

IP MULTICASTING

Concepts and Applications

M. GONCALVES & K. NILES

IP 多信道广播

McGraw-Hill Book Co
麦 哥 律 书 出 版 公 司

IP MULTICASTING

CONCEPTS AND APPLICATIONS



世界图书出版公司
McGraw-Hill Book Co

Library of Congress Cataloging-in-Publication Data

Goncalves, Marcus

IP multicasting / Marcus Goncalves, Kitty Niles.

p. cm.

Includes index.

ISBN 0-07-913791-1

1. Multicasting (Computer networks) 2. Business enterprises—
Computer networks. I. Niles, Kitty. II. Title.

TK5105.887.G66 1999

004.6'6—dc21

98-47432

CIP

McGraw-Hill

A Division of The McGraw-Hill Companies



Copyright ©1999 by The McGraw-Hill Companies, Inc. All rights reserved. Except as permitted under the United States Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a data base or retrieval system, without the prior written permission of the publisher.

P/N 024845-1

PART OF ISBN 0-07-913791-1

Copyright © 1999 by McGraw-Hill Companies, Inc. All Rights reserved. Jointly Published by Beijing World Publishing Corporation/McGraw-Hill. This edition may be sold in the People's Republic of China only. This book cannot be re-exported and is not for sale outside the People's Republic of China.

IE ISBN 0-07-116313-1

Information contained in this work has been obtained by The McGraw-Hill Companies, Inc. ("McGraw-Hill") from sources believed to be reliable. However, neither McGraw-Hill nor its authors guarantee the accuracy or completeness of any information published herein and neither McGraw-Hill nor its authors shall be responsible for any errors, omissions, or damages arising out of use of this information. This work is published with the understanding that McGraw-Hill and its authors are supplying information but are not attempting to render engineering or other professional services. If such services are required, the assistance of an appropriate professional should be sought.

Acknowledgments

We must acknowledge a broad debt of gratitude to all of the authors of the numerous RFC documents referenced and used in the preparation of this book.

We gratefully acknowledge: Junko Higasa and the Internet Initiative Japan Inc. in Tokyo for sharing works in progress with us, and Rob Lee, a consultant for networking and Internet technologies through his firm, Results From Technology, for permission to use material he wrote. We also thank Eliot Abraham and Gilat Satellite Networks, and Phillip Emer, the Associate Director of Advanced Technology Development from North Carolina State University for sharing material with us and granting us permission to use that material.

We thank Stardust Technologies for authorizing us to use a couple of their whitepapers. Ken Miller and StarBurst Communications Corporation, and Laura Spear and Xplex Networks have our thanks. Our thanks also go to Katie Greene and Concentric Corporation, Julie Martin and Backweb Technologies, and Rich Borden and FORE Systems.

