Academia and Education in Information Security: Four Years Later

Matt Bishop
Department of Computer Science
University of California at Davis
One Shields Ave.
Davis, CA 95616-8562
phone: +1 530 752 8060
email: bishop@cs.ucdavis.edu

Introduction

The last four years have seen an explosion in the concern for the security of information. People are becoming aware of how much information is publicly available, as stories in the national news media discuss the ease with which identities are stolen. On a less personal note, compromises of information involving people authorized to access that information show that both companies and governments have problems in securing information. With this awareness has grown an understanding of our dependence on accurate, confidential information, as well as the fragility of the infrastructure we use to secure that information. Of all the questions emerging, the fundamental one is: how can we secure information?

The public perception of the situation is bleak. Woody Allen's statement that he had nothing positive to contribute, so he'd contribute two negatives, summarizes the perception perfectly.

I wish to discuss the role of academia in securing information. First, I will discuss the different forms of education relevant to the problem. Then I will evaluate the importance of each, and suggest ways to strengthen all of them. I will conclude by harkening back to the academic keynote address to the first National Colloquium on Information Systems Security Education, and compare where we are now with where we were then. This leads to several suggestions, which will augment my evaluation.

What "Information Security Education" Is

The terms "education" and "training and education" appear in several places in the National Plan for Information Systems Protection [7], but nowhere are they defined precisely. In fact, the report uses them to cover the gamut of activities from public outreach ([7] p. 73) to support for graduate and undergraduate education ([7] p. 24). The problem is that the report makes cursory mention of the interdependence of the various forms of education, and focuses primarily on training, awareness, and providing support for government and industry, and for academic programs that support the efforts of the government's plan directly.

I believe this approach, which a good first step, is a short-term measure. To understand why, I need to explain the different types of education.

Public Awareness

The most basic form of education is public awareness. Does the public understand there is a problem? How can we communicate the depth of the problem effectively? The public does not want to know details or technologies; they simply want to know how to keep their private information private, and this includes from government entities as well as commercial and academic ones. So education at this level is primarily procedural, and should focus on making the public aware of the threats and of what individuals can do to protect themselves.

As an example, many people are hooking up to the Internet using DSL, the Digital Subscriber Line technology. The marketing literature touts the benefits of speed and lack of busy signals when connecting to the Internet. The obvious conclusion, one that the members of the public do *not* make, is that DSL connects you to the Internet at all times except when your modem (or system). This broadens the interval in which attackers can probe your machine, and increases your exposure to attack.These risks can be ameliorated somewhat by the simple precaution of turning your modem off when not connected to the Internet. The public needs to learn the latter; the reason is not important (unless someone asks, in which case it should *always* be explained, by the principle of open design [11]).

Academic Education

Before we discuss the role of academia in education, we need to define "academia." What, exactly, is "academia"? The dictionary [12] defines it as a Latin word meaning "academy." An "academy" is:

- An institution of higher learning; or
- A private secondary or high school; or
- A school for teaching a particular art or particular sciences

In the spirit of being general, let's look up "academic" also:

- Of schools or colleges and their learning; scholastic, scholarly; or
- Too far from immediate reality; not practical enough; too speculative (Oh, well...)

These definitions identify several types of education as being "academic." The types, broadly stated, are training, undergraduate education, terminal master's education, and doctoral education. The differences between these types of education are illuminating.

Training focuses on particular systems, situations, or both. How do you configure a Windows 2000 computer to be a WWW server in a DMZ separating the protected internal network from the Internet? What happens if you don't use those specific settings, and what does each setting do? Whom do you call if you are a security guard who spots the director of the CIA carrying classified material out of the building? How do you send medical records to a doctor without compromising either the confidentiality or the integrity of those records? The answers to these questions are embodied in procedures and technologies. One need not understand why these procedures are in place, or how the technologies work, in order to use them effectively. (Obviously, the more understanding the trainee has, the better he or she will handle the job. But the understanding is not needed to perform the required tasks; knowing *how* is.)

