# Replicating and Sharing Computer Security Laboratory Environments

*Kara Nance & Brian Hay*
*University of Alaska Fairbanks*
*{ffkln, brian.hay}@uaf.edu*

*Ronald Dodge*
*U.S. Military Academy*
*Ronald.Dodge@usma.edu*

*James Wrubel*
*Carnegie Mellon University*
*jcw@cert.org*

*Steve Burd & Alex Seazzu*
*University of New Mexico*
*{burd, alex}@mgt.unm.edu*

## Abstract

*Many institutions are currently investigating the feasibility of creating Computer Security Laboratory environments for their researchers and students. This paper compares four of the current isolated and remote access labs that institutions could use as models to minimize the effort required to create or access a working computer security lab without investing the years of effort that the original creators did. Laboratory attributes investigated include scalability, access capabilities, teaching environments, time requirements, and cost requirements. Additionally a discussion of the challenges associated with each environment is presented. Finally, a model for sharing remote access laboratory capabilities is delineated as an alternative for programs for which the creation of a local remote access lab would not be cost effective and some future investigation areas are identified.*

## 1. Introduction

The 12th Colloquium for Information Systems Science and Education (CISSE) included a well-attended panel on virtualized computer security laboratory environments. A discussion followed that raised some interesting questions about the creation of lab environments and the time and cost that many institutions were investing in laboratories to bridge the learning curve associated with building a computer security lab. The panel agreed that it has taken years of effort by each institution to arrive at each current laboratory configuration. A follow-on question yielded a more interesting result. When asked about recreating the same environment, the time and cost estimate decreased—from years of effort to weeks of effort, and from six figure costs to four to five figure costs. This paper describes four computer security laboratory environments and is intended to provide examples of laboratory environments that can be replicated or shared remotely to minimize the development efforts and costs generally associated with building new computer security laboratories.

## 2. Laboratory Descriptions

The four laboratory environments (the United States Military Academy, the University of New Mexico, Carnegie Mellon University, and the University of Alaska, Fairbanks) described were selected as they represent a range of solutions that meet the needs of their individual constituencies and can be recreated, recombined, and deployed in new environments to provide other institutions with similar laboratory environments. In some cases, access to these labs can be negotiated with other institutions to minimize the costs associated with creating and maintaining a security laboratory environment for their students. For each environment, the target demographic and driving issues are outlined, as well as a brief description of the lab evolution leading to the current solution being employed. This is followed by a summary discussion of the scalability, access capabilities, teaching environments, time requirements, cost requirements, as well as a discussion of the challenges associated with each environment and alternatives to building a local security lab.

### 2.1. United States Military Academy

The United States Military Academy (USMA) began using virtualization to support information assurance (IA) education and training in 2001 [1, 2]. The requirement to dual boot lab computers between Windows and Linux drove the adoption of VMware workstation. As comfort and experience with virtualization grew, the capabilities of VMware workstation increased, and the lab systems gained processing power, the courses at USMA were completely reworked to leverage the new capabilities introduced with virtualization. Given the nature of the educational environment at USMA, the model does not include distance education. All students are expected to access the lab facilities on site to complete academic

978-0-7695-3450-3/09 $25.00 © 2009 IEEE

requirements. Therefore, the development of virtualization-supported labs focused on local "workstation" solutions.

The primary consideration in the creation of the IA lab (called the Information Warfare Analysis and Research Lab - IWAR) was the isolation of the systems from the campus network. The safest (and adopted solution) was to simply not connect these systems to the campus network. However to ensure the lab could be used for other purposes, a second computer was placed at each workstation that was on the campus network. The two systems are accessed using a KVM switch.

The architecture worked very well for two years as the curriculum and lab infrastructure continued to develop. However, the drawback of the lab that moved the staff to reconfiguration in 2004 was that it could only completely support one class. There were observed time conflicts amongst students, given that the virtual machines being used were resident on a specific machine. To address the problem, USMA integrated an Active Directory (AD) architecture and a storage area network (SAN) into the isolated IA network as shown in Figure 1. The SAN is behind a set of clustered file servers with dual fiber channels on each server.

