



Numerics

- 1+1 redundancy, 12
- 1:1 redundancy, 12
 - Cisco IOS architecture, 39, 62
 - control processor cards, 12
 - control planes, 36
- 1:N redundancy, 12
- 5-9s availability, 6

A

- ABRs (area border routers), 74
- access networks, 293
- address aggregation, 25
- address classes, 24
- Address Family Identifier (AFI) field, 153–154, 169, 174
- address lookup, 24
 - BGP, 147
 - double lookups, 29
 - edge LSRs, 56
 - FECs, 44
 - IP, 26–28
 - label lookup (versus), 50, 53
 - labels, 44
 - LSRs, 164
 - metrics, 26
 - MPLS, 164
 - route caching, 29
 - routers, 26
 - strides, 27
 - temporal versus spatial locality, 29
 - tries, 26–28
- address mapping, 31
- Address messages (LDP), 192
- address prefixes, 25
- Address Resolution Protocol (ARP), 31, 53
- adjacencies, 31, 67
 - ARP, 53
 - broadcast networks, 76, 113
 - cycling through down states, 115
 - flaps, 71, 107
 - graceful restart, 67
 - Hello adjacencies
 - establishing, 189-190
 - maintaining, 190-191
 - packets, 70, 109
 - processing, 81
 - holding timers, 110
 - IS-IS, 107–110
 - LDP, 189–191
 - master/slave status, 76
 - NBMA networks, 76
 - OSPF, 75–76
 - reacquisition, 118–119
 - RouterDeadInterval timers, 75
 - T1 timers, 123
 - three-way handshakes, 109
- Adj-RIBs-In, 148
- Adj-RIBs-Out, 148
- Advertisement messages (LDP), 191
- advertisements (BGP), 144
- AFI (Address Family Identifier) field, 153–154, 169, 174
- aggregates (CIDR), 25
- aggregation (addressing), 25
- ALLSPFRouters addresses, 75
- Any Transport over MPLS (AToM), 55, 275
- APS (Automatic Protection Switch), 9, 12, 17, 280
- area border routers (ABRs), 74
- areas
 - not-so-stubby, 78
 - OSPF, 73
 - stub, 77
- ARP (Address Resolution Protocol), 31, 53
- AS external LSAs (type 5), 77–78
- AS_PATH lists (BGP), 146
- AS_SEQUENCE lists (BGP), 146
- AS_SET lists (BGP), 146
- ASBR summary LSAs (type 4), 77–78
- ASBRs (autonomous system boundary routers), 74
- Asynchronous Transfer Mode (ATM), 48
- AToM (Any Transport over MPLS), 55, 275
- authentication, 19
- Automatic Protection Switch (APS), 9, 12, 17, 280
- autonomous systems, 67, 137, 146
- availability, 6. *See also reliability*
 - carrier-class, 6, 11
 - control/forwarding planes, 35

- DPM, 6
- DPY, 6
- edge routers, 9
- fault tolerance, 7
- five-nines, 6
- MPLS networks, 256
- MTBF/MTTR, 6
- packet-switched networks, 11
- redundancy, 7
- reliability (versus), 6
- service view, 11
- software upgrades (in-service), 295–296
- unavailability (versus), 6

B

- backbone areas (OSPF), 73
- backbone routers, 74
- backup designated routers (BDRs), 76
- backup paths, 7
- backup tunnels, 55
- BDRs (backup designated routers), 76
- BFD (Bidirectional Forwarding Detection), 288–290
 - access networks, 293
 - detecting IGP neighbor liveness, 290
 - FRR, 292–293
 - IP/MPLS networks, 294–295
 - LSPs, 291
 - PWs, 292
- BGP (Border Gateway Protocol), 8, 137
 - advertising labeled routes, 171
 - control processor cards, 138
 - control-plane components, 138
 - End-Of-RIB markers, 153–156
 - external reachability information, 109
 - FIBs, 138
 - forwarding-plane components, 138
 - FSM, 142
 - label distribution, 169
 - label-to-FEC bindings, 46
 - labeled routes (withdrawing), 171
 - LDP, 221
 - line cards, 138
 - messages, 141–142
 - MPLS
 - control planes, 168
 - graceful restart/NSF interoperability, 182–183
 - interrelationship, 168–169
 - NSF, 140, 166, 168
 - path attributes, 143, 146
 - AS_PATH, 146
 - LOCAL_PREF, 149
 - MP_REACH_NLRI, 169
 - MP_UNREACH_NLRI, 169
 - NEXT_HOP, 147
 - protocol extensions, 168
 - recursive paths, 31
 - restarts, 140
 - churns, 140
 - detrimental effects, 139
 - graceful restart, 151–161, 168
 - helper BGP speaker behavior, 155–156
 - mitigating detrimental effects, 150–151
 - MPLS, 166–167, 173–182
 - NSF interoperability, 159–161
 - restarting BGP speaker behavior, 154–155
 - route flaps, 138–140
 - RIBs, 138
 - routing, 140–144
 - fluctuations, 138
 - instability, 138
 - reflection, 149–150, 171
 - selection, 149
 - tables, 25
 - speakers, 137, 141
 - internal/external, 144
 - RIBs, 148
 - RRs, 150
 - states, 142–143
- BGP with an MPLS graceful restart mechanism.
 - See BGP/MPLS networks*
- BGP/MPLS networks
 - advertising labeled BGP routes, 171
 - control planes, 168
 - protocol interrelationship, 168–169
 - restarts, 166–167
 - withdrawing labeled BGP routes, 171
- BGP/MPLS restarts
 - graceful restart, 173–175, 179–180
 - helper LSR behavior, 178
 - network deployment, 181–182

- restarting LSR behavior, 175–178
 - mitigating detrimental effects, 173
- binary tries, 26
- Birectional Forwarding Detection (BFD), 288–290
 - access networks, 293
 - detecting IGP neighbor liveness, 290
 - FRR, 292–293
 - IP/MPLS networks, 294–295
 - LSPs, 291
 - PWs, 292
- black holes, 70, 75, 83, 109, 114
- Border Gateway Protocol. *See* BGP
- bypass tunnels, 55, 284–285

