published as: J. Hu, Christoph Meinel, M. Schmitt: Tele-Lab IT Security: An Architecture for Interactive Lessons for Security Education;
Proceedings of the 35th SIGCSE technical symposium on Computer science education (SIGCSE 2004);
Norfolk (Virginia, USA), 2004, pp. 412-416.

Tele-Lab IT Security: An Architecture for Interactive Lessons for Security Education

Ji Hu
Department of Computer Science
University of Trier, Germany
hu@ti.uni-trier.de

Christoph Meinel
Department of Computer Science
University of Trier, Germany
meinel@ti.uni-trier.de

Michael Schmitt
Department of Computer Science
University of Trier, Germany
michael.schmitt@teststep.org

ABSTRACT
IT security education is an important activity in computer science education. The broad range of existing security threats makes it necessary to teach students the principles of IT security as well as to let them gain hands-on experience. In order to enable students to practice IT security anytime anywhere, a novel tutoring system is being developed at the University of Trier, Germany, which allows them to get familiar with security technologies and tools via the Internet. Based on virtual machine technology, users are able to perform exercises on a Linux system instead of in a restricted simulation environment. This paper describes the user interface of the Tele-Lab IT Security, its system architecture and its functional components.

Categories and Subject Descriptors
K.3.2 *Computer and Information Science Education*: Distance Education, Human-Computer Interface, Lab Environments

General Terms
Design, Human Factors

Keywords
IT Security, Tutoring System, Virtual Machine

1. INTRODUCTION
IT security education is an important activity in computer science education. The broad range of existing security threats makes it necessary to teach students the principles of IT security as well as to let them gain hands-on experience. Recently, many universities have integrated computer security lectures into their curricula and developed academic security laboratories.

In order to enable students to practice IT security anytime anywhere, the Telematics research group of the University of Trier, Germany, is developing an on-line tutoring system, called *Tele-Lab IT Security*. It has two major characteristics: first, it is a web-based tutoring system which introduces students to fundamental IT security concepts. Second, it provides an on-line virtual laboratory in which students are able to gain practical experience. Based on virtual machine technology, a remote machine can be assigned to a student and administrator rights can be granted to him without endangering the stability and security of the tutoring system itself.

The educational content of the Tele-Lab IT Security covers general security topics as well as specific aspects of the Linux operating system. Its major topics include cryptography, digital certificates and secure email, authentication, and security scanning. For every topic, the Tele-Lab IT Security provides a set of exercises. These guided exercises are performed on a real Linux system with standard tools rather than in a restricted simulation environment. This approach allows students to easily apply their knowledge to production systems later.

This paper is structured as follows: Section 2 provides background materials and discusses related work. Section 3 describes the concepts of the Tele-Lab IT Security. The system architecture and functional components are described in Section 4. Section 5 gives a summary and presents future work.

2. BACKGROUND

2.1 Related Work
We have tried to find similar projects concerned with on-line security laboratories but we found only a few simulation systems for security training. The *ID-Tutor* [6] and the intelligent tutoring system (ITS) described by Woo et al. in [8] familiarize students with intrusion detection. ID-Tutor creates audit files with information on user logins and executed commands. The user has to decide whether an intrusion has occurred, and, in case of an intrusion, he/she must resolve the problem. The ITS by Woo et al. is similar to ID-Tutor, but it generates its missions from a knowledge base.

With both tools, the user performs his exercises in a simulation environment. For practical reasons, such a simulator can model a real system only to a very limited degree. As a consequence, a student is not able to apply standard software tools and cannot get feedback in a real-life context.

2.2 Related Techniques
Virtual Machines: Virtual machine (VM) technologies [4] are crucial for the tutoring system to be used remotely. A virtual machine is an abstraction in software of a physical machine that runs as a standard user application. Many virtual machines can be started on a single host system.

A distinct IP address can be assigned to each of them, i.e. they can be integrated into the network and accessed from outside the host. From the user's point of view, a VM looks just like any regular machine on the network. The destruction of a virtual machine does not result in any adverse effect on the underlying host system. Therefore, it is possible to grant administrator rights to an ordinary user, provided that a firewall/trust configuration stops the (super-)user from getting other machines under his control.

