A Survey on Security Aspects of Server Virtualization in Cloud Computing

O Sri Nagesh[1], Tapas Kumar[2], Vedula Venkateswararao[3]
[1,3]Departement of Computer Science and Engineering, Sri Vasavi College of Engineering, Tadepalligudem, India
[2]Departement of Computer science and Engineering, Lingayas University, Faridabad, India

Article Info

Article history:

Received Dec 10, 2016
Revised Mar 15, 2017
Accepted Apr 1, 2017

Keyword:

Cloud computing
Cloud services
ESX
IT infrastructure
Online banking
Security
Service provider
Threats
Virtualization
VMware
Xen

ABSTRACT

Significant exploitation and utilization of cloud computing in industry is come with and in the identical time vulnerable by unease regarding protection of data hold by cloud computing providers. One of the penalties of moving data processing and storage off business site is that organizations have fewer controls over their infrastructure. seeing that, cloud service (CS) providers must hope that the CS provider is capable to protect their data and infrastructure from both exterior and domestic attacks. Presently however, such hope can only rely on organizational procedures stated by the CS provider and cannot be remotely verified and validated by an external party. The central distinction between cloud computing and conventional enterprise internal Information Technology services is that the proprietor and the consumer of cloud Information Technology infrastructures are separated in cloud. This transform requires a safety responsibility severance in cloud computing. Cloud service providers (CSP) should safe the services they propose and cannot surpass the customers' authorities. Virtualization is a buildup utterance in the Information Technology world. With the assure to reduce the ever mounting infrastructure inside data centers connected to other important apprehensions such as ease of use and scalability, virtualization technology has been in advance recognition not only with IT experts yet also among administrators and executives as well. The progressively more growing rate of the approval of this technology has exposed these systems to new protection concerns which in recent history have been unnoticed or merely overlooked. This paper presents an in depth state of art gaze at present most old server virtualization explanations, as well as a writing study on different security matters found inside this virtualization technology. These problems can be practical to all the existing virtualization technologies accessible with no spotlight on a specific answer. Nevertheless, we do susceptibility investigation of two of the mainstream recognized virtualization answers: VMware ESX and Xen. to conclude, we illustrate some clarifications on how to progress the security of online banking and electronic commerce, using virtualization.

Corresponding Author:

O Sri Nagesh,
Departement of Computer Science and Engineering,
Sri Vasavi Engineering College,
Pedatadepalli, Tadepalligudem, WGDT, Andhra Pradesh, India.
Email: osri.pvpsit@gmail.com

1. INTRODUCTION

We are walking into the age of cloud computing. Cloud computing is a new type of Information Technology service in which customers make use of the cloud computing infrastructures such as CPU capability, network and storages provided by cloud service providers (CSP). Cloud computing can assist to condense Information Technology cost of SME's in that they require not to procure their own Information Technology infrastructures and utilize an Information Technology team again. For the meantime, as the price of customary domestic Information Technology communications is fetching a serious load to SME, cloud computing can also assist to build them supplementary aggressive [1]. Present we have three archetypal cloud service delivery models: Software as a Service, Platform as a Service and Infrastructure as a Service The inter rapport and logical margins of these 3 cloud service deployment models were described in the Cloud Reference Architecture as demonstrated in Figure 1 [2]. Because these three service release models offer diverse stages of service, the safety level that CSP promise to consumers should also be dissimilar. CSA believes that CSP should take on the greatest quantity of security accountability in SaaS model and the least amount of security dependability in IaaS model, while in PaaS model, security accountability must be carefully leveraged by the CSP and customers [2]. Virtualization has profoundly changed the information technology (IT) industry in different areas such as network, operating systems, applications or storage. Virtualization is a subject that not only IT people concern now. It has gained space on the administrators and directors vocabulary. Companies have apprehended that most of their systems were operates at share of 10 percent or less of utilization, yet these systems continue to require space, power and cooling system as any other machine. Tumbling these necessities would have a straight crash in dropping the Information Technology budget and atmosphere cares as the carbon footstep. Virtualization technology was the solution found by many companies, moving this technology into the mainstream. According to a recent IDC survey [2], companies that have deployed virtualization could see a return of investment of 472 percent in less than a year. The increased utilization and consolidation of x86 architectures had significant job for this also. Many companies use this architecture since it has lesser cost contrasted with others in the marketplace. Though, this structure had historically hardware support issues for virtualization which significant degrade the performance of the virtual machine weigh against with the same structure operates on a physical host. In order to solve this issue, AMD and Intel executed architectural additions to directly support virtualization in hardware. This conquers the traditional virtualization restrictions of the x86 architecture, improving significant characteristics like performance and scalability creating x86 server virtualization a keystone of most IT consolidation projects. The increasing investment and implementation of virtualization is comparable to the implementation of internet in companies at the end of the last decade. However, in the same way, security was not the top priority, Even though Information Technology administrators are now more rational to this theme. The threat of this latest technology are a little argued almost only at security actions such as Black Hat and it continues to be out of the focus of many allusion and advisor companies that employ this expertise. Virtualization has been presented to companies as an out-of-the-box answer that companies do not have to worry about, as if it was physical machines with the advantages that the hardware virtualized does not "crack" as the physical infrastructure. There is some mythology to split when condensate about virtualization security and this mythology happen because, as other myths, there is not much information about it. Some people assume that having, for instance, 4 virtual machines operates on a physical machine is the same as having 4 physical machines and so the concerns should be the same (e.g. only install patches on the operating system inside the virtual machine). Perhaps this myth exists because there is a common sense that hypervisors are impenetrable, which is false as we are going to see later in this paper. Some IT directors are not aware of the security intensity whereas using virtualization. For example, by means of virtualization it is feasible to pause or take a picture of a virtual machine that has reasonable information as cryptographic keys or password in memory and most companies do not look to these snapshots as critical assets as the running virtual machines. This paper presents a study about virtualization, focusing in security problems as the ones described on the previous paragraph. It covers both server and desktop environments and virtualization software. Regarding servers, we have conducted a vulnerability study, comparing two virtualization solutions, while for the desktops we have made a study about how virtualization can be used to improve security for online banking and e-commerce.

