

©Copyright 2003 AT&T and Lumeta. All Rights Reserved.

Notice: For personal use only. These materials may not be reproduced or distributed in any form or by any means except that they may be downloaded from this source and printed for personal use.

Preface to the Second Edition

But after a time, as Frodo did not show any sign of writing a book on the spot, the hobbits returned to their questions about doings in the Shire.

Lord of the Rings
—J.R.R. TOLKIEN

The first printing of the First Edition appeared at the Las Vegas Interop in May, 1994. At that same show appeared the first of many commercial firewall products. In many ways, the field has matured since then: You can buy a decent firewall off the shelf from many vendors.

The problem of deploying that firewall in a secure and useful manner remains. We have studied many Internet access arrangements in which the only secure component was the firewall itself—it was easily bypassed by attackers going after the “protected” inside machines. Before the trivestiture of AT&T/Lucent/NCR, there were over 300,000 hosts behind at least six firewalls, plus special access arrangements with some 200 business partners.

Our first edition did not discuss the massive sniffing attacks discovered in the spring of 1994. Sniffers had been running on important Internet Service Provider (ISP) machines for months—machines that had access to a major percentage of the ISP’s packet flow. By some estimates, these sniffers captured over a million host name/user name/password sets from passing *telnet*, *ftp*, and *rlogin* sessions. There were also reports of increased hacker activity on military sites. It’s obvious what must have happened: If you are a hacker with a million passwords in your pocket, you are going to look for the most interesting targets, and *.mil* certainly qualifies.

Since the First Edition, we have been slowly losing the Internet *arms race*. The hackers have developed and deployed tools for attacks we had been anticipating for years. IP spoofing [Shimomura, 1996] and TCP hijacking are now quite common, according to the *Computer Emergency Response Team (CERT)*. ISPs report that attacks on the Internet’s infrastructure are increasing.

There was one attack we chose not to include in the First Edition: the SYN-flooding denial-of-service attack that seemed to be unstoppable. Of course, the Bad Guys learned about the attack anyway, making us regret that we had deleted that paragraph in the first place. We still believe that it is better to disseminate this information, informing saints and sinners at the same time. The saints need all the help they can get, and the sinners have their own channels of communication.

Crystal Ball or Bowling Ball?

The first edition made a number of predictions, explicitly or implicitly. Was our foresight accurate?

Our biggest failure was neglecting to foresee how successful the Internet would become. We barely mentioned the Web and declined a suggestion to use some weird syntax when listing software resources. The syntax, of course, was the URL. . .

Concomitant with the growth of the Web, the patterns of Internet connectivity vastly increased. We assumed that a company would have only a few external connections—few enough that they’d be easy to keep track of, and to firewall. Today’s spaghetti topology was a surprise.

We didn’t realize that PCs would become Internet clients as soon as they did. We did, however, warn that as personal machines became more capable, they’d become more vulnerable. Experience has proved us very correct on that point.

We did anticipate high-speed home connections, though we spoke of ISDN, rather than cable modems or DSL. (We had high-speed connectivity even then, though it was slow by today’s standards.) We also warned of issues posed by home LANs, and we warned about the problems caused by roaming laptops.

We were overly optimistic about the deployment of IPv6 (which was called IPng back then, as the choice hadn’t been finalized). It *still* hasn’t been deployed, and its future is still somewhat uncertain.

We were correct, though, about the most fundamental point we made: Buggy host software is a major security issue. In fact, we called it the “fundamental theorem of firewalls”:

Most hosts cannot meet our requirements: they run too many programs that are too large. Therefore, the only solution is to isolate them behind a firewall if you wish to run any programs at all.

If anything, we were too conservative.

Our Approach

This book is nearly a complete rewrite of the first edition. The approach is different, and so are many of the technical details. Most people don’t build their own firewalls anymore. There are far more Internet users, and the economic stakes are higher. The Internet is a factor in warfare.

