TITLE

Virtualization security in Data Centres & cloud

Prof Sarita Dhawale.

Ashoka Center for Business & Computer Studies,Nashik
Head of Department of Computer Science
University of Pune, Maharashtra.

Contact No: 9420890567
E-mail: saritadhawale@gmail.com

Abstract:

In the past decade, with the unprecedented growth in tech companies and advances in cloud computing, it has become increasingly common for companies to incorporate virtualization in their data centers to fully utilize their hardware resources. As a result, virtualization and virtualization security have gone through major transforms in the recent years. Virtualization and its unique architecture have many characteristics and advantages over traditional non-virtualized machines. However, these new characteristics create new vulnerabilities and possible attacks on a virtualized system. In addition, there are new challenges due to the infrastructure of virtualization. Luckily, solutions to the vulnerabilities have been developed or are in the process of being developed. Most of the solutions target either the virtualization architecture itself or the infrastructure. Some virtualization security companies have already utilized many of the solution concepts into their products to combat the vulnerabilities that are present.

Keyword:-

Virtual Machine-VM, Virtual Server, Hypervisor, Virtual Machine Monitor, VMM, Virtualization, VM Sprawl, Security, Emulator, Emulation, VM Sprawl, State Restore, External Monitoring, Denial of Service, VM Escape **OS** Operating System

Introduction

In the past decade, with the unprecedented growth in tech companies and advances in cloud computing, it has become increasingly common for companies to incorporate virtualization in their data centers to fully utilize their hardware resources. According to a research done by Nemertes Research, nearly 93% of the organizations it surveyed in 2012 have deployed virtualization in their servers. However, with the vast benefits that come with adoption of virtualization, new challenges and vulnerabilities also arise at the same time.

This survey paper first provides an overview on the current state of virtualization. Although many forms of virtualization exist, this paper will primarily focus on virtualization techniques that are used in modern data centers and clouds. In addition, this paper will discuss the security vulnerabilities brought about by different virtualization techniques. Specifically, the paper will address the forms of possible attacks on a virtualized machine, the advantages of using virtualization, and some current challenges. Lastly, the paper will present plausible solutions to the security vulnerabilities of virtualization. The solutions will incorporate theoretical defense mechanisms on the architecture and infrastructure, and examples of current virtualization security products developed by security firms. Virtualization is the abstraction of a hardware or software system that lets applications run on top of the virtualized environment without the need of

knowing the underlying resources available. The virtualized environment is otherwise known as the virtual machine (VM). In order to understand the security implications brought about by virtualization and how they can be addressed, this section provides an overview of the principles underlying virtualization.

2.1 Virtualization Architecture

Virtualization comes in different forms. They are distinguished primarily by the layer in the computing system to which virtualization is applied. However, all virtualization forms have an entity called a hypervisor or virtual machine monitor (VMM). It is the central unit that controls how virtualized programs interact with the underlying layer of resources. In a sense, it is the administrator of a virtualized environment. Application virtualization is a virtual implementation of the application programming interface. It enables programs to run on different platforms by providing the common virtual API. Operating system virtualization is a virtual implementation of an operating system (OS) where programs written for that OS can run. Despite the common appearances of the virtualization forms mentioned above, most modern data centers and clouds utilize a form known as full virtualization, which comes in two different types .

Native virtualization is where the hypervisor is directly implemented on the hardware or the computer firmware without any host OS. Each instance that runs on the virtual hardware is called a guest OS or VM. The hypervisor allocates resources between the VMs. Figure 1 shows a high level architecture of native virtualization.

