© 2016 IJSRSET | Volume 2 | Issue 1 | Print ISSN : 2395-1990 | Online ISSN : 2394-4099
Themed Section: Engineering and Technology

Design & Development of a Computational Model using Virtualization and Multi-tenancy Technologies for Cloud Computing Architecture

Prof. Dr. G. Manoj Someswar*, Hemalatha Kalaskar

Department of Computer Science, Mahatma Gandhi Kashi Vidyapith, Varanasi, Uttar Pradesh, India

ABSTRACT

Cloud computing has arrived as a solution to reduce costs in organizations and at the same time offer on-demand resources and computation without requiring to create an IT infrastructure. Services, such as Amazon Web Services (AWS) or Microsoft Azure provide a means for organizations to instantly provision and de-provision virtual machines (VM) depending on their needs, just paying for what they use. In order to make the necessary environment, cloud service providers (CSP) make use of virtualization technologies to maximize the value of their systems. Servers have always needed to run alone in physical machines to avoid other services to interfere with them; but the downside of this was the waste of resources. Virtualization enables the use of all the resources in a physical host by sharing them between the guest operating systems (OS). Many organizations have already deployed private clouds on their own infrastructures or through third parties. However, Public Clouds provide an additional advantage that makes it extremely attractive, cost savings. The resources for a cloud consumer seem to be unlimited by sharing all the host machines between different organizations. At the same time, the CSPs can easily maximize the use of each physical machine. Multi-tenancy is the name that receives this computational model. However, there is a drawback on multi-tenancy and public clouds. Host systems are shared between multiple tenants with different owners and one of them could potentially be a malicious attacker or even a competitor. Now someone trying to compromise an organization's business processes or data will not need to break through their traditional lines of defense. The traditional perimeter in their networks no longer exists. Now an organization's systems coexist shoulder to shoulder with unknown tenants with potential malicious intentions. The virtualization layer adds a new attack surface to be compromised where the hypervisor and the resident VMs can be the target. The alarms have been triggered, stopping many organizations on their path to the Cloud. This research paper aims to provide an overview of the security issues that this new computational model arises. The problem will be aboard from the general cloud computing term, through multi-tenancy, down to virtualization. The main goal is to explore and analyze the different threats that virtualization and multi-tenancy combined bring to the Cloud. More specifically, the venues to compromise a VM or a hypervisor in a physical machine will be analyzed and recommendations will be given on how to mitigate the risks.

Keywords: Virtual machines (VM), Cloud Security Alliance (CSA), Infrastructure-as-a-Service (IaaS), Proofs-of-Concepts (PoC), Distributed Management Task Force (DMTF), European Network and Information Security Agency (ENISA)

I. INTRODUCTION

Cloud computing is here. With its new way to deliver services while reducing ownership, improving responsiveness and agility, and especially by allowing the decision makers to focus their attention on the business rather than their IT infrastructure, there is no organization that has not though about moving to the Cloud. Several surveys from Gartner have shown the importance of cloud computing ranking it as the top

priority for CIOs in 2010, also demonstrating the effort they are doing to adopt it by increasing their expenditures on cloud computing services. There are several benefits, especially on infrastructure costs, but, as with all new technologies, cloud computing also has some drawbacks. The move to the Cloud is a crucial step for any company, but has to be made with a lot of caution because it could turn against users. Organizations need to clearly understand the benefits and challenges, especially for the most critical applications. There are several concerns but, as shown in an IDC survey about the issues of the Cloud, security is the main concern. The question is why security is such a complicated challenge in the decision of moving to the Cloud. The answer is easy: lack of control over their data[1]. When an organization decides to move to the Cloud the data is no longer on their hands. Even if they just use the Cloud for processing and not storage, they are taking the data outside their private perimeter. IT infrastructures have been de-perimeterized, so security needs to be approached from another perspective; but this blurring of the perimeter is not the only issue. Multi-tenancy allows multiple tenants to coexist in the same physical machine sharing its resources (CPU, memory, network...) and, at the same time, creates an isolated environment for each one. Cloud service providers (CSP) can maximize their infrastructures using this architecture by allocating resources from physical machines that are not being full used. Virtualization is the means used to obtain multi-tenancy. Virtualization has been in the IT world for a long time, but it is now when is getting the most attention. Virtualization allows multiple operating systems (OS) to run on the same physical device at the same time [2]. This allows several users to execute their applications on the same physical environment, but isolated from each other.

Multi-tenancy and virtualization enable an efficient computing model. However, the risks associated are something that IT experts have not ignored. Now several tenants coexist on the same physical environment, and the same question comes to all cloud consumers: who is my neighbour? Is it trustworthy? Is it my competency? The trust that traditional IT infrastructure always provided has been broken. There is no perimeter anymore, no firewalls and IDS/IPS at the Internet gateway stopping dishonest people from attack systems. Virtualization has created a new attack surface, the virtualization layer; multi-tenancy and public clouds have made that surface an easy target for attackers.

