

## Numerics

---

- 0.0.0.0/8 addressing, 230
- 16-8-8 data structure, 40
- 256-way mtrie, 40
- 8-8-8-8 mtrie structure, 40

## A

---

- access routers
  - as starting router for network migration, 311
  - migration concerns, 318
  - moving to new peer groups, 319
- ACLs, filtering classless routing updates, 115
- address families (AFs), 28
- address space, provider-based summarization, 237
- addressing
  - anycast, 564
  - IPv6
    - address format, 564–565
    - aggregatable global unicast addresses, 566
    - local addressing, 566
    - special addresses, 567
  - ISP networks, 397
- adjacency tables, 40
- administrative policies, 8
- ADMs (add/drop multiplexers), 583
- advertisement suppression communities, 405
- advertisements
  - AFs, 28
  - aggregates, 230
  - conditional, 123
    - examples, 124, 127–129
    - forms, 124
  - SAFIs, 28
- AFs (address families), 28
- aggregatable global unicast addresses, 566
- aggregation, 230
  - aggregate prefix, default attributes, 130
  - BGP policy control, 130–133
  - remote site (BGP enterprise core), 192
- aggregation layer (ISP networks), 390
- aggregation routers, 391
- Allow-AS, 460
- American Registry for Internet Numbers (ARIN), 226
- anycast, 564
- Anycast RP, 548
- APNIC (Asia Pacific Network Information Centre), 226
- architectures
  - DCNs, 584
  - enterprise network core design, 160
    - eBGP, 168–178
    - iBGP, 161–167
    - internal/external BGP, 178–189, 192
  - migration. *See* migration
  - MPLS VPNs, 449
- are, 599
- ARIN (American Registry for Internet Numbers), 226
- ARP table (Address Resolution Protocol), 33
- AS (autonomous systems), 294
- AS Override, 457–458
- AS path lists, 117
- ASBRs (AS border routers), 461, 476
- AS\_PATH attribute, 18, 21
  - disabling, 460
  - iBGP confederation, 296
  - looping, 458

- manipulation to provide full connectivity
  - among VPN sites, 457
    - Allow-AS, 460
    - AS Override, 458
  - updates with Local AS feature, 135
- Asia Pacific Network Information Centre (APNIC), 226
- ASN
  - generic customer ASN, 398
  - private, 399
- assigning IGP labels, 452
- atom characters, 110
- automatic route filtering, 456
- autosummarization feature (BGP), 396

## B

---

- backbone carriers, 461
- back-to-back VRF, 462–463
- bandwidth, load balancing, 231–232
- baseline convergence, 97
- BGP, 396
  - administrative policies, 8
  - aggregation router configuration template, 392
  - capabilities, 27
  - characteristics, 5
  - CLNS, 595
  - configuring for MPLS VPNs, 452
  - core router configuration template, 390
  - DCN management networks, 583
  - edge router configuration template, 395
  - enterprise core. *See* enterprise core
  - features, 599
  - flexibility, 8
  - GR feature, 88–89
  - internal. *See* iBGP
  - memory use, 16, 48–50
  - multicast NLRI in MP-BGP, 536
  - multiprotocols extensions, 435
    - extended community attribute, 436
    - reachability attributes, 437
  - network command, 31
  - NSAP address family configuration, 594
  - path attributes, 17
  - path selection, 24
  - performance. *See also* performance policies, control techniques, 109
  - processes, 14–15
  - redistributing information into IGP, 23
  - related RFCs, 605
  - reliability, 6
  - routing exchange with an IGP, 31
  - routing policies, 8
  - scalability, 7
  - security, ISPs, 412. *See also* ISP networks
  - stability, 6
  - strengths in enterprise networks, 158
  - summary tables, 317
  - traffic, worst-case input queue values, 69
  - VPN label propagation, 453, 455
  - vs. IGP, 9
  - weaknesses, 159
- bgp always-compare-med command, 19
- BGP-based DCN network design
  - BGP next hop for CLNS prefixes, 587
  - BGP peering relationships, 585
  - IS-IS area layout, 585
- bgp bestpath med-confed command, 26
- bgp bestpath missing-as-worst command, 26

- BGP community attribute
  - dynamic customer policy, 404–405
  - prefix origin tracking, 402
  - QPPB feature, 407, 410
  - static route redistribution, 411
- BGP confederation ID, 326
- BGP confederations, 489
- BGP convergence, 62
  - defined, 62
  - examples, 62–63
  - performance interdependencies, 82
  - queue optimization, 67–68
    - hold queue, 68
    - packet reception process, 67
    - SPD (selective packet discard), 69–70
    - system buffers, 72–73
- TCP protocol concerns, 64
  - max segment size, 65
  - PMTUD (Path MTU Discovery), 65
  - window size, 65
- tuning (case study), 97–100, 103
- update generation, 74
  - BGP dynamic peer group feature, 77
  - BGP read-only mode, 82
  - peer groups, 75
  - peer templates, 77–78
  - update groups, 80–81
  - update packing enhancement, 81
- bgp deterministic-med command, 26
- BGP dynamic peer group feature, 77
- BGP I/O process, 14
- BGP network core design, 160
  - architectures, 160
  - case study, 197
    - core activation, 207
    - design requirements, 195–196
    - design scenario, 194
    - final router configuration, 208–212
    - migration plan, 199–201, 204–206
  - eBGP core architecture, 168
    - administration control, 178
    - failure and recovery, 174–177
    - path selection, 169–173
    - routing policy, 178
  - iBGP core architecture, 161
    - administration control, 167
    - failure and recovery, 165–166
    - path selection, 162–164
    - routing policy, 167
  - internal/external BGP core architecture, 178
    - administrative control, 189
    - failure and recovery, 187
    - path selection, 180–183, 186
    - routing policy, 189, 192
  - remote site aggregation, 192
- BGP network performance features, 83
  - failure mitigation, 83
    - BGP fast external fallover, 84
    - IGP/BGP convergence time deltas, 84–86
    - NSF feature, 87
  - prefix update optimization, 91
    - BGP ORF feature, 96
    - BGP soft reconfiguration, 94
    - route flap dampening, 91–93
    - route refresh, 95
    - transmit (TX) side loop detection, 95
- BGP policy accounting (BPA), 143–144
- BGP Router process, 15–16
- BGP Scanner process, 15
- BGP synchronization feature, 396
- black-holing traffic during migration, 311
- BPA, 144
  - (BGP policy accounting), 143

