

- Part I The Contemporary Internet p. 3
- Chapter 1 Evolution of the Internet p. 5
  - Origins and Recent History of the Internet p. 5
  - From ARPANET to NSFNET p. 7
  - The Internet Today p. 8
  - NSFNET Solicitations p. 10
  - Network Access Points p. 10
  - What Is a NAP? p. 11
  - NAP Manager Solicitation p. 11
  - Federal Internet eXchange p. 12
  - Commercial Internet eXchange p. 12
  - Current Physical Configurations at the NAP p. 13
  - An Alternative to NAPs: Direct Interconnections p. 14
  - Routing Arbiter Project p. 14
  - The Very High Speed Backbone Network Service p. 18
  - Transitioning the Regional Networks from the NSFNET p. 21
  - NSF Solicits NIS Managers p. 22
  - Network Information Services p. 23
  - Creation of the InterNIC p. 23
  - Directory and Database Services p. 23
  - Registration Services p. 25
  - NIC Support Services p. 25
  - Other Internet Registries p. 25
    - ARIN p. 26
    - RIPE NCC p. 26
    - APNIC p. 27
  - Internet Routing Registries p. 27
  - The Once and Future Internet p. 28
  - Next-Generation Internet Initiative p. 28
  - Internet2 p. 30
  - Abilene p. 31
  - Looking Ahead p. 32
  - Frequently Asked Questions p. 34
  - References p. 35
- Chapter 2 ISP Services and Characteristics p. 37
  - ISP Services p. 37
  - Dedicated Internet Access p. 37
  - Frame Relay and ATM Internet Access p. 38
  - Dialup Services p. 39
  - Digital Subscriber Line p. 40
  - Cable Modems p. 41
  - Dedicated Hosting Services p. 41
  - Other ISP Services p. 42
  - ISP Service Pricing, Service-Level Agreements, and Technical Characteristics p. 42
  - ISP Service Pricing p. 42
  - Service-Level Agreements p. 43

- ISP Backbone Selection Criteria p. 43
- Demarcation Point p. 50
- Looking Ahead p. 53
- Frequently Asked Questions p. 54
- Chapter 3 IP Addressing and Allocation Techniques p. 57
  - History of Internet Addressing p. 57
  - Basic IP Addressing p. 57
  - Basic IP Subnetting p. 60
  - VLSMs p. 62
  - IP Address Space Depletion p. 65
  - IP Address Allocation p. 66
  - Classless Interdomain Routing p. 67
  - Private Addressing and Network Address Translation p. 79
  - IP Version 6 p. 82
  - Looking Ahead p. 86
  - Frequently Asked Questions p. 87
  - References p. 89
- Part II Routing Protocol Basics p. 91
- Chapter 4 Interdomain Routing Basics p. 93
  - Overview of Routers and Routing p. 93
  - Basic Routing Example p. 94
  - Routing Protocol Concepts p. 96
  - Distance Vector Routing Protocols p. 96
  - Link-State Routing Protocols p. 99
  - Segregating the World into Autonomous Systems p. 101
  - Static Routing, Default Routing, and Dynamic Routing p. 101
  - Autonomous Systems p. 102
  - Looking Ahead p. 107
  - Frequently Asked Questions p. 108
  - References p. 109
- Chapter 5 Border Gateway Protocol Version 4 p. 111
  - How BGP Works p. 112
  - BGP Message Header Format p. 115
  - BGP Neighbor Negotiation p. 116
  - Finite State Machine Perspective p. 118
  - Notification Message p. 120
  - Keepalive Message p. 122
  - Update Message and Routing Information p. 122
  - BGP Capabilities Negotiation p. 127
  - Multiprotocol Extensions for BGP p. 128
  - TCP MD5 Signature Option p. 129
  - Looking Ahead p. 131
  - Frequently Asked Questions p. 132
  - References p. 133
- Part III Effective Internet Routing Designs p. 135
- Chapter 6 Chapter Tuning BGP Capabilities p. 137

