

## Chapter 2

# About Cloud Security

**Abstract** This chapter starts by discussing how cloud computing security is different from classical network security. The chapter mentions some threats and attacks that apply specifically to cloud computing. The chapter elaborates on most recent real-life attacks to cloud computing in the past few years. The chapter also explains the history of Denial of Service attacks along with other attacks.

**Keywords** Cloud computing · Iaas · Paas · Saas · Security · Cloud security

## 2.1 Introduction

As our lives become more and more connected, network security becomes more and more challenging. Security 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.

The context of network security becomes much more difficult to control when the environment becomes as dynamic and demanding as cloud computing.

The main aim of cloud computing is cost reduction and efficiency improvement. This cost reduction is not only in terms of computing resources, but also in terms of helping its users to focus on the business instead of the information technology enabling this business. Cloud computing is the result of developments in many technology directions such as virtualization, autonomic-computing, grid-computing, and many other technologies as explained earlier in Chap. 1.

As always, with every new technology, new challenges arise. A very important challenge is providing adequate security to that cloud to perform in alliance with business objectives.

After discussing the basics of cloud computing in Chap. 1, in this chapter we will focus on basic security aspects. At the start of our discussion, we must be familiar with three basic concepts: vulnerability, threats, and attack. In the Internet Engineering Task Force (IETF) RFC 2828 [1], a *vulnerability* is defined as a flaw or weakness in a system's design, implementation, or operation and management that could be exploited to violate the system's security policy. A *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 are available only to eligible entities and no unauthorized access to data can be obtained.
- Integrity: Assuring that data have 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.

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 enhance 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 [2].

While computer security is an important concept, network security remains a broader sense. Network security focuses on prevention of unauthorized access to data, software, and hardware at the network level, rather than on host level. It is the proactive detection of active attacks against the resources, the prevention of unnecessary security vulnerabilities, and rapid, appropriate response when a security event takes place. In general, network security consists of three layers: border security, authentication, and authorization. Border security focuses mostly on network layer security devices such as firewalls, intrusion detection systems, and intrusion prevention systems. Authentication helps in assuring identities of users as explained above, while *authorization* is identifying which resources an authenticated users can access, and which type of access this user is granted.

## 2.2 Why Is Cloud Security Different?

As with any other system, cloud computing includes vulnerabilities. These vulnerabilities, when exploited by attackers, can cause service disruptions, data loss, data theft, etc. Given the nature of dynamic resource sharing that take place in the cloud, it is possible that classical attacks and vulnerabilities can cause more harm on a cloud system if it is not protected properly.

The context in which network security can be discussed can identify a long list of threats and attacks. However, the dynamic and unique nature of the cloud can require additional measures and this nature also opens the door for a whole new list of attacks that can be used against the cloud.

Nothing explains this better than an example. One of the unique characteristics of the cloud is availability. The cloud is designed to be available all the time. Whether it is a private or a public cloud, availability is an undeniable feature that many organizations seek. What if attackers target availability of the cloud?

One of the major reasons why organizations decide to switch to a cloud environment is the you-pay-for-what-you-use business model. No one likes paying for resource that are not very well utilized. Hence, when an attack such as Denial-of-Service (DoS) attack happens, not only availability is targeted.

Denial of Service (DoS) attacks aim at making a certain network service unavailable to its legitimate users. In its basic form, these attacks keep the resources busy such that these resources become unavailable to the users this service was aimed to serve.

Using DoS attacks on the cloud, the attacker can cause huge financial implications by consuming high resources in the trial of making the service unavailable. So, for the organization using the cloud, it is a doubled loss.

The organization will be paying a lot of money for the resources consumed by the attack and, after a while, the organization's service will be unavailable due to the DoS attack. This type of attacks is referred to as Fraudulent Resource Consumption (FRC) [3].

