ISSN 2229-5518

Examining Virtualization: A Survey into Fundamentals and Techniques

DebabrataSarddar, Bhavya Deep, Rajesh Bose

Abstract - As a force multiplier, virtualization is an invaluable feature to computer systems and information technology, today. The very act of virtualizing physical resources accords flexibility, robustness, and optimum resource utilization. Going virtual has multiple benefits. These can be witnessed in the form of system virtualization, storage, and network virtualizations. Globally, data centers rely almost exclusively on virtualization methods to increase availability while keeping costs under relative control. The key to virtualization is that almost a wide range of dissimilar software can be run on a single physical machine without interference from one another. Advancements made in technology allows us to use even an array of individual physical machines that together act as a unified platform to virtualize independent software environments. In this paper, we discuss the fundamentals of virtualization and examine prominent virtualization products that are currently in production use.

KeyWords – Virtualization, Data Center, Network, Thin-client, Server.

1. Introduction

Worldwide, data and information multiplies incessantly churning out enormous volumes. Individuals, organizations, and companies have access to almost limitless data. With the sea of information, there been a spurt in the growth of methods and technologies to interpret the data available.

DebabrataSarddar, Assistant Professor in the Department of Computer Science and Engineering, University of Kalyani, Kalyani, Nadia, West Bengal, INDIA. His research interest includes wireless and mobile system and WSN, Cloud computing.
Email - dsarddar@rediffmail.com

Bhavya Deep is Assistant Professor in Department of Computer Science, Bhaskaracharya College of Applied Sciences, University of Delhi, Delhi, India. Currently, he is principal investigator of DU sponsored innovation project titled "Study and Design of framework for cloud implementation in a large university / Higher education. His research interests include cloud computing, networking and internet of things.
Email – bhavya.deep@bcas.du.ac.in

Rajesh Bose is currently pursuing Ph.D. from Kalyani University. He is an IT professional employed as Senior Project Engineer with Simplex Infrastructures Limited, Data Center, Kolkata. His research interests include cloud computing, wireless communication and networking.
Email – bose.raj00028@gmail.com

Virtualization technologies are critical in supporting advancements necessary to support data storage and analysis. It is important to note though, that the genesis of virtualization lies embedded in the days when mainframes were the only means to crunch huge volumes of data. The exponential power of current desktop and server hardware has made it possible forvirtualization techniques to be ported [1] on such. The rebirth of virtualization on more compact hardware that is also cheaper to run and maintain, has invited the attention of the IT industry and rejuvenated the cause for delving in virtualization. In the race to stay more focused and competitive in the long run, several virtualization products [2] now find acceptance in diverse applications. The core focus in virtualization is to delink the hardware from software layers. Such delinking empowers IT administrators with the ability to respond to changing load requirements without being tied down to any single infrastructure pattern. Thus, virtualization has become a standard mantra in almost all data centeridioms and environments. As an example of virtualization, figure 1 is shown which depicts the basic difference between traditional IT deployment model as opposed to virtualized one. The two blocks are labelled "Traditional" and "New" in this figure. An IT infrastructure is essentially one that processes information that flows through it. It constitutes all that includes physical hardware that supports users either directly or indirectly. An IT infrastructure only becomes complete with inclusion of the software that not only regulates passage of information over the hardware, but also managing of information per se.

