




SECURED CLOUD COMPUTING

M Gnanendra Reddy
Global Technical Support
EMC

EMC²

Table of Contents

Acknowledgement	4
Abstract.....	5
Introduction	5
Critical Security threats to Cloud Security	6
Security issues with Multi-tenant Cloud environments	7
Overview of security attacks in Multi-Tenant cloud computing	8
DDoS Attacks.....	8
Real-Time examples of DDoS attacks	10
Man-in-the-Middle Attacks.....	11
Real-Time examples of Man-in-Middle attacks	13
Determining Co-Residence in Multi-tenant environment.....	14
Simulation of Security attacks in test environment.....	15
List of Virtual Machines.....	15
Architecture overview.....	15
Security Attack tools used in this research.....	16
Actual Implementation	17
Step 1: Probing the network	17
Step 2: Probing the ports on Web Server	18
Step 3: Exploiting Web Server through DDoS attack	19
Step 4: Hijacking the communication through MITM attack	21
Step 5: Monitoring and notification of DDoS attack.....	24
Proposed Security measures	26
DDoS	26
MIMT.....	30
Conclusion	31
Recommendations to secure multi-tenant environments	31

List of Figures

Figure 1 : ESxi Server – Multitenant architecture overview	6
Figure 2: Overview of DDoS attacks	9
Figure 4 : ARP poisoning overview	12
Figure 5: Secure communication using HTTPS	12
Figure 6: Hijacked Communication overview – Man-in-the-Middle attack	13
Figure 7: Project architecture overview	16
Figure 8: Low-Orbit-Ion-Cannon – DDoS attack tool	17
Figure 9: Burp Suite – Security testing tool	17
Figure 10: Advanced IP scanner	18
Figure 11: Nmap port scanner	19
Figure 12: DDoS attack on Web Server using Low-Orbit-Ion-Cannon	20
Figure 13: CPU usage of Web Server under DDoS attack	20
Figure 14: Memory usage of Web Server under DDoS attack	21
Figure 3: ARP poisoning tool overview through Nighthawk	22
Figure 15: Request from End User	23
Figure 16: Actual response from Web Server	23
Figure 17: Modified response sent to End User	24
Figure 18: Monitoring Web Server usage and notification through Email/SNMP	24
Figure 19: Monitoring Web Site traffic through Smarter Stats tool	25
Figure 20: Overview of DMZ in a secured environment	27

Disclaimer: The views, processes, or methodologies published in this article are those of the author. They do not necessarily reflect EMC Corporation's views, processes, or methodologies.

Acknowledgement

The data presented and analyzed in this article is part of my academic research - **MSc in Cloud Computing**. The testing and implementation was performed in an EMC lab environment based on the real time scenarios. I used most of the freeware tools to carry out this research. I am grateful for the generous support extended by my professor - Donna 'O' Shea - **Cork Institute of Technology**.

I would like to thank EMC² for providing the infrastructure to carry out my research.

Abstract

Cloud computing enables resources to be dynamically provisioned over the Internet to offer multiple economic advantages and reduce the capital and operational expenditure that an organization can incur. It is well recognized that to fully realize the economic benefits of cloud computing, security and trust issues must be addressed. These new cloud platforms open additional exposure to threats against data, reputation, and systems. These threats arise through attacks such as data-stealing, malware, web attacks, phishing, spam, Trojan, viruses, worms, bots, etc. Virtualization and multi-tenant cloud computing raise new infrastructure issues that security providers must consider when securing the multi-tenant cloud infrastructure against these threats. [1 2]

This article discusses security threats that organizations face while deploying and using multi-tenant cloud computing infrastructures. Provided are real-world examples of DDoS attacks, Man-in-Middle attacks, and attack tools that cyber criminals use to exploit vulnerabilities in multi-tenant cloud environments, as well as recommendations for best security practices.

Introduction

Multi-tenancy is an important part of leading cloud computing offerings. It enables cloud providers to leverage cost effectiveness to be shared and pool resources to be built on a single code base. [3] This model introduces security threats because the same hardware is shared with multiple client machines. In a multi-tenant environment, different organizations may own different Virtual Machines (VM) but these are all located on the same underlying physical infrastructure with only a virtual separation between the organizations. Data and information is one of the largest strategic assets owned by any company; if this data cannot be secured, multi-tenancy will not be a viable solution. It is very important to analyze the security issues with multitenant environments, set up proactive monitoring, and follow best practices to control security attacks. [3 4]

In a multi-tenant environment, communication between the VM's on the same host are not monitored.[4] As shown in Figure 1, all VM's are hosted on a single ESXi server. For example, VM2 can talk to VM3 as the inter-VM traffic will not be monitored and restricted by default.




