

A Highly Scalable and Reduced-Risk Approach to Learning Network Man-in-the- Middle (MITM) and Client-Side Exploitation (CSE)

Dale C. Rowe Ph.D.
dale_rowe@byu.edu

Sarah Cunha B.S.
scunha@byu.edu

Cara Cornel
cara.cornel@byu.edu

Brigham Young University
Cybersecurity & Systems Research Laboratory
Crabtree Technology Building, Provo, UT, 84604

Abstract - Man-in-the-middle attacks are commonly used by penetration testers and malicious hackers to intercept, monitor and manipulate network traffic. MITM attacks may take place at the first four networking layers and are often used in exploit-chains to spread laterally within a target network.

In this paper, we describe the various types of MITM attacks and how they can be used to deliver effective client-side exploits. Various challenges commonly encountered are discussed followed by a novel approach that significantly lowers risks while simultaneously enriching the learning experience for students. We discuss the introduction of this approach into the classroom environment and student feedback.

We conclude with a discussion of future development objectives and a summary of our key findings.

Categories and Subject Descriptors

K.6.5 [Management of Computing and Information Systems]: *Security and Protection – Unauthorized access (e.g., hacking, phreaking)*

General Terms

Security

Keywords

Education, Man-in-the-Middle (MITM), Network Security, Security, Penetration testing

1. INTRODUCTION

A Man-In-The-Middle, or MITM, attack involves the interception and/or manipulation of information between two or more parties by a malicious third party. These attacks frequently take place at layers 1–4 of the OSI network model.

Layer 1 Attack

In a physical layer attack, Mallory uses access to the network media to exploit communications between Alice and Bob. These attacks typically require a level of physical proximity to the victim network. If the physical medium is wireless, attackers may use a variety of techniques to discover wireless encryption keys, perform a denial of service attack or eavesdrop on sensitive communications. For wired networks, a TAP, hub, inline device or SPAN port can allow monitoring of network traffic. Note that only an inline device may actually manipulate network traffic without using higher level network layer attacks to supplement the monitoring capabilities.




Figure 1 – A MITM attack against Alice and Bob using ARP poisoning

Layer 2 Attack

The simplest, and likely most common form of MITM attack focusses on an inherent weakness of the IP protocol within a layer 2 network broadcast domain [1]. Address Resolution Protocol (ARP) provides a lookup capability to determine the Layer 2 (MAC) address of the intended IP destination.

In an ARP poisoning attack, Alice wishes to talk to Bob by means of an IP address. Alice compares the network portion of her own IP address with Bobs and determines that they are located in different networks (**Error! Reference source not found.**). Alice then consults her routing table to identify the IP address of a router in her network capable of reaching Bob (typically this is the same as the default gateway for a client endpoint). Knowing the IP address of the router, Alice broadcasts an ARP request seeking the MAC address of the router, e.g. [[WHO HAS 192.168.0.1]]. Seeing a local broadcast with its own IP address, the router responds with [[192.168.0.1 IS AT 00:11:22:33:44:55]]. Alice can now add the MAC address of the router into the L2 Ethernet destination of her datagram. Alice will remember her routers MAC address for a defined period (typically 60 seconds).

Assuming that the base Layer 2 protocol remains Ethernet, this process repeats between each router until the packet reaches Bob, with the terminal router broadcasting for the MAC address of Bob.

The principal vulnerability in ARP involves sending an unsolicited ARP response, also known as a gratuitous ARP reply; from the attacker to the two parties wishing to communicate on the same L2 broadcast domain. In the above example this may be between any locations identified by a red arrow. The unsolicited ARP response includes the MAC address of the attacker with the IP address of each victim node. In this way, the attacker is sent all packets between the parties.

The successful completion of an ARP poisoning attack places Mallory in a position to bi-directionally intercept *all* communications between Alice and Bob. It should be noted that in the provided example where Alice and Bob are in separate L2 domains, the poisoning of routers will result in *the interception of all traffic passing between the two targeted nodes*. A common tool for ARP poisoning is Ettercap.

Layer 3 Attack

Layer 3 attacks are often sophisticated and thus more rarely used. The manipulation of routing protocols results in the redirection of packets to Mallory. Attacks may take place against a variety of routing protocols such as RIP, RIPv2 and OSPF [2]. Attack methods may involve gaining illegitimate access to a routing device or injecting false routing updates into a network. The latter may require cracking now-obsolete hashing protocols such as MD5, which are used to provide route update integrity [3], [4].