**C**

- cache-based switching, 35
  - advantages over process switching, 35
  - forms of, 36
  - shortcomings, 37
- caches
  - invalidation, 38
  - NetFlow, 37
  - shortcomings of cache-based switching, 38
- capabilities (BGP), 27
  - advertising AFs and SAFIs, 28
  - exchanging between routers, 29
- carrier supporting carrier VPN model. *See* CSC VPN model
- CE (Customer edge) routers, 449, 455
- CEF (Cisco Express Forwarding), 38
  - advantages, 38
  - dCEF, 43
  - distributed CEF. *See* dCEF
  - FIB, 39
  - inconsistency checker, 42
  - load sharing, 44
    - per-packet, 45
    - per-session, 44
- cell-based networks, 439
- Cisco bug ID numbers, 599
- Cisco Express Forwarding. *See* CEF
- Cisco IOS
  - BGP processes, 14–15
  - default label exchange protocol, 443
  - regular expressions, 112
  - RIB lookup operations, 33
  - Site of Origin (SOO), 437
  - supported BGP path attributes, 17
    - switching paths, 33
      - cache-based, 35–37
      - CEF, 38–39, 42–45
      - comparison of, 46
      - process, 33
    - table of BGP features, 599
    - WEIGHT parameter, 19
- Cisco Systems URLs, 606
- Class D and E address space, 230
- classfull lookup operation, 33
- classless lookup operation, 33
- CLNS
  - ISO-IGRP, 584
  - next-hop reachability, 587
  - originating a default from the core network, 586
- clns enable command, 588
- clns filter-set DEFAULT\_OUT command, 595
- CLUSTER\_LIST attribute, 21, 262
- clustering, 259, 288
- commands
  - peer policy templates, 78
  - peer session templates, 77
- communities, iBGP confederation, 297
- community assignments (ISPs), 402
- COMMUNITY attribute, 20
- community lists, 118–119
- conditional advertisement, 123
  - examples, 124, 127–129
  - forms, 124
- conditional injection, 131
- confederation
  - AS\_PATH attribute, 296
  - communities, 297
  - design examples, 298
    - hub-and-spoke, 298
    - setting IGP metrics, 299–303
  - functionality, 294
  - vs. route reflection, 303–304

- confederation-based architecture, 309
  - migrating to from iBGP full mesh (case study), 326
  - migrating to from RR-based (case study), 343
  - migrating to RR-based (case study)
    - migration procedures, 366–377, 379–383
    - starting configurations, 362–365
- configuring
  - aggregation routers, 392
  - BGP core routers, 390
  - conditional route injection, 131
  - edge routers, 395
  - enterprise border routers, 227
  - graded route map dampening, 414
  - MP-BGP for IPv6, 569–570
  - MSDP, 536
  - policy lists, 122
  - routers, 241
- connectivity
  - customer (ISP networks), 398
    - customer BGP peering, 398
    - identifying customer prefixes, 399
    - static route redistribution, 399
  - transit (ISP networks), 400
- consistency checkers (CEF), 43
- control plane
  - RIB, 32
  - vs. forwarding plane, 13
- convergence
  - baseline, 97
  - MPLS VPNs, 499–500
- core routers, 311, 389
- CSC (carrier supporting carrier) VPN model, 481
  - full Internet routes, 482–484
  - hierarchical VPNs, 485, 488
- customer BGP peering, 398
- customer carriers, 461

- customer connectivity (ISP networks), 398
  - customer BGP peering, 398
  - customer route redistribution, 399
  - identifying customer prefixes, 399
- Customer edge routers (CE routers), 449, 455

## D

- dCEF (distributed CEF), 42
- DCNs (Data Communications Network), 583
  - architecture, 584
  - BGP-based design
    - BGP next hop for CLNS prefixes, 587
    - BGP peering relationships, 585
    - IS-IS area layout, 585
  - multiprotocol BGP for CLNS configuration
    - network verification, 590–592
    - router configuration, 592–594
  - scalability, 583
- DDoS attacks, 419–421
- deaggregation, 130–133
- debug ip bgp command, output example, 30
- debug ip bgp update command, output example, 436
- debug ip bgp update in command, 457
- default routing, 416
- default routing only, 221
- defining multicast streams, 517
- deploying
  - MP-BGP for IPv6, 569
  - MPLS VPNs, 490
    - convergence, 499–500
    - route target (RT) design, 497–498
    - scalability, 491–496

