



Third Edition

INTERNETWORKING WITH
TCP/IP

VOLUME I
PRINCIPLES, PROTOCOLS,
AND ARCHITECTURE



DOUGLAS E. COMER

ARRIS EX. 1018

Third Edition

INTERNETWORKING WITH TCP/IP

VOLUME I
PRINCIPLES, PROTOCOLS,
AND ARCHITECTURE

DOUGLAS E. COMER

Over 200,000 Copies Sold

"THE classic text for an introduction to TCP/IP."

—Jon Postel, RFC editor and former Deputy Internet Architect

"Although others have tried, there is no better written or organized explanation of the core of TCP/IP."

—Joel Snyder, *Network Computing*

"As an introduction to the TCP/IP protocol suite and its underpinnings, this is an excellent book. It is also a good reference book to keep around for anyone who is working with TCP/IP."

—George V. Neville-Neil, *USENIX ;login:*

The all-time best-selling TCP/IP book, *Internetworking with TCP/IP*, is still THE reference for anyone who wants to learn about or work with the TCP/IP protocol suite. Volume I of the series by Douglas Comer provides the most up-to-date conceptual introduction to TCP/IP protocols and the latest developments in Internet technology.

Renowned for its clarity and accessibility, this superb text covers wide area (WAN) Internet backbones as well as local area network (LAN) technologies like Ethernet and FDDI. The text explains address binding (ARP), IP connectionless datagram delivery, error detection, multicasting, and routing.

THIS NEW EDITION OF VOLUME I:

- Discusses how to use TCP/IP over an ATM network.
- Covers the latest IPng (next generation) developments and information.
- Describes CIDR (Classless Inter-Domain Routing) and supernetting.
- Discusses security in TCP/IP environments and firewall design.
- Categorizes hundreds of new RFCs and the protocols they describe.

In addition, Volume I:

- Compares the ISO 7-layer reference model to the TCP/IP 5-layer reference model.
- Explains TCP: reliability, acknowledgments, flow control, and sliding windows.
- Details adaptive retransmission, including slow-start and silly window avoidance.
- Describes the socket interface that applications use to access TCP/IP protocols.
- Presents routing architectures for large and small internets.
- Discusses bridges and routers.
- Examines application services:
 - Domain Name System (DNS)
 - Electronic mail (SMTP, MIME)
 - File transfer and access (FTP, TFTP, NFS)
 - Remote login (TELENET, rlogin)
 - Network management (SNMP, MIB, ANS.1)

PRENTICE HALL
Upper Saddle River, NJ 07458

ISBN 0-13-216987-8




90000



9 780132 169875



Internetworking With TCP/IP



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
Upper Saddle River, New Jersey 07458

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

005.2--dc20

95-1830

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
Upper Saddle River, New Jersey 07458

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

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

To Chris

Additional Enthusiastic Comments About *Internetworking With TCP/IP Volume 1*

"Unquestionably THE reference for TCP/IP; both informative and easy to read, this book is liked by both novice and experienced."

– Raj Yavatkar
University of Kentucky
US Editor, Computer Communications

"The third edition maintains Comer's Internetworking with TCP/IP as the acknowledged leader in TCP/IP books by adding up-to-the-minute material on ATM, CIDR, firewalls, DHCP and the next version of IP, IPng."

– Ralph Droms
Bucknell University
IETF Working Group Chair

"Doug Comer remains the first and best voice of Internet technology. Despite the legion of 'Internet carpetbaggers' (the current crop of 'authors' who can barely spell F-T-P) which contributes noise – but no knowledge – on the Internet and its infrastructure, Dr. Comer shines through as the premiere source for lucid explanations and accurate information. He sets a standard for which many strive, but precious few attain."

– Marshall Rose
Dover Beach Consulting
IETF Area Director

"Comer's Volume 1 drastically changed the course of networking history."

– Dan Lynch
Interop Company
IAB Member

"When you need to teach the details of TCP/IP, you need the latest information. Once again, Comer separates the chaff from the wheat with his latest edition of the TCP/IP book that a generation of networkers grew up with."

– Shawn Ostermann
Ohio University

Contents

Foreword **xxi**

Preface **xxiii**

Chapter 1 Introduction And Overview **1**

- 1.1 The Motivation For Internetworking* 1
- 1.2 The TCP/IP Internet* 2
- 1.3 Internet Services* 3
- 1.4 History And Scope Of The Internet* 6
- 1.5 The Internet Architecture Board* 8
- 1.6 The IAB Reorganization* 9
- 1.7 The Internet Society* 11
- 1.8 Internet Request For Comments* 11
- 1.9 Internet Protocols And Standardization* 12
- 1.10 Future Growth And Technology* 12
- 1.11 Organization Of The Text* 13
- 1.12 Summary* 14

Chapter 2 Review Of Underlying Network Technologies **17**

- 2.1 Introduction* 17
- 2.2 Two Approaches To Network Communication* 18
- 2.3 Wide Area And Local Area Networks* 19
- 2.4 Ethernet Technology* 20
- 2.5 Fiber Distributed Data Interconnect (FDDI)* 32
- 2.6 Asynchronous Transfer Mode* 36
- 2.7 ARPANET Technology* 37
- 2.8 National Science Foundation Networking* 39
- 2.9 ANSNET* 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

5.5	<i>Resolution Through Dynamic Binding</i>	75
5.6	<i>The Address Resolution Cache</i>	76
5.7	<i>ARP Refinements</i>	77
5.8	<i>Relationship Of ARP To Other Protocols</i>	77
5.9	<i>ARP Implementation</i>	77
5.10	<i>ARP Encapsulation And Identification</i>	79
5.11	<i>ARP Protocol Format</i>	79
5.12	<i>Summary</i>	81
 Chapter 6 Determining An Internet Address At Startup (RARP)		83
6.1	<i>Introduction</i>	83
6.2	<i>Reverse Address Resolution Protocol (RARP)</i>	84
6.3	<i>Timing RARP Transactions</i>	86
6.4	<i>Primary And Backup RARP Servers</i>	86
6.5	<i>Summary</i>	87
 Chapter 7 Internet Protocol: Connectionless Datagram Delivery		89
7.1	<i>Introduction</i>	89
7.2	<i>A Virtual Network</i>	89
7.3	<i>Internet Architecture And Philosophy</i>	90
7.4	<i>The Concept Of Unreliable Delivery</i>	90
7.5	<i>Connectionless Delivery System</i>	91
7.6	<i>Purpose Of The Internet Protocol</i>	91
7.7	<i>The Internet Datagram</i>	91
7.8	<i>Internet Datagram Options</i>	100
7.9	<i>Summary</i>	106
 Chapter 8 Internet Protocol: Routing IP Datagrams		109
8.1	<i>Introduction</i>	109
8.2	<i>Routing In An Internet</i>	109
8.3	<i>Direct And Indirect Delivery</i>	111
8.4	<i>Table-Driven IP Routing</i>	113
8.5	<i>Next-Hop Routing</i>	113
8.6	<i>Default Routes</i>	115
8.7	<i>Host-Specific Routes</i>	115
8.8	<i>The IP Routing Algorithm</i>	116
8.9	<i>Routing With IP Addresses</i>	116
8.10	<i>Handling Incoming Datagrams</i>	118

- 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

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

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

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

Chapter 17 Internet Multicasting (IGMP) 289

- 17.1 *Introduction* 289
- 17.2 *Hardware Broadcast* 289
- 17.3 *Hardware Multicast* 290
- 17.4 *IP Multicast* 291
- 17.5 *IP Multicast Addresses* 291
- 17.6 *Mapping IP Multicast To Ethernet Multicast* 292
- 17.7 *Extending IP To Handle Multicasting* 293
- 17.8 *Internet Group Management Protocol* 294
- 17.9 *IGMP Implementation* 294
- 17.10 *Group Membership State Transitions* 295
- 17.11 *IGMP Message Format* 296
- 17.12 *Multicast Address Assignment* 297
- 17.13 *Propagating Routing Information* 297
- 17.14 *The Mrouted Program* 298
- 17.15 *Summary* 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) 365

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

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

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	<i>Formal Definitions Using ASN.1</i>	453
26.8	<i>Structure And Representation Of MIB Object Names</i>	453
26.9	<i>Simple Network Management Protocol</i>	458
26.10	<i>SNMP Message Format</i>	460
26.11	<i>Example Encoded SNMP Message</i>	462
26.12	<i>Summary</i>	463
Chapter 27 Summary Of Protocol Dependencies		465
27.1	<i>Introduction</i>	465
27.2	<i>Protocol Dependencies</i>	465
27.3	<i>Application Program Access</i>	467
27.4	<i>Summary</i>	468
Chapter 28 Internet Security And Firewall Design		471
28.1	<i>Introduction</i>	471
28.2	<i>Protecting Resources</i>	472
28.3	<i>The Need For An Information Policy</i>	472
28.4	<i>Communication, Cooperation, And Mutual Mistrust</i>	474
28.5	<i>Mechanisms For Internet Security</i>	475
28.6	<i>Firewalls And Internet Access</i>	476
28.7	<i>Multiple Connections And Weakest Links</i>	477
28.8	<i>Firewall Implementation And High-Speed Hardware</i>	478
28.9	<i>Packet-Level Filters</i>	479
28.10	<i>Security And Packet Filter Specification</i>	480
28.11	<i>The Consequence Of Restricted Access For Clients</i>	481
28.12	<i>Accessing Services Through A Firewall</i>	481
28.13	<i>The Details Of Firewall Architecture</i>	483
28.14	<i>Stub Network</i>	484
28.15	<i>An Alternative Firewall Implementation</i>	484
28.16	<i>Monitoring And Logging</i>	485
28.17	<i>Summary</i>	486
Chapter 29 The Future Of TCP/IP (IPng, IPv6)		489
29.1	<i>Introduction</i>	489
29.2	<i>Why Change TCP/IP And The Internet?</i>	490
29.3	<i>Motivation For Changing IPv4</i>	491
29.4	<i>The Road To A New Version Of IP</i>	492
29.5	<i>The Name Of The Next IP</i>	492

29.6	<i>Features Of IPv6</i>	493
29.7	<i>General Form Of An IPv6 Datagram</i>	494
29.8	<i>IPv6 Base Header Format</i>	494
29.9	<i>IPv6 Extension Headers</i>	496
29.10	<i>Parsing An IPv6 Datagram</i>	497
29.11	<i>IPv6 Fragmentation And Reassembly</i>	498
29.12	<i>The Consequence Of End-To-End Fragmentation</i>	498
29.13	<i>IPv6 Source Routing</i>	500
29.14	<i>IPv6 Options</i>	500
29.15	<i>Size Of The IPv6 Address Space</i>	502
29.16	<i>IPv6 Colon Hexadecimal Notation</i>	502
29.17	<i>Three Basic IPv6 Address Types</i>	503
29.18	<i>The Duality Of Broadcast And Multicast</i>	504
29.19	<i>An Engineering Choice And Simulated Broadcast</i>	504
29.20	<i>Proposed IPv6 Address Space Assignment</i>	504
29.21	<i>IPv4 Address Encoding And Transition</i>	506
29.22	<i>Providers, Subscribers, And Address Hierarchy</i>	506
29.23	<i>Additional Hierarchy</i>	507
29.24	<i>Summary</i>	508
Appendix 1	A Guide To RFCs	511
Appendix 2	Glossary Of Internetworking Terms And Abbreviations	557
Bibliography		591
Index		599



Foreword

Professor Douglas Comer's book has become *the* classic text for an introduction to TCP/IP. Writing an introduction to TCP/IP for the uninitiated is a very difficult task. While combining the explanation of the general principles of computer communication with the specific examples from the TCP/IP protocol suite, Doug Comer has provided a very readable book.

While this book is specifically about the TCP/IP protocol suite, it is a good book for learning about computer communications protocols in general. The principles of architecture, layering, multiplexing, encapsulation, addressing and address mapping, routing, and naming are quite similar in any protocol suite, though, of course, different in detail.

Computer communication protocols do not do anything themselves. Like operating systems, they are in the service of application processes. Processes are the active elements that request communication and are the ultimate senders and receivers of the data transmitted. The various layers of protocols are like the various layers in a computer operating system, especially the file system. Understanding protocol architecture is like understanding operating system architecture. In this book Doug Comer has taken the "bottom up" approach – starting with the physical networks and moving up in levels of abstraction to the applications.

Since application processes are the active elements using the communication supported by the protocols, TCP/IP is an "interprocess communication" (IPC) mechanism. While there are several experiments in progress with operating system style message passing and procedure call types of IPC based on IP, the focus in this book is on more traditional applications that use the UDP datagram or TCP logical connection forms of IPC. Typically in operating systems there is a set of functions provided by the operating system to the application processes. This system call interface usually includes calls for opening, reading, writing, and closing files, among other things. In many systems there are similar system calls for IPC functions including network communication. As an example of such an interface Doug Comer presents an overview of the socket interface.

One of the key ideas inherent in TCP/IP and in the title of this book is "internetworking." The power of a communication system is directly related to the number of entities in that system. The telephone network is very useful because (nearly) all the telephones are connected to one network (as it appears to the users). Computer communication systems and networks are currently separated and fragmented. As more users and enterprises adopt TCP/IP as their network communication technology and are joining the Internet this is becoming less of a problem, but there is still a long way to

go. The goal of interconnection and internetworking, to have a single powerful computer communication network, is fundamental to the design of TCP/IP.

Essential to internetworking is addressing, and a universal protocol – the Internet Protocol. Of course, the individual networks have their own protocols which are used to carry the IP datagrams, and there must be a mapping between the individual network address and the IP address. Over the lifetime of TCP/IP, the nature of these individual networks have changed from the early days of the ARPANET to the recently developed ATM networks. A new chapter in this edition discusses IP over ATM networks. This book now includes recent developments in Dynamic Host Configuration (DHCP) that will ease the administration of networks and the installation of new computers.

To have an internetwork, the individual networks must be connected. The connecting devices are called routers. Further, these routers must have some procedures for forwarding data from one network to the next. The data is in the form of IP datagrams and the destination is specified by an IP address, but the router must make a routing decision based on the IP address and what it knows about the connectivity of the networks making up the Internet. The procedures for distributing the current connectivity information to the routers are called routing algorithms, and these are currently the subject of much study and development. In particular, the recent development of the Classless InterDomain Routing (CIDR) technique to reduce the amount of routing information exchanged is important.

