Journal of Advanced Computer Science and Technology, 3 (2) (2014) 202-213
©Science Publishing Corporation
www.sciencepubco.com/index.php/JACST
doi: 10.14419/jacst.v3i2.3588
Research Paper

Securing the Cloud: Threats, Attacks and Mitigation Techniques

Mohammed M. Alani

Department of Information Technology, Al-Khawarizmi International College, Abu Dhabi, UAE
Email: m@alani.me

Abstract

This paper is aimed to present information about the most current threats and attacks on cloud computing, as well as security measures. The paper discusses threats and attacks that are most effective on cloud computing such as data breach, data loss, service traffic hijacking..etc. The severity and effect of these attacks are discussed along with real-life examples of these attacks. The paper also suggests mitigation techniques that can be used to reduce or eliminate the risk of the threats discussed. In addition, general cloud security recommendations are given.

Keywords: *cloud, cloud computing, security, threats, attacks, SaaS, PaaS, IaaS*

1. Introduction

Network security is a challenging task. It has become an integral part of any network service. With the rapidly increasing number of transactions happening on the Internet, security has become an essential part of everyday life.

Network security becomes much more difficult to control when the environment becomes as dynamic and demanding as cloud computing.

Cloud computing aims at reducing costs. This reduction is not only in terms of computing resource, but also in terms of helping its users to focus on the business instead of the information technology enabling this business. Cloud computing has evolved from many different technologies such as virtualization, autonomic-computing, grid-computing, and many other technologies [1].

With every new technology, new challenges arise. A very important challenge is providing adequate security to that cloud to perform as aimed.

In RFC 2828 [2], threat is identified as A potential for violation of security, which exists when there is a circumstance, capability, action, or event that could breach security and cause harm. On the other hand, the same RFC identifies an attack as an assault on system security that derives from an intelligent threat, i.e., an intelligent act that is a deliberate attempt to evade security services and violate the security policy of a system.

In general, computer security identifies three main objectives:

- Confidentiality: Assuring that data is available only to eligible entities and no unauthorized access to data can be obtained.
- Integrity: Assuring that data has not been altered in any way while it is stored or while its transport over the network.

- Authentication: Assuring the identity of the entity involved in the communication.

However, with the emergence of new technologies and threats, two more objectives can be added to the previous list:

- Availability: Assuring that data and services are always available at the required time.
- Accountability: Assuring that no entity can deny its participation in a data transfer between them [3].

These security objectives require the employment of certain security mechanisms and services to be implemented. We can identify a security mechanism as a process, or a device, aimed to detect, prevent, or recover from a security attack. Security mechanisms like encryption, hashing, steganography, etc. are commonly used in achieving security objectives.

A Security Service can be identified as a processing or communication service aimed to enhances the security of data and the information transfers of an entity. These services help in countering security attacks. Security services usually employ one or more security mechanism to achieve its goals [3].

2. Cloud Computing Service Models

As shown in figure 1, at the base of the cloud computing model, there is a hardware layer that contains the processors, memory, and storage components. Over that layer, there is an abstraction layer of software that realizes the unique characteristics of cloud computing. This is called the hypervisor. Above this layer of abstraction lie three layers Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS), and Software-as-a-Service (SaaS). Each one of these layers represent a service model of cloud computing.

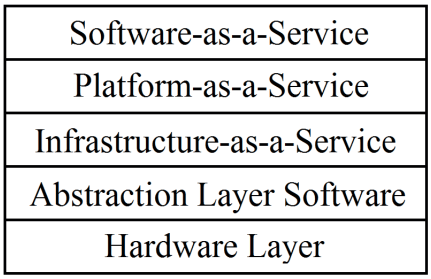

Figure 1: Layers of Cloud Computing

2.1. Infrastructure-as-a-Service

IaaS service model is the lowest level of service provided to the client. In this service model, the cloud computing client is provided with controlled access to the virtual infrastructure. Using this access, the client can install operating system and application software. From the clients point of view, this model is similar to renting the hardware from a service provider and letting the service provider manage the hardware. In this sense, the client does not have control over the physical hardware. On the other hand, the client will have to manage the security aspects from the operating system and up to the applications. This model requires the client to have highly experienced network engineer(s). Handling everything from the operating system and up is a big responsibility that most clients decline to handle, especially because of the security burdens. Thus, this model is not of high preference in the cloud computing clients society [4].

2.2. Platform-as-a-Service

In PaaS, the operating system and all platform-related tools (like compilers) are already installed for the client. These pre-installed components are also managed by the cloud service provider.

Clients have the freedom of installing additional tools based on their needs. However, the control over the infrastructure is retained by the service provider. The client controls applications development, configuration, and deployment. In some aspects, this service model is similar to the traditional webhosting services in which clients rent a remote server with development platform pre-installed on it. The major difference between this model and traditional web hosting is the rapid provisioning. Traditional web hosting is managed manually and requires

human intervention when the demand increases or decreases. On the other hand, provisioning in cloud computing is automatic and rapid. Thus, it does not require any human interventions [4].