Trade organizations such as SANS, USENIX, SHARE, and professional and commercial groups provide this type of training in tutorials. These tutorials are typically intensive, may be hands-on,

and have as their goal that the attendees can walk out of the tutorial and apply what they learned. Having taught many (possibly too many) of these myself, I should also observe that the residual value of these tutorials is having the book of slides and other materials. Often too much is covered to be retained perfectly. But attendees see something, and may remember they have seen it without recalling the details. They can then review the tutorial material and refresh their memories.

One can also acquire this training on the job, provided one is willing to ask questions and is paired with a mentor who is willing to answer them, and show the trainee how the systems are configured and how to reconfigure them. Although usually slow, this technique is effective because it teaches the trainee the problems that the particular site, or system, has, and how to cope with them, rather than the more general knowledge acquired in a tutorial attended by 150 or so people. When combined with a more general tutorial, this training is particularly effective.

The goal of undergraduate education is to learn broad principles, and see how to apply them. Undergraduate education does not focus on any particular situation or system. In practise, the best instructors take case studies and generalize them to exhibit the underlying principles. This helps students acquire a sense of what is principle and what is detail, and how to differentiate them. Subsequent exercises emphasize these principles, and have the students apply the principles in different ways. Throughout, the emphasis is on *what* principles are important, and *how* to apply them.

The advantage of a good undergraduate education is the breadth of application of principles taught. For example, in computer science, classes in algorithms, databases, operating systems, programming languages, architecture, and information systems teach various principles of information security, and how to apply them in the given realm. Political science and history classes teach principles of information security in studies of government and political movements, such as discussed in Sun Tzu's *The Art of War* and Saul Alinsky's *Rules for Radicals*. Literature classes sometimes discuss those principles as they study stories such as "The Purloined Letter" and *Oliver Twist*. The idea that knowledge may come from disciplines other than those naturally allied with information security testifies to the importance of those ideas.

As an example, consider one of Alinsky's rules about tactics for organizers:

> The third rule is; *Whenever possible go outside of the experience of the enemy.* Here you want to cause confusion, fear, and retreat.
>
> General William T. Sherman, whose name still causes a frenzied reaction throughout the South, provided a classic example of going outside the enemy's experience. Until Sherman, military tactics and strategies were based on standard patterns. All armies had fronts, rears, flanks, lines of communication, and lines of supply. Military campaigns were aimed at such standard objectives as rolling up the flanks of the enemy army or cutting the lines of supply or lines of communication, or moving around to attack from the rear. When Sherman cut loose on his famous March to the Sea, he had no front or rear lines of supplies or any other lines. He was on the loose and living on the land. The South, confronted with this new form of military invasion, reacted with confusion, panic, terror, and collapse. Sherman swept on to inevitable victory. It was the same tactic that, years later in the early days of World War II, the Nazi Panzer tank divisions emulated in their far-flung sweeps into enemy territory, as did our own General Patton with the American Third Armored Division. ([1] pp. 127–128)

The relevance to information security is obvious. From the attacker's point of view, look for unexpected openings. Look at the models the defenders have used to secure their information. Find ways to sidestep the mechanisms that the model requires, or—better—invalidate the assumptions that the model makes. Dorothy Denning very eloquently made this point in her National Computer Systems Security Award acceptance speech [8], in which she describes several incidents in which supposedly secure mechanisms were breached by people who went outside the conventional modes of analysis and found forms of attack that the models did not consider. Denning's talk, incidentally, shows the lesson of the Alinsky passage for defenders: expect to be attacked in ways you cannot anticipate, and be prepared for it.

Masters' level education builds on undergraduate education. It requires the student to examine a particular area of the discipline in depth either through additional course work and examinations, or through course work and projects culminating in a master's thesis. Such a thesis typically develops an application of a principle to a specific situation or set of situations, Masters' level education enables one to weigh competing interests and determine how best to apply different technologies to reach the desired balance.

Some examples of typical masters' level work are analyzing a particular network security protocol to determine if it has flaws, and suggesting changes to ameliorate the flaws; designing and implementing a library of specifications for security properties that are to be used with a testing tool; and developing a policy model for an academic institution. The first applies analytic and experimental techniques to a protocol to determine if it works correctly in the Internet. The second uncovers common flaws in programs and show how to abstract from them a description sufficient to identify previously unknown instances of the problems. The third combines technology with an analysis of the needs of the differing organizations making up an academic community, and presents mechanisms to enable the disparate groups to work together.