C

- carrier supporting carrier (CSC), 270–271
- carrier-class networks, 6
 - router expectations, 11
- CEF (Cisco Express Forwarding), 33–34
- cell-mode LSRs, 48
- Checkpoint Procedures for LDP (LDP CDP), 200
- churns (BGP), 140
- CIDR (classless interdomain routing), 25–26
- Cisco Express Forwarding (CEF), 33–34
- Cisco IOS architecture
 - 1:1 redundancy, 39, 62
 - control-plane software, 15
 - label-distribution protocols, 53
 - MPLS SSO/NSF, 61
 - NSF, 39–40
 - SSO, 39–40
 - TFIB/MFI, 50
- Cisco IS-IS restarts, 128–129
 - CSNPs, 129
 - IETF IS-IS restarts (versus), 128
 - interoperability, 132–134
 - operation, 130
 - PSNPs, 129
- classful addressing, 24
 - CIDR (versus), 26
 - IP forwarding, 25
- classless addressing, 25
- classless interdomain routing (CIDR), 25–26
- complete sequence number packets (CSNPs), 111
 - congestion indication (IS-IS), 112–113
 - constraint-based path calculation, 55
 - control capabilities (routers), 132
 - control planes, 8
 - 1:1 redundancy, 36
 - adjacency flaps, 107
 - BGP, 138
 - BGP/MPLS networks, 168
 - Cisco IOS architecture, 15
 - control processor cards, 12
 - FIBs, 23
 - forwarding planes (separating), 19, 36
 - forwarding planes (versus), 35
 - IP, 23, 106
 - restart, 35–36
 - separating forwarding planes, 35
 - SSO, 37
 - tunneling protocols, 8
 - LDP, 197–198
 - maintaining state, 13–14
 - MPLS, 53–54, 61, 163–164
 - applications, 55
 - separating forwarding planes, 54
 - SSO, 57–59
 - MPLS-TE, 258
 - NSF, 37
 - OSPF, 69
 - recovery periods, 35
 - restarts, 13, 15, 132
 - routers, 12
 - RSVP-TE, 236
 - SSO, 15
 - stateful components, 35
 - stateless components, 35
 - unplanned restarts, 13
 - control processor cards, 12
 - 1:1 redundancy, 12
 - active versus standby, 13
 - BGP, 138
 - reducing fault effects, 36
 - restartability, 16
 - RIBs, 69
 - router processors (Cisco), 39, 61
 - CSC (carrier supporting carrier), 270–271
 - CSNPs (complete sequence number packets), 111, 129

D

data encryption, 19
data planes, 277
 APS, 280
 FRR, 281–284
 protection/restoration, 278–280
 IP layer, 280
 MPLS layer, 281–284
 optical layer, 279–280
databases, 70, 110–112
 exchange process (OSPF), 75
 synchronization
 IETF IS-IS restarts, 119–120
 OSPF, 70–71
 timers, 118, 121
Database Description packets, 75
defects per million (DPM), 6
denial-of-service (DoS) attacks, 18
depths (label stacks), 45
designated IS (DIS) routers, 113
designated routers (DRs), 76
Dijkstra algorithm.
 See SPF (Shortest Path First) algorithm
DIS (designated IS) routers, 113
Discovery messages (LDP), 191
disposition (label), 52
distributed forwarding
 CEF, 34
 data rates, 30
 IP forwarding, 30–32
 MPLS, 56
 time criticality, 30
DOD (downstream on-demand) mode, 46, 193
 DU mode (versus), 219–220
 LDP
 graceful restart, 213–219
 nonrestarting LSRs (egress neighbors),
 218
 nonrestarting LSRs (ingress neighbors),
 217–218
 nonrestarting LSRs (transit neighbors),
 218–219
 restarting LSRs (egress), 215–216
 restarting LSRs (ingress), 214–215
 restarting LSRs (transit), 216–217

DoNotAge bits, 90
DoS (denial-of-service) attacks, 18
downstream modes (LDP), 192
downstream on-demand mode. *See DOD (downstream on-demand) mode*
downstream unsolicited mode.
 See DU (downstream unsolicited) mode
downtime, 5–6. *See also network outages*
 carrier-class routers, 11
 causes, 10
 link failures, 17
 software upgrades, 17
downtime per year (DPM), 6
DPM (defects per million), 6
DPY (downtime per year), 6
DRs (designated routers), 76
DU (downstream unsolicited) mode, 46, 194
 DOD mode (versus), 219–220
 LDP graceful restart, 203–204, 208–209, 212

E

eBGP (external BGP) sessions, 144
edge LSRs, 46, 164
 address lookup, 56
 label disposition, 52
 label imposition, 51
 Layer 2/3 services, 164–165
edge routers, 8
EGP (Exterior Gateway Protocol), 140
egress LSRs, 46
element nodes, 26
encoding (labels), 48
encryption (data), 19
End-Of-RIB markers, 153–156
enhanced IS-IS, 108
ERO (EXPLICIT_ROUTE) objects, 233
established state (BGP), 142–143
Ethernet, 48
expansion (prefix), 27
EXPLICIT_ROUTE objects (RSVP-TE), 233
explicitly routed LSPs, 47
Exterior Gateway Protocol (EGP), 140
external BGP (eBGP) sessions, 144
external/internal speakers (BGP), 144