Layer 4 Attack

The second most common form of attack relies on transport-layer protocol vulnerabilities, most commonly targeting DNS. Though DNS attacks generally employ simple impersonation, they may carry out a MITM attack by relaying received traffic to the real destination. As the principal host-resolution protocol of the internet and lacking confidentiality or integrity protection, DNS is a prime target for race conditions, cache poisoning and server-side attacks seeking to misdirect a victim to a malicious target.

Client-Side Exploitation Attacks (CSE)

A client-side attack is an attack that seeks to take advantage of an end-user's working environment such as web browsers, document viewers and media viewing applications. Ed Skoudis, a well-known security expert, stated, "For penetration tests with client-side exploits in scope, compromise is almost always successful". CSE attacks may involve either the misdirection of a user and/or application to a host holding the exploit, or the direct injection of malicious code into legitimate network traffic by means of a MITM attack.

2. CHALLENGES OF MITM STUDENT LEARNING

While few challenges outside of ethical conduct exist in a conceptual approach to MITM attack learning, more hands-on courses such as advanced networking or penetration testing often involve an applied approach to student education. This involves risks at multiple levels.

Prior to 2015, a laboratory assignment in our penetration testing course required students to conduct an ARP poison by creating virtual machines on a lab workstation. With the course targeted to seniors and graduate students, the creation of short-term virtual machines was an unnecessary use of student time given their well-proven capability to provision virtual machines by this stage of their education. Workstation virtualization also posed several risks which are discussed below:

Interference

The most common mistake students make in our experience is the inadvertent use of an incorrect IP address. This is sometimes due to a typographical error, misreading of configuration information, or misinterpretation of instructions. The common result is that two students will now be in a situation of conflict, or 'stacked poisoning' – creating a chain of poisons that often results in an effective denial-of-service to one or more computers.

Network Outages

In extreme circumstances, wildcard poisoning (poisoning an entire subnet) can result in a denial of service when conducted on busy networks, or by multiple students simultaneously. In an extreme instance of this, a student network administrator attempted to visualize large quantities of network data using a tool known as “etherrape”. This tool performs a mass-ARP poison of a subnet and depicts various visualizations of traffic-flows. While the system was left in place for a few days without a serious impact, the network was DoS’ed with extreme prejudice during a network-based operating system reimage over the weekend.

Accidental or Malicious Interception of Legitimate Data

A successful MITM attack is often a means of information-gathering. Many tools, including Ettercap actively identify clear-text credentials, interesting website cookies and other potentially sensitive information. While the use of Layer 4+ encryption technologies such as SSL/TLS are becoming more commonplace, the inadvertent or deliberate capture of fellow student (or faculty) data can be a significant risk.

Delivery of Client-Side Exploit to Out-of-Scope Target

Similar to out-of-scope data interception, it is plausible that a student may incorrectly use a wildcard MITM or host IP to target a host that is out of scope. While every precaution may be taken to keep systems up-to-date, the use of legacy software could easily render a system vulnerable to this form of attack. A recent example of this is the recent discontinuation of Quicktime by Apple [5], or in-the-wild Adobe Flash Player [6] vulnerability. Both of these vulnerabilities are examples of attacks that may be exploited prior to patches becoming available.

Remote Access to Lab Assignment

An increasing request from students is the ability to work remotely. This is more common in senior and graduate classes where many students seek to balance part-time employment with their education. There are inherent dangers to MITM

attacks across remotely accessible networks to both the student's local network and the remote connection mechanisms (e.g. VPN).

Lack of Realism

Many approaches to MITM exploitation involve the student having console level access to the victim machine by means of virtualization software. While this may be desirable in an instructor demonstration, it is unrealistic in a penetration testing exercise. The principal reason for a penetration tester to deliver a client-side exploit is to gain access to a system that is thus far inaccessible.

Scalability

By nature, network attacks are limited in their scalability. The greater the number of attacks, the greater the chance of all aforementioned risks. There are also inherent difficulties in a cybersecurity environment that result in system instability during exploitation. Software exploits carry a high risk of subsequent system instabilities regardless of whether or not the exploit actually succeeds. The greater the number of users attempting to exploit a system, the less likely their chance of success.

Solutions for classroom learning should be scalable and should provide either individual students, or small teams of students with access to their own systems. This approach minimizes frustrations with instabilities caused by peers and provides a more effective learning experience. Consideration must also be given to the required infrastructure and seek to minimize costly increases of capacity for singular short-term exercises. Careful planning can help maintain an efficient use of computing resources.

