


Figure 2-7 *Label Switching with the MPLS Label Stack*



MPLS and VPN Architectures

ISBN: 1-58705-002-1

Copyright © 2001 Cisco Press

Figure 2-8 *Double Lookup in New York POP Router*




Figure 2-9 *Penultimate Hop Popping in the SuperNet Network*




Figure 5-2 *Path MTU Discovery Mechanism*




Figure 5-6 *Path-Vector TLV Loop Prevention Mechanism*




Figure 5-7 *Hop-count Propagation Between ATM-LSRs*




Figure 7-7 *Sample Peer-to-peer VPN*





Figure 7-8 *Peer-to-peer VPN Model: Shared Router Configuration*



MPLS and VPN Architectures

ISBN: 1-58705-002-1

Copyright © 2001 Cisco Press

Figure 7-9 *Peer-to-peer VPN Model: Dedicated Router Configuration*




Figure 9-6 *SOO BGP Extended Community*






Figure 9-7 *BGP Extended Community Attribute Format*



High-order byte of the type field 0x00


Administrator subfield: 2 Octets (AS #)
Assigned Number subfield: 4 Octets



High-order byte of the type field 0x01

Administrator subfield: 4 Octets (IP Address)
Assigned Number subfield: 2 Octets

Figure 12-12 *Route Reflection and ORF Capability*



MPLS and VPN Architectures

ISBN: 1-58705-002-1

Copyright © 2001 Cisco Press

Figure 12-14 \BGP Confederations—Single IGP Environment




Figure 12-15 *BGP Confederations—Multiple IGP Environment*





Figure 12-16 *Resetting of Next-hop at Subconfederation Boundary*



MPLS and VPN Architectures

ISBN: 1-58705-002-1

Copyright © 2001 Cisco Press

Figure 12-17 *Label Exchange Using MP-eBGP and next-hop-self*





Figure 12-25 *Internet Connectivity Through Multiple Exit Points*




MPLS and VPN Architectures

ISBN: 1-58705-002-1

Copyright © 2001 Cisco Press

Figure 13-6 *MPLS Two-Level Label Stack*



MPLS and VPN Architectures

ISBN: 1-58705-002-1

Copyright © 2001 Cisco Press

Figure 13-12 *Route Filtering Using Standard BGP Communities*




Figure 13-13 *Export of Routes with Different Route Targets*




Figure 14-5 *ISP External Traffic—Routing Requirement at PE-router*





Figure 14-6 *Carrier's Carrier—No MPLS Within ISP Sites (Label Assignment)*




MPLS and VPN Architectures

ISBN: 1-58705-002-1

Copyright © 2001 Cisco Press

Figure 14-9 *Local MPLS Deployment—No Label Distribution Between Sites (Label Assignment)*




MPLS and VPN Architectures

ISBN: 1-58705-002-1

Copyright © 2001 Cisco Press

Figure 14-12 *Hierarchical VPN Connectivity (Label Assignment)*




MPLS and VPN Architectures

ISBN: 1-58705-002-1

Copyright © 2001 Cisco Press

Figure 14-16 *Exchange of VPN-IPv4 Routes Between Service Providers (Label Allocation)*




MPLS and VPN Architectures

ISBN: 1-58705-002-1

Copyright © 2001 Cisco Press

Figure 14-19 *Inter-provider VPN—Multihop eBGP Between Customer Sites (Label Allocation)*



MPLS and VPN Architectures

ISBN: 1-58705-002-1

Copyright © 2001 Cisco Press

Tag-switching and MPLS Command Reference

In certain versions of IOS, a number of configuration commands have both MPLS and tag-switching forms and will be supported during the transition from a tag-switching environment to a standards-based MPLS environment. We have used predominantly the tag-switching commands within this publication, but the following table provides a reference which documents some of the more common tag-switching commands and their MPLS equivalent form.

Table A-1 *Tag-switching Versus MPLS Command Structure*

LDP Command	TDP Command	Description
mpls ldp advertise-labels	tag-switching advertise-labels	Controls the distribution of locally assigned (incoming) labels by LDP/TDP
mpls ldp atm control-mode	tag-switching atm allocation-mode	Controls the mode used for handling label-binding requests on LC-ATM interfaces
mpls ldp atm vc-merge	tag-switching atm vc-merge	Controls whether the vc-merge capability is supported for unicast label VCs
mpls ldp maxhops	tag-switching atm maxhops	Limits the number of hops permitted in an LSP established by the downstream on demand method of label distribution
mpls ip (global)	tag-switching ip (global)	Enables MPLS forwarding of IPv4 packets along normally routed paths for the platform
mpls ip (interface)	tag-switching (interface)	Enables MPLS forwarding of IPv4 packets along normally routed paths for a particular interface

Table A-1 *Tag-switching Versus MPLS Command Structure (Continued)*

LDP Command	TDP Command	Description
mpls ldp discovery	tag-switching tdp discovery	Configures the interval between transmission of consecutive LDP/TDP discovery hello messages, the hold time for a discovered LDP/TDP neighbor, or the neighbors from which requests for targeted hellos may be honored
mpls ldp holdtime	tag-switching tdp holdtime	Changes the time for which an LDP/TDP session is maintained in the absence of LDP/TDP messages from the session peer
show mpls atm-ldp bindings	show tag-switching atm-tdp bindings	Displays specified entries from the ATM LDP/TDP label-binding database
show mpls atm-ldp capability	show tag-switching atm-tdp capability	Displays the ATM MPLS capabilities negotiated with LDP/TDP neighbors for LC-ATM interfaces
show mpls interfaces	show tag-switching interfaces	Displays information about one or more interfaces that have been configured for label switching
show mpls ldp bindings	show tag-switching tdp bindings	Displays the contents of the label information base (LIB)
show mpls ldp discovery	show tag-switching tdp discovery	Displays the status of the LDP/TDP discovery process
show mpls ldp neighbor	show tag-switching tdp neighbor	Displays the status of LDP/TDP sessions
show mpls ldp parameters	show tag-switching tdp parameters	Displays current LDP/TDP parameters

