

Ethical hacking

by C. C. Palmer

The explosive growth of the Internet has brought many good things: electronic commerce, easy access to vast stores of reference material, collaborative computing, e-mail, and new avenues for advertising and information distribution, to name a few. As with most technological advances, there is also a dark side: criminal hackers. Governments, companies, and private citizens around the world are anxious to be a part of this revolution, but they are afraid that some hacker will break into their Web server and replace their logo with pornography, read their e-mail, steal their credit card number from an on-line shopping site, or implant software that will secretly transmit their organization's secrets to the open Internet. With these concerns and others, the ethical hacker can help. This paper describes ethical hackers: their skills, their attitudes, and how they go about helping their customers find and plug up security holes. The ethical hacking process is explained, along with many of the problems that the Global Security Analysis Lab has seen during its early years of ethical hacking for IBM clients.

The term “hacker” has a dual usage in the computer industry today. Originally, the term was defined as:

HACKER *noun* 1. A person who enjoys learning the details of computer systems and how to stretch their capabilities—as opposed to most users of computers, who prefer to learn only the minimum amount necessary. 2. One who programs enthusiastically or who enjoys programming rather than just theorizing about programming.¹

This complimentary description was often extended to the verb form “hacking,” which was used to de-

scribe the rapid crafting of a new program or the making of changes to existing, usually complicated software.

As computers became increasingly available at universities, user communities began to extend beyond researchers in engineering or computer science to other individuals who viewed the computer as a curiously flexible tool. Whether they programmed the computers to play games, draw pictures, or to help them with the more mundane aspects of their daily work, once computers were available for use, there was never a lack of individuals wanting to use them.

Because of this increasing popularity of computers and their continued high cost, access to them was usually restricted. When refused access to the computers, some users would challenge the access controls that had been put in place. They would steal passwords or account numbers by looking over someone's shoulder, explore the system for bugs that might get them past the rules, or even take control of the whole system. They would do these things in order to be able to run the programs of their choice, or just to change the limitations under which their programs were running.

Initially these computer intrusions were fairly benign, with the most damage being the theft of computer time. Other times, these recreations would take the

©Copyright 2001 by International Business Machines Corporation. Copying in printed form for private use is permitted without payment of royalty provided that (1) each reproduction is done without alteration and (2) the *Journal* reference and IBM copyright notice are included on the first page. The title and abstract, but no other portions, of this paper may be copied or distributed royalty free without further permission by computer-based and other information-service systems. Permission to *republish* any other portion of this paper must be obtained from the Editor.

form of practical jokes. However, these intrusions did not stay benign for long. Occasionally the less talented, or less careful, intruders would accidentally bring down a system or damage its files, and the system administrators would have to restart it or make repairs. Other times, when these intruders were again denied access once their activities were discovered, they would react with purposefully destructive actions. When the number of these destructive computer intrusions became noticeable, due to the visibility of the system or the extent of the damage inflicted, it became “news” and the news media picked up on the story. Instead of using the more accurate term of “computer criminal,” the media began using the term “hacker” to describe individuals who break into computers for fun, revenge, or profit. Since calling someone a “hacker” was originally meant as a compliment, computer security professionals prefer to use the term “cracker” or “intruder” for those hackers who turn to the dark side of hacking. For clarity, we will use the explicit terms “ethical hacker” and “criminal hacker” for the rest of this paper.

What is ethical hacking?

With the growth of the Internet, computer security has become a major concern for businesses and governments. They want to be able to take advantage of the Internet for electronic commerce, advertising, information distribution and access, and other pursuits, but they are worried about the possibility of being “hacked.” At the same time, the potential customers of these services are worried about maintaining control of personal information that varies from credit card numbers to social security numbers and home addresses.²

In their search for a way to approach the problem, organizations came to realize that one of the best ways to evaluate the intruder threat to their interests would be to have independent computer security professionals attempt to break into their computer systems. This scheme is similar to having independent auditors come into an organization to verify its bookkeeping records. In the case of computer security, these “tiger teams” or “ethical hackers”³ would employ the same tools and techniques as the intruders, but they would neither damage the target systems nor steal information. Instead, they would evaluate the target systems’ security and report back to the owners with the vulnerabilities they found and instructions for how to remedy them.

This method of evaluating the security of a system has been in use from the early days of computers. In one early ethical hack, the United States Air Force conducted a “security evaluation” of the Multics operating systems for “potential use as a two-level (secret/top secret) system.”⁴ Their evaluation found that while Multics was “significantly better than other conventional systems,” it also had “. . . vulnerabilities in hardware security, software security, and procedural security” that could be uncovered with “a relatively low level of effort.” The authors performed their tests under a guideline of realism, so that their results would accurately represent the kinds of access that an intruder could potentially achieve. They performed tests that were simple information-gathering exercises, as well as other tests that were outright attacks upon the system that might damage its integrity. Clearly, their audience wanted to know both results. There are several other now unclassified reports that describe ethical hacking activities within the U.S. military.⁵⁻⁷

With the growth of computer networking, and of the Internet in particular, computer and network vulnerability studies began to appear outside of the military establishment. Most notable of these was the work by Farmer and Venema,⁸ which was originally posted to Usenet⁹ in December of 1993. They discussed publicly, perhaps for the first time,¹⁰ this idea of using the techniques of the hacker to assess the security of a system. With the goal of raising the overall level of security on the Internet and intranets, they proceeded to describe how they were able to gather enough information about their targets to have been able to compromise security if they had chosen to do so. They provided several specific examples of how this information could be gathered and exploited to gain control of the target, and how such an attack could be prevented.

Farmer and Venema elected to share their report freely on the Internet in order that everyone could read and learn from it. However, they realized that the testing at which they had become so adept might be too complex, time-consuming, or just too boring for the typical system administrator to perform on a regular basis. For this reason, they gathered up all the tools that they had used during their work, packaged them in a single, easy-to-use application, and gave it away to anyone who chose to download it.¹¹ Their program, called Security Analysis Tool for Auditing Networks, or SATAN, was met with a great amount of media attention around the world. Most of this early attention was negative, because the tool’s

capabilities were misunderstood. The tool was not an automated hacker program that would bore into systems and steal their secrets. Rather, the tool performed an audit that both identified the vulnerabilities of a system and provided advice on how to eliminate them. Just as banks have regular audits of their accounts and procedures, computer systems also need regular checking. The SATAN tool provided that auditing capability, but it went one step further: it also advised the user on how to correct the problems it discovered. The tool did not tell the user how the vulnerability might be exploited, because there would be no useful point in doing so.

