SECOND EDITION **NETWORK SECURITY** PRIVATE Communication in a PUBLIC World Now includes IPsec, SSL, PKI,

ES, and Web security

Charlie Kaufman - Radia Perlman - Mike Speciner

8888

NETWORK SECURITY PRIVATE Communication in a PUBLIC World



The Radia Perlman Series in Computer Networking and Security *Radia Perlman, Series Editor*

Dusseault	WebDAV: Next-Generation Collaborative Web Authoring
Hellberg, Boyes, and Greene	Broadband Network Architectures: Designing and Deploying Triple Play Services
Kaufman, Perlman, and Speciner	Network Security: Private Communication in a Public World, Second Edition
Liska	The Practice of Network Security: Deployment Strategies for Production Environments
Mancill	Linux Routers: A Primer for Network Administrators, Second Edition
Maufer	A Field Guide to Wireless LANs for Administrators and Power Users
Mirkovic, Dietrich, Dittrich, and Reiher	Internet Denial of Service: Attack and Defense Mechanisms
Skoudis with Liston	Counter Hack Reloaded: A Step-by-Step Guide to Computer Attacks and Effective Defenses, Second Edition
Skoudis with Zeltser	Malware: Fighting Malicious Code
Solomon	Mobile IP: The Internet Unplugged
Syme and Goldie	Optimizing Network Performance with Content Switching: Server, Firewall, and Cache Load Balancing
Tomsu and Schmutzer	Next Generation Optical Networks

NETWORK SECURITY PRIVATE Communication in a PUBLIC World

CHARLIE KAUFMAN • RADIA PERLMAN • MIKE SPECINER



Editorial/Production Supervision: Faye Gemmellaro Executive Editor: Mary Franz Editorial Assistant: Noreen Regina Marketing Manager: Dan DePasquale Manufacturing Buyer: Maura Zaldivar Cover Illustration: Tom Post Cover Design Director: Jerry Votta

> © 2002 by A division of Pearson Education, Inc. Upper Saddle River, New Jersey 07458

For information regarding corporate and government bulk discounts, please contact: Corporate and Government Sales, (800) 382-3419, or corpsales@pearsontechgroup.com

All rights reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the publisher. All trademarks mentioned herein are the properties of their respective owners.

Printed in the United States of America 10 9 8 7 6 5 4 3 2 Second Printing

ISBN 0-13-046019-2

Pearson Education Ltd. Pearson Education Australia PTY, Ltd. Pearson Education Singapore, Pte. Ltd. Pearson Education North Asia Ltd. Pearson Education Canada, Ltd. Pearson Educación de Mexico, S.A. de C.V. Pearson Education—Japan Pearson Education Malaysia, Pte. Ltd. Si spy net work, big fedjaw iog link kyxogy

This page intentionally left blank

CONTENTS

xxv

1.1 Roadmap to the Book 2 1.2 What Type of Book Is This? 3 1.3 Terminology 4 1.4 Notation 6 1.5 Primer on Networking 7 1.5.1 OSI Reference Model 7 1.5.2 IP, UDP, and TCP 8 1.5.3 Directory Service 9 1.5.4 Replicated Services 11 1.5.5 Packet Switching 11 1.5.6 Network Components 12 1.5.7 Destinations: Ultimate and Next-Hop 13 1.5.8 Address Structure 14 1.6 Active vs. Passive Attacks 15 1.7 Layers and Cryptography 15 1.8 Authorization 15 1.9 Tempest 16 1.10 Key Escrow for Law Enforcement 17 1.11 Key Escrow for Careless Users 19 1.12 Viruses, Worms, Trojan Horses 19 1.12.1 Where Do They Come From? 20 1.12.2 Spreading Pests from Machine to Machine	CHAPTER	1 Intr	oduction	1
1.2 What Type of Book Is This? 3 1.3 Terminology 4 1.4 Notation 6 1.5 Primer on Networking 7 1.5.1 OSI Reference Model 7 1.5.2 IP, UDP, and TCP 8 1.5.3 Directory Service 9 1.5.4 Replicated Services 11 1.5.5 Packet Switching 11 1.5.6 Network Components 12 1.5.7 Destinations: Ultimate and Next-Hop 13 1.5.8 Address Structure 14 1.6 Active vs. Passive Attacks 15 1.7 Layers and Cryptography 15 1.8 Authorization 15 1.9 Tempest 16 1.10 Key Escrow for Law Enforcement 17 1.11 Key Escrow for Careless Users 19 1.12.1 Where Do They Come From? 20 1.12.2 Spreading Pests from Machine to Machine 23 1.12.3 Virus Checkers 24 1.12.4 What Can We Do Today? <	1.1	Roadma	p to the Book	2
1.3 Terminology 4 1.4 Notation 6 1.5 Primer on Networking 7 1.5.1 OSI Reference Model 7 1.5.2 IP, UDP, and TCP 8 1.5.3 Directory Service 9 1.5.4 Replicated Services 11 1.5.5 Packet Switching 11 1.5.6 Network Components 12 1.5.7 Destinations: Ultimate and Next-Hop 13 1.5.8 Address Structure 14 1.6 Active vs. Passive Attacks 15 1.7 Layers and Cryptography 15 1.8 Authorization 15 1.9 Tempest 16 1.0 Key Escrow for Law Enforcement 17 1.11 Key Escrow for Careless Users 19 1.12 Viruses, Worms, Trojan Horses 19 1.12.1 Where Do They Come From? 20 1.12.2 Spreading Pests from Machine to Machine 23 1.12.3 Virus Checkers 24 1.12.4 What Can We Do Today?	1.2	What Ty	pe of Book Is This?	3
1.4 Notation	1.3	Termino		4
1.5 Primer on Networking 7 1.5.1 OSI Reference Model 7 1.5.2 IP, UDP, and TCP 8 1.5.3 Directory Service 9 1.5.4 Replicated Services 11 1.5.5 Packet Switching 11 1.5.6 Network Components 12 1.5.7 Destinations: Ultimate and Next-Hop 13 1.5.8 Address Structure 14 1.6 Active vs. Passive Attacks 15 1.7 Layers and Cryptography 15 1.8 Authorization 15 1.9 Tempest 16 1.10 Key Escrow for Law Enforcement 17 1.11 Key Escrow for Careless Users 19 1.12 Viruses, Worms, Trojan Horses 19 1.12.1 Where Do They Come From? 20 1.12.2 Spreading Pests from Machine to Machine 23 1.12.1 Wirus Checkers 24 1.12.3 Virus Checkers 24 1.12.4 What Can We Do Today? 25 1.12.5 Wish	1.4	Notation	 L	6
1.5.1 OSI Reference Model 7 1.5.2 IP, UDP, and TCP 8 1.5.3 Directory Service. 9 1.5.4 Replicated Services 11 1.5.5 Packet Switching 11 1.5.6 Network Components 12 1.5.7 Destinations: Ultimate and Next-Hop 13 1.5.8 Address Structure 14 1.6 Active vs. Passive Attacks 15 1.7 Layers and Cryptography 15 1.8 Authorization 15 1.9 Tempest 16 1.10 Key Escrow for Law Enforcement 17 1.11 Key Escrow for Careless Users 19 1.12 Viruses, Worms, Trojan Horses 19 1.12.1 Where Do They Come From? 20 1.12.2 Spreading Pests from Machine to Machine 23 1.12.3 Viruse Checkers 24 1.12.4 What Can We Do Today? 25 1.12.5 Wish List for the Future 26 1.13 Mandatory (Nondiscretionary) Access Controls 28	1.5	Primer of	n Networking	7
1.5.2 IP, UDP, and TCP 8 1.5.3 Directory Service. 9 1.5.4 Replicated Services 11 1.5.5 Packet Switching 11 1.5.6 Network Components 12 1.5.7 Destinations: Ultimate and Next-Hop 13 1.5.8 Address Structure 14 1.6 Active vs. Passive Attacks 15 1.7 Layers and Cryptography 15 1.8 Authorization 15 1.9 Tempest 16 1.10 Key Escrow for Law Enforcement 17 1.11 Key Escrow for Careless Users 19 1.12.1 Where Do They Come From? 20 1.12.2 Spreading Pests from Machine to Machine 23 1.12.1 What Can We Do Today? 25 1.12.5 Wish List for the Future 26 1.13 Mandatory (Nondiscretionary) Access Controls 28 1.13.1 Mandatory Access Control Rules 29 1.13.4 Covert Channels 30 1.13.5 The Orange Book 32 <tr< td=""><th></th><td>1.5.1</td><td>OSI Reference Model</td><td>7</td></tr<>		1.5.1	OSI Reference Model	7
1.5.3Directory Service91.5.4Replicated Services111.5.5Packet Switching111.5.6Network Components121.5.7Destinations: Ultimate and Next-Hop131.5.8Address Structure141.6Active vs. Passive Attacks151.7Layers and Cryptography151.8Authorization151.9Tempest161.10Key Escrow for Law Enforcement171.11Key Escrow for Careless Users191.12.1Where Do They Come From?201.12.2Spreading Pests from Machine to Machine231.12.3Virus Checkers241.12.4What Can We Do Today?251.12.5Wish List for the Future261.13The Multi-level Model of Security271.13.1Mandatory (Nondiscretionary) Access Controls281.13.2Levels of Security291.13.3Mandatory Access Control Rules291.13.4Covert Channels301.13.5The Orange Book321.13.6Successors to the Orange Book35		1.5.2	IP, UDP, and TCP	8
1.5.4 Replicated Services 11 1.5.5 Packet Switching 11 1.5.6 Network Components 12 1.5.7 Destinations: Ultimate and Next-Hop 13 1.5.8 Address Structure 14 1.6 Active vs. Passive Attacks 15 1.7 Layers and Cryptography 15 1.8 Authorization 16 1.10 Key Escrow for Law Enforcement 17 1.11 Key Escrow for Careless Users 19 1.12.1 Where Do They Come From? 20 1.12.2 Spreading Pests from Machine to Machine 23 1.12.3 Virus Checkers 24 1.12.4 What Can We Do Today? 25 1.12.5 Wish List for the Future 26 1.13 The Multi-level Model of Security 27 1.13.1 Mandatory (Nondiscretionary) Access Controls 28 1.13.2 Levels of Security 29 1.13.3 Mandatory Access Control Rules 29 1.13.4 Covert Channels 30 1.13.5 The Orange Book		1.5.3	Directory Service	9
1.5.5 Packet Switching 11 1.5.6 Network Components 12 1.5.7 Destinations: Ultimate and Next-Hop 13 1.5.8 Address Structure 14 1.6 Active vs. Passive Attacks 15 1.7 Layers and Cryptography 15 1.8 Authorization 15 1.9 Tempest 16 1.10 Key Escrow for Law Enforcement 17 1.11 Key Escrow for Careless Users 19 1.12.1 Where Do They Come From? 20 1.12.2 Spreading Pests from Machine to Machine 23 1.12.3 Virus Checkers 24 1.12.4 What Can We Do Today? 25 1.12.5 Wish List for the Future 26 1.13 The Multi-level Model of Security 27 1.13.1 Mandatory (Nondiscretionary) Access Controls 28 1.13.2 Levels of Security 29 1.13.4 Covert Channels 30 1.13.5 The Orange Book 32 1.13.6 Successors to the Orange Book 35		1.5.4	Replicated Services	11
1.5.6Network Components121.5.7Destinations: Ultimate and Next-Hop131.5.8Address Structure141.6Active vs. Passive Attacks151.7Layers and Cryptography151.8Authorization151.9Tempest161.10Key Escrow for Law Enforcement171.11Key Escrow for Careless Users191.12.1Where Do They Come From?201.12.2Spreading Pests from Machine to Machine231.12.3Virus Checkers241.12.4What Can We Do Today?251.12.5Wish List for the Future261.13The Multi-level Model of Security271.13.1Mandatory (Nondiscretionary) Access Controls281.13.2Levels of Security291.13.3Mandatory Access Control Rules291.13.4Covert Channels301.13.5The Orange Book321.13.6Successors to the Orange Book35		1.5.5	Packet Switching	11
1.5.7Destinations: Ultimate and Next-Hop131.5.8Address Structure141.6Active vs. Passive Attacks151.7Layers and Cryptography151.8Authorization151.9Tempest161.10Key Escrow for Law Enforcement171.11Key Escrow for Careless Users191.12Viruses, Worms, Trojan Horses191.12.1Where Do They Come From?201.12.2Spreading Pests from Machine to Machine231.12.3Virus Checkers241.12.4What Can We Do Today?251.12.5Wish List for the Future261.13The Multi-level Model of Security271.13.1Mandatory (Nondiscretionary) Access Controls281.13.2Levels of Security291.13.3Mandatory Access Control Rules291.13.4Covert Channels301.13.5The Orange Book321.13.6Successors to the Orange Book35		1.5.6	Network Components	12
1.5.8Address Structure141.6Active vs. Passive Attacks151.7Layers and Cryptography151.8Authorization151.9Tempest161.10Key Escrow for Law Enforcement171.11Key Escrow for Careless Users191.12Viruses, Worms, Trojan Horses191.12.1Where Do They Come From?201.12.2Spreading Pests from Machine to Machine231.12.3Virus Checkers241.12.4What Can We Do Today?251.12.5Wish List for the Future261.13The Multi-level Model of Security271.13.1Mandatory (Nondiscretionary) Access Controls281.13.2Levels of Security291.13.3Mandatory Access Control Rules291.13.4Covert Channels301.13.5The Orange Book321.13.6Successors to the Orange Book35		1.5.7	Destinations: Ultimate and Next-Hop	13
1.6Active vs. Passive Attacks		1.5.8	Address Structure	14
1.7Layers and Cryptography151.8Authorization151.9Tempest161.10Key Escrow for Law Enforcement171.11Key Escrow for Careless Users191.12Viruses, Worms, Trojan Horses191.12.1Where Do They Come From?201.12.2Spreading Pests from Machine to Machine231.12.3Virus Checkers241.12.4What Can We Do Today?251.12.5Wish List for the Future261.13The Multi-level Model of Security271.13.1Mandatory (Nondiscretionary) Access Controls281.13.2Levels of Security291.13.3Mandatory Access Control Rules291.13.4Covert Channels301.13.5The Orange Book321.13.6Successors to the Orange Book35	1.6	Active v	s. Passive Attacks	15
1.8Authorization151.9Tempest161.10Key Escrow for Law Enforcement171.11Key Escrow for Careless Users191.12Viruses, Worms, Trojan Horses191.12.1Where Do They Come From?201.12.2Spreading Pests from Machine to Machine231.12.3Virus Checkers241.12.4What Can We Do Today?251.12.5Wish List for the Future261.13The Multi-level Model of Security271.13.1Mandatory (Nondiscretionary) Access Controls281.13.2Levels of Security291.13.3Mandatory Access Control Rules291.13.4Covert Channels301.13.5The Orange Book321.13.6Successors to the Orange Book35	1.7	Layers a	nd Cryptography	15
1.9Tempest161.10Key Escrow for Law Enforcement171.11Key Escrow for Careless Users191.12Viruses, Worms, Trojan Horses191.12Viruses, Worms, Trojan Horses191.12.1Where Do They Come From?201.12.2Spreading Pests from Machine to Machine231.12.3Virus Checkers241.12.4What Can We Do Today?251.12.5Wish List for the Future261.13The Multi-level Model of Security271.13.1Mandatory (Nondiscretionary) Access Controls281.13.2Levels of Security291.13.3Mandatory Access Control Rules291.13.4Covert Channels301.13.5The Orange Book321.13.6Successors to the Orange Book35	1.8	Authoriz	zation	15
1.10Key Escrow for Law Enforcement171.11Key Escrow for Careless Users191.12Viruses, Worms, Trojan Horses191.12.1Where Do They Come From?201.12.2Spreading Pests from Machine to Machine231.12.3Virus Checkers241.12.4What Can We Do Today?251.12.5Wish List for the Future261.13The Multi-level Model of Security271.13.1Mandatory (Nondiscretionary) Access Controls281.13.2Levels of Security291.13.3Mandatory Access Control Rules291.13.4Covert Channels301.13.5The Orange Book321.13.6Successors to the Orange Book35	1.9	Tempest		16
1.11Key Escrow for Careless Users191.12Viruses, Worms, Trojan Horses191.12.1Where Do They Come From?201.12.2Spreading Pests from Machine to Machine231.12.3Virus Checkers241.12.4What Can We Do Today?251.12.5Wish List for the Future261.13The Multi-level Model of Security271.13.1Mandatory (Nondiscretionary) Access Controls281.13.2Levels of Security291.13.3Mandatory Access Control Rules291.13.4Covert Channels301.13.5The Orange Book321.13.6Successors to the Orange Book35	1.10	Key Esc	row for Law Enforcement	17
1.12Viruses, Worms, Trojan Horses191.12.1Where Do They Come From?201.12.2Spreading Pests from Machine to Machine231.12.3Virus Checkers241.12.4What Can We Do Today?251.12.5Wish List for the Future261.13The Multi-level Model of Security271.13.1Mandatory (Nondiscretionary) Access Controls281.13.2Levels of Security291.13.3Mandatory Access Control Rules291.13.4Covert Channels301.13.5The Orange Book321.13.6Successors to the Orange Book35	1.11	Key Esc	row for Careless Users	19
1.12.1Where Do They Come From?201.12.2Spreading Pests from Machine to Machine231.12.3Virus Checkers241.12.4What Can We Do Today?251.12.5Wish List for the Future261.13The Multi-level Model of Security271.13.1Mandatory (Nondiscretionary) Access Controls281.13.2Levels of Security291.13.3Mandatory Access Control Rules291.13.4Covert Channels301.13.5The Orange Book321.13.6Successors to the Orange Book35	1.12	Viruses,	Worms, Trojan Horses	19
1.12.2Spreading Pests from Machine to Machine231.12.3Virus Checkers241.12.4What Can We Do Today?251.12.5Wish List for the Future261.13The Multi-level Model of Security271.13.1Mandatory (Nondiscretionary) Access Controls281.13.2Levels of Security291.13.3Mandatory Access Control Rules291.13.4Covert Channels301.13.5The Orange Book321.13.6Successors to the Orange Book35		1.12.1	Where Do They Come From?	20
1.12.3Virus Checkers241.12.4What Can We Do Today?251.12.5Wish List for the Future261.13The Multi-level Model of Security271.13.1Mandatory (Nondiscretionary) Access Controls281.13.2Levels of Security291.13.3Mandatory Access Control Rules291.13.4Covert Channels301.13.5The Orange Book321.13.6Successors to the Orange Book35		1.12.2	Spreading Pests from Machine to Machine	23
1.12.4What Can We Do Today?251.12.5Wish List for the Future261.13The Multi-level Model of Security271.13.1Mandatory (Nondiscretionary) Access Controls281.13.2Levels of Security291.13.3Mandatory Access Control Rules291.13.4Covert Channels301.13.5The Orange Book321.13.6Successors to the Orange Book35		1.12.3	Virus Checkers	24
1.12.5Wish List for the Future261.13The Multi-level Model of Security271.13.1Mandatory (Nondiscretionary) Access Controls281.13.2Levels of Security291.13.3Mandatory Access Control Rules291.13.4Covert Channels301.13.5The Orange Book321.13.6Successors to the Orange Book35		1.12.4	What Can We Do Today?	25
1.13The Multi-level Model of Security		1.12.5	Wish List for the Future	26
1.13.1Mandatory (Nondiscretionary) Access Controls281.13.2Levels of Security291.13.3Mandatory Access Control Rules291.13.4Covert Channels301.13.5The Orange Book321.13.6Successors to the Orange Book35	1.13	The Mul	ti-level Model of Security	27
1.13.2Levels of Security		1.13.1	Mandatory (Nondiscretionary) Access Controls	28
1.13.3Mandatory Access Control Rules291.13.4Covert Channels301.13.5The Orange Book321.13.6Successors to the Orange Book35		1.13.2	Levels of Security	29
1.13.4Covert Channels301.13.5The Orange Book321.13.6Successors to the Orange Book35		1.13.3	Mandatory Access Control Rules	29
1.13.5The Orange Book321.13.6Successors to the Orange Book35		1.13.4	Covert Channels	30
1.13.6 Successors to the Orange Book 35		1.13.5	The Orange Book	32
		1.13.6	Successors to the Orange Book	35