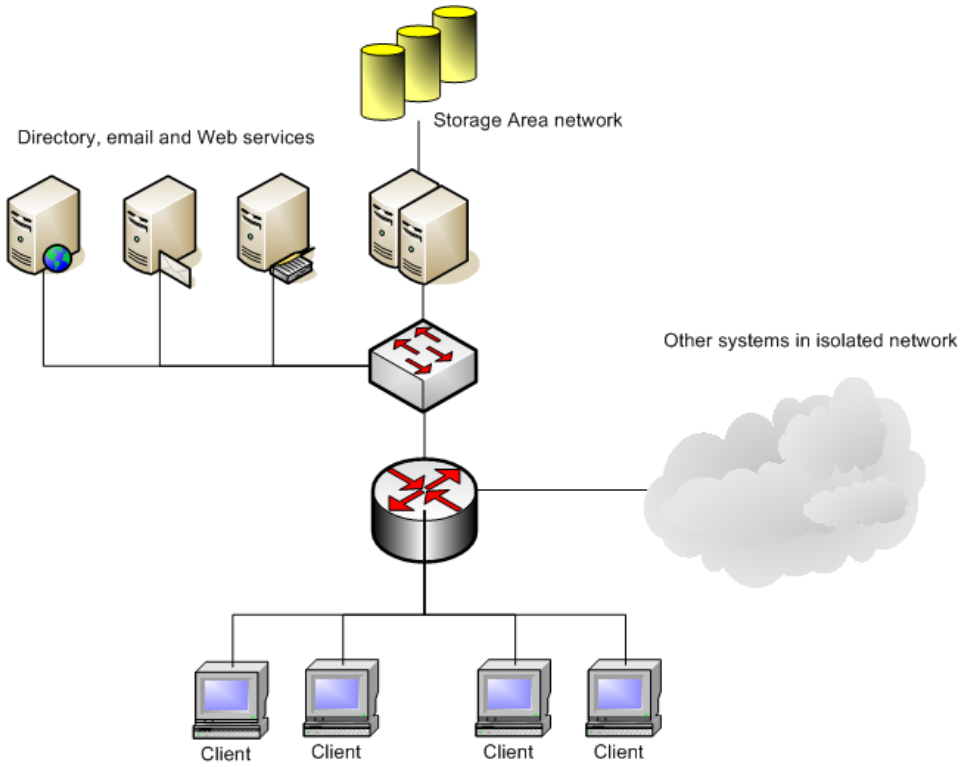


**Figure 1: USMA Virtualization Network Topology**

The AD environment utilizes roaming profiles to enable a student to sit at any machine in the lab and have full access to the required virtual machines. In this architecture, the virtual machine files (.vmx, .vmdk, etc.) all reside on the student's file share on the SAN. This allows all the processing to be done locally on the student's workstation while the virtual machine files remain on the SAN. Additionally, email and web services for dissemination of course material are included. The lab environment also includes several other networks that can be reached by the student for specific course objectives. These networks are similarly not connected to the campus network or the Internet.

## 2.2. University of New Mexico

In late 2004 the Anderson School at the University of New Mexico (UNM) undertook a strategic initiative that expanded the scope of its computing capabilities beyond the then existing physical lab (PLAB) to all school classrooms and common areas. Key initiative components included extensive use of computer-based pedagogy, wireless access throughout school buildings, student laptop computers, and a virtual computer lab (VLAB) with remote access from any Internet-connected computer. VLAB requirements included providing the same application suite as the PLAB

computers, consistent look-and-feel across all lab and classroom computers, and flexibility to support a variety of learning environments.

To extend the computing resources of the PLAB the following project parameters were articulated for the VLAB:

1. Resources must be available outside regular university business hours (evenings, nights and weekends).
2. Resources must be accessible from inside and outside the Anderson School classrooms through Internet connectivity.
3. The same portfolio of applications and services found in the PLAB must be made available.
4. The student and faculty must have a consistent experience when using the resources in terms of connectivity, software resources and configuration.

The first step to create the VLAB was to allocate part of the space occupied by the former PLAB to host rack-mounted workstations. Forty two existing PLAB workstations running Windows XP were repurposed as VLAB workstations. The workstations were mounted on racks and connected to KVM switches, network switches, and uninterruptible power supplies. A SAN was installed to store VMware operating system images in support of different information systems environments including Information Security, Database Administration, and System and Network Administration.

The VLAB workstations are supported and managed by servers implementing Microsoft Active Directory, Microsoft Remote Installation Services, Symantec Antivirus, and a Nokia Checkpoint firewall. The servers control various aspects of workstation configuration and operation including:

- Operating system configuration including available utilities, patch installation, desktop settings, and user ability (or lack thereof) to perform functions such as accessing printers, installing device drivers, and executing command line functions
- Application software installation and configuration
- Security settings including file system permissions, antiviral scans, and protection from malware