Tenants potentially coexist with malicious virtual machines looking for ways to compromise sensitive data or processes. Organizations need to be aware of this and understand the attack vectors coming from malicious virtual machines (VM).[3]

A. Problem Statement

Security is the key for the Cloud success. As many surveys show security in the cloud is now the main challenge of cloud computing. Until a few years ago all the business processes of organizations were on their private infrastructure and, though it was possible to outsource services, it was usually non-critical data/applications on private infrastructures. Now with cloud computing, the story has changed. The traditional network perimeter is broken, and organizations feel they have lost control over their data. New attack vectors have appeared, and the benefit of being accessible from anywhere becomes a big threat. Many of the cloud computing issues are similar to the old ones but in a new setting. This requires re-assessing the risks related to each of the critical areas considering the new hazardous environment.[4] The Cloud Security Alliance (CSA) defines 12 areas of concern for cloud computing divided into two broad categories: governance and operations. All of these areas are critical and should be taking in consideration when evaluating the security of a cloud environment. Amongst resilience and agility, the low costs that provide cloud computing is a real hook for companies trying to reduce costs. Start-ups looking for a place in the market pray for an economic solution that allows them to focus on their business without worrying on maintain an IT infrastructure.

With multi-tenancy resources are shared by multiple users. For example, two or more tenants could have their OSs running on the same server or two or running an instance of the same application with different data. Depending on the cloud deployment model the level of importance and sharing of multi-tenancy would be different but without any doubt Infrastructure-as-a-Service (IaaS) in public clouds creates the most risks off all[5].

It is necessary to adopt virtualization technologies to

allow the use of multi-tenant environments. Both give place to a new set of challenges when put together. Virtualization adds a new layer that can be targeted, and multi-tenancy facilitates the process to reach the layer. Virtualization security issues need to be reviewed from a new point of view not seebefore, coexisting with possible malicious tenants.

Several proofs-of-concepts (PoC) have been demonstrated. Examples like [KOR09] where a VM escapes from isolation and compromise its host, or where a rootkit subjugates a hypervisor, are enough evidence of the importance of security.

Hyde defines three classes of attacks on VMs:
A malicious VM compromises another VM on the same physical machine A malicious VM compromises the hypervisor on the same physical machine A malicious VM performs a denial of service (DoS) on the resources of the physical machine.

B. Research Aims

This research work aims to provide an understanding of the different attack vectors created by multi-tenancy and virtualization in a public IaaS cloud. The vectors will be explored, focusing on the threats arisen from different tenants coexisting in the same physical host.[6]

A critical analysis of the different vectors will be provided along with guidance on how to approach them. This analysis will be performed using previous works from different entities and authors, along with personal knowledge obtained from experience.

As part of the aim of this research, a strong foundation will be provided on the terms of cloud computing, multi-tenancy and virtualization. All these areas will be explored giving a strong definition. The different security issues will be also explored in order to provide an introduction to the main focus of the research.

C. Objectives of Research

General Objectives

Identify research and analyze the threats exposed by coexisting with unknown tenants on a public IaaS cloud environment.

Specific Objectives

The following are the specific objectives of the thesis :
Understand the concepts of cloud computing, multi-tenancy, virtualization, and their security issues. Identity the unique threats that multi-tenancy and virtualization create on public IaaS clouds. Identify and research possible attacks performed from malicious VMs to other VMs on the same physical host. Provide guidance and recommendations on how to mitigate the risks identified. analyze the issues identified using previous work and personal understanding of the problems.

D. Research Strategy

The strategy used for this research work was to identify and analyze the potential attack vectors and vulnerabilities that virtualization technologies create on cloud computing. Firstly, the different architectures and technologies will be explored analyzing the security issues arisen from each
of them, including:

1) Cloud computing.
2) Multi-tenancy.
3) Virtualization.

Virtualization creates some problems when deployed on traditional IT infrastructures, but some other security issues are unique when used in cloud computing. More precisely, the multi-tenancy architecture is what increases the risks from virtualization.[7]

Once identified the different security problems from all the areas, the threats and vulnerabilities to compromise the cloud will be explored. Most of these issues have already been largely covered; therefore, this thesis will focus on those coming from potential malicious VMs. With these main goals already defined, the next step will be to analyze the potential attacks, exploring the different approaches and stages of each one. More precisely the following will be studied:

1. Threats coming from a malicious VM to a VM.
2. Threats coming from a malicious VM to a hypervisor.

Finally, having analyzed these threats, recommendations and guidelines on how to mitigate

them will be provided. Additionally to the analysis, different areas for future research will be explored.

E. Research Limitations

Cloud computing and virtualization are large topics that had required establishing some limitations for the research of this thesis:

The cloud scenario for this thesis will be a public IaaS cloud. Other service models have also security issues related with multi-tenancy; however, in an IaaS environment the tenants have the most responsibilities and control over their systems, thereby the risks are higher. The budget for the research was limited so it was not responsible to perform some practical scenarios in a cloud environment.

There exist a large set of security issues arisen from the use of virtualization in the cloud which are out of the scope. In this research paper, the risks created by a malicious VM to compromise co-resident VMs and the hypervisor will be explored.

An in depth analysis of the different methods to exploit each of the three attack vectors specified is out of the scope since it would require a longer document than allowed.[8] The aim is to provide an understanding of each attack surface with the different steps and possible exploits.