## designing

- iBGP confederation architecture, 298
  - hub-and-spoke, 298
  - setting IGP metrics, 299–303
- ISP networks, 388
  - aggregation layer, 390
  - network core layer, 389
  - network edge layer, 393–395
- network core design, 158–160
- networks
  - BGP in the enterprise core, 157
  - determining solution, 158
- RDs for MPLS VPNs, 495–496
- route reflection systems, 266–276
  - clustering design, 288
  - enabling deterministic-MED, 278
  - enforcing IGP metric settings, 286
  - iBGP mesh, 277
  - resetting next hop, 290
  - RR with peer groups, 292
  - setting IGP metrics, 279–285
- RR-based MPLS VPNs, 492
- disabling AS\_PATH, 460
- discard adjacency, 41
- distributed CEF, 42
- distributed optimum switching, 36
- downstream distribution method, 441
- downstream on demand, 441
- downtime, minimizing during migration, 318
- drop adjacency, 41
- dual stacks, 568
- dynamic black hole routing, ISP security case study, 420

---

**E**

---

## eBGP

- BGP network core design, 160
  - multihop feature, 233
  - multipath feature, 235
  - multiprotocol, 465
  - paths, 26
  - routers, 232, 311
- eBGP core architecture, 168
- administration control, 178
  - failure and recovery, 174–177
  - path selection, 169–173
  - routing policy, 178
- ebgp-multihop command, 235
- eBGP-neighbor routes, 591
- Edge Label Switch Router (LER), 438
- edge routers, 394
- eBGP vs. iBGP, 24
- egress LERs, 446
- end node autoconfiguration block, 229
- end systems (ESs), 583
- End System-to-Intermediate System (ES-IS)
  - adjacency, 587
- end-of-RIB marker, 88
- end-to-end LSPs
  - carrying VPN traffic, 471
  - full Internet routes, 482
- enforcing policies, 8
- enterprise core
  - BGP use, 157
  - designing, 160
    - architectures, 160
    - BGP core design (case study), 194
    - core design, 197
    - design requirements, 195–196
    - eBGP core architecture, 168–178

- final router configuration (case study), 208–212
- iBGP core architecture, 161–167
- internal/external BGP core architecture, 178–183, 186–189, 192
- migration plan, 199–201, 204–207
- remote site aggregation, 192
- Enterprise Networks
  - designing
    - BGP strengths, 158
    - BGP weaknesses, 159
  - Internet connectivity, 221
    - default plus partial routes, 222
    - default routing only, 221
    - full Internet tables, 222
    - load balancing, 231–235, 240
    - peering filters, 238
    - provider-based summarization, 237
    - route filtering, 230
- entries, 40
- ES-IS (End System-to-Intermediate System)
  - adjacency, 587, 590
- ESs (end systems), 583
- exception adjacencies, 41
- Exp field (frame-based label header), 440
- Explicit Null, 441
- Explicit Null labels, 448
- extended community attribute, 436–437
- extensions (multiprotocol), 435
  - extended community attribute, 436
  - reachability attributes, 437
- extranet VPNs, 497

## F

- failure mitigation, 83
  - BGP fast external fallover, 84
  - IGP/BGP convergence time deltas, 84–86
  - NSF feature, 87
- fast switching, 35
- FIBs (forwarding information bases), 14, 39–40
  - adjacency tables, 40
  - data structures, 39
  - dCEF, 42
  - policy parameters, 138
  - routes, 448
  - setting policy entries based on BGP tagging, 139
- fields (frame-based label header), 439–440
- filter lists
  - AS path, 117
  - community, 118–119
  - prefix lists, 114–115
- filter processing order, 123
- filtering
  - classless routing updates, 115
  - inbound, 229
  - Martian address space, 230
  - outbound, 230
  - peer filtering (ISP networks), 413
  - peering, 238
  - prefixes for IPv6, 570
  - route (automatic), 456
  - updates (route maps), 120
  - upstream prefix advertisements, 403
- flexibility, 8
- format, 435
- forwarding
  - labeled packets, 446–448
  - packets, 479

forwarding information bases. *See* FIBs  
forwarding loops, 334  
forwarding paths, 450  
forwarding plane  
    RIB, 32  
    vs. control plane, 13  
frame-based networks, 439

## G

---

gateway network elements (GNEs), 584  
generic customer ASN, 398  
Gigabit RP (GRP), 42  
glean adjacency, 41  
GNEs (gateway network element), 584  
GR (Graceful Restart) feature, 88–89  
graded route map dampening, 416  
GRE tunneling, 418  
groups (multicast), 517  
GRP (Gigabit RP), 42

## H

---

hierarchical route reflection, 264–265  
hierarchical VPNs, 485, 488  
hierarchy (ISP networks)  
    aggregation layer, 390  
    network core layer, 389  
    network edge layer, 393–395  
hub-and-spoke topologies  
    iBGP confederation system, 298  
    VPNs, 497