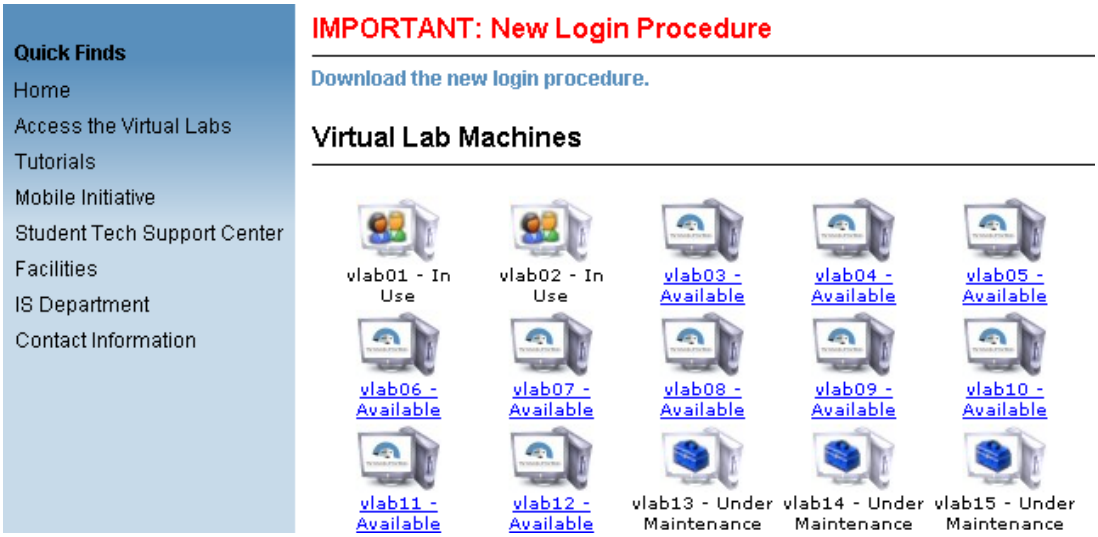


**Figure 2: Partial view of UNM web page for accessing VLAB computers.**

To simplify access to the VLAB workstations, an in-house developer created a Web-based interface on a small dedicated web server [3]. The interface includes embedded programs written in C and VBScript to generate the web page from text configuration files and to update its content. One VBScript polls the workstations regularly to determine which are available and which are in use and updates the web page icons and text as appropriate as shown in Figure 2. When a user clicks on an available system, a script initiates a Microsoft Remote Desktop (RDP) connection from the user's computer directly to the chosen workstation.

The VLAB supports general-purpose computing and provides software to specifically support courses in areas such as marketing research, business strategy, and IA. VMware workstation is installed on all VLAB machines and it is heavily used within IA courses for classroom exercises and student projects.

## 2.3. Carnegie Mellon University

In 2004, in response to the needs of the student population and to lessen the administrative burden on instructors, the Software Engineering Institute (SEI) at Carnegie Mellon University transitioned its hands-on training component away from physical servers in classrooms to a remote server solution based on virtualization technology, allowing students to access the hands-on training materials asynchronously. The SEI had three goals for the new environment:

1. The online on-demand experience had to be exactly equivalent to the classroom experience.
2. With our small teaching staff, students must be able to access the material without an instructor present, but the environment should support both instructor-led and self-paced uses.
3. The environment had to 'just work'. No resource scheduling, no client configuration changes or software installs, and it had to be available anywhere the Internet reaches. It had to work in 'locked down' environments and only on well-known ports.

To fulfill these requirements, the SEI created a hands-on training lab capability to enhance its Virtual Training Environment [4]. The hands-on lab system contains the following components:

- A rack of servers with maxed CPU and memory and a SAN with a library of disk images.
- A database of lab configurations, which are combinations of images and network interconnects.
- A .NET interface to VMware Virtual Center to deploy lab configurations on demand.
- Manuals for each lab configuration that walks the student through a task or tasks in the lab.
- A Web application and Java client to allow students access to their deployed lab environment (and only that environment).

Students access lab environments by logging on to a Web portal that lists all the environments that the student can access. The student selects an environment and clicks the 'Launch' button – The VTE application selects a machine on which to house the student's environment, provisions the virtual machines, then loads the Java applet allowing the student to access their lab through their browser. A clock running on the server and visible to the student counts down the reservation time – when the clock runs out the environment is recycled for use by other students. This environment fulfills all of the goals for the online on-demand lab.

## 2.4. University of Alaska Fairbanks

There were four major driving forces behind the development of the Advanced Systems Security Education Research, and Training (ASSERT) remote access virtualization lab including:

1. More effective use of physical space in the existing computer science building. (The space allocated to the existing security lab was sufficient to hold small numbers of students comfortably, but as interest in the IA field grew, and interest in the use of virtualization grew into mainstream computer science classes, the physical limitations of the lab were problematic.)
2. The ability to safely use lab resources from faculty offices, home, and, most importantly, larger general-purpose labs around campus was a primary concern.
3. The ability to provide students in distance education courses with high-quality lab experiences. UAF is charged with providing distance education throughout much of the State of Alaska, and currently serves many students in geographically isolated locations.
4. The development of a proof of concept virtual lab environment that could be easily scaled to serve as a national resource, allowing other institutions to focus on the use, rather than the development, of virtual lab environments.

The current instantiation of the ASSERT virtual lab is the culmination of five years of effort, beginning with a small scale, completely isolated virtual lab (with no remote access) in 2003. While the lab has evolved through many iterations, each aimed at addressing the limitations of the previous version, it has been incredibly valuable to students and researchers throughout its lifetime [5].

