Using Virtual Machines to Improve Learning and Save Resources in an Introductory IT Course

Geoff Stoker, Todd Arnold, and Paul Maxwell
Department of Electrical Engineering and Computer Science
United States Military Academy
West Point, NY
(845) 938 - 2193
{geoffrey.stoker, todd.arnold, paul.maxwell}@usma.edu

ABSTRACT
Information technology courses often require the use of software and hardware to support classroom learning. These systems can assist in achieving the learning objectives for a course through classroom problems and laboratory exercises. The procurement and maintenance of these systems can be a challenge even for well resourced organizations. In this paper we discuss how virtual machines can relieve organizations of some of their resource burdens while effectively achieving course learning objectives and provide examples of how that is currently done at the United States Military Academy.

Categories and Subject Descriptors
K.3.2 [**Computer and Information Science Education**]: Computer Science Education – *abstract data types*

Keywords
Information technology, virtual machines

1. INTRODUCTION
Teaching Information Technology (IT) often requires the use of software packages (e.g., web page servers, software development kits) and hardware systems (e.g., client computers, network switches) to reinforce learning objectives. To effectively teach IT subjects, the use of multiple operating systems, heterogeneous computers, and networks with various architectures is often required. The procurement, configuration, and maintenance of these packages and systems can present a significant expense in terms of money and personnel hours. In resource constrained organizations, this can be an insurmountable obstacle. Using *virtual machines* (*VM*), course designers can incorporate these important IT systems into their lesson plans while adhering to the resource constraints that are imposed upon them.

In IT courses, it is frequently useful to develop classroom or laboratory environments with a variety of configurations. Those configuration choices include items such as, operating systems, processor/machine architectures, networking architectures, and installed software. Manually configuring a set of machines to meet the needs of an IT course or courses consumes hours of technician time. Additionally, certain IT learning objectives may not be implementable using standard classroom equipment. Activities such as hacking and experimental networking configurations are often not acceptable to system administrators due to the risks involved for the rest of the network infrastructure. To be able to achieve the learning objectives in these situations, your choice is to have a specially configured and isolated classroom/lab or to use virtual machines. For introductory-level students, tasks such as installing and configuring software can be challenging. Additionally, control over a student's choice of processor and operating system for their personal computer is usually infeasible. If a particular software package is required for classroom learning then there exists the risk that the package will not function properly on the student's hardware platform with their choice of operating system. These two software issues create an extra burden for the instructor and the technical staff as they try to resolve the problem for the student.

One method to overcome many of these obstacles is to use virtual machines such as VMware Workstation and VMware Player [1][1]. Virtual machines provide a flexible and relatively inexpensive way to design and execute classroom and laboratory exercises without the purchase of hardware and the time cost for IT support personnel. Using virtual machines, classroom and laboratory exercises can be developed that have heterogeneous operating systems and processor architectures, relieves the requirement to purchase additional processors, and alleviates the vulnerabilities presented by executing IT lessons on standard machines. Additionally, virtual machines can allow an instructor to quickly reset or change the machine environment in a time constrained setting (e.g., between class hours) in a way that would be infeasible for a technician to reconfigure physical equipment.

In this paper we describe how virtual machines can be used in an entry-level IT course to support the course learning objectives. The use of virtual machines allows for the exploration of various environments found in introductory IT courses ranging from hacking labs to basic web site creation. We compare how the core Information Technology course at the United States Military Academy conducted its lessons prior to the incorporation of virtual machines to the current and future lesson plans. Included is a discussion on the advantages and disadvantages of using virtual machines in the IT classroom.

The remainder of the paper begins with a discussion of other efforts using virtual machines in section II. In section III, we

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
RIIT'13, October 10–12, 2013, Orlando, Florida, USA.
Copyright © 2013 ACM 978-1-4503-2494-6/13/10…$15.00.
http://dx.doi.org/10.1145/2512276.2512287

[1] The views expressed in this article are those of the authors and do not reflect the official policy or position of the Department of the Army, Department of Defense, or the U.S. Government.

discuss how virtual machines can be implemented in an entry level IT course. We then provide conclusions and suggestions for future work in section IV.

