TAKING THE NETWORK ON THE ROAD:
Portable Network Solutions for Computer Security Educators[*]

Timothy Rosenberg[†]
Lance J. Hoffman[‡]

Abstract

Educational institutions that wish to implement a specialized teaching laboratory often have a variety of obstacles to overcome. Some of these are related to adapting existing classroom infrastructure, building the laboratory, and the demands for multiple-use classrooms to maximize the return on investment. In some cases, such as computer security, they must also maintain a controlled environment.

Portable educational networks of computers allow teaching specialized topics that have heretofore required specialized laboratories in existing classrooms. Their flexibility allows them to be used for a variety of content -- operating systems, networking, security and forensics. These systems are in use today supporting undergraduate, graduate and professional education.

Issues in Teaching Computer Security

Adequate teaching of information security often requires some form of laboratory environment where students are able to experiment with new and untested computer programs. In the case of computer security, they may view malicious code, execute network attacks, and penetrate enterprise servers. The very nature of these requirements dictates a need for a separate, dedicated network, in order to minimize potential damage to production systems or the public computing community at large. This network must be unconnected to the Internet, campus, and corporate networks.

Traditional approaches to this problem have involved retrofitting existing classrooms with a computer network and server infrastructure (typically housed in a large equipment rack cabinet in the front of the room). Workstations are then supplied at tables for student use. While this solution works, it requires a significant investment in hardware that is tied physically to an individual room. Any maintenance and laboratory preparation must be scheduled around the existing class schedule. The more the room is used (which is good for return on investment), the less time is available to maintain and upgrade the environment for various classes in the lab.

Another problem with the permanent classroom solution is flexibility. One could simply take a full sized equipment rack and fill it with many servers. Then, servers can be segmented or 'carved off' for use by different courses, thus preventing the damaging of systems that are to be used for other courses. Using this solution, the class content is

---

[*] Work supported in part by grants from the U. S. Defense Department and the National Science Foundation.

[†] trosenbe@gwu.edu, Computer Science Dept., The George Washington University, Washington DC 20052.
[‡] lanceh@gwu.edu, Computer Science Dept., The George Washington University, Washington DC 20052.

limited to the systems in the rack and the instructor must ensure that all are using the proper system for each course. We discuss below a more efficient approach that allows for easily configured systems that can be easily rebuilt or modified. Furthermore, the systems are physically segmented preventing any possible 'cross-contamination' of courses.

A further complication is scaling. Ideally, the classroom/lab will be large enough to teach as many students as possible. Unfortunately, this means dedicating a large amount of floor space in the hopes that there are sufficient classes to use the lab and a high enough student enrollment to offset the expense of building and furnishing it. The lead time on "building out" a classroom often requires that some or all of the equipment in it be purchased and installed before any realistic estimates of class size are available, often before registration for the course is opened.

A Portable Solution

We here describe the Portable Educational Network (PEN), a solution to the problem that is portable, scalable, and flexible.

1. The system is portable. A single person can wheel this lab into any classroom, open the front and back of the case, connect the power supplies into the wall and turn everything on. Everything is pre-wired and pre-cabled. The only requirement is access to power.
2. The system is scalable. The laptops are standard systems from any major provider and can be quickly staged out for the classroom using any commercial hard drive imaging software (e.g., Symantec Ghost Enterprise). It can quickly contract or expand, depending on class size. Unused equipment can be sent back to a central pool and additional equipment can be added (up to a maximum) to accommodate more students.
3. The system, in a standard computer equipment rack, is modular. Any rack-mountable piece of hardware can be installed into the unit. The generic system architecture (Figure 1) is very close to that of the Internet or a standard local area network with Domains for the Internet, Enterprise and Administration.
    a. The Enterprise Domain is a representative architecture placeholder for an organization's presence on the Internet. In the real world, this is typified by a collection of servers providing web content, email and other network services. Likewise in the PEN, this Domain is where the instructor would place servers that represent a hypothetical internet presence.
    b. The Internet Domain is a generic descriptor for any system or network outside an organization's control which does not reside in the Enterprise Domain. In the real world, this is where the home user, remote office and the 'un-trusted' internet reside. In the classroom, this would be populated with student workstations with client-side software such as web browsers, email programs and VPN clients.
    c. The Admin Domain is where the necessary support features reside. This includes system log collection, server backups and source files and monitoring and reporting software is installed and used. In the lab

environment this is typically one or two very large servers that are used to store the data needed to maintain and support the overall environment.