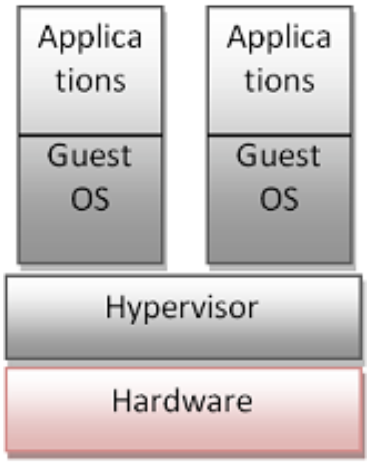

Figure 1: Native Virtualization Architecture

Hosted virtualization is where the hypervisor runs on a host OS that manages the hardware resources. The hypervisor still manages the guest OSs or VM, except the hypervisor is treated as an application on the host OS. Figure 2 shows a high level architecture of hosted virtualization

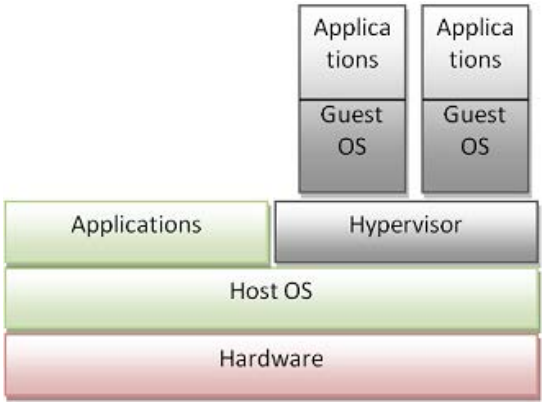

Figure 2: Hosted Virtualization Architecture

However, the basic architecture of virtualization does not sufficiently explain the security implications of virtualization. New characteristics present in virtualization that were not in traditional machines are the keys to obtain security for virtualized data centers and clouds.

2.2 Virtualization Characteristics

With the introduction of the virtualization architecture, new fundamental characteristics are present in virtualized systems. The four fundamental characteristics that affect security in virtualized systems are new management layer, concentration, variable state, and mobility .

The new management layer is essentially a layer created by the hypervisor. Since the hypervisor manages all VMs that run on the physical machine, it possesses administrative rights to all the virtualized components. For a virtualized data center or cloud, the information attackers seek is almost always in the VMs. Thus, by taking control of the hypervisor or the host OS that the hypervisor runs on, the attacker will be able to compromise most if not all of the VMs, posing significant threat to the entire data center or cloud.

Concentration is the characteristic that a plethora of VMs will run on the same physical machine, since the main purpose of virtualization is to fully utilize the physical resources or hardware available. It is an issue directly related to the new management layer created by virtualization. If one physical machine only runs one VM, then compromising the machine is no different than damaging one server that runs on the VM.

However, due to concentration, taking control of one VM on the machine can also potentially let the attacker gain access to other VMs as they run on the same physical machine, thus greatly increase the damage done.

Variable state is how each VM can be on, off, suspended, or in some customized state despite that the underlying hardware is still running. Due to this characteristic, many new security complexities are introduced dealing with virtualized systems. Some of the examples are access control to VMs in different states, data integrity of the VMs, and policies to change the state of VMs.

Mobility is the unique trait of virtualization that allows VMs to move from one physical machine to another without moving any hardware. Due to the ease at which VMs are transported across machines, security issues regarding networking and integrity when transporting VMs will become prevalent as VMs can be moved between machines in the same data center, another data center, or even clouds. This is an issue not present in non-virtualized environments, as movement of non-virtualized systems only requires moving the physical media. Also, the security boundary for each VM is very difficult to maintain as they can easily move between different infrastructures .Despite the issues that arise with the new characteristics in virtualization, there are many advantages to virtualization.

2.3 Virtualization advantages

The new characteristics in virtualization bring forth many implications to security, but deploying virtualization also offers substantial advantages. Five of the most prominent advantages are discussed here.

Cost effective operation is the primary advantage of virtualization that is caused by concentration.

Since virtualization utilizes one physical machine to the fullest by running multiple VMs on it, the processors and storage up time will be much greater than if the machine is only running one instance of an OS. Thus, the cost of operation for data centers and clouds is greatly reduced as less physical capital will need to be purchased.

Cost effective security is another advantage that spins off concentration. Due to concentration, perimeter defenses only need to be applied to one physical machine rather than multiple, which reduces the cost. Also, security appliances can be applied to each VM with software rather than securing each physical server, again decreasing the cost for security.