VMware [7], a commercial VM product, and User Mode Linux (UML) [1], an open source project, are two major VM implementations. VMware creates a software PC machine on a physical machine. Any operating system can be installed on such a virtual machine. The User Mode Linux implements a Linux virtual machine that runs on a Linux host. Due to its resource-friendly nature, many virtual machines can share one PC with decent performance.

Remote Access: In order to run graphics applications remotely on a virtual machine, a client needs a remote access program to display output from and send user input to the remote machine.

It is possible to embed a special Java applet into a browser as a remote access client. This feature is supported, e.g. by WeirdX [3] or by a VNC applet client [5]. WeirdX supports the X protocol and can be integrated into web applications seamlessly.

The VNC applet is a desktop viewer that is able to interact with a VNC server on the remote machine through the remote frame buffer protocol (RFB). This protocol is used to collect user input from an applet, encode desktop displays on the server side, and send them to the applet. Experience has shown that the VNC protocol results in a better performance than WeirdX.

3. CONCEPTS

The Tele-Lab IT Security is based on a standalone computer-based tutoring system, called *E-learning platform IT security* (LPF), which was also developed at the University of Trier [2]. The LPF is designed to equip a security laboratory. Its concepts are illustrated in Figure 1.

Each LPF system is installed on a Linux machine and integrated with a web interface. A registered student can log into the system and access security information via a browser. When the student is requested to perform exercises, such as cracking passwords and security scanning, she or he completes the tasks through a shell terminal or an X-based interface. The results can be either answers submitted to the web server or some direct changes to the Linux system. LPF usually needs to execute scripts to evaluate these results. The evaluation must be done by invoking real operations to compare answers or to trace the changes made to the system by the student because the concrete exercise contents are created on the fly.

However, this scheme introduces a potential risk, i.e. some exercises require a student to perform system operations in the underlying Linux system. The student would corrupt the complete system if she or he misuses the privilege right that we assign her or him temporarily. When a partition is corrupted, we have to restore it by a backup of the hard-disk partition, or reboot the entire system from a pre-built CD-ROM. Unfortunately recovery from failures is quite inconvenient for practical uses since it interrupts learning processes and spoils the student's enthusiasm.

From the point of view of e-learning or tele-teaching, we need to move the tutoring system on-line. The current development of the Tele-Lab IT Security adopts a novel scheme which leverages virtual machine technology to enhance the reliability of the system and to degrade its maintenance cost. The (simplified) structure is shown in Figure 2.

The web server, called "tutor", is the portal to the system. If a student needs to perform exercises, she or he will be redirected to a virtual "exercise server" on the host. Afterwards, the student can get privilege rights and perform his/her exercises. After the student has finished his exercises, the tutor takes over the interaction with the student and terminates the user access to the virtual machine.

If a crash is detected, the tutor assigns a new available virtual machine to the student and continues the interaction. The failed virtual machine can be refreshed by a backup copy and restarted in the background. Multiple virtual machines can be installed on the same host PC where the tutor is located, or on a separate host which acts as an exercise server cluster. In addition, there should be a remote access interface (including a remote access agent and its client) for a user to log into the virtual machine and perform exercises.

Though the idea does not sound complex, the approach requires resolving a series of technical problems in practice. For example, virtual machine management, connection transitions and the interface for the Internet users must be considered carefully.

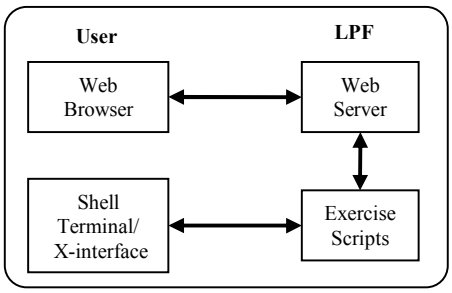

Figure 1. Standalone tutoring system.