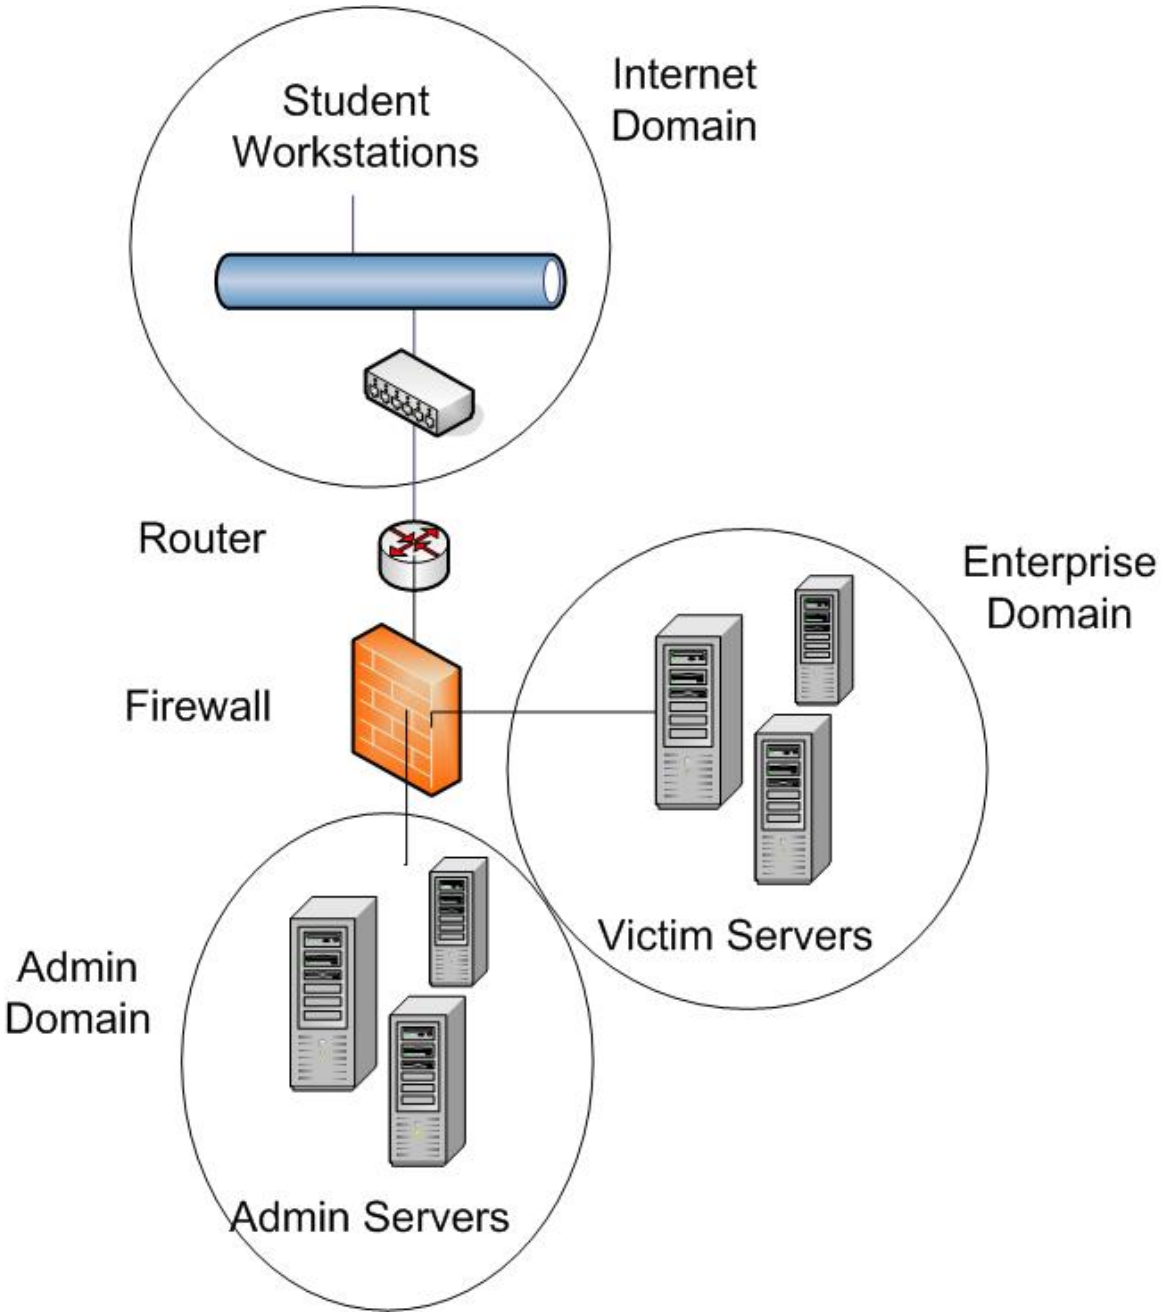

Figure 1. Portable Educational Network Architecture (Generic)

Network Architecture

As an example, at The George Washington University, the portable laboratory shown in Figures 2a and 2b (the GW-PEN) is used to teach computer network defense. It ameliorates a shortage of classroom and laboratory space by residing in a 15RU* network rack housed in a hardened, wheeled case.

Figure 2a. PEN in its case

Figure 2b. PEN on the move in D.C.

The GW-PEN uses standard enterprise grade components to simulate a standard Internet connection. The overall theory of this design for computer security educators is "students as attacker and instructor as defender". The instructor can display the results of the students' attacks and show in real time the impact on the network. Furthermore, the instructor can walk the class through the appropriate countermeasures and show their effectiveness.

Figure 3 shows an actual implementation of the generic Figure 1 schematic. The Internet and Enterprise Domains have been changed to the Attacking and Victim Domains, while the Admin Domain has stayed the same. These Domains are used to teach computer security, and could be easily changed to implement any type of domain necessary to teach a variety of computer lab based courses. The Attacking Domain is where the students reside for the entirety of the class. Under ideal financial and space conditions, each student will be assigned to a desktop or laptop. From this domain, the students launch attacks against the Victim Domain, a collection of production operating systems that are available for the Attacking Domain to attempt to penetrate. The Attacking Domain is protected using an enterprise firewall and Intrusion Detection System.

* An RU or 'Rack Unit' is a standard method for measuring the height of rack mountable equipment. A single RU is 1.7 inches in height.