2. RELATED WORKS

The idea of using virtual machines for IT education is not new. In the last several years, many authors have discussed techniques for incorporating virtual machines into the framework of their courses. In general, these ideas have fallen into two main categories: using virtual machines to allow distance learning and using virtual machines to create networked environments for higher-level (i.e., junior, senior) IT courses.

Examples of work in the first category can be found in [2, 3, 4]. These papers describe using virtual machines to create Virtual Networked Labs (VNL) that students can access remotely using networked client systems. The main focus of these works is to enable students enrolled in distance learning programs to have a similar experience as those on-campus students who have access to a physical lab. The authors describe how they were able to establish virtual labs with similar capabilities as their physical labs and thus facilitate the remote students' learning. Our work does not fall into this category. We are promoting the advantages that can be obtained by using virtual machines in the on-campus environment.

In the work of [5] and [6], the authors describe using virtual machines in an on-campus environment. In their environment, virtual machines were developed to provide a laboratory for advanced courses in IT, such as, system and network administration, and information security and assurance. This virtual lab clearly enables students at the higher levels of IT education the opportunity to learn advanced topics without the need for expensive hardware and software. In our work, we extend this concept into the lower levels of IT education where the majority of students are not IT majors. We seek to leverage the benefits of virtual machines in that domain to conserve resources and augment student learning.

3. VIRTUAL MACHINE IN AN INFORMATION TECHNOLOGY COURSE

3.1 Background

At the United States Military Academy, every student is required to take a set of courses known as the core curriculum. The core curriculum is designed to support the Academy's educational goals of providing students with a breadth of knowledge to draw from after graduation. One of the Academy's goals focuses on educating our graduates to understand information technology and states: "Graduates understand and apply Information Technology concepts to acquire, manage, communicate and defend information, solve problems, and adapt to technological change." [7]

To support the Academy's IT goal, the Department of Electrical Engineering and Computer Science teaches two IT courses within the core curriculum. The first, IT105: Intro to Computing and Information Technology, is taken during freshman year, is taught to 99% of students, and covers basic computer functionality and programming concepts. The second course, IT305: Theory and Practice of Military Information Technology Systems, is most often taken during the students' third year which is after our students choose their major. Those who choose an engineering major meet the IT goal through courses taken within their major. As a result, IT305 is taken by students who select a non-engineering major, which is about 80% of each graduating class.

The goal of IT305 is to provide these non-engineering students with an understanding of the underlying theory behind the types of hardware and software systems they will see upon graduation. Specifically, upon completing the course the students should be able to:

1. Understand the underlying physical and mathematical concepts relevant to IT
2. Understand the ways in which IT systems function
3. Articulate the methods for successfully employing IT systems
4. Demonstrate the effective use of IT to solve problems and make decisions
5. Understand the importance and implications of IT

To support these outcomes, the students are introduced to Web 2.0, databases, networking, and cyber security, in that order.

3.2 Using Virtualization in the Classroom

Accomplishing the goals of IT305 requires the use of specific software packages and hardware systems. To standardize these tools and teaching methods, we require all of our students to install and use the same development tools. For the web page portion of the course, we use Microsoft Expression Web 4.0.

Expression Web is not part of the standard software package provided to each student, so the first course challenge is getting all of the students to download and properly install the software. Since the population we are targeting is the non-technologically inclined, there are sometimes significant challenges in getting everyone to a stable state where we can begin the course work. Every semester we encounter students who are challenged by the installation process or who experience software or hardware compatibility issues. The result is our instructors and computer support technicians spend numerous hours each semester attempting to resolve these issues so the students can commence with the learning process. One particularly difficult challenge occurred when attempting to install Expression Web on the standard student Windows Vista image and a specific hardware platform (which was issued to the entire cohort group). Expression Web functioned incorrectly on the combination of hardware and operating system. The only way we were able to resolve the issue was when the entire year group was upgraded to Windows 7. The upgrade eliminated the Vista specific issue though the specific cause of the problem was never fully resolved. Avoiding this type of problem would save instructor and computer support technician time and alleviate student frustration.