1.14	Legal Iss	ues	36
	1.14.1	Patents	36
	1.14.2	Export Controls	37
		CRYPTOGRAPHY	
CHAPTER	2 Intro	oduction to Cryptography	41
2.1	What Is (Cryptography?	41
	2.1.1	Computational Difficulty	42
	2.1.2	To Publish or Not to Publish	43
	2.1.3	Secret Codes	44
2.2	Breaking	an Encryption Scheme	45
	2.2.1	Ciphertext Only	45
	2.2.2	Known Plaintext	46
	2.2.3	Chosen Plaintext	46
2.3	Types of	Cryptographic Functions	47
2.4	Secret K	ey Cryptography	47
	2.4.1	Security Uses of Secret Key Cryptography	47
	2.4.2	Transmitting Over an Insecure Channel	48
	2.4.3	Secure Storage on Insecure Media	48
	2.4.4	Authentication	48
	2.4.5	Integrity Check	49
2.5	Public K	ey Cryptography	50
	2.5.1	Security Uses of Public Key Cryptography	52
	2.5.2	Transmitting Over an Insecure Channel	52
	2.5.3	Secure Storage on Insecure Media	52
	2.5.4	Authentication	53
	2.5.5	Digital Signatures	54
2.6	Hash Alg	gorithms	54
	2.6.1	Password Hashing	55
	2.6.2	Message Integrity	56
	2.6.3	Message Fingerprint	56
	2.6.4	Downline Load Security	57
	2.6.5	Digital Signature Efficiency	57
2.7	Homewo	rk	57
CHAPTER	3 Sec	ret Key Cryptography	59
3.1	Introduct	ion	59
3.2	Generic	Block Encryption	59
3.3	Data Enc	ryption Standard (DES)	62
	3.3.1	DES Overview	64
	3.3.2	The Permutations of the Data	66
	3.3.3	Generating the Per-Round Keys	67

	3.3.4	A DES Round	69
	3.3.5	The Mangler Function	70
	3.3.6	Weak and Semi-Weak Keys	74
	3.3.7	What's So Special About DES?	74
3.4	Internat	ional Data Encryption Algorithm (IDEA)	75
	3.4.1	Primitive Operations	75
	3.4.2	Key Expansion	77
	3.4.3	One Round	78
		3.4.3.1 Odd Round	78
		3.4.3.2 Even Round	79
	3.4.4	Inverse Keys for Decryption	80
	3.4.5	Does IDEA Work?	81
3.5	Advanc	ed Encryption Standard (AES)	81
	3.5.1	Basic Structure	
	3.5.2	Primitive Operations	84
		3.5.2.1 What about the inverse cipher?	87
	3.5.3	Key Expansion	89
	3.5.4	Rounds	90
	3.5.5	Inverse Rounds	91
	3.5.6	Optimization	91
3.6	RC4		92
3.7	Homew	70rk	92
CHAPTER	4 Mo	odes of Operation	95
4.1	Introduc	ction	
4.2	Encrypt	ting a Large Message	
	4.2.1	Electronic Code Book (ECB)	
	4.2.2	Cipher Block Chaining (CBC)	
		4.2.2.1 CBC Threat 1—Modifying Ciphertext Blocks	
		4.2.2.2 CBC Threat 2—Rearranging Ciphertext Blocks	
	4.2.3	Output Feedback Mode (OFB)	
	4.2.4	Cipher Feedback Mode (CFB)	
	4.2.5	Counter Mode (CTR)	104
4.3	Generat	ting MACs	
	4.3.1	Ensuring Privacy and Integrity Together	
	4.3.2	CBC with a Weak Cryptographic Checksum	107
	4.3.3	CBC Encryption and CBC Residue with Related Keys	108
	4.3.4	CBC with a Cryptographic Hash	108
	4.3.5	Offset Codebook Mode (OCB)	108
4.4	Multiple	e Encryption DES	
	manupr		
	4.4.1	How Many Encryptions?	

	4	.4.1.2	Encrypting Twice with Two Keys	
	4	.4.1.3	Triple Encryption with only Two Keys	
	4.4.2 C	CBC Out	side vs. Inside	
4.5	Homework			
CHAPTEF	85 Hashe	es and	Message Digests	117
5.1	Introductio	n		
5.2	Nifty Thing	gs to Do	with a Hash	
	5.2.1 A	Authenti	cation	
	5.2.2 C	Computi	ng a MAC with a Hash	
	5.2.3 E	Encryptio	on with a Message Digest	
	5	.2.3.1	Generating a One-Time Pad	
	5	.2.3.2	Mixing In the Plaintext	
	5.2.4 U	Jsing Se	cret Key for a Hash	
	5	.2.4.1	UNIX Password Hash	
	5	.2.4.2	Hashing Large Messages	
5.3	MD2			
	5.3.1 N	AD2 Pac	lding	
	5.3.2 N	AD2 Ch	ecksum Computation	
	5.3.3 N	AD2 Fin	al Pass	
5.4	MD4			
	5.4.1 N	AD4 Me	ssage Padding	
	5.4.2 C	Overviev	v of MD4 Message Digest Computation	
	5.4.3 N	AD4 Me	ssage Digest Pass 1	
	5.4.4 N	AD4 Me	ssage Digest Pass 2	
	5.4.5 N	AD4 Me	ssage Digest Pass 3	
5.5	MD5			
	5.5.1 N	AD5 Me	ssage Padding	
	5.5.2 C	Overviev	v of MD5 Message Digest Computation	
	5.5.3 N	AD5 Me	ssage Digest Pass 1	
	5.5.4 N	AD5 Me	ssage Digest Pass 2	
	5.5.5 N	AD5 Me	ssage Digest Pass 3	
	5.5.6 N	AD5 Me	ssage Digest Pass 4	
5.6	SHA-1			
	5.6.1 S	HA-1 M	lessage Padding	
	5.6.2 C	Overviev	v of SHA-1 Message Digest Computation	
	5.6.3 S	HA-1 C	peration on a 512-bit Block	
5.7	HMAC		-	
5.8	Homework			
CHAPTER	6 Public	c Key A	llgorithms	147
6.1	Introductio	n		

6.2	Modula	ar Arithmeti	ic	148
	6.2.1	Modular	Addition	
	6.2.2	Modular	Multiplication	149
	6.2.3	Modular	Exponentiation	151
6.3	RSA		-	152
	6.3.1	RSA Alg	gorithm	152
	6.3.2	Why Do	es RSA Work?	153
	6.3.3	Why Is F	RSA Secure?	153
	6.3.4	How Eff	icient Are the RSA Operations?	154
		6.3.4.1	Exponentiating with Big Numbers	154
		6.3.4.2	Generating RSA Keys	156
			Finding Big Primes p and q	156
			Finding d and e	158
		6.3.4.3	Having a Small Constant e	158
		6.3.4.4	Optimizing RSA Private Key Operations	160
	6.3.5	Arcane F	RSA Threats	161
		6.3.5.1	Smooth Numbers	161
		6.3.5.2	The Cube Root Problem	162
	6.3.6	Public-K	Ley Cryptography Standard (PKCS)	163
		6.3.6.1	Encryption	163
		6.3.6.2	Encryption—Take 2	164
		6.3.6.3	Signing	165
6.4	Diffie-l	Hellman		166
	6.4.1	The Buc	ket Brigade/Man-in-the-Middle Attack	167
	6.4.2	Defenses	s Against Man-in-the-Middle Attack	169
		6.4.2.1	Published Diffie-Hellman Numbers	169
		6.4.2.2	Authenticated Diffie-Hellman	169
	6.4.3	Encrypti	on with Diffie-Hellman	170
	6.4.4	ElGamal	Signatures	170
	6.4.5	Diffie-H	ellman Details—Safe Primes	171
6.5	Digital	Signature S	Standard (DSS)	172
	6.5.1	The DSS	S Algorithm	172
	6.5.2	Why Do	es the Verification Procedure Work?	174
	6.5.3	Why Is 7	This Secure?	174
	6.5.4	The DSS	Controversy	175
	6.5.5	Per-Mess	sage Secret Number	176
6.6	How Se	ecure Are R	SA and Diffie-Hellman?	177
6.7	Elliptic	Curve Cry	ptography (ECC)	178
6.8	Zero K	nowledge P	Proof Systems	179
	6.8.1	Zero Kno	owledge Signatures	181
6.9	Homew	ork Problem	ms	

9.7.2

CHAPTE	R 7 Number Theory	185
7.1	Introduction	
7.2	Modular Arithmetic	
7.3	Primes	
7.4	Euclid's Algorithm	
	7.4.1 Finding Multiplicative Inverses in Modular Arithmetic	
7.5	Chinese Remainder Theorem	
7.6	Z _n *	
7.7	Euler's Totient Function	
7.8	Euler's Theorem	194
	7.8.1 A Generalization of Euler's Theorem	195
7.9	Homework Problems	195
CHAPTE	R 8 Math with AES and Elliptic Curves	197
8.1	Introduction	
8.2	Notation	
8.3	Groups	
8.4	Fields	
	8.4.1 Polynomials	
	8.4.2 Finite Fields	
	8.4.2.1 What Sizes Can Finite Fields Be?	
	8.4.2.2 Representing a Field	
8.5	Mathematics of Rijndael	
	8.5.1 A Rijndael Round	
8.6	Elliptic Curve Cryptography	
8.7	Homework	
	AUTHENTICATION	
CHAPTE	R 9 Overview of Authentication Systems	215
9.1	Password-Based Authentication	
	9.1.1 Off- vs. On-Line Password Guessing	
	9.1.2 Storing User Passwords	
9.2	Address-Based Authentication	
	9.2.1 Network Address Impersonation	
9.3	Cryptographic Authentication Protocols	
9.4	Who Is Being Authenticated?	
9.5	Passwords as Cryptographic Keys	
9.6	Eavesdropping and Server Database Reading	
9.7	Trusted Intermediaries	
	9.7.1 KDCs	

	9.7.3	Certificate Revocation	
	9.7.4	Multiple Trusted Intermediaries	
		9.7.4.1 Multiple KDC Domains	
		9.7.4.2 Multiple CA Domains	
9.8	Session	Key Establishment	
9.9	Delegati	ion	
9.10	Homew	ork	
CHAPTER	R 10 Au	thentication of People	237
10.1	Passwor	rds	
10.2	On-Line	Password Guessing	
10.3	Off-Line	e Password Guessing	
10.4	How Bi	g Should a Secret Be?	
10.5	Eavesdr	opping	
10.6	Passwor	ds and Careless Users	
	10.6.1	Using a Password in Multiple Places	
	10.6.2	Requiring Frequent Password Changes	
	10.6.3	A Login Trojan Horse to Capture Passwords	
	10.6.4	Non-Login Use of Passwords	
10.7	Initial P	assword Distribution	
10.8	Authent	ication Tokens	250
10.9	Physical	Access	
10.10) Biometr	ics	253
10.11	l Homew	ork	
CHAPTER	R 11 See	curity Handshake Pitfalls	257
11.1	Login O	nly	
	11.1.1	Shared Secret	
	11.1.2	One-Way Public Key	
11.2	Mutual	Authentication	
	11.2.1	Reflection Attack	
	11.2.2	Password Guessing	
	11.2.3	Public Keys	
	11.2.4	Timestamps	
11.3	Integrity	/Encryption for Data	
	11.3.1	Shared Secret	
	11.3.2	Two-Way Public Key Based Authentication	
	11.3.3	One-Way Public Key Based Authentication	
	11.3.4	Privacy and Integrity	
11.4	Mediate	d Authentication (with KDC)	
	11.4.1	Needham-Schroeder	275
	11.4.2	Expanded Needham-Schroeder	

	11.4.3 Otway-Rees	
11.5	Nonce Types	
11.6	Picking Random Numbers	
11.7	Performance Considerations	
11.8	Authentication Protocol Checklist	
11.9	Homework	
CHAPTEF	12 Strong Password Protocols	291
12.1	Introduction	
12.2	Lamport's Hash	
12.3	Strong Password Protocols	
	12.3.1 The Basic Form	
	12.3.2 Subtle Details	
	12.3.3 Augmented Strong Password Protocols	
	12.3.4 SRP (Secure Remote Password)	
12.4	Strong Password Credentials Download Protocols	
12.5	Homework	

STANDARDS

307

CHAPTER 13 Kerberos V4

13.1	Introduct	tion		
13.2	Tickets and Ticket-Granting Tickets			
13.3	Configuration			
13.4	Logging Into the Network			
	13.4.1	Obtaining a Session Key and TGT		
	13.4.2	Alice Asks to Talk to a Remote Node		
13.5	Replicate	ed KDCs		
13.6	Realms.			
13.7	Interreal	m Authentication		
13.8	Key Ver	sion Numbers		
13.9	Encryption for Privacy and Integrity			
13.10	Encryption for Integrity Only			
13.11	Network Layer Addresses in Tickets			
13.12	Message Formats			
	13.12.1	Tickets		
	13.12.2	Authenticators		
	13.12.3	Credentials		
	13.12.4	AS_REQ		
	13.12.5	TGS_REQ		
	13.12.6	AS_REP and TGS_REP		
	13.12.7	Error Reply from KDC		
	13.12.8	AP REQ.		

	13.12.9 AP_REP	
	13.12.10 Encrypted Data (KRB_PRV)	
	13.12.11 Integrity-Checked Data (SAFE)	
	13.12.12 AP_ERR	
13.13	Homework	
CHAPTER	14 Kerberos V5	337
14.1	ASN.1	
14.2	Names	
14.3	Delegation of Rights	
14.4	Ticket Lifetimes	
	14.4.1 Renewable Tickets	
	14.4.2 Postdated Tickets	
14.5	Key Versions	
14.6	Making Master Keys in Different Realms Different	
14.7	Optimizations	
14.8	Cryptographic Algorithms	
	14.8.1 Integrity-Only Algorithms	
	14.8.1.1 rsa-md5-des	
	14.8.1.2 des-mac	
	14.8.1.3 des-mac-k	
	14.8.1.4 rsa-md4-des	
	14.8.1.5 rsa-md4-des-k	
	14.8.2 Encryption for Privacy and Integrity	
14.9	Hierarchy of Realms	
14.10	Evading Password-Guessing Attacks	
14.11	Key Inside Authenticator	
14.12	Double TGT Authentication	
14.13	PKINII—Public Keys for Users	
14.14	KDC Database	
14.15	Kerberos V 5 Messages	
	14.15.1 Authenticator	
	14.15.2 IICKet	
	14.15.5 AS_KEQ	
	14.15.4 IOS_KEQ	
	14.15.5 AS_KEP	
	14.13.0 105_NEF	
	14.13.7 ALCU	
	14.15.0 AF_KEF	
	14.15.7 KKD_SATE	
	14.15.11 KDR CPED	
	14.1J.11 KKD_CKED	

	14.15.12	2 KRB_ERROR	
14.16	Homewo	ork	
CHAPTER	15 PK	l (Public Key Infrastructure)	371
15.1	Introduc	tion	
15.2	Some Te	erminology	
15.3	PKI Tru	st Models	
	15.3.1	Monopoly Model	
	15.3.2	Monopoly plus Registration Authorities (RAs)	
	15.3.3	Delegated CAs	
	15.3.4	Oligarchy	
	15.3.5	Anarchy Model	
	15.3.6	Name Constraints	
	15.3.7	Top-Down with Name Constraints	
	15.3.8	Bottom-Up with Name Constraints	
	15.3.9	Relative Names	
	15.3.10	Name Constraints in Certificates	
	15.3.11	Policies in Certificates	
15.4	Revocat	ion	
	15.4.1	Revocation Mechanisms	
		15.4.1.1 Delta CRLs	
		15.4.1.2 First Valid Certificate	
	15.4.2	OLRS Schemes	
	15.4.3	Good-lists vs. Bad-lists	
15.5	Director	ies and PKI	
	15.5.1	Store Certificates with Subject or Issuer?	
	15.5.2	Finding Certificate Chains	
15.6	PKIX an	nd X.509	
	15.6.1	Names	
	15.6.2	OIDs	
	15.6.3	Specification of Time	
15.7	X.509 ar	nd PKIX Certificates	
	15.7.1	X.509 and PKIX CRLs	
15.8	Authoriz	zation Futures	
	15.8.1	ACL (Access Control List)	
	15.8.2	Central Administration/Capabilities	
	15.8.3	Groups	
		15.8.3.1 Cross-Organizational and Nested Groups	
	15.8.4	Roles	
	15.8.5	Anonymous Groups	
15.9	Homewo	ork	401

