Proceedings of the 10th Colloquium for Information Systems Security Education
University of Maryland, University College
Adelphi, MD June 5-8, 2006

Evolution of the ASSERT Computer Security Lab

Brian Hay, Kara L. Nance, *ASSERT Center, University of Alaska Fairbanks*

Abstract – These times of declining academic budgets coupled with increased demand for information assurance professionals presents unique challenges for academic departments wishing to build capacity in information assurance. This paper discusses the evolution of the Advanced System Security Education, Research, and Training (ASSERT) Lab at the University of Alaska Fairbanks. The effort began with the low cost construction of a proof-of-concept dedicated information assurance lab that was then used to leverage additional funding to build a high capacity research and educational environment to meet the needs of the students, faculty, and researchers who now utilize this vital facility.

Index Terms – Information Assurance Laboratory, Computer Security, Educational Environment, ASSERT

I. Introduction

As is true of many disciplines, hands-on experience is a vital component of a thorough computer security education [HDR03]. Although computer science programs typically have lab facilities available for student use, these labs are generally not well suited to IA activities which often involve the deliberate use of tools and malicious code which the administrators of production networks go to great lengths to prohibit. As such, alternate, often dedicated, lab environments are extremely valuable for IA teaching and research. This paper will describe the successful evolutionary efforts undertaken by the University of Alaska Fairbanks (UAF) over the past 4 years to create and refine the Advanced System Security Education, Research, and Training (ASSERT) Lab.

II. Motivation

Alaska faces challenges as a result of its geographic isolation at two levels, namely that the state as a whole is geographically unconnected to the rest of the country, and that many communities within the state, including the state capital, are not on the road system. As such, the Internet as a provider of access to commercial, governmental, and educational services is of vital importance to Alaskans. Any disruption to networked services can have a much more severe impact in Alaska than other states, as the option for alternative methods of content delivery are severely limited. The University of Alaska statewide system is the major provider of higher education in the state, and relies heavily on various networked systems to assist in fulfilling that responsibility. This includes connectivity to research partners around the world, and to students in remote villages who communicate via its extensive distance education programs.

In addition to networked educational opportunities, reliable computer and network systems allow businesses in Alaskan communities to serve statewide, national, and global customers on a level playing field, providing opportunities for employment within the state, even in rural communities. As such, it is imperative that UAF provide students with the skills to support the IT infrastructure within the state.

Therefore, UAF, as the main research and computer science institution in the state and a Center of Academic Excellence in Information Assurance Education (CAEIAE), has undertaken the development of a program which will provide education and research opportunities in the IA domain, which is so vital to the future of Alaska. In support of this effort, the Computer Science Department currently houses the ASSERT Lab, which provides an isolated testing and research environment.

III. Surplus Equipment – The First Step

The first steps towards creating a lab environment were taken in 2001, although at that time virtually no funding was available for the purchase of equipment. Like many institutions, UAF maintains a warehouse where surplus equipment is stored and reassigned to other departments. Although such equipment is rarely state-of-the-art, it is not uncommon to find computer systems in surplus which are relatively modern, as faculty and general use lab systems are usually replaced every three years. While such systems do not feature the latest processors or largest hard drives, they are frequently suitable for use as IA lab computers. In our case we were fortunate to acquire 10 homogeneous systems, which made management of the lab easier, but this is not a requirement for a successful lab. The surplus facility also proved to be a useful source of other components, such as spare Ethernet cards and additional RAM, which were used to augment the surplus computers to enhance their capabilities. The surplus warehouse also provided as many CRT monitors as could be used, as many faculty and labs had recently transitioned to LCD panels in their offices.