F. Literature Review

Several books and entities have covered for the last years the concept of cloud computing. It is a hot topic nowadays in the technology and business world; thus there are multiple definitions. The National Institute of Standards and Technology (NIST) provide a well-recognized description for cloud computing including its characteristics, service models and deployments models. The NIST also contributes with another document with security guidelines for cloud computing. Cloud Security Alliance (CSA) has gained a lot of renown in the past few years for its significant contribution to the Cloud. The document for security guidance provides an excellent overview of cloud computing, supporting NIST's definition. The document also defines 12 areas of security concerns with an overview and recommendations for each one. Additional, CSA also has contributed with other documents like the top threats to cloud computing. [9]

Other relevant entities which contribute with definitions, guidelines and recommendations about cloud computing, are the European Network and Information Security Agency (ENISA), Gartner or the Jericho Forum. In these documents, different aspects of cloud computing are explored, along with the benefits, risks, and recommendations for securing the Cloud. Velte et al., Reese and Krutz et al., provide a general overview of cloud computing. These books cover the definitions, benefits, security, compliance, services, and all the other areas and aspects related to cloud computing. Other additional sources were used to complement the concept of cloud computing and its security.

Multi-tenancy is not an area considerably covered in any book or article. Thus, the bibliography used was extracted from the sources used for cloud computing and virtualization.

The document presented by Almond et al. was used as the main source. The CSA and ENISA provide material related to the security concerns of the multitenant architecture used in the Cloud, especially in public clouds.

Virtualization technologies have long been in the IT world. The material is large and well established. Marshal et al. explore the essentials of virtualization and provide an in- deep view of the VMware ESX server. Haletky provides a clear definition of virtualization and explores the necessary steps to deploy a secure virtualization using VMware ESX server.

Velte et al. and Reese include an overview of virtualization technology and its security issues focused on virtualization in cloud environments. The PCI Security Standards Council contributed with a document providing virtualization guidelines to comply with PCI DSS standard [11]

In support with many of the most established virtualization organizations, the Distributed Management Task Force (DMTF) provided an open virtualization format for packaging and distributing virtual appliances, with the intent to achieve a standard. A document released by NIST provides an analysis of

virtualization and an overview of the security concerns with guidelines to secure full virtualization technologies. Several conference proceedings focus on the area of virtualization, and more specifically in the security of virtualization. Ristenpart et al. explore the different steps to perform an attack on a VM in Amazon Elastic Compute Cloud (EC2). The research analyzes how to locate a target and achieve co-residence, and some possible compromises. Jarabek provides an analysis of several documents exploring topics related to virtualization and more specifically on side-channel attacks. Hyde gives an overview of the security of VMs including the potential threats.

The Center for Internet Security (CIS) explores the problem areas and vulnerabilities of virtual environments, including the threats arisen from malicious VMs to co-resident VMs.[12] Additional articles, conference proceedings, and electronic resources were used to complement the previous literature.

As with VM-to-VM attacks, many articles and conference proceedings cover the aspect of the hypervisor, its security, and the potential threats. The Burton Group provide an overview of the typical threats to a hypervisor from a malicious VM. Ormandy explores the security exposures of host machines in virtualized environments. Kortchinsky researches a PoC of a VM escape in IBM's Cloudburst. Ferri and Shelton provide other ways to perform VM escapes. Kato research discovered vulnerability on VMware that allows the use of a backdoor to perform a VM escape. King and Rutkowska develop some VM-based rootkits as PoCs to subjugate the hypervisor and the host machine.

Comparison against previous work During the research, several sources were founded that covered some aspect of the areas explored in this thesis. However, none of them covered the full spectrum.

G. VM – TO – VM

Ristenpart et al. research is one of the most complete about VM to VM attacks. They explore how to perform an attack using the Amazon EC2 as a case study. As a PoC, they show how to locate successfully a target VM on EC2 cloud using different network techniques. In-deep information is provided with a detailed explanation of the process. After locating the target, two methods are provided to achieve co-residence with the target on the same host machine. Though they provide some ways to determine if the VM is located successfully with the target, other methods explored in different researches can be used as well.

In the last part, they explore side-channel attacks in order to leak information and compromise the target. The following methods are proposed as examples:
1) Measuring cache usage.
2) Load-based co-residence detection.
3) Estimating traffic rates.
4) Keystroke timing attack.

For each part, a few recommendations are provided as a mean to mitigate or solve the vulnerabilities. This research, though complete on the information related to locate as the best PoC so far, did not cover the threat of side channel attacks with enough depth. Other materials were used to supplement this part of the attack. In relation with VM-to-VM attacks, the main goal of this thesis was to define and explore each of the steps to perform the attack, taking particular consideration on the last step where the compromise of another VM is achieved.

H. VM – TO – HYPERVISOR

No research was founded that fully covered the different aspects of this type of attack. The Burton Group provided a broad overview of the possible venues to perform this attack; however, the information given was not highly detailed. Ormandy provides several PoCs on how to perform VM escapes on different virtualization platforms. The research provides strong evidences about the viability of this type of attacks with a detailed explanation of the processes used in each platform. SubVirt and Blue Pill are some examples of VM-based rootkits PoCs. As a complex topic, each of them was studied on concrete scenarios. However, they both provide in-depth information of the exploitation. This research paper aims to provide an understanding of the threats to the hypervisor from a malicious VM, but none of the previous work fully covered all the venues for this attack.

I. Virtualization

Background

Like multi-tenancy, virtualization technologies are not mentioned as an essential characteristic by NIST. However, when present, it helps to strengthen the characteristics that the Cloud provides. Virtualization has been in the IT world for a long time. It was IBM the first that introduced the idea in the early 1960's with the term 'Time Sharing'. Virtualization technologies are already established in traditional IT environments, being deployed in many infrastructures according to a Gartner research.