Figure 1: ESXi Server – Multitenant architecture overview

Before choosing the cloud technology it is important to understand the risks involved in moving to cloud. Risk assessment should be carried out before handing over control to the cloud provider. A major impediment to cloud computing is calculating added risk from all known and unknown sources. Understanding the amount of risk an organization can tolerate depends on assessing the security requirements and how information assets such as data, applications, and processes are valued. The most important risks in multi-tenant cloud environment are: [8, 15]

- Information Assets and Risk
- Privacy and Confidentiality Concerns
- Data Governance
- Customer's security expectations
- Malicious insider

Multi-tenancy makes use of virtualization technologies to increase load balancing, resource utilization, scalability, and reliability. This allows cloud service providers to maximize use of their infrastructure by multiplexing their physical machines with virtualization and then assigning the VM's to different clients when required. This will lead to different users using the same infrastructure. There is a possibility that an attacker can rent one of the VMs and pose several potential threats to the infrastructure. [10]

Critical Security threats to Cloud Security

Multi-tenancy is a technology where more than one organization's applications and data are hosted on the same infrastructure. [11] Multi-tenant technology is widely used in cloud computing due to its economic benefits. Per CSA guidelines, the top security threats to cloud security are: [1]

Insecure API's and Interface Cloud providers will provide a set of software interfaces or API's to customers for managing and interacting with cloud services. Security on the cloud services depends on these API's and should be protected from malicious attempts. [1]

Malicious Insiders This is a well-known threat to most organizations. The roles and responsibilities should be clearly defined for cloud management users. There should be a process that defines strong authentication and authorization mechanisms. Transparency is required for overall information security, management practices, and compliance reporting. [1]

Data Loss/Leakage Data can be compromised at many stages. For example, deletion or alteration of data without records or backup is considered as data loss. Loss of encoding key will result in effective destruction. The most important aspect is to prevent unauthorized users to gain access to sensitive information. [1]

Account, Service, and Traffic Hijacking User credentials are often reused which increases the impact of these attacks. If an intruder gains access to user credentials they can monitor user activities, transactions, access sensitive data and return false information, and redirect users to different sites. Strong two-level authentication techniques need to be implemented wherever possible. Additionally, proactive monitoring is required to detect unauthorized activity. [1]

Abuse and Nefarious Use of Cloud Computing A few IaaS providers offer customers unlimited network, compute, and storage capacity where any user can register with a valid credit/debit card and use the cloud services immediately. Some cloud providers offer free limited trial packages. Spammers, malicious code authors, and other cyber criminals are able to hack into the systems using such trial packages and perform password and key cracking, Distributed Denial of Service (DDoS) attacks, host malicious data, build rainbow tables, CAPTCHA codes, etc. [1]

Unknown Risk Profile Organizations using cloud computing are less involved in hardware and software ownership. Still, organizations should be aware of security practices, security compliance and configuration hardening, patching, auditing and logging. Security should always be at the top of the priority list. [1]

Shared Technology Issues: Cloud technology is based on shared infrastructure. Flaws may enable guest operating systems to gain unauthorized access. To ensure that customers do not tread on another's territory, a strong isolation mechanism and monitoring is required. [1]

Security issues with Multi-tenant Cloud environments

Multi-tenant architecture allows multiple VMs to run on a single host (ESXi). In multi-tenant environments, hackers and the attacked machine can be present in the same location. This provides the attacker with unparalleled access to other VMs. The dependency of all VMs on a single hypervisor will bring new security challenges. With multi-tenant technology gaining more popularity,

the attack vector on the ESX host is increasing considerably. Hackers use DDoS and MIMT attacks to bring down a service or steal user data. Physical separation and hardware-based security cannot stop these types of attacks between VMs on the same ESXi server. [12]

Overview of security attacks in Multi-Tenant cloud computing

DDoS Attacks

A DDoS attack is an attempt to make a resource unavailable to its end users. This type of attack typically targets the services hosted on high profile web servers, i.e. credit card payment gateways, bank sites, etc. In a DDoS attack, incoming traffic floods the victim from many different sources, potentially hundreds, thousands, or even more. [13 14] This effectively makes it impossible to stop the DDOS attack simply by restricting a single IP address. As well, it is very difficult to distinguish normal user traffic from attack traffic when spread across so many points of origin. With DDoS attacks, hackers target bandwidth, processing power, and storage capacities of a cloud network which brings down the hosted web services on the target host. DDoS attacks typically target one of the following:

- Exhaust bandwidth, disk space, processor usage and memory usage
- Disruptions of services hosted on a high profile Web Server
- Disrupting communication media between end user and victim, so that they can no longer communicate with each other

Launching a DDoS attack from multiple sources will bombard the target with multiple requests, overloading the target. It is difficult to track the DDoS attack as it will be launched from multiple locations. DDoS attacks are quite scary; they are easy to execute and difficult to find the source. [14]