ISBN 1-933510-98-6/$15.00 © 2006 CISSE

The surplus systems were connected using a small SOHO switch and a few hundred feet of Ethernet cable, and for a cost of less than $200 a basic but functional lab was created. The surplus warehouse also provided a set of thirty identical 1GB IDE hard drives, which are still used in the lab to this date. While such drives serve almost no useful purpose for storage in modern computer systems, they have been very useful in computer forensics exercises. For example, the disks can be preloaded with a given configuration, and then distributed to students for forensic analysis. The small size of these disks allows the principles of forensic investigation to be experienced without subjecting the student to overwhelming large volumes of data. The ASSERT faculty continue to visit the surplus warehouse to supplement the ASSERT Lab capabilities.

IV. Network Isolation

During the initial design and development of the lab it was decided to physically isolate the lab network. During discussion with UAF administrators and attorneys the faculty made a considerable effort to ensure that any view of the ASSERT lab as a “hacker training lab” was corrected, but the fact that such misunderstandings were present made it clear that any breach of the university networks, no matter how accidental, could jeopardize the continued operation of the lab, which was deemed vital to successful IA education and research. As such, the network connections from the lab room to the campus network were disconnected both at the switch and physically at the wall ports in the lab, so that no possible configuration of the lab systems could result in an impact on the external networks.

This isolation continues to be a vital part of the current lab design, which does result in some minor inconveniences for lab management. However, the removal of the possibility of a headline such as “Hacker Lab Shuts Down Campus Network” in the local newspaper, and the resulting impact to the program, seems to make such inconveniences justified.

One negative impact of the network isolation is that no internet-based research is possible while in the lab, which was a frequent issue raised by students. As a compromise, a single workstation in the lab was configured with a wireless card and connected to the campus wireless network. The use of a wireless connection was chosen so that there was no need to have a live network cable running to the lab. This workstation is physically separated from the other systems in the lab, and has had its Ethernet card removed so that it cannot be connected to the lab network, and currently provides students with access to the Internet without jeopardizing the isolation of the lab network.

The second issue that arose was the inability to get tools and software patches from external sources onto the lab systems. This became particularly important as the lab evolved to include various administrative systems. The current solution to this issue is to script the retrieval of updated software using an external system, followed by the creation of a CD-R disc containing the newly acquired software. This CD is then brought into the lab, and the various packages are installed in the appropriate locations, which is again a scripted activity. These updates are typically performed on a weekly basis, and through the use of scripting can be performed with minimal human intervention, other than the physical task of carrying the disk to the lab and initiating the update process.

The network isolation does not offer any direct capability for students or researchers to get files out of the lab, but this has not been a significant issue. Student assignments can be submitted to instructors within the lab environment, which currently hosts an email system, or alternatively a CD can be created within the lab and removed if necessary, provided that the individual understands that an appropriate level of care is taken to ensure that the CD does not carry malicious logic, for example, to external systems. In general, the results of work performed in the lab are used within the lab and can therefore be stored in the user’s home directory in the lab, so there is usually no need to remove data.

V. Virtualization – Host Based Images

Although the initial lab was very useful, it did suffer from some limitations. The most problematic was that of recreating a known configuration, either for a given system, or for a group of systems. In a lab in which users generally required full control of the machines it was very difficult to determine the state of a system at any given time. To partially address this issue the first administrative server was added to the lab, which was again acquired from the surplus warehouse, although several additional hard disks were added to its initial configuration. This server was controlled by the lab administrator, and it provided several network accessible disk images which could be easily loaded onto the lab systems, allowing multiple base configurations to be created and then loaded onto any of the lab systems within 10 or 15 minutes.

After the initial lab configuration was created and proof-of-concept demonstrated, two sources of funding (an NSF capacity building grant and UAF internal funding) were available for new lab equipment. At that time eight new workstations were purchased for the lab, in addition to KVM devices, two Cisco 2950 Series switches, and a Cisco 2600 Series Router. This allowed the number of physical systems in the lab to be expanded, and the Cisco network devices allow more complex network topologies to be created.