Numerics

802.1q VLANs, 337

A

address families, 191
 routing context, configuring, 193–194
 VPN-IPv4, configuring, 192
 address space, overlapping, 148
 advertisements
 AS-override, 180
 Extended Community attribute, 192–193
 PE routers, requirements for, 167
 aggregates, 29
 penultimate hop popping, 41
 allocating labels
 across ATM-LSR domains, 56–61
 distribution methods, 34–36
 independent control, 57
 VC merge, 58–61
 AllowAS-in feature, deploying, 245–247
 any-to-any connectivity
 intranet model, 168–169
 extranet model, 138
 application hosting, 240
 applications, control protocols, 19–20
 architecture
 Cell-mode MPLS, control plane, 51–53
 Edge-LSRs, 18
 MPLS/VPN, 13, 145–146
 Extended Communities ORF-type
 capability, 203
 FECs, 16
 inter-PE router routing, 156–157
 intranet topology, 168–169
 overlapping addresses, 174
 route distinguishers, 157
 route targets, 155–156
 scalability, 157
 sites, 153
 SOO, 177, 180
 area 0 (OSPF), PE-to-CE connectivity, 213,
 220–224

arguments, tag-switching advertise-tags
 command, 75
 AS Override, 232–234
 ASN (Autonomous System Number), assigning to
 route distinguishers, 175
 assigning
 interfaces to VRFs, 188
 labels. See bindings
 loopback addresses to PE-routers, 328
 route distinguishers to VRFs, 391–392
 unique numbers to route distinguishers,
 175–176
 associating VRF with global routing table, 310–311, 314
 ATM (Asynchronous Transfer Mode)
 backbone networks
 migration to Cell-mode MPLS, 108–110
 migration to Frame-mode MPLS,
 106–108
 Cell-mode MPLS, 49, 69
 control-plane connectivity, 52–53
 limitation, 50
 Frame-mode MPLS, 67–69
 heterogeneous MPLS-mode operation, 69
 Layer 3 lookup, 50
 LSRs, 13–15
 convergence, 61–62
 domains, 50
 Edge-LSRs, 13–15, 50
 IP connectivity, 51
 labeled packet forwarding, 55–56
 loopback interfaces, 53
 service provider backbone, convergence, 254
 switches
 control-plane implementation, 54
 ordered control, 57
 VCs, 49
 attributes (BGP), 171
 community, removing, 351
 Extended Community, 177
 advertising, 192–193
 format, 182
 OSPF route propagation, 220
 route target, 178–179
 MP_REACH_NLRI, 189
 MP_UNREACH_NLRI, 189

Automatic Route Filtering, 245
 PE-routers, 197
avoiding
 routing loops in VPNs, 247
 summarization on PE-routers, 332

B

backbone
 migration to Frame-mode MPLS, 97, 99
 MPLS/VPN
 customer access, 207–208
 deploying BGP, 324–329, 331–333, 335
 OSPF route propagation, 217–220
 route advertisements, 260–265
 traceroute, 351, 353–354
 reachability, 322
 service provider, convergence, 253–254
best path, selecting, 271
BGP (Border Gateway Protocol), 257
 attributes, 171
 community attribute, removing, 351
 confederations
 deploying, 280–283, 285
 multiple IGP environment, 286–289
 single IGP environment, 285–286
 deploying in MPLS/VPN backbone, 324–329,
 331–335
 export policies, configuring, 179–180
 extended community attribute, 177
 format, 182
 OSPF route propagation, 220
 route target, 178–179
 import policies, configuring, 179–180
 internal sessions, 188
 label bindings, 34
 next-hops
 forwarding, 325–328
 summarizing, 332
 OPEN messages, Capabilities parameter, 189
 PE-to-CE connectivity, 209–211
 peer groups, leader selection, 325
 peering sessions, scaling, 325
 peering structure, 101–103
 route reflectors, hierarchical structure, 268–269
 Route Refresh, 199, 201

 scanner process, 258–259
 selection process, 244
 standard community extension, filtering on PE
 routers, 274–277
 unnecessary sessions, removing, 105
bindings, 12, 34
 conservative retention mode, 36
 liberal retention mode, 253
Bottom-of-Stack bit, 27
BPX switches, LSC (Label Switch Controller), 54
business problem-based VPN classification,
 118–120

C

calculating VCs per full-mesh networks, 133
Capabilities parameter (OPEN messages), 189
carrier's carrier MPLS/VPN, 358–360
 routes, 359
 topologies, 361–372
case studies
 intranet service, 168–169
 migration
 ATM backbone to Cell-mode MPLS,
 108–110
 ATM backbone to Frame-mode MPLS,
 97–99, 106–107
 CEF requirements, 100
 internal links to MPLS, 103–105
 premigration infrastructure checks,
 99–100
 SuperCom network, VPNs, 146–148
 VRF configuration, 170
CE (Customer Edge) devices, 117
 links, troubleshooting, 326
 routers
 PE-to-CE connectivity, provisioning,
 291–292
 tunneling to PE-routers, 335
CEF switching, assigning interfaces to VRFs, 188
Cell-mode MPLS, 49
 ATM-LSRs, 50
 loopback interfaces, 53
control-plane
 connectivity, 52–53
 implementing in ATM switches, 54