3. SANDBOXING APPROACHES TO MITM EDUCATION

Client-Side Virtual Networks

The most common approach to sandboxing MITM learning environments involves the use of virtualization software on a student workstation. Students are

typically instructed to use host-only or NAT networking to provide a “safer” network environment from which to conduct their work.

This approach relies entirely on the students correctly understanding and following instructions to configure the sandbox environment. If implemented incorrectly, all the aforementioned risks may be present. It has been our finding that the learning experience is diminished by the student’s time spent on setting up the environment and ability to access the victim console.

Isolated Classroom

In discussions with other institutions, we have found that most programs will completely isolate a network to ensure students remain in-scope. While an effective method to prevent mistaken targeting outside the classroom network, students may still cause a detrimental experience to each other. Furthermore, the realism of the experience is diminished by the inability of victim machines to have regular contact with the internet.

Network Separation

A hybrid approach is to separate network subnets from each other by means of a layer-3 routing topology. While this provides protection against most MITM attacks outside the classroom, it can still impact students engaged in a learning exercise. This can become much more significant in a shared-classroom environment where other scheduled classes may be affected by both successful and failed MITM attempts.

Architecture of a Lower-Risk Approach

We have taken a multi-layered approach to MITM and CSE education. It has been our experience [citation will be added after review] that many penetration testers use MITM attacks as a pivoting mechanism within an organizations network. An initial network penetration to a target provides the attacker with the capability to launch MITM and CSE attacks to a wide range of victims and maneuver both horizontally and vertically within an organization’s network.

Our approach is based on double separation of the target network from the classroom. Students conduct their attack from a remote virtual host in an isolated IP subnet. This host is configured with two network interfaces: One to receive the incoming connection and the other in a second subnetwork shared with the victim. All networks are routable to the Internet allowing the manipulation of realistic traffic from the target host.

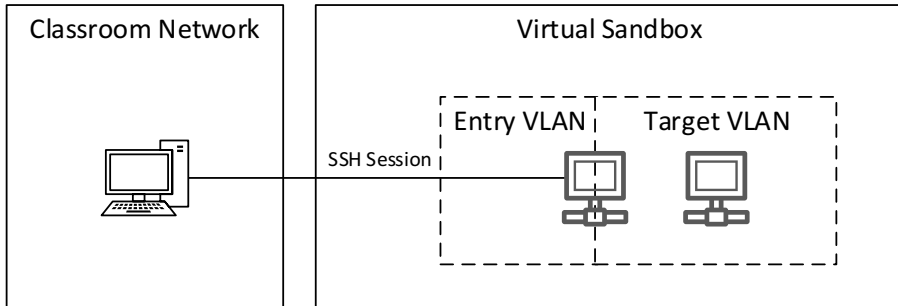


Figure 2 – Proposed Approach to MITM Lab Exercise

The scalability of this system revolves around the ability to create multiple isolated instances of both the entry and target networks. This prevents accidental peer-interference, risks of wild-card ARP poisoning, and inadvertent collection of private data. This is achieved while facilitating remote access using existing methods. Our process for the deployment and management of network segments and workstations is fully automated and tested to cope with over 40 simultaneous users (80 VM's) on surplus (9+ years old) hardware (6 HP BL460C G1 Blades, Dual Xeon 5160 CPU's with 32Gb RAM).

After completion of a lab exercise, students are shown the architecture of the lab environment. This architecture is used to demonstrate various defense-in-depth principles such as network layering, traffic separation, intrusion detection and full lifecycle security planning.

Example Student Laboratory Assignments

Our penetration testing course includes two exercises involving a MITM delivered Client-Side Attack. In the first, students are provided with SSH access to an internally hosted Ubuntu 14.04 LTS server as depicted in **Error! Reference source not found.** The target host is connected to a second network interface on the Ubuntu 14.04 host with a script refreshing a standard HTTP website in Firefox every 5 minutes. The target is set to auto-open PDF files if they are received by the browser. Each student must exploit a vulnerability in Adobe Acrobat Reader 9.3.4 such as CVE-2010-2883 (Cooltype SING vulnerability) by means of manipulating the HTTP traffic in a MITM attack. If successful, students retrieve a flag from the compromised host as proof of success.