The other major upgrade that was possible at that time was the purchase of virtualization software, specifically VMware workstation [Vmw06b]. When added to the new workstations with their 2.5GB of RAM, the virtualization software allowed several virtual machines to run on each of the new physical workstations. The configuration of the new workstations was modified so that each lab user had minimal permissions, and as such was incapable of modifying the workstation configuration. Once logged into a workstation, each user could then run one or more virtual machines over which they had full root/administrative control, which could therefore be configured to meet their particular needs.

Virtualization brought several major benefits to the lab, including the ability to create standard configurations for virtual machines, which could then be essentially cloned and used by lab users by simply copying the contents of a virtual machine directory, as opposed to actually installing an operating system and software on a physical machine. The *snapshot* feature, in which the state of a virtual machine can be saved and later restored, also proved to be very useful to lab users.

While the inclusion of virtualization was a major improvement to the lab, the initial lab configuration stored the virtual machine (VM) images locally on a given workstation, meaning that a user was essentially bound to a particular physical machine. Since there were more lab users than workstations, it was often the case that there would be free workstations but a given a user would be unable to access their virtual machine images because the physical workstation on which his or her assigned VM images were stored was in use.

VI. VIRTUALIZATION – NETWORK BASED IMAGES

The purchase of a new server for the lab, in addition to several hard drives, provided the option to make the VM images available over a network, rather than storing them locally on a given workstation, allowing users to operate their virtual machines on any available physical workstation in the lab. This approach provided flexibility in the use of the lab, but also resulted in reduced performance of the virtual machines, which was particularly significant during the startup, shutdown, and snapshot of the virtual machines. When multiple users attempted to start or snapshot VMs simultaneously, it could take up to 10 minutes to complete the operation, during which time the user simply had to wait. Since the use of network based images was intended to allow the lab to be utilized by more simultaneous users, this excessive wait time in large part defeated the purpose of the network accessible images.

While we suspected that the 100Mb Ethernet network was the primary bottleneck in this process, it was important to confirm that to be the case prior to attempting to solve the problem. The *ntop* tool [Nto06] was particularly useful in that regard, and did indeed confirm that the network quickly became saturated when more than two or three virtual machines were performing power on, power off, or snapshot operations.

The solution to this issue was to upgrade to Gigabit Ethernet, which required the purchase of a Gigabit switch. While the workstations and the server each already had Gigabit Ethernet NICs, two additional Gigabit NICs were added to the server. The three server NICs were then bonded into a single channel to the Gigabit switch to reduce congestion to and from the server. The resulting configuration, as shown in figure 1, provides a 3Gb/s channel from the server to/from the switch, and a 1Gb/s channel from the switch to each workstation. The use of this network configuration resulted in greatly reduced delays during the power on, power off, and snapshot operations for virtual machines, all of which are now commonly completed within 30 seconds, even during periods where all workstations are in use. In addition, the performance of the virtual machines during normal use (other than the three particularly network intensive operations previously identified) does not appear to be significantly different than running from a local image.

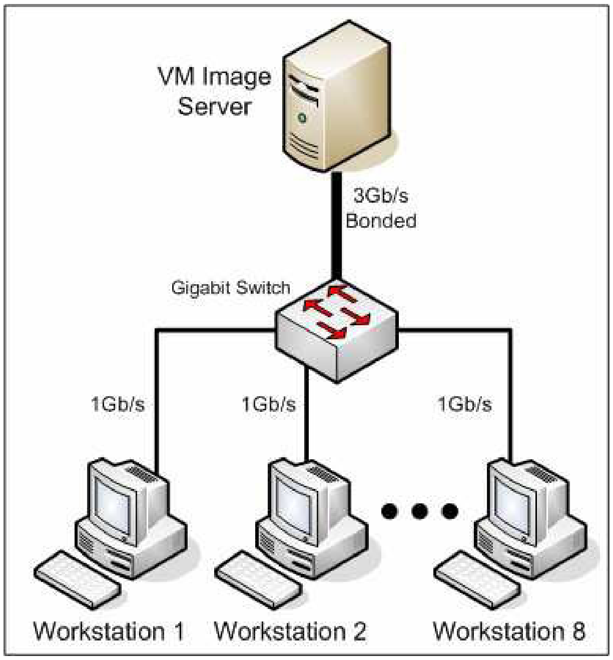

Figure 1: Configuration of the Gigabit Ethernet connection between the VM image server and the workstations.