Who are ethical hackers?

These early efforts provide good examples of ethical hackers. Successful ethical hackers possess a variety of skills. First and foremost, they must be completely trustworthy. While testing the security of a client's systems, the ethical hacker may discover information about the client that should remain secret. In many cases, this information, if publicized, could lead to real intruders breaking into the systems, possibly leading to financial losses. During an evaluation, the ethical hacker often holds the "keys to the company," and therefore must be trusted to exercise tight control over any information about a target that could be misused. The sensitivity of the information gathered during an evaluation requires that strong measures be taken to ensure the security of the systems being employed by the ethical hackers themselves: limited-access labs with physical security protection and full ceiling-to-floor walls, multiple secure Internet connections, a safe to hold paper documentation from clients, strong cryptography to protect electronic results, and isolated networks for testing.

Ethical hackers typically have very strong programming and computer networking skills and have been in the computer and networking business for several years. They are also adept at installing and maintaining systems that use the more popular operating systems (e.g., UNIX** or Windows NT**) used on target systems. These base skills are augmented with detailed knowledge of the hardware and software provided by the more popular computer and networking hardware vendors. It should be noted that an additional specialization in security is not always necessary, as strong skills in the other areas imply a very good understanding of how the security on various systems is maintained. These systems management skills are necessary for the actual vulner-

ability testing, but are equally important when preparing the report for the client after the test.

Finally, good candidates for ethical hacking have more drive and patience than most people. Unlike the way someone breaks into a computer in the mov-

**Just as in sports or warfare,
knowledge of the skills
and techniques of your opponent
is vital to your success.**

ies, the work that ethical hackers do demands a lot of time and persistence. This is a critical trait, since criminal hackers are known to be extremely patient and willing to monitor systems for days or weeks while waiting for an opportunity. A typical evaluation may require several days of tedious work that is difficult to automate. Some portions of the evaluations must be done outside of normal working hours to avoid interfering with production at "live" targets or to simulate the timing of a real attack. When they encounter a system with which they are unfamiliar, ethical hackers will spend the time to learn about the system and try to find its weaknesses. Finally, keeping up with the ever-changing world of computer and network security requires continuous education and review.

One might observe that the skills we have described could just as easily belong to a criminal hacker as to an ethical hacker. Just as in sports or warfare, knowledge of the skills and techniques of your opponent is vital to your success. In the computer security realm, the ethical hacker's task is the harder one. With traditional crime anyone can become a shoplifter, graffiti artist, or a mugger. Their potential targets are usually easy to identify and tend to be localized. The local law enforcement agents must know how the criminals ply their trade and how to stop them. On the Internet anyone can download criminal hacker tools and use them to attempt to break into computers anywhere in the world. Ethical hackers have to know the techniques of the criminal hackers, how their activities might be detected, and how to stop them.

Given these qualifications, how does one go about finding such individuals? The best ethical hacker candidates will have successfully published research papers or released popular open-source security software.¹² The computer security community is strongly self-policing, given the importance of its work. Most ethical hackers, and many of the better computer and network security experts, did not set out to focus on these issues. Most of them were computer users from various disciplines, such as astronomy and physics, mathematics, computer science, philosophy, or liberal arts, who took it personally when someone disrupted their work with a hack.

One rule that IBM's ethical hacking effort had from the very beginning was that we would not hire ex-hackers. While some will argue that only a "real hacker" would have the skill to actually do the work, we feel that the requirement for absolute trust eliminated such candidates. We likened the decision to that of hiring a fire marshal for a school district: while a gifted ex-arsonist might indeed know everything about setting and putting out fires, would the parents of the students really feel comfortable with such a choice? This decision was further justified when the service was initially offered: the customers themselves asked that such a restriction be observed. Since IBM's ethical hacking group was formed, there have been numerous ex-hackers who have become security consultants and spokespersons for the news media. While they may very well have turned away from the "dark side," there will always be a doubt.

What do ethical hackers do?

An ethical hacker's evaluation of a system's security seeks answers to three basic questions:

- What can an intruder see on the target systems?
- What can an intruder do with that information?
- Does anyone at the target notice the intruder's attempts or successes?

While the first and second of these are clearly important, the third is even more important: If the owners or operators of the target systems do not notice when someone is trying to break in, the intruders can, and will, spend weeks or months trying and will usually eventually succeed.

When the client requests an evaluation, there is quite a bit of discussion and paperwork that must be done up front. The discussion begins with the client's an-