The previous example shows us how the same attack can have different effect on different technology. For example, DoS attack on a classic server would render the service unavailable. If the same attack happens on a mobile ad hoc network, it would make the service unavailable and consume valuable battery life [4]. On the other hand, DoS on the cloud would render the service unavailable and cost the organization a lot of money for the consumed resources. This is why the uniqueness of the cloud technology open the door for unique attacks or at least unique effects of old common attacks.

Having the multiple layers discussed in Chap. 1, cloud computing can be target for attacks at any of these levels. We will see in the coming chapters that threats exist at virtually any lever of the cloud computing system. As you will see, there are threats at the hypervisor level, threats at the platform level, threats at the software level, etc. All of these attacks are unique to cloud computing alone and cannot be used on classical network security model.

Given the dynamic nature and the huge processing power of the cloud, it can also be used by attackers as a powerful attacking tool. The attacker can benefit from the on-demand processing power and employ it in performing DoS attacks among other attacking choices.

## 2.3 Famous Attacks on Cloud

Although cloud computing, in its current definition, is not that old, the cloud has got its good share of attacks.


### 2.3.1 History of Denial of Service Attacks on the Cloud

With rapidly increasing applications on the Internet, people rely more and more on Internet-based services in their regular daily actions. Availability of these services has grown to be one of the biggest concerns for both clients and service providers. During the past few years, many attacks have targeted availability of Internet-based services. As network security research grew stronger, simple DoS attacks were not as important as they were once. DoS attacks became less effective and easily detectable. Since the regular DoS attack comes from a single source, it becomes less effective once the source is detected by security appliances, or software, and blocked. A more sophisticated and harder to detect version evolved, namely Distributed Denial of Service (DDoS) attack.

A *DDoS attack* is an attack that targets availability of a system and is launched from multiple locations at the same time. The idea behind launching the attack from multiple location is to make detection much harder. Figure 2.1 shows how a DDoS attack works.

Figure 2.1 shows how an attacker implants malicious code in computers or servers that are not well protected. At the attack time, the attacker gives orders to the bots to

**Fig. 2.1** Anatomy of DDoS attack



start the attacks. Sometimes even giving orders is not necessary as malicious codes can work based on a specific action time preprogrammed. Usually, after the attack, the attacker cleans up all log files and other traces that can lead back to him/her. This would make it nearly impossible to trace back the source of the attack.

As explained earlier, in DDoS attacks, a DoS attack is launched from multiple sources, usually tens or hundreds, at the same time. Because it is coming from multiple source, DDoS attack is harder to detect and deter as compared to simple DoS attacks.

Reports show that only six DDoS attacks took place in the year 1988. The year 2000 witnessed DDoS attacks on large websites such as CNN, Yahoo, and Amazon. At that time, reports shown that DDoS attack rates reached 1 GBps.

In 2007, DDoS attacks reached the rates of 70 GBps. In 2013, a huge attack took place on Spamhaus spam detection service that reached the huge rate of 300 GBps. This attack aimed at bringing down Spamhaus so that the blacklisted sources of spam can go undetected on the Internet.

In February 2014, the largest DDoS attack in history took place with the rate of 400 GBps which is the largest known DDoS attack known until now.

The largest DDoS attack in history mentioned earlier 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. Regular Internet users experienced noticeable slowness in their Internet services around the world that day.

DDoS attacks can be launched in many different mechanisms. However, the purpose is the same, prevent the legitimate users from using the system. DDoS can be done through flooding, amplification, malformed packets, or exploiting a vulnerability in a networking protocol [5]. The mechanism of attack used against CloudFlare was Network Time Protocol (NTP) amplification.

NTP is a simple protocol that helps computers set their time accurately by contacting an NTP server. NTP protocol uses UDP at the transport layer, which means that it does not require any handshake, nor acknowledgment like TCP [6]. The attack uses a very old attacking technique called spoofing. *IP address spoofing* is sending an IP packet with a fake source IP address.

