ISSN 2320-9984

Volume 4 Issue 4 | December 2016 ISSN: 2320-9984 (Online)

International Journal of Modern Engineering & Management Research

Website: www.ijmemr.org

A Review on Security Issues for Virtualization and Cloud Computing

Abhiruchi Pillai
M.Tech. Research Scholar
Gyan Ganga Institute of Technology and Science
Jabalpur (M.P.), [INDIA]
Email: abhiruchip@gmail.com

Dr. Mukta Bhatele
Associate Professor
Department of Computer Science & Engineering
Gyan Ganga Institute of Technology & Sciences
Jabalpur (M.P.), [INDIA]
Email: mukta_bhatele@rediffmail.com

***Abstract*—*Cloud computing is the fastest growing technology in the IT world. The technology offers reduced IT costs and provides on the demand services to the individual users as well as organizations over the internet. The means of cloud computing is obtained by the virtualization of the resources such as hardware, platform, operating system and storages devices. Virtualization permits multiple operating systems to run on the same physical machine. Multiple tenants are unaware of the presence of the other tenant with whom they are sharing the resources. The co-existence of multiple virtual machines can be exploited to gain the access over other tenant's data or attack to deny of services. The significant concern is insuring the network security and providing isolation between multiple operating systems. The paper explores various kinds of vulnerabilities and attacks associated with the virtualization.*

Keywords:— *Cloud computing, Virtualization, Attacks, Network Security*

1. INTRODUCTION

Virtualization of servers in the cloud operates by adding a new layer to the software stack known as the hypervisor [1] or Virtual Machine Monitor (VMM) [2]. The hypervisor encapsulates the hardware, allowing it to be used by multiple operating system instances concurrently. This flexibility, coupled with the cost and performance advantages of sharing the underlying hardware, has revolutionized the computing industry: large numbers (i.e. hundreds of thousands) of generic hardware platforms, using multi-core blade technology, are now coupled through high-performance networking to produce a generic computing surface. Any subset of this collection can be combined to operate in tandem for a particular application using a multitude of operating systems. Conceptually, the hypervisor presents a virtual machine abstraction that restricts malicious code, executing within one instance of an operating system, from affecting a different instance. Unfortunately, hypervisors have introduced their own new security challenges: Adversaries now actively attempt to detect the presence of an operating hypervisor in order to tailor attacks accordingly [3]. A wide range of hypervisor detection techniques have already appeared against popular systems such as VMWare, VirtualPC, Bochs, Hydra, Xen, and QEMU [4]. Often, these techniques operate by exploiting timing differences between virtualized and non-virtualized operations [5]. Alternatively, they detect unusual memory locations associated with key operating system data structures [6].

The presence of a hypervisor has no impact on the vulnerabilities associated with the operating system. As a result, any exploit that leverages a known vulnerability will still operate successfully [9]. Although, a remote exploit gives the adversary control of a single virtual machine, by using the exploit in a virus the entire cloud could be compromised. It is this vulnerability amplification that poses the most significant threat to the future of cloud computing. Direct attacks against a vSwitch may undermine the operation of multiple virtual machines on a single host by denying connectivity to all of them simultaneously. The vSwitch provides the same functionality as a physical switch and in consequence exhibits the same vulnerabilities, enabling the same exploits [10]. For example, Address Resolution Protocol (ARP) spoofing involves the interception of valid network packets by sending fake ARP packets to a switch [11].

Hypervisor attacks involve the direct exploitation of vulnerabilities in the hypervisor. All virtual machines executing on a hypervisor have distinct data structures, separated in hardware. This separation forms a semantic gap [12] that prevents virtual machines from having visibility or impact upon each other's data structures. Direct Kernel Structure Manipulation (DKSM) bridges the semantic gap by patching virtual machine data structures and redirecting hypervisor accesses to shadow copies. This allows the virtual machine to present false information to the hypervisor regarding the virtual machine state, allowing implants, such as rootkits, to persist without detection.

Virtualization provides inherent redundancy and appears to provide robust, large-scale, cost-effective availability of shared resources. However, this perception is tempered by the known risk of vulnerability amplification and the paucity of knowledge regarding zero-day exploitation in clouds: history has shown that lack of detection does not imply lack of infection. Current mitigation techniques reviewed by this paper have already evolved based on malware detection and prevention, secure virtual machine managers, and cloud resilience. These three categories and their roles in preventing an attacker from gaining access to the cloud is illustrated in Figure 1.

