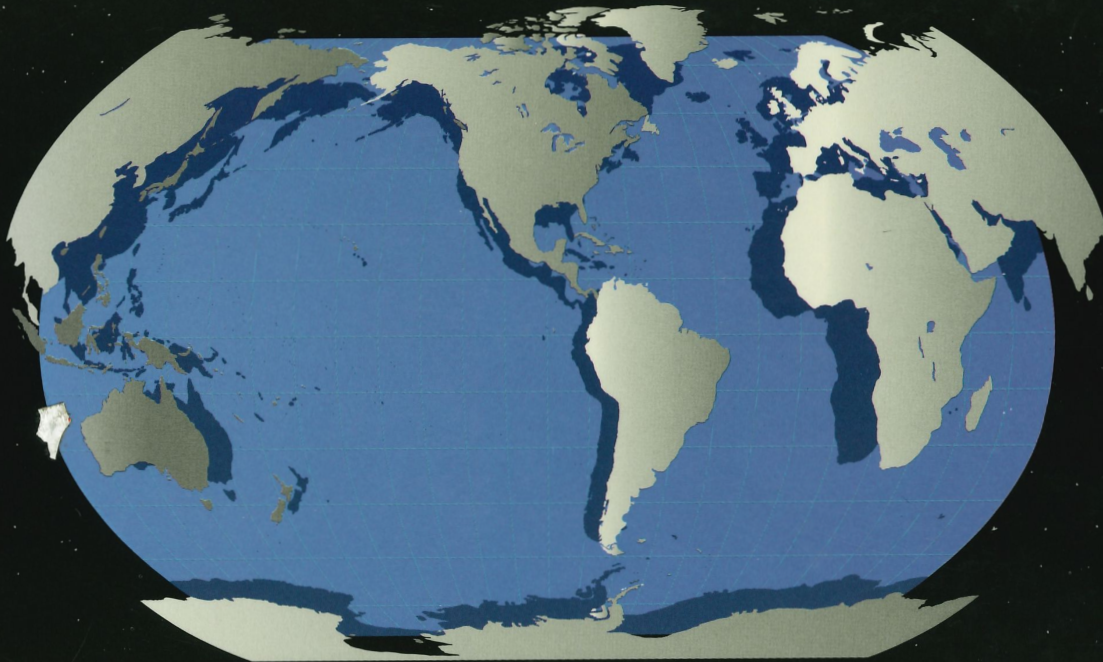


Third Edition

INTERNETWORKING WITH
TCP/IP

VOLUME I
PRINCIPLES, PROTOCOLS,
AND ARCHITECTURE

Includes
*IPng and IP
over ATM*



DOUGLAS E. COMER

Internetworking With TCP/IP
Vol I:
Principles, Protocols, and Architecture
Third Edition

DOUGLAS E. COMER

*Department of Computer Sciences
Purdue University
West Lafayette, IN 47907*



PRENTICE HALL
Englewood Cliffs, New Jersey 07632

Library of Congress Cataloging-in-Publication Data

Comer, Douglas
Internetworking with TCP/IP / Douglas E. Comer. -- 3rd ed.
p. cm.
Includes bibliographical references and index.
Contents: v. 1. Principles, protocols, and architecture
ISBN 0-13-216987-8 (v. 1)
1. TCP/IP (Computer network protocol) 2. Client/server computing.
3. Internetworking (Telecommunication) I. Title.
TK5105.585.C66 1995 95-1830
005.2--dc20 CIP

Acquisitions editor: ALAN APT
Production editor: IRWIN ZUCKER
Cover designer: WENDY ALLING JUDY
Buyer: LORI BULWIN
Editorial assistant: SHIRLEY MCGUIRE



©1995 by Prentice-Hall, Inc.
A Simon & Schuster Company
Englewood Cliffs, New Jersey 07632

All rights reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the publisher.

The author and publisher of this book have used their best efforts in preparing this book. These efforts include the development, research, and testing of the theories and programs to determine their effectiveness. The author and publisher make no warranty of any kind, expressed or implied, with regard to these programs or the documentation contained in this book. The author and publisher shall not be liable in any event for incidental or consequential damages in connection with, or arising out of, the furnishing, performance, or use of these programs.

UNIX is a registered trademark of UNIX System Laboratories, Incorporated
proNET-10 is a trademark of Proteon Corporation
LSI 11 is a trademark of Digital Equipment Corporation
Microsoft Windows is a trademark of Microsoft Corporation

Printed in the United States of America

10 9 8 7 6 5 4 3

ISBN 0-13-216987-8

Prentice-Hall International (UK) Limited, London
Prentice-Hall of Australia Pty. Limited, Sydney
Prentice-Hall Canada Inc., Toronto
Prentice-Hall Hispanoamericana, S.A., Mexico
Prentice-Hall of India Private Limited, New Delhi
Prentice-Hall of Japan, Inc., Tokyo
Simon & Schuster Asia Pte. Ltd., Singapore
Editora Prentice-Hall do Brasil, Ltda., Rio de Janeiro

Contents

Foreword	xxi
Preface	xxiii
Chapter 1 Introduction And Overview	1
1.1 <i>The Motivation For Internetworking</i>	1
1.2 <i>The TCP/IP Internet</i>	2
1.3 <i>Internet Services</i>	3
1.4 <i>History And Scope Of The Internet</i>	6
1.5 <i>The Internet Architecture Board</i>	8
1.6 <i>The IAB Reorganization</i>	9
1.7 <i>The Internet Society</i>	11
1.8 <i>Internet Request For Comments</i>	11
1.9 <i>Internet Protocols And Standardization</i>	12
1.10 <i>Future Growth And Technology</i>	12
1.11 <i>Organization Of The Text</i>	13
1.12 <i>Summary</i>	14
Chapter 2 Review Of Underlying Network Technologies	17
2.1 <i>Introduction</i>	17
2.2 <i>Two Approaches To Network Communication</i>	18
2.3 <i>Wide Area And Local Area Networks</i>	19
2.4 <i>Ethernet Technology</i>	20
2.5 <i>Fiber Distributed Data Interconnect (FDDI)</i>	32
2.6 <i>Asynchronous Transfer Mode</i>	36
2.7 <i>ARPANET Technology</i>	37
2.8 <i>National Science Foundation Networking</i>	39
2.9 <i>ANSNET</i>	44

- 2.10 *A Planned Wide Area Backbone* 44
- 2.11 *Other Technologies Over Which TCP/IP Has Been Used* 44
- 2.12 *Summary And Conclusion* 47

Chapter 3 Internetworking Concept And Architectural Model 49

- 3.1 *Introduction* 49
- 3.2 *Application-Level Interconnection* 49
- 3.3 *Network-Level Interconnection* 50
- 3.4 *Properties Of The Internet* 51
- 3.5 *Internet Architecture* 52
- 3.6 *Interconnection Through IP Routers* 52
- 3.7 *The User's View* 54
- 3.8 *All Networks Are Equal* 54
- 3.9 *The Unanswered Questions* 55
- 3.10 *Summary* 56

Chapter 4 Internet Addresses 59

- 4.1 *Introduction* 59
- 4.2 *Universal Identifiers* 59
- 4.3 *Three Primary Classes Of IP Addresses* 60
- 4.4 *Addresses Specify Network Connections* 61
- 4.5 *Network And Broadcast Addresses* 61
- 4.6 *Limited Broadcast* 62
- 4.7 *Interpreting Zero To Mean "This"* 62
- 4.8 *Weaknesses In Internet Addressing* 63
- 4.9 *Dotted Decimal Notation* 65
- 4.10 *Loopback Address* 65
- 4.11 *Summary Of Special Address Conventions* 66
- 4.12 *Internet Addressing Authority* 66
- 4.13 *An Example* 67
- 4.14 *Network Byte Order* 69
- 4.15 *Summary* 70

Chapter 5 Mapping Internet Addresses To Physical Addresses (ARP) 73

- 5.1 *Introduction* 73
- 5.2 *The Address Resolution Problem* 73
- 5.3 *Two Types Of Physical Addresses* 74
- 5.4 *Resolution Through Direct Mapping* 74

