

VYATTA, INC.



Vyatta System

BGP

REFERENCE GUIDE



Vyatta
Suite 200
1301 Shoreway Road
Belmont, CA 94002
vyatta.com
650 413 7200
1 888 VYATTA 1 (US and Canada)

COPYRIGHT

Copyright © 2005–2012 Vyatta, Inc. All rights reserved.

Vyatta reserves the right to make changes to software, hardware, and documentation without notice. For the most recent version of documentation, visit the Vyatta web site at vyatta.com.

PROPRIETARY NOTICES

Vyatta is a registered trademark of Vyatta, Inc.

Hyper-V is a registered trademark of Microsoft Corporation.

VMware, VMware ESX, and VMware server are trademarks of VMware, Inc.

XenServer, and XenCenter are trademarks of Citrix Systems, Inc.

All other trademarks are the property of their respective owners.

RELEASE DATE: October 2012

DOCUMENT REVISION. 6.5R1 v01

RELEASED WITH: 6.5R1

PART NO. A0-0216-10-0016

Contents

Quick List of Commands	xii
List of Examples	xix
Preface	xx
Intended Audience	xxi
Organization of This Guide	xxi
Document Conventions	xxiii
Vyatta Publications	xxiv
Chapter 1 BGP Configuration	1
BGP Overview	2
iBGP and eBGP	3
iBGP	3
eBGP	4
BGP ID Selection Process	5
BGP Path Selection Process	6
Scalability of BGP	7
Confederations	7
Route Reflection	8
Route Flapping and Flap Damping	10
AS Paths	11
BGP Communities	12
Peer Groups	13
BGP Multipath Support	13
IPv4 and IPv6 Support	13
Supported Standards	14
Configuring BGP	14
Basic iBGP Configuration	15
Verifying the iBGP Configuration	24
R1: show ip bgp summary	24
R1: show ip bgp	25
Basic eBGP Configuration	25
Verifying the eBGP Configuration	27
R1: show ip bgp summary	27
R1: show ip bgp	28
Originating a Route to eBGP Neighbors	28

Verifying the Route Origination	30
R1: show ip bgp summary	30
R1: show ip bgp	31
R1: show ip route bgp	32
R4: show ip bgp summary	33
R4: show ip bgp	33
Inbound Route Filtering	34
Verifying the Inbound Filter	39
R1: show ip bgp	39
R1: show ip bgp	39
R4: show ip bgp	40
R4: show ip bgp	41
Outbound Route Filtering	41
Verifying the Outbound Filter	45
AS 200: show ip bgp	45
AS 200: show ip bgp	46
Confederations	46
Verifying the Confederation	54
R1: show ip bgp summary	54
R1: show ip bgp	55
R2: show ip bgp summary	56
R2: show ip bgp	56
R3: show ip bgp summary	57
R3: show ip bgp	57
R4: show ip bgp summary	58
R4: show ip bgp	58
Route Reflectors	59
Verifying the Route Reflector	65
R1: show ip bgp summary	65
R1: show ip bgp	66
R2: show ip bgp summary	66
R2: show ip bgp	67
R3: show ip bgp summary	67
R3: show ip bgp	68
R4: show ip bgp summary	69
R4: show ip bgp	69
Route Redirection	70
Configuring IPv6 Routing with BGP	70
Enable Forwarding on R1 and R2	71
Configure IPv6 BGP Peer	71
Advertise Connected Networks	73

Confirm Advertised Routes	73
Chapter 2 Global and Router-Specific Configuration	75
Global and Router-Specific Commands	76
monitor protocol bgp disable	81
monitor protocol bgp disable all	82
monitor protocol bgp disable events	83
monitor protocol bgp disable filters	84
monitor protocol bgp disable fsm	85
monitor protocol bgp disable keepalives	86
monitor protocol bgp disable rib	87
monitor protocol bgp disable updates	88
monitor protocol bgp enable	89
monitor protocol bgp enable events	90
monitor protocol bgp enable filters	91
monitor protocol bgp enable fsm	92
monitor protocol bgp enable keepalives	93
monitor protocol bgp enable rib	94
monitor protocol bgp enable updates	95
protocols bgp <asn>	96
protocols bgp <asn> address-family ipv6-unicast	98
protocols bgp <asn> address-family ipv6-unicast aggregate-address <ipv6net>	99
protocols bgp <asn> address-family ipv6-unicast network <ipv6net>	101
protocols bgp <asn> aggregate-address <ipv4net>	103
protocols bgp <asn> maximum-paths ebgp <max-paths>	105
protocols bgp <asn> maximum-paths ibgp <max-paths>	107
protocols bgp <asn> network <ipv4net>	109
protocols bgp <asn> parameters always-compare-med	111
protocols bgp <asn> parameters bestpath as-path	113
protocols bgp <asn> parameters bestpath compare-routerid	115
protocols bgp <asn> parameters bestpath med	117
protocols bgp <asn> parameters dampening	119
protocols bgp <asn> parameters default	121
protocols bgp <asn> parameters deterministic-med	123
protocols bgp <asn> parameters distance global	124
protocols bgp <asn> parameters distance prefix <ipv4net> distance <dist>	126
protocols bgp <asn> parameters disable-network-import-check	128
protocols bgp <asn> parameters enforce-first-as	130
protocols bgp <asn> parameters graceful-restart	132
protocols bgp <asn> parameters log-neighbor-changes	134
protocols bgp <asn> parameters no-fast-external-failover	136

protocols bgp <asn> parameters router-id <id>	138
protocols bgp <asn> parameters scan-time <seconds>	140
protocols bgp <asn> timers.	142
reset ip bgp <address>	144
reset ip bgp <address> ipv4 unicast	146
reset ip bgp all	148
reset ip bgp dampening.	149
show ip bgp	150
show ip bgp attribute-info	151
show ip bgp cidr-only.	152
show ip bgp community-info	153
show ip bgp community <community>	154
show ip bgp community-list <list-name>.	155
show ip bgp dampened-paths	156
show ip bgp filter-list <as-path-list-name>	157
show ip bgp flap-statistics.	158
show ip bgp flap-statistics cidr-only.	159
show ip bgp flap-statistics filter-list <as-path-list-name>.	160
show ip bgp flap-statistics prefix-list <list-name>.	161
show ip bgp flap-statistics regexp <expr>	162
show ip bgp flap-statistics route-map <map-name>	163
show ip bgp ipv4 unicast	164
show ip bgp ipv4 unicast cidr-only.	165
show ip bgp ipv4 unicast community <community>	166
show ip bgp ipv4 unicast community-list <list-name>	167
show ip bgp ipv4 unicast filter-list <as-path-list-name>.	168
show ip bgp ipv4 unicast paths.	169
show ip bgp ipv4 unicast prefix-list <list-name>	170
show ip bgp ipv4 unicast regexp <regexp>	171
show ip bgp ipv4 unicast route-map <map-name>	172
show ip bgp ipv4 unicast statistics.	173
show ip bgp ipv4 unicast summary	174
show ip bgp memory	175
show ip bgp paths	176
show ip bgp prefix-list <list-name>	177
show ip bgp regexp <regexp>	178
show ip bgp route-map <map-name>	179
show ip bgp scan	180
show ip route bgp	181
show ipv6 bgp	183
show ipv6 bgp community <community>	184

show ipv6 bgp community-list <list-name>	185
show ipv6 bgp filter-list <as-path-list-name>	186
show ipv6 bgp prefix-list <list-name>	187
show ipv6 bgp regexp <regexp>	188
show monitoring protocols bgp	189
Chapter 3 Route Reflection	190
Route Reflection Commands	191
protocols bgp <asn> neighbor <id> address-family ipv6-unicast route-reflector-client	192
protocols bgp <asn> neighbor <id> route-reflector-client	194
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast route-reflector-client	197
protocols bgp <asn> peer-group <group-name> route-reflector-client	199
protocols bgp <asn> parameters cluster-id <id>	201
protocols bgp <asn> parameters no-client-to-client-reflection	203
Chapter 4 Confederations	205
Confederation Commands	206
protocols bgp <asn> parameters confederation identifier <asn>	207
protocols bgp <asn> parameters confederation peers <asn>	209
Chapter 5 Neighbors	211
Neighbor Commands	212
protocols bgp <asn> neighbor <id>	217
protocols bgp <asn> neighbor <id> address-family ipv6-unicast	219
protocols bgp <asn> neighbor <id> address-family ipv6-unicast allowas-in	220
protocols bgp <asn> neighbor <id> address-family ipv6-unicast attribute-unchanged	222
protocols bgp <asn> neighbor <id> address-family ipv6-unicast capability dynamic	224
protocols bgp <asn> neighbor <id> address-family ipv6-unicast capability orf	226
protocols bgp <asn> neighbor <id> address-family ipv6-unicast default-originate	228
protocols bgp <asn> neighbor <id> address-family ipv6-unicast disable-send-community	230
protocols bgp <asn> neighbor <id> address-family ipv6-unicast distribute-list export <access-list6-name>	232
protocols bgp <asn> neighbor <id> address-family ipv6-unicast distribute-list import <access-list6-name>	234
protocols bgp <asn> neighbor <id> address-family ipv6-unicast filter-list export <as-path-list-name>	236
protocols bgp <asn> neighbor <id> address-family ipv6-unicast filter-list import <as-path-list-name>	238
protocols bgp <asn> neighbor <id> address-family ipv6-unicast maximum-prefix <max-num>	240
protocols bgp <asn> neighbor <id> address-family ipv6-unicast nexthop-local unchanged	242
protocols bgp <asn> neighbor <id> address-family ipv6-unicast nexthop-self	244
protocols bgp <asn> neighbor <id> address-family ipv6-unicast peer-group <group-name>	246
protocols bgp <asn> neighbor <id> address-family ipv6-unicast prefix-list export <prefix-list6-name>	248
protocols bgp <asn> neighbor <id> address-family ipv6-unicast prefix-list import <prefix-list6-name>	250
protocols bgp <asn> neighbor <id> address-family ipv6-unicast remove-private-as	252

protocols bgp <asn> neighbor <id> address-family ipv6-unicast route-map export <map-name>	254
protocols bgp <asn> neighbor <id> address-family ipv6-unicast route-map import <map-name>	256
protocols bgp <asn> neighbor <id> address-family ipv6-unicast soft-reconfiguration inbound	258
protocols bgp <asn> neighbor <id> address-family ipv6-unicast unsuppress-map <map-name>	260
protocols bgp <asn> neighbor <id> advertisement-interval <seconds>	262
protocols bgp <asn> neighbor <id> allowas-in	264
protocols bgp <asn> neighbor <id> attribute-unchanged	266
protocols bgp <asn> neighbor <id> capability dynamic	268
protocols bgp <asn> neighbor <id> capability orf	270
protocols bgp <asn> neighbor <id> default-originate	272
protocols bgp <asn> neighbor <id> description <desc>	274
protocols bgp <asn> neighbor <id> disable-capability-negotiation	276
protocols bgp <asn> neighbor <id> disable-connected-check	278
protocols bgp <asn> neighbor <id> disable-send-community	280
protocols bgp <asn> neighbor <id> distribute-list export <acl-num>	282
protocols bgp <asn> neighbor <id> distribute-list import <acl-num>	284
protocols bgp <asn> neighbor <id> ebgp-multihop <ttl>	286
protocols bgp <asn> neighbor <id> filter-list export <as-path-list-name>	288
protocols bgp <asn> neighbor <id> filter-list import <as-path-list-name>	290
protocols bgp <asn> neighbor <id> local-as <asn>	292
protocols bgp <asn> neighbor <id> maximum-prefix <max-num>	294
protocols bgp <asn> neighbor <id> nexthop-self	296
protocols bgp <asn> neighbor <id> override-capability	298
protocols bgp <asn> neighbor <id> passive	300
protocols bgp <asn> neighbor <id> password <pwd>	302
protocols bgp <asn> neighbor <id> peer-group <group-name>	304
protocols bgp <asn> neighbor <id> port <port-num>	306
protocols bgp <asn> neighbor <id> prefix-list export <list-name>	308
protocols bgp <asn> neighbor <id> prefix-list import <list-name>	310
protocols bgp <asn> neighbor <id> remote-as <asn>	312
protocols bgp <asn> neighbor <id> remove-private-as	314
protocols bgp <asn> neighbor <id> route-map export <map-name>	316
protocols bgp <asn> neighbor <id> route-map import <map-name>	318
protocols bgp <asn> neighbor <id> shutdown	320
protocols bgp <asn> neighbor <id> soft-reconfiguration inbound	322
protocols bgp <asn> neighbor <id> strict-capability-match	324
protocols bgp <asn> neighbor <id> timers	326
protocols bgp <asn> neighbor <id> ttl-security hops <hops>	328
protocols bgp <asn> neighbor <id> unsuppress-map <map-name>	330
protocols bgp <asn> neighbor <id> update-source <source>	332
protocols bgp <asn> neighbor <id> weight <weight>	334