People with this kind of experience know how to weigh the conflicting needs of policy requirements, capabilities of technology, and human factors. They can analyze problems, look for solutions, and bring the two together. Sometimes no solutions are possible; this leads to approximations, and a good analyst can determine what the potential problems with the approximation are. In any case, these people can bring their experience in technology, principles, and analysis together to formulate guidelines that describe the needed protection. They then can design mechanisms to provide that protection.

Graduate education at the doctoral level also builds on the undergraduate education. Unlike a masters' education, though, doctoral level work analyzes the principles, extends the principles, changes them, improves them, or derives new principles. The goal is to deepen the student's understanding of systems in such a way as to enable that student to add to the body of knowledge. From this, we glean fundamental views of how to improve the state of the art and science of information security, and indeed what is, and is not, possible.

The difference between doctoral level work and masters level work lies in the nature of the concepts studied. Masters' level work typically emphasizes applications or applied research in some form. Doctoral level work emphasizes fundamental results and research, often called "basic research." Doctoral work pushes the boundaries of knowledge. The results may not be applicable immediately, or even in the short or medium term. But they help us better understand the technology, its limits, and its uses, and for that reason is critical.

Doctoral study also provides the credential needed to be hired by a research university: testament to the ability of the student to perform original, significant research. At research universities, teaching is not only a classroom exercise. Professors work with students in their research. Students learn how to conduct research, how to ask meaningful questions, and how to design experiments to demonstrate problems and solutions. In addition, students acquire an understanding of how to abstract problems into mathematical realms where they can be analyzed formally. With any luck, they also learn how to relate the formalism back to the problem to use whatever light their abstract analysis sheds on the problem.

The notion of "academic education" covers all of these forms: training, undergraduate education, masters' education, and doctoral education. It is imperative to understand that there is no hierarchy of importance or merit; someone with a doctorate is not better educated for a particular problem than someone with training to handle that problem. But someone with a doctorate can analyze that problem, abstract the problem, work with the abstraction, and suggest potential lines of research to eliminate the problem, and ones similar to it. People with Ph.D.s tend to generalize and try to solve classes of problems; people with training tend to focus on the particular problem at hand.Which is better? If you're a university trying to add faculty, the people with doctorates. If you're a Chief Executive Officer with a major computer incident on your hands, the people trained to handle the situation. (Of course, you might also want to find out why the incident occurred and how to prevent such incidents in the future. This would typically call for people with experience beyond training.)

Contrasts

Let's spend a few minutes contrasting the uses of these different educations.

Academics emphasize the principles underlying computer security. These range from the theoretical (such as the HRU result [9]) to the applied (such as Saltzer's and Schroeder's Design Principles for security mechanisms [11]). The goal is to be able to apply those principles to situations; in other words, to practice the science, and art, of computer security.

Good instructors use exercises to drive the ubiquity of these principles into the students. This type of teaching requires equipment and software that reflects the principles being taught, or to which the students can apply the principles and achieve an improvement, or visible alteration, to the system being modified. The students then see that they understand the principles well enough to apply them.

Industry needs to protect its investments in people, equipment, and its intangibles – bank balances, availability of services, proprietary information, etc. The security mechanisms must do this effectively. The principles they embody are less important.

In this realm, computer security is applied and practical. The goal of this type of computer security education is to be able to analyze a site, balance (internal and external) threats to the company with costs of implementing security measures, and achieving a balance between the two, with a minimum cost in training to the company. Understanding principles helps develop and implement policies and mechanisms, but the results are what matter.

Government uses computer security as one of many tools to protect the national interest (we assume this is well defined). The threats arise from external attackers and from government employees who act against the best interests of the citizenry or who abuse their authority. The spe-

cific protections are legally mandated, and not subject to the same cost-benefit analysis industry can afford. Hence computer security education focuses on developing policies and systems to implement laws and regulations, and less on cost balancing.

This points out the need for education at many levels. Each level has something to contribute. Most importantly, people at each level help educate each other. For this reason, all levels must be supported, and must play a part in protecting information.

The Role of Basic Research

That broad principles underlie information security is apparent. This leads to two questions:

1. What are these basic principles?
2. Given a particular system or situation, how do we apply them?

We have discussed the role of education in the latter. But the research involved in determining the principles, and understanding them *per se*, is a crucial component of the educational infrastructure needed to support education in applying the basic principles. The latter is all too often overlooked.