Virtualization concerns have been long analyzed and studied; but, with the arrival of the Cloud, it has taken the spotlight because the many benefits and challenges it brings. In traditional IT environments, the risk of an attack was low. Now tenants do not know with whom they are sharing the infrastructure so, even if virtualization isolates each tenant, VMs must stay alert against their neighbours.

Defining Virtualization

There are many definitions, with almost the same meaning, but it is necessary to put it in context to give a more exact one. Although there are several forms of virtualization, virtualized OSs are the most used according to CSA, therefore it will be the main focus of this research work.

Virtualization of operating systems, also called server virtualization, is defined as "a way of making a physical computer function as if it were two or more computers where each non-physical or virtualized computer is provided with the same basic architecture as that of a generic physical computer. Virtualization technology therefore allows the installation of an operating system on hardware that does not really exist." A software layer called abstraction recreates the hardware characteristics of the physical computer to make it function as more than one computer. An example of this layer of abstraction is the Windows Hardware Abstraction Layer (HAL) which provides a common way for all drivers and software to talk to the hardware in a standardized format.

With virtualization, resources can be divided or shared through multiple environments, where those environments may be aware of not of the others. These environments are known as virtual machines (VMs), and usually host an OS, which are usually referred as guest Oss.

When a VM needs to interact with the hardware, instructions are passed directly to the physical hardware in order to decrease the latency and to operate more efficiently. However, there are some instructions that require being inspected and analyzed before executed in order to assure compatibility with the hardware. According to Velte et al., there are two virtualization types that concern cloud computing:

Full Virtualization: In this type of virtualization, a complete installation of one machine is run on another.

Para Virtualization: This type of virtualization allows multiple modified OSs to run on a single hardware device at the same time by more efficiently using system resources. The main difference between them is that in full virtualization the entire system needs to be emulated (BIOS, drive...); but in para virtualization, the OSs has been modified to work more efficiently with the hypervisor. Para virtualization usually runs better because fewer elements need to be emulated. Also allows for better scaling since a guest instance requires less processor time, so it is possible to host more guest OSs. However, Velte et al. also point some trade-offs. The use of para virtualization reduces flexibility since OSs need to be properly modified to run, which means that probably new Oss will need some time before being available on this type of virtualization. Also, there is an increased security impact since the modified OSs have more control over the underlying hardware which can impact on the other virtualized systems and the host OS.

There are also two main types of virtualization architectures: Hosted Architecture: In this approach, the host OS has a virtualization platform (hypervisor) installed into which one or more VMs run.

Hypervisor Architecture: In this approach, the virtualization layer sits on top of the hardware exporting the virtual machine abstraction.

Figure VII represents both architectures and a typical computer architecture.

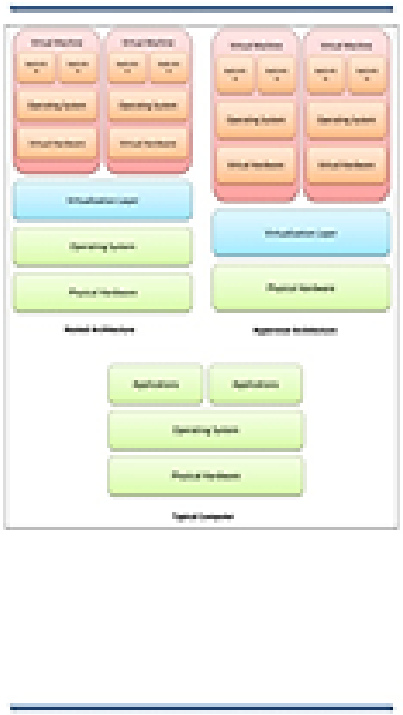

Figure 1: Virtualization architectures

According to the term workload is increasingly being used to describe the vast array of virtualized resources, like a VM.

The following are the components that form a virtualized environment. A virtual machine monitor (VMM) and one or more VMs that interact with either a hypervisor or a host OS in order to access hardware, local input/output, and networking resources. These components along with others will be defined on the following paragraphs.

Hypervisor

There is always a lot of confusion about what is the hypervisor because it is use indistinctly for both components, the VMM and the hypervisor. The reasoning is that they form the virtualization layer so in some way it makes sense treat them like one. The hypervisor is defined as a thin layer of software or firmware that provides access to hardware resources and provides virtual partitioning capabilities. It gives a single physical system, or host server, the ability to distribute resources to one or more VMs at a single time. The hypervisor is directly responsible for hosting and managing VMs running on the host. It provides a virtual hardware that is comprised of a configuration file, some virtual disk files, and various other files such as non-volatile RAM.

There exist two types of hypervisor:

Type 1 Hypervisor: A Type 1 hypervisor (also known as 'bare metal') runs directly on the hardware and is responsible for coordinating access to hardware resources, as well as hosting and managing VMs.
Type 2 Hypervisor: A Type 2 hypervisor (also known as 'hosted') runs as an application on an existing OS. This type of hypervisor emulates the physical resources required by each VM and is considered just another application, as far as the underlying OS is concerned.