swers to questions similar to those posed by Garfinkel and Spafford:¹³

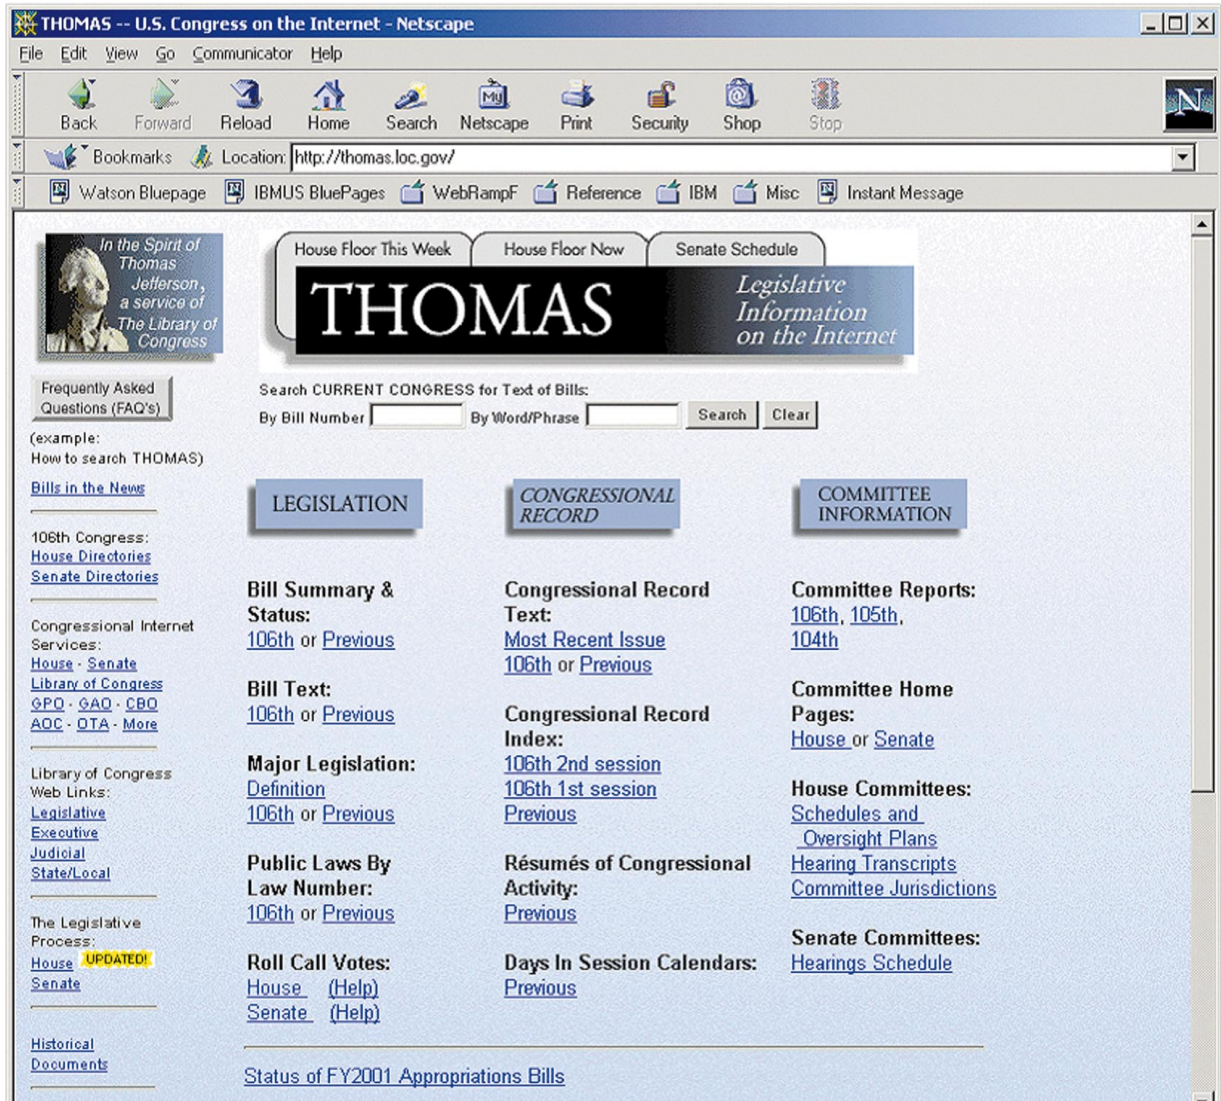
1. What are you trying to protect?
2. What are you trying to protect against?
3. How much time, effort, and money are you willing to expend to obtain adequate protection?

A surprising number of clients have difficulty precisely answering the first question: a medical center might say "our patient information," an engineering firm might answer "our new product designs," and a Web retailer might answer "our customer database."

All of these answers fall short, since they only describe targets in a general way. The client usually has to be guided to succinctly describe all of the critical information assets for which loss could adversely affect the organization or its clients. These assets should also include secondary information sources, such as employee names and addresses (which are privacy and safety risks), computer and network information (which could provide assistance to an intruder), and other organizations with which this organization collaborates (which provide alternate paths into the target systems through a possibly less secure partner's system).

A complete answer to (2) specifies more than just the loss of the things listed in answer to (1). There are also the issues of system availability, wherein a denial-of-service attack could cost the client actual revenue and customer loss because systems were unavailable. The world became quite familiar with denial-of-service attacks in February of 2000 when attacks were launched against eBay**, Yahoo!**, E*TRADE**, CNN**, and other popular Web sites. During the attacks, customers were unable to reach these Web sites, resulting in loss of revenue and "mind share." The answers to (1) should contain more than just a list of information assets on the organization's computer. The level of damage to an organization's good image resulting from a successful criminal hack can range from merely embarrassing to a serious threat to revenue. As an example of a hack affecting an organization's image, on January 17, 2000, a U.S. Library of Congress Web site was attacked. The original initial screen is shown in Figure 1, whereas the hacked screen is shown in Figure 2. As is often done, the criminal hacker left his or her nickname, or handle, near the top of the page in order to guarantee credit for the break-in.