Fundamental research is the abstraction of ideas that underlie tools, protocols, methodologies, and procedures. For example, many institutions have several sub-units with differing security policies. The applied research question ism how does one derive an institution-wide security policy that meets the requirements of both sub-units? The fundamental research question this suggests is, under what conditions can one create a security policy that meets all requirements of two or more security policies? The utility of the second question is that it is very general and applies to many more situations than the first. The problem is that it is far more difficult to answer.

As a simple example of the utility of fundamental research, consider the question of testing a system for security. Is there a way to test a given system to determine if the information on it is secure? Yes; auditors do it all the time. Now, say we have an algorithm, or technique, to do this for one system. Can we apply that technique to any system? It seems logical that, with some modification, we could. So it seems as though we should try to find that technique; this would solve our testing problems.

Unfortunately, it has been shown that the general question of information security is undecidable; that is, no such single algorithm, or family of algorithms, exits [9]. So time spent looking for such an algorithm is fruitless. This is one benefit of fundamental research; it tells us what we cannot do. As another, more research has suggested characterizations of systems that make the security question decidable. For obvious reasons, systems with this characterization are simpler to understand and easier to demonstrate security for. (In fact, almost all systems do meet the characterization.)

The second benefit of fundamental research is to enable us to understand the technologies used in information security. As an example, how good is a cryptosystem? A number of measures such as key length, resistance to various cryptanalytic techniques, and the presence or absence of specific properties give insight into the strength of a cryptosystem. But none of these *prove* the system is resistant to cryptanalysis. Fundamental research in provably secure cryptosystems aims to develop techniques, and cryptosystems, that provide such assurance. The emphasis here is on *proof*, not opinion.

A second example of this is Denning's lattice model representation of the Bell-LaPadula model. Denning's analysis provided a framework to make the model very simple to understand, and is considered seminal, foundational work. Others used her techniques to describe security models in terms of lattices for information flow.

The third benefit of fundamental research is its breadth of applicability. The best example of this is the ubiquity of Saltzer's and Schroeder's design principles for security mechanisms [11]. The eight principles underlie all security procedures and mechanisms:

- The *principle of least privilege* says that a process should have only those right necessary to complete the task. In government and industry, this is the "need to know" principle.
- The *principle of fail-safe defaults* says that when a security mechanism or system fails, the system should revert to a known, secure state. This essentially says to deny access to sensitive information unless access is explicitly granted, again a standard when dealing with sensitive information.
- The *principle of economy of mechanism* says that security mechanisms and procedures should be as simple as possible, because as system and mechanism become more complex, more can go wrong. Further, the more complicated a mechanism, the harder it is to convince people that the mechanism works as needed. This is a general rule, born from human nature and experience. Arthur Clarke's marvelous short story "Superiority" casts this in terms of science fiction, in which the desire to develop complex, powerful weapons leads defeat at the hands of simpler, less powerful, but functional weapons.
- The *principle of complete mediation* means the mechanism cannot be evaded. Denning made the importance of this principle explicit in her talk before the National Information Systems Security Conference. She pointed out that attackers often evade controls designed to stop them. The controls are never invoked, so they are completely ineffective.
- The *principle of separation of privilege* says that multiple properties must hold for access to be granted. In financial circles, this is called "separation of duty." Two people must sign checks over $10,000. Two soldier must insert keys to launch missiles. One person is easier to compromise than two who must work in concert. Again, mathematically this is a fallacy, but humans are not mathematical.
- The *principle of open design* says that the security of a system should not be based upon hiding the details of how the system functions. Hiding specific information such as passwords does not violate this principle, but hiding the general design of a security policy or system does. Attackers can construct the details of systems in a variety of ways. For security procedures, dumpster diving is effective. Woodward and Bernstein determined the lines of reporting in the highly secretive Committee to Re-Elect the President by examining telephone numbers and seeing who had phone numbers "close" to whom.
- Finally, the *principle of psychological acceptability* says that security procedures and mechanisms must be as easy to use as ignore. This principle is usually watered down to say that using the security mechanisms must not be too onerous. Passwords and badges are generally acceptable. In high-security institutions, fingerprints provide a high degree of authentication. But requiring fingerprints for authentication to enter a university laboratory would be unacceptable, at least at our university. The students, staff, and faculty would simply refuse to tolerate it.

