ASEE AMERICAN SOCIETY FOR ENGINEERING EDUCATION

2017 ASEE Mid Atlantic Section
Spring Conference: Morgan State University, Baltimore, Maryland Apr 7 Paper ID #20827

A Lightweight Collaborative Virtual Computer Laboratory for Cybersecurity Education

Dr. Abdullah Konak, Penn State Berks

Abdullah Konak is a Professor of Information Sciences and Technology at the Pennsylvania State University Berks. Dr. Konak received his degrees in Industrial Engineering, B.S. from Yildiz Technical University, Turkey, M.S. from Bradley University, and Ph.D. from the University of Pittsburgh. Dr. Konak's current research interest is in the application of Operations Research techniques to complex problems, including such topics as network design, network reliability, facilities design, and data mining. Dr. Konak has published papers in journals such as IIE Transactions, Operations Research Letters, Informs Journal on Computing, IEEE Transactions on Reliability, International Journal of Production Research, and Production Economics. He has been a principle investigator in sponsored projects from the National Science Foundation, the US Department of Labor, and the National Collegiate Inventors and Innovators Alliance. Dr. Konak currently teaches courses on Database Management Systems, Information Security, and Technology-based Entrepreneurship. He is a member of IIE, IIIE and INFORMS.

Mr. Anuvrat Sheoran, Pennsylvania State University - Berks Campus

I am a senior at The Pennsylvania State University taking Security Risk Analysis with the Cyber Security option as my major and Information Sciences & Technology as my minor. I am certified by The National Security Agency with the certificate of Achievement.

A Lightweight Collaborative Virtual Computer Laboratory for Cybersecurity Education

Abdullah Konak and Anuvrat Sheoran
Penn State Berks

Abstract

As more and more cybersecurity related threats emerge, it is imperative that cybersecurity students are trained to deal with these threats swiftly and efficiently. Pennsylvania State University - Berks Campus have been using a virtual computer laboratory called Collaborative Virtual Computer Laboratory (CVCLAB) over the last decade to provide students with hands-on experiences in cybersecurity topics. The CVCLAB is quite functional and has made a significant impact on student learning. However, the CVCLAB can be resource intensive in particularly when many students use it simultaneously. Therefore, it can also be costly to maintain. In this paper, we present the blueprint of a new virtual computer laboratory (Lightweight (L)-CVCLAB) based on the command prompt interface. The LCVCLAB has been designed for teaching technical skills related to host and network penetration testing. Students gain access to the LCVCLAB through an SSH client and then connect to multiple Linux-based virtual computers specifically configured for cybersecurity training. Without the overhead of a desktop or web-based client, the LCVCLAB has a better response time and can scale well for large classes. The paper also presents sample hands-on activities that can be performed in the LCVCLAB.

Keywords

Cybersecurity Education, Virtual Computer Laboratories, Penetration Testing

I. Introduction

In recent years, many private and government organizations have been victims of cybersecurity attacks such as denial of service, losses/theft of personally identifiable information (PII), phishing attacks, and social engineering frauds. The average security breach can cost a company between $90 and $305 per lost record, according to a study by Forrester Research.[1] Not surprisingly, cybersecurity tops the list of information technology priorities at many organizations. As a result, the rate of growth for jobs in information security is projected at 18% from 2014 to 2024, which is higher than the average job growth, according to the Bureau of Labor Statistics.[2]

Technical skills required for cybersecurity jobs include incident response and management, firewall/router administration skills, OS analytics, intrusion detection, penetration testing, access management, web security, application security, malware prevention, and cloud computing/virtualization. Cybersecurity graduates are expected to not only learn these skills conceptually but also practice them. However, providing students with hands-on practical learning experiences in cybersecurity topics has been difficult for many reasons. Major problems are: *i*) lack of financial resources to establish and maintain sufficient number of hardware and software so that each student can have adequate access to the learning environment; *ii*) inaccessibility of computer laboratories outside scheduled class times; *iii*) lack of physical space; *iv*) lack of dedicated staff to support computer labs; *v*) information technology policies restricting students' privileges on laboratory computers in order to protect the other services provided by the laboratory; and *vi*) lack of secure networking environment to test critical tasks. In summary, providing students with exposure to a broad range of practical cybersecurity experiences is a costly and complex operation, which many education institutions may not undertake.