The field of study is also much larger—there is too much to cover in a single book. One reviewer suggested that Chapters 2 and 3 could be a six-volume set. (They were originally one mammoth chapter.) Our goal, as always, is to teach an approach to security. We took far too long to write this edition, but one of the reasons why the first edition survived as long as it did was that we concentrated on the concepts, rather than details specific to a particular product at a particular time. The right frame of mind goes a long way toward understanding security issues and making reasonable security decisions. We’ve tried to include anecdotes, stories, and comments to make our points.

Some complain that our approach is too academic, or too UNIX-centric, that we are too idealistic, and don’t describe many of the most common computing tools. We are trying to teach

attitudes here more than specific bits and bytes. Most people have hideously poor computing habits and network hygiene. We try to use a safer world ourselves, and are trying to convey how we think it should be.

The chapter outline follows, but we want to emphasize the following:

It is OK to skip the hard parts.

If we dive into detail that is not useful to you, feel free to move on.

The introduction covers the overall philosophy of security, with a variety of time-tested maxims. As in the first edition, Chapter 2 discusses most of the important protocols, from a security point of view. We moved material about higher-layer protocols to Chapter 3. The Web merits a chapter of its own.

The next part discusses the threats we are dealing with: the kinds of attacks in Chapter 5, and some of the tools and techniques used to attack hosts and networks in Chapter 6.

Part III covers some of the tools and techniques we can use to make our networking world safer. We cover authentication tools in Chapter 7, and safer network servicing software in Chapter 8.

Part IV covers firewalls and *virtual private networks (VPNs)*. Chapter 9 introduces various types of firewalls and filtering techniques, and Chapter 10 summarizes some reasonable policies for filtering some of the more essential services discussed in Chapter 2. If you don't find advice about filtering a service you like, we probably think it is too dangerous (refer to Chapter 2).

Chapter 11 covers a lot of the deep details of firewalls, including their configuration, administration, and design. It is certainly not a complete discussion of the subject, but should give readers a good start. VPN tunnels, including holes through firewalls, are covered in some detail in Chapter 12. There is more detail in Chapter 18.

In Part V, we apply these tools and lessons to organizations. Chapter 13 examines the problems and practices on modern intranets. See Chapter 15 for information about deploying a hacking-resistant host, which is useful in any part of an intranet. Though we don't especially like *intrusion detection systems (IDSs)* very much, they do play a role in security, and are discussed in Chapter 15.

The last part offers a couple of stories and some further details. The Berferd chapter is largely unchanged, and we have added "The Taking of Clark," a real-life story about a minor break-in that taught useful lessons.

Chapter 18 discusses secure communications over insecure networks, in quite some detail. For even further detail, Appendix A has a short introduction to cryptography.

The conclusion offers some predictions by the authors, with justifications. If the predictions are wrong, perhaps the justifications will be instructive. (We don't have a great track record as prophets.) Appendix B provides a number of resources for keeping up in this rapidly changing field.

Errata and Updates

Everyone and every thing seems to have a Web site these days; this book is no exception. Our "official" Web site is <http://www.wilyhacker.com>. We'll post an errata list there; we'll

also keep an up-to-date list of other useful Web resources. If you find any errors—we hope there aren't many—please let us know via e-mail at firewall-book@wilyhacker.com.

Acknowledgments

For many kindnesses, we'd like to thank Joe Bigler, Steve “Hollywood” Branigan, Hal Burch, Brian Clapper, David Crocker, Tom Dow, Phil Edwards and the Internet Public Library, Anja Feldmann, Karen Gettman, Brian Kernighan, David Korman, Tom Limoncelli, Norma Loquendi, Cat Okita, Robert Oliver, Vern Paxson, Marcus Ranum, Eric Rescorla, Guido van Rooij, Luann Rouff (a most excellent copy editor), Abba Rubin, Peter Salus, Glenn Sieb, Karl Siil (we'll always have Boston), Irina Strizhevskaya, Rob Thomas, Win Treese, Dan Wallach, Avishai Wool, Karen Yannetta, and Michal Zalewski, among many others.

BILL CHESWICK
ches@cheswick.com

STEVE BELLOVIN
smb@stevebellovin.com

AVI RUBIN
avi@rubin.net