Isolation is a strong inherent defense of virtualization in that each VM runs without the knowledge of other VMs. Besides some instances that allow communication between VMs, when a VM is compromised, it is generally difficult for an attacker to access other VMs as only the hypervisor knows the existence of other VMs. However, bypassing isolation is possible and can be exploited.

Fast Recovery is a significant advantage of virtualization that comes due to its mobility. Since VMs can be easily transported, moved, and copied to other locations due to its nature, backup images of VMs can be effortlessly made. Thus, if an attack compromises a VM, it can simply be restored to a previous state. Likewise, any erroneous changes can be easily restored. In some cases, the entire data center may be backed up at a backup data center to prevent a large scale attack.

Variable state introduces complexities in the security architecture, but it is perhaps the best defense of virtualization. Since the state of a VM can change with a few simple operations from an administrator, the attackers can easily lose their progress due to a VM being turned off. After all, the best way to defend against attacks online is to not be online at all, which is what variable state provides.

With the architecture, characteristics, and advantages of virtualization in mind, the vulnerabilities of virtualization will be explored in the following section .Although virtualization is not a new topic, its prevalence in its areas of applications such as data centers and clouds has grown unprecedentedly. As a result, the need for security for virtualization and its new infrastructure has become increasingly important. This section will discuss the current and possible vulnerabilities in virtualized data centers and clouds.

3.1 Attacks on Hypervisor

Not surprisingly, the most obvious way to attack a virtualized data center or cloud is to gain access to the hypervisor, which controls all the VMs running in the data center or cloud. For the native virtualization architecture, there have been no known attacks on a hypervisor due to its nature of being embedded in the hardware. Otherwise, two types of attacks on the hypervisor exist: attack on hypervisor through the host OS and attack on hypervisor through a guest OS.

Attacks on hypervisor through host OS is to exploit vulnerabilities of the host OS on which the hypervisor runs. Due to native virtualization architecture requires specially configured hardware; most virtualization deployments are done with the hosted architecture. With vulnerabilities and security holes in most modern OSs, attacks can be done to gain control of the host OS. Since the hypervisor is simply a layer running on top of the host OS, once the attacker has control of the host OS, the hypervisor is essentially compromised. Thus, the administrative privileges of the hypervisor will enable the attacker to perform any malicious activities on any of the VMs hosted by the hypervisor. This propagation

of attacks from the hosted OS to the hypervisor then to the VMs is shown in Figure 3.

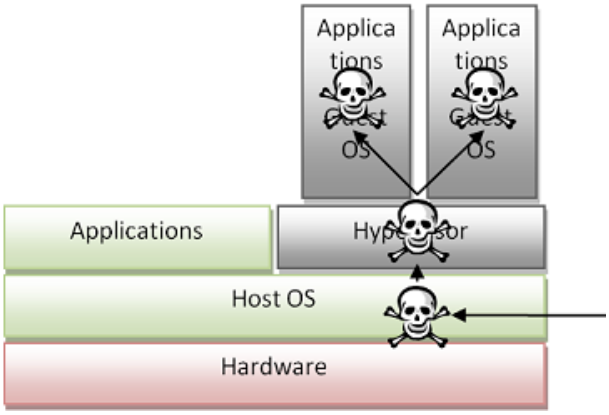

Figure 3: Attack on Hypervisor through Host OS

Attacks on hypervisor through guest OS is to use a guest OS to gain unauthorized access to other VMs or the hypervisor. This is also known as VM escapes or jailbreak attacks as the attacker essentially "escapes" the confinement of the VM into layers that are otherwise unknown to the VM . This is the most plausible attack on the hypervisor, since usually an attacker can only compromise a VM remotely as the underlying host OS is invisible. However, since many VMs share the same physical resources, if the attacker can find how his VM's virtual resources map to the physical resources, he will be able to conduct attacks directly on the real physical resources. By modifying his virtual memory in a way that exploits how the physical resources are mapped to each VM, the attacker can affect all the VMs, the hypervisor, and potentially other programs on that machine. Figure 4 shows the relationship between the virtual resources and the physical resources, and how the attacker can attack the hypervisor and other VMs.