2. EXISTING SYSTEM AND RELATED WORK

2.1. Virtualization

Defining virtualization is not an easy task because as we will see later, there are different types of virtualization and a definition that would be adequate for all is not easy to achieve. Singh [13] describes virtualization as "framework or methodology of dividing the resources of a computer into multiple implementation environments, by concerning one or more notions or technologies such as software partitioning, hardware division, time-sharing, partial or complete machine simulation, emulation, worth of service, and many others". While, this description leaves out bags as virtualization in the network, application virtualization or storage virtualization. Kiyanclar [14] describes virtualization as "the consistent duplicate of a complete architecture in software that presents the illusion of a real machine to all software running above it". Most of the definitions are correct if we only consider server virtualization; nevertheless, I adapted Singh's definition saying that: Definition. Virtualization is as a framework dividing the resources of the device from the execution environment, allowing environment plurality by using one or more techniques such as time-sharing, emulation, partitioning.

2.2. CPU Virtualization

The x86 architecture is the majority used CPU design in enterprise data-centers today. The virtualization can take benefits of that. The Intel 80286 chip-set, introduced on February 1982, was the first of the x86 family to provide two main methods of addressing memory: real mode and protected mode. Later, in 1985, with the 80386 chipset, a third mode was introduced called virtual 8086 mode (also called virtual real mode, V86-mode or VM86). The VM86 allowed multiple real mode processes to be run simultaneously while taking full advantage of the 80386 protection mechanism. Real mode soon became obsolete because it had some disadvantages, such as it was limited to a one megabyte of memory and only one program can be run at a time. The same way, virtual mode was locked in at 16-bit and became obsolete with the high use of 32-bit operating system. Protected mode, by the other hand, is the natural 32-bit environment of the 80386 processor providing many features in order to support multitasking, such as hardware support for virtual memory and segmenting processor. Protected mode in the x86 family uses 4 privilege levels, numbered from 0 to 3. Sometimes these levels are designated as rings, and the term comes from the MULTICS system [17], in which privilege levels were illustrated as a set of concentric rings. We are about to employ the word "ring" as level, because it is a vocabulary is old. System memory is divided into segments and each segment is assigned and dedicated to a particular ring. The processor uses the privilege ring to decide what actions can be done with the code or data within a segment. As it shows in the Figure 2, Ring 0 is considered the innermost ring, which has total control of the hardware while Ring 3 is the outermost ring and has restricted access.

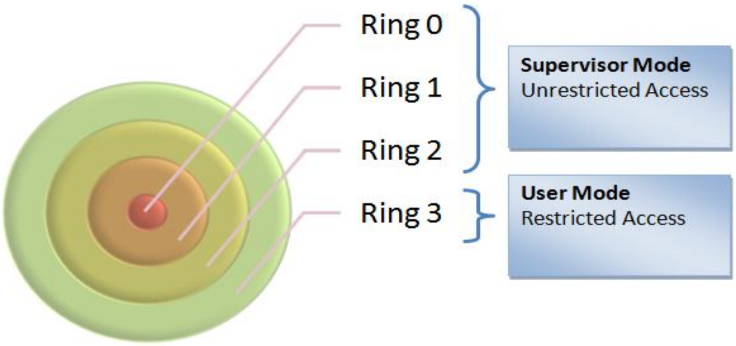

Figure 2. Privilege rings of the x86 architecture. High privilege:0; Low privilege: 3

The supervisor mode is the execution mode on an x86 processor with unrestricted access, which enables the executions of all instruction, including I/O and memory management operations, which are privileged instructions. Operating system runs on this supervisor mode, normally on the ring 0. However if this ring is compromised, it will have direct impacts on the ring 3 (user mode). The idea of having isolated ring 0 for each virtualized guest is that if one of the ring 0 of a virtualized guest is affected by, for instance, a failure it will not have impact the ring 0 of others virtualized guest. In order to do this, it is necessary to make this ring 0 closer to the guest, residing in either ring 1 or ring 2 for x86 architectures. However, the further it goes from the real ring 0, the more distant is from executing direct hardware operations, resulting in a loss of performance and independence. Virtualization moves ring 0 up one level in the privilege rings model and places the virtual machine monitor in the next higher privilege ring. This will be the ring 0 and it is upon this the guest operating systems runs, while the Virtual Machine Monitor controls the interaction with the underlying hardware platform. VMMs can be classified in two types:

a. Type 1: This type is also called as native or bare metal since the hypervisor software operates on top of the host's hardware on the real ring 0 (Figure 3(a)). A guest operating system thus runs on another level above the hypervisor, allowing for true isolation of each virtual machine. This is the classic VM architecture. An example of this implementation is the VMware ESX Server and Xen.

b. Type 2: This type is also called hosted VMM because the hypervisor software runs within a normal host operating system previously mounted, usually in ring 3 (Figure 3(b)). This type of VMM has a lower performance than the other type because factors as calls to the hardware must traverse many diverse layers before the operations are returned to the guest operating system. Examples of this implementation include VMware Workstation, Sun Virtual Box and Parallels Workstations.

