

THIRD EDITION

# NETWORK SECURITY

PRIVATE Communication in a PUBLIC World

**NEW  
CONTENT**

Quantum computing,  
post-quantum algorithms,  
multiparty computation,  
fully homomorphic  
encryption,  
and more!



CHARLIE KAUFMAN • RADIA PERLMAN  
MIKE SPECINER • RAY PERLNER



FREE SAMPLE CHAPTER |



# NETWORK SECURITY

## PRIVATE Communication in a PUBLIC World

THIRD EDITION

CHARLIE KAUFMAN • RADIA PERLMAN  
MIKE SPECINER • RAY PERLNER

◆ Addison-Wesley

Boston • Columbus • New York • San Francisco • Amsterdam • Cape Town  
Dubai • London • Madrid • Milan • Munich • Paris • Montreal • Toronto • Delhi • Mexico City  
São Paulo • Sydney • Hong Kong • Seoul • Singapore • Taipei • Tokyo

Many of the designations used by manufacturers and sellers to distinguish their products are claimed as trademarks. Where those designations appear in this book, and the publisher was aware of a trademark claim, the designations have been printed with initial capital letters or in all capitals.

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

For information about buying this title in bulk quantities, or for special sales opportunities (which may include electronic versions; custom cover designs; and content particular to your business, training goals, marketing focus, or branding interests), please contact our corporate sales department at [corpsales@pearsoned.com](mailto:corpsales@pearsoned.com) or (800) 382-3419.

For government sales inquiries, please contact [governmentsales@pearsoned.com](mailto:governmentsales@pearsoned.com).

For questions about sales outside the U.S., please contact [intlcs@pearson.com](mailto:intlcs@pearson.com).

Visit us on the Web: [informit.com/aw](http://informit.com/aw)

Library of Congress Control Number: 2022941173

Copyright © 2023 Pearson Education, Inc.

All rights reserved. This publication is protected by copyright, and permission must be obtained from the publisher prior to any prohibited reproduction, storage in a retrieval system, or transmission in any form or by any means, electronic, mechanical, photocopying, recording, or likewise. For information regarding permissions, request forms and the appropriate contacts within the Pearson Education Global Rights & Permissions Department, please visit [www.pearsoned.com/permissions/](http://www.pearsoned.com/permissions/).

ISBN-13: 978-0-13-664360-9

ISBN-10: 0-13-664360-4

ScoutAutomatedPrintCode

## ***Pearson's Commitment to Diversity, Equity, and Inclusion***

Pearson is dedicated to creating bias-free content that reflects the diversity of all learners. We embrace the many dimensions of diversity, including but not limited to race, ethnicity, gender, socioeconomic status, ability, age, sexual orientation, and religious or political beliefs.

Education is a powerful force for equity and change in our world. It has the potential to deliver opportunities that improve lives and enable economic mobility. As we work with authors to create content for every product and service, we acknowledge our responsibility to demonstrate inclusivity and incorporate diverse scholarship so that everyone can achieve their potential through learning. As the world's leading learning company, we have a duty to help drive change and live up to our purpose to help more people create better lives for themselves and to create a better world.

Our ambition is to purposefully contribute to a world where:

- Everyone has an equitable and lifelong opportunity to succeed through learning
- Our educational products and services are inclusive and represent the rich diversity of learners
- Our educational content accurately reflects the histories and experiences of the learners we serve
- Our educational content prompts deeper discussions with learners and motivates them to expand their own learning (and worldview)

While we work hard to present unbiased content, we want to hear from you about any concerns or needs with this Pearson product so that we can investigate and address them.

Please contact us with concerns about any potential bias at <https://www.pearson.com/report-bias.html>.

*Si spy net work, big fedjaw iog link kyxogy*

# CONTENTS

<b>Acknowledgments</b>	<b>xix</b>
<b>About the Authors</b>	<b>xx</b>
<b>CHAPTER 1 Introduction</b>	<b>1</b>
1.1 Opinions, Products .....	2
1.2 Roadmap to the Book .....	2
1.3 Terminology .....	4
1.4 Notation .....	6
1.5 Cryptographically Protected Sessions .....	6
1.6 Active and Passive Attacks .....	7
1.7 Legal Issues .....	8
1.7.1 Patents.....	8
1.7.2 Government Regulations .....	9
1.8 Some Network Basics.....	9
1.8.1 Network Layers .....	10
1.8.2 TCP and UDP Ports.....	11
1.8.3 DNS (Domain Name System) .....	11
1.8.4 HTTP and URLs.....	12
1.8.5 Web Cookies .....	12
1.9 Names for Humans .....	13
1.10 Authentication and Authorization .....	14
1.10.1 ACL (Access Control List).....	14
1.10.2 Central Administration/Capabilities .....	14
1.10.3 Groups .....	15
1.10.4 Cross-Organizational and Nested Groups .....	15
1.10.5 Roles.....	17
1.11 Malware: Viruses, Worms, Trojan Horses .....	18
1.11.1 Where Does Malware Come From? .....	19
1.11.2 Virus Checkers .....	20
1.12 Security Gateway.....	21
1.12.1 Firewall.....	21
1.12.2 Application-Level Gateway/Proxy .....	23
1.12.3 Secure Tunnels .....	24
1.12.4 Why Firewalls Don't Work .....	25

---

---

1.13	Denial-of-Service (DoS) Attacks .....	26
1.14	NAT (Network Address Translation).....	27
1.14.1	Summary .....	28
<b>CHAPTER 2</b>	<b>Introduction to Cryptography</b> .....	<b>31</b>
2.1	Introduction .....	31
2.1.1	The Fundamental Tenet of Cryptography .....	32
2.1.2	Keys.....	32
2.1.3	Computational Difficulty .....	32
2.1.4	To Publish or Not to Publish .....	34
2.1.5	Earliest Encryption.....	34
2.1.6	One-Time Pad (OTP) .....	35
2.2	Secret Key Cryptography .....	36
2.2.1	Transmitting Over an Insecure Channel.....	37
2.2.2	Secure Storage on Insecure Media .....	37
2.2.3	Authentication .....	37
2.2.4	Integrity Check.....	38
2.3	Public Key Cryptography.....	39
2.3.1	Transmitting Over an Insecure Channel.....	40
2.3.2	Secure Storage on Insecure Media .....	41
2.3.3	Authentication .....	41
2.3.4	Digital Signatures .....	42
2.4	Hash Algorithms .....	43
2.4.1	Password Hashing .....	43
2.4.2	Message Integrity .....	43
2.4.3	Message Fingerprint .....	44
2.4.4	Efficient Digital Signatures.....	44
2.5	Breaking an Encryption Scheme .....	45
2.5.1	Ciphertext Only .....	45
2.5.2	Known Plaintext.....	46
2.5.3	Chosen Plaintext.....	46
2.5.4	Chosen Ciphertext .....	47
2.5.5	Side-Channel Attacks .....	48
2.6	Random Numbers.....	48
2.6.1	Gathering Entropy .....	49
2.6.2	Generating Random Seeds .....	50
2.6.3	Calculating a Pseudorandom Stream from the Seed .....	50
2.6.4	Periodic Reseeding .....	51
2.6.5	Types of Random Numbers.....	51
2.6.6	Noteworthy Mistakes .....	52
2.7	Numbers .....	53
2.7.1	Finite Fields.....	54

---



---

2.7.2	Exponentiation.....	55
2.7.3	Avoiding a Side-Channel Attack.....	55
2.7.4	Types of Elements used in Cryptography.....	56
2.7.5	Euclidean Algorithm.....	56
2.7.6	Chinese Remainder Theorem.....	57
2.8	Homework.....	58
<b>CHAPTER 3</b>	<b>Secret Key Cryptography</b>	<b>61</b>
3.1	Introduction.....	61
3.2	Generic Block Cipher Issues.....	61
3.2.1	Blocksize, Keysize.....	61
3.2.2	Completely General Mapping.....	62
3.2.3	Looking Random.....	63
3.3	Constructing a Practical Block Cipher.....	64
3.3.1	Per-Round Keys.....	64
3.3.2	S-boxes and Bit Shuffles.....	64
3.3.3	Feistel Ciphers.....	65
3.4	Choosing Constants.....	67
3.5	Data Encryption Standard (DES).....	67
3.5.1	DES Overview.....	68
3.5.2	The Mangler Function.....	69
3.5.3	Undesirable Symmetries.....	70
3.5.4	What's So Special About DES?.....	71
3.6	3DES (Multiple Encryption DES).....	71
3.6.1	How Many Encryptions?.....	72
3.6.1.1	Encrypting Twice with the Same Key.....	73
3.6.1.2	Encrypting Twice with Two Keys.....	73
3.6.1.3	Triple Encryption with Only Two Keys.....	74
3.6.2	Why EDE Rather Than EEE?.....	75
3.7	Advanced Encryption Standard (AES).....	75
3.7.1	Origins of AES.....	75
3.7.2	Broad Overview.....	76
3.7.3	AES Overview.....	78
3.7.4	Key Expansion.....	80
3.7.5	Inverse Rounds.....	80
3.7.6	Software Implementations of AES.....	81
3.8	RC4.....	81
3.9	Homework.....	83
<b>CHAPTER 4</b>	<b>Modes of Operation</b>	<b>85</b>
4.1	Introduction.....	85
4.2	Encrypting a Large Message.....	86



4.2.1	ECB (Electronic Code Book).....	86
4.2.2	CBC (Cipher Block Chaining).....	88
4.2.2.1	Randomized ECB .....	88
4.2.2.2	CBC .....	88
4.2.2.3	CBC Threat—Modifying Ciphertext Blocks.....	90
4.2.3	CTR (Counter Mode).....	91
4.2.3.1	Choosing IVs for CTR Mode .....	92
4.2.4	XEX (XOR Encrypt XOR) .....	93
4.2.5	XTS (XEX with Ciphertext Stealing) .....	94
4.3	Generating MACs .....	96
4.3.1	CBC-MAC .....	96
4.3.1.1	CBC Forgery Attack.....	96
4.3.2	CMAC .....	97
4.3.3	GMAC.....	98
4.3.3.1	GHASH .....	98
4.3.3.2	Transforming GHASH into GMAC .....	99
4.4	Ensuring Privacy and Integrity Together .....	100
4.4.1	CCM (Counter with CBC-MAC).....	100
4.4.2	GCM (Galois/Counter Mode) .....	101
4.5	Performance Issues.....	102
4.6	Homework.....	102
<b>CHAPTER 5 Cryptographic Hashes</b>		<b>105</b>
5.1	Introduction .....	105
5.2	The Birthday Problem .....	108
5.3	A Brief History of Hash Functions .....	108
5.4	Nifty Things to Do with a Hash .....	110
5.4.1	Digital Signatures .....	110
5.4.2	Password Database.....	111
5.4.3	Secure Shorthand of Larger Piece of Data .....	111
5.4.4	Hash Chains.....	111
5.4.5	Blockchain.....	112
5.4.6	Puzzles.....	112
5.4.7	Bit Commitment.....	113
5.4.8	Hash Trees.....	113
5.4.9	Authentication .....	114
5.4.10	Computing a MAC with a Hash.....	114
5.4.11	HMAC .....	116
5.4.12	Encryption with a Secret and a Hash Algorithm.....	118
5.5	Creating a Hash Using a Block Cipher .....	119
5.6	Construction of Hash Functions.....	119
5.6.1	Construction of MD4, MD5, SHA-1 and SHA-2.....	120

5.6.2	Construction of SHA-3.....	122
5.7	Padding.....	124
5.7.1	MD4, MD5, SHA-1, and SHA2-256 Message Padding.....	124
5.7.2	SHA-3 Padding Rule.....	126
5.8	The Internal Encryption Algorithms.....	127
5.8.1	SHA-1 Internal Encryption Algorithm.....	127
5.8.2	SHA-2 Internal Encryption Algorithm.....	128
5.9	SHA-3 $f$ Function (Also Known as KECCAK- $f$ ).....	129
5.10	Homework.....	132
<b>CHAPTER 6 First-Generation Public Key Algorithms</b>		<b>135</b>
6.1	Introduction.....	135
6.2	Modular Arithmetic.....	136
6.2.1	Modular Addition.....	136
6.2.2	Modular Multiplication.....	137
6.2.3	Modular Exponentiation.....	138
6.2.4	Fermat's Theorem and Euler's Theorem.....	139
6.3	RSA.....	140
6.3.1	RSA Algorithm.....	140
6.3.2	Why Does RSA Work?.....	141
6.3.3	Why Is RSA Secure?.....	141
6.3.4	How Efficient Are the RSA Operations?.....	142
6.3.4.1	Exponentiating with Big Numbers.....	142
6.3.4.2	Generating RSA Keys.....	144
6.3.4.3	Why a Non-Prime Has Multiple Square Roots of One.....	146
6.3.4.4	Having a Small Constant $e$ .....	147
6.3.4.5	Optimizing RSA Private Key Operations.....	148
6.3.5	Arcane RSA Threats.....	149
6.3.5.1	Smooth Numbers.....	150
6.3.5.2	The Cube Root Problem.....	151
6.3.6	Public-Key Cryptography Standard (PKCS).....	151
6.3.6.1	Encryption.....	152
6.3.6.2	The Million-Message Attack.....	152
6.3.6.3	Signing.....	153
6.4	Diffie-Hellman.....	154
6.4.1	MITM (Meddler-in-the-Middle) Attack.....	155
6.4.2	Defenses Against MITM Attack.....	156
6.4.3	Safe Primes and the Small-Subgroup Attack.....	157
6.4.4	ElGamal Signatures.....	159
6.5	Digital Signature Algorithm (DSA).....	160
6.5.1	The DSA Algorithm.....	161
6.5.2	Why Is This Secure?.....	162