The attacker starts the NTP amplification attack using a rogue NTP server that is controlled by the attacker on a network that does not prevent spoofing. A large number of UDP segments with spoofed source IP address putting the target's IP address as the source of all of those packets. The spoofed packets are directed to a large number of NTP servers (on standard NTP port 123). These requests are not sent using simple NTP commands rather using MONLIST command. MONLIST command, in NTP, results in a response around 206 times the request size. The command MONLIST sends a list of up to the last 600 IP addresses that last accessed the NTP server [7]. In theory, an attacker with 1 Gbps link can generate more than 200 Gbps of DDoS traffic.

According to details published by CloudFlare in [8], the attacker(s) used 4,529 NTP servers running on 1,298 different networks. On average, each server sent around 87 Mbps of DDoS traffic to CloudFlare's network. Again, in theory, the attacker could

have used a single attacking server to start the attack. All the attacker needed was a network that does not prevent spoofing to initiate the requests.

NTP amplification seem to be much more dangerous as compared to DNS amplification attacks the were used to attack Spamhaus in 2013. The attacker needed about 1/7th the number of compromised servers that were used in the Spamhaus attack, yet produced 33 % more traffic.

The dynamic nature of cloud computing can be beneficial to counter DDoS attacks in some scenarios. However, the different levels at which DDoS attacks can be performed on a cloud-based service can make defense more complex, as we will see in the next chapters.

### **2.3.2 Other Attacks**

Amazon provides a user interface for its clients through which the clients can start new instances of a machine, terminating an instance among other control actions. In 2008, a vulnerability was found in this control service. The vulnerability was about using a special type of Signature Wrapping Attack, the attacker would be able to modify an eavesdropped message despite the digital signature of the message. Through this modification, the attacker would be able to execute any code they want on the server machine. This attack can be split into two separate activities; attacking the cloud control interface to get control of the cloud system, then attacking the service instances using the service-to-cloud attack surface [9]. In 2011, the year of cloud data breaches, the website TripAdvisor, which was hosted on the cloud, was compromised and users e-mail addresses were stolen. The attack is thought to be a possible SQL-injection type of attack [10].

In 2015, a security vulnerability named venom was detected in many cloud-based data centers. This vulnerability can allow hackers to take over entire data centers, presumably even those owned by Amazon, RackSpace, and Oracle. Venom stands for Virtualized Environment Neglected Operations Manipulation.

Another massive breach that took place in 2011 was Sony attacks. Sony had around a dozen data breaches in 2011 that hit its Sony PlayStation Network, Sony Pictures, Sony Online Entertainment, along a few other Sony-owned websites. The attacks compromised around 100 million user accounts. The compromise included password and other private data. Given the frequency of password reuse by users, there is a very good probability that user account in other online services can be targeted using the obtained information from Sony breaches. In this attack, attackers gained access to one server that was not well protected and started escalating from there to obtain access to many other servers. Sony's network was not layered well enough to isolate breaches happening in one part of the network to proceed to the rest of the network. The clear weakness in Sony's security was evident when the attackers published server certificates that used the password "password." These certificates were later used to distribute malware. Many weaknesses contributed to the Sony breach including the use of weak passwords, lack of individual server hardening,

lack of proper network isolation, not having proper security controls to set off alerts, not responding to alerts properly, in adequate logging and monitoring, and lack of security awareness [11, 12].

In 2011 as well, attackers targeted a cloud-based Nasdaq system named “Directors Desk.” The system facilitates communication between 10,000 senior executives and company directors. By having access to this system, attacks can eavesdrop on private conversations between executives that can be used as stock market-leaked information to benefit competitors. While attackers had not directly attacked trading servers, they were able to install malware on sensitive systems, which enabled them to spy on dozens of company directors. This kind of attack has a long-term effect that cannot be easily calculated now [13].