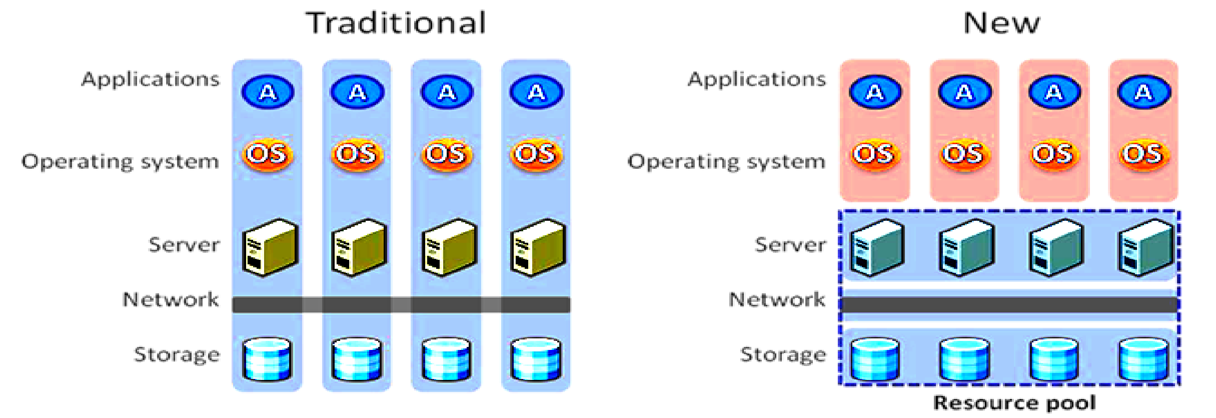

Fig.1. Traditional and new IT infrastructure

The two separate blocks as shown in figure 1 consist of identical components that includes storage, network, server, operating systems, and application layer. In the traditional model, components of a layer are connected individually to corresponding components of the layer adjacent. Each server contains an operating system with a set of applications running, and served by a network and a single storage device. The network connects the servers to the storage systems, and also connects servers to the outside world. The blue-coloured columns represent a specific combination of storage system, network, server, operating system, and applications. Unsurprisingly, therefore, a failure that affects any one of these components is sufficient to cause the entire column to be destabilized. As an example, a server hardware failure is sufficient for a business process, relying on the particular column, to collapse. Unless the hardware failure is fixed, the operating system and the applications hosted on that particular server, remain down. The block on the right in the figure, the blue band encompassing storage units, network, and servers represent a common resource pool. The combination of operating systems and application software are shown in an orange band sitting on this common resource pool. This indicates that neither the operating systems nor the application software are tied down to any specific server or hardware asset. The advantage of being connected to a resource pool is that failure of any hardware component has almost little or no effect at all on the overall functioning of a business process. This is a simplification of some of the more intricate and sophisticated virtualization technologies that are currently in the market today[4]. It should be borne in mind that virtualization technologies are as different from one another as the different layers are in an IT infrastructure. Combined these technologies morph a traditional IT infrastructure into one that is virtual. There are several players in the market today who offer similar virtualization technologies using diverse approach and processes. It is not unknown for a virtualization service provider to offer a basket containing a heterogeneous mix of virtualization technologies from a broad range of software manufacturers. This practice is often followed to provide bespoke IT solutions that are relevant and unique in many cases where large enterprises are involved.

2. VIRTUALIZATION APPROACHES [5]

In traditional environments, physical servers are often directly connected to one or more physical switches. It, therefore, becomes possible for IT administrators to obtain detailed management reports, by monitoring the switch, on the traffic that emanates from the servers. On the other hand, a similar level of data cannot be gleaned from a virtual switch. In fundamental terms, a virtual switch is linked to the physical switch using a physical NIC that is connected to virtual machines. Consequently, the hindrance of not being able to oversee the flow of traffic among connected virtual machines on a similar physical plane, limits security and diminishes the ability to gauge performance. Several approaches exist to virtualization, each with its own set of features that help control virtual machine performance. The underlying architectures of such approaches are shown in Figure 2.