---



---

6.5.3	Per-Message Secret Number .....	162
6.6	How Secure Are RSA and Diffie-Hellman? .....	163
6.7	Elliptic Curve Cryptography (ECC).....	164
6.7.1	Elliptic Curve Diffie-Hellman (ECDH).....	166
6.7.2	Elliptic Curve Digital Signature Algorithm (ECDSA) .....	167
6.8	Homework.....	168
<b>CHAPTER 7</b>	<b>Quantum Computing</b>	<b>169</b>
7.1	What Is a Quantum Computer?.....	169
7.1.1	A Preview of the Conclusions.....	169
7.1.2	First, What Is a Classical Computer?.....	170
7.1.3	Qubits and Superposition .....	171
7.1.3.1	Example of a Qubit.....	173
7.1.3.2	Multi-Qubit States and Entanglement .....	173
7.1.4	States and Gates as Vectors and Matrices.....	174
7.1.5	Becoming Superposed and Entangled.....	175
7.1.6	Linearity .....	176
7.1.6.1	No Cloning Theorem .....	176
7.1.7	Operating on Entangled Qubits .....	177
7.1.8	Unitarity .....	177
7.1.9	Doing Irreversible Operations by Measurement .....	178
7.1.10	Making Irreversible Classical Operations Reversible .....	178
7.1.11	Universal Gate Sets .....	179
7.2	Grover's Algorithm .....	180
7.2.1	Geometric Description .....	182
7.2.2	How to Negate the Amplitude of $ k\rangle$ .....	183
7.2.3	How to Reflect All the Amplitudes Across the Mean.....	184
7.2.4	Parallelizing Grover's Algorithm.....	186
7.3	Shor's Algorithm.....	186
7.3.1	Why Exponentiation mod $n$ Is a Periodic Function .....	187
7.3.2	How Finding the Period of $a^x \bmod n$ Lets You Factor $n$ .....	187
7.3.3	Overview of Shor's Algorithm.....	188
7.3.4	Converting to the Frequency Graph—Introduction .....	190
7.3.5	The Mechanics of Converting to the Frequency Graph .....	191
7.3.6	Calculating the Period .....	194
7.3.7	Quantum Fourier Transform .....	194
7.4	Quantum Key Distribution (QKD).....	195
7.4.1	Why It's Sometimes Called Quantum Encryption.....	197
7.4.2	Is Quantum Key Distribution Important?.....	197
7.5	How Hard Are Quantum Computers to Build?.....	197
7.6	Quantum Error Correction .....	199
7.7	Homework.....	201

---



---

<b>CHAPTER 8</b>	<b>Post-Quantum Cryptography</b>	<b>203</b>
8.1	Signature and/or Encryption Schemes.....	204
8.1.1	NIST Criteria for Security Levels .....	205
8.1.2	Authentication .....	205
8.1.3	Defense Against Dishonest Ciphertext.....	206
8.2	Hash-based Signatures.....	207
8.2.1	Simplest Scheme – Signing a Single Bit .....	208
8.2.2	Signing an Arbitrary-sized Message .....	208
8.2.3	Signing Lots of Messages.....	209
8.2.4	Deterministic Tree Generation .....	211
8.2.5	Short Hashes.....	212
8.2.6	Hash Chains.....	213
8.2.7	Standardized Schemes .....	214
8.2.7.1	Stateless Schemes .....	214
8.3	Lattice-Based Cryptography.....	216
8.3.1	A Lattice Problem.....	217
8.3.2	Optimization: Matrices with Structure .....	218
8.3.3	NTRU-Encryption Family of Lattice Encryption Schemes .....	219
8.3.3.1	Bob Computes a (Public, Private) Key Pair .....	220
8.3.3.2	How Bob Decrypts to Find $m$ .....	220
8.3.3.3	How Does this Relate to Lattices?.....	221
8.3.4	Lattice-Based Signatures .....	223
8.3.4.1	Basic Idea.....	223
8.3.4.2	Insecure Scheme .....	223
8.3.4.3	Fixing the Scheme .....	224
8.3.5	Learning with Errors (LWE) .....	225
8.3.5.1	LWE Optimizations .....	226
8.3.5.2	LWE-based NIST Submissions .....	229
8.4	Code-based Schemes .....	229
8.4.1	Non-cryptographic Error-correcting Codes.....	230
8.4.1.1	Invention Step .....	231
8.4.1.2	Codeword Creation Step .....	231
8.4.1.3	Misfortune Step.....	231
8.4.1.4	Diagnosis Step .....	232
8.4.2	The Parity-Check Matrix .....	234
8.4.3	Cryptographic Public Key Code-based Scheme.....	235
8.4.3.1	Neiderreiter Optimization .....	236
8.4.3.2	Generating a Public Key Pair.....	236
8.4.3.3	Using Circulant Matrices .....	239
8.5	Multivariate Cryptography .....	239
8.5.1	Solving Linear Equations .....	240

8.5.2	Quadratic Polynomials .....	240
8.5.3	Polynomial Systems .....	240
8.5.4	Multivariate Signature Systems .....	241
8.5.4.1	Multivariate Public Key Signatures.....	241
8.6	Homework.....	244
<b>CHAPTER 9</b>	<b>Authentication of People</b>	<b>249</b>
9.1	Password-based Authentication .....	251
9.1.1	Challenge-Response Based on Password.....	251
9.1.2	Verifying Passwords .....	252
9.2	Address-based Authentication.....	253
9.2.1	Network Address Impersonation.....	254
9.3	Biometrics .....	255
9.4	Cryptographic Authentication Protocols.....	255
9.5	Who Is Being Authenticated? .....	256
9.6	Passwords as Cryptographic Keys .....	256
9.7	On-Line Password Guessing .....	257
9.8	Off-Line Password Guessing .....	260
9.9	Using the Same Password in Multiple Places .....	261
9.10	Requiring Frequent Password Changes.....	261
9.11	Tricking Users into Divulging Passwords.....	262
9.12	Lamport’s Hash.....	263
9.13	Password Managers.....	265
9.14	Web Cookies .....	266
9.15	Identity Providers (IDPs) .....	267
9.16	Authentication Tokens .....	268
9.16.1	Disconnected Tokens .....	268
9.16.2	Public Key Tokens .....	270
9.17	Strong Password Protocols.....	272
9.17.1	Subtle Details .....	274
9.17.2	Augmented Strong Password Protocols .....	275
9.17.3	SRP (Secure Remote Password) .....	276
9.18	Credentials Download Protocols.....	277
9.19	Homework.....	278
<b>CHAPTER 10</b>	<b>Trusted Intermediaries</b>	<b>281</b>
10.1	Introduction .....	281
10.2	Functional Comparison .....	282
10.3	Kerberos .....	283
10.3.1	KDC Introduces Alice to Bob .....	283
10.3.2	Alice Contacts Bob.....	284
10.3.3	Ticket Granting Ticket (TGT).....	285

---

---

10.3.4	Interrealm Authentication.....	286
10.3.5	Making Password-Guessing Attacks Difficult .....	287
10.3.6	Double TGT Protocol .....	288
10.3.7	Authorization Information.....	288
10.3.8	Delegation.....	289
10.4	PKI.....	289
10.4.1	Some Terminology .....	290
10.4.2	Names in Certificates .....	290
10.5	Website Gets a DNS Name and Certificate.....	291
10.6	PKI Trust Models .....	292
10.6.1	Monopoly Model.....	292
10.6.2	Monopoly plus Registration Authorities (RAs) .....	293
10.6.3	Delegated CAs.....	294
10.6.4	Oligarchy .....	294
10.6.5	Anarchy Model.....	295
10.6.6	Name Constraints .....	296
10.6.7	Top-Down with Name Constraints.....	297
10.6.8	Multiple CAs for Any Namespace Node.....	297
10.6.9	Bottom-Up with Name Constraints .....	297
	10.6.9.1 Functionality of Up-Links. ....	298
	10.6.9.2 Functionality of Cross-Links .....	299
10.6.10	Name Constraints in PKIX Certificates .....	300
10.7	Building Certificate Chains .....	301
10.8	Revocation.....	302
10.8.1	CRL (Certificate Revocation list).....	302
10.8.2	Online Certificate Status Protocol (OCSP) .....	304
10.8.3	Good-Lists vs. Bad-Lists .....	304
10.9	Other Information in a PKIX Certificate.....	305
10.10	Issues with Expired Certificates .....	306
10.11	DNSSEC (DNS Security Extensions) .....	307
10.12	Homework.....	309
<b>CHAPTER 11</b>	<b>Communication Session Establishment</b>	<b>313</b>
11.1	One-way Authentication of Alice.....	314
	11.1.1 Timestamps vs. Challenges .....	316
	11.1.2 One-Way Authentication of Alice using a Public Key .....	318
11.2	Mutual Authentication.....	320
	11.2.1 Reflection Attack.....	320
	11.2.2 Timestamps for Mutual Authentication.....	322
11.3	Integrity/Encryption for Data .....	323
	11.3.1 Session Key Based on Shared Secret Credentials .....	324
	11.3.2 Session Key Based on Public Key Credentials .....	325

11.3.3	Session Key Based on One-Party Public Keys .....	326
11.4	Nonce Types.....	326
11.5	Intentional MITM.....	328
11.6	Detecting MITM .....	329
11.7	What Layer?.....	330
11.8	Perfect Forward Secrecy .....	333
11.9	Preventing Forged Source Addresses.....	335
11.9.1	Allowing Bob to Be Stateless in TCP .....	335
11.9.2	Allowing Bob to Be Stateless in IPsec.....	336
11.10	Endpoint Identifier Hiding .....	337
11.11	Live Partner Reassurance.....	339
11.12	Arranging for Parallel Computation.....	340
11.13	Session Resumption/Multiple Sessions.....	341
11.14	Plausible Deniability .....	343
11.15	Negotiating Crypto Parameters .....	343
11.15.1	Suites vs. à la Carte .....	344
11.15.2	Downgrade Attack.....	345
11.16	Homework.....	345
<b>CHAPTER 12</b>	<b>IPsec</b>	<b>349</b>
12.1	IPsec Security Associations .....	349
12.1.1	Security Association Database.....	350
12.1.2	Security Policy Database.....	351
12.1.3	IKE-SAs and Child-SAs.....	351
12.2	IKE (Internet Key Exchange Protocol) .....	353
12.3	Creating a Child-SA .....	355
12.4	AH and ESP .....	356
12.4.1	ESP Integrity Protection.....	356
12.4.2	Why Protect the IP Header?.....	357
12.4.3	Tunnel, Transport Mode.....	358
12.4.4	IPv4 Header.....	360
12.4.5	IPv6 Header.....	360
12.5	AH (Authentication Header) .....	361
12.6	ESP (Encapsulating Security Payload) .....	362
12.7	Comparison of Encodings .....	364
12.8	Homework.....	364
<b>CHAPTER 13</b>	<b>SSL/TLS and SSH</b>	<b>367</b>
13.1	Using TCP.....	367
13.2	StartTLS .....	368
13.3	Functions in the TLS Handshake .....	369
13.4	TLS 1.2 (and Earlier) Basic Protocol.....	369

---

---

13.5	TLS 1.3 .....	371
13.6	Session Resumption.....	372
13.7	PKI as Deployed by TLS.....	373
13.8	SSH (Secure Shell).....	374
	13.8.1 SSH Authentication.....	376
	13.8.2 SSH Port Forwarding .....	376
13.9	Homework .....	377
<b>CHAPTER 14 Electronic Mail Security</b>		<b>379</b>
14.1	Distribution Lists .....	380
14.2	Store and Forward .....	382
14.3	Disguising Binary as Text .....	383
14.4	HTML-Formatted Email .....	384
14.5	Attachments.....	385
14.6	Non-cryptographic Security Features.....	385
	14.6.1 Spam Defenses .....	385
14.7	Malicious Links in Email .....	387
14.8	Data Loss Prevention (DLP) .....	387
14.9	Knowing Bob's Email Address .....	388
14.10	Self-Destruct, Do-Not-Forward, .....	388
14.11	Preventing Spoofing of From Field.....	389
14.12	In-Flight Encryption .....	390
14.13	End-to-End Signed and Encrypted Email.....	390
14.14	Encryption by a Server .....	392
14.15	Message Integrity .....	393
14.16	Non-Repudiation .....	394
14.17	Plausible Deniability .....	394
14.18	Message Flow Confidentiality.....	395
14.19	Anonymity .....	397
14.20	Homework .....	398
<b>CHAPTER 15 Electronic Money</b>		<b>401</b>
15.1	ECASH .....	402
15.2	Offline eCash.....	403
	15.2.1 Practical Attacks.....	405
15.3	Bitcoin .....	406
	15.3.1 Transactions.....	407
	15.3.2 Bitcoin Addresses.....	407
	15.3.3 Blockchain.....	408
	15.3.4 The Ledger.....	408
	15.3.5 Mining .....	410
	15.3.6 Blockchain Forks.....	411