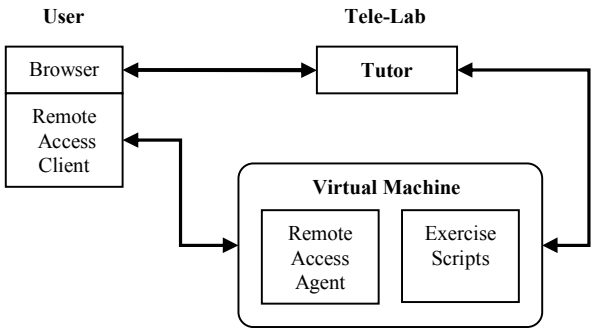

Figure 2. Virtual machine based Tele-Lab.

4. SYSTEM ARCHITECTURE

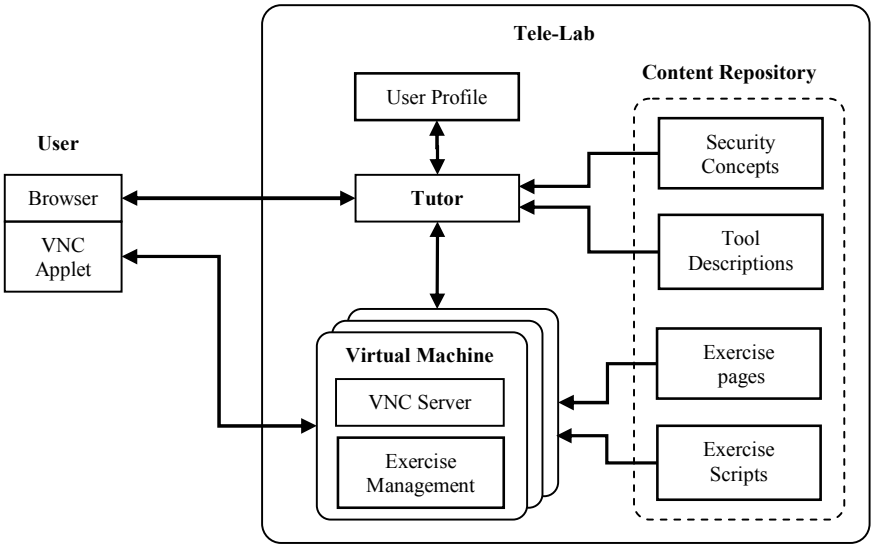

Figure 3. The Tele-Lab system architecture.

The architecture of the Tele-Lab IT Security concerns the conceptual knowledge of the security domain, procedural knowledge for exercises, user performance in learning, and pedagogical strategies of the tutor. Furthermore, the development must focus on the creation of a convenient user interface.

The main components are shown in Figure 3. The *content repository* is an IT security knowledge base. It consists of a collection of web pages and scripts. The *tutor* navigates students through security topics in the content repository, assigns exercises to them, and manages their performance. It is also responsible for connection transitions and the maintenance of the virtual machines. The *virtual machine* is a simplified virtual Linux system that is equipped with a tiny web server. It interacts with a user via a VNC applet embedded in the browser window. It also evaluates the user's results. The *user profile* keeps track of the user's knowledge at every stage in the learning process. The web browser provides a uniform and user-friendly interface. In order to perform exercises, the user can easily get access to the system via the applet.

4.1 The Content Repository

The content repository is a knowledge base that consists of three types of contents:

1. Descriptions of IT security concepts
2. Descriptions of security tools
3. Security exercises

The first two types represent declarative knowledge that is presented to a user in the form of web pages. The security exercises represent procedural knowledge. Its implementation is a complicated task because it must be given in terms of scenarios that require step-by-step interactions with the user.

The content repository is organized as follows. The basic unit is the section. It covers a complete learning item or several closely related items. A section is also the basic unit for measuring the performance of the user. Completing a section means that the user has succeeded in acquiring some given knowledge or skills. The Tele-Lab IT Security has three types of sections including concepts, tool usages and exercises. Every section consists of one or more pages. The exercise sections comprise some additional scripts. The exercise pages and scripts are shared by all virtual machines. It is up to the tutor to decide which virtual machine is used for each individual section.