A way to minimize the impact of introducing additional software to the student computers for each portion of the course would be to virtualize the software. This could be accomplished by the creation of a VM using a tool such as VMware Workstation [1] with the software pre-installed in a stable configuration. Students could then utilize the software regardless of their underlying system using VMWare Player or VirtualBox.

Another method to avoid installation issues is to virtualize the application. Tools such as VMware's ThinApp [8] can create a virtual executable that will run the desired software on a machine yet is more light weight than a full VM. The students simply download an executable package that runs the software. No

matter the technique, the students could be provided a course virtual solution at the beginning of the semester that has all of the software packages required for the semester.

In addition to the administrative difficulties of software installation and usage, IT courses are often required to use specialized hardware and software to achieve their learning objectives. For IT305, a major learning experience occurs during the networking block in three hands-on lab periods. In our profession, each of our graduates will experience using or operating a network in a military tactical environment. To simulate this they are required to construct the user portion of a tactical network (Figure 1), that is similar to a small office environment of about 40 employees. The devices include switches, laptops, and Voice over IP (VoIP) phones. We also introduce them to some network monitoring tools (e.g., SNMPc, Solar Winds Orion) they may use on their networks. Errors are purposely introduced into the network during the labs for experience with troubleshooting.

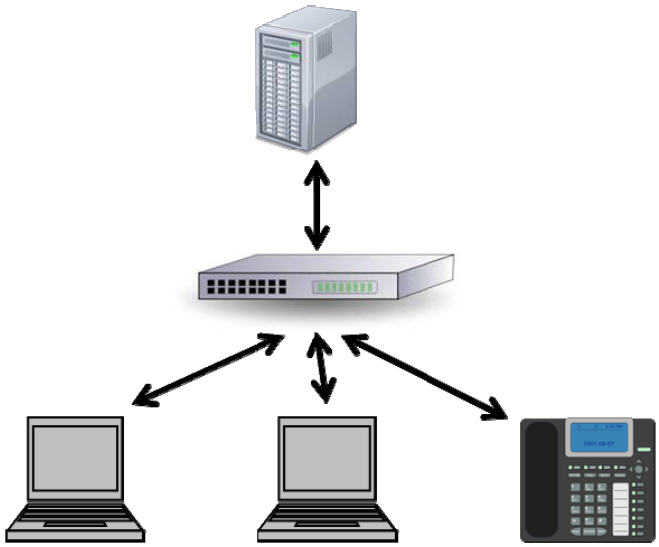

Figure 1. Example configuration of an IT305 networking lab. The lab connects different client systems to a server using network switches.

While valuable as a teaching tool, these hands-on lab experiences are costly, both financially and physically (i.e., scheduling classroom resources). Similar to many institutions, we do not have enough dedicated networking lab space for constructing these networks, so we must reserve eight additional classrooms each semester just for the labs. This means other courses cannot use the spaces during those class hours for the entire semester thus constraining the academy. Additionally, to support the labs, each classroom requires dedicated networking equipment. This is a costly investment, especially because the equipment is used only during these three lab periods each semester.

Instructor time, another constrained resource, is also impacted with this lab configuration. Because the networking equipment is only used a few times each semester, every instructor must construct the lab just as the students will to identify faulty equipment before the lab period. Before each lab hour, the instructors must also place the equipment in the room because the equipment cannot be left unsecured in the classrooms. Furthermore, the lab rooms are often used by other courses during the hours immediately preceding the lab times and thus unavailable for conducting this set-up procedure. Many times, instructors are only allowed the 15 minutes between classes to set up the entire lab worth of equipment. This constraint can result in lab exercises that do not fully meet their desired learning objectives.

Virtualization of the hardware and software used in these lab exercises could reduce or eliminate many of these issues. The basic networking labs could be taught with a focus on basic functionality, configuration, and connectivity. The clients and servers used in the labs, to include the phones, can be virtualized and run on the students' laptops in their regular classroom thus freeing the institutional resources and requiring fewer pieces of dedicated hardware. Lab preparation can then be as simple as connecting networking cables between student laptops and a classroom switch. Using virtual machines and virtual network simulators, such as Dynamips [9], students would be responsible for configuring network devices during the labs. This virtualization would allow each student to connect, configure, and troubleshoot the entire network while doing so in a less resource intensive environment.