---

iBGP (internal BGP), 23  
    BGP network core design, 160  
    core architecture, 161  
        administration control, 167  
        failure and recovery, 165–166  
        path selection, 162, 164  
        routing policy, 167  
    full mesh topology, 308  
    migrating to confederation (case study),  
        326–331, 334–343  
    migrating to route reflection (case study),  
        312–322, 325  
    loop-prevention mechanism, 253  
    paths, 26  
    scalability, 253  
        confederation, 294–303  
        method comparison, 303–304  
        route reflection, 254–267, 269–279,  
            281–286, 288–292. *See also* route  
            reflection  
    transporting prefix information, 23  
    vs. eGBP, 24  
    vs. IGP, 23  
IGP (Interior Gateway Protocol), 9  
    design goal, 22  
    ISP networks, 388  
    label assignment, 452  
    network command, 31  
    routing exchange with BGP, 31  
    vs. BGP, 9  
    vs. iBGP, 23  
IGP/BGP convergence time deltas, 84, 86  
Implicit Null, 442  
Implicit Null labels, 448



- import map command, 499
- imposition (LSRs), 446
- in labels, 444
- inbound filtering, 229
- inbound load balancing, 231
- input hold queue, 67–69
- integrating two AS systems (case study), 145–151
- inter-AS VPNs, 461
  - back-to-back VRF, 462–464
  - multihop eBGP between RAs and IPv4 labels (casestudy), 501, 504
  - multihop multiprotocol eBGP, 470, 474
    - carrying remote addresses in iBGP, 475–476
    - label exchange, 472
  - next-hop-self command, 470
  - non-VPN transit providers, 476–478
  - option comparison, 480–481
  - single-hop multiprotocol eBGP, 465–469
- interdomain multicast. *See* MSDP
- Interior Gateway Protocol. *See* IGP
- intermediate systems (ISs), 583
- internal BGP (iBGP), 23
- internal/external BGP core architecture, 178
  - administrative control, 189
  - failure and recovery, 187
  - path selection, 180–183, 186
  - routing policy, 189
- INTERNET community, 20
- Internet connectivity
  - accepting information, 221
    - default plus partial routes, 222
    - default route only, 221
    - full Internet tables, 222
  - load balancing, 231
    - case study, 240
    - inbound traffic, 231–232
    - multiple sessions to the same provider, 233–235
  - multihoming, 222
    - standard multihomed network design, 226–229
    - stub network design, 223–225
  - peering filters, 238
  - provider-based summarization, 237
  - route filtering, 230
  - stub network single-homed design, 223
- Internet Service Providers. *See* ISP networks
- interrupt contact switching, 35
- intraconfederation eBGP sessions, 329
- invalidation (cache), 38
- ip bgp-community new-format command, 20
- IP multicast. *See* multicast delivery model
- ip policy-list command, 122
- IP to IP path, 446
- IP to label path, 446
- ip vrf forwarding command, 451
- IPv4
  - address family, 28
  - addressing capability, 562
  - deploying a dual stack IPv6 configuration (case study), 572–576
  - header format, 563
  - historical overview, 561
  - QoS, 563
  - reserved labels, 441
  - security, 563
- IPv6
  - addressing capability, 562–564
    - address format, 564–565
  - aggregatable global unicast addresses, 566
  - local addressing, 566
  - special addresses, 567

- autoconfiguration, 562
- deploying dual-stack configuration (case study), 572–576
- development goals, 561
- enhancements over IPv4, 561
- header format, 563
- MP-BGP extensions, 568
  - configuring, 569–570
  - deployment considerations, 569
  - dual-stack deployment, 568
- security, 563
- ipv6 unicast-routing command, 569
- IS-IS protocol
  - BGP-based DCN network design, 585
  - Overload bit (OL-bit), 85
- ISO-IGRP, 584
- ISP (Internet Service Providers) networks, 387
  - addressing methodology, 397
  - architecture, 387
    - IGP layout, 388
    - network design, 388–390, 393–395
  - BGP community attribute
    - dynamic customer policy, 404–405
    - prefix origin tracking, 402
    - QPPB feature, 407, 410
    - static route redistribution, 411
  - BGP peering with customers, 398
  - customer connectivity, 398
    - customer BGP peering, 398
    - identifying customer prefixes, 399
    - static route redistribution, 399
  - peering, 400
  - security, 412
    - DDoS defense (case study), 419–421
    - graded route map dampening, 416
    - peer filtering, 413
    - public peering risks, 416–418
    - TCP MD5 signatures, 412

- service levels, 407
- tiers, 401
- transit connectivity, 400
- ISs (intermediate systems), 583

---

## J–K

---

- Keepalive messages, 27

---

## L

---

- label disposition, 446
- Label Distribution Protocol (LDP), 440
- Label Edge Router (LER), 438
- label exchange protocols, 443
- Label Forwarding Information Base (LFIB), 446
- Label Switch Router (LSR), 438
- label to IP path, 446
- label to label path, 446
- labels
  - binding information, 444
  - distribution, 486
  - Explicit Null, 448
  - FIB routes, 448
  - forwarding labeled packets, 446–448
  - Implicit Null, 448
  - installed in LFIBs, 447
  - label exchange, 440, 444–445
  - local/in, 444
  - MPLS, 439
  - propagation, 453–455
  - received, 444
  - reserved, 441
  - sending, neighbor send-label command, 471
  - shim headers, 440

