

# IT Security Research and Education in Synergy<sup>†</sup>

Stefan Lindskog<sup>‡</sup>, Ulf Lindqvist, and Erland Jonsson

*Department of Computer Engineering, Chalmers University of Technology,  
SE-412 96 Göteborg, Sweden  
{stefanl, ulfl, erland.jonsson}@ce.chalmers.se*

**Keywords:** IT Security Education, Computer Security, Intrusion Experiments, Intrusion Analysis, Remediation, Intrusion Detection.

**Abstract:** This paper presents experience from laboratory projects performed by students in Applied Computer Security, a course given at Chalmers University of Technology. From this, we conclude that the combination of security research and education results in a synergy that can be very beneficial to both fields. The paper describes in detail three different types of laboratory projects: intrusion experiments, intrusion analysis and remediation, and intrusion detection. Furthermore, we discuss the outcome and impact of the student projects with respect to education, research, and synergy between the two, as well as ethical issues. Among the benefits of the close connection between research and education in the projects, we find that the students were very motivated by this research connection and that they produced unique research data with natural diversity.

<sup>†</sup> In Proceedings of the 1st World Conference on Information Security Education (WISE1), Stockholm, Sweden, June 17-19, 1999.

<sup>‡</sup> The author is also with the Department of Computer Science, Karlstad University, SE-651 88 Karlstad, Sweden.

## 1. INTRODUCTION

Most networked computer systems are continuously exposed to intrusion attempts. Even worse, some of the attacks succeed because people do not know how to implement system security. Unfortunately, the number of skilled practitioners of computer security today is small [10], making it necessary to incorporate security studies into educational programs. Furthermore, students must be prepared for real situations. Thus, practical laboratory experience is essential [11].

Since 1993, a course in Applied Computer Security has been given at Chalmers University of Technology, offered in the final year of the Master's Degree Program. Topics covered in the course are: access control, authentication, cryptography basics, database security, network security, operating system security, physical security, privacy and ethical issues, program security, risk analysis, security management, security models, and security policies. An important component of the course is a laboratory project. During the first years of the course, the students were asked to evaluate the security of a target system by means of attacking it. The results of those experiments have been presented in many research papers, including [14].

The synergy between our research and education in security has been very fruitful for both fields. It is obvious that our research would not have been what it is today without the Applied Computer Security course. This elective has been much appreciated, and students have claimed that it has given them a "hands-on" feeling for computer security and contact with the research frontier.

In [12], Jonsson and Janczewski presented a survey of practical security experiments, conducted at a number of universities on most continents, including Europe, North America, Africa, Asia, and Australia/Oceania. The aim of the study was to investigate the role of experiments in information security education. They also suggested a taxonomy for such experiments. They found that many of the exercises and projects reported were different in terms of degree of applicability, degree of innovation and generalization. Moreover, the duration and effort (in man-hours) of performing an activity ranges from a duration of two hours with an effort of two hours to a duration of three months with an effort of 100 hours.

In the following, section 2 gives some background information. The feasibility study is further described in section 3, the first full-scale experiment is presented in section 4 and other intrusion experiments conducted at Chalmers are briefly discussed in section 5. Section 6 introduces a new approach called intrusion analysis and remediation, and section 7 focuses on the intrusion detection assignment conducted in 1998. Section 8 discusses ethical issues, synergy effects, and educational and research aspects. Section 9 concludes the paper.

## 2. BACKGROUND

In 1992, an EU research project was started with the intent of finding quantitative measures of operational security. For that purpose, empirical data was needed. To gather such data, practical intrusion experiments were conducted.

For the first experiment [2], which is referred to in this paper as the feasibility study, the primary goal was to investigate whether it would be possible to gather data for security modeling. This study was successful in many respects. First, we found that students are indeed able to break into a standard system within a limited time period and that they can be used in this type of experiment. Second, we learned how to set up such experiments and how to act as coordinators.

The first full-scale experiment [19] was done half a year after the feasibility study. In this case, our goal was to gather enough data for a quantitative security-modeling attempt. Final year students were engaged as attackers. All were at the same time taking Applied Computer Security, and the experiment was performed as a project in the course.

In both the feasibility study and the first full-scale experiment, the target system was a networked UNIX system. Three full-scale experiments have followed. On one occasion, a PC network was used as the target system [5, 6].