F

- F (Forwarding State) bits, 154
- failures (network/services), 5
- Fast ReRoute. *See* *FRR (Fast ReRoute)*
- fault tolerance, 7, 35
 - availability, 7
 - data-plane protection, 277
 - MPLS networks, 257
 - MPLS-based recovery, 277–285
 - OAM mechanisms, 287–290
 - network design, 11
 - restarts, 201–203
 - security, 18–19
 - software upgrades (in-service), 295–296
- FEC (forwarding equivalence class), 44, 164
 - address lookups, 44
 - label mappings, 45–46, 54, 176
 - LSP associations, 188–189
 - Layer 2 circuit IDs, 197
 - NHLFE (FTN) entries, 197
 - NHLFE (FTN) maps, 50
 - next-hop mappings, 54
- FIBs (Forwarding Information Bases), 23, 31, 106
 - BGP, 138
 - building, 32, 56
 - line cards, 69
 - m-tries, 31
 - RIBs (versus), 30
 - routers, 31
 - RSVP-TE, 238
 - SPF algorithm, 73, 111
 - synchronizing with RIBs, 32, 56
- filtering (packets), 19
- FILTERSPEC objects (RSVP), 229
- finite state machines (FSMs)
 - OSPF, 85, 93
 - BGP, 142
- firewalls, 19
- five-nines availability, 6
- flags, 112
- flaps (adjacency/route), 107, 138–140
- flooding, 75, 105
 - IS-IS, 112
 - Link-State Acknowledgment packets, 75
 - OSPF, 68–70
- FLOWSPEC objects (RSVP), 229
- forwarding
 - ATM versus MPLS, 48
 - FEC, 44
 - FIBs, 23
 - IP, 23–24, 35–36
 - loops, 114
 - routers, 132
 - states
 - IP, 12, 39
 - LDP, 197–198
 - MPLS, 12, 166
 - tables
 - FECs, 44
 - IP forwarding plane, 23
 - L3VPNs, 266–267
 - route caching (versus), 29
 - VPNs (166
- forwarding planes, 12
 - adjacency flaps, 107
 - BGP, 138
 - control planes, 19, 35–36
 - FIBs, 23, 31, 106
 - IP, 23
 - control planes (separating), 35
 - NSF, 36
 - operations, 24
 - tasks, 23
 - LDP, 197
 - MPLS, 53–54, 61, 163–164
 - control planes (separating), 54
 - label-stack operations, 51–54
 - NSF, 43, 58
 - MPLS-TE, 258
 - NSF, 37
 - OSPF, 69
 - RSVP-TE, 236
 - scalability, 35, 54
- Forwarding State (F) bits, 154
- FR (Frame Relay), 48
- FRR (Fast ReRoute), 17, 50, 55, 281–284
 - BFD, 292–293
 - link failures, 17
- FSMs (finite-state machines)
 - BGP, 142
 - OSPF, 85, 93
- FT ACK TLV, 202
- FT Protection TLV, 202

FT Session TLV, 201, 204–205
 FTN (FEC-to-NHLFE) entries, 197
 FTN (FEC-to-NHLFE) maps, 50

G–H

Generic Routing Encapsulation (GRE), 8
 GR. *See graceful restarts (GRs)*
 grace LSAs, 89

- planned/unplanned restarts, 92–93
- reason TLV, 90

 graceful restarts (GRs), 67

- BFD, 294–295
- BGP, 151–161
- BGP/MPLS, 173–175, 179–180
 - helper LSR behavior, 178
 - network deployment, 181–182
 - restarting LSR behavior, 175–178
- IETF OSPF Working Group, 100
- LDP, 201–203
 - DOD mode, 213–219
 - DU mode, 203–204, 208–209, 212
 - exchanging initial capability, 204
 - GR versus FT, 203
 - interoperability, 221–222
 - network deployment, 220–221
 - nonrestarting peers, 206–207
 - restarting peers, 207–208
 - session failure, 205
 - session reestablishment, 205
 - state recovery, 205
- OSPF, 89–100
- reason TLV, 90
- RSVP-TE, 237–239, 245–247
 - control-plane restart, 239
 - exchanging initial capability, 239
 - network deployment, 248
 - nonrestarting LSRs, 244–245
 - recovery procedure, 242
 - reestablishing Hello communication, 240
 - restarting LSRs, 241, 243–244

 GRE (Generic Routing Encapsulation), 8

hardware redundancy, 11
 head-end restarts (RSVP-TE), 243
 Hello adjancies

- establishing, 189–190
- maintaining, 190–191
- packets, 70, 109
- processing, 81

 Hello packets

- OSPF, 70
- RS bits, 80

 HELLO_ACK object messages (RSVP-TE), 235
 HELLO_REQUEST object messages (RSVP-TE), 235
 helper nodes, 113
 helper speakers (BGP), 154
 hierarchical routing (IS-IS), 108–109
 holding timers, 110
 hop-by-hop routed LSPs, 47

I

I (Init) bits, 82
 iBGP (internal BGP) sessions, 144
 idle state (BGP), 142–143
 IETF (Internet Engineering Task Force)

- LDP, 187, 200
- MPLS-TE, 264
- RSVP-TE, 238

 IETF IS-IS restarts, 115–116

- adjacency reacquisition, 118–119
- Cisco IS-IS restarts (versus), 128
- database synchronization, 119–120
- interoperability, 132
- LSP generation/flooding, 120–121
- nonrestarting routers, 121–122
- operation, 124
 - restarting routers, 126–128
 - starting routers, 124–125
- restart TLV, 116–117
- restarting routers (with preserved FIB), 118–121
- SPF computation, 121
- starting routers (without preserved FIB), 122–123
- timers (T1–T3), 118

- IETF OSPF Working Group, 100
- IGPs (interior gateway protocols), 72
- IIHs (IS-IS Hellos), 109, 116
- ILMs (incoming label maps), 50, 197
- Implicit NULL labels, 52
- imposition (label), 51
- incoming label maps (ILMs), 50, 197
- incoming labels, 45
- ingress LSRs, 46
- Init (I) bits, 82
- in-service software upgrades, 295–296
- Integrated IS-IS, 108
- interarea paths (OSPF), 78
- interior gateway protocols (IGPs), 13
- Intermediate System-to-Intermediate System.
 - See IS-IS*
- internal BGP (iBGP) sessions, 144
- internal/external speakers (BGP), 144
- Internet Protocol (IP), 23
- Internet routing architecture, 67–69
- interoperability
 - BGP, 159–161
 - Cisco IS-IS restarts, 134
 - Cisco/IETF IS-IS restart, 132
 - OSPF restarts, 100–102
 - RSVP-TE, 248–249
- interprocessor communication (IPC), 39
- intra-area paths (OSPF), 78
- intrusion detection/prevention, 19
- IOS. *See Cisco IOS architecture*
- IP (Internet Protocol), 23
- IP forwarding, 23–24
 - addressing (classful/classless), 25
 - address lookup, 26–28
 - architectures, 28
 - distributed, 30–32
 - route caching, 29
 - CEF, 33–34
 - NSF, 36
- IP networks
 - addressing, 24–25
 - control-plane protocols, 37
 - forwarding, 44
 - forwarding planes, 12, 24
 - forwarding tables, 28
 - IS-IS, 107
 - load balancing, 31
 - NSF, 36–40, 106
 - routing protocols, 8, 12
 - SSO, 35–36, 40
- IP over MPLS, 55
- IP/MPLS networks, 5
 - BFD, 294–295
 - control-plane protocols, 12
 - control-plane restart, 13
 - control-plane software, 8
 - forwarding planes, 12
 - forwarding states, 12
 - IP routing protocols, 8
 - IP tunneling protocols, 8
 - LSRs, 12
 - migrating legacy services, 5
 - MPLS signaling protocols, 8
 - network design, 19
 - OAM mechanisms, 17
 - protocols, 8
 - security, 18
 - service/network views, 11
- IPC (interprocessor communication), 39
- IS-IS (Intermediate System-to-Intermediate System), 8
 - dual routing, 107
 - flags, 116
 - Hello packets, 109
 - IP networks, 107
 - LDP, 221
 - levels, 68
 - LSPs, 68
 - restart TLV, 116–117
 - sequence number packets, 111
 - three-way handshakes, 109
- IS-IS Hellos (IIHs), 109
- IS-IS restarts, 113–114
 - adjacency flaps, 107
 - Cisco, 128–131
 - detrimental effects, 105, 113
 - IETF restarts, 124–128
 - IETF versus Cisco, 131
 - mitigating detrimental effects, 113
 - Cisco restarts, 128–129
 - IETF IS-IS restarts, 115–128
 - network deployment, 132
 - original behavior, 105–107