Another advantage of network based images versus images stored locally on a workstation is that it is easier for an instructor to propagate a virtual machine to each of the students in a course. In the ASSERT Lab, instructors can create a VM images in their accounts, which are saved across the network to the server. Using a predefined script they can then create copies of the new images in the accounts of each of their students, providing each member of the class with access to a machine with a known configuration without any required action on the part of the student. As the lab use increases, it has become common to find students with several VMs in their accounts, each of which may be targeted towards a particular course, or even an assignment within a course.

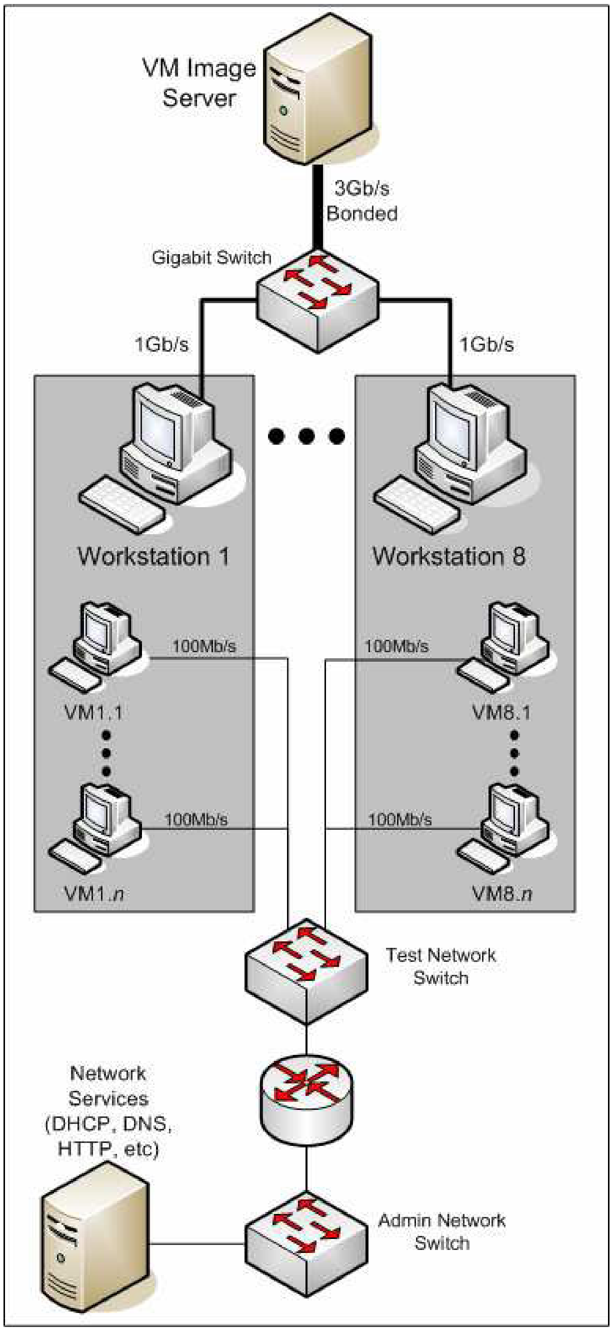

Figure 2: Intermediate Network Configuration.