CHAPTER	16 Real-time Communication Security	403
16.1	What Layer?	
16.2	Session Key Establishment	
16.3	Perfect Forward Secrecy	
16.4	PFS-Foilage	
16.5	Denial-of-Service/Clogging Protection	410
	16.5.1 Cookies	410
	16.5.2 Puzzles	411
16.6	Endpoint Identifier Hiding	
16.7	Live Partner Reassurance	
16.8	Arranging for Parallel Computation	415
16.9	Session Resumption	416
16.10	Plausible Deniability	416
16.11	Data Stream Protection	
16.12	Negotiating Crypto Parameters	
16.13	Easy Homework	
16.14	Homework	
CHAPTER	17 IPsec: AH and ESP	423
17.1	Overview of IPsec	
	17.1.1 Security Associations	
	17.1.2 Security Association Database	
	17.1.3 Security Policy Database	
	17.1.4 AH and ESP	
	17.1.5 Tunnel, Transport Mode	
	17.1.6 Why Protect the IP Header?	
17.2	IP and IPv6	
	17.2.1 NAT (Network Address Translation)	
	17.2.2 Firewalls	
	17.2.3 IPv4 Header	
	17.2.4 IPv6 Header	
17.3	AH (Authentication Header)	
	17.3.1 Mutable, Immutable	
	17.3.2 Mutable but Predictable	
17.4	ESP (Encapsulating Security Payload)	
17.5	So, Do We Need AH?	
17.6	Comparison of Encodings	
17.7	Easy Homework	
17.8	Homework	
CHAPTER	18 IPsec: IKE	441
18.1	Photuris	

18.2	SKIP		443
18.3	History of	of IKE	444
18.4	IKE Pha	ses	445
18.5	Phase 1	IKE	446
	18.5.1	Aggressive Mode and Main Mode	446
	18.5.2	Key Types	448
	18.5.3	Proof of Identity	449
	18.5.4	Cookie Issues	450
	18.5.5	Negotiating Cryptographic Parameters	451
	18.5.6	Session Keys	
	18.5.7	Message IDs	454
	18.5.8	Phase 2/Quick Mode	454
	18.5.9	Traffic Selectors	454
	18.5.10	The IKE Phase 1 Protocols	455
		18.5.10.1 Public Signature Keys, Main Mode	455
		18.5.10.2 Public Signature Keys, Aggressive Mode	456
		18.5.10.3 Public Encryption Key, Main Mode, Original	457
		18.5.10.4 Public Encryption Key, Aggressive Mode, Original	458
		18.5.10.5 Public Encryption Key, Main Mode, Revised	458
		18.5.10.6 Public Encryption Key, Aggressive Mode, Revised	459
		18.5.10.7 Shared Secret Key, Main Mode	459
		18.5.10.8 Shared Secret Key, Aggressive Mode	460
18.6	Phase-2	IKE: Setting up IPsec SAs	
18.7	ISAKMI	P/IKE Encoding	463
	18.7.1	Fixed Header	465
	18.7.2	Payload Portion of ISAKMP Messages	467
	18.7.3	SA Payload	467
		18.7.3.1 Ps and Ts within the SA Payload	468
		18.7.3.2 Payload Length in SA, P, and T Payloads	468
		18.7.3.3 Type of Next Payload	468
		18.7.3.4 SA Payload Fields	469
	18.7.4	P Payload	
	18.7.5	T Payload	471
	18.7.6	KE Payload	
	18.7.7	ID Payload	
	18.7.8	Cert Payload	473
	18.7.9	Certificate Request Payload	474
	18.7.10	Hash/Signature/Nonce Payloads	474
	18.7.11	Notify Payload	474
	18.7.12	Vendor ID Payload	475
18.8	Homewo	ork	476

CHAPTER 19 SSL/TLS

19.1	Introduction	7	
19.2	Using TCP		
19.3	Quick History	7	
19.4	SSL/TLS Basic Protocol	3	
19.5	Session Resumption)	
19.6	Computing the Keys	1	
19.7	Client Authentication	2	
19.8	PKI as Deployed by SSL	2	
19.9	Version Numbers	3	
19.10	Negotiating Cipher Suites	1	
	19.10.1 Who Makes the Decision?	5	
	19.10.2 Cipher Suite Names	5	
19.11	Negotiating Compression Method	5	
19.12	Attacks Fixed in v3	5	
	19.12.1 Downgrade Attack	5	
	19.12.2 Truncation Attack	5	
19.13	Exportability	7	
	19.13.1 Exportability in SSLv2	7	
	19.13.2 Exportability in SSLv3	3	
	19.13.3 Server Gated Cryptography/Step-Up)	
19.14	Encoding)	
	19.14.1 Encrypted Records	1	
	19.14.2 Handshake Messages	2	
	19.14.2.1 ClientHello	3	
	19.14.2.2 ServerHello	3	
	19.14.2.3 ServerHelloDone	3	
	19.14.2.4 ClientKeyExchange	1	
	19.14.2.5 ServerKeyExchange	1	
	19.14.2.6 CertificateRequest	5	
	19.14.2.7 Certificate	5	
	19.14.2.8 CertificateVerify	5	
	19.14.2.9 HandshakeFinished	5	
	19.14.3 ChangeCipherSpec	5	
	19.14.4 Alerts	7	
19.15	Further Reading	7	
19.16	Easy Homework	7	
19.17	Homework	3	

ELECTRONIC MAIL

CHAPTER 20	Electronic Mail Security
------------	--------------------------

	20.1	Distribution Lists	501
	20.2	Store and Forward	504
	20.3	Security Services for Electronic Mail	505
	20.4	Establishing Keys	506
		20.4.1 Establishing Public Keys	507
		20.4.2 Establishing Secret Keys	507
	20.5	Privacy	508
		20.5.1 End-to-End Privacy	508
		20.5.2 Privacy with Distribution List Exploders	509
	20.6	Authentication of the Source	510
		20.6.1 Source Authentication Based on Public Key Technology	510
		20.6.2 Source Authentication Based on Secret Keys	511
		20.6.3 Source Authentication with Distribution Lists	512
	20.7	Message Integrity	512
		20.7.1 Message Integrity without Source Authentication	513
	20.8	Non-Repudiation	514
		20.8.1 Non-Repudiation Based on Public Key Technology	514
		20.8.2 Plausible Deniability Based on Public Key Technology	514
		20.8.3 Non-Repudiation with Secret Keys	515
	20.9	Proof of Submission	516
	20.10	Proof of Delivery	516
	20.11	Message Flow Confidentiality	517
	20.12	Anonymity	517
	20.13	Containment	519
	20.14	Annoying Text Format Issues	519
		20.14.1 Disguising Data as Text	521
	20.15	Names and Addresses	523
	20.16	Verifying When a Message Was Really Sent	524
		20.16.1 Preventing Backdating	524
		20.16.2 Preventing Postdating	525
	20.17	Homework	525
CHA	PTER	21 PEM & S/MIME 5	529
	21.1	Introduction	529
	21.2	Structure of a PEM Message	530
	21.3	Establishing Keys	533
	21.4	Some PEM History	534
	21.5	PEM Certificate Hierarchy	536
	21.6	Certificate Revocation Lists (CRLs)	538
	21.7	Reformatting Data to Get Through Mailers	539
	21.8	General Structure of a PEM Message	540
	21.9	Encryption	541

21.10	Source Authentication and Integrity Protection	
21.11	Multiple Recipients	543
21.12	Bracketing PEM Messages	544
21.13	Forwarding and Enclosures	547
	21.13.1 Forwarding a Message	547
21.14	Unprotected Information	549
21.15	Message Formats	550
	21.15.1 ENCRYPTED, Public Key Variant	551
	21.15.2 ENCRYPTED, Secret Key Variant	554
	21.15.3 MIC-ONLY or MIC-CLEAR, Public Key Variant	556
	21.15.4 MIC-ONLY and MIC-CLEAR, Secret Key Variant	557
	21.15.5 CRL-RETRIEVAL-REQUEST	558
	21.15.6 CRL	558
21.16	DES-CBC as MIC Doesn't Work	558
21.17	Differences in S/MIME	
21.18	S/MIME Certificate Hierarchy	
	21.18.1 S/MIME with a Public Certifier	
	21.18.2 S/MIME with an Organizational Certifier	
	•	
	21.18.3 S/MIME with Certificates from Any Old CA	564
21.19	21.18.3 S/MIME with Certificates from Any Old CA Homework	564 565
21.19	 21.18.3 S/MIME with Certificates from Any Old CA Homework	564 565 567
21.19 CHAPTER	 21.18.3 S/MIME with Certificates from Any Old CA Homework	
21.19 CHAPTER 22.1	 21.18.3 S/MIME with Certificates from Any Old CA Homework	564 565 567 567
21.19 CHAPTER 22.1 22.2	 21.18.3 S/MIME with Certificates from Any Old CA Homework 22 PGP (Pretty Good Privacy) Introduction Overview 	
21.19 CHAPTER 22.1 22.2 22.3	 21.18.3 S/MIME with Certificates from Any Old CA Homework	
21.19 CHAPTER 22.1 22.2 22.3 22.4 22.5	21.18.3 S/MIME with Certificates from Any Old CA Homework	
21.19 CHAPTER 22.1 22.2 22.3 22.4 22.5	21.18.3 S/MIME with Certificates from Any Old CA Homework	
21.19 CHAPTER 22.1 22.2 22.3 22.4 22.5 22.6 22.6	 21.18.3 S/MIME with Certificates from Any Old CA	
21.19 CHAPTER 22.1 22.2 22.3 22.4 22.5 22.6 22.7	21.18.3 S/MIME with Certificates from Any Old CA Homework	
21.19 CHAPTER 22.1 22.2 22.3 22.4 22.5 22.6 22.7 22.8	21.18.3 S/MIME with Certificates from Any Old CA Homework	
21.19 CHAPTER 22.1 22.2 22.3 22.4 22.5 22.6 22.7 22.8 22.9	21.18.3 S/MIME with Certificates from Any Old CA Homework	
21.19 CHAPTER 22.1 22.2 22.3 22.4 22.5 22.6 22.7 22.8 22.9	21.18.3 S/MIME with Certificates from Any Old CA Homework	
21.19 CHAPTER 22.1 22.2 22.3 22.4 22.5 22.6 22.7 22.8 22.9	21.18.3 S/MIME with Certificates from Any Old CA Homework	
21.19 CHAPTER 22.1 22.2 22.3 22.4 22.5 22.6 22.7 22.8 22.9 22.10	21.18.3 S/MIME with Certificates from Any Old CA Homework	
21.19 CHAPTER 22.1 22.2 22.3 22.4 22.5 22.6 22.7 22.8 22.9 22.10	21.18.3 S/MIME with Certificates from Any Old CA Homework	

LEFTOVERS

23.1 Packet Filters	588
23.2 Application Level Gateway	589
23.3 Encrypted Tunnels	591

	23.4	Compariso	ons	592
	23.5	Why Firev	walls Don't Work	592
	23.6	Denial-of-	-Service Attacks	593
	23.7	Should Fin	rewalls Go Away?	594
СНА	PTER	24 More	e Security Systems	595
	24.1	NetWare '	V3	595
	24.2	NetWare `	V4	597
		24.2.1	NetWare's Guillou-Quisquater Authentication Scheme	600
	24.3	KryptoKn	ight	602
		24.3.1	KryptoKnight Tickets	603
		24.3.2	Authenticators	604
		24.3.3	Nonces vs. Timestamps	604
		24.3.4	Data Encryption	605
	24.4	DASS/SP	X	605
		24.4.1	DASS Certification Hierarchy	605
		24.4.2	Login Key	606
		24.4.3	DASS Authentication Handshake	606
		24.4.4	DASS Authenticators	608
		24.4.5	DASS Delegation	608
		24.4.6	Saving Bits	609
	24.5	Lotus Not	es Security	609
		24.5.1	ID Files	610
		24.5.2	Coping with Export Controls	611
		24.5.3	Certificates for Hierarchical Names	612
		24.5.4	Certificates for Flat Names	613
		24.5.5	Lotus Notes Authentication	614
		24.5.6	The Authentication Long-Term Secret	616
		24.5.7	Mail	616
		24.5.8	Certification Revocation	617
	24.6	DCE Secu	ırity	617
	24.7	Microsoft	Windows Security	622
		24.7.1	LAN Manager and NTLM	622
		24.7.2	Windows 2000 Kerberos	624
	24.8	Network I	Denial of Service	626
		24.8.1	Robust Broadcast	626
		24.8.2	Robust Packet Delivery	628
	24.9	Clipper	•	629
		24.9.1	Key Escrow	632
	24.10	Homewor	k	633

25	5.1 Introduction	635
25	5.2 URLs/URIs	636
25	5.3 HTTP	638
25	5.4 HTTP Digest Authentication	639
25	5.5 Cookies	641
	25.5.1 Alternatives to Cookies	641
	25.5.2 Cookie Rules	642
	25.5.3 Tracking Users	643
25	5.6 Other Web Security Problems	645
	25.6.1 Spoofing a Site to a User	645
	25.6.2 Merchants Unclear on the Concept	646
	25.6.3 Getting Impersonated by a Subsequent User	646
	25.6.4 Cross-Site Scripting	647
	25.6.5 Poisoning Cookies	649
	25.6.6 Other Misuse of Cookies	649
25	5.7 Homework	650
CHAPT	ER 26 Folklore	653
26	5.1 Perfect Forward Secrecy	
26	5.2 Change Keys Periodically	
26	5.3 Multiplexing Flows over a Single SA	
	26.3.1 The Splicing Attack	
	26.3.2 Service Classes	
	26.3.3 Different Cryptographic Algorithms	
26	5.4 Use Different Kevs in the Two Directions	657
26	5.5 Use Different Secret Keys for Encryption vs. Integrity Protection	657
26	5.6 Use Different Keys for Different Purposes	658
26	5.7 Use Different Keys for Signing vs. Encryption	658
26	5.8 Have Both Sides Contribute to the Master Key	659
26	5.9 Don't Let One Side Determine the Key	659
26	5.10 Hash in a Constant When Hashing a Password	
26	5.11 HMAC Rather than Simple MD	661
26	5.12 Key Expansion	661
26	5.13 Randomly Chosen IVs	
26	5.14 Use of Nonces in Protocols	663
26	6.15 Don't Let Encrypted Data Begin with a Constant	663
26	5.16 Don't Let Encrypted Data Begin with a Predictable Value	664
26	6.17 Compress Data Before Encrypting It	664
26	5.18 Don't Do Encryption Only	665
26	5.19 Avoiding Weak Keys	665
26	5.20 Minimal vs. Redundant Designs	666
26	5.21 Overestimate the Size of Key	

26.22	Hardware Random Number Generators	
26.23	Timing Attacks	
26.24	Put Checksums at the End of Data	
26.25	Forward Compatibility	669
	26.25.1 Options	669
	26.25.2 Version Numbers	670
	26.25.2.1 Version Number Field Must Not Move	670
	26.25.2.2 Negotiating Highest Version Supported	670
	26.25.2.3 Minor Version Number Field	671
	26.25.3 Vendor Options	
26.26	Negotiating Parameters	
26.27	Homework	
Bibliograp	hy	675
Glossary		685
Index		703

ACKNOWLEDGMENTS

Despite the controversies that crop up around security issues, it has been our experience that people in the security community are generally generous with their wisdom and time. It's always a little scary thanking specific people, for fear we'll leave someone out, but leaving everyone out seems wrong. It's not even the fair thing to do, since some people would be more egregiously wronged by being left out than others.

Eric Rescorla and Hilarie Orman have been particularly helpful with answering questions and reviewing chapters for this edition. Other reviewers, and people who have been helpful answering questions, include Tom Wu, Kevin Fu, Marshall Rose, Joe Tardo, Joe Pato, Seth Proctor, Timothy Spiller, Tom Rice, Kristen McIntyre, Gary Winiger, Dan Harkins, Peter Memishian, Jeff Schiller, Burt Kaliski, Tony Lauck, Phil Karn, Ron Rivest, Steve Crocker, Steve Kent, John Linn, Steve Hanna, Jim Bidzos, Dave Jablon, Ted Ts'o, Matthew Barnes, Keith McCloughrie, Jeffrey Case, Kathrin Winkler, Philippe Auphelle, Sig Handelman, Phillip Hallam-Baker, Uri Blumenthal, Serge Vaudenay, and Boyd Roberts.

We could not have done Chapter 24 *More Security Systems* without help from the various companies involved, since for the most part the security systems were previously undocumented. We'd like to thank Al Eldridge from Iris (Lotus Notes), Amir Herzberg and Mark Davis from IBM (KryptoKnight), Walt Tuvell from OSF, and Cliff Van Dyke from Microsoft (LAN Manager and Windows NT security) for explaining their systems to us, doing timely reviews of what we wrote, and being enthusiastic and supportive of the project. Although nearly 67% of us work for companies that have products in this area, the opinions we offer are ours alone, and not those of our companies.

Mary Franz, our editor at Prentice Hall, has been enthusiastic and optimistic and patient with us throughout. She's shown good judgment about when to be helpful, when to keep out of the way, when to nag, and when to just look soulful so we feel guilty enough to meet a deadline.

Despite the fact that this book has kept both of his parents busy for a significant part of his life, Ray Perlner has kept us inspired with his wholehearted and unselfish enthusiasm for the project. He's shown genuine interest in the subject matter, offered useful advice during interauthor arguments, helped search for quotes, reviewed part of the book, and particularly liked the subscripted pronouns. If we overdo those, it's just because it's fun to see him giggle. Dawn Perlner has also been a great supporter of the project, and manages to convince a surprising number of her friends, as well as complete strangers, to buy the book.