reset ip bgp external	336
reset ip bgp external ipv4 unicast.	338
show ip bgp ipv4 unicast neighbors	340
show ip bgp ipv4 unicast neighbors <id> advertised-routes	341
show ip bgp ipv4 unicast neighbors <id> prefix-counts	342
show ip bgp ipv4 unicast neighbors <id> received prefix-filter	343
show ip bgp ipv4 unicast neighbors <id> received-routes	344
show ip bgp ipv4 unicast neighbors <id> routes	345
show ip bgp neighbors.	346
show ip bgp neighbors <id> advertised-routes	347
show ip bgp neighbors <id> dampened-routes.	348
show ip bgp neighbors <id> flap-statistics.	349
show ip bgp neighbors <id> prefix-counts.	350
show ip bgp neighbors <id> received prefix-filter.	351
show ip bgp neighbors <id> received-routes.	352
show ip bgp neighbors <id> routes	353
show ipv6 bgp neighbors.	354
show ipv6 bgp neighbors <ipv6> advertised-routes	355
show ipv6 bgp neighbors <ipv6> received-routes	356
show ipv6 bgp neighbors <ipv6> routes	357
Chapter 6 Peer Groups	358
Peer Group Commands	359
protocols bgp <asn> peer-group <group-name>	363
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast.	365
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast allowas-in.	366
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast attribute-unchanged.	368
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast capability dynamic	370
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast capability orf	372
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast default-originate	374
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast disable-send-community	376
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast distribute-list export <access-list6-name>	378
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast distribute-list import <access-list6-name>	380
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast filter-list export <as-path-list-name> 382	
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast filter-list import <as-path-list-name>.	384
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast maximum-prefix <max-num>.	386
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast nexthop-local unchanged.	388

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast nexthop-self	390
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast prefix-list export <prefix-list6-name>	392
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast prefix-list import <prefix-list6-name>	394
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast remove-private-as	396
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast route-map export <map-name> . . 398	
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast route-map import <map-name> . . 400	
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast soft-reconfiguration inbound .	402
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast unsuppress-map <map-name>	404
protocols bgp <asn> peer-group <group-name> allowas-in	406
protocols bgp <asn> peer-group <group-name> attribute-unchanged	408
protocols bgp <asn> peer-group <group-name> capability dynamic	410
protocols bgp <asn> peer-group <group-name> capability orf	412
protocols bgp <asn> peer-group <group-name> default-originate	414
protocols bgp <asn> peer-group <group-name> description <desc>	416
protocols bgp <asn> peer-group <group-name> disable-capability-negotiation	418
protocols bgp <asn> peer-group <group-name> disable-connected-check	420
protocols bgp <asn> peer-group <group-name> disable-send-community	422
protocols bgp <asn> peer-group <group-name> distribute-list export <acl-num>	424
protocols bgp <asn> peer-group <group-name> distribute-list import <acl-num>	426
protocols bgp <asn> peer-group <group-name> ebgp-multihop <ttl>	428
protocols bgp <asn> peer-group <group-name> filter-list export <as-path-list-name>	430
protocols bgp <asn> peer-group <group-name> filter-list import <as-path-list-name>	432
protocols bgp <asn> peer-group <group-name> local-as <asn>	434
protocols bgp <asn> peer-group <group-name> maximum-prefix <max-num>	436
protocols bgp <asn> peer-group <group-name> nexthop-self	438
protocols bgp <asn> peer-group <group-name> override-capability	440
protocols bgp <asn> peer-group <group-name> passive	442
protocols bgp <asn> peer-group <group-name> password <pwd>	444
protocols bgp <asn> peer-group <group-name> prefix-list export <list-name>	446
protocols bgp <asn> peer-group <group-name> prefix-list import <list-name>	448
protocols bgp <asn> peer-group <group-name> remote-as <asn>	450
protocols bgp <asn> peer-group <group-name> remove-private-as	452
protocols bgp <asn> peer-group <group-name> route-map export <map-name>	454
protocols bgp <asn> peer-group <group-name> route-map import <map-name>	456
protocols bgp <asn> peer-group <group-name> shutdown	458
protocols bgp <asn> peer-group <group-name> soft-reconfiguration inbound	460
protocols bgp <asn> peer-group <group-name> ttl-security hops <hops>	462

protocols bgp <asn> peer-group <group-name> unsuppress-map <map-name>	464
protocols bgp <asn> peer-group <group-name> update-source <source>	466
protocols bgp <asn> peer-group <group-name> weight <weight>	468
reset ip bgp peer-group <group-name>	470
reset ip bgp peer-group <group-name> ipv4 unicast	472
Chapter 7 Route Redistribution	474
Route Redistribution Commands	475
protocols bgp <asn> address-family ipv6-unicast redistribute connected	476
protocols bgp <asn> address-family ipv6-unicast redistribute kernel	478
protocols bgp <asn> address-family ipv6-unicast redistribute ospfv3	480
protocols bgp <asn> address-family ipv6-unicast redistribute ripng	482
protocols bgp <asn> address-family ipv6-unicast redistribute static	484
protocols bgp <asn> redistribute connected	486
protocols bgp <asn> redistribute kernel	488
protocols bgp <asn> redistribute ospf	490
protocols bgp <asn> redistribute rip	492
protocols bgp <asn> redistribute static	494
Chapter 8 Route Server	496
Route Server Commands	497
protocols bgp <asn> neighbor <id> address-family ipv6-unicast route-server-client	498
protocols bgp <asn> neighbor <id> route-server-client	500
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast route-server-client	502
protocols bgp <asn> peer-group <group-name> route-server-client	504
show ip bgp ipv4 unicast rsclient summary	506
show ip bgp rsclient <address>	507
Glossary of Acronyms	508

Quick List of Commands

Use this list to help you quickly locate commands.

monitor protocol bgp disable all	82
monitor protocol bgp disable events	83
monitor protocol bgp disable filters	84
monitor protocol bgp disable fsm	85
monitor protocol bgp disable keepalives	86
monitor protocol bgp disable rib	87
monitor protocol bgp disable updates	88
monitor protocol bgp disable	81
monitor protocol bgp enable events	90
monitor protocol bgp enable filters	91
monitor protocol bgp enable fsm	92
monitor protocol bgp enable keepalives	93
monitor protocol bgp enable rib	94
monitor protocol bgp enable updates	95
monitor protocol bgp enable	89
protocols bgp <asn> address-family ipv6-unicast aggregate-address <ipv6net>	99
protocols bgp <asn> address-family ipv6-unicast network <ipv6net>	101
protocols bgp <asn> address-family ipv6-unicast redistribute connected	476
protocols bgp <asn> address-family ipv6-unicast redistribute kernel	478
protocols bgp <asn> address-family ipv6-unicast redistribute ospfv3	480
protocols bgp <asn> address-family ipv6-unicast redistribute ripng	482
protocols bgp <asn> address-family ipv6-unicast redistribute static	484
protocols bgp <asn> address-family ipv6-unicast	98
protocols bgp <asn> aggregate-address <ipv4net>	103
protocols bgp <asn> maximum-paths ebgp <max-paths>	105
protocols bgp <asn> maximum-paths ibgp <max-paths>	107
protocols bgp <asn> neighbor <id> address-family ipv6-unicast allowas-in	220
protocols bgp <asn> neighbor <id> address-family ipv6-unicast attribute-unchanged	222
protocols bgp <asn> neighbor <id> address-family ipv6-unicast capability dynamic	224
protocols bgp <asn> neighbor <id> address-family ipv6-unicast capability orf	226
protocols bgp <asn> neighbor <id> address-family ipv6-unicast default-originate	228
protocols bgp <asn> neighbor <id> address-family ipv6-unicast disable-send-community	230

protocols bgp <asn> neighbor <id> address-family ipv6-unicast distribute-list export <access-list6-name>	232
protocols bgp <asn> neighbor <id> address-family ipv6-unicast distribute-list import <access-list6-name>	234
protocols bgp <asn> neighbor <id> address-family ipv6-unicast filter-list export <as-path-list-name>	236
protocols bgp <asn> neighbor <id> address-family ipv6-unicast filter-list import <as-path-list-name>	238
protocols bgp <asn> neighbor <id> address-family ipv6-unicast maximum-prefix <max-num>	240
protocols bgp <asn> neighbor <id> address-family ipv6-unicast nexthop-local unchanged	242
protocols bgp <asn> neighbor <id> address-family ipv6-unicast nexthop-self	244
protocols bgp <asn> neighbor <id> address-family ipv6-unicast peer-group <group-name>	246
protocols bgp <asn> neighbor <id> address-family ipv6-unicast prefix-list export <prefix-list6-name>	248
protocols bgp <asn> neighbor <id> address-family ipv6-unicast prefix-list import <prefix-list6-name>	250
protocols bgp <asn> neighbor <id> address-family ipv6-unicast remove-private-as	252
protocols bgp <asn> neighbor <id> address-family ipv6-unicast route-map export <map-name>	254
protocols bgp <asn> neighbor <id> address-family ipv6-unicast route-map import <map-name>	256
protocols bgp <asn> neighbor <id> address-family ipv6-unicast route-reflector-client	192
protocols bgp <asn> neighbor <id> address-family ipv6-unicast route-server-client	498
protocols bgp <asn> neighbor <id> address-family ipv6-unicast soft-reconfiguration inbound	258
protocols bgp <asn> neighbor <id> address-family ipv6-unicast unsuppress-map <map-name>	260
protocols bgp <asn> neighbor <id> address-family ipv6-unicast	219
protocols bgp <asn> neighbor <id> advertisement-interval <seconds>	262
protocols bgp <asn> neighbor <id> allowas-in	264
protocols bgp <asn> neighbor <id> attribute-unchanged	266
protocols bgp <asn> neighbor <id> capability dynamic	268
protocols bgp <asn> neighbor <id> capability orf	270
protocols bgp <asn> neighbor <id> default-originate	272
protocols bgp <asn> neighbor <id> description <desc>	274
protocols bgp <asn> neighbor <id> disable-capability-negotiation	276
protocols bgp <asn> neighbor <id> disable-connected-check	278
protocols bgp <asn> neighbor <id> disable-send-community	280
protocols bgp <asn> neighbor <id> distribute-list export <acl-num>	282
protocols bgp <asn> neighbor <id> distribute-list import <acl-num>	284
protocols bgp <asn> neighbor <id> ebgp-multihop <ttl>	286
protocols bgp <asn> neighbor <id> filter-list export <as-path-list-name>	288
protocols bgp <asn> neighbor <id> filter-list import <as-path-list-name>	290
protocols bgp <asn> neighbor <id> local-as <asn>	292
protocols bgp <asn> neighbor <id> maximum-prefix <max-num>	294
protocols bgp <asn> neighbor <id> nexthop-self	296