Like all communication systems, the TCP/IP protocol suite is an unfinished system. It is evolving to meet changing requirements and new opportunities. Thus, this book is, in a sense, a snapshot of TCP/IP. And, as Doug Comer points out, there are many loose ends. With the recent rapid growth of the Internet there is concern about it outgrowing the capabilities of the TCP/IP protocols, particularly the address space. In response the research and engineering community has developed a “next generation” version of the Internet Protocol called IPng. Many of the enterprises now joining the Internet have concerns about security. A new chapter in this edition discusses the security and firewalls.

Most chapters end with a few pointers to material “for further study.” Many of these refer to memos of the RFC series of notes. This series of notes is the result of a policy of making the working ideas and the protocol specifications developed by the TCP/IP research and development community widely available. This availability of the basic and detailed information about these protocols, and the availability of the early implementations of them, has had much to do with their current widespread use. This commitment to public documentation at this level of detail is unusual for a research effort, and has had significant benefits for the development of computer communication.

This book brings together information about the various parts of the TCP/IP architecture and protocols and makes it accessible. Its publication is a very significant milestone in the evolution of computer communications.

Jon Postel,
Associate Director for Networking
Information Sciences Institute
University of Southern California

January 1995

Preface

The world has changed dramatically since the second edition of this book was published. It hardly seems possible only four years have elapsed. When I began the second edition in the summer of 1990, the Internet had grown to nearly 300,000 host computers, up from 5,000 hosts when the book was first written. At the time, we marveled at how large an obscure research project had become. Cynics predicted that continued growth would lead to a complete collapse by 1993. Instead of collapsing, the Internet has continued its explosive expansion; the “large” Internet of 1990 is only 7% of the current Internet.

TCP/IP and the Internet have accommodated change well. The basic technology has survived over a decade of exponential growth and the associated increases in traffic. The protocols have worked over new high-speed network technologies, and the design has handled applications that could not be imagined a decade ago. Of course, the entire protocol suite has not remained static. New protocols have been deployed, and new techniques have been developed to adapt existing protocols to new network technologies. Changes are documented in RFCs, which have increased by over 50 percent.

This edition contains updated information throughout the text (including use of the commercially popular term *IP router* in place of the traditional scientific term *IP gateway*) as well as new material that describes technical advances and changes. The chapter on subnet addressing now describes supernetting as well as subnetting, and shows how the two techniques are motivated by the same goal. The chapter on bootstrapping explains a significant advance that will eliminate the need for manual configuration of host computers and allow a computer to obtain an IP address automatically: the Dynamic Host Configuration Protocol (DHCP). The chapter on TCP includes a description of Silly Window Syndrome and an explanation of the heuristics TCP uses to prevent the problem. The chapter on electronic mail includes a description of the Multipurpose Internet Mail Extensions (MIME), which permit non-ASCII data to be sent in a standard e-mail message.

Three new chapters contain detailed information about significant developments. Chapter 18 explains how TCP/IP is being used over ATM networks. The chapter discusses the organization of ATM hardware, the purpose of adaptation layer protocols, IP encapsulation, address binding, routing, and virtual circuit management. The chapter illustrates how a connectionless protocol like IP can use the connection-oriented interface that ATM provides. Chapter 28 covers a topic that is crucial to many organizations as they contemplate connecting to the global Internet – security. The chapter describes the internet firewall concept, and shows how a firewall architecture can be

used to protect networks and computers inside an organization from unwanted access. The chapter also discusses the principles underlying a two-level firewall design, and considers outside access from a secure computer. Finally, a new chapter is devoted to what may be the most significant change in TCP/IP since its inception: the imminent adoption of a next generation Internet Protocol (IPng). Chapter 29 describes the protocol that the IETF has developed to serve as IPng. Although it has not been thoroughly tested or approved as a permanent standard, the new design appears to be the consensus choice. The chapter presents the proposed design and address assignment scheme.

The third edition retains the same general contents and overall organization as the second edition. The entire text focuses on the concept of internetworking in general and the TCP/IP internet technology in particular. Internetworking is a powerful abstraction that allows us to deal with the complexity of multiple underlying communication technologies. It hides the details of network hardware and provides a high level communication environment. The text reviews both the architecture of network interconnections and the principles underlying protocols that make such interconnected networks function as a single, unified communication system. It also shows how an internet communication system can be used for distributed computation.

After reading this book, you will understand how it is possible to interconnect multiple physical networks into a coordinated system, how internet protocols operate in that environment, and how application programs use the resulting system. As a specific example, you will learn the details of the global TCP/IP Internet, including the architecture of its router system and the application protocols it supports. In addition, you will understand some of the limitations of the internet approach.

Designed as both a college text and as a professional reference, the book is written at an advanced undergraduate or graduate level. For professionals, the book provides a comprehensive introduction to the TCP/IP technology and the architecture of the Internet. Although it is not intended to replace protocol standards, the book is an excellent starting point for learning about internetworking because it provides a uniform overview that emphasizes principles. Moreover, it gives the reader perspective that can be extremely difficult to obtain from individual protocol documents.

When used in the classroom, the text provides more than sufficient material for a single semester network course at either the undergraduate or graduate level. Such a course can be extended to a two-semester sequence if accompanied by programming projects and readings from the literature. For undergraduate courses, many of the details are unnecessary. Students should be expected to grasp the basic concepts described in the text, and they should be able to describe or use them. At the graduate level, students should be expected to use the material here as a basis for further exploration. They should understand the details well enough to answer exercises or solve problems that require them to explore extensions and subtleties. Many of the exercises suggest such subtleties; solving them often requires students to read protocol standards and apply creative energy to comprehend consequences.

At all levels, hands-on experience sharpens the concepts and helps students gain intuition. Thus, I encourage instructors to invent projects that force students to use Internet services and protocols. The semester project in my graduate Internetworking

course at Purdue requires students to build an IP router. We supply hardware and the source code for an operating system, including device drivers for network interfaces; students build a working router that interconnects three networks with different MTUs. The course is extremely rigorous, students work in teams, and the results have been impressive (many industries recruit graduates from the course). Although such experimentation is safest when the instructional laboratory network is isolated from production computing facilities, we have found that students exhibit the most enthusiasm, and benefit the most, when they have access to a functional TCP/IP internet.

The book is organized into four main parts. Chapters 1 and 2 form an introduction that provides an overview and discusses existing network technologies. In particular, Chapter 2 reviews physical network hardware. The intention is to provide basic intuition about what is possible, not to spend inordinate time on hardware details. Chapters 3-13 describe the TCP/IP Internet from the viewpoint of a single host, showing the protocols a host contains and how they operate. They cover the basics of Internet addressing and routing as well as the notion of protocol layering. Chapters 14-18 and 28 describe the architecture of an internet when viewed globally. They explore routing architecture and the protocols routers use to exchange routing information. Finally, Chapters 19-27 discuss application level services available in the Internet. They present the client-server model of interaction, and give several examples of client and server software.

The chapters have been organized bottom up. They begin with an overview of hardware and continue to build new functionality on top of it. This view will appeal to anyone who has developed Internet software because it follows the same pattern one uses in implementation. The concept of layering does not appear until Chapter 11. The discussion of layering emphasizes the distinction between conceptual layers of functionality and the reality of layered protocol software in which multiple objects appear at each layer.

A modest background is required to understand the material. The reader is expected to have a basic understanding of computer systems, and to be familiar with data structures like stacks, queues, and trees. Readers need basic intuition about the organization of computer software into an operating system that supports concurrent programming and application programs that users invoke to perform computation. Readers do not need sophisticated mathematics, nor do they need to know information theory or theorems from data communications; the book describes the physical network as a black box around which an internetwork can be built. It states design principles in English and discusses motivations and consequences.

I thank all the people who have contributed to versions of this book. John Lin provided extensive assistance with this edition, including classifying RFCs. Ralph Droms reviewed the chapter on bootstrapping, and Sandeep Kumar, Steve Lodin, and Christoph Schuba, from the COAST security project at Purdue, commented on the security chapter. Special thanks go to my wife, Chris, whose careful editing made many improvements in wording.

FOR FURTHER STUDY

The protocols described in this chapter are all specified in Internet RFCs. Postel [RFC 821] describes the Simple Mail Transfer Protocol and gives many examples. The exact format of mail messages is given by Crocker [RFC 822]. Borenstein and Freed [RFC 1521] specifies the standard for MIME, including the syntax of header declarations, the interpretation of content types, and the *base64* encoding mentioned in this chapter. Moore [RFC 1522] defines MIME header extensions for non-ASCII text, and Postel [RFC 1590] describes the procedure for registration of new content and encoding types. Partridge [RFC 974] discusses the relationship between mail routing and the domain name system. Horton [RFC 976] proposes a standard for the UNIX UUCP mail system.

EXERCISES

- 25.1 Some mail systems force the user to specify a sequence of machines through which the message should travel to reach its destination. The mail protocol in each machine merely passes the message on to the next machine. List three disadvantages of such a scheme.
- 25.2 Find out if your computing system allows you to invoke SMTP directly.
- 25.3 Build an SMTP client and use it to deliver a mail message.
- 25.4 See if you can send mail through a mail gateway and back to yourself.
- 25.5 Make a list of mail address forms that your site handles and write a set of rules for parsing them.
- 25.6 Find out how the UNIX *sendmail* program can be used to implement a mail gateway.
- 25.7 Find out how often your local mail system attempts delivery and how long it will continue before giving up.
- 25.8 Many mail systems allow users to direct incoming mail to a program instead of storing it in a mailbox. Build a program that accepts your incoming mail, places your mail in a file, and then sends a reply to tell the sender you are on vacation.
- 25.9 Read the SMTP standard carefully. Then use TELNET to connect to the SMTP port on a remote machine and ask the remote SMTP server to expand a mail alias.
- 25.10 A user receives mail in which the *To* field specifies the string *important-people*. The mail was sent from a computer on which the alias *important-people* includes no valid mailbox identifiers. Read the SMTP specification carefully to see how such a situation is possible.
- 25.11 Read the MIME standard carefully. What servers can be specified in a MIME external reference?

Applications: Internet Management (SNMP, SNMPv2)

26.1 Introduction

In addition to protocols that provide network level services and application programs that use those services, an internet needs software that allows managers to debug problems, control routing, and find computers that violate protocol standards. We refer to such activities as *internet management*. This chapter considers the ideas behind TCP/IP internet management software, and describes an internet management protocol.

26.2 The Level Of Management Protocols

Originally, many wide area networks included management protocols as part of their link level protocols. If a packet switch began misbehaving, the network manager could instruct a neighboring packet switch to send it a special *control packet*. Control packets caused the receiver to suspend normal operation and respond to commands from the manager. The manager could interrogate the packet switch to identify problems, examine or change routes, test one of the communication interfaces, or reboot the switch. Once managers repaired the problem, they could instruct the switch to resume normal operations. Because management tools were part of the lowest level protocol, managers were often able to control switches even if higher level protocols failed.

Unlike a homogeneous wide area network, a TCP/IP internet does not have a single link level protocol. Instead, the internet consists of multiple physical networks interconnected by IP routers. As a result, internet management differs from network management. First, a single manager can control heterogeneous routers[†]. Second, the controlled entities may not share a common link level protocol. Third, the set of machines a manager controls may lie at arbitrary points in an internet. In particular, a manager may need to control one or more machines that do not attach to the same physical network as the manager's computer. Thus, it may not be possible for a manager to communicate with machines being controlled unless the management software uses protocols that provide end-to-end connectivity across an internet. As a consequence, the internet management protocol used with TCP/IP operates above the transport level:

In a TCP/IP internet, IP routers form the active switches that managers need to examine and control. Because routers attach to heterogeneous networks, protocols for internet management operate at the application level and communicate using TCP/IP transport-level protocols.

Designing internet management software to operate at the application level has several advantages. Because the protocols can be designed without regard to the underlying network hardware, one set of protocols can be used for all networks. Because the protocols can be designed without regard to the hardware on the managed machine, the same protocols can be used for all managed devices. From a manager's point of view, having a single set of management protocols means uniformity – all routers respond to exactly the same set of commands. Furthermore, because the management software uses IP for communication, a manager can control the routers across an entire TCP/IP internet without having direct attachment to every physical network or router.

Of course, building management software at the application level also has disadvantages. Unless the operating system, IP software, and transport protocol software work correctly, the manager may not be able to contact the router. For example, if the router's routing table becomes damaged, it may be impossible to correct the table or reboot the machine from a remote site. If the operating system on a router crashes, it will be impossible to reach the application program that implements the internet management protocols even if the router can still field hardware interrupts and route packets.

26.3 Architectural Model

Despite the potential disadvantages, having TCP/IP management software operate at the application level has worked well in practice. The most significant advantage of placing network management protocols at a high level becomes apparent when one considers a large internet, where a manager's computer does not need to attach directly to all physical networks that contain managed entities. Figure 26.1 shows an example of the architecture.

[†]Although managers can control both routers and hosts, we will focus on control of routers because they present the most complexity.

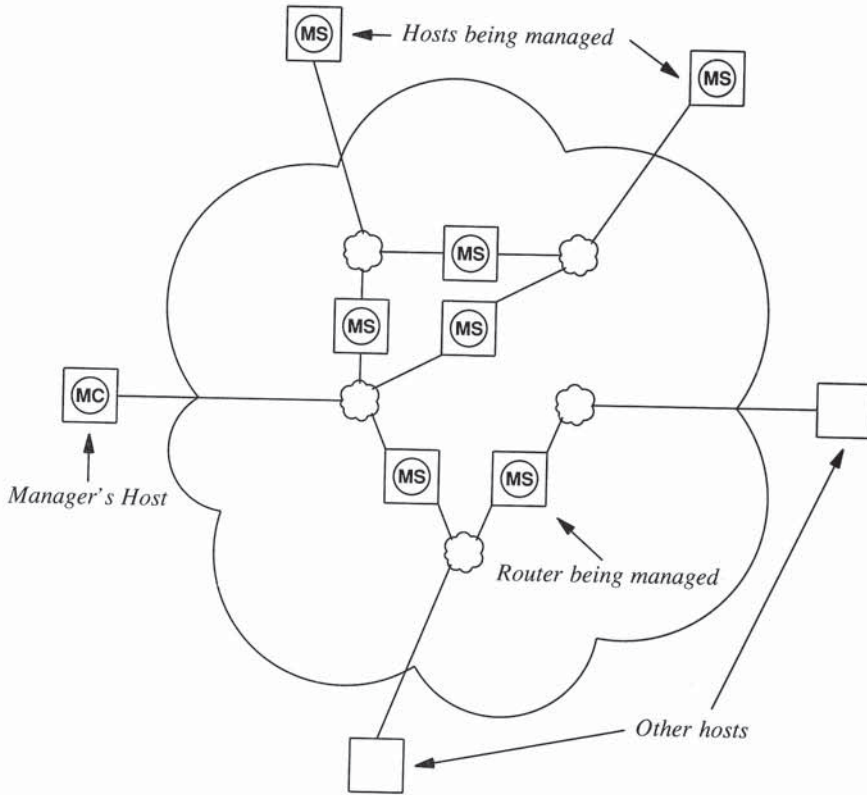


Figure 26.1 Example of network management. A manager invokes management client (MC) software that contacts management server (MS) software on routers throughout the internet.