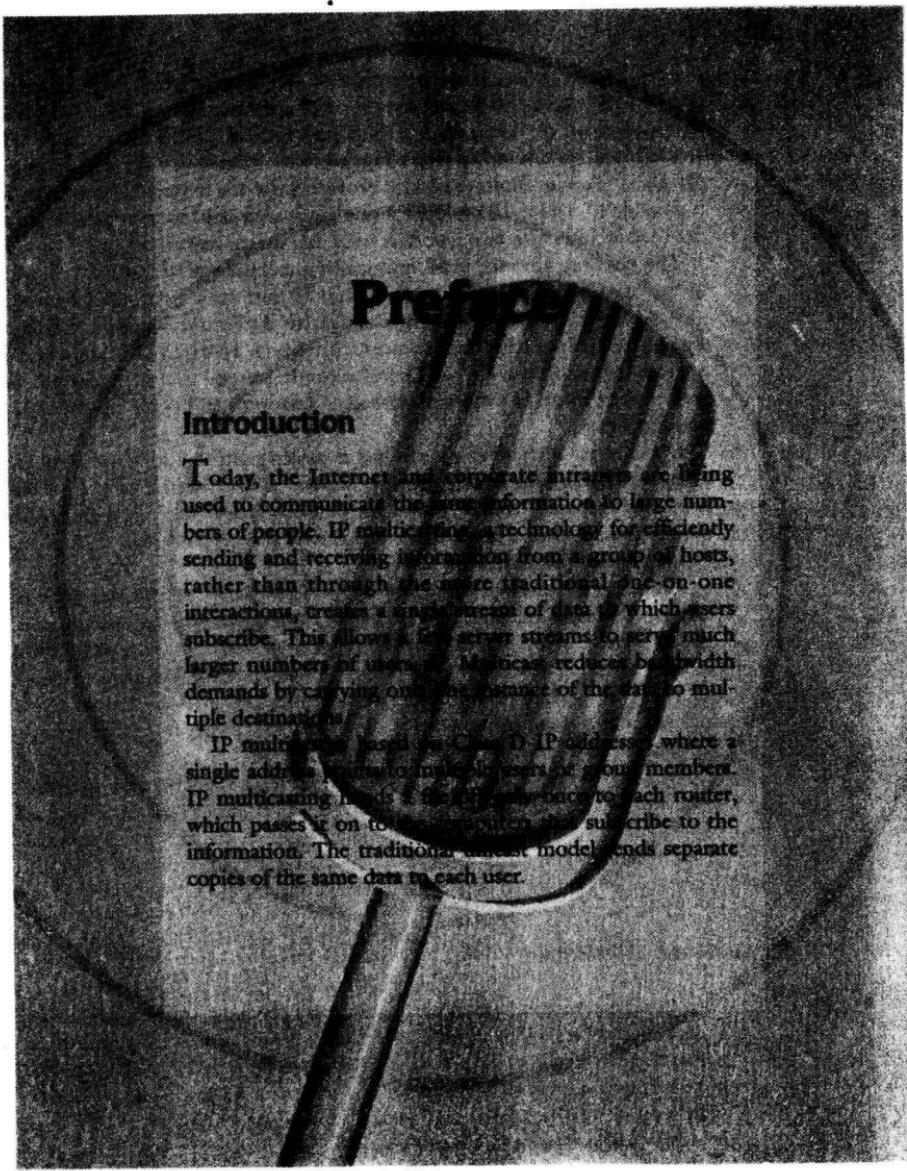
Thanks to Simon Oosthoek, Jorma Paananen, Matti Tella, and Juha Luoma for their cooperation and input and for so graciously providing material for us to use at our discretion.

We gratefully acknowledge the assistance of Anup Rao, Henning Schulzrinne, Rob Lanphier, and Kevin Epstein.

Special thanks go to Sue Michener, Jill and Don Eastlake for their understanding and support during the process of writing this book.

We would like to thank Simon Yates of McGraw-Hill and his confidence in us, entrusting us with this task. We would also like to thank all the members of the production team at McGraw-Hill and in particular, Patricia Wallenburg for the high quality of her work, her commitment, her professionalism, and her cheerfulness throughout.

Glory be to God, for allowing us to contribute to a better multimedia virtual world in this way.



Preface

Introduction

Today, the Internet and corporate intranets are being used to communicate business information to large numbers of people. IP multicasting technology for efficiently sending and receiving information from a group of hosts, rather than through the more traditional one-on-one interactions, creates a new form of data in which users subscribe. This allows a few server streams to serve much larger numbers of users. This model reduces bandwidth demands by carrying only the amount of traffic to multiple destinations.

IP multicasting uses group IP addresses, where a single address can be used to describe group members. IP multicasting needs a delivery service to each router, which passes it on to the hosts that have subscribed to the information. The traditional unicast model sends separate copies of the same data to each user.

IP multicast relies on the existence of an underlying multicast delivery system to forward data from a sender to all the intended receivers. Until all the routers on the Internet are multicast-enabled, multicast through the Internet will be limited to tunneling, services offered by multicast-enabled islands, and some Internet service providers. In the meantime, however, it is possible to deploy IP multicast using a private virtual network based on a network architecture or network service that is multicast-enabled.

The next big advance in IP multicasting technology will probably come in the area of reliable delivery of data. IP multicast protocol does not include any acknowledgment mechanism to determine whether the data sent to recipients have arrived correctly. Much research is underway to address this issue and the IETF has recently formed a working group to sort out the maze of proposed protocols and tools and come up with recommendations.

IP multicast is eminently useful for deploying information from applications based on push technologies where bandwidth consumption is a problem. For example, the IP multicast protocol uses only the bandwidth required to update a single desktop in order to update several hundred desktops requesting the same data.

Multicast technology is being heavily used for experimentation into audio/video conferencing over IP. Some commercial applications are coming into their own. Big organizations such as Toys 'R Us, Ford Motor Company, General Motors, The Ohio Company, and the American Stock Exchange are already using multicasting technology to solve software and information distribution problems. Some satellite delivery companies also are using IP multicast over their delivery systems.