Figure 1 Library of Congress Web page before attack

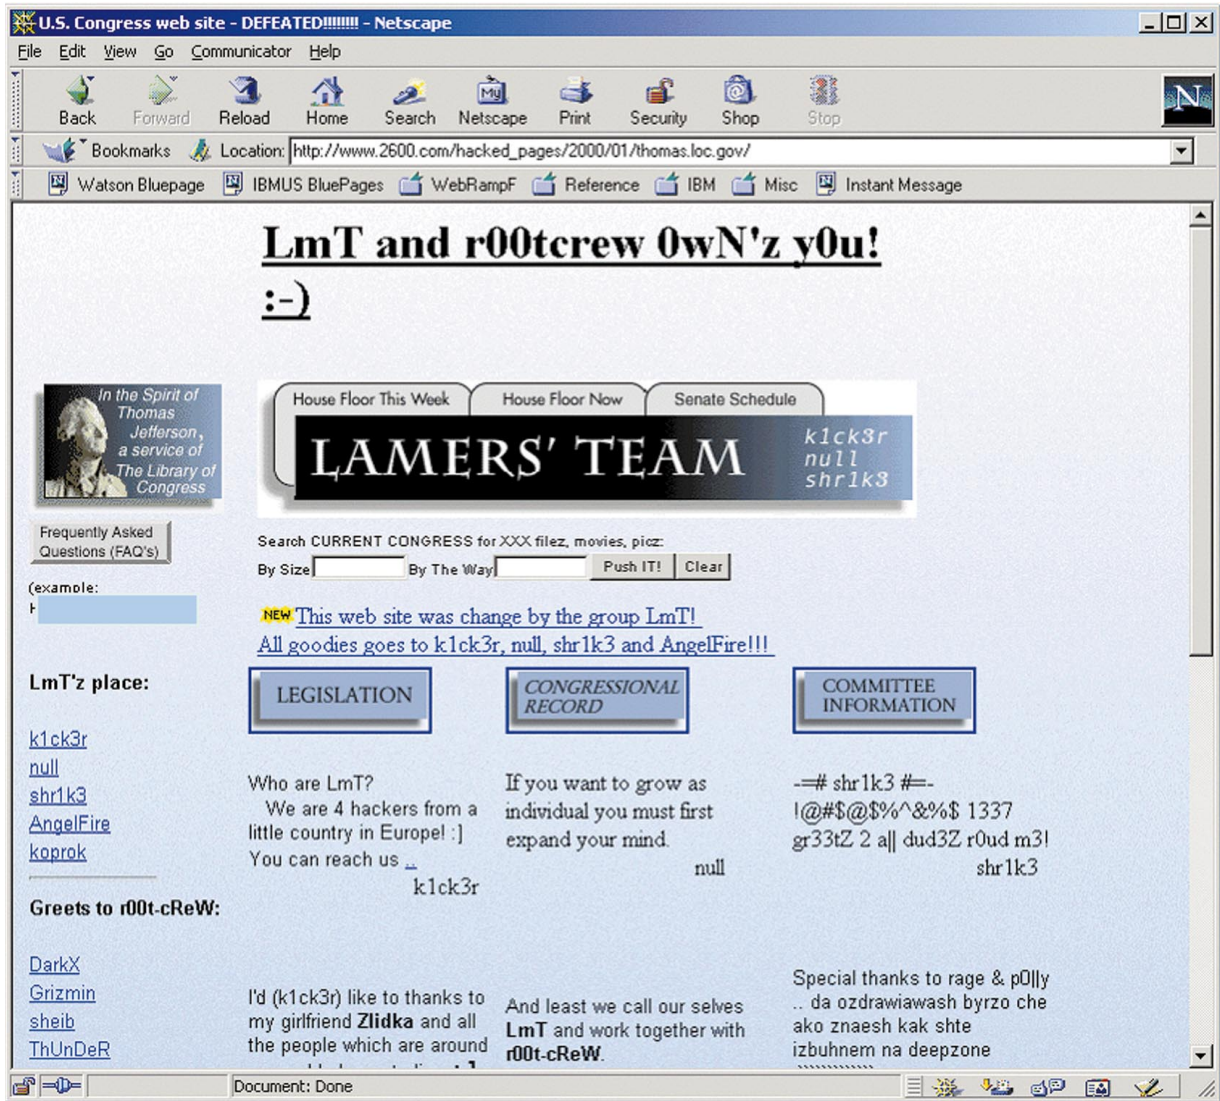


Some clients are under the mistaken impression that their Web site would not be a target. They cite numerous reasons, such as “it has nothing interesting on it” or “hackers have never heard of my company.” What these clients do not realize is that *every Web site is a target*. The goal of many criminal hackers is simple: Do something spectacular and then make sure that all of your pals know that you did it. Another rebuttal is that many hackers simply do not care who your company or organization is; they hack your Web site *because they can*. For example, Web

administrators at UNICEF (United Nations Children’s Fund) might very well have thought that no hacker would attack them. However, in January of 1998, their page was defaced as shown in Figures 3 and 4. Many other examples of hacked Web pages can be found at archival sites around the Web.¹⁴

Answers to the third question are complicated by the fact that computer and network security costs come in three forms. First there are the real monetary costs incurred when obtaining security consulting, hiring

Figure 2 Hacked Library of Congress Web page



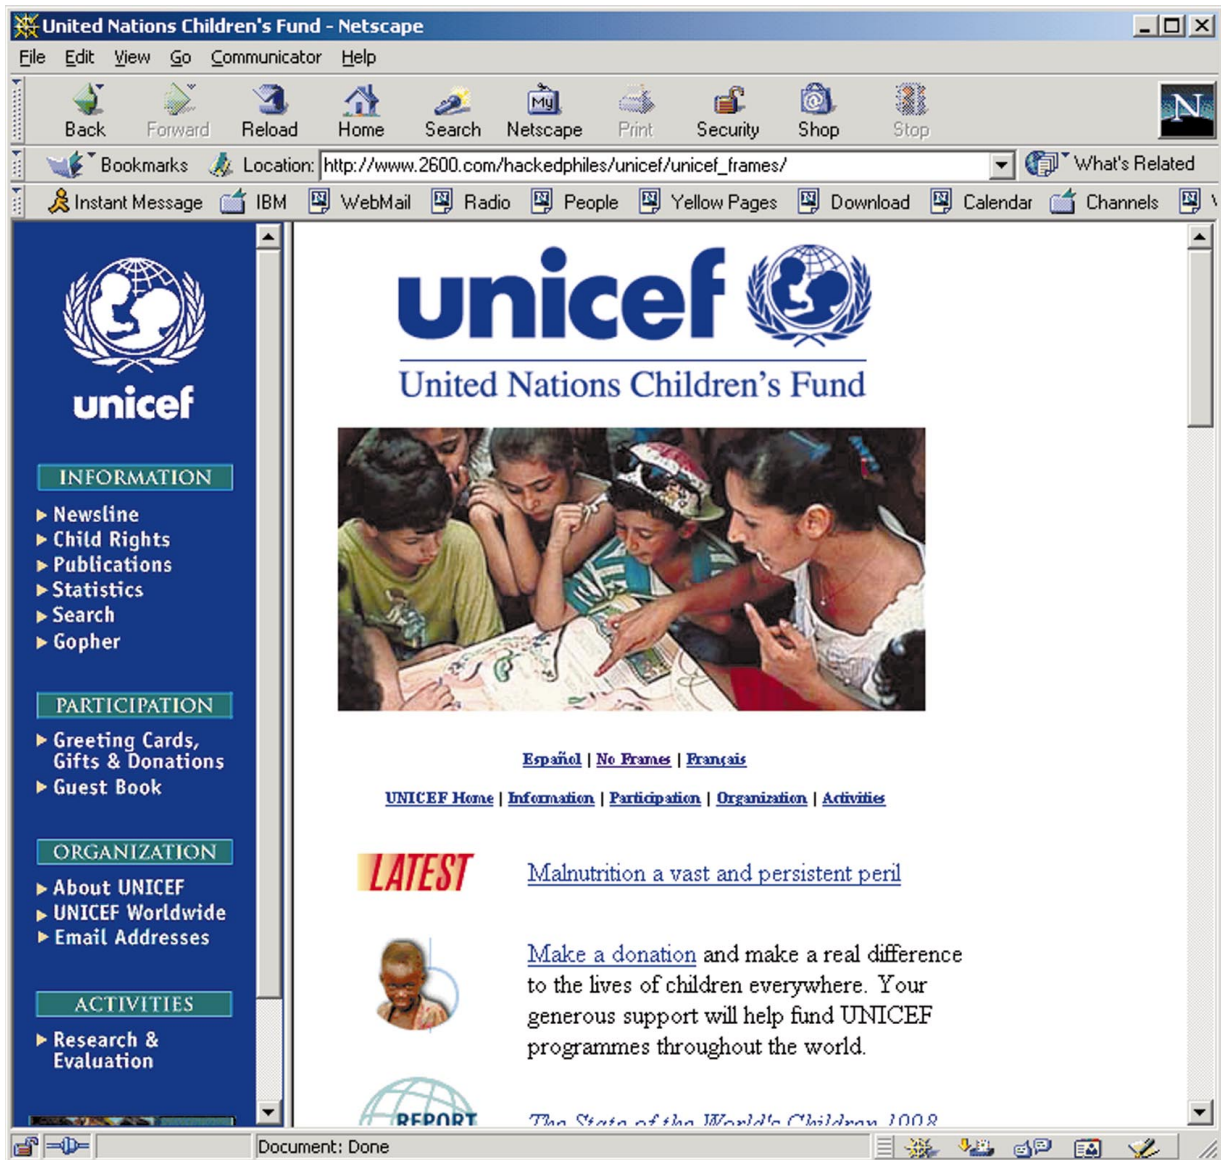
personnel, and deploying hardware and software to support security needs. Second, there is the cost of usability: the more secure a system is, the more difficult it can be to make it easy to use. The difficulty can take the form of obscure password selection rules, strict system configuration rules, and limited remote access. Third, there is the cost of computer and network performance. The more time a computer or network spends on security needs, such as strong cryptography and detailed system activity logging, the less time it has to work on user problems.