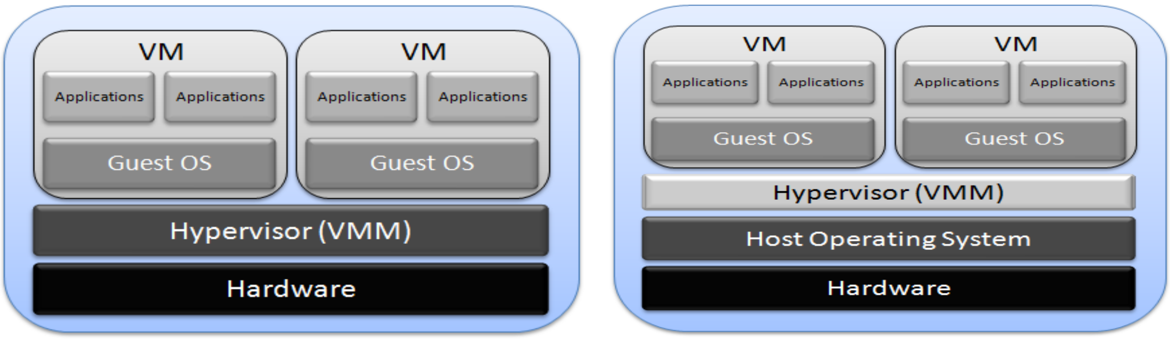

(a) (b)

Figure 3. (a) Type 1 Hypervisor/VMM, (b) Type 2 Hypervisor/VMM

2.3. Memory Virtualization

Standard operating system employs page-tables to converts addresses which are virtual and converted into addresses which physical data is provided. Virtual machines gives raise to new observations about memory virtualization since memory is going to be shared although isolation as to be guaranteed. We can consider three classes of addresses on a virtualized system:

a. Virtual addresses, which are the same as the ones used by a conventional OS • Guest physical address
b. Machine memory

Operating systems which are guests will maintain page tables that translate from virtual addresses to pseudo-physical addresses, and hypervisor preserves individual shadow page tables that translate from virtual addresses to device or system addresses [24].

The current x86 CPUs hold memory in hardware. Translation from virtual to physical addresses is performed by the memory management unit and the most used parts of the page tables are cached in the translation look aside buffer (TLB). Guest OS sees page tables, which run on an emulated MMU. These tables provide the guest OS with the false impression that it can interpret the virtual guest OS addresses into physical addresses of the machine, but it is the hypervisor that deals with it. The real page table is the shadow page table used to translate the virtual addresses of the guest OS into the real physical pages.

The classic implementation of hypervisors maintains a shadow page table, which allows to control what page of the machine's memory is available to a virtual machine. Just like in a traditional operating system's virtual memory subsystem, when the memory allocated to VMs exceed the host physical memory size, the hypervisor can page the VM to the disk. This way, the hypervisor can dynamically control how much memory each VM receives. Figure 4 shows ESX server memory mapping.

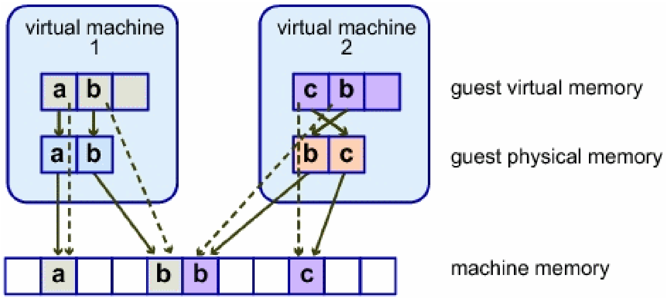

Figure 4. ESX Server Memory Mapping

2.4. Device and I/O Virtualization

The VMM virtualizes the physical hardware and allows each virtual machine a set of customizable virtual devices. Most of this virtualized I/O requires software drivers that run on the host operating system to access the real hardware. If it is a type 2 hypervisor, then it will use the device drivers already in the host OS, otherwise it may be necessary to develop its own device drivers for the hardware on the machine, like in the case of VMware ESX. Emulation is normally used for a VMM to handle I/O devices, and it is the VMM the responsible to implement a software model of the I/O device, making believe the guest OS that it is communicating to a hardware device, when is communicating with a software model. The I/O virtualization may provide to the guest, virtual hardware that does not exist in the real hardware, for instance, emulating an IDE hard disk when the real hardware is SATA. The direct memory access (DMA) has problems when used with virtual machines. The DMA controller can write to the entire physical memory instead of only the memory assigned to the guest OS. In order to deal with this problem, Intel and AMD added I/O Memory Management Unit (IOMMU). With IOMMU it is possible to restrict which physical address a device may access.

2.5. Types of Virtualization

When people talk about virtualization, normally they are talking about server virtualization.

However, information technology has other forms of virtualization commonly known and used by other groups of people. For some, virtualization means storage virtualization, or network virtualization or even application virtualization. Although my thesis will only concern about server virtualization, I will do a brief explanation of each one.