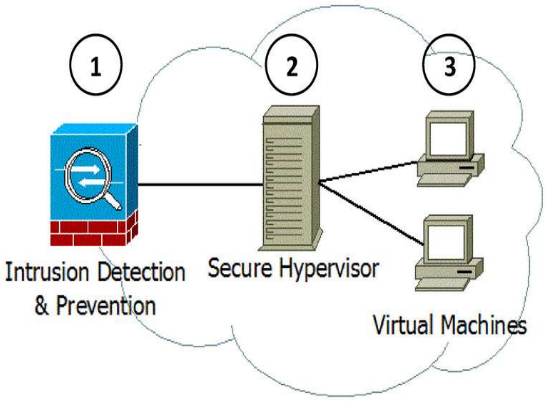

Figure 1: The three cloud security techniques reviewed by this paper: intrusion detection & prevention, secure hypervisors, and virtual machines.

2. VIRTUALIZATION SECURITY ISSUES

A. Software-as-a-service(SaaS) security issues

SaaS provides application services on demand such as email, conferencing software, and business applications such as ERP, CRM, and SCM. SaaS users have less control over security among the three fundamental delivery models in the cloud. The adoption of SaaS applications may raise some security concerns.

1) Application security

These applications are typically delivered via the Internet through a Web browser [12]. However, flaws in web applications may create vulnerabilities for the SaaS applications. Attackers have been using the web to compromise user's computers and perform malicious activities such as steal sensitive data [1]. Security challenges in SaaS applications are not different from any web application technology, but traditional security solutions do not effectively protect it from attacks, so new approaches are necessary [2]. The Open Web Application Security Project (OWASP) has identified the ten most critical web applications security threats [5]. There are more security issues, but it is a good start for securing web applications.

2) Multi-tenancy

SaaS applications can be grouped into maturity models that are determined by the following characteristics: scalability, configurability via metadata, and multi-tenancy [8]. In the first maturity model, each customer has his own customized instance of the software. This model has drawbacks, but security issues are not so bad compared with the other models. In the second model, the vendor also provides different instances of the applications for each customer, but all instances use the same application code. In this model, customers can change some configuration options to meet their needs. In the third maturity model multi-tenancy is added, so a single instance serves all customers. This approach enables more efficient use of the resources but scalability is limited. Since data from multiple tenants is likely to be stored in the same database, the risk of data leakage between these tenants is high. Security policies are needed to ensure that customer's data are kept separate from other customers. For the final model, applications can be scaled up by moving the application to a more powerful server if needed.

3) Data security

Data security is a common concern for any technology, but it becomes a major challenge when SaaS users have to rely on their providers for proper security [12]. In SaaS, organizational data is often processed in plaintext and stored in the cloud. The SaaS provider is the one responsible for the security of the data while is being processed and stored. Also, data backup is a critical aspect in order to facilitate recovery in case of disaster, but it introduces security concerns as well [6]. Also cloud providers can subcontract other services such as backup from third-party service providers, which may raise concerns. Moreover, most compliance standards do not envision compliance with regulations in a world of Cloud Computing. In the world of SaaS, the process of compliance is complex because data is located in the provider's data centres, which may introduce regulatory compliance issues such as data privacy, segregation, and security, that must be enforced by the provider.

4) Accessibility

Accessing applications over the internet via web browser makes access from any network device easier, including public computers and mobile devices. However, it also exposes the service to additional security risks. The Cloud Security Alliance [7] has released a document that describes the current state of mobile computing and the top threats in this area such as information stealing mobile malware, insecure networks (WiFi), vulnerabilities found in the device OS and official applications, insecure marketplaces, and proximity-based hacking.

B. Platform-as-a-service (PaaS) security issues

PaaS facilitates deployment of cloud-based applications without the cost of buying and maintaining the underlying hardware and software layers [9]. As with SaaS and IaaS, PaaS depends on a secure and reliable network and secure web browser. PaaS application security comprises two software layers: Security of the PaaS platform itself (i.e., runtime engine), and Security of customer applications deployed on a PaaS platform [10]. PaaS providers are responsible for securing the platform software stack that includes the runtime engine that runs the customer applications. Same as SaaS, PaaS also brings data security issues and other challenges that are described as follows:

1) Third-party relationships

Moreover, PaaS does not only provide traditional programming languages, but also does it offer third-party web services components such as mashups [10]. Mashups combine more than one source element into a single integrated unit. Thus, PaaS models also inherit security issues related to mashups such as data and network security [9]. Also, PaaS users have to depend on both the security of web-hosted development tools and third-party services.