In the last decade, virtual computer laboratories have been utilized effectively to address the problems listed above. A comprehensive list of virtual computer laboratories for cybersecurity education and training is given by Richards et al.[3] At Penn State University's Berks Campus, we have been using a virtual computer laboratory, called Collaborative Virtual Computer Laboratory (CVCLAB), in several face-to-face and online information security courses[4, 5] since 2009. The CVCLAB has made a significant impact on student learning as documented in several empirical studies.[6, 7, 8, 9,10] However, scaling up the CVCLAB for a large number of concurrent users has become prohibitory expensive as an increasing number of courses have started using it.

In this paper, we present a simplified model of virtual computer laboratories for cybersecurity education. While the CVCLAB is a general-purpose computer laboratory that can be used for the instruction of many different computer sciences and information security courses, the model presented in this paper focuses on only limited application areas, in particular, penetration testing and intrusion detection. Therefore, we refer to this new virtual computer laboratory as Lightweight-CVCLAB (LCVCLAB). We have implemented a prototype of the LCVCLAB based on open source technologies and a free version of the VMware vSphere Hypervisor. Our initial trials show that the LCVCLAB is a promising approach for providing students with hands-on experiences in a secure environment with a minimum deployment cost.

II. Infrastructure of the CVCLAB and Justifications for the LCVCLAB

Currently, the CVCLAB is implemented using VMWare vSphere 6.0 technology and hosted by several Intel(R) Xeon(R) CPU E5-2650 servers, each with 32 cores, 64GB Memory, and 8 network interface cards. The CVCLAB includes several virtual computer laboratories to support

various learning objectives. Figure 1 illustrates the logical topology of the CVCLAB used in our information security courses. In the CVCLAB, a student controls four different types of virtual machines: Windows 7, Windows Server 2008, Backtrack Linux, and Kali Linux. Student virtual machines are organized into two separate Local Area Networks (LAN), and a third LAN includes target virtual machines. All virtual machines and user accounts are administered centrally via a VMWare vSphere Manager.

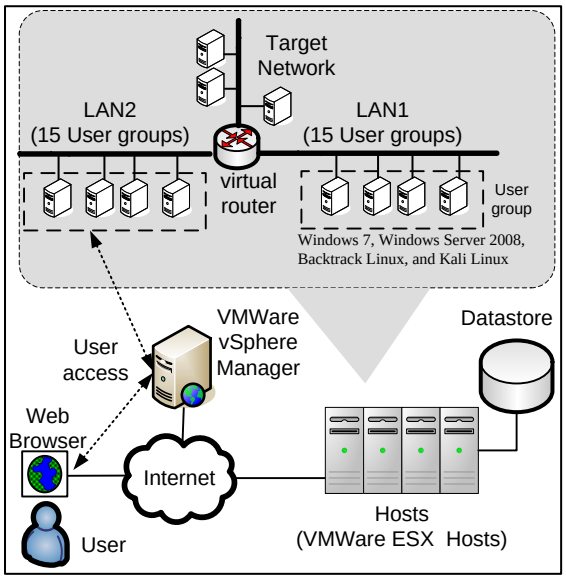

Figure 1. Logical Diagram of the CVCLAB

We studied the effectiveness of the CVCLAB in several controlled empirical field experiments. For example in an introductory-level database the class, 97 students performed a database security hands-on activity and were asked to evaluate the CVCLAB. To understand the problems that students faced in using the CVCLAB, they were asked: "What did you like the least about the CVCLAB?" We first performed a sentiment analysis to identify common themes and patterns emerged in the student responses to this question. Based on the extracted themes, we created several categories and assigned individual student responses into these categories. Figure 2 presents the percentage of the responses in each identified category. As seen in the figure, the slow response time of virtual machines emerged as the third major concern by the students. In addition, the other two concerns, "Long time to complete the activity" and "Technical problems", are also somewhat related to the slow response time of virtual machines. In two other empirical studies[8,9], we also observed similar student concerns.

