecture are:

- **ESX Server virtualization layer:** Separates the underlying physical resources from the virtual machines.
- **Resource manager:** Creates virtual machines and delivers processing units, memory, network bandwidth, and disk bandwidth to them. It efficiently maps the physical resources to the virtual resources.
- **Service console:** Controls the installation, configuration, administration, troubleshooting, and maintenance of the ESX Server. The service console resides in its own virtual machine. ESX Server automatically configures the service console virtual machine when you install ESX Server. The service console also provides a place to install systems software.
- **Hardware interface components (including device drivers):** Delivers hardware-specific services while hiding hardware differences from other parts of the system.

1.4.2 Current Players of Hypervisor in Market:

There are many organizations which, now a days, building their own virtualization products such as, Microsoft, VMware inc. etc. some of them can be describe briefly below.

a) VMware: VMware's virtual machine (VM) approach creates a uniform hardware image — implemented in software— on which operating systems and applications run. On top of this platform, [11] [15] VMware's VirtualCenter provides management and provisioning of virtual machines, continuous workload consolidation across physical servers and VMotion technology for virtual machine mobility. VMware offers a wide range of products for data center infrastructure, workstations, enterprise desktops,

virtualization accelerators, as well as several free virtualization products. **VMware Server** is a hosted hypervisor that is available free of charge. Installation of the hypervisor must be done on top of either Windows or Linux host operating system. Range of supported operating systems is quite good and VMware Server also supports 64-bit guests. VMware Server is a lightweight solution that cannot compete with bare-metal hypervisor solutions in performance but is useful, for example as a testing environment or legacy hardware elimination. VMware Server is a great product to get started with virtualization as it is free and does not set any hardware requirements [15] [11].

VMware ESX 5.x will be the last major version of ESX with the service console included; many are starting to look toward the future and the role of ESXi in their infrastructures. The market provides several organizations and companies that each offer their own approach to x86 server virtualization.

VMware ESX Server is a bare-metal hypervisor that provides a wide range of supported host hardware and guest operating systems. VMware ESX Server products have a very rich set of features and guest operating systems can run at near native speed. Many different versions of ESX Server are available, from free of charge ESXi to enterprise level Virtual Infrastructure Enterprise. Basic instructions are the same in all versions but features like high availability and live migration of guest operating systems are only available in the Enterprise version. ESX Server has the widest range of supported guest operating systems including most of the 64-bit operating systems available. ESX Server can be used in enterprise production environments for testing, development, enterprise server consolidation and high availability systems. Different versions of VMware ESX Server products and additional features will be discussed in chapter 2.

b) Xen: Xen is a virtual machine monitor (VMM) for x86-compatible computers. Xen is built to securely execute multiple virtual machines, each running its own operating system, on a single physical system with close-to-native performance. Xen is open source, and is released under terms of the GNU General Public License. Xen is a Type 1 hypervisor that runs directly on the system hardware. Xen originated as a research project at the University of Cambridge. Citrix XenServer is bare-metal hypervisor that is based on an

open-source Xen hypervisor. XenServer can host Microsoft Windows and Linux operating systems, both 32-bit and 64-bit, but will require Intel VT or AMD-V on the underlying hardware in order to host any Microsoft Windows operating system. Like VMware ESX Server, Citrix XenServer offers many different editions for different situations. Performance of guests hosted by XenServer is near native. Similarly to its competitors, XenServer can be implemented on enterprise production as well as testing and development environments. [15] [16].

c) Microsoft Virtual Server: Windows Server virtualization, as a part of Microsoft's "Longhorn" server, takes a big step forward in bringing some of the advanced capabilities of virtualization to bear and providing customers with a scalable, secure and highly available virtualization platform.

Microsoft Virtual Server 2005 R2 is cost-effective server virtualization technology engineered for the Windows Server System platform.

Microsoft Hyper-V Server 2008 R2 is the hypervisor-based server virtualization product that allows you to consolidate workloads onto a single physical server. It is a stand-alone product that provides a reliable and optimized virtualization solution enabling organizations to improve server utilization and reduce costs. [17] Microsoft Hyper-V is a bare-metal hypervisor that competes with VMware ESX Server and Citrix XenServer.

Hyper-V comes shipped together with Microsoft Server 2008 and it targets networks from workgroups to enterprises. Hyper-V offers a moderately good range of supported guest operating systems, including 64-bit versions of Windows operating systems and a few Linux platforms. Hyper-V cannot be installed on old systems as it requires hardware that supports Intel VT or AMD-V. It offers a near native performance to the supported guest operating systems. [18] Hyper-V can be used in workgroups as well as in enterprise production environments, for example in large scale testing, enterprise server consolidation or high availability systems. As Hyper-V comes free with the new [19] Microsoft Server 2008, it is a notable alternative in the bare-metal hypervisor market although it does not provide the flexibility or performance of its more expensive competitors. [20][15].

d) Virtual Desktop Infrastructure (VDI): The VMs that Hyper-V provides can be used in many different ways. Using an approach called Virtual Desktop Infrastructure, for example, Hyper-V can be used to run client desktops on a server.
VDI runs an instance of Windows Vista in each of Hyper-V's child partitions (i.e. its VMs). Vista has built-in support for the Remote Desktop Protocol (RDP), which allows its user interface to be accessed remotely. The client machine can be anything that supports RDP, such as a thin client, a Macintosh, or a Windows system.

e) Intel & AMD Offer Virtualization: Industry heavyweights like Intel and AMD have also given virtualization a huge credibility boost. For example, the two chip vendors are building virtualization capabilities into their chip architectures — Intel with its Intel Virtualization Technology (VT) and AMD with its AMD-Virtualization (AMD-V), on Xeon and Opteron processors, respectively. From an application standpoint, the hardware will enable applications that have previously been hard to virtualized (e.g., I/O-intensive apps like database applications for which the overhead has been too high) and to be virtualized much more successfully. As Intel's and AMD's technologies are introduced, it will offer a choice of which virtualization path to explore Windows Server that of VMware, the market leader; open source software from Xen-Source.

1.5 Security Vulnerabilities in Virtualization:

Most of the security flaws identified in a virtual machine environment are very similar to the security flaws associated with any physical system. Some of the security flaws that are unique to the virtual environment [21] :

1.5.1 Communication between VMs or between VMs and Hosts:

One of the primary benefits that virtualization bring is isolation. This benefit, if not carefully deployed become a threat to the environment. Isolation should be carefully configured and maintained in a virtual environment to ensure that the applications running in one VM do not have access to the application running in another VM. Isolation should be strongly maintained that break-in into one virtual machine should not provide access either to virtual machines

in the same environment or to the underlying host machines. In some VM technologies, the VM layer is able to log keystrokes and screen updates across the virtual terminals, provided that the host operating system kernel has given necessary permission. These captured logs are stored out in the host, which creates an opportunity to the host to monitor even the logs of the encrypted terminal connections inside the VMs. Some virtualization avoids isolation, in order to support applications designed for one operating system to be operated on another operating system, this solution completely exploits the security bearers in both the operating systems. This kind of system, where there is no isolation between the host and the VMs gives the virtual machines an unlimited access to the host's resources, such as file system and networking devices in which case the host's file system become vulnerable.

1.5.2 VM monitoring from the host: Host machine in the virtual environment is considered to be the control point and there are implications that enable the host to monitors and communicate with the VM applications up running. Therefore it is more necessary to strictly protect the host machines than protecting distinctive VMs. Different virtualization technologies have different implications for the host machine to influence the VMs up running in the system. Following are the possible ways for the host to influence the VMs [8]:

- The host can start shutdown, pause and restart the VMs.
- The host can able to monitor and modify the resources available for the virtual machines.
- The hosts if given enough rights can monitor the applications running inside the VMs.
- The host can view, copy, and likely to modify the data stored in the virtual disks assigned to the VMs.

In general all the network traffic to/from the VMs pass through the host, this enables the host to monitor all the network traffic for all its VMs.

1.5.3 VM monitoring from another VM: As isolation plays a vital role in virtualization, it is considered as a threat when one VM without any difficult

may be allowed to monitor resources of another VM and moreover, VMs does not have the possibility to directly access the file system of the host machine, so it is impossible for a VM to access the virtual disk allocated to the another VM on the host. When comes to the network traffic, isolation completely depends on the connection (network) setup of the virtualized environment. If the host machine is connected to the guest machine by means of physical dedicated channel, then it's unlikely that the guest machine can sniff the packets to the host and vice-versa. However in reality, the VMs are linked to the host machine by means "virtual hub" or by a virtual switch.

1.5.4 Guest to Guest attack: It is important to prevent the host machine than the individual VMs. If the attacker gain the administrator privileges of the hardware then it's likely that the attacker can break-in into the virtual machines. It is termed as guest to guest attack because the attacker can able to hop from one virtual machine to another virtual machine provided that the underlying security framework is already broken.

1.5.5 External modification of a VM: There are some sensitive applications exist which rely on the infrastructure of the VM environment. These applications running inside a virtual machine requires the virtual machine to be trusted environment to execute that application. If the VM is modified for some reason, the applications can still be able to run on the VM but the truest is broken.

1.5.6 External modification of the Hypervisor: Hypervisor is responsible for providing isolation between the guest machines. The VMs are said to be completely isolated or "self protected" only if the underlying hypervisor behaves well. A badly behaved hypervisor will break the security model of the system. There are several solutions exists for this problem, one of them is to use secure hypervisor Hyper-V [8] [9] to ensure security in hypervisor. Another solution is to protect the hypervisor from unauthorized modifications or enable the guest machines to validate the hypervisor.

1.6 Benefits of virtualization

In today's world, the main reason why IT organizations are considering virtualization of some, or all, of their computing infrastructures is that the technology helps reduce costs drastically.

Consider, for example, the case of the server infrastructure where utilization of servers increases from 15% to 80%, thus eliminating the need of extra physical servers, which are costly to run and maintain. In this scenario, consolidation of multiple physical servers takes place, onto one machine running a number of virtual servers. Such a setup brings about several cost and productivity benefits that are described below.[22]

a) **Cheaper implementation:** The case of the server infrastructure where instead of purchasing 5 Windows Server 2003 licenses, which would cost something in the region of $10,000 – $15,000 in licensing fees, and can simply purchase 1 license and host the remaining 4 servers without any extra cost. Furthermore, having fewer physical servers saves money on power bills, maintenance fees and datacenter office space and fees.

b) **Business doesn't stop:** Consolidating software applications, operating systems and hardware platforms, leads to fewer redundant physical devices needed to serve as primary machines. Conventional high-availability setups often require a 1:1 ratio of software-to-hardware while in the virtualized environment multiple servers can fail over to a set of backup servers. This therefore allows for a many-to-one configuration ratio, which increases service uptime and availability.

c) **Higher availability and uptime:** One of the main benefits of virtual servers is that they are completely isolated from one another, running as though they rely on separate hardware, which decreases downtime during maintenance periods. This means that changes can be made to one virtual server without affecting others sharing the same hardware. This means that maintenance can be done in a production environment without affecting business and causing disruptions.

d) **Speedy Installations:** Virtual devices allow fast installations of new server applications or router and switch software services, because no longer purchase of equipment take days or weeks to get it ordered, delivered and set up. Instead, simply configure a new virtual machine, router, switch or storage drive using the special virtualization management software tool used. This process generally consists of simply copying an image, thereby significantly reducing setup times.

e) **Corporate directives:** Corporate control lead organizations to use consolidated computing and networking infrastructures. Virtualization helps support higher management, security and tracking that directives required, in a profitable way. It is essentially easier to manage and enforce policies and to configure software from a central, common console than in a distributed manner, which makes it difficult to remain up to date with software updates.

1.7 Advantages and Disadvantages

There are many advantages and disadvantages of Virtualized environment, which helps an organization to choose the product.

1.7.1 Advantages of Virtualization

- **Flexibility:** is given in several ways. It is added because one can run more than one instances of an Operating system on a single computer, it is possible to migrate a virtualized instance to another physical computer and the virtual instances can be graceful from the host operating system with features like “pause”, “resume”, “shutdown” and “boot” [23]. It is also possible to change the specifications of virtual computers while they are running, for example the amount of ram, hard disc size and more [18].
- **Availability:** is added because one can keep the virtualized instances running even though the physical node has to be shut down, i.e. for hardware upgrade or maintenance. This is done by temporarily migrating the virtual instances to another computer, and migrate them back when the maintenance is finished and the primary computer is ready to serve. Hardware can be changed, upgraded, maintained and repaired without downtime in the services.

- **Scalability:** is added because is very easy to add or remove nodes. If the demand for capacity increases over time, it is very easy to insert a physical node with the basic cluster installation, and it will contribute in running the existing virtual machines that run services. This way, cluster will scale the company as it expands.
- **Hardware utilization:** is most likely increased if more than one operating system is hosted simultaneously. This is because virtual machines utilize hardware resources that are left idle by the host operating system [23].
- **Security:** is added because greater separation of services is introduced. Using multiple virtual machines it is possible to separate services by running one service on each virtual machine. If one service is compromised, the others are unaffected [24]. Using virtualization, the server would contain a minimal install that could host several virtual machines. Each virtual machine consists of a minimal operating system install and one service, for example, the web server is being compromised. The web pages hosted will be unreliable, but the break in does not affect the remaining services – the database server, mail server and the file server.
- **Cost:** it is possible to achieve cost reductions by consolidation smaller servers into powerful servers. Cost reductions stem from hardware cost reduction, operations cost reductions in terms of personnel, and software licenses.
- **Adaptability to workload variations:** changes in workload intensity levels can be easily taken care of by shifting resources and priority allocations among virtual machines. Autonomic computing based resource allocation techniques can be used to dynamically move processors from one virtual machine to another.
- **Load balancing:** since the software state of an entire virtual machine is completely encapsulated by the VMM, it is relatively easy to migrate virtual machines to other platforms in order to improve performance through better load balancing.
- **Legacy applications:** even if an organization decides to migrate to a different operating system, it is possible to continue to run legacy applications

on the old operating system running on a guest OS within a VM. This reduces the migration cost.