49

- 5.5 *Resolution Through Dynamic Binding* 75
- 5.6 *The Address Resolution Cache* 76
- 5.7 *ARP Refinements* 77
- 5.8 *Relationship Of ARP To Other Protocols* 77
- 5.9 *ARP Implementation* 77
- 5.10 *ARP Encapsulation And Identification* 79
- 5.11 *ARP Protocol Format* 79
- 5.12 *Summary* 81

Chapter 6 Determining An Internet Address At Startup (RARP) 83

- 6.1 *Introduction* 83
- 6.2 *Reverse Address Resolution Protocol (RARP)* 84
- 6.3 *Timing RARP Transactions* 86
- 6.4 *Primary And Backup RARP Servers* 86
- 6.5 *Summary* 87

59

Chapter 7 Internet Protocol: Connectionless Datagram Delivery 89

- 7.1 *Introduction* 89
- 7.2 *A Virtual Network* 89
- 7.3 *Internet Architecture And Philosophy* 90
- 7.4 *The Concept Of Unreliable Delivery* 90
- 7.5 *Connectionless Delivery System* 91
- 7.6 *Purpose Of The Internet Protocol* 91
- 7.7 *The Internet Datagram* 91
- 7.8 *Internet Datagram Options* 100
- 7.9 *Summary* 106

Chapter 8 Internet Protocol: Routing IP Datagrams 109

- 8.1 *Introduction* 109
- 8.2 *Routing In An Internet* 109
- 8.3 *Direct And Indirect Delivery* 111
- 8.4 *Table-Driven IP Routing* 113
- 8.5 *Next-Hop Routing* 113
- 8.6 *Default Routes* 115
- 8.7 *Host-Specific Routes* 115
- 8.8 *The IP Routing Algorithm* 116
- 8.9 *Routing With IP Addresses* 116
- 8.10 *Handling Incoming Datagrams* 118

(ARP) 73

- 8.11 *Establishing Routing Tables* 119
- 8.12 *Summary* 119

Chapter 9 Internet Protocol: Error And Control Messages (ICMP) 123

- 9.1 *Introduction* 123
- 9.2 *The Internet Control Message Protocol* 123
- 9.3 *Error Reporting vs. Error Correction* 124
- 9.4 *ICMP Message Delivery* 125
- 9.5 *ICMP Message Format* 126
- 9.6 *Testing Destination Reachability And Status (Ping)* 127
- 9.7 *Echo Request And Reply Message Format* 128
- 9.8 *Reports Of Unreachable Destinations* 128
- 9.9 *Congestion And Datagram Flow Control* 130
- 9.10 *Source Quench Format* 130
- 9.11 *Route Change Requests From Routers* 131
- 9.12 *Detecting Circular Or Excessively Long Routes* 133
- 9.13 *Reporting Other Problems* 134
- 9.14 *Clock Synchronization And Transit Time Estimation* 134
- 9.15 *Information Request And Reply Messages* 136
- 9.16 *Obtaining A Subnet Mask* 136
- 9.17 *Summary* 137

Chapter 10 Subnet And Supernet Address Extensions 139

- 10.1 *Introduction* 139
- 10.2 *Review Of Relevant Facts* 139
- 10.3 *Minimizing Network Numbers* 140
- 10.4 *Transparent Routers* 141
- 10.5 *Proxy ARP* 142
- 10.6 *Subnet Addressing* 143
- 10.7 *Flexibility In Subnet Address Assignment* 146
- 10.8 *Implementation Of Subnets With Masks* 147
- 10.9 *Subnet Mask Representation* 148
- 10.10 *Routing In The Presence Of Subnets* 149
- 10.11 *The Subnet Routing Algorithm* 150
- 10.12 *A Unified Routing Algorithm* 151
- 10.13 *Maintenance Of Subnet Masks* 152
- 10.14 *Broadcasting To Subnets* 152
- 10.15 *Supernet Addressing* 153
- 10.16 *The Effect Of Supernetting On Routing* 154
- 10.17 *Summary* 155

) 123

139

Chapter 11 Protocol Layering**159**

- 11.1 Introduction 159
- 11.2 The Need For Multiple Protocols 159
- 11.3 The Conceptual Layers Of Protocol Software 160
- 11.4 Functionality Of The Layers 163
- 11.5 X.25 And Its Relation To The ISO Model 164
- 11.6 Differences Between X.25 And Internet Layering 167
- 11.7 The Protocol Layering Principle 169
- 11.8 Layering In The Presence Of Network Substructure 171
- 11.9 Two Important Boundaries In The TCP/IP Model 173
- 11.10 The Disadvantage Of Layering 174
- 11.11 The Basic Idea Behind Multiplexing And Demultiplexing 174
- 11.12 Summary 176

Chapter 12 User Datagram Protocol (UDP)**179**

- 12.1 Introduction 179
- 12.2 Identifying The Ultimate Destination 179
- 12.3 The User Datagram Protocol 180
- 12.4 Format Of UDP Messages 181
- 12.5 UDP Pseudo-Header 182
- 12.6 UDP Encapsulation And Protocol Layering 183
- 12.7 Layering And The UDP Checksum Computation 185
- 12.8 UDP Multiplexing, Demultiplexing, And Ports 185
- 12.9 Reserved And Available UDP Port Numbers 186
- 12.10 Summary 188

Chapter 13 Reliable Stream Transport Service (TCP)**191**

- 13.1 Introduction 191
- 13.2 The Need For Stream Delivery 191
- 13.3 Properties Of The Reliable Delivery Service 192
- 13.4 Providing Reliability 193
- 13.5 The Idea Behind Sliding Windows 195
- 13.6 The Transmission Control Protocol 198
- 13.7 Ports, Connections, And Endpoints 199
- 13.8 Passive And Active Opens 201
- 13.9 Segments, Streams, And Sequence Numbers 201
- 13.10 Variable Window Size And Flow Control 202
- 13.11 TCP Segment Format 203

13.12	<i>Out Of Band Data</i>	205
13.13	<i>Maximum Segment Size Option</i>	206
13.14	<i>TCP Checksum Computation</i>	207
13.15	<i>Acknowledgements And Retransmission</i>	208
13.16	<i>Timeout And Retransmission</i>	209
13.17	<i>Accurate Measurement Of Round Trip Samples</i>	211
13.18	<i>Karn's Algorithm And Timer Backoff</i>	212
13.19	<i>Responding To High Variance In Delay</i>	213
13.20	<i>Response To Congestion</i>	214
13.21	<i>Establishing A TCP Connection</i>	216
13.22	<i>Initial Sequence Numbers</i>	217
13.23	<i>Closing a TCP Connection</i>	217
13.24	<i>TCP Connection Reset</i>	219
13.25	<i>TCP State Machine</i>	219
13.26	<i>Forcing Data Delivery</i>	221
13.27	<i>Reserved TCP Port Numbers</i>	221
13.28	<i>TCP Performance</i>	221
13.29	<i>Silly Window Syndrome And Small Packets</i>	223
13.30	<i>Avoiding Silly Window Syndrome</i>	224
13.31	<i>Summary</i>	227

Chapter 14 Routing: Cores, Peers, And Algorithms (GGP) 231

14.1	<i>Introduction</i>	231
14.2	<i>The Origin Of Routing Tables</i>	232
14.3	<i>Routing With Partial Information</i>	233
14.4	<i>Original Internet Architecture And Cores</i>	234
14.5	<i>Core Routers</i>	235
14.6	<i>Beyond The Core Architecture To Peer Backbones</i>	238
14.7	<i>Automatic Route Propagation</i>	240
14.8	<i>Vector Distance (Bellman-Ford) Routing</i>	240
14.9	<i>Gateway-To-Gateway Protocol (GGP)</i>	242
14.10	<i>GGP Message Formats</i>	243
14.11	<i>Link-State (SPF) Routing</i>	245
14.12	<i>SPF Protocols</i>	246
14.13	<i>Summary</i>	246

