

Designing a Flexible, Multipurpose Remote Lab for the IT Curriculum

Steven Rigby
BYU-Idaho
Smith 418
Rexburg, ID 83440
rigbys@byui.edu

Melissa Dark
Purdue University
Knoy Hall
West Lafayette, IN 47906
dark@purdue.edu

ABSTRACT

A key inhibitor to effective distance education in Information Technology is providing a “hands on” laboratory experience that allows students to acquire the application and problem solving skills expected of IT graduates. While there are instances of universities developing and deploying remote labs where students are able to perform labs through the Internet using “virtual machines” and other technologies, many have found the complexity and time required to maintain labs problematic and therefore prohibitive. This paper analyzes current trends in remote lab design and explores a design that intends to increase utilization between courses, lower costs, ease management, and reduce the time needed to implement remote labs.

Categories and Subject Descriptors

K.3.1 [Computers and Education]: Computer Uses in Education
–Distance learning

General Terms

Experimentation, Security, Human Factors, Standardization.

Keywords

IT education, curriculum, remote labs.

1. TRENDS IN DISTANCE EDUCATION

Enrollment in higher education in the United States continues to increase. In Fall 2003 there were 17.3 million students enrolled in higher education (NCES, 2003); this increased 4% for a total enrollment of 18 million in 2004 (NCES, 2004). At the same time that enrollment in higher education is growing modestly, enrollment in distance education courses and programs is also growing. According to the National Center for Education Statistics (1999), in the 1997-1998 academic year 33% of Title IV eligible degree granting institutions were offering courses and/or programs using distance education. By the 2000-2001 academic year (the last year for which data are currently

available), this had increased to 56%, an increase of 23% in 3 years (NCES, 2002). Furthermore, in 1997-1998 there were 1,363,670 students enrolled in credit granting courses at these institutions (NCES, 1999) compared to 2,876,000 students in 2000-2001 (NCES, 2002). A closer look at these trends shows that in 1997-1998 62% of 2 year institutions were offering some courses or programs via distance education; in 2000-2001 this had increased to 90%. Similarly, in 1997-1998 78% of 4 year institutions were participating in distance education; this had increased to 89% in 2000-2001. Clearly distance education is increasing.

At the same time that distance education is growing in terms of institutional participation and student enrollments, there are factors preventing growth in distance education. Some of the more frequently cited preventive factors include: lack of fit with institutional mission, concerns about faculty workload, concerns about course quality, the costs associated with developing distance education courses and programs, the costs of maintaining equipment, and limited technological infrastructure (NCES, 2002).

The preventive factors are especially noteworthy when the courses/programs to be offered are technical in nature. This is due to the increased reliance on equipment in technical courses/programs and the difficulty in designing hands-on laboratory experiences at a distance that are as effective as the on campus laboratory. Despite these obstacles, solutions are being found to offer remote labs in technical education.

2. CURRENT TRENDS IN IT EDUCATION

The advent of virtual machine software has created a whole new world of possibilities for creating remote labs (Brown and Lahoud, 2005; Leitner and Cane, 2005; Stockman, 2003). What was technically difficult and cost prohibitive just a few years ago is now commonplace. For example, a student can create an entire network of computers virtually on one physical computer. This virtual environment allows for labs to better imitate real world situations. The two main virtual machines software packages used are VMware and Microsoft Virtual PC. Both of these packages allow for multiple OS's to interact with each other through virtual network cards thus allowing for virtual networks. While virtual networks reduce reliance on hardware infrastructure, they can be time consuming to manage.

Some have found a less time consuming and more manageable solution is to outsource the infrastructure to a third party (Brown

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. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SIGITE'06, October 19–21, 2006, Minneapolis, Minnesota, USA.
Copyright 2006 ACM 1-59593-521-5/06/0010...\$5.00.

and Lahoud, 2005). This allows institutions to focus more time on creating labs and instruction and less on maintaining the network environment. Vendors can provide almost all of the pieces needed for a successful implementation. However, the down side to having a 3rd party vendor manage the environment is the increased costs. Tens of thousands of dollars can be easily spent for hosting of networks and labs.

Another approach is to build remote lab environments in-house (Brown and Lahoud, 2005). This can be done by the individual institution at a single site or at multiple sites in using a consortium approach (Toderick, Mohammed, and Tabrizi, 2005; Yang, Yue, Liaw, Collins, Venkatraman, Achar, Sadasivam, and Chen, 2004). This approach has the potential of being lower cost, especially when that cost is shared among institutions or sites. The advantages of “in-house” implementations are that many of the required pieces are already built into the existing campus infrastructure. Remote labs can easily piggy-back onto existing Internet networks and authentication could be centrally managed through a campus RADIUS or directory server. Additionally, VPN solutions may already be in place that allows off-campus access.