1.7.2 Disadvantages of virtualization

- **Overhead** causing decreased performance has been the biggest con with virtualization. Performance is often being compromised due to flexibility. The developers have worked hard to decrease the overhead, to bring it very close to the performance of the standalone physical computer.

- **SPOF (Single Point of failure)** in the hardware is still an issue. Even though the virtual machine is decoupled from the hardware, it is still dependent on the hardware working. Failure in the hardware will most likely lead to failure in the virtual machine, which will force a reboot.

- **The management interface** is closely linked to the virtualization platform. This can be a problem as it encumbers consolidation of several platforms into the same environment.

There are also some security risks related to virtualization. One of them is interactive virtualization-related risk, say, for example; when there is a virtualized server and virtualized network, are also critical security issue. In this case, the total risk exceeds the sum of the individual risk. Another security-related issue is ‘hyper-jacking’ in which an attacker crafts and then runs a very thin hypervisor that takes complete control of underlying operating system.

2. Literature Review

In 1999[25], VMware introduced virtualization to x86 systems, which VMware points out were not designed for virtualization in the way mainframes were. The problem entered on nearly 20 instructions that could cause application termination or system crash when they were virtualized. VMware addressed this with what it called an adaptive virtualization routine that contained the instructions as they're generated and allows other instructions to be passed through without intervention. With this breakthrough they were first on the market with a product that slowly began to attract attention, and then accelerated to the point where by 2008, a significant percentage of companies had begun to virtualized a small portion of their not-business-critical applications, and they began carrying out Windows Server backup on their new virtual machines.

2.1 Survey Analysis: When market analysts say "most companies have virtualized", some technical context should be considered. Virtualization can be said to have been "adopted" by a firm that has only virtualized at the shallowest levels. Virtualization is often adopted incrementally.

It was clear in 2007 that x 86-based server deployment patterns are changing dramatically in the market. The rapid emergence of multicore architectures and virtualization technologies was significantly restricting worldwide x86 server shipments. According to IDC's updated forecast [1], multicore and virtualization will cost the x86 market more than 4.5 million shipments and $2.4 billion in customer spending between 2006-2010. Overall, x86 shipments that were once projected to increase 61% by 2010 are now facing just 39% growth during that same period.

In 2011[26], x86 (also known as "volume class" or "commodity") server virtualization is the highest-priority agenda item for a majority of enterprises. Increased efficiency, hardware reduction, control over spending, and infrastructure agility continue to be the main reasons for investing in virtualization.

The SMB (Small and Medium Business) sector represents an emerging market for virtualization. A number of these smaller organizations have not yet begun to virtualized; they take the wait-and-watch approach instead, citing high costs and incompatible applications. Utility support is needed to help SMBs make the transition. Utility incentives alone will not turn the market, but they can serve as a "tipping point agent". Incentives, when available, have accelerated project timelines or moved them to the top of long to-do lists. Typically, incentives cover 15% to 30% of a project's cost, including, for example, some but not all of the professional services and hardware necessary for the project. According to IT research company Gartner, only 18% of the workloads on enterprise x86 servers that could be virtualized actually have been virtualized. As the global economy starts to recover, adoption should increase at a faster rate than in the last 24 months. In fact, 38% of the technology's existing users have suggested that they will increase their investment in virtualization.

In an effort to continue gathering data from to better understand this growing and ever-changing IT environment, enterprise and small-to-medium organizations, we were able to determine some key characteristics and opportunities for improvement in management within virtual environments.

The survey responses into **five categories:**

The virtual environment – what the make-up is of the respondent environments.

Determining applications to migrate – how respondent organizations identify systems and applications for migration to the virtual environment?

Ensuring performance and availability – what management tools are deployed currently by the respondents with respect to hypervisor, the hybrid environment and application performance management?

Reporting – what level of reporting respondents currently have for their virtual environment?

Securing the virtual environment – how respondents are planning to address securing the systems, applications and data migrated to virtual machines (a relatively new discussion for virtualization)?

2.1.1 The virtual environment

To understand what organizations are doing to manage virtualization technologies, there will be a clearer picture of what the virtual environments look like. What percent of the entire IT environment is virtualized and which hypervisor technologies are most widely deployed? We took this one step further and asked specifically about VMware ESX Server – how many hosts are deployed on average in the virtualized environment and approximately how many virtual machines are relied upon day to day?

Most organizations who are deploying virtual technology are taking their time in virtualizing their production environments. This was not surprising because the technology remains relatively new for many and only becoming a more mainstream technology within the last two to three years. The percentage of virtualization can be shown in Fig. 2.1 below.

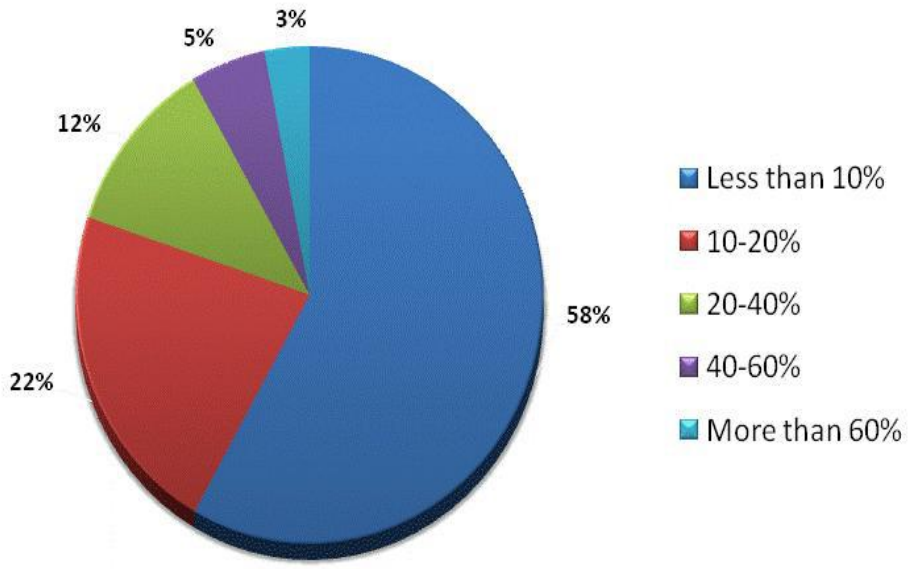

Fig. 2.1 Percentage of Virtualization in the environment

In IT that seems like a long period of time, but consider the enormity of this shift to the underlying technology. Many organizations will virtualize as their previous technology investments begin to end of life. This will continue to

lengthen the time it takes for some companies to have a significant portion of their IT deployment virtualized.

The most deployed hypervisors are VMware and Hyper-V. In 2008 [27], hypervisors in production is shown in Fig.2.2. VMware is in production more than any other individual hypervisor. But unlike most other industry reports, VMware's share is nearly 20-30 percent less and Microsoft Hyper-V is on the rise. With the official introduction of Hyper-V coming after this survey, the question raised how many respondents were choosing to deploy the Hyper-V beta into production?

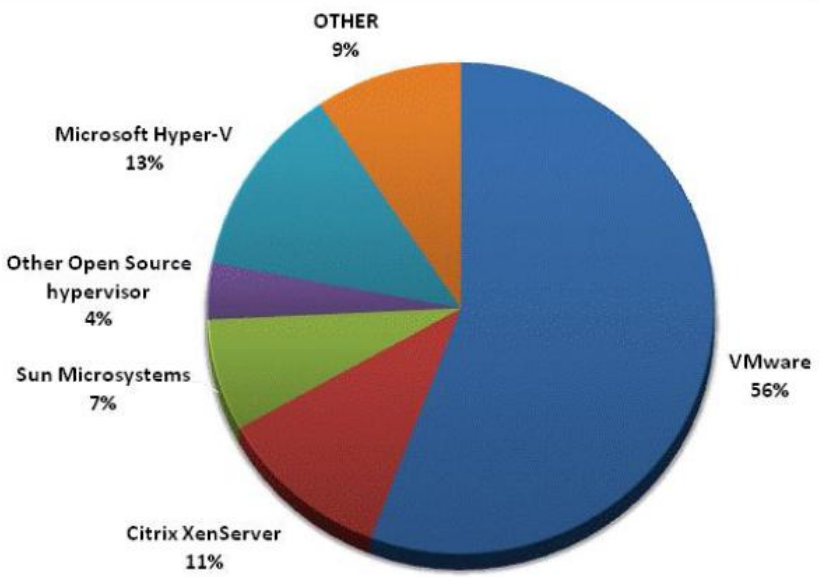

Fig. 2.2 Hypervisors in production

What was more interesting though, what is being virtualized to the hosts and VMs, although not an exhaustive list of applications [28] deployed throughout the datacenter, Fig.2.3 shows the "usual suspects" (Web services and file/print servers) as popular systems to virtualized. But there were also a large number of respondents virtualizing databases. The lower numbers for Microsoft Exchange Server and the systems management tools are in line with other reports from around the industry, and still noticeable was the 8% that are virtualizing Microsoft Exchange Server.

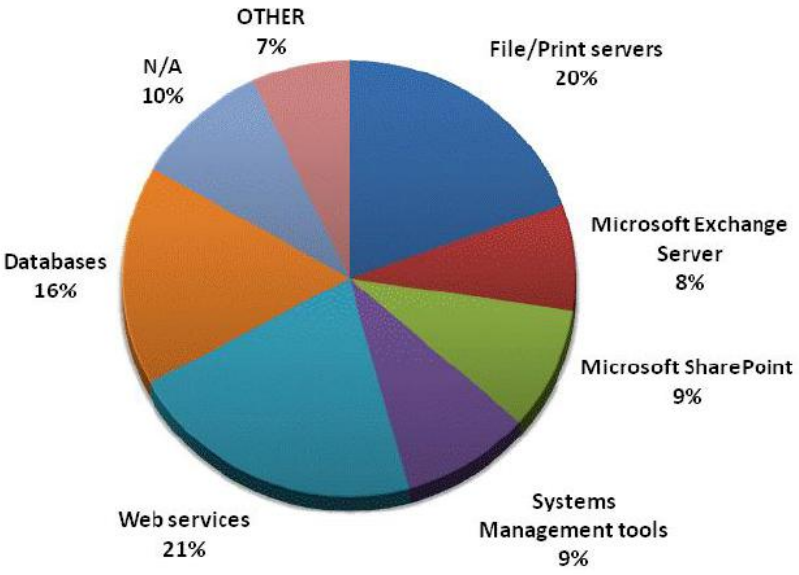

Fig.2.3 virtualized applications

2.1.2 Primary reason for deploying a virtualization solution

Survey 2011 shows [28] that, there are many reasons to deploy a virtualisation solution. But the question is to find out what was the primary reason that virtualisation was deployed, shown in Fig. 2.4.

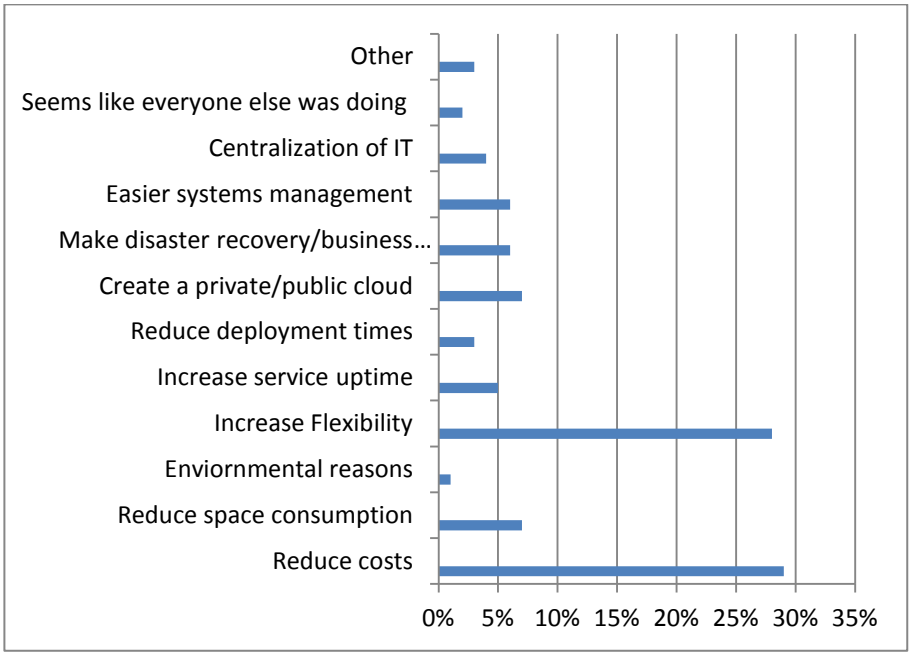

Fig. 2.4 Primary Reasons for virtualization solution

Reducing costs would feature highly in the results but we expect increased flexibility as high a selection. This could point to a solution such as private cloud being very popular in the future. This is proof that money talks; virtualisation does give reduced costs, mainly by reduced electricity consumption, and therefore greener computing is a happy side effect. When deploying Hyper-V we obviously need to focus on cost and flexibility.

Organisations that are new to virtualisation or don't quite grasp it, traditionally have had several opinions of how it can be successfully used. Some think that it should not be used for production. Some believe only lightweight server applications should be run on it. It is well known that VMware had captured 100% of the Fortune 1000 market and much of the mid-market by the time Windows Server 2008 R2 Hyper-V was released. Microsoft recommended a "use Hyper-V for lighter or sales systems" approach to get a foot in the door.

According to survey 2011[28], It appears that Hyper-V customers appreciate the potential of the solution. Nearly 80% (79.90%) of Hyper-V customers will run any production system on the Microsoft enterprise hypervisor. It will be interesting to see if (or should that be how) this will close in on 100% when "Windows 8" Hyper-V adds support for more than 16 vCPUs in a virtual machine.