- IP connectivity, 51
- loop prevention, 82–87
- migration from ATM backbone networks, 108–110
- over ATM, 69
- central services topology, 137, 240–242
- central sites
 - extranet VPN connectivity, 238
 - IP tunneling, migration to VPN/MPLS, 395–398
- Cisco MPLS, 79
- Cisco Secure Integrated Software, 295
- Cisco VPN Solutions Centre, 338
- classification of VPNs
 - business problem-based, 118–120
 - implementation models, 121–128
 - topologies, 129–141
- class-of-service information, 26
- clear ip bgp command, 201
- clusters, route reflection between PE-routers, 272
- commands
 - clear ip bgp, 201
 - debug tag tdp transport, 31
 - ip ospf network point-to-point, 328
 - ip route, 185
 - ip vrf, 170
 - ip vrf forwarding, 188
 - maximum routes, 291
 - mpls ip, 31
 - mpls label-distribution both, 31
 - neighbor, 192, 327
 - neighbor advertisement-interval, 257
 - neighbor maximum-prefix, 291
 - network, 186
 - next-hop self, 324, 326
 - next-hop-self, 289
 - propagate-ttl, 353
 - redistribute, 185, 217
 - redistribute bgp metric transparent, 187
 - route-maps, 185
 - router ospf, 215
 - route-target, 179
 - route-target export, 351
 - rr-group, 203
 - set extcommunity rt, 351
 - show atm vc, 54
 - show ip bgp neighbor, 200, 263–264, 325
 - show ip ospf, 216
 - show tag tdp neighbor, 33
 - show tag-switching forwarding-table, 204
 - show tag-switching tdp bindings, 36
 - show tag-switching tdp neighbor, 67, 75
 - tag tdp router-id, 33
 - tag-switching, MPLS equivalents, 405–406
 - tag-switching advertise-tags, 75–76
 - tag-switching ip, 31
 - tag-switching mtu, 80
- common services VPNs, route distinguishers, 175
- community attributes (BGP)
 - comparing Standard and Extended, 344–346
 - removing, 351
- comparing
 - Extended and Standard BGP Communities, 344–346
 - IP routing and IP forwarding tables, 150
 - VPN peer-to-peer models, 128
- complexity of modern VPNs, 118
- confederations
 - configuring, 282–283, 285
 - deploying, 280–283, 285
 - multiple IGP environment, 286–289
 - single IGP environment, 285–286
 - IGP, designing, 282
- configuring
 - AS Override, 233–234
 - ATM-LSRs, loopback interfaces, 53
 - BGP, import/export policies, 179–180
 - confederations, 282–283, 285
 - Internet access, 292–293
 - dynamic default routing, 308–314
 - on separate service provider, 315
 - performing additional global routing table lookup, 314
 - PE-to-CE connectivity, 300–308
 - static default routing, 295–300
 - MP-iBGP, 191–193
 - PE-routers
 - loopback addresses, 328
 - MP-iBGP, 400
 - summarization on loopback addresses, 329–334
 - PE-CE links, 184
 - route distinguishers, 175–176

- route reflectors, 203
 - ORF capability, 279–280
 - SOO Extended Community, 181–182
 - VRFs, 170
 - import/export policies, 179–180
 - connectivity, 394
 - Cell-mode MPLS, control-plane, 51
 - central services topology (VPNs), requirements, 240
 - control-plane in Cisco IOS software, 52–53
 - hierarchical VPNs, 372–374
 - Internet access
 - configuring, 292–293
 - dynamic default routing, 308–314
 - firewalls, 293–295
 - performing additional global routing table lookup, 314
 - static default routing, 295–300
 - through separate service provider, 315
 - inter-provider VPNs, 374–383
 - leased lines, 116
 - PE-to-CE, 207–211
 - configuring, 184–187
 - convergence, 256–257
 - links, troubleshooting, 326
 - OSPF, 212–213, 220–224
 - OSPF without site area 0, 224–225
 - provisioning, 291–292
 - redistribution, 208
 - running separate BGP sessions, 300–301
 - tunneling, 335
 - VPN customer network topologies, 226–230
 - VRFs, separating, 214–217
 - SuperCom network example, requirements, 151–153
 - conservative retention mode, 36
 - control modes, 73
 - control-plane, 12
 - Cell-mode MPLS, IP connectivity, 51
 - Cisco IOS software, 52–53
 - implementing in ATM switches, 54
 - control protocols
 - MPLS applications, 19–20
 - MPLSC, 27
 - convergence, 251
 - ATM-LSR domains, 61–62
 - between VPN sites, 255–260
 - Frame-mode MPLS, 37–39
 - within service provider backbone, 253–254
 - Core LSRs, 25
 - core redundancy, WANs, 7
 - CPE (Customer Premises Equipment) devices, 117
 - VCs, Frame-mode MPLS operation over Frame Relay, 65–67
 - customer links
 - network management, 338–347
 - export maps, 348–351
 - importing selected customer routes, 339, 341–343
 - route filtering, 343–347
 - customer routes
 - advertising, requirements, 167
 - IGP, migrating to BGP, 319–324
 - next-hop forwarding, 325–328
- ## D
-
- datagrams, Maximum Initially Labeled IP Datagram Size, 79
 - data-plane, 12
 - Cell-mode, loop prevention, 86–88
 - packet forwarding, 203–205
 - debug tag tdp transport command, 31
 - dedicated router approach, peer-to-peer VPN model, 126–127
 - default behavior, ATM service provider backbone, 254
 - defining VPNs, IP tunneling migration case study, 390
 - delays, convergence over ATM-LSR domains, 61–62
 - deploying
 - AllowAS-in, 245, 247
 - BGP in MPLS/VPN backbone, 324–329, 331–335
 - confederations, 280–285
 - multiple IGP environment, 286, 288–289
 - single IGP environment, 285–286
 - MPLS/VPN on LAN interfaces, 335, 337
 - PBR, 10

- designing Frame-mode MPLS domains,
 - convergence, 37–39
- detecting routing loops
 - Cell-mode MPLS, 82–88
 - Frame-mode MPLS, 81–82
- devices
 - CE (Customer Edge), 117
 - P-devices, 117
 - PE (Provider Edge), 117
 - routers, LSRs, 13
- dial-up lines, 115–116
- differentiated services, IP packet forwarding, 9
- distribution (labels)
 - across ATM-LSR domains, 56–61
 - downstream on-demand, 57
 - on-demand, 35
 - unsolicited, 35
- domains, ATM-LSRs
 - convergence, 61–62
 - Frame-mode, convergence, 37–39
 - label allocation and distribution, 56–61
 - labeled packet forwarding, 55–56
- down-bit, summary LSAs, 223
- downstream label distribution
 - unsolicited, 35
 - on-demand, 57
- draft-eitf-mpls-framework document, 11
- dynamic default routing, configuring Internet access, 308–314