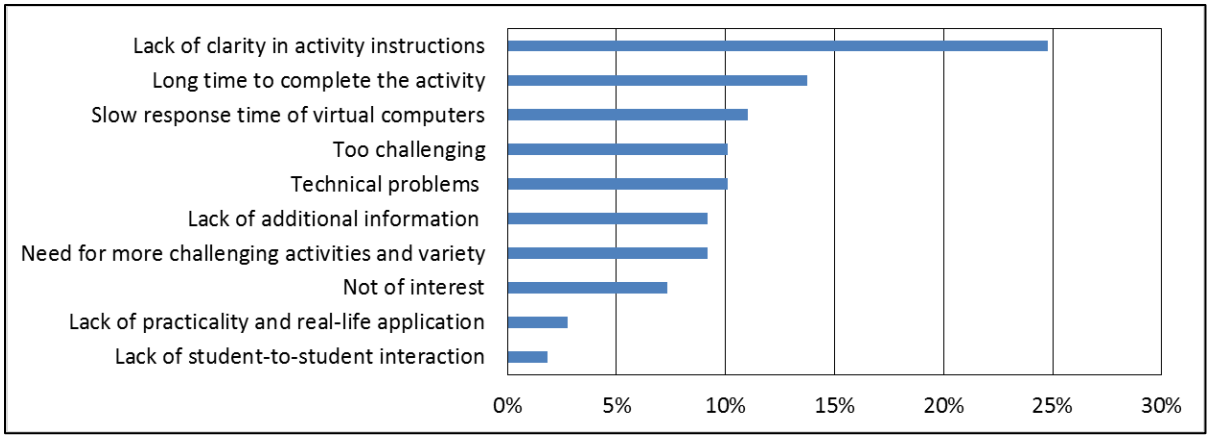

Figure 2. The frequency of student responses to the question: "What did you like the least about the CVCLAB?"

The hardware infrastructure hosting the CVCLAB has a total resource pool of 215 GHz CPU, 281 GB Memory, and 10TB network attached storage. The allocated resource pool is quite adequate to support 30 concurrent users. Therefore, the hardware capacity could not possibility be the main cause of the slow response time of virtual machines experienced by the students. Our preliminary analysis also showed that the slow response time was caused by connecting the desktops of virtual machines over the Internet. Slow Internet connection speeds and network outages significantly affect students' experiences with the CVCLAB. In fact, students reported increasing technical problems and experienced slower response times when we upgraded to the vSphere Web Client from the vSphere Desktop Client.

Although students can access to the desktops of their virtual machines, the majority of our cybersecurity hands-on activities are performed through a terminal window. In order words, such hands-on activities do not require a graphical user interface or accessing the desktop of a virtual machine directly. Furthermore, it is also important for information security students to get familiar with using the command prompt interface.

III. Description of the LCVCLAB

Figure 3 illustrates the logical diagram of the LCVCLAB. Unlike the CVCLAB, the LCVCLAB is hosted by a single VMWare ESX server and does not require on a VMWare vSphere Manager for the management of user accounts and virtual machines. The Access Manager, which is also a virtual machine (Ubuntu Server 16.1), is configured as a perimeter router/firewall that controls all accesses to the LCVCLAB. The Access Manager has two interfaces, one connected to the Internet and the other connected to the virtual machine Local Area Network (LAN) as shown in Figure 3. Students connect to the Access Manager via an SSH client using a stripped-down user account with very limited privileges. After a student gains access to the Access Manager, the next step is connecting a virtual machine residing on the virtual machine LAN. Students connect

to these virtual machines via SSH as well, but this time using the root account so that they are able to use all cybersecurity applications available in these virtual machines without any restriction. Currently, all student virtual machines are Kali Linux, which is a distribution of the Linux Operating System specifically configured for digital forensics and penetration testing with hundreds of pre-installed applications.