Padman, Meemon, Frankl and Naumovich (2003) suggested creating a configurable lab environment using swappable hard drives (Padman, etal, 2003). This approach provided a solution for isolating the campus network and allowing for easy configuration of different hard drives for specific labs. Using interchangeable hard drives becomes more challenging in a remote environment where the student will not have physical access to hard drives or network equipment. The goal for remote lab design is to be as unobtrusive as possible. Time spent on infrastructure setup and configuration should be as minimal as possible.

3. THE BUILDING BLOCKS OF REMOTE LABS

In order to visualize the desired goal for remote labs, it is important to understand the critical pieces that “make up” the infrastructure. In this section we discuss the interdependence of the critical pieces as well as services including: scheduling, remote access, an operating environment, connected devices/systems, reporting, and assessment.

A key factor to success is the ability for students to schedule when they will perform the lab. Work schedules and geographic locations necessitate the flexibility of allowing students to decide when they can fit labs into their schedules. This piece of the remote lab puzzle can be difficult. Not only will the scheduling software need to allot time when students are allowed to perform labs, but the software will also need to provide an enforcement mechanism to ensure students are done when their time is up.

The next piece to be considered is how the remote students will make a connection into the lab environment. One way to provide this connection is to implement remote access protocols and virtual private networks (VPN’s). VPN’s create an encrypted tunnel where the traffic between the student and the VPN server are private. This encrypted tunnel can be

accomplished through the use of network devices and/or servers, and will allow students to make secure connections to campus lab environments.

After a student connects to a lab environment, they will need an operating environment to perform the lab. This is where the deviations among the different institutions and even courses take place. For networking students learning about routers and switches, networking equipment is required. For students taking operating systems courses, single or multiple Linux/Windows systems may be needed. Security courses may require the use of multiple operating systems and networking devices in order to perform a lab. If remote labs include networking equipment, a console server is necessary to connect and configure the device. A console server allows users to SSH or Telnet into the console ports of the networking devices. This ability is essential for configuring and recovering routers, switches, firewalls, access points, and other network and security equipment. The diversity between these requirements has led to the creation of distance labs for a specific course or curriculum. The limitation is that economies of scale of resources are not realized. We address this limitation more fully in the next section.

The final piece of any lab experience is reporting and assessment. Was the objective of the lab realized and did the students master the skills? Did the student’s mental dots of “theory” connect with the dots of “hands-on application?” These questions can be better understood through reporting of lab time by the students and of assessment that shows the students mastery of the objectives.

4. A NEW DESIGN

Our goal is to introduce a remote lab design that is simple, scaleable, and flexible enough to use for the entire IT curriculum to use. This design incorporates the latest in virtual machine, virtual private networks (vpn), console server, and virtual patch panel technology.

An exciting new technology that will improve the usability of remote labs is virtual patch panels (also referred to as physical layer switches). Physical layer switches allow for network connections to be made among different types of equipment without having to do any “cable-swapping.” The virtual patch panels allow for many different implementations including preset configurations that create the necessary required connections for each lab. By combining the features of virtual machines, console servers, and virtual patch panels, new ways of developing, maintaining, and offering remote labs are not only possible, but synergistic. This design is referred to as the Integrated Virtual Remote Lab (IVRL).

In this design, students can access the scheduling server through any Internet connection and schedule what time they would like to perform the lab. The scheduling server may include a RADIUS server for authorizing students and has the ability to open and close vpn sessions based on the student’s lab start and stop time.