2.7.2 Development Life Cycle

From the perspective of the application development, developers face the complexity of building secure applications that may be hosted in the cloud. The speed at which applications will change in the cloud will affect both the System Development Life Cycle (SDLC) and security [12]. Developers have to keep in mind that PaaS applications should be upgraded frequently, so they have to ensure that their application development processes are flexible enough to keep up with changes. However, developers also have to understand that any changes in PaaS components can compromise the security of their applications. Besides secure development techniques, developers need to be educated about data legal issues as well, so that data is not stored in inappropriate locations. Data may be stored on different places with different legal regimes that can compromise its privacy and security.

2.7.3 Underlying infrastructure security

In PaaS, developers do not usually have access to the underlying layers, so providers are responsible for securing the underlying infrastructure as well as the applications services. Even when developers are in control of the security of their applications, they do not have the assurance that the development environment tools provided by a PaaS provider are secure.In conclusion, there is less material in the literature about security issues in PaaS. SaaS provides software delivered over the web while PaaS offers development tools to create SaaS applications. However, both of them may use multi-tenant architecture so multiple concurrent users utilize the same software. Also, PaaS applications and user's data are also stored in cloud servers which can be a security concern as discussed on the previous section. In both SaaS and PaaS, data is associated with an application running in the cloud. The security of this data while it is being processed, transferred, and stored depends on the provider.

C. Infrastructure-as-a-service (IaaS) security issues

IaaS provides a pool of resources such as servers, storage, networks, and other computing resources in the form of virtualized systems, which are accessed through the Internet [4]. Users are entitled to run any software with full control and management on the resources allocated to them [8]. With IaaS, cloud users have better control over the security compared to the other models as long there is no security hole in the virtual machine monitor. They control the software running in their virtual machines, and they are responsible to configure security policies correctly. However, the underlying compute, network, and storage infrastructure is controlled by cloud providers. IaaS providers must undertake a substantial effort to secure their systems in order to minimize these threats that result from creation, communication, monitoring, modification, and mobility [11]. Here are some of the security issues associated to IaaS.

1) Virtualization

Virtualization allows users to create copy, share, migrate, and roll back virtual machines, which may allow them to run a variety of applications [2]. However, it also introduces new opportunities for attackers because of the extra layer that must be secured. Virtual machine security becomes as important as physical machine security, and any flaw in either one may affect the other. Virtualized environments are vulnerable to all types of attacks for normal infrastructures; however, security is a greater challenge as virtualization adds more points of entry and more interconnection complexity. Unlike physical servers, VMs have two boundaries: physical and virtual.

2) Virtual machine monitor

The Virtual Machine Monitor (VMM) or hypervisor is responsible for virtual machines isolation; therefore, if the VMM is compromised, its virtual machines may potentially be compromised as well. The

VMM is a low-level software that controls and monitors its virtual machines, so as any traditional software it entails security flaws. Keeping the VMM as simple and small as possible reduces the risk of security vulnerabilities, since it will be easier to find and fix any vulnerability.

3) Shared resource

VMs located on the same server can share CPU, memory, I/O, and others. Sharing resources between VMs may decrease the security of each VM. For example, a malicious VM can infer some information about other VMs through shared memory or other shared resources without need of compromising the hypervisor. Using covert channels, two VMs can communicate bypassing all the rules defined by the security module of the VMM [9]. Thus, a malicious Virtual Machine can monitor shared resources without being noticed by its VMM, so the attacker can infer some information about other virtual machines.

4) Public VM image repository

In IaaS environments, a VM image is a pre-packaged software template containing the configurations files that are used to create VMs. Thus, these images are fundamental for the overall security of the cloud [9]. One can either create her own VM image from scratch, or one can use any image stored in the provider's repository. For example, Amazon offers a public image repository where legitimate users can download or upload a VM image. Malicious users can store images containing malicious code into public repositories compromising other users or even the cloud system. For example, an attacker with a valid account can create an image containing malicious code such as a Trojan horse. If another customer uses this image, the virtual machine that this customer creates will be infected with the hidden malware. Moreover, unintentionally data leakage can be introduced by VM replication. Some confidential information such as passwords or cryptographic keys can be recorded while an image is being created. If the image is not "cleaned", this sensitive information can be exposed to other users. VM images are dormant artifacts that are hard to patch while they are offline [5].