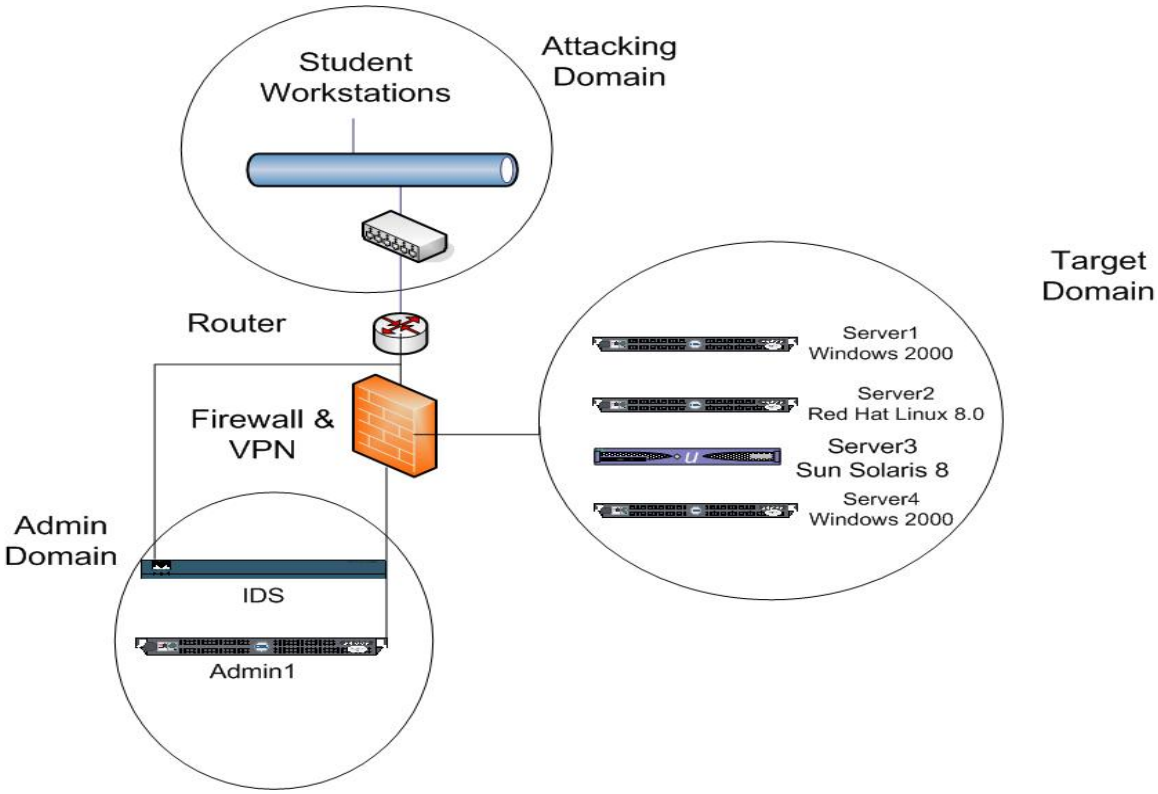

Figure 3. PEN Architecture for Computer Security Class

Furthermore, a Virtual Private Network (VPN) concentrator is available to offer secure access to the systems behind the firewall. Finally, the Admin Domain is used for configuring the network devices, collecting log information, and miscellaneous administrative functions. In our experience, it has been useful and necessary as a repository for code, operating system images and network device configuration files.

Lastly, for the purposes of teaching network security a Network Intrusion Detection System has been added to the Admin domains. Readers will note the second connection from the IDS appliance to the network segment between the router and the firewall. This line represents the capture or sniffing interface of the IDS. Most network IDS installations have one interface that is used for management and one or more interfaces that are used for capturing network traffic for analysis. Depending on the IDS solution that is installed, the line denoting the capturing interface may be in one or several locations. This is completely dependant on a specific PEN build and in some cases actually changes through out the course to reflect the different alerts that result in IDS placement.

PEN Rack Components

Figure 4 shows some rack details of the GW-PEN. From the top of the rack to the bottom, the systems in the PEN are: (1) A 15 inch flat panel monitor, keyboard, touchpad and 8 port KVM switch all in one RU space. The next three systems (2, 3, 4) are Dell PowerEdge 650s running a variety of production operating systems including Windows 2000 Server and Red Hat Linux. These are three of the four servers in the