15.3.7	Why Is Bitcoin So Energy-Intensive? .....	412
15.3.8	Integrity Checks: Proof of Work vs. Digital Signatures .....	412
15.3.9	Concerns .....	413
15.4	Wallets for Electronic Currency .....	414
15.5	Homework .....	414
<b>CHAPTER 16</b>	<b>Cryptographic Tricks</b>	<b>417</b>
16.1	Secret Sharing .....	417
16.2	Blind Signature .....	418
16.3	Blind Decryption .....	419
16.4	Zero-Knowledge Proofs .....	420
16.4.1	Graph Isomorphism ZKP .....	420
16.4.2	Proving Knowledge of a Square Root .....	421
16.4.3	Noninteractive ZKP .....	422
16.5	Group Signatures .....	423
16.5.1	Trivial Group Signature Schemes .....	424
16.5.1.1	Single Shared Key .....	424
16.5.1.2	Group Membership Certificate .....	424
16.5.1.3	Multiple Group Membership Certificates .....	425
16.5.1.4	Blindly Signed Multiple Group Membership Certificates .....	425
16.5.2	Ring Signatures .....	425
16.5.3	DAA (Direct Anonymous Attestation) .....	427
16.5.4	EPID (Enhanced Privacy ID) .....	428
16.6	Circuit Model .....	428
16.7	Secure Multiparty Computation (MPC) .....	429
16.8	Fully Homomorphic Encryption (FHE) .....	431
16.8.1	Bootstrapping .....	432
16.8.2	Easy-to-Understand Scheme .....	433
16.9	Homework .....	434
<b>CHAPTER 17</b>	<b>Folklore</b>	<b>437</b>
17.1	Misconceptions .....	437
17.2	Perfect Forward Secrecy .....	438
17.3	Change Encryption Keys Periodically .....	439
17.4	Don't Encrypt without Integrity Protection .....	440
17.5	Multiplexing Flows over One Secure Session .....	441
17.5.1	The Splicing Attack .....	441
17.5.2	Service Classes .....	442
17.5.3	Different Cryptographic Algorithms .....	442
17.6	Using Different Secret Keys .....	443
17.6.1	For Initiator and Responder in Handshake .....	443
17.6.2	For Encryption and Integrity .....	443

17.6.3	In Each Direction of a Secure Session .....	444
17.7	Using Different Public Keys .....	444
17.7.1	Use Different Keys for Different Purposes .....	444
17.7.2	Different Keys for Signing and Encryption.....	445
17.8	Establishing Session Keys.....	446
17.8.1	Have Both Sides Contribute to the Master Key .....	446
17.8.2	Don't Let One Side Determine the Key .....	446
17.9	Hash in a Constant When Hashing a Password.....	447
17.10	HMAC Rather than Simple Keyed Hash.....	447
17.11	Key Derivation .....	448
17.12	Use of Nonces in Protocols .....	449
17.13	Creating an Unpredictable Nonce .....	449
17.14	Compression .....	450
17.15	Minimal vs. Redundant Designs.....	451
17.16	Overestimate the Size of Key.....	451
17.17	Hardware Random Number Generators .....	452
17.18	Put Checksums at the End of Data .....	452
17.19	Forward Compatibility .....	453
17.19.1	Options .....	453
17.19.2	Version Numbers.....	454
17.19.2.1	Version Number Field Must Not Move .....	454
17.19.2.2	Negotiating Highest Version Supported .....	455
17.19.2.3	Minor Version Number Field .....	455

**Glossary**

**G-1**

**Math**

**M-1**

M.1	Introduction .....	M-1
M.2	Some definitions and notation.....	M-2
M.3	Arithmetic.....	M-3
M.4	Abstract Algebra.....	M-4
M.5	Modular Arithmetic .....	M-5
M.5.1	How Do Computers Do Arithmetic?.....	M-6
M.5.2	Computing Inverses in Modular Arithmetic.....	M-6
M.5.2.1	The Euclidean Algorithm.....	M-7
M.5.2.2	The Chinese Remainder Theorem .....	M-8
M.5.3	How Fast Can We Do Arithmetic?.....	M-8
M.6	Groups .....	M-9
M.7	Fields .....	M-10
M.7.1	Polynomials .....	M-11
M.7.2	Finite Fields .....	M-13
M.7.2.1	What Sizes Can Finite Fields Be?.....	M-14

M.7.2.2	Representing a Field .....	M-14
M.8	Mathematics of Rijndael .....	M-16
M.8.1	A Rijndael Round.....	M-17
M.9	Elliptic Curve Cryptography .....	M-18
M.10	Rings.....	M-19
M.11	Linear Transformations .....	M-20
M.12	Matrix Arithmetic.....	M-21
M.12.1	Permutations.....	M-22
M.12.2	Matrix Inverses.....	M-22
M.12.2.1	Gaussian Elimination.....	M-22
M.13	Determinants .....	M-23
M.13.1	Properties of Determinants.....	M-24
M.13.1.1	Adjugate of a Matrix.....	M-24
M.13.2	Proof: Determinant of Product is Product of Determinants .....	M-25
M.14	Homework.....	M-26
<b>Bibliography</b>		<b>B-1</b>
<b>Index</b>		<b>I-1</b>

# 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.

We'd like to thank the reviewers of various chapters of this book, which include Luís T. A. N. Brandão, Bill Burr, Lisa Carnahan, Lily Chen, David Cooper, Joan Daemen, Michael Davidson, Morrie Dworkin, Ned Freed, Simson Garfinkel, Jonathan Katz, John Kelsey, Yi-Kai Liu, Kerry McKay, Carl Miller, Dustin Moody, Nicky Mouha, Rene Peralta, Andy Regenscheid, Scott Rose, Seth Rothschild, Matt Scholl, Katherine Sharpless, Jim St. Pierrez, Kevin Stine, Sheryl Taylor, Meltem Sönmez Turan, and Tatu Ylönen.

Other people who have been helpful with answering questions, especially on previous editions of this book, include Eric Rescorla, Al Eldridge, Hilarie Orman, Tom Wu, Kevin Fu, Amir Herzberg, Mark Davis, 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, Walt Tuvell, Cliff Van Dyke, Kathrin Winkler, Philippe Auphelle, Sig Handelman, Phillip Hallam-Baker, Uri Blumenthal, Serge Vaudenay, and Boyd Roberts.

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 [www.informit.com/title/9780136643609](http://www.informit.com/title/9780136643609).

Our current email addresses are [charliekaufman@outlook.com](mailto:charliekaufman@outlook.com), [radia@alum.mit.edu](mailto:radia@alum.mit.edu), [ms@alum.mit.edu](mailto:ms@alum.mit.edu), and [rperlner@gmail.com](mailto:rperlner@gmail.com).

## ***ABOUT THE AUTHORS***

**Charlie Kaufman** is currently Security Architect for Dell Storage Systems. Previously, he was the Security Architect for Microsoft Azure and before that for Lotus Notes. He has contributed to a number of IETF standards efforts including IPsec, S/MIME, and DNSSEC and served as a member of the Internet Architecture Board. He served on the National Academy of Sciences expert panel that wrote the book *Trust In Cyberspace*.

**Radia Perlman** is currently a Fellow at Dell Technologies. She is known for her contributions to bridging (spanning tree algorithm), routing (link state routing), and security (distributed systems robust despite malicious participants). She's the author of *Interconnections: Bridges, Routers, Switches, and Internetworking Protocol*. She's been elected to the National Academy of Engineering, the National Inventors Hall of Fame, the Internet Hall of Fame, and awarded lifetime achievement awards from Usenix and ACM's SIGCOMM. She has a PhD in computer science from MIT.

**Mike Speciner** is an MIT-trained technologist with expertise in mathematics, physics, and computer science. He currently serves as CTO and cofounder of The Singing Torah. His hobby is writing software for educational purposes in various common and obscure programming languages.

**Ray Perlner** is a Mathematician in the Cryptographic Technology Group of the National Institute of Standards and Technology. He has over a dozen research papers focusing primarily on post-quantum cryptography. He has degrees in both physics and mathematics from MIT.

# 1

## *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, household appliance, automobile, child's toy, and embedded medical device, all into some wonderful global internetwork. How wonderful! You'd be able to control your nuclear power plant with simple commands sent across the network while you were vacationing in Fiji. Or sunny Havana. Or historic Pyongyang. 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 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 you can verify but not generate? A person's handwritten signature can (more or less) only be generated by that person, though it can (more or less) 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 digital 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 (theoretically) has, in that it can only be generated by one person but can be verified by lots of people. 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 OPINIONS, PRODUCTS

Opinions expressed are those of the authors alone (and possibly not even agreed upon by all the authors). Opinions do not necessarily represent the views of any of the authors' past, current, or future employers. Any mention of commercial products or reference to commercial organizations is for information only. It does not imply recommendation or endorsement by NIST or any of the current, future, or prior organizations employing any of the authors.

## 1.2 ROADMAP TO THE BOOK

We aim to make this book comprehensible to engineers, giving intuition about designs. But readability doesn't mean lack of technical depth. We try to go beyond the information one might find in specifications to give insight into the designs. Given that specifications are easily available on the web today, we do not give exact packet formats.

This book should be usable as a textbook at either the undergraduate or graduate level. Most of the chapters have homework problems at the end. And to make life easier for professors who want to use the book, we will provide slides and an answer manual (to the professors using the

---

---

book) for many of the chapters. Even if you are not taking a class, you might want to do the homework problems. This book should be understandable to anyone with technical curiosity, a sense of humor\*, and a good night's sleep in the recent past. The chapters are:

- **Chapter 1 *Introduction***: This gives an overview of the chapters and some network basics.
- **Chapter 2 *Introduction to Cryptography***: This explains the cryptographic principles which will be covered in more detail in later chapters.
- **Chapter 3 *Secret Key Cryptography***: This chapter explains the uses of secret key cryptographic algorithms and describes how cryptographers create these algorithm.
- **Chapter 4 *Modes of Operation***: Since most secret key algorithms encrypt a fairly small (*e.g.*, 128 bits) block, this chapter explains various algorithms to efficiently and securely encrypt arbitrarily large amounts of data.
- **Chapter 5 *Cryptographic Hashes***: This chapter explains what hashes are used for and intuition into the methods by which cryptographers create secure and efficient hashes.
- **Chapter 6 *First-Generation Public Key Algorithms***: This chapter describes the designs of the current widely deployed public key algorithms. Unfortunately, these would be insecure if the world were able to create a sufficiently large quantum computer. So the world will migrate to different public key algorithms that we'll describe in Chapter 8 *Post-Quantum Cryptography*.
- **Chapter 7 *Quantum Computing***: This chapter gives an intuitive understanding of how a quantum computer differs from a classical computer, as well as explaining the intuition behind the two major cryptography-relevant quantum algorithms (Grover's and Shor's).
- **Chapter 8 *Post-Quantum Cryptography***: This describes the types of math problems that would remain difficult to solve even if there were quantum computers, how to turn them into public key algorithms, and various optimizations that can make them efficient.
- **Chapter 9 *Authentication of People***: This chapter describes the challenges involved in authenticating humans, and various types of technology that are or could be deployed.
- **Chapter 10 *Trusted Intermediaries***: This chapter describes technologies for distributing cryptographic keys. It also talks about trust model issues in today's deployed designs.
- **Chapter 11 *Communication Session Establishment***: This chapter describes conceptual issues in doing mutual authentication handshakes and establishing secure sessions.
- **Chapter 12 *IPsec***: This chapter goes into detail about the design of IPsec.
- **Chapter 13 *SSL/TLS and SSH***: This chapter goes into detail about the design of SSL/TLS and SSH.

---

\*Although a sense of humor is not strictly necessary for understanding the book, it is an important characteristic to have in general.



- **Chapter 14 *Electronic Mail Security***: This chapter describes various issues and solutions involved in electronic mail.
- **Chapter 15 *Electronic Money***: This chapter describes various goals of electronic money and various technologies to address them. It covers cryptocurrencies and anonymous cash.
- **Chapter 16 *Cryptographic Tricks***: This describes various exotic technologies such as secure multiparty computation and homomorphic encryption, as well as widely used technologies such as secret sharing.
- **Chapter 17 *Folklore***: This chapter gives a summary of some of the design lessons discussed in the rest of the book and also describes some common misconceptions.
- **Glossary**: We define many of the terms we use in the book.
- **Math**: This provides more in-depth coverage of the mathematics used in the rest of the book. The appendix is written solely by Mike Speciner. It's not essential for appreciating the rest of the book. It's a sample of the content in Mike's github repositories <https://github.com/ms0/>.

## 1.3 TERMINOLOGY

Computer science is filled with ill-defined terminology used by different authors in conflicting ways. 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*.