The hypervisor may also include a VMM. The VMM is a software component that implements and manages VM hardware abstraction. The VMM manages the system's resources to allocate what each VM guest OS requires. From a security point of view, there is more risk when type 2 hypervisors are added because they add more complexity and vulnerabilities to the host. On the other hand, type 1 hypervisors are less complex than host OS, providing a smaller target and improved security.

The hypervisors use different techniques to provide VMs with a virtual environment. A virtual disk drive is achieved by using a disk image file that looks to the guest OS like a disk drive. If the hypervisor is type 2, the disk image appears in the host OS as a file. On type 1 hypervisors, all the components of the guest OS are stored into a single logical unit called image. Images are stored on hard drives and can be moved to other systems. The Open Virtualization Format (OVF) is a virtualization image metadata standard that provides interoperability between virtualization solutions.

Virtual Machine

A virtual machine (VM) is a virtualized representation of a physical machine operated and maintained by the virtualization software. Each VM is a self-contained operation environment that behaves as a separate computer, emulating the processor, memory, network adapter, removable drives and peripheral devices. In the same physical machine, several VMs with different OSs can be operated simultaneously; but, the only hardware present for each guest OS is the one presented by the

hypervisor. VMs provide some benefits over physical machines. VMs are usually compromised by a single or group of files that are read and executed by the virtualization platform. This means that they can be easily migrated from one system to another, copied, or backed up.

Virtual Appliance

A virtual appliance (VA) is described by as "a pre-packaged software image designed to run inside a virtual machine". Each VA is intended to deliver specific functions and typically consist of a basic OS and a single application. Examples of VAs are the virtualized forms of physical network devices such as routers or switches.
There is a special type of VAs called virtual security appliance (VSA). A VSA consists of a hardened OS and a single security application, and are usually assigned a higher level of trust to access the hypervisor and other resources like virtual networks running inside the hypervisor. This higher privilege allows the VSA to perform system and management functions. Examples of VSAs are firewalls, anti-virus, or IDS/IPS.

Virtualization Security

Virtualization security has been a hot topic for a long time. Before companies started to move to the Cloud, a lot of them already were implementing virtualization technologies on their IT infrastructures. Security concerns were already in the managements' minds, and a lot of effort was put in order to solve these problems.
With the coming of cloud computing, virtualization security is again on the mouth of security practitioners. As a recent study by Gartner indicates, in 2012 around 60% of the virtualized servers will be less secure than the physical servers they replace, hopefully dropping to 30% by 2015.

Virtualization is now deployed not in the private and perimeterized physical IT environments of organizations, but in the new environment defined by the Cloud were the perimeter is almost impossible to define, coexisting with unknown neighbours. As Haletky points that the security of a VM is dependent upon the OS in use; therefore, it should follow the security practices as if the VM was a physical host. From a security point of view, a VM and a physical server do not differ. There are two main ways to access a VM. One is through the hypervisor, and the other is through the network connections. So it is of utmost importance to secure both. Host servers, for example, should not be placed on an internet-facing connection unless necessary in order to minimize the risk. If the host server ends up being compromised, then all the VMs residing in the server will be at risk.

A compromised VM can be used to affect the host servers and other VMs in the same virtual or physical network. Attacks could be launched against these VMs or a DoS attack could be performed in the host server. In the case of Cloud environments, the risk increases since an attacker does not need to compromise a VM in order to attack other VMs or the network. The attacker just needs to pay for a cloud service and, as a consumer, start the attack avoiding the traditional security network devices.

Lindstrom provides an interesting approach listing five immutable laws of virtualization security:

Law 1: All existing OS-level attacks work in the exact same way.
Law 2: The hypervisor attack surface is additive to a system's risk profile.
Law 3: Separating functionality and/or content into VMs will reduce risk.
Law 4: Aggregating functions and resources onto a physical platform will increase risk.
Law 5: A system containing a 'trusted' VM on an 'untrusted' host has a higher risk level than a system containing a 'trusted' host with an 'untrusted' VM.

Lindstrom continues and explains that, in a broad sense, the vulnerability level of a system is a measure of the attack surface. An attack surface can be defined as the nature and extent of resources on a system that are exposed and, therefore, attackable. Virtualization increases the vulnerability by adding the attack surface of the hypervisor and the VMM.

There are several documents out there summarizing the general virtualization security concerns. In cloud computing, virtualization technologies still share the same security issues, but those are increased by the multi-tenant architecture and the erosion of the perimeter.

The traditional measures for security should be no longer an option. Organizations do not realize that using their existing physical server security in virtual environments actually limits their ability to maximize their use of virtualization and cloud technologies. This also leaves organizations exposed in ways not thought, causing significant security gaps.

CSA understands the importance with a specific security domain to deal with the problems created by virtualization in multi-tenant environments security guidance. In this guidance, CSA is primarily concern about the impact that virtualization has on network security. Because VMs can now communicate through the hypervisor instead of through the physical network, the traditional network security controls become useless; and express the necessity of these controls to take a new form in the virtual environment. Another important aspect of the security is the sharing of resources between VMs with different sensitivities, security, and owners. Unless a new security architecture is developed that does not require any network dependency for protection, this risk will always be present.