VII. Network Components

In addition to faculty and staff researchers, the ASSERT lab was utilized by an increasing number of students, and this increased usage made an additional round of funding available from UAF, with the aim of providing a flexible networked environment within the ASSERT Lab. The network configuration at the time consisted of two Cisco 2950 Series switches, a Gigabit Ethernet switch, and a Cisco 2600 Series router, connected as shown in figure 2. While this configuration did allow VM images to be retrieved quickly, and isolated all VM network traffic to its own network, all of the network components were essentially administrative devices, in that no user reconfiguration was allowed as it would have disrupted the network for other lab users. While VMware does offer some level of virtual networking, in that virtual machines on a given workstation can interact with each other on one of seven virtual networks, it did not offer the opportunity to experience network configuration of the complexity seen in even low end commercial networking hardware.

The UAF funds obtained at that time were used to purchase six Cisco 2600 series routers, six Cisco 2950 series switches, two 48 port patch panels, and several boxes of Ethernet cable, with the intent of providing access to network devices which could configured by lab users.

Using 100Mb/s network cards obtained from the UAF surplus warehouse, each of the lab workstations was configured with three network cards, as follows:

1. A Gigabit Ethernet card, which was used to connect the workstation to the lab administrative systems, primarily for user authentication and delivery of the VM images.
2. A 100Mb/s card, which was connected to one of the existing Cisco 2950 switches, which form an administratively configured network. Any virtual machine can be configured to use this network for any purpose, but the network equipment itself is may not be a target of attack, nor may lab users alter its configuration. The intent of this connection is to provide a relatively stable network environment for VM use, although users of this network should assume that it can, and likely will, contain malicious traffic which may be harmful to their virtual machines.
3. A 100Mb/s card which was connected to a port on one of the two patch panels, which are mounted in a separate rack from the administratively controlled network equipment. The newly purchased network equipment is also mounted in this rack, allowing lab users to

quickly create a connection between any of their virtual machines and the user-controlled network components simply by connecting a short cable between the patch panel and the network device. Two basic workstations, obtained from the UAF surplus warehouse, also have been placed adjacent to the network equipment, allowing serial port based configuration of the user-controlled network devices.

The complete network configuration is shown in figure 3. Through the configuration of the workstations and virtualization software, no VM traffic is permitted on the Gigabit network, nor do the workstations directly utilize the 100Mb connections, allowing VM users to undertake network activities with little risk of compromise to the physical workstations, or the to the administrative servers in the lab.

VIII. Usage Policy

Prior to being granted access to the lab, any prospective user is required to read, sign, and return a usage policy which defines their responsibilities in the lab and the types of actions which are allowed, and those which remain prohibited. There are many actions which are permissible in the lab, but would result in academic or criminal punishment if performed on external networks. However, it is vital that lab users understand that there are limits to their activity even within the lab. For example, there are several systems in the lab, such as the servers and the administratively-controlled network devices, which should not be seen as targets for attack under any circumstances. In addition, users of the lab are often exposed to techniques or tools which could adversely affect production networks, so the usage policy also makes clear the responsibility of the user to apply this knowledge wisely and carefully beyond the bounds of the lab. To reinforce the policy and assess student comprehension of the policy, many IA instructors include questions regarding the policy on their course exams.

IX. Lab Usage

The ASSERT Lab has proved to be an extremely valuable tool for UAF, both in the development of its IA curriculum, and as part of the outreach effort to the local community.

Within the Computer Science Department, several courses uses the lab to a significant extent, including the four undergraduate and three graduate IA-specific courses, in which the lab has made a practical exploration of the theoretical concepts possible. However, students are exposed to the lab environment at several other points throughout the CS program through the integration of modules into existing classes. The following are examples of such modules:

1. An exploration of the difference between plaintext and encrypted network traffic, using examples such as Telnet versus SSH, and HTTP versus HTTPS traffic.
2. An introduction to data recovery, using JPEG images on a disk as an example. Students in CS1 are required to write a program to search a disk image for deleted JPEG images, and extract them for review. This task is one of the most frequently requests for assistance that ASSERT staff receive from the local community, often as a result of digital camera images being deleted from flash memory.
3. An example of the use of steganography, in which students experiment with steganographic techniques and tools, and learn the circumstances in which steganography may be useful.
4. A discussion of strong password selection, including a demonstration of password cracking, a discussion of why password selection is important, and how to select, and remember, strong passwords.