E

- Edge-LSRs, 13–14
 - actions, 15
 - architecture, 18
 - egress, 17
 - ingress, 25
 - push action, 15–16
- EIGRP (Enhanced Interior Gateway Routing Protocol), equal-cost load sharing, 6
- encapsulation, 27–28
 - across Ethernet links, 77
- end-to-end convergence delay between VPN sites, 255
- equal-cost load sharing (EIGRP), 6

- establishing LDP/TDP sessions, 31, 33–34
- Ethernet
 - MPLS encapsulation, 77
 - PE-routers, tunneling to CE routers, 335
- evolution of VPNs, 115–117
- examples
 - migration
 - ATM backbone to Cell-mode MPLS, 108
 - ATM backbone to Frame-mode MPLS, 106–108
 - backbone network to Frame-mode MPLS, 97–99
 - CEF requirements, 100
 - internal links to MPLS, 103–104
 - premigration infrastructure checks, 99–100
 - PE-to-CE connectivity, 208
 - SuperCom network, connectivity requirements, 151–153
 - SuperNet, loopback addresses, 24
- experimental bits, 27
- export maps, 348–351
- export policies (BGP)
 - configuring, 179–180
 - IP tunneling migration case study, 393
 - route target, 178–179
- exporting routes to specific targets, 347–351
- Extended Communities attribute
 - advertising, 192–193
 - comparing to standard community attribute, 344–346
 - format, 182
 - OPSF route propagation, 220
 - route targets, 178–179
 - SOO, 180–182
- Extended Communities ORF-type, 203
- Extended Community attribute (BGP), 177
- external controllers, LSC (Label Switch Controller), 54
- external routes, 359
 - information propagation, 10
- extranet topology, 135–137
 - central services, 137–139
 - intranet integration, 237–238
 - route distinguishers, 175

F

- Fast Ethernet, mapping VLANs to VPNs, 336–337
- FECs (Forwarding Equivalence Classes), 16
 - label binding, 74
 - unicast routing, 34
- filtering
 - customer links, 343–347
 - label mappings, 74–77
 - ORF, 201
 - VPN-IPv4 routes, 197
- firewalls
 - Cisco Secure Integrated software, 295
 - Internet access, 293–295
- flexibility of IP forwarding, 5
- format, extended community attribute (BGP), 182
- forwarding
 - packets
 - label imposition, 16
 - label swapping, 12
 - traffic, 10
 - VRF, 150
- fragmenting packets, Path MTU discovery, 78
- Frame-based LSRs, 50
- Frame-mode MPLS, 65–67
 - convergence, 37–39
 - label lookup, 28–29
 - loop detection
 - control plane, 82
 - data plane, 81
 - LSR capabilities, 29
 - migrating from backbone network, 97–99
 - migration from ATM backbone networks, 106–108
- full mesh topology
 - route reflectors, 270
 - VCs, calculating, 133
 - VPNs, provisioning, 134

G

- Giants, fragmentation, 78–80
- Gigabit Ethernet, restrictions on MPLS
 - encapsulation, 77
- global routing table
 - associating with VRF, 310–311
 - performing additional lookup, 314
- goal of MPLS, 11
- GRE tunneling, migration to MPLS/VPN solution, 387–389, 393, 400
 - configuring route reflectors, 398–399
 - defining VPNs, 390
 - routing policies, 390, 392–394
 - staging and execution, 394–398
 - VPN site migration, 401
 - VRF definitions, 391–392

H

- heterogeneous MPLS-modes, over ATM, 69
- hierarchical structure
 - route reflectors, 268–271
 - VPNs, 372–374
- hop-by-hop destination-based unicast routing, 5
- hub-and-spoke topology
 - AllowAS-in feature, 245–247
 - route distinguishers, 175
 - SOO (Site of Origin) attribute, 246
 - VPNs, 129–131, 242–244
- hybrid topology, overlay model, 134–135

I

- iBGP
 - migrating to MPLS/VPN, 230–231
 - peering sessions, scaling, 325
- IETF (Internet Engineering Task Force)
 - draft-eitf-mpls-framework document, 11
 - Web site, 177

IGPs (Interior Gateway Protocols)

confederations

multiple process configuration, 286–289

single process configuration, 285–286

customer routes, migrating to BGP, 319–320, 322–324

implementation models (VPNs)

overlay model, 121–123

peer-to-peer model, 123–125

dedicated router approach, 126–127

shared router approach, 125–126

implementing

BGP confederations, 280–285

multiple IGP environment, 286–289

single IGP environment, 285–286

control-plane in ATM switches, 54

Multiprotocol BGP in SuperCom network, example, 158–161

implicit-null value, 41–42

import policies (BGP)

configuring, 179–180

IP tunneling migration, case study, 393

route target, 178–179

importing selected customer routes, 339–343

in-band management virtual circuit, Cell-mode MPLS, 51

independent control mode, 57, 73, 254

Frame-mode MPLS over Frame Relay, 65–67

label assignment, 35

infrastructure of VPDNs, 140

Ingress Edge-LSRs, 25

interfaces

assigning to VRFs, 188

LC-ATM

control-plane connectivity, 51

downstream on demand, 57

status, verifying, 53

loopback, route summarization, 333

partitioning, 335

internal BGP sessions, 188

internal links, migration to MPLS, 103–104

internal routes, 359

Internet access

configuring, 292–293

dynamic default routing, 308–314

global routing table, additional lookup, 314

PE-to-CE connectivity, 300–308

firewalls, 293–295

global routing table, associating with VRF, 310–314

separate service provider, configuring, 315 static default routing, 295–300

inter-PE-router routing, 156–157

inter-provider VPN solutions, 374–383

initiating packet labeling, 31

Intranet topology, MPLS/VPN, 168–169

route distinguishers, 175

extranet integration, 237–238

IOS software, control-plane connectivity, 52–53

IP addressing

connectivity between ATM-LSRs, 51

overlapping address space, 148

packet forwarding, 5

differentiated services, 9

external routing information propagation, 10 VRP (VPN routing and forwarding), 150