The Cyber security portion of the course is a block of five lessons, with four of the lessons providing hands-on experience. The first block introduces the students to the Army's Blue Force Tracking (BFT) system [10]. Blue Force Tracking is used by the Army to digitally track its assets (e.g., tanks, helicopters) on the battlefield, providing units with real-time location information and messaging capability. This system is fielded Army-wide and is a critical Army IT system to the deployed force.

Rather than just providing the students with an overview of the system, we want experiential learning to occur. However, each system is a stand-alone Linux-based device and is only allowed on a tactical network. We are not authorized to connect the devices to the classroom network. In addition, each of these BFT systems is costly and requires access to a satellite network for proper operation.

To provide the students with access to the system, we created VMs with the BFT software pre-installed and were given an exception to connect the VMs to the network on a temporary basis. Using these virtual machines, each classroom is able to construct a tactical network and simulate a deployed unit as it moves across the battlefield using the virtual BFT interface. The VM provides the same functionality and appearance of the real system without the expensive hardware.

The other three lessons of the cyber security block provide lessons on reconnaissance, defensive, and offensive cyber operations. Based on end of course feedback each semester, the hacking labs really capture the imagination of many of the students and is at least interesting to all of them. The context of the hacking block is designed to illustrate important points about the risks associated with operating personal computing devices on public networks like those found at airports or public restaurants.

Once again, restrictions on connecting either vulnerable systems to the network or on installing 'hacking' tools onto networked computers poses a challenge for IT courses. Institutional network administrators are responsible for the entire network and therefore are rightfully resistant to allowing these types of systems on the network. This reasonable constraint poses a challenge for those attempting to impart IT lessons to students. A solution to this challenge is to virtualize the lab. As a result, we distribute two

VMs to the students that contain virtual attacking and defending machines connected over an isolated virtual network.

In the first lesson of this virtual lab, we familiarize students with the concept of a virtual machine and spend some time emphasizing the difference between applications that are running on an installed operating system, such as Windows 7 which is installed on their individual laptops, and the virtual machines that are complete operating systems running logically within their laptops. Each student makes use of an Attacker machine and a Defender machine, both of which are running an installation of un-patched Windows XP Pro. We use the image in Figure 2 to help them understand abstractly how the three machines co-exist.

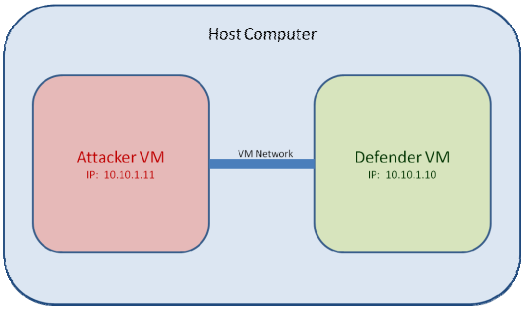

Figure 2. System configuration for IT course hacking laboratory exercise.

Once the students start the VMs, they see something similar to Figure 3. The VM on the left with the red desktop and many shortcut icons is the Attacker and the VM on the right with the blue desktop is the Defender.

During the first hacking lab we expose the students to the concepts of vulnerabilities, threats, and risks and set the stage for a cyber attack with reconnaissance of the defender machine. Because the scenario is of a migrant user temporarily making use of a public network, we stress that whatever protection a user can reasonably rely on must be present on their individual PC. Simple steps, like having a firewall and having it enabled are important ones. We use the specific tools SuperScan v4.0 (McAfee) and Winfingerprint 0.6.2 (available SourceForge) to show students that information about their computer name, the current Windows version and patch level, usernames, open ports, etc., are readily available over the network from any machine that is open to probing. This type of information is covertly collected by the Attacker about the Defender for use in subsequent lessons.