Figure 2: Overview of DDoS attacks

Attack Name	Description
TCP SYN flood attack	The attacker sends a request to the victim using packets with the source address that is unreachable for the victim. The victim server will try to respond to the connection request and use its resources to process the request, but the connection never completes.
Smurf IP attack	The attacker sends forged ICMP packets to broadcast addresses of vulnerable networks. All the systems on the network send a reply, exhausting the available bandwidth available in the network.
UDP flood attack	The attacker sends a UDP packet to the victim on a random port or a range of ports. When the attacker receives the reply, it will determine which application is listening on the victim's destination port When the attacker realizes that there is no application waiting on the random port, it will generate an ICMP destination unreachable request and send it to the fake source address.
Ping of Death	The attacker sends an ICMP packet to the victim that is larger than the size of normal packet. When the victim receives the packet, it will attempt to re-assemble it. Since the packet size is too large than the actual size, the victim will not be able to process the request, thus causing the victim to crash or reboot.
Teardrop	The attacker sends two fragments of packets that cannot be re-assembled properly by manipulating the offset value of the packet, thus the victim will crash or reboot.
Land	The attacker sends a forged packet to the victim using the victim's address as both the destination and source IP addresses. The victim will be confused and will crash or reboot.

Table 1: Common DDoS attack types: [14]

Real-Time examples of DDoS attacks

1. Hackers hit the Nasdaq Stock Exchange with a DDoS attack. As a result, NASDAQ was down for three hours. An Iranian hacking group—Cyber Fighters—claimed credit for orchestrating sophisticated attacks that have overwhelmed the expensive security systems U.S. banks have put in place to keep their online banking services up and secure. Complete details about this attack can be found in [hacker news bulletin](#).

2. Attackers hit one American bank after another. As a result of so many attacks, dozens of online banking sites slowed down, hiccupped, or ground to a halt before recovering several minutes later. These clouds are run mostly by Amazon and Google. Many smaller players have also started renting cloud services to other companies. It appears the hackers remotely hijacked some of these clouds and used the computing power to take down American banking sites. Complete details about this attack can be found in [NY Times report](#).

3. In the US, six major American banks – Wells Fargo, Bank of America, JP Morgan Chase, US Bank, CITI Group, and PNC were hit by a wave of computer attacks in 2012. A group claimed responsibility that has ties to the Middle East. At the time of the attack, users were unable to access their accounts, pay bills online, and became upset because the banks did not provide information on what was happening. The banks were under DDoS attacks, in which the hackers generated huge traffic that hit the web site until it was flooded and shut down. Complete details about this attack can be found in [NY Times Report](#).

Man-in-the-Middle Attacks

With Man-in-the-Middle (MITM), attackers intrude into an existing communication and inject false information by eavesdropping, intruding into a connection, intercepting messages, and selectively modifying data. [16]

The MITM attack uses [ARP spoofing](#) to trick System A into thinking that it is communicating with System B and System B into thinking that it is communicating with System A. The attack initially starts with sniffing and eavesdropping on a network stream, and ends with trying to alter, forge, or reroute the intercepted data.

There are 2 types of MITM attack; Active and Passive. [17] In an active attack, the communication is entirely controlled by the attacker. In a passive attack, the attacker can view the user's traffic and steal the most valuable user data and session cookies.

Along with ARP spoofing, a MITM attack uses DNS spoofing and SSL strip to hijack the communication. There are many tools for ARP spoofing, which can inject false information about the user MAC address and can poison the ARP. The attacker sends an ARP packet to update the MAC address of the attacker host as default gateway. Consequently, all traffic will be redirected to the attacker machine. [16]

```
Administrator: Command Prompt
C:\Users\Administrator>
C:\Users\Administrator>arp -a

Interface: 10.31.137.230 --- 0xb
Internet Address      Physical Address      Type
10.31.137.149         00-50-56-a3-2a-26    dynamic
10.31.138.110         00-50-56-a3-2a-26    dynamic
224.0.0.252          01-00-5e-00-00-fc    static

C:\Users\Administrator>
```

Figure 3: ARP poisoning overview

Figure 3 shows the MAC address of two hosts are the same. The attacker forces the victim to update the wrong MAC address through a duplicate ARP packet, resulting in the attacker receiving all the traffic intended to a different target. DNS spoofing is another technique in which the attacker supplies false DNS information to a victim in order to redirect the entire traffic to the attacker machine. [17]

Sensitive information such as bank account numbers and passwords are exchanged over the internet. The protocol used to protect such information is HTTPS which provides encryption between the web server and browser. When a secure site is visited, HTTPS will appear in the browser with a lock icon in the browser's address field. This mechanism works well most times, but it can be compromised by a method called eavesdropping. [17]

The certificate provided by a bank will be authorized by a Certificate authority (CA) to validate that the certificate belongs to the bank. In a few cases, the connection will be breached by telling the browser to accept a new certificate. In some organizations, employers install special certificates that enable IT departments to intercept HTTPS traffic and monitor employee activity.