The virtual machine LAN is a virtual network created within the virtual machine host and is not connected to the campus network. In other words, students can access Kali Linux virtual machines remotely from anywhere with full administrative privileges on them, but they cannot use them to gain access to the campus network. Therefore, they do not pose a security threat to the campus network.

The third part of the LCVCLAB is the target network. Students do not directly access and use the virtual machines on the target network. These virtual machines are configured with security flaws so that students can analyze their vulnerabilities and exploit them in several hands-on activities. The target and virtual machine LANs are connected through a virtual router/firewall that also blocks some types of traffic between the networks. This topology allows us to emulate more realistic network configurations and enables more advanced hands-on activity capabilities.

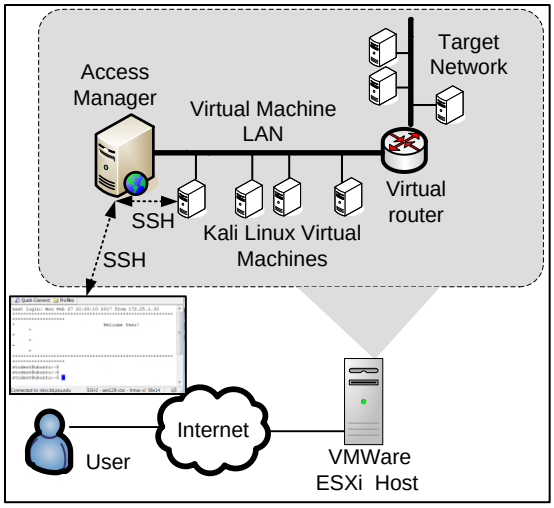

Figure 3. Logical Diagram of the CVCLAB

A major advantage of the LCVCLAB is its relatively low cost of building and maintaining the laboratory with respect to the CVCLAB where students use a client to access the desktops of virtual machines directly. When a remote client is used to gain the access to the desktop of a virtual machine, only a single user can use the virtual machine at a time. Therefore, an individual virtual machine should be assigned to each user in the CVCLAB, which increases the number of virtual machines required as well as the hardware resources to support them. On the other hand, the SSH interface allows multiple users to log into the same virtual machine

simultaneously. Therefore, we do not need to create as many virtual machines in the case of LCVCLAB as the number of concurrent users. This reduces the need for hosting the virtual laboratory by multiple powerful servers. In addition, the user account management is simplified since the Access Manager handles the user account management. There is no need for an additional server to manage the virtual machines or user accounts. Reducing the number of required virtual machines also simplifies the management and the storage requirement of the laboratory. Finally, the SSH protocol requires limited network bandwidth. Therefore, students even with very slow Internet connections can still use the LCVCLAB effectively.

The SSH interface is also the major drawback of the LCVCLAB because it limits types of hands-on activities to ones that can be performed only through the terminal window. Therefore, GUI-based cybersecurity applications are not available in the LCVCLAB or they are used through the terminal window. However, this drawback also motives students to get familiar with Linux command prompt tools, which is an important skill sought by many cybersecurity organizations. Table 1 illustrates some of the hands-on activities available for the LCVLAB.

Table 1. The list of the CVCLAB hands-activities that can be performed in the LCVCLAB

Topic	**Hands-on Educational Material Topics**
Foundation of Computer Networking	Network Addressing (IP, MAC, Ports, ARP); Network Management and Diagnostic Tools;
Cryptography	OpenSSL; Symmetric and Asymmetric Algorithms; Public Key Infrastructure; Key Exchange; Hashing;
Internet Security Technologies	Digital Signatures; Digital Certificates; Authentication and Authorization
Penetration Testing	Target Discovery and Enumeration; Vulnerability Assessment;
Network Attacks	Spoofing: IP and MAC Addresses; Denial of Service Attacks; Password Attacks; Privilege Escalation;
Protecting Linux Hosts	Linux Accounts and Permissions; Linux Services and Processes; TCP Wrapper; IP Tables;
Intrusion Detection	Honeypot; Host-Based Intrusion Detection System; Network-Based Intrusion Detection System;