protocols bgp <asn> neighbor <id> override-capability	298
protocols bgp <asn> neighbor <id> passive	300
protocols bgp <asn> neighbor <id> password <pwd>	302
protocols bgp <asn> neighbor <id> peer-group <group-name>	304
protocols bgp <asn> neighbor <id> port <port-num>	306
protocols bgp <asn> neighbor <id> prefix-list export <list-name>	308
protocols bgp <asn> neighbor <id> prefix-list import <list-name>	310
protocols bgp <asn> neighbor <id> remote-as <asn>	312
protocols bgp <asn> neighbor <id> remove-private-as	314
protocols bgp <asn> neighbor <id> route-map export <map-name>	316
protocols bgp <asn> neighbor <id> route-map import <map-name>	318
protocols bgp <asn> neighbor <id> route-reflector-client	194
protocols bgp <asn> neighbor <id> route-server-client	500
protocols bgp <asn> neighbor <id> shutdown	320
protocols bgp <asn> neighbor <id> soft-reconfiguration inbound	322
protocols bgp <asn> neighbor <id> strict-capability-match	324
protocols bgp <asn> neighbor <id> timers	326
protocols bgp <asn> neighbor <id> ttl-security hops <hops>	328
protocols bgp <asn> neighbor <id> unsuppress-map <map-name>	330
protocols bgp <asn> neighbor <id> update-source <source>	332
protocols bgp <asn> neighbor <id> weight <weight>	334
protocols bgp <asn> neighbor <id>	217
protocols bgp <asn> network <ipv4net>	109
protocols bgp <asn> parameters always-compare-med	111
protocols bgp <asn> parameters bestpath as-path	113
protocols bgp <asn> parameters bestpath compare-routerid	115
protocols bgp <asn> parameters bestpath med	117
protocols bgp <asn> parameters cluster-id <id>	201
protocols bgp <asn> parameters confederation identifier <asn>	207
protocols bgp <asn> parameters confederation peers <asn>	209
protocols bgp <asn> parameters dampening	119
protocols bgp <asn> parameters default	121
protocols bgp <asn> parameters deterministic-med	123
protocols bgp <asn> parameters disable-network-import-check	128
protocols bgp <asn> parameters distance global	124
protocols bgp <asn> parameters distance prefix <ipv4net> distance <dist>	126

protocols bgp <asn> parameters enforce-first-as	130
protocols bgp <asn> parameters graceful-restart	132
protocols bgp <asn> parameters log-neighbor-changes	134
protocols bgp <asn> parameters no-client-to-client-reflection	203
protocols bgp <asn> parameters no-fast-external-failover	136
protocols bgp <asn> parameters router-id <id>	138
protocols bgp <asn> parameters scan-time <seconds>	140
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast allowas-in	366
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast attribute-unchanged	368
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast capability dynamic	370
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast capability orf	372
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast default-originate	374
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast disable-send-community	376
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast distribute-list export <access-list6-name>	378
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast distribute-list import <access-list6-name>	380
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast filter-list export <as-path-list-name> ..	382
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast filter-list import <as-path-list-name> ..	384
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast maximum-prefix <max-num>	386
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast nexthop-local unchanged	388
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast nexthop-self	390
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast prefix-list export <prefix-list6-name> ..	392
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast prefix-list import <prefix-list6-name> ..	394
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast remove-private-as	396
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast route-map export <map-name> ..	398
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast route-map import <map-name> ..	400
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast route-reflector-client	197
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast route-server-client	502
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast soft-reconfiguration inbound	402
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast unsuppress-map <map-name> ...	404
protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast	365
protocols bgp <asn> peer-group <group-name> allowas-in	406

protocols bgp <asn> peer-group <group-name> attribute-unchanged	408
protocols bgp <asn> peer-group <group-name> capability dynamic	410
protocols bgp <asn> peer-group <group-name> capability orf	412
protocols bgp <asn> peer-group <group-name> default-originate	414
protocols bgp <asn> peer-group <group-name> description <desc>	416
protocols bgp <asn> peer-group <group-name> disable-capability-negotiation	418
protocols bgp <asn> peer-group <group-name> disable-connected-check	420
protocols bgp <asn> peer-group <group-name> disable-send-community	422
protocols bgp <asn> peer-group <group-name> distribute-list export <acl-num>	424
protocols bgp <asn> peer-group <group-name> distribute-list import <acl-num>	426
protocols bgp <asn> peer-group <group-name> ebgp-multihop <ttl>	428
protocols bgp <asn> peer-group <group-name> filter-list export <as-path-list-name>	430
protocols bgp <asn> peer-group <group-name> filter-list import <as-path-list-name>	432
protocols bgp <asn> peer-group <group-name> local-as <asn>	434
protocols bgp <asn> peer-group <group-name> maximum-prefix <max-num>	436
protocols bgp <asn> peer-group <group-name> nexthop-self	438
protocols bgp <asn> peer-group <group-name> override-capability	440
protocols bgp <asn> peer-group <group-name> passive	442
protocols bgp <asn> peer-group <group-name> password <pwd>	444
protocols bgp <asn> peer-group <group-name> prefix-list export <list-name>	446
protocols bgp <asn> peer-group <group-name> prefix-list import <list-name>	448
protocols bgp <asn> peer-group <group-name> remote-as <asn>	450
protocols bgp <asn> peer-group <group-name> remove-private-as	452
protocols bgp <asn> peer-group <group-name> route-map export <map-name>	454
protocols bgp <asn> peer-group <group-name> route-map import <map-name>	456
protocols bgp <asn> peer-group <group-name> route-reflector-client	199
protocols bgp <asn> peer-group <group-name> route-server-client	504
protocols bgp <asn> peer-group <group-name> shutdown	458
protocols bgp <asn> peer-group <group-name> soft-reconfiguration inbound	460
protocols bgp <asn> peer-group <group-name> ttl-security hops <hops>	462
protocols bgp <asn> peer-group <group-name> unsuppress-map <map-name>	464
protocols bgp <asn> peer-group <group-name> update-source <source>	466
protocols bgp <asn> peer-group <group-name> weight <weight>	468
protocols bgp <asn> peer-group <group-name>	363
protocols bgp <asn> redistribute connected	486
protocols bgp <asn> redistribute kernel	488

protocols bgp <asn> redistribute ospf	490
protocols bgp <asn> redistribute rip	492
protocols bgp <asn> redistribute static	494
protocols bgp <asn> timers	142
protocols bgp <asn>	96
reset ip bgp <address> ipv4 unicast	146
reset ip bgp <address>	144
reset ip bgp all	148
reset ip bgp dampening	149
reset ip bgp external ipv4 unicast	338
reset ip bgp external	336
reset ip bgp peer-group <group-name> ipv4 unicast	472
reset ip bgp peer-group <group-name>	470
show ip bgp attribute-info	151
show ip bgp cidr-only	152
show ip bgp community <community>	154
show ip bgp community-info	153
show ip bgp community-list <list-name>	155
show ip bgp dampened-paths	156
show ip bgp filter-list <as-path-list-name>	157
show ip bgp flap-statistics cidr-only	159
show ip bgp flap-statistics filter-list <as-path-list-name>	160
show ip bgp flap-statistics prefix-list <list-name>	161
show ip bgp flap-statistics regexp <expr>	162
show ip bgp flap-statistics route-map <map-name>	163
show ip bgp flap-statistics	158
show ip bgp ipv4 unicast cidr-only	165
show ip bgp ipv4 unicast community <community>	166
show ip bgp ipv4 unicast community-list <list-name>	167
show ip bgp ipv4 unicast filter-list <as-path-list-name>	168
show ip bgp ipv4 unicast neighbors <id> advertised-routes	341
show ip bgp ipv4 unicast neighbors <id> prefix-counts	342
show ip bgp ipv4 unicast neighbors <id> received prefix-filter	343
show ip bgp ipv4 unicast neighbors <id> received-routes	344
show ip bgp ipv4 unicast neighbors <id> routes	345
show ip bgp ipv4 unicast neighbors	340

show ip bgp ipv4 unicast paths	169
show ip bgp ipv4 unicast prefix-list <list-name>	170
show ip bgp ipv4 unicast regexp <regexp>	171
show ip bgp ipv4 unicast route-map <map-name>	172
show ip bgp ipv4 unicast rsclient summary	506
show ip bgp ipv4 unicast statistics	173
show ip bgp ipv4 unicast summary	174
show ip bgp ipv4 unicast	164
show ip bgp memory	175
show ip bgp neighbors <id> advertised-routes	347
show ip bgp neighbors <id> dampened-routes	348
show ip bgp neighbors <id> flap-statistics	349
show ip bgp neighbors <id> prefix-counts	350
show ip bgp neighbors <id> received prefix-filter	351
show ip bgp neighbors <id> received-routes	352
show ip bgp neighbors <id> routes	353
show ip bgp neighbors	346
show ip bgp paths	176
show ip bgp prefix-list <list-name>	177
show ip bgp regexp <regexp>	178
show ip bgp route-map <map-name>	179
show ip bgp rsclient <address>	507
show ip bgp scan	180
show ip bgp	150
show ip route bgp	181
show ipv6 bgp community <community>	184
show ipv6 bgp community-list <list-name>	185
show ipv6 bgp filter-list <as-path-list-name>	186
show ipv6 bgp neighbors <ipv6> advertised-routes	355
show ipv6 bgp neighbors <ipv6> received-routes	356
show ipv6 bgp neighbors <ipv6> routes	357
show ipv6 bgp neighbors	354
show ipv6 bgp prefix-list <list-name>	187
show ipv6 bgp regexp <regexp>	188
show ipv6 bgp	183
show monitoring protocols bgp	189

List of Examples

Use this list to help you locate examples you'd like to look at or try.

Example 1-4 Verifying iBGP on R1: “show ip bgp”	25
Example 1-7 Verifying eBGP on R1: “show ip bgp”	28

Preface

This document describes the various deployment, installation, and upgrade options for Vyatta software.

This preface provides information about using this guide. The following topics are presented:

- [Intended Audience](#)
- [Organization of This Guide](#)
- [Document Conventions](#)
- [Vyatta Publications](#)

Intended Audience

This guide is intended for experienced system and network administrators. Depending on the functionality to be used, readers should have specific knowledge in the following areas:

- Networking and data communications
- TCP/IP protocols
- General router configuration
- Routing protocols
- Network administration
- Network security
- IP services

Organization of This Guide

This guide has the following aid to help you find the information you are looking for:

- [Quick List of Commands](#)
Use this list to help you quickly locate commands.
- [List of Examples](#)
Use this list to help you locate examples you'd like to try or look at.

This guide has the following chapters:

Chapter	Description	Page
Chapter 1: BGP Configuration	This chapter describes how to configure the Border Gateway Protocol on the Vyatta System.	1
Chapter 2: Global and Router-Specific Configuration	This chapter describes commands for global and router-specific configuration for BGP.	75
Chapter 3: Route Reflection	This chapter describes commands for for BGP route reflection.	190
Chapter 4: Confederations	This chapter describes commands for BGP confederations.	205
Chapter 5: Neighbors	This chapter describes commands for BGP neighbors .	211

Chapter 6: Peer Groups	This chapter describes commands for BGP peer groups.	358
Chapter 7: Route Redistribution	This chapter describes commands for BGP route redistribution.	474
Chapter 8: Route Server	This chapter describes commands for BGP route server.	496
Glossary of Acronyms		508

Document Conventions

This guide uses the following advisory paragraphs, as follows.