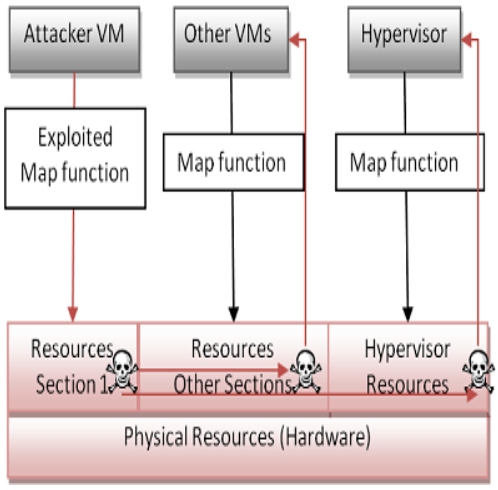

Figure 4: Attack on Hypervisor through Guest OS

These two types of attacks are the most obvious vulnerabilities in virtualization, but there are other potential ways to exploit a virtualized data center or cloud.

3.2 Other Forms of Attack

The two attacks on the hypervisor mainly exploit the architecture of virtualization. Other forms of attack such as virtual library check-out, migration attack, and encryption attack are exploits on the characteristics and infrastructure of virtualization.

Virtual library check-out is when a checked-out VM image becomes infected on another VMM and later readmitted to its original virtual library . This type of attack exploits on the fact that the guest VMM may not be as secure as the original virtual library. When the VM image becomes infected on the guest VMM and readmitted into the virtual library, the infection can potentially spread through the entire virtual library to other VMs and hypervisors in the data center or cloud. An example of virtual library check-out is described as follows. A large company hosts many virtual servers in its data center. This virtual library of servers consists of all the VMs controlled by the company. Periodically, employees of the company check out images of the VMs to perform software upgrades and maintenance on their own machines. Because the employee machines are not as secure as the company data center, attackers compromised one of the VMs on an employee machine and planted an infectious seed in the VM image. When the employee checks the VM back into the company data center, the infection migrates with the VM into the data center, giving the attacker access to everything in the data center.

Migration attack is an attack on the network during VM migration from one place to another. This attack is an exploit on the mobility of virtualization. Since VM images are easily moved between physical machines through the network, enterprises constantly move VMs to various places based on their usage.

For example, VMs from a canceled customer may be moved to a backup data center, and VMs that need maintenance may be moved to a testing data center for changes. Thus, when VMs are on the network between secured perimeters, attackers can exploit the network vulnerability to gain unauthorized access to VMs. Similarly, the attackers can plant malicious code in the VM images to plant attacks on data centers that VMs. Similarly, the attackers can plant malicious code in the VM images to plant attacks on data centers that

Encryption attack is an attack used to retrieve unauthorized information from VMs by exploiting security vulnerabilities in the virtualization software. This type of attack is not usually done in practice due to its complexity. Some of the possible exploits with

encryption attack include gaining access information to a VM, gaining session keys between VM transfers (like a migration attack), and obtaining other storage information by gaining the encryption key used to store VMs.

In addition to the types of vulnerabilities presented in this section, virtualization security also faces challenges in several areas.

3.3 Challenges

With the fast growth in virtualization and virtualization security the past few years, many problems have been solved by new and existing companies. However, the field of virtualization continues to face challenges in many areas such as monitoring, visibility, and infrastructure.

Monitoring is the ability for data centers and clouds to log trustworthy data on activities in VMs or the hosts. Usually, a company only Imposes strong defense and monitoring on the perimeter networks, which is an insufficient strategy as they have little to no protection against internal threats. However, even for companies that provide extensive internal monitoring, the characteristics of virtualization make monitoring very difficult. The new management layer created in virtualization is intended to abstract away the underlying resources from the VMs, but due to this new layer, some information may be abstracted away from a monitor, which will generate insufficient data to determine potential threats. Moreover, the variable states and mobility of each VM makes implementing a monitor that oversees every VM in the data center very difficult, as the monitors do not control the VMs. Also, additional monitoring would be needed to ensure host OSs and hypervisors are not potential threats.