As the figure shows, each participating host or router runs a server program. Technically, the server is called a *management agent*. A manager invokes client software on the local host computer and specifies an agent with which it communicates. After the client contacts the agent, it sends queries to obtain information or it sends commands to change conditions in the router. Of course, not all routers in a large internet fall under a single manager. Most managers only control a few routers at their local sites.

Internet management software uses an authentication mechanism to ensure only authorized managers can access or control a particular router. Some management protocols support multiple levels of authorization, allowing a manager specific privileges on each router. For example, a specific router could be configured to allow several managers to obtain information while only allowing a select subset of them to change information or control the router.

26.4 Protocol Architecture

TCP/IP network† management protocols divide the management problem into two parts and specify separate standards for each part. The first part concerns communication of information. A protocol specifies how client software running on a manager's host communicates with an agent. The protocol defines the format and meaning of messages clients and servers exchange as well as the form of names and addresses. The second part concerns the data being managed. A protocol specifies which data items a router must keep as well as the name of each data item and the syntax used to express the name.

26.4.1 A Standard Network Management Protocol

The current standard TCP/IP network management protocol is the *Simple Network Management Protocol (SNMP)*. A second version has been approved, but is not widely in use at the time this is being written. Known as *SNMPv2*, the new version adds new capabilities, including stronger security.

26.4.2 A Standard For Managed Information

A router being managed must keep control and status information that the manager can access. For example, a router keeps statistics on the status of its network interfaces, incoming and outgoing traffic, dropped datagrams, and error messages generated. Although it allows a manager to access these statistics, SNMP does not specify exactly which data can be accessed. Instead, a separate standard specifies the details. Known as a *Management Information Base (MIB)*, the standard specifies the data items a host or router must keep and the operations allowed on each. For example, the MIB specifies that IP software must keep a count of all octets that arrive over each network interface, and it specifies that network management software can only read those values.

The MIB for TCP/IP divides management information into eight categories as Figure 26.2 shows. The choice of categories is important because identifiers used to specify items include a code for the category.

†Technically, there is a distinction between internet management protocols and network management protocols. Historically, however, TCP/IP internet management protocols are known as *network management protocols*; we will follow the accepted terminology.

MIB category	Includes Information About
system	The host or router operating system
interfaces	Individual network interfaces
addr. trans.	Address translation (e.g., ARP mappings)
ip	Internet Protocol software
icmp	Internet Control Message Protocol software
tcp	Transmission Control Protocol software
udp	User Datagram Protocol software
egp	Exterior Gateway Protocol software

Figure 26.2 Categories of information in the MIB. The category is encoded in the identifier used to specify an object.

Keeping the MIB definition independent of the network management protocol has advantages for both vendors and users. A vendor can include SNMP agent software in a product such as a router, with the guarantee that the software will continue to adhere to the standard after new MIB items are defined. A customer can use the same network management client software to manage multiple routers that have different versions of a MIB. Of course, a router that does not have new MIB items cannot provide the information in those items. However, because all routers use the same language for communication, they can all parse a query and either provide the requested information or send an error message explaining that they do not have the requested item.

26.5 Examples of MIB Variables

In addition to the standard TCP/IP MIB, which is known as *MIB-II*, many RFCs document MIB variables for specific devices. Examining a few of the data items the standard MIB includes will help clarify the contents. Figure 26.3 lists example MIB variables along with their categories.

MIB Variable	Category	Meaning
sysUpTime	system	Time since last reboot
ifNumber	interfaces	Number of network interfaces
ifMtu	interfaces	MTU for a particular interface
ipDefaultTTL	ip	Value IP uses in time-to-live field
ipInReceives	ip	Number of datagrams received
ipForwDatagrams	ip	Number of datagrams forwarded
ipOutNoRoutes	ip	Number of routing failures
ipReasmOKs	ip	Number of datagrams reassembled
ipFragOKs	ip	Number of datagrams fragmented
ipRoutingTable	ip	IP Routing table
icmpInEchos	icmp	Number of ICMP Echo Requests received
tcpRtoMin	tcp	Minimum retransmission time TCP allows
tcpMaxConn	tcp	Maximum TCP connections allowed
tcpInSegs	tcp	Number of segments TCP has received
udpInDatagrams	udp	Number of UDP datagrams received
egpInMsgs	egp	Number of EGP messages received

Figure 26.3 Examples of MIB variables along with their categories.

Values for most of the items listed in Figure 26.3 can be stored in a single integer. However, the MIB also defines more complex structures. For example, the MIB variable *ipRoutingTable* refers to a router's routing table. Additional MIB variables define the contents of a routing table entry, and allow the network management protocols to reference the data for individual entries. Of course, MIB variables present only a logical definition of each data item – the internal data structures a router uses may differ from the MIB definition. When a query arrives, software in the agent on the router is responsible for mapping between the MIB variable and the data structure the router uses to store the information.

26.6 The Structure Of Management Information

In addition to the MIB standard, which specifies network management variables and their meanings, a separate standard specifies a set of rules used to define and identify MIB variables. The rules are known as the *Structure of Management Information (SMI)* specification. To keep network management protocols simple, the SMI places restrictions on the types of variables allowed in the MIB, specifies the rules for naming those variables, and creates rules for defining variable types. For example, the SMI standard includes definitions of terms like *IpAddress* (defining it to be a 4-octet string) and *Counter* (defining it to be an integer in the range of 0 to $2^{32} - 1$), and specifies that they are the terms used to define MIB variables. More important, the rules in the SMI describe how the MIB refers to tables of values (e.g., the IP routing table).

26.7 Formal Definitions Using ASN.1

The SMI standard specifies that all MIB variables must be defined and referenced using ISO's *Abstract Syntax Notation 1 (ASN.1)*[†]. ASN.1 is a formal language that has two main features: a notation used in documents that humans read, and a compact encoded representation of the same information used in communication protocols. In both cases, the precise, formal notation removes any possible ambiguities from both the representation and meaning. For example, instead of saying that a variable contains an integer value, a protocol designer who uses ASN.1 must state the exact form and range of numeric values. Such precision is especially important when implementations include heterogeneous computers that do not all use the same representations for data items.

Besides keeping standards documents unambiguous, ASN.1 also helps simplify the implementation of network management protocols and guarantees interoperability. It defines precisely how to encode both names and data items in a message. Thus, once the documentation of a MIB has been expressed using ASN.1, the human readable form can be translated directly and mechanically into the encoded form used in messages. In summary:

The TCP/IP network management protocols use a formal notation called ASN.1 to define names and types for variables in the management information base. The precise notation makes the form and contents of variables unambiguous.

26.8 Structure And Representation Of MIB Object Names

We said that ASN.1 specifies how to represent both data items and names. However, understanding the names used for MIB variables requires us to know about the underlying namespace. Names used for MIB variables are taken from the *object identifier* namespace administered by ISO and ITU. The key idea behind the object identifier namespace is that it provides a namespace in which all possible objects can be named. The namespace is not restricted to variables used in network management – it includes names for arbitrary objects (e.g., each international protocol standard document has a name).

The object identifier namespace is *absolute (global)*, meaning that names are structured to make them globally unique. Like most namespaces that are large and absolute, the object identifier namespace is hierarchical. Authority for parts of the namespace is subdivided at each level, allowing individual groups to obtain authority to assign some of the names without consulting a central authority for each assignment[‡].

The root of the object identifier hierarchy is unnamed, but has three direct descendants managed by: ISO, ITU, and jointly by ISO and ITU. The descendants are assigned both short text strings and integers to identify them (the text strings are used

[†]ASN.1 is usually pronounced by reading the dot: 'A-S-N dot 1'.

[‡]Readers should recall from the Domain Name System discussion in Chapter 22 how authority for a hierarchical namespace is subdivided.

when humans need to understand object names; computer software uses the integers to form compact, encoded representations of the names). ISO has allocated one subtree for use by other national or international standards organizations (including U.S. standards organizations), and the U.S. National Institute for Standards and Technology† has allocated a subtree for the U.S. Department of Defense. Finally, the IAB has petitioned the Department of Defense to allocate it a subtree in the namespace.

Figure 26.4 illustrates pertinent parts of the object identifier hierarchy and shows the position of the node used by TCP/IP network management protocols.

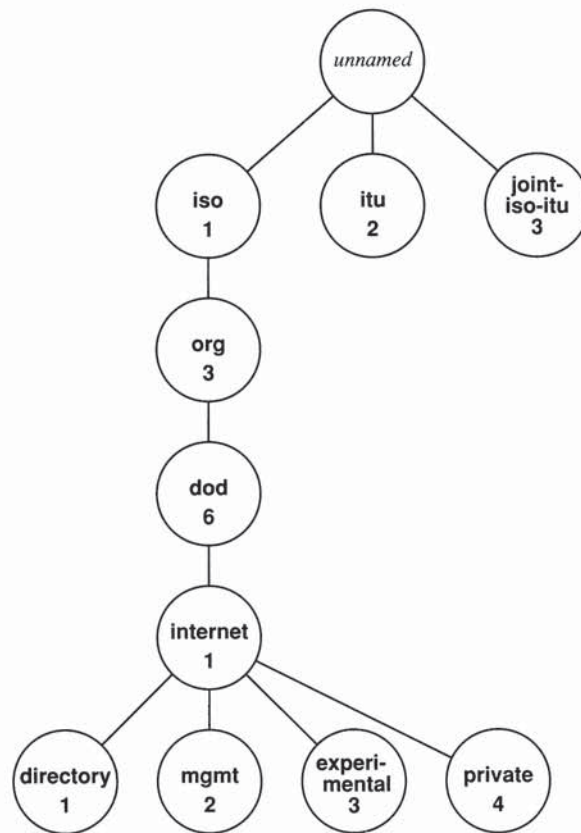


Figure 26.4 Part of the hierarchical object identifier namespace used to name MIB variables. An object's name consists of the numeric labels along a path from the root to the object.

†NIST was formerly the National Bureau of Standards.

The name of an object in the hierarchy is the sequence of numeric labels on the nodes along a path from the root to the object. The sequence is written with periods separating the individual components. For example, the name *1.3.6.1.1* denotes the node labeled *directory*. The MIB has been assigned a node under the *internet mgmt* subtree with label *mib* and numeric value *1*. Because all MIB variables fall under that node, they all have names beginning with the prefix *1.3.6.1.2.1*.

Earlier we said that the MIB groups all variables into eight categories. The exact meaning of the categories can now be explained: they are the eight subtrees of the *mib* node of the object identifier namespace. Figure 26.5 illustrates the idea by showing part of the naming subtree under the *mib* node.

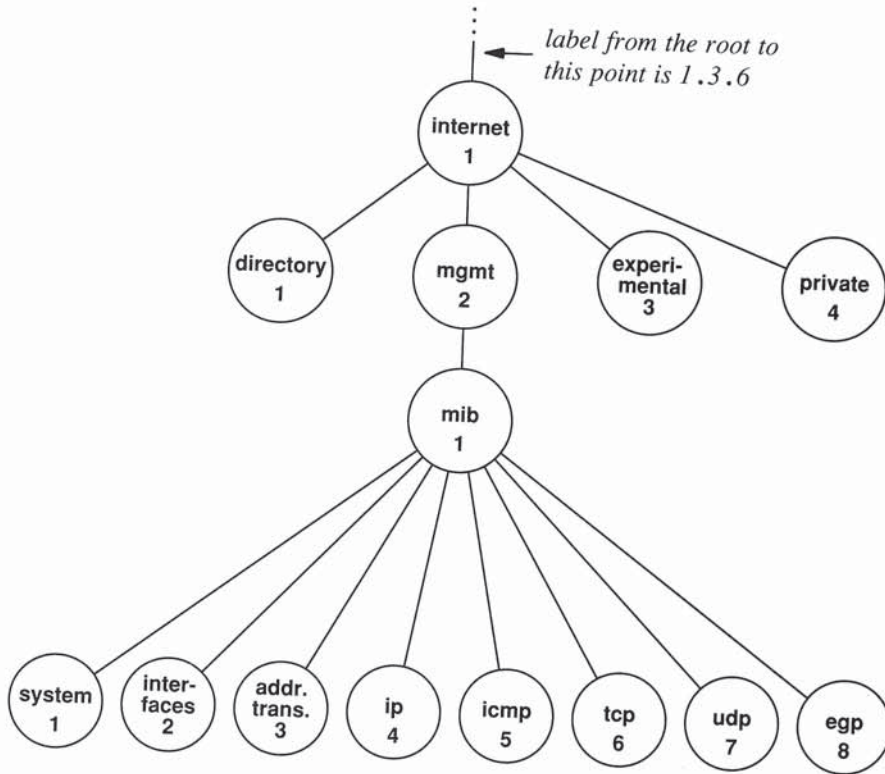


Figure 26.5 The object identifier namespace under the IAB *mib* node. Each subtree corresponds to one of the eight categories of MIB variables.