Because of Moore's Law,¹⁵ this may be less of an issue for mainframe, desktop, and laptop machines. Yet, it still remains a concern for mobile computing.

The "get out of jail free card"

Once answers to these three questions have been determined, a security evaluation plan is drawn up that identifies the systems to be tested, how they should be tested, and any limitations on that testing. Commonly referred to as a "get out of jail free card," this

Figure 3 UNICEF Web page before attack

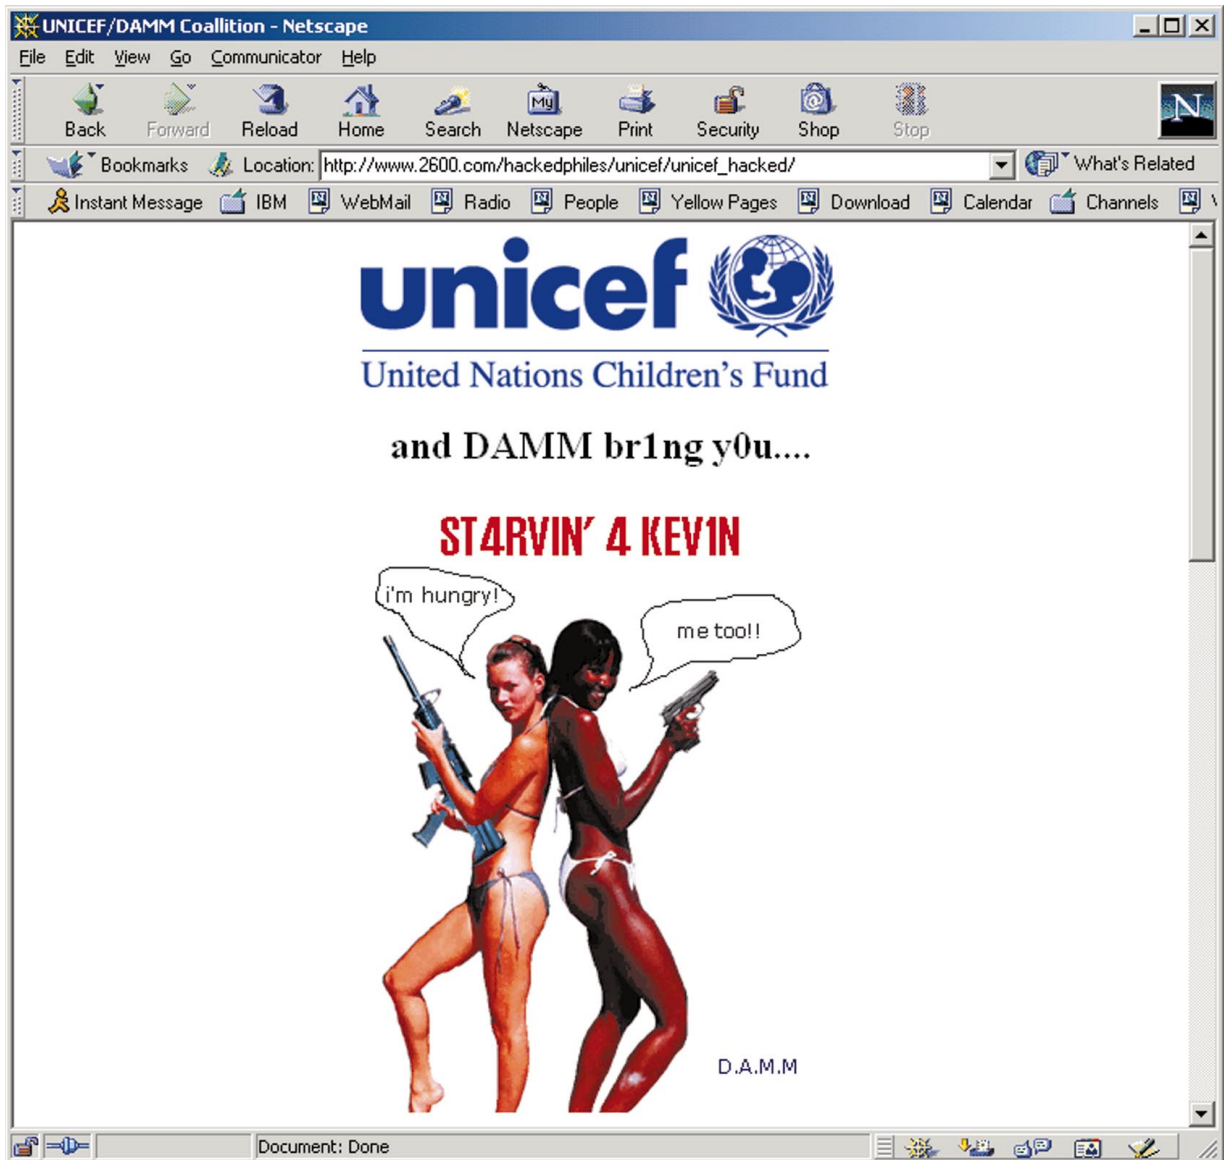


is the contractual agreement between the client and the ethical hackers, who typically write it together. This agreement also protects the ethical hackers against prosecution, since much of what they do during the course of an evaluation would be illegal in most countries. The agreement provides a precise description, usually in the form of network addresses or modem telephone numbers, of the systems to be evaluated. Precision on this point is of the utmost

importance, since a minor mistake could lead to the evaluation of the wrong system at the client's installation or, in the worst case, the evaluation of some other organization's system.

Once the target systems are identified, the agreement must describe how they should be tested. The best evaluation is done under a "no-holds-barred" approach. This means that the ethical hacker can try

Figure 4 Hacked UNICEF Web page



anything he or she can think of to attempt to gain access to or disrupt the target system. While this is the most realistic and useful, some clients balk at this level of testing. Clients have several reasons for this, the most common of which is that the target systems are “in production” and interference with their operation could be damaging to the organization’s interests. However, it should be pointed out to such clients that these very reasons are precisely why a

“no-holds-barred” approach should be employed. An intruder will not be playing by the client’s rules. If the systems are that important to the organization’s well-being, they should be tested as thoroughly as possible. In either case, the client should be made fully aware of the risks inherent to ethical hacker evaluations. These risks include alarmed staff and unintentional system crashes, degraded network or system performance, denial of service, and log-file size explosions.

Some clients insist that as soon as the ethical hackers gain access to their network or to one of their systems, the evaluation should halt and the client be notified. This sort of ruling should be discouraged, because it prevents the client from learning all that the ethical hackers might discover about their systems. It can also lead to the client's having a false sense of security by thinking that the first security hole found is the only one present. The evaluation should be allowed to proceed, since where there is one exposure there are probably others.