ISSN 2229-5518

(a) Operating system-based Virtualization

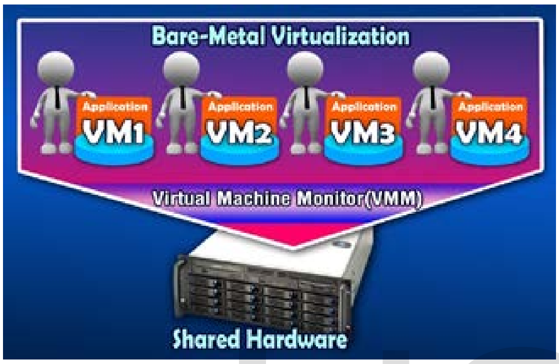

(b) Application-based Virtualization

(c) Hypervisor-based Virtualization

Fig.2. Virtualization Approaches[5]

2.1. Operating system-based virtualization

By far, one of the more recognizable virtualization technology is the one that has been depicted in figure 2(a). Using a single host operating system, several individual guest operating systems are hosted on a single physical machine. Each of the guest operating systems run independent of each other. However, being on the same host operating system, a common kernel retains control on the underlying hardware. The operating systems on each of the virtual machines so running on the host operating system, can be monitored and managed by the latter. This approach has several drawbacks – chief of which being compromise in performance metrics. All input / output operations from active virtual machines have to be handled simultaneously by the host operating system. This, therefore, leads to an automatic increase in overheads, and thus, results in lowering of performance.

2.2. Application -based virtualization

Application-based virtualization is no different from the previous discussed method. In this approach, individual virtual machines are created with their own guest operating systems and related applications. It is the virtual machine monitor (VMM) that emulates each virtual machine. Commercially, this form of virtualization has not found as many takers as other competing designs. As in the previous model, here too, security concerns are key factors that prevent its widespread adoption and use.

2.3. Hypervisor -based virtualization

Virtualization using hypervisor entails direct interaction with hardware resources of a physical machine. As depicted in figure 1.c, a hypervisor is the platform on which individual virtual machines are setup with their respective operating systems. The hypervisor acts to keep the individual virtual machines in a compartmentalized environment. Each virtual machine remains unaffected by the performance of the other. Since each virtual machine runs its own operating system software, it is possible to have one or more virtual machines running dissimilar operating systems. One of the crucial functions of the hypervisor is to continuously monitor hardware resources of the physical machine. It is the job of the hypervisor to allocate hardware resources to the virtual machine running a resource-intensive application. Since, a hypervisor itself requires some amount of hardware resources to function, the physical machine supporting it must set aside adequate resources for proper functioning of the hypervisor itself. Consequently, it is not unknown for applications to experience high latency at times. Physical machine performance also diminishes as a result.

3. Virtualization Categories:

The taxonomy of virtualization can be stated in terms of three broad categories. These are described in the following sections. It is important to note that the objective of virtualization is to utilize physical resources as much as possible without incurring additional overheads to achieve the objective.

3.1. Hardware Virtualization

By far, hardware virtualization is the most popular form of virtualization that has gained worldwide support and acceptance. A hardware virtualization typically involves setting up a virtualization software on a physical machine. The job of this virtualization software is to create a digital representation that records all the characteristics of the physical host machine. Such a digital representation is termed as "image". The purpose of creating this "image" is to setup an emulated environment or a "virtual server" that can run with identical characteristics as that of the original physical host. For the "image" to be ready for use, it has to be converted into a "virtual server". This is accomplished when this "image" is deployed on an existing array of virtual servers using commercial technologies such as VMware, Hyper-V, etc. Converting physical servers to virtual machines in this manner offer several advantages. These advantages are discussed next.

1. Upgrades to virtual servers become effortless as RAM and CPU resources can be adjusted at runtime.
2. Physical machines that are deemed old can be repurposed to function as additional hardware resources in a cluster of physical servers that support virtual server farms.
3. Physical component failures while being important, are rather relegated to a status wherein these are deemed to be less of a concern. Hardware can now be utilized beyond a point where such would have been deemed outdated.
4. Electric power consumption can be regulated to a finer degree. As a result, power distribution and consumption efficiency can be raised. Thus, a smaller carbon footprint becomes an achievable target [7].