2.3. Software-as-a-Service

SaaS model focuses on the application level and abstracts the user away from infrastructure and platform details. Usually, applications are provisioned via thin client interfaces such as web browsers or even mobile phone apps [4].

Microsofts Outlook.com is a clear example of this. An organization can adopt Outlook.com electronic mail service and never bother with hardware maintenance, service uptime, security, or even operating system management. The client is given the control over certain parameters in the software configuration, for example, creating and deleting mail boxes. These parameters can be controlled through the interface of the application.

3. Public And Private Clouds

Before we start discussing security threats and attacks, we need to identify private and public clouds. A public cloud can be identified as a cloud service provided and managed by an organization (such as Google, and Microsoft) to the general public. This cloud can provide services to multiple clients.

On the other hand, a private cloud can be identified as a cloud service provided and managed by an organization for the use of the organizations private use. This cloud provides services only to this particular organization.

Another type can be a hybrid cloud where a cloud is partially public, and partially private [4].

4. Threats

In addition to the regular threats to network security, the unique nature of cloud computing creates a different type of threats that are available only in a cloud environment. For example, attacks on cloud components like hypervisors are not available in the classic network security terminology.

In their The Notorious Nine report, Cloud Security Alliance (CSA) have identified nine threats that represent most important threats to cloud computing security in the year 2013 [5]. In the coming sections, these threats will be discussed.

4.1. Data Breaches

In many organizations, the worst scenario that can happen is to have your sensitive information falling in the hands of your competitors.

In the case of a poorly designed multitenant cloud service database, a flaw in one clients application could allow an attacker access the data of that client and all other clients [6].

In 2012, researchers introduced a side-channel attack by which one Virtual Machine (VM) can extract private cryptographic keys on the same physical machine [7]. It threats the models of IaaS, SaaS, and PaaS. Mitigation of this threat is not a simple task. The interaction between Data Breaches and Data Loss is delicate and the emergency plans need to be crafted carefully. One way of eliminating data breaches is to encrypt all of the clients data. However, if the encryption key is lost, the client would have a complete data loss. Thus, the client would need to have a backup copy of the data, somewhere else, or even offline backup. The client should keep in mind that having more copies of the data would potentially increase the probability of data breaches.

4.2. Data Loss

Data loss is not always caused by a malicious attacker. In addition to natural catastrophes such as floods, fires, and earthquakes, data loss can be a result of accidental erasure by the cloud service provider. In some scenarios, data loss can be cause by the client, and not the service provider. An example of this case is the one mentioned in the previous sub-section; if the client encrypts the data before uploading to the cloud and loses the encryption key, the data would be lost. Thus, the burden of data loss does not fall on the shoulders of the service provider only, but on the client as well.

In many countries, the organizations are required to keep complete audit logs of their work. If these logs were stored on the cloud and lost, this can jeopardize the existence of the organization and cause many legal issues.

It is considered a threat to the IaaS, SaaS, and PaaS models as well.

Mitigation of this threat can be done through backups. Regular (daily or even hourly) offline backups can be used to restore data with minimum loss. For services that have zero-tolerance for data loss, on-line backups with a different service provider can be a costly, but safe, solution.

4.3. Account or Service Traffic Hijacking

Old attacks like fraud, phishing, and exploiting software vulnerabilities are still in action. These attacks can still achieve the intended result for a malicious attacker. Reusing of usernames and passwords magnifies the severity of this threat.

A new scope is added to these attacks in cloud computing. The attacker, after gaining access to the clients credentials, can eavesdrop on the client transactions, return falsified information, manipulate data, and even redirect the users to illegitimate sites. In addition to that, the attacker can use the instances of the client as attacking bases to attack other people. Such access, can compromise confidentiality, availability, and integrity.

In 2009, Amazon had a large number of their cloud systems hijacked and were used to run Zeus botnet nodes [8]. According to [9], in 2010 Amazon.com had a cross-site scripting (XSS) bug that allowed attackers to hijack credentials from the site.

This threat exists in IaaS, SaaS, and PaaS models of the could system.

The first mitigation technique is to increase the awareness. Increasing the awareness of the organizations employees can be crucial in cancelling the dangers of attacks like phishing and other social engineering techniques.

Keeping all systems up-to-date and always patching an updating operating systems and protection software is also vital.

Another important aspect of mitigation is to prohibit account credentials sharing between users and services. The use of two-factor authentication like password and fingerprint, password and voiceprint, ..etc. can eliminate this threat to a great extent.

The client must be aware that there is no way to fully proof your system against these types of attacks. Instead, there are many ways to reduce their probability severely and scenarios of handling these breaches after they happen.

4.4. Insecure Interfaces and APIs

In order the client can manage and interact with the cloud services, the cloud service provider needs to provide a set of Application Programming Interfaces (APIs). These APIs are used for provisioning, management, orchestration, and monitoring. Availability and security of the cloud service is heavily dependent on the security of these APIs.