There are several other modules, including those aimed towards the upper level networking and operating systems classes, which utilize the ASSERT Lab. While each of the modules provide valuable and relevant information, both for the class in which they are taught and for computer users in general, they also serve as a good introduction to computer security concepts, which often spawns an interest in the more specific IA classes at the undergraduate and graduate levels.

X. Use of Existing Labs

During the course of the ASSERT Lab development, several other configurations for the lab were considered. The most obvious involved the temporary conversion of the existing CS lab through the use of bootable Linux CDs, often called Live CDs [Kno06, Ubu06]. However, this raised several concerns, including that once booted from a CD, the user typically has full access to the disk on the host machine, and as such can manipulate it at will, with the likely result that they could bypass any administrative control on the workstation once it returned to normal use. In addition, there is the issue of providing an isolated networked environment, which is possible to provide in a general purpose lab only if the uplink from the switch is accessible. As a result of these concerns, it was felt that the use of one of the existing CS lab was impractical for UAF.

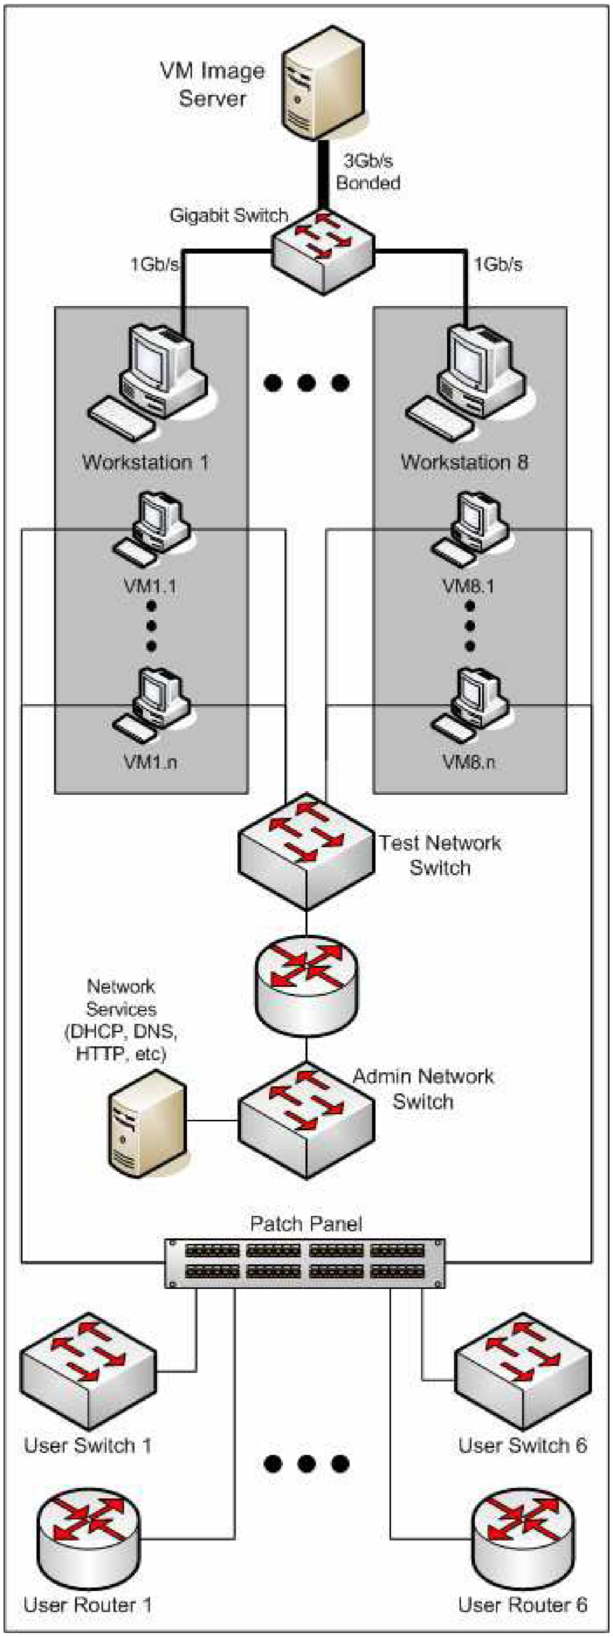

Figure 3: Current Network Configuration.

XI. CONCLUSIONS

Traditional computer labs are typically unsuitable for computer security, information assurance, and networking research and classwork, but it is important that students and researcher have access to environments in which they can gain practical experience. As shown in this paper, there are a variety of ways to provide such opportunities, even when there is little or no budget available for software and equipment. If additional funding is available, more complex environments can be constructed, particularly when virtualization technologies are applied to leverage the full capabilities of modern workstations and servers.

XII. FUTURE CONSIDERATIONS

While the isolated lab does solve several issues, it presents the problem that each lab user must be physically present in the building in order to use the lab facilities.

While this can be a minor to medium inconvenience for those users physically located in Fairbanks, it can present some real challenges for more remote students. One of the major goals at UAF is the provision of distance education opportunities to students at the 12 rural campuses around the state. For students involved in such distance education programs, travel to Fairbanks to use the ASSERT Lab presents at minimum a serious problem, and at times a virtual