2.5.1. Server Virtualization

There are many different implementations of server virtualization on, and for a big range of CPU platforms and architectures. Informally, server virtualization can be seen as creating many virtual systems within a single physical system. To accomplish this, we can take three approaches: physical layer, virtualization layer and OS layer. Hardware partitioning divides a single physical server into partitions where each partition is able to run an operating system while hypervisor places a layer of software between the physical hardware and the multiple operating systems that will share the same physical hardware.

Physical Layer:

• Hardware partitioning: The server is physically segmented into distinct smaller systems that will act as a physically independent and self-contained server. Normally each of these smaller systems has their own CPUs, OS, boot area, memory and network resources. The implementation of this technique includes Static Hard Partitioning, vPar, nPar among others [27].

Virtualization layer: Hypervisor technology can be organized in some distinct categories:

• Full virtualization: Allows virtual infrastructures to run unmodified operating systems in separation. Present operating system (OS) executing inside the virtual machine is called guest operating system. This approach was pioneered in 1967 with IBM CP-40 and CP-67, predecessors of VM family. In order to implement full virtualization, it is necessary a full combination of hardware and software, however not all architectures have hardware to support virtualization. It was not achievable on IBM System/370 until 1972 and it was not natively possible in the x86 architecture [28] until 2005 when Intel and AMD added the hardware virtualization extensions (Intel VT and AMD-V respectively). Nevertheless, many companies tried to accomplish full virtualization on x86 architecture even before Intel VT and AMD-V additions. VMware uses a combination of direct implementation with binary translation techniques [29] to accomplish full virtualization of an x86 system. Figure 5 shows the binary translation approach to x86 virtualization used by VMware

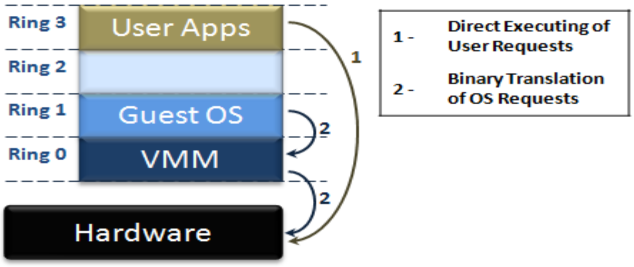
Ring 3
Ring 2
Ring 1
Ring 0
User Apps
Guest OS
VMM
Hardware
1
2
2
1 - Direct Executing of User Requests
2 - Binary Translation of OS Requests

Figure 5. The Binary Translation Approach to x86 Virtualization used by VMware

• Para virtualization: modifies the guest kernel system in order to purge the necessity of binary translation. It has the advantage of higher performance but has the drawback of needing a modified operating system kernel. The fact the virtual platform is not identical to the real hardware, it makes necessary for the operating system to be ported to the abstracted machine interface. This could be seen as a violation of the Goldberg's equivalence requirements, because the architecture-dependent part of the operating system kernel needs to be changed [30]. The non virtualizable instructions are replaced with hypercalls that communicate directly with the virtualization layer hypervisor. The architecture independent part and the entire user mode software stack stay unmodified.

• Emulation: Occasionally community confounds emulation with full virtualization. Even though both operate unmodified guest operating systems, they are both very different. In emulation, the virtual machine simulates the entire hardware set needed to run the unmodified guest OS normally for a completely different hardware architecture. There are some utilities for this technique.

Operating System layer:

• Operating System-Level Virtualization: This is a technology that virtualizes servers at the OS (kernel) layer. The physical server and instance of the OS is virtualized into multiple isolated partitions. Each of them will look like a real server, from the point of view of its owner. The OS kernel will run a single OS and provide its functionality to each of the partitions.

2.5.2. Storage Virtualization

Storage virtualization has been around for a number of years. It has beginning with the use of redundant array of independent disks (RAID). Using RAID it is possible to logically group physical disks and present those groupings as a virtual disk to the OS. Using storage virtualization it is possible to merge physical storage from many devices which will appear as a single storage pool. This storage can be classified as direct attached storage (DAS), network attached storage (NAS) and storage area network (SAN). They can be linked using Fibre Channel, Fibre Channel and Internet Small Computer Systems Interface (iSCSI), Fibre Channel on Ethernet or Network File System (NFS). Storage virtualization it is not a requirement for server virtualization but its use provides benefits since it can rely on the assignation of a logical unit (LUN) of storage, but provisioning it only when needed. For instance, if we have a LUN of 500 GB but we are only using 20GB, then only 20GB of actual storage is provisioned. This reduces the cost of storage, since we only use what is needed. Storage virtualization brings also help to the storage administrator, since it is easier to manage tasks as backup, archiving or recovery.

2.5.3. Network Virtualization

When people talk about network virtualization, probably the first thing that comes to their minds is Virtual Private Network (VPN) or perhaps Virtual Local Area Networks (VLAN).

The most used network virtualizations are:

a. Virtual LAN (VLAN): Defined in the IEEE 802.1Q standard, is a method of creating independent networks using a shared physical network. They are used to logically segment broadcast domains and control the interaction between different network segments. VLANS is a common feature in all modern Ethernet switches, allowing to create multiple virtual networks, which isolates each segment from the others. All the available resources are segments and allocated to each of these segments. Therefore, VLAN is a safe method of creating independent or isolate logical networks within a shared physical network.
b. Virtual IP (VIP): A VIP is an IP address that is not associated to a specific computer or network interface, but is normally assigned to a network device that is in-path of the network traffic. Incoming packets are sent to the VIP but are redirected to the actual network interface of the receiving host or hosts. It is used in solutions like High-Available and Load-Balancing, where multiple systems have a common application, and they are able to receiving the traffic as redirected by the network device.
c. Virtual Private Network (VPN): It is a private communication network that uses public network, such as Internet. Its purpose is to guarantee confidentiality on an unsecured network channel, from one site to another.