For a full-blown implementation of IP multicast to work, both the hardware and software in the data path need to support IP multicast. The good news is that IP multicast is already supported in the newest equipment from networking vendors such as Cisco, and a few ISPs such as MCI and Uunet are beginning to deploy IP multicast services.

Who Should Read This Book

This book is an introduction to IP multicast concepts and features and is intended for users with varying degrees of experience with the IP multicast technologies, decisionmakers wrestling with the problems of corporate communication, system administrators and operators managing their own workstations, and programmers. Finally, this book is intended for communications specialists, field system engineers, and the general audience who wants or needs to know about IP multicasting.

This book can be used by those who design or develop applications that run on the network or work in technical support. It can be used to gain a basic understanding of IP multicasting and as a source for information about IP multicasting. Information about ISPs, suppliers and vendors, current and future IP multicasting applications, implementations, and supporting software is here.

For an engineer interested in evaluating or implementing IP multicast for an organization, product or service, this book will facilitate understanding of how to deploy IP multicast over these network infrastructures.

Organization

The book is organized into 24 chapters and appendices, which are broadly separated into four parts:

PART I TCP/IP and IP Multicasting Fundamentals

This part comprises Chapters 1 through 10 and is about the technology around IP multicasting. We provide a brief history of the Internet and introduce the different types of network infrastructures upon which IP multicast can run. This part also includes fundamentals about IP multicasting, describes how it works, introduces multicast routing concepts, algorithms, and protocols, presents interoperability, tunneling, and implementation issues, and touches on the challenges facing IP multicasting.

PART II The Mbone and Other Multicast Networks

Part two consists of Chapters 11 through 15 and is about preparing for and multicasting on different network infrastructures. We discuss Asynchronous Transfer Mode technology and the Mbone, we briefly look at the higher-level protocols used with IP multicast, we investigate writing IP multicast applications, and introduce some standard IP multicast application programming interfaces.

PART III IP Multicasting Applications

This consists of Chapters 16 through 21 and is about IP multicast applications, equipment, and suppliers. We look at some of the major players, IP multicasting application software for servers and clients, and IP multicasting content providers, management and middleware software products, and IP multicasting network service providers. Lastly we describe available IP multicasting network equipment and software development kits.

PART IV IP Multicasting Resource Center

Chapters 22-24 are about techniques used for developing IP multicast-enabled applications and how to deploy IP multicast in a diverse network platform. They also cover the main IP multicast switch equipment, including both companies and their products.

Appendices provide a comprehensive list of IP Multicast acronyms, a bibliography, and an extensive glossary.

Typographic Conventions

Code and data structures examples included in the book are expressed in Courier fonts.

A Note on Trademarks and Standards

Product names mentioned herein may be trademarks and/or registered trademarks of their respective companies.

We have made reference to many RFCs. This is a very fast-moving area and we have attempted to be as accurate as possible. Part of this book probably became out of date as we wrote it, and other parts may have changed as the book was being published. The standards and the code and applications themselves are, as always, the best source of information.

A great many RFCs were used in the preparation of this document and we gratefully acknowledge all the authors. The bibliography at the end of this book lists the RFCs and their respective authors. The following copyright notice applies to all RFCs:

Copyright © The Internet Society (1998). All Rights Reserved.—This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

Table of Contents

Preface

Introduction	xiii
Who Should Read This Book	xiv
Organization	xv
Typographic Conventions	xvi
A Note on Trademarks and Standards	xvi

PART I TCP/IP and IP Multicasting Fundamentals

CHAPTER 1 The Topology and Internal Structure of a Network	3
Some Basics	4
Local Area Networks	4
Wide Area Networks	5
Metropolitan Area Networks	5
The Internet	5
Virtual Private Networks or Extranets	10
Intranets	11
A Basic Network	11
A Network and its Topology	12
A Network and its Protocol Stack	14
IP Addressing Summary	23
General IPv4 Addressing Architecture	23
IP Subnetting	24
ARP Protocol	26
Reverse Address Resolution Protocol (RARP)	27
Internet Control Message Protocol (ICMP)	27
Internet Group Management Protocol (IGMP)	28
Network Components	29
CHAPTER 2 Introducing the Network Infrastructure	31
An Overview of Backbone LANs and WANs	32
Transport Technologies	34
Switching Techniques	35
Campus and LAN Infrastructures	37
Ethernet Technologies	37
Token Ring	41
Fiber Distributed Data Interface	43
Comparison of Ethernet, Token Ring, and FDDI	44
Backbone WAN Infrastructure	45
Satellite	45
Switched Multimegabit Data Service	46
Frame Relay	48
Asynchronous Transfer Mode	50
Comparison of Packet Switching Technologies	53
Integrated Services Digital Network	53
Operation	54
Custom ISDN and National ISDN	55
Connections	55