Path MTU discovery, 78–80

route summarization, 92–93

tunneling, migration to MPLS/VPN solution, 387–389, 393, 400

configuring route reflectors, 398–399

defining VPNs, 390

routing policies, 392–394

staging and execution, 394–398

VPN site migration, 401

VRF definitions, 391–392

VPDNs, 140

ip ospf network point-to-point command, 328

ip route command, 185

ip vrf command, 170

ip vrf forwarding command, 188

ISPs (Internet Service Providers), Managed Network service, 123

L

labels

- allocation/distribution
 - across ATM-LSR domains, 56–61
 - on-demand distribution, 35
 - VC merge, 58–61
 - binding, 34
 - conservative retention mode, 36
 - Frame-mode MPLS, 28–29
 - imposition, Edge-LSRs, 15–16
 - independent control, 57
 - liberal retention mode, 36
 - mappings, filtering, 74–77
 - packet forwarding, 10
 - across ATM-LSR domains, 55–56
 - frame-based label forwarding, 56
 - penultimate hop popping, 40
 - implicit-null value, 41–42
 - stack header, 26, 30
 - Bottom of Stack bit, 27
 - two-level stack, 332
 - swapping, 11
 - TTL propagation, 353–354
 - two-level stacks, 334
 - untagging, 29
- label-to-prefix bindings, 34
- LANs
- encapsulation, 27–28
 - Layer 2 paths, selecting, 6–7
 - MPLS/VPN deployment, 335–37
- large-scale networks, route reflection, 269
- Layer 2 paths, selecting, 6–7
- Layer 3 lookup, ATM, 50
- LC-ATM interfaces, 50
- control-plane connectivity, 51
 - downstream on demand, 57
 - status, verifying, 53
 - VCs, minimizing, 56
- LDP (Label Distribution Protocol), 31
- sessions, establishing, 33–34
- leased lines, 115–116
- LFIB (Label Forwarding Information Base) entries,
 - Frame-mode MPLS label lookup, 28–29
- LIB (Label Information Base), 17

liberal retention mode, 36, 253

Frame-mode convergence, 37–39

limitations, Cell-mode MPLS, 50

limiting label mapping distribution, 73–77

links

customer, managing, 338–351

PE-CE

configuring, 184

convergence, 256–257

RIP v2, 185–187

running separate BGP sessions, 300–301

static routing, 184–185

link-state routing protocols, OSPF. *See* OSPF

load sharing, equal-cost (EIGRP), 6

lookup, VPI/VCI, 49

loopback addresses

configuring on PE routers, 328

route summarization, 333

SuperNet example, 24

TDP/LDP identifiers, 75

loopback interfaces, ATM-LSRs, 53

loops, preventing

Cell-mode MPLS, 82–87

Frame-mode MPLS, 81–82

SOO, 180

LSAs (link-state advertisements), down-bits, 223

LSC (Label Switch Controller), 54

LSPs (Label Switched Paths), 17

LSRs (Label Switch Routers), 13

actions, 15

ATM-LSRs, 14, 50

ATM edge-LSRs, actions, 15

IP connectivity, 51

conservative retention mode, 36

Core LSRs, 25

Edge-LSRs, 14

actions, 15

architecture, 18

Ingress, 25

push action, 15–16

Frame-based, 50

Frame-mode, capabilities, 29

label stack, 30

liberal retention mode, 36

loopback interfaces, 53

pop action, 13

push action, 13

M

- Managed Network service, 123
- managed network VPN topology, 141
- managing customer links, 338–347
 - export maps, 348–351
 - importing selected customer routes, 339, 341–343
 - route filtering, 343–347
- mapping VPNs to VLANs, 336–337
- mappings (label), filtering, 74–77
- Maximum Initially Labeled IP Datagram Size, 79
- maximum routes command, 291
- maximum session handling, route reflectors, 268
- messages, Route Refresh, 202
- migration strategies
 - ATM backbone
 - case study, 97–99
 - to Cell-mode MPLS, 108
 - to Frame-mode MPLS, 106–107
 - BGP customer routes to MPLS/VPN, 230–231
 - IGP customer routes to BGP, 319–324
 - internal links to MPLS, 103–105
 - IP tunneling to MPLS/VPN solution, 387–389, 393, 400
 - configuring route reflectors, 398–399
 - defining VPNs, 390
 - routing policies, 390–394
 - staging and execution, 394–398
 - VPN site migration, 401
 - VRF definitions, 391–392
 - partial migration, 104
 - premigration infrastructure checks, 99–100
- minimizing VCs over LC-ATM interfaces, 56
- MP_REACH_NLRI attribute, 189
- MP_UNREACH_NLRI attribute, 189
- MP-iBGP, 171
 - configuring, 191–193
- PE-routers
 - configuring, 400
 - Route Refresh, 199–201
 - route distinguishers, 172
 - routing context, configuring, 193–194
 - static routes, redistribution, 185
- mpls ip command, 31
- mpls label-distribution both command, 31

- MPLS/VPN architecture, 11–13, 145
- AS Override, 232–234
- backbone
 - convergence between VPN sites, 255–260
 - customer access, 207–208
 - deploying BGP, 324–329, 331–333, 335
 - OSPF route propagation, 217–220
 - route advertisements, 260–265
 - traceroute, 351, 353–354
- carrier's carrier, 358–360
 - topologies, 361–372
- control protocols, 19–20
- customer routes, next-hop forwarding, 325–328
- data plane, packet forwarding, 203, 205
- deploying on LAN interfaces, 335, 337
- Extended Communities ORF-type
 - capability, 203
- FECs, 16
- goal of, 11
- inter-PE-router routing, 156–157
- intranet topology, 168–169
- migration from iBGP, 230–231
- migration from IP tunneling, 387, 389, 393, 400
 - configuring route reflectors, 398–399
 - defining VPNs, 390
 - routing policies, 390–394
 - staging and execution, 394–398
 - VPN site migration, 401
 - VRF definitions, 391–392
- MP-iBGP, 171
- overlapping addresses, 174
- PE-routers
 - Automatic Route Filtering, 197
 - requirements for advertisements, 167
- PE-to-CE connectivity
 - BGP-4, 209–211
 - OSPF, 212–213
 - RIP v2, 185–187
 - static routing, 184–185
- route distinguishers, 157, 172
 - configuring, 175–176
- route reflectors, configuring, 203
- routes, importing, 179
- SOO, 177–180
 - Extended Community, configuring, 181–182