2.1.3 Competitors of Hypervisors in market place:

There are many hypervisor products in the market available today. Graphical representation of competition is shown in Fig.2.5

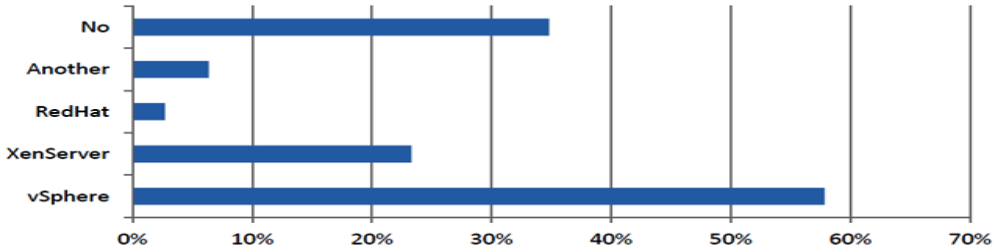

Fig. 2.5 Hypervisor competitors

Not surprisingly, VMware tops the chart, but it only scores at 57.84%. What is very surprising is that 34.8% say that no other product was considered.

[28] Survey strongly recommends that an assessment is performed before starting any project related to hypervisor. It provides the data required for an accurate proposal by consultants, or hardware sizing by anyone else.

Unfortunately survey shows that very few people, either internal IT or consultants, are performing an assessment. The numbers reflect that nearly 50% "knew" their requirements, nearly 5% didn't know what an assessment was, Only 42.89% have done an assessment. The remainder leave their deployments at risk of chronic over- or under-sizing their infrastructure, and most certainly have not considered supported configurations for the applications that are being virtualised on their infrastructure.

2.2 VMware vSphere Hypervisor (ESXi) [29]: VMware vSphere virtualizes and aggregates the underlying physical hardware resources across multiple systems and provides pools of virtual resources to the datacenter. Virtualization is a process that breaks the hard connection between the physical hardware and the operating system and applications running on it. After being virtualized in a vSphere virtual machine, the operating system and applications are no longer constrained by the limits imposed by residing on a single physical machine. Virtual equivalents of physical elements such as switches and storage operate within a virtual infrastructure that can span the enterprise.

2.2.1 A platform for virtualization and cloud infrastructure: VMware vSphere manages large collections of infrastructure, such as CPUs, storage, and networking, as a seamless and dynamic operating environment, and also manages the complexity of a datacenter. The VMware vSphere software stack is composed of the virtualization, management, and interface layers [30] as shown in Fig.2.6

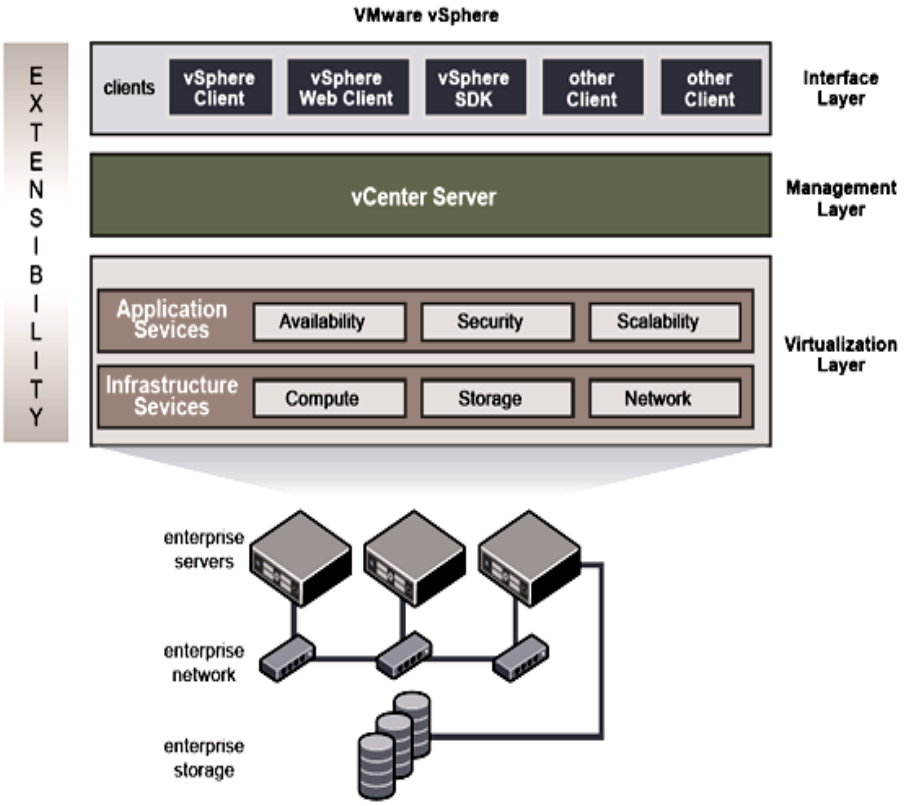

Fig.2.6 Relationship between the Component Layers of VMware vSphere.

a) Virtualization Layer: The virtualization layer of VMware vSphere includes infrastructure services and application services. Infrastructure services such as compute, storage, and network services abstract, aggregate, and allocate hardware or infrastructure resources. Infrastructure [30] include the following types in Table 2.1:

Table 2.1 Services at virtualization layer

Compute services	Includes the VMware capabilities that abstract away from underlying disparate server resources. Compute services aggregate these resources across many discrete servers and assign them to applications.
Storage services	The set of technologies that enables the most efficient use and management of storage in virtual environments.
Network Services	The set of technologies that simplify and enhance networking in virtual environments.

Application services are the set of services provided to ensure availability, security, and scalability for applications. Examples are: vSphere High Availability and Fault Tolerance.

b) Management Layer: VMware vCenter Server is the central point for configuring, provisioning, and managing virtualized IT environments.

c) Interface Layer: Users can access the VMware vSphere datacenter through GUI clients such as the vSphere Client or the vSphere Web Client. Additionally, users can access the datacenter through client machines that use command line interfaces and SDKs for automated management.

2.2.2 ESX for Consoling Database Systems

VMware ESX allows hardware to be partitioned, providing applications such as databases enough resources to keep utilization high while using the remaining resources for other workloads. Along with this resource partitioning, ESX provides the scalability required for database workloads. The following are some of the many features that make ESX ideal for consolidating database systems[31] on a single computing platform:

- **High-performance I/O:** ESX can drive over 100,000 database I/O accesses per second—more than enough to accommodate the requirements of even the largest databases.
- **CPU scalability:** ESX can make full use of the increasingly large number of cores in today's high-performance servers and offers two types of CPU scalability:
  - Scaling out by supporting multiple virtual machines on a single physical host
  - Scaling up by supporting up to four virtual processors in each guest virtual machine
- **Memory scalability:** Databases benefit greatly from large amounts of memory. ESX allows each virtual machine to be configured with up to 64GB of memory. Because consolidation of workloads allows higher processor

utilization, the average memory requirement per processor is higher—often about twice that of non-virtualized systems. To accommodate these growing requirements, the memory scalability curve has been pushed considerably, with ESX now supporting up to 256GB of physical memory.

- **Large pages:** Databases have for some time used large memory pages in the CPU's memory management unit (MMU) to optimize memory performance. This large-page feature can be enabled in many operating systems. ESX provides large-page support, allowing the database to fully utilize this feature.

2.2.3 VMware vSphere Components and Features

An introduction to the components and features of VMware vSphere helps you to understand the parts and how they interact. VMware vSphere includes the following components and features as shown in Table 2.2

Table 2.2 vsphere features [46]

VMware ESXi [32]	A virtualization layer runs on physical servers that abstract processor, memory, storage, and resources into multiple virtual machines.
VMware vCenter Server	The central point for configuring, provisioning, and managing virtualized IT environments. It provides essential datacenter services such as access control, performance monitoring, and alarm management.[33]
VMware vSphere Client [34]	An interface that enables users to connect remotely to vCenter Server or ESXi from any Windows PC.
VMware vSphere Web Client [35]	A Web interface that enables users to connect remotely to vCenter Server from a variety of Web browsers and operating systems.
VMware vSphere SDKs [36]	Feature that provides standard interfaces for VMware and third-party Solutions to access VMware vSphere.

VMFS (vSphere Virtual Machine File System) [37]	A high performance cluster files system for ESXi virtual machines.
vSphere Virtual SMP [38]	Enables a single virtual machine to use multiple physical processors Simultaneously.
vSphere vMotion [39]	Enables the migration of powered-on virtual machines from one physical server to another with zero down time, continuous service availability, and complete transaction integrity. Migration with vMotion cannot be used to move virtual machines from one datacenter to another.
vSphere Storage vMotion [40]	Enables the migration of virtual machine files from one datastore to another without service interruption. You can place the virtual machine and all its disks in a single location, or select separate locations for the virtual machine configuration file and each virtual disk. The virtual machine remains on the same host during Storage vMotion.
vSphere High Availability (HA) [41]	A feature that provides high availability for virtual machines. If a server fails, affected virtual machines are restarted on other available servers that have spare capacity.
vSphere Distributed Resource Scheduler (DRS) [42]	Allocates and balances computing capacity dynamically across collections of hardware resources for virtual machines. This feature includes distributed power management (DPM) capabilities that enable a datacenter to significantly reduce its power consumption.
vSphere Storage DRS [43]	Allocates and balances storage capacity and I/O dynamically across collections of datastores. This feature includes management capabilities that minimize the risk of running out of space and the risk of I/O bottlenecks slowing the performance of virtual machines.
vSphere Fault Tolerance [44]	Provides continuous availability by protecting a virtual machine with a copy. When this feature is enabled for a virtual machine, a secondary copy of the original, or primary, virtual machine is created. All actions completed on the primary virtual machine are also applied to the

	secondary virtual machine. If the primary virtual machine becomes unavailable, the secondary machine becomes immediately active.
vSphere Distributed Switch (VDS) [45]	A virtual switch that can span multiple ESXi hosts, enabling significant reduction of on-going network maintenance activities and increasing network capacity.

2.2.4 vSphere Client and vSphere Web Client:

All administrative functions are available through the vSphere Client. A subset of those functions is available through the vSphere Web Client. Comparisons of two clients are as shown in below Table 2.3.

Table 2.3 Comparison of Two vSphere clients

vSphere Client [34]	**vSphere Web Client [35]**
(For infrastructure configuration and day to day operations)	**(For day to day operations)**
Locally installed application	Web applications
Windows Operating System only	Cross Platform
Can connect to ESXi or directly to a host	Can connect to only ESXi
Full range of administrative functionality	Subset of full functionality, focused on virtual machine deployment and basic monitoring functions. Cannot configure hosts, clusters, networks, datastores, or datastore clusters.
	Extensible plug-in based architecture
Users: Virtual infrastructure administrators for specialized functions.	Users: Virtual infrastructure administrators, help desk, network operations center operators, virtual machine owners.

2.2.5 Direct Virtual Machine Console Access

If the virtual machine is running and the user knows the IP address of the virtual machine, the user can directly access the virtual machine console [47] by using standard tools, such as Windows Terminal Services. Only physical host administrators in special circumstances should directly access hosts. All relevant functions that can be done on the host can also be done in vCenter Server.

2.3 esxtop and esxplot

For testing the performance of server “esxtop” is used; same as linux, becoz ESXi server is using Linux-based Kernel. For analyzing the result and display the data in graphical form esxplot is used.

2.3.1 ESXTOP: Esxtop is VMware's version of popular "top" command in Linux or Unix, that run on an ESX or ESXi server. Both top and esxtop run only at the command line of a server. Access to these commands can be done either by going directly to a server console or by connecting to a server remotely via SSH (or telnet, if enabled).

Esxtop [48] is used to analyze real-time performance data from an individual ESX or ESXi server. It can also be done by entering to the physical server console or remote console via SSH [49], logging in, and typing esxtop. Esxtop analyzes CPU, Memory, Disk, and Network statistics. More specifically, esxtop has 8 different **"displays"** that show CPU, interrupt, memory, network, disk adapter, disk interface, disk VM, and power management as shown in Fig.2.7

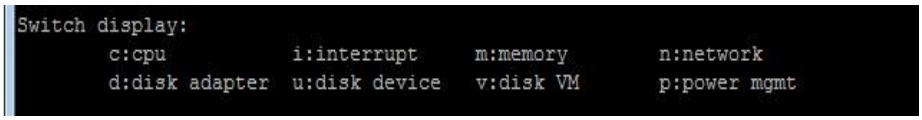

Fig.2.7: esxtop display

In the graphic above, different displays are brought up by pressing the corresponding letter for each display (i.e., press "d" for the "disk adapter" display). Once in a display, it always have uptime and CPU trending info at the top of each screen. From there, column view shows stats for different objects. For example, in the CPU view, processes running on the host down

the left side (which could correspond to a particular VM) and columns across the top with different stats for each of those processes.

Command for esxtop[50] in a batch mode is given by using

esxtop –a –b -d (x) –n (y) > filename.csv

Where, -a: ALL; -b: BATCH MODE; –d: DELAY; -n: NUMBER of iterations.

The resulting CSV file then can be analyzed using ESXPLOT for graphical representation.