3.2. Desktop/Application Virtualization

Building on the traditional thin-client model, a new paradigm in virtualization has gained mainstream acceptance. Desktop virtualization, which is also known as Virtual Desktop Infrastructure, was designed to offer a twin advantage. The first of which was an advantage that IT administrators can use to deploy applications from within data centers. These applications are projected directly to the desktops of users. The actual desktops need not be full-featured machines. These only need be light-weight without the complexities of information processing and storage. The second advantage is that the users are still able to enjoy an experience of comparable to that of using a full-featured desktop [8]. Globally, desktop virtualization is fast gaining traction as a means to replace costly traditional desktop deployments. The underlying principle in desktop virtualization is the establishment of operating system and application software execution to a data center. The local device is freed from the burden of conducting any form of software execution. A data center, in any remote location, assumes those functionalities. As a result, the actual end-user device is transitioned to a mere lightweight computer that needs to manage its immediate peripherals and accessories. Such include monitor, mouse, keyboard and may extend to scanners and printers at the most. Remote protocols [9] [10] are used to maintain connectivity between data center that extends desktop virtualization on to the actual end-user device. The advantages of adopting desktop virtualization manifest themselves in lower management costs in terms of infrastructural deployment, as well as reduced labour costs for installation and maintenance. Further, local troubleshooting and diagnosis becomes simpler [11] as the end-user device is no more than a terminal with no need to retain user interactions by itself. Desktop virtualization offers several advantages. These also include a number of those which are found in application virtualization.

High Availability

Downtime can be significantly reduced using techniques that invoke replication and fault-tolerance.

Extended Refresh Cycles

Life span of physical hardware of client PCs can be enhanced. As a result of lower processing demands placed by deploying larger capacity servers, client PCs do not have to operate at high performance levels.

Multiple Desktops

A unique advantage of using this form of virtualization is the ability to access a wide range of multiple desktops from the same end-user PC. As such, a multitude of disparate tasks can be accomplished without the need to upgrade existing PC hardware. A significant number of proposed desktop virtualization systems [12] [13] [14] rely on technologies that produce whole desktops. These are tuned to extend superior quality display effects, improve user interactions, and/or reduce delays in transmitting the virtual desktops to clients. However, in doing so, almost all these protocols bypass issues that are related to mitigating granularity of virtual desktop displays. Nor do these protocols achieve anything of significance to optimize the performance of the protocol used to establish communication between the servers and the clients. The latter feature being extremely important in the context of supporting flexibility and ensuring extendibility of the system as a whole. By itself, the term "application virtualization" is a nomenclature that encompasses software that are designed to make legacy applications accessible beyond the respective operating systems which host these. An application that has been virtualized, is not installed as is the norm. Instead, it is still operated as if it has been installed. In operating system virtualization, the whole of the operating system is virtualized. This is different from application virtualization where only specific applications are virtualized [8]. Utilizing local and stream application virtualization approach, an application can be installed on-demand. With streaming enabled, select portions are sent first which are required for the application to start. These optimize the start-up times. It is not unknown for virtual application to create virtual registries and use local file systems to enable compartmentalization within user's physical machine. Prominent local application virtualization solutions are sold by Citrix and Microsoft. Virtual appliances, as those which can be consumed using VMware's VMware Player, also fall into this category of application virtualization. Following are the benefits of application virtualization:

Security

By running virtual applications in user mode, these are isolated from the operating system.

Management

Virtual machines can be easily managed and upgraded when required, from a central location.

Legacy Support

Legacy applications that were not originally designed to run on modern operating systems, can be made to operate using virtualization technologies.

Access

Virtual applications can be installed from locations that allow failover and replication. As a result, virtual applications when deployed on demand, are seldom at risk of being unavailable to users. A fully featured PC, on the other hand, may fail at any time leaving a user with no access at all to the applications installed in it locally.