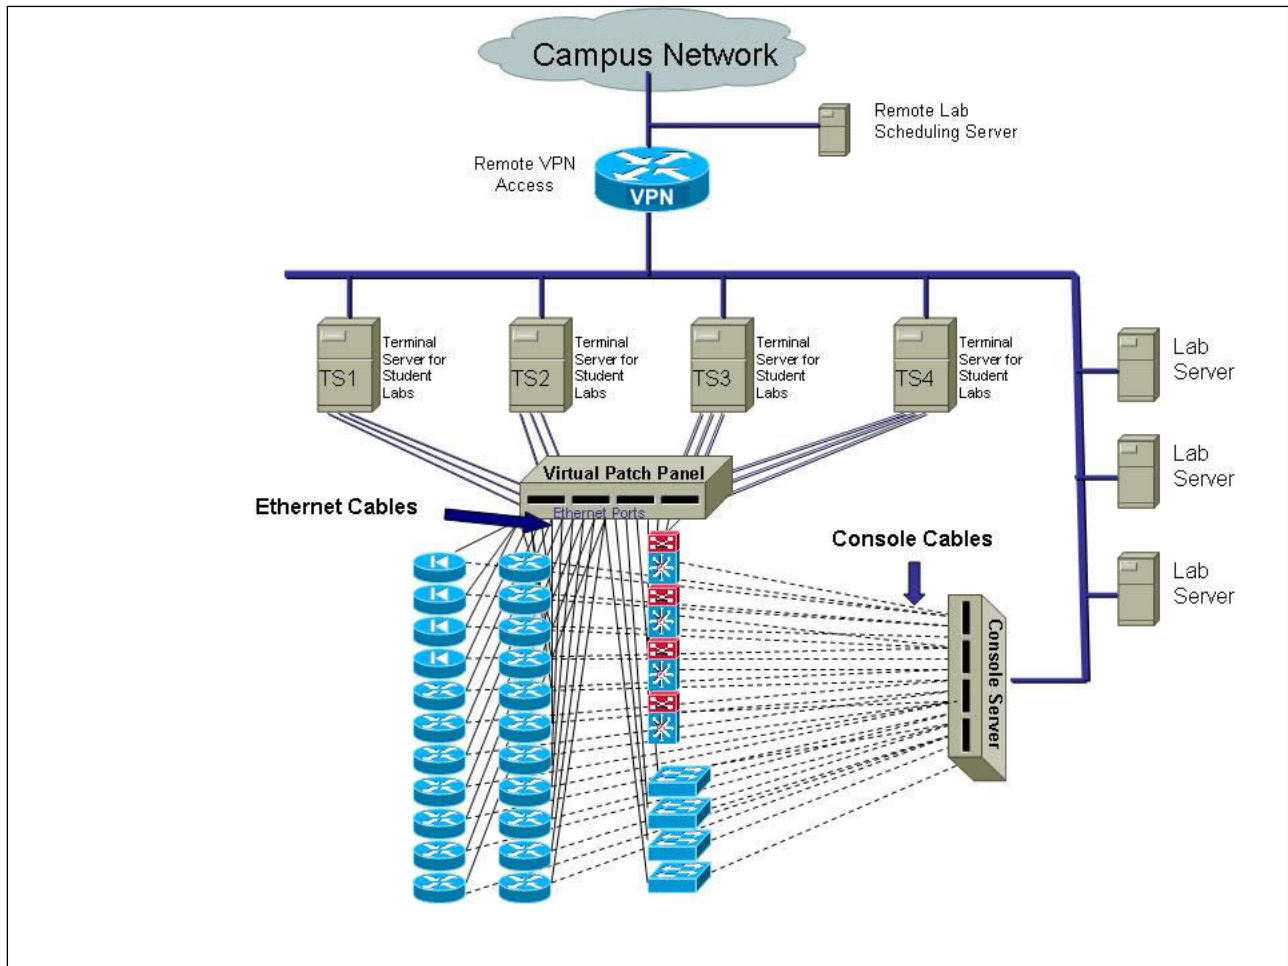


Figure 1. Integrated Virtual Remote Lab (IVRL)

To understand how this lab environment could be utilized we will provide an example. A security student desiring to accomplish a firewall lab opens a web browser and enters the URL of the scheduling server. The web page allows the student to view available times that the firewall lab could be accomplished. The student selects an open time that fits into his/her schedule and then submits the selection. When the appropriate time comes, the student creates a vpn connection to the vpn server. The vpn server communicates with the scheduling server to verify that this student is allowed onto the network and then creates the connection.

Using a web browser or remote desktop software, the student creates a remote session to one of the terminal servers and logs in. The student then proceeds to perform the Pix firewall lab, which utilizes 3 windows virtual machines and a Cisco Pix firewall shown in Figure 2. One virtual machine represents a computer on the Internet, one represents a computer in the DMZ, and the third represents a computer on the internal network. The pix firewall has 3 Ethernet interfaces that need to be connected to the appropriate virtual machines, one for the Internet computer, one for the DMZ computer, and the third for the internal computer. The virtual patch panel has a preset

configuration that makes these connections for the student. The student can now configure the pix firewall using the console server and test the configuration with the virtual machines. An example of these tests would be to configure the pix firewall to only allow HTTP traffic into the DMZ. After making the configuration, the student will be able to test using the virtual machine on the Internet interface and trying to browse to the virtual machine on the DMZ interface. If the student was able to see the web page, the HTTP port was opened correctly. The next step the student would take is to verify that all remaining ports are closed by using a port scanner on the virtual machine connected to the Internet interface on the pix firewall. After finishing the lab, the student logs off the terminal server and closes his/her vpn connection into the lab network.

The environment is then refreshed so that the next student who has scheduled time on the IVRL is ready to go. This is accomplished by downloading the correct images from the lab servers or running the virtual machines off of the network. Vmware has a "snapshot" option that allows the virtual machine to revert back to a predefined image.