Victim Domain. The next two devices (5 & 6) are switches for use in the Target and Administrative Domains. The next three devices (7, 8, 9) are security countermeasures -- Virtual Private Network (VPN) concentrator, network Intrusion Detection System (IDS) appliance, and a multi-interface firewall. The final systems are (10) a SUNFire server for the victim domain, a server for the Administrative Domain (11) and 2RU of storage space (12).

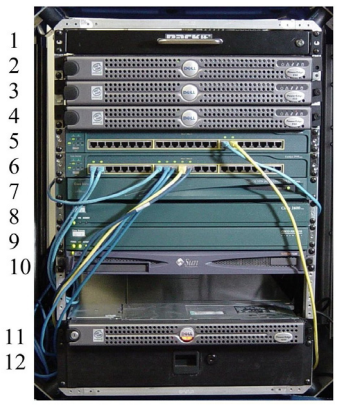

Figure 4. PEN rack details

Costs and Tradeoffs

The table below summarizes some of the server options and their advantages and disadvantages in their deployment.

Server	Cost[†]	Form factor (RUs)	Advantages	Disadvantages
Any enterprise server[‡]	Around $2,000 per unit, one server per unit	1RU in height	Support contracts (4 hour on site, 3 year, etc)	Cost , weight, white noise generation, possibly proprietary rack rails
High density collocation systems[§]	Around $650 per unit, one server per unit	1RU in height	Generally quieter than enterprise grade hardware, cost	Support contracts are limited, for the cost, these are 'disposable' systems
High density industrial rackmount systems[**]	Around $2,000 per unit, three servers per unit	1RU in height	Quiet, three servers using the space of one	On site support contracts are generally limited, support relegated to shipping back to manufacturer

[†] This cost is as of Spring 2004, and does not include the shipping cost or the cost of the operating system.
[‡] An example of a standard enterprise server would be any single or dual processor server in a 1RU form factor.
[§] An example of a typical high density collocation system is a 1RU form factor server designed for low cost services such as web hosting. These systems are distinguished from an enterprise server by their cost and limited feature set.
[**] A typical high density industrial rackmount system is any platform using SBC's.