This security task becomes more complicated when many organizations build on these APIs to provide value-added services to their customers. This moves that APIs to a layered model which also increases risk as the organization might have to give their credentials to a third-party to enable them to create or use these new APIs. This threat exists on IaaS, SaaS, and PaaS models.

It is essential that the clients understand the security implications that come with the usage, management, orchestration and monitoring of cloud services.

It is also essential to select a cloud service provider that provides authentication and access control, encryption and activity monitoring APIs that are designed to protect against accidental as well as malicious attempts to circumvent the policy.

Depending on poorly designed APIs can compromise the confidentiality integrity, availability and accountability. Thus, secure and properly design APIs must be seeked by the client when selecting the cloud service provider.

4.5. Threats to Availability

Threats to availability exist in almost all networking services. In general, these threats aim at preventing the service from being provided to its intended audience. This can be through preventing website visitors from viewing the website, blocking legitimate user access to a Voice-over-IP (VoIP) server, ..etc.

In cloud computing, the situation is slightly worse. Denial of Service (DoS) attacks would not only render the service unavailable, but cause huge additional financial implications. Since cloud service providers charge their clients based on the amount of resources they consume, the attacker can cause a huge increase in the bill even if the attacker did not succeed in taking the clients system completely down. Another point that makes this threat even more dangerous in cloud computing is that cloud computing clients share the same infrastructure. Hence, a heavy DoS attack on one client can bring down the whole cloud. This threat exists in IaaS, SaaS, and PaaS models.

Being at the receiving end of a DoS attack is analogous to being caught in traffic lock; you cant get to your destination and you can do nothing about it except waiting. The service outage becomes very frustrating to clients and they start re-considering the reasons why they moved their data to the cloud [5].

There is no clearly-identified cure to this threat. However, service providers tend to use security appliances like firewalls, intrusion detection, and intrusion prevention systems that can help in reducing the risk and early detection of the attacks.

4.6. Malicious Insiders

A malicious insider, although the probability of occurrence is very low, can have a very high magnitude of impact. In [7], it is considered one of the highest possible risks on a cloud computing service. The reason behind that is that cloud architectures necessitate certain roles which are considered of the highest possible risk. An example of these roles is CP system administrators and auditors and managed security service providers dealing with intrusion detection reports and incident response.

Organizations that depend solely on the service provider in security are at great risk due to malicious insiders.

From IaaS to PaaS and SaaS models the malicious insider can have increasing access levels to more critical data [5]. Encrypting the client data will not completely mitigate this threat. If the encryption keys are not stored with the client and are only available at data-usage time, the system is still vulnerable to malicious insider attack. Thus, it is advisable that all client data is encrypted and the keys should be kept with the client.

In [10], a detailed review of the existing trust management research in cloud environments was introduced. The paper also provided a detailed assessment of the research on trust management. This can be useful in the avoidance and detection of insider attacks.

A novel method of detecting malicious insiders and malicious administrators was proposed in [11]. This paper discusses the threats to client data privacy and integrity posed by a malicious cloud administrator. The paper also explains how the existing cloud architecture needs to be extended to accommodate the suggested method. Abuse of Cloud Services The idea behind cloud computing is to provide low-cost high-resource solutions. These low-cost solutions can be very beneficial to small companies that require high computing resources for a short period of time. On the other hand, these services can be used by malicious attackers. The access to these huge computing resources can be abused and these resources can be directed towards attacking other systems.

The imminence of this threat has dropped over the past few years due to stricter policies followed by cloud service providers. This threat applies to IaaS, and PaaS models.

The only possible way to mitigate this threat is to use a cloud service provider that has strict policies related to service abuse with a quick response time to violations of these policies. There is nothing technical to be done as far as the client is concerned.

4.7. Insufficient Due Diligence

Cloud computing has become the big thing that every organization would like to use. Some companies jump into using cloud computing with maturation of the concept and without being completely prepared and comprehend the model fully.

The quick adoption of cloud computing can be a realistic goal for organizations that have the resources required to implement cloud computing properly. Organizations that does not fully comprehend the requirements of proper implementation of cloud computing will have many issues with operational responsibility, incident response, and many other aspects [5].

Organizations with weak understanding of cloud computing can have contractual issues over obligations on liability, response, or transparency. This is caused by mismatch of expectations between the client and the cloud service provider. Security issues might arise by pushing applications that are dependent on internal network-level security controls to the cloud. Organizations might also face unknown operational and architectural issues when they use designers and architects that are unfamiliar with cloud technologies for their applications. This threat applies to IaaS, SaaS, and PaaS models.

The only way to mitigate this threat is that the organization should not migrate to using cloud computing unless they are fully aware of their capabilities and certain that the have the human and information technology resources required.

4.8. Shared Technology Vulnerabilities

Cloud computing relies on many backend technologies to provide its services. Any vulnerability existent in the backend can lead to exploitation in all clients. There are cases where the underlying architecture (such as CPU caches, GPUs, etc.) does not provide complete isolation properties.