- restarting routers, 115
- starting routers, 115
- state information, 128
- IS-IS routing, 107–108
 - adjacencies, 110
 - congestion indication, 112–113
 - designated routers, 113
 - discovering neighbors, 109–110
 - establishing adjacencies, 109–110
 - hierarchical, 108–109
 - LSPs, 110–112

K–L

- Keepalive messages (BGP), 141–142
- keepalive timers (LDP), 190
- L bits, 80
- L2TP (Layer 2 Tunneling Protocol), 8
- L2VPNs (Layer 2 virtual private networks), 50, 55, 164, 197, 272–273
 - control-plane failures, 274
 - forwarding tables, 166
 - VPLS, 275–277
 - VPWS, 274–275
- L3VPNs (Layer 3 virtual private networks), 55, 164, 265
 - control-plane failures, 268
 - CSC, 270–271
 - multi-AS backbone, 269–270
 - single-AS backbone, 268–269
 - forwarding tables, 166
 - IPv6-based services, 268
 - label distribution, 268
 - MPLS, 50
 - PE-based, 265–266
 - PE-to-PE tunnels, 267
 - routing/forwarding tables, 266–267
- Label Abort messages (LDP), 192
- Label Distribution Protocol. *See* LDP
- Label Mapping messages (LDP), 192
- Label Request messages (LDP), 192
- Label Withdraw messages (LDP), 192
- label-advertisement mode, 192
- label-distribution control mode, 192
- label-encoding techniques, 48
- Label Information Bases. *See* LIBs
- label-retention mode, 194
- labels, 44
 - distribution (BGP), 169
 - DOD mode, 46
 - DU mode, 46
 - encoding techniques, 48
 - Implicit NULL, 52
 - imposition/disposition, 46
 - incoming versus outgoing, 45
 - label space, 50
 - label stacks, 45–46
 - label-to-FEC binding, 45
 - lookup, 50, 53, 163
 - MPLS forwarding, 51
 - label disposition, 46, 52
 - label imposition, 46, 51
 - label swapping, 51
 - recovering labels, 176
- label-switched paths. *See* LSPs
- Label-Switching Information Bases. *See* LFIBs
- label-switching routers. *See* LSRs
- label-to-FEC mappings, 176
- Layer 2 Tunneling Protocol (L2TP), 8
- Layer 2 virtual private networks. *See* L2VPNs
- Layer 3 virtual private networks. *See* L3VPNs
- LC-ATM (label-switching controlled ATM)
 - interfaces, 48
- LDP (Label Distribution Protocol), 8, 46, 187–188
 - active versus standby, 59
 - BGP, 221
 - Cisco IOS architecture, 53
 - control planes, 197–198
 - DOD mode, 193, 214
 - DU mode (versus), 219–220
 - graceful restarts, 213–219
 - restarting LSRs, 214
 - restarting LSRs (egress neighbors), 218
 - restarting LSRs (engress), 215–216
 - restarting LSRs (ingress neighbors), 217–218
 - restarting LSRs (ingress), 215

- restarting LSRs (transit neighbors), 218–219
 - restarting LSRs (transit), 216–217
- DU mode, 194
 - graceful restarts, 203–204, 208–209, 212
 - DOD mode (versus), 219–220
- FEC-to-label mappings, 54
- FEC-to-LSP associations, 188–189
- FIBs, 198
- forwarding planes, 197
- forwarding states, 197–198
- graceful restarts
 - DOD, 214–219
 - exchanging initial capability, 204
 - GR versus FT, 203
 - interoperability, 221–222
 - network deployment, 220–221
 - nonrestarting peers, 206–207
 - restarting peers, 207–208
 - session failures, 205
 - session reestablishment, 205
 - state recovery, 205
- Hello adjacencies
 - establishing, 189–190
 - maintaining, 190–191
- hop-by-hop routed LSPs, 48
- IETF, 187, 200
- ILMs, 197
- Implicit NULL labels, 52
- IS-IS, 221
- keepalive timers, 190
- label-to-FEC bindings, 46
- label-advertisement mode, 192
- label-distribution control mode, 192
- label-retention mode, 194
- LDP CKP, 200
- LFIBs, 194
- LIB/LFIB/routing interactions, 194–196
- LSPs, 187, 197–198
- Maximum Recovery Time, 206
- messages, 191–192
- MPLS, 43
- MPLS forwarding state, 198
- Neighbor Liveness Timer, 206
- NHLFEs, 197
- NSF, 197
- OSPF, 221
- peers, 189
- PWs, establishing, 196–197
- Recovery Time, 205
- restarts, 198–199
 - methods, 201–203
 - mitigating detrimental effects, 200
- routing, 194
- TCP, 189
- LDP CKP (Checkpoint Procedures for LDP), 200
- leaf nodes, 26
- level 1/2 LAN Hello (IS-IS), 109
- LFIBs (Label-Switching Information Bases), 50, 163
 - building, 56
 - LDP, 194–196
 - RSVP-TE, 236
 - synchronizing with the LIB, 56
- LIBs (Label Information Bases)
 - LDP, 194–196
 - RIBs (versus), 53
- line cards
 - BGP, 138
 - CEF, 33
 - distributed forwarding, 30
 - FIBs, 69
 - forwarding-plane functions, 12
 - IPC, 39
 - MPLS networks, 8
 - synchronizing RIBs to FIBs, 32
- link failures (mitigating), 17
- link local signaling (LLS), 80
- Link-State Acknowledgment (LSA) packets, 75
- link-state advertisements. *See* LSAs
- link-state databases. *See* LSDBs
- link-state protocols, 68, 72, 105
- Link-State Request packets, 75
- Link-State Update packets, 75
- LLS (link-local signaling), 80
- load balancing, 31
- locality (temporal versus spatial), 29
- Loc-RIB, 148
- LR (LSDB Resynchronization) bits, 82
- LSAs (link-state advertisements), 13, 68
 - flooding, 70
 - grace LSAs, 89, 92–93
 - MaxAge, 77
 - OSPF, 76–79
 - self-originated LSAs, 85
 - types, 77