WARNING Warnings alert you to situations that may pose a threat to personal safety.



CAUTION Cautions alert you to situations that might cause harm to your system or damage to equipment, or that may affect service.

NOTE Notes provide information you might need to avoid problems or configuration errors.

This document uses the following typographic conventions.

Monospace	Examples, command-line output, and representations of configuration nodes.
bold Monospace	Your input: something you type at a command line.
bold	Commands, keywords, and file names, when mentioned inline. Objects in the user interface, such as tabs, buttons, screens, and panes.
<i>italics</i>	An argument or variable where you supply a value.
<key>	A key on your keyboard, such as <Enter>. Combinations of keys are joined by plus signs (“+”), as in <Ctrl>+c.
[key1 key2]	Enumerated options for completing a syntax. An example is [enable disable].
<i>num1–numN</i>	A inclusive range of numbers. An example is 1–65535, which means 1 through 65535, inclusive.
<i>arg1..argN</i>	A range of enumerated values. An example is eth0..eth3, which means eth0, eth1, eth2, or eth3.
<i>arg[arg...]</i> <i>arg[,arg...]</i>	A value that can optionally represent a list of elements (a space-separated list and a comma-separated list, respectively).

Vyatta Publications

Full product documentation is provided in the Vyatta technical library. To see what documentation is available for your release, see the *Guide to Vyatta Documentation*. This guide is posted with every release of Vyatta software and provides a great starting point for finding the information you need.

Additional information is available on www.vyatta.com and www.vyatta.org.

Chapter 1: BGP Configuration

This chapter describes how to configure the Border Gateway Protocol on the Vyatta System.

The following topics are covered:

- [BGP Overview](#)
- [Supported Standards](#)
- [Configuring BGP](#)
- [Configuring IPv6 Routing with BGP](#)

BGP Overview

This section presents the following topics:

- [iBGP and eBGP](#)
- [BGP ID Selection Process](#)
- [BGP Path Selection Process](#)
- [Scalability of BGP](#)
- [Route Flapping and Flap Damping](#)
- [AS Paths](#)
- [BGP Communities](#)
- [Peer Groups](#)
- [BGP Multipath Support](#)
- [IPv4 and IPv6 Support](#)

Border Gateway Protocol (BGP) is the principal inter-domain routing protocol used on the Internet. BGP version 4 is specified in RFC 4271, which obsoletes the original BGPv4 specification defined in RFC 1771.

The principal concept of BGP is that of the Autonomous System (AS). An AS is a routing domain that is under one administrative authority, and which implements its own routing policies. For example, one Internet Service Provider (ISP) would have its own AS, while another would have its own, different, AS. Many large enterprises also have their own AS, particularly if they are multi-homed (that is, connected to multiple ISPs). The BGP routing protocol is used to convey network reachability information between ASs.

Routers that are configured to run BGP between one another are known as BGP peers or BGP neighbors. BGP uses a TCP connection on the well-known port 179 to exchange routing information between peers. BGP peers that are configured within the same AS are referred to as internal BGP (iBGP) peers. BGP peers that are configured in different ASs are referred to as external BGP (eBGP) peers.

There are two basic types of BGP route exchanges that occur between peers: route announcements and route withdrawals.

- A route announcement tells a peer that it can reach a particular network via the announcing router, and includes attributes associated with that path.
- A route withdrawal tells a peer that a previously announced route is no longer reachable via this peer.

All valid route announcements that are received on a BGP router are placed into the router's BGP table. (These routes are typically referred to as BGP paths.) This means that, for a particular network prefix—for example, 10.0.0.0/8—the local BGP router might have recorded multiple available paths: one through any of its BGP peers. For

each prefix, the BGP process uses a path selection algorithm to select the best available path from all those learned from its peers. Once the best path has been selected, that path becomes the candidate route from the BGP protocol for inserting into the active routing table.

Each BGP path includes several attributes that are used by the BGP path selection process to determine which path is the best available path. These attributes can also be used in user-defined routing policies applied to BGP; these can allow the router to perform additional actions on a matching path, such as determining whether to accept or reject a route announcement.

One of the most commonly used BGP path attributes is the AS path. The AS path lists each of the ASs by which the prefix has been announced, and is used to prevent routing loops. The AS path is read from right to left, where the right-most AS is the AS that originated the network prefix (that is, it was the first AS to announce reachability for this prefix). This AS is known as the origin AS.

As a network prefix is advertised between ASs, each AS prepends its own AS number to the AS path. For example, the AS path “4 3 2 1” indicates that AS 1 originated the network prefix. The prefix was advertised from AS 1 to AS 2, then from AS 2 to AS 3, and finally from AS 3 to AS 4.

Other BGP path attributes include origin, next hop, multi-exit discriminator (“med”), local preference (“local pref”), atomic aggregate, and aggregator. These attributes are described in more detail in another section of this document.

iBGP and eBGP

A BGP peer can be one of two types:

- Internal BGP (iBGP) peers are peers that are configured with the same AS number.
- External BGP (eBGP) peers are peers that are configured with different AS numbers.

iBGP

The BGP protocol requires that all iBGP peers within an AS have a connection to one another, creating a full-mesh of iBGP peering connections. (The exception to this is route reflection.) When a prefix is announced from one iBGP peer to another, the AS path is not changed. Due to the full-mesh requirement, all iBGP peers should have the same view of the BGP table, unless different routing policies have been applied to some of the peers.

When a router receives an iBGP announcement, the BGP process uses the BGP best path selection algorithm to determine whether the received announcement is the best available path for that prefix. If it is the best available path, then the BGP process uses this route as the BGP candidate route for insertion into the routing table, and

the BGP process announces this path to all its peers, both iBGP and eBGP peers. If it is not the best available path, then the BGP process keeps a copy of this path in its BGP table, so that it can be used to calculate the best available path when path information for that prefix changes (for example, if the current best available path is withdrawn).

The BGP ID is a unique identifier in the format of an IP address used to identify a peer. The peering IP address is the actual IP address used for the BGP connection.

For iBGP peerings, the BGP ID and peering IP is frequently the IP address bound to that router's loopback interface. An iBGP session is usually contained within a local LAN, with multiple redundant physical links between the iBGP devices. For iBGP routes, reachability is all that is necessary, and the loopback interface is reachable so long as at least one physical interface is operational. Because of the physical and/or logical redundancy that exists between iBGP peers, iBGP peering on the loopback interface works well.

Since BGP does not provide reachability information, you must make sure that each iBGP peer knows how to reach other peers. To be able to reach one another, each peer must have some sort of Interior Gateway Protocol (IGP) route, such as a connected route, a static route, or a route through a dynamic routing protocol such as RIP or OSPF, which tells them how to reach the opposite router.

eBGP

External BGP is the method that different Autonomous Systems (ASs) use to interconnect with one another. eBGP usually takes place over WAN links, where there may be a single physical path between eBGP peers. Alternatively, they may have multiple eBGP peer connections to provide redundancy and/or traffic load balancing. Redundant peers use distinct BGP sessions so that, if one session fails, another can take over.

BGP uses an AS path to track the path of a prefix through the various ASs that send or receive the prefix announcement. When a prefix is announced to an eBGP peer, the local AS number is prepended to the AS path. This helps to prevent routing loops by rejecting any prefix announcements that include the local AS number in the AS path. Prefix announcements learned via eBGP are also analyzed using the BGP best path selection process.

For eBGP peerings, the BGP ID and peering IP address is typically the local IP address of the interface that is being used to connect to the eBGP peers. However if more than one physical interface is being used for eBGP peering it is also common to use a loopback IP address as the BGP ID, but still use the physical interface IP address as the peering IP address.

BGP ID Selection Process

It is best practice to configure the BGP ID explicitly using the `protocols bgp <asn> parameters router-id <id> command`. If explicitly assigned, then this is the BGP ID that will be used. If it is not assigned explicitly then the system will choose one automatically. If a loopback address is configured as is not 127.0.0.1, it will be used. If the loopback address is not used then the largest IP address assigned to an interface in the system configuration will be used.

So, best practice suggests that the following configuration commands be run on each BGP router to explicitly set the BGP ID. In this example a loopback address of 10.0.0.11/32 is used.

Example 1-1 Best practice for setting the BGP ID

Router	Step	Command(s)
R1	Set the loopback address and prefix length (note the prefix length of 32). In this case 10.0.0.11/32.	<pre>vyatta@R1# set interfaces loopback lo address 10.0.0.11/32</pre>
R1	Set the router ID to the previously defined loopback address.	<pre>vyatta@R1# set protocols bgp 100 parameters router-id 10.0.0.11</pre>
R1	Commit the configuration.	<pre>vyatta@R1# commit</pre>
R1	Verify the configuration.	<pre>vyatta@R1# show interfaces loopback loopback lo { address 10.0.0.11/32 } vyatta@R1# show protocols bgp bgp 100 { parameters { router-id 10.0.0.11 } }</pre>

BGP Path Selection Process

The BGP process may receive advertisements from multiple peers for the same network prefix. Each of these announcements from a peer for a prefix is called a path. The BGP process selects the “best” path from all available paths and this path becomes the candidate route announced by BGP for inclusion in the Routing Information Base (RIB).

Depending on what other protocols also have candidate routes for this network prefix, the BGP route may or may not be added to the RIB. For instance if the RIB has candidate routes from both BGP and static routing for the same prefix, the static route, not the BGP route, will be included in the RIB. This is because the static route process has a lower administrative cost than the BGP process.

It is important to note that BGP will not consider paths that contain a NEXT_HOP value that is not reachable via an entry in the RIB. For all valid paths, the Vyatta system uses a BGP path selection process based on decision process described in RFC 4271, section 9.1. BGP paths are preferred based on the following:

- **LOCAL PREFERENCE:** Prefer the path with the highest LOCAL_PREF
- **AS PATH LENGTH:** Prefer the path with the shortest AS_PATH.
- **ORIGIN:** Prefer the path with the lowest ORIGIN type.
- **MULTI_EXIT_DISC:** Prefer the path with the lowest MED.
- **PEER TYPE:** Prefer paths learned via eBGP over paths learned via iBGP.
- **IGP METRIC.** Prefer paths with lower IGP metric for the path’s NEXT_HOP address.
- **BGP ID:** Prefer the path with the lowest BGP ID.
- **PEER IP:** Prefer the path with the lowest peer IP address.

The best path selection process is performed as “first match and out.” This means that two paths will be compared until the first difference in preference criteria. For example, two paths for the same network prefix may have the same LOCAL_PREF value, but different AS path lengths. In this case, the path with the shortest AS path would be the “best” path. If the peer IP address is being used to select the best path, this means that all other path criteria were the same for the available paths.

You can use the **show ip bgp** command to see the current best paths in the RIB.

Scalability of BGP

The Border Gateway Protocol 4 specification (RFC 4271) requires that iBGP peers be fully meshed; that is, every iBGP peer must have a connection to every other iBGP peer. A full mesh of iBGP peers does not scale well to large ASs, which can have hundreds of iBGP routers. To overcome scalability issues, two enhancements have been developed for BGP:

- BGP Confederations (RFC 3065)
- Route Reflection (RFC 2796)

Confederations

Confederations enable you to reduce the size and complexity of the iBGP mesh. In a BGP confederation, a single AS is divided into multiple internal sub-ASs to help keep the number of iBGP peer connections manageable. Each sub-AS is assigned its own AS number; this is typically assigned from the private AS number space, which ranges from 64512 to 65534. Within a sub-AS, all the standard iBGP rules, including full-mesh peering, apply. The connections between confederation sub-ASs use eBGP peering. One or more eBGP connections can be made between each sub-AS. The sub-ASs are grouped as a confederation, which advertises as a single AS to external peers.