IV. A Hands-on Exercise Example

In this section, we present a hands-on exercise to demonstrate the capabilities and the use of the LCVCLAB. The exercise is about computer penetration testing, which is systematically evaluating the vulnerabilities of a computer system or a network by simulating attacks from internal and external threats. Penetration testing is a time demanding and tedious process as illustrated in Figure 4. The penetration testing process starts by creating an inventory of systems

on a target network (target discovery). The next step is collecting information about the target systems (target enumeration). The final step is analyzing the target systems for potential vulnerabilities (target assessment). The exercise mimics the penetration testing process and can be performed in both face-to-face and online courses. In each phase of the exercise, students first learn about the tools and methods utilized and then put their learning into practice by performing a penetration testing of the target network.

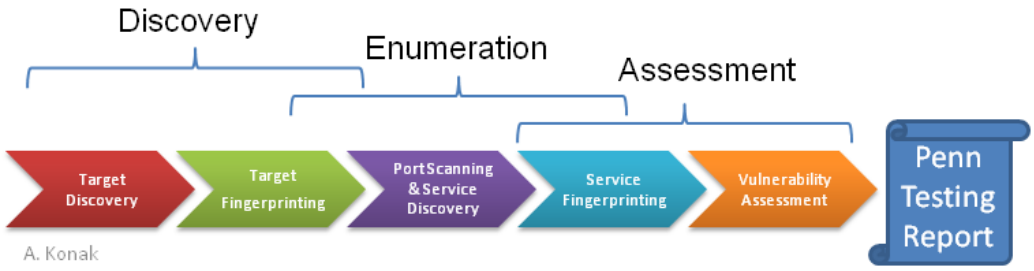

Figure 4. The process of penetration testing applied in the exercise.

Phase 1. Target Discovery

In this phase of the exercise, students are expected to create an inventory of the systems that they have discovered in the target network (discovery) and collect data about their operating systems (fingerprinting). The virtual machines on the target network include various types of OS and configured in a way that it is not always straightforward to discover them (i.e., they might not respond to ping requests). In addition, students have no pre-knowledge about the target systems. Students are first introduced to several tools and methods for target discovery and fingerprinting. This part of the exercise is formatted as a step-by-step tutorial demonstrating the use of the following tools:

- arping - a tool is used to discover a host on a LAN by sending Address Resolution Protocol (ARP) requests.
- fping- a program that can send ping request to multiple systems simultaneously.
- netdiscover- a program that can discover all alive hosts using ARP on a LAN.
- nbtscan- This is a command-line tool that scans for NETBIOS nameservers on a network.
- p0f – a passive OS fingerprinting tool that analyzes the network traffic without creating any extra packets.
- xprobe2- an active OS fingerprinting tool uses probabilistic approaches to guess the target machine's OS.

After completing the tutorial, students are instructed to discover systems on the target network. In this stage, no specific directions are provided, and students are allowed to use any set of the tools that they have learned.

Phase 2. Target Enumeration

This phase of the exercise involves gathering information about the services running on the target systems. At the end of the phase, students are expected to create a report summarizing the status of the ports and the services running on each system that they have discovered in the previous phase. Nmap, which is a powerful open-source network scanner, is the main tool used in this phase. First, students complete two tutorials to prepare them for the required tasks. The first tutorial is about basic port scanning using Nmap. The learning objective of the first tutorial is to describe various states of TCP ports as they are reported by Nmap (e.g., open, closed, filtered, unfiltered etc.). The second tutorial introduces different scan types that can be performed by Nmap, including TCP SYN, TCP Connect, UDP, TCP Xmas, and TCP Ack. The learning objective of the second tutorial is to identify which scan types are appropriate for different types of scenarios.