3.3. Network Virtualization

Network virtualization may be thought of a single software structure that combines hardware and software networks into one. Network virtualization combines platform and resource virtualization. It is classified as being either external, or internal [8]. In its external form, network virtualization combines several networks, or even portions of networks, into a virtual one. In the internal variant, functionality akin to actual networks is extended to software in a single system [8]. The latter definition enables server and desktop virtualization solutions to extend network access between a host and its guest, or among several guests. Virtual switches are being readily accepted as an integral part of server virtualization stack. The external form, or definition, of network virtualization is more popular of the two. Virtual Private Network (VPNs) are indispensable to network administrators' repository of tools. Virtual LANs (VLANs) are not altogether uncommon. VPNs are almost always used across many organizations to facilitate communications. Advancements in network technologies, such as 10 Gigabit Ethernet, networks are no longer needed to be structured on geographical topologies. Companies such as Cisco and 3Leaf have already forayed with products that offer the following benefits of network virtualization.

Tailored access

Network administrators can modify network parameters that lead to immediate impact on available bandwidth and quality of service.

Strengthening of networks

A single virtual network can be built by federating physical networks. This significantly improves network manageability.

However, as in the case of server virtualization, network virtualization also invokes some amount of performance overheads. It is also a complex process, and as such requires administrators to have broader and, preferably, in-depth knowledge and experience in handling such.

3.4. Storage Virtualization

Storage virtualization defines the process of forming logical storage from physical ones [8]. To a limited extent, RAID offers an almost similar functionality. However, storage virtualization encompasses features such as caching and data migration. It is not easy to pin down the definition of storage virtualization as there are a number of ways in which it can be offered. Storage virtualization is an amalgam of the following:

- Special device drivers that rely on the host.
- Array controllers.

- Network switches
- Standalone network appliances.

Almost every vendor differs in their approach to offering storage virtualization.

Storage virtualization is also classified in terms of symmetric and asymmetric virtualization modes. The former supports caching. The latter uses special device drivers that reside in the host to query the meta data. The meta data indicates physical location of a given file. This allows the host to access the file within the storage. This approach does not, however, support caching. Benefits of storage virtualization are enumerated as follows.

Migration

Easy migration of data between storage locations without affecting performance.

Resource optimization

Manageability of storage devices can be enhanced, thus, leading to offset of either over- or under-utilization.

Manageability

A single physical device can be centrally controlled to assist host machines to leverage storage space.

4. VMWARE

With virtualization, the capabilities of a single desktop can be vastly amplified. Most desktop PCs which are powered by advanced microprocessors, are seldom fully utilized. Virtualization allows a single physical machine to host multiple virtual ones. Thus, the sophisticated hardware resources that cannot be harnessed sufficiently by a single machine alone, can now be exploited to the hilt. Several virtual machines running atop a single physical one can simultaneously utilize the available resources to the fullest possible extent. This is made possible as a result of compartmentalized and tightly integrated structure of a virtual machine. Each virtual machine is a separate digital entity which acts like its physical counterpart. VMware has designed its Workstation software to link virtual machine resources to corresponding physical resources. VMware Workstation enables each virtual machine to operate with its own physical resources which include CPU, memory, storage, and network devices. Further, a virtual machine is also able to connect directly to physical resources of the host machine. Besides emulating hardware, VMware Workstation offers the ability to configure virtual network adapters to use Network Address Translation (NAT). Thus, individual virtual machines can be allocated IP addresses [5] to communicate with the outside world.

Flagship products

The IT industry, as a whole, recognizes ESX, or ESXi, as VMware's hypervisor. This product has been branded by VMware as vSphere Hypervisor. Although offered as a free version, it has been meant to be used as such for limited testing and evaluation purposes. VMware also offers a free version of VMware Server for testing and evaluation. This product partitions a physical server to form several virtual machines. It can be run off of any server hardware. VMware, however, markets VMware vSphere in several editions, as its leading product. With vSphere 4.1, VMware has packaged advanced memory management functionalities, active directory, and high availability features [15].

Standout features

VMware offers several class-leading functionalities that are aimed at enhancing manageability of virtual machines. With the introduction of these features, VMware offers IT administrators to effect seamless transitions from physical machines to virtual ones. The following salient points also illustrate the possible weakness of virtualization, and how VMware has progressed in addressing these on a broad front with its bouquet of virtualization products.