Figure 1-1 shows the large number of iBGP connections that must be configured in even a moderately sized AS. In this example, 14 routers are participating in iBGP.

Figure 1-1 iBGP full mesh

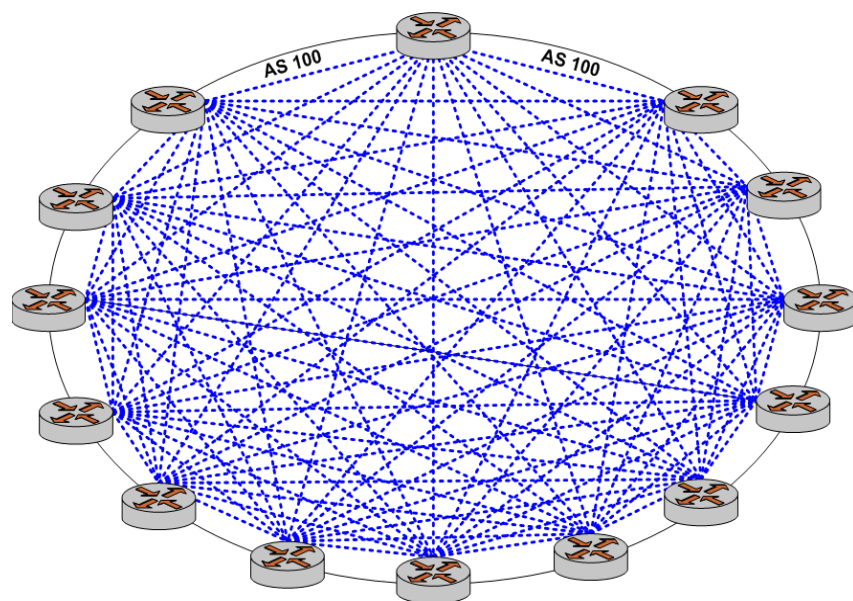
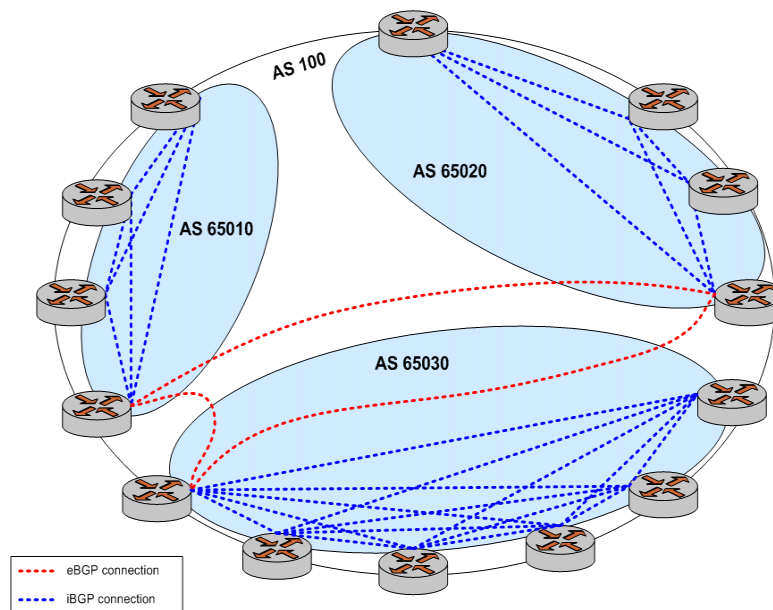


Figure 1-2 shows a BGP confederation that splits the single AS shown in Figure 1-1 into three sub-ASs, which each use private AS numbers. Within each sub-AS, all of the iBGP peers are fully meshed. The sub-ASs are connected to one another using an eBGP connection.

Figure 1-2 BGP confederation



Route Reflection

Another technology designed to help ASs with large numbers of iBGP peers is route reflection. In a standard BGP implementation, all iBGP peers must be fully meshed. Because of this requirement, when an iBGP peer learns a route from another iBGP peer, the receiving router does not forward the route to any of its iBGP peers, since these routers should have learned the route directly from the announcing router.

In a route reflector environment the iBGP peers are no longer fully meshed. Instead, each iBGP peer has an iBGP connection to one or more route reflectors (RRs). Routers configured with a connection to an RR server are referred to as RR clients. Only the RR server is configured to be aware that the RR client is part of an RR configuration; from the RR client's point of view, it is configured normally, and does not have any awareness that it is part of a RR configuration.

In route reflection, internal peers of an RR server are categorized into two types:

- **Client peers.** The RR server and its client peers form a cluster. Within a cluster, client peers need not be fully meshed, but must have an iBGP connection to at least one RR in the cluster.
- **Non-client peers.** Non-client peers, including the RR server, must be fully meshed.

An RR environment is unlike a regular environment, where iBGP peers never forward a route update to other iBGP peers (which is the reason why each iBGP peer must peer with all other peers). When an RR server receives an iBGP update from an RR client, these route updates can also be sent to all other RR clients. When an RR server receives a route update from a peer, it selects the best path based on its path selection rule. After the best path is selected, the RR server chooses its action depending on the type of the peer from which it learned the best path.

- If the route was learned from a client peer, the RR reflects the route to both client and non-client peers. All iBGP updates from client peers are reflected to all other client peers in the cluster. This is done regardless of whether the update was the best path for the RR itself.
- If the route was learned from a non-client iBGP peer, it is reflected out to all RR client peers.
- If the route was learned from an eBGP peer, the route is reflected to all RR clients and all non-clients.

Figure 1-3 shows again the full mesh of iBGP connections in even a moderately sized AS.

Figure 1-3 iBGP full mesh

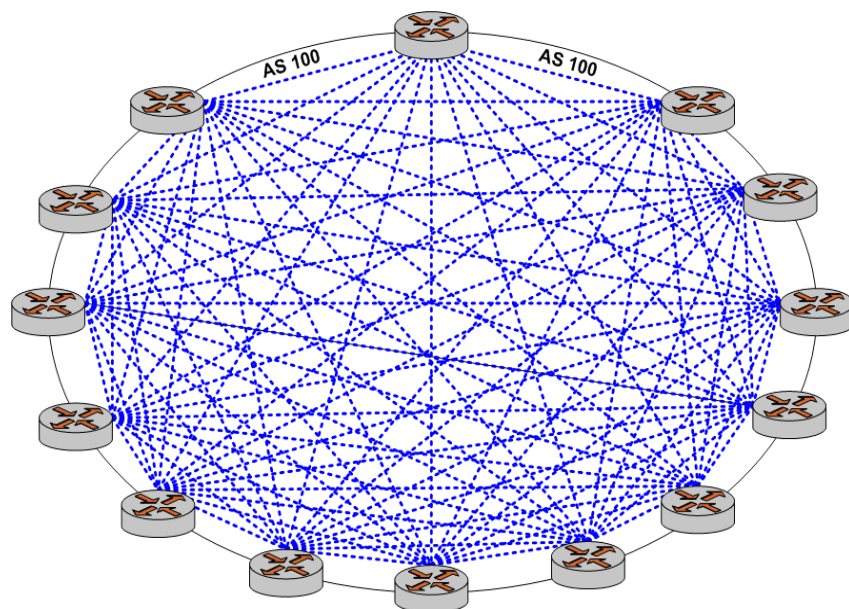
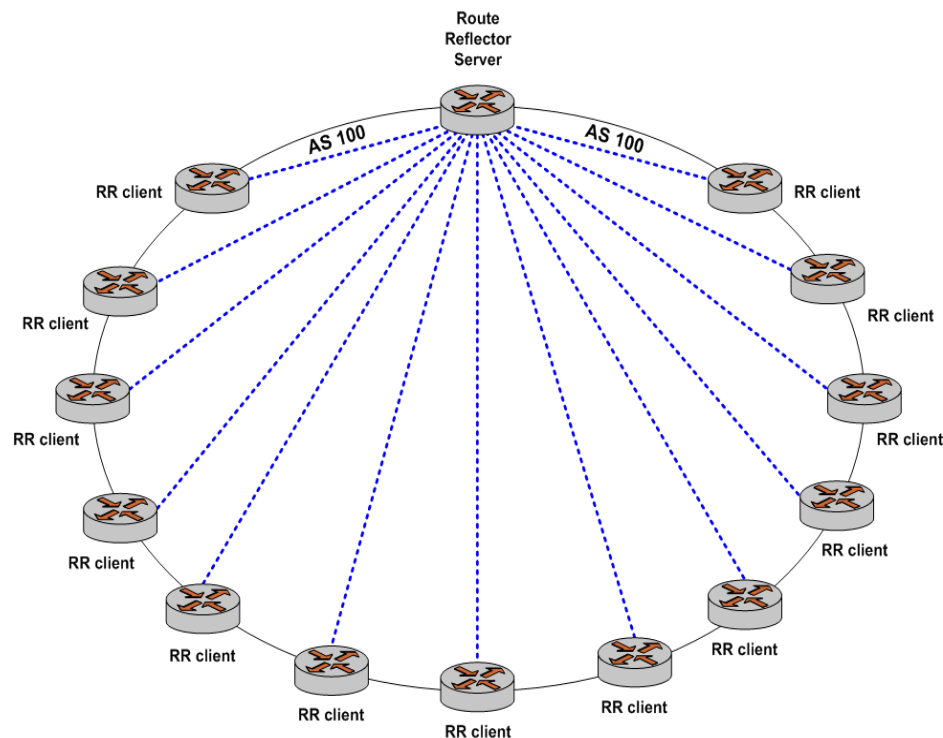


Figure 1-4 shows how introducing route reflection into the AS dramatically reduces the number of iBGP connections required within the AS.

Figure 1-4 iBGP route reflection



Note that to prevent looping, clients must not peer with RRs outside of the cluster.

To achieve redundancy, more than one RR server can be configured within a cluster. Also, to scale to very large networks, a large AS can be configured to have multiple clusters with redundant RR servers, where the RR servers are all configured with a full mesh of iBGP connections between the RR servers.

Route Flapping and Flap Damping

Route flapping is a situation where a route fluctuates repeatedly between being announced, then withdrawn, then announced, then withdrawn, and so on. In this situation, a BGP system will send an excessive number of update messages advertising network reachability information.

Route flapping can cause several different issues. First, each time a new route is learned or withdrawn the BGP best path selection process for that prefix must be executed, which can result in high CPU utilization. If a large enough number of routes are flapping, the BGP process may not be able to converge sufficiently quickly. Second, the route flapping issue can become amplified as it passes from peer to peer. For example, if a router with two peers flaps a route, and those two peers each have 10 peers, the flapping route affects 20 BGP routers.

Route damping is intended to minimize the propagation of update messages between BGP peers for flapping routes. This reduces the load on these devices without unduly impacting the route convergence time for stable routes.

When route damping is enabled, a route is assigned a penalty each time it “flaps” (that is, each time it is announced and then withdrawn within a short interval). If the penalty exceeds 1000 (its *suppress* value) the route is suppressed.

After the route has been stable for a configured interval (its *half-life*) the penalty is reduced by half. Subsequently, the penalty is reduced every five seconds. When the penalty falls below a configured value (its *reuse* value), the route is unsuppressed.

The penalty applied to a route will never exceed the *maximum penalty*, which is computed from configured attributes as follows:

$$\text{Maximum penalty} = \text{reuse} * 2^{(\text{max-suppress-time}/\text{half-life})}$$

While the route is being “damped,” updates and withdrawals for this route from a peer are ignored. This helps to localize the route flapping to a particular peering connection.

AS Paths

An AS path is a path to a destination in the Border Gateway Protocol (BGP). The path is represented as a sequence of AS numbers, which are the numbers uniquely identifying BGP autonomous systems. Each AS number represents an autonomous system (which may be comprised of multiple networks) that a packet traverses if it takes the associated route to the destination.