Visibility is how much intrusion detection and prevention systems can see into a virtualized network. It is an issue closely related to monitoring; since with no monitoring, there will be no detection or prevention. Visibility is also an issue for the virtualization software vendors. A very limited view is provided into the host OS and virtual network with the current virtualization software by leading companies such as VMware. Their implementation is written with the intention of protecting the hosts and network from infected VMs. However, this characteristic also causes the visibility on the hosts OSs and the virtual networks to lower, making it harder to detect infected VMs and to prevent malicious intrusions. Again, there currently lacks a balanced solution between visibility and inherent security for virtualization.

Infrastructure is the way virtualization is set up in a data center or cloud. Many companies use virtualization software and security software from various vendors. Their data centers or clouds' setups largely depend on which vendors' software they used. As a result, the security structure within a virtualized data center or cloud needs to be highly specific to the particular data center or cloud. This in turn causes security between data bases and clouds to weaken due to misconfiguration, incompatibility or other potential issues. Also, different security measures for VMs and hosts within a data center can cause unforeseen problems. All of these problems come from the many ways a virtualization infrastructure can be set up. Although many vulnerabilities and challenges still exist, countless effective solutions have been developed by virtualization security firms. Some of the leading solutions and techniques of virtualization security will be examined next. With the growth of virtualization and problems in virtualization security, many firms and researchers have developed ways to combat the potential vulnerabilities. This section will examine the prominent approaches to virtualization security and present some existing products based on these approaches.

4.1 Solutions Based on Virtualization Architecture

The solutions based on virtualization architecture aim to solve security vulnerabilities by employing security measures on the virtualization components and characteristics. The three major approaches are hypervisor security, guest OS security, and image management security.

Hypervisor security is the application of traditional security measures to the hypervisor .This is a principle component of virtualization security. The hypervisor is the entire management layer for a virtualized system. Thus, if the hypervisor is compromised, then so are all the VMs created or controlled by the hypervisor. As long as the security of the hypervisor is strong enough, compromising all the VMs will be difficult for the attacker. For native virtualization architecture, there are currently many physical ways to ensure access control to the hypervisor. An example would be a hardware token possessed by the administrator in order to launch the hypervisor. However, as noted before, attacks on hypervisor in a native virtualization architecture is currently not known, thus making hypervisor security on such architecture almost irrelevant. For hosted virtualization architecture, traditional ways to protect running processes on an OS are currently implemented to protect the hypervisor. Security measures such as access control, automatic updating, networking, and introspection on guest OSs are all ways to protect the hypervisor from unauthorized access. These elements of security are generally implemented in software and can be easily updated to keep the security features of the hypervisor up to date.

Guest OS security is the application of traditional security measures to the guest OSs .This may sound like a redundant process to hypervisor security, but in virtualization, every component must be secure in order for the virtualized system to be secure. Since guest OSs running on a VM behave just like a real OS on physical machine, important security measures for single instance OSs are deployed on each guest OS.

Also, each guest OS must have sufficient isolation so one VM being compromised does not lead to other VMs on the same machine being compromised. More importantly, since guest OSs can use physical peripherals available on the machine, the communication between guest OSs and the hypervisor must be secure and the abstraction provided by the hypervisor must be enforced. Currently, many virtualization security firms are using guest OS monitoring to detect and quarantine infected guest OSs or revert them to a previous stage with stored guest OS images.

Image management security is the securing of how VM images are stored, transported, and managed in a virtualized data center or cloud. This is an important aspect of security in virtualization due to mobility and variable state in each VM, and how attackers exploit the fact that security measures are weaker on the network or backup data centers. Thus, to achieve image management security, strong storage encryption must be applied so sensitive data does not leak from the images; strong network security must be in place to ensure safe transportation of VM images. Another fact to consider is that VM images can be created quickly and easily. This can generate many unnecessary distributions of the same VM, and this vulnerability is generally called VM sprawl . In order to control the unnecessary distribution of VM images, a strong access control on the image management facility must be in place. VM software companies generally implement different levels of authority to control how each image can be managed to ensure image management security.