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*. When we say *secret key*, we mean a key that is used both for encryption and decryption. When we say *public key*, we are referring to a key pair consisting of a public key (used for encryption or signature verifica-

tion) and a private key (used for decryption or signing). Using the terms *public key* and *private key* is occasionally regrettable because both the words *public* and *private* start with “p”.

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* that state that citizens have the right to see records kept about them. *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 book title 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?*

Heckler<sub>1</sub>: *You mean confidentiality. Get your terms straight.*

Heckler<sub>2</sub>: *Why do security types insist on inventing their own language?*

Heckler<sub>3</sub>: *It’s a denial-of-service attack.*

—Overheard at 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 descriptions unambiguous 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-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 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 always 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.

When we need a name for a bad guy, we usually choose *Trudy* (since it sounds like *intruder*) or *Eve* (since it sounds like *eavesdropper*) or *Mallory* (since it sounds like *malice*). Everyone would assume Alice, Eve, and Trudy are she, and Bob is he. For inclusivity, we wanted at least one of the evil characters to be male, and we chose Mallory as the name of a male evildoer. Mallory can be

used for either gender, and is gaining more popularity as a female name, but when we use Mallory we will assume Mallory is male and use the pronoun *he*.

*With a name like yours, you might be any shape, almost.*

—Humpty Dumpty to Alice (in *Through the Looking Glass*)

Occasionally, one of the four 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 me<sub>3</sub> (the rest of us are wimpier), we use the term *we*.

## 1.4 NOTATION

We use the symbol  $\oplus$  (pronounced *ex-or*) for bitwise-exclusive-or. We use the symbol  $|$  for concatenation. We denote encryption with curly brackets followed by the key with which something was encrypted, as in  $\{message\}K$ , which means *message* is encrypted with *K*. We denote a signature with square brackets followed by the key, as in  $[message]_{Bob}$ . Sometimes the key is a subscript, and sometimes not. There is no deep meaning to that. Honestly, it's that sometimes we are using a key that has subscripts, such as  $K_{Alice}$ , and the formatting tool we are using makes it very difficult to have a subscripted subscript.

## 1.5 CRYPTOGRAPHICALLY PROTECTED SESSIONS

When Alice and Bob use modern cryptography and protocols, such as IPsec (Chapter 12) or TLS (Chapter 13), they first exchange a few messages in which they establish session secrets. These session secrets allow them to encrypt and integrity-protect their conversation. Although their physical connectivity is a path across the Internet, once Alice and Bob create the protected session, data that they send to each other is as trustworthy as if they had a private physically protected link.

There are various terms for this type of protected session. We will usually refer to it as a **secure session**. It is considered good security practice to use several cryptographic keys in a secure session between Alice and Bob. For example, there might be different session keys for

- encryption of Alice to Bob traffic,
- encryption of Bob to Alice traffic,

- integrity protection of Alice to Bob traffic, and
- integrity protection of Bob to Alice traffic.

Alice and Bob will each have a database describing their current secure sessions. The information in the database will include information such as how the session will be identified on incoming data, who is on the other end of the session, which cryptographic algorithms are to be used, and the sequence numbers for data to be sent or received on the session.


## 1.6 ACTIVE AND 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 or delay selected messages in transit. Passive attacks are less risky for the attacker, because it is tricky to detect or prove someone has eavesdropped. If the attacker is not on the path between Alice and Bob, the passive attack can be done by having an accomplice router make copies of the traffic and send them to the attacker, who can then analyze the data later, in private.

A typical active attack is one in which an intruder impersonates one end of the conversation, or acts as a **meddler-in-the-middle (MITM)**. (Note that the acronym MITM used to be expanded to be man-in-the-middle, but the industry is trying to move to more inclusive language. In this case, it is acknowledging that being annoying is not gender-specific.) A MITM attack is where an active attacker, say, Trudy, acts as a relay between two parties (Alice and Bob), and rather than simply forwarding messages between Alice and Bob, modifies, deletes, or inserts messages. If Trudy were faithfully forwarding messages, she could be acting as a passive eavesdropper, or she could be a correctly functioning router.

If Alice communicates with Bob using a secure session protocol with strong cryptographic protection, Trudy would gain no information by eavesdropping and would not be able to modify messages without being detected.

However, Trudy might be able to impersonate Bob's IP address to Alice, tricking Alice into establishing a secure session between Alice and Trudy. Then Trudy can simultaneously impersonate Alice to Bob, and establish a secure session between Trudy and Bob. (See Figure 1-1.)



**Figure 1-1.** MITM Attack

Alice and Bob will think they are talking to each other, but in fact they are each talking to Trudy. Data sent by Alice to Bob will be decrypted by Trudy using the session secret for the Alice-Trudy secure session, and encrypted for Bob with the session secret for the Trudy-Bob secure session. It is difficult for Alice and Bob to know that they have a MITM. Alice could attempt to make sure she's really talking to Bob by asking questions such as "What did I order when we first met for dinner?", but Trudy can forward the questions and answers. We will explain in §11.6 *Detecting MITM* how Alice and Bob can detect a MITM. And as we will explain in later chapters, if Alice has credentials for Bob that Trudy cannot impersonate, Alice and Bob can prevent a MITM.

## 1.7 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 build anything involving cryptography, talk to a lawyer. The combination of patents and export controls slowed down deployment of cryptographically secure networking, and caused strange technical choices.

### 1.7.1 Patents

One legal issue that affects the choice of security mechanisms is patents. Most cryptographic techniques were covered by patents and historically this has slowed their deployment. One of the important criteria for NIST's selection of algorithms (such as AES [§3.7 *Advanced Encryption Standard (AES)*], SHA-3 [§5.6.2 *Construction of SHA-3*], and post-quantum algorithms [Chapter 8 *Post-Quantum Cryptography*]) is whether they are royalty-free.

The widely deployed RSA algorithm (see §6.3) was developed at MIT, and under the terms of MIT's funding at the time, there were 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 20 September 2000. There were many parties on that day.

*"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*

To avoid large licensing fees, many protocol standards used DSA (see §6.5) instead of RSA. Although in most respects DSA is technically inferior to RSA, when first announced it was advertised that DSA would be freely licensable so 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 DSA. Until the patents expired (and luckily the relevant patents have expired), the situation was murky.

## 1.7.2 Government Regulations

*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 (as well as other governments) 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. Although U.S. companies still need permission from the Department of Commerce to export products containing cryptography, since around the year 2000 it has been easy to get products approved.

Additionally, even today, some countries have usage controls, so even if it were legal to export a product, it might not be legally usable inside some other country. For instance, some countries have developed their own cryptographic algorithms, and they want all their citizens to use those. Most of the reason for these sorts of rules is so that a government can't be prevented from accessing data, for instance, for law enforcement purposes.

## 1.8 SOME NETWORK BASICS

Although I<sub>2</sub> get a bit frustrated with teaching network concepts as if TCP/IP is the only way, or the best way to design a network, this is what today's Internet is built with, so we need to understand some of the details. Here is a brief introduction.

### 1.8.1 Network Layers

A good way of thinking about networking concepts is with *layers*. The concept of a layer is that inside a node, there are interfaces to adjacent layers (the layer above or the layer below). Between nodes, there are protocols for talking to peer layers. The actual protocol inside a layer can in theory be replaced by a layer that gives similar functionality to the adjacent layers. Although layers are a good way to learn about networks, deployed networks do not cleanly follow a layering model. Layers often use data associated with layers other than peer layers or adjacent layers. Layers are often subdivided into more layers, and an implementation might merge layers. ISO (International Organization for Standardization) defined a model with seven layers. The bottom layers look like this:

- Layer 1, **physical layer**. Defines how to send a stream of bits to a neighbor node (neighbors reside on the same link).
- Layer 2, **data link layer**. Defines how to structure a string of bits (provided by layer 1) into packets between neighbor nodes. This requires using the stream of bits to signal information such as “this is the beginning of a packet”, “this is the end of a packet”, and an integrity check.
- Layer 3, **network layer**. This allows a source node to send a packet of information across many links. The source adds header information to a packet to let the network know where to deliver the packet. This is analogous to putting a postal message inside an envelope, and writing the destination on the envelope. A network will consist of many links. Nodes known as *routers* or *switches* forward between links. Such nodes are connected to two or more links. They have a table known as a *forwarding table* that tells them which link to forward on, to get closer to the destination. Usually, network addresses are assigned hierarchically, so that a bunch of addresses can be summarized in one forwarding entry. This is analogous to the post office only needing to look at the destination country, and then once inside that country, forwarding towards the state, and once inside the state, forwarding to the destination city, etc. The usual protocol deployed in the Internet today for layer 3 is IP (Internet Protocol), which basically consists of adding a header to a packet identifying the source and destination, a hop count (so the network can discard packets that are looping), and other information. There are two versions of IP. IPv4 has 32-bit addresses. IPv6 has 128-bit addresses. One extra piece of information in the IP header is the 16-bit “protocol type”, which indicates which layer 4 protocol is sending the data.
- Layer 4, **transport layer**. This is information that is put in by the source, and interpreted at the destination. The service provided by TCP (Transmission Control Protocol, RFC 793) to the layer above it consists of accepting a stream of bytes at the source, and delivering the stream of bytes to the layer above TCP at the destination, without loss or duplication. To accomplish this, TCP at the sender numbers bytes; TCP at the destination uses the sequence numbers to acknowledge receipt of data, reorder data that has arrived out of sequence, and

ask for retransmission of lost data. UDP (User Datagram Protocol, RFC 768) is another layer 4 protocol that does not worry about lost or reordered data. Many processes in the layer above TCP or UDP will be reachable at the same IP address, so both UDP and TCP headers include **ports** (one for source, and one for destination), which tell the destination which process should receive the data.

## 1.8.2 TCP and UDP Ports

There are two 16-bit fields in TCP and UDP—a source port and a destination port. Typically an application on a server will be reachable at a “well-known port”, meaning that the port is specified in the protocol. If a client wants to reach that application at a server, the protocol type field in the IP header will be either TCP (6) or UDP (17) and the destination port field in the layer 4 header (TCP or UDP in this case) will be the well-known port for that application. For example, HTTP is at port 80, and HTTPS is at port 443. The source port will usually be a dynamically assigned port (49152 through 65535).

## 1.8.3 DNS (Domain Name System)

Another aspect of Internet networking we will be discussing is DNS. It is basically a distributed directory that maps DNS names (*e.g.*, `example.com`) to IP addresses. DNS names are hierarchical. A simple way to think of DNS is that for each level in the DNS name (*e.g.*, `root`, `.org`, `.com`, `example.com`) there is a server that keeps a directory associated with names in that level. The root would have a directory that allows looking up servers for each of the top-level domains (TLDs) (*e.g.*, `.org`, `.com`, `.gov`, `.tv`). There are currently over a thousand TLDs, so the root would have information associated with each of those TLDs in its database. In general, to find a DNS name, a node starts at the root, finds the server that holds the directory for the next level down, and keeps going until it gets to the server that stores information about the actual name. There are several advantages to DNS being hierarchical.

- Someone that wishes to purchase a DNS name has a choice of organizations from which to purchase a name. If a name is purchased from the organization managing names in the TLD `.org`, the purchased name will be of the form `example.org`. If you purchase the name `example.org`, you can then name anything that would be below that name in the DNS hierarchy, such as `xyz.example.com` or `labs.xyz.example.com`.
- The DNS database will not become unmanageably large, because no organization needs to keep the entire DNS database. In fact, nobody knows how many names are in the DNS database.



- It is fine to have the same lower level name in multiple databases. For instance, there is no problem with there being DNS names `example.com` and `example.org`.

### 1.8.4 HTTP and URLs

When we access things on the web, we use a protocol known as HTTP (hypertext transfer protocol). HTTP allows specifying more than a DNS name; it allows specifying a particular web page at the service with a DNS name. The URL (uniform resource locator) is the address of the web page. The URL contains a DNS name of the service, followed by additional information that is interpreted solely by the server that receives the request. The additional information might be, for instance, the directory path at the destination server that finds the information to construct the page being requested.

Sometimes humans type URLs, but usually URLs are displayed as links in a webpage that can be clicked on. It is common for people to do an Internet search (*e.g.*, using Google or Bing) for something, and then click on choices. URLs can be very long and ugly, and people usually don't look at the URL they click on. Often the web page that displays a link does not display the actual URL. Mousing over the link will sometimes show the human a URL. Unfortunately, the web page can choose what to display on the page as the link, and what to display the mouse-over link as. These can be different from the actual URL that will be followed if the link is clicked on. For example, on a web page, a clickable link (which is usually displayed in a different color), might display as “click here for information”, and if a suspicious user moused-over the link, it might display “`http://www.example.com/information`”, but if the user clicks on the link, the malicious webpage could send them to any URL, *e.g.*, `http://www.rentahitman.com`.

The two main HTTP request types are GET and POST. GET is for reading a web page and POST is for sending information to a web server. The response contains information such as the content requested and status information (such as “OK” or “not found” or “unauthorized”). One status that might be included in a response is a redirect. This informs the browser that it should go to a different URL. The browser will then go to the new URL, as if the user had clicked on a link.

