

Survey Paper on Alleviation of Cloud Internal Denial of Service Attacks

Nikhita Nerkar¹, Vina M Lomte²

Assistant Professor, Computer Department RMDSSOE, Warje, Pune, India

Abstract: Cloud computing security is now becoming one of the main concerns preventing the adoption of the cloud by many organizations. This paper surveys the attacks done on cloud and their mitigation strategies to defend the cloud specific CIDoS class of attacks (Cloud-Internal Denial of Service). The alleviation approaches are based on techniques used in signals processing field. The main strategy to detect the attack is the calculation of correlations measurement and distances between attackers workload patters; we use DCT (Discrete Cosine Transform) to accomplish this task. This paper also suggests some prevention and response strategies.

Keywords: Cloud DoS attacks, Cloud Attack Mitigation, IaaS Cloud Security, CIDoS attack detection, Cloud Computing Security

1. Introduction

Cloud computing is a general term for anything that involves delivering hosted services over the Internet. Cloud computing is shown in Fig. 3.1. These services are broadly divided into three categories: Infrastructure-as-a-Service (IaaS), Platform-as-a-Service (PaaS) and Software-as-a-Service (SaaS). The name cloud computing was inspired by the cloud symbol that's often used to represent the Internet in flow charts and diagrams [2]. The development in computer and communication technologies has led to a new computing paradigm called Cloud computing, which delivers computing services to users as utilities in a pay-as-you-go manner. A cloud service has three distinct characteristics that differentiate it from traditional hosting. It is sold on demand, typically by the minute or the hour; it is elastic a user can have as much or as little of a service as they want at any given time; and the service is fully managed by the provider (the consumer needs nothing but a personal computer and Internet access). Significant innovations in virtualization and distributed computing, as well as improved access to high-speed Internet and a weak economy, have accelerated interest in cloud computing. A cloud can be private or public. A public cloud sells services to anyone on the Internet. (Currently, Amazon Web Services is the largest public cloud provider.) A private cloud is a proprietary network or a data center that supplies hosted services to a limited number of people. When a service provider uses public cloud resources to create their private cloud, the result is called a virtual private cloud. Private or public, the goal of cloud computing is to provide easy, scalable access to computing resources and IT services. Service providers make use of IaaS and PaaS to deploy their services without concerns about managing physical resources. Under the Cloud computing model, users can access on-demand and pay-per-use services anywhere in the world. The pay-as-you-go mechanism in Cloud computing assures Service Level Agreements (SLAs) between customers and Cloud providers. SLAs specify the negotiated agreements on the Quality of Service (QoS), such as deadline constraints. Cloud computing according to definition as shown in Fig. 1 refers to applications and services that run on a distributed network using virtualized resources and accessed by common Internet protocols and networking standards. It is distinguished by the notion that

resources are virtual and limitless and that details of the physical systems on which software runs are abstracted from the user. In an effort to better describe cloud computing, a number of cloud types have been defined.




Figure 1: Cloud computing




Figure 2: Layers of Cloud Computing

2. Network Issues in Cloud Computing

There are number of Network Issues in Cloud Computing:

1) Denial Of Service

Denial of Service attacks has existed since the Internet first started. It's a malicious attempt with a single person or group to cause a site, victim, or node to deny services to some customer.

2) Man in Middle Attack

A man-in-the-middle (MITM) attack is a form of eavesdropping where communication between two users is monitored and modified by an unauthorized party. The MITM intercepts communications between two systems and is performed when the attacker is in control of a router along normal point of traffic.

3) Network Sniffing

It's this idea of actually watching what's happening on the network and having some device, whether it's a computer or a hardware-based specifically designed device, doing the network watching. You're paying attention to all the data, packets, and traffic going across the network and capturing all of that data, and traffic that's going around, and using it for statistical analysis.

4) Port Scanning

The act of systematically scanning a computer's ports. Since a port is a place where information goes into and out of a computer, port scanning identifies open doors to a computer.

5) SQL Injection attack

An SQL injection is a computer attack in which malicious code is embedded in a poorly-designed application and then passed to the backend database. The malicious data then produces database query results or actions that should never have been executed.

6) Cross Site Scripting

Cross-site scripting vulnerabilities arise when data is copied from a request and echoed into the application's immediate response in an unsafe way. An attacker can use the vulnerability to construct a request which, if issued by another application user, will cause JavaScript code supplied by the attacker to execute within the user's browser in the context of that user's session with the application.