Phase 3. Vulnerability Assessment

Vulnerability assessment involves finding system flaws that have the potential to be exploited by attackers. These flaws must be identified and fixed before attackers discover them. A system might have applications running on its open ports. However, having open ports does not mean that a system can be exploited. In most cases, vulnerabilities are caused by outdated versions of applications or misconfigurations. In this phase, students are expected to figure out whether the discovered systems have such vulnerabilities. Students use host-based (Lynis) and client/server-based (OpenVas) vulnerability scanners to evaluate the vulnerability of the discovered systems. They create a final report summarizing all critical vulnerabilities that they identify and make recommendations to fix them.

V. Conclusions

In this paper, we presented our approach to building a cybersecurity virtual computer laboratory (LCVCLAB) in which students can practice skills and tools that are required for cybersecurity jobs. Primarily, the LCVCLAB aims to support penetration testing hands-on activities, which are difficult to perform in traditional computer laboratories due to strict information technology policies imposed by universities. The proposed approach is independent of the virtualization technology utilized, and it can be implemented by any virtualization technology. The user interface depends on the command prompt interface and SSH. Therefore, the LCVCLAB can be effectively used even with a poor Internet connection and requires less computing resources for the server hosting the virtual environment. Therefore, the LCVCLAB is cost effective and efficient. The main disadvantage of the system is being limited to the command prompt tools.

References

1. Gaudin, S. (2007). Security Breaches Cost $90 To $305 Per Lost Record. *InformationWeek*. April 2007.

2. Bureau of Labor Statistics (2017). Information Security Analysts. Retrieved from https://www.bls.gov/ooh/computer-and-information-technology/information-security-analysts.htm.

3. Richards, R., Konak, A., Bartolacci, M. R. and Nasereddin, M. (2015). *Collaborative Learning in Virtual Computer Laboratory Exercises*. Spring 2015 Mid-Atlantic ASEE Conference, 2015 Villanova University,1-13.

4. Konak, A. and Bartolacci, M. R. (2016). Using a Virtual Computing Laboratory to Foster Collaborative Learning for Information Security and Information Technology Education. *Journal of Cybersecurity Education, Research and Practice*, Vol. 2016: No. 1, Article 2, 1-25.

5. Konak, A. and Ryoo, J., and Kulturel-Konak, S. (2014). *Student Perceptions of a Hands-on Delivery Model for Asynchronous Online Courses in Information Security*. ASEE Mid-Atlantic Section Fall 2014 Conference, Swarthmore College, Swarthmore, PA,1-7.

6. Konak, A., Bartolacci, M. and Huff, H. (2012). *An Exploratory Factor Analysis of Student Learning in a Collaborative Virtual Computer Laboratory*. Proceedings of AMCIS 2012 Seattle, WA,1-8.

7. Konak, A., Clark, T. and Nasereddin, M. (2014). Using Kolb's Experiential Learning Cycle to Improve Student Learning in Virtual Computer Laboratories. *Computers & Education*, 72, 11-22.

8. Konak, A., Kulturel-Konak, S., Nasereddin, M. and Bartolacci, M. R. (2017). Impact of Collaborative Work on Technology Acceptance: A Case Study from Virtual Computing. *Journal of Information Technology Education: Research* ,16, 15-29.

9. Wagner, K. G., Myers, M. C. and Konak, A. (2013). *Fostering Student Learning in Information Security Fields through Collaborative Learning in Virtual Computer Laboratories*. The Third Integrated STEM Education (ISEC), IEEE,1-7.

10. Konak, A., Bartolacci, M. R., Kulturel-Konak, S. and Nasereddin, M. (2016). *Impact of collaborative learning on student perception of virtual computer laboratories*. IEEE Frontiers in Education Conference (FIE),1-4.

Abdullah Konak, PhD., is a Professor of Information Sciences and Technology at Penn State Berks. Dr. Konak teaches courses in information security and data sciences.

Anuvrat Sheoran is a fourth-year student in the Security and Risk Analysis program at Penn State Berks.