When an integral part is compromised, such as the hypervisor, it exposes not only the compromised client, but the entire environment. This type of vulnerability is considered extremely dangerous because it can affect a complete cloud all at once [5].

The severity of this threat has dropped over the past few years. This drop is due to more accurate configuration and isolation by the hardware manufacturers and the cloud service providers. This threat exists in IaaS, SaaS, and PaaS models.

The mitigation of this threat is to be done by the cloud service provider. Keeping systems updated and giving high attention to configuration can reduce the probability of exploiting such vulnerability.

4.9. Other Threats

There are many other threats to the cloud computing systems that have varying severity. The following list provides some of the other threats of less importance than the ones discussed earlier.

- Lock-in: No clear rules governing the portability of data and services. If a client wishes to change provider, many issues may arise [12].
- Insecure or Incomplete Data Deletion: when a request to delete a cloud resource is made, as with most operating systems, this may not result in true wiping of the data [12].
- Loss of Governance: There are some cases in which the Service-Level Agreements (SLA) do not cover all security aspects such the client is unclear about which measures need to be taken by the client and which by the cloud service provider.
- Acquisition of the cloud provider: this could increase the likelihood of a strategic shift and may put non-binding agreements at risk [12].

5. Attacks On The Cloud

Cloud computing, as any other platform, is a target for many attacks. These attacks have different aims starting from eavesdropping, to complete system failure. It is vital that these attacks are identified clearly and their mitigation techniques.

5.1. Denial-of-Service

DoS attacks try to render the service unavailable to its users. The attack consumes large amounts of system resources such as processing power, memory, and bandwidth. This consumption will leave the service inaccessible to the users or intolerably slow.

DoS attacks, and their variant Distributed Denial of Service (DDoS), grab high media attention mostly because of their magnitude. In 1988, only six DDoS attacks took place, as the reports show. DDoS attacks targeted large websites like CNN, Yahoo, and Amazon in the year 2000 with an attack rate of about 1GBps. DDoS attacks reached the rates of 70GBps in the year 2007. In 2013, a huge attack took place on Spamhaus spam detection service that reached highest rate at that time of 300 GBps [13]. Recently, in February 2014, the largest DDoS attack known until now with the rate of 400 GBps took place. This attack targeted a public cloud service provider called CloudFlare. Attacks of such magnitude affect not only their targets, but affect the overall Internet in the area. As mentioned in [14], regular Internet users experienced noticeable slowness in their Internet services.

DoS attacks can be operated on multiple layers in the network. An attacker can operate a DoS attack at the network level to render the whole server unreachable. This is done by getting the Network Interface Card (NIC) of the server completely occupied with useless packets in such a way that no more bandwidth is available for legitimate users. A DoS attack can be launched in the transport layer using the very old, but still effective, SYN Flood technique. In a SYN flood attack the attacker sends a flood of TCP SYN requests that gets the server busy without actually completing the three-way handshake procedure used in the setup of TCP sessions. DoS attacks can also be launched at the application level by sending fake requests to the application layer protocol to consume

the servers memory and processing power. Sending a flood of fake SMTP requests to an electronic-mail server is a clear example.

Many security appliances are currently capable of detecting simple DoS attacks that come from a single attacking node. Thus, DoS attacks have evolved into a more complicated attack called Distributed DoS (DDoS). In DDoS, the DoS attack is performed from multiple sources around the Internet such that it would be harder to trace and block the attacker. More information about DDoS in cloud computing can be found in [15].

Although DDoS take major attention in the media, it is not the only threatening type of DoS attacks. Another type of DoS is Asymmetric application-level DoS attacks that exploit the vulnerabilities in web-servers, databases, as well as other cloud resources. This type of attack allows a malicious attacker to bring down an application using very small attack payload, sometimes as small as 100 bytes.

In [16], a method that depends on covariance-matrix was used to detect DoS attacks. This method was proven to be highly effective in detecting DoS attacks that are based on flooding.

A new DoS attack along with its counter-measure was introduced in [17]. This attack operates on the application level to detect a network bottle-neck in one of the links and focus on flooding this link. One research direction has employed Game Theory defense mechanisms in defending the cloud infrastructure against a special type of DoS attacks called Co-Resident DoS attack. In co-resident DoS attack, the attacker rents a VM inside the public cloud and conducts the DoS from the rented VM onto another VM within the same node. The attacker uses simple tools (like nmap and hping) to deduce the exact location of the VM in the cloud and conduct a DoS attack on the bottle-neck network channel shared among the VMs. In [18] a detection method that is based on game theory defense mechanisms was introduced.

A model was suggested to prevent flooding attacks was introduced in [19]. This model can provide the foundation of further study in the topic of DoS flooding attacks prevention.

More DoS and DDoS detection and prevention methods can be found in [13, 20, 21, 22, 23]

5.2. Attacks on Hypervisor

As explained earlier, a hypervisor is the abstraction layer software that sits between the hardware and the VMs that comprise the cloud. Although not many attacks were conducted on hypervisors, any compromise in the hypervisor security can bring the whole cloud down.