3. Related Work

Techniques of Secure Live Migration of Virtual Machine [2] live migration are an essential feature of virtualization that allows transfer of virtual machine from one physical server to another without interrupting the services running in virtual machine. Live migration facilitates workload balancing, fault tolerance, online system maintenance, consolidation of virtual machines etc. Unfortunately the disclosed vulnerabilities with the live migration pose significant security risks. Because of these security risks the industry is hesitant to adapt the technology for sensitive applications. This paper is an investigation of attacks on live migration of virtual machine and discusses the key proposed and implemented approaches to secure live migration.

Hey, You, Get Off of My Cloud: Exploring Information Leakage in Third-Party Compute Clouds [3] Third-party cloud computing represents the promise of out-sourcing as

applied to computation. Services, such as Microsoft's Azure and Amazon's EC2, allow users to instantiate virtual machines (VMs) on demand and thus purchase precisely the capacity they require when they require it. In turn, the use of virtualization allows third-party cloud providers to maximize the utilization of their sunk capital costs by multiplexing many customer VMs across a shared physical infrastructure. However, in this paper, we show that this approach can also introduce new vulnerabilities. Using the Amazon EC2 service as a case study, we show that it is possible to map the internal cloud infrastructure, identify where a particular target VM is likely to reside, and then instantiate new VMs until one is placed co-resident with the target. We explore how such placement can then be used to mount cross-VM side-channel attacks to extract information from a target VM on the same machine.

Robust Group Key Agreement Using Short Broadcasts [4] A group key agreement protocol (GKA) allows a set of players to establish a shared secret key which can be used to secure a subsequent communication. Several efficient constant-round GKAs have been proposed. However, their performance degrades if some players fail during protocol execution. This is a problem in practice, e.g. for mobile nodes communicating over wireless media, which can lose connectivity during the protocol execution. Current constant-round GKA protocols are either efficient and non-robust or robust but not efficient: Assuming a reliable broadcast communication medium, the standard encryption-based group key agreement protocol can be robust against arbitrary number of node faults, but the size of the messages broadcast by every player is proportional to the number of players. In contrast, non-robust group key agreement can be achieved with each player broadcasting just constant-sized messages. We propose a novel 2-round group key agreement protocol which tolerates up to T node failures using $O(T)$ -sized messages, for any T . To exemplify the usefulness of this flexible tradeoff between message size and fault tolerance, we show that the new protocol implies a fully-robust group key agreement with $O(\log n)$ -sized messages and expected round complexity close to 2, assuming random node faults. The proposed protocol is secure under the (standard) Decisional Square Diffie-Hellman assumption.

Secure Cloud Computing Environment against DDOS and EDOS Attacks [5] Cloud computing is becoming one of the fastest growing field in the information technology. Cloud computing allows us to scale our servers in magnitude and availability in order to provide service to greater number of end users. Moreover, cloud service model are charged based on a pay-per-use basis of the cloud's server and network resource. In cloud computing where infrastructure is shared by potentially millions of users, Distributed Denial of Service (DDoS) attacks have the potential to have much greater impact than against single tenanted architectures. With this model, a conventional DDoS attack on server and network resources is transformed in a cloud environment to a new breed of attack that targets the cloud user's economic resource, namely Economic Denial of Service attacks. In this paper, we propose a novel solution, named DDoS and EDoS- Shield, to avoid the Denial of service and Economic Denial of Sustainability (EDoS) attack in the cloud computing systems.

4. DOS attack

4.1 What is DOS Attack?

- A denial of service (DoS) attack is an incident in which a user or organization is deprived of the services of a resource they would normally expect to have.
- In a distributed denial-of-service, large numbers of compromised systems attack a single target.
- A DoS attack does not usually result in the theft of information or other security loss.
- Typically, the loss of service is the inability of a particular network service, such as e-mail, to be available or the temporary loss of all network connectivity and services.
- A denial of service attack can also destroy programming and files in affected computer systems.

4.2 DOS Attack Example

- Recently the website Ebay was a victim of DOS Attack. By just one targeted DoS attack millions of people were taken offline. Just imagine the amount of damage it must have done to the business.
- Recently, a massive DoS attack took place at CloudFlare, with hackers exploiting the vulnerabilities in the Network Time Protocol Servers (NTP).