Chapter 15 Routing: Autonomous Systems (EGP) 249

15.1	<i>Introduction</i>	249
15.2	<i>Adding Complexity To The Architectural Model</i>	249
15.3	<i>A Fundamental Idea: Extra Hops</i>	250

15.4	<i>Autonomous System Concept</i>	252
15.5	<i>Exterior Gateway Protocol (EGP)</i>	254
15.6	<i>EGP Message Header</i>	255
15.7	<i>EGP Neighbor Acquisition Messages</i>	256
15.8	<i>EGP Neighbor Reachability Messages</i>	257
15.9	<i>EGP Poll Request Messages</i>	258
15.10	<i>EGP Routing Update Messages</i>	259
15.11	<i>Measuring From The Receiver's Perspective</i>	261
15.12	<i>The Key Restriction Of EGP</i>	262
15.13	<i>Technical Problems</i>	264
15.14	<i>Decentralization Of Internet Architecture</i>	264
15.15	<i>Beyond Autonomous Systems</i>	264
15.16	<i>Summary</i>	265

Chapter 16 Routing: In An Autonomous System (RIP, OSPF, HELLO) 267

16.1	<i>Introduction</i>	267
16.2	<i>Static Vs. Dynamic Interior Routes</i>	267
16.3	<i>Routing Information Protocol (RIP)</i>	270
16.4	<i>The Hello Protocol</i>	276
16.5	<i>Combining RIP, Hello, And EGP</i>	278
16.6	<i>The Open SPF Protocol (OSPF)</i>	279
16.7	<i>Routing With Partial Information</i>	286
16.8	<i>Summary</i>	286

Chapter 17 Internet Multicasting (IGMP) 289

17.1	<i>Introduction</i>	289
17.2	<i>Hardware Broadcast</i>	289
17.3	<i>Hardware Multicast</i>	290
17.4	<i>IP Multicast</i>	291
17.5	<i>IP Multicast Addresses</i>	291
17.6	<i>Mapping IP Multicast To Ethernet Multicast</i>	292
17.7	<i>Extending IP To Handle Multicasting</i>	293
17.8	<i>Internet Group Management Protocol</i>	294
17.9	<i>IGMP Implementation</i>	294
17.10	<i>Group Membership State Transitions</i>	295
17.11	<i>IGMP Message Format</i>	296
17.12	<i>Multicast Address Assignment</i>	297
17.13	<i>Propagating Routing Information</i>	297
17.14	<i>The Mrouted Program</i>	298
17.15	<i>Summary</i>	300

Chapter 18 TCP/IP Over ATM Networks 303

- 18.1 Introduction 303
- 18.2 ATM Hardware 304
- 18.3 Large ATM Networks 304
- 18.4 The Logical View Of An ATM Network 305
- 18.5 The Two ATM Connection Paradigms 306
- 18.6 Paths, Circuits, And Identifiers 307
- 18.7 ATM Cell Transport 308
- 18.8 ATM Adaptation Layers 308
- 18.9 AAL5 Convergence, Segmentation, And Reassembly 311
- 18.10 Datagram Encapsulation And IP MTU Size 311
- 18.11 Packet Type And Multiplexing 312
- 18.12 IP Address Binding In An ATM Network 313
- 18.13 Logical IP Subnet Concept 314
- 18.14 Connection Management 315
- 18.15 Address Binding Within An LIS 316
- 18.16 ATMARP Packet Format 316
- 18.17 Using ATMARP Packets To Determine An Address 318
- 18.18 Obtaining Entries For A Server Database 320
- 18.19 Timing Out ATMARP Information In A Server 320
- 18.20 Timing Out ATMARP Information In A Host Or Router 320
- 18.21 Summary 321

Chapter 19 Client-Server Model Of Interaction 325

- 19.1 Introduction 325
- 19.2 The Client-Server Model 325
- 19.3 A Simple Example: UDP Echo Server 326
- 19.4 Time And Date Service 328
- 19.5 The Complexity of Servers 329
- 19.6 RARP Server 330
- 19.7 Alternatives To The Client-Server Model 331
- 19.8 Summary 332

Chapter 20 The Socket Interface 335

- 20.1 Introduction 335
- 20.2 The UNIX I/O Paradigm And Network I/O 336
- 20.3 Adding Network I/O to UNIX 336
- 20.4 The Socket Abstraction 337

20.5	<i>Creating A Socket</i>	337
20.6	<i>Socket Inheritance And Termination</i>	338
20.7	<i>Specifying A Local Address</i>	339
20.8	<i>Connecting Sockets To Destination Addresses</i>	340
20.9	<i>Sending Data Through A Socket</i>	341
20.10	<i>Receiving Data Through A Socket</i>	343
20.11	<i>Obtaining Local And Remote Socket Addresses</i>	344
20.12	<i>Obtaining And Setting Socket Options</i>	345
20.13	<i>Specifying A Queue Length For A Server</i>	346
20.14	<i>How A Server Accepts Connections</i>	346
20.15	<i>Servers That Handle Multiple Services</i>	347
20.16	<i>Obtaining And Setting Host Names</i>	348
20.17	<i>Obtaining And Setting The Internal Host Domain</i>	349
20.18	<i>BSD UNIX Network Library Calls</i>	349
20.19	<i>Network Byte Order Conversion Routines</i>	350
20.20	<i>IP Address Manipulation Routines</i>	351
20.21	<i>Accessing The Domain Name System</i>	352
20.22	<i>Obtaining Information About Hosts</i>	354
20.23	<i>Obtaining Information About Networks</i>	355
20.24	<i>Obtaining Information About Protocols</i>	355
20.25	<i>Obtaining Information About Network Services</i>	356
20.26	<i>An Example Client</i>	357
20.27	<i>An Example Server</i>	359
20.28	<i>Summary</i>	362

Chapter 21 Bootstrap And Autoconfiguration (BOOTP, DHCP)

21.1	<i>Introduction</i>	365
21.2	<i>The Need For An Alternative To RARP</i>	366
21.3	<i>Using IP To Determine An IP Address</i>	366
21.4	<i>The BOOTP Retransmission Policy</i>	367
21.5	<i>The BOOTP Message Format</i>	368
21.6	<i>The Two-Step Bootstrap Procedure</i>	369
21.7	<i>Vendor-Specific Field</i>	370
21.8	<i>The Need For Dynamic Configuration</i>	370
21.9	<i>Dynamic Host Configuration</i>	372
21.10	<i>Dynamic IP Address Assignment</i>	372
21.11	<i>Obtaining Multiple Addresses</i>	373
21.12	<i>Address Acquisition States</i>	374
21.13	<i>Early Lease Termination</i>	374
21.14	<i>Lease Renewal States</i>	376
21.15	<i>DHCP Message Format</i>	377
21.16	<i>DHCP Options And Message Type</i>	378

- 21.17 *Option Overload* 379
- 21.18 *DHCP And Domain Names* 379
- 21.19 *Summary* 380

Chapter 22 The Domain Name System (DNS)

383

- 22.1 *Introduction* 383
- 22.2 *Names For Machines* 384
- 22.3 *Flat Namespace* 384
- 22.4 *Hierarchical Names* 385
- 22.5 *Delegation Of Authority For Names* 386
- 22.6 *Subset Authority* 386
- 22.7 *TCP/IP Internet Domain Names* 387
- 22.8 *Official And Unofficial Internet Domain Names* 388
- 22.9 *Items Named And Syntax Of Names* 390
- 22.10 *Mapping Domain Names To Addresses* 391
- 22.11 *Domain Name Resolution* 393
- 22.12 *Efficient Translation* 394
- 22.13 *Caching: The Key To Efficiency* 395
- 22.14 *Domain Server Message Format* 396
- 22.15 *Compressed Name Format* 399
- 22.16 *Abbreviation Of Domain Names* 399
- 22.17 *Inverse Mappings* 400
- 22.18 *Pointer Queries* 401
- 22.19 *Object Types And Resource Record Contents* 401
- 22.20 *Obtaining Authority For A Subdomain* 402
- 22.21 *Summary* 403

Chapter 23 Applications: Remote Login (TELNET, Rlogin)