### 1.8.5 Web Cookies

If a client is browsing content that requires authentication and access control, or is accumulating information such as items in a virtual shopping basket to be purchased when the user is finished browsing the on-line catalog, the information for that session needs to be kept somewhere. But HTTP is stateless. Each request/response interaction is allowed to take place over a fresh TCP connection. The cookie mechanism enables the server to maintain context across many request/response interactions. A **cookie** is a piece of data sent to the client by the server in response

to an HTTP request. The cookie need not be interpreted by the client. Instead, the client keeps a list of DNS names and cookies it has received from a server with that DNS name. If the client made a request at `example.com`, and `example.com` sent a cookie, the client will remember (`example.com: cookie`) in its cookie list. When the client next makes an HTTP request to `example.com`, it searches its cookie database for any cookies received from `example.com`, and includes those cookies in its HTTP request.

The cookie might contain all the relevant information about a user, or the server might keep a database with this information. In that case, the cookie only needs to contain the user's identity (so the server can locate that user in its database), along with proof that the user has already authenticated to the server. For example, if Alice has authenticated to Bob, Bob could send Alice a cookie consisting of some function of the name "Alice" and a secret that only Bob knows. A cookie will be cryptographically protected by the server in various ways. It might be encrypted with a key that only the server knows. It might contain information that only allows the cookie to be used from a specific machine. And it is almost always protected when transmitted across the network because the client and server will be communicating over a secure session.

## 1.9 NAMES FOR HUMANS

It is sometimes important that an identifier for a human be unique, but it is not important for a human to have a *single* unique name. Humans have many unique identifiers. For example, an email address, a telephone number, or a username specific to a website. In theory, nobody but the human needs to know that the various identities refer to the same human, but, unfortunately, it has become easy for organizations to correlate various identities. Also, it is not uncommon for family members or close friends to share an account, so an email address or username at a website might actually be multiple humans, but we will ignore that issue.

Human names are problematic. Consider email addresses. 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. Then, to send to your colleague John Smith, you have to do the best you can to figure out which email address in the company directory is the one you want, based on various attributes (such as their location or their job title, if you are lucky enough to have this information included in the directory). There will be lots of confusion when one John Smith gets messages intended for a different John Smith. This is a problem for both the John Smith that is mistakenly sent the email, as well as the John Smith who was the intended recipient of the email. Usually, a person can quickly delete spam, but with a name like John Smith, irrelevant-looking email might actually be important email for a different John Smith in the company, so must be carefully

read, and forwarded just in case. And the unfortunate John Smith who received the email has the problem of figuring out which John Smith he should forward the email to.

One way of solving this problem is that once a company hired someone with a particular name, they just wouldn't hire another. I<sub>2</sub> (with the name Radia Perlman, which is probably unique in the entire world) 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

## 1.10 AUTHENTICATION AND AUTHORIZATION

Authentication is when Alice proves to Bob that she is Alice. Authorization is having something decide what a requester is allowed to do at service Bob.

### 1.10.1 ACL (Access Control List)

Typically the way a server decides whether a user 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 George and Carol are allowed to read and write it. This database is often referred to as an **ACL (access control list)**.

### 1.10.2 Central Administration/Capabilities

Another model of authorization is that instead of listing, with each resource, the set of authorized users and their rights (*e.g., read, write, execute*), service Bob might have a database that listed, for each user, everything she was allowed to do. If everything were a single application, then the ACL model and the central administration model would be basically the same, since in both cases there would be a database that listed all the authorized users and what rights each had. But in a world in which there are many resources, not all under control of the same organization, it would be difficult to have a central database listing what each user was allowed to do. This model would have scaling problems if there were many resources each user was allowed to access, especially if resources

were created and deleted at a high rate. And if resources are under control of different organizations, there wouldn't be a single organization trusted to manage the authorization information.

Some people worry that ACLs don't scale well if there are many users allowed access to each resource. But the concept of groups helps the scaling issue.

### 1.10.3 Groups

Suppose there were a file that should be accessible to, say, any Dell employee. It would be tedious to type all the employee names into that file's ACL. And for a set of employees such as "all Dell employees", there would likely be many resources with the same set of authorized users. And it would take a lot of storage to have such huge ACLs on that many resources, and whenever anyone joined or left the company, all those ACLs would have to be modified.

The concept of a group was invented to make ACLs more scalable. It is possible to include a group name on an ACL, which means that any member of the group is allowed access to the resource. Then, group membership can be managed in one place, rather than needing to update ACLs at every resource when the membership changes.

Traditionally, a server that protected a resource with a group named on the ACL needed to know all the members of the group, but if there are many servers that store resources for that group, this would be inefficient and inconvenient. Also, that model would preclude more flexible group mechanisms; for example:

- cross-organizational groups, where no one server is allowed to know all the members;
- anonymous groups, where someone can prove membership in the group without having to divulge their identity.

Traditionally, groups were centrally administered, so it was easy to know all the groups to which a user belonged, and the user would not belong to many groups. But in many situations, it is useful for any user to be able to create a group (such as *Alice's friends*, or *students who have already turned in their exams in my course*), and have anyone be able to name such a group on an ACL.

Scaling up this simple concept of users, groups, and ACLs to a distributed environment has not been solved in practice. This section describes various ways that it might be done and challenges with truly general approaches.

### 1.10.4 Cross-Organizational and Nested Groups

It would be desirable for an ACL to allow any Boolean combination of groups and individuals. For instance, the ACL might be the union of six named individuals and two named groups. If someone is one of the named individuals, or in one of the groups, the ACL would grant them permission. It

---

---

would also be desirable to have the ACL be something like Group A and NOT group B, *e.g.*, U.S. citizen and not a felon.

Likewise, it would be desirable for group membership to be any Boolean combination of groups and individuals, *e.g.*, the members of Alliance-executives might be CompanyA-execs, CompanyB-execs, and John Smith. Each of the groups Alliance-executives, CompanyA-execs, and CompanyB-execs is likely to be managed by a different organization, and the membership is likely to be stored on different servers. How, then, can a server (Bob) that protects a resource that has the group Alliance-executives on the ACL know whether to allow Alice access? If she's not explicitly listed on the ACL, she might be a member of one of the groups on the ACL. But Bob does not necessarily know all the members of the group. Let's assume that the group name (Alliance-executives) can be looked up in a directory to find out information such as its network address and its public key. Or perhaps the group name would contain the DNS name of the server that manages that group, so the group name might be `example.com/Alliance-executives`.

- Bob could periodically find every group on any ACL on any resource it protects, and attempt to collect the complete membership. This means looking up all the members of all subgroups, and subgroups of subgroups. This has scaling problems (the group memberships might be very large), performance problems (there might be a lot of traffic with servers querying group membership servers for membership lists), and cache staleness problems. How often would this be done? Once a day is a lot of traffic, but a day is a lot of time to elapse for Alice's group membership to take effect, and for revocations to take effect.
- When Alice requests access, Bob could then ask the on-line group server associated with the group whether Alice is a member of the group. This could also be a performance nightmare with many queries to the group server, especially in the case of unauthorized users creating a denial of service attack by requesting access to services. At the least, once Alice is discovered to either belong or not belong, Bob could cache this information. But again, if the cache is held for a long time, it means that membership can take a long time to take effect, and revocation can also take a long time to take effect.
- When Alice requests access to a resource, Bob could reply, "You are not on the ACL, but here are a bunch of groups that are on the ACL, so if you could prove you are a member of one of those groups, you can access it." Alice could then contact the relevant group servers and attempt to get certification of group membership. Then, she can reply to Bob with some sort of proof that she is a member of one of the groups on the ACL. Again, this could have a cache staleness problem, and a performance problem if many users (including unauthorized troublemakers) contact group servers asking for proof of group membership.

### 1.10.5 Roles

The term *role* is used in many different ways. Authorization based on roles is referred to as **RBAC** (**role-based access control**). In most usage today, a **role** is the same as what we described as a group. In some cases, Alice will need to log in as the role rather than as the individual. For example, she might log in as `admin`. However, many users might need to log in as `admin`, and for auditing purposes, it is important to know which user was invoking the `admin` role. Therefore, for auditing purposes, a user might be simultaneously logged in as the role (`admin`) and as the individual (Alice). In many environments, a role is given a name such as “`admin`”, and being able to invoke the `admin` role on your own machine should not authorize you to invoke the `admin` role on controlling a nuclear power plant. So, both groups and roles should have names that specify the domain, *e.g.*, the DNS name of the service that manages the group or role membership.

Although a lot of systems implement roles as the same thing as groups, it is confusing to have two words that mean the same thing. We recommend that the difference between a group and a role is that a role needs to be consciously invoked by a user, often requiring additional authentication such as typing a different password. In contrast, we recommend that with a group, all members automatically have all rights of the group, without even needing to know which groups they are a member of. With roles, users may or may not be allowed to simultaneously act in multiple roles, and perhaps multiple users may or may not be allowed to simultaneously act in a particular role (like *President of the United States*). Again, we are discussing various ways things could work. Deployed systems make different choices. Some things people would like to see roles solve:

- When a user is acting in a particular role, the application presents a different user interface. For instance, when a user is acting as *manager*, the expense reporting utility might present commands for approving expense reports, whereas when the user is acting as *employee*, the application might present commands for reporting expenses.
- Having roles enables a user to be granted a subset of all the permissions they might have. If the user has to explicitly invoke the privileged role, and only keeps themselves authenticated as that role for as long as necessary, it is less likely that a typo will cause them to inadvertently do an undesirable privileged operation.
- Allowing a user to be able to run with a subset of her rights (not invoking her most privileged role except when necessary) gives some protection from malicious code. While running untrusted code, the user should be careful to run in an unprivileged role.
- Sometimes there are complex policies, such as that you are allowed to read either file A or file B but not both. Somehow, proponents of roles claim roles will solve this problem. This sort of policy is called a **Chinese wall**.

## 1.11 MALWARE: 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, so in the book we'll refer to all kinds of malicious software generically as **malware**. However, here are 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.
- **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. Often people forget to take these out when the product ships. However, sometimes these undocumented “features” are intentionally put into software by an employee who thinks he might at some point become disgruntled.
- **bot**—a machine that has been infected with malicious code that can be activated to do some malicious task by some controller machine across the Internet. The controller is often referred to as a **bot herder**. The intention of the disease-infected rodents (apologies if we offend actual disease-infected rodents who might be reading this book) who turn a computer into a bot is that the owner of the machine should not be aware that their machine is infected. To the owner of the machine, it continues operating as usual. However, the controller can at some point rally all the bots under its control to do some sort of mischief like sending out spam or flooding some service with nuisance messages. The bots under its control are often referred to as a **bot army**. There are price lists on the dark web for renting a bot army, priced according to the size of the army and the amount of time you would like to rent them. (The **dark web** is a subset of the Internet that is not visible to search engines, and requires special access mechanisms. It is an ideal place for purchasing illegal goods.)
- **logic bomb**—malicious instructions that trigger on some event in the future, such as a particular time occurring. The delay might advantage criminals that create a logic bomb, because

they can first deploy the logic bomb in many places. Or, they can make the bad event occur long after they have left the company, so they won't be under suspicion.

- **ransomware**—malicious code that encrypts the user's data, and then the helpful criminal offers to help get the data back, for a small fee (such as several million dollars).

Most of the items above exploit bugs in operating systems or applications. There are also vulnerabilities that exploit bugs in humans. One example is **phishing**, the act of sending email to a huge unstructured list of email addresses, that will trick some small percentage of the recipients into infecting their own machines, or divulging information such as a credit card number. **Spear phishing** means sending custom messages to a more focused group of people.

### 1.11.1 Where Does Malware Come From?

Where do these nasties come from? Originally, it was hobbyists just experimenting with what they could do. Today, big money can be made with malware. A bot army can be rented to someone that wants to damage a competitor, or a political organization they disagree with. And malware can be used for demanding ransoms. It is also used by sophisticated spy organizations. A notable example is Stuxnet. This was malware aimed specifically at damaging a certain type of equipment (centrifuges) used in Iran, to slow Iran's ability to produce nuclear weapons.

How could an implementer get away with intentionally writing malware into a program? Wouldn't someone notice by looking at the program? A good example of how hard it can be to decipher what a program is doing, even given the source code, is the following nice short program, written by Ian Phillipps (see Figure 1-2).

It 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 ( $I_1$  have analyzed the thing carefully and  $I_2$  have complete faith in me<sub>1</sub>), but we doubt many people would take the time to understand this program before running it

But also, 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 claimed advantage of 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



```

/* 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,
"@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#'1 \
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}'/{1,+K {rw' iK{:[nl]'/w#q#n'wk nw' \
iwk{KK{nl}!/w{%l###w#' i; :{nl}'/*{q#'ld;r'}{nlwb}/*de}'c \
; ;{nl}-{r}w}'+,}##*')#nc, '#nw}'/+kd'+e}+;#'rdq#w! nr'/ ' )}{xl}'{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-O;m .vpbks,fxntdCeghiry"),a+1);
}

```

Figure 1-2. Christmas Card?

- 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.

## 1.11.2 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 oldest form of virus checker knows instruction sequences that have appeared in known viruses, but are believed not to occur in benign code. It 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 include the newest viruses.