- VPN-IPv4 addresses, 191
- VRFs
 - interface assignment, 188
 - configuring, 170
- See also VPNs
- MPLSCP (MPLS Control Protocol), 27
- MSS (maximum segment size), 78
- MTU, Path MTU discovery, 78–80
- Multiprotocol extensions (BGP), 189
- SuperCom network example, 158–161

N

- NBMA (non-broadcast multi-access) networks,
 - OSPF, PE-to-CE connectivity, 213
- neighbor advertisement-interval command, 257
- neighbor command, 192, 327
- neighbor maximum-prefix command, 291
- network command, 186
- network layer packet forwarding, 5–6
- network management, customer links, 338–347
 - export maps, 348–351
 - importing selected customer routes, 339–343
 - route filtering, 343–347
- networks (large-scale), route reflection, 269
- next hops
 - BGP, summarizing, 332
 - penultimate hop popping, 40
 - aggregates, 41
 - implicit-null value, 41–42
 - resetting at subconfederation boundary, 289
- next-hop forwarding, 325–328
- next-hop-self command, 289, 324–326
- NLRI (Network Layer Reachability Information), 189
- numbering route distinguishers, 175–176

O

- on-demand distribution, 35
- one-site-one-VRF model, 154
- OPEN messages (BGP), Capabilities parameter, 189
- optimization techniques (labels), VC merge, 58–61

- optional, non-transitive attributes (BGP), 189
- ordered control mode, 57, 73
- ORF capability, 201
 - configuring on route reflectors, 279–280
 - type definitions, 202
- OSPF (Open Shortest Path First)
 - PE-to-CE connectivity, 212–213
 - site area 0 support, 220–224
 - VRFs, separating, 214–217
 - without site area 0, 224–225
 - route propagation over MPLS/VPN backbone, 217–220
- overlapping address space, 148, 174
- overlapping VPNs, 151–154
- overlay VPN model, 121–123
 - convergence, 255–260
 - full-mesh, 133–134
 - hub-and-spoke topology, 133
 - hybrid topology, 134–135
 - partial-mesh, 133
 - PE-to-CE connectivity, OSPF, 212–213

P

- packets
 - Ethernet encapsulation, 77
 - forwarding
 - encapsulation, 27–28
 - external routing information
 - propagation, 10
 - label imposition, 16
 - label swapping, 12
 - MPLS/VPN data plane, 203–205
 - Path MTU discovery, 78
- partial-mesh VPNs, 133
- partial migration strategy, 104
- partitioning
 - LAN interfaces, 335
 - route reflectors, 272–274
- Path MTU
 - discovery, 78–80
 - MSS, 78
- PBR (policy-based routing), deploying, 10
- PDBs (protocol descriptor blocks), 292

- PE (Provider Edge) routers, 117
 - advertisements, requirements for, 167
 - Automatic Route Filtering, 197
 - BGP
 - route selection, memory requirements, 195
 - sessions, activating, 191
 - CE links, configuring, 184
 - RIP v2, 185–187
 - static routes, 184–185
 - loopback addresses
 - configuring, 328
 - summarization, 329–332, 334
 - MP-iBGP, configuring, 400
 - ORF, 201
 - PDBs, 292
 - provisioning, 291–292
 - route filtering based on community value, 274
 - route reflection, 270–271
 - Route Refresh, 199, 201
 - routing policies, IP tunneling migration case
 - study, 390
 - scaling PDBs, 292
 - standard community filtering, 276–277
 - tunneling to CE-router, 335
 - VRFs, numbering route distinguishers, 175–176
- PE-ASBR routers, 376
- PE-to-CE connectivity, 207
 - BGP-4, 209–211
 - iBGP customers, migrating to MPLS/VPN, 230–231
 - convergence, 256–257
 - MPLS/VPN backbone, customer access, 207–208
 - OSPF, 212–213
 - route propagation over MPLS/VPN backbone, 217–220
 - site area 0 support, 220–224
 - without site area 0, 224–225
 - redistribution, 208
 - running separate BGP sessions, 300–301
 - VPN customer network topologies, 226–230
 - VRFs, separating, 214–217
- peering sessions, iBGP, 325
- peer-to-peer VPN model, 123–125
 - comparing, 128
 - dedicated router approach, 126–127
 - shared router approach, 125–126
- penultimate hop popping, 40–42
 - aggregates, 41
 - implicit-null value, 41–42
- per-VPN routing protocols, 158
 - redistribution, 160–161
- PHP (penultimate hop popping), 16, 163
- physical interfaces, partitioning, 335
- P-network, 117
- policies (import/export), VRF configuration, 179–180
- pop actions, 13
- pop tag, 29
- popping the label stack, 40
- POPs (points of presence), migrating route reflection
 - to MPLS-based infrastructure, 98–99
- populating VRFs with customer routes, 207–208
- prefixes
 - label mappings, liberal retention mode, 36
 - route summarization, 92–93
 - unlabeled, 37
- premigration infrastructure checks, 99–100
- preparatory steps, VPN packet forwarding, 161
- preventing
 - labeled packet fragmentation, 79
 - routing loops
 - Cell-mode MPLS, 82–87
 - Frame-mode MPLS, 81–82
 - SOO, 180
- propagate-ttl command, 353
- protocols, ships-in-the-night, 54
- provisioning
 - full-mesh VPNs, 134
 - PE-to-CE connectivity, PE-router
 - configuration, 291–292
 - virtual circuits, 8
- push action, 13
 - Edge-LSRs, 15–16
- push tag, 29
- PVCs
 - ATM, Frame-mode MPLS operation, 67–69
 - shadow PVCs, 131