5) Virtual machine rollback

Furthermore, virtual machines are able to be rolled back to their previous states if an error happens. But rolling back virtual machines can re-expose them to security vulnerabilities that were patched or re-enable previously disabled accounts or passwords. In order to provide rollbacks, we need to make a "copy" (snapshot) of the virtual machine, which can result in the propagation of configuration errors and other vulnerabilities [10].

6) Virtual machine life cycle

Additionally, it is important to understand the lifecycle of the VMs and their changes in states as they move through the environment. VMs can be on, off, or suspended which makes it harder to detect malware. Also, even when virtual machines are offline, they can be vulnerable [8]; that is, a virtual machine can be instantiated using an image that may contain malicious code. These malicious images can be the starting point of the proliferation of malware by injecting malicious code within other virtual machines in the creation process.

7) Virtual networks

Network components are shared by different tenants due to resource pooling. As mentioned before, sharing resources allows attackers to launch cross-tenant attacks. Virtual Networks increase the VMs interconnectivity, an important security challenge in Cloud Computing. The most secure way is to hook each VM with its host by using dedicated physical channels. However, most hypervisors use virtual networks to link VMs to communicate more directly and efficiently. For instance, most virtualization platforms such as Xen provide two ways to configure virtual networks: bridged and routed, but these techniques increase the possibility to perform some attacks such as sniffing and spoofing virtual network [4].

3. THREAT MODEL

The security implementation analyzed in this survey address the threat model for intrusions employing remote control outlined in Figure 2. It may involve several steps including surveillance to determine if vulnerability exists, use of an appropriate exploit or other access method, privilege escalation, removing exploit artifacts, and hiding behaviour [11]. Surveillance may involve obtaining a copy of the binary code and using reverse engineering or fuzzing [2] to facilitate a broad range of attack vectors including return oriented programming. The implant then persists for a time sufficient enough to carry out some malicious effect, obtain useful information, or propagate intrusion to other systems.

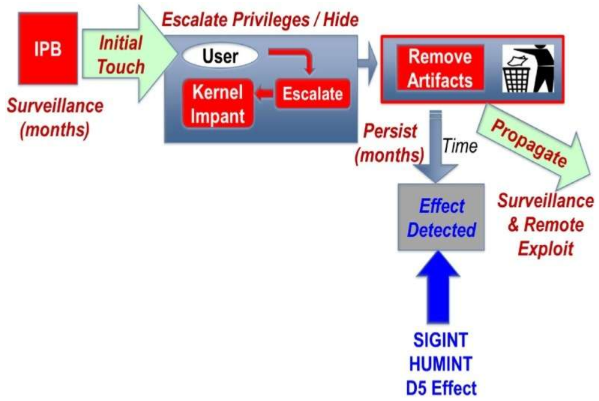

Figure 2: The threat model, detailing the process from surveillance to exploitation in the cloud.

Nevertheless, each cloud security technique represents an integral building block in the multilayered defence of the cloud. Malware detection and prevention systems are the initial line of defence in preventing an attacker from gaining a foothold on a cloud. The secure hypervisors present a hardened code base that restricts access to hardware to all, but the most privileged operations. Lastly, cloud resilient solutions are present to protect against the unknown exploits, which may allow an attacker to operate on a cloud indefinitely.

4. SECURITY ATTACKS IN VIRTUALIZATION

A. VM to VM attacks

The co-existence of multiple VMs on a single piece of hardware presents malicious VM owners to glean potentially sensitive information from victim VMs sharing the same hardware resources. An attacker may use guest OS (Virtual Machine) try to communicate and compromise other Virtual Machines on the same physical host, therefore breaking the isolation characteristic of VMs [12]. The attacker VM tries to attack over other VM exiting in the virtualized environment as the vulnerabilities in the LAN.

B. Sniffing Attack:

All the Virtual machines share a common network on a host machine. There is a possibility of traffic sniffing when the VMs communicate with each other. The hypervisors offer virtual network to link VMs using virtual bridge and route. In virtualized environments, virtual Hubs are created to share the same network in the bride mode. If these hubs are not properly configured, malicious attacker can try to sniff traffic to its own VM that is directed to other VM on the network. The malicious VM can use sniffing tools like “wire shark” to sniff the virtual network traffic. By using these tools attacker can sniff IP address of the other VM that is available neighbour to it. The attacker can perform packet sniffing attack over the victim. As a result, isolation is easy to be broken[8].