The solutions discussed above are all generic approaches to achieving security in virtualization. The actual implementations of these approaches can differ significantly, and it is outside the scope of this paper to discuss them. In addition to securing the components in virtualization, security measures in the infrastructure itself can greatly reduce the possibility of attacks.

4.2 Solutions Based on Virtualization Infrastructure

The solutions based on virtualization infrastructure aim to solve security vulnerabilities by creating secure gateways in the virtualization infrastructure. This set of solutions is predominantly for data centers and clouds as infrastructure is an integral part in the construction process. The two dominant areas are security on the virtual layer and security on the physical layer.

Security on virtual layer is achieved by securing how VMs and hypervisors talk to each other in a virtual network. In order to take full advantage of the virtualization infrastructure, virtual private networks (VPNs) are commonly created to manage different levels of authority in VMs. Because of the virtual nature of the network, features such as monitoring, access controls, integrity, encryption, authentication, and transportability of VMs can be implemented directly into the network. This will solve many of the vulnerabilities present in a virtualization as the security on the virtual layer will isolate different virtual management networks and bring ease to deployment and operation of VMs across different authorities or data centers.

Security on physical layer is the design of the structure of the physical systems that brings about security in a virtualized environment. One of the most notable features in this area is host-based intrusion detection and prevention. It allows the system to ensure that at least the physical layer will not be compromised easily through other means. The structure of the data center or the cloud also plays an important role. How the machines that are running the VMs interconnected physically can determine the possible security measures that can be used. Also, routine inspection for hardware failures and outdated systems is part of the security on the physical infrastructure that plays a large role in determining how secure the virtualized environment is. Although it is outside the scope of this paper to discuss how each solution is implemented, the next part will present some of the current products that utilize these general solutions.

4.3 Example Virtualization Security Products

The growth of virtualization recently has brought about many virtualization security products. It is not within the scope of this paper to discuss all of them, but a couple simple examples will be provided to show how the generic solutions are being used in the industry.

Trend Micro Virtualization Security solution is based on infrastructure of virtualization. They believe that malware can enter from various levels of the virtualization infrastructure such as apps running on the guest OSs, apps on the host OSs, or even the OSs themselves. Therefore, they deployed various levels of security on the virtual layer to protect the entire system. For example, their solution consists of a watch dog on the hypervisor level that solves the issues of monitoring and attacks on the hypervisor from the host OS.

Their solution also consists of intrusion detection modules on each of the VMs. This will solve the vulnerability issue due to concentration since each VM has its own self defense mechanisms from the hypervisor. The infrastructure of the solution is shown

in Figure 5. The security components are shaded in green.

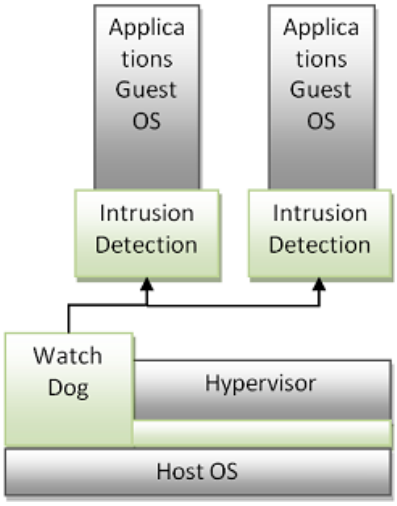

figure 5:

Trend Micro Virtualization Security Infrastructure **VBlock©** is another infrastructure based solution developed by EMC. It is a combination of security on the virtual layer and physical layer. EMC designed VBlock with the intention that it would be easy for companies to set up an already secured virtualization infrastructure. In each VBlock, EMC heavily integrated its own hardware security features with Cisco's networking equipments, and on top of that, the VMware is integrated as part of the VBlock. Because of the heavy integration between the physical layer, the

networking layer, and the virtual layer, and each having their own sets of security and defense mechanisms, a VBlock comes as a pre-integrated and secure infrastructure for any data center or cloud to use. This implementation solves the internal vulnerabilities that most companies face since most already have a strong perimeter defense against intrusions. The infrastructure of VBlocks in a virtualized environment is shown in Figure 6.