The current lab hardware is very similar to that found at the SEI, with several multi-core rack servers accessing virtual machine images on a fibre-channel connected SAN, with another server providing the management, authentication, and authorization services. The software used is different, in that UAF uses VMware's ESX product, which runs directly on the server with no intervening operating system, as opposed to the VMware Server product used at SEI, which runs as an application in a general purpose operating system.

The user interface for the UAF lab is VMware Virtual Center itself. This UI requires the user to download and install a client program (VMware Infrastructure Client), which is free, but does require the user to connect from a Windows-based system. Administrative users (typically instructors) assign virtual machines, which are generally derived from templates, to users. Virtual Center is integrated with a Windows Active Directory, against which authentication is performed, and virtual machines can be assigned to single or multiple users, either explicitly or based on AD group membership. While virtual networks are sufficient for many tasks, several physical network components are also integrated into the lab environment, allowing interaction with the real components such as a Cisco IOS.

## 3. Observations

While each lab meets the needs of a particular demographic, there are some overarching issues that provide interesting analysis of the environments individually and collectively. These include the scalability of the solution, access capabilities, teaching environments, time investment, and costs associated with deploying and maintaining the environments. In addition to these attributes, there are challenges associated with the various environments that should be the focus of future research and development.

### 3.1. Scalability

Scalability refers to how the lab environments grow to meet increasing demand. In order to increase capacity in the USMA and UNM labs, an additional physical workstation needs to be added for every additional concurrent user. In the case of USMA this actually involves the addition of two workstations in a traditional computer lab environment. At UNM it requires installation of a workstation in a rack, which requires significantly less space.

Both the UAF and SEI labs can be scaled much more easily, and with far less reliance on physical space, to meet the demands of additional users by adding capacity at five locations:

1. Memory – each current server has 16-20GB of RAM, although they have the option for 32GB per server. While many virtual machines in these environments can be configured to use relatively little RAM (e.g., a Linux router with 32-64MB of RAM, or a Windows XP systems with 128MB), some systems (such as a VM running as a heavily loaded server, or Windows Vista in general) greatly benefit from more RAM. RAM has been the most common limiting factor in adding virtual machines to a physical server in the ASSERT Lab. In early testing of the UAF environment, a single server with 8GB of RAM was capable of supporting 15-20 concurrent users, each using two virtual machines (Windows XP and Centos 5, configured with 128-256MB or RAM per VM). Subsequent increases in the RAM allocation in each server have resulted in a linear increase in the number of virtual machines and concurrent session that each server can support.
2. CPU – the current servers in the ASSERT Lab are dual processor, quad core (a total of 8-cores per physical server), and this does not seem to be a limiting factor for the usage seen at UAF. While high CPU load is commonly seen if many VMs are booted simultaneously (such as at the start of a class or lab session), CPU utilization on the physical hosts is generally well within capacity during normal lab use. In the early UAF testing scenario described in the RAM section, CPU utilization on the servers was rarely above 25% of the total capacity.
3. Physical Servers – additional physical servers can be easily added to the environment with little more than the installation of VMware ESX, physical installation in the rack, and the very quick process of integration into the Virtual Center system. At some point, power and cooling may become an issue.
4. Bandwidth – as all of these labs require some level of network access, sufficient bandwidth must be available. Testing at UAF shows that the bandwidth requirements for the Virtual Center server are very low (less than 250Kb/s for 15 on-campus connected users accessing 30 virtual machines, each with 1280x1024 displays), as shown in Figure 3. The ESX server running the 30 virtual machines required a peak data transfer of 20Mb/s, as shown in Figure 4, while running an interactive lab exercise (i.e., interactive access to the Windows and Linux graphical user interfaces). The server to client (i.e., receive) direction

accounts for the vast majority of the traffic, which is well suited to the needs of off-campus users accessing the lab on asymmetric broadband connections.