3. VMWARE

Founded in 1998 by Diane Greene and Dr. Mendel Rosenblum along with two students from Stanford University and a colleague from Berkley, VMware is a well known company on the x86 virtualization market. In October of the same year, these five founders filed for a patent regarding new virtualization techniques. These techniques were supported by a project called Sim-OS carried out at Stanford University. The U.S. Patent 6,397,242 was awarded on May 28, 2002 [42].

Their first product was VMware Workstation with the first version being released on February 8, 1999 for Windows and Linux. This is one of the mainly victorious products from VMware for desktop and stays as a commercial product with its current version 6.5.x. VMware Workstation is a type 2 hypervisor, supported on top of a Host OS – either Windows or Linux – and able to create virtual machines for a variety of guests OS, such as Solaris x86, Netware, FreeBSD, Windows and Linux. In late 2000, they released their first version of the server virtualization platform called VMware GSX Server. In 2006, VMware GSX Server was renamed to VMware Server and it is now released as freeware. In 2001, VMware release their first version of Elastic Sky X (ESX) [43].

3.1. ESX Platform

The core of VMware ESX has three main modules capable of regulating CPU affinity, memory allocation and oversubscription, network bandwidth throttling and I/O bandwidth control. Along with these, Virtual Machine File System completes the VMware ESX base platform. The three primary components of VMware ESX are: Physical Host Server: This is related with the physical host server where VMware ESX runs on.

VMkernel: The VMkernel is the center of the VMware ESX hypervisor, and it is a high performance operating system developed by VMware to run on the ESX host server.

The Console Operating System (COS): The service relieve has been promoted from being based on a variant of Red Hat version 7.2 with ESX 2.x to Red Hat Enterprise Linux 3, Update 6 for ESX 3.0 [45] and Update 8 for ESX 3.5 [46].

Virtual Machine File System (VMFS): The VMFS is a high performance cluster file system created by VMware. VMFS has many advantages compared to conventional file system.

VirtualCenter: The VMware VirtualCenter is the management console used to control the virtualized enterprise environments.

3.2. VMware ESXi

VMware ESXi was announced during VMworld 2007 and it is an integrated version of VMware ESX but without the COS. This is important in terms of security. VMware ESXi had RHEL-based COS replaced with BusyBox, which is a single binary that provides a minimal set of services. Many of the security patches for VMware ESX where related with security vulnerabilities on the Service Console (e.g. CVE-2009-1185, CVE-2009- 0034, CVE-2009-0846).

3.3. Xen

The Xen project was first described in the paper “Xen and the Art of Virtualization” presented at SOSP in 2003 [49]. It was a project originated by the System Reseach Group at the University of Cambridge Computer Laboratory and was part of the XenoServers projects [50] which had the goal of build a public infrastructure for global-scale service deployment. XenSource, Inc. was a company founded by Ian Pratt, senior lecturer at Cambridge and lead of the Xen project, with the goal of supporting and developing the open source Xen project and to create a commercial enterprise version of the software. In 2005, XenSource release Xen 3.0.

Xen is an open-source hypervisor for both 32 and 64 bit process architecture that runs on top of the bare-metal. It allows to securely and efficiently run several virtual guest OS on the same host at the same time. Xen as many features as:

a. Near native performance on the virtual machines
b. Full support on x86 (32-bit) with and without Physical Address Extension (PAE)
c. Full support on x86 with 64-bit expansion
d. Support for almost all hardware with Linux drivers available
e. Live migration of running virtual machines between two physical hosts with zero downtime
f. Support of Hardware Virtualization extensions from Intel (Intel-VT) andAMD(AMDV), allowing unmodified guest operating systems.

3.4. KVM

KVM stands for Kernel-based Virtual Machine and is full virtualization tool for Linux on x86 machine, on the basis of virtualization in hardware extensions (Intel VT-X and AMDV) and an adopted version of QEMU, using the Linux Kernel as the Hypervisor [53]. KVM cannot operte on CPUs without the hardware virtualization extensions. KVM consists of two modules:

a. kvm.ko: A loadable kernel module that provides the core virtualization infrastructure
b. kvm-[intel|amd].ko: A processor specific module for Intel or ADM.

4. SECURITY OF VIRTUAL MACHINES

Virtual Machine technology is mainstream, implemented on every branch of the industry (e.g. Telecom, Finance), running critical services, which were previously implemented in isolated servers. One important issue is security. According to Gartner [60], "through 2009, 60 percent of production VMs will be less secure than their physical counterparts" and that "30 percent of deployments [will be associated] with a VM-related security incident". With virtualization, a new scenario is presented, with all the servers consolidated on a single server running three virtual machines. In this case, a new threat model (Figure 6(b) on page 46) allows not only the same attacks as the traditional model but also opens the possibility to do new attacks and explore new vulnerabilities. Since each VM can be “exactly the same” as the real one (in the sense that it can have, for instance, the same number of CPUs, amount of memory, patches installed and configuration), an attacker can explore a vulnerability in an application the same way as it was done on the traditional threat model (arrows #1). Besides that, she has new targets to check for vulnerabilities. If the virtualization layer have vulnerabilities, she could launch an attack from a guest OS against the others VMs in the same host (arrows #2), or could attack the host by doing a denial of service exploring the host vulnerability in a virtual device . Doing this attack, she could attack the other guests in the same host (arrows #3). The impact of compromise the virtualization layer raises the risk’s level, since it is a critical asset shared by all the virtual machines. However, the attack can be remote too (arrow #4).