During the second hacking lab lesson we illustrate how the information gathered about a target machine can be used to find an entry point into it. We start with simple ideas like trying to guess a password for one of the accounts we discovered and show how we can leverage the power of computers to automate and accelerate the guessing process. During this task, which uses a batch file and a dictionary file, we are able to expose students to things very few have previously experienced, like using command line tools and MS-DOS commands to gain access across the network to another machine's C: drive. We extend the entry point search to using pwdump and John the Ripper to demonstrate how the process can be quickly expanded once access to a machine is gained. Just prior to this task we direct them to reset a password for one of the Defender machine's accounts to test how easy or hard it is for automated tools to guess students' passwords. We also talk about and demonstrate how tools like Nessus Vulnerability Scanner are available to help us find weaknesses and focus our defensive efforts.

The third and final lesson of this block is quite successful at exciting and alarming most of the students. Using the tools Sub7 by mobman – sent via Trojan attachment – NetBus, and Metasploit, students are able to have the Attacker machine take total control of the Defender machine. The impact of these three hands-on lessons is immense and would not be nearly as effective, or even possible, without the ability to use VMs to gain exposure to all these tools.

Conservatively, to conduct these hacking labs with real equipment would require 108 computers, 54 cross-over cables, and about 40 person-hours of technician support each semester to properly resource the labs and deal with any failures. These computers could be relatively low-powered ones that cost less than $500 each or be repurposed from life-cycle requirements. The cost in instructor or technician time to deal with student-caused configuration problems, or the loss for subsequent classes in instruction quality due to unusable computers is hard to gauge; we just note that it would be non-zero.

3.3 Implementation Challenges

There are many advantages to using VMs. Virtualizing labs and lessons can reduce the need for additional equipment; the space in which it must be set-up and/or stored; the instructor time required to set it up and tear it down if not permanently set-up; and the technician time needed to configure, maintain, and troubleshoot it. Given the skill-level of our student population, using VMs greatly reduces the complexity of teaching topics like the hacking labs by eliminating the time required for installation and configuration of the necessary software. The use of VMs also reduces the security risks to the primary school network and the risks of students creating problems with the starting configuration of the labs for each subsequent course section.

Of course, use of VMs brings its own challenges. Someone on staff must be competent at building and configuring the baseline VM required for lessons or tools that are being virtualized. The file size of VMs is not small, for example our Attacker machine is about 5 GB and the Defender machine about 3 GB. The issue of software licenses also cannot be ignored. Virtualization methods that rely on network connectivity, such as Citrix, are subject to network availability and throughput. Access to these networked virtual systems can pose a problem when students are off campus and require access to the software.

However, overall the advantages of using VMs greatly outweigh the disadvantages. As we move through the process of reducing and eventually eliminating dedicated computer labs, the ability to use virtualization to present key IT topics will be important in providing our students a very worthwhile hands-on experience while reducing the amount of resources required.

4. CONCLUSIONS AND FUTURE WORK

In conclusion, the use of virtual machines in introductory information technology courses can help students realize the course learning objectives. The virtual machines provide