Figure 4: Secure communication using HTTPS




Figure 5: Hijacked Communication overview – Man-in-the-Middle attack

When an HTTP session is breached, the user may access what they believe to be a "secure" site and will be fooled into thinking that there is an end-to-end secure HTTPS session. In fact, the browser is just creating a secure session with an intermediate "MITM" server, where all activities are monitored. [17]

Real-Time examples of Man-in-Middle attacks

1. European Parliament Shuts Its Public Wi-Fi After Discovering Man-in-the-Middle Attack (November 27 & 28, 2013):

European Parliament turned off its public Wi-Fi in Strasbourg, France after discovering that a Man-in-the-Middle attack captured communications between mobile phones and Wi-Fi network. The entire parliament Wi-Fi network was shut down to install the certificates by IT. This attack took place in November 27/28, 2013. Complete details about this attack can be found in [SANS Newsletter](#).

2. Bitcointalk.org Forum Targeted by DNS Redirect and DDoS Attacks (December 2, 2013):

Bitcointalk organization users were urged not to log in to their accounts because of a redirect attack. The hacker managed to redirect the traffic to a different address using DNS attacks and captured user login information of those who logged in to their accounts on December 1 and 2. Bitcointalk organization took appropriate steps and cautioned users not to log in to their accounts as they are moving their accounts to a different registrar. The forum was targeted with DDoS attacks as well. Complete details about this attack can be found in [SANS Newsletter](#).

Determining Co-Residence in a Multi-tenant environment

Multiple organizations share same hardware—hypervisor, physical server, network, and storage—to manage their workloads and data. All workloads run on a virtualized infrastructure in cloud computing. The cloud service provider will have full access to all components such as storage, network, and data. Data breached in cloud will result in financial loss, damage to enterprise brand value, and will impose heavy fines on the cloud service provider as per their country's law. [18]

Co-existing on a public cloud

Cloud providers maximize the use of their infrastructure by multiplexing the physical machines with a virtualization concept known as Multi-Tenant technology, which enables assignment of different VM's to different clients when required. This leads to different user's co-existing in the same environment, which poses a potential threat. [19] An attacker can rent one of these VMs and compromise other machines on the same server by selecting the target, placing malicious code on it, and attacking the other machines.

3 steps to attack multi-tenant cloud environments:

1. Locate the target VM and place a malicious VM next to it.
2. Gather information about the target VM.
3. Compromise the target VM using various attack mechanisms.

Step 1: Locate and Place: Cloud service providers usually provide an internal DNS service to map public to private IP addresses. First, it is necessary to enumerate the public service using external probes. Second, map the IP addresses of the public services to their internal IP addresses. To determine the location of a target, attackers can use network-based checks, such as performing a route trace to the target and checking the number of hops, or using side-channel vulnerabilities to analyze possible co-residence. [19]

Step 2: Gather Information: Once the malicious VM is placed near the target VM, information about the target VM is gathered through different side-channels such as measuring the cache usage or estimating traffic rate. This could provide valuable information for accomplishing the last phase of the attack; compromising the target. [19]

Step 3: Compromise: Finally, the target VM can be compromised using either of two vectors. The attacker might decide to attack the VM directly through the compromise of side-channels such as memory or virtual networks. An attacker could compromise the hypervisor, thereby gaining access to all the resources and obtaining full control of the physical machine. [19]

In multi-tenant environments, the attacker first analyzes the infrastructure details, places a malicious VM, and attacks other virtual machines. This is explained in the next section by real time testing in a lab environment.

Simulation of Security attacks in test environment

A sample multi-tenant environment is deployed to simulate DDoS and MITM attacks. The attacks are initiated from CITINTRUDER by collecting all the required details about the web server. The infrastructure details are listed below.

Name	IP address	Type	Deployed Applications
CIT-WEBSERVER	10.x.x.x	Web server	Simple Spring Web Application
CIT-ENDUSER	10.x.x.x	End User	Smarter Stats
CIT-INTRUDER	10.x.x.x	Hacker	Low Orbit Ion Canon, Burp Proxy

1. List of Virtual Machines

- 1. CIT-Webserver:** An insecure http web application with simple authentication is deployed on this server using Spring Suite software. The URL used to access this web application is: <http://citwebserver:8080/SampleSpring>. Once the user is authenticated successfully, the user is redirected to a different page the user is entitled to. The Webserver listens on port 8080.
- 2. CIT-Enduser:** This virtual machine is used to access the Web Application deployed on the server. Smarter Stats is installed on this machine. This tool will analyze the Web Server traffic and provide Web site statistics such as visitor details based on GEO and used bandwidth, notifying administrators if there is suspicious activity on the website.
- 3. CIT-Intruder:** This virtual machine is used to simulate DDoS and MIMT attacks present on the same ESXi server. To simulate a DDoS attack, Low-Orbit-Ion-Cannon tool is used; to simulate a MIMT attack, Burp Suite is used. Advanced IP scanner and nmap (port scanner) is used to find the available resources in the network.

Architecture overview

All the VMs are hosted on a single ESXi host. By default, the traffic between these VMs are not monitored or filtered. CITWEBSERVER is hosted on Guest 1, CITENDUSER is hosted on Guest 2,

and CITINTRUDER is hosted on Guest 3. All three hosts are discovered via Watch4Net monitoring tool. This tool will collect and report real-time statistics i.e. CPU usage, memory usage and disk usage, etc. for all the VMs.