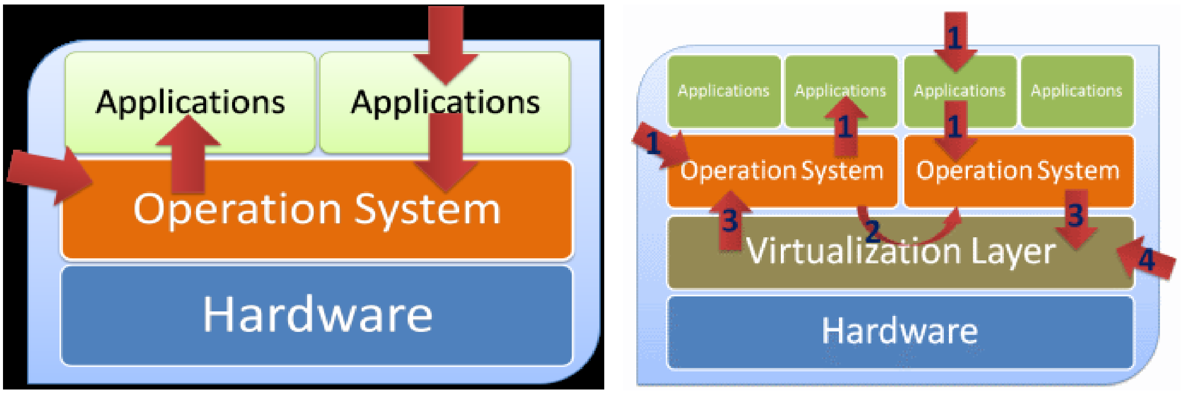

(a) (b)

Figure 6. (a)Traditional threat model, (b) New threat model

5. ANALYZE OF SECURITY VULNERABILITIES IN VIRTUALIZATION

5.1. Attackes from the Guest to the Host

There can be some external or internal factors that can compromise isolation as miss implementation configuration or some bug in the virtualization software. Any attacks can be made if isolation between the Virtual Machine and the host which may be dangerous. Any one may be compromised causes dangerous effects.That attack is called ”VM escape” and happens when a program can bypass the virtual machine layer from inside a VM and get access to the host machine. The host machine is the root of all the VMs, and so if a program escapes from the virtual machines privileges it will get root, allowing to control all the VMs and the traffic between them. This kind of attack are normally possible by exploiting bugs on the VMM combined with improperly configuration of the host/guest interaction. However, current VMMs do not offer perfect isolation, although they claim to. Many bugs have been found in all popular VMMs, some of them allowing ”VM escape”, like the VMware Workstation 6 CVE-2007-4496 [62] bug discovered by Rafal Wojtczuk which allows authenticated users with administrative privileges on a guest operating system to corrupt memory and possibly execute arbitrary code on the host operating system via unspecified vectors.

5.2. Remote Management Vulnerabilities

It is common in the current VM Environments to have a management consoles that manages the virtual machines. Normally, commercial products have their own. For instance, VMware uses VMware vSphere to manage the Hypervisor, while Citrix XenServer can use XenCenter. These consoles bring new facilities for administrators to manage their machines, but also open new vulnerabilities. Compromising a management console allows an attacker to control all the virtual machines managed by it. This kind of technology generally related with the VMM using HTTP/HTTPS whose meaning is the VMM is having service running accepting HTTP connection. Xen, for instance, has the XenAPI HTTP Interface that had a

Cross-site scripting (XSS) susceptibility, which permitted running a script code in a user's browser session in context of an affected site. HyperVM [66] is a multi-tiered, multi-server, multi-virtualization software which allows to create and manage different Virtual Machines (Xen or OpenVZ) each with each Virtual Private Server (VPS) having its own operating system.

5.3. Denial of Service

A Denial of Service (DoS) has the goal to make a computer resource not available to its intended users. In virtual machine architecture, resources as CPU, memory, storage and network are common among the host and the guests. It is then probable for a guest machine to force a denial of service (intentional or not) to others guest which would also influence the host by captivating all the potential resources of the system. When other guests attempt to request a resource, the system will deny that access since there is no resource available. VMware has been shown to suffer from several DoS vulnerabilities(4.2). A good approach to prevent this attack from a guest is to limit the resources VMs can access. Most of the current virtualization technologies have the mechanisms necessary to limit the resources allocated to each guest machine. With the correct configuration of the host virtualization, this attack can be minimized.

6. CONCLUSION

This paper is a result of study work about virtual machines, their principal characteristics and differences, and the security impact of their utilization. We initially studied the history of the virtual machines, following their evolution from the origins, in 1960s until nowadays with its implementation in the x86 architecture. We also presented the different components of virtualization, focusing on the problem of the x86 architecture that natively did not support virtualization and how some virtualization software companies worked around this problem. We have made a study about the current state of art of the main server virtualization products, describing their principal characteristics and systems supported.