And of course we thank you, our reader. We welcome your comments and suggestions. Compliments are always welcome. We hope to update the book periodically, so if there are topics you wish we'd covered or errors you'd like us to correct, let us know. Errata can be found at http://www.phptr.com/networksecurity.

Our current email addresses are ckaufman@us.ibm.com, radia@alum.mit.edu, and ms@alum.mit.edu. But we've found that email is not always reliable and email addresses change. If all else fails, you can contact the publisher, Prentice Hall, particularly our editor Mary Franz (mfranz@prenhall.com), and find out our current addresses.

We wish to thank the following for their permission to use their quotes in this book:

- Quote on page 10 from *The Hollywood Book of Quotes*, Omnibus Press.
- Quotes on page 17 and page 25 Copyright © 1994 Newsweek, Inc. All rights reserved. Reprinted by permission.
- Quote on page 43 reprinted by permission of Singer Media Corporation.
- Quote on page 19 reprinted by permission of Turner Entertainment.
- Quote on page 117 courtesy of Donald Knuth.
- Quote on page 253 reprinted by permission of *The Wall Street Journal*, Copyright © 1992 Dow Jones & Company, Inc.

INTRODUCTION

It was a dark and stormy night. Somewhere in the distance a dog howled. A shiny object caught Alice's eye. A diamond cufflink! Only one person in the household could afford diamond cufflinks! So it was the butler, after all! Alice had to warn Bob. But how could she get a message to him without alerting the butler? If she phoned Bob, the butler might listen on an extension. If she sent a carrier pigeon out the window with the message taped to its foot, how would Bob know it was Alice that was sending the message and not Trudy attempting to frame the butler because he spurned her advances?

That's what this book is about. Not much character development for Alice and Bob, we're afraid; nor do we really get to know the butler. But we do discuss how to communicate securely over an insecure medium.

What do we mean by communicating securely? Alice should be able to send a message to Bob that only Bob can understand, even though Alice can't avoid having others see what she sends. When Bob receives a message, he should be able to know for certain that it was Alice who sent the message, and that nobody tampered with the contents of the message in the time between when Alice launched the message and Bob received it.

What do we mean by an insecure medium? Well, in some dictionary or another, under the definition of "insecure medium" should be a picture of the Internet. The world is evolving towards interconnecting every computer, and people talk about connecting household appliances as well, all into some wonderful global internetwork. How wonderful! You'd be able to send electronic mail to anyone in the world. You'd also be able to control your nuclear power plant with simple commands sent across the network while you were vacationing in Fiji. Or sunny Libya. Or historic Iraq. Inside the network the world is scary. There are links that eavesdroppers can listen in on. Information needs to be forwarded through packet switches, and these switches can be reprogrammed to listen to or modify data in transit.

The situation might seem hopeless, but we may yet be saved by the magic of mathematics, and in particular cryptography, which can take a message and transform it into a bunch of numbers known as ciphertext. The ciphertext is unintelligible gibberish except to someone who knows the secret to reversing the transformation. Cryptography allows us to disguise our data so that eavesdroppers gain no information from listening to the information as transmitted. Cryptography also allows us to create an unforgeable message and detect if it has been modified in transit. One method of accomplishing this is with a **digital signature**, a number associated with a message and its sender that can be verified as authentic by others, but can only be generated by the sender. This should seem astonishing. How can there be a number which you can verify but not generate? A person's handwritten signature can (more or less) only be generated by that person, though it can be verified by others. But it would seem as if a number shouldn't be hard to generate, especially if it can be verified. Theoretically, you could generate someone's signature by trying lots of numbers and testing each one until one passed the verification test. But with the size of the numbers used, it would take too much compute time (for instance, several universe lifetimes) to generate the signature that way. So a digital signature has the same property as a handwritten signature, in that it can only be generated by one person. But a digital signature does more than a handwritten signature. Since the digital signature depends on the contents of the message, if someone alters the message the signature will no longer be correct and the tampering will be detected. This will all become clear if you read Chapter 2 *Introduction to Cryptography*.

Cryptography is a major theme in this book, not because cryptography is intrinsically interesting (which it is), but because many of the security features people want in a computer network can best be provided through cryptography.

1.1 ROADMAP TO THE BOOK

After this introductory chapter, there are five main sections in the book:

- Part 1 CRYPTOGRAPHY Chapter 2 Introduction to Cryptography is the only part of the cryptography section of the book essential for understanding the rest of the book, since it explains the generic properties of secret key, message digest, and public key algorithms, and how each is used. We've tried our best to make the descriptions of the actual cryptographic algorithms nonthreatening yet thorough, and to give intuition into why they work. It's intended to be readable by anyone, not just graduate students in mathematics. Never once do we use the term *lemma*. We do hope you read Chapter 3 Secret Key Cryptography, Chapter 4 Modes of Operation, Chapter 5 Hashes and Message Digests, and Chapter 6 Public Key Algorithms which give the details of the popular standards, but it's also OK to skip them and save them for later, or just for reference. Chapter 7 Number Theory and Chapter 8 Math with AES and Elliptic Curves gives a deeper treatment of the mathematics behind the cryptography. Reading them is not necessary for understanding the rest of the book.
- **Part 2 AUTHENTICATION** Chapter 9 Overview of Authentication Systems introduces the general issues involved in proving your identity across a network. Chapter 10 Authentication of People deals with the special circumstances when the device proving its identity is a

human being. Chapter 11 *Security Handshake Pitfalls* deals with the details of authentication handshakes. There are many security flaws that keep getting designed into protocols. This chapter attempts to describe variations of authentication handshakes and their relative security and performance strengths. We end the chapter with a checklist of security attacks so that someone designing a protocol can specifically check their protocol for these flaws.

- **Part 3 STANDARDS** This portion of the book describes the standards: Kerberos versions 4 and 5, certificate and PKI standards, IPsec, and SSL. We hope that our descriptions will be much more readable than the standards themselves. And aside from just describing the standards, we give intuition behind the various choices, and criticisms where they are overly complex or have flaws. We hope that our commentary will make the descriptions more interesting and provide a deeper understanding of the design decisions. Our descriptions are not meant to, and cannot, replace reading the standards themselves, since the standards are subject to change. But we hope that after reading our description, it will be much easier to understand the standards.
- Part 4 ELECTRONIC MAIL Chapter 20 *Electronic Mail Security* describes the various types of security features one might want, and how they might be provided. Chapter 21 *PEM* & *S/MIME* and Chapter 22 *PGP (Pretty Good Privacy)* describe the specifics of PEM, S/MIME, and PGP.
- Part 5 LEFTOVERS Chapter 23 *Firewalls* talks about what firewalls are, what problems they solve, and what problems they do not solve. Chapter 24 *More Security Systems*, describes a variety of security systems, including Novell NetWare (Versions 3 and 4), Lotus Notes, DCE, KryptoKnight/NetSP, Clipper, SNMP, DASS/SPX, Microsoft (LAN Manager and Windows NT), and sabotage-proof routing protocols. Chapter 25 *Web Issues* talks about the protocols involved in web surfing: URLs, HTTP, HTML, cookies, etc., and the security issues these raise. We close with Chapter 26 *Folklore*, which describes the reasoning behind some of the advice you will hear from cryptographers.

1.2 WHAT TYPE OF BOOK IS THIS?

We believe the reason most computer science is hard to understand is because of jargon and irrelevant details. When people work with something long enough they invent their own language, come up with some meta-architectural framework or other, and forget that the rest of the world doesn't talk or think that way. We intend this book to be reader-friendly. We try to extract the concepts and ignore the meta-architectural framework, since whatever a meta-architectural framework is, it's irrelevant to what something does and how it works.

We believe someone who is a relative novice to the field ought to be able to read this book. But readability doesn't mean lack of technical depth. We try to go beyond the information one might find in specifications. The goal is not just to describe exactly how the various standards and de facto standards work, but to explain why they are the way they are, why some protocols designed for similar purposes are different, and the implications of the design decisions. Sometimes engineering tradeoffs were made. Sometimes the designers could have made better choices (they are human after all), in which case we explain how the protocol could have been better. This analysis should make it easier to understand the current protocols, and aid in design of future protocols.

The primary audience for this book is engineers, especially those who might need to evaluate the security of, or add security features to, a distributed system; but the book is also intended to be usable as a textbook, either on the advanced undergraduate or graduate level. Most of the chapters have homework problems at the end.

Not all the chapters will be of interest to all readers. In some cases we describe and critique a standard in great detail. These chapters might not be of interest to students or people trying to get a conceptual understanding of the field. But in many cases the standards are written fairly unintelligibly. People who need to understand the standard, perhaps to implement it, or maybe even to use it, need to have a place where it is described in a readable way (and we strive for readability), but also a place in which mistakes in the standard are pointed out as such. It's very difficult to understand why, for instance, two fields are included which both give the same information. Sometimes it is because the designers of the protocol made a mistake. Once something like that is pointed out as a simple mistake, it's much easier to understand the specification. We hope that reading the descriptions in the book will make the specifications more intelligible.

1.3 TERMINOLOGY

Computer science is filled with ill-defined terminology used by different authors in conflicting ways, often by the same author in conflicting ways. We apologize in advance for probably being guilty sometimes ourselves. Some people take terminology very seriously, and once they start to use a certain word in a certain way, are extremely offended if the rest of the world does not follow.

When I use a word, it means just what I choose it to mean—neither more nor less.

-Humpty Dumpty (in Through the Looking Glass)

Some terminology we feel fairly strongly about. We do *not* use the term *hacker* to describe the vandals that break into computer systems. These criminals call themselves hackers, and that is how they got the name. But they do not deserve the name. True hackers are master programmers, incorruptibly honest, unmotivated by money, and careful not to harm anyone. The criminals termed "hackers" are not brilliant and accomplished. It is really too bad that they not only steal money, people's time, and worse, but they've also stolen a beautiful word that had been used to describe some remarkable and wonderful people. We instead use words like *intruder*, *bad guy*, and *impostor*. When we need a name for a bad guy, we usually choose *Trudy* (since it sounds like *intruder*).

We grappled with the terms *secret key* and *public key* cryptography. Often in the security literature the terms *symmetric* and *asymmetric* are used instead of *secret* and *public*. We found the terms *symmetric* and *asymmetric* intimidating and sometimes confusing, so opted instead for *secret key* and *public key*. We occasionally regretted our decision to avoid the words *symmetric* and *asymmetric* when we found ourselves writing things like *secret key based interchange keys* rather than *symmetric interchange keys*.

We use the term *privacy* when referring to the desire to keep communication from being seen by anyone other than the intended recipients. Some people in the security community avoid the term *privacy* because they feel its meaning has been corrupted to mean *the right to know*, because in some countries there are laws known as *privacy laws* which state that citizens have the right to see records kept about themselves. *Privacy* also tends to be used when referring to keeping personal information about people from being collected and misused. The security community also avoids the use of the word *secrecy*, because *secret* has special meaning within the military context, and they feel it would be confusing to talk about the secrecy of a message that was not actually labeled top secret or secret. The term most commonly used in the security community for keeping communication from being seen is confidentiality. We find that strange because confidential, like secret, is a security label, and the security community should have scorned use of *confidential*, too. In the first edition, we chose not to use *confidentiality* because we felt it had too many syllables, and saw no reason not to use *privacy*. For the second edition we reconsidered this decision, and were about to change all use of *privacy* to *confidentiality* until one of us pointed out we'd have to change the title of the book to something like Network Security: Confidential Communication in a Non-Confidential World, at which point we decided to stick with privacy.

> Speaker: Isn't it terrifying that on the Internet we have no privacy? Heckler₁: You mean confidentiality. Get your terms straight. Heckler₂: Why do security types insist on inventing their own language? Heckler₃: It's a denial-of-service attack. —Overheard at recent gathering of security types

We often refer to things involved in a conversation by name, for instance, *Alice* and *Bob*, whether the things are people or computers. This is a convenient way of making things unambigu-

ous with relatively few words, since the pronoun *she* can be used for Alice and *he* can be used for Bob. It also avoids lengthy inter- (and even intra-) author arguments about whether to use the politically incorrect *he*, a confusing *she*, an awkward *he/she* or (*s*)*he*, an ungrammatical *they*, an impersonal *it*, or an incredibly awkward rewriting to avoid the problem. We remain slightly worried that people will assume when we've named things with human names that we are referring to people. Assume Alice, Bob, and the rest of the gang may be computers unless we specifically say something like *the user Alice*, in which case we're talking about a human.

With a name like yours, you might be any shape, almost. —Humpty Dumpty to Alice (in Through the Looking Glass)

Occasionally, one of the three of us authors will want to make a personal comment. In that case we use I or *me* with a subscript. When it's a comment that we all agree with, or that we managed to slip past me₃ (the rest of us are wimpier), we use the term *we*.

1.4 NOTATION

We use the symbol \oplus (pronounced *ex-or*) for the bitwise-exclusive-or operation. We use the symbol | for concatenation. We denote secret key encryption with curly brackets preceded by the key with which something was encrypted, as in *K*{*message*}, which means *message* is secret key encrypted with *K*. Public key encryption we denote with curly braces, and the name of the owner of the public key subscripting the close brace, as in {*message*}_{Bob}. Signing (which means using the private key), we denote with square brackets, with the name of the owner of the key subscripting the close bracket, as in [*message*]_{Bob}.

Table of Notation

\oplus	bitwise exclusive or (pronounced ex-or)
I	concatenation (pronounced concatenated with)
K{message}	message encrypted with secret key K
{message} _{Bob}	message encrypted with Bob's public key
[message] _{Bob}	message signed with Bob's private key

1.5 PRIMER ON NETWORKING

You have to know something about computer networks to understand computer network security, so we're including this primer. For a more detailed understanding, we recommend PERL99, TANE96, COME00, STEV94, KURO00.

Networks today need to be very easy to use and configure. Networks are no longer an expensive educational toy for researchers, but instead are being used by real people. Most sites with networks will not be able to hire a full-time person with networking expertise to start and keep the network running.

1.5.1 OSI Reference Model

Somehow, a book about computer networks would seem incomplete without a picture of the OSI (Open Systems Interconnection) Reference Model, so here it is.

application layer
presentation layer
session layer
transport layer
network layer
data link layer
physical layer

Figure 1-1. OSI Reference Model

The OSI Reference Model is useful because it gives some commonly used terminology, though it might mislead you into thinking that there is only one way to construct a network. The reference model was designed by an organization known as the International Standards Organization (ISO). The ISO decided it would be a good idea to standardize computer networking. Since that was too big a task for a single committee, they decided to subdivide the problem among several committees. They somewhat arbitrarily chose seven, each responsible for one layer. The basic idea is that each layer uses the services of the layer below, adds functionality, and provides a service to the layer above. When you start looking at real networks, they seldom neatly fit into the seven-layer model, but for basic understanding of networking, the OSI Reference Model is a good place to start.

- 1. **physical layer**. This layer delivers an unstructured stream of bits across a link of some sort.
- 2. **data link layer**. This layer delivers a piece of information across a single link. It organizes the physical layer's bits into packets and controls who on a shared link gets each packet.

- 3. **network layer**. This layer computes paths across an interconnected mesh of links and packet switches, and forwards packets over multiple links from source to destination.
- 4. **transport layer**. This layer establishes a reliable communication stream between a pair of systems across a network by putting sequence numbers in packets, holding packets at the destination until they can be delivered in order, and retransmitting lost packets.
- session layer. The OSI session layer adds extra functions to the reliable pair-wise communication provided by the transport layer. Most network architectures do not have or need the functionality in this layer, and it is not of concern to security, so for the purposes of this book we can ignore it.
- 6. **presentation layer**. This layer encodes application data into a canonical (system-independent) format and decodes it into a system-dependent format at the receiving end.
- 7. **application layer**. This is where the applications that use the network, such as web surfing, file transfer, and electronic mail, reside.

A layer communicates with the equivalent layer in a different node. In order to get data to a peer layer, though, the layer at the transmitting node gives the data to the layer below it (on the same node), which adds a header containing additional information if necessary, and that layer in turn gives it to the layer below. As the packet is received by the destination node, each layer reads and strips off its own header, so that the packet received by layer n looks to that layer just like it did when it was sent down to layer n-1 for transmission.

This seven-layer model is the basis for a lot of the terminology in networking, and a good first step towards understanding networks, but today's network protocols do not neatly fit this model. Throughout the book we sometimes use the OSI terminology by discussing things such as encryption at layer 2 vs. layer 3 vs. layer 4, or use the terms *data link layer*, or *transport layer*.

1.5.2 IP, UDP, and TCP

Today the most common protocols are the ones standardized by the IETF (Internet Engineering Task Force). All the IETF documents are on-line and freely available from the web site www.ietf.org. The protocols are specified in documents known as RFCs. (RFC is an abbreviation for "Request for Comments", but the time to comment is when the documents are in the more preliminary "internet draft" stage. Nobody wants to hear your comments on RFCs.)

The IETF's protocol suite is usually referred to as the "TCP/IP suite", after the most common layer 3 (IP) and layer 4 (TCP) protocols at the time the suite was being nicknamed. IP (Internet Protocol), the layer 3 protocol, is defined in RFC 791. Its job is to deliver data across a network. To get a letter mailed with the postal service, you put it in an envelope that specifies the source and desti-

nation address. Similarly, the IP layer adds an envelope (a header) to the data that specifies the source and destination addresses.

But the IP address only specifies the destination machine. There might be multiple processes at the destination machine all communicating across the network, so it's necessary to also specify which process should receive the data. This is similar to putting an apartment number on the envelope in addition to the street address. IP doesn't identify the processes, but instead has a 1-octet field that specifies which protocol should receive the packet, and the rest of the information necessary to identify the destination process is contained in the layer 4 header, in the PORT fields.

The two most important layer 4 protocols in the IETF suite are TCP (Transmission Control Protocol, defined in RFC 793) and UDP (User Datagram Protocol, defined in RFC 768). TCP sends an unlimited size stream of data, reliably (either all data is delivered to the other end without loss, duplication, or misordering, or the connection is terminated). UDP sends limited-sized individual chunks, with best-effort service. Both TCP and UDP have fields for SOURCE PORT and DESTINA-TION PORT, which specify the process to whom the data belongs. TCP additionally has sequence numbers and acknowledgments to ensure the data arrives reliably.