R

- RD (route distinguisher), 161
- reachability
 - backbone networks, 322
 - BGP, NLRI (Network Layer Reachability Information), 189
- recursive VPNs, 372–374
- redistribute bgp metric transparent command, 187
- redistribute command, 185, 217
- redistribution, 208
 - into OSPF, 212
 - per-VPN routing, 160–161
 - VRF OSPF routes into MP-iBGP, 217–218
- redundancy
 - hub-and-spoke VPN topology, 131
 - route reflectors, 269–270
- reflected routes, 268
 - best path selection, 271
 - improving convergence time, 257
 - partitioning, 272
 - route filtering based on route target attribute, 277
- removing
 - BGP community attribute, 351
 - unnneeded BGP peering sessions, 105
- requirements
 - BGP configuration in MPLS/VPN environment, 209
 - connectivity, SuperCom network example, 151–153
 - PE router advertisements, 167
 - VPNs
 - central services topology, 240
 - unique addresses, 149
- restricting label mapping distribution, 73–77
- restrictions
 - Cell-mode MPLS, 50
 - Ethernet encapsulation, 77
- retention modes (labels), 36
- reusing AS numbers, 232–234
- RIP v2, PE-to-CE link configuration, 185–187
- route distinguishers, 157, 171–172
 - assigning to VRFs, 391–392
 - configuring, 175–176
 - VPN-Ipv4, 172
- route filtering
 - automatic, 245
 - based on route target attribute, 277
 - customer links, 343, 345–347
- route reflectors, 268, 272
 - hierarchical structure, 270
 - PE routes, 270–271
 - improving convergence time, 257
 - IP-iBGP, configuring for IP tunneling migration, 398–399
 - maximum session handling, 268
 - migration to MPLS-based infrastructure, 98–99
 - MPLS/VPN architecture, configuring, 203
 - ORF capability, 279–280
 - partitioning, 272
 - redundancy, 269–270
 - route target attribute-based filtering, 277
- Route Refresh, 199–201
 - ORF entries, 202
 - See also ORF capability
- route summarization, 92–93
 - loopback addresses, 333
- route target attribute, 155–156, 178–179
 - extended community, format, 183
 - route filtering on reflected routes, 277
- route-maps command, 185
- router ospf command, 215
- router-on-a-stick, 268
- routers
 - ATM-LSRs, labeled packet forwarding, 55–56
 - independent control, 57
 - LSRs, 13
 - actions, 15
 - ATM edge-LSRs, 15
 - ATM-LSRs, 14–15, 50
 - Edge-LSRs, 14–15
 - Frame-based, 50
 - loopback interfaces, 53
 - pop action, 13
 - push action, 13
 - PE
 - MP-iBGP, configuring, 400
 - Standard Community filtering, 274–277
 - virtual routers, 149

- routes
 - exporting to specific targets, 347–351
 - per-VPN redistribution, 160–161
 - reflected. *See* reflected routes
 - VPN-IPv4, selecting, 194–196
- route-target command, 179
- route-target export commands, 351
- routing context, 159, 184
 - MP-iBGP configuration, 193–194
- routing loops, 270
 - preventing
 - Cell-mode MPLS, 82–87
 - Frame-mode MPLS, 81–82
 - SOO, 180
 - VPNs, avoiding, 247
- routing policies, IP tunneling migration case study, 390–394
- routing protocols
 - per-VPN, 158
 - VPN-aware, 184
- routing tables (BGP), scanner process, 258–259
- rr-group command, 203

S

- scalability
 - IP forwarding, 5
 - MPLS/VPN architecture, 157
 - PBR, 10
 - PE-routers, PDBs, 292
 - WANs, 7
- scanner process, 258–259
- security
 - central services, 137
 - firewalls
 - Cisco Secure Integrated software, 295
 - Internet connectivity, 293–295
- selecting
 - best path
 - BGP, 244
 - reflected routes, 271
 - BGP peer group leaders, 325
 - customer routes to import, 339–343
 - VPN-IPv4 routes, 194–196

- service providers
 - application hosting, 240
 - backbone, convergence, 253–254
- services, central services topology, 240
- sessions
 - BGP
 - activating between PE-routers, 191
 - internal, 188
 - PE-to-CE connectivity, 300–301
 - LDP, establishing, 31, 33–34
 - maximum handling capability, route reflectors, 268
 - MP-iBGP, scaling, 268
 - TDP, establishing, 31, 33–34
- set extcommunity rt command, 351
- shadow PVCs, 131
- shared router approach, peer-to-peer VPN model, 125–126
- shim headers, 27
- ships-in-the-night, 54
- show atm vc command, 54
- show ip bgp neighbor command, 200, 263–264, 325
- show ip ospf command, 216
- show tag tdp neighbor command, 33
- show tag-switching forwarding-table command, 204
- show tag-switching tdp bindings command, 36
- show tag-switching tdp neighbor command, 67, 75
- site area 0 (OSPF), 220–224
- sites, 117, 153, 393
 - stub sites, 184
 - Web, IETF, 11, 177
- software
 - Cisco IOS, control-plane connectivity, 52–53
 - Cisco Secure Integrated software, 295
- SOO (Site of Origin) attribute, 177, 180, 246
 - configuring, 181–182
 - Extended Community format, 183
- specifying route targets for export, 347–351
- stack header, 26
 - Bottom of Stack bit, 27
- staging IP tunneling migration, 394–398
- standard community attribute (BGP)
 - comparing to extended community attribute, 344–346
 - filtering on PE-routers, 274–277
 - route filtering on PE-routers, 274
- starting MPLS packet labeling, 31

