Calls for more education in computer security increased after 9/11. But what does it take to effectively educate the workforce?

Computer Security Education: Training, Scholarship, and Research

Matt Bishop, Department of Computer Science, University of California, Davis

Since 9/11, we are increasingly aware of threats to security and computer system vulnerabilities. We are also more aware of the need to educate the workforce quickly and effectively.

Traditionally, computer security education falls into two distinct classes. The first is *training*, marked by an emphasis on particular systems, situations, or environments rather than broad principles. The second is *scholarly (or scholarship)*, marked by an emphasis on underlying principles, concepts, and their application.

Research in computer security provides the needed breakthroughs enabling us to meet new and evolving threats. But research done within the context of training differs from scholarly research. Effectively preparing the workforce to meet the challenges today means using the strengths of both.

Training education

Training classes serve a specific purpose. Some training classes help prepare students for a particular examination, such as a certification exam. Training courses emphasize information specific to the particular system or environment rather than general information. Their goal: prepare students to apply what they have learned when they encounter that specific system, environment, or situation.

Training classes emphasize the results of applying the principles and concepts, not the principles and concepts themselves. The course might not even identify the principles underlying the mechanisms. But the students will know what to do in specific situations, and have a good idea of how to handle different, but similar, situations.

A training course on Unix security

Consider a training course on Unix security. While mentioning general principles such as the Principle of Least Privilege and the Principle of Psychological Acceptability,[1] the training course will emphasize their application to the Unix system. The Principle of Least Privilege will lead quickly into a discussion of the creation of subordinate system management accounts so system administrators can avoid using the omnipotent root account. Limiting the number and power of network servers that the system runs will be another component of the discussion.

The Principle of Psychological Acceptability will lead to a discussion of the suitability of various authentication mechanisms for different environments, as well as methods to provide users with timely assistance and working with them to find ways they can be productive without violating security policies.

As another example, a discussion of network security on Unix systems might emphasize the need to use mechanisms (such as wrappers) to log information about the origin of each request. A good course will discuss the lack of a general authentication mechanism in IPv4, the most common Internet protocol, and describe the difficulties of tying connections to their point of origin in the absence of cryptographically based authentication mechanisms. But the discussion will focus on what these mechanisms can provide in the context of authentication the origin of connections, and not the broader issues of the difficulty of providing an Internet infrastructure to support cryptographic authentication, or the mechanisms in IPv6 that will ameliorate this problem.

The role of research in training education

The goals of training guide the research that training institutions and associated research groups perform. This research seeks to uncover the information that the training courses should provide and to learn new ways to apply existing technology to meet new trends. A security course should provide information about effective prevention and detection mechanisms and countermeasures. This enables the students to take precautions in advance of an attack. The precautions are intended to stop or ameliorate the attack and to warn the system administrators that the attack is taking place. The administrators can then take the appropriate steps to thwart the attack, contain it, and report the attack to the appropriate authorities.

The rise of denial of service attacks is a good example. Several years ago, few training courses discussed how to block access to a site, because so few such attacks occurred. As the attacks became more common, training courses began to add information describing the attacks, ways to detect them, and how to counter them or ameliorate the damage. Now, training courses dealing with networks routinely describe not just

the simple denial of service attacks, but also the more complex distributed denial of service attacks.

One part of research in this area studies trends in attacks. The organizations gather information from many sources. The data lets the researchers analyze the most recent attacks they encounter, giving them information on the nature of current attacks. By combining this with data on past attacks, the researchers might be able to establish trends. The instructors take this information and modify their courses to provide up-to-date information on these attacks, the threats they realize, and the vulnerabilities they exploit.

A second part of this research is to use existing technologies in new ways. The goal is to provide insight into improving the security of an existing system without requiring untested, unknown mechanisms. A good example of this is Bastille Linux, a project to develop a set of tools to harden an off-the-shelf Linux system.[2] The changes make the Linux system more difficult to break into in many environments. The tools are simple shell scripts and programs that do not present any new ideas or technology, but make effective use of existing ideas and technologies. This approach contributes to the use of more secure systems.