Some port numbers are "well-known", i.e., permanently assigned to a particular service, whereas others are dynamically assigned. Being able to contact someone at a well-known port makes it easy to establish communication. In contrast, if Alice and Bob were going to attempt to communicate by going to public telephones wherever they happened to be, they'd never be able to communicate, since neither one would know what number to call. But if one of them were listening at a well-known telephone number, then the other could call from anywhere. This is very similar to the use of well-known ports.

To communicate with a particular service, say the telnet service, at some machine at IP address x, you'd know that telnet uses TCP, and is always assigned to port 23. So in the IP header, you'd specify x as the destination address, and 6 (which means TCP) as the protocol type. In the TCP header, you'd specify port 23 as the destination port. Your process would be at a dynamically assigned port, but the recipient process at node x would know which port to reply to by copying the port from the source port in the received TCP header.

This will all become much more relevant when we discuss firewalls in Chapter 23 *Firewalls*, and how they can distinguish telnet packets (which firewall administrators would usually like to block) from, say, email packets (which firewall administrators would usually like to allow).

1.5.3 Directory Service

Having a telephone line into your house means you can access any phone in the world, if you know the telephone number. The same thing is true, more or less, in computer networks. If you know the network layer address of a node on the network, you should be able to communicate with that node. (This isn't always true because of security gateways, which we'll discuss in Chapter 23 *Firewalls*.)
But how do you find out another node's network layer address? Network layer addresses are not the kind of things that people will be able to remember, or type. People instead will want to access something using a name such as File-Server-3.

This is a similar problem to finding someone's telephone number. Typically you start out by knowing the name of the person or service you want to talk to, and then look the name up in a telephone book. In a computer network there is a service which stores information about a name, including its network layer address. Anything that needs to be found is listed in the service. Anything that needs to find something searches the service.

We call such a service a **directory**, though some people like to reserve the term "directory" for something in which you search based on an attribute (e.g., "find all the people who live on Main Street") rather than look up something based on knowing its name. Those people would call a simple service in which you look up information (rather than do complex searches) a **naming** service. We see no reason to make that distinction. It might be nice to search for all items that match a certain attribute, but usually the name will be known, and the attributes of that name will be fetched.

Rather than keeping all names in one directory, the directory service is typically structured as a tree of directories. Usually a name is hierarchical, so that the directory in which the name can be found is obvious from the name. For example, an Internet name looks like radia@east.sun.com. The top level consists of pointers to the directories com for commercial enterprises, edu for educational institutions, gov for U.S. government, and various country names. Under com, there are various company names.

Having multiple directories rather than keeping all names in one directory serves two purposes. One is to prevent the directory from getting unreasonably large. The other reason is to reduce name collisions (more than one object with the same name). For instance, when you're looking up a telephone number for your friend John Smith, it's bad enough trying to figure out which John Smith is the one you want if you know which town he lives in and the telephone company has separate directories for each town, but imagine if the telephone company didn't have separate books for each town and simply had a list of names and telephone numbers!

Ideally, with a hierarchy of directories, name collisions could be prevented. Once a company hired one Radia Perlman, they just wouldn't hire another. I_2 think that's reasonable, but someone with a name like John Smith might start having problems finding a company that could hire him.

Now why did you name your baby John? Every Tom, Dick, and Harry is named John.

-Sam Goldwyn

For electronic mail addresses, conflicts must be prevented. Typically, companies let the first John Smith use the name John@companyname for his email address, and then perhaps the next one will be Smith@companyname, and the next one JSmith@companyname, and the next one has to start using middle initials. But for directories of names, there is usually no way to avoid name

collisions within a directory. In other words, both John Smiths will use the same name within the company. Then, just like with a telephone book and multiple John Smiths, you have to do the best you can to figure out which one you want based on various attributes (such as in the telephone directory, using the street address). And just like in "real life," there will be lots of confusion where one John Smith gets messages intended for a different John Smith.

The directory service is very important to security. It is assumed to be widely available and convenient to access—otherwise large-scale networking really is too inconvenient to be practical. The directory service is a convenient place to put information, such as a user's public cryptographic key. But the directory service, although convenient, is not likely to be very secure. An intruder might tamper with the information. The magic of cryptography will help us detect such tampering so that it will not be necessary to physically secure all locations that store directory service information. If the information is tampered with, good guys will detect this. It might prevent good guys from accessing the network, since they won't be able to find information they can trust, but it will not allow bad guys unauthorized access.

1.5.4 Replicated Services

Sometimes it is convenient to have two or more computers performing the same function. One reason is performance. A single server might become overloaded, or might not be sufficiently close to all users on a large network. Another reason is availability. If the service is replicated, it does not matter if some of the replicas are down or unavailable. When someone wants to access the service provided, it doesn't matter which of the computers they reach. Often the user can't even tell whether there's a single copy of the service or there are replicas.

What are the security issues with a replicated service? You'd want the user to have the same authentication information regardless of which replica was authenticating the user. If authentication information is stored at each replica, then coordinating the databases, for example after a **change password** command, can be tricky. And if the identical exchange will work with any of the replicas, then having an eavesdropper repeat the authentication handshake with a different replica might be a security problem.

1.5.5 Packet Switching

A really naive assumption would be that if people wanted computer A to talk to computer B, they'd string a wire between A and B. This doesn't work if networks get large, either in number of nodes $(n^2 \text{ wires})$ or physical distance (it takes a lot of wire to connect each of 10000 nodes in North America with each of 10000 nodes in Asia). So in a network, messages do not go directly from sender to recipient, but rather have to be forwarded by various computers along the way. These

message forwarders are referred to as packet switches, routers, gateways, bridges, and probably lots of other names as well.

A message is generally broken into smaller chunks as it is sent through the network. There are various reasons for this.

- Messages from various sources can be interleaved on the same link. You wouldn't want your
 message to have to wait until someone else finished sending a huge message, so messages are
 sent a small chunk at a time. If the link is in the process of sending the huge message when
 your little single-chunk message arrives, your message only has to wait until the link finishes
 sending a chunk of the large message.
- Error recovery is done on the chunk. If you find out that one little chunk got mangled in transmission, only that chunk needs to be retransmitted.
- Buffer management in the routers is simpler if the size of packets has a reasonable upper limit.

1.5.6 Network Components

The network is a collection of packet switches (usually called routers) and links. A link can either be a wire between two computers or a multi-access link such as a LAN (local area network). A multi-access link has interesting security implications. Whatever is transmitted on the link can be seen by all the other nodes on that link. Multi-access links with this property include Ethernet (also known as CSMA/CD), token rings, and packet radio networks.

Connected to the backbone of the network are various types of nodes. A common categorization of the nodes is into *clients*, which are workstations that allow humans to access the resources on the network, and *servers*, which are typically dedicated machines that provide services such as file storage and printing. It should be possible to deploy a new service and have users be able to conveniently find the service. Users should be able to access the network from various locations, such as a public workstation a company makes available for visitors. If a person has a dedicated workstation located in one location, such as an office, it should be possible with a minimum of configuration for the user to plug the workstation into the network.

Historically, another method for users to access a network is through a *dumb terminal*. A dumb terminal is not a general-purpose computer and does not have the compute power to do cryp-tographic operations. Usually a dumb terminal hooks directly into a host machine, or into a *terminal server* which relays the terminal's keystrokes via a network protocol across the network to the host machine (the machine the user logs into). Very few dumb terminals remain today, but their legacy lives on in the form of software-based terminal emulators implemented in most PCs and workstations. Even though these devices are capable of complex calculations, for backward compatibility, they don't do them.

1.5.7 Destinations: Ultimate and Next-Hop

A network is something to which multiple systems can attach. We draw it as a cloud since, from the point of view of the systems connected to it, exactly what goes on inside is not relevant. If two systems are on the same cloud, one can send a message to the other by attaching a header that contains a source address and a destination address, much like putting a letter into an envelope for delivery by the postal service.



Figure 1-2. A Network

But how do you connect to the network? With a point-to-point link to a packet switch inside the network, things are reasonably simple. If A wants to send a message to B, A will put A as source address and B as destination address and send the message on the point-to-point link. But what if A is connected on a LAN? In that case, in order to transmit the packet through the network, A has to specify which of its neighbors should receive the message. For example:



Figure 1-3. Network Connections

If A wants to send a message to D it has to know (somehow—if you care how, you can read my_2 book [PERL99]) that the appropriate neighbor for forwarding the packet is R_2 . So when A transmits the message there are two destinations: R_2 as the next recipient and D as the ultimate recipient. A reasonably simple way of thinking about this is that the data link layer worries about transmission across a single link. The data link header has a source address and a destination address which indicate the transmitter on that link and the receiver on that link. The network layer worries about transmission across a multi-hop network. It has a header that carries the original source and ultimate destination. The data link header is removed each time a message is received, and a new data link header is tacked onto the message when it is forwarded to the next hop.

When A transmits the packet, the network header has source A, destination D. The data link header has source A, destination R_2 . R_2 forwards the packet to R_5 . Since R_2 is connected to R_5 with a point-to-point link, the data link header will not have addresses. But when R_5 forwards the packet to R_6 across the LAN, the network layer header will (still) be source A, destination D. The data link header will be source R_5 , destination R_6 . When R_6 forwards it (across the token ring LAN) the network header is still the same, and the data link header has source R_6 , destination D.

Most likely A's data link address will look different from its network layer address, so it's a bit sloppy to say source A in both the data link header and network header. But this is all irrelevant to security. Fascinating in its own right, but irrelevant to this book.

The network layer header can be thought of as an envelope for the message. The data link header is an outer envelope. We've described the case of two envelopes—a network header inside a data link header. The world can be even more complicated than this. In fact, the "data link layer" might be a multi-hop network with multi-hop networks inside it as well. So a message might wind up with several envelopes. Again this is fascinating stuff but irrelevant to this book.

1.5.8 Address Structure

What do addresses look like? In terms of security, the main issue is how difficult it is to forge a source address, and how easy it is to arrange for the network to deliver packets to you when they are addressed to someone other than you. For instance, think of a letter as having a source address (the return address, it's called in paper mail) and a destination address. It's easy to send a letter to any-one and put President, White House, USA as the source address. It's harder to arrange to receive mail sent to President, White House, USA if you are not the U.S. President, especially if you don't live in the White House, and most likely more difficult the further you live from the address you'd like to impersonate. Network addresses are usually hierarchical, just like a postal address. If we think of the address as specifying country/state/city/person, then in general it will be easier to arrange to receive someone else's messages if you reside in the same city (for instance by bribing a postal employee), and most difficult if they're in a different country.

Forging source addresses is easy in most network layers today. Routers can be built more defensively and do a sanity check on the source address, based on where they receive the packet from. After some highly publicized denial of service attacks, where vandals overwhelmed victim sites with nuisance traffic, many routers are now deployed with this feature of checking source addresses and discarding traffic received from an unexpected direction. It's not a perfect solution, though. As typically implemented, it requires extra configuration (so the routers will know what source addresses to expect from which directions), somewhat violates my_2 philosophy (as a layer 3 specialist) that routers should be self-configuring and adapt to topological changes, and slows down the router because it has to make an extra check when forwarding a packet.

1.6 ACTIVE VS. PASSIVE ATTACKS

A passive attack is one in which the intruder eavesdrops but does not modify the message stream in any way. An active attack is one in which the intruder may transmit messages, replay old messages, modify messages in transit, or delete selected messages from the wire. A typical active attack is one in which an intruder impersonates one end of the conversation, or acts as a man-in-the-middle (see §6.4.1 *The Bucket Brigade/Man-in-the-Middle Attack*).

1.7 LAYERS AND CRYPTOGRAPHY

Encryption and integrity protection are sometimes done on the original message or on each chunk of the message, and if on each chunk, it might be done end-to-end or hop-by-hop. There are interesting tradeoffs and implications of these choices. If done on the original message, it can be protected while being stored, and the infrastructure does not need to even know whether the data it is moving is cryptographically protected. This means that the location where the cryptographically protected message is kept, and the infrastructure for transmitting the message, need not be trusted.

Encryption hop-by-hop can foil traffic analysis, i.e., it hides from eavesdroppers the information about which parties are communicating. Thus it is useful even if encryption is being done at other layers. If done hop-by-hop, the packet switches must be trusted, because by definition of hopby-hop, the packet switches will see the plaintext.

If done end-to-end as the data is being transmitted, if individual chunks are separately encrypted and integrity protected, then the data that arrives intact can be used, whereas if there's only a single integrity check for the entire message, then any corruption or loss will require retransmitting the entire thing, since (by definition of cryptographically protecting the data as a whole instead of individual chunks) there will be no way to know where the loss/corruption occurred.

1.8 AUTHORIZATION

Network security basically attempts to answer two questions: "who are you?" and "should you be doing that?" **Authentication** proves who you are. **Authorization** defines what you're allowed to do. Typically the way a server decides whether someone should have access to a resource is by first authenticating the user, and then consulting a database associated with the resource that indicates

who is allowed to do what with that resource. For instance, the database associated with a file might say that Alice can read it and Bob and Carol can both read and write it. This database is often referred to as an **ACL** (access control list).

Another model of authorization is known as the **capability model**. Instead of listing, with each resource, the set of authorized users and their rights (e.g., read, write, execute), you would have a database that listed, for each user, everything she was allowed to do.

If there were only a single resource, then the ACL model and the capability model would be basically the same, since in both cases there would be a database that lists all the authorized users and what rights each has to that resource. But in a world in which there are many resources, not all under control of one organization, it would be difficult to have a central database listing what each user was allowed to do (for instance, all the files that user is allowed to read), and it would have scaling problems if there were many resources each user was allowed to access.

Some people worry that ACLs don't scale well if there are many users allowed access to each resource. But the concept of a **group** answers that concern. A very basic form of group implemented in some systems is that each user is a member of one group, and someone with special privileges assigns users to groups. There is a special group known as "world", which includes everyone. Alice would be allowed to read a file if her name was listed on the ACL with read access, or if her group was listed on the ACL with read access.

Extensions to the idea of groups that might be useful:

- allow a user to be in multiple groups (researchers, security experts, U.S. citizens)
- allow anyone (not just someone with special privileges) to create a group. Allow anyone to name that group on an ACL they are authorized to administer.
- allow a group for which the user explicitly invokes his membership. This type of group is known as a **role**. The main difference between what people think of as a role and what people think of as a group is that the user always has all the rights of all the groups he is a member of, but only has the rights of the role he has explicitly invoked. Some people would claim that if the user is allowed to assert multiple roles, he can have only one of them active at any time.

We discuss ways of implementing very flexible notions of groups in §15.8.3 Groups.

1.9 TEMPEST

One security concern is having intruders tap into a wire, giving them the ability to eavesdrop and possibly modify or inject messages. Another security concern is electronic emanation, whereby through the magic of physics, the movement of electrons can be measured from a surprising dis-

tance away. This means that intruders can sometimes eavesdrop without even needing to physically access the link. The U.S. military Tempest program measures how far away an intruder must be before eavesdropping is impossible. That distance is known as the device's **control zone**. The control zone is the region that must be physically guarded to keep out intruders that might be attempting to eavesdrop. A well-shielded device will have a smaller control zone. I₁ remember being told in 1979 of a tape drive that had a control zone over two miles. Unfortunately, most control zone information is classified, and I₂ couldn't get me₁ to be very specific about them, other than that they're usually expressed in metric. Since it is necessary to keep intruders away from the control zone, it's certainly better to have something with a control zone on the order of a few inches rather than a few miles (oh yeah, kilometers).

CIA eavesdroppers could not intercept the radio transmissions used by Somali warlord Mohammed Farah Aidid; his radios, intelligence officials explained, were too "low tech."

-Douglas Waller & Evan Thomas, Newsweek, October 10, 1994, page 32

1.10 KEY ESCROW FOR LAW ENFORCEMENT

Law enforcement would like to preserve its ability to wiretap otherwise secure communication. (Also, sometimes companies want to be able to read all data of their employees, either to enforce company policies, or to ensure data is not lost when an employee forgets a password or leaves the company.)

In order for the government to ensure it can always wiretap, it must prevent use of encryption, break the codes used for encryption (as it did in a military context during World War II), or somehow learn everyone's cryptographic keys. The Clipper proposal was proposed in the mid-90's and attempted the third option. It allows the government to reconstruct your key (only upon court order and with legitimate cause of course). This is made possible through the use of a device known as the Clipper chip. A lot about Clipper was classified by the government as secret (and classified by a lot of other people as evil). We describe the basic technical design of Clipper in \$24.9 *Clipper*. Although the Clipper proposal appears to have been a failure, and the government appears to have for the moment at least given up on attempting to control cryptography, the Clipper design was fascinating, and is worth learning about. The simple concept is that encryption is done with a special chip (the Clipper chip). Each chip manufactured has a unique key, and the government keeps a record of the serial number/encryption key correspondence of every chip manufactured. Because not all people have complete trust in the government, rather than keeping the key in one place, each key is broken into two quantities which must be \oplus 'd in order to obtain the actual key. Each piece is

completely useless without the other. Since each piece is kept with a separate government agency, it would require two U.S. government agencies to cooperate in order to cheat and obtain the key for your Clipper chip without a valid court order. The government assures us, and evidence of past experience supports its claim, that cooperation between U.S. government agencies is unlikely.

The Clipper proposal was always controversial, starting with its name (which violated someone's trademark on something unrelated). Why would anyone use Clipper when alternative methods should be cheaper and more secure? The reason alternatives would be cheaper is that enforcing the ability of the U.S. government to wiretap adds a lot of complexity over a design that simply encrypts the data. Proponents of Clipper gave several answers to this question:

- The government would buy a lot of Clipper chips, bringing the cost down because of volume production, so Clipper would wind up being the most cost-effective solution.
- Encryption technology is only useful if both parties have compatible equipment. Since the U.S. government would use Clipper, to talk securely to the U.S. government, you would have to use Clipper. So any other mechanism would have to be implemented *in addition* to Clipper.
- Again, since encryption technology is only useful if both parties have compatible equipment, if Clipper took over enough market share, it would essentially own the market (just like VHS, a technically inferior standard supposedly, beat out Beta in the VCR marketplace). Since Clipper would be one of the earliest standards, it might take over the marketplace before any other standards have an opportunity to become entrenched. The argument was that most people wouldn't care that Clipper enables wiretapping, because they'll assume they have nothing to fear from the U.S. government wiretapping them.
- The government claimed that the cryptographic algorithm in Clipper was stronger than you could get from a commercial source.