A non-secure web application is hosted on the web server for testing. In this research, the sample web site DDoS and MITM attacks are generated using the available tools to demonstrate the security threats in a multi-tenant environment. Guest 3, the attacker machine in this design, uses various tools to find the resources around it and exploit the target by DDoS and MITM attacks.




Figure 6: Project architecture overview

Security Attack tools used in this research

Low-Orbit-Ion-Cannon: An open source network stress testing and distributed denial-of-service attack application, this tool can flood simultaneous ping requests based on hostname or IP address and port which makes the target unresponsive at the time of [attack](#).

To start a DDoS attack, the IP address of the host is required. If the complete web site address is available, it can be used to target a web server directly based on that port.




Figure 7 Low-Orbit-Ion-Cannon – DDoS attack tool

Burp Suite: A Java application that can be used to secure or crack web applications. This tool can listen in on the specified port and is capable of intercepting/modifying the requests received on that port.




Figure 8: Burp Suite – Security testing tool

Actual Implementation

Step 1: Probing the network

Scan the network and find the IP addresses using Advanced IP scanner that are reachable from the intruder machine (CIT-INTRUDER). This tool will send anonymous ping requests based on the specified subnet and find the IPs that are reachable and the resources that are shared from that IP.

A depiction of Advanced IP Scanner's ability to find the other 2 IPs that are reachable from this machine are shown in Figure 9.




Figure 9: Advanced IP scanner

Step 2: Probing the ports on Web Server

Find the open ports on target host (citwebserver) using Network Mapper (Nmap) - Port scanner.

Nmap is a free and open source utility for network discovery and security auditing. As per the scan results from NMAP many ports are open on citwebserver host. By default, Web Server uses 80/8080 ports unless it is changed by the administrator.




Figure 10: Nmap port scanner

Step 3: Exploiting Web Server through DDoS attack

After collecting all the details of the Webserver, it is attacked (DDoS) on port 8080 using the Low-Orbit Ion Cannon tool. This tool will send thousands of requests simultaneously to Web Server on port 8080, making the resources hosted on the Web Server unavailable to its end users. Although the motive of DDoS attack may vary, it is targeted to temporarily or indefinitely interrupt or suspend services hosted on a high profile Web Server. In this research project, the website experiences 5 to 10 seconds of delay when under attack. The end-user will be completely unaware about the delay in website access.




Figure 11: DDoS attack on Web Server using Low-Orbit Ion Cannon

Situations under DDoS attack: All VMs are monitored using the EMC Watch4Net Monitoring tool. Watch4Net will provide real time monitoring results of all discovered VMs. Figure 12 and 13 show CPU and memory usage peaks at the time of attack. In this situation, the administrator will be notified through an email/SNMP trap.

While the administrator will have no idea about the peak utilization of resources, the exact reason can be found by analyzing Webserver logs. It's determined that the attacker sends thousands of requests to the website and the Webserver will try to respond to all the requests initiated by this tool which increases resource utilization.



Figure 12: CPU usage of Web Server under DDoS attack

Figure 13 shows memory usage is at its peak when the DDoS attack takes place.



Figure 13: Memory usage of Web Server under DDoS attack

Step 4: Hijacking the communication through MITM attack

If the end user tries to access the website, the first step is to send an ARP request asking “who is web server”. The Intruder may reply “I’m the webserver” as ARP does not have any authentication. For example, if server A wanted to communicate with server B, which has the IP address of 10.x.x.x and the MAC address of 00-0A-CC-xx-xx-xx, server A would send out an ARP request asking, “Who is 10.x.x.x?” Then the switch or the operating system would respond, replying with its MAC address, which is 00-0A-CC-69-89-74. The issue with ARP is that any malicious user could send out an ARP request instead of the actual server. [18]




Figure 14: ARP poisoning tool overview through Nighthawk

In this research project, the end user is notified that the actual webserver is an intruder by making an entry in etc/hosts file, instead of ARP spoofing which is used by ethical hackers. By using Burp Suite, the traffic is intercepted between webserver and end user from intruder machine. The intruder can control and intercept or modify the actual data, using this technique to steal valuable data. When the end user tries to access the website, the request is forwarded to intruder due to the entry in hosts file.

Phase 1 - Request from End user: If the end user tries to access the Web Page as per the design in this research project, the request should be forwarded to Web Server. Since the intruder posted itself as Web Server, the request will be forwarded to it.




Figure 15: Request from End User

Phase 2 – Actual response: If the end user is authenticated by the web server, the request will be redirected to another page to which the end user is entitled. The desired response is shown in Figure 16.




Figure 16: Actual response from Web Server

Phase 3- Modification of Data: The Intruder on CITINTRUDER machine will have complete control of the communication between the end user and web server. Intruder can modify the data through Burp Suite and send a new message to the end user as shown in Figure 17.