CHAPTER 3 ATM's Role in Multicasting	57
Describing ATM Services and Support	58
Connection-Oriented Support	59
Integrated Services Support	60
Quality-of-Service Support	62
Summary	66
ATM Networking	66
The End-to-end ATM Model	67
The ATM Desktop Model	68
The Campus Backbone Model	69
The WAN Model	70
ATM Carrier Model	71
IP Over ATM Requirements	73
ATM Network Services	75
ATM Client/Server Accessibility	76
ATM CIP Services	79
The Next Hop Resolution Protocol (NHRP)	81
ATM Forum Multiprotocol Over ATM (MPOA)	85
IP Switching and Multiprotocol Label Switching (MPLS)	88
Summary	90
What's Next	90
CHAPTER 4 Multicasting Overview	91
Introducing Multicasting	92
Transmission Types	94
Summary of Multicast Operation	95
IP Multicasting Uses and Benefits	99
Reducing Network Load	100
Uses	100
IP Multicasting within Workgroups	103
Benefits Check List	104
Introducing the Multicast Backbone	104
Understanding Multicasting Capabilities	106
Multicasting Capabilities on the Host	106
Multicasting Capabilities on Routers	107
Multicast Routing in the Internet Protocol	109
Considerations When Adding Multicast Support to the IP Layer	110
IP Multicast-Enabled Application Development	111
Multimedia Applications	111
Information Distribution Applications	112
Evaluation and Implementation	112
Recommendations From the IP Multicast Initiative	113
After Deciding to Implement IP Multicasting	113
A Checklist for Deploying IP Multicasting	113
The IP Multicast Initiative	115
Security and IP Multicasting	115
Differences Between IPv4 and IPv6 Multicast Addressing	116
CHAPTER 5 Multicast Operation	119
A Quick Description	120

Multicast Requirements	121
Industry Needs	121
Network Requirements to Support Multicasting	123
IPv4 Multicast Addressing	125
IPv4 Addressing Architecture	125
Impact of Multicast Addressing	128
Class D Address Format	129
Layer 3-to-Layer 2 Translation	132
Transmission and Delivery of Multicast Datagrams	134
IPv6 Multicast Addressing	135
Non-Permanently and Permanently Assigned Multicast Addresses	136
Predefined Well-Known Multicast Addresses	137
Building a Multicast Group	138
Overview of Multicast Forwarding Methods	140
Source-Rooted Tree Protocols	141
Shared Trees	143
Time-to-Live Setting	144
Role of Quality of Service and the Resource Reservation Protocol	146
Quality of Service	146
Integrated Services	150
CHAPTER 6 Under the Hood of IP Multicasting Protocols	155
Host Extensions for IP Multicasting	156
Introduction	156
Conformance	157
Host Group Addressing Issues	157
Sample Host IP Multicast Implementation	158
Sending Multicast IP Datagrams	158
Receiving Multicast IP Datagrams	161
Internet Group Management Protocol Basic Operation	165
IGMP Messages	166
IGMP Operation	171
Compatibility with IGMPv1 Routers	173
Compatibility with IGMPv1 Hosts	174
IGMP Version 3 in Progress	174
Host State Diagram	176
Events That Cause a Host State Transition	177
Responses to Events Causing a Host State Transition	177
Router State Diagrams	178
Router States With Respect to Any Single Attached Network	178
Router State with Respect to an IP Multicast Group on Any Single Attached Network	180
List of Timers and Default Values	183
CHAPTER 7 Multicast Routing Concepts and Algorithms	185
Introduction to Routing in an IP Environment	187
IP Routing	188
Routing Protocols	190
Routing Protocol Metrics	194
Multicast Routing	195
Multicast Routing Choices	196