Two examples will make the naming syntax clear. Figure 26.5 shows that the category labeled *ip* has been assigned the numeric value 4. Thus, the names of all MIB variables corresponding to IP have an identifier that begins with the prefix *1.3.6.1.2.1.4*. If one wanted to write out the textual labels instead of the numeric representation, the name would be:

iso.org.dod.internet.mgmt.mib.ip

A MIB variable named *ipInReceives* has been assigned numeric identifier 3 under the *ip* node in the namespace, so its name is:

iso.org.dod.internet.mgmt.mib.ip.ipInReceives

and the corresponding numeric representation is:

1.3.6.1.2.1.4.3

When network management protocols use names of MIB variables in messages, each name has a suffix appended. For simple variables, the suffix 0 refers to the instance of the variable with that name. So, when it appears in a message sent to a router, the numeric representation of *ipInReceives* is:

1.3.6.1.2.1.4.3.0

which refers to the instance of *ipInReceives* on that router. Note that there is no way to guess the numeric value or suffix assigned to a variable. One must consult the published standards to find which numeric values have been assigned to each object type. Thus, programs that provide mappings between the textual form and underlying numeric values do so entirely by consulting tables of equivalences – there is no closed-form computation that performs the transformation.

As a second, more complex example, consider the MIB variable *ipAddrTable*, which contains a list of the IP addresses for each network interface. The variable exists in the namespace as a subtree under *ip*, and has been assigned the numeric value 20. Therefore, a reference to it has the prefix:

iso.org.dod.internet.mgmt.mib.ip.ipAddrTable

with a numeric equivalent:

1.3.6.1.2.1.4.20

In programming language terms, we think of the IP address table as a one-dimensional array, where each element of the array consists of a structure (record) that contains five items: an IP address, the integer index of an interface corresponding to the entry, an IP subnet mask, an IP broadcast address, and an integer that specifies the maximum

datagram size that the router will reassemble. Of course, not all routers have such an array in memory. The router may keep this information in many variables or may need to follow pointers to find it. However, the MIB provides a name for the array as if it existed, and allows network management software on individual routers to map table references into appropriate internal variables.

Using ASN.1 style notation, we can define *ipAddrTable*:

```
ipAddrTable ::= SEQUENCE OF IpAddrEntry
```

where *SEQUENCE* and *OF* are keywords that define an *ipAddrTable* to be a one-dimensional array of *IpAddrEntry*s. Each entry in the array is defined to consist of five fields (the definition assumes that *IpAddress* has already been defined).

```
IpAddrEntry ::= SEQUENCE {
    ipAdEntAddr
        IpAddress,
    ipAdEntIfIndex
        INTEGER,
    ipAdEntNetMask
        IpAddress,
    ipAdEntBcastAddr
        IpAddress,
    ipAdEntReasmMaxSize
        INTEGER (0..65535)
}
```

Further definitions must be given to assign numeric values to *ipAddrEntry* and to each item in the *IpAddrEntry* sequence. For example, the definition:

```
ipAddrEntry { ipAddrTable 1 }
```

specifies that *ipAddrEntry* falls under *ipAddrTable* and has numeric value 1. Similarly, the definition:

```
ipAdEntNetMask { ipAddrEntry 3 }
```

assigns *ipAdEntNetMask* numeric value 3 under *ipAddrEntry*.

We said that *ipAddrTable* was like a one-dimensional array. However, there is a significant difference in the way programmers use arrays and the way network management software uses tables in the MIB. Programmers think of an array as a set of elements that have an index used to select a specific element. For example, the programmer might write *xyz[3]* to select the third element from array *xyz*. ASN.1 syntax does not use integer indices. Instead, MIB tables append a suffix onto the name to select a specific element in the table. For our example of an IP address table, the standard specifies that the suffix used to select an item consists of an IP address. Syntactically,

the IP address (in dotted decimal notation) is concatenated onto the end of the object name to form the reference. Thus, to specify the network mask field in the IP address table entry corresponding to address 128.10.2.3, one uses the name:

iso.org.dod.internet.mgmt.mib.ip.ipAddrTable.ipAddrEntry.ipAdEntNetMask.128.10.2.3

which, in numeric form, becomes:

1.3.6.1.2.1.4.20.1.3.128.10.2.3

Although concatenating an index to the end of a name may seem awkward, it provides a powerful tool that allows clients to search tables without knowing the number of items or the type of data used as an index. The next section shows how network management protocols use this feature to step through a table one element at a time.

26.9 Simple Network Management Protocol

Network management protocols specify communication between the network management client program a manager invokes and a network management server program executing on a host or router. In addition to defining the form and meaning of messages exchanged and the representation of names and values in those messages, network management protocols also define administrative relationships among routers being managed. That is, they provide for authentication of managers.

One might expect network management protocols to contain a large number of commands. Some early protocols, for example, supported commands that allowed the manager to: *reboot* the system, *add* or *delete* routes, *disable* or *enable* a particular network interface, or *remove cached address bindings*. The main disadvantage of building management protocols around commands arises from the resulting complexity. The protocol requires a separate command for each operation on a data item. For example, the command to delete a routing table entry differs from the command to disable an interface. As a result, the protocol must change to accommodate new data items.

SNMP takes an interesting alternative approach to network management. Instead of defining a large set of commands, SNMP casts all operations in a *fetch-store paradigm*[†]. Conceptually, SNMP contains only two commands that allow a manager to fetch a value from a data item or store a value into a data item. All other operations are defined as side-effects of these two operations. For example, although SNMP does not have an explicit *reboot* operation, an equivalent operation can be defined by declaring a data item that gives the time until the next reboot and allowing the manager to assign the item a value (including zero).

The chief advantages of using a fetch-store paradigm are stability, simplicity, and flexibility. SNMP is especially stable because its definition remains fixed, even though new data items are added to the MIB and new operations are defined as side-effects of storing into those items. SNMP is simple to implement, understand, and debug because

[†]The fetch-store paradigm comes from a management protocol system known as HEMS. See Partridge and Trewitt [RFCs 1021, 1022, 1023, and 1024] for details.

it avoids the complexity of having special cases for each command. Finally, SNMP is especially flexible because it can accommodate arbitrary commands in an elegant framework.

From a manager's point of view, of course, SNMP remains hidden. The user interface to network management software can phrase operations as imperative commands (e.g., *reboot*). Thus, there is little visible difference between the way a manager uses SNMP and other network management protocols. In fact, vendors have begun to sell network management software that offers a graphical user interface. Such software displays diagrams of network connectivity, and uses a point-and-click style of interaction.

As Figure 26.6 shows, SNMP offers more than the two operations we have described.

Command	Meaning
get-request	Fetch a value from a specific variable
get-next-request	Fetch a value without knowing its exact name
get-response	Reply to a fetch operation
set-request	Store a value in a specific variable
trap	Reply triggered by an event

Figure 26.6 The set of possible SNMP operations†. *Get-next-request* allows the manager to iterate through a table of items.

Operations *get-request*, *get-response*, and *set-request* provide the basic fetch and store operations (as well as replies to those operations). SNMP specifies that operations must be *atomic*, meaning that if a single SNMP message specifies operations on multiple variables, the server either performs all operations or none of them. In particular, no assignments will be made if any of them are in error. The *trap* operation allows managers to program servers to send information when an event occurs. For example, an SNMP server can be programmed to send a manager a *trap* message whenever one of the attached networks becomes unusable (i.e., an interface goes down).

26.9.1 Searching Tables Using Names

We said that ASN.1 does not provide mechanisms for declaring arrays or indexing them in the usual sense. However, it is possible to denote individual elements of a table by appending a suffix to the object identifier for the table. Unfortunately, a client program may wish to examine entries in a table for which it does not know all valid suffixes. The *get-next-request* operation allows a client to iterate through a table without knowing how many items the table contains. The rules are quite simple. When sending a *get-next-request*, the client supplies a prefix of a valid object identifier, *P*. The server examines the set of object identifiers for all variables it controls, and responds by sending a *get-response* command for the one that has an object identifier lexicographically

†SNMPv2 adds a *get-bulk* operation that permits a manager to fetch multiple values with a single request.

greater than P . Because the MIB uses suffixes to index tables, a client can send the prefix of an object identifier corresponding to a table and receive the first element in the table. The client can send the name of the first element in a table and receive the second, and so on.

Consider an example search. Recall that the *ipAddrTable* uses IP addresses to identify entries in the table. A client that does not know which IP addresses are in the table on a given router cannot form a complete object identifier. However, the client can still use the *get-next-request* operation to search the table by sending the prefix:

```
iso.org.dod.internet.mgmt.mib.ip.ipAddrTable.ipAddrEntry.ipAdEntNetMask
```

which, in numeric form, is:

```
1.3.6.1.2.1.4.20.1.3
```

The server returns the network mask field of the first entry in *ipAddrTable*. The client uses the full object identifier returned by the server to request the next item in the table.

26.10 SNMP Message Format

Unlike most TCP/IP protocols, SNMP messages do not have fixed fields. Instead, they use the standard ASN.1 encoding. Thus, they can be difficult for humans to decode and understand. After examining the SNMP message definition in ASN.1 notation, we will review the ASN.1 encoding scheme briefly, and see an example of an encoded SNMP message.

An SNMP message consists of three main parts: a protocol *version*, an SNMP *community* identifier (used to group together the routers managed by a given manager), and a *data* area. The data area is divided into *protocol data units (PDUs)*. Each PDU consists of a request (sent by client) or a response (sent by server). Figure 26.7 shows how the message can be described in ASN.1 notation.

```
SNMP-Message ::=
  SEQUENCE {
    version INTEGER {
      version-1 (0)
    },
    community
      OCTET STRING,
    data
      ANY
  }
```

Figure 26.7 The SNMP message format in ASN.1 notation. The *data* area contains one or more protocol data units.

The five types of protocol data units are further described in ASN.1 notation in Figure 26.8.

```

SNMP-PDUs ::=
  CHOICE {
    get-request
      GetRequest-PDU,
    get-next-request-PDU
      GetNextRequest-PDU,
    get-response
      GetResponse-PDU,
    set-request
      SetRequest-PDU,
    trap
      Trap-PDU,
  }

```

Figure 26.8 The ASN.1 definitions of an SNMP PDU. The syntax for each request type must be specified further.

The definition specifies that each protocol data unit consists of one of the five request or response types. To complete the definition of an SNMP message, we must further specify the syntax of the five individual types. For example, Figure 26.9 shows the definition of a *get-request*.

```

GetRequest-PDU ::= [0]
  IMPLICIT SEQUENCE {
    request-id
      RequestID,
    error-status
      ErrorStatus,
    error-index
      ErrorIndex,
    variable-bindings
      VarBindList
  }

```

Figure 26.9 The ASN.1 definition of a *get-request* message. Formally, the message is defined to be a *GetRequest-PDU*.

Further definitions in the standard specify the remaining undefined terms. *Request-ID* is defined to be a 4-octet integer (used to match responses to queries). Both *Error-Status* and *ErrorIndex* are single octet integers which contain the value zero in a re-

quest. Finally, *VarBindList* contains a list of object identifiers for which the client seeks values. In ASN.1 terms, the definitions specify that *VarBindList* is a sequence of pairs of object name and value. ASN.1 represents the pairs as a sequence of two items. Thus, in the simplest possible request, *VarBindList* is a sequence of two items: a name and a *null*.

26.11 Example Encoded SNMP Message

The encoded form of ASN.1 uses variable-length fields to represent items. In general, each field begins with a header that specifies the type of object and its length in bytes. For example, Figure 26.10 shows the string of encoded octets in a *get-request* message for data item *sysDescr* (numeric object identifier *1.3.6.1.2.1.1.1*).

```

    30      29      02      01      00
SEQUENCE len=41 INTEGER len=1  vers=0

    04      06      70      75      62      6C      69      63
string  len=6   p      u      b      l      i      c

    A0      1C      02      04      05      AE      56      02
getreq. len=28 INTEGER len=4  ----- request ID -----

    02      01      00      02      01      00
INTEGER len=1  status INTEGER len=1 error index

    30      0E      30      0C      06      08
SEQUENCE len=14 SEQUENCE len=12 objectid len=8

    2B      06      01      02      01      01      01      00
1.3 . 6 . 1 . 2 . 1 . 1 . 0

    05      00
null  len=0

```

Figure 26.10 The encoded form of a *get-request* for data item *sysDescr* with octets shown in hexadecimal and their meanings below. Related octets have been grouped onto lines; they are contiguous in the message.

As Figure 26.10 shows, the message starts with a code for *SEQUENCE* which has a length of 41 octets. The first item in the sequence is a 1-octet integer that specifies the protocol *version*. The *community* field is stored in a character string, which in the example, is a 6-octet string that contains the word *public*.

The *GetRequest-PDU* occupies the remainder of the message. The initial code specifies a *get-Request* operation. Because the high-order bit is turned on, the interpretation is *context specific*. That is, the hexadecimal value *A0* only specifies a *GetRequest-PDU* when used in an SNMP message; it is not a universally reserved value. Following the request octet, the length octet specifies the request is 28 octets long. The request ID is 4 octets, but each of the error status and error index are one octet. Finally, the sequence of pairs contains one binding, a single object identifier bound to a *null* value. The identifier is encoded as expected except that the first two numeric labels are combined into a single octet.

26.12 Summary

Network management protocols allow a manager to monitor and control routers and hosts. A network management client program executing on the manager's workstation contacts one or more servers, called agents, running on the computers to be controlled. Because an internet consists of heterogeneous machines and networks, TCP/IP management software executes as application programs and uses internet transport protocols (e.g., UDP) for communication between clients and servers.

The standard TCP/IP network management protocol is SNMP, the Simple Network Management Protocol. SNMP defines a low-level management protocol that provides two basic operations: fetch a value from a variable or store a value into a variable. In SNMP, all operations occur as side-effects of storing values into variables. SNMP defines the format of messages that travel between a manager's computer and a managed entity.

A companion standard to SNMP defines the set of variables that a managed entity maintains. The standard is known as a Management Information Base, or MIB. MIB variables are described using ASN.1, a formal language that provides a concise encoded form as well as a precise human-readable notation for names and objects. ASN.1 uses a hierarchical namespace to guarantee that all MIB names are globally unique while still allowing subgroups to assign parts of the namespace.

FOR FURTHER STUDY

Schoffstall, Fedor, Davin, and Case [RFC 1157] contains the standard for SNMP. ISO [May 87a] and [May 87b] contain the standard for ASN.1 and specify the encoding. McCloghrie and Rose [RFC 1213] defines the variables that comprise MIB-II, while McCloghrie and Rose [RFC 1211] contains the SMI rules for naming MIB variables.