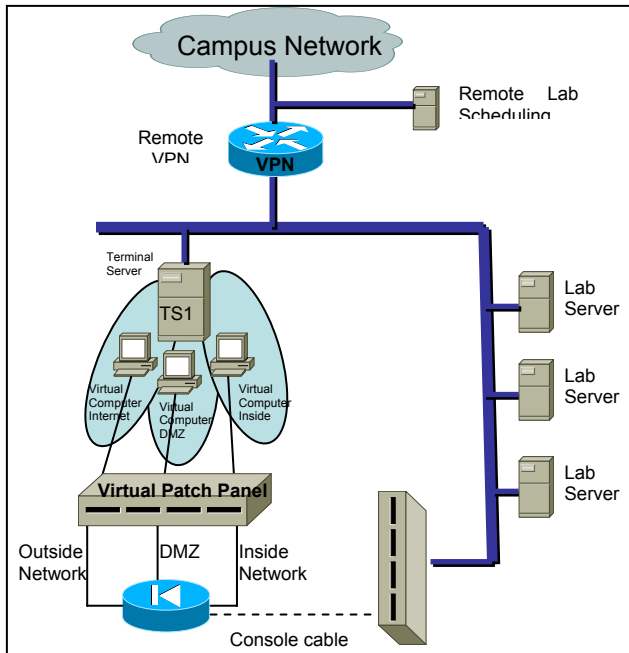


Figure 2

In the above example, a student needed to accomplish a lab for a security course, however; one of the main advantages of IVRL is that other courses can also utilize this environment. For example, an operating system instructor could create a lab to have students configure DHCP, DNS, Email, Directory Services, user administration, and many others. Networking instructors could use IVRL to configure routers and switches. Web design and development courses would also benefit from the IVRL environment where labs could be created to setup web services and learn how to deploy applications.

The number of simultaneous students performing labs depends on the number of terminal servers and networking equipment available. In Figure 1 there are 4 terminal server computers that allow 4 students to perform labs at the same time. In this scenario, assuming a two hour lab time for completing a lab, and those students would be scheduling labs between 8:00am and midnight, this would allow for a maximum of 224 labs performed each week.

4.1 CONCLUSION

This paper discusses the increasing need for remote labs for distance education along with the necessary pieces for implementation. A new design is proposed that utilizes scheduling software, virtual private networks, virtual machines, virtual patch panels and console server technology that is flexible and scaleable.

The vital key to success in a remote lab implementation is to keep flexibility and scalability as the foremost goals. Courses change frequently in the technology fields and labs become obsolete. The ability for an instructor to develop a lab and

deploy it quickly will ease the burden for the instructor and foster new lab creations at a faster rate. In addition to being flexible in design, the environment requires scalability to ensure that additional resources can be added during periods of increased demand.

The IVRL design was implemented for a networking course as a proof of concept. The final project required for this course was to setup and configure seven multilayer switches using the IVRL. Students were able to schedule lab time that was convenient for them and then vpn into the lab network from any Internet location to configure the switches.

The scheduling software that was incorporated into this course was provided by the Center for Systems Security and Information Assurance CSSIA. This software was developed through a grant from the NSF and included a RADIUS server for vpn access.

The technical details for the setup and configuration of this design are available upon request.

5. REFERENCES

Brown, S., & Lahoud, H. (2005). An Examination of Online Lab Technologies. Proceedings of the ACM SIGITE 2005 Conference, Newark, NJ.

Leitner, L., & Cane, J. (2005). A Virtual Laboratory Environment for Online IT Education. Proceedings of the ACM SIGITE 2005 Conference, Newark, NJ.

NCES (1999). Post-Secondary Education Quick Information System. Survey on Distance Education at Higher Education Institutions: 1997-1998.

NCES (2002). Post-Secondary Education Quick Information System. Survey on Distance Education at Higher Education Institutions: 2000-2001.

NCES (2003). Enrollment in Postsecondary Institutions, Fall 2003.

NCES (2004). Enrollment in Postsecondary Institutions, Fall 2004.

Padman, V., Memon, N., Frankl, P., and Naumovich, G. (2003). Design and Implementation of an Information Security Laboratory. Proceedings of World Conference on Information Security Education.

Stockman, M. (2003). Creating Remotely Accessible "Virtual Networks: on a Single PC to Teach Computer Networking and Operating Systems. Proceedings of the 2003 CITC, Lafayette, IN.

Toderick, L., Mohammed, T., Tabrizi, M. (2005). A Consortium of Secure Remote Access Labs for Information Technology Education, Proceedings of the ACM SIGITE 2005 Conference, Newark, NJ.

Yang, T., Yue, K., Liaw, M., Collins, G., Venkatraman, J., Achar, S., Sadasivam, K., Chen, P. (2004). Design of a distributed computer security lab. Journal of Computing Sciences in Colleges, 20 (1), pp. 332-347.