Routing Standards	197
Basic Approaches to IP Multicast Routing	197
Multicast Algorithms	198
Flooding	198
Spanning Trees	199
Reverse-Path Broadcasting	203
Truncated Reverse-Path Broadcasting	205
Reverse-Path Multicasting (RPM)	207
Shortest-Path Trees	209
Core-Based Trees	209
CHAPTER 8 Multicast Routing Protocols	213
Group Membership Protocol	214
Sparse-Mode Multicast Routing	215
Dense-Mode Multicast Routing	215
Distance Vector Multicast Routing Protocol	215
Multicast Open Shortest-Path First	224
Protocol-Independent Multicast Protocols	235
PIM Operation	235
Protocol-Independent Multicast-Dense Mode	237
Protocol-Independent Multicast-Sparse Mode	238
Core-Based Trees (CBT)	240
CHAPTER 9 Implementing IP Multicasting	243
Reliability Issues	244
Joining and Leaving a Group	244
Join- and Leave-Group Operations	245
Join and Leave Group Actions	245
IP Tunnels and Multicasting	246
Truncated and Pruned Tunneling	247
Sending Datagrams Over Tunnels	248
Things to Note About Tunnels	248
Multicasting Methods	249
IPv4 Option for Sender-Directed Multidestination Delivery	249
Multicast Transport Protocol (MTP)	249
MARS-Based ATM Multicasting	250
CLNP Multicasting	250
Alternative Types of Network Architectures	251
Frame Relay	251
Switched Multimegabit Data Service	253
Satellite Networks	254
Multicasting and the Internet, Campus and LAN Infrastructures	256
Integrated Services Digital Network	257
CHAPTER 10 IP Multicasting Challenges	259
Multicasting Problems	260
Lack of Support	261
Interaction with Firewalling Strategies	262
Quality of Service	263
Putting Multicasting to Work	265
A Word About Mbone	265

The Reliable Multicast Transport Protocol (RMTP)	267
The Future of IP Multicasting	268
PART II The Mbone and Other Multicast Networks	
CHAPTER 11 An Overview of Asynchronous Transfer Mode Technology	273
ATM Technical Specifications	276
ATM's Mission: the Delivery of Multimedia Services	278
ATM Network Design and Solutions	279
ATM Network Client Models	279
ATM is Connection-Oriented	285
Basic Understanding of ATM Networks	287
CHAPTER 12 Understanding Mbone	295
Mbhone Background	296
Mbhone Fundamentals	296
Multicasting a Conference on Mbhone	297
Configuration Checklist	302
Security: Mbhone and Firewalls	304
CHAPTER 13 Higher-Level Protocols Used with IP Multicast	305
Real-time Transport Protocol	306
Real-time Control Protocol	308
Real-time Streaming Protocol	309
Properties of RTSP	310
Resource Reservation Protocol (RSVP)	311
Understanding RSVP	312
RSVP Challenges	315
Reliable Multicasting	315
Application and Network Characteristics Affecting Reliable Multicasting	315
Problems for Reliable Multicast Error Detection	316
Reliable Multicast Protocols	316
Multicast Implementation Study (MIST)	317
Reliable Multicasting Protocol	318
RMP Major Features	318
RMP Status	319
Reliable Multicast Framework Protocol (RMF)	320
Reliable Adaptive Multicast Protocol (RAMP)	321
RAMP Features	322
RAMP Operating Modes	322
Reliable Multicast Transport Protocol (RMTP)	323
How e-cast Reliable Multicast FTP Works	323
Advantages of e-cast RMFTP Over Multiple Unicast	324
CU-SeeMe	325
CHAPTER 14 Writing IP Multicast-Enabled Applications	327
Applications That Can Benefit From Multicast	328
Design Considerations	329
Implementing Multicast-Enabled Applications	329
Defining Network Layer Parameters	330
Sending and Receiving Data from Multicast Groups	330
Multicast Groups and Addresses	331

Time-To-Live (TTL) for Multicast Packets	331
Sending Multicast Datagrams	332
Receiving Multicast Datagrams	332
An Example of Multicast Programming	333
Application Requirements	335
Multicasting and Transport Protocols	335
Adapting Unicast Datagram Applications To Multicast	336
Things to Keep in Mind	337
CHAPTER 15 Standard Multicast Application Programming Interfaces	339
Typical APIs	340
Berkeley Sockets Multicast APIs	340
Windows Sockets Multicast APIs	343
Open Transport/TCP Multicast APIs	345
Reliable Multicast Protocol Application Programming Interface	346
RMP Fault Detection	346
Types and Objects	346
API Limitations	348
PART III IP Multicasting Applications	
CHAPTER 16 Major IP Multicasting Players	351
The ISPs	352
MCI Communications Corp.	352
GTE BBN Planet Corp.	354
UUNet Technologies Inc.	355
@Home Network Inc.	356
The Suppliers	357
Stardust Forums: Managing IP Multicasting Initiative	357
Novell's Initiative	357
3Com: Information Bus Publish/Subscribe Model	358
Lucent Technologies' e-cast Reliable Multicast Software	360
Microsoft's NetShow	361
Cisco/Precept Software's IP/TV	364
Starlight Networks	368
VDOnet	368
RealNetworks, Inc./VIVO Software	369
Layer 3 Switch Vendors	371
Satellite Providers	376
Wireless Vendors	380
CHAPTER 17 IP Multicasting Application Software/Client	383
EMULive Imaging Corporation's Active Theater	384
EMULive Server	385
ICAST Corporation's ICAST Viewer	385
IVS (INRIA Videoconferencing System)	386
Fantastic Corporation's MediaSurfer	386
Microsoft's NetShow Theater Server	388
Live Networks, Inc.'s Multikit	390
Precept IP/TV	391
Intel's Proshare	392
RealNetworks RealSystems G2	393