For a packet to reach a destination using this route, it traverses the listed ASs from the leftmost AS number to the rightmost, where the rightmost is the AS immediately preceding its destination.

Using policies, match conditions can be defined based on all or portions of the AS path. To do this, you can either specify the AS path directly in a policy command using a regular expression in the **as-path** attribute, or create a named AS path regular expression using the **as-path-list** attribute and including the name in a policy command.

BGP Communities

All BGP updates include a BGP attribute called the communities path attribute. The communities path attribute allows ASs to “tag” prefix announcements. This tag can then be used by routing policies to modify the normal behavior for that prefix announcement. For example, an AS could choose to filter out all prefix announcements containing a community value that identifies the prefix as a customer-generated prefix, instead only announcing the summary prefix for all customer prefixes. It is important to note that the community path attribute is carried in BGP update messages, which allows ASs not directly connected to each other to share information about a prefix.

The format for community identifiers is defined in RFC 1997: “BGP Communities Attribute.” The community identifier is a 32-bit value, where the first two bytes of the value are the AS number and the second two bytes are an arbitrary value defined by the AS. This format can be represented as *AA:NN*, where *AA* is the AS number of the AS adding the community identifier to the community path attribute, and *NN* represents a user-defined policy value.

There are two types of BGP communities: “well-known” communities and user-defined or private communities.

The Vyatta system recognizes the following BGP well-known communities as per RFC 1997:

NO_EXPORT: Allows a BGP speaker to specify that redistribution should extend only to the neighbouring AS. All routes received carrying a communities attribute containing this value are not advertised outside a BGP confederation boundary (a stand-alone autonomous system that is not part of a confederation should be considered a confederation itself).

NO_ADVERTISE: All routes received carrying a communities attribute containing this value are not advertised to other BGP peers.

LOCAL_AS: All routes received carrying a communities attribute containing this value are not advertised to other BGP peers.

In addition, the Vyatta system recognizes the **INTERNET** community. The **INTERNET** community is not, strictly speaking, a well-known community. (It is not, for example, mentioned in RFC 1997.) You can use the **INTERNET** community when you would use a “permit all” statement; for example, at the end of a community list. (Like access lists, community lists have an implicit “deny all” statement at their end.) The “permit **INTERNET**” statement is equivalent to the “permit any” statement, matching any community in the IP community list.

Peer Groups

In order to simplify the configuration of multiple similar peer sessions, “peer groups” are available. Configuring a peer group is the same as configuring an individual peer. The peer group configuration applies to all members of the peer group. Peer groups are created and configured using “[protocols bgp <asn> peer-group <group-name>](#)” on page 363. Peers are added to a peer group using “[protocols bgp <asn> neighbor <id> peer-group <group-name>](#)” on page 304.

BGP Multipath Support

In general, BGP selects a single best path to a given destination. BGP multipath support allows for multiple paths to a given destination. Having multiple paths to the same destination enables the traffic load to be shared across the multiple paths, and makes use of network resources that might normally be used only as backup.

The BGP multipath feature does not affect the BGP path selection process. Instead, the best path as well as the additional multipaths are all added to the routing table.

For more information on configuring BGP multipath support, see “[protocols bgp <asn> maximum-paths ebgp <max-paths>](#)” on page 105 and “[protocols bgp <asn> maximum-paths ibgp <max-paths>](#)” on page 107.

IPv4 and IPv6 Support

The Vyatta BGP implementation supports the following:

- Peering sessions between IPv4 peers
- Peering sessions between IPv6 peers
- IPv4 routing information can be carried over either IPv4 or IPv6 peering sessions
- IPv6 routing information can be carried over either IPv4 or IPv6 peering sessions
- Both IPv4 and IPv6 routing information can be carried over a single IPv4 or IPv6 peering session

NOTE *IPv4 routes over IPv6 sessions and IPv6 routes over IPv4 sessions cannot currently be seen via `show` commands.*

IPv4 routes can be exchanged after BGP has been enabled on the system (using the [show monitoring protocols bgp](#) command).

IPv6 routes can be exchanged once either you have enabled any of the IPv6 unicast address family (using the [protocols bgp <asn> address-family ipv6-unicast](#) command), the IPv6 unicast neighbor address family (using the [protocols bgp <asn> neighbor <id> address-family ipv6-unicast](#) command), or the IPv6 unicast peer group address family (using the [protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast](#) command).

Supported Standards

The Vyatta implementation of BGP complies with the following standards:

- RFC 4271: A Border Gateway Protocol (BGP-4)
- RFC 4273: Definitions of Managed Objects for BGP-4
- RFC 2842: Capabilities Advertisement with BGP-4
- RFC 4760: Multiprotocol Extensions for BGP-4
- RFC 1997: BGP Communities Attribute
- RFC 3065: Autonomous System Confederations for BGP
- RFC 2796: BGP Route Reflection - An Alternative to Full Mesh IBGP

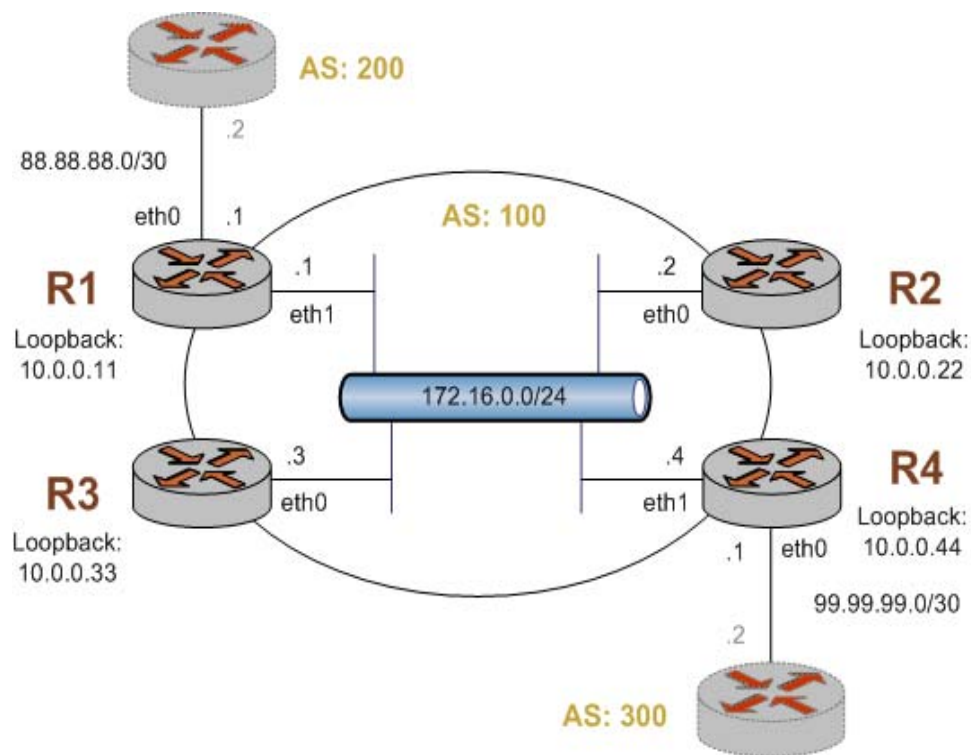
Configuring BGP

This section presents the following topics:

- [Basic iBGP Configuration](#)
- [Verifying the iBGP Configuration](#)
- [Basic eBGP Configuration](#)
- [Verifying the eBGP Configuration](#)
- [Originating a Route to eBGP Neighbors](#)
- [Verifying the Route Origination](#)
- [Inbound Route Filtering](#)
- [Verifying the Inbound Filter](#)
- [Outbound Route Filtering](#)
- [Verifying the Outbound Filter](#)
- [Confederations](#)
- [Verifying the Confederation](#)
- [Route Reflectors](#)
- [Verifying the Route Reflector](#)
- [Route Redirection](#)

This section presents sample configurations for BGP. The configuration examples are based on the reference diagram in [Figure 1-5](#).

Figure 1-5 BGP configuration reference diagram



Basic iBGP Configuration

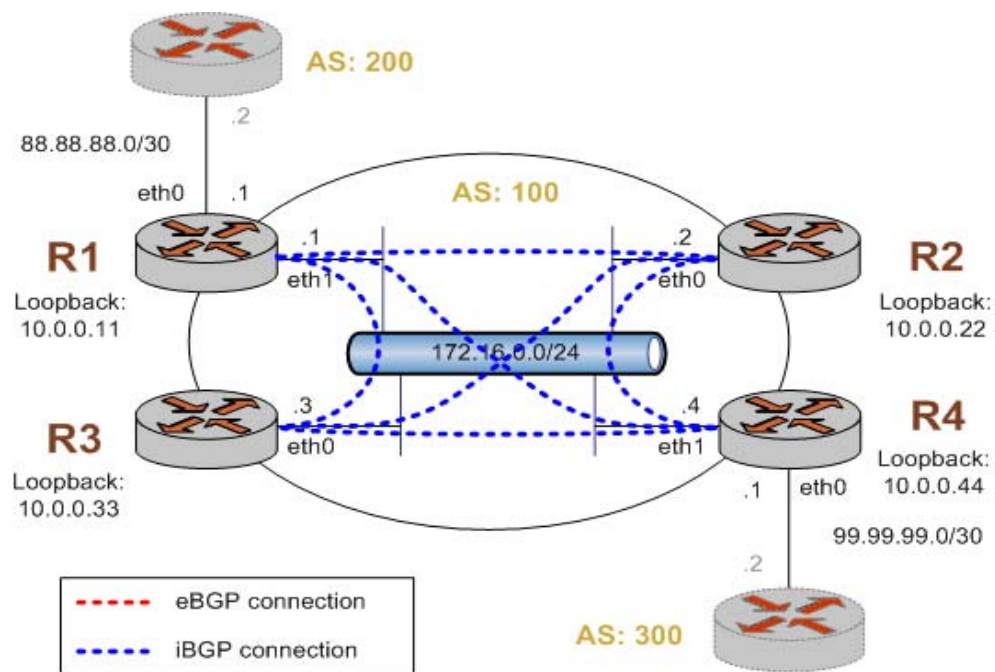
In this section, you configure iBGP on the routers labeled R1, R2, R3, and R4 in the reference network diagram. Each router has an iBGP peering connection to each of the other iBGP routers in the network, satisfying the full mesh iBGP peering requirement.

In the example the iBGP peering connections are established between iBGP neighbors using the loopback interface IP addresses. This is typical practice, particularly when there are redundant connections between the iBGP routers.

In order for the routers to be able to reach each other using the loopback IP address, the loopback IP addresses must be reachable via an entry in the router's routing table. This requires some form of Internal Gateway Protocol (IGP). In the example here, we will use a basic Open Shortest Path First (OSPF) configuration to announce the loopback addresses between neighbors.

Figure 1-6 shows the BGP connections after you have completed the iBGP configuration.

Figure 1-6 Basic iBGP configuration



This example assumes that you have already configured the router interfaces; only the steps required to implement BGP are shown.

To create a basic iBGP configuration, perform the following steps in configuration mode.