The second exercise is part of the final lab and involves a similar method of attack, this time to exploit a Java vulnerability in an Ubuntu 10.04 host. This exercise is significantly more difficult due to the need to pivot through two hosts, and the deliberate sharing of the victim network with other students. This was performed to educate students on the risks of MITM attacks while containing their impact to an isolated network. The setup for this and is shown in **Error! Reference source not found.**

The progression of these two labs provides students with a chance to become familiar with the steps involved in a successful man-in-the-middle attack without interference from their peers. This ceteris paribus approach is particularly useful in letting the student research their own solutions, as failure conditions are nearly always the result of an incorrect attempt to carry out the exploit. By providing a shared target network in a later exercise, students can also become familiar with the indicators of a failed ARP poison.

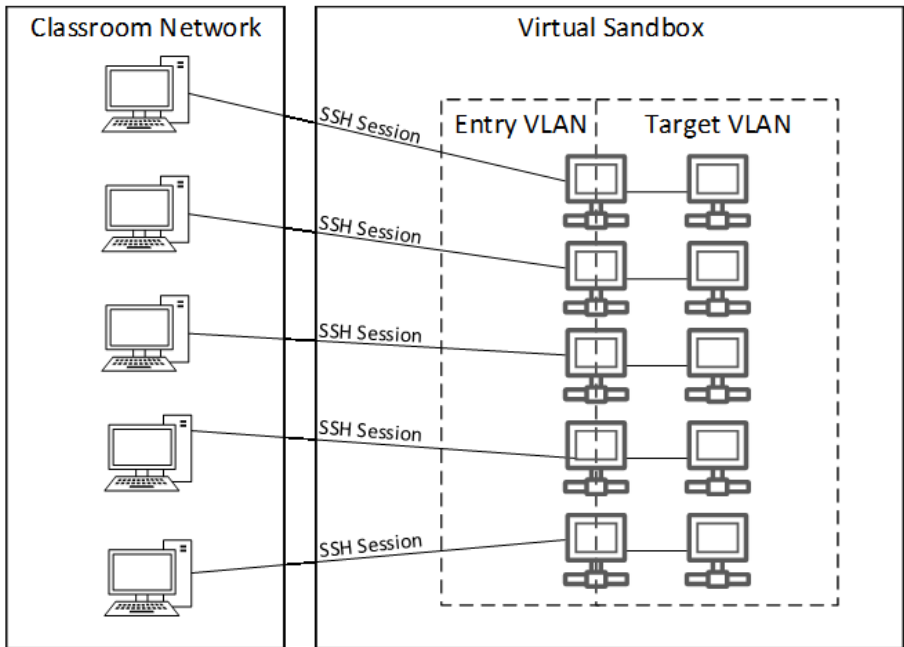


Figure 3 - Proposed MITM Lab Architecture (shared target network)

4. FINDINGS

This approach was deployed in 2015 and has been used in four class sections totaling approximately 160 students. Feedback for the laboratory has been positive although initial issues with cloning target systems proved problematic due to the Windows Firewall detecting a new network and blocking the network exploits from Acrobat extensions. It should be noted that there have been no network problems that are attributable to these exercises since migrating to this approach.

Our findings on the lowered-risk approach can be summarized as follows:

- Risk avoidance by containment of malicious ARP traffic in sandboxed network.
 - No risk to confidentiality of sensitive network data such as real login credentials.
 - No risk to out-of-scope host integrity in client-side exploitation.
 - No risk to availability from network denial-of-service.
 - No risk of availability from peer interference. (Risk is reduced in second stage where students share a target environment).
- Improved learning experience from:
 - Per-user network sandboxing avoids the attribution of failure to a peer's attempt.
 - Realistic pivoting scenario where students must execute their MITM attack attempts from a remote system.
 - Realistic end-goal (client-side exploit) shows a complete exploit-chain on the final lab (breach machine A in entry VLAN to MITM CSE attack victim in target VLAN).
 - Removal of ability to 'cheat' in obtaining flags directly from virtual-disks by hosting target VM remotely.
 - Ability to use the system's architecture as an example of defense-in-depth design.

5. FUTURE EFFORTS

We have identified a few areas that could increase the learning experience for students and are working to implement these for upcoming course offerings. These are discussed forthwith.

Increasing realism of the overall scenario

A key part of an applied learning experience is providing a realistic environment in which to learn [7]–[10]. One thing absent from our current environment is network ‘chatter’ from other hosts and users. Our goal is to add this to the second stage exercise once students have already learned how to successfully launch such an attack.

Wherever possible in our curriculum, offensive activities are discussed alongside defensive strategies to mitigate or eliminate the involved risk. A third exercise is under development that involves researching solutions to mitigate the MITM CSE attack vector at various levels. Strategies for this include client-side patch management, dynamic ARP inspection or static ARP.