- Layer 3, 435
- Layer 3 topology, 584
- Layer 3 VPNs, 435
- lc-detect, 43
- LDP (Label Distribution Protocol), 440, 443
- leaf, 40
- LER (Label Edge Router), 438
- LFIB (Label Forwarding Information Base), 446, 448
- link addressing, 397
- load balancing, 231
  - inbound traffic, 231, 243
  - multihomed networks, 240
  - multiple sessions to the same provider
    - eBGP multihop feature, 233
    - eBGP multipath feature, 235
  - outbound traffic, 232, 245
- load sharing
  - CEF, 44
  - LFIB, 448
  - per-packet, 45
- local addressing (IPv6), 566
- Local AS feature, 135–136
  - AS\_PATH updates, 135
  - integration (case study), 145–151
- local labels, 444
- LOCAL\_AS community, 20
- LOCAL\_PREF attribute, 19
- loopback addressing, 397, 567
- loop-free topology, 21
- looping
  - AS Override, 459
  - AS\_PATH, 458
  - forwarding, 334
  - loop-prevention mechanism (iBGP), 260
    - CLUSTER\_LIST attribute, 262–263
    - ORIGINATOR\_ID attribute, 261

- looping traffic during migration, 311
- loop-prevention mechanism (iBGP), 260
  - CLUSTER\_LIST attribute, 262–263
  - ORIGINATOR\_ID attribute, 261
- LSPs, 450
- LSRs (Label Switch Routers), 438
  - exchanging label information, 442
  - swap operation, 446

## M

- management VPNs, 498
- Martian address space, 229
- mBGP, interaction with MSDP, 537
- MDTs (Multicast Distribution Trees), 517
  - building, 519, 531
  - components, 520
  - shared trees, 518
- member autonomous systems, 294
- memory
  - BGP, 16, 48–50
  - BGP Router process, 16
  - process switching, 34
  - use estimation formulas, 51
- mesh groups, 546
- messages
  - BGP supported (overview), 27
  - SA (MSDP), 535
- migration, 324
  - confederation to route reflection (case study)
    - migration procedures, 366–377, 379–383
    - starting configuration, 362–365
  - iBGP full mesh to confederating (case study), 326–331, 334–343

- iBGP full mesh to route reflection (case study),
  - 312
  - configurations and RIBs, 312–319
  - final BGP configurations, 325
  - migration procedures, 318–322
- identifying starting and final network topologies, 308–309
- identifying starting router, 311
- preparatory steps, 307
- route reflection to confederation (case study)
  - migration procedures, 347–350, 353–361
  - starting configurations, 343–346
- traffic loss
  - concerns, 311
  - troubleshooting, 320
- migration plan (BGP network core design),
  - 199–201, 204–206
- MIX (Multicast Internet Exchange), 534
- MPLS (Multiprotocol Label Switching), 435
  - devices, 438
  - downstream distribution method, 441
  - forwarding labeled packets, 446–448
  - label exchange, 440, 444–445
  - labels, 439–441
  - LSP setup, 440, 444–445
  - overview, 438
  - shim headers, 440
  - terminology use in command outputs, 445
  - VPN services. *See* MPLS VPNs
- mpls label protocol both command, 441
- MPLS VPNs, 448
  - AS\_PATH manipulation, 457
    - Allow-AS, 460
    - AS Override, 458
  - automatic route filtering, 456
  - BGP confederations, 489
  - CSC VPN model, 481
    - full Internet routes, 482–484
    - hierarchical VPNs, 485, 488
  - deployment considerations, 490
    - convergence, 499–500
    - route target (RT) design, 497–498
    - scalability, 491–496
  - devices, 449
  - forwarding path, 450
  - inter-AS VPNs, 461
    - back-to-back VRF, 462–464
    - multihop multiprotocol eBGP, 470–476
    - non-VPN transit providers, 476–478
    - option comparison, 480–481
    - single-hop multiprotocol eBGP, 465–469
  - routing/forwarding instance, 451
  - VPNv4 and label propagation, 453–455
- MSDP (Multicast Source Discovery Protocol), 515
  - Anycast RP, 548
  - case study, 548–554
  - configuring, 536
  - deployment, 536
  - interaction with mBGP, 537
  - peer-RPF, 535
  - peer-RPF rules for incoming SA messages, 537
    - e(m)BGP sessions, 540
    - i(m)BGP sessions, 538
  - purpose of, 534
  - route reflection, 547
  - SA, 535
- mtree data structure, 39
- mtree
  - components, 40
  - data structure, 39
- MULTI\_EXIT\_DISC (MED) attribute, 19