High Availability

VMware High Availability feature allows failovers to take place in case of physical machine failures. In the event of an actual physical machine failure, VMware HA triggers a rapid response that sets in motion restarts of all affected virtual machines on available servers which have sufficient resources to spare. If the failure is only an operating system one, VMware HA automatically restarts the affected virtual machine on the same server. With the incorporation of VMware HA features to the vSphere platform, IT administrators are able to control the degree of availability required by important applications individually or as federation [16].

vMotion

One of the primary requirements when implementing virtualization, is reduction of downtime. VMware's vMotion offers live migration wherein actively running virtual machines can be moved from one physical server to another without the need to shut down any. This, therefore, automatically ensures service availability. vMotion has become an integral part for creating a data center that is not only dynamic, but is also self-regulating and automated [17].

vCenter

VMware positions its vCenter Server as its leading tool in managing virtual machines through a single interface. This enables IT administrators an in-depth view into the functioning of individual components of each and every virtual machine running in a data center environment. The functionality enhances productivity in context of physical infrastructure management [18].

5. Microsoft

Microsoft offers Hyper-V as its answer to virtualization challenges. From the perspective of the masses, Hyper-V offers familiarity bundled with Microsoft's proven administrative tools and management consoles. Realizing that the market share has been cornered by VMware, it has made Hyper-V R2 an integrated component within Windows Server 2008 R2. This allows for Hyper-V to start out as a less costly solution than comparable products in the market. It also goes on to provide as many as four distinct layers of management covering physical hardware, and application and services in virtual machines. The pivot of Hyper-V's allure is the ease of licensing which allows potential users to build virtual machines without sacrificing features. Microsoft has realized that for some, a single vendor for such solutions can be effective and considerably more manageable. Hyper-V R2 shares with Windows Server 2008 R2, the same driver for devices that are attached to host machines [15].

Microsoft's hypervisor

Plunging into the virtualization market, Microsoft has made Hyper-V a part of its Windows Server 2008 R2 operating system. The prominent features of this hypervisor from Microsoft are those which can be evidenced in the form of its live migration and shared VM storage features. Advanced functionalities such as enhanced memory support, failover, host clustering, and dynamic allocation of resources are provided in the Enterprise and Data Center editions of Hyper-V that come bundled with Windows 2008 R2. Microsoft provides licensing options that make adoption of Hyper-V seemingly effortless and attractive.

The features on offer with Hyper-V

A host of features and benefits have been packed by Microsoft in its hypervisor product. These are described below in seriatim:

1. Native hypervisor virtualization for 64 bit machines.
2. Support for concurrent operations of 32- and 64-bit machines.
3. Support for uniprocessor and multiprocessor virtual machines.
4. Creation and restoration of virtual machine snapshots to enable reinstatement to a restoration point.
5. Support for large virtual machine memory.
6. Support for virtual local area network (VLAN).
7. Console snap-in for Microsoft Management Console.
8. Interfaces for scripting and management for Windows Management Instrumentation.
9. Ability to migrate virtual machines without incurring downtime.
10. Dynamically resize virtual machine storage.

6. Citrix

Along with Microsoft, Citrix offers a cheaper virtualization alternative to VMware. However, the quality of its products come powered by industry-strength open-source products. This gives Citrix the edge by virtue of a very large developer base. Citrix is mostly preferred by customers because it offers desktop virtualization [15] alongside server virtualization. XenServer 5.6 is one of its class-leading products. Despite being free like some of its counterparts, XenServer 5.6 incorporates robust features that make it reliable in production environments. Citrix nudges ahead of its rivals when it comes to cloud computing. Combined with Citrix CloudPlatform, the software manufacturer offers XenDesktop and Citrix XenApp that enables delivery of application to any desktop from the cloud [20]. Translated in enterprise terms, most organizations favour Citrix product as means to convert legacy software into cloud-based virtualization. In addition to desktop virtualization, CloudPlatform offers mobile extensibility and usage.