Civil libertarians feared Clipper was a first step towards outlawing untappable cryptography. Clipper proponents say it was not. It's true that outlawing alternatives was not part of the Clipper proposal. However, there have been independent efforts to outlaw cryptography. Those efforts have been thwarted in part with the argument that industry needs security. But if Clipper were deployed, that argument would have gone away.

Clipper was designed for telephones, fax, and other low-speed applications, and in some sense is not relevant to computer networking. Many people regard it, however, as a first step and a model for taking the same approach for computer networks.

The Clipper proposal was a commercial failure, and export controls are currently relaxed. However, the technical aspects of such designs are fascinating, laws can change at any time, and export controls have created other fascinating and arcane designs that we will describe throughout the book, for instance, §19.13 *Exportability*.

1.11 KEY ESCROW FOR CARELESS USERS

It is prudent to keep your key in a safe place so that when you misplace your own key you can retrieve a copy of the key rather than conceding that all your encrypted files are irretrievably lost. It would be a security risk to have all users' keys stored unencrypted somewhere. The database of keys could be stored encrypted with a key known to the server that was storing the database, but this would mean that someone who had access to that machine could access all the user keys. Another possibility is to encrypt the key in a way that can only be reconstructed with the cooperation of several independent machines. This is feasible, and we'll discuss it more in §24.9.1 *Key Escrow*.

Some applications don't require recoverable keys. An example of such an application is login. If a user loses the key required for login, the user can be issued a new key. A user may therefore want different keys for different uses, where only some of the keys are escrowed. For applications that do require recoverable keys, protection from compromise can be traded off against protection from loss.

1.12 VIRUSES, WORMS, TROJAN HORSES

Lions and tigers and bears, oh my!

—Dorothy (in the movie *The Wizard of Oz*)

People like to categorize different types of malicious software and assign them cute biological terms (if one is inclined to think of worms as cute). We don't think it's terribly important to distinguish between these things, but will define some of the terms that seem to be infecting the literature.

- **Trojan horse**—instructions hidden inside an otherwise useful program that do bad things. Usually the term Trojan horse is used when the malicious instructions are installed at the time the program is written (and the term *virus* is used if the instructions get added to the program later).
- virus—a set of instructions that, when executed, inserts copies of itself into other programs. More recently, the term has been applied to instructions in email messages that, when executed, cause the malicious code to be sent in email to other users.
- worm—a program that replicates itself by installing copies of itself on other machines across a network.

- **trapdoor**—an undocumented entry point intentionally written into a program, often for debugging purposes, which can be exploited as a security flaw.
- logic bomb—malicious instructions that trigger on some event in the future, such as a particular time occurring.
- **zombie**—malicious instructions installed on a system that can be remotely triggered to carry out some attack with less traceability because the attack comes from another victim. Often the attacker installs large numbers of zombies in order to be able to generate large bursts of network traffic.

We do not think it's useful to take these categories seriously. As with most categorizations (e.g., plants vs. animals, intelligence vs. instinct), there are things that don't fit neatly within these categories. So we'll refer to all kinds of malicious software generically as **digital pests**.

1.12.1 Where Do They Come From?

Where do these nasties come from? Except for trapdoors, which may be intentionally installed to facilitate troubleshooting, they are written by bad guys with nothing better to do with their lives than annoy people.

How could an implementer get away with writing a digital pest into a program? Wouldn't someone notice by looking at the program? One of the most famous results in computer science is that it is provably impossible to be able to tell what an arbitrary program will do by looking at it*, so certainly it would be impossible to tell, in general, whether the program had any unpleasant side effects besides its intended purpose. But that's not the real problem. The real problem is that nobody looks. Often when you buy a program you do not have access to the source code, and even if you did, you probably wouldn't bother reading it all, or reading it very carefully. Many programs that run have never been reviewed by anybody. A major advantage offered by the "open source" movement (where all software is made available in source code format) is that even if you don't review it carefully, there is a better chance that someone else will.

What does a virus look like? A virus can be installed in just about any program by doing the following:

- replace any instruction, say the instruction at location *x*, by a jump to some free place in memory, say location *y*; then
- write the virus program starting at location *y*; then

* This is known in the literature as *the halting problem*, which states that it is impossible in general to tell whether a given program will halt or not. In fact it is impossible in general to discern any nontrivial property of a program.

 place the instruction that was originally at location x at the end of the virus program, followed by a jump to x+1.

Besides doing whatever damage the virus program does, it might replicate itself by looking for any executable files in any directory and infecting them. Once an infected program is run, the virus is executed again, to do more damage and to replicate itself to more programs. Most viruses spread silently until some triggering event causes them to wake up and do their dastardly deeds. If they did their dastardly deeds all the time, they wouldn't spread as far.

How does a digital pest originally appear on your computer? All it takes is running a single infected program. A program posted on a bulletin board might certainly be infected. But even a program bought legitimately might conceivably have a digital pest. It might have been planted by a disgruntled employee or a saboteur who had broken into the computers of the company and installed the pest into the software before it was shipped. There have been cases where commercial programs were infected because some employee ran a program gotten from a friend or a bulletin board.

Often at holiday times people send email messages with attached programs and instructions to run them. While this used to require extracting the email message to a file and possibly processing it first, modern email systems make it very convenient to run such things... often just by clicking on an icon in the message. Often the result is some sort of cute holiday-specific thing, like displaying a picture of a turkey or playing a Christmas carol. It could certainly also contain a virus. Few people will scan the program before running it, especially if the message arrives from a friend. And if you were to run such a program and it did something cute, you might very well forward it to a friend, not realizing that in addition to doing the cute thing it might very well have installed a virus that will start destroying your directory a week after the virus is first run. A good example of a Christmas card email message is a program written by Ian Phillipps, which was a winner of the 1988 International Obfuscated C Code Contest. It is delightful as a Christmas card. It does nothing other than its intended purpose (I₁ have analyzed the thing carefully and I₂ have complete faith in me₁), but we doubt many people would take the time to understand this program before running it (see Figure 1-4).

Sometimes you don't realize you're running a program. PostScript is a complete programming language. It is possible to embed a Trojan horse into a PostScript file that infects files with viruses and does other damage. Someone could send you a file and tell you to display it. You wouldn't realize you were running a program by merely displaying the file. And if there was any hope of scanning a C program to find suspicious instructions, there are very few people who could scan a PostScript file and look for suspicious PostScript commands. PostScript is, for all practical purposes, a write-only language.

As mail systems get more clever in their attempt to make things more convenient for users, the risk becomes greater. If you receive a PostScript file and you are running a non-clever mail program, you will see delightful crud like Figure 1-5

```
/* Have yourself an obfuscated Christmas! */
#include <stdio.h>
main(t, ,a)
char *a;
{
return!0<t?t<3?main(-79,-13,a+main(-87,1- ,main(-86,0,a+1)+a)):
1,t< ?main(t+1, ,a):3,main(-94,-27+t,a)&&t==2? <13?
main(2, +1,"%s %d %d\n"):9:16:t<0?t<-72?main( ,t,</pre>
"@n'+,#'/*{}w+/w#cdnr/+,{}r/*de}+,/*{*+,/w{%+,/w#q#n+,/#{1,+,/n{n+,/+#n+,/#}
;#q#n+,/+k#;*+,/'r :'d*'3,}{w+K w'K:'+}e#';dq#'l \
q#'+d'K#!/+k#;q#'r}eKK#}w'r}eKK{nl]'/#;#q#n'){)#}w'){){nl]'/+#n';d}rw' i;# \
){nl]!/n{n#'; r{#w'r nc{nl]'/#{l,+'K {rw' iK{;[{nl]'/w#q#n'wk nw' \
iwk{KK{nl]!/w{%'l##w#' i; :{nl]'/*{q#'ld;r'}{nlwb!/*de}'c \
;;{nl'-{}rw]'/+,}##'*}#nc,',#nw]'/+kd'+e}+;#'rdq#w! nr'/ ') }+}{rl#'{n' ')# \
}'+}##(!!/")
:t<-50? ==*a?putchar(31[a]):main(-65, ,a+1):main((*a=='/')+t, ,a+1)
:0<t?main(2,2,"%s"):*a=='/'| main(0,main(-61,*a,
"!ek;dc i@bK'(q)-[w]*%n+r3#1,{}:\nuwloca-0;m .vpbks,fxntdCeghiry"),a+1);
}
```

Figure 1-4. Christmas Card?

```
%!PS-Adobe-3.0
%%Creator: Windows PSCRIPT
%%Title: c:\book\spcaut.doc
%BoundingBox: 18 9 593 784
%%DocumentNeededResources: (atend)
%%DocumentSuppliedResources: (atend)
%%Pages: (atend)
%%BeginResource: procset Win35Dict 3 1
/Win35Dict 290 dict def Win35Dict begin/bd{bind def}bind def/in{72
mul}bd/ed{exch def}bd/ld{load def}bd/tr/translate ld/gs/gsave ld/gr
/grestore ld/M/moveto ld/L/lineto ld/rmt/rmoveto ld/rlt/rlineto ld
/rct/rcurveto ld/st/stroke ld/n/newpath ld/sm/setmatrix ld/cm/currentmatrix
ld/cp/closepath ld/ARC/arcn ld/TR{65536 div}bd/lj/setlinejoin ld/lc
/setlinecap ld/ml/setmiterlimit ld/sl/setlinewidth ld/scignore false
def/sc{scignore{pop pop pop}{0 index 2 index eq 2 index 4 index eq
and {pop pop 255 div setgray } {3 { 255 div 3 1 roll } repeat setrgbcolor } if else } bd
```

Figure 1-5. Typical PostScript Code

If you wanted to display the file, you'd have to extract the mail message and send it to the printer, or input it into a special program that displays PostScript on your screen. However, a clever mail program might look at the message, recognize that it was PostScript, and automatically run the PostScript code to display the message. Although this is convenient, it is dangerous.

There are various other clever features being added to mail programs. Some mail programs allow the sender to send a program along with the mail message. Usually the mail message will arrive and display some sort of icon. If the receiver clicks on the icon, the transmitted program is executed. Someone, to illustrate this point, sent such a mail message. It displayed only as a button

that said *push me*. When the person who received the message clicked on the box, it came up with a message that said, *I could have just reformatted your hard drive*.

In the first edition of this book, we said "Before the technology for clever mail goes much further, we ought to consider how we can reap the benefits of such cleverness while minimizing the security risks". We can now confidently say that it has gone much further, no one has worried at all about security, and as anyone who has been stung by the email virus of the week can attest, the situation is a disaster.

There remain dangers associated with booting from floppy disks. If the hard drive of a computer were completely wiped clean, there has to be some way to come up again, so machines provide the feature that if there is a floppy in the drive when the machine is powered on, the machine boots off the floppy. This can be disabled, but it rarely is. Even if system software becomes sophisticated about security, it won't be able to protect against Trojan horses on the boot device. When you turn on many machines with a floppy disk inserted into the drive (intentionally or accidentally), they execute the code on the floppy in a privileged mode. If there is a virus in that code, it can infect the system.

Most PCs are configured to detect a CD-ROM placed in the drive and will execute a startup program automatically when a new one is inserted. This may have been a relatively safe thing to do when writable CD-Rs were a rarity and most CD-ROMs were commercially manufactured, but today it is likely viruses will spread this way. With CD-RWs becoming common, viruses can spread across your CDs the way they once did across floppies.

1.12.2 Spreading Pests from Machine to Machine

How might a virus or worm spread from machine to machine? An infected program might be copied to a floppy or other medium and moved to another machine. Or, as we said, a mail message might carry the infection. But it is also possible for the pest to explore the network and send itself to other machines, once it is running in one machine on a network. Such a pest is called a worm.

One method a worm can employ to transmit itself to another machine is to actually log into the other machine. It can do this by guessing passwords. Sometimes things are set up to make it really easy; for instance, account name/password pairs might be stored in script files so that a naive user can access remote resources automatically. If the worm manages to find that information, it can easily log into other machines and install itself. Sometimes machines have trapdoor debugging features, such as the ability to run a command remotely. Sometimes it isn't even necessary to log in to run such a command. And if intruders can execute commands, they can do anything.

> The only reason all the computers in the world haven't crashed at the same time is that they're not all connected together yet.

> > -Dave Cheriton

In the first widely publicized example, a Christmas-card-type email message displayed a pretty animated tree on the screen, but also scanned the user's directory for all possible correspondents, including distribution lists, and mailed itself to them. Each time a recipient read a copy of the message, it sprang into life and mailed itself off to all the mailboxes it could locate from that node. This was written by someone who was not trying to attack systems—he was just trying to be "cute" but didn't consider the likely consequences of his action. Although this program did not destroy data, like a classic worm it completely overloaded and disabled the corporation's electronic mail.

Email borne viruses have become the most common kind in the last few years, taking advantage both of bugs in email systems and naive users. The "ILoveYou" virus was particularly virulent because users were willing to take a chance and run the attachment when the message had a seductive subject line.

Worms have also become more effective at spreading as more machines are continuously connected to the Internet via cable modems and DSL, but not carefully configured for security. In several cases, insecure default configurations or bugs in services that were enabled by default provided fertile ground for worms to spread.

1.12.3 Virus Checkers

How can a program check for viruses? There's rather a race between the brave and gallant people who analyze the viruses and write clever programs to detect and eliminate them, and the foul-smelling scum who devise new types of viruses that will escape detection by all the current virus checkers.

The most common form of virus checker knows the instruction sequence for lots of types of viruses, checks all the files on disk and instructions in memory for those patterns of commands, and raises a warning if it finds a match hidden somewhere inside some file. Once you own such a virus checker, you need to periodically get updates of the patterns file that includes the newest viruses. To evade detection of their viruses, virus creators have devised what is known as a **polymorphic virus** which changes the order of its instructions, or changes to functionally similar instructions, each time it copies itself. Such a virus may still be detectable, but it takes more work and typically requires a coding change and not just a new pattern file. Modern virus checkers don't just periodically scan the disk. They actually hook into the operating system and inspect files before they are written to disk.

Another type of virus checker takes a snapshot of disk storage by recording the information in the directories, such as file lengths. It might even take message digests of the files. It is designed to run, store the information, and then run again at a future time. It will warn you if there are suspicious changes. One virus, wary of changing the length of a file by adding itself to the program, compressed the program so that the infected program would wind up being the same length as the

original. When the program was executed, the uncompressed portion containing the virus decompressed the rest of the program, so (other than the virus portion) the program could run normally.

It would be natural for viruses to attack the virus checkers rather than just trying to elude them. One could even imagine a virus spread by the mechanism a virus checker uses to keep its pattern files up to date. Such a virus has not been seen at the time of this writing, but it is something to look forward to.

1.12.4 What Can We Do Today?

The world was a scary place before computer viruses came along, and will most likely continue to be scary. How can you know whether you can trust a program not to do bad things when you run it? Wouldn't it be nice to have the equivalent of a lie-detector test for programs?

Ames slipped by a lie-detector test because no one had told the polygrapher he was under suspicion.

—Douglas Waller & Evan Thomas Newsweek article on CIA and traitor Aldrich Ames, October 10, 1994

Given that there is no infallible method to test a program for hidden bad side effects, we can't be completely safe, but there are some precautions that are worth taking:

- Don't run software from suspicious sources, like bulletin boards or people who aren't as careful as you are.
- Frequently run virus checkers. Have the industry employ people whose job it is to keep up with virus technology and come up with vaccines.
- Try to run programs in the most limited possible environments. For instance, if you have a PC in order to get real work done, and you also want to play games, sometimes using shareware or games copied from bulletin boards, have two machines. If you run a game with a virus, you'll only wipe out your games. A somewhat more practical way to accomplish this is to have a machine with multiple disks and a physical switch that connects only one of them at a time.
- When your system puts up a warning saying that something is dangerous, don't do it!
- Do frequent backups, and save old backups for a long time.
- Don't boot off floppies, except in an extreme circumstance, such as the first time you unpack your machine and turn it on. In those circumstances, be extremely careful about what floppy you boot from.

But mostly, the situation is pretty bleak, given the design of the operating systems for PCs. Maybe in the future, some of our suggestions in the next section might be implemented.

1.12.5 Wish List for the Future

 I_2 always assumed computers were designed in such a way that no program that ran on the machine could possibly injure the machine. Likewise, it should not be possible to send a piece of information to a machine that might damage the machine. People are designed properly that way, aren't they?

Sticks and stones may break my bones but words will never hurt me. (chant designed to encourage bullies to get physical)

But one of my_2 first programs consisted of something that just sat around and waited for interrupts, and processed the interrupts. The program itself, not counting the instructions for processing the interrupts, was **HERE:** JUMP HERE (that's not the exact instruction, because I_2 don't want to divulge the brand of computer). I_2 discovered (the hard way) that you weren't supposed to do that, because it burned out core at the instruction that kept getting executed over and over. Gerry Sussman, while a high school student, wrote a program that broke magnetic tapes. The guru who guarded and ran the mainframe didn't believe Gerry when he boasted that he could write a program to break tapes, so Gerry wrote his program to go down the entire row of tape drives, breaking the tape in each one. Another example was a machine designed with a small amount of nonvolatile memory, but the type that wore out after a finite number of writes, say 10000. That kind of memory is fine for something like saving terminal settings, since a human has to type the key sequence to cause a write, and a human won't do it very often. But if the same kind of memory is used for storing parameter settings received by a network management message, then it is possible to break the machine within seconds by sending it parameter settings over and over.

In an ideal world, it should be possible to load a floppy and examine the contents without fear. You should be able to receive any email message without fear. If it is a multimedia message, it should be possible to play the audio, display the video, print the text, or waft the odors without damage to either the machine or files stored on the machine. A file is just a bunch of bits. If the file claims to be audio, it should be possible, without risk of any type of harm, to play the file. Likewise, if the file claims to be something worthy of printing, it should not cause any harm to print the file.