These principles permeate all of information security. They were developed as a set of coherent, cohesive concepts underlying the development of security mechanisms. The principles provide a measuring stick (granted, of the most basic sort) that speaks to the effectiveness of security mechanisms for all of information security. This is a mark of fundamental research.

Fundamental research is the glue that binds all these disparate parts together and makes information security a discipline, not an amalgamation of *ad hoc* techniques. Without fundamental research, everything will be specific, and work done for one system or situation will need to be redone and re-validated for other systems and situations. This is the key: by improving our ability to carry out fundamental research, we save ourselves the cost of developing and redeveloping solutions. We learn our theoretical limits, and how to work right up to those limits.

How Are We Doing?

We can now answer the question, "What is the state of INFOSEC education in academia and how does that compare to the state four years ago?" Be warned, it isn't pretty.

First, the good news. There is interest, and discussion, on improving the state of information system security education. The desired improvements include establishing core curricula and integrating computer security into more aspects of computer science education. Specifically, the Centers of Academic Excellence in Information Assurance Education has as one evaluation criterion that "[t]he academic program demonstrates [information security] is not treated as a separate discipline, but as a multidisciplinary science with the body of [information assurance] knowledge incorporated into various disciplines" [6]). This program recognizes institutions that are teaching students about information security, even when the student's primary interest is not information security. The recognition of the seven original Centers of Excellence, and the others that will be recognized later today, is a first step.

It is, however, only a first step. The designation involved no support and no benefits other than being able to say that the institution was a Center of Excellence designated by the National Security Agency. To be fair, the NSA has always said that this is the only reward, but they hoped that the "Centers for Academic Excellence may become focal points for recruiting and may create a climate to encourage independent research in Information Assurance." Perhaps that will happen soon.

The contents of the *National Plan for Information Systems Protection*, with its statement of support for education and research, is heartening. But, I personally view it with joy tinged with cynicism. From the document it is clear that the "education" being discussed is primarily training, although support for undergraduate work is provided through Scholarship for Service ([7], p. 67-68), and recognition of programs through the INFOSECURITY Centers of Excellence is discussed (this seems to be an extension of the NSA's program;[1] [7], p. 24). Where is the support we need for basic research? This money is essentially money for band-aids and patching broken systems and infrastructure. While that is critical, equally critical is that we move to developing more robust systems, and learn how to design protocols, infrastructure, and systems that are more

1. The report says 8 universities have been designated as Centers of Excellence in Information Assurance, but the NSA press announcements and web site name only 7 (James Madison University, George Mason University, Idaho State University, Iowa State University, Purdue University, University of California at Davis, and University of Idaho).

secure (for some definition of "secure") than our current ones. The National Plan does not discuss any funding for basic research in these areas. If this is indeed a lack of support for basic research, the lack is disastrous.

Our second metric was the development of models that accurately reflect the systems to which they are applied. This issue arises because models that do not reflect reality cannot be the basis for effective computer security controls. A perfect example of this is a recent model to use an analysis technique called conservation of flow to detect malicious network infrastructure components [5]. Conservation of flow requires that every packet entering a network be accounted for. In theory, the model is superb, and can detect rouge infrastructure systems. The problem, of course, is that in reality, conservation of flow does not apply to the Internet. The network loses packets for a variety of reasons. When you add factors to compensate for this loss, the precision of the suggested protocol evaporates, leaving only an approximation technique that an attacker can defeat [10].

We are making progress in this area. A number of models describing specific, real-world situations, are being applied to protect information security. Ross Anderson's medical information protection model is an excellent example in this regard [2]. The model was developed with the advice of clinicians and the medical community, was implemented and tested, and some hospitals in the United Kingdom are using it. But once the model was put into practise, the researchers found that access control was the easiest of the three main problems. The other two are inference control--how to prevent the association of sanitized records with individuals--and medical information that is passed to groups outside "professional control," such as insurers [3]. My point is that the model reflects much of the practice of the medical profession's requirements on medical data security, at least in the United Kingdom, and the researchers understand the problems with applying the model. All too often, that perceptiveness is lost.