C. Spoofing Attack

In the route mode, virtual switches are used to connect the virtual machines to the host machine. The virtual switches need a dedicated interface to connect each VM. Media access control (MAC) address is assigned to each virtual machine. As address resolution protocol (ARP) is necessary to implement to redirect VMs’ traffic over the network. The routing table is maintained by sending an ARP command to each VM in boot time. A common vulnerability of ARP is ARP spoofing attack because ARP does not require proof-of origin. It is possible for the attacker to claim any MAC address by issuing ARP reply message with his IP. Hence the attacker can use ARP spoofing attack to redirect all the traffic of a victim VM to his VM [8].

D. Denial of Service (DoS)

In the virtualized environment all the virtual machines are sharing common resources such as storage space, network bandwidth, CPU usage. The denial of service attack is aimed to exhaust the common resources in order deny the services over the other guest virtual machines[3].In Denial of service attack one victim virtual machine receives more request than its capacity and other end users requests cannot be served. In the cloud environment, DoS attack is more dangerous than unclouded environment because of VMs are sharing their resources with other virtual machines over the same physical machine. One virtual machine can perform denial of service attack to another virtual machine in the virtualized environment.

The Transmission Control Protocol (TCP) provides reliable delivery of data over the internet. TCP can be exploited to perform denial of service attack known as TCP SYN flood attack. As a TCP connection is established a by 3-way handshake and the attacker takes advantage of this. An attacker overloads the victim with so many TCP connection requests that it will not be able to respond the legitimate requests. This is done through sending too many TCP SYN packets to the victim virtual machine. The victim allocates buffers for each new TCP connection and transmits a SYN-ACK in response to the connection request. The attacker does not replyto the SYN-ACK packets. Flooding based attacks can also exhaust other resources of the system [19].An attacker tries to identify the vulnerabilities in the hypervisor so that he can potentially access to the host OS and shared hardwires [4]

E. VM To Hypervisor Attacks

As the hypervisor is responsible for managing a virtualized cloud, the attackers target it to access the Guest VMs and the physical hardware that are shared among virtual machines. As the hypervisor resides between VMs and hardware, so the attack on the hypervisor can damage the VMs and hardware. It is identified that security of VMs is compromised.

5. THREATS IN CLOUD COMPUTING

Account or service hijacking

An account theft can be performed by different ways such as social engineering and weak credentials. If an attacker gains access to a user’s credential, he can perform malicious activities such as access sensitive data, manipulate data, and redirect any transaction.

Data scavenging

Since data cannot be completely removed from unless the device is destroyed, attackers may be able to recover this data.

Data leakage

Data leakage happens when the data gets into the wrong hands while it is being transferred, stored, audited or processed.

Denial of Service

It is possible that a malicious user will take all the possible resources. Thus, the system cannot satisfy any request from other legitimate users due to resources being unavailable.

Customer-data manipulation

Users attack web applications by manipulating data sent from their application component to the server’s application.

VM escape

It is designed to exploit the hypervisor in order to take control of the underlying infrastructure.

VM hopping

It happens when a VM is able to gain access to another VM (i.e. by exploiting some hypervisor vulnerability).

Malicious VM creation

An attacker who creates a valid account can create a VM image containing malicious code such as a Trojan horse and store it in the provider repository.

Insecure VM migration

Live migration of virtual machines exposes the contents of the VM state files to the network. An attacker can do the following actions:

(a) Access data illegally during migration.

(b) Transfer a VM to an untrusted host.