Future Work

When using enterprise servers, the amount of white noise that is generated makes it difficult to hear students beyond approximately six feet. Some of our recent work has focused on a smaller machine with increased flexibility and portability and reduced cost and white noise. We can now build a mini-PEN (see Figure 5) using Single Board Computers (SBC), in a hardened, mobile case that is one third the overall height of the original PEN. An SBC is a complete computer on a single board; CPU, memory and I/O (drives, keyboard, mouse, video, etc.) are all on a single board. The advantage of the SBC is that each server has its own motherboard, CD, and hard drive and is easily connected to any standard KVM switch. The net result is to put three servers in a 1RU space in the rack. Furthermore, the cost of the mini-PEN is half that of the full-size PEN.

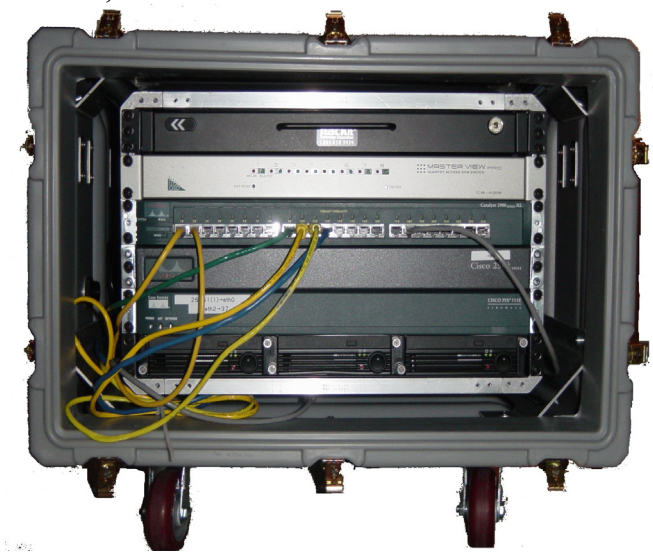

Figure 5. The MINI-PEN in its case

Another way to maximize the flexibility and minimize the size is by combining devices. Manufacturers of enterprise routers offer modules to extend the functionality of their devices without using more space in the rack. For example, it is possible to equip a Cisco 2600XM series router with a Virtual Private Network accelerator card and a Network Intrusion Detection card, thereby removing two more devices from the rack without sacrificing feature functionality.

This architecture modification may change some ways of teaching the class. For example, with two mini-PENs, the overall focus of the network security class changes from students attacking and instructor defending to a head to head environment, in which the class is divided into two teams and each is given a mini-PEN. Their responsibility during the class is to attack the victim servers on the other mini-PEN, while their own servers are being attacked, and then learn from these exercises how to defend computer systems from such attacks.

Summary

PENs provide several advantages over the building of a permanent, fixed classroom and are today often less expensive than the costs of physically configuring and wiring a special-purpose laboratory. They can be safely used to teach several courses over a given time period, using where necessary course-specific supplemental equipment. Everything that the instructor sees can be projected on a screen or wall in the room. PENs were first built commercially in 2001[††] and first used, to our knowledge, at The George Washington University in 2003. There are currently several versions of this lab at a variety of educational and governmental institutions.

References

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.

L. J. Hoffman, R. Dodge, T Rosenberg, and D. J. Ragsdale, "Novel Approaches for Information Assurance Laboratories," the Journal of Information Security, May 2003.

J. M. D. Hill, C. A. Carver, Jr., J. W. Humphries, and U. W. Pooch. Using an isolated network laboratory to teach advanced networks and security. In *Proceedings of the 32nd SIGCSE Technical Symposium on Computer Science Education*, pages 36–40, Charlotte, NC, USA, February 2001.

J. Mayo and P. Kearns. A secure unrestricted advanced systems laboratory. In *Proceedings of the 30th SIGCSE Technical Symposium on Computer Science Education*, pages 165–169, New Orleans, USA, March 24-28 1999.

Using an Instructional Operating System in Teaching Computer Security Courses □ Wenliang Du, Systems Assurance Institute, Department of Electrical Engineering and Computer Science, Syracuse University, http://www.sai.syr.edu/facultypapers/Instructional%20Operating%20System.pdf

[††] For references to commercial use, see http://www.whitewolfsecurity.com.