So that's the good news. But we are still failing in what I view as the critical areas of information security.

The ILOVEYOU virus[1] is a perfect example. In 1988, before the Internet virus appeared, the CHRISTMA EXEC worm threaded its way through several IBM networks. People received a letter telling them to save the body of the letter as a file, and then execute the file, to get a pleasant Christmas greeting. When they did this, they saw a Christmas tree with blinking lights drawn on their screens. What they did not see was the rest of the program. It then looked in the NAMES and NETLOG files to get names of other correspondents, to whom it would forward itself. The resulting E-mail storm made several IBM networks unusable until the worm was cleaned out. Does this sound familiar?

The ILOVEYOU worm used almost exactly the same techniques. The only differences were that the recipient had to click on a button, rather than save the file and execute it, and the ILOVEYOU worm downloaded a second program which harvested passwords from the Windows system's cache. These are reasonable updates given the changes in our world over the previous 11 years.

In other communities, software still suffers from buffer overflows. Privileges are not constrained properly. Race conditions allow unscrupulous users to acquire control of systems. There is nothing new under the sun. What has happened to us will happen again, and we're not learning from these mistakes.

---

1. Technically, it is a worm, not a virus, but I'll use the common term here.

Nor have we improved how we design systems and programs to account for security problems. I'll use Windows 2000 as an example, because it is a system that is, or will soon be, very widely used, and has been introduced after 1997. Microsoft's security mechanisms, in concept, are excellent. But their implementation and integration into the system seem to lack coherency and cohesiveness. Further, some subsystems have design problems and implementation problems. Microsoft has released several patches for both systems and application software, and still has numerous security-related issues pending. One gets the impression that, while security was somewhat of a consideration in Windows 2000, the ideas of backwards compatibility, "gee-whiz" features, and completeness were more important. To be blunt, I understand there are roughly 33 *million* lines of code in the Windows kernel (down from 47.5 million lines of code). I simply do not believe that a kernel of that size, and intended for general purpose use, can be made secure.

Again, to be fair, other companies have the same problem; their products are just not so popular as Microsoft's, so they don't get the publicity. A good example of this is your favorite UNIX vendor. Many versions of UNIX systems have equally grave security problems. I'm not as conversant with other systems, but I'd speculate they do too. The problem is that we do not design with security as an integral part of the design. We patch. We add security above the kernel, or retrofit it. This causes problems.

The third criterion is to understand how humans interact with systems, how security problems arise from this interaction, and using this knowledge to build systems that minimize the possibility and effects of errors. The politest statement is, "forget it." We don't have that knowledge, and we don't seem interested in acquiring it. We are still hoping that technology can solve our security problems, and have not yet realized that the technology is intimately bound up with the human beings who use it. I'll say more about this when we discuss the different types of computer security education.

A quick review of the goals and our progress towards them is:

- We are making (minimal) progress in integrating information security into other parts of our curricula.
- We have not learned from our mistakes, and continue repeat the errors from the past.
- We have not improved how we design systems and programs to account for security constraints, nor have we reduced the number of security patches necessary.
- We are learning how to abstract the requisite characteristics of a system towards this end, but we have much to learn.
- We do not understand how humans interact with systems, how security problems arise from this interaction, and therefore cannot use this knowledge to build systems that minimize the possibility and effects of errors.

The 1997 keynote made several suggestions for improving the state of information security education that bear repeating.

- Academia needs long-term funding to provide a stable base for our research. Short-term resources are not enough, because the burden of trying to find new funding to keep our work going, and to build a long-term research program, is a drain on our resources. Most importantly, a stable funding base would give industry, government, and the nation a set of resources upon which they could draw without having to start from scratch. The importance of this cannot be underestimated. This base of research and knowledge can provide help and

research results to deal with the crisis we face in information security, and to solve the problems causing the crisis.