Programs are a bit trickier. It should be possible to run a program and have confidence that it will not affect the files stored on your machine or the basic integrity of the operating system.

We know none of these "shoulds" are true today. The files on your machine can be virusinfected through email, by displaying a PostScript file, or simply by inserting a floppy disk or CD-ROM (on some machines). How could systems be designed more defensively? One simple feature would be a write-protect switch on your hard disk. Sometimes you run programs that you know should not be writing to your hard disk. A game program shouldn't be writing to your hard disk. Perhaps it wants to record highest score, but consumers might be willing to do without that frill if it means that they can run any game they want without fear of wiping out their life's work. Legitimate game manufacturers could design their games to work with write-protected hard disks. There is still the risk that the social misfits who design games that spread viruses could design their games to work with write-protected hard disks, but have their program check to see if some user has forgotten to write-protect the disk, and then launch the virus.

Timesharing systems used to be much more defensive. You could not run a program that would write into some other user's memory, or modify any portion of the disk that you were not authorized to write into. But PCs make the assumption that there is only one human on the machine at a time, and that human ought to be able to do anything it wants. To make matters worse, to enable snazzy maximally flexible features, there are all sorts of surprising places where someone can insert a program that will be executed.

The operating system ought to be built more defensively. PCs should have accounts just like timesharing systems so that you can set up a game account that can't affect the rest of the system. Likewise, your normal account shouldn't be able to alter system software. You should be able to easily run a program with the right to access only a single directory (its own). The world is moving in this direction, but progress is slow.

1.13 THE MULTI-LEVEL MODEL OF SECURITY

Computer security has become sufficiently important that it was inevitable that governments would decide they needed to "do something about it". And when governments want to know something about security, they turn to the experts: the military. And they develop standards and measurement tools by which security can be measured that are unbiased so as not to favor any one organization. And they mandate that anyone they can influence buy products meeting those standards.

The problem is that *secure* is not as simple to define as, say, *flame-retardant*. The security threats in different environments are very different, as are the best ways to counter them. The military has traditionally focussed on keeping their data secret (and learning the secrets of the other side). They are less concerned (though probably shouldn't be) about data getting corrupted or forged. In a paper world, forgeries are so difficult and so likely to expose the spies placing them that this threat takes a back seat. In the computerized environment, modification or corruption of data is a more likely threat.

1.13.1 Mandatory (Nondiscretionary) Access Controls

O negligence! ... what cross devil made me put this main secret in the packet I sent the king? Is there no way to cure this? No new device to beat this from his brains?

-Shakespeare's King Henry VIII, act 3, scene 2

Discretionary means that someone who owns a resource can make a decision as to who is allowed to use (access) it. **Nondiscretionary access controls** enforce a policy where users might be allowed to use information themselves but might not be allowed to make a copy of it available to someone else. Strict rules are automatically enforced about who is allowed access to certain resources based on the attributes of the resource, and even the owners of the resources cannot change those attributes. The analogy in the paper world is that you might be given a book full of confidential information, but you are not allowed to take the book out of the building. In the military, information often has a security classification, and just because you have access to secret information does not mean you can forward it as you see fit. Only someone with the proper clearance can see it, and clearance levels are decided by a separate organization based on background investigations—not based on whether someone seems like a nice guy at lunch.

The basic philosophy behind discretionary controls is that the users and the programs they run are good guys, and it is up to the operating system to trust them and protect each user from outsiders and other users. The basic philosophy behind nondiscretionary controls is that users are careless and the programs they run can't be presumed to be carrying out their wishes. The system must be ever vigilant to prevent the users from accidentally or intentionally giving information to someone who shouldn't have it. Careless users might accidentally type the wrong file name when including a file in a mail message, or might leave a message world-readable. The concept is to confine information within a **security perimeter**, and thus not allow any information to move from a more secure environment to a less secure environment. A secure system would have both discretionary and nondiscretionary access controls, with the latter serving as a backup mechanism with less granularity.

Now of course, if you allow the users out of the building, there is an avenue for information to leak out of a secure environment, since a user can remember the information and tell someone once the user gets out of the security perimeter. There really is no way for a computer system to prevent that. But the designers wanted to ensure that no Trojan horse in software could transmit any information out of the perimeter, that nothing a user did inadvertently could leak information, and that users couldn't spirit out larger amounts of information than they could memorize.

1.13.2 Levels of Security

What does it mean for something to be "more sensitive" than something else? We will use a somewhat simplified description of the U.S. Department of Defense (DoD) definitions of levels of security as an example. It is a reasonably general model and similar to what is done in other contexts. It is sufficient to understand the security mechanisms we'll describe.

The security label of something consists of two components:

- A security level (also known as classification), which might be an integer in some range, but in the U.S. DoD consists of one of the four ratings unclassified, confidential, secret, and top secret, where unclassified < confidential < secret < top secret.
- A set of zero or more **categories** (also known as **compartments**), which describe kinds of information. For instance, the name CRYPTO might mean information about cryptographic algorithms, INTEL might mean information about military intelligence, COMSEC might mean information about communications security, or NUCLEAR might mean information about types of families.

Documents (or computer files) are marked with a security label saying how sensitive the information is, and people are issued security clearances according to how trustworthy they are perceived to be and what information they have demonstrated a "need to know."

A clearance might therefore be (SECRET;{COMSEC,CRYPTO}), which would indicate someone who was allowed to know information classified unclassified, confidential, or secret (but not top secret) dealing with cryptographic algorithms or communications security.

Given two security labels, (X, S_1) and (Y, S_2) , (X, S_1) is defined as being "at least as sensitive as" (Y, S_2) iff $X \ge Y$ and $S_2 \subseteq S_1$. For example,

(TOP SECRET, {CRYPTO, COMSEC}) > (SECRET, {CRYPTO})

where ">" means "more sensitive than".

It is possible for two labels to be incomparable in the sense that neither is more sensitive than the other. For example, neither of the following are comparable to each other:

(TOP SECRET, {CRYPTO, COMSEC})

(SECRET, {NUCLEAR, CRYPTO})

1.13.3 Mandatory Access Control Rules

Every person, process, and piece of information has a security label. A person cannot run a process with a label higher than the person's label, but may run one with a lower label. Information is only allowed to be read by a process that has at least as high a rating as the information. The terminology used for having a process read something with a higher rating than the process is **read-up**. Read-up is illegal and must be prevented. A process cannot write a piece of information with a rating lower than the process's rating. The terminology used for a process writing something with a lower rating than the process is **write-down**. Write-down is illegal and must be prevented.

The rules are:

- A human can only run a process that has a security label below or equal to that of the human's label.
- A process can only read information marked with a security label below or equal to that of the process.
- A process can only write information marked with a security label above or equal to that of the process. Note that if a process writes information marked with a security label above that of the process, the process can't subsequently read that information.

The prevention of read-up and write-down is the central idea behind mandatory access controls. The concepts of confinement within a security perimeter and a generalized hierarchy of security classes were given a mathematical basis by Bell and La Padula in 1973 [BELL74]. There is significant complexity associated with the details of actually making them work. There has been significant subsequent research on more complex models that capture both the trustworthiness and the confidentiality of data and programs.

1.13.4 Covert Channels

A **covert channel** is a method for a Trojan horse to circumvent the automatic confinement of information within a security perimeter. Let's assume an operating system has enforced the rules in the previous section. Let's assume also that a bad guy has successfully tricked someone with a TOP SECRET clearance into running a program with a Trojan horse. The program has access to some sensitive data, and wants to pass the data to the bad guy. We're assuming the operating system prevents the process from doing this straightforwardly, but there are diabolical methods that theoretically could be employed to get information out.

The Trojan horse program cannot directly pass data, but all it needs is for there to be anything it can do that can be detected by something outside the security perimeter. As long as information can be passed one bit at a time, anything can be transmitted, given enough time.

One kind of covert channel is a **timing channel**. The Trojan horse program alternately loops and waits, in cycles of, say, one minute per bit. When the next bit is a 1, the program loops for one minute. When the next bit is a 0, the program waits for a minute. The bad guy's program running on the same computer but without access to the sensitive data constantly tests the loading of the sys-

tem. If the system is sluggish, its conspirator inside the perimeter is looping, and therefore transmitting a 1. Otherwise, the conspirator is waiting, and therefore transmitting a 0.

This assumes those two processes are the only ones running on the machine. What happens if there are other processes running and stopping at seemingly random times (from the point of view of the program trying to read the covert channel)? That introduces noise into the channel. But communications people can deal with a noisy channel; it just lowers the potential bandwidth, depending on the signal to noise ratio.

Another kind of covert channel called a **storage channel** involves the use of shared resources other than processor cycles. For instance, suppose there were a queue of finite size, say the print queue. The Trojan horse program could fill the queue to transmit a 1, and delete some jobs to transmit a 0. The covert channel reader would attempt to print something and note whether the request was accepted. Other possible shared resources that might be exploited for passing information include physical memory, disk space, network ports, and I/O buffers.

Yet another example depends on how clever the operating system is about not divulging information in error messages. For instance, suppose the operating system says file does not exist when a file really does not exist, but says insufficient privilege for requested operation when the file does exist, but inside a security perimeter off limits to the process requesting to read the file. Then the Trojan horse can alternately create and delete a file of some name known to the other process. The conspirator process periodically attempts to read the file and uses the information about which error message it gets to determine the setting of the next bit of information.

There is no general way to prevent all covert channels. Instead, people imagine all the different ways they can think of, and specifically attempt to plug those holes. For instance, the timing channel can be eliminated by giving each security level a fixed percentage of the processor cycles. This is wasteful and impractical in general, because there can be an immense number of distinct classifications (in our model of (*one of four levels*, {*categories*}), the number of possible security perimeters is $4 \cdot 2^n$, where *n* is the number of categories).

Most covert channels have very low bandwidth. In many cases, instead of attempting to eliminate a covert channel, it is more practical to introduce enough noise into the system so that the bandwidth becomes too low to be useful to an enemy. It's also possible to look for jobs that appear to be attempting to exploit covert channels (a job that alternately submitted enough print jobs to fill the queue and then deleted them would be suspicious indeed if someone knew to watch). If the bandwidth is low *and* the secret data is large, and knowing only a small subset of the secret data is not of much use to an enemy, the threat is minimized.

How much secret data must be leaked before serious damage is done can vary considerably. For example, assume there is a file with 100 megabytes of secret data. The file has been transmitted, encrypted, on an insecure network. The enemy therefore has the ciphertext, but the cryptographic algorithm used makes it impossible for the enemy to decrypt the data without knowing the key. A Trojan horse with access to the file and a covert channel with a bandwidth of 1 bit every 10 seconds would require 250 years to leak the data (by which time it's hard to believe the divulging of the

information could be damaging to anyone). However, if the Trojan horse had access to the 56-bit key, it could leak that information across the covert channel in less than 10 minutes. That information would allow the enemy to decrypt the 100-megabyte file. For this reason, many secure systems go to great pains to keep cryptographic keys out of the hands of the programs that use them.

1.13.5 The Orange Book

The National Computer Security Center (NCSC), an agency of the U.S. government, has published an official standard called "Trusted Computer System Evaluation Criteria", universally known as "the Orange Book" (guess what color the cover is). The Orange Book defines a series of ratings a computer system can have based on its security features and the care that went into its design, documentation, and testing. This rating system is intended to give government agencies and commercial enterprises an objective assessment of a system's security and to goad computer manufacturers into placing more emphasis on security.

The official categories are D, C1, C2, B1, B2, B3, and A1, which range from least secure to most secure. In reality, of course, there is no way to place all the possible properties in a linear scale. Different threats are more or less important in different environments. The authors of the Orange Book made an attempt to linearize these concerns given their priorities. But the results can be misleading. An otherwise A1 system that is missing some single feature might have a D rating. Systems not designed with the Orange Book in mind are likely to get low ratings even if they are in fact very secure.

The other problem with the Orange Book rating scheme is that the designers focused on the security priorities of military security people—keeping data secret. A rating of B1 or better requires implementation of multi-level security and mandatory access controls. In the commercial world, data integrity is at least as important as data confidentiality. Mandatory access controls, even if available, are not suitable for most commercial environments because they make some of the most common operations, such as having a highly privileged user send mail to an unprivileged user, very cumbersome.

Mandatory access controls do not by themselves protect the system from infection by viruses. Mandatory access controls allow write-up, so if some unprivileged account became infected by having someone carelessly run, say, a game program loaded from a bulletin board, the virus could spread to more secure areas. Ironically, if it was a very secure area that first got infected, the mandatory access control features would prevent the infection from spreading to the less secure environments.

The following is a summary of what properties a system must have to qualify for each rating.

D – Minimal Protection. This simply means the system did not qualify for any of the higher ratings; it might actually be very secure. No system is ever going to brag about the fact that it was awarded a D rating.

C1 – Discretionary Security Protection. The requirements at this level correspond roughly to what one might expect from a classic timesharing system. It requires

- The operating system must prevent unprivileged user programs from overwriting critical portions of its memory. (Note that many PC operating systems do not satisfy this condition.)
- Resources must be protected with access controls. Those access controls need not be sophisticated; classic owner/group/world controls would be sufficient.
- The system must authenticate users by a password or some similar mechanism, and the password database must be protected so that it cannot be accessed by unauthorized users.

There are additional requirements around testing and documentation, which become more detailed at each successive rating.

C2 – Controlled Access Protection. This level corresponds roughly to a timesharing system where security is an important concern but users are responsible for their own fates; an example might be a commercial timesharing system. The additional requirements (over those required for C1) for a C2 rating are

- access control at a per user granularity—It must be possible to permit access to any selected subset of the user community, probably via ACLs. An ACL is a data structure attached to a resource that specifies the resource's authorized users. C2 does not explicitly require ACLs, but they are the most convenient way to provide the granularity of protection that C2 requires.
- clearing of allocated memory—The operating system must ensure that freshly allocated disk space and memory does not contain "left-over" data deleted by some previous user. It can do that by writing to the space or by requiring processes to write to the space before they can read it.
- auditing—The operating system must be capable of recording security-relevant events, including authentication and object access. The audit log must be protected from tampering and must record date, time, user, object, and event. Auditing must be selective based on user and object.

It is reasonable to expect that C2-rateable systems will become ubiquitous, since they contain features that are commonly desired and do not represent an unacceptable overhead. It is somewhat surprising that such systems are not the norm. **B1 – Labeled Security Protection.** Additional requirements at this level are essentially those required to implement Mandatory Access Controls for secrecy (not integrity) except that little attention is given to covert channels. Requirements for B1 above those for C2 include

- Security Labels: Sensitivity labels must be maintained for all users, processes, and files, and read-up and write-down must be prevented by the operating system.
- Attached devices must either themselves be labeled as accepting only a single level of information, or they must accept and know how to process security labels.
- Attached printers must have a mechanism for ensuring that there is a human-readable sensitivity label printed on the top and bottom of each page corresponding to the sensitivity label of the information being printed. The operating system must enforce this correspondence.

B2 – Structured Protection. Beyond B1, there are few new features introduced; rather, the operating system must be structured to greater levels of assurance that it behaves correctly (i.e., has no bugs). Additional requirements for B2 include

- trusted path to user—There must be some mechanism to allow a user at a terminal to reliably distinguish between talking to the legitimate operating system and talking to a Trojan horse password-capturing program.
- security level changes—A terminal user must be notified when any process started by that user changes its security level.
- security kernel—The operating system must be structured so that only a minimal portion of it is security-sensitive, i.e., that bugs in the bulk of the O/S cannot cause sensitive data to leak. This is typically done by running the bulk of the O/S in the processor's user mode and having a secure-kernel mini-O/S which enforces the mandatory access controls.
- Covert channels must be identified and their bandwidth estimated, but there is no requirement that they be eliminated.
- Strict procedures must be used in the maintenance of the security-sensitive portion of the operating system. For instance, anyone modifying any portion must document what they changed, when they changed it, and why, and some set of other people should compare the updated section with the previous version.

B3 – Security Domains. Additional requirements for B3 mostly involve greater assurance that the operating system will not have bugs that might allow something to circumvent mandatory access controls. Additional requirements include

- ACLs must be able to explicitly deny access to named individuals even if they are members of groups that are otherwise allowed access. It is only at this level that ACLs must be able to separately enforce modes of access (i.e., read vs. write) to a file.
- active audit—There must be mechanisms to detect selected audited events or thresholds of audited events and immediately trigger notification of a security administrator.
- secure crashing—The system must ensure that the crashing and restarting of the system introduces no security policy violations.

A1 – Verified Design. There are no additional features in an A1 system over a B3 system. Rather, there are formal procedures for the analysis of the design of the system and more rigorous controls on its implementation.

1.13.6 Successors to the Orange Book

Gee, I wish we had one of them Doomsday machines —Turgidson, in Dr. Strangelove

Governments are rarely willing to adopt one another's ideas, especially if they didn't contribute. They would rather develop their own. The publication of the Orange Book in 1983 set off a series of efforts in various countries to come up with their own standards and classifications. These efforts eventually merged in 1990 into a single non-U.S. standard called ITSEC, followed by a reconciliation with the U.S. and the development of a single worldwide standard called the Common Criteria in 1994. Version 2.1 of the Common Criteria became an international standard in 1999.

The details of the various rating systems are all different, so passing the bureaucratic hurdles to qualify for a rating under one system would not be much of a head start toward getting an equivalent rating in another (much as different countries don't recognize each other's credentials for practicing medicine). The following table is an oversimplification that we hope won't be too offensive to the advocates of these rating systems, but it does allow the novice to judge what is being claimed about a system. In a partial acknowledgment of the multifaceted nature of security, many of the systems gave two ratings: one for the features provided and one for the degree of assurance that the system implements those features correctly. In practice, like the artistic and technical merit scores at the Olympics, the two scores tend to be closely correlated.

TCSEC (Orange Book)	German (Green Book)	British CLEF	ITSEC	Common Criteria
D	Q0		E0	EAL 0
				EAL 1
C1	Q1/F1	L1	C1/E1	EAL 2
C2	Q2/F2	L2	C2/E2	CAPP/EAL 3
B1	Q3/F3	L3	B1/E3	CSPP/EAL 4
B2	Q4/F4	L4	B2/E4	EAL 5
B3	Q5/F5	L5	B3/E5	EAL 6
A1	Q6/F5	L6	B3/E6	EAL 7

1.14 LEGAL ISSUES

The legal aspects of cryptography are fascinating, but the picture changes quickly, and we are certainly not experts in law. Although it pains us to say it, if you're going to do anything involving cryptography, talk to a lawyer.