(c) Create and migrate several VM causing disruptions or DoS

Sniffing/Spoofing virtual networks

A malicious VM can listen to the virtual network or even use ARP spoofing to redirect packets from/to other VMs.[13]

6. RELATIONSHIPS BETWEEN THREATS AND VULNERABILITIES

INCIDENTS	COUNTERMEASURES
An attacker can use the victim's account to get access to the target's resources.	Identity and Access Management Guidance.
	Dynamic credential.
Data from hard drives that are shared by several customers cannot be completely removed.	Specify destruction strategies on Service-level Agreements (SLAs)
Authors in illustrated the steps necessary to gain confidential information from other VMs co-located in the same server as the attacker.	FRS techniques.
	Digital Signatures.
Side channel	Encryption
	Homomorphism encryption
An attacker can request more computational resources, so other legal users are not able to get additional capacity.	Cloud providers can force policies to offer limited computational resources
Some examples are described in such as SQL, command injection, and cross-site scripting	Web application scanners
A zero-day exploit in the HyperVM virtualization application that destroyed about 100,000 websites	HyperSafe
	TVDc (Trusted Virtual Datacenter)
An attacker can create a VM image containing malware and publish it in a public repository.	Mirage
Has empirically showed attacks against the migration functionality of the latest version of the Xen and VMware virtualization products.	PALM
Sniffing and spoofing virtual networks	Virtual network framework based on Xen network modes: "bridged" and "routed"

7. CONCLUSIONS

Cloud Computing is a relatively new concept that presents a good number of benefits for its users; however, it also raises some security problems which may slow down its use. Understanding what vulnerabilities exist in Cloud Computing will help organizations to make the shift towards the Cloud. Since Cloud Computing leverages many technologies, it also inherits their security issues. Traditional web applications, data hosting, and virtualization have been looked over, but some of the solutions offered are immature or inexistent. We have presented security issues for cloud models: IaaS, PaaS, and IaaS, which vary depending on the model. As described in this paper, storage, virtualization, and networks are the biggest security concerns in Cloud Computing. Virtualization which allows multiple users to share a physical server is one of the major concerns for cloud users. Also, another challenge is that there are different types of virtualization technologies, and each type may approach security mechanisms in different ways. Virtual networks are also target for some attacks especially when communicating with remote virtual machines. Some surveys have discussed security issues about clouds without making any difference between vulnerabilities and threats. We have focused on this distinction, where we consider important to understand these issues.

REFERENCES:

[1] Gartner, "The Top 10 Strategic Technology Trends for 2015", 2014. 08.

[2] CSA, "The Notorious Nine – Cloud Computing Top Threats in 2013", 2013. 02.

[3] Michael Jarschel, Simon Oechsner, Daniel Schlosser, Rastin Pries, Sebastian Goll, Phuoc Tran Gia, "Modeling and Performance evaluation of an Open Flow architecture", ITC '11 Proceedings of the 23rdInternational Teletraffic Congress, pp.1-7, 2011.

[4] Rasib Hassan Khan, JukkaYlitalo, Abu Shohel Ahmed, "Open ID authentication as a service in Open Stack", 2011 7th Information Assurance and Security (IAS), pp.372 -377, 2011.

[5] A Guide for Blocking the Outside Internet for Information Communication Service Providers, KISA, 2013.

[6] R. McRee, "Microsoft threat modeling tool 2014: identify &mitigate," ISSA Journal, pp. 39–42, 2014.

[7] M. Y. Yun, Y. S. Sin, S. G. Gi, H. C. Jung, and Y. J. Won, "Security requirements analysis studies for the construction of secure cloud services," in Proceedings of the Autumn Conference, pp. 453–457, The Korea Society of Management Information Systems, October 2012.

[8] S. J. Jung and Y. M. Bae, "Trends analysis of treats and technologies for cloud security," Journal of Security Engineering, vol. 10, no. 2, pp. 199–212, 2013.

[9] C. W. Lee, S. K. Kim, Y. M. Yeo, and J. S. Moon, "A studyon information requirements considering the security technical aspects in cloud service," Journal of Security Engineering, vol. 10,no. 3, pp. 355–370, 2013.

[10] C. S. Kim, B. I. Jang, and H. K. Jung, "A study on the security technology for the introduction of the secure cloud computing service," Journal of Security Engineering, no. 5,pp. 568–579, 2013.

[11] Khan, A., 2012. Access Control in

Cloud Computing Environment, In ARPN Journal of Engineering and Applied Sciences, vol-7, no-5.,pp.613 -615.

[12] Afoulk i, Z., et al. 2012. MAC protection of the Open Nebula Cloud environment, High Performance Computing and Simulation (HPCS), 2012 International Conference on, vol., no., pp.85

[13] Modi, C., Patel, D., Borisaniya, B., Patel, A. & Rajarajan, M. (2013). A survey on security issues and solutions at different layers of Cloud computing. The Journal of Supercomputing, 63(2), pp.561-592. doi: 10.1007/s11227-012-0831-5