Cabletron's SecureFast SmartSwitch 9000	395
The SmartSwitch 9000	396
StarBurst Communications' StarBurst Multicast	398
University College London's Robust-Audio Tool (RAT)	400
New Features	401
Lawrence Berkeley National Laboratory's Vic	402
MMosaic	402
CHAPTER 18 IP Multicasting Application Software/Server	405
EMULive Imaging Corporation's EMULive Server	406
The Fantastic Corp.'s Channel Editing Center (CEC)	407
ICAST Broadcaster	409
Microsoft's NetShow Services	409
Probita's Proclaim Digital Distribution System	411
Starburst's Streamworks Server	412
CHAPTER 19 IP Multicasting Content Providers	415
PointCast Network	416
Broadcasting Benefits	418
PointCast: Not a Push Technology Product	418
Intranet Broadcast Solution	419
Kiosk Client	420
CHAPTER 20 IP Multicasting Middleware Software Products	421
Overview of StarBurst Multicasting	422
StarBurst Multicast Sender for Microsoft's SMS	423
StarBurst's MFTP	425
StarBurst OmniCast	425
Benefits	426
Features	426
Distributed Architecture	427
PeerLogic's PIPES1	429
Understanding PIPES	430
PIPES Platform and Messaging Communications	431
Technical Overview	432
The PIPES Platform API (PAPI)	440
A PIPES Platform Example	440
CHAPTER 21 IP Multicasting Network Management Tools	445
DWTNDA-Dr. Watson, The Network Detective's Assistant	446
Using Dr. Watson	447
Network Detective and the Toolbox	447
Features	451
FORE Systems' ForeThought	453
Features	454
PART IV IP Multicasting Resource Center	
CHAPTER 22 Implementing IP Multicast in Different Network Infrastructures	459
Introduction	460
Satellite Networks	461
Mechanisms for Enabling IP Multicast	461
Special Considerations	462

Frame Relay Networks	462
Mechanisms for Enabling IP Multicast	463
Special Considerations	465
Switched Multimegabit Data Service (SMDS)	465
Mechanisms for Enabling IP Multicast	466
Special Considerations	466
ATM Networks	467
Mechanisms for Enabling IP Multicast	467
Special Considerations	468
The Internet	469
Mechanisms for Enabling IP Multicast	469
Campus and Premise LAN Infrastructures	469
Mechanisms for Enabling IP Multicast	470
Special Considerations	470
Dial-up Connections	470
Special Considerations	470
ISDN (Integrated Services Digital Network)	471
Special Considerations	472
CHAPTER 23 Writing IP Multicast-enabled Applications	473
Applications That Can Benefit from Multicast	474
Application Requirements for Multicast-Enabling	474
Adapting Unicast Datagram Applications to Multicast	475
Standard Multicast APIs	477
Berkeley Sockets Multicast APIs	477
Windows Sockets Multicast APIs	478
Open Transport/TCP Multicast APIs	479
Multicasting With Care	479
Port and Address Considerations	480
Multicast API Limitations	481
Future Application Needs—Quality of Service	482
Code Samples	483
Application Requirements	484
CHAPTER 24 IP Multicasting Switch Equipment	493
Ascend Technologies	494
Ascend GRF IP Switch	498
Dynamic Switch Circuit Routing	501
ISDN and Videoconferencing	503
3Com CoreBuilder 3500 Layer 3 High-Function Switch	505
Cabletron SmartSwitch 9000	511
HP ProCurve Switch 8000M	517
APPENDIX A Bibliography	521
Books	522
RFCs522	525
Web	525
Glossaries	529
APPENDIX B List of Abbreviations and Acronyms	531
APPENDIX C Glossary	547
INDEX	589

PART I

**TCP/IP and
IP Multicasting
Fundamentals**