5. Cloud-Internal Denial of Service attack

A class of DoS attacks that works in IaaS clouds. It is based on misusing two of the main features of the cloud, which are 'migration' and 'over-commitment'. The idea of the attack is coordination, between group of attackers, to increase the consumption of cloud resources by building on the existed wave of consumption (riding the workload wave). This increase might trigger the very expensive migration process for false reasons which might have negative effect on the service. The scale of this attack has the potential of causing a complete paralysis of the cloud instance. The attack is called 'Cloud-Internal Denial of Service attack', CIDoS, and requires performing the following steps. First, it requires a number of malicious VMs (over a specific threshold) in the same physical host (called co-residency). Second, the attack class relies on the use of covert channels for communication and coordination of the attack between participating VMs. Covert channels are communication channels which are not designed to carry information, as such, and typically beyond authorisation of access control mediums. In covert channels, flaws in isolation between resources are exploited to enable communication. The severity of the covert channel is measured by channel capacity, eliminating all covert channels is extremely difficult. Third, the attack class operates a protocol through the covert channel. When attackers are ready to attack (by the end of a successful protocol operation), each of the participating VMs will increase their utilisation following the existed workload pattern of the host and building over it in an attempt to break the host ability to cope with the increased stress (a technique similar to jamming signal). The attack is hard to detect because the behaviour of attackers looks similar to normal workload. The attack deceives the cloud by falsely showing the host very busy, which might trigger severe consequences.

There are three parties; the attackers (us), other attackers (not us), the defenders (security systems and administrators in the cloud). There are several motivations for CIDoS attacks. First, attackers can be competitors of some of the cloud consumers. In this scenario, attackers try to affect the cloud service to make the targeted consumer fail to use its virtual machines thus losing money or reputation. Several services are time-critical and rely on quality of service, so even delayed response will be make a difference to such customers (i.e. retail or financial services). In this case the attacker can establish the attack at critical peak times in the competitor's cloud provider, which may affect the service, the victim may fail to satisfy customers' requests and hence lose business. Other motivation is when the attacker is a competitor cloud provider. He or she might try to affect the service of the targeted cloud provider and force a breach of service level agreements (each year cloud providers announce how many minutes of downtime they have registered, and it affects their reliability heavily thus reducing revenue). The attack class described here is a public IaaS cloud specific DoS attack; it is targeting availability in the cloud and thus reliability. The attacker need to:

- 1) Increase probability of co-residency.
- 2) Determine if instances are co-resident on the same 'physical' machine
- 3) Coordinate the attack and communicate using covert channels
- 4) Coordinate and trigger the attack

6. Conclusion

As Malicious VM is used to attack a cloud host, we can use Discrete Wavelet Transform for early detection of the attack. We considered here the scenario where the attackers had leader, Also other scenarios can be considered, consensus agreement where there is no leader.

References

- [1] S. Alarifi and S. D. Wolthusen, "Robust coordination of cloud-internal denial of service attacks," in Cloud and Green Computing (CGC), 2013 Third International Conference on, 2013, pp. 135-142.
- [2] A. Mohan and S. S., "Survey on live vm migration techniques," International Journal of Advanced Research in Computer Engineering and Technology, vol. 2, 2013.
- [3] T. Ristenpart, E. Tromer, H. Shacham, and S. Savage, "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, ser. CCS '09. New York, NY, USA: ACM, 2009, pp. 199-212. [Online]. Available: <http://doi.acm.org/10.1145/1653662.1653687>
- [4] S. Jarecki, J. Kim, and G. Tsudik, "Robust Group Key Agreement using Short Broadcasts," in Proceedings of the 14th ACM Conference on Computer and Communications Security (CCS 2007), S. De Capitani di Vimercati and P. Syverson, Eds. Alexandria, VA, USA: ACM Press, Oct. 2007, pp. 411-420.
- [5] P. Chimento, "IP Packet Delay Variation Metric for IP Performance Metrics (IPPM)," 2002. [Online]. Available: <http://tools.ietf.org/html/rfc3393>
- [6] F. Zhou, M. Goel, P. Desnoyers, and R. Sundaram, "Scheduler vulnerabilities and attacks in cloud computing," CoRR, vol. abs/1103.0759, 2011.