Figure 17: Modified response sent to End User

Step 5: Monitoring and notification of DDoS attack

The virtual machine on which Web site is running is discovered by Watch4Net monitoring tool. If the CPU, memory, or disk usage of Web Server reaches a certain threshold, the administrator will be notified through email, SNMP trap, etc. The screenshot below (Figure 18) is an overview of the alerting configuration in Watch4Net.



Figure 18: Monitoring Web Server usage and notification through Email/SNMP

As the DDoS attack is carried out on port 8080, malicious traffic is logged in Web Server logs which can be analyzed by Smarter Stats, Weblog Expert etc. to find visitor details such as IP address,

number of visits, and bandwidth used to access the website in a graphical format. Screenshots are available below.

The Web Server logs can be imported to the Smarter Stats tool manually or through FTP links. It analyzes the data and displays the report in GUI format. Figure 19 shows the number of hits on the web site and Figure 20 shows the IP address details that accessed the web site.




Figure 19: Web server traffic analysis through Smarter Stats – # of hits




Figure 20: Web server traffic analysis through Smarter Stats – IP address details

Site administrators can customize Smarter Stats to send scheduled email reports. This feature is available in the full version of Smarter Stats.

Proposed Security measures

It is important to monitor all the critical resources in a multi-tenant cloud environment. Prevention is always better than the cure. With the security measures proposed below, one can easily detect when a DDoS attack occurs by monitoring the infrastructure. On the other hand, it is difficult to detect MITM attacks. This can be prevented only by following security measures. The following are security measures to follow to detect/prevent these attacks in multi-tenant environment.

DDoS

In a DDoS attack, a single IP address is bombarded with a large amount of traffic. If the IP address points to Web Server, it may be flooded with the incoming requests resulting in normal traffic unable to reach the Web Server. Below are the security principles that should be followed to prevent DDoS attacks in a multi-tenant environment. [13]

1. **Use of Firewalls:** By default, inter-VM traffic is not monitored in multi-tenant environments. The Web Server should be placed behind the firewall to ensure that the malicious traffic cannot reach the Web Server directly. All traffic should be monitored, scanned, and

authorized before reaching the Web Server. The Web Server should be placed in a DMZ zone and allow LAN/WAN traffic only after authorization.




Figure 21: Overview of DMZ in a secured environment

There are many tools available in the market such as Pfsense and VMware Vshield. Pfsense can be deployed as a VM on ESXi server and is an open source firewall.

Features of Pfsense:

- Filter traffic by source and destination IP, IP protocol, source and destination port for TCP and UDP traffic
- Able to limit simultaneous connections on a per-rule basis
- Option to log or not log traffic matching each rule
- Aliases allow grouping and naming of IPs, networks, and ports. This helps keep your firewall rule set clean and easy to understand, especially in environments with multiple IPs and numerous servers
- Pfsense provides RRDtool graphs, visually displaying every operational process in the box, including WAN/LAN traffic and system processes




Figure 22: Features of PfSense




Figure 23: Traffic overview in PfSense

2. Defend the attack at network perimeter

There are a few technical security measures that can be taken to partially minimize the effect of attack. DDoS attackers use the network as a medium to generate the attack. In multi-tenant environments, a firewall or IPS can be used to enable access control lists which can help control or minimize the impact of a DDoS attack.

3. Use of IDS

It is important to monitor malicious traffic like dropped packets, unauthorized access attempts, and so on. This can be achieved by deploying IDS on a VM and enable promiscuous mode on the port that is connected to IDS. A promiscuous-enabled port will listen to all traffic on the wire. [Catbird](#) is an example of a security device that can be used to protect the hypervisor.




Figure 24: IDS overview in multitenant environment

4. Bandwidth Consumption

An intruder may consume all available network bandwidth by generating a large number of packets directed to the network. While these packets are ICMP ECHO packets, in principle they may be anything.

5. Review Web Server logs

It is important to review the Web Server logs to find visitor details. There are many tools, such as Smarter Stats and others to review Web Server logs in real time and send reports to the administrator at scheduled intervals.

6. Physical Security

It is important to protect the physical servers in a data center environment. Attackers who have access to the physical infrastructure can easily launch attacks and steal the most valuable data.

7. Contact DDoS Specialist

Sometimes, DDoS attacks will have heavy impact and revenue loss. For example, if a bank site is affected, revenue loss is high. The best choice to prevent this attack is to contact a DDoS Mitigation Company. These organizations use a variety of technologies to prevent DDoS attacks and keep the web site online. DDoS mitigation specialists include:

- [Arbor Networks](#)
- [Black Lotus](#)
- [DOSarrest](#)
- [Prolexic](#)
- [VeriSign](#)

MIMT

Although it is difficult to detect a MITM attack, it can be identified by analyzing the network traffic. However, it is very difficult to analyze thousands of network packets. Preventive measures that can be used to control MIMT include: [21]