In the last couple of years, new ideas have been developed primarily to the benefit of education and secondarily to research. In 1997, for example, we introduced the remedy dimension in the project work. That time, the students were asked to suggest how a given vulnerability, which students found in the earlier experiments, could be eliminated from the system. In 1998, our increased research interest in intrusion detection was reflected in the course. The students would construct a small intrusion detection system using an expert system tool.

### **3. FEASIBILITY STUDY**

The feasibility study [2] (sometimes referred to as the pilot experiment) was intended as a pure research activity and conducted during the spring of 1993. The goal was to determine whether it would be possible for students to break into a normal system and how data from these attacks could be collected for quantitative security modeling.

#### **3.1 Experimental Set-up**

The target system was a networked UNIX system running SunOS 4.1.2, consisting of 22 disk-less workstations and a file server, all configured as recommended by the vendor. Standard accounting was enabled on all computers for monitoring purposes and, when the experiment was performed, the system was in regular use by other students. Finally, the security of the system was that of a default installation and was not improved in any way.

Three different types of actors were identified: the attackers, an experimental coordinator and the normal system administrator. Members of the Chalmers University Computer Club adopted the role of attackers, as expected these students have greater knowledge of security vulnerabilities than other students. Only 13 students participated in the first experiment, and each attacker was a legal user of the system, i.e. an insider. One researcher acted as experimental coordinator, monitoring and coordinating all activities in the experiment. Finally, the normal system administrator represented the system owner. He was instructed to behave as normally as possible.

The attackers were required to follow a few rules, which were presented in a briefing before the experiment began. The main reason for these was to encourage the attackers to behave as realistically as possible. For example, they were not allowed to cooperate with other attackers during the experiment or to cause physical damage to the hardware.

Before the experiment took place, each attacker was required to fill in a questionnaire that described his/her formal education, prior experience in the area etc. In addition, for each intrusion attempt, an activity report was written. Thus, working time, on-line time, resources used and the occasion at which the activity was performed were documented. At the completion of the experiment, every attacker filled in an evaluation report. The activity and evaluation reports were both standardized, fill-in forms. Finally, the attackers were asked to write a final report to give their views of the experiment, summarize their activities, describe successful and unsuccessful attacks and ideas for other attacks etc., and to describe the outcome of the experiment from a personal point of view.

#### **3.2 Results**

With respect to data collection methodology, the most important result was that students could indeed be used in experiments such as this. However, the amount of data received was, as expected, too sparse to be used for any statistical modeling.

This experiment also taught us a great deal about the intrusion process. Most notable is that, by measuring the *effort* needed to break into a system, we might in the future be able to quantitatively determine the security level of a particular type of computer system. Another interesting conclusion is that it is even more difficult to measure the *reward*, i.e. the motivating force, of a particular successful intrusion than it is the *effort*. Finally, data collected in this study have been used in several research papers, including [2].

Problems that occurred were that attackers tended to leave the experiment for a lack of ideas and/or motivation and that the required reporting turned out to be too extensive. Still, this experiment was highly appreciated by the participants.

## 4. FULL-SCALE INTRUSION EXPERIMENT

While the feasibility study was a pure research activity, the first full-scale intrusion experiment was a combination of research and education. The research goals were to collect enough data to make a quantitative modeling attempt possible and to learn more about the intrusion process. The educational goal, however, was to offer a realistic project in the Applied Computer Security course, which was given for the first time. The experiment was conducted in November and December, 1993.

### 4.1 Experimental Set-up

We used the same target system as in the feasibility study, with one security improvement. The workstations were set up to require a password to successfully perform a so-called single user boot-up sequence.

In this experiment, we engaged as attackers Master of Science students in Computer Science and Engineering in their final year. They worked in groups of two students, and a total of 12 groups attacked the system at the same time. We expected each group to spend approximately 40 hours of effective working time. The attackers were required to contact the coordinator once a week to discuss their progress. The same rules as described in section 3 were to be followed.

In the feasibility study, the participants had complained about the excessive reporting, and several simplifications were thus made here. At the end, 65 security breaches of various kinds were reported, to be compared with the 25 breaches in the feasibility study.