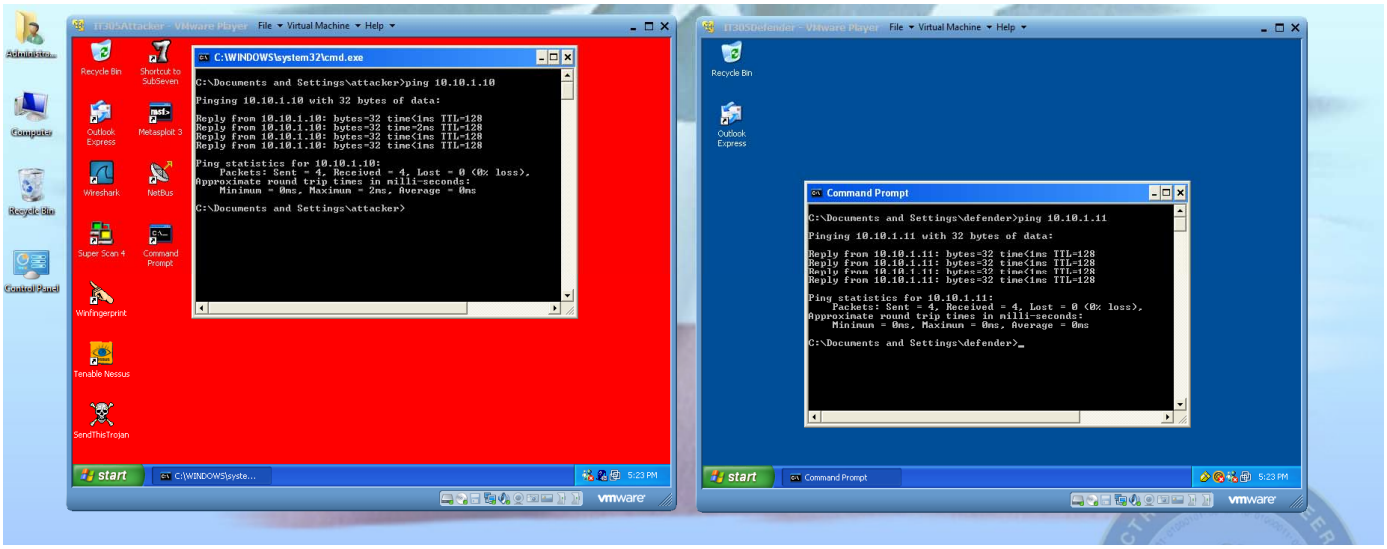

Figure 3. Screen capture of IT hacking laboratory exercise. Left window shows the attacking machine and the right window shows the defending machine.

educational tools to assist with the students' learning while alleviating some resource requirements, such as, heterogeneous machines, heterogeneous operating systems, and technical support hours. Additionally, the VMs can reduce student stress and save them time by removing issues that arise from incompatible software and hardware combinations. Using the United States Military Academy IT305 course, we have demonstrated how virtual machines can be integrated into the introductory course material. The use of virtual machines in introductory courses can also facilitate their use in advanced IT courses due to student familiarity with the tool.

Future work in this domain could include a detailed comparison of courses using traditional teaching methods versus a virtual machine based course. Additionally, an analysis on which IT topics are best suited for VMs would be useful to allow course designers to focus their efforts in those areas.

5. REFERENCES

[1] VMware, http://www.vmware.com/, accessed 30 April 2013.

[2] C. Border, "The development and deployment of a multi-user, remote access virtualization system for networking, security, and system administration classes," *Proceedings of the 38th SIGCSE Technical Symposium on Computer Science Education* (SIGCSE '07), ACM, New York, NY, pp. 576 – 580.

[3] L. Leitner and J. Cane, "A Virtual Laboratory Environment for Online IT Education," *Proceedings of the 6th conference on Information Technology Education* (SIGITE '05), ACM, New York, NY, pp. 283 – 289.

[4] P. Li, L. Toderick, and P. Lunsford, "Experiencing virtual computing lab in information technology education," *Proceedings of the 10th conference on Information Technology Education* (SIGITE '09), ACM, New York, NY, pp. 55 – 59.

[5] W. Bullers, S. Burd, and A. Seazzu, "Virtual machines – an idea whose time has returned: Application to network, security, and database courses," *Proceedings of the 37th SIGCSE Technical Symposium on Computer Science Education (SIGSCE '06)*, ACM, New York, NY, pp. 102-106, 2006.

[6] M. Stockman, "Creating remotely accessible 'virtual networks' on a single PC to teach computer networking and operating systems," *Proceedings of the 4th conference on Information Technology Curriculum (CITC4 '03)*, ACM, New York, NY, pp. 67 – 71.

[7] Office of the Dean, "Educating Future Army Officers for a Changing World," 3rd ed., West Point, NY.

[8] VMware, *Thinapp User's Guide*, http://pubs.vmware.com/thinapp4/help/wwhelp/wwhimpl/js/html/wwhelp.htm, accessed 24 May 2013.

[9] Graphical Network Simulator, *Dynamips*, http://www.gns3.net/dynamips/, accessed 27 May 2013.

[10] FBCB2 – Blue Force Tracker, http://en.wikipedia.org/wiki/Blue_Force_Tracking, accessed 27 May 2013.