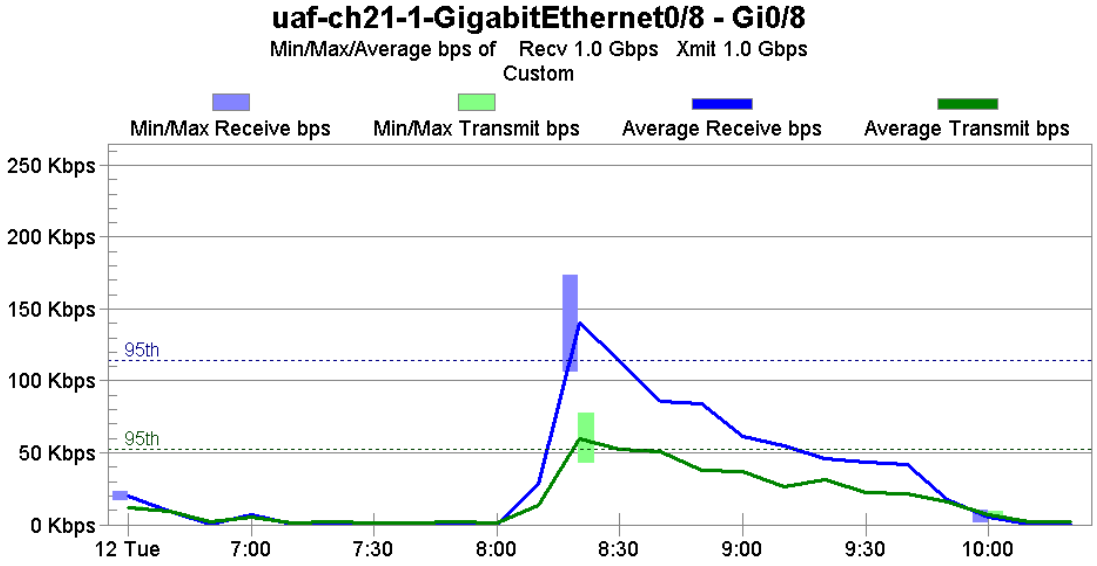


**Figure 3: Network bandwidth utilization for VMware Virtual Center in a lab (8.15pm-10.00pm) with a peak of 15 students accessing 30 virtual machines.**

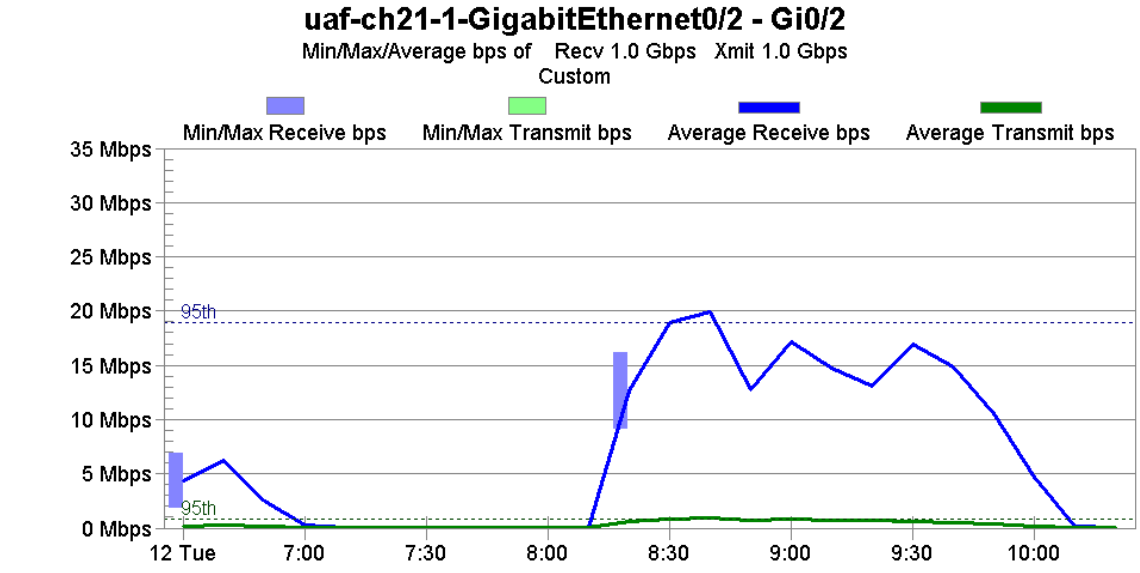


| | | | | | | | | | | | | | | | | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Min/Max Receive bps | 6.95 | 6.19 | 2.49 | 0.23 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.10 | 16.22 | 18.94 | 19.98 | 12.78 | 17.19 | 14.78 | 13.08 | 16.93 | 14.82 | 10.58 | 4.67 | 0.18 | 0.00 |
| Min/Max Transmit bps | 0.31 | 0.30 | 0.13 | 0.01 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.01 | 0.70 | 0.88 | 0.92 | 0.69 | 0.80 | 0.76 | 0.69 | 0.60 | 0.53 | 0.37 | 0.17 | 0.01 | 0.00 |
| Average Receive bps | 4.40 | 6.19 | 2.49 | 0.23 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.10 | 12.67 | 18.94 | 19.98 | 12.78 | 17.19 | 14.78 | 13.08 | 16.93 | 14.82 | 10.58 | 4.67 | 0.18 | 0.00 |
| Average Transmit bps | 0.21 | 0.30 | 0.13 | 0.01 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.00 | 0.01 | 0.59 | 0.88 | 0.92 | 0.69 | 0.80 | 0.76 | 0.69 | 0.60 | 0.53 | 0.37 | 0.17 | 0.01 | 0.00 |