CSA goes further with the importance of virtualization in the Cloud in their document "Top Threats to Cloud Computing" where the 'shared technologies issues' is mentioned as one of the most significant threats. Trends provide a list of security challenges of virtualization in the Cloud that summarize almost all the problems:
Inter-VM Attacks: The new communication channel created between VMs cannot be monitored using traditional network security controls. instant-on gaps: Provide up-to-date security to dormant VMs becomes a difficult task.
A compromised image of a VM could potentially create a security breach when instanced.
Mixed Trust level VMs: Several VMs with different security levels could potentially be placed on the same host machine. This is especially concerning when coexisting with unknown tenants.
Resource contention: Accidental or unauthorized use of shared resources can potentially lead to a denial of service. Complexity of management: Management of the VMs becomes harder than before, requiring more complex patching and configuration policies.
Multi tenancy: VMs now coexist with other unknown and potentially malicious VMs.
Lack of audit trail: The process of monitoring and log VMs activities becomes more difficult on virtualization environments Other concerns like laws, regulations, and standards that govern the IT infrastructure were not designed with virtualization in mind. Furthermore, some of them even predate the acceptance of virtualization technologies, like a standard that requires certain data to be stored on a different server than other system logic. In a virtualization environment, both systems could exist on the same host, and be isolated from each other, so, is it or not the environment compliant with the standard? Several issues arise from virtualization in cloud environments, but this can actually become an advantage for organizations. The absence of a security perimeter and the highly volatile nature of VMs will force organizations to adopt robust security processes which can result in a high-security computing infrastructure according to Reese. This thesis will focus on the threats exposed by a malicious tenant coexisting in the same host system with other tenants in a public IaaS Cloud. More precisely the following threats will be analyzed:
1)Virtual machine to virtual machine attacks (VM-to-VM).
2)Virtual machine to hypervisor attacks (VM-to-Hypervisor).

J. Analysis & Interpretation of Threats

Cloud computing with its scalable, agile, and on-demand services provides many benefits to organizations; but at the same time introduces a new range of risks to information security.

Most of them are created from the new trust relationship between the CSP and the consumers. However, there are other risks related to the infrastructures that could go unnoticed. In order to maximize their infrastructures, CSPs enable multi-tenancy thus allowing different tenants to coexist on the same physical host. VMs will share the resources (memory, CPU, network…) of the physical machine and, at the same time, each VM will be separated from the others creating a false state of 'isolation'.

Tenants with different security contexts and unknown owners will be sharing the same environment believing they are alone. Now a consumer's VM could be sharing resources with their adversaries or with malicious VM

that will wait for the opportunity to penetrate this 'isolation' barrier in order to glean sensitive information and violate customer confidentiality.
Multi-tenancy introduces many risks in all the cloud service models, but especially IaaS clouds where the consumers have a lot of control. However, it needs to be considered that some CSPs are actually hosted in IaaS clouds. For example, some SaaS providers like twitter make use of services from IaaS providers like Amazon Web Services (AWS). Thus, the risk also extends to the users of those SaaS. Broadly speaking there are three types of attacks on VMs :

VM-to-VM: An attacker may use a VM to try to communicate and compromise other VM on the same physical host; therefore breaking the isolation characteristic of VMs.

Denial of service (DoS): An attacker will try to exhaust the resources from a physical host in order to deny service of the other VMs in the machine. As the source of the attack is a VM and the target are the co-resident VMs, DoS will be considered as a VM-to-VM attack.
VM-to-Hypervisor: An attacker tries to escape from the isolation created by the hypervisor in order to compromise it, which can potentially give access to the host OS and hardware. In this section, we will look at the first two attacks. An overview will be provided with a critical analysis of the issue, consequences and solutions.

K. The Problem

Traditionally, hosts were secured by placing security controls in the network monitoring the communications. After years of experience, these controls were secure enough to provide a reasonable sense of security. Now cloud computing and virtualization have broken that confidence.

The long-established network security controls are no longer effective since communications between VMs (on the same physical host) are through virtual networks provided by hypervisors; and the hypervisor, in order to reduce its complexity, does not have the proper capability to monitor and analyze the communications. The traffic between VMs is now 'off the radar', which gives opportunity to several threats.

There is another cause for these new threats. The CSP creates some level of 'trust' between VMs allowing VMs from different consumers to coexist without considering the risks that could potentially create between them in case a malicious VM. In a VM-to-VM attack, also referred as guest-to-guest attack, an attacker makes use of a VM to try to access, control, or gain information of other VMs on the same host machine. To do this, the malicious VM makes use of the shared resources namely shared memory, network connections etc. to compromise the other VMs, without compromising the hypervisor layer. For example, an attacker could try to interfere in a VM's operation by determining where the victim's allocated memory lies and later writing over it. Figure VIII shows an attack from a malicious VM to other VMs hosted on the same hypervisor.

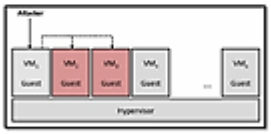

Figure 2: VM to VM attack

Provide full isolation is not always possible, and sometimes it is desirable to allow VMs to communicate with each other. There are several applications that require this such as dedicated monitoring using VSAs, or implementing a network technology that requires multiple peers. There are few documented VM-to-VM attacks on IaaS public clouds, but an attack is generally placed in three main stages:

1. Locate and place the attacker VM on the same physical machine of the target VM.
2. Obtain information about the target VM.
3. Exploit a shared resource to compromise the target VM.

An attacker aims to initialize a VM on the same physical host of the target VM. When placed on the same host is

necessary to check that the malicious VM coexists with the target VM. After that, information is gathered about the target. Finally, the attacker exploits a resource shared between the co-resident VMs in order to compromise the target.