### 4.2 Educational Results

The project was well received by the students. We are convinced that the results were very surprising to many of the participants when they discovered that it was not as difficult as they thought to break into a standard UNIX system. Indeed, most students claimed that the experiment gave them a “hands-on” feeling for computer security and raised their awareness of the problems involved.

### 4.3 Research Outcome

This study taught us even more about the intrusion process. Publicly available tools for searching and utilizing known vulnerabilities in the specific system, i.e. SunOS 4.1.2, were used in several intrusion attempts. The groups downloaded these tools from the Internet. Further, more events were reported and the recording of data was more complete in this second experiment. The data collected has been used for research presented in several papers, including [13].

## 5. FURTHER EXPERIMENTS

Three additional full-scale intrusion experiments have been conducted. Thus, a total of five intrusion experiments in which students have adopted the role as attackers has been carried out at the Department of Computer Engineering. We summarize these in table 1.

Table 1. Summary of experiments

| Year        | Experiment                                      | Operating system    | # of students |
|-------------|---|---------------------|---------------|
| Spring 1993 | Feasibility study                               | SunOS 4.1.2         | 13            |
| Autumn 1993 | 1st full-scale intrusion experiment             | SunOS 4.1.2         | 24            |
| Autumn 1994 | 2nd full-scale intrusion experiment             | Novell Netware 3.12 | 14            |
| Autumn 1995 | 3rd full-scale intrusion experiment             | SunOS 4.1.3         | 32            |
| Autumn 1996 | 4th full-scale intrusion experiment             | SunOS 4.1.3         | 42            |
| Autumn 1997 | Intrusion analysis and remediation <sup>†</sup> | SunOS 4.1.3         | 59            |
| Autumn 1998 | Intrusion detection <sup>‡</sup>                | SunOS 4.1.3         | 87            |

<sup>†</sup> Intrusion analysis and remediation are described in section 6.

<sup>‡</sup> Intrusion detection is covered in section 7.

In 1994, the target system was a PC network system. Insecure clients were running DOS 6.2 and Windows 3.1, and the network file server was configured with Novell Netware 3.12. Despite the fact that only one student was familiar with Novell Netware, the target system was easily and successfully attacked and breached by all students except one. In most cases, they exploited the insecurity of the clients in the sense that any user had full access to everything at the client. The experiment is further described in [6], and [5] presents an interpretation of the intrusion data collected.

In the third and fourth full-scale experiments, a networked UNIX system, this time with SunOS 4.1.3, was again used as the target system. In table 1, we can see that the number of attackers is increasing, and our experience is that our approach to intrusion experiments scales well to the number of attackers in terms of the educational resources required. However, when the number of attackers increases, the compilation of recorded data for research purposes becomes much more difficult because most of the work needs to be done manually.

Two other experiments should also be mentioned. First, procedures inspired by the intrusion experiments were used in an analysis of a secure system based on trusted components. See [16, 17] for a detailed description of this experiment. Second, a similar procedure was used to evaluate the security of Windows NT, see [7, 8] for an in-depth description. These two experiments differ from those described earlier in one important aspect: we did not engage students as attackers. Instead, we adopted the role as attackers, presumably more of *tiger team* attackers [1, 4, 9], ourselves. In the latter study, we found that Windows NT can be penetrated quite easily. This implies that Windows NT can be used as a candidate target-system in future intrusion experiments in which students act as attackers. One such experiment will be conducted at Karlstad University in the spring of 2000.

## 6. INTRUSION ANALYSIS AND REMEDIATION

In the previous experiments, students were asked to find as many weaknesses as possible in the target system by attacking it. In 1997, after four full-scale experiments, we were interested in trying something new. The new approach was to let the students act as though they were on the other side, i.e. to act as the system owner. This experiment was primarily an educational activity in the sense that we did not foresee any specific research outcome.