A series of RFCs defines SNMPv2, which is a proposed standard at the time of this writing. Case, McCloghrie, Rose, Waldbusser [RFC 1441] contains an introduction to SNMPv2. Case, McCloghrie, Rose, and Waldbusser [RFC 1450] defines the

SNMPv2 MIB. Galvin and McCloghrie [RFC 1446] discusses SNMPv2 security protocols. Case, McCloghrie, Rose, Waldbusser [RFC 1448] specifies protocol operations.

An older proposal for a network management protocol called HEMS can be found in Trewitt and Partridge [RFCs 1021, 1022, 1023, and 1024]. Davin, Case, Fedor, and Schoffstall [RFC 1028] specifies a predecessor to SNMP known as the Simple Gateway Monitoring Protocol (SGMP).

EXERCISES

- 26.1 Capture an SNMP packet with a network analyzer and decode the fields.
- 26.2 Read the standard to find out how ASN.1 encodes the first two numeric values from an object identifier in a single octet. Why does it do so?
- 26.3 Read the specification for CMIP. How many commands does it support?
- 26.4 Suppose the MIB designers needed to define a variable that corresponded to a two-dimensional array. How can ASN.1 notation accommodate references to such a variable?
- 26.5 What are the advantages and disadvantages of defining globally unique ASN.1 names for MIB variables?
- 26.6 If you have SNMP client code available, try using it to read MIB variables in a local router. What is the advantage of allowing arbitrary managers to read variables in all routers?
- 26.7 Read the MIB specification to find the definition of variable *ipRoutingTable* that corresponds to an IP routing table. Design a program that will use SNMP to contact multiple routers, and see if any entries in their routing tables cause a routing loop. Exactly what ASN.1 names should such a program generate?

Bibliography

- ABRAMSON, N. [1970], The ALOHA System – Another Alternative for Computer Communications, *Proceedings of the Fall Joint Computer Conference*.
- ABRAMSON, N. and F. KUO (EDS.) [1973], *Computer Communication Networks*, Prentice Hall, Englewood Cliffs, New Jersey.
- ANDREWS, D. W., and G. D. SHULTZ [1982], A Token-Ring Architecture for Local Area Networks: An Update, *Proceedings of Fall 82 COMPCON*, IEEE.
- BALL, J. E., E. J. BURKE, I. GERTNER, K. A. LANTZ, and R. F. RASHID [1979], Perspectives on Message-Based Distributed Computing, *IEEE Computing Networking Symposium*, 46-51.
- BBN [1981], A History of the ARPANET: The First Decade, *Technical Report* Bolt, Beranek, and Newman, Inc.
- BBN [December 1981], Specification for the Interconnection of a Host and an IMP (revised), *Technical Report 1822*, Bolt, Beranek, and Newman, Inc.
- BIAGIONI E., E. COOPER, and R. SANSOM [March 1993], Designing a Practical ATM LAN, *IEEE Network*, 32-39.
- BERTSEKAS D. and R. GALLAGER [1987], *Data Networks*, Prentice-Hall, Englewood Cliffs, New Jersey.
- BIRRELL, A., and B. NELSON [February 1984], Implementing Remote Procedure Calls, *ACM Transactions on Computer Systems*, 2(1), 39-59.
- BOGGS, D., J. SHOCH, E. TAFT, and R. METCALFE [April 1980], Pup: An Internetwork Architecture, *IEEE Transactions on Communications*.
- BORMAN, D., [April 1989], Implementing TCP/IP on a Cray Computer, *Computer Communication Review*, 19(2), 11-15.
- BROWN, M., K. KOLLING, and E. TAFT [November 1985], The Alpine File System, *ACM Transactions on Computer Systems*, 3(4), 261-293.
- BROWNBRIDGE, D., L. MARSHALL, and B. RANDELL [December 1982], The Newcastle Connections or UNIXes of the World Unite!, *Software – Practice and Experience*, 12(12), 1147-1162.
- CASNER, S., and S. DEERING [July 1992], First IETF Internet Audiocast, *Computer Communications Review*, 22(3), 92-97.

- CERF, V., and E. CAIN [October 1983], The DOD Internet Architecture Model, *Computer Networks*.
- CERF, V., and R. KAHN [May 1974], A Protocol for Packet Network Interconnection, *IEEE Transactions of Communications*, Com-22(5).
- CERF, V. [October 1989], A History of the ARPANET, *ConneXions, The Interoperability Report*, 480 San Antonio Rd, Suite 100, Mountain View, California.
- CHERITON, D. R. [1983], Local Networking and Internetworking in the V-System, *Proceedings of the Eighth Data Communications Symposium*.
- CHERITON, D. R. [April 1984], The V Kernel: A Software Base for Distributed Systems, *IEEE Software*, 1(2), 19-42.
- CHERITON, D. [August 1986], VMTP: A Transport Protocol for the Next Generation of Communication Systems, *Proceedings of ACM SIGCOMM '86*, 406-415.
- CHERITON, D., and T. MANN [May 1984], Uniform Access to Distributed Name Interpretation in the V-System, *Proceedings IEEE Fourth International Conference on Distributed Computing Systems*, 290-297.
- CHESSON, G. [June 1987], Protocol Engine Design, *Proceedings of the 1987 Summer USENIX Conference*, Phoenix, AZ.
- CHESWICK, W., and S. BELLOVIN [1994], *Firewalls And Internet Security: Repelling the Wiley Hacker*, Addison-Wesley, Reading, Massachusetts.
- CLARK, D. [December 1985], The structure of Systems Using Upcalls, *Proceedings of the Tenth ACM Symposium on Operating Systems Principles*, 171-180.
- CLARK, D., M. LAMBERT, and L. ZHANG [August 1987], NETBLT: A High Throughput Transport Protocol, *Proceedings of ACM SIGCOMM '87*.
- CLARK, D., V. JACOBSON, J. ROMKEY, and H. SALWEN [June 1989], An Analysis of TCP Processing Overhead, *IEEE Communications*, 23-29.
- COHEN, D., [1981], On Holy Wars and a Plea for Peace, *IEEE Computer*, 48-54.
- COMER, D. E. and J. T. KORB [1983], CSNET Protocol Software: The IP-to-X25 Interface, *Computer Communications Review*, 13(2).
- COMER, D. E. [1984], *Operating System Design - The XINU Approach*, Prentice-Hall, Englewood Cliffs, New Jersey.
- COMER, D. E. [1987], *Operating System Design Vol II. - Internetworking With XINU*, Prentice-Hall, Englewood Cliffs, New Jersey.
- COMER, D. E. and D. L. STEVENS [1994] *Internetworking With TCP/IP Volume II - Design, Implementation, and Internals*, 2nd edition, Prentice-Hall, Englewood Cliffs, New Jersey.
- COMER, D. E. and D. L. STEVENS [1993] *Internetworking With TCP/IP Volume III - Client-Server Programming And Applications, BSD socket version*, Prentice-Hall, Englewood Cliffs, New Jersey.
- COMER, D. E. and D. L. STEVENS [1994] *Internetworking With TCP/IP Volume III - Client-Server Programming And Applications, AT&T TLI version*, Prentice-Hall, Englewood Cliffs, New Jersey.

- COMER, D. E., T. NARTEN, and R. YAVATKAR [April 1987], The Cypress Network: A Low-Cost Internet Connection Technology, *Technical Report TR-653*, Purdue University, West Lafayette, IN.
- COMER, D. E., T. NARTEN, and R. YAVATKAR [1987], The Cypress Coaxial Packet Switch, *Computer Networks and ISDN Systems*, vol. 14:2-5, 383-388.
- COTTON, I. [1979], Technologies for Local Area Computer Networks, *Proceedings of the Local Area Communications Network Symposium*.
- CROWLEY, T., H. FORSDICK, M. LANDAU, and V. TRAVERS [June 1987], The Diamond Multimedia Editor, *Proceedings of the 1987 Summer USENIX Conference*, Phoenix, AZ.
- DALAL Y. K., and R. S. PRINTIS [1981], 48-Bit Absolute Internet and Ethernet Host Numbers, *Proceedings of the Seventh Data Communications Symposium*.
- DEERING S. E., and D. R. CHERITON [May 1990], Multicast Routing in Datagram Internetworks and Extended LANs, *ACM Transactions on Computer Systems*, 8(2), 85-110.
- DEERING, S., D. ESTRIN, D. FARINACCI, V. JACOBSON, C-G LIU, and L. WEI [August 1994], An Architecture for Wide-Area Multicasting Routing, *Proceedings of ACM SIGCOMM '94*, 126-135.
- DENNING P. J., [September-October 1989], *The Science of Computing: Worldnet*, in American Scientist, 432-434.
- DENNING P. J., [November-December 1989], *The Science of Computing: The ARPANET After Twenty Years*, in American Scientist, 530-534.
- DE PRYCKER, M. [1993] *Asynchronous Transfer Mode Solution for Broadband ISDN*, 2nd edition, Ellis Horwood, UK.
- DIGITAL EQUIPMENT CORPORATION., INTEL CORPORATION, and XEROX CORPORATION [September 1980], *The Ethernet: A Local Area Network Data Link Layer and Physical Layer Specification*.
- DION, J. [Oct. 1980], The Cambridge File Server, *Operating Systems Review*, 14(4), 26-35.
- DRIVER, H., H. HOPEWELL, and J. IAQUINTO [September 1979], How the Gateway Regulates Information Control, *Data Communications*.
- EDGE, S. W. [1979], Comparison of the Hop-by-Hop and Endpoint Approaches to Network Interconnection, in *Flow Control in Computer Networks*, J-L. GRANGE and M. GIEN (EDS.), North-Holland, Amsterdam, 359-373.
- EDGE, S. [1983], An Adaptive Timeout Algorithm for Retransmission Across a Packet Switching Network, *Proceedings of ACM SIGCOMM '83*.
- ENSLOW, P. [January 1978], What is a 'Distributed' Data Processing System? *Computer*, 13-21.
- ERIKSSON, H. [August 1994] MBONE: The Multicast Backbone, *Communications of the ACM*, 37(8), 54-60.
- FALK, G. [1983], The Structure and Function of Network Protocols, in *Computer Communications, Volume 1: Principles*, CHOU, W. (ED.), Prentice-Hall, Englewood Cliffs, New Jersey.
- FARMER, W. D., and E. E. NEWHALL [1969], An Experimental Distributed Switching System to Handle Bursty Computer Traffic, *Proceedings of the ACM Symposium on Probabilistic Optimization of Data Communication Systems*, 1-33.

- FCCSET [November 1987], A Research and Development Strategy for High Performance Computing, *Report from the Executive Office of the President and Office of Science and Technology Policy*.
- FEDOR, M. [June 1988], GATED: A Multi-Routing Protocol Daemon for UNIX, *Proceedings of the 1988 Summer USENIX conference*, San Francisco, California.
- FEINLER, J., O. J. JACOBSEN, and M. STAHL [December 1985], *DDN Protocol Handbook Volume Two, DARPA Internet Protocols*, DDN Network Information Center, SRI International, 333 Ravenswood Avenue, Room EJ291, Menlo Park, California.
- FLOYD, S. and V. JACOBSON [August 1993], Random Early Detection Gateways for Congestion Avoidance, *IEEE/ACM Transactions on Networking*, 1(4).
- FRANK, H., and W. CHOU [1971], Routing in Computer Networks, *Networks*, 1(1), 99-112.
- FRANK, H., and J. FRISCH [1971], *Communication, Transmission, and Transportation Networks*, Addison-Wesley, Reading, Massachusetts.
- FRANTA, W. R., and I. CHLAMTAC [1981], *Local Networks*, Lexington Books, Lexington, Massachusetts.
- FRICC [May 1989], *Program Plan for the National Research and Education Network*, Federal Research Internet Coordinating Committee, US Department of Energy, Office of Scientific Computing report ER-7.
- FRIDRICH, M., and W. OLDER [December 1981], The Felix File Server, *Proceedings of the Eighth Symposium on Operating Systems Principles*, 37-46.
- FULTZ, G. L., and L. KLEINROCK, [June 14-16, 1971], Adaptive Routing Techniques for Store-and-Forward Computer Communication Networks, presented at *IEEE International Conference on Communications*, Montreal, Canada.
- GERLA, M., and L. KLEINROCK [April 1980], Flow Control: A Comparative Survey, *IEEE Transactions on Communications*.
- GOSIP [April 1989], U.S. Government Open Systems Interconnection Profile (GOSIP) version 2.0, GOSIP Advanced Requirements Group, National Institute of Standards and Technology (NIST).
- GRANGE, J-L., and M. GIEN (EDS.) [1979], *Flow Control in Computer Networks*, North-Holland, Amsterdam.
- GREEN, P. E. (ED.) [1982], *Computer Network Architectures and Protocols*, Plenum Press, New York.
- HINDEN, R., J. HAVERTY, and A. SHELTZER [September 1983], The DARPA Internet: Interconnecting Heterogeneous Computer Networks with Gateways, *Computer*.
- INTERNATIONAL ORGANIZATION FOR STANDARDIZATION [June 1986a], Information processing systems — Open Systems Interconnection — *Transport Service Definition*, International Standard number 8072, ISO, Switzerland.
- INTERNATIONAL ORGANIZATION FOR STANDARDIZATION [July 1986b], Information processing systems — Open Systems Interconnection — *Connection Oriented Transport Protocol Specification*, International Standard number 8073, ISO, Switzerland.