- static routing
 - default routes, 295–300
 - PE to CE link configuration, 184–185
- stripping BGP community attribute, 351
- stub sites, 184
- sub-ASs, label exchange, 289
- subnet-to-label bindings, 34
- summarization, 92–93
 - on PE loopback addresses, 329–334
- SuperCom network example
 - Multiprotocol BGP, implementing, 158–161
 - VPNs, case study, 146–148
- swap tag, 29
- switches
 - ATM
 - Cell-mode MPLS operation, 69
 - control-plane implementation, 54
 - Frame-mode MPLS operation, 67–69
 - Layer 3 lookup, 50
 - ordered control, 57
 - Ethernet, MPLS MTU, 80
 - Frame Relay, Frame-mode MPLS operation, 65–67
 - VSI (Virtual Switch Interface), 54
- syntax
 - route target command, 179
 - tag-switching advertise-tags command, 75

T

- tables (IP forwarding), comparing to IP routing tables, 150
- tag tdp router-id command, 33
- tag-switching advertise-tags command, 75–76
- tag-switching commands, MPLS equivalents, 405–406
- tag-switching ip command, 31
- tag-switching mtu command, 80
- TCP (Transmission Control Protocol), MSS, 78
- TDP (Tag Distribution Protocol), 31
 - sessions, establishing, 33–34
- TFIB (Tag Forwarding Information Base), 18
- TIB (Tag Information Base), 17
 - liberal retention mode, 36
- tools, traceroute, running between VPN sites, 351–354

- topologies
 - BGP
 - confederations, deploying, 280–283, 285
 - peering structure, 101–103
 - carrier's carrier MPLS/VPN, 361–372
 - convergence
 - across ATM-LSR domains, 61–62
 - between VPN sites, 255–260
 - within service provider backbone, 253–254
 - MPLS/VPN Intranet, 168–169
 - PE-to-CE connectivity options, 208
 - route reflection
 - between PE-routers, 270–271
 - hierarchical, 269–270
 - large-scale networks, 269
 - VPNs
 - central services, 240–242
 - central-services extranet, 137–139
 - customer networks, 226–230
 - extranet, 135, 137
 - hub-and-spoke, 129, 131, 242–244
 - managed network, 141
 - VPDN, 139
- traceroute, 88–92
 - between VPN sites, 351–354
- traffic, forwarding, 10
- TransitNet, migration examples, 97
- troubleshooting
 - CE-to-PE links, 326
 - traceroute utility, 88–92
- TTL (Time-To-Live) field, 26
 - propagation, 353–354
- tuning BGP scanner process, 259
- tunneling, PE-to-CE routers, 335
- two-level label stacks, 332–334

U

- unicast routing, FECs, 34
- unique addresses, VPNs, 149
- unique numbers, assigning to route distinguishers, 175–176
- unlabeled prefixes, 37
- unsolicited downstream label distribution, 35
 - Frame-mode MPLS over Frame Relay, 65–67

- untag command, 29
- update-source parameter, neighbor command, 327
- utilities, traceroute, 88–89, 91–92
 - running between VPN sites, 351, 353–354

V

- VC merge, 58–61
- VCI (virtual circuit identifier), 49
- VCs (Virtual Circuits), 117
 - Frame-mode MPLS over Frame Relay, 65–67
 - in-band management, 51
 - lookup operation, 49
 - minimizing on LC-ATM interfaces, 56
 - provisioning, 8
 - shadow PVCs, 131
- verifying
 - imported customer routes, 346–347
 - LC-ATM interface status, 53
- virtual routers, 149
- VLANs
 - 802.1q, running MPLS/VPN, 337
 - mapping to VPNs, 336–337
- VPDN (Virtual Private Dialup Network)
 - topology, 139
- VPI (virtual path identifier), 49
- VPN-aware routing protocols, 184
- VPN-IPv4, 172
 - addresses, MP-iBGP configuration, 191
 - Automatic Route Filtering, 197
 - route exchange, configuring, 192
 - route selection, 194–196
- VPNs
 - business problem-based classification, 118–120
 - central-services extranet topology, 137–139
 - Cisco VPN Solutions Centre, 338
 - complexity of, 118
 - defining IP tunneling migration case study, 390
 - evolution of, 115–117
 - extranet/intranet integration, 237–238
 - hierarchical, 372–374
 - hub-and-spoke, 129–131, 242–244
 - AllowAS-in feature, 245–247
 - redundancy, 131
 - inter-provider, 374–383
 - label allocation, 162

- mapping to VLANs, 336–337
- overlapping, 151, 153–154
- overlay model, 121–123
 - full-mesh, 133–134
 - hybrid topology, 134–135
 - partial-mesh, 133
 - PE-to-CE connectivity, OSPF, 212–213
- packet forwarding, 161–163
- peer-to-peer models, 123–125
 - comparing, 128
 - dedicated router approach, 126–127
 - shared router approach, 125–126
- per-VPN routing, redistribution, 160–161
- route distinguishers, 157
- route targets, 155–156
- routes, importing, 179
- routing loops, avoiding, 247
- sites, 153
- SuperCom network example, 146–148
- topologies
 - central services, 240–242
 - customer network, 226–230
 - extranet, 135–137
 - full-mesh, provisioning, 134
 - managed network, 141
 - VPDN, 139
 - unique addresses, requirements, 149
- VRF (VPN routing and forwarding), 149–150
 - associating with global routing table, 310–311, 314
 - configuring, 170
 - import/export policies, 179–180
 - customer routes, importing, 339, 341–343
 - defining for IP tunneling migration, 391–392
 - interfaces, assigning, 188
 - one-site-one-VRF model, 154
 - populating with customer routes, 207–208
 - route distinguishers
 - configuring, 175–176
 - numbering, 175–176
 - route distinguishers, assigning, 391–392
 - route filtering, 343–347
 - separating on PE-to-CE links, 214–217
 - table, 159
- VSI (Virtual Switch Interface), 54

W

WANS

ATM

- Cell-mode MPLS operation, 69

- heterogeneous MPLS-mode operation, 69

- MPLS operation, 67, 69

- core redundancy, 7

- Frame Relay, MPLS operation, 65–67

- Layer 2 paths, selecting, 6–7

- virtual circuits, provisioning, 8

- Web sites, IETF, 11, 177