Another attack that happened in 2011 was the Epsilon attack. Epsilon is a cloud-based e-mail service provider that went under attack in April, 2011. The attack was a spear phishing attack. While a *phishing* attack is a type of social engineering attack in which attackers use spoofed e-mail messages to trick victims into sharing sensitive information or installing malware on their computers, spear phishing is a more complex type of phishing. Spear phishing attacks use specific knowledge of individuals and their organizations to make the spoofed e-mail look more legitimate and targeted to a specific person [14]. In the attack on Epsilon, data of 75 business organizations were beached and the list was growing. Although Epsilon did not disclose the names of companies affected by the attack, it is estimated that around 60 million customer e-mails were breached.

A few other cloud-based breaches happened in 2011. This have slowed down the pace of adoption of the cloud by businesses in the next year. However, cloud adoption recovered from these breaches because most of these breaches were not cloud specific and could have happened with the classic server model.

In 2012, another cloud-related breach took place at the Institute of Electrical and Electronics Engineers (IEEE). The breach was discovered by an independent security researcher who was searching for some literature on IEEE File Transfer Protocol (FTP) servers [15]. The researcher found a cache of 100,000 usernames and passwords in plaintext just sitting there waiting to be grabbed. The credentials included those of members working in Google, Apple, NASA, Stanford University, among others. In addition to login credentials, the researcher was able to access more than 100 GB of web server log data containing detailed information on 350 million-plus HTTP requests made by IEEE members over one month. The FTP server this researcher was surfing was a public one. The researcher said that these files were easily accessible on the public servers for over a month. An institution as large and known as IEEE should not have fallen in these naive security errors.

According to the cloud security report, issued by Alert Logic, attacks on cloud computing became almost equal in number to attacks on classic computer systems in 2014 [16]. The report says that brute-force attacks have increase 30 % of all cloud attacks to 44 %. A brute-force attack is an attack that involves large number of trials of credentials to access a certain resource. At the same time, malware attacks on clouds have climbed from 5 % to 11 %.

The *venom* vulnerability can be identified as vulnerability in the code of the virtual floppy that can be exploited to access sensitive and personal data stored on the cloud. Security researchers said that systems run be Microsoft Hyper-V, VMware, and Bochs were not affected by the vulnerability.

Although most organizations affected by this vulnerability said that this vulnerability was never used by any attacker, this vulnerability sends a clear message that cloud security is not perfect. Actually, it is far from perfect. Just like any other component in the field of information security [17].

In 2015, a security services company named Elasticity revealed that Salesforce.com servers were vulnerable to a Cross-Site Scripting (XSS) attack. This vulnerability, although it was patched two days before the announcement of its existence, would enable the attacker to run JavaScript to steal the cookies that contain session credentials and hijack the session. The vulnerability is also said to enable the attacker to force Salesforce.com client to phishing websites that can enable the attacker to get credential information through social engineering. The attacker would also be able to force the client to download and run malicious code from the infected application [18].

Later in 2015, a vulnerability was discovered in Amazon cloud storage platform by a group of scientists [19]. The study revealed that a sophisticated CPU cache attack against an Amazon EC2 instance could have given a hacker complete access to a 2048-bit RSA key used in a separate instance hosted at the same physical server. The attacks starts by co-location identification and verification. Then, the attacker would perform a sort of cache attack to detect keys used in RSA encryption. Luckily, this implementation vulnerability was detected on in Libgcrypt RSA implementation libraries.

Another vulnerability, called DROWN, was recently discovered in 2016. This vulnerability was discovered in OpenSSL by a group of 15 scientists [20]. OpenSSL is an open-source implementation of Secure Socket Layer (SSL) security protocol. According to the report, attacker could break HTTPS traffic by leveraging an older attack method from 1998 against SSLv2, even if the server's traffic was protected with newer and more secure Transport Layer Security (TLS) certificates. The report says that around one-third of all websites using HTTPS. When the study was first published in March 1, 2016, Skyhigh Networks, a cloud security company, explains that, during its scans, it detected 653 cloud service providers susceptible to DROWN attacks. Seven days later, 620 out of the 653 cloud services were still unpatched. This shows an example of the lack of timely response by some cloud service providers operating in the market.