2.3.2 ESXPLOT: [51] esxplot is a GUI application that explore the data collected by esxtop in batch mode. The program takes a single command line argument which is the esxtop batch mode output file. It can also be done by simply start esxplot without any arguments, and enter a dataset file via the File attribute of the menu bar. Esxplot loads the data in this file and presents the metrics as a hierarchical tree where the values are selectable in the left panel. In the right panel, a graph is plotted (value over time) of the selected metric, in this way, you can "browse" the contents of these somewhat unwieldy files as shown in Fig. 2.8

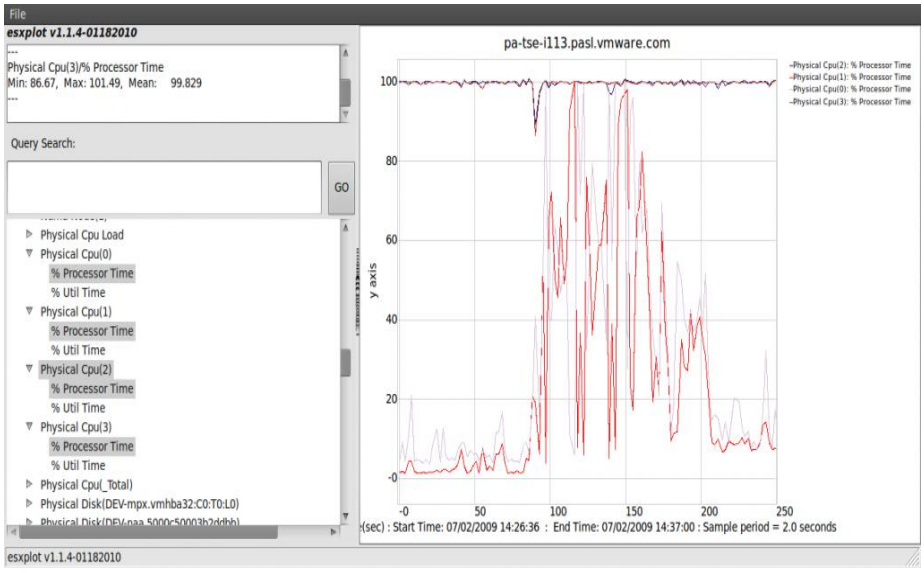

Fig. 2.8 ESXPLOT Display

2.4 Benchmark applications

A benchmark is the act of running a computer program, a set of programs, or other operations, in order to assess the relative performance of an object, normally by running a number of standard tests and trials against it. The term 'benchmark' is also mostly utilized for the purposes of elaborately-designed benchmarking programs themselves. Benchmarking is usually associated with assessing performance characteristics of computer hardware, for example, the floating point operation performance of a CPU, but there are circumstances when the technique is also applicable to software. Software benchmarks are, for example, run against compilers or database management systems. Benchmarks provide a method of comparing the performance of various subsystems across different chip/system architectures. Test suites are a type of system intended to assess the correctness of software.

Benchmarking [52] focuses on company-to-company comparisons of how well basic functions and processes are performed. Benchmarking enables managers to determine what the best practice is, to prioritize opportunities for improvement, to enhance performance relative to customer expectations. It also helps managers to understand the most accurate and efficient means of performing an activity, to learn how lower costs are actually achieved, and to take action to improve a company's cost competitiveness. As a result, benchmarking has been used in many companies as a tool for obtaining a competitive advantage. Companies usually undertake benchmarking with a view towards the many improvements that it may offer. These benefits include reducing labor cost, streamlining the work flow through reengineered business processes and common administrative systems, improving data center operations through consolidation and downsizing, cooperative business and information technology planning, implementing new technology, outsourcing some assignments and functions, redesigning the development and support processes, and restructuring and reorganizing the information technology functions.

2.4.1 Goal of benchmarking

To identify the weaknesses within an organization and improve upon them, with the idea of becoming the "best of the best." The benchmarking process helps managers to find gaps in performance and turn them into opportunities for improvement. Benchmarking enables companies to identify the most successful strategies used by other companies of comparable size, type, or regional location, and then adopt relevant measures to make their own programs more efficient. Most companies apply benchmarking as part of a broad strategic process. For example, companies use benchmarking in order to find breakthrough ideas for improving processes, to support quality improvement programs, to motivate staffs to improve performance, and to satisfy management's need for competitive assessments. Benchmarking targets roles, processes, and critical success factors. “Roles” are what define the job or function that a person fulfils. “Processes” are what consume a company's resources.

Benchmarking is an organizational tool to drive continuous improvements using best practices which results into increased efficiency and create competitive advantages. Performance metrics measure and report an organization's performance, both business practices require willingness and a high commitment to change [53] in developing a benchmarking program, managers review strategic business goals, collect data related to best practices and use data analysis to develop performance metrics and plan improvements. This process serves as the basis for effective employee goal-setting.

a) Business Processes: Examines core business processes to establish what it will measure. Flow charts assist in identifying the inputs and outputs required for a key process. These are the specific activities involved in its function and drive its productivity. Activities required for a core business process are generally similar across different organizations. Time and cost are two significant factors that create performance differences.

b) Functional Benchmarking: Identify best practices to help employees improve performance. Management can motivate staffs using functional benchmarking, which illustrates what is possible at the highest level of performance for a particular function within an organization. This requires

data collection that can use for internal and external comparisons. Statistics, reports and case studies are examples of the type of data and information used in functional benchmarking.

c) Performance Metrics: A performance metric might include 15 to 20 core measurements. This serves as a type of scorecard linked to specific performance goals and activities. Measurements can relate to financial, customer service, operational and innovation perspectives.

d) Time Limited Goals: Employee performance goals require a start date and end date to be effective. Timetables should be long enough to make goals attainable, but short enough to convey urgency. Time frames should also include milestones that allow an employee and management to track progress towards meeting the goal. It also include time frame data in the performance metric.

e) Employee Resources: An organization should be prepared to assist employees in building the skills required to meet benchmark standards. While employee goal-setting encourages them to push beyond their existing performance, it might also require additional resources to meet performance goals. For example, additional employee training helps develop technical and soft skills.

f) Monitoring Performance: In assessing employees, provide honest evaluations of failures and generous praise for successes. Assessment and monitoring performance should include milestones that allow employees time to take corrective actions before the end of the period set for completing a specific goal. This will serve employees in the self-assessment process and provide the platform for setting new targets

2.4.2 Types of Benchmark

There are a number of different types of bench-marking [54], which are driven by different motivating factors and thus involve different comparisons. Some of the major types of benchmarking are as follows:

Metric benchmarking is the use of quantitative measures as reference points for comparisons. Best-practice benchmarking focuses on identifying outstanding techniques.

Information technology benchmarking includes data processing, systems analysis, programming, end-user support, and networks.

Infrastructure benchmarking includes data centers, networks, data/information, end-user support, and distribution remote centers.

Application benchmarking includes system analysis, development and maintenance programming, and functionality.

Strategy benchmarking includes skills assessment, information technology strategy, business-technology alignment, and delineation of roles and responsibilities.

2.4.3 Benchmark tools are different for windows and Linux. [54] [55] Some of the OS based commonly used tools are:

a) **SuperPi:** It is free benchmarking testing software for windows, which is capable of calculating Pi up to 32 million digits after the decimal point. This complex math significantly tasks your computer's processor, and SuperPi keeps track of how quickly it takes for your computer to complete the calculation. SuperPi is focused on your processor's speed, not the speed of other components, so it only useful when judging changes to your computer's processors. For example, overclockers consider SuperPi to be one of the best free benchmarks around because it provides a processor-focused test which can help them judge how much extra performance their overclocking has gained them.

b) **3DMark06 / PCMark05:** Futuremark is a software company which offers a wide variety of computer benchmarks which are considered among the best in the world. Chances are that any review of a desktop or laptop computer you've ever read has, at some point, quoted 3DMark or PCMark. The latest versions of 3DMark and PCMark aren't free, but older versions with limited options can be used an unlimited number of times without charge. While these benchmarks are over four years old, they are still among the best free benchmarks available. Futuremark's benchmarks are very demanding, and there are many modern computers which will achieve very low scores when running these benchmarks. **3DMark** focuses on gamers, has

it only tests the power of your video card. **PCMark** is for more general use and tests numerous computer components. Both can be available at Futuremark website.

c) **SiSoft Sandra:** Sandra stands for System Analyzer, Diagnostic and Reporting Assistant. It is a fully-featured benchmark suite which is aimed at users who are very well informed about the inner workings of their computers and for businesses which need to perform a detailed analysis on multiple computers. SiSoft Sandra kindly offers a free version of the software. As shown in Fig.2.9, you can test your computer's memory bandwidth, network performance, computer's power efficiency. Chances are that you'll come across one or two benchmarks in SiSoft Sandra which benchmark hardware you didn't even know existed.

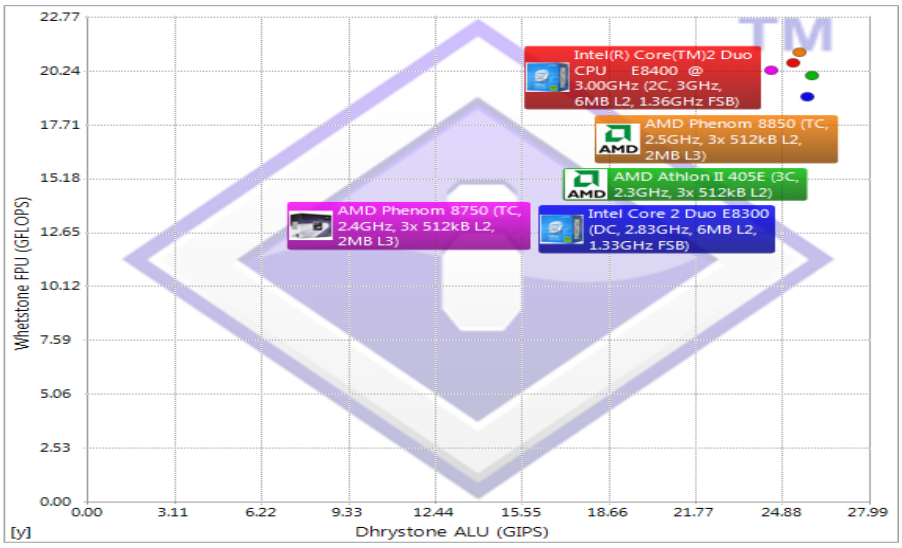

Fig. 2.9 SiSoft Sandra analysis

Another useful feature of SiSoft Sandra is the inclusion of references. SiSoft Sandra will benchmark processor and then compare performance to five other similar processors to gives a better idea of how an upgrade may or may not help. This is something that only SiSoft Sandra offers, and it is incredibly useful.

d) **Phoronix:**[56] The Phoronix Test Suite is the most comprehensive testing and benchmarking platform available that provides an extensible framework for which new tests can be easily added. The software is designed to effectively carry out both qualitative and quantitative benchmarks in a clean, reproducible, and easy-to-use manner. The Phoronix Test Suite is based upon the extensive testing and internal tools developed by Phoronix.com since 2004, along with support from leading tier-one computer hardware and software vendors. This software is open-source and licensed under the GNU. Originally developed for automated Linux testing, support to the Phoronix Test Suite has since been added for OpenSolaris, Apple Mac OS X, Microsoft Windows, and BSD operating systems. The Phoronix Test Suite consists of a lightweight processing core with each benchmark consisting of an XML-based profile and related resource scripts. The process from the benchmark installation, to the actual benchmarking, to the parsing of important hardware and software components is heavily automated and completely repeatable, asking users only for confirmation of actions.

e) **Sysbench:** SysBench [57] is a modular, cross-platform and multi-threaded benchmark tool [56] for evaluating OS parameters that are important for a system running a database under intensive load.

- **Features of SysBench:** Current features allow to test the following system parameters:

a. file I/O performance
b. scheduler performance
c. memory allocation and transfer speed
d. POSIX threads implementation performance
e. database server performance

The idea of this benchmark suite is to quickly get an impression about system performance without setting up complex database benchmarks or even without installing a database at all.

- **General syntax for SysBench is:**

  sysbench [common-options] --test=name [test-options] command

  where,

common options are as listed in Table 2.4 below

Table 2.4 Command option table

Options	**Description**	**Default Value**
--num-threads	The total number of worker threads to create	1
--max-requests	Limit for total number of requests. 0 means unlimited	10000
--max-time	Limit for total execution time in seconds. 0 (default) means unlimited	0
--forced-shutdown	Amount of time to wait after --max-time before forcing shutdown. The value can be either an absolute number of seconds or as a percentage of the --max-time value by specifying a number of percents followed by the '%' sign. “off” (the default value) means that no forced shutdown will be performed.	off
--thread-stack-size	Size of stack for each thread	32K
--init-rng	Specifies if random numbers generator should be initialized from timer before the test start	off
--test	Name of the test mode to run	Required
--debug	Print more debug info	Off
--validate	Perform validation if result tests where possible	Off
--help	Print help on general syntax or on a test mode specified with --test, and exit	Off

--verbosity	Verbosity level (0 - only critical messages, 5 - debug)	4
--percentile	SysBench measures execution times for all processed requests to display statistical information like minimal, average and maximum execution time. For most benchmarks it is also useful to know a request execution time value matching some percentile This option allows to specify a percentile rank of query execution times to count	95
--batch	Dump current results periodically	off
--batch-delay	Delay between batch dumps in seconds	300
--validate	Perform validation of test results where possible	Off

Commands are as:

prepare: Performs preparative actions for those tests which need them, e.g. creating the necessary files on disk for the ***fileio*** test, or filling the test database for the oltp test.

run: Runs the actual test specified with the ***--test=name*** option.

cleanup: Removes temporary data after the test run in those tests which create one.

help: Displays usage information for a test specified with the
--test=name option.

Batchmode: In some cases it is useful to have not only the final benchmarks statistics, but also periodical dumps of current stats to see how they change over the test run. For this purpose SysBench has a batch execution mode

which is turned on by the ***--batch*** option. It may specify the delay in seconds between the consequent dumps with the ***--batch-delay*** option.

Example: *sysbench --batch --batch-delay=5 --test=threads run*

This will run SysBench in a threads test mode, with the current values of minimum, average, maximum and percentile for request execution times printed every 5 seconds.