407

- 23.1 *Introduction* 407
- 23.2 *Remote Interactive Computing* 407
- 23.3 *TELNET Protocol* 408
- 23.4 *Accommodating Heterogeneity* 410
- 23.5 *Passing Commands That Control The Remote Side* 412
- 23.6 *Forcing The Server To Read A Control Function* 414
- 23.7 *TELNET Options* 414
- 23.8 *TELNET Option Negotiation* 415
- 23.9 *Rlogin (BSD UNIX)* 416
- 23.10 *Summary* 417

383

Chapter 24 Applications: File Transfer And Access (FTP, TFTP, NFS) 419

- 24.1 Introduction 419
- 24.2 File Access And Transfer 419
- 24.3 On-line Shared Access 420
- 24.4 Sharing By File Transfer 421
- 24.5 FTP: The Major TCP/IP File Transfer Protocol 421
- 24.6 FTP Features 422
- 24.7 FTP Process Model 422
- 24.8 TCP Port Number Assignment 424
- 24.9 The User's View Of FTP 424
- 24.10 An Example Anonymous FTP Session 426
- 24.11 TFTP 427
- 24.12 NFS 429
- 24.13 NFS Implementation 429
- 24.14 Remote Procedure Call (RPC) 430
- 24.15 Summary 431

Chapter 25 Applications: Electronic Mail (822, SMTP, MIME) 433

- 25.1 Introduction 433
- 25.2 Electronic Mail 433
- 25.3 Mailbox Names And Aliases 435
- 25.4 Alias Expansion And Mail Forwarding 435
- 25.5 The Relationship Of Internetworking And Mail 436
- 25.6 TCP/IP Standards For Electronic Mail Service 438
- 25.7 Electronic Mail Addresses 438
- 25.8 Pseudo Domain Addresses 440
- 25.9 Simple Mail Transfer Protocol (SMTP) 440
- 25.10 The MIME Extension For Non-ASCII Data 443
- 25.11 MIME Multipart Messages 444
- 25.12 Summary 445

407

Chapter 26 Applications: Internet Management (SNMP, SNMPv2) 447

- 26.1 Introduction 447
- 26.2 The Level Of Management Protocols 447
- 26.3 Architectural Model 448
- 26.4 Protocol Architecture 450
- 26.5 Examples of MIB Variables 451
- 26.6 The Structure Of Management Information 452

- 26.7 *Formal Definitions Using ASN.1* 453
- 26.8 *Structure And Representation Of MIB Object Names* 453
- 26.9 *Simple Network Management Protocol* 458
- 26.10 *SNMP Message Format* 460
- 26.11 *Example Encoded SNMP Message* 462
- 26.12 *Summary* 463

Chapter 27 Summary Of Protocol Dependencies

465

- 27.1 *Introduction* 465
- 27.2 *Protocol Dependencies* 465
- 27.3 *Application Program Access* 467
- 27.4 *Summary* 468

Chapter 28 Internet Security And Firewall Design

471

- 28.1 *Introduction* 471
- 28.2 *Protecting Resources* 472
- 28.3 *The Need For An Information Policy* 472
- 28.4 *Communication, Cooperation, And Mutual Mistrust* 474
- 28.5 *Mechanisms For Internet Security* 475
- 28.6 *Firewalls And Internet Access* 476
- 28.7 *Multiple Connections And Weakest Links* 477
- 28.8 *Firewall Implementation And High-Speed Hardware* 478
- 28.9 *Packet-Level Filters* 479
- 28.10 *Security And Packet Filter Specification* 480
- 28.11 *The Consequence Of Restricted Access For Clients* 481
- 28.12 *Accessing Services Through A Firewall* 481
- 28.13 *The Details Of Firewall Architecture* 483
- 28.14 *Stub Network* 484
- 28.15 *An Alternative Firewall Implementation* 484
- 28.16 *Monitoring And Logging* 485
- 28.17 *Summary* 486

Chapter 29 The Future Of TCP/IP (IPng, IPv6)

489

- 29.1 *Introduction* 489
- 29.2 *Why Change TCP/IP And The Internet?* 490
- 29.3 *Motivation For Changing IPv4* 491
- 29.4 *The Road To A New Version Of IP* 492
- 29.5 *The Name Of The Next IP* 492

465

471

- 29.6 *Features Of IPv6* 493
- 29.7 *General Form Of An IPv6 Datagram* 494
- 29.8 *IPv6 Base Header Format* 494
- 29.9 *IPv6 Extension Headers* 496
- 29.10 *Parsing An IPv6 Datagram* 497
- 29.11 *IPv6 Fragmentation And Reassembly* 498
- 29.12 *The Consequence Of End-To-End Fragmentation* 498
- 29.13 *IPv6 Source Routing* 500
- 29.14 *IPv6 Options* 500
- 29.15 *Size Of The IPv6 Address Space* 502
- 29.16 *IPv6 Colon Hexadecimal Notation* 502
- 29.17 *Three Basic IPv6 Address Types* 503
- 29.18 *The Duality Of Broadcast And Multicast* 504
- 29.19 *An Engineering Choice And Simulated Broadcast* 504
- 29.20 *Proposed IPv6 Address Space Assignment* 504
- 29.21 *IPv4 Address Encoding And Transition* 506
- 29.22 *Providers, Subscribers, And Address Hierarchy* 506
- 29.23 *Additional Hierarchy* 507
- 29.24 *Summary* 508

Appendix 1 A Guide To RFCs 511

Appendix 2 Glossary Of Internetworking Terms And Abbreviations 557

Bibliography 591

Index 599

489

1

Introduction And Overview

1.1 The Motivation For Internetworking

Data communication has become a fundamental part of computing. World-wide networks gather data about such diverse subjects as atmospheric conditions, crop production, and airline traffic. Groups establish electronic mailing lists so they can share information of common interest. Hobbyists exchange programs for their home computers. In the scientific world, data networks are essential because they allow scientists to send programs and data to remote supercomputers for processing, to retrieve the results, and to exchange information with colleagues.

Unfortunately, most networks are independent entities, established to serve the needs of a single group. The users choose a hardware technology appropriate to their communication problems. More important, it is impossible to build a universal network from a single hardware technology because no single network suffices for all uses. Some users need a high-speed network to connect machines, but such networks cannot be expanded to span large distances. Others settle for a slower speed network that connects machines thousands of miles apart.

In the past 15 years, a new technology has evolved that makes it possible to interconnect many disparate physical networks and make them function as a coordinated unit. The technology, called *internetworking*, accommodates multiple, diverse underlying hardware technologies by providing a way to interconnect heterogeneous networks and a set of communication conventions. The internet technology hides the details of network hardware and permits computers to communicate independent of their physical network connections.

The internet technology described in this book is an example of *open system interconnection*. It is called an *open system* because, unlike proprietary communication systems available from one specific vendor, the specifications are publicly available. Thus,

anyone can build the software needed to communicate across an internet. More important, the entire technology has been designed to foster communication between machines with diverse hardware architectures, to use almost any packet switched network hardware, and to accommodate multiple computer operating systems.

To appreciate internet technology, think of how it affects a professional group. Consider, for example, the effect of interconnecting the computers used by scientists. Any scientist can exchange data resulting from an experiment with any other scientist. National centers can collect data from natural phenomena and make the data available to all scientists. Computer services and programs available at one location can be used by scientists at other locations. As a result, the speed with which scientific investigations proceed increases; the changes are dramatic.

1.2 The TCP/IP Internet

U.S. government agencies have realized the importance and potential of internet technology for many years and have been funding research that has made possible a global internet. This book discusses principles and ideas underlying the internet technology that has resulted from research funded by the *Advanced Research Projects Agency (ARPA)*†. The ARPA technology includes a set of network standards that specify the details of how computers communicate, as well as a set of conventions for interconnecting networks and routing traffic. Officially named the TCP/IP Internet Protocol Suite and commonly referred to as *TCP/IP* (after the names of its two main standards), it can be used to communicate across any set of interconnected networks. For example, some corporations use TCP/IP to interconnect all networks within their corporation, even though the corporation has no connection to outside networks. Other groups use TCP/IP for communication among geographically distant sites.