## References

1. R. Shirey, Rfc 2828: Internet security glossary, in *The Internet Society* (2000), p. 13
2. W. Stallings, *Cryptography and Network Security, 4/E* (Pearson Education, Upper Saddle River, 2006)



3. J. Idziorrek, M. Tannian, D. Jacobson, Attribution of fraudulent resource consumption in the cloud, in *Proceedings of the IEEE 5th International Conference on Cloud Computing (CLOUD)*, 2012, pp. 99–106
4. M.M. Alani, Manet security: a survey, in *Proceedings of the IEEE International Conference on Control System, Computing and Engineering (ICCSCE)*, 2014, pp. 559–564
5. R.V. Deshmukh, K.K. Devadkar, Understanding ddos attack & its effect in cloud environment. *Procedia Comput. Sci.* **49**, 202–210 (2015)
6. M.M. Alani, *Guide to OSI and TCP/IP models* (Springer, Berlin, 2014)
7. J. Graham-Cumming, Understanding and mitigating ntp-based ddos attacks, vol. 9 (Cloudflare Inc, California, 2014)
8. M. Prince, Technical details behind a 400gbps ntp amplification ddos attack, vol. 13 (Cloudflare Inc, California, 2014)
9. N. Gruschka, M. Jensen, Attack surfaces: a taxonomy for attacks on cloud services, in *Proceedings of the IEEE 3rd International Conference on Cloud Computing*, 2010, pp. 276–279
10. Tripadvisor: E-mail addresses stolen in data breach, <http://www.cnet.com/news/tripadvisor-e-mail-addresses-stolen-in-data-breach/>. Accessed 27 March 2016
11. 6 worst data breaches of 2011, <http://www.darkreading.com/attacks-and-breaches/6-worst-data-breaches-of-2011/d/d-id/1102001?> Accessed 29 March 2016
12. The sony hack what happened, how did it happen.what did we learn? <http://blogs.umb.edu/itnews/2015/01/06/the-sony-hack/>. Accessed 29 March 2016
13. Nasdaq server breach: 3 expected findings, <http://www.darkreading.com/attacks-and-breaches/nasdaq-server-breach-3-expected-findings/d/d-id/1100934?> Accessed 29 March 2016
14. J. Hong, The state of phishing attacks. *Commun. ACM* **55**(1), 74–81 (2012)
15. Data breach at ieee.org: 100k plaintext passwords, <http://ieeelog.dragusin.ro/init/default/log>. Accessed 29 March 2016
16. A. Logic, Cloud security report-spring 2014, 2014
17. J.-M. Brook, R. Brooks, A decade of lessons learned: Transforming the enterprise for todays cloud architecture, in *Proceedings of the ICCSM2015 3rd International Conference on Cloud Security and Management: ICCSM2015*, Academic Conferences and publishing limited, 2015, p. 16
18. Salesforce accounts susceptible to hijacking using xss flaw, <https://www.elastica.net/salesforce-accounts-susceptible-to-hijacking-using-xss-flaw>. Accessed 21 March 2016
19. M.S. Inci, B. Gulmezoglu, G. Irazoqui, T. Eisenbarth, B. Sunar, Seriously, get off my cloud! cross-vm rsa key recovery in a public cloud (Technical report, IACR Cryptology ePrint Archive, 2015)
20. N. Aviram, S. Schinzel, J. Somorovsky, N. Heninger, M. Dankel, J. Steube, L. Valenta, D. Adrian, J.A. Halderman, V. Dukhovni et al., Drown: Breaking tls using sslv2



<http://www.springer.com/978-3-319-41410-2>

Elements of Cloud Computing Security

A Survey of Key Practicalities

Alani, M.M.

2016, XI, 55 p. 7 illus., Softcover

ISBN: 978-3-319-41410-2