- **Strong encryption:** Using encryption between server and client makes it difficult to hijack the session. The server can authenticate itself by presenting a certificate and then the client and server can establish an encrypted channel to send sensitive data.
- **Use of IDS:** The IDS will monitor the network traffic, and if someone tries to hijack the traffic flow, IDS will notify the administrator. Tools which use the advanced address resolution protocol (i.e. XARP or ARPOn) and measures such as implementing dynamic host configuration protocol (DHCP) snooping on switches can limit or prevent ARP spoofing. This in turn can help prevent man in the middle attacks.
- **Use of certificate checker:** While browsing, the user can check if the certificate is issued by a legitimate CA or if it's a fake certificate issued by some local CA. In Mozilla browser, [Cert Patrol](#) and [Perspective](#) are plugins that can check and validate certificates. These add-ons cannot detect a MITM attack, but can detect if something is odd about the website certificates.
- **Configure HTTPS:** With HTTPS, it is difficult for hackers to view network traffic. However, some hackers are able to hack HTTPS-configured Web sites as well. If HTTPS is configured, the browser will detect and warn, unless the hacker already compromised the systems.
- **Session ID's:** It is important to configure session ID's for all transactions. Hackers cannot hack the session until they can intercept the session ID's.

Conclusion

As cloud computing gains popularity, its widespread use raises the issue of cloud security. Due to its flexibility and cost efficiency, multi-tenant architecture widely used in cloud computing is prone to security threats, notably Distributed Denial of Service (DDoS) attacks and Man-in-Middle (MITM) attacks. These threats are tested in this Knowledge Sharing article with real time scenarios. To improve resource availability and data integrity in multi-tenant cloud environments it is essential to set up proactive monitoring and implement strong security measures as proposed in this article. Doing so will greatly reduce the impact of an attack.

Recommendations to secure multi-tenant environments

There are no established guidelines to control and monitor the system in cloud under an attack. In the research discussed in this article, various freeware tools are used to initiate the security attack and monitor the systems under DDoS attack. However, there are no fixed security principles to control or monitor a MITM attack. IDS devices can be used to some extent to discover a MITM attack which can come from various channels. Thus, it is important to consider the recommendations below to strengthen the security of ESX server in a multi-tenant environment and secure all of the attack channels.

1. **Patch Management:** This is an important process to protect the infrastructure components from various attacks. Patching and malware protection is very important. Critical patches should be applied to ESX servers and VMs which is the backend to cloud technology. These patches must be tested before being applied to the infrastructure. The vulnerability and patch management system should be up to date in order to tackle the latest security attacks. vCenter Update Manager enables centralized, automated patch and version management for VMware vSphere and offers support for VMware ESX/ESXi hosts, virtual machines, and virtual appliances as well. As per industry standards, the timeframe in which patches can be applied is shown in the table below.

Criticality	Critical	Important	Moderate/Low
Time Frame	Pilot Testing: 24-72 hours Primary Update: < 14 days Secondary/Final Update: 15-60 days	Pilot Testing: 7-10 days Primary Update: 10-30 days Secondary/Final Update: 30-90 days	Pilot Testing: 30/45 days Primary Update: 45-90 days Secondary/Final Update: 90-180 days

Patches overview

2. **Protection at various network levels:** As attackers will use the network to attack the systems, it is recommended to harden the network security at various levels:

- Layer 2: VLANs
- Layer 3: IP-based ACLs and Firewall rules
- Layer 4: TCP, UDP, and ICMP rules
- Layer 5-7: Application- and Session-based access rules

3. **Admin and Network access:** Strictly control the vCenter administrator privileges to secure the system. By default the local Windows administrator user will have access to vCenter. To rectify this, do the following:

1. Create a local admin group with full admin access to VCenter and delete the local admin group.
2. Create a domain global group for all VC admins and add domain global group to it.

The vCenter server should not be placed in any network other than the management network. By limiting network connectivity, certain types of network attacks can be controlled. Block the network ports that are not in use, by using the Windows firewall or any external firewall.

4. **vCenter certificates:** Configure SSL between vCenter, vSphere, and ESXi hosts. Install new certificates that are signed by a valid internal certificate authority or purchase from a trusted security authority.
5. **Time Sync:** All hosts in the network should be configured to use NTP server. If time is not synchronized between the network machines, SSL certificates will not be recognized and can result in authentication problems. This will also affect the log analysis, forensic analysis, and troubleshooting.