impossibility, but these students are also from communities that may rely most heavily on a reliable and safe computing infrastructure for access to commercial, governmental, and educational services, including the distance education classes themselves. As such, UAF is currently exploring options for allowing remote access to an environment similar to the ASSERT Lab. The provision of such an environment would not only allow students in remote areas to fully participate in the IA program, but could also allow local users to access the lab resources from around campus or town, reducing the need to expand the physical size of the lab to meet the growing class sizes.

One approach to this remote access that is currently being tested involves the use of one of the VMware server products [Vmw06a]. In this configuration, the virtual machines are executed on the server, but the user is provided access to the VM through several methods, including a interface with a similar appearance to the VMware Workstation product. One advantage of this approach is that the client, which is the remote user in this case, does not execute the virtual machine locally, therefore reducing their need for high performance network connections or computing hardware. There are also several other virtualization products available which may be more suited to providing a remotely accessible lab. Of these, Xen [Xen06, Cam06] is also being tested by ASSERT as a possible solution. Like the VMware products, Xen also provides the ability to run multiple virtual machines on a single server, although it provides that functionality using a different approach, which typically requires patching the VM operating system.

XII. References

[Cam06] University of Cambridge Xen Virtual Machine Monitor. Retrieved February 15, 2006 from http://www.cl.cam.ac.uk/Research/SRG/netos/xen/

[HDR03] L. J. Hoffman, R. Dodge, T Rosenberg and D. J. Ragsdale "Information Assurance Laboratory Innovations," 7th Colloquium for Information Systems Security Education Washington, DC, June 2-6, 2003

[Kno06] Knoppix Homepage. Retrieved December 13, 2006 from http://www.knoppix.org/.

[Nto06] ntop.org Home page. Retrieved February 15, 2006 from http://www.ntop.org

[Ubu06] Ubuntu Homepage. Retrieved December 13, 2005 from http://www.ubuntu.com

[Vmw06a] VMware Home Page. Retrieved February 15, 2006 from http://www.vmware.com/products/ws/.

[Vmw06b] VMware Work Stations. Retrieved February 15, 2006 from http://www.vmware.com/products/ws/.

[Xen06] Xen Source Home page. Retrieved February 15, 2006 from http://www.xensource.com/

Acknowledgements

This research effort was sponsored through generous grants from the National Science Foundation and the University of Alaska Technology Advisory Board.