- multicast delivery model, 517
  - groups (notation), 517
  - MDTs, 519, 531
  - overview, 515
  - shared trees, 518
  - source trees, 519
- Multicast Distribution Trees. *See* MDTs
- Multicast Internet Exchange (MIX), 534
- multicast NLRI in MP-BGP, 536
- Multicast Source Discovery Protocol. *See* MSDP
- multihoming, 222
  - multiple providers, 231
  - standard multihomed network design, 226
    - multiple border router, 229
    - single border router, 226–227
  - stub network design, 223
    - multiple border routers, 225
    - single border router, 224
- multihop multiprotocol eBGP, 470, 474
  - carrying remote addresses in iBGP, 475–476
  - label exchange, 472
- multiplier characters, 110–111
- multiprotocol BGP, 590–594
- multiprotocol eBGP, 466
- multiprotocol extensions, 435
  - extended community attribute, 436
    - Route Origin extended community, 437
    - RT extended community, 436
  - reachability attributes, 437
- Multiprotocol Label Switching. *See* MPLS
- neighbor next-hop unchanged command, 471
- neighbor send-label command, 471
- NEs (network elements), 583–584
- NetFlow switching, 37
- network command, 31
- network core layer (ISP networks), 389
- Network Descriptor Blocks (NDBs), 16
- network edge layer (ISP networks), 393–395
- network elements (NEs), 583–584
- Network Layer Reachability Information (NLRI), 437
- networks
  - designing
    - determining solution, 158
    - network core design. *See* BGP network core design
    - problem definition, 158
  - ISPs, 387. *See also* ISP networks
  - migration
    - confederation to route reflection (case study), 362–371, 373–383
    - iBGP full mesh to confederation (case study), 326–331, 334–343
    - iBGP full mesh to route reflection (case study), 312–322, 325
    - identifying starting and final network topologies, 308–309
    - identifying starting router, 311
    - preparatory steps, 307
    - route reflection to confederation (case study), 343–355, 357–361
    - traffic loss concerns, 311
  - MPLS, 439
  - MPLS VPNs, 448
    - AS\_PATH manipulation, 457–460
    - automatic route filtering, 456
    - BGP confederations, 489
    - CSC VPN model, 481–485, 488

## N

- 
- NDBs (Network Descriptor Blocks), 16, 32
  - neighbor advertisement-interval command, 500
  - neighbor CE4 allowas-in command, 460

- deployment considerations,
  - 490–492, 494–500
- devices, 449
- inter-AS VPNs, 461–467, 469–476,
  - 478–481
- routing/forwarding instance, 451
- VPNv4 route and label propagation,
  - 453–455
- MSDP (case study), 548–554
- PIM-SM, 531
- NEXT\_HOP attribute, 18
- next-hop reachability, 591
- next-hop-self command, 19, 329, 469
- next hop unchanged, 468
- NLRI (Network Layer Reachability Information),
  - 437
- NO\_ADVERTISE community, 20
- no bgp default route-target filter command, 465
- NO\_EXPORT community, 20
- no ip cef table consistency-check command, 42
- no ip prefix-list seq command, 114
- non-packet-based networks, 439
- non-VPN transit providers, 476–478
- Notification messages, 27
- NSAP address family configuration, 594
- NSF (Non-Stop Forwarding) feature, 87
- null adjacency, 41

## O

---

- Open messages, 27
- optimum switching, 36
- ORF (Outbound Route Filtering), 91, 96, 123
- ORIGIN attribute, 17

- ORIGINATOR\_ID attribute, 20, 261
- OSPF, max-metric feature, 86
- outbound filtering, 230
- outbound load balancing, 232
- Outbound Route Filtering (ORF), 91, 96, 123
- Overload bit (OL-bit), 85

## P

---

- P (Provider) routers, 449
- packet forwarding, 479
- packet-forwarding databases, 33
- packets
  - BGP reception process, 67
  - forwarding labeled, 446–448
- packet-switching, 446
- partial routing, 222, 229
- partitioning, 493–494
- path attributes, 17
- Path MTU Discovery (PMTUD), 65
- paths, 446
  - best-path selection, 25
  - BGP, memory use, 53
  - eBGP vs. iBGP, 26
  - forwarding, 450
  - selection, 19, 24
  - switching, 33
    - cache-based, 35–37
    - CEF, 38–39, 42–45
    - comparison of, 46
    - process, 33
- PE (Provider edge) routers, 449
  - AS\_PATH manipulation, 457
    - Allow-AS, 460
    - AS Override, 458

- automatic route filtering, 456
  - carrying loopback addresses in iBGP, 475
  - layers, 450
  - loopback addresses, 478
  - VRF, 451
- peer filtering (ISP networks), 413
- peer groups, 7
  - BGP dynamic peer group feature, 77
  - in route reflection environment, 292
  - peer templates, 77–78
  - update groups, 80–81
  - update replication, 75
- peer policy templates, 78
- peer session templates, 77
- peering, 400
  - filters, 238
  - private, 401
- peering point abuse
  - default routing, 416
  - GRE tunneling, 418
  - third-party next-hop traffic manipulation, 417
- peer-RPF rules
  - e(m)BGP sessions, 540–542
  - i(m)BGP sessions, 538
  - mesh groups, 546
  - no (m)BGP session, 543
- penultimate-hop LSRs, 446
- performance
  - BGP convergence, 62
    - defined, 62
    - examples, 62–63
    - optimization interdependencies, 82
    - queue optimization, 67–70, 72–73
    - TCP protocol concerns, 64–65
    - tuning (case study), 97–100, 103
    - update generation, 74–82
  - BGP network features, 83
    - failure mitigation, 83–87
    - prefix update optimization, 91–96
  - CEF, 38
  - MPLS VPNs, 492
  - process switching effects, 34
- per-packet load sharing, 44–45
- per-session, 44
- per-session load sharing, 44
- PIM (Protocol-Independent Multicast), 519, 531
- PIM-SM networks, 531
- PMTUD (Path MTU Discovery), 65
- policies
  - aggregation and deaggregation, 130–133
  - BPA (BGP policy accounting), 143–144
  - conditional advertisement, 123
    - examples, 124, 127–129
    - forms, 124
  - control techniques, 109
    - filter lists, 114–119
    - filter processing order, 123
    - policy lists, 122
    - regular expressions, 109–112
    - route maps, 120–121
  - customer (ISPs)
    - advertisement suppression communities, 405
    - local preference manipulation, 404
  - enforcing, 8
  - inbound traffic, 242
  - Local AS feature, 135–136
  - outbound traffic, 245
  - QoS (ISPs), 407, 410
  - QoS Policy Propagation, 138
    - configuring traffic lookup, 140
    - enforcing policing on an interface, 140
    - example, 141–142
    - setting FIB policy entries based on BGP tagging, 139