To evade detection of their viruses, virus creators have devised what are known as a **polymorphic viruses**. When they copy themselves, they change the order of their instructions or change instructions to functionally similar instructions. A polymorphic virus may still be detectable, but it takes more work, 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.

Some viruses attack the virus checkers rather than just trying to elude them. If an attacker can penetrate a company that disseminates code, such as virus signatures or program patches, they can spread their malware to all the customers of that company.

## 1.12 SECURITY GATEWAY

A security gateway (see Figure 1-3) sits between your internal network and the rest of the network and provides various services. Such a box usually provides many functions, such as firewall (§1.12.1), web proxy (§1.12.2), and network address translation (§1.14). We will describe these features in the next few sections.




Figure 1-3. Security Gateway

### 1.12.1 Firewall

There was a time when people assumed that their internal network and all the users and systems on the internal network were trustworthy. The only challenge was connectivity to the Internet. Users

need to communicate with publicly available services. The company needs to make some of its own services available to customers located outside the corporate network. So the thought was to install a box between your network and the scary Internet that can keep things secure, somehow.

The belief that a firewall between your internal network and the scary Internet is all you need to protect you has long been discredited. Even if you had the most secure firewall, *i.e.*, a box that blocked any traffic between your network and the Internet, malware inside your network could do arbitrary damage. And even carefully configured firewalls often break legitimate usage.

A recent buzzword is **zero trust**. It basically means the opposite of the previous thinking. Although firewalls can add some amount of extra security, all applications must protect themselves. They must authenticate everything they are talking to and enforce access control rules. The buzzword **defense-in-depth** means having multiple stages of security, so if one fails, hopefully another will protect you. Although people pretty much agree that the old firewall-will-protect-you model is no longer the right way to think about security, a lot of the mechanisms supporting that model are still deployed.

Firewalls centrally manage access to services in ways that individual systems should, but often don't. Firewalls can enforce policies such as *systems outside the firewall can't access file services on any systems inside the firewall*. With such a restriction, even if the internal systems are more open than they should be, or have bugs, they can't be attacked directly from systems outside the firewall.

Note that a firewall need not be a physical box that the company buys. It could instead be software on each machine that does the same sorts of filtering that a firewall box would do. This concept is sometimes called a **distributed firewall**. Both "firewall" and "distributed firewall" are buzzwords, and, as with most buzzwords, the definitions evolve and are used by different vendors in different ways. Usually *distributed firewall* implies that all the components doing "firewall stuff" are centrally managed. Otherwise, it would just be multiple firewalls.

The simplest form of firewall selectively discards packets based on configurable criteria, such as addresses in the IP header, and does not keep state about ongoing connections. For example, it might be configured to only allow some systems on your network to communicate outside, or some addresses outside your network to communicate into your network. For each direction, the firewall might be configured with a set of legal source and destination addresses, and it drops any packets that don't conform. This is known as **address filtering**.

Packet filters usually look at more than the addresses. A typical security policy is that for certain types of traffic (*e.g.*, email, web surfing), the rewards outweigh the risks, so those types of traffic should be allowed through the firewall, whereas other types of traffic (say, remote terminal access), should not be allowed through.

To allow certain types of traffic between host A and B while disallowing others, a firewall can look at the protocol type in the IP header, at the ports in the layer 4 (TCP or UDP) header, and at anything at any fixed offset in the packet. For web traffic, either the source or destination port will likely be 80 (http) or 443 (https). For email, either the source or destination port will likely be 25.

Firewalls can be even fancier. Perhaps the policy is to allow connections initiated by machines inside the firewall, but disallow connections initiated by machines outside the firewall. Suppose machine A inside the firewall initiates a connection to machine B outside the firewall. During the conversation, the firewall will see packets from both directions (A to B as well as B to A), so it can't simply disallow packets from B to A. The way it manages to enforce only connections initiated by A, is to look at the TCP header. TCP has a flag (called ACK) that is set on all but the first packet, the one that establishes the connection. So if the firewall disallows packets from B without ACK set in the TCP header, then it will usually have the desired effect.

Another approach is a **stateful packet filter**, *i.e.*, a packet filter that remembers what has happened in the recent past and changes its filtering rules dynamically as a result. A stateful packet filter could, for instance, note that a connection was initiated from inside using IP address  $s$ , to IP address  $d$ , and then allow (for some period of time) connections from IP address  $d$  to IP address  $s$ .

## 1.12.2 Application-Level Gateway/Proxy

An application-level gateway, otherwise known as a **proxy**, acts as an intermediary between a client and the server providing a service. The gateway could have two network adaptors and act as a router, but more often it is placed between two packet-filtering firewalls, using three boxes (see Figure 1-4). The two firewalls are routers that refuse to forward anything unless it is to or from the gateway. Firewall  $F_2$  refuses to forward anything from the global net unless the destination address is the gateway and refuses to forward anything to the global net unless the source is the gateway. Firewall  $F_1$  refuses to forward anything from your network unless the destination address is the gateway, and refuses to forward anything to your network unless the source address is the gateway. To transfer a file from your network to the global network, you could have someone from inside transfer the file to the gateway machine, and then the file is accessible to be read by the outside world. Similarly, to read a file into your network, a user can arrange for it to first get copied to the gateway machine. To log into a machine in the global network, you could first log into the gateway machine, and from there you could access machines in the remote network. An application-level gateway is sometimes known as a **bastion host**. It must be implemented and configured to be very




Figure 1-4. Application-Level Gateway

secure. The portion of the network between the two firewalls is known as the **DMZ** (demilitarized zone).

The gateway need not support every possible application. If the gateway does not support an application, either the firewalls on either side of the gateway should block that application (based on layer 4 port), or both firewalls would need to allow that application, depending on whether you want that application to work through the firewall. An example strategy is to allow only electronic mail to pass between your corporate network and the outside world. The intention is to specifically disallow other applications, such as file transfer and remote login. However, electronic mail can certainly be used to transfer files. Sometimes a firewall might specifically disallow very large electronic mail messages, on the theory that this will limit the ability to transfer files. But often, large electronic mail messages are perfectly legitimate, and any file can be broken down into small pieces.


Sometimes an application might specifically be aware of proxies running on an application gateway. For example, browsers can be configured with the address of a proxy, and then all requests will be directed to the proxy. The proxy might be configured with which connections are allowed and which are disallowed, and could even inspect the data passing through to check for malware.

### 1.12.3 Secure Tunnels

A *tunnel* (see Figure 1-5) is a point-to-point connection created between nodes A and B, where the connection is a path across a network. A and B can treat the tunnel as a direct link between each other. An *encrypted tunnel* is when A and B establish a secure session. An encrypted tunnel is sometimes called a **VPN** (virtual private network). We think that's a really bad term, and a more accurate term would be **VPL** (virtual private link). Whoever decided to call it a VPN imagined that the encrypted tunnel becomes an extra link in your private network, so a network consisting of a combination of private links and encrypted tunnels across the Internet becomes your private network. If we just use the acronym VPN for the encrypted tunnel (like the industry seems to) and don't think about what VPN expands to, I<sub>2</sub> guess we can live with the term.

Suppose the only reason you've hooked into the Internet is to connect disconnected pieces of your own network to each other. Instead of the configuration in Figure 1-5, you could have bought dedicated links between  $G_1$ ,  $G_2$ , and  $G_3$ , and trusted those links as part of your corporate network because you owned them. But it's likely to be cheaper to have the  $G$ s pass data across the Internet. How can you trust your corporate data crossing over the Internet? You do this by configuring  $G_1$ ,  $G_2$ , and  $G_3$  with security information about each other (such as cryptographic keys), and creating secure tunnels between them.

The mechanics of the tunnel between  $G_1$  and  $G_2$ , from the IP (network layer 3) point of view, is that when A sends a packet to C, A will launch it with an IP header that has source=A, destination=C. When  $G_1$  sends it across the tunnel, it puts it into another envelope, *i.e.*, it adds an



**Figure 1-5.** Connecting a Private Network over a Public Internet

additional IP header, treating the inner header as data. The outer IP header will contain source= $G_1$  and destination= $G_2$ . And all the contents (including the inner IP header) will be encrypted and integrity protected, so it is safe to traverse the Internet.

### 1.12.4 Why Firewalls Don't Work

Firewalls alone (without also doing end-to-end security) assume that all the bad guys are on the outside, and everyone inside can be completely trusted. This is, of course, an unwarranted assumption. Should employees have access (read and write) to the salary database, for instance?

Even if the company is so careful about hiring that no employee would ever do anything bad intentionally, firewalls can be defeated if an attacker can inject malicious code into a machine on the corporate network. This can be done by tricking someone into downloading something from the Internet or launching an executable from an email message. It is quite common for an attacker to break into one system inside your firewall, and then use that system as a platform for attacking other systems. Someone once described firewall-protected networks as “hard and crunchy on the outside; soft and chewy on the inside.”

Firewalls often make it difficult for legitimate users to get their work done. The firewall might be configured incorrectly or might not recognize a new legitimate application. And if the firewall allows one application through (say email or http), people figure out how to do what they need to do by disguising it as traffic that the firewall is configured to allow. The ironic term for disguising traffic in order to fool a firewall is to carry your traffic in a **firewall-friendly** protocol. Since firewalls commonly allow http traffic (since it's the protocol used for browsing the web), there are many proposals for doing things over http. The most extreme example is to carry IP over http, which would allow any traffic through! Firewall-friendly? The whole point is to defeat the best efforts of the firewall administrator to disallow what you are doing! It isn't somehow “easier” for the firewall

to carry http traffic than any other. The easiest thing for the firewall would be to allow everything or nothing through!

Just as breaking a large file into lots of pieces to be individually carried in separate emails is inefficient, having protocols run on top of *http* rather than simply on top of IP is also inefficient in terms of bandwidth and computation.

## 1.13 DENIAL-OF-SERVICE (DOS) ATTACKS

A **denial-of-service attack (DoS)** is one in which an attacker prevents good guys from accessing a service, but does not enable unauthorized access to any services. In the naive old days, security people dismissed the prospect of denial-of-service attacks as unlikely, since the attacker had nothing to gain. Of course, that turned out to be faulty reasoning. There are terrorists, disgruntled employees, and people who delight in causing mischief for no good reason.

In the earliest types of denial-of-service attacks, the attacker repeatedly sent messages to the victim machine. Most machines at the time were vulnerable to this sort of attack since they had resources that could easily be depleted. For instance, the storage area for keeping track of pending TCP connections tended to be very limited, on the order of, say, ten connections. The probability of ten legitimate users connecting during a single network round trip time was sufficiently small that ten was a reasonable number. But it was easy for the attacking machine to fill up this table on a server, even if the attacking machine was attached to the Internet with a low-speed link.

To avoid being caught at this mischief, it was common for the attacker to send these malicious packets from forged source addresses. This made it difficult to find (and prosecute) the attacker, and it made it difficult to recognize packets from the malicious machine and filter them at a firewall.

As a defense, people advocated having routers have the capability of doing sanity checks on the source address. These routers could be configured to drop packets with a source address that could not have legitimately come from the direction from which the packet was received. Routers might be configured with which source addresses to expect on each of their ports, or they might infer the expected direction from their forwarding tables. This concept was not deployed because it would cause problems. If sanity checks are based on configured information, topological changes in the Internet (such as links going down and alternative routes being used) could cause the routers to make incorrect assumptions. And Mobile IP (RFC 5944) allows a node to move around in the Internet and keep its IP address, which would confuse routers attempting sanity checks.

A deployed defense against a single malicious node attempting to swamp the resources of a server was to increase resources at the server so that a single attacker, at the speeds at which such attackers were typically connected to the Internet, could not fill the pending TCP connection table.

Another level of DoS escalation was to send a single packet that caused a lot of legitimate machines to send messages to the victim machine. An example of such a packet is a packet transmitted to the broadcast address, with the packet's source address forged to the address of the victim's machine, asking for all receivers to respond. All the machines that receive the broadcast will send a response to the victim's machine. Such a mechanism magnifies the effect the attacker can have from his single machine, since each packet he creates turns into  $n$  packets directed at the victim machine.

As a defense against machines sending packets from forged IP addresses, protocols such as TCP, IPsec, and TLS have been designed to avoid requiring a receiver, Bob, to keep state or do significant computation if requests are arriving from forged IP source addresses. Unless the requester can receive packets at the IP address they claim to be coming from, Bob will not need to keep state about the request. Only when the requester returns something that Bob sent to its claimed IP address will Bob pay attention to this request.

But then came the next level of escalation, which is known as a **distributed-denial-of-service attack (DDoS)**. In this form of attack, the attacker breaks into a lot of innocent machines, and installs software on them to have them all attack the victim machine. These innocent machines are called **zombies** or **drones** or **bots**. With enough bots attacking it, any machine can be made inaccessible, since even if the machine itself can process packets as fast as they can possibly arrive, the links or routers in front of that machine can be overwhelmed. The defenses in TCP, IPsec, and TLS will not help, since the bot machines are using their own IP addresses in the requests. Since requests are coming from hundreds or thousands of innocent machines, it is hard to distinguish these packets from packets coming from legitimate users.