### 6.1 Experimental Set-up

The target system was the same networked UNIX system as was used in the third and fourth full-scale intrusion experiments described in section 5. A list of 13 carefully selected intrusions

found in them was presented to the students. The class of 59 students was divided into groups of two students each, i.e. 30 project groups were constructed. This implies that several groups analyzed the same intrusion. Each project group selected one intrusion and was asked to:

1. Reproduce the selected intrusion, verifying that it can be used to circumvent the security of the target system
2. Describe the weakness and explain why the intrusion works
3. Suggest remedies that would make the system secure against this kind of intrusion

The only intrusion descriptions they were given were those written by the previous students in intrusion experiments. Furthermore, each intrusion was assigned an estimated level of difficulty. Roughly, three levels were used—simple, medium and complicated. The groups were required to document their work. In fact, two reports were compulsory: a complete project report and a summary report of approximately two pages. They had to describe the vulnerability according to the taxonomy suggested by Landwehr *et al.* [15]. Thus, they needed to answer the following questions:

- How did the vulnerability enter the system?
- When did the vulnerability enter the system?
- Where in the system does the vulnerability manifest?

When the exercise was finished, copies of the summary reports were distributed to all participants to give the students a broader view of intrusions than the specific case they had studied. In addition, each group presented its work to the class.

## 6.2 Results

The students were put in a situation like one many of them will probably experience in their future careers as either software engineers or system administrators. Most students reported that they found the exercise interesting and valuable. It is also of interest to note that some groups had great difficulty in reproducing the intrusion. In some cases, we had been overly optimistic as to the programming skills and computer literacy of the students and, in other cases, the reason was probably that it was impossible to perform the intrusion on our system. The exercise required far more supervision than we had expected.

## 7. INTRUSION DETECTION

In 1998, the steadily increasing number of students in our course had forced us to give all groups identical assignments, as in a traditional laboratory exercise. To retain the connection with our research, we decided to let the students use a software tool that had previously been used only by experts to construct research prototype intrusion-detection systems. The tool, called P-BEST, is further described in [18].

### 7.1 Experimental Set-up

The assignment was to construct a system that could be used to automatically detect attacks against a file transfer (FTP) server. The students were given a tool designed for building expert system components of intrusion detection systems. The tool is the same one we use in our research, and the exercise would support or contradict our hypothesis that the tool is easy to use for beginners. It was required of the students that they include in their lab reports a discussion of their experiences with using the tool.

For evaluation of the system they had designed, the students were given a very large data file containing recorded network data representing actual FTP transactions. A small number of real

and synthetic intrusions was mixed with a large number of normal transactions, and the students used their system to find those intrusions.

## 7.2 Results

A total of 87 students participated in the assignment, with a few exceptions working in pairs, making a total of 46 groups. Of those 46 groups, 25 had constructed a system that gave a completely correct answer. An additional eight groups would most likely have obtained the correct result if they had not misinterpreted a vaguely formulated part of the instructions. Only a handful of groups failed to hand in a report before the given deadline. Most students reported that they found the exercise interesting, and some even took the time to give detailed suggestions for improvements to the tool.

## 8. DISCUSSION

The outcome and impact of our student projects can be viewed from several different perspectives. This section discusses aspects of education, research and synergy between the two, as well as ethical issues.

### 8.1 Educational Aspects

Probably the most important result from an educational perspective is that the laboratory exercises, described in sections 4 through 7, give a practical view of security. The students have learned among other things how an intruder thinks and become aware of common threats to today's networked computer systems. We therefore believe that the assignments have prepared them for forthcoming real situations. In addition, comments by the students about the exercises have been very positive. Two key words often found in the evaluation reports are "fun" and "exciting". Several students claimed this to be the best course they have taken. Furthermore, students usually have very little contact with research during their studies, and this has made our course different and interesting.

### 8.2 Research Aspects

There are several concerns as regards the validity and accuracy of the data obtained from student exercises. The first question is whether the students in our experiments constitute a good approximation of real attackers. In reasoning about security, the discussion often quickly reaches the point where the goal is to protect the system against the most skilled and powerful attacker in the world. The term "übercracker" [3] has been used for this picture of a diabolic, omnipotent adversary. We did not wish to investigate the übercracker, partly because most people in the security community are already doing that. We wanted to see how ordinary computer users with academic training in computer science and engineering—but no previous attacker experience—would operate when they suddenly had a reason to attack the system on which they were working. We found that unskilled attackers can indeed perform technically advanced attacks, thanks to so-called *exploit scripts* which they download from the Internet [16]. An exploit script is a program designed and published by a skilled attacker. When executed, it will automatically carry out an attack on the target system. We claim that we have successfully modeled the infamous *insider* threat and that our results are of general value [14].