6. **Logging:** By default, the logs on ESXi server are stored in the in-memory file systems. Only 24-hour data is saved and this will be lost when the host is rebooted. It is recommended to forward the ESX logs to a remote syslog server. The important logs that should be collected are:
 - Service console logs with level info or greater
 - Authentication logs – “Auth Priv”
 - Vmkernel warnings with facility “local6” and level “warning” or greater
 - Vmkernel logs with facility “local6” and level “notice” or greater
7. **Warning Banners:** This should be configured to communicate Security policy and to provide legal protection. Banner messages will alert users logging in to the system accidentally to log out. Different types of warning banners are Console login, Remote login, Emergency Console, After login, and SSH login.
8. **Disable unneeded services:** Some services that exist on the ESXi hosts might not be necessary. Disable the services that are not in use and prevent users without privileges from mounting CD's and file systems through the console.
9. **VMware guest security:** OS hardening should be implemented by treating the guest OS as a real OS by using hardening guides. Disable unnecessary services on the VM and use templates to deploy virtual machines. Disable Guest <-> Host Copy & Paste as this is enabled by default. Disable unauthorized devices and prevent device connection and removal.
10. **MAC Address changes:** The MAC address change options will affect the traffic that a virtual machine receives. If this option is set to Accept, ESXi accepts requests to change the effective MAC address to other than the initial MAC address. When the option is set to Reject, ESXi does not accept requests to change the effective MAC address, protecting the host against MAC impersonation. It is recommended to set the “Forged Transmission” option to reject on ESXi so that it compares source and effective MAC addresses which restricts forged transmits.
11. **DMZ Virtualization:** A Demilitarized Zone is a physical or logical sub network that contains and exposes an organization's external-facing services to a larger untrusted network, usually the Internet. DMZ adds a layer of security to an organization's LAN; an external attacker only has access to equipment in the DMZ, rather than any other part of the network. The biggest risk to a DMZ in a virtual environment is

misconfiguration, not the technology. (Gartner) Thus, strong audit controls are required to avoid misconfiguration, either accidental or malicious.

12. vCenter Security: vCenter is a critical component of vSphere and it needs to be secured. vCenter generally runs on a Windows server. The host should be protected against vulnerabilities and attacks ensuring that the host is as secure as possible. Strictly control vCenter server administrator privileges to increase the security of the system. vCenter server should be placed only in management network. Use a local firewall on the Windows system or network firewall to limit the connectivity to vCenter server.

13. Minimize Loss of Control (Monitoring): The most important part is to pro-actively monitor the entire infrastructure such that the administrators will be notified if there is something wrong. There are many monitoring and security tools available to monitor or prevent security attacks. All the tools should be integrated and develop a solution that helps find problems easily.

References

1. Cloud Security Alliance, "Security Guidance for Critical Areas of Focus in Cloud Computing, ver. 2.1, 2009" – White paper
2. Trend Micro Corporation, "Security Threats to Evolving data centers". White paper.
3. The Force.com, "[Multitenant architecture, Understanding the design and Salesforce.com's internet application development platform](#)" – White paper
4. iXia, "[The Virtual Blind Spot Best Practices for Monitoring Virtual Environments](#)" – White Paper.
5. Wipro, "[Multi-Tenant enabling a Single-Tenant Application](#)" – White paper.
6. Ristenpart T. et al. (2009) "Hey You Get Off My Cloud," Proceedings of the 16th ACM conference on Computer and communications security, Chicago, Illinois, USA
7. Nils Gruschka and Meiko Jensen "Attack Surfaces: A Taxonomy for Attacks on Cloud Computing", 3rd International Conference on Cloud Computing, 2010.
8. Winkler, Vic (2011). [Securing the Cloud: Cloud Computer Security Techniques and Tactics](#). Waltham, MA USA: Elsevier. pp. 65, 68, 72, 81, 218–219, 231, 240. [ISBN 978-1-59749-592-9](#).

9. Clark, C., Franser, K., Hand, S., Hansen, J. G., Jul, E., Limpach, C., Pratt, I. and Warfield, A.: Live Migration of Virtual Machines, Proc. Symp. Networked Systems Design and Implementation, pp. 273–286 (2005).
10. Dave Shackelford, "[Next-Generation Datacenters](#)" – SANS white paper.
11. Trusted Computing Group: TPM Main Specification Version 1.2, <http://www.trustedcomputinggroup.org/>.
12. AFORE, "[Protecting Data In Multi-Tenant Clouds](#)" - White Paper
13. Lonea A.M, Popescu D.E, Tianfield H, "Detecting DDoS Attacks in Cloud Computing Environment"
14. Cart Tim, "Seven Deadliest Social Network attacks"
15. Cloud Computing Security Risk Assessment - The European Network and Information Security Agency (ENISA)
16. Dhananjay Sakhalkar, "[How to Identify and Mitigate Man-in-the-Middle attacks](#)" – Cognizant white paper
17. Christopher Shields M., Matthew Toussain M., "[The MITM Framework](#)" – Subterfuge White Paper
18. Ronald Krutz L., Russell Dean Vines, "Cloud security: a comprehensive guide to secure cloud computing"
19. Trend Micro White Paper - <http://la.trendmicro.com/media/wp/cloud-computing-security-en.pdf>
20. Heiser J and Nicolett M, "Assessing the Security Risks of Cloud Computing," Gartner 2008.
21. Mell P. and Grance T., "Effectively and Securely Using the Cloud Computing Paradigm," ed: National Institute of Standards and Technology, Information Technology Laboratory, 2009.
22. McDermott J., (2009) "Security Requirements for Virtualization in Cloud Computing," presented at the ACSAC Cloud Security Workshop, Honolulu, Hawaii, USA, 2009.

EMC believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED "AS IS." EMC CORPORATION MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.