Although the TCP/IP technology is noteworthy by itself, it is especially interesting because its viability has been demonstrated on a large scale. It forms the base technology for a global internet that connects homes, university campuses and other schools, corporations, and government labs in 61 countries. In the U.S., The *National Science Foundation (NSF)*, the *Department of Energy (DOE)*, the *Department of Defense (DOD)*, the *Health and Human Services Agency (HHS)* and the *National Aeronautics and Space Administration (NASA)* have all participated in funding the Internet, and use TCP/IP to connect many of their research sites. Known as the *ARPA/NSF Internet*, the *TCP/IP Internet*, the *global Internet*, or just the *Internet*‡, the resulting internet allows researchers at connected institutions to share information with colleagues around the world as easily as they share it with researchers in the next room. An outstanding success, the Internet demonstrates the viability of the TCP/IP technology and shows how it can accommodate a wide variety of underlying network technologies.

Most of the material in this book applies to any internet that uses TCP/IP, but some chapters refer specifically to the global Internet. Readers interested only in the technology should be careful to watch for the distinction between the Internet architecture as it exists and general TCP/IP internets as they might exist. It would be a mis-

†ARPA was called the *Defense Advanced Research Projects Agency* for several years during the 1980s.

‡We will follow the usual convention of capitalizing *Internet* when referring specifically to the global Internet, and use lower case to refer to private internets that use TCP/IP.

More impor-
 ation between
 switched net-
 s.
 sional group.
 d by scientists.
 other scientist.
 s data available
 on can be used
 atific investiga-

ntial of internet
 made possible a
 e internet tech-
 i Projects Agen-
 that specify the
 for interconnect-
 t Protocol Suite
 standards), it can
 r example, some
 orporation, even
 ther groups use

cially interesting
 e base technolo-
 id other schools,
*National Science
 ment of Defense
 onal Aeronautics
 Internet*, and use
NSF Internet, the
 ig internet allows
 agues around the
 outstanding suc-
 and shows how it

uses TCP/IP, but
 rested only in the
 Internet architec-
 t would be a mis-

ars during the 1980s.
 ecifically to the global

take, however, to ignore completely sections of the text that describe the global Internet – many corporate networks are already more complex than the global Internet of ten years ago, and many of the problems they face have already been solved in the global Internet.

1.3 Internet Services

One cannot appreciate the technical details underlying TCP/IP without understanding the services it provides. This section reviews internet services briefly, highlighting the services most users access, and leaves to later chapters the discussion of how computers connect to a TCP/IP internet and how the functionality is implemented.

Much of our discussion of services will focus on standards called *protocols*. Protocols like TCP and IP provide the rules for communication. They contain the details of message formats, describe how a computer responds when a message arrives, and specify how a computer handles errors or other abnormal conditions. Most important, they allow us to discuss computer communication independent of any particular vendor's network hardware. In a sense, protocols are to communication what algorithms are to computation. An algorithm allows one to specify or understand a computation without knowing the details of a particular CPU instruction set. Similarly, a communication protocol allows one to specify or understand data communication without depending on detailed knowledge of a particular vendor's network hardware.

Hiding the low-level details of communication helps improve productivity in several ways. First, because programmers deal with higher-level protocol abstractions, they do not need to learn or remember as many details about a given hardware configuration. They can create new programs quickly. Second, because programs built using higher-level abstractions are not restricted to a particular machine architecture or a particular network hardware, they do not need to be changed when machines or networks are reconfigured. Third, because application programs built using higher-level protocols are independent of the underlying hardware, they can provide direct communication for an arbitrary pair of machines. Programmers do not need to build special versions of application software to move and translate data between each possible pair of machine types.

We will see that all network services are described by protocols. The next sections refer to protocols used to specify application-level services as well as those used to define network-level services. Later chapters explain each of these protocols in more detail.

1.3.1 Application Level Internet Services

From the user's point of view, a TCP/IP internet appears to be a set of application programs that use the network to carry out useful communication tasks. We use the term *interoperability* to refer to the ability of diverse computing systems to cooperate in solving computational problems. Internet application programs exhibit a high degree of

interoperability. Most users that access the Internet do so merely by running application programs without understanding the TCP/IP technology, the structure of the underlying internet, or even the path the data travels to its destination; they rely on the application programs and the underlying network software to handle such details. Only programmers who write network application programs need to view the internet as a network and need to understand some of the technology.

The most popular and widespread Internet application services include:

- *Electronic mail.* Electronic mail allows a user to compose memos and send them to individuals or groups. Another part of the mail application allows users to read memos that they have received. Electronic mail has been so successful that many Internet users depend on it for normal business correspondence. Although many electronic mail systems exist, using TCP/IP makes mail delivery more reliable because it does not rely on intermediate computers to relay mail messages. A TCP/IP mail delivery system operates by having the sender's machine contact the receiver's machine directly. Thus, the sender knows that once the message leaves the local machine, it has been successfully received at the destination site.
- *File transfer.* Although users sometimes transfer files using electronic mail, mail is designed primarily for short text messages. The TCP/IP protocols include a file transfer application program that allows users to send or receive arbitrarily large files of programs or data. For example, using the file transfer program, one can copy from one machine to another a large data base containing satellite images, a program written in Pascal or C++, or an English dictionary. The system provides a way to check for authorized users, or even to prevent all access. Like mail, file transfer across a TCP/IP internet is reliable because the two machines involved communicate directly, without relying on intermediate machines to make copies of the file along the way.
- *Remote login.* Remote login allows a user sitting at one computer to connect to a remote machine and establish an interactive login session. The remote login makes it appear that a window on the user's screen connects directly to the remote machine by sending each keystroke from the user's keyboard to the remote machine and displaying each character the remote computer prints in the user's window. When the remote login session terminates, the application returns the user to the local system.

We will return to these and other applications in later chapters to examine them in more detail. We will see exactly how they use the underlying TCP/IP protocols, and why having standards for application protocols has helped ensure that they are widespread.

1.3.2 Network-Level Internet Services

A programmer who writes application programs that use TCP/IP protocols has an entirely different view of an internet than a user who merely executes applications like electronic mail. At the network level, an internet provides two broad types of service

ring application
the underlying
the application
Only program-
et as a network

ide:

s and send them
ws users to read
essful that many

Although many
more reliable be-
il messages. A
chine contact the
e message leaves
ion site.

tronic mail, mail
ols include a file
arbitrarily large
program, one can
atellite images, a
system provides
s. Like mail, file
achines involved
s to make copies

er to connect to a
The remote login
ctly to the remote
rd to the remote
rints in the user's
ication returns the

nine them in more
rotocols, and why
are widespread.

P protocols has an
es applications like
ad types of service

that all application programs use. While it is unimportant at this time to understand the details of these services, they cannot be omitted from any overview of TCP/IP:

- *Connectionless Packet Delivery Service.* This service, explained in detail throughout the text, forms the basis for all other internet services. Connectionless delivery is an abstraction of the service that most packet-switching networks offer. It means simply that a TCP/IP internet routes small messages from one machine to another based on address information carried in the message. Because the connectionless service routes each packet separately, it does not guarantee reliable, in-order delivery. Because it usually maps directly onto the underlying hardware, the connectionless service is extremely efficient. More important, having connectionless packet delivery as the basis for all internet services makes the TCP/IP protocols adaptable to a wide range of network hardware.
- *Reliable Stream Transport Service.* Most applications need much more than packet delivery because they require the communication software to recover automatically from transmission errors, lost packets, or failures of intermediate switches along the path between sender and receiver. The reliable transport service handles such problems. It allows an application on one computer to establish a "connection" with an application on another computer, and then to send a large volume of data across the connection as if it were a permanent, direct hardware connection. Underneath, of course, the communication protocols divide the stream of data into small messages and send them, one at a time, waiting for the receiver to acknowledge reception.