- INTERNATIONAL ORGANIZATION FOR STANDARDIZATION [May 1987a], Information processing systems — Open Systems Interconnection — *Specification of Basic Specification of Abstract Syntax Notation One (ASN.1)*, International Standard number 8824, ISO, Switzerland.
- INTERNATIONAL ORGANIZATION FOR STANDARDIZATION [May 1987b], Information processing systems — Open Systems Interconnection — *Specification of Basic Encoding Rules for Abstract Syntax Notation One (ASN.1)*, International Standard number 8825, ISO, Switzerland.
- INTERNATIONAL ORGANIZATION FOR STANDARDIZATION [May 1988a], Information processing systems — Open Systems Interconnection — *Management Information Service Definition, Part 2: Common Management Information Service*, Draft International Standard number 9595-2, ISO, Switzerland.
- INTERNATIONAL ORGANIZATION FOR STANDARDIZATION [May 1988a], Information processing systems — Open Systems Interconnection — *Management Information Protocol Definition, Part 2: Common Management Information Protocol*, Draft International Standard number 9596-2.
- JACOBSEN, O. J. (PUBLISHER) [1987-], *ConneXions, The Interoperability Report*, Interop Company, a division of Softbank Exposition and Conference Company, Foster City, California.
- JACOBSON, V. [August 1988], Congestion Avoidance and Control, *Proceedings ACM SIGCOMM '88*.
- JAIN, R. [January 1985], On Caching Out-of-Order Packets in Window Flow Controlled Networks, *Technical Report*, DEC-TR-342, Digital Equipment Corporation.
- JAIN, R. [March 1986], Divergence of Timeout Algorithms for Packet Retransmissions, *Proceedings Fifth Annual International Phoenix Conference on Computers and Communications*, Scottsdale, AZ.
- JAIN, R. [October 1986], A Timeout-Based Congestion Control Scheme for Window Flow-Controlled Networks, *IEEE Journal on Selected Areas in Communications*, Vol. SAC-4, no. 7.
- JAIN, R. [May 1992], Myths About Congestion Management in High-speed Networks, *Internetworking: Research and Experience*, 3(3), 101-113.
- JENNINGS, D. M., L. H. LANDWEBER, and I. H. FUCHS [February 28, 1986], Computer Networking for Scientists and Engineers, *Science* vol 231, 941-950.
- JUBIN, J. and J. TORNOW [January 1987], The DARPA Packet Radio Network Protocols, *IEEE Proceedings*.
- KAHN, R. [November 1972], Resource-Sharing Computer Communications Networks, *Proceedings of the IEEE*, 60(11), 1397-1407.
- KARN, P., H. PRICE, and R. DIERSING [May 1985], Packet Radio in the Amateur Service, *IEEE Journal on Selected Areas in Communications*,
- KARN, P., and C. PARTRIDGE [August 1987], Improving Round-Trip Time Estimates in Reliable Transport Protocols, *Proceedings of ACM SIGCOMM '87*.
- KENT, C., and J. MOGUL [August 1987], Fragmentation Considered Harmful, *Proceedings of ACM SIGCOMM '87*.

- KLINE, C. [August 1987], Supercomputers on the Internet: A Case Study, *Proceedings of ACM SIGCOMM '87*.
- KOCHAN, S. G., and P. H. WOODS [1989], *UNIX Networking*, Hayden Books, Indianapolis, IN.
- LABARRE, L. (ED.) [December 1989], OSI Internet Management: Management Information Base, *Internet Draft <IETF.DRAFTS>DRAFT-IETF-SNMP-MIB2-01.TXT*, DDN Network Information Center, SRI International, Ravenswood, CA.
- LAMPSON, B. W., M. PAUL, and H. J. SIEGERT (EDS.) [1981], *Distributed Systems - Architecture and Implementation (An Advanced Course)*, Springer-Verlag, Berlin.
- LANZILLO, A. L., and C. PARTRIDGE [January 1989], Implementation of Dial-up IP for UNIX Systems, *Proceedings 1989 Winter USENIX Technical Conference*, San Diego, CA.
- LAQUEY, T. L., [July 1989], *User's Directory of Computer Networks*, Digital Press, Bedford, MA.
- LAZAR, A. [November 1983], Optimal Flow Control of a Class of Queuing Networks in Equilibrium. *IEEE Transactions on Automatic Control*, Vol. AC-28:11.
- LEFFLER, S., M. MCKUSICK, M. KARELS, and J. QUARTERMAN [1989], *The Design and Implementation of the 4.3BSD UNIX Operating System*, Addison-Wesley, Reading, Massachusetts.
- LYNCH, D. C., (FOUNDER) [1987-], The NETWORLD+INTEROP Conference, *Interop Company*, a division of Softbank Exposition and Conference Company, Foster City, California.
- MCNAMARA, J. [1982], *Technical Aspects of Data Communications*, Digital Press, Digital Equipment Corporation, Bedford, Massachusetts.
- MCQUILLAN, J. M., I. RICHER, and E. ROSEN [May 1980], The New Routing Algorithm for the ARPANET, *IEEE Transactions on Communications*, (COM-28), 711-719.
- MERIT [November 1987], Management and Operation of the NSFNET Backbone Network: A Proposal Funded by the National Science Foundation and the State of Michigan, *MERIT Incorporated*, Ann Arbor, Michigan.
- METCALFE, R. M., and D. R. BOGGS [July 1976], Ethernet: Distributed Packet Switching for Local Computer Networks, *Communications of the ACM*, 19(7), 395-404.
- MILLER, C. K., and D. M. THOMPSON [March 1982], Making a Case for Token Passing in Local Networks, *Data Communications*.
- MILLS, D., and H-W. BRAUN [August 1987], The NSFNET Backbone Network, *Proceedings of ACM SIGCOMM '87*.
- MITCHELL, J., and J. DION [April 1982], A Comparison of Two Network-Based File Servers, *Communications of the ACM*, 25(4), 233-245.
- MORRIS, R. [1979], Fixing Timeout Intervals for Lost Packet Detection in Computer Communication Networks, *Proceedings AFIPS National Computer Conference*, AFIPS Press, Montvale, New Jersey.
- NAGLE, J. [April 1987], On Packet Switches With Infinite Storage, *IEEE Transactions on Communications*, Vol. COM-35:4.

- NARTEN, T. [Sept. 1989], Internet Routing, *Proceedings ACM SIGCOMM '89*.
- NEEDHAM, R. M. [1979], System Aspects of the Cambridge Ring, *Proceedings of the ACM Seventh Symposium on Operating System Principles*, 82-85.
- NELSON, J. [September 1983], 802: A Progress Report, *Datamation*.
- OPPEN, D., and Y. DALAL [October 1981], The Clearinghouse: A Decentralized Agent for Locating Named Objects, Office Products Division, XEROX Corporation.
- PARTRIDGE, C. [June 1986], Mail Routing Using Domain Names: An Informal Tour, *Proceedings of the 1986 Summer USENIX Conference*, Atlanta, GA.
- PARTRIDGE, C. [June 1987], Implementing the Reliable Data Protocol (RDP), *Proceedings of the 1987 Summer USENIX Conference*, Phoenix, Arizona.
- PARTRIDGE, C. [1994], *Gigabit Networking*, Addison-Wesley, Reading, Massachusetts.
- PETERSON, L. [1985], *Defining and Naming the Fundamental Objects in a Distributed Message System*, Ph.D. Dissertation, Purdue University, West Lafayette, Indiana.
- PIERCE, J. R. [1972], Networks for Block Switching of Data, *Bell System Technical Journal*, 51.
- POSTEL, J. B. [April 1980], Internetwork Protocol Approaches, *IEEE Transactions on Communications*, COM-28, 604-611.
- POSTEL, J. B., C. A. SUNSHINE, and D. CHEN [1981], The ARPA Internet Protocol, *Computer Networks*.
- QUARTERMAN, J. S. [1990], *The Matrix: Computer Networks and Conferencing Systems Worldwide*, Digital Press, Digital Equipment Corporation, Maynard, MA.
- QUARTERMAN, J. S., and J. C. HOSKINS [October 1986], Notable Computer Networks, *Communications of the ACM*, 29(10).
- RAMAKRISHNAN, K. and R. JAIN [May 1990], A Binary Feedback Scheme For Congestion Avoidance In Computer Networks, *ACM Transactions on Computer Systems*, 8(2), 158-181.
- REYNOLDS, J., J. POSTEL, A. R. KATZ, G. G. FINN, and A. L. DESCHON [October 1985], The DARPA Experimental Multimedia Mail System, *IEEE Computer*.
- RITCHIE, D. M., and K. THOMPSON [July 1974], The UNIX Time-Sharing System, *Communications of the ACM*, 17(7), 365-375; revised and reprinted in *Bell System Technical Journal*, 57(6), [July-August 1978], 1905-1929.
- ROSE, M. (ED.) [October 1989], Management Information Base for Network Management of TCP/IP-based Internets, *Internet Draft <IETF.DRAFTS>DRAFT-IETF-OIM-MIB2-00.TXT*, DDN Network Information Center, SRI International, Ravenswood, CA.
- ROSENTHAL, R. (ED.) [November 1982], *The Selection of Local Area Computer Networks*, National Bureau of Standards Special Publication 500-96.
- SALTZER, J. [1978], Naming and Binding of Objects, *Operating Systems, An Advanced Course*, Springer-Verlag, 99-208.
- SALTZER, J. [April 1982], Naming and Binding of Network Destinations, *International Symposium on Local Computer Networks*, IFIP/T.C.6, 311-317.
- SALTZER, J., D. REED, and D. CLARK [November 1984], End-to-End Arguments in System Design, *ACM Transactions on Computer Systems*, 2(4), 277-288.

- SCHWARTZ, M., and T. STERN [April 1980], *IEEE Transactions on Communications*, COM-28(4), 539-552.
- SHOCH, J. F. [1978], Internetwork Naming, Addressing, and Routing, *Proceedings of COMPCON*.
- SHOCH, J. F., Y. DALAL, and D. REDELL [August 1982], Evolution of the Ethernet Local Computer Network, *Computer*.
- SNA [1975], *IBM System Network Architecture - General Information*, IBM System Development Division, Publications Center, Department E01, P.O. Box 12195, Research Triangle Park, North Carolina, 27709.
- SOLOMON, M., L. LANDWEBER, and D. NEUHEGEN [1982], The CSNET Name Server, *Computer Networks* (6), 161-172.
- STALLINGS, W. [1984], *Local Networks: An Introduction*, Macmillan Publishing Company, New York.
- STALLINGS, W. [1985], *Data and Computer Communications*, Macmillan Publishing Company, New York.
- SWINEHART, D., G. MCDANIEL, and D. R. BOGGS [December 1979], WFS: A Simple Shared File System for a Distributed Environment, *Proceedings of the Seventh Symposium on Operating System Principles*, 9-17.
- TANENBAUM, A. [1981], *Computer Networks: Toward Distributed Processing Systems*, Prentice-Hall, Englewood Cliffs, New Jersey.
- TICHY, W., and Z. RUAN [June 1984], Towards a Distributed File System, *Proceedings of Summer 84 USENIX Conference*, Salt Lake City, Utah, 87-97.
- TOMLINSON, R. S. [1975], Selecting Sequence Numbers, *Proceedings ACM SIGOPS/SIGCOMM Interprocess Communication Workshop*, 11-23, 1975.
- WARD, A. A. [1980], TRIX: A Network-Oriented Operating System, *Proceedings of COMPCON*, 344-349.
- WATSON, R. [1981], Timer-Based Mechanisms in Reliable Transport Protocol Connection Management, *Computer Networks*, North-Holland Publishing Company.
- WEINBERGER, P. J. [1985], The UNIX Eighth Edition Network File System, *Proceedings 1985 ACM Computer Science Conference*, 299-301.
- WELCH, B., and J. OSTERHAUT [May 1986], Prefix Tables: A Simple Mechanism for Locating Files in a Distributed System, *Proceedings IEEE Sixth International Conference on Distributed Computing Systems*, 1845-189.
- WILKES, M. V., and D. J. WHEELER [May 1979], The Cambridge Digital Communication Ring, *Proceedings Local Area Computer Network Symposium*.
- XEROX [1981], Internet Transport Protocols, *Report X SIS 028112*, Xerox Corporation, Office Products Division, Network Systems Administration Office, 3333 Coyote Hill Road, Palo Alto, California.
- ZHANG, L. [August 1986], Why TCP Timers Don't Work Well, *Proceedings of ACM SIGCOMM '86*.

Index

Constants and numeric items

10Base-T 25, 558
1822 38
220 441
221 443
250 441
576 96, 558
802.3 20, 558
822 438, 446, 513, 558
9180 312

A

AAL 558
AAL1 310
AAL5 310
abort 205
absolute name 453
Abstract Syntax Notation 1 165, 453
accept system call 346
access control 475, 476
ACK 193, 558
acknowledgement 193, 208, 558
 ambiguity 211
 cumulative 208
 delayed 225
active 271
active monitoring 486
active open 201, 558
adapter 21, 26
adaptive bridge 31
adaptive retransmission algorithm 209

address 5, 51, 59, 384
 ARPANET 39
 Ethernet 28
 IP 60
 X.121 45
 X.25 45
 broadcast 29
 class 60
 class D 291
 hardware 20, 28
 internet 60, 73, 83
 mail 435, 438
 multicast 29
 network 20
 physical 28
 resolution 73, 74
 supernet 153
 unicast 29
address boundary 118
address lease 373
address mask 136, 558
address resolution 559
address resolution problem 74
Address Resolution Protocol 75
address-to-name translation 383
Advanced Networks and Services 44,
 559
advertise routes 242
agent 449, 559
algorithm
 routing 116
 shortest path 246

- alias
 - mail 435
 - all hosts group 292
 - all hosts multicast 504
 - all nodes multicast 504
 - all routers multicast 504
 - alternative subtype (MIME) 444
 - ambiguity of acknowledgements 211
 - anonymous FTP 426, 559
 - ANS 44, 559
 - ANSI 559
 - ANSNET 44, 559
 - application program 179
 - area 280
 - ARP 75, 559
 - encapsulation 79
 - hack 142
 - implementation 77
 - protocol 73
 - ARPA 2, 37, 559
 - ARPA/NSF Internet 2
 - ARPANET 37, 559
 - ARPANET address 39
 - ARPANET port 38
 - ARQ 560
 - ASN.1 453, 467, 560
 - Assigned Numbers 560
 - Asynchronous Transfer Mode 36, 303
 - ATM 36, 303, 560
 - NNI 304
 - UNI 304
 - ATM Adaptation Layer 308, 560
 - ATMARP 316, 560
 - atomic assignment 459
 - attachment unit interface 21
 - AUI 21, 560
 - authentication 280, 475
 - authority zone 403, 560
 - authorization 475
 - automatic configuration 372
 - autonomous confederation 265
 - autonomous system 252, 560
 - autonomous system number 253
 - availability 475
- B**
- backbone network 39, 560
 - backoff 27
 - base header 494, 561
 - base64 443
 - baseband 561
 - bastion host 481, 561
 - baud 561
 - BBN 37
 - Bellman 240
 - Bellman-Ford 240
 - Berkeley broadcast 561
 - Berkeley Software Distribution 6
 - Berkeley UNIX 6
 - best-effort delivery 27, 91, 291, 561
 - BGP 266, 561
 - big endian 69, 561
 - bind system call 339
 - BISYNC 562
 - block 479
 - BNC 562
 - BNC connector 24
 - BOOTP 136, 365, 366, 562
 - BOOTP protocol 365
 - bootstrap 365, 427
 - BOOTstrap Protocol 366
 - Border Gateway Protocol 562
 - bps 562
 - bridge 109, 562
 - mail 437
 - broadband 562
 - broadcast 26, 152, 504, 562
 - broadcast address 62, 289
 - broadcasting 289
 - brouter 563
 - BSC 563
 - BSD UNIX 6, 563
 - buffer 180
 - bus 26
 - Butterfly 246
 - byte 29
 - byte order 69