The use of virtualization brings many advantages, such as reduction of the hardware resources needed with direct impact on cost efficiency, but also security advantages. The latter benefit is commonly used to spread the word on virtualization, but we wanted to demystify this myth, presenting some of the security problems that server virtualization brought and their impact. We choose two of the main server virtualization products commonly used by companies and universities and conducted a vulnerability analysis, using as reference the CVEs reported. The conclusion from that analysis is that both products show a similar security risk, and they require an extra attention to have their security patches applied.

Desktop security always brought some concerns to IT administrators and also end users, which had their computers infected by virus, worms, or their credentials stolen by some key logger or phishing attack. We have described some solutions to this problem, using virtualization. The first solution presented is based on type 2 hypervisor. The second solution requires the user to reboot its machine and run a read-only bootable media. In this scenario, virtualization would be used to maintain a master version of the bootable media installed within a virtual machine, which allows to be updated, and this way create new versions of the read-only bootable media with the last updates applied. It was done a security analysis, describing how the solution can mitigate some of the security problems faced by a user. I believe that, in a near future, virtualization will start being regarded as a desktop security enabling technique rather than just a server workload consolidation mechanism, as it is by most of the IT community presently.

REFERENCES

[1] Joanna Rutkowska, "IsGame Over (), Anyone? Black Hat Conference", 2007. Retrieved July 26, 2009.

[2] Randy Perry Al Gillen, "Tim Grieser. Business Value of Virtualization: Realizing the benefits of Integrated Solutions", Technical report, IDC, July 2008.

[3] OpenDNS. Phishtank, June 2009. URL http://www.phishtank.com/stats/2009/05/. Retrieved November 10, 2009.

[4] Robert P. Goldberg, "Survey of Virtual Machine Research", Computer, pages 34–45,1974.

[5] Christopher Strachey, *"Time Sharing in large fast Computers"*, In International Conference on Information Processing, pages 336–341. UNESCO, June 1959.

[6] John McCarthy, "Reminiscences on the History of Time-Sharing", volume 14, pages 19–24, Piscataway, NJ, USA, 1992. IEEE Educational Activities Department.

[7] J. Howlett, "The Atlas Computer Laboratory", *Annals of the History of Computing*, IEEE, 21(1):17–23, Jan-Mar 1999. ISSN 1058-6180.

[8] Derrick Morris, Frank H. Sumner, Michael T. Wyld, *"An appraisal of the Atlas Supervisor"*, In Proceedings of the 1967 22nd national conference, pages 67–75, New York, NY, USA, 1967. ACM.

[9] Barbara S. Brawn, Frances G. Gustavson, Efrem S. Mankin, "Sorting in a Paging Environment", *Commun. ACM*, 13(8):483–494, 1970. ISSN 0001-0782.

[10] Peter J. Denning, "Performance Evaluation: Experimental Computer Science at its best", In SIGMETRICS '81: ACM Press. ISBN 0897910516.93

[11] Stuart E. Madnick, John J. Donovan, *"Application and Analysis of the Virtual Machine Approach to Information System Security and Isolation",* The Workshop Proceedings on Virtual Computer Systems, pages 210–224, New York, NY, USA, 1973. ACM.

[12] VMware, Inc., Vmware milestones, 2009. URL http://www.vmware.com/company/mediaresource/milestones.html. Retrieved November 22, 2009.

[13] Nadir Kiyanclar, "A Survey of Virtualization Techniques Focusing on Secure on-Demand Cluster Computing", ArXiv Computer Science e-prints, November 2005. Provided by the SAO/NASA Astrophysics Data System.

[14] Jim Smith, Ravi Nair, "Virtual Machines: Versatile Platforms for Systems and Processes (The Morgan Kaufmann Series in Computer Architecture and Design)", Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2005. ISBN 1558609105.

[15] Mendel Rosenblum, Tal Garfinkel, "Virtual Machine Monitors: Current Technology and Future Trends", *Computer*, 38(5):39–47, 2005. ISSN 0018-9162.

[16] Rich Uhlig, Gil Neiger, Dion Rodgers, Amy L. Santoni, Fernando C.M. Martins, Andrew V. Anderson, Steven M. Bennett, Alain K?gi, Felix H. Leung, Larry Smith, "Intel Virtualization Technology", *Computer*, 38(5):48–56, 2005. ISSN 0018-9162.

[17] AMD, AMD64 Architecture Programmer's Manual Volume 2: System Programming. Number 24593. September 2007. URL http://www.amd.com/us--en/assets/content_type/white_papers_and_tech_docs/24593.pdf. Retrieved July 23, 2009.

[18] Carl A, "Waldspurger. Memory Resource Management in VMware ESX Server", SIGOPS Oper. Syst. Rev., 36(SI):181–194, 2002. ISSN 0163-5980.

[19] VMware, Inc, Resource Management Guide Update 2 and later for ESX Server 3.5, ESX Server 3i version 3.5, VirtualCenter 2.5, VMware, Inc., 2009.

[20] G. Milos, D. G. Murray, S. Hand, M. Fetterman, "Satori: Enlightened Page Sharing", In Usenix, 2009.

[21] Oracle Corporation, Partitioning, Technical report, Oracle Corporation, 2002, URL

[22] http://www.oracle.com/corporate/pricing/ partitioning.pdf. Retrieved November 15, 2009.

[23] John Scott Robin, Cynthia E Irvine, *"Analysis of the Intel Pentium's Capability to bear a Secure Virtual Machine monitor"*, In SSYM'00: Proceedings of the 9th conference on USENIX Security Symposium, pages 10–10, Berkeley, CA, USA, 2000. USENIX Association.