Many networks provide basic services similar to those outlined above, so one might wonder what distinguishes TCP/IP services from others. The primary distinguishing features are:

- *Network Technology Independence.* While TCP/IP is based on conventional packet switching technology, it is independent of any particular vendor's hardware. The global Internet includes a variety of network technologies ranging from networks designed to operate within a single building to those designed to span large distances. TCP/IP protocols define the unit of data transmission, called a *datagram*, and specify how to transmit datagrams on a particular network.
- *Universal Interconnection.* A TCP/IP internet allows any pair of computers to which it attaches to communicate. Each computer is assigned an *address* that is universally recognized throughout the internet. Every datagram carries the addresses of its source and destination. Intermediate switching computers use the destination address to make routing decisions.
- *End-to-End Acknowledgements.* The TCP/IP internet protocols provide acknowledgements between the source and ultimate destination instead of between successive machines along the path, even when the two machines do not connect to a common physical network.
- *Application Protocol Standards.* In addition to the basic transport-level services (like reliable stream connections), the TCP/IP protocols include standards for

many common applications including electronic mail, file transfer, and remote login. Thus, when designing application programs that use TCP/IP, programmers often find that existing software provides the communication services they need.

Later chapters will discuss the details of the services provided to the programmer as well as many of the application protocol standards.

1.4 History And Scope Of The Internet

Part of what makes the TCP/IP technology so exciting is its almost universal adoption as well as the size and growth rate of the global Internet. ARPA began working toward an internet technology in the mid 1970s, with the architecture and protocols taking their current form around 1977-79. At that time, ARPA was known as the primary funding agency for packet-switched network research and had pioneered many ideas in packet-switching with its well-known *ARPANET*. The *ARPANET* used conventional point-to-point leased line interconnection, but ARPA had also funded exploration of packet-switching over radio networks and satellite communication channels. Indeed, the growing diversity of network hardware technologies helped force ARPA to study network interconnection, and pushed internetworking forward.

The availability of research funding from ARPA caught the attention and imagination of several research groups, especially those researchers who had previous experience using packet switching on the *ARPANET*. ARPA scheduled informal meetings of researchers to share ideas and discuss results of experiments. By 1979, so many researchers were involved in the TCP/IP effort that ARPA formed an informal committee to coordinate and guide the design of the protocols and architecture of the emerging Internet. Called the Internet Control and Configuration Board (ICCB), the group met regularly until 1983, when it was reorganized.

The global Internet began around 1980 when ARPA started converting machines attached to its research networks to the new TCP/IP protocols. The *ARPANET*, already in place, quickly became the backbone of the new Internet and was used for many of the early experiments with TCP/IP. The transition to Internet technology became complete in January 1983 when the Office of the Secretary of Defense mandated that all computers connected to long-haul networks use TCP/IP. At the same time, the *Defense Communication Agency* (DCA) split the *ARPANET* into two separate networks, one for further research and one for military communication. The research part retained the name *ARPANET*; the military part, which was somewhat larger, became known as the *military network*, *MILNET*.

To encourage university researchers to adopt and use the new protocols, ARPA made an implementation available at low cost. At that time, most university computer science departments were running a version of the UNIX operating system available in the University of California's *Berkeley Software Distribution*, commonly called *Berkeley UNIX* or *BSD UNIX*. By funding Bolt Beranek and Newman, Inc. (BBN) to implement its TCP/IP protocols for use with UNIX, and funding Berkeley to integrate the

and remote lo-
gins, programmers
as they need.

programmer as

universal adop-
tion began working to-
wards protocols taking
as the primary
of many ideas in
of conventional
of exploration of
of protocols. Indeed, the
of A to study net-

tion and imagina-
tion of previous experi-
mental meetings of
1979, so many
of informal commit-
tees of the emerging
of), the group met

of inverting machines
of ARPANET, already
used for many of
of technology became com-
mandated that all
of time, the *Defense*
of networks, one for
of part retained the
of name known as the

of protocols, ARPA
of university computer
of system available in
of only called *Berke-*
of c. (BBN) to imple-
of ment to integrate the

protocols with its software distribution, ARPA was able to reach over 90% of the university computer science departments. The new protocol software came at a particularly significant time because many departments were just acquiring second or third computers and connecting them together with local area networks. The departments needed communication protocols and no others were generally available.

The Berkeley software distribution became popular because it offered more than basic TCP/IP protocols. In addition to standard TCP/IP application programs, Berkeley offered a set of utilities for network services that resembled the UNIX services used on a single machine. The chief advantage of the Berkeley utilities lies in their similarity to standard UNIX. For example, an experienced UNIX user can quickly learn how to use Berkeley's remote file copy utility (*rcp*) because it behaves exactly like the UNIX file copy utility except that it allows users to copy files to or from remote machines.

Besides a set of utility programs, Berkeley UNIX provided a new operating system abstraction known as a *socket* that allows application programs to access communication protocols. A generalization of the UNIX mechanism for I/O, the socket has options for several types of network protocols in addition to TCP/IP. Its design has been debated since its introduction, and many operating systems researchers have proposed alternatives. Independent of its overall merits, however, the introduction of the socket abstraction was important because it allowed programmers to use TCP/IP protocols with little effort. Thus, it encouraged researchers to experiment with TCP/IP.

The success of the TCP/IP technology and the Internet among computer science researchers led other groups to adopt it. Realizing that network communication would soon be a crucial part of scientific research, the National Science Foundation took an active role in expanding the TCP/IP Internet to reach as many scientists as possible. Starting in 1985, it began a program to establish access networks centered around its six supercomputer centers. In 1986 it expanded networking efforts by funding a new wide area backbone network, called the *NSFNET*[†], that eventually reached all its supercomputer centers and tied them to the ARPANET. Finally, in 1986 NSF provided seed money for many regional networks, each of which now connects major scientific research institutions in a given area. All the NSF-funded networks use TCP/IP protocols, and all are part of the global Internet.

Within seven years of its inception, the Internet had grown to span hundreds of individual networks located throughout the United States and Europe. It connected nearly 20,000 computers at universities, government, and corporate research laboratories. Both the size and the use of the Internet continued to grow much faster than anticipated. By late 1987, it was estimated that the growth had reached 15% per month. By 1994, the global Internet reached over 3 million computers in 61 countries.

Adoption of TCP/IP protocols and growth of the Internet has not been limited to government-funded projects. Major computer corporations connected to the Internet as did many other large corporations including: oil companies, the auto industry, electronics firms, pharmaceutical companies, and telecommunications carriers. Medium and small companies began connecting in the 1990s. In addition, many companies have used the TCP/IP protocols on their internal corporate internets even though they choose not to be part of the global Internet.

[†]The term *NSFNET* is sometimes used loosely to mean all the NSF-funded networking activities, but we will use it to refer to the backbone. The next chapter gives more details about the technology.

Rapid expansion introduced problems of scale unanticipated in the original design and motivated researchers to find techniques for managing large, distributed resources. In the original design, for example, the names and addresses of all computers attached to the Internet were kept in a single file that was edited by hand and then distributed to every site on the Internet. By the mid 1980s, it became apparent that a central database would not suffice. First, requests to update the file would soon exceed the personnel available to process them. Second, even if a correct central file existed, network capacity was insufficient to allow either frequent distribution to every site or on-line access by each site.

New protocols were developed and a naming system was put in place across the global Internet that allows any user to resolve the name of a remote machine automatically. Known as the *Domain Name System*, the mechanism relies on machines called *name servers* to answer queries about names. No single machine contains the entire domain name database. Instead, data is distributed among a set of machines that use TCP/IP protocols to communicate among themselves when answering a query.

1.5 The Internet Architecture Board

Because the TCP/IP internet protocol suite did not arise from a specific vendor or from a recognized professional society, it is natural to ask, "who sets the technical direction and decides when protocols become standard?" The answer is a group known as the *Internet Architecture Board (IAB)*[†]. The IAB provides the focus and coordination for much of the research and development underlying the TCP/IP protocols, and guides the evolution of the Internet. It decides which protocols are a required part of the TCP/IP suite and sets official policies.