LSDB Resynchronization (LR) bits, 82

LSDBs (link-state databases), 68

- OSPF, 70–71
- purging LSAs, 77
- resynchronization, 81

LSP_TUNNEL objects (RSVP-TE), 231

LSPs (label-switched paths), 8, 106, 166

- adjacency flaps, 107
- BFD, 291
- congestion indication, 112–113
- CSNPs, 111
- database synchronization, 111, 114
- explicitly routed/traffic engineered, 47
- flags, 112
- flooding, 112
- fragmenting, 112
- generation/flooding, 120–121
- hop-by-hop routed, 47, 188
- IS-IS, 109–112
- LDP, 187, 197–198
- MPLS, 61, 197–198
- MPLS-TE, 226
- nesting, 46
- tunnels, 231
- zeroth LSPs, 112

LSRefreshTime timer, 90

LSRs (label-switching routers), 45, 163

- address lookup, 164
- ATM, 48
- cell mode, 48
- DOD mode (LDP), 214–219
- edge LSRs, 46, 164–165
- frame mode, 48
- label-distribution protocols, 46
- LFIBs, 50
- penultimate LSP LSR, 46
- MPLS domains, 46
- NSF, 58
- recovery periods, 58
- RSVP-TE
 - nonrestarting LSRs, 244–245
 - restarting LSRs, 241–244
- transit LSRs, 46, 164

M

M (More) bits, 82

M/S (Master/Slave) bits, 82

make-before-break approach (MPLS-TE), 262

Master/Slave (MS) bits, 82

Maximum Recovery Time (LDP), 206

mean time between failures (MTBF), 6

mean time to repair (MTTR), 6

mesh connectivity, 8

MESSAGE_ACK objects (RSVP-TE), 235

messages

- BGP, 141–142
- LDP, 191–192

MFI (MPLS Forwarding Infrastructure), 50

midpoint restarts (RSVP-TE), 243

modularity, 15, 17

More (M) bits, 82

MP_REACH_NLRI attributes, 169

MP_UNREACH_NLRI attributes, 169

MPLS (Multiprotocol Label Switching), 7, 43–50

- applications, 55
- BGP, 168–169
- control-plane components, 163–164
- FECs, 164
- forwarding planes
 - components, 163–164
 - NSF, 43
- forwarding state, 166
 - entries, 166
 - LSPs, 197
- IP (versus), 163–164
- label lookup, 50, 53
- labels, 45
- LDP, 187, 198
- networks. *See MPLS networks*
- SONET/SDH (versus), 17

MPLS Forwarding Infrastructure (MFI), 50

MPLS Forwarding State Holding Timer, 205

MPLS networks, 46

- applications, 50
- ATM, 48
- backbone, 7
- components, 7, 9
- control planes, 59, 61
- distributed forwarding, 56
- domains, 46

- edge LSRs, 164
 - Ethernet, 48
 - FEC-to-label mappings, 46
 - forwarding, 50, 53
 - architecture, 56
 - ATM (versus), 48
 - forwarding planes, 12, 61
 - IP (versus), 50
 - label operations, 51
 - state, 166
 - FR, 48
 - FRR, 50
 - IP forwarding (versus), 50–53
 - L2VPNs, 50
 - L3VPNs, 50
 - Layer 2/3 services, 254
 - availability, 256
 - fault-tolerance techniques, 257
 - provider-provisioned VPNs, 254
 - service attributes, 256
 - VPN tunnels, 255–256
 - line cards, 8
 - LSPs, 47, 61
 - LSRs, 45
 - NSF, 43, 58–62
 - PPP, 48
 - protection layer interaction, 285
 - recovery, 277–285
 - routers, 7
 - separating control/forwarding planes, 53–54
 - signaling protocols, 12
 - SSO, 57–58, 62
 - MPLS/BGP networks
 - graceful restart/NSF interoperability, 182–183
 - MPLS-TE (Multiprotocol Label Switching traffic engineering), 50, 55, 226, 257
 - control-plane failures, 262–263
 - establishing tunnels, 259–261
 - functional modules, 258
 - IETF, 264
 - inter-area/intra-AS, 264
 - inter-AS, 264
 - intra-area, 263
 - reoptimizing tunnels, 262
 - using RSVP, 230
 - MS (Master/Slave) bits, 82
 - MSE (multiservice edge) routers, 7
 - MTBF (mean time between failures), 6
 - m-tries (multibit-tries), 27
 - MTTR (mean time to repair), 6
 - multibit-tries (m-tries), 27
 - Multiprotocol Label Switching. *See MPLS*
 - multiservice edge (MSE) routers, 7
- ## N
-
- Neighbor Liveness Timer, 206
 - nesting (LSPs), 46
 - network design
 - core networks, 8
 - edge networks, 8
 - fault tolerance, 11
 - IP/MPLS networks, 19
 - redundancy, 7
 - network layer reachability information (NLRI), 141
 - network LSAs (type 2), 77–78
 - network outages, 10. *See also outages*
 - causes, 10–11
 - downtime, 5
 - mitigating
 - hardware outages, 11–12
 - link/node failures, 17
 - node-level outages, 17
 - via OAM mechanisms, 17
 - software outages, , 12–17
 - planned, 10, 17
 - unplanned, 10–17
 - network partitioning, 108
 - network summary LSAs (type 3), 77–78
 - NEXT_HOP attributes (BGP), 147
 - next-hop label forwarding entries (NHLFEs), 50, 197, 238
 - NLRI (network layer reachability information), 141
 - NMBA (nonbroadcast multiaccess) networks, 75–76
 - node failures (mitigating), 17
 - nonrestarting routers (IETF IS-IS restarts), 121–122
 - nonstop forwarding. *See NSF*
 - Notification messages
 - BGP, 141–142
 - LDP, 191
 - not-so-stubby-areas (NSSAs), 78
 - NSF (nonstop forwarding), 15, 69, 106
 - BGP, 140, 159–161, 166–168