- Building Peer Sessions p. 137
- Physical Versus Logical Connections p. 139
- Obtaining an IP Address p. 140
- Authenticating the BGP Session p. 140
- BGP Continuity Inside an AS p. 141
- Synchronization Within an AS p. 142
- Sources of Routing Updates p. 144
- Injecting Information Dynamically into BGP p. 144
- Injecting Information Statically into BGP p. 147
- Origin of Routes p. 148
- An Example of Static Versus Dynamic Routing: Mobile Networks p. 150
- Overlapping Protocols: Backdoors p. 150
- The Routing Process Simplified p. 152
- BGP Routes: Advertisement and Storage p. 153
- The BGP Routing Information Bases p. 154
- Routes Received from Peers p. 155
- Input Policy Engine p. 155
- Routes Used by the Router p. 155
- Output Policy Engine p. 156
- Routes Advertised to Peers p. 156
- Sample Routing Environment p. 156
- BGP Decision Process Summary p. 158
- Controlling BGP Routes p. 159
- BGP Path Attributes p. 160
- NEXT\_HOP Behavior on Multiaccess Media p. 172
- NEXT\_HOP Behavior Over Nonbroadcast Multiaccess Media p. 173
- Use of next-hop-self versus Advertising DMZ p. 174
- Using Private ASs p. 175
- AS\_PATH and Route Aggregation Issues p. 177
- AS\_PATH Manipulation p. 178
- Route Filtering and Attribute Manipulation p. 180
- Inbound and Outbound Filtering p. 181
- The Route Filtering and Manipulation Process p. 182
- Peer Groups p. 190
- BGP-4 Aggregation p. 192
- Aggregate Only, Suppressing the More-Specific Routes p. 192
- Aggregate Plus More-Specific Routes p. 193
- Aggregate with a Subset of the More-Specific Routes p. 195
- Loss of Information Inside Aggregates p. 196
- Changing the Attributes of the Aggregate p. 196
- Forming the Aggregate Based on a Subset of the More-Specific Routes p. 196
- Looking Ahead p. 197
- Frequently Asked Questions p. 199
- References p. 201
- Chapter 7 Redundancy, Symmetry, and Load Balancing p. 203
- Redundancy p. 203

- Geographical Restrictions Pressure p. 204
- Setting Default Routes p. 205
- Symmetry p. 210
- Load Balancing p. 210
- Specific Scenarios: Designing Redundancy, Symmetry, and Load Balancing p. 212
- Scenario 1 Single-Homing p. 213
- Scenario 2 Multihoming to a Single Provider p. 213
- Scenario 3 Multihoming to Different Providers p. 223
- Scenario 4 Customers of the Same Provider with a Backup Link p. 228
- Scenario 5 Customers of Different Providers with a Backup Link p. 231
- Looking Ahead p. 236
- Frequently Asked Questions p. 237
- References p. 239
- Chapter 8 Controlling Routing Inside the Autonomous System p. 241
- Interaction of Non-BGP Routers with BGP Routers p. 241
- Injecting BGP into the IGP p. 241
- Following Defaults Inside an AS p. 242
- BGP Policies Conflicting with Internal Defaults p. 244
- Defaults Inside the AS: Primary/Backup BGP Policy p. 244
- Defaults Inside the AS: Other BGP Policies p. 250
- Policy Routing p. 252
- Policy Routing Based on Traffic Source p. 252
- Policy Routing Based on Traffic Source/Destination p. 253
- Policy Routing Defaults to Dynamic Routing p. 254
- Other Applications of Policy Routing p. 255
- Looking Ahead p. 257
- Frequently Asked Questions p. 258
- Chapter 9 Controlling Large-Scale Autonomous Systems p. 261
- Route Reflectors p. 261
- Internal Peers Without Route Reflectors p. 262
- Internal Peers with Route Reflectors p. 263
- Naming Conventions and Rules of Operation p. 264
- Redundancy Issues and Multiple Route Reflectors in an AS p. 265
- Route Reflection Topology Models p. 266
- Route Reflectors and Peer Groups p. 269
- Confederations p. 271
- Confederation Drawbacks p. 273
- Route Exchange and BGP Decisions with Confederations p. 274
- Recommended Confederation Design p. 274
- Confederations Versus Route Reflectors p. 275
- Controlling IGP Expansion p. 275
- Segmenting the AS with Multiple Regions Separated by IBGP p. 277
- Segmenting the AS with Multiple Regions Separated by EBGp p. 279
- Looking Ahead p. 283
- Frequently Asked Questions p. 284
- References p. 285