- We need more industry and government participation in selecting research topics. Nothing is more frustrating than solving a problem, only to find it is not really a problem, or the "real world" version of the problem has additional constraints that change the approach drastically. Ways to do this are through partnerships with industry in which we discuss problems and possible approaches, and work together to solve them; through internships, where members of industry come to academic institutions for a period of time to teach and work on projects with students, and where faculty and students go to industry for periods of time to work on problems of interest to the industry. One of the most common complaints of students is the lack of "real world" experience, and of industry and government is that the students lack "real world" experience. These measures would provide them. The National Plan calls for many of these steps, such as internships, co-ops, and partnerships.
- Industry and government should fund "blue sky" research and long term, directed research. Blue sky research is speculative; it may succeed, it may fail, but the body of knowledge that comes out of it will advance the field in some manner. Remember, failure can be just as strong a result as success. Long-term research would allow us to turn our academic resources to problems that we could study thoroughly and attempt to solve in a number of different ways. Both these suggestions would produce an immeasurable amount of research and scholarship, upon which short-term projects could be built.

A Parable

The use of real-world examples in class makes the learning come alive for students and imprints the lessons in their mind much more strongly than a theoretical discussion. This openness, to both failure and success, is critical to the development of students' analytical abilities. Industry is often unable to do this, because of proprietary matters. It would be unreasonable to ask, for example, Hewlett-Packard to publish the source code to their system to be used as a pedagogic tool. The federal government is often constrained in a similar fashion, but for national security reasons.

But the government has a tendency to be secretive when not needed, and were this overcome it would be a boon to education. I had a very unpleasant example of this, which shows how the ambivalence can hinder education. Early this year a series of distributed denial of service attacks were launched. The NIPC subsequently posted on its web page two programs to analyze your system and determine whether you had tools to launch the attack. When we got the tools, we noticed they were executable files, not source code. This immediately raised several questions:

1. Had the web site been compromised and fake tools substituted? Attackers have compromised several US Government web sites, including the Department of Justice web site, and the NIPC web site appears to be part of the FBI's site, so we thought this was possible.

2. Did the tools do something beyond what the documentation says? While we doubted the authors would do this deliberately (because the attendant publicity would destroy the FBI's credibility in the computer security community), we were concerned about coding errors and the thoroughness of the scan.

One of my graduate students and I decided to treat the software as we would any other unknown piece of code. We downloaded the Linux version onto a Sun system (so it could not be executed

even by accident) and scanned it looking for suspicious data. We found several strings that appeared to be help messages for attack tools, such as "4 to bind a rootshell (specify port)", "<iplist> contains a list of numerical hosts that are ready to flood", and the immortal "[t]his is the 'trinoo" AKA DoS Project master server.". We immediately moved the program to an isolated network and began a full-scale analysis. Needless to say, we didn't run it anywhere—and I know of commercial groups that saw similar strings and declined to run the code as well.

Here's the sad part. After intensive analysis, my graduate student concluded that this tool was a signature scanner, and the phrases were parts of the attack tools it would look for. If this is true (and although I still have not seen the source, I have reason to believe it to be true), I suspect the reason the source was not distributed is because it contains strings from the attack tools that someone feared would make those tools easier to implement. But a number of techniques, such as hashing the code being looked for and comparing that to hashes generated from files, could have performed the task equally well, been released freely without revealing any part of the attack tool, and demonstrated the way the scanning worked. I would have loved to use this source code in my class. That's an example of how government can contribute to classroom studies. Money is good and necessary, of course; but from a pedagogic viewpoint, actually showing how one checks for the rogue code would have made the concepts come alive for the students. The discussion we had on how the tool "probably" worked was nowhere near as satisfying to them.

Conclusion

I have made many specific criticisms in this talk. It's easy to be a critic; as Sam Rayburn once said, "Any jackass can kick a barn down." So let me close with a suggestion.

All forms of education, from basic research to training, are critical to responding effectively to the information security crisis we face now. In addition to focusing our efforts on training, we should focus our efforts on basic research and higher education. The latter two will provide the teachers and researchers we need to train system administrators, business executives, and management in the intricacies of information security that affect them and their organizations. Further, the emphasis on basic research will lead to more research faculty in the area of information security, thereby seeding more universities and academic institutions with people who can teach and do research in that area.