The second question is how well the students' reports correspond to reality. The set of actions actually carried out by the attackers in our experiments is largely in agreement with the set of actions documented by the attackers in their reports, but the sets are not identical. There are indeed actions reported that we seriously doubt were ever performed on the system, for example types of intrusions to which we know that the system was not vulnerable. This became painfully

clear to the students who later tried to confirm the reports in the intrusion analysis exercise. It is also likely that there are actions that the attackers performed but never reported to us. Why is there not a perfect match between the two sets? For the first case, we do not really know what reason the students have for exaggerating their results. The grading of the reports was not based on their successes but on their efforts and their analysis of their work. We had also made it clear that their reports would be used as research material, and it is therefore disappointing and worrisome that some of them tried to polish their results. For the second case, we can only hope that actions left out of the reports were not reported because the students considered them insignificant and not because they wanted to hide something from us.

Finally, there is the problem of using a relatively large group of humans as “guinea pigs” at an engineering department. Our expertise is in technology, not in the behavioral sciences. Factors not related to technology can greatly affect the results, for example, how instructions are interpreted, what motivation the participants have, whether they obey the rules of the experiment, whether they tell the truth etc. This is probably an area in which we could benefit from consulting external expertise.

### 8.3 Synergy Effects

To understand the benefits of conducting this kind of research and education together, we can reason hypothetically about having one without the other. First, we consider having education without research. Before our research group was formed, security was only briefly introduced to students as a part of other courses, for example, operating systems, computer communication, and computer networks. The growing general interest in security had probably led to the introduction of a separate course in security sooner or later. However, security is a demanding topic for teachers because it is a rapidly moving field in which it is difficult to remain up-to-date—not only because of the sheer amount of information but also because of obscurity, misunderstandings, and a lack of perspective in many bulletins. Teachers must also have a thorough knowledge of the fundamental principles and lessons often forgotten today. Without the genuine interest and devotion to the topic that active members of the research community have, it is much more difficult to give high quality education in security.

Second, we can consider research without education. We have already argued how valuable the data produced by the students has been for our research. Computer science students are often considered a homogeneous group, but our student experiments have actually helped us to create a diversity in our data that we would not have been able to achieve on our own. Students are also naturally willing to learn and therefore ask questions that less motivated participants would not bother with.

### 8.4 Ethical Issues

In the remainder of this section, we take an external view and discuss our research methods and their consequences as they are seen outside the security research community.

#### 8.4.1 The hacker school

One concern that was voiced by several sources when our experiments came to public knowledge was that we were training university students to become computer criminals. This was amplified by some reports in newspapers and television. Fortunately, the spectacular idea of a television talk show host to broadcast attacking activity live from our computer lab was dropped a couple of days before it was to go on the air for reasons not having to do with our activities. A message that did reach the public, and still haunts us several years later, was that Chalmers had attacked hospital computers. The journalistic logic behind this conclusion was that, because we had performed our first experiment on a UNIX system and the hospital admini-

stration in a nearby county had this type of system, hospitals could probably be attacked in a similar way and, *ergo*, Chalmers had trained students to attack hospitals.

What are the facts in this case? In the intrusion experiments, we refrained entirely from training students in attacking the system, which was necessary to ensure that the data collected was valid for security modeling. Some students complained about the lack of hints and instructions from us, since we simply gave them access to a computer network which they were allowed to attack for a limited period of time. We also gave them access to the Internet, where they could seek information. In 1993, this was not something that students could normally arrange on their own. Today, any student can afford to set up a computer network in his or her home, consisting for example of a number of “old” PCs with an Intel 486 processor running Linux. Internet access is something that most people take for granted today. This change is reflected in our approach to the remedy analysis, where we in fact did tell the students how they could perform attacks. In 1997, this was a much less controversial issue because all the attacks had been discussed publicly on the Internet and because anyone could try it at home, as explained above. All along, we have informed the students about what constitutes computer crime according to national laws, and why certain behavior is inappropriate or illegal.

#### 8.4.2 Informed consent

An important concept in research ethics is *informed consent*, which means that experiment participants should be informed of all possible consequences and risks of the experiment, should understand that information, and should be given the possibility to refuse participation. In the intrusion experiments, the students who played the role of attackers were of course informed about the experiment but were not informed that they could refuse and be given another assignment instead, simply because we assumed that they all wanted to participate.