Hyperjacking was identified in [25] as the attackers attempt to craft and run a very thin hypervisor that takes complete control of the underlying operating system.

Once the attacker gains full control of the operating system, the whole cloud is compromised. The attacker will be able to eavesdrop, manipulate clients data, disrupt or even shut down the complete cloud service. Although the probability of this attack succeeding is very low, it is still a source of concern.

Another hypervisor vulnerability was previously reported in [26]. This vulnerability was found in many commercial VM and cloud computing products.

In [27] a turn-around was suggested. The paper suggested the elimination of the hypervisor attack surface by enabling the guest VMs to run natively on the underlying hardware while maintaining the capability of running multiple VMs at the same time.

Another solution was suggested in [28] through employing a hierarchical secure virtualization model. The suggested hierarchical model employs a technique of threat quarantine and conquer in addition to complete control on virtualization.

In [29] a hardware protection scheme for VMs was suggested. This hardware protection scheme would protect the VMs from a compromised hypervisor while keeping the flexibility to manage the cloud environment required by the hypervisor. The proposed method applies to multi-core multiprocessor systems. More details on hypervisor vulnerabilities can be found in [30].

5.3. Resource-Freeing Attacks

When multiple VMs share the same physical node in a cloud, the performance of any given VM will degrade if another VM is over-using the resources. Research conducted in [31] has shown that the performance of a cache-sensitive benchmark can degrade by more than 80% because of interference from another VM.

The same paper, [31], discusses how an attack can be conducted by a cloud user to free up resources used by other users. By doing this, the user can have more free resources at his/her disposal.

5.4. Side-Channel Attacks

In a side-channel attack, the attacker gains information about the cryptographic technique used by analysing physical characteristics of the cryptosystem implementation. The attacker uses information about the timing, power consumption, electromagnetic leaks,..etc. to exploit the system. This collected information can be employed in finding sensitive information about the cryptographic system in use. For example, information about power consumption can result in knowing the key used in encryption.

These attacks, despite their relative easiness of implementation, can result in dangerous exploitations that can render the whole cryptosystem worthless. More information about side-channel attacks can be found in [32].

In cloud computing, side-channels attacks are conducted through gaining access to the physical node hosting the target VM. This access can be available through creating a VM in the same physical node that is hosting the target VM. This is particularly possible in public clouds. The attacker can keep creating VMs in the cloud until one VM is created in the same physicl node of the target VM. Afterwards, the attacker can start collecting information necessary to conduct the attack. A method to do this was introduced in [33].

Many researchers focused in their side-channel attacks research on cache memory. In [34], the research focused on threats on the L2 covert channels and how these threats can be exploited or countered.

Research was also done on attacking AES encryption in a virtualized environment. In this attack, the attacker was able to extract sensitive keying material from an isolated trusted execution domain using Bernsteins correlation [35].

Research on timing channels and their determination in cloud computing was conducted in [36]. This research proposed using provider-enforced deterministic execution as a replacement of resource partitioning to eliminate timing channels within a shared cloud domain.

Many researchers tackled countering side-channel attacks. Papers such as [37] discussed mitigating side-channel attacks in different environments. This paper proposed a general mitigation strategy that focuses on the infrastructure used to measure side channel leaks rather than the source of leaks. The proposed technique can be applied to all known and unknown microarchitectural side-channel leaks.

More cloud-computing oriented research on the mitigation of side-channel attacks was introduced in [38]. This paper proposed an approach that leverages dynamic cache coloring. The proposed dynamic cache coloring is by notifying the VM management software to swap the process data to a safe and isolated cache line whenever the application is handling security-sensitive data. The performance degradation caused by this method is less when compared to other techniques to mitigate cache side-channel attacks.

A similar method was later introduced in [39]. STEALTHMEM that was introduced in this paper is a system-level protection mechanism designed to counter cache side-channel attacks in a cloud computing environment. The suggested system contains a set of locked cache lines for each core. These cache lines are never evicted from the cache and the system efficiently multiplexes them so that each VM would be capable of loading its own sensitive data into the locked cache. In this way, any VM on the physical node can hide memory access patterns of confidential data from other VMs.

More cloud-based side-channel attacks mitigation techniques were introduced in [40, 41, 42, 43, 44]

5.5. Attacks on Confidentiality

It is a major concern for all cloud computing clients to secure their data. The confidentiality intended by clients is not only to protect their data from public attacks, but to protect their data from their cloud service provider. Clients would not accept that their service provider is capable of accessing their private data whenever they want. Thus, clients use encryption.

Confidentiality has always been a target for security attacks since the start of computers. In cloud computing, confidentiality is not only about client data confidentiality. Confidentiality is required in the cloud infrastructure as well. Exploiting private cloud information like encryption-keys, VM locations, or operating system information can lead to more dangerous attacks.

A non-technical attack can be conducted through social engineering. In a such attack, that attacker can get private information like encryption-keys, passwords, usernames,..etc. by tricking privileged users into giving access to their accounts. There is no single specific form this attack takes. It can be mostly done by the attacker impersonating the identity of an IT-support technician, system administrator, or any other person that can have access to private information. The best way to counter this attack through educating users about the nature and shape of these attacks and sometimes additional institutional policies can reduce the probability of such attacks occurrence.