The drawback of this type of research lies in its short-term effects. Many systems are left unpatched against these attacks. The research does not explain the more fundamental problems of why the attacks worked in the first place and why these attacks still work. It does not help develop new systems resistant to attacks in general—only to particular attacks. The research reveals the exploited vulnerabilities. This is, of course, an important contribution, because it provides data that can guide the dissemination of information and help an organization establish priorities for patching vulnerabilities. But it does not speak to the issue of how to prevent the introduction of these vulnerabilities in the first place. Similarly, the use of existing technologies to harden existing systems does not eliminate inherent vulnerabilities. For example, Bastille Linux does not eliminate the root user on Linux systems. Fixing that requires the development of new technology.

Scholarly education

Scholarly education focuses on principles and concepts and their application. A scholarly course in computer security would not provide the focused attention to topics likely to be asked on a certification exam. The course would provide an introduction to the ideas underlying computer security, and how those ideas are translated into policies, procedures, and mechanisms. A key feature is the focus on the ideas themselves.

Once general concepts are understood, the student examines how to apply them to specific situations. However, the applications are drawn from a wide variety of environments and systems, rather than from one particular environment and system. This distinguishes the applications from those offered in training courses. The applications are chosen to teach the student how to generalize and to work with a variety of environments.

Consequently, after completing a scholarly course on computer security, students typically need some training for specific environments and systems because, in all probability, they haven't learned about those systems. However, the student will probably require less training than other students who never studied the material taught in a scholarly class.

Furthermore, once trained, the students can apply their deeper understanding of concepts and principles to find improvements to current techniques and to suggest alternate approaches. In general, students with a background in scholarly education are more flexible and have a broader, deeper understanding than students with training limited to specific systems.

A scholarly course on computer security

Consider an example of how a scholarly course would tackle the following issue. Michael Harrison and his colleagues demonstrated the impossibility of devising a generic algorithm to determine if a generic systems is secure.[3] A scholarly class would examine this result, and explore restrictions on the definition of "secure" and the nature of the systems that make devising such algorithms possible. Similarly, the class would study computer viruses in the context of malicious logic, and examine how defenses against computer viruses relate to defenses against other forms of malicious logic. The emphasis is to understand the general concepts, and not how to set configuration parameters on a particular system to prevent a computer virus from damaging that system.

The role of research in scholarly education

Scholarly research supports the notion of generality over specific technologies. A good scholarly research project seeks to prove or disprove something, or to understand a phenomenon better. The researchers better understand the limits to technology and to ideas, and create new approaches to solving problems.

For example, many organizations have several subordinate groups, each having its own security policy. The organization wants to derive a global security policy that meets the needs of all subordinate groups. This speaks to a more fundamental question: what are the conditions under which one can create a security policy that combines the security policies of the subordinate groups. Answering the fundamental question will provide an answer that many groups and organizations can use. It speaks to the broader issue of composition of security policies.

As a second example, consider how to test a system for security. For any particular system, and for a given definition of "secure," auditors can determine if the information on it is secure. It seems reasonable to believe we could modify these techniques to apply to any system, for any definition of "secure." This would solve the problem of determining whether any system is secure. Unfortunately, research demonstrates that the general question of information security is unanswerable.[1] No such algorithm, or family of algorithms, exists. This is an example of the benefit of scholarly research.

A different benefit is to understand the limits of technologies. For example, we want to know how resistant to attack a cryptosystem is. The strength of a cipher can be measured in many ways. Key length, resistance to carious cryptanalytic techniques, and the presence or absence of specific properties all give insight into the strength of a cryptosystem. But none

Continued on pg. 30 ➢

9. R. Bace and P. Mell, "Special Publication on Intrusion Detection Systems," Tech. Report SP 800-31, National Institute of Standards and Technology, Gaithersburg, Md., Nov. 2001.

10. G. Vigna, R.A. Kemmerer, and P. Blix, "Designing a Web of Highly Configurable Intrusion Detection Sensors," *Proc. Fourth Int'l Symp. Recent Advances in Intrusion Detection (RAID 2001)*, Lecture Notes in Computer Science, vol. 2212, Springer Verlag, New York, 2001, pp. 69–84.