The timing of the evaluations may also be important to the client. The client may wish to avoid affecting systems and networks during regular working hours. While this restriction is not recommended, it reduces the accuracy of the evaluation only somewhat, since most intruders do their work outside of the local regular working hours. However, attacks done during regular working hours may be more easily hidden. Alerts from intrusion detection systems may even be disabled or less carefully monitored during the day. Whatever timing is agreed to, the client should provide contacts within the organization who can respond to calls from the ethical hackers if a system or network appears to have been adversely affected by the evaluation or if an extremely dangerous vulnerability is found that should be immediately corrected.

It is common for potential clients to delay the evaluation of their systems until only a few weeks or days before the systems need to go on-line. Such last-minute evaluations are of little use, since implementations of corrections for discovered security problems might take more time than is available and may introduce new system problems.

In order for the client to receive a valid evaluation, the client must be cautioned to limit prior knowledge of the test as much as possible. Otherwise, the ethical hackers might encounter the electronic equivalent of the client's employees running ahead of them, locking doors and windows. By limiting the number of people at the target organization who know of the impending evaluation, the likelihood that the evaluation will reflect the organization's actual security posture is increased. A related issue that the client must be prepared to address is the relationship of the ethical hackers to the target organization's employees. Employees may view this "surprise inspection" as a threat to their jobs, so the organization's management team must be prepared to take steps to reassure them.

The ethical hack itself

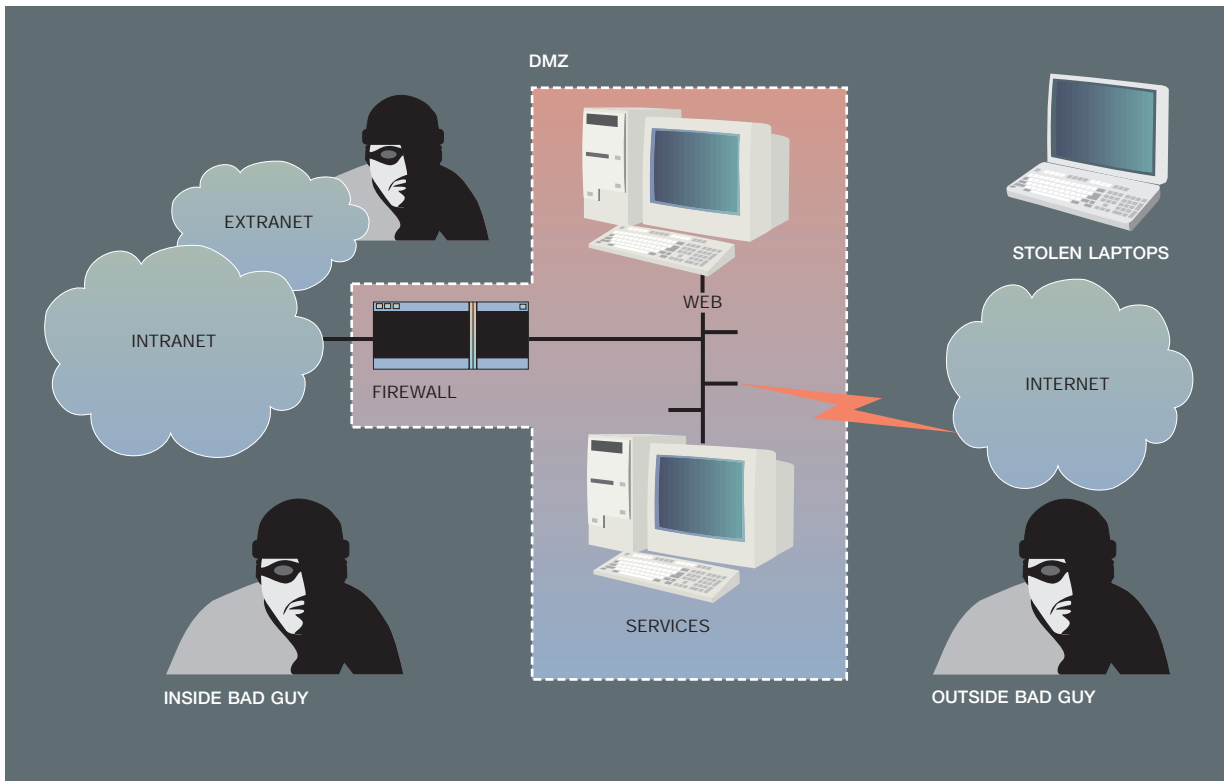
Once the contractual agreement is in place, the testing may begin as defined in the agreement. It should be noted that the testing itself poses some risk to the client, since a criminal hacker monitoring the transmissions of the ethical hackers could learn the same information. If the ethical hackers identify a weakness in the client's security, the criminal hacker could potentially attempt to exploit that vulnerability. This is especially vexing since the activities of the ethical hackers might mask those of the criminal hackers. The best approach to this dilemma is to maintain several addresses around the Internet from which the ethical hacker's transmissions will emanate, and to switch origin addresses often. Complete logs of the tests performed by the ethical hackers are always maintained, both for the final report and in the event that something unusual occurs. In extreme cases, additional intrusion monitoring software can be deployed at the target to ensure that all the tests are coming from the ethical hacker's machines. However, this is difficult to do without tipping off the client's staff and may require the cooperation of the client's Internet service provider.