Many other attacks on confidentiality can be orchestrated through side-channel attacks. In [45], an attack was conducted by mapping the internal cloud infrastructure, identifying where a particular VM is likely to reside, and then initiating VMs until on of the is co-resident with the target VM in the same physical machine. Afterwards, the attacker used this co-resident placement to mount cross-virtual-machine side channel attacks to extract information from the target VM that lies within the same physical machine. Amazon EC2 public cloud service was used for case study to conduct the attack. The first part of the procedure is similar to the approach used in [18] to conduct DoS attack.

Side channel attacks were also employed in [46] but this time it was used to extract private keys used in client data encryption. This attack was the first of its kind that was demonstrated on a symmetric multiprocessing system with Xen virtualization. The attacker was able to retrieve encryption keys of ElGamal [47, 48] encryption algorithm use in the target VM.

5.6. Other Attacks

In addition to the attacks described earlier, other attacks exist. More information about other attacks can be found in [49, 50, 51, 52, 53]

6. General Cloud Security Recommendations

- Install and maintain a firewall configuration. A firewall should be placed at each external network interface and between each security zone within the cloud [54].
- Do not use vendor supplied defaults for passwords and other security parameters.
- Research into standardized SLAs and liability provisions could lead to greater accountability [55].
- Ensure that no unnecessary functions or processes are active.
- Ensure patch management.
- Protect encryption keys from misuse or disclosure.
- Promptly revoke access for terminated users [54].

More general discussions and recommendations about cloud security can be found in [56, 57, 58, 59, 60]

7. Conclusions

With the rapid increase in the adoption of cloud computing by many organizations, security issues arise. In this paper we have discussed in details the most important threats and attacks on cloud security. The paper also included suggestions to mitigate these threats and attacks. The paper provided general cloud security recommendations as well.

References

[1] Buyya, R., Broberg, J., Goscinski, A.M.: Cloud computing: Principles and paradigms, vol. 87. *John Wiley & Sons*,2010.

[2] Shirey, R.: Rfc 2828: Internet security glossary. *The Internet Society*, 2000.

[3] Stallings, W.: Cryptography and Network Security, 4/E. *Pearson Education India*, 2006.

[4] Hill, R., Hirsch, L., Lake, P., Moshiri, S.: Guide to cloud computing: principles and practice. *Springer*, 2012.

[5] Top-Threats-Working-Group: The notorious nine: Cloud computing top threats in 2013. *Cloud Security Alliance* , 2013.

[6] Chong, F., Carraro, G., Wolter, R.: Multi-tenant data architecture. *MSDN Library, Microsoft Corporation* , 2006.

[7] Zhang, Y., Juels, A., Reiter, M.K., Ristenpart, T.: Cross-vm side channels and their use to extract private keys. *In: Proceedings of the 2012 ACM conference on Computer and communications security, pp. 305–316. ACM*, 2012.

[8] Godin, D.: The register: Zeus bot found using amazons ec2 as cc server. , 2009.. http://www.theregister.co.uk/2009/12/09/amazon_ec2_bot_control_channel/

[9] Godin, D.: The register: Amazon purges account hijacking threat from site , 2010.. http://www.theregister.co.uk/2010/04/20/amazon_website_treat/

[10] Noor, T.H., Sheng, Q.Z., Zeadally, S., Yu, J.: Trust management of services in cloud environments: Obstacles and solutions. *ACM Computing Surveys, CSUR)* ***46****(1), 12*, 2013.

[11] Bleikertz, S., Kurmus, A., Nagy, Z.A., Schunter, M.: Secure cloud maintenance: protecting workloads against insider attacks. *In: Proceedings of the 7th ACM Symposium on Information, Computer and Communications Security, pp. 83–84. ACM*, 2012.

[12] Catteddu, D.: Cloud Computing: benefits, risks and recommendations for information security. *Springer*, 2010.

[13] Yu, S.: Distributed Denial of Service Attack and Defense. *Springer*, 2014.

[14] Lenon, M.: Cloudflare infrastructure hit with 400gbs ntp-based ddos attack , 2014.. http://www.securityweek.com/cloudflare-infrastructure-hit-400gbs-ntp-based-ddos-attack

[15] Kumar, N., Sharma, S.: Study of intrusion detection system for ddos attacks in cloud computing. *In: Wireless and Optical Communications Networks (WOCN), 2013 Tenth International Conference on, pp. 1–5. IEEE*, 2013.

[16] Ismail, M.N., Aborujilah, A., Musa, S., Shahzad, A.: Detecting flooding based dos attack in cloud computing environment using covariance matrix approach. *In: Proceedings of the 7th International Conference on Ubiquitous Information Management and Communication, p. 36. ACM*, 2013.