Snapshot of Victim Terminal

We have noticed that some students experience difficulty in the debugging cycle as they develop their client-side attack. To assist in this debugging, we are creating a simple agent that will post a per-minute screen capture of the victim’s user desktop environment. This will assist students in determining if elements of their attack are successful.

SSL/TLS

In 2010, instructors at East Carolina University described a MITM laboratory exercise against the TLS protocol [11]. The authors describe a series of exercises designed to help students understand the weaknesses of asymmetric PKI trust relationships. We are in the process of extending the exploit-chain to include attacks against encryption channels such as certificate trust manipulation and SSL/TLS interception.

Detection and Reporting Capabilities

MITM attacks are relatively noisy, availing themselves to multiple detection methods from both hosts and network devices. We are developing a logging mechanism to allow students to see such attacks in real-time from a reporting standpoint. This will not only assist in exploitation attempts, but also provide students with tangible data to demonstrate the ease of which such attacks can be detected. The latter is an attempt to dissuade any deliberate abuse of the knowledge acquired during these learning experiences. We have found sharing real-time detections of penetration test activities to be effective in raising awareness of effective system monitoring.

Use of Additional Network Layers

Although the focus has been on Layer 2 attacks, the presented approach to low-risk MITM learning is well-suited to attacks against layers 1, 3 and 4 of the OSI layer. We are developing additional scenarios that require students to target additional network protocols in order to carry out a successful attack.

6. CONCLUSIONS

In this paper, we have described how an automated virtualization environment has been designed to provide a powerful, hands-on experience of a high-risk MITM exploit chain. The design effectively eliminates and/or minimizes common risks typically encountered in the learning process. The design is also suitable for hands-on instruction of various defense-in-depth principles that can be used to supplement student learning. Student feedback has been noted as positive even with larger class sizes (up to 40 students). While the configuration of the learning environment is more complex than workstation-based virtualization, equipment costs are minimal and surplus legacy hardware is often more than adequate.

Finally, we have identified various areas for future improvements to further enhance the learning experience and have discussed these at a high-level.

ACKNOWLEDGEMENTS

The authors wish to express their gratitude to BYU's Office of Information Technology for the donation of surplus infrastructure used in this project.

BIBLIOGRAPHY

- [1] R. Wagner, "Address resolution protocol spoofing and man-in-the-middle attacks.," 2001.
- [2] A. Wong and A. Yeung, "Network Infrastructure Security," Boston, MA: Springer US, 2009, pp. 181–218.
- [3] Y. Wang, J. Chen, and D. He, "A new collision attack on MD5," in *Proceedings - International Conference on Networks Security, Wireless Communications and Trusted Computing, NSWCTC 2009*, 2009, vol. 2, pp. 767–770.
- [4] M. Stevens, A. K. Lenstra, and B. De Weger, "Chosen-prefix collisions for MD5 and applications," *Int. J. Appl. Cryptogr.*, vol. 2, no. 4, pp. 322–359, 2012.
- [5] US-CERT, "Apple Ends Support for QuickTime for Windows; New Vulnerabilities Announced," *Alert (TA16-105A)*, 2016. [Online]. Available: <https://www.us-cert.gov/ncas/alerts/TA16-105A>. [Accessed: 15-Apr-2016].
- [6] Adobe, "Security Advisory for Adobe Flash Player," *CVE-2016-1019*, 2016. [Online]. Available: <https://helpx.adobe.com/security/products/flash-player/apsa16-01.html>. [Accessed: 16-Apr-2016].
- [7] G. Vigna, "Teaching Hands-On Network Security: Testbeds and Live Exercises," *J. Inf. Warf.*, vol. 2, pp. 8–24, 2003.
- [8] R. L. Fanelli and T. J. O, Connor, "Experiences with Practice-focused Undergraduate Security Education," in *Proceedings of the 3rd International Conference on Cyber Security Experimentation and Test (CSET'10)*, 2010, pp. 1–8.
- [9] P. M. Stohr-Hunt, "An Analysis of Frequency of Hands-On Experience and Science Achievement," *J. Res. Sci. Teach.*, vol. 33, no. 1, pp. 101–109, 1996.
- [10] K. E. Kercher and D. C. Rowe, "Improving the Learning Experience for the Deaf through Augmented Reality Innovations," in *18th International ICE-Conference on Engineering, Technology and Innovation*, 2012.
- [11] J. Lewis and P. Lunsford, "TLS man-in-the-middle laboratory exercise for network security education," in *Proceedings of the 2010 ACM conference on Information technology education - SIGITE '10*, 2010, p. 117.