Test modes are as : Tests can be performed in five modes as:

1. cpu
2. threads
3. mutex
4. fileio
5. memory

Detailed description available for each test in Sysbench is as:

1. cpu: The cpu is one of the most simple benchmarks in SysBench. In this mode each request consists in calculation of prime numbers up to a value specified by the ***--cpu-max-primes*** option. All calculations are performed using 64-bit integers. Each thread executes the requests concurrently until either the total number of requests or the total execution time exceeds the limits specified with the common command line options.

Example: *sysbench --test=cpu --cpu-max-prime=20000 run*

2. **threads:** This test mode was written to benchmark scheduler performance, more specifically the cases when a scheduler has a large number of threads competing for some set of mutexes. SysBench creates a specified number of threads and a specified number of mutexes. Then each thread starts running the requests consisting of locking the mutex, yielding the CPU, so the thread is placed in the run queue by the scheduler, then unlocking the mutex when the thread is rescheduled back to execution. For each request, the above actions are run several times in a loop, so the more

iterations is performed, the more concurrency is placed on each mutex. The following options are available in this test mode:

- ***--threads-yields:*** number of lock/yield/unlock loops to execute per each request. Its default value is 1000

- ***–threads-lock:*** number of mutexes to create. Its default 8.

Example: *sysbench --num-threads=64 --test=threads --thread-yields=100 --thread-locks=2 run*

3. **memory:** This test mode can be used to benchmark sequential memory reads or writes. Depending on command line options each thread can access either a global or a local block for all memory operations. The following options are available in this test mode:

- ***--memory-block-size:*** Size of memory block to use. Its default value is 1K.
- ***–memory-scope:*** Possible values: global, local. Specifies whether each thread will use a globally allocated memory block, or a local one. Its default values is global
- ***–memory-total-size:*** total size of data to transfer. Its default value is 100G.
- ***–memory-oper:*** type of memory operations. Possible values: read, write. Its default value is 100G.

4. **fileio:** This test mode can be used to produce various kinds of file I/O workloads. At the prepare stage SysBench creates a specified number of files with a specified total size, then at the run stage, each thread performs specified I/O operations on this set of files. When the global --validate option is used with the fileio test mode, SysBench performs checksums validation on all data read from the disk. On each write operation the block is filled with random values, and then the checksum is calculated and stored in the block along with the offset of this block within a file. On each read operation the

block is validated by comparing the stored offset with the real offset, and the stored checksum with the real calculated checksum.

The following I/O operations are supported:

- ***seqwr:*** sequential write
- ***seqrewr:*** sequential rewrite
- ***seqrd:*** sequential read
- ***rndrd:*** random read
- ***rndwr:*** random write
- ***rndrw:*** combined random read/write

list of ***test-specific*** option for fileio mode:

- ***--file-num:*** number of files to create. Its default value is 128.

- ***--file-block-size:*** Block size to use in all I/O operations. Its default value is 16K

- ***--file-total-size:*** Total size of files. Its default size is 2G

- ***--file-test-mode:*** Type of workload to produce. Possible values: *seqwr, seqrewr, seqrd, rndrd, rndwr, rndwr*. It is required.
- ***--file-io-mode:*** I/O mode. Possible values: *sync*, *async*, *fastmmap*, *slowmmap* (only if supported by the platform). Its default value is sync

- ***--file-async-backlog:*** Number of asynchronous operations to queue per thread (only for --file-io-mode=async). Its default value is 128

- ***--file-extra-flags***: Additional flags to use with *open(2)*

- ***--file-fsync-freq:*** Do *fsync()* after this number of requests (0 - don't use *fsync( )*). Its default value is 100.

- ***--file-fsync***-all : Do *fsync()* after each write operation. It has no default values.

- ***--file-fsync***-end: Do *fsync()* at the end of the test.

- ***--file-fsync***-mode: Which method to use for synchronization. Possible values: *fsync, fdatasync*. Its default value is *fsync.*

- ***--file-merged-request:*** Merge at most this number of I/O requests if possible (0 - don't merge). Its default value is 0.

- ***--file-rw-ratio:*** reads/writes ration for combined random read/write test. Its default value is 1.5.

Usage example:

$ sysbench --num-threads=x --test=fileio --file-total-size=y --file-test-mode=rndrw prepare

$ sysbench --num-threads=x --test=fileio --file-total-size=y --file-test-mode=rndrw run

$ sysbench --num-threads=x --test=fileio --file-total-size=y --file-test-mode=rndrw cleanup

First command creates *x* number of files with the total size of *y* in GB in the current directory, **Second command** runs the actual benchmark and displays the results upon completion,

Third one removes the files used for the test.

3. Problem Statement

The main objective of this thesis is to test the system performance in the field of virtualization.

In the field of virtualization, first major problem is to investigate the system performance i.e. memory usage, CPU usage which may not bring immediate focus among other things, as CPU resources are difficult to interpret at the very first glance, and it always look like, that there are still sufficient resources available.

But it may come at some point of utilization at “hang” state for a short time or virtual systems may “freeze”, even though the resource utilization is below 20%.

4. Set Up and Implementation

4.1 Installation of VMware ESX/ESXi on Local Storage

Basic hardware requirements [58]

To successfully install ESX/ESXi, proper hardware requirement must be there:

- 2 GB of RAM
- One Gigabit network interface
- 4 GB of hard drive space
- 2 processors

Once the hardware and version requirement has been met for a specific environment, you must obtain the installation media.

Installation Steps

- ESXi can be installed from a downloaded ISO file, or
- Install it from any bootable media either from bootable USB or from bootable [59] CD/DVD.

There are also three ways to run the installer:

- Interactive Graphical (Default)
- Interactive Text
- Scripted

VMware ESXi Server comes in two versions:

Embedded and **Installable**

- The **embedded** [60] version is "installed" directly on the server by the hardware vendor.
- The **installable** [61] ESXi version is installed in a similar fashion as ESX, but the resulting server will not have a fully functional Service Console.

Installation ESXi 5.0 from the direct console [32]

1. Install the boot DVD and choose the install in text mode.
2. Once the kernel and some additional modules have started, a welcome screen for installation appears, click Enter for Continue.
3. At the welcome screen, press F11 to accept the EULA (End User License Agreement).
4. Installer will then scan the host for available Disks and displays a list of valid Disks. Here,
    - Disks are separated into : local: and remote:
    - VMFS partition will be displayed with an asterisk (*) next to them.
    - An option to press F1 key to get additional details about the disks displayed.
    - If you press F1 to get the details of the disks then, note two things:
        - ESX(i) Found : No and Datastores: (none). Then press enter to return to the selected disk screen and then press Enter to continue installation.
5. On the select Keyboard Screen, choose U.S English (the default).
6. Enter password for Root, and then Enter to continue.
7. An error/warning screen will appear and can ignore these errors since they are in a virtualized environment. Press Enter to continue.
8. To confirm installation, press F11.
9. Now, by pressing ALT-F1 sequence in this stage will appear a root login Prompt. Now, login to ESXi 5.0 at this stage and have a look around at the logs if required. And at this stage ESXi5.0 root password is blank.
10. Press ALT-F2 sequence to see the VMkernel logs.
11. Press ALT-F2 sequence to return to the installer screen.
12. Once the installer has completed, click Enter for Reboot.
13. Once the reboot completes, the main server screen window shows the host IP address of the server.

14. Now, by pressing ALT-F11, warning appears on server Screen as:

/vmfs/devices/char/vmkdriver/usbpassthrough not found (Which shows that, a server is running in a virtualized environment.)

ESXi 5.0 Configuration [58]

1. Once the ESXi 5.0 host installed; now will continue and configure the network settings.
2. Move to ALT-F2 screen and press F2 key to Customize System/View Logs.
3. On the root login screen, enter root as login and 1234567 as password.
4. System Customization Screen will appears, note that the Configure Password is set.
5. Configuring the ESXi 5.0 networking and the DNS settings. The settings which are using to configure the host is shown in Table 4.1 below:

Table 4.1 Configuring ESXi

IP Configuration:	Static or Dynamic
IP addresses:	172.31.5.66
Subnet Mask:	255.255.255.0
Default Gateway:	172.31.5.1

- From the System Customization menu, select Configure Management Network and press Enter.
- Select IP configuration and then Enter. Select Set static IP address and network configuration (use the arrow keys to move down and press the spacebar to select) and change the IP Address, Subnet Mask and Default Gateway settings to the entries that have outlined in the table above.
- Once these values have been entered, press Enter.

6. Configuring the DNS settings[62] as:
   - From the Configure Management Network menu, select DNS Configuration and then press Enter.
   - The settings will use to configure the host is shown in Table 4.2 below:

Table 4.2: Details for network in DNS settings

Primary DNS Server:	172.31.1.6
Secondary DNS Server:	Optional (for static)
Hostname:	localhost

- Once these values have been entered, press Enter.

7. Press ESC to return to the Customize System menu.
8. Since, changes to the host management network are done, confirmation windows appear and restart the management network.
9. Press Y to accept the changes and ESC to return to the main ESXi 5.0 console screen.

4.2 Installation of vSphere Client [29]

1. In a browser, enter the IP address of vSphere Client host. This is where the VMware vSphere Hypervisor Client resides.
2. Click Download vSphere Client and Save the file. After the file downloads, installation wizard of the VMware vSphere Client appears.
3. Click Next to begin going through the wizard.
4. Click install, when the task is complete, the installation completed screen appears.
5. Click Finish.

Prepare the VMware host

Before deploying the OVF template, the VMware host [63] must meet system requirements.

1. On your computer desktop, double-click the shortcut for the VMware vSphere Client. The client login screen appears as shown in Fig.4.1 below.
2. Enter your IP address/ Name field i.e. Enter the same IP address and password you used for the VMware vSphere host and continue for login.
3. The vSphere Client screen appears. Then Double-click the inventory icon and getting started.

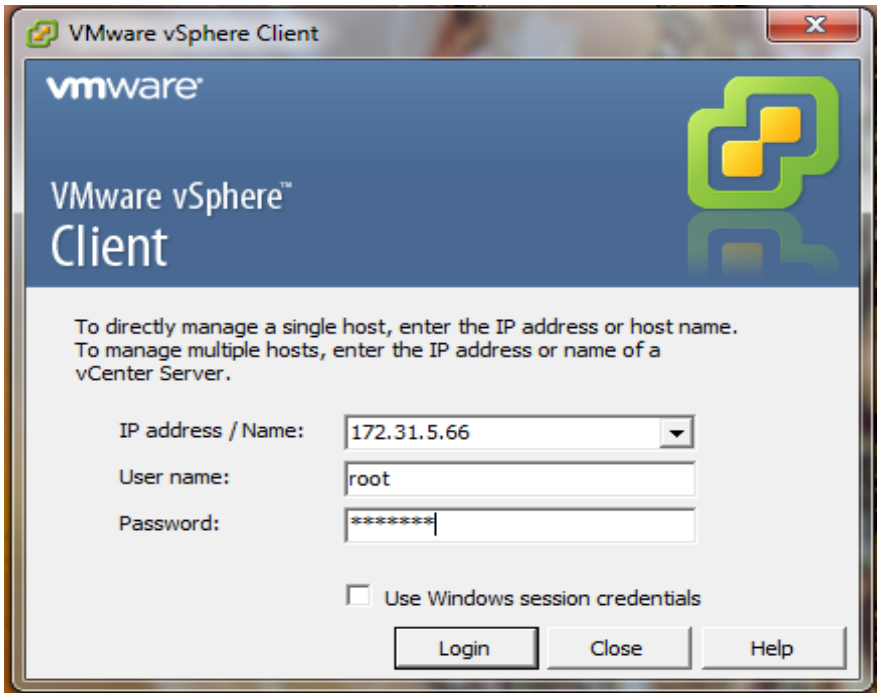

Fig. 4.1 Login screen of vSphere Client

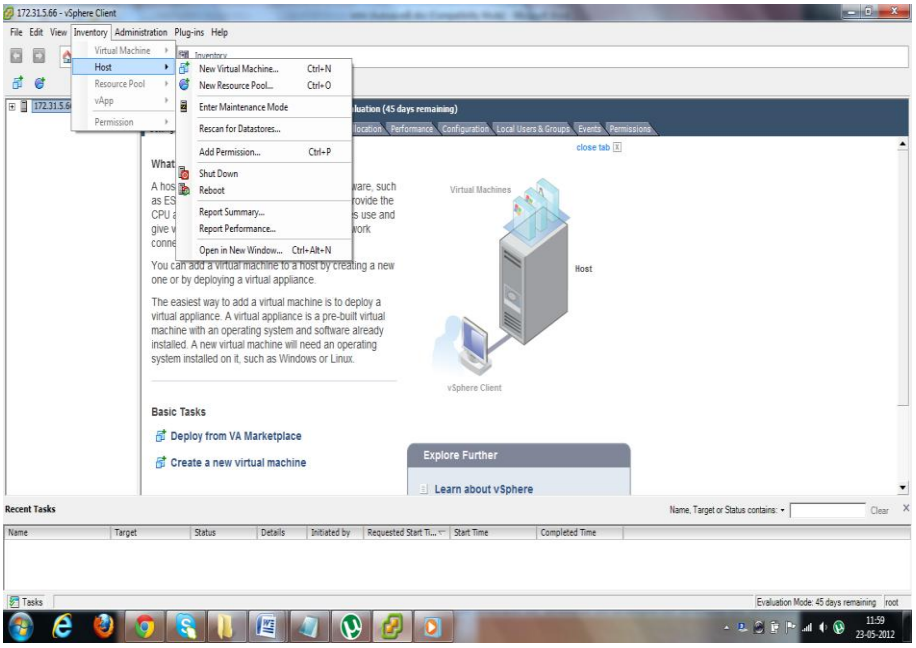

Fig. 4.2 Inventory of vSphere

Configure ESX/ESXi

To configure ESXi, follow these steps:

1. Once logged into the host with the vSphere Client, Home view window appears.

2. By clicking the Inventory icon in the Inventory panel, ESX server can be seen as in Fig. 4.2
3. By clicking on the ESX server, several tabs in the panel on the right are presented. Click the Configuration tab for Licensing as in Fig.4.3

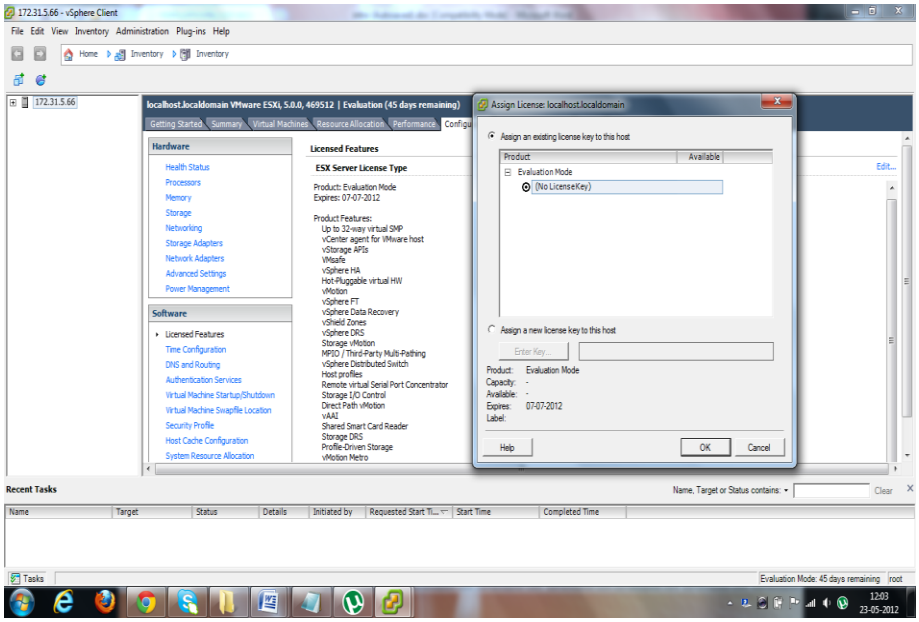

Fig. 4.3 Licensing Tab of vSphere

4. Click on Edit to add the license, and then Ok.

4.3 Installation steps of Sysbench:

1. Download the system performance benchmark i.e. sysbench-0.4.12.tar.gz file [57].
2. Open the terminal as a root:

./configure

3. Type ./configure --without-mysql (to compile sysbench without MYSQL support)
4. Type ./configure –with-mysql (to compile sysbench with MYSQL support)
5. A screen appears showing Sysbench installed with its usage options.
6. Now, run the sysbench commands for testing our performance.

4.4 Running esxtop [48]

1. Open console session by ALT-F1or ssh [49] to ESX(i) and type:
   esxtop
2. By default, the screen will be refreshed in every 5 seconds, changes can be done by pressing *Shift* key
3. Changing views is easy. Type the following keys for the associated views:
   c = cpu m = memory n = network i = interrupts d = disk adapter
   u = disk device v = disk VM p = power states
4. Enter the f command, then a field select page appears.

4.5 Capturing ESXTOP result:

1. Configured esxtop as needed, run it in batchmode and save results in .csv file as: :

esxtop –a –b –d (x) –n (y) > filename.csv

Where, -a: for all; -b: batch mode; -d (x): delay for x sec and –n (y): y iterations.

2. To retrieve the file from server, run commands as:

cd /vmfs/volumes/datastore1

cp /filename.csv filename.csv

4.6 Analyzing results:

You can use multiple tools to analyze the captured data.

1) Perfmon [64]
2) excel
3) esxplot [51]

In this thesis work, esxplot is using for analyzing results.

5. Experimentation and result analysis

5.1 Running esxtop

Running the esxtop on ESX directly or by using SSH for login into the server. Below are the given steps for experimentation

1. Start esxtop (Screen appears as in Fig. 5.1)

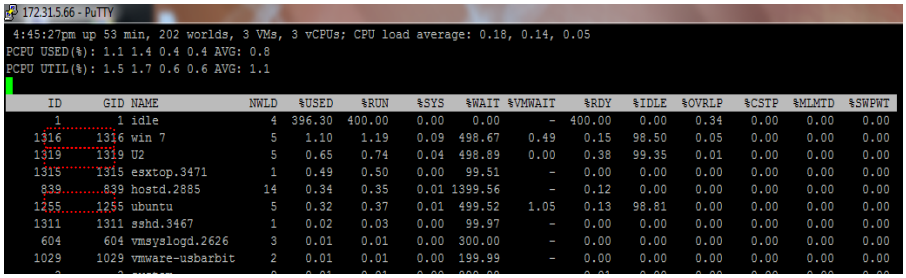

Fig.5.1 esxtop main screen

2. By typing *h* , help screen will appear, and will display all the possible switch which can be used in interactive mode.

3. To switch between different set of metrics use key from **Switch Display** section.

To monitor the memory metrics, then type *m*, and esxtop with memory metrics screen appears.

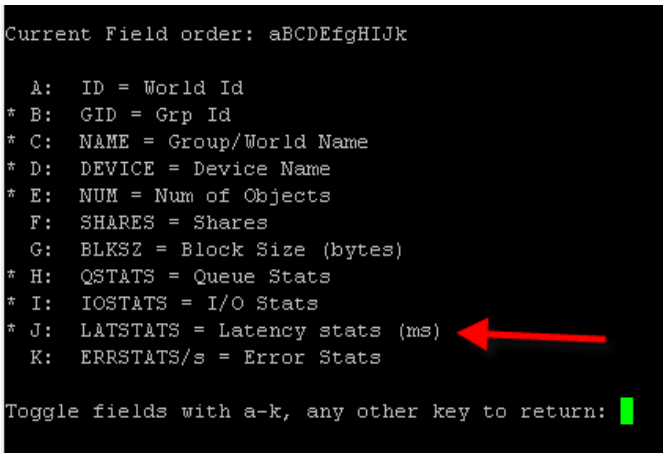

Fig. 5.2 Currently Monitored Metrics

4. To see list of metrics which are available and currently monitored type F or f, * next to metrics indicates that metrics is currently monitored as in Fig.5.2

5. To save the metrics in a file type :

esxtop –a –b –d 10 –n 550 > hiteshi1.csv (as shown in Fig. 5.3 below)

this will save the metrics of the load generated on the esx server by client for 550 seconds with a delay of 10seconds.

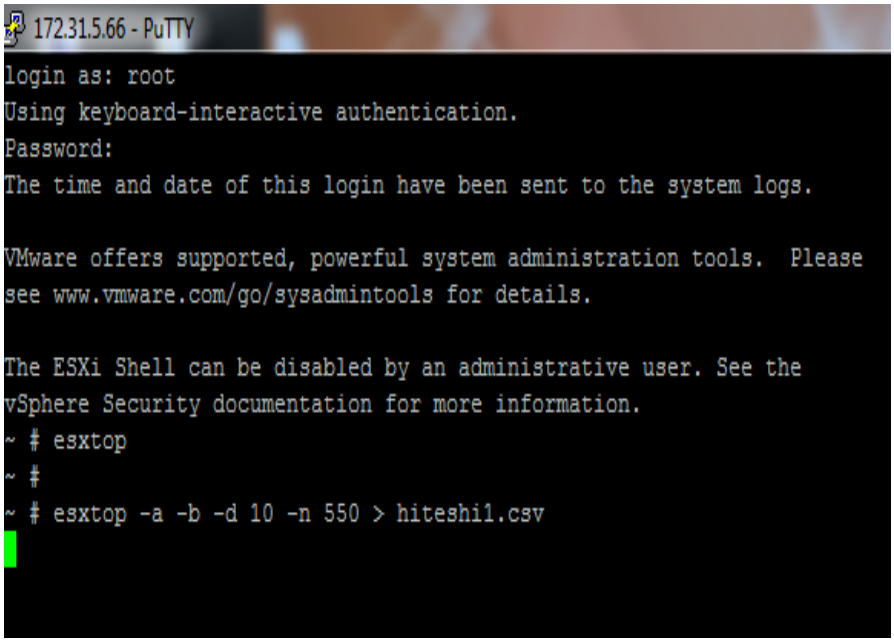

Fig. 5.3 Save metrics in .csv file

6. For generating load on server, here a benchmark application named "Sysbench" is used.

7. Run sysbench on linux VM and testing the cpu by sysbench in one terminal by using command as:

$ sysbench –test=cpu –cpu-max-prime=10000 run

And the resulting screen appears as in Fig. 5.4 below

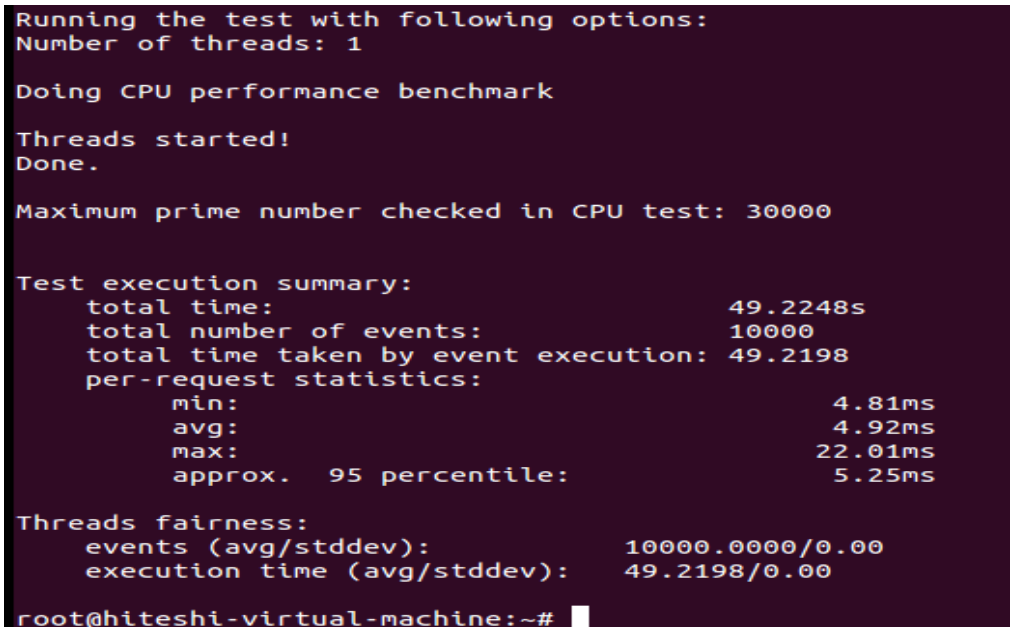

Fig. 5.4 Resulting Screen of CPU test.

8. Side by side, we are also running commands for *fileio* in another terminal , so that the load can be easily ganerated and recorded in a file by commands as:

$ sysbench --num-threads=64--test=fileio --file-total-size=4G --file-test-mode=rndrw prepare

$ sysbench --num-threads=64 --test=fileio --file-total-size=4G --file-test- mode=rndrw run

$ sysbench --num-threads=64 --test=fileio --file-total-size=4G --file-test-mode=rndrw cleanup

9. When we get the output in CSV file [65] then we use ESXPLOT for graphical representation of our metrics.

5.2 Transfering file to the client:

To retrieve the file from server, run commands as:

cd /vmfs/volumes/datastore1

cp /filename.csv filename.csv

5.3 Analyzing results:

For analyzing this experiment, esxplot is using

.

% CPU latency Graph:

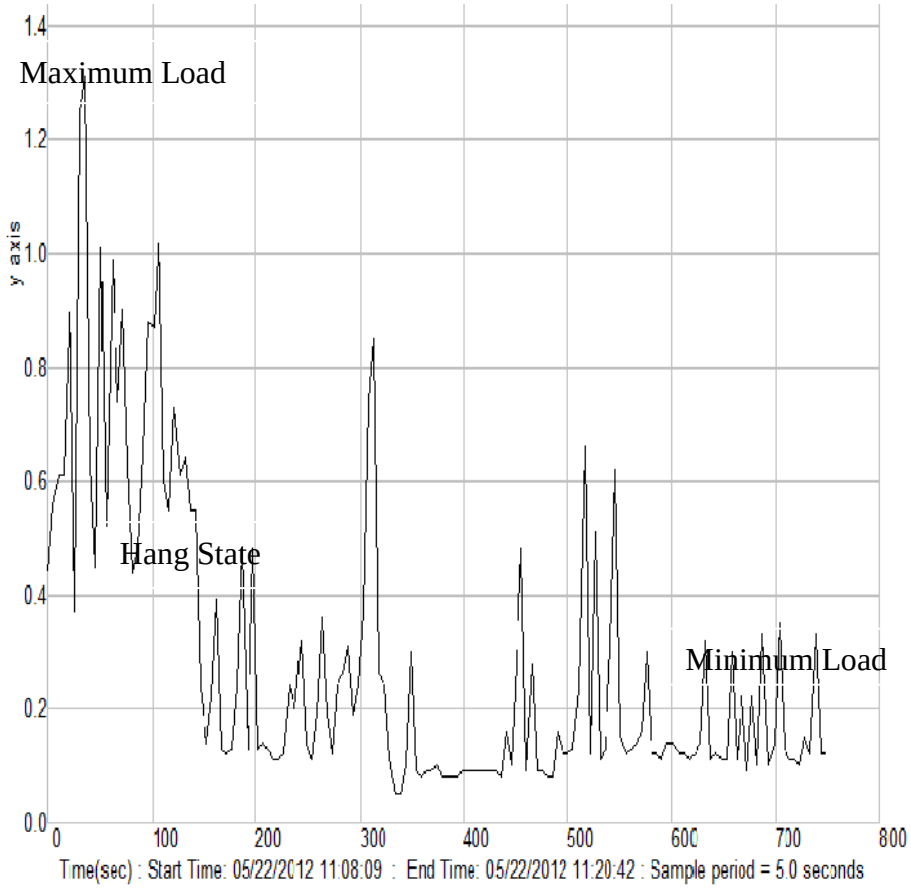

Fig. 5.5 % CPU latency

%Overlap Graph:

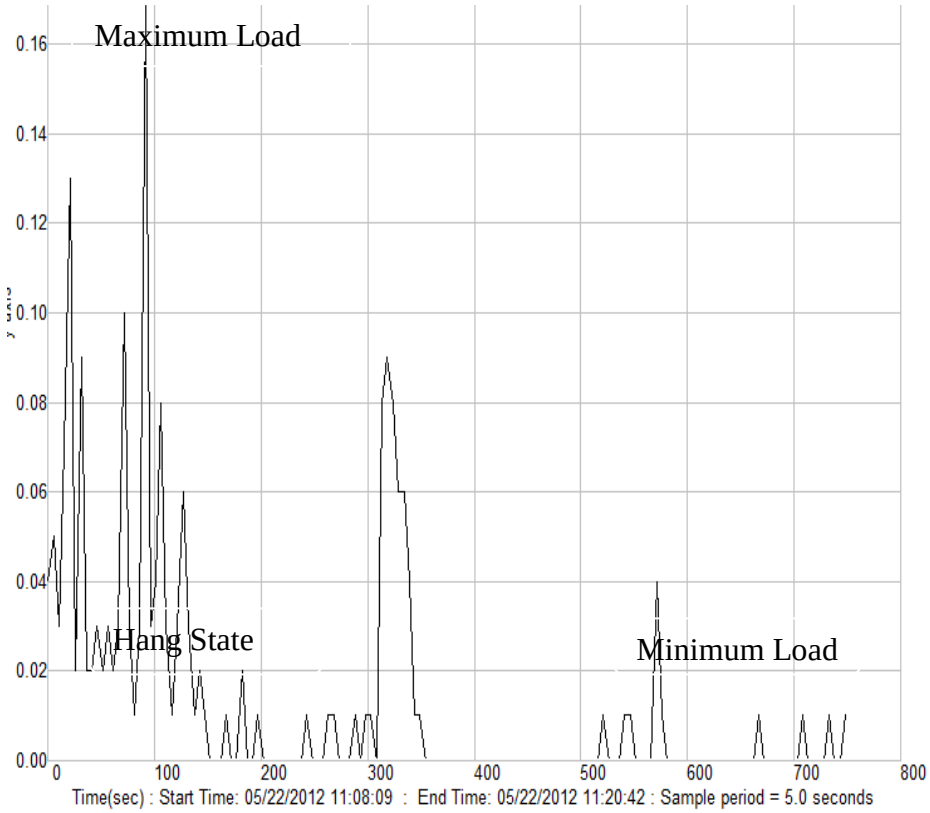

Fig. 5.6 % overlap Graph

The percentage of time spent by system services on behalf of other worlds.

%Used

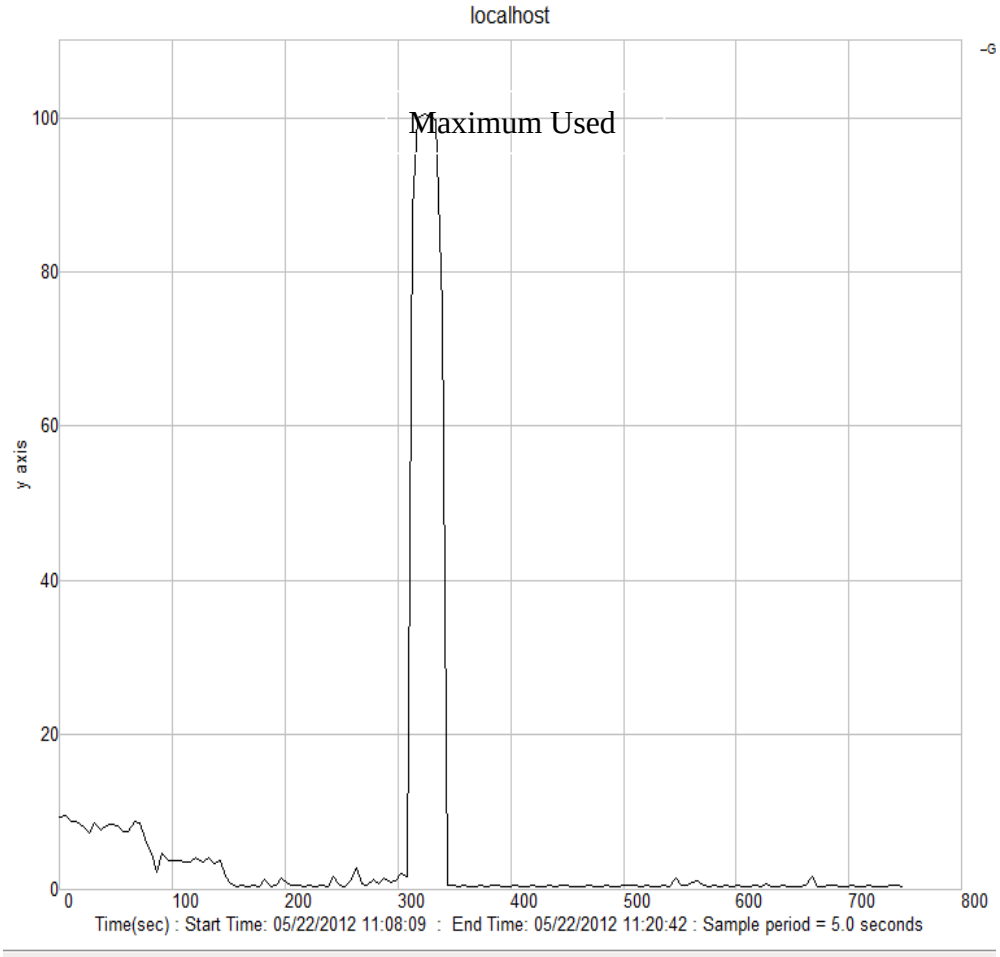

Fig. 5.7 % Used graph

The percentage physical CPU time accounted to the world i.e. group of statistics. If a system service runs on behalf of this world, the time spent by that service (i.e. %SYS) should be charged to this world. If not, the time spent (i.e. %OVRLP) should not be charged against this world.

%USED = %RUN + %SYS - %OVRLP

%Run:

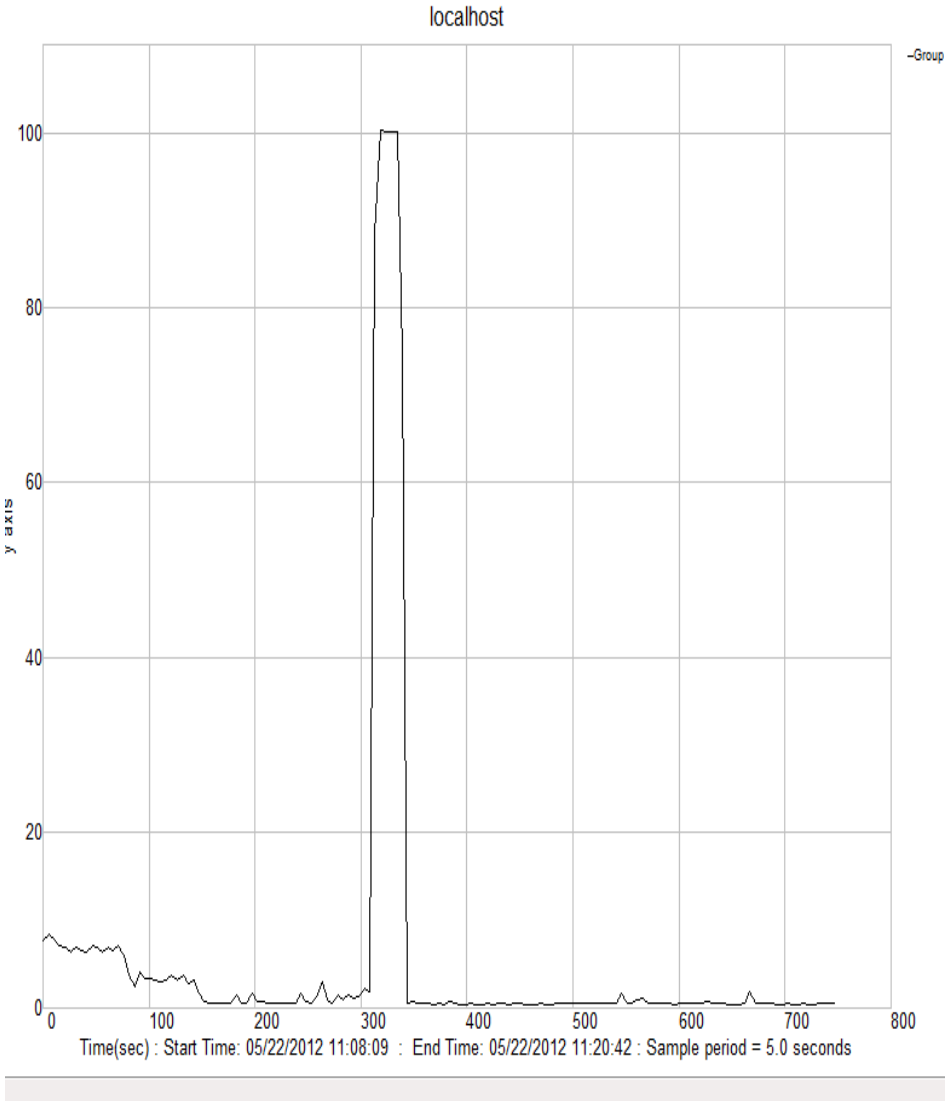

Fig. 5.8 %run

The percentage of total scheduled time for the world to run.

% System

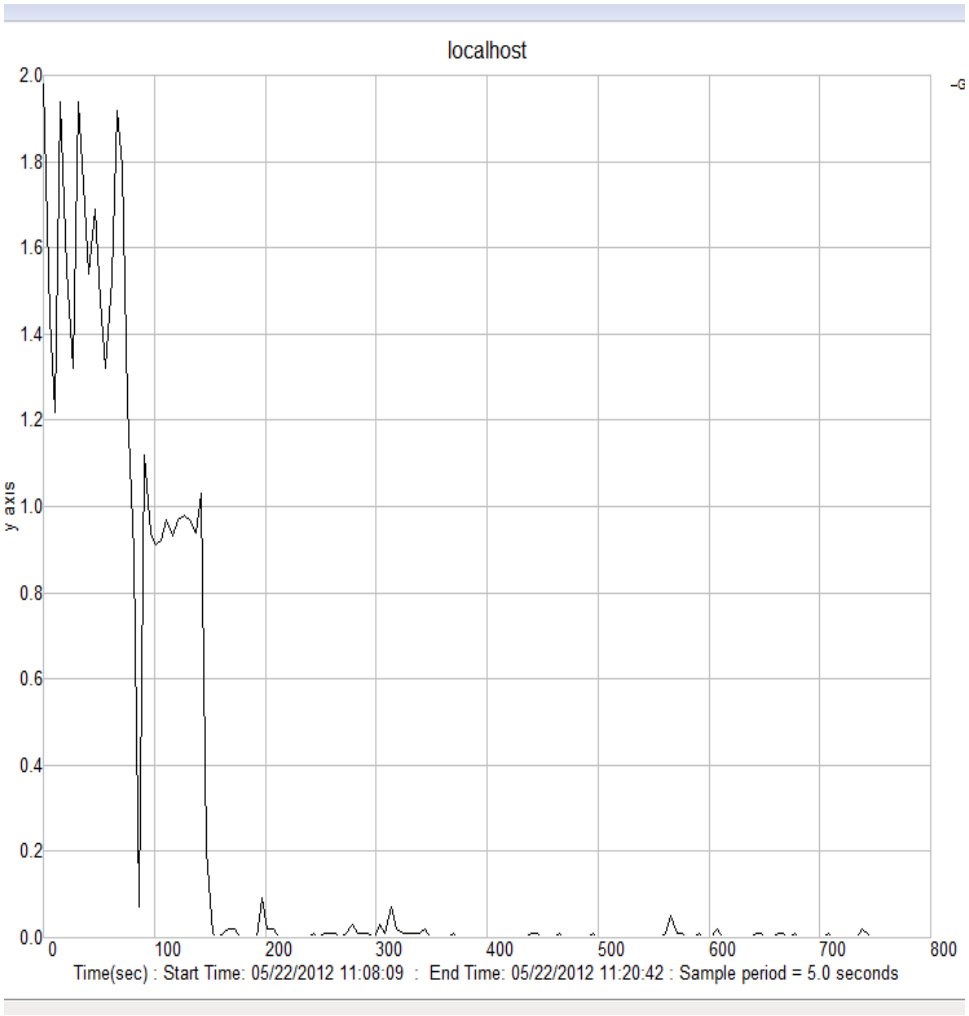

Fig.5.9 % System

Individual VM CPU performance

The below graph shows the CPU performance of virtual machines by using vSphere client.