- policy lists, 122
- pop operation, 446
- prefix lists, 114–115
- prefixes
  - advertising, 256–257
  - communities, 20
  - distribution (hierarchical VPNs), 486
  - exception adjacencies, 41
  - FIB entries, 40
  - filtering for IPv6, 570
  - injecting, 18, 200
  - Martian address space, 229
  - outbound filtering, 230
  - standard peering filters, 239
  - update optimization, 91
    - BGP ORF feature, 96
    - BGP soft reconfiguration, 94
    - route flap dampening, 91–93
    - route refresh, 95
    - transmit (TX) side loop detection, 95
- private ASN, 399
- private communities, 20
- private peering, 400–401
- process switching, 33–34
- processing VPN packets, 455
- Protocol-Independent Multicast (PIM), 519
- protocols
  - label exchange, 443
  - MSDP
    - case study, 548–554
    - configuring, 536
    - SA, 535
  - stability, 6
- provider backbone convergence, 500
- Provider edge routers. *See* PE routers
- Provider routers. *See* P routers

- provider-based summarization, 237
- public peering, 400
- public peering security risks, 416
  - default routing, 416
  - GRE tunneling, 418
  - third-party next-hop traffic manipulation, 417
- punt adjacency, 41
- push operation, 446

---

## Q

---

- QoS
  - IPv4 vs. IPv6, 563
  - QoS Policy Propagation, 138–139
    - configuring traffic lookup, 140
    - enforcing policing on an interface, 140
    - example, 141–142
  - QoS Policy Propagation via BGP (QPPB), 407
  - QPPB (QoS Policy Propagation via BGP), 138, 407
    - configuring traffic lookup, 140
    - enforcing policing on an interface, 140
    - example, 141–142
    - setting FIB policy entries based on BGP tagging, 139

---

## R

---

- range characters, 110–112
- RDBs (Routing Descriptor Blocks), 16, 32
- RDs (Route Distinguishers), 435
  - designing for MPLS VPNs, 495
  - VPN-v4, 435



- reachability attributes, 437
- reachability between RRs and PE routers, 502
- read-only mode, 82
- received labels, 444
- redistribute command, 31
- redistribution (ASBR receiving behavior), 467
- redundancy
  - affected by migration, 333
  - clustering in RR environments, 259
- regular expressions, 109
  - AS\_PATH pattern matching, 113
  - components, 110–111
  - use in Cisco IOS software, 112
- reliability, 6, 222
- remote site aggregation, 192
- replication, 81
- RFC 1918 addressing, 229
- RIB (Routing Information Base), 9, 13, 32
  - control plane, 13
  - failures, 32
  - forwarding information base (FIB), 14
  - lookup operations supported by Cisco IOS, 33
- RIPv2, configuration example, 452
- RO\_Limit, 82
- route aggregation, 394
- route dampening, 394
- route flap dampening, 91–93
- route flapping, 87
- route maps, 120–122. *See also* policy lists
- Route Origin extended community, 437
- route processors. *See* RPs
- route reflection, 254
  - clustering, 259
  - vs. confederation, 303–304
  - design examples, 266–276
    - clustering design, 288
    - enabling deterministic-MED, 278
    - enforcing IGP metric settings, 286
    - iBGP mesh, 277
    - resetting next hop, 290
    - RR with peer groups, 292
    - setting IGP metrics, 279–285
  - functionality, 255
  - hierarchical route reflection, 264–265
  - loop-prevention mechanisms, 260–263
  - MSDP, 547
  - rules for prefix advertisement, 256–257
- route reflection-based network, 362
- route reflectors (RRs), 470
  - BGP-based DCN network design, 585
  - establishing eBGP connectivity, 473
- route refresh, 91, 95
- routers
  - aggregation, 391
  - ASBRs, 461
  - BGP-IGP routing exchange, 31
  - control plane, 13
  - core, 389
  - edge, 394
  - exchanging BGP capabilities, 29
  - forwarding plane, 13
  - MPLS networks, 438
  - MPLS VPNs, 449
- routes
  - aggregate, 455
  - automatic filtering, 456
  - ORIGIN attribute, 25
  - untagged, 455
- route-target import command, 499
- routing
  - default, 221, 416
  - filtering, 230
  - full Internet tables, 222
  - partial routes, 222
  - policies, 8