95th Percentile : Average Receive bps is 18 Mbps, Average Transmit bps is 855 Kbps
SolarWinds.Net Orion NPM Web Engine Version 8.5.1

**Figure 4: Network bandwidth utilization for VMware ESX server in a lab (8.15pm-10.00pm) with a peak of 15 students accessing 30 virtual machines.**

The SEI configuration is far less bandwidth intensive by utilizing the Remote Desktop Protocol over SLL to allow client interaction with the virtual machines, with an F5 Load Balancer/SSL Accelerator acting as the external gateway for clients. When a lab environment is

assigned to a student, the configuration includes a special VM called LaunchPad, which is dual-homed between the physical network (i.e., the connection from the F5 device) and the virtual network in which the students perform their lab activities. This results in bandwidth utilization of ~1Mb/s for 15 concurrent users because the RDP protocol has been carefully optimized for remote access. A minor tradeoff in this case is that there is an indirect network pathway between the virtual machines and the client machine, although it can be managed if the LaunchPad VM and F5 are carefully considered to restrict any outbound traffic other than the RDP connections.

In addition to network bandwidth, network latency is an important consideration for remote access to the lab environments. While it is easy to connect on-campus labs to the remote access labs on low-latency links, of-campus users must also be able to access the lab without an unreasonable delay, and as the remote lab scales it should be able to handle connections from locations around the country, or even around the world. Testing at UAF from a variety of locations around the United States has shown that connections to the lab have acceptable GUI performance with ping times as high as 200ms. Latency below this threshold has been possible in connections to Alaska from cities around the United States, including recent tests in Colorado, Florida, Texas, and California. Preliminary, and somewhat informal, testing has shown that Windows VMs tend to perform better than Linux VMs on higher latency networks, and that Windows Vista is the worst performer in the Windows family. However, better response can be achieved from all guests if certain graphical features are disabled, or if the display resolution is reduced.

5. Storage – in each case the virtual machines (configuration files and virtual disks) are stored on a SAN, which is accessed by the workstations or servers on demand. It is vital that this storage is sufficiently large to store all of the virtual machines, and that access to these files is available at high throughput and low latency. While storage is relatively cheap (e.g., new hard disks for a SAN), the SAN itself, and the network (e.g., fibre channel) that connects it to the workstations or severs can be expensive components of the lab.

## 3.2. Concurrent Access

Concurrent access to the USMA and UNM labs is limited almost solely by the number of physical workstations in the lab environment, although in both cases students can utilize any free workstation (i.e., user virtual machines are not tied to a specific physical workstation). Adding additional workstations to either environment increases the number of possible concurrent users, and the actions of one user have little impact on other lab users (provided that sufficient network bandwidth both to the virtual machine images and to the clients is available)

Limitations on concurrent access in the UAF and SEI labs are based on a more complex formula, as each server can accept multiple concurrent connections. In the USMA and UNM environments each user receives exclusive use of specific CPU and memory resources supported by shared use of network and some storage resources. Thus, individual user actions have limited impact on the resources available to others. This is not true in the UAF and SEI scenarios by default. For example, a single user running 20 virtual machines, or even 2 virtual machines that are performing processor intensive tasks, is consuming resources that are no longer available to other users. As such, the limitations on concurrent access are based on the portion of the resources consumed by each user, which is no longer a fixed value. It is possible to configure the ESX lab environment to allocate resources pools to users or virtual machines, in which the allocation of processor time or RAM is limited, thereby trading flexibility for a more defined limit on the number of concurrent users Virtual machine templates can be created which are optimized for the lab by reducing RAM allocations, turning off unnecessary services or programs, and disabling the "flashier" features of graphical user environments. There is also benefit to knowing potential demand and planning hardware capacity in anticipation of associated resource requirements. In addition, adding servers can also increase the number of concurrent users.

## 3.3. Flexibility and Realism

One of the great advantages of all of the virtual labs over their physical counterparts is that they provide a standard computing environment for students, faculty, and researchers, than can be duplicated easily, and in which it is easy to return (or even advance) to known states.

There are some differences in the extent to which other systems are integrated into the lab environments. In both the SEI and UNM labs there are few additional components beyond the virtual machines, which severely limits the range of accessible devices (essentially only x86 devices are supported in these labs). At UAF, a small number of additional

components are currently integrated into the lab, including physical Cisco switches, routers, and firewalls. The USMA IWAR configuration is by far the most flexible in this regard, in that it provides an environment where the virtual machines that the student uses for lab work are not restricted to the "virtual world" only. The virtual machines can connect to physical machines that are part of the IWAR network (although still physically separate from the University Network and the Internet). This capability allow for the inclusion of any number of different machine architectures, operating systems, and network topologies. This resource however would still be available if the IWAR were to include a remote access capability.