11. D. Curry and H. Debar, "Intrusion Detection Message Exchange Format: Extensible Markup Language (XML) Document Type Definition," http://www.ietf.org/internet-drafts/ draft-ietf-idwg-idmef-xml-06.txt, Dec. 2001.

Richard A. Kemmerer is a professor in and past chair of the Department of Computer Science at the University of California, Santa Barbara. His research interests include computer security, formal specification and verification, and software engineering. Contact him at kemm@cs.ucsb.edu.

Giovanni Vigna is an assistant professor in the Department of Computer Science at the University of California, Santa Barbara. His research interests include network and computer security, intrusion detection systems, security of mobile code systems, penetration testing, and distributed systems. Contact him at vigna@cs.ucsb.edu.

Research vs. Practical Security, *continued from pg. 32*

of these prove the resistance of the cryptosystem to attack. Scholarly research into provably secure cryptosystems aims to develop techniques, and cryptosystems, that provide such assurance. The emphasis here is on proof, not opinion.

Scholarly research explores new technologies as well. The goal is to go beyond existing technologies to find new ways and new mechanisms to improve the state of computer security. The Eros operating system is a good example.[4] Eros is a capability-based system exploring an unusual approach to building systems designed with security in mind. While capability-based systems are not new, current operating system technology generally focuses on access control list technology. So Eros is exploring an approach that could prove fruitful. The approach could fail. But, whether it succeeds or fails, our knowledge and understanding of protection and the technologies that support it will increase. This is the mark of a scholarly research project.

The drawback of this type of research lies in its uncertainty and its long range. Scholarly research explores new avenues of ideas and principles, and any particular research project might not produce useful new results. In some cases, the usefulness of the results might not be apparent for years, decades, or even centuries. But in other cases, the usefulness could be obvious; intrusion detection systems, first proposed in 1986,[5] were being marketed by 1989 and are now a very popular technology). All this work uncovers material that can be incorporated into scholarly education to foster a deeper understanding of computer security.

A common perception is that training is superior to (or inferior to) scholarly education. The truth of this perception depends on the education's purpose. If the goal is to train someone how to use specific systems, work in specific environments, or perform specific tasks, training will achieve this goal more quickly than scholarly education. If the goal is to train someone to understand the general principles and ideas underlying a subject, or a technology, scholarly education will achieve this goal more quickly than training. But each form of learning enhances the other. A student of scholarly education sees how current technology applies the principles when she attends a training class. A student of training who takes a scholarly class will learn about the principles that guide the methodologies and technologies she learned. The ideal student will have both training and scholarly education.

Similarly, a common perception is that research used to support training classes is more beneficial (or less beneficial) than scholarly research. Again, the goal of the research determines its usefulness. If the goal is to provide short-term results leading to improving existing technology, the research supported by (or influenced by) training organizations is appropriate. If the goal is to foster a deeper understanding of the problems, to obtain more effective and long-term solutions, or to explore new approaches or technologies, scholarly research is the more appropriate venue.

Scholarly education and training complement each other. Frequently, data from the research of training organizations provides information that scholarly research projects can use to test their ideas or techniques, or to suggest alternate paths of research. Similarly, training organizations can build research, and education, around the results of scholarly research projects to better understand new technologies and how those can be used to improve the state of security.

References

1. J. Saltzer and M. Schroeder, "The Protection of Information in Computer Systems," *Proc. IEEE*, vol. 63, no. 9, 1975, pp. 1278–1308.

2. Bastille Unix Web page, http://www.bastille-linux.org/ (current Mar. 2002).

3. M. Harrison, W. Ruzzo, and J. Ullman, "Protection in Operating Systems," *Comm. ACM*, vol. 19, no. 8, Aug. 1976, pp. 461–471.

4. EROS: The Extremely Reliable Operating System Web page, http://www.eros-os.org (current Mar. 2002).

5. D. Denning, "An Intrusion Detection System," *Proc. Symp. Security and Privacy*, IEEE Computer Soc. Press, Los Alamitos, Calif., 1986, pp. 118–131.

Matt Bishop is an associate professor of computer science at the Department of Computer Science, the University of California at Davis. He is a charter member of the National Colloquium on Information Systems Security Education. Contact him at bishop@cs.ucdavis.edu.