[17] Liu, H.: A new form of dos attack in a cloud and its avoidance mechanism. *In: Proceedings of the 2010 ACM workshop on Cloud computing security workshop, pp. 65–76. ACM*, 2010.

[18] Bedi, H.S., Shiva, S.: Securing cloud infrastructure against co-resident dos attacks using game theoretic defense mechanisms. *In: Proceedings of the International Conference on Advances in Computing, Communications and Informatics, pp. 463–469. ACM*, 2012.

[19] Zunnurhain, K.: Fapa: a model to prevent flooding attacks in clouds. *In: Proceedings of the 50th Annual Southeast Regional Conference, pp. 395–396. ACM*, 2012.

[20] Karnwal, T., Thandapanii, S., Gnanasekaran, A.: A filter tree approach to protect cloud computing against xml ddos and http ddos attack. *In: Intelligent Informatics, pp. 459–469. Springer*, 2013.

[21] Chouhan, V., Peddoju, S.K.: Hierarchical storage technique for maintaining hop-count to prevent ddos attack in cloud computing. *In: Proceedings of International Conference on Advances in Computing, pp. 511–518. Springer*, 2012.

[22] Gupta, S., Kumar, P.: Vm profile based optimized network attack pattern detection scheme for ddos attacks in cloud. *In: Security in Computing and Communications, pp. 255–261. Springer* , 2013.

[23] Contractor, D., Patel, D.R.: Trust management framework for attenuation of application layer ddos attack in cloud computing. *In: Trust Management VI, pp. 201–208. Springer*, 2012.

[24] Yu, S.: Distributed Denial of Service Attack and Defense. *Springer*, London, 2014.

[25] Ray, E., Schultz, E.: Virtualization security. *In: Proceedings of the 5th Annual Workshop on Cyber Security and Information Intelligence Research: Cyber Security and Information Intelligence Challenges and Strategies, p. 42. ACM*, 2009.

[26] VU, C.V.N.: 649219

[27] Szefer, J., Keller, E., Lee, R.B., Rexford, J.: Eliminating the hypervisor attack surface for a more secure cloud. *In: Proceedings of the 18th ACM conference on Computer and communications security, pp. 401–412. ACM*, 2011.

[28] Manavi, S., Mohammadalian, S., Udzir, N.I., Abdullah, A.: Hierarchical secure virtualization model for cloud. *In: Cyber Security, Cyber Warfare and Digital Forensic (CyberSec), 2012 International Conference on, pp. 219–224. IEEE*, 2012.

[29] Szefer, J., Lee, R.B.: A case for hardware protection of guest vms from compromised hypervisors in cloud computing. *In: Distributed Computing Systems Workshops (ICDCSW), 2011 31st International Conference on, pp. 248–252. IEEE*, 2011.

[30] Perez-Botero, D., Szefer, J., Lee, R.B.: Characterizing hypervisor vulnerabilities in cloud computing servers. *In: Proceedings of the 2013 international workshop on Security in cloud computing, pp. 3–10. ACM*, 2013.

[31] Varadarajan, V., Kooburat, T., Farley, B., Ristenpart, T., Swift, M.M.: Resource-freeing attacks: improve your cloud performance (at your neighbor's expense). *In: Proceedings of the 2012 ACM conference on Computer and communications security, pp. 281–292. ACM*, 2012.

[32] Zhou, Y., Feng, D.: Side-channel attacks: Ten years after its publication and the impacts on cryptographic module security testing. *IACR Cryptology ePrint Archive **2005**, 388*, 2005.

[33] Zhang, Y., Juels, A., Oprea, A., Reiter, M.K.: Homealone: Co-residency detection in the cloud via side-channel analysis. *In: Security and Privacy (SP), 2011 IEEE Symposium on, pp. 313–328. IEEE*, 2011.

[34] Xu, Y., Bailey, M., Jahanian, F., Joshi, K., Hiltunen, M., Schlichting, R.: An exploration of l2 cache covert channels in virtualized environments. *In: Proceedings of the 3rd ACM workshop on Cloud computing security workshop, pp. 29–40. ACM*, 2011.

[35] Weiß, M., Heinz, B., Stumpf, F.: A cache timing attack on aes in virtualization environments. *In: Financial Cryptography and Data Security, pp. 314–328. Springer* , 2012.

[36] Aviram, A., Hu, S., Ford, B., Gummadi, R.: Determinating timing channels in compute clouds. *In: Proceedings of the 2010 ACM workshop on Cloud computing security workshop, pp. 103–108. ACM*, 2010.

[37] Martin, R., Demme, J., Sethumadhavan, S.: Timewarp: Rethinking timekeeping and performance monitoring mechanisms to mitigate side-channel attacks. *In: ACM SIGARCH Computer Architecture News, vol. 40, pp. 118–129. IEEE Computer Society*, 2012.

[38] Shi, J., Song, X., Chen, H., Zang, B.: Limiting cache-based side-channel in multi-tenant cloud using dynamic page coloring. *In: Dependable Systems and Networks Workshops (DSN-W), 2011 IEEE/IFIP 41st International Conference on, pp. 194–199. IEEE*, 2011.