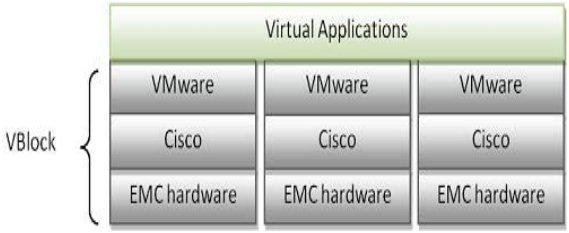

Figure 6: Infrastructure with VBlock

There exist countless other solutions developed by virtualization security companies in the past few years. However, most of them generally employ traditional security mechanism on the virtualization architecture and develop new infrastructures to achieve stronger data center or cloud wide security.

Conclusion

Virtualization has become an increasingly important field in the recent years due to the growth of data centers and cloud computing. The most common two types of virtualization architecture are native virtualization and hosted virtualization, where in native virtualization the hypervisor runs directly on the hardware and in hosted virtualization the hypervisor runs on a hosted OS. With the implementation of virtualization in a data center with the growth of virtualization and problems in virtualization security, many firms and researchers have developed ways to combat the potential vulnerabilities. Some basic forms of attack include attack on the hypervisor through either the guest OS or host OS. Other forms of attacks that can potentially compromise a virtualized system are virtual library check-out, migration attack, and encryption attack. These types of attacks, unlike the basic forms, do not attack the architecture of virtualization directly. Also, some new challenges in security are faced with the implementation of virtualization. Monitoring is harder due to lower visibility in a virtualized environment. This is also caused by the abstraction that virtualization brings. In addition, the infrastructure of virtualization is an on-going challenge for modern data centers and clouds.

With the vulnerabilities in virtualization, many solutions have been developed to combat them. The most basic forms of security involve implementing traditional security mechanisms such as intrusion detection software and firewall on components of virtualization such as the hypervisor and the guest OS. Also, security on how images of VMs are transported, stored and managed is very important due to mobility of VMs. To add additional layer of security, infrastructure security of virtualization is used. This form of security usually involves securing the virtual infrastructure, the physical infrastructure or bot. currently; sufficient virtualization security can be achieved by employing enough measures discussed in the paper. However, the continued growth of virtualization in data centers, clouds, and everywhere bring about new vulnerabilities and challenges to be solved. But at the same time, there is a growth in the virtualization security companies that continues to solve these problems.

Virtualization introduces an entirely new environment to the data center, with unique characteristics associated with management, concentration, variable state and movement. IT staffs must evaluate each one of these characteristics for potential compliance risk. Moreover, as virtualization deployments mature and expand, the dynamics of virtualization will force organizations to rearchitect their security defenses from a strong perimeter defense to defense-in-depth.

References:

1 Karen Scarfone, Murugiah Souppaya, Paul Hoffma, "Guide to Security for Full Virtualization Technologies", 2011 http://csrc.nist.gov/publications/nistpubs/800-125/SP800-125-final.pdf

2. Ted Ritter, "Virtualization Security", 2009 http://www.gtsi.com/eblast/corporate/cn/02_25_2010/PDFs/Nemertes%20Virtualization%20Security%20Key%20Trends.pdf

3 Mike Lococo, "Virtualization and Security Boundaries", 2009 http://mikelococo.com/files/2009/virtualization-and-security-boundaries.pdf

4.Alan Murphy, "Security Implications of the Virtualized Datacenter", 2007 http://www.f5.com/pdf/white-papers/virtual-data-center-security-wp.pdf

5. Rob Randell, "Virtualization Security and Best Practices", 2006 http://www.cpd.iit.edu/netsecure08/ROBERT_RANDELL.pdf

6. Trend Micro Virtualization Security, "Meeting the Challenges of Virtualization Security", 2009 http://trendedge.trendmicro.com/pr/tm/te/document/wp02_virtsec_090812us.pdf

7. Grzegorz Mucha, "RSA Security Solutions for Virtualization", 2009 http://poland.emc.com/collateral/campaign/global/forums/C1.pdf