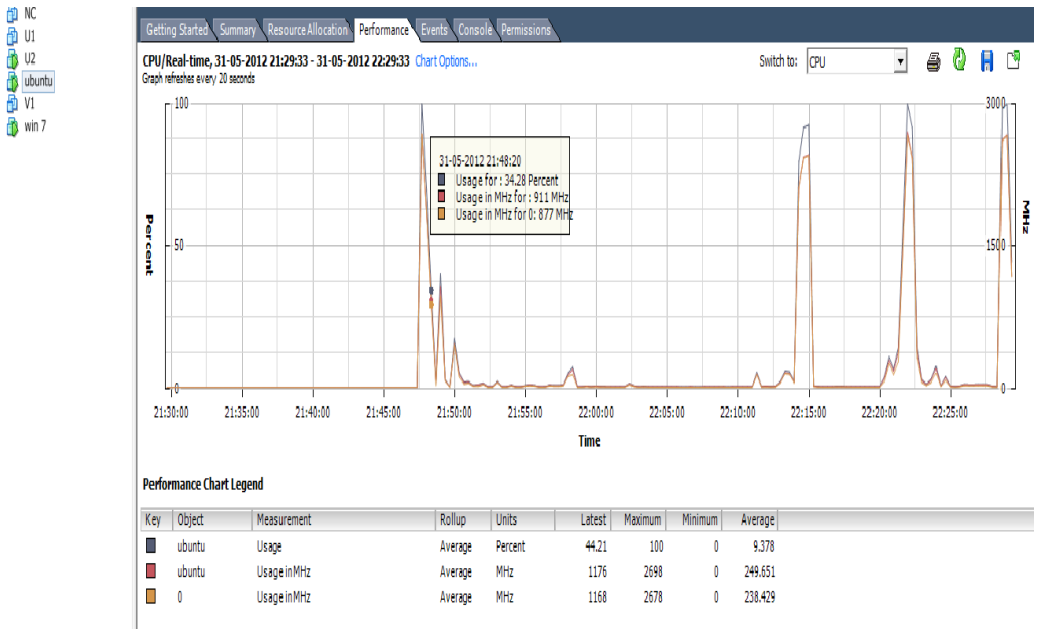

Fig. 5.10 Individual VM performance

This shows the usage of the running virtual machines for generating the load with time.

5.4 Comparisons of Hypervisors

Here, final comparisons of hypervisor are shown in the Table 5.1 below.
A Comparison is performing between the Xen hypervisor and VMware ESXi 5.0.

Table 5.1 Comparisons of Hypervisors

Hypervisor Attributes	**VMware ESXi 5.0**	**Xen Server**
Small Disk Footprint	144MB disk footprint	>1GB
OS independence	Not reliable on general Purpose OS	Relies on Linux in Dom0 management partition
Advanced Storage Management	VMware vStorage VMFS, Storage vMotion, Storage DRS.	Lacks an integrated cluster file system, no live storage migration, storage features support very few arrays.
Host Resource Management	Network traffic shaping, per-VM resource shares, set quality of service priorities for storage and network I/O	Lacks similar capabilities
Performance Enhancements	AMD RVI, Intel EPT large memory pages, universal 32-way vSMP, VMI paravirtualization, VMDirectPath I/O, PV guest SCSI driver	No large memory pages, no paravirt guest SCSI device, Requires inflexible SR-IOV

6. Conclusion and Future Work

6.1. Conclusion

This thesis work concludes different types of hypervisors and their architectural differences.

It has been shown that virtualization can play different roles in a computing system, be it for security, consolidation, etc. Once the role is known a proper VM mechanism can be chosen and evaluated allowing minimal performance impact to be achieved.

Each of the virtualization methodologies presents its own benefits. In cases where the guest operating system cannot be modified for enhanced virtualization performance, only native virtualization and emulation methods may be implemented. In such cases where modification to the guest operating systems is available, the enhancements afforded by para-virtualization can and should be taken advantage of. All the mechanisms and methods must still be secured, both for each virtual machine on its own as well as the physical host. The final analysis must be then that the selection of a virtualization method for a single physical host must be based on application based policies and reassessed as new technologies become available.

6.2. Future work

It is proposed that future work be done in developing a methodology for selecting a virtualization solution based on a flow model. The taxonomy combined with a selection methodology would allow for standardized approach to choosing virtual machine software solution.

There are also many hypervisors available in market which is new to virtualization and can be used for further research. Some of the hypervisor security products can be as ***HookSafe*** is hypervisor based kernel root-kit protection system proposed by Wang [66]. ***sHype*** is a secure hypervisor system developed by IBM research.

References

[1] “Enterprise Virtualization Software Consulting.” IDC, July 12, 2007. <http://virtualization.info/en/news/2007/07/idc-predicts-virtualization-services.html >.

[2] R. J. Adair, R. U. Bayles, L.W. Comeau, R. J. Creasy, “A Virtual Machine System for the 360/40,” IBM Corporation, Cambridge Scientific Center Report No. 320-2007, 1996.

[3] “Virtual Machine security guidelines, Version 1.0,” Copyrighted; The Centre of Internet Security, September, 2007.
<http://cisecurity.org>

[4] J. Brodken, “With long history of Virtualization behind it, IBM looks to the future,” Network world. UNESCO, April 30, 2009. <http://www.networkworld.com/news/2009/043009-ibm-virtualization.html>.

[5] Al-Rabayah, “Virtualization concepts and history.” Remote IT Services. January 24, 2010. <http://www.remoteitservices.com/content/virtualization-concept-and-history>.

[6] Popek, J. Gerald ; Goldberg, P. Robert : “Formal Requirements for Virtualizable Third Generation Architectures, “ Communications of the ACM, vol. 17, pp. 412-421, July 17, 1974.

[7] M. Jones, “Discover the Linux Kernel Virtual Machine.” IBM. <http://www.-128.ibm.com/developerworks/linux/library/l-linux-kvm/>.

[8] J. Krich, “Virtual Machine security guidelines, the center for internet security,” September 2007.
<http://www.cisecurity.org/tools2/vm/CIS_VM_Benchmark_v1.0.pdf>

[9] A. Mann, “The pros and cons of virtualization,” BTQ, 2007. <http://www.btquarterly.com/?mc=pros-cons-virtualization\&page=virt-view%research>

[10] Chaudhary, V. Minsuk Cha, Walters, J.P., Guercio, S., Gallo, S. “A Comparisons of Virtualization Technologies for HPC.” Advanced Information Networking and Applications, AINA, 22nd International Conference, page no. 861-868, March 25, 2008.

[11] Rule, David-Dittner, Rogier, “The Best Damn Virtualization Book Period. Burlington, Massachusetts. Syngress Publishing, Inc.” 2007.

[12] “Introduction to Real Time Virtualuzation, NI Technical Symposium, 2009. <http://www.ieee.li/pdf/viewgraphs/ni_real-time_hypervisor.pdf>

[13] Massino, “ESX Architecture.” IT 2.0, Next Generation IT Infrastructure, June 17, 2007. <http://it20.info/2007/06/a-brief-architecture-overview-of-vmware-esx-xen-and-ms-viridian/>

[14] I.Ahmad, Anderson, J.M. “An analysis of disk performance in VMware ESX Server virtual machines.” Workload Characterization, WWC-6, IEEE International Workshop, October 27, 2003.

[15] S. Stephen David, Marshall, Beaver, McCarthy, W. Jason, “VMware ESX Essentials in the Virtual Data Center, Boca Raton, Florida. Taylor & Francis Group, LLC,” 2009.

[16] N. Matthews Jeanna, M. Dow Eli, F. Wilbur Patrick, Deshane Todd, Hu Wenjin, Bongio Jeremy, Johnson Brendan, “Running Xen : A Hands-On Guide to the Art of Virtualization 1,” Prentice Hall, 2008.

[17] “Microsoft Hyper-V Server: Overview,” March 02, 2010. <http://trendsbuzz.com/q/Hyper-V>

[18] “Hyper-V.” <http://www.microsoft.com/download/en/details.aspx?id=2428>

[19] “Microsoft Server 2008.” <http://www.microsoft.com/en-us/server-cloud/windows-server/default.aspx>

[20] Chappell, David 2008; “Virtualization for Windows: A Technology Overview,” March 03, 2009.

[21] D. Marshell. Whitepaper: “Virtual Machine Security Guidelines. InfoWorld,” September 2007.

[22] “Virtualize Your IT Infrastructure,” IBM Systems. < http://www-03.ibm.com/systems/virtualization/infrastructure/>

[23] M. Rosenblum and T. Garfinkel, “Virtual Machine Monitors: Current Technology and Future Trends,” Computer, 38(5): 39-47, May 2005.

[24] J. Renato Figueiredo, A. Peter Dinda, and J. Fortes, “A Case for Grid Computing on Virtual Machines,” In ICDCS’03, Proceedings of the 23rd International Conference on Distributed Computing Systems, page 550, Washington, DC, USA, 2003.

[25] Campbell Seam, Jeronimo Micheal, “Applied Virtualization Technology, Usage Models for IT Professionals and Software Development.” Copyrighted; Intel Corporation, 2006.

[26] “An Analysis of server virtualization, 2011.” <http://www.gartner.com/id=1382327>

[27] “A survey on virtualization technologies, 2008.” <http://www.esg-global.com/lab-reports/microsoft-sql-server-2008-r2-and-hyper-v-r2-sp1-performance-analysis>

[28] “The great big Hyper-V survey, 2011.” <http://www.hyper-v.nu/archives/tag/survey/>

[29] “VMware vSphere hypervisor.” <http://vmware.com/products/vsphere-hypervisor/overview.html>

[30] “Services at virtualization layer,” <http://blogs.cisco.com/datacenter/enabling-the-network-service-layer-for-virtualized-and-cloud-infrastructure/>

[31] Agarwal, El Abbadi, A. Das, S. Elmor, AJ “Database Systems for Advanced Applications, Proceedings of the 16th International Conference, Dasfaa, 2011.

[32] “ESXi basics,” <http://www.vmware.com/products/vsphere-hypervisor>

[33] “VMware vCenter Server.” <http://www.vmware.com/products/vcenter-server/overview.html>

[34] “Using the vSphere Client.” <http://pubs.vmware.com/vsp40_e/admin/c_using_the_vi_client.html>

[35] “Using the vSphere Web Client.” <http://pubs.vmware.com/vsphere-50/>

[36] “VMware vSphere Web Services SDK Documentation.” <http://www.vmware.com/support/developer/vc-sdk/>

[37] VMFS, “VMware vStorage Virtual Machine File System,” A VMware Technical White Paper for VMware vSphere. <http://vmware.com/files/pdf/techpaper/VMware-VMFS-Tech-Overview.pdf>

[38] “Best practices Using VMware Virtual SMP.” <http://www.vmware.com/resources/techresources/240>

[39] “VMware vMotion.” <http://www.vmware.com/files/pdf/VMware-VMotion-DS-EN,pdf>

[40] “Migration with Storage vMotion.” <http://pubs.vmware.com/vmware.vsphere.vcenterhost.doc_50/>

[41] “How vSphere HA works.” <http://pubs.vmware.com/topic/com.vmware.vsphere.avail.doc_50/GUI33A65FF7-DA22-4DC5-8B18-5A7F97CCA536.htm>

[42] “Using vSphere HA and DRS Together.” <http://pubs.vmware.com/topic/com.vmware.vsphere.avail.doc_50/GUID-1D8B1384-59A4-41E2-AF05-697FC06D9EF9.htm>

[43] “vSphere 5.0 Storage Features, Storage DRS-Initial Placement.” <http://blogs.vmware.com.vsphere/2011/07/vsphere-50-storage-features-part-5-storage-drs-initial-placement.html>

[44] “vSphere Fault Tolerance.” <http://www.vmware.com/products/fault-tolerance/overview.html>

[45] “vSphere Distributed Switch.” <http://pubs.vmware.com/vsphere.com/vsphere50/com.vmware.vsphere.privileges.doc_50>

[46] “vSphere components and features.”
<http://pubs.vmware.com/vsphere-4-esx-vcenter/com.vmware.vsphere.intro.doc_41/c_vmware_infrastructure_components.html>

[47] “Viewing the console of a virtual machine.”
<http://www.symantec.com/docs/HOWTO14611>

[48] “ESXTOP.”
<http://www.Yellowbricks.com>

[49] “SSH.”
<http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>

[50] “VMware ESXTOP commands for Storage Monitoring-VM install.”
<http://www.vminstall.com/vmware-esxtop-commands-for-storage-monitoring/>

[51] “ESXPLOT.”
<http://labs.vmware.com/flings/esxplot>

[52] Tirbutt, Edmund, “Brimming with confidence; Benchmarking your Perks against your Rivals Can Provide HR with Added Reassurance.” Employee Benefits, November 2004.

[53] “Benchmarking Fundamentals.”
<http://reliabilityweb.com/excerpts/WiremanChap2PDF.pdf>

[54] “Bechmark Tools.”
<http://majorgeeks.com/downloads4.html>

[55] “Tools for Bechmarking.”
<http://benchmarkhq.ru/english.html?/be_cpu.html>

[56] “Phoronix benchmark tool.”
<http://ubuntuusers.de/Phoronix_Test_Suite>

[57] “Sysbench download.”
<http://sf.net/projects/sysbench/>

[58] “VMware Compatibility Guide.”
<http://www.vmware.com/resources/compatibility>

[59] “Steps for UNetbootin.”
<http://www.webudp8.org/2009/04/4-ways-to-create-bootable-live-usb.html>

[60] “Getting started with ESXi Embedded.”
<http://www.vmware.com/pdf/vsphere4/r40/vsp_40_esxi_e_e_eget_start.pdf>

[61] “VMware ESXi step-by-step Installation Guide with Screenshots,” TechRepublic, November 19, 2010.

[62] “DNS Settings.”
<http://www.ntchosting.com/dns/settings.html>

[63] “Host System Requirements- VMware.”
<http://www.vmware.com/support/ws5/doc/intro_hostreq_ws.html>

[64] “Using Perfmon for esxtop-based Performance analysis.”
<http://communities.vmware.com/docs/DOCS-5100?tstart=870>

[65] “How to use CSV files.”
<http://www.imf.org/external/help/csv.htm>

[66]Z. Wang, X. Jiang, W. Cui, and P. Ning. “Countering kernel rootkits with lightweighthook protection,” Pages 545-554, 2009.

List of Publications

Hiteshi, Dr. V.P. Singh, "An analysis of system performance VMware ESXi Server virtual machines" communicated in Journal of Computer Technology and Applications, in June 2012 **(Communicated).**