Formed in 1983 when ARPA reorganized the Internet Control and Configuration Board, the IAB inherited much of its charter from the earlier group. Its initial goals were to encourage the exchange of ideas among the principals involved in research related to TCP/IP and the Internet, and to keep researchers focused on common objectives. Through the first six years, the IAB evolved from an ARPA-specific research group into an autonomous organization. During these years, each member of the IAB chaired an *Internet Task Force* charged with investigating a problem or set of issues deemed to be important. The IAB consisted of approximately ten task forces, with charters ranging from one that investigated how the traffic load from various applications affects the Internet to one that handled short term Internet engineering problems. The IAB met several times each year to hear status reports from each task force, review and revise technical directions, discuss policies, and exchange information with representatives from agencies like ARPA and NSF who funded Internet operations and research.

[†]IAB originally stood for *Internet Activities Board*.

The original design distributed resources. Computers attached when distributed to a central database feed the personnel, network capacity on-line access by

in place across the machine automation machines called contains the entire machines that use a query.

specific vendor or sets the technical is a group known focus and coordination IP protocols, and a required part of

and Configuration up. Its initial goals involved in research on common objectives-specific research member of the IAB item or set of issues on task forces, with from various applications-engineering problems. Each task force, review ge information with internet operations and

The chairman of the IAB had the title *Internet Architect* and was responsible for suggesting technical directions and coordinating the activities of the various task forces. The IAB chairman established new task forces on the advice of the IAB and also represented the IAB to others.

Newcomers to TCP/IP are sometimes surprised to learn that the IAB did not manage a large budget; although it set direction, it did not fund most of the research and engineering it envisioned. Instead, volunteers performed much of the work. Members of the IAB were each responsible for recruiting volunteers to serve on their task forces, for calling and running task force meetings, and for reporting progress to the IAB. Usually, volunteers came from the research community or from commercial organizations that produced or used TCP/IP. Active researchers participated in Internet task force activities for two reasons. On one hand, serving on a task force provided opportunities to learn about new research problems. On the other hand, because new ideas and problem solutions designed and tested by task forces often became part of the TCP/IP Internet technology, members realized that their work had a direct, positive influence on the field.

1.6 The IAB Reorganization

By the summer of 1989, both the TCP/IP technology and the Internet had grown beyond the initial research project into production facilities on which thousands of people depended for daily business. It was no longer possible to introduce new ideas by changing a few installations overnight. To a large extent, the literally hundreds of commercial companies that offer TCP/IP products determined whether products would interoperate by deciding when to incorporate changes in their software. Researchers who drafted specifications and tested new ideas in laboratories could no longer expect instant acceptance and use of the ideas. It was ironic that the researchers who designed and watched TCP/IP develop found themselves overcome by the commercial success of their brainchild. In short, TCP/IP became a successful, production technology and the market place began to dominate its evolution.

To reflect the political and commercial realities of both TCP/IP and the Internet, the IAB was reorganized in the summer of 1989. The chairmanship changed. Researchers were moved from the IAB itself to a subsidiary group and a new IAB board was constituted to include representatives from the wider community.

Figure 1.1 illustrates the new IAB organization and the relationship of subgroups.

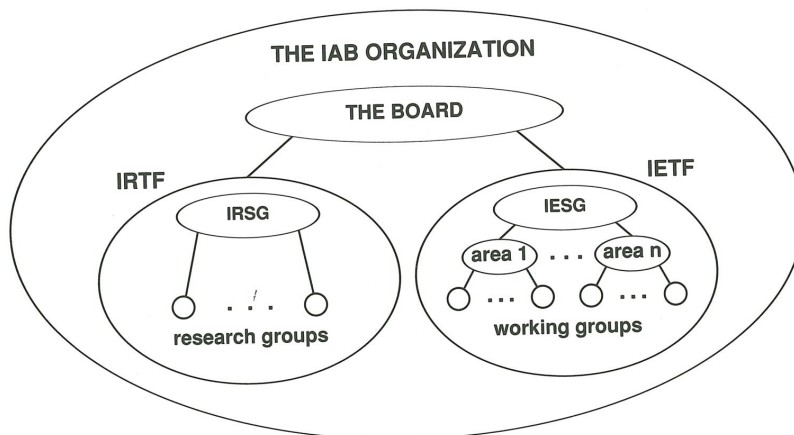
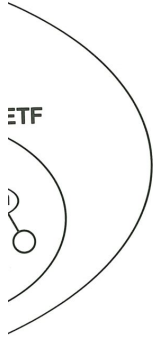


Figure 1.1 The structure of the IAB after the 1989 reorganization.

As Figure 1.1 shows, in addition to the board itself, the IAB organization contains two major groups: the *Internet Research Task Force (IRTF)* and the *Internet Engineering Task Force (IETF)*.

As its name implies, the IETF concentrates on short-term or medium-term engineering problems. The IETF existed in the original IAB structure, and its success provided part of the motivation for reorganization. Unlike most IAB task forces, which were limited to a few individuals who focused on one specific issue, the IETF grew to include dozens of active members who worked on many problems concurrently. Before the reorganization, the IETF was divided into over 20 *working groups*, each focusing on a specific problem. Working groups held individual meetings to formulate problem solutions. In addition, the entire IETF met regularly to hear reports from working groups and discuss proposed changes or additions to the TCP/IP technology. Usually held three times annually, full IETF meetings attracted hundreds of participants and spectators. The IETF had become too large for the chairman to manage.

Because the IETF was known throughout the Internet, and because its meetings were widely recognized and attended, the reorganized IAB structure retains the IETF, but splits it into approximately a dozen areas, each with its own manager. The IETF chairman and the area managers comprise the *Internet Engineering Steering Group (IESG)*, the individuals responsible for coordinating the efforts of IETF working groups. The name "IETF" now refers to the entire body, including the chairman, area managers, and all members of working groups.



on.

nization contains
Internet Engineer-

medium-term en-
s, and its success
ask forces, which
the IETF grew to
currently. Before
each focusing on
ormulate problem
rts from working
hology. Usually
f participants and
ge.

ause its meetings
retains the IETF,
anager. The IETF
g Steering Group
F working groups.
ie chairman, area

Created during the reorganization, the Internet Research Task Force is the research counterpart to the IETF. The IRTF coordinates research activities related to TCP/IP protocols or internet architecture in general. Like the IETF, the IRTF has a small group called the *Internet Research Steering Group* or *IRSG*, that sets priorities and coordinates research activities. Unlike the IETF, the IRTF is currently a much smaller and less active organization. Each member of the IRSG chairs a volunteer *Internet Research Group* analogous to the IETF working groups; the IRTF is not divided into areas.

1.7 The Internet Society

In 1992, as the Internet moved away from its U.S. government roots, a society was formed to encourage participation in the Internet. Called *The Internet Society*, the group is an international organization inspired by the National Geographic Society. The host for the IAB, the Internet Society continues to help people join and use the Internet around the world.

1.8 Internet Request For Comments

We have said that no vendor owns the TCP/IP technology nor does any professional society or standards body. Thus, the documentation of protocols, standards, and policies cannot be obtained from a vendor. Instead, the National Science Foundation funds a group at AT&T to maintain and distribute information about TCP/IP and the global Internet. Known as the *Internet Network Information Center (INTERNIC)*†, the INTERNIC handles many administrative details for the Internet in addition to distributing documentation.

Documentation of work on the Internet, proposals for new or revised protocols, and TCP/IP protocol standards all appear in a series of technical reports called *Internet Requests For Comments*, or *RFCs*. (Preliminary versions of RFCs are known as *Internet drafts*.) RFCs can be short or long, can cover broad concepts or details, and can be standards or merely proposals for new protocols‡. The RFC editor is a member of the IAB. While RFCs are edited, they are not refereed in the same way as academic research papers. Also, some reports pertinent to the Internet were published in an earlier, parallel series of reports called *Internet Engineering Notes*, or *IENs*. Although the IEN series is no longer active, not all IENs appear in the RFC series. There are references to RFCs and a few IENs throughout the text.

The RFC series is numbered sequentially in the chronological order RFCs are written. Each new or revised RFC is assigned a new number, so readers must be careful to obtain the highest numbered version of a document; an index is available to help identify the correct version.