Conclusion

The paper elaborates salient features of virtualization. It highlights the advantages and offers insights as to what is possible by going virtual. More than this, the paper surveys the areas of virtualization and the products that can help achieve the goals of virtualization. Of the three software manufacturers, VMware retains its position by a solid margin. With a desktop operating system user base that is unrivalled, Microsoft is not too far behind. Microsoft's server operating systems are none too far behind either in terms of popularity. To capitalize on this, it has incorporated its hypervisor product at almost zero initial cost. Although this seems advantageous, licensing costs in the long-run remain applicable. Citrix, on the other hand, approaches the issue of virtualization from three different angles. From the perspective of porting legacy applications to cloud, it offers two class-leading products from server and desktop points of view. Enterprise scalability and multi-tenancy seem to be handled rather well. In spite of the advantages offered by Citrix and Microsoft, VMware remains competitive by staying closely aligned to the very essence of virtualization. In trying to do so, it has created products that handle the issues of live migrations and central point management rather well. To a potential user, the merits and demerits of the three competing virtualization solutions become readily apparent through this study. It remains, though, the ultimate selection of the consumer to choose from these the one that appears to be more closely aligned with the benefits of virtualization discussed in this paper.

REFERENCES

[1] IDC. (2006). DOI=http://www.networkworld.com/news/2006/102506-idc - virtualization.html

[2] Bittman, T.J. (2009). Virtual Machines and Market Share Through 2012. Gartner. Publication date: 7 October 2009

[3] Searchdatacenter.com. Definition of infrastructure. DOI=http://searchdatacenter.techtarget.com/ sDefinition/0,,sid80_gci212346,00.html

[4] Dynamic quest. DOI=http://www.dynamicquest.com/virtualization.php

[5] DebabrataSarddar , Rajesh Bose, " Architecture of Server Virtualization Technique Based on VMware ESXI server in the Private Cloud for an Organization", International Journal of Innovation and Scientific Research ISSN 2351-8014 Vol. 12 No. 1 Nov. 2014, pp. 284-294

[6] Rajesh Bose, DebabrataSarddar, "A Secure Hypervisor-based Technology Create a Secure Cloud Environment", International Journal of Emerging Research in Management &Technology ISSN: 2278-9359 (Volume-4, Issue-2).

[7] DebabrataSarddar, Rajesh Bose, " Secure Co-processor and Billboard Manager Based Architecture Help to Protect & Store the Citrix Xenserver Based Virtual Data", COMPUSOFT, An international journal of advanced computer technology, 3 (1), January-2014 (Volume-III, Issue-I)

[8] Wikipedia

[9] Real VNC, "Vnc," http://www.realvnc.com/.

[10] Microsoft Corporat ion, "Remote Desktop Protocol," http://msdn2.microsoft.com/en-us/library/aa383015.aspx.

[11]Guangda Lai, Hua Song, XiaolaLin,A Service Based Lightweight Desktop Virtualizat ion System,© 2010 IEEE DOI 10.1109/ICSS.2010.44

[12] Citrix Corporat ion, "CitrixApplicat ion Delivery Infrastructure."

[13] "Kernel Virtual Machines," http://sourceforge.net /projects/kvm.

[14] Desktone, "Desktop as a Service," http://desktone.com/, 2008.

[15] Server Virtualization,White paper,

[16] VMware High Availability,http://www.vmware.com/files/pdf/VMware-High-Availability-DS-EN.pdf

[17] VMware Vmotion Live Migration for Virtual Machines Without Service Interruption,http://www.vmware.com/files/pdf/VMware-VMotion-DS-EN.pdf

[18] VMware vCenter Server Centrally Managed Virtual Infrastructure Delivered with Confidence,https://www.vmware.com/files/pdf/products/vCenter /VMware-vCenter-Server-Datasheet.pdf

[19]Overview of Hyper V,https://technet.microsoft.com/enus/library/cc816638%28v=ws.10 %29.aspx

[20] http://cloudplatform.accelerite.com/pdf/xendesktop-cloud-provisioning-with-cloudplatform-white-paper.pdf