[24] Keith Adams, Ole Agesen, *"The differentiation between Software and Hardware Techniques for x86 Virtualization"*, In ASPLOS-XII: Proceedings of the 12th international conference on Architectural support for programming languages and operating systems, pages 2–13, New York, NY, USA, 2006. ACM. ISBN 1-59593-451-0.

[25] Adam Lackorzynski Björn Döbel Alexander Böttcher Hermann Härtig, Michael Roitzsch, "L4 - virtualization and beyond. Korean Information Science Society Review", 2008. 95

[26] Jenni Susan Reuben, "A Survey on Virtual Machine Security", In Jukka Manner and Laura Takkinen, editors, Security of the End Hosts on the Internet, Seminar on Network Security Autumn 2007.

[27] Abraham Silberschatz, Peter Baer Galvin, Greg Gagne, "Operating System Concepts", John Wiley & Sons, Inc., New York, NY, USA, 2001. ISBN 0471417432.

[28] K. G. Anagnostakis, S. Sidiroglou, P. Akritidis, K. Xinidis, E. Markatos, A. D. Keromytis, *"Detecting Targeted Attacks using Shadow Honeypots"*, In SSYM'05: Proceedings of the 14th conference on USENIX Security Symposium, pages 9–9, Berkeley, CA, USA, 2005. USENIX Association.

[29] Xuxian Jiang, Dongyan Xu, *"Collapsar: a VM-based Architecture for Network Attack Detention Center"*, In SSYM'04: Proceedings of the 13th conference on USENIX Security Symposium, pages 2–2, Berkeley, CA, USA, 2004, USENIX Association.

[30] Tal Garfinkel, Mendel Rosenblum, *"A Virtual Machine Introspection based Architecture for Intrusion Detection"*, In In Proc. Network and Distributed Systems Security Symposium, pages 191–206, 2003.

[31] Lionel Litty, H. Andrés Lagar-Cavilla, David Lie, *"Hypervisor Support for Identifying Covertly Executing Binaries",* In SS'08: Proceedings of the 17th conference on Security symposium, pages 243–258, Berkeley, CA, USA, 2008. USENIX Association.

[32] Xuxian Jiang, Xinyuan Wang, Dongyan Xu, *"Stealthy Malware Detection through vmm-based "out-of-the-box" Semantic View Reconstruction"*, In CCS '07: Proceedings 96 of the 14th ACM conference on Computer and communications security, pages 128–138, New York, NY, USA, 2007. ACM. ISBN 978-1-59593-703-2. [41] Thomas J. Bittman. Virtualization with VMware or Hyper-V: What You Need To Know. Gartner Webinar, August 2009.

[33] Scott W. Devine, Edouard Bugnion, Mendel Rosenblum, "Virtualization System Including a Virtual Machine Monitor for a Computer with a Segmented Architecture", February 2002.

[34] Eric Siebert, "A Brief History of VMware", IT Knowledge Exchange, February 2009. URL http://itknowledgeexchange.techtarget.com/ virtualization-pro/a-brief-history-of-vmware-2/. Retrieved August 6, 2009.

[35] Edward Haletky, "VMware ESX Server in the Enterprise: Designing and Securing Virtualization Servers", Pearson Education, Inc., 1st ed. edition, 2008.

[36] VMware, Inc. Best Practices for VMware ESX Server 3, June 2006. URL www.vmware.com/pdf/esx3_best_practices.pdf. Retrieved September 01, 2009.

[37] VMware, Inc., “Details of What’s New and Improved in VMware Infrastructure 3 version 3.5”, URL www.vmware.com/support/vi3/doc/whatsnew_esx35vc25.html. Retrieved September 01, 2009.
[38] VMware, Inc., “Configuration Maximums - VMware Infrastructure 3”, January 2009. URL http://www.vmware.com/pdf/vi3_301_201_config_max.pdf. Retrieved September 09, 2009.

BIOGRAPHIES OF AUTHORS

O Sri Nagesh/osri.pvpsit@gmail.com, Ogirala Sri Nagesh is Research Scholar at Lingayya University and working as Assistant Professor in CSE Department at Sri Vasavi Engineering College, Tadepalligudem, Andhra Pradesh, India. He received Masters Degree in Computer Science Engineering from Jawaharlal Nehru Technological University Kakinada, His research interests include Cloud Computing, Security in cloud servers and Big Data. He published several papers in International conferences and journals.

Dr. Tapas Kumar / kumartapas534@gmail.com, Dr. TAPAS KUMAR, Working as a Professor, Dean & H.O.D in School of Computer Science & Engineering, Lingaya’s University, Faridabad. He holds a Doctorate in Computer Science & Engineering. He has more than experience of 15 years in Academics & Administration. He has published various Research papers in various National & International Journals of Reputed

Vedula Venkateswararao/ venkatvedula2012@gmail.com, Vedula Venkateswara Rao working as Associate Professor in Department of Computer Science Engineering at Sri Vasavi Engineering College, Tadepalligudem, Andhra Pradesh, India. He received Masters Degree in Computer Science Engineering from Jawaharlal Nehru Technological University Kakinada, Masters Degree in Information Technology from Punjabi University, Patiayala, India. His research interests include Cloud Computing and Distributed Systems, Data Mining, Big Data and Image Processing. He published several papers in International conferences and journals.