The line between criminal hacking and computer virus writing is becoming increasingly blurred. When requested by the client, the ethical hacker can perform testing to determine the client's vulnerability to e-mail or Web-based virus vectors. However, it is far better for the client to deploy strong antivirus software, keep it up to date, and have a clear and simple policy in place for the reporting of incidents. IBM's Immune System for Cyberspace^{16,17} is another approach that provides the additional capability of recognizing new viruses and reporting them to a central lab that automatically analyzes the virus and provides an immediate vaccine.

As dramatized in Figure 5, there are several kinds of testing. Any combination of the following may be called for:

- *Remote network.* This test simulates the intruder launching an attack across the Internet. The primary defenses that must be defeated here are border firewalls, filtering routers, and Web servers.
- *Remote dial-up network.* This test simulates the intruder launching an attack against the client's modem pools. The primary defenses that must be defeated here are user authentication schemes. These kinds of tests should be coordinated with the local telephone company.

Figure 5 Different ways to attack computer security



- *Local network.* This test simulates an employee or other authorized person who has a legal connection to the organization's network. The primary defenses that must be defeated here are intranet firewalls, internal Web servers, server security measures, and e-mail systems.
- *Stolen laptop computer.* In this test, the laptop computer of a key employee, such as an upper-level manager or strategist, is taken by the client without warning and given to the ethical hackers. They examine the computer for passwords stored in dial-up software, corporate information assets, personnel information, and the like. Since many busy users will store their passwords on their machine, it is common for the ethical hackers to be able to use this laptop computer to dial into the corporate intranet with the owner's full privileges.
- *Social engineering.* This test evaluates the target organization's staff as to whether it would leak information to someone. A typical example of this would be an intruder calling the organization's computer help line and asking for the external tele-

phone numbers of the modem pool. Defending against this kind of attack is the hardest, because people and personalities are involved. Most people are basically helpful, so it seems harmless to tell someone who appears to be lost where the computer room is located, or to let someone into the building who "forgot" his or her badge. The only defense against this is to raise security awareness.

- *Physical entry.* This test acts out a physical penetration of the organization's building. Special arrangements must be made for this, since security guards or police could become involved if the ethical hackers fail to avoid detection. Once inside the building, it is important that the tester not be detected. One technique is for the tester to carry a document with the target company's logo on it. Such a document could be found by digging through trash cans before the ethical hack or by casually picking up a document from a trash can or desk once the tester is inside. The primary defenses here are a strong security policy, security

guards, access controls and monitoring, and security awareness.

Each of these kinds of testing can be performed from three perspectives: as a total outsider, a “semi-outsider,” or a valid user.

A total outsider has very limited knowledge about the target systems. The only information used is available through public sources on the Internet. This test represents the most commonly perceived threat. A well-defended system should not allow this kind of intruder to do anything.

A semi-outsider has limited access to one or more of the organization’s computers or networks. This tests scenarios such as a bank allowing its depositors to use special software and a modem to access information about their accounts. A well-defended system should only allow this kind of intruder to access his or her own account information.

A valid user has valid access to at least some of the organization’s computers and networks. This tests whether or not insiders with some access can extend that access beyond what has been prescribed. A well-defended system should allow an insider to access only the areas and resources that the system administrator has assigned to the insider.

The actual evaluation of the client’s systems proceeds through several phases, as described previously by Boulanger.¹⁸

The final report

The final report is a collection of all of the ethical hacker’s discoveries made during the evaluation. Vulnerabilities that were found to exist are explained and avoidance procedures specified. If the ethical hacker’s activities were noticed at all, the response of the client’s staff is described and suggestions for improvements are made. If social engineering testing exposed problems, advice is offered on how to raise awareness. This is the main point of the whole exercise: it does clients no good just to tell them that they have problems. The report must include specific advice on how to close the vulnerabilities and keep them closed. The actual techniques employed by the testers are never revealed. This is because the person delivering the report can never be sure just who will have access to that report once it is in the client’s hands. For example, an employee might want to try out some of the techniques for himself or her-

self. He or she might choose to test the company’s systems, possibly annoying system administrators or even inadvertently hiding a real attack. The employee might also choose to test the systems of another organization, which is a felony in the United States when done without permission.

The actual delivery of the report is also a sensitive issue. If vulnerabilities were found, the report could be extremely dangerous if it fell into the wrong hands. A competitor might use it for corporate espionage, a hacker might use it to break into the client’s computers, or a prankster might just post the report’s contents on the Web as a joke. The final report is typically delivered directly to an officer of the client organization in hard-copy form. The ethical hackers would have an ongoing responsibility to ensure the safety of any information they retain, so in most cases all information related to the work is destroyed at the end of the contract.