## 1.14 NAT (NETWORK ADDRESS TRANSLATION)

NAT was designed, out of necessity, because IPv4 addresses were too small (four bytes) to give unique addresses to every node on the Internet. With NAT, a piece of the Internet (say a corporate network) can use IP addresses that are not globally unique, and, in fact, these addresses are reused in many other networks. This means that a node inside such a network cannot be contacted from outside that network. However, if the security gateway provides NAT functionality, it will have a pool of globally unique IP addresses that can be assigned as needed.

The NAT box will almost certainly not have a large enough pool of IP addresses to give a globally reachable IP address to every internal node communicating to outside nodes. So the NAT box also translates the TCP or UDP port as well. So a NAT implementation might have a mapping of  $\langle \text{internal IP, internal port} \rangle$  maps to  $\langle \text{external IP, external port} \rangle$ . When internal node Alice sends a packet to external destination Bob, the NAT box will replace the source IP and port to the external



---

IP and port in the NAT box's table. Likewise, when packets arrive from Bob for Alice's assigned external  $\langle \text{IP}, \text{port} \rangle$ , the NAT box replaces these fields in the destination fields in Bob's packet before forwarding the packet on the internal network.

There is somewhat of a security problem with this approach if the NAT box simply translates tuples of  $\langle \text{IP address}, \text{port} \rangle$ . Suppose Alice starts a connection to Bob, and the NAT box then creates an entry to translate Alice's internal address and port to, for instance, globally reachable tuple  $\langle \text{IP}_{\text{Alice}}, \text{Port}_{\text{Alice}} \rangle$ . If the NAT box will forward anything to Alice that is addressed to Alice's temporarily assigned global address and port, then any node on the Internet could send a packet to Alice by addressing it to  $\langle \text{IP}_{\text{Alice}}, \text{Port}_{\text{Alice}} \rangle$ . Sometimes this is the desired behavior. For example, assume there is a conferencing system. Alice, George, and Carol (all behind NATs) join the conference by contacting the central server, but it is not desirable for all of the conference communication to go through the central server. If the NAT box allows anyone that knows Alice's temporary global address to contact Alice, then the conference coordinator can tell all of the members the other members' global addresses, and they can then directly communicate with each other.

To create the behavior that only the node that Alice has initiated a connection to, to be able to reach Alice, then the NAT box will keep a mapping of 4-tuples to  $\langle \text{external IP}, \text{port} \rangle$  pairs. For example, if Alice, at internal  $\text{IP}=\text{a}$ , internal  $\text{port}=\text{p}$  initiates a connection to external Bob, at IP address= $\text{B}$ ,  $\text{port}=\text{P}_\text{B}$ , the NAT box might assign Alice, for this connection, the external address and port  $\langle \text{IP}_\text{A}, \text{Port}_\text{A} \rangle$ . The NAT table would include Bob's address in the mapping, and only allow packets from Bob's address and port to be forwarded to Alice. So the NAT entry would have a six-tuple  $\langle \text{B}, \text{P}_\text{B}, \text{IP}_\text{A}, \text{Port}_\text{A}, \text{a}, \text{p} \rangle$ , meaning that only packets from Bob (at  $\langle \text{B}, \text{P}_\text{B} \rangle$ ) would be translated and forwarded to Alice.

Another fortuitous use of NATs is for all the devices inside a home to have the same IP address. This is useful because some ISPs (Internet Service Providers) charged for Internet connectivity per device, and the NAT box made it appear to the ISP as if there were only a single device in the house. The ISP's answer to this threat to their pricing model was to include in the 74-page EULA (end user license agreement) that everyone has to click on (but nobody reads), an agreement by the user not to use a NAT box. Now, if an actual human read the 74-page agreement, they would most likely think "What's a NAT box?"

### 1.14.1 Summary

These are the main concepts in the Internet. We will give more details about these as they come up in the book.

The Internet has evolved from the original design, and the design was never the only or best way to design a network, but the industry has made it work. An analogy is the English language. It might be overly complicated, with all the spelling and grammar exceptions, but it does the job. And

each year some mysterious panel of people decide which new words should officially be added to English, and how the grammar rules should change.

Similarly, for the Internet. If there is anything it can't do, at least so far, the world has figured out how to evolve the Internet to do what is needed.

*This page intentionally left blank*

# INDEX

## Numerics

3DES, 61, 72

## A

Abel, Niels Henrik, M-4

Abelian group, M-4

absorb phase, 123

Abstract Syntax Notation 1. *See* ASN.1

access control list, 14

ACL, 14

active attack, 7

address filter, 22

address-based authentication, 253–255

adjugate, M-24

Adleman, Leonard, 140

Advanced Encryption Standard, 75–81,  
M-16–M-18

AES, 61, 75–81, M-16–M-18

AES-GCM, 344

AH, 349, 356–365

alternating group, M-27

amplitude, 171

ancilla, 184

anonymity, 397

append attack, 115

application level gateway, 23

ASN.1, 151

associated data, 101

associativity, M-3

asymmetric, 4

attack

active, 7

chosen-ciphertext, 206

chosen-plaintext, 46, 85

ciphertext-only, 45

denial-of-service, 26–27, 331, 335

dictionary, 260

distributed-denial-of-service, 27

downgrade, 345

known-plaintext, 46

meet-in-the-middle, 73

million-message, 152, 206

multi-target, 212

off-line password-guessing, 260–261

on-line password-guessing, 257–260

passive, 7

recognizable-plaintext, 45

reflection, 320–322

side-channel, 55

small- $n$ , 264

small-subgroup, 157

splicing, 441

attestation, 271, 427

augmented strong password protocol, 275–276

authenticated Diffie-Hellman exchange, 156

authentication, 31, 37

address-based, 253–255

cryptographic, 255

multifactor, 250

mutual, 320–322

one-way, 314–319

password-based, 251–253

public key, 41

strong, 37

using hash, 114

authentication facilitator node, 252

authentication storage node, 252

authentication token, 268–272

## B

back door, 67

bad guy, 32

bad-list, 304

BASE64, 384

BASE85, 384

basis, 216

- bastion host, 23
- Bellovin, Steven, 272
- BGP, 254
- big-endian, 149
- Biham, Eli, 71, 108
- BIKE, 230
- biometrics, 255
- birthday problem, 108, 439
- bit shuffle, 65
- Bitcoin, 401, 406–413
- Bitcoin wallet, 407
- bit-slicing, 131
- Bleichenbacher, Daniel, 152
- blind decryption, 419
- blind signature, 418
- block cipher, 61
- block modifier, 93
- blockchain, 112, 406
- Bmod, 93
- Boneh, Dan, 392
- bootstrapping, 432
- Bosselaers, Antoon, 108
- bot, 18, 27
- C**
- CA, 281
- CAC, 270
- Caesar cipher, 34
- Captain Midnight secret decoder rings, 34
- Carmichael number, 145
- CBC residue, 96
- CBC-MAC, 96
- CCM, 100
- cellular phone, 251
- certificate, 281
- certificate authority, 281
- challenge, 37
- channel binding, 329
- CHAP, 251
- characteristic, M-14
- Chaum, David, 397, 401, 418
- checksum, 38, 452
- Chinese remainder theorem, 57, 147, 187, M-8
- Chinese wall, 17
- chosen-ciphertext attack, 206
- chosen-plaintext attack, 46, 85
- ciphertext, 31
- ciphertext stealing, 94
- ciphertext-only attack, 45
- circuit, 428
- circuit privacy, 431
- circulant matrix, 218, 239
- class of service, 442
- cleartext. *See* plaintext
- clone, 251
- CMAC, 97
- CNOT, 174, 175
- CNOT gate, 179
- code-based cryptography, 229–239
- codeword, 229
- coefficient, 54, 171, 174, 218, 219, 241, 418, M-11
- collision resistance, 105
- commitment, 420
- commutative group, M-4
- commutative ring, M-4
- commutativity, M-3
- composition, M-9
- compression, 450
- compression function, 120
- computational difficulty, 32
- confidentiality, 5
- confusion layer, 77
- conjugate transpose, M-33
- constant, M-11
- constrained quadratic polynomial, 242
- controlled not, 174
- conventional cryptography, 36
- cookie, 12, 266, 336
- coprime, M-5
- counter mode, 91
- cover traffic, 398
- credential, 277
- credentials download protocol, 277–278
- critical bit, 454

- 
- 
- CRL, 302  
cryptogram, 35  
cryptographic checksum. *See* message integrity check  
cryptographic hash, 105–134  
cryptography, 31–56  
CRYSTALS-Kyber, 228  
CTR, 91  
cyclic group, M-4, M-9
- D**
- DAA, 427  
Daemen, Joan, 76  
DANE, 390  
Danish sugar beet auction, 429  
dark web, 18  
data loss prevention, 387  
data variable, 51  
Davies-Meyer, 121  
DDoS, 27  
decryption, 31  
decryption failure rate, 206  
deduplication, 111  
degree, M-11  
 $\delta_{ij}$ . *See* Kronecker delta  
delta CRL, 303  
den Boer, Bert, 108  
denial-of-service attack, 26–27, 331, 335  
derivative, M-30  
DES, 61, 67–75, 78  
    mangler, 69  
    multiple key, 72  
    substitution, 70  
destination port., 11  
determinant, M-23–M-25  
DFT, M-34  
diagonal matrix, M-24  
dictionary attack, 260  
differential cryptanalysis, 71  
Diffie-Hellman, 154–158, 163, 333  
diffusion layer, 77  
Digicash, 401  
digital rights management, 389  
digital signature, 39, 42  
Digital Signature Standard. *See* DSS  
Dilithium, 224  
dimension, M-20  
direct anonymous attestation, 427  
discrete Fourier transform, 189–195, M-34  
discrete logarithm, 158, 163, 165, M-4, M-19  
dishonest ciphertext, 206  
distributed firewall, 22  
distributed-denial-of-service attack, 27  
distribution list, 381  
distributivity, M-3  
divide, M-11  
division ring, M-4  
divisor, M-5  
DKIM, 389  
DLP, 387  
DMZ, 24  
DNS, 11  
DNS-Based Authentication of Named Entities, 390  
DNSSEC, 389  
domain separator, 126  
DomainKeys Identified Mail, 389  
DoS, 26–27  
double TGT authentication, 288  
downgrade attack, 345  
drone, 27  
DSA. *See* DSS  
DSS, 9, 160–163  
DTLS, 332
- E**
- eCash, 401, 402–405  
ECB, 86  
ECC, 164, M-18  
ECDH, 164, 166  
ECDSA, 164, 167  
EDE, 72  
EKE, 272–275  
Eldridge, Al, 340  
electronic code book, 86  
electronic mail, 379–399

ElGamal, 159, 162  
ElGamal, Taher, 159  
elliptic curve cryptography, 164, M-18  
encryption, 31  
    large message, 86  
    using hash, 118  
endpoint identifier hiding, 337  
ephemeral public key, 334  
EPID, 428  
error vector, 225  
error-correcting code, 229  
escrow-foilage, 334  
ESP, 349, 356–365  
Euclidean algorithm, 56, 139, M-7, M-12  
Euler totient, 138, M-7, M-26  
Euler’s formula, 195  
Euler’s theorem, 139, 187  
evaluate, M-13  
explicit rejection, 207  
exponentiation, 55, 138, 142, M-8  
export control, 9, 443  
extendable output function, 122

**F**

factor, 141, 146, 168, 187–188, M-11  
Falcon, 224  
Federal Information Processing Standard, 76  
Feistel cipher, 65  
Fermat’s theorem, 139, 145  
FHE, 431–434  
Fiat-Shamir construction, 422  
FIDO, 270  
field, M-4, M-10–M-16  
field extension, M-11  
finite field, 54, M-13–M-16  
FIPS, 76  
firewall, 376  
firewall friendly, 25  
floating point, M-6  
forward compatibility, 453  
Fourier transform, 189–195, M-34  
Franklin, Matthew, 392  
frequency, 189

Frodo, 229  
fully homomorphic encryption, 431–434  
Fundamental Tenet of Cryptography, 32, 141,  
155, 162

## G

Galois field, 54, M-13–M-16  
Galois, Évariste, 54, M-7  
gate, 170, 172  
gateway, 23  
Gaussian elimination, M-22, M-32  
gcd, 56, M-7, M-12  
GCM, 101  
generator, 157, 274, M-9  
generator matrix, 231  
Gentry, Craig, 224, 432  
 $GF(p^n)$ . *See* Galois field  
GHASH, 98  
GMAC, 98, 99  
Goldreich, Goldwasser, Halevi, 217, 223  
Goldschlag, David, 397  
good guy, 32  
good-list, 305  
Goppa code, 233  
GPG, 379  
graph isomorphism, 420  
greatest common divisor, 56, M-7, M-12  
group, 15, M-4, M-9–M-10, M-28  
    Abelian, M-4  
    alternating, M-27  
    commutative, M-4  
    cyclic, M-4, M-9  
    symmetric, M-27  
group manager, 423  
Grover’s algorithm, 170, 180

## H

hacker, 4  
Hadamard gate, 175, 179  
Hamilton, William Rowan, M-19  
hash, 43, 105–134  
    *See also* message digest  
    keyed, 44, 116