C

capacity 28
carriage control 410
carrier sense 27
CCIRN 563
CCITT 39, 45, 164, 563
CDDI 32
cell 36, 308, 563
checksum 100, 126, 563
Chernobylgram 563
CIDR 154, 563
circuit switching 18
class A address 61
class B address 61
class C address 61
class D address 291
class of address 60, 563
class of name 390
Classless Inter-Domain Routing 154, 564
client 325, 326, 368
 example 357
client-server 325, 564
 see Volume III
CLNS 492
close 217
closing connections 217
clumping 226
cluster 503
coaxial cable 20
collision 27
colon hexadecimal notation 502
community 460
congestion 130, 214
congestion avoidance 215
congestion collapse 214
congestion control 203
congestion window 214
congestion window limit 214
connect system call 340
connected socket 340
connection 5, 192, 216, 306, 564
 closing 217
 reset 219
connection abstraction 199

connection endpoint 200
connection oriented 18, 36
connectionless 18, 91
connectionless service 91, 564
content type 443
context specific 463
control connection 422
control message 123
control packet 447
convergence 311
Copper Distributed Data Interface 32
core gateway 564
core router 235
cosmic significance 418
count to infinity 272
counter rotating 33
CRC 29, 564
CSMA 27
CSMA/CD 27, 564
CSNET 44
cumulative acknowledgement 208

D

DARPA 564
data availability 472
data field 29
data mark 414
data stream 201
data transfer connection 422
datagram 5, 91, 92, 564
 MTU 95
 UDP 181
 fragmentation control 98
 size 95
 time to live 99
 type of service 93
datagram format 92
datagram options 100
date service 328
DCA 37
DCE 565
DDCMP 565
DDN 37, 565
default route 233, 276

- default router 115
 - Defense Communication Agency 6
 - delay 19
 - delayed acknowledgement 225
 - demultiplex 174, 565
 - designated router 280
 - destination port 180
 - destination unreachable 128
 - DHCP 365, 366, 372, 565
 - DHCP lease 373
 - DHCP protocol 365
 - dial-up IP 46
 - digest subtype (MIME) 444
 - Dijkstra shortest path algorithm 246
 - direct delivery 111
 - directed broadcast address 62, 565
 - distance metric 242
 - Distance Vector Multicast Routing Protocol 297
 - DNS 8, 387, 467, 565
 - dn_comp procedure 354
 - dn_expand procedure 353
 - DO (TELNET) 415
 - do not fragment 98, 367
 - DOD 2
 - DOE 2
 - domain 565
 - domain class 390
 - domain name 383, 387
 - pointer query 400
 - recursive resolution 393
 - server 391
 - zone 403
 - Domain Name System 8, 383
 - domain name system 387
 - domain suffix list 400
 - domain type 390
 - DON'T (TELNET) 415
 - dotted decimal notation 65, 565
 - dotted hexadecimal 290
 - dotted quad notation 65
 - draft standard 516
 - dropping packets 129
 - DS3 44, 566
 - DTE 566
 - DVMRP 297, 566
 - dynamic configuration 372
 - Dynamic Host Configuration Protocol 366, 372
- E**
- e-mail
 - see* electronic mail
 - E.164 316, 566
 - EACK 566
 - echo
 - GGP request/reply 245
 - ICMP request/reply 127
 - UDP request/reply 326
 - UDP server 326
 - echo port 326
 - echo service 326
 - EGP 176, 254, 566
 - message header 255
 - neighbor 255
 - neighbor acquisition 256
 - neighbor reachability 257
 - peer 255
 - poll request 258
 - protocol 249
 - routing update 255, 259
 - third party restriction 259
 - EGP2 264
 - EGP3 264
 - EIA 566
 - electronic mail 4, 433
 - destination 435
 - list 435
 - spool 434
 - encapsulation 94, 566
 - ICMP 125
 - IP 94
 - IP datagram 95
 - RARP 85
 - enclosures 444
 - encoding type 443
 - encryption 475
 - end-of-packet bit 311

end-to-end 166, 168
endhostent procedure 355
endnetent procedure 355
endpoint 200
endprotoent procedure 356
endservent procedure 356
extension header 494
epoch date 328, 567
error reporting mechanism 124
escape 408
escape sequence 413
establishing a connection 216
ether 20
Ethernet 20, 567
 AUI 21
 CRC 29
 address 28
 broadcast 26
 collision 27
 data field 29
 frame 29
 host adapter 21
 host interface 21
 hub 25
 preamble 29
 repeater 30
 transceiver 21
 type 95
 type field 29
Ethernet meltdown 566
Ethernet multicast 290
exchanger (e-mail) 439
exec system call 338
exploder 435
exponential backoff 27
Exterior Gateway Protocol 254
exterior neighbor 254
exterior router 254
eXternal Data Representation 430, 567
extra hop problem 251

F

fair queueing 567
FCCSET 15, 567

FDDI 32, 567
 frame 35
 symbol 35
FDM 567
fetch-store paradigm 458
file descriptor 336
file server 326, 419, 567
file transfer 4, 421
File Transfer Protocol 421
filter 479
fingerd 330
finite state machine 219
firewall 471, 568
flat namespace 384, 568
flow 495
flow control 130, 201, 568
FLOW LABEL 494
Ford Fulkerson 240
fork system call 338
forwarding
 mail 435
fragment bit 98
Fragment Extension Header 498
fragmentation 95, 498, 568
fragmentation control 98
fragmentation needed 129
frame 29, 35, 164, 568
 self-identifying 30
FTP 65, 421, 467, 568
full duplex 193
Fuzzball 41
FYI 568

G

gated 279, 568
gateway 52, 109, 242, 254, 569
 VAN 45
 designated 280
 mail 437
gateway requirements 569
gateway-to-gateway protocol 242
Gbps 36
getdomainname system call 349
gethostbyaddr procedure 355

gethostbyname procedure 354
gethostent procedure 355
gethostname system call 348
getnetbyaddr procedure 355
getnetbyname procedure 355
getnetent procedure 355
getpeername system call 344
getprotobyname procedure 355
getprotobynumber procedure 356
getprotoent procedure 356
getservbyname procedure 356
getservbyport procedure 356
getservent procedure 356
getsockname system call 345
GGP 242, 569
GIF 443
gif 444
global Internet 2
global name 453
gopher 12, 465, 569
GOSIP 569
graceful shutdown 217

H

half duplex 193
hardware address 20, 28, 74, 569
HDLC 39, 164
header length field 93
HELLO 267, 569
hello (OSPF) 281
hello interval (EGP) 256
HELO 441, 569
HHS 2
hidden network 251
hierarchical addressing 145
hierarchical routing 145, 569
high-level name 384
historic 516
history 6
hold down 273
hop count 133, 242, 271, 569
hop count metric 271
HOP LIMIT 494
host 38, 570

host adapter 21
host interface 21
host requirements 570
host table 71
htonl procedure 351
htons procedure 351
hub 25, 570

I

IAB 8, 570
IANA 66, 516, 570
IBM token ring 46
ICCB 6, 570
ICMP 123, 124, 176, 570
 address mask 136
 checksum 126
 code 126
 destination unreachable 128
 echo request/reply 127
 encapsulation 126
 information request/reply 136
 message format 125
 message types 127
 parameter problem 134
 protocol 123
 redirect 131
 redirect message 131
 source quench 130
 subnet mask 136
 time exceeded 133
 timestamp 134
 type 126
IEEE 28
IEN 11, 570
IESG 10, 571
IETF 10, 570
 area manager 10
 working group 10
IGMP 289, 294, 571
IGP 269, 571
IMP 37
implementation
 see Volume II
InATMARP 319

- indirect delivery 111
- inet_addr procedure 351
- inet_inaof procedure 352
- inet_makeaddr procedure 352
- inet_netof procedure 352
- inet_network procedure 351
- inet_ntoa procedure 352
- infinity 373
- infinity (small) 272
- information request 136
- information security 472
- INOC 571
- inside 477
- integrated 420
- integrity 475
- interface 21
- Interface Message Processor 37
- interior
 - router or gateway 267
- interior gateway protocol 269
- interior neighbor 254
- internals
 - see* Volume II
- International Organization for Standardization 571
- International Telecommunications Union 571
- Internet 571
- internet 50, 51, 89, 571
 - address 60, 73, 83
 - control message 123
 - properties 51
 - router 52
 - routing table 113
- internet access 476
- Internet Activities Board 8
- Internet address 571
- internet address 74
 - dotted decimal notation 65
- Internet Architect 9
- Internet Architecture Board 8
- Internet Assigned Number Authority 66
- Internet Assigned Numbers Authority 516
- Internet Control Message Protocol 124
- Internet datagram 91
- Internet draft 11, 12
- Internet Engineering Notes 11
- Internet Engineering Steering Group 10
- Internet Engineering Task Force 10
- internet firewall 476
- internet gateway 52
- Internet Group Management Protocol 294
- internet layer 167
- internet management 447
- Internet Network Information Center 66
- Internet Protocol 89, 91, 571
 - version 4 491
- Internet Research Group 11
- Internet Research Steering Group 11
- Internet Research Task Force 10
- internet router 52
- internet routing 110
- internet security 472
- internet services 3
- Internet Society 11, 572
- Internet standard 516
- Internet Task Force 8
- Internet worm 37, 330, 572
- internetwork
 - see* internet
- internetworking 1
- INTERNIC 11, 66, 384, 513, 572
- interoperability 3, 572
- interpret as command 413
- interrupt 205
- Inverse ARP 87
- Inverse ATMARP 319
- inverse query 400
- IP 572
 - address 60
 - checksum 100
 - data 100
 - destination address 100
 - dial-up 46
 - encapsulation 94, 95
 - header length 93

- option code 100
 - precedence 93
 - protocol field 494
 - purpose 91
 - reassembly 96
 - record route 102
 - router 52
 - source address 100
 - source route 103
 - time to live 99
 - timestamp 104
 - type of service 93
 - version 93
 - IP address 60, 91, 572
 - dotted decimal notation 65
 - IP datagram 92, 572
 - IP forwarding 110
 - IP gateway 242, 254
 - IP multicasting 291
 - IP next generation 492
 - IP options 100
 - IP router 109
 - IP routing 110
 - IP switching 110
 - IP-based technology 91
 - IP6 492
 - ipAddrTable 456
 - ipInReceives 456
 - IPng 492, 572
 - see IPv6
 - IPv4 491, 492, 573
 - IPv6 492, 573
 - broadcast 504
 - cluster 503
 - destination 495
 - end-to-end header 500
 - fragmentation 498
 - hop limit 495
 - hop-by-hop header 500
 - multicast 503, 504
 - path MTU 498
 - payload length 495
 - provider prefix 507
 - source route 500
 - subnet prefix 507
 - subscriber prefix 507
 - unicast 503
 - version 495
 - IRSG 11
 - IRTF 10, 573
 - ISDN 573
 - ISO 163, 573
 - ISO model 163
 - ISOC 573
 - ISODE 573
 - iterative name resolution 393
 - ITU-TS 45, 164, 573
- J**
- jpeg 444
- K**
- k-out-of-n rule 246, 257
 - Karn's Algorithm 212, 573
 - Kbps 573
 - kerberos 487
 - Kramer 573
- L**
- label 387
 - LAN 19, 574
 - LAPA 164
 - LAPB 39, 164
 - layering 159, 166
 - ISO 163
 - TCP/IP 165
 - layering principle 169
 - learning bridge 31
 - lease 373
 - level 1 574
 - level 2 574
 - level 3 574
 - limited broadcast address 62
 - linefeed 410
 - link-state 245
 - LIS 314, 574
 - listen system call 346

- little endian 69, 574
 - LLC 312, 574
 - load balancing 280
 - local area network 19
 - local network broadcast address 62
 - logging 485
 - Logical IP Subnet 314
 - Logical Link Control 312
 - long haul network 19
 - longest-match 155
 - loopback 65
 - loose source routing 104
 - low-level name 384
- M**
- MAC 574
 - machine status
 - see* ruptime
 - magic cookie 370
 - mail alias expansion 435
 - mail bridge 437, 574
 - mail destination 435
 - mail exchanger 390, 439, 574
 - mail exploder 435, 575
 - mail forwarding 435
 - mail gateway 437, 575
 - mail processing 435
 - mail queue 434
 - mail relay 437
 - mail spool area 434
 - mailbox address 435
 - mailing list 435
 - MAN 575
 - management agent 449
 - Management Information Base 450, 575
 - manual bypass 482
 - manual configuration 372
 - martians 575
 - maximum segment lifetime 219, 575
 - maximum segment size 206, 575
 - maximum transfer unit 95, 575
 - MBONE 300, 575
 - Mbps 575
 - metric transformation 265
 - MIB 450, 576
 - MIB-II 451
 - mid-level network 39
 - military network 6
 - MILNET 6, 37, 576
 - MIME 576
 - mixed subtype (MIME) 444
 - monitoring 485
 - Mosaic 469, 576
 - MOTIS 165
 - mrouted 298, 576
 - MSS 206, 576
 - MTP 440
 - MTU 95, 576
 - datagram 95
 - multi-homed host 61, 576
 - multicast 289, 503, 504, 576
 - all hosts 504
 - all nodes 504
 - all routers 504
 - multicast address 29
 - Multicast Backbone 300
 - multicast group 29
 - multicast kernel 298
 - multicast routers 291
 - multicasting 63, 290
 - multimode 304
 - multipart type (MIME) 444
 - multiplex 174
 - multiplicative decrease 214
 - Multipurpose Internet Mail Extensions 443
 - mutual trust 474
- N**
- Nagle algorithm 226, 576
 - NAK 577
 - name 51, 59, 384
 - abbreviation 399
 - domain 383, 387
 - recursive resolution 393
 - resolver 391
 - name caching 395
 - name resolution 393, 577