To aid the INTERNIC and make document retrieval quicker, many sites around the world store copies of RFCs and make them available to the community. One can obtain RFCs by postal mail, by electronic mail, or directly across the Internet using a file

†Pronounced "Inter-Nick" after its acronym, the organization is a successor to the original Network Information Center (NIC).

‡Appendix 1 contains an introduction to RFCs that examines the diversity of RFCs, including jokes that have appeared.

transfer program. In addition, the INTERNIC and other organizations make available preliminary versions of RFC documents, known as *Internet drafts*. Ask a local network expert how to obtain RFCs or Internet drafts at your site, or refer to Appendix 1 for further instructions on how to retrieve them.

1.9 Internet Protocols And Standardization

Readers familiar with data communication networks realize that many communication protocol standards exist. Many of them precede the Internet, so the question arises, "Why did the Internet designers invent new protocols when so many international standards already existed?" The answer is complex, but follows a simple maxim:

Use existing protocol standards whenever such standards apply; invent new protocols only when existing standards are insufficient, and be prepared to use new standards when they become available and provide equivalent functionality.

So, despite appearances to the contrary, the TCP/IP Internet Protocol Suite was not intended to ignore or avoid extant standards. It came about merely because none of the existing protocols satisfied the need for an interoperable internetworking communication system.

1.10 Future Growth And Technology

Both the TCP/IP technology and the Internet continue to evolve. New protocols are being proposed; old ones are being revised. NSF added considerable complexity to the system by introducing a backbone network, regional networks, and hundreds of campus networks. Other groups around the world continue to connect to the Internet as well. The most significant change comes not from added network connections, however, but from additional traffic. As new users connect to the Internet and new applications appear, traffic patterns change. When physicists, chemists, and biologists began to use the Internet, they exchanged files of data collected from experiments. Such files seemed large compared to electronic mail messages. As the Internet became popular and users began to browse information using services like *gopher* and the *World Wide Web*, traffic increased again.

To accommodate growth in traffic, the capacity of the NSFNET backbone has already been increased three times, making the current capacity approximately 840 times larger than the original; an additional increase by another factor of 3 is scheduled for 1995. At the current time, it is difficult to foresee an end to the need for more capacity.

make available a local network Appendix 1 for

any communication question arises, international standard:

ply; internet, and able and

col Suite was not cause none of the communication

e. New protocols able complexity to and hundreds of t to the Internet as connections, however t and new application biologists began to iments. Such files et became popular nd the World Wide

T backbone has approximately 840 times 3 is scheduled for 1 for more capacity.

Growth in demands for networking should not be unexpected. The computer industry has enjoyed a continual demand for increased processing power and larger data storage for many years. Users have only begun to understand how to use networks. In the future we can expect continual increases in the demand for communications. Thus, higher-capacity communication technologies will be needed to accommodate the growth.

Figure 1.2 summarizes expansion of the Internet and illustrates an important component of growth: the change in complexity arises because multiple autonomous groups manage parts of the global Internet. The initial designs for many subsystems depended on centralized management. Much effort is needed to extend those designs to accommodate decentralized management.

	number of networks	number of computers	number of managers
1980	10	10 ²	10 ⁰
1990	10 ³	10 ⁵	10 ¹
1997	10 ⁶	10 ⁸	10 ²

Figure 1.2 Growth of the connected Internet. In addition to traffic increases that result from increased size, the Internet faces complexity that results from decentralized management of both development and operations.

1.11 Organization Of The Text

The material on TCP/IP has been written in three volumes. This volume presents the TCP/IP technology, applications that use it, and the architecture of the global Internet in more detail. It discusses the fundamentals of protocols like TCP and IP, and shows how they fit together in an internet. In addition to giving details, the text highlights the general principles underlying network protocols and explains why the TCP/IP protocols adapt easily to so many underlying physical network technologies. Volume II discusses in depth the internal details of the TCP/IP protocols and shows how they are implemented. It presents code from a working system to illustrate how the individual protocols work together, and contains details useful to people responsible for building a corporate internet. Volume III shows how distributed applications use TCP/IP for communication. It focuses on the client-server paradigm, the basis for all distributed programming. It discusses the interface between programs and protocols†, and shows how clients and server programs are organized. In addition, Volume III describes the remote procedure concept, and shows how programmers use tools to build client and server software.

†Volume III is available in two versions: one uses the socket interface and the other uses the Transport Layer Interface.

So far we have talked about the TCP/IP technology and the Internet in general terms, summarizing the services provided and the history of their development. The next chapter provides a brief summary of the type of network hardware used throughout the Internet. Its purpose is not to illuminate nuances of a particular vendor's hardware, but to focus on the features of each technology that are of primary importance to an internet architect. Later chapters delve into the protocols and the Internet, fulfilling three purposes: they explore general concepts and review the Internet architectural model, they examine the details of TCP/IP protocols, and they look at standards for high-level services like electronic mail and electronic file transfer. Chapters 3 through 12 review fundamental principles and describe the network protocol software found in any machine that uses TCP/IP. Later chapters describe services that span multiple machines, including the propagation of routing information, name resolution, and applications like electronic mail.

Two appendices follow the main text. The first appendix contains a guide to RFCs. It expands on the description of RFCs found in this chapter, and gives examples of information that can be found in RFCs. It describes in detail how to obtain RFCs by electronic mail, postal mail, and file transfer. Finally, because the standard RFC index comes in chronological order, the appendix presents a list of RFCs organized by topic to make it easier for beginners to find RFCs pertinent to a given subject.

The second appendix contains an alphabetical list of terms and abbreviations used throughout the literature and the text. Because beginners often find the new terminology overwhelming and difficult to remember, they are encouraged to use the alphabetical list instead of scanning back through the text.

1.12 Summary

An internet consists of a set of connected networks that act as a coordinated whole. The chief advantage of an internet is that it provides universal interconnection while allowing individual groups to use whatever network hardware is best suited to their needs. We will examine principles underlying internet communication in general and the details of one internet protocol suite in particular. We will also discuss how internet protocols are used in an internet. Our example technology, called TCP/IP after its two main protocols, was developed by the Advanced Research Projects Agency. It provides the basis for the global Internet, a large, operational internet that connects universities, corporations, and government departments in many countries around the world. The global Internet is expanding rapidly.

ernet in general
elopment. The
sed throughout
dor's hardware,
rtance to an in-
, fulfilling three
itectural model,
ls for high-level
ough 12 review
found in any
span multiple
tion, and appli-

ains a guide to
i gives examples
obtain RFCs by
rdard RFC index
ganized by topic
t.
breviations used
e new terminolo-
e the alphabetical

ordinated whole.
nnection while al-
ted to their needs.
eneral and the de-
how internet pro-
P/IP after its two
gency. It provides
nects universities,
d the world. The

FOR FURTHER STUDY

Cerf's *A History Of The ARPANET* [1989] and *History of the Internet Activities Board* [RFC 1160] provide fascinating reading and point the reader to early research papers on TCP/IP and internetworking. Denning [Nov-Dec 1989] provides a different perspective on the history of the ARPANET. Jennings et. al. [1986] discusses the importance of computer networking for scientists. Denning [Sept-Oct 1989] also points out the importance of internetworking and gives one possible scenario for a world-wide internet. The Federal Coordinating Committee for Science, Engineering and Technology [FCCSET] suggests networking should be a national priority.

The IETF publishes minutes from its regular meetings; these are available from the Corporation for National Research Initiatives in Reston, VA. The *Journal of Internetworking: Research and Experience* reports on internetworking research, with emphasis on experimental validation of ideas. The periodical *Connexions* [Jacobsen 1987-] contains articles about TCP/IP and the Internet as well as official statements of policy from the IAB. Finally, the reader is encouraged to remember that the TCP/IP protocol suite and the Internet continue to change; new information can be found in RFCs and at conferences such as the annual ACM SIGCOMM Symposium and Interop Company's NETWORLD+INTEROP events.

EXERCISES

- 1.1 Explore application programs at your site that use TCP/IP.
- 1.2 Find out whether your site connects to the Internet.
- 1.3 TCP/IP products account for over a billion dollars per year in gross revenue. Read trade publications to find a list of vendors offering such products.