

Firewalls and Internet Security Second Edition

Repelling the Wily Hacker

William R. Cheswick
Steven M. Bellovin
Aviel D. Rubin



claim, in designations have been printed in initial capital letters or in all capitals.

The authors and publisher have taken care in the preparation of this book, but make no expressed or implied warranty of any kind and assume no responsibility for errors or omissions. No liability is assumed for incidental or consequential damages in connection with or arising out of the use of the information or programs contained herein.

The publisher offers discounts on this book when ordered in quantity for bulk purchases and special rates. For more information, please contact:

U.S. Corporate and Government Sales
(800) 382-3419
corpsales@pearsontechgroup.com

For sales outside of the U.S., please contact:

International Sales
(317) 581-3793
international@pearsontechgroup.com

Visit Addison-Wesley on the Web: www.awprofessional.com

Library of Congress Cataloging-in-Publication Data

Cheswick, William R.

Firewalls and Internet security : repelling the wily hacker / William R. Cheswick, Steven M. Bellovin and Aviel D. Rubin.— 2nd ed.
p. cm.

Includes bibliographical references and index.

ISBN 020163466X

1. Firewalls (Computer security) I. Bellovin, Steven M. II. Rubin, Aviel D. III. Title.

TK5105.875.I57C44 2003
005.8—dc21

2003000644

Copyright © 2003 by AT&T and Lumeta Corporation

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, transmitted, in any form, or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior consent of the publisher. Printed in the United States of America. Published simultaneously in Canada.

For information on obtaining permission for use of material from this work, please submit a request to:

Pearson Education, Inc.
Rights and Contracts Department
501 Boylston Street, Suite 900
Boston, MA 02116
Fax: (617) 671-3447

ISBN: 0-201-63466-X

This product is printed digitally on demand.
1 2 3 4 5 6 7 8 9 10—CRS—0706050403

First printing, February 2003

Trend Micro, Inc.
EX1011-2

things. They find them useful when trying to recruit people with particular skills. Nor is such information entirely benign at universities. Privacy considerations (and often legal strictures) dictate some care about what information can be released. Examples of this are the *Family Educational Rights and Privacy Act (FERPA)* and the EU Privacy Directives.

Another fruitful source of data is the DNS. We have already described the wealth of data that can be gathered from it, ranging from organizational details to target lists. Controlling the outflow is hard; often, the only solution is to limit the externally visible DNS to list gateway machines only.

Sophisticated hackers know this, of course, and don't take you at your word about what machines exist. They do port number and address space scans, looking for interesting services and hidden hosts. The best defense here is a good firewall; if they can't send packets to a machine, it's much less likely to be penetrated.

5.7 Exponential Attacks—Viruses and Worms

Exponential attacks use programs to spread themselves, multiplying their numbers quickly. When the programs travel by themselves, they are *worms*. When they attach to other programs, they are *viruses*. The mathematics of their spread is similar, and the distinction not that important. The epidemiology of such programs is quite similar to biological infectious agents.

These programs succeed by exploiting common bugs or behaviors found in a large population of susceptible programs or users. They can spread around the world within hours, and potentially in a few minutes [Staniford *et al.*, 2002; Rubin, 2001]. They can cause vast economic harm spread over a large community. The Melissa worm clogged the Microsoft-based e-mail in some companies for five days. Various worms have added substantial load to the entire Internet. (Nor is this threat new, or restricted to the Internet. The "IBM Christmas Card virus" clogged IBM's internal bisync network in 1987. See *RISKS Digest*, Vol. 5, Issue 81.)

These programs tend to infect "targets of opportunity," rather than specific individuals or organizations. But their payloads can and do attack popular political and commercial targets.

There are several ways to minimize the chance of getting a virus. By definition, the least popular way is to stay out of the popular monoculture. If you write your own operating system and applications, you are unlikely to be infectible. Microsoft Windows systems have traditionally hosted the vast majority of viruses, which means that Macintosh and UNIX users have suffered less. But this is changing, especially for Linux users. We are now seeing Linux worms, as well as cross-platform worms that can spread through several monocultures, and by direct network access as well as via Web pages and e-mail.

If you don't communicate with an affected host, you can't get the virus. Careful control of network access and the files obtained from foreign sources can greatly reduce the risk of infection. Note that there are also a number of human-propagated viruses, where people forward messages (often containing urban legends) to all of their friends, with instructions to send to all of their friends. These mostly serve as an annoyance. However, they can cause panic in individuals with less computer knowledge. Some contain incorrect messages that the recipient's computer has been infected. In one instance, this was accompanied by instructions to remove a crucial system file. Many people damaged their own computers by following these instructions.

Virus-scanning software is popular and quite effective against known viruses. The software must be updated constantly, as there is an arms race between virus writers and virus detection software companies. The viruses are becoming fantastically effective at hiding their presence and activities. Virus scanners can no longer be content looking for certain strings in the executable code: They have to emulate the code and look for viral behavior. As the viruses get more sophisticated, virus detection software will probably have to take more time examining each file, perhaps eventually taking too long. It is possible that virus writers may eventually be able to make code that cannot be identified in a reasonable amount of time.

Finally, it would be nice to execute only approved, unmodified programs. There are cryptographic technologies that can work here, but the entire approach is tied up with the political furor over copyright protection mechanisms and privacy.

5.8 Denial-of-Service Attacks

Hello! Hello! Are you there? Hello! I called you up to say hello. I said hello. Can you hear me, Joe?

Oh, no. I can not hear your call. I can not hear your call at all. This is not good and I know why. A mouse has cut the wire. Good-by!

One Fish, Two Fish, Red Fish, Blue Fish
—DR. SEUSS

We've discussed a wide variety of popular attacks on Internet hosts. These attacks rely on such things as protocol weaknesses, programming bugs in servers, and even inappropriately helpful humans. *Denial-of-Service (DOS)* attacks are a different beast. They are the simple overuse of a service—straining software, hardware, or network links beyond their intended capacity. The intent is to shut down or degrade the quality of a service, and that is generally a modest goal.

These attacks are different because they are obvious, not subtle. Shutting down a service should be easy to detect. Though the attack is usually easy to spot, the source of the attack may not be. They often involve generated packets with spoofed, random (and useless) return addresses.

Distributed Denial-of-Service (DDoS) attacks use many hosts on the Internet. More often than not, the participating hosts are unwitting accomplices to the attack, having been compromised in some way and outfitted with some malicious code. DDoS attacks are more difficult to recover from because the attacks come from all over. They are discussed further in Section 5.8.3.

39 There is no absolute remedy for a denial-of-service attack. As long as there is a public service, the public can abuse it. It is possible to make a perfectly secure site unavailable to the general public for a fair amount of time, and do this anonymously.

It is easy to compute a conservative value for the cost of a DOS attack. If a Web server is down for several days, a business should have a fairly good idea of what that cost them. If it doesn't, probably didn't have a good business plan for the Web service in the first place.

Companies may try to recover some of these losses through lawsuits, if a culprit can be located. This potential may force intermediate parties to take more care.