L. Reaching the Target

The first step to compromise a VM is locating the target and instance a malicious VM on the same host machine. In an IaaS cloud, it is hard to conceal the network topology. This, added to the transparency of the network topology, allows an attacker to discover the location of the target VM on the cloud, instance a VM on the same host machine, and determine co-residence with the target. Therefore, the process of locating the target can be divided into two steps:

1) Mapping the Cloud.
2) Achieving co-residence.
3) Mapping the Cloud (Cloud Cartography)

In a cloud infrastructure, network topology is supposed to be kept private. CSPs try to achieve this by just allowing users to specify the rough geographic region (e.g. US, Europe) where create their VMs. However, a research performed by UCSD and MIT demonstrated a way to map the cloud and further compromise a VM on Amazon's EC2. So far this work has been the most accurate on this area; therefore, will be used as a model. Although the attack was performed on EC2, the researchers claim that this vulnerability would probably arise too on other IaaS clouds since functionalities provided are similar. The process to map a cloud is composed of two steps.

M. Enumerating Public Services

In the case of EC2, VMs can only be co-resident if they have the same creation parameters (e.g. region, instance type). The research also identified that a strong correlation exists between the cloud internal IP addresses assigned to VM and their creation parameters. To achieve this conclusion they enumerated the public EC2-based web servers using external probes (nmap, hping, and wget).

N. Mapping external to internal IPs

The next step is to translate responsive public IPs to internal IPs. CSPs usually provide a DNS service to map public IP addresses to private IP addresses which means any user can run a DNS query, even a malicious attacker. The information acquired so far can then be used to determine if an application is hosted on EC2 and the creation parameters of its VM. This drastically reduces the number of instances needed before a co-resident placement is achieved .

O. Determining Co-Residence

Before placing a malicious VM with the target VM, it is necessary to have a process to determine the correct co-residence. This process widely varies depending of the hypervisor implemented and the network policy on the IaaS cloud. The Xen hypervisor runs a privileged VM (called dom0) along with the other VMs on the host machine. In virtualized environments, all the network traffic to the physical network goes through the hypervisor, and the Xen hypervisor run as a VM which means it has an IP address. This means that the next hop of a VM's network traffic is the hypervisor's IP address, which can be used to identify the server.

A trace route analysis between the malicious VM and the target VM returns the IP of the routers between them, which can be used to confirm co-residence. The process could be even simpler. If both VMs are on the same host machine, just the IP of the hypervisor should be returned from the trace route; therefore, if just one IP is returned it means that co-residence is achieved.

Other network-based method to determine co-residence were identified on the research, including packet round-trip times, and numerically close internal IP addresses. However, as the method already described, all are implemented with some restrictions (Xen hypervisor, EC2 cloud structure). The goal of this section is to describe the necessity of this step as part of an attack, considering that depending on the environment conditions different techniques will be required.

There are other approaches to determine co-residence. Home Alone is a system that verifies if a VM has exclusive use of a host machine. Although the purpose of this tool is to assure exclusivity for a VM, it can be used as a method to confirm co-residence with other VMs. Instead of a network-based approach, Home Alone makes use of the vulnerabilities of side- channels to analyze possible co-residence.

P. Achieving Co-Residence

After having determined the creation parameters for initialize the malicious VM, and have a reliable test to determine co-residence, the next step is to create a VM on the same physical host as the target. After achieving co-residence, the malicious VM would be able to start compromising the target VM.

Two approaches are tested to achieve co-residence. Brute forcing placement This approach is the simplest possible: launch VM instances over a relative long period of time, with the appropriate creation parameters, checking co-residence with the target. If it is not co-resident, then the VM is terminated and another one is launched. This process is repeated until co-residence is achieved.

A moderate success rate was obtained for this method on. A target group set of 1,686 servers were selected and 1,785 probes instances were launched over 18 days, each checking co-residence against all the targets. The prove VMs achieved at least 8.4% coverage of the target set.

Q. Instance Flooding

The instance flooding approach aims to "take advantage of the parallel placement locality exhibited by the EC2 placement algorithms" [RIS09]. These algorithms tend to initialize newly VMs on the same server. So, when a target VM is launched, the attacker launches simultaneously as much VMs as possible.

One of the characteristics of cloud computing is to run servers only when needed. When the VMs are not needed they are simply stopped, and later resumed when needed. Therefore, an attacker can monitor a server's state until it stops, and as soon as it is resumed, engage in instance flooding.

This approach is better for individual or small sets of targets. However, it exploits a feature of EC2 (auto-scaling and EC2 placement algorithms). Depending on the IaaS environment, this type of approach will need to be executed using different methods, though auto-scaling will probably be a feature of any IaaS cloud.

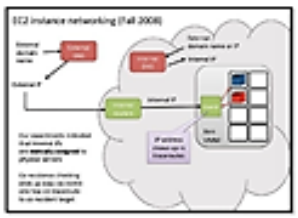

Figure 3 : EC2 instance networking Obtaining Target's Information

Once a malicious VM is placed on the target's host machine, an attacker can use side-channels to leak information about the target. When enough information is gathered the attacker will be able to compromise the VM.

The main goal of side-channel attacks, also known as cross-VM when between VMs, is to extract confidential information from neighbor VMs. Because of resource allocation (e.g. CPU's cache data, network...) between co-resident VMs, it is possible to leak information about co-resident VM's activity levels using several techniques.