- capable versus incapable neighbors, 40
- Cisco IOS architecture, 39–40
- defensive techniques, 18
- external view, 40
- IP, 36–39
 - control-plane SSO, 37
 - forwarding, 36
 - neighboring-node perspective, 40
 - separating control/forwarding, 37
- LDP, 197
- MPLS
 - control-plane SSO, 59
 - neighboring-node perspective, 62
- routers, 36
- RSVP-TE, 238
- NSSAs (not-so-stubby areas), 78

O

- OAM (operation, administration, and maintenance)
 - mechanisms, 17, 287–290
- OL (overload) bit, 112–113
- one-for-N (1:N) redundancy, 12
- one-for-one (1:1) redundancy, 12
- one-plus-one (1+1) redundancy, 12
- OOB (out-of-band) resynchronization, 82
- OOBResynch bits, 82
- Open messages (BGP), 141–142, 151
- Open Shortest Path First. *See OSPF*
- operation, administration, and maintenance (OAM)
 - mechanisms, 17, 287–290
- optional nontransitive (BGP path attribute), 146
- optional transitive (BGP path attribute), 146
- OSPF (Open Shortest Path First), 8
 - areas, 73
 - control-plane components, 69
 - databases
 - Database Description packets, 75
 - exchange process, 75
 - synchronizing, 70–71
 - discovering neighbors, 75
 - flooding, 70
 - forwarding-plane components, 69
 - FSM, 85, 93
 - Hello packets, 70, 81
 - inter-area paths, 78
 - intra-area paths, 78
 - LDP, 221
 - link-state protocols, 68, 72
 - Link-State Request packets, 75
 - Link-State Update packets, 75
 - LSAs, 75, 77
 - restarts. *See OSPF restarts*
 - routing, 72–73
 - adjacencies, establishing, 75–76
 - hierarchical, 73–74
 - link-state databases, synchronizing, 75–76
 - LSAs, 76–79
 - paths, selecting, 78–79
 - router classification, 74
 - tables, 73–74
 - SPF algorithm, 72
 - type 1/2 external paths, 79
- OSPF restarts, 69–70
 - control-plane restarts, 69–71
 - detrimental effects, 70–71
 - graceful restarts
 - neighboring router behavior, 95–96
 - restarting router behavior, 93–94
 - interoperability, 100–102
 - mitigating detrimental effects, 79
 - graceful restarts, 89–92, 96, 99
 - graceful restarts versus restart signaling, 99–100
 - restart signaling, 80–81
 - neighbor awareness, 70
 - nonrestarting routers, 85–86
 - planned/unplanned, 79, 92–93
 - protocol extensions, 70
 - restart signaling, 86–89
 - restarting routers, 84
 - SPF calculations, 85
- outages
 - causes, 10–11
 - downtime, 5
 - mitigating
 - hardware outages, 11–12
 - link/node failures, 17
 - node-level outages, 17
 - via OAM mechanisms, 17
 - software outages, 12–17

- planned, 10, 17
- unplanned, 10–17
- node/link, 277
- services, 256
- outgoing labels, 45
- out-of-band (OOB) resynchronization, 82
- Overload (OL) bits, 112–113

P

- packet-scheduling algorithms, 31
- packet-switched networks, 11
- packets
 - filtering, 19
 - labeled versus unlabeled, 45
- partial sequence number packets (PSNPs), 111
- path attributes (BGP), 146
 - LOCAL_PREF, 149
 - MP_REACH_NLRI, 169
 - MP_UNREACH_NLRI, 169
- Path messages (RSVP), 227
- path state block (PSB) messages, 227
- PathErr (path error) messages, 227
- paths, 105
 - backup, 7
 - mesh connectivity, 8
 - next hop, 30
 - selecting (OSPF), 78–79
- peer relationships. *See adjacencies*
- penultimate LSRs, 164
- penultimate-hop LSRs, 52
- penultimate-hop popping, 46, 164
- point-to-point Hello (IS-IS), 109
- Point-to-Point Protocol (PPP), 48
- point-to-point three-way handshakes, 109
- portioning (networks), 108
- PPP (Point-to-Point Protocol), 48
- PPVPNs (provider-provisioned VPNs), 254
- prefixes
 - addresses, 25–28
 - expansion, 27
 - CIDR, 25–26
 - supernets, 25
- protection (data-plane failure), 278
- provider-provisioned VPNs (PPVPNs), 254

- PSB (path state block) messages, 227
- pseudowires (PWs), 196–197, 292
- PSNPs (partial sequence number packets), 111, 129
- PWs (pseudowires), 196–197, 292

R

- R (Restart State) bits, 82, 153
- RA (Restart Acknowledgment) flags, 116
- radix (tries), 26
- reason TLV (graceful restarts), 90
- RECORD_ROUTE objects (RSVP-TE), 233–234
- recovery periods
 - MPLS, 58
 - routers, 35
- RECOVERY_LABEL objects (RSVP-TE), 239
- recursive paths, 30
- redundancy, 7, 257
 - APS, 12
 - edge routers, 9
 - fault tolerance, 7
 - hardware, 11
 - IP control-plane protocols, 37
 - MPLS control-plane SSO, 59
 - schemes, 12
 - software approaches, 14
 - switchovers, 7, 15
- refresh timers, 84
- reliability. *See also availability*
 - availability (versus), 6–7
 - MTBF/MTTR, 6
 - routers, 9
 - service view, 11
- reservation state block (RSB) messages, 229
- resource exhaustion DoS attacks, 18
- Resource Reservation Protocol. *See RSVP*
- resources (networks), 18
- Restart Acknowledgment (RA) flags, 116
- Restart Request (RR) flags, 116
- restart signaling
 - OSPF, 80–81, 86–89
 - graceful restart (versus), 99–100
- Restart Signaling (RS) bits, 80
- Restart State (R) bits, 153
- restart TLV, 116–117
- RESTART_CAP objects (RSVP-TE), 239