My greatest fear is that we will forget the dreamers, the people with long-range vision. To date, we are focusing on short-range or medium-range planning. We must do that, but we cannot forsake the long term. Our technologies will change; our systems will become obsolete; our infrastructure will evolve in ways we cannot anticipate. The dreamers will provide the vision. Ted Nelson conceived of hypertext in the mid-1970s as he studied how computers and books would work together; can anyone imaging the World Wide Web without clickable links, or—more properly—hypertext? Nelson was a dreamer, but he had a technologically sound vision. We must treasure people like that, because they guide the way. If we focus on the immediate, and near-term, present, we run the risk of becoming General Carpenter in Alfred Bester's story "Disappearing Act."

In that story, America is involved in a war, and has become a nation of experts. "Every man and woman must be a specific tool for a specific job, hardened and sharpened by your training and education to win the fight for the American Dream." But wounds have caused some injured sol-

diers in a hospital to vanish and reappear at will. Investigation convinces the general that the casualties are going back into time, and he asks a historian (who is released from his prison sentence for questioning the war) to see how they do it. The historian quickly realizes that the casualties are travelling elsewhere, "back into a time of their own imagination." He continues:

> "The concept is almost beyond understanding. These people have discovered how to turn dreams into reality. They know how to enter their dream realities. They can stay there, live there, perhaps forever. My God, Carpenter, *this* is your American dream. It's miracle working, immortality, Godlike creation, mind over matter ... It must be explored. It must be studied. it must be given to the world."
>
> "Can you do it, Scrim?"
>
> "No, I cannot. I'm an historian. I'm non-creative, so it's beyond me. You need a poet ..."
>
> ...
>
> Carpenter snapped up his intercom. "Send me a poet," he said.
>
> He waited and waited ... and waited ... while America sorted through its two hundred and ninety millions of hardened and sharpened experts, its specialized tools to defend the American Dream of Beauty and Poetry and the Better Things in Life. He waited for them to find a poet, not understanding the endless delay, the fruitless search; not understanding why Bradley Scrim laughed and laughed and laughed at this final, fatal disappearance. [4]

The worst catastrophe that could befall us is having a "cyberspace" of hardened, sharpened tools trained and educated for a specific job, and no-one who knows how to ask if there is another approach to the task, or how to look for it.

References

1. S. Alinsky, *Rules for Radical*s, Random House, Inc., New York, NY (1972).
2. R. Anderson, "A Security Policy Model for Clinical Information systems," *Proceedings of the 1996 IEEE Symposium on Security and Privacy* pp. 30-43 (May 1996).
3. R. Anderson, "Privacy Technology Lessons from Healthcare," *Proceedings of the 2000 IEEE Symposium on Security and Privacy* pp. 78-79 (May 2000).
4. A. Bester, "Disappearing Act" (1953); in *Virtual Unrealities: The Short Fiction of Alfred Bester*, Vintage Books (New York, NY (1997)
5. K. Bradley, S. Cheung, N. Puketza, B. Mukherjee, and R. Olsson, "Detecting Disruptive Routers: A Distributed Network Monitoring Approach," *Proceedings of the 1998 IEEE Symposium on Security and Privacy* pp. 115-124 (May 1998).
6. *Centers of Academic Excellence in Information Assurance Education (Graduate and Undergraduate Levels): Criteria for Measurement*, http://www.nsa.gov/isso/programs/coeiae/measure.html (Oct. 1999).
7. *Defending America's Cyberspace: National Plan for Information Systems Protection: An Invitation to a Dialogue*, Version 1.0, The White House (2000)

8. D. Denning, “The Limits of Formal Security Models,” National Computer Systems Security Award Acceptance Speech, http://www.cs.georgetown.edu/~denning/infosec/award.html (Oct. 1999)

9. M. Harrison, W. Ruzzo, and J. Ullman, “Protection in Operating Systems,” *Communications of the ACM* **19**(8) pp. 461-471 (Aug. 1976).

10. J. Hughes, T. Aura, and M. Bishop, “Using Conservation of Flow as a Security Mechanism in Network Protocols,” *Proceedings of the 2000 IEEE Symposium on Security and Privacy* pp. 132-141 (May 2000).

11. J Saltzer, and M. Schroeder, “The Protection of Information in Computer Systems,” *Proceedings of the IEEE*, **63**(9) pp. 1278-1308 (1975).

12. *Webster’s New Twentieth Century Dictionary,* Second Edition, Simon and Schuster, New York, NY (1979).