Side-channel attacks have long been studied, especially in multi-process environments where have been used to extract cryptographic keys. These attacks demonstrate that even though it may not seem particularly useful to monitor the resource usage, it actually is if obtained by a clever attacker.

The side-channel attacks used to extract cryptographic keys mainly compromise the data cache memory. However, any resource shareable between different tenants can be used as a side-channel: CPU branch predictors and instruction cache, CPU pipelines, DRAM memory bus etc.

Some complications arise when trying to run cross-VM attacks on the cloud. Several factors like core migration, coarser scheduling algorithms, or double indirection of memory addresses, can make the attacks more complex. In the case of EC2, other factors like unknown load from other instances and the CPU configuration (no hyper threading), affect too. In order to overcome these issues, it is better opt for more coarse-grained attacks. Less information is obtained, but the implementation is much easier and robust in noisy environments like the Cloud.
Our research work explores cross-VM methods analyzed in as possible means to leak information about the co-resident VMs on an IaaS cloud.

R. Measure Cache Usage

The goal of measuring cache usage is to see how busy a server is. By measuring the CPU cache utilization on the physical machine, a malicious VM can try to estimate the current load. If the attacker VM is co-resident with the target VM, then a high load will indicate activity on the target. Despite the several methods available, most of them exploit the timing difference between the cache and the main memory. A well-known method is PRIME+PROBE. In a nutshell, the malicious VM fills an entire cache set by reading a memory region M from its own memory space. Then the attacker waits a specified interval while the cache is utilized by the target VM. Finally the malicious VM times the reading of the same memory region M to learn the target's cache activity on the cache set. If the target has a high cache activity, then most of the attacker's data will be erased from the cache. This will result in a higher timing measurement when the attacker reads again the memory after the specified interval.

S. Estimating Traffic Rates

This method could sound harmless; but, if properly analyzed, it can be used to deduce activity patterns, peak trading times for maximize the effect of DoS attacks, and other uses. Additionally, information about the number of visitors or most frequently visited pages can be estimated. This information might not be public, so a competitor could take advantage of it. We estimate the web-traffic of a co-resident web server in our research. However, there is a significant limitation to their method since it requires the malicious VM and the co-resident target VM to be the only VMs on the physical machine.

II. REFERENCES

[1] O. Acıiçmez, C. Kaya Koç, and J.P. Seifert, On the power of simple branch prediction analysis, IACR Cryptology ePrint Archive, report 2006/351, 2006.

[2] Acıiçmez, and J.P. Seifert, Cheap hardware parallelism implies cheap security, Workshop on Fault Diagnosis and Tolerance in Cryptography – FDTC '07, pp. 80–91, IEEE, 2007.

[3] O. Acıiçmez, Ç. Kaya Koç, and J.P. Seifert, Predicting secret keys via branch prediction, RSA Conference Cryptographers Track – CT-RSA '07, LNCS vol. 4377, pp. 225–242, Springer, 2007.

[4] C. Almond, P. C. Chiquito, C. H. Fachim, S. Kim, M. Okajima and P. Rämö, Multitenant Utility Computing on IBM Power Systems Running AIX, IBM Redbooks, February 2009, http://www.redbooks.ibm.com/redbooks/pdfs/sg247681.pdf

[5] Amazon Web Services, Zeus Botnet Controller, Accessed on July 2011, http://aws.amazon.com/es/security/zeus-botnet-controller.

[6] Amazon Elastic Compute Cloud (EC2), http://aws.amazon.com/en/ec2, A. M. Azab, P. Ning, Z. Wang, X. Jiang, X. Zhang and N. C. Skalsky, HyperSentry: Enabling Stealthy In-context Measurement of Hypervisor Integrity, Proceedings of the 17th ACM Conference on Computer and Communications Security (CCS 2010) Chicago, Il, October 2010.

[7] S. Bozidar, Hacking Virtual Machines Part 1 – Sniffing, Accesed on July 2011, http://www.shortinfosec.net/2010/10/hacking-virtual-machines-part-1.html

[8] A. Cargile, Hypervisor Security Concerns, December,2009,http://thecoffeedesk.com/news/index.php/2009/12/01/hypervisor-security-concerns.

[9] R. Chow, P. Golle, M. Jakobsson, E. Shi, J. Staddon, R. Masuoka and J. Molina, Controlling Data in the Cloud: Outsourcing Computation without Outsourcing Control, Proceedings of the 2009 ACM Workshop on Cloud Computing Security (CCSW 2009), 2009 November 13, Chicago, IL. NY: ACM; 2009; 85-90.

[10] The Center for Internet Security, Security Configuration Benchmark for VMware ESX 3.5, December2009,http://benchmarks.cisecurity.org/tools2/vm/CIS_VMware_ESX_Server_3.5_Benchmark_v1.2.0.pdf

[11] P. Cox, Top virtualization security risks and how to prevent them, April 2011, http://searchcloudsecurity.techtarget.com/tip/Top-virtualization-security-risks-and-how-to-prevent them.

[12] Cloud Security Alliance, Security Guidance for Critical Areas of Focus in Cloud Computing V2.1, December 2009, https://cloudsecurityalliance.org/wp-content/uploads/2011/07/csaguide.v2.1.pdf