- Chapter 10 Designing Stable Internets p. 287
- Route Instabilities on the Internet p. 287
- IGP Instability p. 287
- Faulty Hardware p. 288
- Software Problems p. 288
- Insufficient CPU Power p. 288
- Insufficient Memory p. 289
- Network Upgrades and Routine Maintenance p. 289
- Human Error p. 290
- Link Congestion p. 290
- BGP Stability Features p. 290
- Controlling Route and Cache Invalidation p. 291
- BGP Route Refresh p. 291
- Route Dampening p. 292
- Looking Ahead p. 296
- Frequently Asked Questions p. 297
- Part IV Internet Routing Device Configuration p. 299
- Chapter 11 Configuring Basic BGP Functions and Attributes p. 301
- Building Peering Sessions p. 301
- Route Filtering and Attribute Manipulation p. 308
- BGP Route Maps p. 308
- Prefix Lists p. 310
- Identifying and Filtering Routes Based on the NLRI p. 312
- Identifying and Filtering Routes Based on the AS\_PATH p. 315
- Peer Groups p. 316
- Sources of Routing Updates p. 318
- Injecting Information Dynamically into BGP p. 318
- Injecting Information Statically into BGP p. 325
- Overlapping Protocols: Backdoors p. 326
- BGP Attributes p. 328
- The NEXT\_HOP Attribute p. 331
- The AS\_PATH Attribute p. 332
- The LOCAL\_PREF Attribute p. 335
- The MULTI\_EXIT\_DISC Attribute p. 337
- The COMMUNITY Attribute p. 340
- BGP-4 Aggregation p. 342
- Aggregate Only, Suppressing the More-Specific p. 343
- Aggregate Plus More-Specific Routes p. 346
- Aggregate with a Subset of the More-Specific Routes p. 350
- Loss of Information Inside Aggregates p. 354
- Changing the Aggregate's Attributes p. 357
- Forming the Aggregate Based on a Subset of Specific Routes p. 359
- Looking Ahead p. 361
- Chapter 12 Configuring Effective Internet Routing Policies p. 365
- Redundancy, Symmetry, and Load Balancing p. 365
- Dynamically Learned Defaults p. 365

- Statically Set Defaults p. 367
- Multihoming to a Single Provider p. 370
- Multihoming to Different Providers p. 384
- Customers of the Same Provider with a Backup Link p. 388
- Customers of Different Providers with a Backup Link p. 391
- Following Defaults Inside an AS p. 395
- BGP Policies Conflicting with the Internal Default p. 398
- Policy Routing p. 411
- Route Reflectors p. 415
- Confederations p. 419
- Controlling Route and Cache Invalidation p. 424
- BGP Soft Reconfiguration p. 425
- Outbound Soft Reconfiguration p. 425
- Inbound Soft Reconfiguration p. 425
- BGP Route Refresh p. 429
- BGP Outbound Request Filter Capability p. 431
- Route Dampening p. 432
- Looking Ahead p. 435
- Part V Appendixes p. 439
- Appendix A BGP Command Reference p. 441
- Appendix B References for Further Study p. 449
  - Interesting Organizations p. 449
  - Research and Education p. 449
  - Miscellaneous p. 449
  - Books p. 450
  - TCP/IP-Related Sources p. 450
  - Routing-Related Sources p. 450
  - Internet Request For Comments p. 450
- Appendix C BGP Outbound Route Filter (ORF) p. 455
  - When to Use BGP ORF p. 455
  - Configuration p. 456
  - Enabling the BGP ORF Capability as Send-Mode p. 456
  - Enabling the BGP ORF Capability as Receive-Mode p. 456
  - Ensuring Backward Compatibility of the Old Knobs p. 457
  - EXEC Commands p. 457
  - Pushing Out A Prefix List and Receiving a Route Refresh from a Neighbor p. 457
  - Displaying the Prefix List Received from a Neighbor p. 458
  - Displaying Changes to the Neighbor BGP Table p. 458
  - Closing Remarks p. 458
- Appendix D Multiprotocol BGP (MBGP) p. 461
  - The Motivation Behind the New Command-Line Interface p. 461
  - Organizing Command Groups in the New Configuration p. 462
  - Activate p. 464
  - Old Style p. 464
  - AF Style p. 464
  - Network p. 465

- Old Style p. 465
- AF Style p. 465
- Peer Groups p. 465
- Old Style p. 466
- AF Style p. 466
- Route Maps p. 466
- Old Style p. 466
- AF Style p. 467
- Redistribution p. 468
- Old Style p. 468
- AF Style p. 468
- Route Reflector p. 469
- Old Style p. 469
- AF Style p. 469
- Aggregation p. 469
- Old Style p. 470
- AF Style p. 470
- List of BGP Commands p. 470
- Upgrading to the AF Style p. 472
- References p. 473
- Index p. 475