Example 1-2 Basic iBGP configuration

Router	Step	Command(s)
R1	Advertise the loopback address within the OSPF area. This is needed for iBGP.	<code>vyatta@R1# set protocols ospf area 0.0.0.0 network 10.0.0.11/32</code>
R1	Advertise the local network within the OSPF area.	<code>vyatta@R1# set protocols ospf area 0.0.0.0 network 172.16.0.0/24</code>
R1	Advertise the external network within the OSPF area.	<code>vyatta@R1# set protocols ospf area 0.0.0.0 network 88.88.88.0/30</code>
R1	Set the router ID to be the loopback IP address.	<code>vyatta@R1# set protocols ospf parameters router-id 10.0.0.11</code>

Example 1-2 Basic iBGP configuration

R1	Configure the eth0 interface to be a passive interface (i.e. don't advertize our internal network routes over the external network).	<pre>vyatta@R1# set protocols ospf passive-interface eth0</pre>
R1	Create an iBGP peer for R2. The peer is an iBGP peer because it resides within the same AS as this router.	<pre>vyatta@R1# set protocols bgp 100 neighbor 10.0.0.22 remote-as 100</pre>
R1	Define the IP address on the local R1 router that is used to peer with the R2 router.	<pre>vyatta@R1# set protocols bgp 100 neighbor 10.0.0.22 update-source 10.0.0.11</pre>
R1	Create an iBGP peer for R3. The peer is an iBGP peer because it resides within the same AS as this router.	<pre>vyatta@R1# set protocols bgp 100 neighbor 10.0.0.33 remote-as 100</pre>
R1	Define the IP address on the local R1 router that is used to peer with the R3 router.	<pre>vyatta@R1# set protocols bgp 100 neighbor 10.0.0.33 update-source 10.0.0.11</pre>
R1	Create an iBGP peer for R4. The peer is an iBGP peer because it resides within the same AS as this router.	<pre>vyatta@R1# set protocols bgp 100 neighbor 10.0.0.44 remote-as 100</pre>
R1	Define the IP address on the local R1 router that is used to peer with the R4 router.	<pre>vyatta@R1# set protocols bgp 100 neighbor 10.0.0.44 update-source 10.0.0.11</pre>
R1	Set the router ID to the loopback address, which on R1 is 10.0.0.11.	<pre>vyatta@R1# set protocols bgp 100 parameters router-id 10.0.0.11</pre>
R1	Commit the configuration.	<pre>vyatta@R1# commit</pre>

Example 1-2 Basic iBGP configuration

R1	Display the configuration.	<pre> vyatta@R1# show protocols bgp 100 { neighbor 10.0.0.22 { remote-as 100 update-source 10.0.0.11 } neighbor 10.0.0.33 { remote-as 100 update-source 10.0.0.11 } neighbor 10.0.0.44 { remote-as 100 update-source 10.0.0.11 } parameters { router-id 10.0.0.11 } } ospf { area 0.0.0.0 { network 172.16.0.0/24 network 88.88.88.0/30 network 10.0.0.11/32 } parameters { router-id 10.0.0.11 } passive-interface eth0 } </pre>
R2	Advertize the loopback address within the OSPF area. This is needed for iBGP.	<pre> vyatta@R2# set protocols ospf area 0.0.0.0 network 10.0.0.22/32 </pre>
R2	Advertize the local network within the OSPF area.	<pre> vyatta@R2# set protocols ospf area 0.0.0.0 network 172.16.0.0/24 </pre>
R2	Set the router ID to be the loopback IP address.	<pre> vyatta@R2# set protocols ospf parameters router-id 10.0.0.22 </pre>
R2	Create an iBGP peer for R1. The peer is an iBGP peer because it resides within the same AS as this router.	<pre> vyatta@R2# set protocols bgp 100 neighbor 10.0.0.11 remote-as 100 </pre>

Example 1-2 Basic iBGP configuration

R2	Define the IP address on the local R2 router that is used to peer with the R1 router.	<pre>vyatta@R2# set protocols bgp 100 neighbor 10.0.0.11 update-source 10.0.0.22</pre>
R2	Create an iBGP peer for R3. The peer is an iBGP peer because it resides within the same AS as this router.	<pre>vyatta@R2# set protocols bgp 100 neighbor 10.0.0.33 remote-as 100</pre>
R2	Define the IP address on the local R2 router that is used to peer with the R3 router.	<pre>vyatta@R2# set protocols bgp 100 neighbor 10.0.0.33 update-source 10.0.0.22</pre>
R2	Create an iBGP peer for R4. The peer is an iBGP peer because it resides within the same AS as this router.	<pre>vyatta@R2# set protocols bgp 100 neighbor 10.0.0.44 remote-as 100</pre>
R2	Define the IP address on the local R2 router that is used to peer with the R4 router.	<pre>vyatta@R2# set protocols bgp 100 neighbor 10.0.0.44 update-source 10.0.0.22</pre>
R2	Set the router ID to the loopback address, which on R2 is 10.0.0.22.	<pre>vyatta@R2# set protocols bgp 100 parameters router-id 10.0.0.22</pre>
R2	Commit the configuration.	<pre>vyatta@R2# commit</pre>

Example 1-2 Basic iBGP configuration

R2	Display the configuration.	<pre> vyatta@R2# show protocols bgp 100 { neighbor 10.0.0.11 { remote-as 100 update-source 10.0.0.22 } neighbor 10.0.0.33 { remote-as 100 update-source 10.0.0.22 } neighbor 10.0.0.44 { remote-as 100 update-source 10.0.0.22 } parameters { router-id 10.0.0.22 } } ospf { area 0.0.0.0 { network 172.16.0.0/24 network 10.0.0.22/32 } parameters { router-id 10.0.0.22 } } </pre>
R3	Advertize the loopback address within the OSPF area. This is needed for iBGP.	<pre> vyatta@R3# set protocols ospf area 0.0.0.0 network 10.0.0.33/32 </pre>
R3	Advertize the local network within the OSPF area.	<pre> vyatta@R3# set protocols ospf area 0.0.0.0 network 172.16.0.0/24 </pre>
R3	Set the router ID to be the loopback IP address.	<pre> vyatta@R3# set protocols ospf parameters router-id 10.0.0.33 </pre>
R3	Create an iBGP peer for R1. The peer is an iBGP peer because it resides within the same AS as this router.	<pre> vyatta@R3# set protocols bgp 100 neighbor 10.0.0.11 remote-as 100 </pre>

Example 1-2 Basic iBGP configuration

R3	Define the IP address on the local R3 router that is used to peer with the R1 router.	<pre>vyatta@R3# set protocols bgp 100 neighbor 10.0.0.11 update-source 10.0.0.33</pre>
R3	Create an iBGP peer for R2. The peer is an iBGP peer because it resides within the same AS as this router.	<pre>vyatta@R3# set protocols bgp 100 neighbor 10.0.0.22 remote-as 100</pre>
R3	Define the IP address on the local R3 router that is used to peer with the R2 router.	<pre>vyatta@R3# set protocols bgp 100 neighbor 10.0.0.22 update-source 10.0.0.33</pre>
R3	Create an iBGP peer for R4. The peer is an iBGP peer because it resides within the same AS as this router.	<pre>vyatta@R3# set protocols bgp 100 neighbor 10.0.0.44 remote-as 100</pre>
R3	Define the IP address on the local R3 router that is used to peer with the R4 router.	<pre>vyatta@R3# set protocols bgp 100 neighbor 10.0.0.44 update-source 10.0.0.33</pre>
R3	Set the router ID to the loopback address, which on R3 is 10.0.0.33.	<pre>vyatta@R3# set protocols bgp 100 parameters router-id 10.0.0.33</pre>
R3	Commit the configuration.	<pre>vyatta@R3# commit</pre>

Example 1-2 Basic iBGP configuration

R3	Display the configuration.	<pre> vyatta@R3# show protocols bgp 100 { neighbor 10.0.0.11 { remote-as 100 update-source 10.0.0.33 } neighbor 10.0.0.22 { remote-as 100 update-source 10.0.0.33 } neighbor 10.0.0.44 { remote-as 100 update-source 10.0.0.33 } parameters { router-id 10.0.0.33 } } ospf { area 0.0.0.0 { network 172.16.0.0/24 network 10.0.0.33/32 } parameters { router-id 10.0.0.33 } } </pre>
R4	Advertize the loopback address within the OSPF area. This is needed for iBGP.	<pre> vyatta@R4# set protocols ospf area 0.0.0.0 network 10.0.0.44/32 </pre>
R4	Advertize the local network within the OSPF area.	<pre> vyatta@R4# set protocols ospf area 0.0.0.0 network 172.16.0.0/24 </pre>
R4	Advertize the external network within the OSPF area.	<pre> vyatta@R4# set protocols ospf area 0.0.0.0 network 99.99.99.0/30 </pre>
R4	Set the router ID to be the loopback IP address.	<pre> vyatta@R4# set protocols ospf parameters router-id 10.0.0.44 </pre>

Example 1-2 Basic iBGP configuration

R4	Configure the eth0 interface to be a passive interface (i.e. don't advertize our internal network routes over the external network).	<pre>vyatta@R4# set protocols ospf passive-interface eth0</pre>
R4	Create an iBGP peer for R1. The peer is an iBGP peer because it resides within the same AS as this router.	<pre>vyatta@R4# set protocols bgp 100 neighbor 10.0.0.11 remote-as 100</pre>
R4	Define the IP address on the local R4 router that is used to peer with the R1 router.	<pre>vyatta@R4# set protocols bgp 100 neighbor 10.0.0.11 update-source 10.0.0.44</pre>
R4	Create an iBGP peer for R2. The peer is an iBGP peer because it resides within the same AS as this router.	<pre>vyatta@R4# set protocols bgp 100 neighbor 10.0.0.22 remote-as 100</pre>
R4	Define the IP address on the local R4 router that is used to peer with the R2 router.	<pre>vyatta@R4# set protocols bgp 100 neighbor 10.0.0.22 update-source 10.0.0.44</pre>
R4	Create an iBGP peer for R3. The peer is an iBGP peer because it resides within the same AS as this router.	<pre>vyatta@R4# set protocols bgp 100 neighbor 10.0.0.33 remote-as 100</pre>
R4	Define the IP address on the local R4 router that is used to peer with the R3 router.	<pre>vyatta@R4# set protocols bgp 100 neighbor 10.0.0.33 update-source 10.0.0.44</pre>
R4	Set the router ID to the loopback address, which on R4 is 10.0.0.44.	<pre>vyatta@R4# set protocols bgp 100 parameters router-id 10.0.0.44</pre>
R4	Commit the configuration.	<pre>vyatta@R4# commit</pre>

Example 1-2 Basic iBGP configuration

```
R4      Display the configuration.  vyatta@R4# show protocols
                                          bgp 100 {
                                          neighbor 10.0.0.11 {
                                          remote-as 100
                                          update-source 10.0.0.44
                                          }
                                          neighbor 10.0.0.22 {
                                          remote-as 100
                                          update-source 10.0.0.44
                                          }
                                          neighbor 10.0.0.33 {
                                          remote-as 100
                                          update-source 10.0.0.44
                                          }
                                          parameters {
                                          router-id 10.0.0.44
                                          }
                                          }
                                          ospf {
                                          area 0.0.0.0 {
                                          network 172.16.0.0/24
                                          network 99.99.99.0/30
                                          network 10.0.0.44/32
                                          }
                                          parameters {
                                          router-id 10.0.0.44
                                          }
                                          }
                                          passive-interface eth0
                                          }
```

Verifying the iBGP Configuration

The following commands can be used to verify the iBGP configuration.

R1: show ip bgp summary

[Example 1-3](#) shows the output of the `show ip bgp summary` command for router R1 at this stage of the configuration.

Example 1-3 Verifying iBGP on R1: "show ip bgp summary"

```
vyatta@R1:~$ show ip bgp summary
BGP router identifier 10.0.0.11, local AS number 100
RIB entries 1, using 64 bytes of memory
Peers 3, using 7560 bytes of memory
```



```

Neighbor      V    AS  MsgRcvd  MsgSent    TblVer  InQ  OutQ  Up/Down  State/PfxRcd
10.0.0.22     4   100      6       10         0    0    0 00:04:18      0
10.0.0.33     4   100      6       10         0    0    0 00:04:14      0
10.0.0.44     4   100      5        6         0    0    0 00:02:55      0

```

```

Total number of neighbors 3
vyatta@R1:~$

```

The most important fields in the output for **show ip bgp summary** are the **Up/Down** and **State** fields. All the iBGP peers for R1 show times in the **Up/Down** field, which means the connection has been “established” for that period of time. The established state indicates that the peers have successfully created a BGP connection between one another, and are now able to send and receive BGP update messages.