- restartability, 15–17
- restarting routers, 115
 - adjacency reacquisition, 119
 - CSNP, 120
 - IETF IS-IS, 118–121
 - timers, 118
- restarting speakers (BGP), 154
- restarts
 - BGP, 140
 - graceful restarts, 151–154
 - mitigating detrimental effects, 150–151
 - route flaps, 138–140
 - fault-tolerant restarts, 201–203
 - grace LSAs, 92–93
 - head-end restarts (RSVP-TE), 243
 - IS-IS, 113–114
 - detrimental effects, 105
 - IETF versus Cisco, 131
 - mitigating detrimental effects, 113–116
 - network deployment, 132
 - original behavior, 105–107
 - LDP
 - detrimental effects, 198–199
 - methods, 201–203
 - mitigating detrimental effects, 200
 - midpoint restarts (RSVP-TE), 243
 - OSPF, 69–70
 - detrimental effects, 70–71
 - mitigating detrimental effects, 79
 - graceful restart, 89–92, 96, 99
 - graceful restart versus restart signaling, 99–100
 - restart signaling, 80–81
 - preserving state information, 131
 - protocol extensions, 69–70
 - RSVP-TE, 236–245
 - tail-end restarts, 244
- RestartState bits, 81
- restoration (data-plane failure), 278
- Resv messages (RSVP), 228–229
- ResvErr (reservation error) messages, 229
- Resynctimeout timer, 82
- RIBs (Routing Information Bases), 30, 69
 - BGP, 138
 - BGP speakers, 148
 - control processor cards, 69
 - End-of-RIB markers, 153, 155–156
 - LIBs (versus), 53
 - m-tries, 31
 - SPF algorithm, 73
 - updates, 31
- route caching
 - FIBs (versus), 31
 - hit ratio, 29
 - IP forwarding, 29
- route flaps, 138, 140, 167
- route reflectors (RRs), 149–150, 171
- route selection, 149
- router LSAs (type 1), 77–78
- router processors (RPs), 39, 61
- RouterDeadInterval timers, 75, 88
- routers, 23
 - ABRs, 74
 - address lookup, 26, 44
 - adjacencies, 76
 - ASBRs, 74
 - ATM LSRs, 48
 - autonomous systems (AS), 137
 - backbone routers, 74
 - BDRs, 76
 - BGP, 25
 - BGP speakers, 137, 144
 - bottlenecks, 23
 - carrier-class routers, 10–11
 - cell-mode LSRs, 48
 - components, 8
 - control/forwarding capabilities, 132
 - control-plane availability, 16
 - control-plane protocols, 12
 - control-plane software, 12
 - designated routers (IS-IS), 113
 - distributed forwarding, 30
 - downtime, 10
 - DRs, 76
 - edge routers, 8
 - FIB, 31
 - flags, 116
 - flooding, 105, 112
 - forwarding process, 24
 - fragmenting LSPs, 112
 - frame-mode LSRs, 48
 - IGPs, 13

- internal routers, 74
- IP routers, 23
- IS-IS restart TLV, 116–117
- IS-IS restarts
 - nonrestarting routers, 121–122
 - restarting routers (with preserved FIB), 118–121
 - starting routers (without preserved FIB), 122–123
- link-state protocols (LSPs), 105
- LSRs, 12, 45
- MPLS networks, 7
- MSE routers, 7
- network partitioning, 108
- NSF, 15, 36, 69
- OSPF classifications, 74
- OSPF restarts
 - nonrestarting routers, 85–86
 - restarting routers, 84–85
- outages, 11
- recovery periods, 35
- reliability, 9
- restarting routers, 115
 - adjacency reacquisition, 119
 - CSNP, 120
 - IETF IS-IS, 126–128
 - OSPF, 71
 - T3 timers, 118
- SPF algorithm, 111
- SSO, 16
- starting routers, 115, 124–125
- traditional forwarding architectures, 28
- routing
 - BGP, 137, 140–144
 - domains, 67, 72
 - dual routing, 108
 - IGP, 72
 - interdomain routing, 137
 - Internet architecture, 67–69
 - IS-IS, 107–108
 - congestion indication, 112–113
 - designated routers, 113
 - discovering neighbors, 109–110
 - establishing adjacencies, 109–110
 - hierarchical, 108–109
 - LSPs, 110, 112
 - maintaining adjacencies, 110
 - LDP, 194–196
 - OSPF, 72–73
 - establishing adjacencies, 75–76
 - hierarchical, 73–74
 - LSAs, 76–79
 - selecting paths, 78–79
 - synchronizing LSDBs, 75–76
- Routing Information Bases. *See RIBs*
- routing protocols, 8, 12
 - active versus standby, 37, 59
 - control processor cards, 36
 - control software, 106
 - FEC-to-next hop mappings, 54
 - RIBs, 30
- routing tables
 - BGP, 25, 143
 - L3VPNs, 266–267
 - OSPF, 73–74
 - reducing the size of, 25
- RPs (router processors), 39, 61
- RR (Restart Request) flags, 116
- RRO (RECORD_ROUTE) objects, 233
- RRs (route reflectors), 149–150, 171
- RS (Restart Signaling) bits, 80
- RSB (reservation state block) messages, 229
- RSVP (Resource Reservation Protocol), 8, 227
 - label-to-FEC bindings, 46
 - MPLS-TE, 230
 - Path messages, 227
 - PathErr messages, 227
 - Resv messages, 228–229
 - ResvErr messages, 229
 - soft state, 229
- RSVP traffic engineered (RSVP-TE), 48, 187, 225
 - control/forwarding-plane components, 236
 - detecting failures, 235–236
 - EXPLICIT_ROUTE objects, 233
 - FIBs, 238
 - graceful restart, 237–247
 - Hello extensions, 235
 - IETF, 238
 - interoperability, 248–249
 - LFIBs, 236
 - LSP tunnels, 231
 - LSP_TUNNEL objects, 231–232
 - MESSAGE_ACK objects, 235
 - new objects, 230–232, 239

- NHLFE, 238
- RECORD_ROUTE objects, 233–234
- restarts, 236–245
- SESSION_ATTRIBUTE objects, 232
- soft state, 234
- state lifetime, 234–235
- SUMMARY_REFRESH messages, 235
- traffic flows, 230
- RSVP_HOP object (RSVP), 227, 229