None of the labs allow the VMs to access production networks (i.e., the VMs are confined to isolated networks) in order to ensure that lab activities cannot accidentally or deliberately impact production systems. However, this makes the transfer of data into or out of the lab environments for legitimate purposes much more challenging than it would be in a physical lab. Such uses include students saving work from the lab for submission as homework, or the transfer of new software packages into the lab. The SEI lab is primarily based on pre-configured lab exercises, and these can be submitted to the instructor for grading entirely within the virtual environment, reducing the need to remove content from the lab. In the UNM lab, instructors have access to student VMs stored on the SAN to support consultation and grading. The UAF environment includes an FTP server to which students can submit assignments or contents, and which can either be graded by the instructor from another VM, or, if necessary, can be burned to CD or DVD on a physical workstation connected to the remote lab. As for moving data into the environment, UAF provides mirrors of various Internet resources (such as operating system patch and software download sites) as a VM within the lab itself. In addition, the Virtual Center software enables students to make read-only connection from their VMs to CD or DVD images on their client workstations, although this is not a particularly efficient method of transferring large volumes of data into the lab.

The lab environment at USMA includes two primary mechanisms for students to move data from the virtual networks to the physical one. First, each physical host and virtual machine has access to network printers on the respective networks. Second, the students can mount USB flash drives to shuttle data. The distribution of software (patches, tools, etc) is accomplished done through enterprise services in the IWAR. Web, FTP, and file servers available on both the host and virtual networks are used.

The UNM design enables lab workstations to establish network connections to university online services in order to download lab exercises, including approved applications embedded in the lab package. The package can then be dragged into the required VMs to complete the lab exercise. The evaluation of these labs by the instructors does not require the analysis of new or malicious code. The student's performance is measured by their answers to lab questions. During this process the VMs remain isolated within their "host-only" network configuration

### 3.4. Teaching Environment

The SEI environment is deployed as a set of pre-configured exercises, which the student can "check out" and work through in some predetermined time limit. Instructors create the exercises, and can then make them available to all students, or only specific groups. An instruction document for each lab exercises is included as an integral component of the environment.

The UNM and UAF lab takes the approach of assigning virtual machines to students (or groups of students), and these VMs are available in the students' accounts. It is common for VMs to be assigned for an entire semester, and even longer in the case of research projects or special group projects, such as the Collegiate Cyber Defense Competition team. At UAF, several scripts have been written to augment the native Virtual Center capabilities in this area, allowing an instructor to, for example, create new VMs from a template for each student in a specified class list. This approach is far less structured than that chosen by the SEI, and while it offers increased flexibility to students in terms of the use of the VMs, it also generally requires more experience on the part of the lab user and the instructor.

USMA follows a similar model as used at UAF. An identical network of virtual machines is provided to each user. Throughout the semester, different virtual machines from the network are used to demonstrate given learning objectives. Students are able to add additional virtual machines with instructor coordination.

### 3.5. Time Investment

All of the labs described have developed into their current forms over several years, but the time required to replicate them at another institution, is likely to be on the order of a few days to weeks. However, lab administration is still time consuming, even if we accept the unlikely assumption that the lab capabilities

are expected to remain static. Resources within the labs, like patch and software repositories, require frequent updates, as does the virtualization software itself. For example, the USMA IWAR facility is administered by one full-time employee.

Virtualization is a dynamic field, with new tools being introduced, and existing tools being upgraded with new or improved capabilities. While the labs featured in this paper all currently use some VMware product for virtualization, advances in the Xen hypervisor and associated management tools, the introduction of Microsoft's Hyper-V hypervisor, and the evolution of virtualization specific functionality in Intel and AMD processors are all important trends in the x86 virtualization market which lab designers and administrators must carefully monitor to determine how best to provide the next generation of virtual labs.

Although virtualization can simplify the task of deploying lab exercises to students, it does little to reduce the burden of creating lab assignments, which is typically placed on the instructor, and which of course varies by the course and the complexity of the assignment. For example, a USMA estimate for the creation of a moderately involved exercise that uses the virtual machine network shown in Figure 5 to increase student understanding of the concept of a botnet and the security measures associated with managing this thread is 30 hours of instructor time. This includes building the virtual machine as well as documenting the steps in the exercise. However, once an exercise has been created it can be replicated very quickly to large groups of students within an institution. From a technical perspective, cross-institutional sharing of exercises is almost as simple, even if the actual data transfer of the typically large virtual machine disk files takes place using DVDs sent via the postal service. However, the largest impediment to this sharing is the licensing of proprietary operating systems and software packages, even between institutions that both have site licenses for the software in question. Compatibility of VM files across software packages and versions can also be an impediment.