The situation was different for the students who were ordinary users of the system. They were not informed about the experiments because we did not want them to be more concerned about security than they normally are. Unknowingly, they played the role of victims. Some were deceived by forged e-mail to send the attackers their passwords, others had their passwords revealed in other ways and, worst of all, some passwords to other systems were monitored by the attackers when lab computers were used as terminals for remote access. This is perhaps the most questionable part of the experiment, with a trade-off between realism and ethical concerns. The users were not as easily fooled in the later experiments because they had heard about our previous activities. In fact, any computer malfunction was blamed on the security experimenters.

In the intrusion detection exercise, we wanted to use primarily real recorded data to make the students feel the realism and relevance of their assignment. Users accessing the monitored FTP server had been warned through a login banner message that their activities were monitored, and the passwords for non-anonymous users were never recorded. Still, there can be innocent users whose transactions appear suspicious in the log file or other privacy concerns. Again, there is a trade-off between realism and ethics.

#### 8.4.3 Related discussions on ethics

In a discussion of the ethical aspects of spreading information on methods for computer crime, computer crime expert Donn Parker claims that the intent of the publisher is what matters [20]. The director of security research at Purdue University, Eugene Spafford, shares this view [21]. If the intent is to raise awareness and protect systems, then it is ethical (and legal). If the intent is to encourage people to attack systems, then it is unethical and probably illegal.

At the 21st National Information Systems Security Conference in Arlington, Virginia, in October of 1998, there was a panel discussion entitled “Do attack/defend exercises belong in the classroom?”. It is interesting to note that all panelists from academia were in favor of such exercises as part of security education and that it was difficult for the organizers to find a panelist

with the opposite opinion. One concern brought up in the discussion was that students could use skills acquired in the exercise for evil purposes. This is however a general concern that is not restricted to computer security; any tool or skill can be used to do evil. A similar view is presented by White and Nordstrom [22], who claim that it would be more dangerous *not* to educate future system administrators in the details of attacking techniques, because they would otherwise be “sitting ducks” for attackers who possess these skills.

## 9. CONCLUSIONS

We have presented how we have successfully used students to produce security research data while at the same time educating the students in applied computer security. The close connection between research and education has been a motivating factor for the students, and our view of the students as a valuable research resource has increased our motivation in our role as teachers. At the same time, we have collected valuable research data that could not easily have been acquired in other ways. In many cases, the dual role as teacher and researcher at a university leads to conflicts of interest, but we have managed to turn that duality into a fruitful synergy.

## ACKNOWLEDGMENTS

We are grateful to all the students that have participated in the experiments throughout the years.

## REFERENCES

- [1] C. R. Attanasio, P. Markstein, and R. J. Phillips. Penetrating an operating system: A study of VM/370 integrity. *IBM Systems Journal*, 15(1):102–116, 1976.
- [2] Sarah Brocklehurst, Bev Littlewood, Tomas Olovsson, and Erland Jonsson. On measurement of operational security. In *Proceedings of the Ninth Annual Conference on Computer Assurance, COMPASS'94*, pages 257–266, Gaithersburg, MD, USA, June 27-July 1 1994. IEEE.
- [3] Dan Farmer and Wietse Venema. Improving the security of your site by breaking into it. Posted on comp.security.unix and several other Usenet newsgroups, December 1993.
- [4] Peter D. Goldis. Questions and answers about tiger teams. *The EDP Audit, Control, and Security Newsletter*, XVII(4):1–10, October 1989.
- [5] Ulf Gustafson, Erland Jonsson, and Tomas Olovsson. On the modelling of preventive security based on a PC network intrusion experiment. In *Proceedings of the Australasian Conference on Information Security and Privacy, ACISP'96*, volume 1172 of LNCS, pages 242–252, Wollongong, Australia, June 24-26 1996. Springer-Verlag.
- [6] Ulf Gustafson, Erland Jonsson, and Tomas Olovsson. Security evaluation of a PC network based on intrusion experiments. In *Proceedings of the 14th International Congress on Computer and Communications Security, SECURICOM'96*, pages 187–202, Paris, France, June 5-6 1996.
- [7] Hans Hedbom, Stefan Lindskog, Stefan Axelsson, and Erland Jonsson. Analysis of the security of Windows NT. Technical Report 98:04, Department of Computer Science, Karlstad University, SE-651 88 Karlstad, Sweden, 1998.