[39] Kim, T., Peinado, M., Mainar-Ruiz, G.: Stealthmem: system-level protection against cache-based side channel attacks in the cloud. *In: Proceedings of the 21st USENIX conference on Security symposium, pp. 11–11. USENIX Association*, 2012.

[40] Stefan, D., Buiras, P., Yang, E.Z., Levy, A., Terei, D., Russo, A., Mazières, D.: Eliminating cache-based timing attacks with instruction-based scheduling. *In: Computer Security–ESORICS 2013, pp. 718–735. Springer*, 2013.

[41] Vattikonda, B.C., Das, S., Shacham, H.: Eliminating fine grained timers in xen. *In: Proceedings of the 3rd ACM workshop on Cloud computing security workshop, pp. 41–46. ACM*, 2011.

[42] Zhang, Y., Li, M., Bai, K., Yu, M., Zang, W.: Incentive compatible moving target defense against vm-colocation attacks in clouds. *In: Information Security and Privacy Research, pp. 388–399. Springer* , 2012.

[43] Godfrey, M., Zulkernine, M.: A server-side solution to cache-based side-channel attacks in the cloud. *In: Cloud Computing (CLOUD), 2013 IEEE Sixth International Conference on, pp. 163–170. IEEE*, 2013.

[44] Atici, A.C., Yilmaz, C., Savas, E.: An approach for isolating the sources of information leakage exploited in cache-based side-channel attacks. *In: Software Security and Reliability-Companion (SERE-C), 2013 IEEE 7th International Conference on, pp. 74–83. IEEE*, 2013.

[45] Ristenpart, T., Tromer, E., Shacham, H., Savage, S.: Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. *In: Proceedings of the 16th ACM conference on Computer and communications security, pp. 199–212. ACM*, 2009.

[46] Zhang, Y., Juels, A., Reiter, M.K., Ristenpart, T.: Cross-vm side channels and their use to extract private keys. *In: Proceedings of the 2012 ACM conference on Computer and communications security, pp. 305–316. ACM*, 2012.

[47] ElGamal, T.: A public key cryptosystem and a signature scheme based on discrete logarithms. *In: Advances in Cryptology, pp. 10–18. Springer*, 1984.

[48] Elgamal, T.: Side-channel attacks: Ten years after its publication and the impacts on cryptographic module security testing. *IEEE Transactions on Information Theory* ***31****, 469–472*, 1985.

[49] Ayala, I.D.C.L., Vega, M., Vargas-Lombardo, M.: Emerging threats, risk and attacks in distributed systems: Cloud computing. *In: Innovations and Advances in Computer, Information, Systems Sciences, and Engineering, pp. 37–51. Springer*, 2013.

[50] Varadharajan, V., Tupakula, U.: Counteracting security attacks in virtual machines in the cloud using property based attestation. *Journal of Network and Computer Applications* , 2013.

[51] Noor, T.H., Sheng, Q.Z., Alfazi, A.: Detecting occasional reputation attacks on cloud services. *In: Web Engineering, pp. 416–423. Springer*, 2013.

[52] Patel, A., Taghavi, M., Bakhtiyari, K., Celestino JúNior, J.: An intrusion detection and prevention system in cloud computing: A systematic review. *Journal of Network and Computer Applications* ***36****(1), 25–41* , 2013.

[53] Noor, T.H., Sheng, Q.Z., Zeadally, S., Yu, J.: Trust management of services in cloud environments: Obstacles and solutions. *ACM Computing Surveys (CSUR)* ***46****(1), 12*, 2013.

[54] Buecker, A., Lodewijkx, K., Moss, H., Skapinetz, K., Waidne, M.: Cloud security guidance. *IBM Red paper* ***2009****, 12*, 2009.

[55] Robinson, N., Valeri, L., Cave, J., Starkey, T., Graux, H., Creese, S., Hopkins, P.P.: The cloud: understanding the security, privacy and trust challenges. *Privacy and Trust Challenges (November 30, 2010)* , 2010.

[56] Wei, L., Zhu, H., Cao, Z., Dong, X., Jia, W., Chen, Y., Vasilakos, A.V.: Security and privacy for storage and computation in cloud computing. *Information Sciences* ***258****, 371–386*, 2014.

[57] He, X., Chomsiri, T., Nanda, P., Tan, Z.: Improving cloud network security using the tree-rule firewall. *Future Generation Computer Systems* ***30****, 116–126*, 2014.

[58] Koushik, S., Patil, A.P.: Open security system for cloud architecture. *In: ICT and Critical Infrastructure: Proceedings of the 48th Annual Convention of Computer Society of India-Vol I, pp. 467–471. Springer*, 2014.

[59] Winkler, V.J.: Securing the Cloud: Cloud computer Security techniques and tactics. *Elsevier*, 2011.

[60] Xiangyi, H., Zhanguo, M., Yu, L.: The research of the cloud security architecture. *In: Instrumentation, Measurement, Circuits and Systems, pp. 379–385. Springer*, 2012.