Once the ethical hack is done and the report delivered, the client might ask “So, if I fix these things I’ll have perfect security, right?” Unfortunately, this is not the case. People operate the client’s computers and networks, and people make mistakes. The longer it has been since the testing was performed, the less can be reliably said about the state of a client’s security. A portion of the final report includes recommendations for steps the client should continue to follow in order to reduce the impact of these mistakes in the future.

Conclusions

The idea of testing the security of a system by trying to break into it is not new. Whether an automobile company is crash-testing cars, or an individual is testing his or her skill at martial arts by sparring with a partner, evaluation by testing under attack from a real adversary is widely accepted as prudent. It is, however, not sufficient by itself. As Roger Schell observed nearly 30 years ago:

From a practical standpoint the security problem will remain as long as manufacturers remain committed to current system architectures, produced without a firm requirement for security. *As long as there is support for ad hoc fixes and security packages for these inadequate designs and as long as the illusory results of penetration teams are accepted as demonstrations of a computer system security, proper security will not be a reality.*¹⁹

Regular auditing, vigilant intrusion detection, good system administration practice, and computer security awareness are all essential parts of an organization's security efforts. A single failure in any of these areas could very well expose an organization to cyber-vandalism, embarrassment, loss of revenue or mind share, or worse. Any new technology has its benefits and its risks. While ethical hackers can help clients better understand their security needs, it is up to the clients to keep their guards in place.

Acknowledgments

The author would like to thank several people: the members of the Global Security Analysis Lab at IBM Research for sharing their amazing expertise and their ability to make just about anyone understand more about security; Chip Coy and Nick Simicich for their trailblazing work in defining IBM's Security Consulting Practice at the very beginning; and Paul Karger for his encyclopedic knowledge of computer security research and for his amazing ability to produce copies of every notable paper on the subject that was ever published.

**Trademark or registered trademark of the Open Group, Microsoft Corporation, eBay Inc., Yahoo! Inc., E*TRADE Securities, Inc., or Cable News Network LP, LLLP.

Cited references and notes

1. E. S. Raymond, *The New Hacker's Dictionary*, MIT Press, Cambridge, MA (1991).
2. S. Garfinkel, *Database Nation*, O'Reilly & Associates, Cambridge, MA (2000).
3. The first use of the term "ethical hackers" appears to have been in an interview with John Patrick of IBM by Gary Antheus that appeared in a June 1995 issue of *ComputerWorld*.
4. P. A. Karger and R. R. Schell, *Multics Security Evaluation: Vulnerability Analysis*, ESD-TR-74-193, Vol. II, Headquarters Electronic Systems Division, Hanscom Air Force Base, MA (June 1974).
5. S. M. Goheen and R. S. Fiske, *OS/360 Computer Security Penetration Exercise*, WP-4467, The MITRE Corporation, Bedford, MA (October 16, 1972).
6. R. P. Abbott, J. S. Chen, J. E. Donnelly, W. L. Konigsford, and S. T. Tokubo, *Security Analysis and Enhancements of Computer Operating Systems*, NBSIR 76-1041, National Bureau of Standards, Washington, DC (April 1976).
7. W. M. Inglis, *Security Problems in the WWMCCS GCOS System*, Joint Technical Support Activity Operating System Technical Bulletin 730S-12, Defense Communications Agency (August 2, 1973).
8. D. Farmer and W. Z. Venema, "Improving the Security of Your Site by Breaking into It," originally posted to Usenet (December 1993); it has since been updated and is now available at <ftp://ftp.porcupine.org/pub/security/index.html#documents>.
9. See <http://www.faqs.org/usenet/>.
10. Who can really determine who said something first on the Internet?
11. See <http://www.cs.ruu.nl/cert-uu/satan.html>.
12. This strategy is based on the ideal of raising the security of the whole Internet by giving security software away. Thus, no one will have any excuse not to take action to improve security.
13. S. Garfinkel and E. Spafford, *Practical Unix Security*, First Edition, O'Reilly & Associates, Cambridge, MA (1996).
14. For a collection of previously hacked Web sites, see http://www.2600.com/hacked_pages/ or <http://defaced.alldes.de>. Be forewarned, however, that some of the hacked pages may contain pornographic images.
15. In 1965, Intel cofounder Gordon Moore was preparing a speech and made a memorable observation. When he started to graph data about the growth in memory chip performance, he realized there was a striking trend. Each new chip contained roughly twice as much capacity as its predecessor, and each chip was released within 18–24 months of the previous chip. In subsequent years, the pace slowed down a bit, but data density has doubled approximately every 18 months, and this is the current definition of Moore's Law.
16. J. O. Kephart, G. B. Sorkin, D. M. Chess, and S. R. White, "Fighting Computer Viruses," *Scientific American* 277, No. 5, 88–93 (November 1997).
17. See <http://www.research.ibm.com/antivirus/SciPapers.htm> for additional antivirus research papers.
18. A. Boulanger, "Catapults and Grappling Hooks: The Tools and Techniques of Information Warfare," *IBM Systems Journal* 37, No. 1, 106–114 (1998).
19. R. R. Schell, P. J. Downey, and G. J. Popek, *Preliminary Notes on the Design of Secure Military Computer Systems*, MCI-73-1, ESD/AFSC, Hanscom Air Force Base, Bedford, MA (January 1973).

Accepted for publication April 13, 2001.

Charles C. Palmer IBM Research Division, Thomas J. Watson Research Center, P.O. Box 218, Yorktown Heights, New York 10598 (electronic mail: ccpalmer@us.ibm.com). Dr. Palmer manages the Network Security and Cryptography department at the IBM Thomas J. Watson Research Center. His teams work in the areas of cryptography research, Internet security technologies, Java™ security, privacy, and the Global Security Analysis Lab (GSAL), which he cofounded in 1995. As part of the GSAL, Dr. Palmer worked with IBM Global Services to start IBM's ethical hacking practice. He frequently speaks on the topics of computer and network security at conferences around the world. He was also an adjunct professor of computer science at Polytechnic University, Hawthorne, New York, from 1993 to 1997. He holds four patents and has several publications from his work at IBM and Polytechnic.