S

- SA (Suppress Adjacency Advertisement) flags, 117
- SAFI (Subsequent Address Family Identifier) field, 153, 169, 173–174
- scheduling, 24
- security, 18
 - DoS attacks, 18
 - improving via fault-tolerant mechanisms, 18–19
- send routing message flags. *See SRMflags*
- send sequence number flags (SSNflags), 112
- SENDER_TEMPLATE objects (RSVP), 227
- SENDER_TSPEC objects (RSVP), 227
- sequence number packets (IS-IS), 111
- service-level agreements (SLAs), 11
- services (networks), 10, 256
 - outages, 6–7
- Session message (LDP), 191
- SESSION objects (RSVP), 227, 229
- SESSION_ATTRIBUTE objects (RSVP-TE), 232
- Shortest Path First, 13
- Shortest Path First algorithm.
 - See SPF (Shortest Path First) algorithm*
- signaling protocols, 8
- SLAs (service-level agreements), 11
- soft state
 - RSVP, 229
 - RSVP-TE, 234
- software upgrades, 17, 295–296
- SONET (Synchronous Optical Network), 280
- SONET/SDH (Synchronous Optical Network/ Synchronous Digital Hierarchy), 7
 - APS, 9, 12
 - MPLS (versus), 17
- speakers (BGP), 137, 141
 - internal/external, 144
 - RIBs, 148
 - RRs, 150
- special locality, 29
- SPF (Shortest Path First) algorithm, 13, 68, 105
 - LSPs, 111, 120
 - OSPF, 72
 - OSPF restarts, 85
 - RIBs, 73
 - zeroth LSP, 112
- SRMflags (send routing message flags), 112
 - Cisco IS-IS restarts, 129
 - IS-IS restarts, 114
- SSNflags (send sequence number flags), 112
- SSO (stateful switchover), 15
 - capable components, 37
 - Cisco IOS architecture, 39–40
 - defensive techniques, 18
 - external view, 40
 - IP
 - control plane, 35–37
 - neighboring-node perspective, 40
 - MPLS
 - control plane, 59
 - neighboring-node perspective, 62
 - restartability, 16
- starting routers, 115
 - IETF IS-IS restarts, 122–123
 - timers, 118
- state transitions, 93
- stateful components, 35–37, 57
- stateful switchover. *See SSO*
- stateless components, 35
- states (BGP), 142–143
- strides
 - address lookup, 27
 - patterns, 28
- stub areas, 77
- STYLE objects (RSVP), 229
- Subsequent Address Family Identifier (SAFI) field, 153, 169, 173–174
- SUGGESTED_LABEL objects (RSVP-TE), 239
- SUMMARY_REFRESH messages (RSVP-TE), 235
- supernets, 25
- Suppress Adjacency Advertisement (SA) flags, 117
- survivability. *See fault tolerance*
- swapping (labels), 51
- switching, 24
- switching fabric (routers), 8

switchovers
 fault control processor cards, 36
 redundancy, 7
 restartability, 15
 SSO, 35–36
 synchronizing databases, 114, 118
 IETF IS-IS restarts, 119–120
 OSPF, 70–71
 timers, 118, 121
 Synchronous Optical Network/Synchronous Digital Hierarchy. *See SONET/SDH*

T

T1 timers, 118
 adjacency reacquisition, 119
 restarting versus starting routers, 123
 T2 timers, 118
 adjacency reacquisitions, 119
 database synchronization, 120
 SPF computations, 121
 start-capable routers, 122
 T3 timers, 118
 expiring before T2 timers, 121
 IETF IS-IS restart routers, 118
 inhibiting IIIH transmissions, 119
 Tag Forwarding Information Bases (TFIBs), 50
 tail-end restarts (RSVP-TE), 244
 TCP (Transmission Control Protocol), 189
 TE (traffic engineering), 50, 225–226, 257
 temporal locality, 29
 TFIB (Tag Forwarding Information Bases), 50
 three-way handshakes, 109
 timers
 adjacencies, 110
 Hello adjacencies, 190
 IETF IS-IS restarts, 118
 LDP, 206
 LSP generation/flooding, 120–121
 LSRefreshTime, 90
 MaxAge, 77
 MPLS Forwarding State Holding Timer, 205
 refresh timers, 84
 ResyncTimeout, 82
 RouterDeadInterval timer, 88

TLVs (time length values)
 FT ACK, 202
 FT Protecton, 202
 FT Session, 201, 204
 reason TLV, 90
 restart TLV, 116–117
 traffic engineering (TE), 50, 225–226, 257
 traffic trunks, 226
 traffic-engineered LSPs, 47
 transit LSRs, 46, 51, 164
 tries, 26, 28
 tunnels
 bypass tunnels, 284–285
 MPLS-TE
 establishing, 259–261
 protecting, 262–263
 reoptimizing, 262
 RSVP, 231
 VPNs, 255
 demultiplexing, 255
 signaling labels, 256
 tuples (CIDR), 25
 type 1 LSAs, 77–78
 type 1/2 external paths (OSPF), 79
 type 2 LSAs, 77
 type 3 LSAs, 77
 type 4 LSAs, 77
 type 5 LSAs, 77
 type 7 LSAs, 77

U–W

unavailability, 6. *See also downtime; outages*
 Update messages (BGP), 141–144
 distributing routing information, 171
 path attributes, 146–147
 Withdrawn Routes field, 171
 VC labels, 165, 168
 VCs (virtual circuit identifiers), 45
 Virtual Private LAN Service (VPLS), 275–277
 virtual private networks. *See VPNs*
 Virtual Private Wire Service (VPWS), 274–275
 VPLS (Virtual Private LAN Service), 275–277

VPNs (virtual private networks)

- backbones, 164

- forwarding tables, 166

- L2VPNs, 272–273

 - control-plane failures, 274

 - VPLS, 275–277

 - VPWS, 274–275

- L3VPNs, 265

 - control-plane failures, 268–271

 - IPv6 services, 268

 - label distribution, 268

 - PE-based, 265–266

 - PE-to-PE tunnels, 267

 - routing/forwarding tables, 266–267

- MPLS, 164–165

- provider provisioned, 254

- tunnels, 168, 255

 - demultiplexing, 255

 - signaling labels, 256

VPWS (Virtual Private Wire Service), 274–275

wavelength-division multiplexing (WDM),

- 7, 279–280

well-known discretionary (BGP path attribute), 146

well-known mandatory (BGP path attribute), 146

- AS_PATH, 146

- NEXT_HOP, 147