- name server 8
 - name-to-address translation 383
 - namespace partition 386
 - NAP 506, 577
 - NASA 2
 - National Institute for Standards and Technology 454
 - National Science Foundation 39
 - NBS 454
 - neighbor 242
 - neighbor acquisition 255
 - neighbor router 254
 - NetBIOS 577
 - netstat 65, 333
 - network 18
 - address 20, 59
 - capacity 28
 - Network Access Provider 506
 - network byte order 577
 - Network File System 429
 - network information center 11
 - network interface 166, 167
 - network management 450, 577
 - network MTU 95
 - network security 472
 - Network Service Access Point 316
 - Network Service Provider 154
 - network services 4
 - network standard byte order 69
 - Network to Network Interface 304
 - network virtual terminal 408, 410, 424
 - next generation 492
 - NEXT HEADER 497
 - next hop 113, 117, 150
 - NFS 429, 577
 - NIC 577
 - NIST 577
 - NNI 304
 - NOC 37, 577
 - non-selfreferential 608
 - nonauthoritative 395
 - noncore router 235
 - nonrouting router 235
 - NSAP 578
 - NSAP address 317
 - NSF 2, 39, 578
 - NSFNET 7, 40, 578
 - nslookup 404
 - ntohl procedure 351
 - ntohs procedure 351
 - null 463
 - number of hops 242, 271
 - NVT 410
- O**
- object identifier 453
 - OC3 578
 - octet 29
 - on-line access 420
 - Open SPF protocol 279
 - open standard 492
 - open system interconnection 1
 - open-read-write-close 336
 - Option Overload 379
 - options 100
 - Organizationally Unique Identifier 313
 - OSI 578
 - OSPF 267, 279, 578
 - OUI 313
 - out of band 205, 414
 - outside 477
 - overlapping segment problem 229
- P**
- packet 18, 578
 - packet filter 467, 479
 - packet radio 47
 - packet switch 19
 - packet switching 18
 - Packet Switching Node 37
 - PAD 165
 - parallel subtype (MIME) 444
 - parameter problem 134
 - parent domain 394
 - passive 271
 - passive monitoring 486
 - passive open 201

- Path MTU 498
 - PDN 45, 578
 - PDU 460
 - peer backbone networks 238
 - PEM 487, 578
 - permanent virtual circuit 306
 - PF_INET 337, 340
 - physical address 28, 74
 - piggybacking 193
 - PING 127, 138, 578
 - pipe 338
 - Point to Point Protocol 171
 - pointer query 400, 401
 - poison reverse 274
 - polling interval (EGP) 256
 - port 129, 326, 578
 - ARPANET 38
 - PORT command (FTP) 427
 - port unreachable 186
 - positive acknowledgement 193, 578
 - PPP 171, 579
 - preamble 29
 - primary server 86
 - privacy 472, 475, 476
 - privacy enhanced mail 487
 - process 179, 326
 - promiscuous ARP 142, 579
 - proNET 46
 - proposed standard 516
 - protocol 3, 579
 - ARP 73
 - BOOTP 365
 - DHCP 365
 - EGP 249, 254
 - GGP 242
 - HELLO 267, 276
 - ICMP 123
 - IGMP 289
 - IGP 269
 - IP 89
 - IPng 492
 - IPv4 492
 - IPv6 492
 - Internet 89
 - MTP 440
 - OSPF 267, 279
 - RARP 83, 84
 - RIP 267, 270
 - SMTP 440
 - SNMP 450
 - ST 492
 - TCP 191, 198
 - TELNET 408
 - UDP 179, 180
 - application 166
 - data link 166
 - datagram 179
 - internet 166
 - layering 159, 166
 - network management 450
 - port 180
 - standards 12
 - stream 191
 - protocol data unit 460
 - protocol family 159
 - protocol port 199, 221, 326, 579
 - protocol standards 8
 - protocol suite 159
 - provider prefix 507
 - proxy ARP 142, 579
 - pseudo header 182, 207, 506, 579
 - pseudo terminal 410
 - PSN 37
 - Public Data Networks 45
 - public key encryption 475, 579
 - PUP 579
 - purpose of IP 91
 - push 193, 221, 580
 - PVC 306
- R**
- RARP 83, 84, 136, 580
 - RARP server 85
 - rcp 7
 - RDP 580
 - reachability 262
 - Read Only Memory 83
 - read request 428

- read system call 343
 - readv system call 343
 - reassembly 96, 97, 311, 498, 580
 - reassembly timer 97
 - receiver SWS avoidance 224
 - recommended 516
 - record route option 102
 - recursive name resolution 393
 - recvfrom system call 344
 - recvmsg system call 344
 - redirect 131, 580
 - redirect message 131
 - reference model 163
 - regional network 39
 - relay
 - mail 437
 - relay agent 373
 - reliable stream service 90
 - reliable transfer 193
 - remote login 4, 408
 - Remote Procedure Call 430, 480
 - repeater 30, 580
 - Request For Comments 11, 511
 - required 516
 - reset 219
 - resolving addresses 74
 - resolving names 393
 - resource records 398
 - res_init procedure 353
 - res_mkquery procedure 353
 - res_send 353
 - retransmission 193, 208, 209
 - retransmit 194, 208
 - Reverse Address Resolution Protocol 84
 - reverse path forwarding 152, 580
 - RFC 11, 511, 580
 - RFNM 38
 - ring 33
 - RIP 270, 580
 - RIP protocol 267
 - RJE 581
 - rlogin 416, 581
 - ROADS 153, 581
 - ROM 83
 - round trip sample 209
 - route 53, 59, 131, 581
 - route advertisement 242
 - routed 270, 298, 581
 - router 52, 109, 112, 131, 581
 - designated 280
 - exterior 254
 - interior 254
 - nonrouting 235
 - stub 235
 - router hops 242
 - routing 91, 109
 - routing cycle 133
 - routing hierarchically 145
 - Routing Information Protocol 270
 - routing loops 239
 - routing table 113
 - routing update 243, 255, 259
 - RPC 430, 467, 581
 - RS232 581
 - rsh 416
 - RTO 581
 - RTT 582
 - runtime 331
- S**
- SACK 582
 - sample round trip time 209
 - SAR (ATM) 311
 - security 115, 471, 472
 - security perimeter 477
 - segment 201, 203, 582
 - segmentation 311
 - selective acknowledgement 582
 - self clocking 226
 - self-healing 33
 - self-identifying frame 30, 38, 312, 582
 - send system call 342
 - sender SWS avoidance 225
 - sendmsg system call 342
 - sendto system call 342
 - Serial Line IP 171
 - server 84, 325, 368
 - RARP 85

- example 359
- file 326
- primary 86
- time of day 326
- service
 - connectionless service 5
 - reliable stream service 5
 - reliable stream transport 191
 - unreliable packet delivery 91
- SERVICE TYPE 494
- setdomainname system call 349
- sethostent procedure 355
- sethostname system call 348
- setnetent procedure 355
- setprotoent procedure 356
- setservent procedure 356
- SGMP 464, 582
- shortest path algorithm 246
- Shortest Path First 245
- signaling 306, 582
- silly window syndrome 224, 582
- Simple IP 492
- Simple IP Plus 492
- Simple Network Management Protocol 450
- SIP 492, 582
- SIPP 492, 582
- size
 - datagram 95
- sliding window 195, 583
- SLIP 171, 583
- slow convergence 272
- slow-start 214, 215, 583
- small infinity 272
- SMDS 312, 583
- SMTP 440, 583
- SNA 583
- SNAP 313, 583
- SNMP 450, 467, 583
- SNMPv2 450
- SOA 583
- sockaddr 340
- sockaddr_in 340
- socket 7, 337, 584
- socket library 363
- socket system call 337
- soft-start 215
- Sorcerer's Apprentice Bug 429
- source port 180
- source quench 130, 584
- source route 103, 129, 500, 584
- source route option 103
- SPF 245, 584
- split horizon update 273
- spoofing 143
- spooling 434
- SPREAD 246
- ST protocol 492
- standard byte order 69
- standard error 417
- standard input 417
- standard output 417
- standardization 12
- STD 584
- stream 192
- strict source routing 104
- Structure of Management Information 452
- stub network 484
- stub router 235
- subnet address 143
- subnet addressing 584
- subnet broadcast 152
- subnet mask 136, 147, 150
- subnet prefix 507
- subnet route 143
- subnet routing 149
- subnet rule 149
- subnetting 143
- SubNetwork Attachment Point 313, 584
- subscriber 506
- subscriber prefix 507
- subtype 443
- supernet addressing 153, 584
- supernetting 153
- SVC 306
- Switched Multimegabit Data Service 312

switched virtual circuit 306
SWS 584
SWS avoidance 224
symbol 35
symmetric 415
SYN 584
SYNCH 414
system call 336

T

T3 44, 585
tap 21
task 179
TCP 176, 191, 198, 585
 protocol port 221
TCP header 204
TCP protocol 191
TCP/IP 2
TCP/IP Internet 2
TCP/IP Internet Protocol Suite 585
TDM 585
TDMA 585
technology independence 5
telephone system 386
TELNET 65, 408, 424, 467, 585
TFTP 427, 585
thick Ethernet 24
thicknet 24, 585
thin-wire Ethernet 23, 24
thinnet 23, 24, 586
time exceeded message 133
time service 328
time to live 99, 133, 171, 291, 298, 395,
 494
time-of-day server 326
timeout 208, 209
timeout and retransmission 367
timer backoff 212
timestamp option 104
timestamp reply 135
timestamp request 135
TLI 363, 586
TLV encoding 501, 586
tn3270 418, 586

token 33
token ring 33, 46, 586
TOS 93, 586
TP-4 586
traceroute 138, 586
traffic class 496
trailers 586
transceiver 21, 586
transceiver cable 21
transient multicast groups 291
Transmission Control Protocol 191, 198
transparent access 420
transparent router 141
transparent service 408
transport layer 166
Transport Layer Interface 363
triggered updates 274
Trivial File Transfer Protocol 427
TRPB 298, 587
truncated reverse path broadcast 298
trust 474
trusted hosts 408
TTL 99, 298, 587
tunnel 298, 299
tunneling 45, 480, 587
twisted pair Ethernet 25, 587
two-stage oscillation 277
type field 29
type of name 390
Type Of Service 93
type of service routing 280, 587

U

UART 587
UCBCAST 587
UDP 176, 180, 587
 echo server 326
 encapsulation 183
 message format 181
 port 180
 protocol 179
 pseudo header 182
unacknowledged packet 196
unconnected socket 340

UNI 304
unicast 29, 503, 588
unicast address 290
universal assignment 186
universal communication service 59
universal interconnection 5, 50
universal time 105, 328, 588
unreachable destination 128
unreliable 91
unreliable packet delivery 90, 91
urgent data 205, 414, 418, 588
URL 588
user datagram 179, 181, 365
User Datagram Protocol 180
user level process 179
user process 326
User to Network Interface 304
UUCP 440, 588

V

VAN gateway 45
vBNS 44, 588
VCI 307
vector-distance 240, 588
very high speed Backbone Network Service 44, 588
virtual circuit 45, 192, 588
virtual circuit identifier 307
virtual path identifier 307
VPI 307
VPI/VCI 307, 588

W

WAN 19, 589
weak authentication 475
weakest link axiom 477
well-known address 291
well-known port 186, 221, 589
whole-file copying 420
wide area network 19
wildcard 346
WILL (TELNET) 415
window 196, 589

congestion 214
window advertisement 202
window size 196, 582
Winsock 335, 363, 589
wireless network 47
wiring 26
WON'T (TELNET) 415
working group 10, 589
World Wide Web 12, 465, 589
worm 37, 330
write 226
write request 428
write system call 341
writev system call 341
WWW 589

X

X 589
X-Window System 590
X.25 39, 45, 589
X.400 165, 589
X25NET 44, 589
XDR 430, 467, 590
XNS 107

Z

zero compression 502
zone of authority 403, 590

The Series Of Internetworking Books from Douglas Comer and Prentice Hall

Internetworking With TCP/IP Volume I: Principles, Protocols, and Architecture 3rd ed., 1995, ISBN 0-13-216987-8

The classic reference in the field for anyone who wants to understand Internet technology, Volume I surveys the TCP/IP and describes each component. The highly accessible text presents the scientific principles used in the construction of TCP/IP, and shows how the components were designed to work together. It covers details of each protocol, including ARP, RARP, IP, TCP, UDP, RIP, OSPF and others.

Internetworking With TCP/IP Volume II: Design, Implementation, and Internals (with D. Stevens) 2nd ed: 1994, ISBN 0-13-125527-4

Ideal for implementors, Volume II continues the discussion of Volume I by using code from a running implementation of TCP/IP to illustrate all the details. The text shows, for example, how TCP's slow start algorithm interacts with the Partridge-Karn exponential retransmission backoff, and how routing updates interact with datagram forwarding.

Internetworking With TCP/IP Volume III: Client-Server Programming and Applications (with D. Stevens)

BSD Socket Version: 1992, ISBN 0-13-474222-2

AT&T TLI Version: 1993, ISBN 0-13-474230-3

Volume III describes the fundamental concept of client-server computing used to build all distributed computing systems. The text discusses various server designs as well as the tools and techniques used to build clients and servers, including Remote Procedure Call (RPC). It contains examples of running programs that illustrate each of the designs and tools. Two versions of volume III are available for the widely used BSD socket and AT&T TLI interfaces.

The Internet Book: Everything you need to know about computer networking and how the Internet works

Paperback: 1995, ISBN 0-13-151565-9

Study Guide: 1995, ISBN 0-13-188012-8

Book plus Study Guide: 1995, ISBN 0-13-400029-3

A gentle introduction to networking and the Internet, *The Internet Book* does not assume the reader has a technical background. It explains the Internet, how it works, and services available in general terms, without focusing on a particular computer or a particular brand of software. Ideal for someone who wants to become Internet and computer networking literate, *The Internet Book* explains the terminology as well as the concepts; an extensive glossary of terms and abbreviations is included. A separate *Study Guide* provides suggestions for readers, review questions, and exercises).

**To order from North America, contact your local bookstore,
call 1-515-284-6751 or send a FAX 1-515-284-2607
Outside North America, contact your Prentice Hall representative**