Multiple sections are combined into a chapter that represents a security topic. In most cases, a chapter introduces security concepts first. Then it explains related tools or commands. Afterwards, the user is asked to perform some practical exercises.

4.2 Exercise Management

Exercises are specified as Perl or PHP scripts in the content repository. Typically an exercise takes place in the following three phases:

In the first phase, the working environment, i.e. the Linux operating system, is configured. For example, if a user has to perform a security scan, some services are activated so that the user can get effective results. The next phase deals with generating questions or tasks and passing them to the user. Where possible, these tasks are created dynamically. I.e. all users do the same type of exercise but with different detailed content. For example, for password cracking, a UNIX *passwd* file (that the user must decrypt) is generated at run-time. After the user completes the tasks, Tele-Lab evaluates his or her result in the third phase.

The preparation, execution and result analysis of exercises in a real environment require intensive efforts. Here we describe the necessary procedures in detail using a real example, the secure email exercise. This exercise is intended to make users familiar

with digital certificates so that they can sign and encrypt emails. The user interface for this exercise is shown in Figure 4.

Phase 1: Preparation of the working environment:

- Clear previous settings and prepare necessary working directories.
- Set up a local mail server.
- Create a virtual partner, called *Alice*, which can communicate with the user by email. Her Linux account and email settings must also be configured.
- Create a certificate authority by *OpenSSL* commands.
- Issue certificates to Alice and the user, install the certificate for Alice, and guide the user to import his or her certificate to a mail client (e.g. a Mozilla mail-client).

Phase 2: Generation of tasks:

- Create a random message with a signature on behalf of Alice.
- Send the message to the user.
- Ask the user to
  1. verify the signature attached to the message, accept and trust Alice's certificate.
  2. reply to Alice's message with his or her own signature.
  3. send Alice a message encrypted in the public key contained in the Alice's certificate.

Phase 3: Evaluation of the answers:

- Fetch a message from the Alice's mailbox.
- Verify its signature to see whether it has been sent from the user and matches with the original.
- Try to decrypt the message using the Alice's private key which corresponds to her certificate.
- Record the (un-)successful completion of the exercise in the user profile.

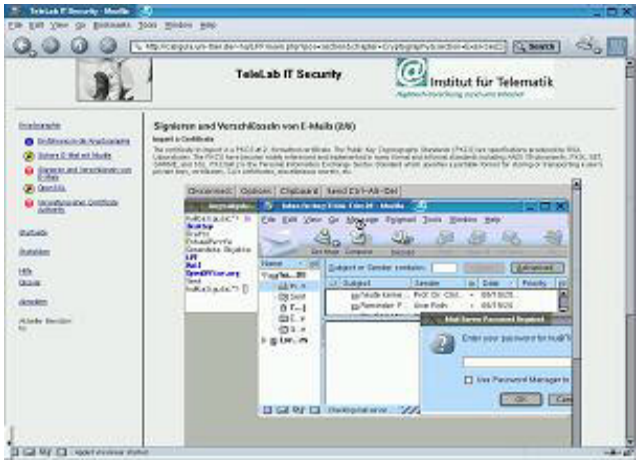

Figure 4. The user interface for the secure email exercise.

4.3 The User Profile

It is helpful if a tutoring system "knows" individual students and keeps track of their performance. This feature in Tele-Lab is supported by the maintenance of a user profile. The user profile contains two types of data: the first one is personal information, such as accounts and profiles. The profile classifies users into three categories: administrators, ordinary users and IT students. Tele-Lab defines different sets of security topics for each category. The second type of data includes records of the completed sections, the time spent on each exercise and so on. The data contained in the user profile can be used to analyze a user's performance and to present statistics on the current status. The user profile also stores the Linux home directory of the user. The directory is kept on the host PC and accessible from virtual machines. Even if a virtual machine is corrupted, the user can still use his or her profile to continue learning on another machine.