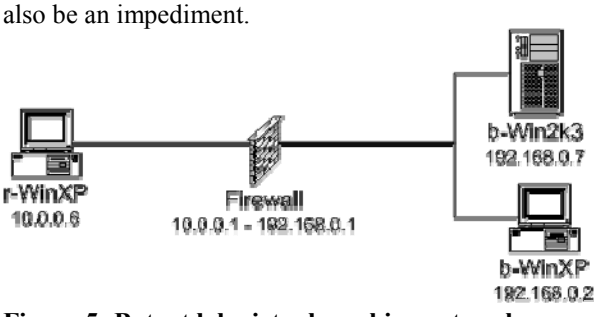


**Figure 5: Botnet lab virtual machine network**

### 3.7. Challenges and Limitations

While each laboratory configuration is currently meeting the needs of the associated users, there are challenges associated with each. The IWAR lab at USMA has continued to grow to meet the IA needs of the students studying computer science and IT related topics; however, the requirement to have some degree of IA education and training for every student is looming. As this expansion of the requirements becomes a reality, the current architecture will require modification to include potentially less capable, but more accessible functionality.

The limitations associated with both SEI and UAF configurations are primarily with respect to diversity. It is difficult to include anything other than x86 systems, such as hardware devices, wireless networks, and VoIP. These types of applications are critical for security education but difficult if not impossible to virtualize. In addition, both institutions are facing software licensing challenges. Many commercial software packages do not offer a licensing model for virtualization, which limits the scenarios that can be distributed and shared.

SEI and UNM have also experienced some challenges associated with students becoming disoriented with multiple *layers* of desktop interfaces. "Which Start menu do I click?" is a common question. This requires training for students and faculty to understand layers, particularly when using virtual environments on a remote host.

All of the virtual environments are also faced with the potential for single points of failure such as the primary firewall/router, website, or lab repository. In the case of UAF, the limited infrastructure to support connectivity itself is a single point of failure. These challenges and limitations are being evaluated and addressed as the laboratory environments continue to evolve to meet the needs of their growing constituencies.

## 4. Future Considerations

The institutions featured here have been at the forefront of virtual lab evolution, but are by no means the only institutions pursuing this effort. Virtual lab environments are the focus of many conference discussions, and many schools either have a virtual lab at some level of sophistication from basic to the more advanced examples presented here. Many more institutions are exploring virtualization technology with a view towards deploying it in the future. The current approach, in which every institution climbs the

virtualization learning curve, fails to really leverage the power of this technology, and to exploit the economies of scale that it offers. The need for virtual security labs is analogous to the need for supercomputing resources in the United States. Some institutions need frequent access to supercomputers, while others have infrequent high-performance computing classes, or cover supercomputing as a module in a more general class, and as such require far more sporadic access to supercomputing resources. The approach to providing these resources in the U.S. is to fund a small number of supercomputing centers, each of which provides resources to students and researchers across the nation on the basis of need. This reduces the number of people who are required to have the knowledge necessary to design and administer supercomputing centers, while ensuring that users can focus on the use, rather than the management, of the resource.

The current state in virtual labs has every institution painfully acquiring the knowledge necessary to design, implement, and administer virtual labs, which reduces the time they have to actual use and develop content for their labs. Research into the feasibility of developing a new model or following a model similar to the supercomputing model is needed. For example, a small number of high-capability virtual labs could be deployed in the nation, each of which would then serve as resources for institutions around the country. Two of the lab environments presented here are highly scalable, and while the effort required to administer virtual labs increases linearly with the number of institutions that host them (as is the current case), the effort required to administer each national or regional lab does not increase significantly as additional servers are added to increase capacity.

By taking such an approach, virtualization lab environments could be made available nationwide to all institutions, with no requirement on the part of client institutions to focus on the mechanics of virtualization, but rather on the use of virtual labs in the curriculum and research projects. In addition, schools which cannot currently justify labs dedicated to topics outside their focus areas, such as computer security, could use these national resources to easily integrate high-quality hand-on lab exercises as occasional modules during their existing classes.

By drawing on the proof of concept experience of the institutions leading the virtual lab evolution, as presented here and elsewhere, and using software and hardware that is currently available, opportunities for resource sharing now support distributed models of virtualization laboratories. Future research into models of sharing these resources is needed to continue to more efficiently and effectively meet the needs of our user groups.

## 5. References

[1] J. Schafer, D. J. Ragsdale, J. R. Surdu, and C. A. Carver, Jr., " The IWAR Range: A Laboratory for Undergraduate Information Assurance Education," Proceedings of the 6th Annual CCSCNC, Middlebury, VT, April 20-21, 2001.

[2] 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.

[3] VLAB Website. Retrieved August 25, 2008 from http://averia.mgt.unm.edu/ASM_VLab_Design_EXTERNAL.pdf.

[4] CERT Virtual Training Center. Retrieved June 5, 2008 from https://www.vte.cert.org/vteweb/

[5] Hay, B. and K. Nance. Evolution of the ASSERT Computer Security Lab. Proceedings of the 10th Colloquium for Information Systems Security Education. Adelphi, Maryland. June 2006.