- 
- 
- Lamport. *See* Lamport hash
- hash chain, 111
- hash tree, 113
- hash-based signature, 207–216
- Hermitian transpose, M-33
- HMAC, 116, 127–132, 447
- homomorphic encryption, 431–434
- honest verifier zero knowledge, 422
- HTTP, 11, 25, 266, 367, 372
- HTTPS, 11
- I**
- IBE, 392
- IBM, 67
- ideal cipher model, 63
- ideal lattice, 218
- identity, M-3, M-21
- identity-based encryption, 392
- IKE, 349
- immutable, 357, 361
- implicit rejection, 206
- index, M-10
- information-theoretic security, 35, 197
- initialization vector. *See* IV
- integrity. *See* message integrity
- inverse, 137, M-3, M-6, M-12, M-15, M-21
- InvMixColumns, M-16
- IPsec, 313, 330–332, 336, 339, 341, 344, 345, 349–365
- IPv4, 360
- IPv6, 360
- irreducible, 54, M-15
- IRS, 36, 382, 399, G-7
- isogeny-based cryptography, 204
- issuer, 290
- IV, 89
- J**
- Jablon, David, 273
- K**
- Karn, Phil, 265
- KDC, 281
- KDF, 51
- KECCAK, 110, 122
- Kelsey. John, 215
- KEM, 205
- Kerberos, 283–289
- ket, 171
- key, 32
- key derivation, 448
- key distribution, 390–393
- key distribution center, 281
- key encapsulation mechanism, 205
- key escrow, 282
- key expansion, 64
- key handle, 271
- key recovery, 282
- key rollover, 323, 333, 439
- key schedule, 64
- key seed, 448
- key-derivation function, 51
- keyed hash, 44, 116
- KMAC, 117
- known-plaintext attack, 46
- Kronecker delta, M-21
- L**
- $\lambda(n)$ . *See* reduced totient
- Lamport hash, 263–265
- lattice, 216
- lattice-based cryptography, 216–229
- leap of faith, 289
- learning with errors, 225–229
- Lightning Network, 413
- linear code, 234
- linear cryptanalysis, 71
- linear transformation, 76, M-20, M-31
- linearity, 176
- LMS, 214
- local exploder, 381
- logic bomb, 18
- logical qubit, 198
- Lotus Notes, 340, 379
- Lucifer cipher, 67
- LWE, 225–229



Lyubashevsky, Vadim, 224

## M

MAC, 38, 42, 85, 96  
    CBC residue, 96  
    using hash, 114  
mail reflector, 381  
mangler  
    DES, 69  
master key, 448  
master secret, 370  
matrix, 174, 216, 218, 225, 231, M-20–M-25  
McEliece, 230, 233  
McEliece, Robert, 233  
MD, 108  
MD2, 108  
MD3, 108  
MD4, 108, 120  
MD5, 108, 120  
MDPC, 230  
meddler-in-the-middle, 7, 155–156, 271  
meet-in-the-middle attack, 73, 121  
Merkle tree, 113  
Merkle, Ralph, 108  
Merkle–Damgård, 120  
Merritt, Michael, 272  
message authentication code, 38  
message digest, 43, 105–134  
    *See also* hash  
message flow confidentiality, 395  
message integrity, 43, 393  
    and privacy, 100  
message integrity check, 31  
    using secret key, 38  
message integrity code, 38  
message transfer agent, 383  
MIC, 38  
Miller-Rabin test, 145  
million-message attack, 152, 206  
miner (cryptocurrency), 406  
Misoczki, Rafael, 230  
MIT, 8  
MITM, 7, 155–156, 271, 328, 329, 354, 355

MixColumns, 80, M-16  
Mixmaster, 397  
Mixminion, 397  
mobile IP, 26, 254  
modes of operation, 61, 85–104  
modular arithmetic, 136–139, M-5–M-8  
modulus, 54, 56, 94, 140, 154, M-12  
monic, M-11  
monoalphabetic cipher, 35  
MPC, 429  
MTA, 383  
MTA Strict Transport Security, 390  
MTA-STS, 390  
multicast, 349  
multifactor authentication, 250  
multi-target attack, 212  
multivariate cryptography, 239–244  
mutable, 357, 361  
mutual authentication, 320–322

## N

name constraint, 296, 300  
NAT, 27  
National Bureau of Standards, 67  
National Institute of Standards and Technology,  
    67, 75  
NBS, 67  
Needham, Roger, 283  
negation, M-23  
Neiderreiter, Harald, 236  
network address translation, 27  
Nguyen, Phong, 224  
NIST, 67, 75, 160, 203  
NIST security strength category, 205  
nonce, 326–328, 449  
noninteractive zero knowledge proof, 422  
nonlinear transformation, 77  
non-repudiation, 42, 394  
notation, 6, M-2  
nothing up my sleeve, 67  
NP-hard, 170  
NSA, 67, 68, 76, 162  
NTRU, 219–223

**O**

OAEP, 153  
OAuth, 267  
obfuscated C code contest, 19  
oblivious transfer, 419  
OCSP, 302, 304  
off-line password-guessing attack, 260–261  
oil and vinegar, 242  
OLRS, 302, 304  
one-time pad, 35, 81  
one-time password, 263  
one-to-one, 62  
one-way authentication, 314–319  
onion routing, 397  
on-line certificate status protocol, 302  
on-line password-guessing attack, 257–260  
on-line revocation service, 302  
OpenID, 267  
option, 453  
oracle, 63, 105  
order, 157, M-9  
Ovaltine, 34

**P**

padding, 124–127  
PAKE, 272  
parity-check matrix, 234  
passive attack, 7  
password, 251–276  
    conversion to key, 256  
password equivalent, 253  
password hash, 43, 253  
password-based authentication, 251–253  
patent, 8  
P-box, 65  
PDM, 273, 275  
Peikert, Chris, 224  
PEM, 379  
per message secret  
    in DSS and ElGamal, 162  
perfect forward secrecy, 326, 333  
permutation, 65, M-9, M-22, M-26, M-32

PFS, 333  
PGP, 379  
phase, 172, 190, 193, 195, 200  
 $\phi(n)$ . *See* totient  
Phillipps, Ian, 19  
phishing, 19  
Photuris, 336  
 $\pi/8$  gate, 180  
PIV, 270  
PKCS, 151, 319  
PKG, 392  
PKI, 281–311  
PKIX, 290, 300, 367  
plaintext, 31  
plausible deniability, 343, 394, 395  
polarization, 173, 196  
polymorphic virus, 21  
polynomial, M-11–M-13  
polynomial system, 240  
pornographic screen saver, 398  
port forwarding, 376  
Post-it, 107  
post-quantum cryptography, 203–247  
preauthentication, 287  
preimage resistance, 105  
pre-master secret, 370  
PRF, 51  
prime, M-5  
    finding, 144  
    safe, 157  
    Sophie Germain, 157  
    testing, 145  
primitive, M-15  
principal, 281, 290  
principle of least privilege, 289  
privacy, 5  
private key, 162  
private key generator, 392  
PRNG, 50  
probability amplitude, 171  
proof of work, 412  
proxy, 23  
PRP, 63

pseudonymity, 398  
pseudorandom, 82, 118  
pseudorandom function, 51  
pseudorandom number generator, 50  
pseudorandom permutation, 63  
public key, 4, 39, 135–168  
    creation from password, 256  
public key infrastructure, 281–311  
Public-Key Cryptography Standard. *See* PKCS  
puzzle, 112  
Pyramid, 215

## Q

QKD, 195–197  
quadratic equation, 240  
quadratic polynomial, 240  
quantum computing, 169–195  
quantum encryption, 195–197  
quantum error correction, 199  
quantum gate, 171, 172  
quantum key distribution, 195–197  
quantum mechanics, 169  
quaternion, M-4, M-19  
quotient, M-5, M-11

## R

RA, 293  
RADIUS, 252  
Rainbow, 239  
rainbow table, 261  
random, 81, 452  
random number, 452  
random oracle model, 63  
ransomware, 19  
RBAC, 17  
RC4, 81  
realm, 283  
recognizable-plaintext attack, 45  
redirect, 12  
reduced totient, 140, 187, M-7, M-28  
Reed, Michael, 397  
reflection attack, 320–322  
Regev, Oded, 224

registration authority, 293  
related keys, 81  
relatively prime, 138, M-5  
relying party, 290  
remainder, M-5, M-11  
remote exploder, 381  
repudiation, 394  
response, 37  
revocation, 302–305  
Rijmen, Vincent, 76  
Rijndael, 75–81, M-16–M-18  
ring, M-4, M-19  
ring signature, 425  
ring with identity, M-4  
RIPEMD, 110  
Rivest, Ron, 9, 108, 140  
rock paper scissors, 132  
Rogaway, Phillip, 93  
role, 17  
root  
    of namespace, 11  
    of polynomial, M-13  
    of tree, 113  
    of trust, 297  
round, 64  
RSA, 8, 140–153, 163

## S

S/Key, 265  
S/MIME, 379  
SA, 349–352  
safe prime, 157, 275  
salt, 119, 213, 260, 264  
SAML, 267  
S-box, 64, 70, M-16  
Schroeder, Michael, 283  
second preimage resistance, 105  
secret key, 4, 36, 61–83  
secret sharing, 417  
secure enclave, 431  
secure hash algorithm. *See* SHA  
secure multiparty computation, 429  
secure session, 6

- 
- 
- security association, 349–352
  - security chain, 290
  - security parameter index, 350
  - security strength, 33
  - seed, 36, 49, 50, 51
  - semi-weak key, 70
  - Sender Policy Framework, 389
  - service class, 442
  - session key, 323–326
  - session resumption, 341, 372
  - SEV, 431
  - SGX, 431
  - SHA-0, 109
  - SHA-1, 109, 120
  - SHA-2, 109, 120
  - SHA-3, 109, 122–124
  - SHAKE, 122, 123
  - Shamir, Adi, 71, 108, 140, 392
  - share, 417
  - ShiftRows, 80
  - Shor’s algorithm, 166, 170, 186
  - side-channel attack, 55
  - signature, 1, 39, 159–163
    - hash-based, 207–216
  - SIKE, 204
  - skew field, M-4
  - small- $n$  attack, 264
  - small-subgroup attack, 157
  - smooth number, 150, 157
  - SMTP, 383
  - snake oil, 448
  - SNEFRU, 108
  - Snowden, Edward, 53
  - Sophie Germain prime, 157
  - source port, 11
  - spam, 385
  - spear phishing, 19
  - SPEKE, 273, 274
  - SPF, 389
  - SPHINCS+, 216
  - SPI, 350
  - SPKI, 291
  - splicing attack, 441
  - splitting field, M-14
  - SPN, 77
  - Spock, 132
  - SPONGE, 122
  - squeeze phase, 123
  - SRP, 273, 276
  - SSH, 313
  - SSL, 330, 367–377
  - STARTTLS, 390
  - stateful packet filter, 23
  - stream cipher, 36, 61, 81
  - strong authentication, 37
  - strong password protocol, 272–276
  - Stuxnet, 19
  - SubBytes, 80
  - subfield, M-11
  - subgroup, 157, M-9
  - substitution, 65
    - DES, 70
  - substitution-permutation network, 77
  - superposition, 171
  - symmetric, 4
  - symmetric cryptography, 36
  - symmetric group, M-27
  - syndrome, 232
  - systematic form, 230, 234
  - Syverson, Paul, 397
  - Szydlo, Michael, 224
- ## T
- TCG, 427
  - TCP, 10, 367
  - TDEA, 72
  - TDES, 72
  - terminology, 4
  - TGT, 285
  - ticket, 284
  - ticket granting ticket, 285
  - timestamp, 316, 327
  - TLD, 11
  - TLS, 313, 330, 367–377
  - TLV, 361, 453

token  
    challenge-response, 269  
    time-based, 268  
top-level domain, 11  
TOR, 397  
totient, 138, 187, M-7, M-26, M-28  
TPM, 427  
transmission control protocol, 10  
transport mode, 358  
trapdoor, 18  
triangular matrix, M-24  
triple data encryption, 72  
Trojan horse, 18  
trust anchor, 289, 290  
trust on first use, 289  
Trusted Computing Group, 427  
trusted platform module, 427  
trusted repeater, 197  
trusted third party, 281  
tunnel mode, 358  
two's complement, M-6

## U

UA, 383  
UDP, 11, 367  
unbalanced oil and vinegar, 239  
uniform resource locator, 12, 266  
unitary, 177, M-33  
universal quantum gate set, 179  
unspent transaction output, 409  
UOV, 239  
URL, 12, 266  
user agent, 383  
user datagram protocol, 11  
UTF8, 384  
UTXO, 409

## V

Vaikuntanathan, Vinod, 224  
Vandermonde matrix, M-33  
vector, 174, M-20  
verifier, 290  
version number, 454–455

virus, 18  
virus checker, 20  
VPN, 24

## W

weak key, 70  
Winternitz, Robert, 214  
work factor, 33  
worm, 18  
Wu, Tom, 276

## X

X.509, 290, 303, 305  
XEX, 93  
XMSS, 214  
XOF, 122  
XTS, 94

## Y

You-Tarzan-Me-Jane, 339

## Z

zero knowledge, 422  
zero trust, 22  
Zimmermann, Phil, 379  
zombie, 27