- Routing Descriptor Blocks (RDBs), 16
- Routing Information Base. *See* RIB
- routing/forwarding instance (VRF), 451
- RPs (route processors), 42
  - Anycast RP, 548
  - consistency checking, 43
- RR-based network, migrating to confederation-based networks (case study)
  - migration procedures, 347–350, 353–361
  - starting configurations, 343–346
- RRs (route reflectors), 470
  - carrying loopback addresses in iBGP, 475
  - designing RR-based MPLS VPNs, 492
  - establishing eBGP connectivity, 473
  - in MPLS VPNs, 494
  - partitioning, 493–494
  - tuning, 495
- RTs (route targets), 497
  - design examples for MPLS, 498
  - design examples for MPLS VPNs, 497
  - extended community, 436
- MPLS VPNs
  - RD design, 495–496
  - resource consumption on PE devices, 491
  - RR design, 492–494
- scan-lc, 43
- scan-rib, 43
- scan-rp, 43
- SDP, peer-RPF rules for incoming SA messages
  - e(m)BGP sessions, 542
  - mesh groups, 546
  - no (m)BGP session, 543
- security
  - IPv4 vs. IPv6, 563
  - ISP networks, 412
    - DDoS defense (case study), 419–421
    - graded route map dampening, 416
    - peer filtering, 413
    - public peering risks, 416–418
    - TCP MD5 signatures, 412
- selective packet discard (SPD), 69–70
- set metric-type internal command, 19
- shared trees, 518
- shim headers, 440
- show buffers command, 72
- show cef interface command, 47
- show interface statistics command, 47
- show ip bgp command, 25
- show ip bgp neighbor command, 28
- show ip bgp summary command, 16
- show ip cache verbose command, 35
- show ip cef command, 40
- show mpls forwarding command, 447
- show mpls ldp bindings command, 444
- show mpls ldp neighbor command, 443
- signatures (TCP MD5), 412
- single-hop multiprotocol eBGP, 465
  - next hop reset by next-hop set command, 469
- redistribution, 467

## S

---

- S field (frame-based label header), 440
- SA (Source Active) messages, 535–537
- SAFIs (subsequent address family identifiers), 28
- scalability, 7
  - CEF, 38
  - DCNs, 583
  - iBGP, 253
    - confederation, 294–303
    - method comparison, 303–304
    - route reflection. *See* route reflection

Site of Origin (SOO), 437  
 site-to-site convergence, 500  
 soft reconfiguration, 91, 94  
 SOO (Site of Origin), 437  
 source trees (SPTs), 519  
 sparse mode, 531  
 SPD (selective packet discard), 69–70  
 SPD headroom, 67  
 SPTs (source trees), 519  
 stability, 6  
 standard multihomed network, 226
     multiple border router, 229
     single border router, 226–227  
 static route redistribution, 399  
 streaming (multicast), 517  
 stub network multihomed design, 223
     multiple border routers, 225
     single border router, 224  
 stub network single-homed design, 223  
 subautonomous systems, 294  
 subsequent address family identifiers (SAFIs), 28  
 summarization (provider-based), 237  
 swap operation, 446  
 switching (MPLS networks), 438  
 switching paths, 33
     cache-based, 35–37
     CEF, 38–39, 42–45
     comparison of, 46
     process, 33–35  
 synchronization feature (BGP), 396  
 system buffers, 67
     default values, 73
     show buffers command, 72  
 system local addressing, 229

## T

TCP, 65  
 TCP MD5 signatures, 412  
 TDP (Tag Distribution Protocol), 440  
 test network addressing, 230  
 third-party next hop, 18  
 third-party next-hop traffic manipulation, 417  
 tiers (ISPs), 401  
 topologies
     back-to-back VRF, 463
     IPv6 network, 574
     loop-free, 21
     PIM-SM networks, 527
     VPNs, 497  
 traffic
     BGP, worst-case input queue values, 69
     inbound traffic policies, 242
     load balancing, 231
         inbound, 231
         multiple sessions to the same provider, 233–235
         outbound, 232
     loss
         avoiding during migration, 320
         concerns during network migration, 311
     multicast, 515
     outbound traffic policies, 245  
 transit connectivity (ISP networks), 400  
 transmit (TX) side loop detection, 91, 95  
 troubleshooting
     avoiding traffic loss during migration, 311
     bug toolkit, 599
     traffic loss during migration, 320
     tuning RRs, 495  
 TTL field (frame-based label header), 440

## U

---

- unsolicited downstream, 441
- unspecified addresses, 567
- untagged routes, 455
- update groups, 7, 80–81
- Update messages, 27
- update packing enhancement, 81

## V

---

- VCI (virtual circuit identifier), 440
- verifying
  - BGP reachability for all prefixes (network migration), 324
  - ES-IS adjacency, 590
  - routing information during migration, 360
- verifying BGP reachability for all prefixes, 324
- VIPs (Versatile Interface Processors), 36
- VPI (virtual path identifier), 440
- VPNs, 435
  - extranet, 497
  - hierarchical, 485, 488
  - hub sites, 460
  - management, 498
- VPNv4, 435
  - addresses, receiving, 436
  - label distribution, 476
  - multihop multiprotocol eBGP, 470
- VRF (routing/forwarding instance), 451, 462–464

## W–Y

---

- Web sites
  - Cisco bug toolkit, 599
  - Cisco Systems, 606
  - regional registries, 226
- WEIGHT parameter, 19, 25
- well-known communities, 20

## Z

---

- zeroing BGP MEDs, 394