4.4 The Tutor

The tutor is a core functional component in the Tele-Lab. Its first function is navigation. The tutor presents educational materials in a structured manner to a user and updates the user profile in time. Its second function is to assign the user exercises and manage connection transitions between the tutor and virtual machines.

Before starting a Tele-Lab session, a user must register or log into the system with a valid account. At the same time, his or her user profile is restored by the tutor and the home directory becomes ready for use. Then the tutor provides a list of available chapters from which the user can choose one topic. When the user enters a chapter, the tutor creates a navigation bar on the left side of the web page that lists all sections and represents their type by a small icon (see Figure 4). When the user is going to perform an exercise, a series of steps are required for virtual machine maintenance and connection transitions. This procedure is described as follows.

The first thing that the tutor should do is to save the current user profile into the user browser's cookie file. This cookie is accessible by virtual machines in the same network domain. In this way, the user record and user authentication can be passed to a virtual machine. Next, the tutor chooses an available virtual machine in a list and assigns it to user, and redirect the connection to it. When the user opens the first exercise page, an applet is sent to the browser and launched as a remote desktop viewer. The virtual machine opens a system session with the user account and dumps its display to the applet. Thus, the user can do all work in a browser window. The user profile can also be updated at every step in an exercise.

The reverse process is similar to the procedure above. Firstly, the virtual machine returns the user profile to the tutor via the cookie. It also closes the user session and notifies the browser to close the applet window. Then the connection is redirected to the tutor. If the user used a privilege access to the virtual machine, or a failure is detected, the tutor will refresh this machine, that is, restart it with a fresh copy of its file system.

5. SUMMARY AND OUTLOOK

In this paper, we have presented an on-line tutoring system, Tele-Lab IT security. Tele-Lab improves existing security education activities in the following ways: it offers users a real system environment instead of a limited simulation environment; it has a navigation mechanism that presents contents and creates exercises dynamically; its user interface is based on a pure web browser interface. The tools and programs needed in exercises are available via the browser; the system is based on virtual machine

technology which enhances the reliability of Tele-Lab and degrades its maintenance cost.

Up to now, the Tele-Lab IT Security is still under development. Tests indicate that the User-Mode-Linux is an excellent virtual machine implementation. The repository is located on a shared NFS server which also stores the user profiles, the user home directories and all exercise scripts. For security reasons, we grant write access to the NFS directories only to the tutor. A VNC applet is integrated into the user interface, so that students are enabled to perform exercise via a browser. The behavior and the performance of the user are recorded in the user profile. The evaluation of this information for adapting the teaching contents dynamically feature of the student modeling is considered as a direction for future research.

6. REFERENCES

[1] Dike, J. A user-mode port of the Linux kernel. In *Proceedings of the 4th Annual Linux Showcase and Conference (Usenix 2000)*, Atlanta, GA, 2000.

[2] Hu, J., Schmitt, M., Willems, Ch. and Meinel, Ch. A Tutoring System for IT Security. In *Proceedings of the 3rd World Conference in Information Security Education*, Monterey, CA, 2003, 51-60. Kluwer Academic Publishers, 2003.

[3] JCraft Inc. WeirdX. Available at: http://www.jcraft.com/weirdc/index.html, 2002.

[4] McEwan, W. Virtual Machine Technologies and Their Application in the Delivery of ICT. In *Proceedings of the 15th Annual NACCQ*, Hamilton, New Zealand, 2003.

[5] Richardson, T., Stafford-Fraser, Q., Wood, K. R., and Hopper, A. Virtual Network Computing. *IEEE Internet Computing, Vol.2,* No.1 (Jan/Feb 1998), 33-38.

[6] Rowe, N. C. and Schiavo, S. An Intelligent Tutor for Intrusion Detection on Computer System. *Computers and Education*, 1998, 395-404.

[7] VMware Inc. VMware. Available at: http://www.vmware.com/, 2002.

[8] Woo, Ch., Choi, J. and Evens, M. Web-based ITS for Training System Managers on the Computer Intrusion. In *Proceedings of the 6th International Conference on Intelligent Tutoring System*, Biarritz, France and San Sebastian, Spain, 2002, 311-31.