- [8] Hans Hedbom, Stefan Lindskog, and Erland Jonsson. A preliminary evaluation of the security of a non-distributed version of Windows NT. In Arto Karila and Timo Aalto, editors, *Proceedings of the Second Nordic Workshop on Secure Computer Systems*, Espoo, Finland, November 6-7 1997. Helsinki University of Technology.
- [9] Israel Samuel Herschberg. Make the tigers hunt for you. *Computers & Security*, 7(2):197–203, 1988.
- [10] Cynthia Irvine, Shiu-Kai Chin, and Deborah Frincke. Integrating security into the curriculum. *Computer*, 31(12):25–30, December 1998.
- [11] Cynthia E. Irvine, Daniel F. Warren, and Paul C. Clark. The NPS CISR graduate program in infosec: Six years of experience. In *Proceedings of the 20th National Information Systems Security Conference*, pages 22–30, Baltimore, MD, USA, October 1997. National Institute of Standards and Technology/National Computer Security Center.
- [12] Erland Jonsson and Lech J. Janczewski. A taxonomy and overview of information security experiments. In Louise Yngström and Jan Carlsen, editors, *Proceedings of the 13th International Information Systems Security, IFIP/SEC'97*, pages 139–150, Copenhagen, Denmark, May 14-16 1997. Chapman & Hall.
- [13] Erland Jonsson and Tomas Olovsson. An empirical model of the security intrusion process. In *Proceedings of the Eleventh Annual Conference on Computer Assurance, COMPASS'96*, pages 176–186, Gaithersburg, MD, USA, June 17-21 1996. IEEE.
- [14] Erland Jonsson and Tomas Olovsson. A quantitative model of the security intrusion process based on attacker behavior. *IEEE Transactions on Software Engineering*, 2(4):235–245, April 1997.
- [15] Carl E. Landwehr, Alan R. Bull, John P. McDermott, and William S. Choi. A taxonomy of computer program security flaws, with examples. *ACM Computing Surveys*, 26(3):211–254, September 1994.
- [16] Ulf Lindqvist and Erland Jonsson. A map of security risks associated with using COTS. *Computer*, 31(6):60–66, June 1998.
- [17] Ulf Lindqvist, Tomas Olovsson, and Erland Jonsson. An analysis of a secure system based on trusted components. In *Proceedings of the Eleventh Annual Conference on Computer Assurance, COMPASS'96*, pages 213–223, Gaithersburg, MD, USA, June 17-21 1996. IEEE.
- [18] Ulf Lindqvist and Phillip A Porras. Detecting computer and network misuse through the production-based expert system toolset (P-BEST). In *Proceedings of the 1999 IEEE Symposium on Security and Privacy*, Oakland, CA, USA, May 9–12, 1999. IEEE Computer Society Press. To appear.
- [19] Tomas Olovsson, Erland Jonsson, Sarah Brocklehurst, and Bev Littlewood. Towards operational measures of computer security: Experimentation and modelling. In Brian Randell, Jean-Claude Laprie, Hermann Kopetz, and Bev Littlewood, editors, *Predictably Dependable Computing Systems*, chapter VIII, pages 555–572. Springer-Verlag, 1995.
- [20] Donn B Parker. Colleagues debate Denning's comments. *Communications of the ACM*, 34(3):33–41, March 1991. Reprinted in K. W. Bowyer, *Ethics and Computing*, IEEE Computer Society Press, 1996.
- [21] Eugene H Spafford. Are computer hacker break-ins ethical? In Deborah G Johnson and Helen Nissenbaum, editors, *Computers, Ethics & Social Values*, pages 125–135. Prentice-Hall, 1995.
- [22] Gregory White and Gregory Nordstrom. Security across the curriculum: Using computer security to teach computer science principles. In *Proceedings of the 19th National Information Systems Security Conference*, pages 483–488, Baltimore, Maryland, October 22–25, 1996. National Institute of Standards and Technology/National Computer Security Center.