1.14.1 Patents

One legal issue that affects the choice of security mechanisms is patents. Most cryptographic techniques are covered by patents and historically this has slowed their deployment. One of the important criteria for selection of the AES algorithm (see §3.5 *Advanced Encryption Standard (AES)*) was that it be royalty free.

The most popular public key algorithm is RSA (see §6.3 *RSA*). RSA was developed at MIT, and under the terms of MIT's funding at the time there are no license fees for U.S. government use. It was only patented in the U.S., and licensing was controlled by one company which claimed that the Hellman-Merkle patent also covered RSA, and that patent is international. Interpretation of patent rights varies by country, so the legal issues were complex. At any rate, the last patent on RSA ran out on September 20, 2000. There were many parties on that day.

To avoid large licensing fees, many standards attempted to use DSS (see §6.5 *Digital Signature Standard (DSS)*) instead of RSA. Although in most respects DSS is technically inferior to RSA, when first announced it was advertised that DSS would be freely licensable, i.e., it would not be necessary to reach agreement with the RSA licensing company. But the company claimed Hellman-Merkle covered all public key cryptography, and strengthened its position by acquiring rights to a patent by Schnorr that was closely related to DSS. Until the patents expired, the situation was murky.

"I don't know what you mean by your way," said the Queen: "all the ways about here belong to me..."

-Through the Looking Glass

Some of the relevant patents, luckily all expired, are:

- Diffie-Hellman: Patent #4,200,770, issued 1980. This covers the Diffie-Hellman key exchange described in §6.4 *Diffie-Hellman*.
- Hellman-Merkle: Patent #4,218,582, issued 1980. This is claimed to cover all public key systems. There is some controversy over whether this patent should be valid. The specific public key mechanisms described in the patent (*knapsack* systems) were subsequently broken.
- Rivest-Shamir-Adleman: Patent #4,405,829, issued 1983. This covers the RSA algorithm described in §6.3 *RSA*.
- Hellman-Pohlig: Patent #4,424,414, issued 1984. This is related to the Diffie-Hellman key exchange.

1.14.2 Export Controls

Mary had a little key (It's all she could export) And all the email that she sent Was opened at the Fort. —Ron Rivest

The U.S. government used to impose severe restrictions on export of encryption. This caused much bitterness in the computer industry and led to some fascinating technical designs so that domestic products, which were legally allowed to use strong encryption, could use strong encryption where possible, and yet interoperate with exportable products that were not allowed to use strong encryption. We describe such designs in this book (see §19.13 *Exportability*). Luckily, export controls seem to be pretty much lifted. For historical reasons, and in case they become relevant again (we sure hope not), we couldn't bear to delete the following section. It was written for the first edition of this book, and was timely in 1995 when that edition was published:

The U.S. government considers encryption to be a dangerous technology, like germ warfare and nuclear weapons. If a U.S. corporation would like to sell to other countries (and the proceeds are not going to be funding the Contras), it needs export approval. The export control laws around encryption are not clear, and their interpretation changes over time. The general principle is that the U.S. government does not want you to give out technology that would make it more difficult for them to spy. Sometimes companies get so discouraged that they leave encryption out of their products altogether. Sometimes they generate products that, when sold overseas, have the encryption mechanisms removed. It is usually possible to get export approval for encryption if the key lengths are short enough for the government to brute-force check all possible keys to decrypt a message. So sometimes companies just use short keys, or sometimes they have the capability of varying the key length, and they fix the key length to be shorter when a system is sold outside the U.S.

Even if you aren't in the business of selling software abroad, you can run afoul of export controls. If you install encryption software on your laptop and take it along with you on an international trip, you may be breaking the law. If you distribute encryption software within the U.S. without adequate warnings, you are doing your customers a disservice. And the legality of posting encryption software on a public network is questionable.

CRYPTOGRAPHY

PART 1

This page intentionally left blank

INTRODUCTION TO CRYPTOGRAPHY

2.1 WHAT IS CRYPTOGRAPHY?

The word *cryptography* comes from the Greek words $\kappa \rho \upsilon \pi \tau \sigma$ (*hidden* or *secret*) and $\gamma \rho \alpha \phi \eta$ (*writing*). So, cryptography is the art of secret writing. More generally, people think of cryptography as the art of mangling information into apparent unintelligibility in a manner allowing a secret method of unmangling. The basic service provided by cryptography is the ability to send information between participants in a way that prevents others from reading it. In this book we will concentrate on the kind of cryptography that is based on representing information as numbers and mathematically manipulating those numbers. This kind of cryptography can provide other services, such as

- integrity checking—reassuring the recipient of a message that the message has not been altered since it was generated by a legitimate source
- authentication—verifying someone's (or something's) identity

But back to the traditional use of cryptography. A message in its original form is known as **plaintext** or **cleartext**. The mangled information is known as **ciphertext**. The process for producing ciphertext from plaintext is known as **encryption**. The reverse of encryption is called **decryption**.

plaintext ______ ciphertext ______ plaintext

While cryptographers invent clever secret codes, cryptanalysts attempt to break these codes. These two disciplines constantly try to keep ahead of each other. Ultimately, the success of the cryptographers rests on the

Fundamental Tenet of Cryptography

If lots of smart people have failed to solve a problem, then it probably won't be solved (soon). Cryptographic systems tend to involve both an algorithm and a secret value. The secret value is known as the **key**. The reason for having a key in addition to an algorithm is that it is difficult to keep devising new algorithms that will allow reversible scrambling of information, and it is difficult to quickly explain a newly devised algorithm to the person with whom you'd like to start communicating securely. With a good cryptographic scheme it is perfectly OK to have everyone, including the bad guys (and the cryptanalysts) know the algorithm because knowledge of the algorithm without the key does not help unmangle the information.

The concept of a key is analogous to the combination for a combination lock. Although the concept of a combination lock is well known (you dial in the secret numbers in the correct sequence and the lock opens), you can't open a combination lock easily without knowing the combination.

2.1.1 Computational Difficulty

It is important for cryptographic algorithms to be reasonably efficient for the good guys to compute. The good guys are the ones with knowledge of the keys.* Cryptographic algorithms are not impossible to break without the key. A bad guy can simply try all possible keys until one works. The security of a cryptographic scheme depends on how much work it is for the bad guy to break it. If the best possible scheme will take 10 million years to break using all of the computers in the world, then it can be considered reasonably secure.

Going back to the combination lock example, a typical combination might consist of three numbers, each a number between 1 and 40. Let's say it takes 10 seconds to dial in a combination. That's reasonably convenient for the good guy. How much work is it for the bad guy? There are 40^3 possible combinations, which is 64000. At 10 seconds per try, it would take a week to try all combinations, though on average it would only take half that long (even though the right number is always the last one you try!).

Often a scheme can be made more secure by making the key longer. In the combination lock analogy, making the key longer would consist of requiring four numbers to be dialed in. This would make a little more work for the good guy. It might now take 13 seconds to dial in the combination. But the bad guy has 40 times as many combinations to try, at 13 seconds each, so it would take a year to try all combinations. (And if it took that long, he might want to stop to eat or sleep).

With cryptography, computers can be used to exhaustively try keys. Computers are a lot faster than people, and they don't get tired, so thousands or millions of keys can be tried per second. Also, lots of keys can be tried in parallel if you have multiple computers, so time can be saved by spending money on more computers.

Sometimes a cryptographic algorithm has a variable-length key. It can be made more secure by increasing the length of the key. Increasing the length of the key by one bit makes the good guy's

^{*}We're using the terms *good guys* for the cryptographers, and *bad guys* for the cryptanalysts. This is a convenient shorthand and not a moral judgment—in any given situation, which side you consider *good* or *bad* depends on your point of view.

job just a little bit harder, but makes the bad guy's job up to twice as hard (because the number of possible keys doubles). Some cryptographic algorithms have a fixed-length key, but a similar algorithm with a longer key can be devised if necessary. If computers get 1000 times faster, so that the bad guy's job becomes reasonably practical, making the key 10 bits longer will make the bad guy's job as hard as it was before the advance in computer speed. However, it will be much easier for the good guys (because their computer speed increase far outweighs the increment in key length). So the faster computers get, the better life gets for the good guys.

Keep in mind that breaking the cryptographic scheme is often only one way of getting what you want. For instance, a bolt cutter works no matter how many digits are in the combination.

You can get further with a kind word and a gun than you can with a kind word alone.

-Willy Sutton, bank robber

2.1.2 To Publish or Not to Publish

Some people believe that keeping a cryptographic algorithm as secret as possible will enhance its security. Others argue that publishing the algorithm, so that it is widely known, will enhance its security. On the one hand, it would seem that keeping the algorithm secret must be more secure—it makes for more work for the cryptanalyst to try to figure out what the algorithm is.

The argument for publishing the algorithm is that the bad guys will probably find out about it eventually anyway, so it's better to tell a lot of nonmalicious people about the algorithm so that in case there are weaknesses, a good guy will discover them rather than a bad guy. A good guy who discovers a weakness will warn people that the system has a weakness. Publication provides an enormous amount of free consulting from the academic community as cryptanalysts look for weaknesses so they can publish papers about them. A bad guy who discovers a weakness will exploit it for doing bad-guy things like embezzling money or stealing trade secrets.

It is difficult to keep the algorithm secret because if an algorithm is to be widely used, it is highly likely that determined attackers will manage to learn the algorithm by reverse engineering whatever implementation is distributed, or just because the more people who know something the more likely it is for the information to leak to the wrong places. In the past, "good" cryptosystems were not economically feasible, so keeping the algorithms secret was needed extra protection. We believe (we hope?) today's algorithms are sufficiently secure that this is not necessary.

Common practice today is for most commercial cryptosystems to be published and for military cryptosystems to be kept secret. If a commercial algorithm is unpublished today, it's probably for trade secret reasons or because this makes it easier to get export approval rather than to enhance its security. We suspect the military ciphers are unpublished mainly to keep good cryptographic methods out of the hands of the enemy rather than to keep them from cryptanalyzing our codes.

2.1.3 Secret Codes

We use the terms *secret code* and *cipher* interchangeably to mean any method of encrypting data. Some people draw a subtle distinction between these terms that we don't find useful.

The earliest documented cipher is attributed to Julius Caesar. The way the **Caesar cipher** would work if the message were in English is as follows. Substitute for each letter of the message, the letter which is 3 letters later in the alphabet (and wrap around to A from Z). Thus an A would become a D, and so forth. For instance, DOZEN would become GRCHQ. Once you figure out what's going on, it is very easy to read messages encrypted this way (unless, of course, the original message was in Greek).

A slight enhancement to the Caesar cipher was distributed as a premium with Ovaltine in the 1940s as *Captain Midnight Secret Decoder rings*. (There were times when this might have been a violation of export controls for distributing cryptographic hardware!) The variant is to pick a secret number *n* between 1 and 25, instead of always using 3. Substitute for each letter of the message, the letter which is *n* higher (and wrap around to A from Z of course). Thus if the secret number was 1, an A would become a B, and so forth. For instance HAL would become IBM. If the secret number was 25, then IBM would become HAL. Regardless of the value of *n*, since there are only 26 possible *ns* to try, it is still very easy to break this cipher if you know it's being used and you can recognize a message once it's decrypted.

The next type of cryptographic system developed is known as a **monoalphabetic cipher**, which consists of an arbitrary mapping of one letter to another letter. There are 26! possible pairings of letters, which is approximately 4×10^{26} . [Remember, *n*!, which reads "*n* factorial", means $n(n-1)(n-2)\cdots 1$.] This might seem secure, because to try all possibilities, if it took 1 microsecond to try each one, would take about 10 trillion years. However, by statistical analysis of language (knowing that certain letters and letter combinations are more common than others), it turns out to be fairly easy to break. For instance, many daily newspapers have a daily cryptogram, which is a monoalphabetic cipher, and can be broken by people who enjoy that sort of thing during their subway ride to work. An example is

Cf lqr'xs xsnyctm n eqxxqgsy iqul qf wdcp eqqh, erl lqrx qgt iqul!

Computers have made much more complex cryptographic schemes both necessary and possible. Necessary because computers can try keys at a rate that would exhaust an army of clerks; and possible because computers can execute the complex algorithms quickly and without errors.

2.2 BREAKING AN ENCRYPTION SCHEME

What do we mean when we speak of a bad guy Fred *breaking* an encryption scheme? The three basic attacks are known as **ciphertext only**, **known plaintext**, and **chosen plaintext**.

2.2.1 Ciphertext Only

In a ciphertext only attack, Fred has seen (and presumably stored) some ciphertext that he can analyze at leisure. Typically it is not difficult for a bad guy to obtain ciphertext. (If a bad guy can't access the encrypted data, then there would have been no need to encrypt the data in the first place!)

How can Fred figure out the plaintext if all he can see is the ciphertext? One possible strategy is to search through all the keys. Fred tries the decrypt operation with each key in turn. It is essential for this attack that Fred be able to recognize when he has succeeded. For instance, if the message was English text, then it is highly unlikely that a decryption operation with an incorrect key could produce something that looked like intelligible text. Because it is important for Fred to be able to differentiate plaintext from gibberish, this attack is sometimes known as a **recognizable plaintext** attack.

It is also essential that Fred have enough ciphertext. For instance, using the example of a monoalphabetic cipher, if the only ciphertext available to Fred were XYZ, then there is not enough information. There are many possible letter substitutions that would lead to a legal three-letter English word. There is no way for Fred to know whether the plaintext corresponding to XYZ is THE or CAT or HAT. As a matter of fact, in the following sentence, any of the words could be the plaintext for XYZ:

The hot cat was sad but you may now sit and use her big red pen.

[Don't worry-we've found a lovely sanatorium for the coauthor who wrote that.

-the other coauthors]

Often it isn't necessary to search through a lot of keys. For instance, the authentication scheme Kerberos (see §13.4 *Logging Into the Network*) assigns to user Alice a DES key derived from Alice's password according to a straightforward, published algorithm. If Alice chooses her password unwisely (say a word in the dictionary), then Fred does not need to search through all 2⁵⁶ possible DES keys—instead he only needs to try the derived keys of the 10000 or so common English words.

A cryptographic algorithm has to be secure against a ciphertext only attack because of the accessibility of the ciphertext to cryptanalysts. But in many cases cryptanalysts can obtain additional information, so it is important to design cryptographic systems to withstand the next two attacks as well.
2.2.2 Known Plaintext

Sometimes life is easier for the attacker. Suppose Fred has somehow obtained some (plaintext, ciphertext) pairs. How might he have obtained these? One possibility is that secret data might not remain secret forever. For instance, the data might consist of specifying the next city to be attacked. Once the attack occurs, the plaintext to the previous day's ciphertext is now known.

With a monoalphabetic cipher, a small amount of known plaintext would be a bonanza. From it, the attacker would learn the mappings of a substantial fraction of the most common letters (every letter that was used in the plaintext Fred obtained). Some cryptographic schemes might be good enough to be secure against ciphertext only attacks but not good enough against known plaintext attacks. In these cases, it becomes important to design the systems that use such a cryptographic algorithm to minimize the possibility that a bad guy will ever be able to obtain $\langle plaintext, ciphertext \rangle$ pairs.

2.2.3 Chosen Plaintext

On rare occasions, life may be easier still for the attacker. In a chosen plaintext attack, Fred can choose any plaintext he wants, and get the system to tell him what the corresponding ciphertext is. How could such a thing be possible?

Suppose the telegraph company offered a service in which they encrypt and transmit messages for you. Suppose Fred had eavesdropped on Alice's encrypted message. Now he'd like to break the telegraph company's encryption scheme so that he can decrypt Alice's message.

He can obtain the corresponding ciphertext to any message he chooses by paying the telegraph company to send the message for him, encrypted. For instance, if Fred knew they were using a monoalphabetic cipher, he might send the message

The quick brown fox jumps over the lazy dog.

knowing that he would thereby get all the letters of the alphabet encrypted and then be able to decrypt with certainty any encrypted message.

It is possible that a cryptosystem secure against ciphertext only and known plaintext attacks might still be susceptible to chosen plaintext attacks. For instance, if Fred knows that Alice's message is either **Surrender** or **Fight on**, then no matter how wonderful an encryption scheme the telegraph company is using, all he has to do is send the two messages and see which one looks like the encrypted data he saw when Alice's message was transmitted.

A cryptosystem should resist all three sorts of attacks. That way its users don't need to worry about whether there are any opportunities for attackers to know or choose plaintext. Like wearing both a belt and suspenders, many systems that use cryptographic algorithms will also go out of their way to prevent any chance of chosen plaintext attacks.

2.3 TYPES OF CRYPTOGRAPHIC FUNCTIONS

There are three kinds of cryptographic functions: hash functions, secret key functions, and public key functions. We will describe what each kind is, and what it is useful for. Public key cryptography involves the use of two keys. Secret key cryptography involves the use of one key. Hash functions involve the use of zero keys! Try to imagine what that could possibly mean, and what use it could possibly have—an algorithm everyone knows with no secret key, and yet it has uses in security.

Since secret key cryptography is probably the most intuitive, we'll describe that first.

2.4 SECRET KEY CRYPTOGRAPHY

Secret key cryptography involves the use of a single key. Given a message (called plaintext) and the key, encryption produces unintelligible data (called an IRS Publication—no! no! that was just a finger slip, we meant to say "ciphertext"), which is about the same length as the plaintext was. Decryption is the reverse of encryption, and uses the same key as encryption.



Secret key cryptography is sometimes referred to as **conventional cryptography** or **symmetric cryptography**. The Captain Midnight code and the monoalphabetic cipher are both examples of secret key algorithms, though both are easy to break. In this chapter we describe the functionality of cryptographic algorithms, but not the details of particular algorithms. In Chapter 3 *Secret Key Cryptography* we describe the details of some popular secret key cryptographic algorithms.

2.4.1 Security Uses of Secret Key Cryptography

The next few sections describe the types of things one might do with secret key cryptography.