If a peer shows in either **Active** or **Idle** in the **State** field, it means there is some issue that is keeping the BGP peers from forming an adjacency.

- The **Active** state identifies that the local router is actively trying to establish a TCP connection to the remote peer. You may see this if the local peer has been configured, but the remote peer is unreachable or has not been configured.
- The **Idle** state indicates that the local router has not allocated any resources for that peer connection, so any incoming connection requests will be refused.

R1: show ip bgp

Because we have not configured any routing announcements yet, the BGP table is currently empty. This can be seen by the output of **show ip bgp** for R1, which is shown in [Example 1-4](#).

Example 1-4 Verifying iBGP on R1: “show ip bgp”

```

vyatta@R1:~$ show ip bgp
No BGP network exists
vyatta@R1:~$

```

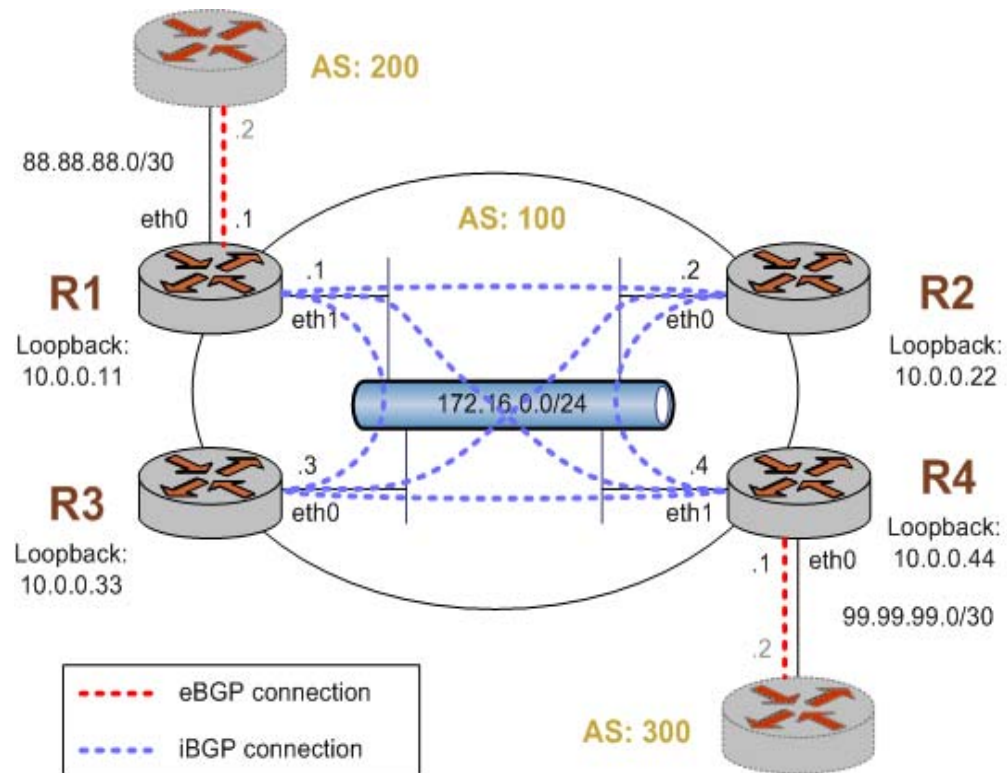
Basic eBGP Configuration

In this section, you configure eBGP on the routers labeled R1 and R4 in the reference network diagram. Router R1 is peering with an eBGP neighbor that is configured to be in AS 200 and router R4 is peering with an eBGP neighbor in AS 300.

In this example, the eBGP peering connections are established between eBGP neighbors using the physical interface IP addresses. This is a common configuration for eBGP peers. If the link between the peers goes down, the peering relationship should also go down since there is no redundancy.

After the basic eBGP configuration has been completed, the network should look as shown in [Figure 1-7](#).

Figure 1-7 Basic eBGP configuration



This example assumes the following:

- The configuration in [Example 1-6](#) has already been performed.
- The eBGP peers connecting to R1 and R4 have been properly configured for BGP.

To create a basic eBGP configuration, perform the following steps in configuration mode.

Example 1-5 Basic eBGP configuration

Router	Step	Command(s)
R1	Create an eBGP peer for R1. The peer is an eBGP peer because it resides in a different AS than this router.	vyatta@R1# set protocols bgp neighbor 88.88.88.2 remote-as 200
R1	Commit the configuration.	vyatta@R1# commit
R4	Create an eBGP peer for R4. The peer is an eBGP peer because it resides in a different AS than this router.	vyatta@R4# set protocols bgp neighbor 99.99.99.2 remote-as 300
R4	Commit the configuration.	vyatta@R4# commit

Verifying the eBGP Configuration

The following commands can be used to verify the eBGP configuration. Note that the output shown for these commands would be obtained *after* the configuration for both router R1 and router R4 has been completed.

R1: show ip bgp summary

[Example 1-6](#) shows the output of the `show ip bgp summary` command for router R1 at this stage of the configuration.

Example 1-6 Verifying eBGP on R1: “show ip bgp summary”

```
vyatta@R1:~$ show ip bgp summary
BGP router identifier 10.0.0.11, local AS number 100
RIB entries 23, using 1472 bytes of memory
Peers 4, using 10080 bytes of memory

Neighbor      V    AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.0.22     4   100     40     44      0    0    0 00:38:23    0
10.0.0.33     4   100     40     44      0    0    0 00:38:22    0
10.0.0.44     4   100     43     47      0    0    0 00:38:22    0
88.88.88.2    4   200      4      5      0    0    0 00:01:22    0
```

```
Total number of neighbors 4
vyatta@R1:~$
```

After adding the eBGP peer 88.88.88.2 (the BGP ID configured for the router connected to AS 200) to R1 we can see that the connection to the new peer is established (in the **Up/Down** field). This indicates that the peer was properly preconfigured for this connection.

Additionally you may notice that the **MsgRcvd** and **MsgSent** fields for peer 88.88.88.2 shows “4” and “5” respectively. This shows that R1 has received four BGP messages from 88.88.88.2 and sent it five, which are associated with the connection to the peer (the **show ip bgp neighbors** command can be used to provide additional detail). Also, you can see that router R1 has received no route prefixes (**PfxRcd** column) from either 10.0.0.44 or from 88.88.88.2. This indicates that, though the BGP neighbors are communicating, they are not sending any prefixes. This is the way it should be at this point in the example as we have not yet configured routing announcements on any of the BGP peers.

R1: show ip bgp

Again, because we have not configured any routing announcements, the BGP table is currently empty. This can be seen by the output of **show ip bgp** for R1, which is shown in [Example 1-7](#).

Example 1-7 Verifying eBGP on R1: “show ip bgp”

```
vyatta@R1:~$ show ip bgp
No BGP network exists
vyatta@R1:~$
```

Originating a Route to eBGP Neighbors

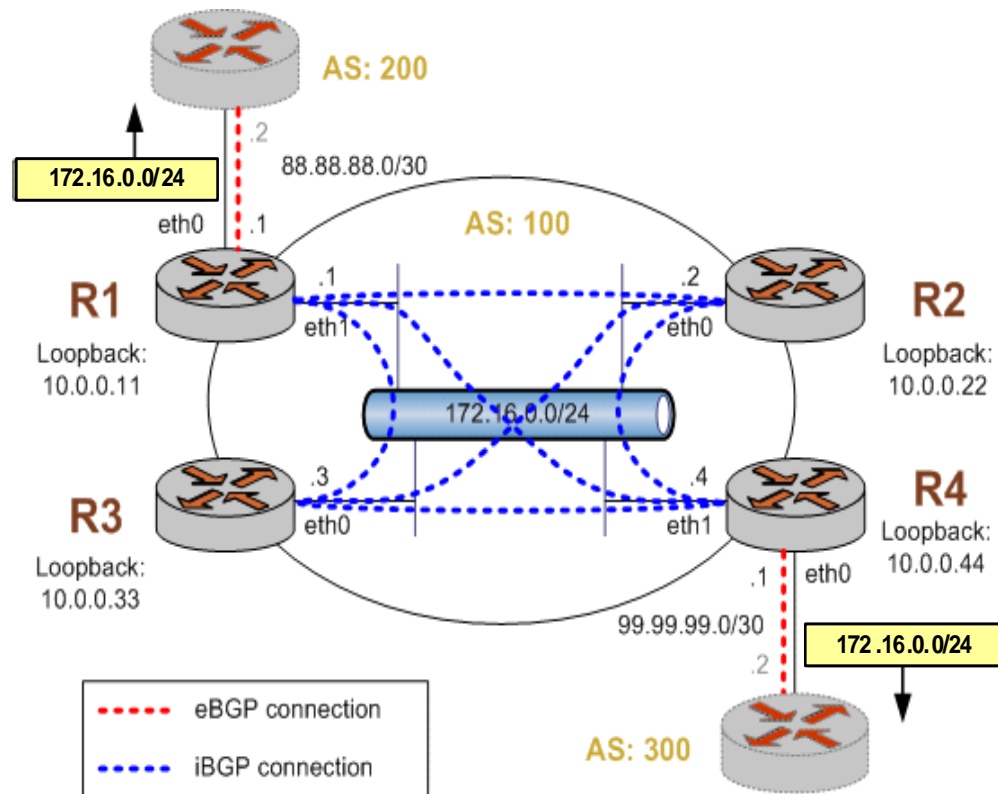
One of the main requirements for BGP configurations is to originate network prefixes to BGP peers. On the Vyatta router this is accomplished using the **network** option within the BGP configuration.

In this section, you originate the network prefix from both the R1 and R4 routers. This is shown in [Figure 1-8](#).

NOTE We assume that the routers in AS200 and AS300 are configured appropriately as eBGP peers and have routes to advertise as well.

NOTE The example in this section assumes that the desired network to originate to our BGP peers is 172.16.0.0/24, which is a private RFC 1918 network address. Typically, the router would be originating a public IP network assigned by an Internet routing registry such as the American Registry for Internet Numbers (ARIN).

Figure 1-8 Originating a route to eBGP neighbors



This example assumes that the configurations in previous sections have been performed.

To originate a route to eBGP neighbors, perform the following steps in configuration mode.

Example 1-8 Originating routes to eBGP neighbors

Router	Step	Command(s)
R1	Advertize the local network to BGP.	<code>vyatta@R1# set protocols bgp 100 network 172.16.0.0/24</code>
R1	Commit the configuration.	<code>vyatta@R1# commit</code>

Example 1-8 Originating routes to eBGP neighbors

R1	Verify the configuration.	<pre>vyatta@R1# show protocols bgp bgp 100 { neighbor 10.0.0.22 { remote-as 100 update-source 10.0.0.11 } neighbor 10.0.0.33 { remote-as 100 update-source 10.0.0.11 } neighbor 10.0.0.44 { remote-as 100 update-source 10.0.0.11 } network 172.16.0.0/24 { } parameters { router-id 10.0.0.11 } }</pre>
R4	Advertize the local network to BGP.	<pre>vyatta@R4# set protocols bgp 100 network 172.16.0.0/24</pre>
R4	Commit the configuration.	<pre>vyatta@R4# commit</pre>

Verifying the Route Origination

The following commands can be used to verify the route origination configuration.

R1: show ip bgp summary

[Example 1-9](#) shows the output of the `show ip bgp summary` command for router R1 at this stage of the configuration. The **MsgSent** column indicates that the router has been sending BGP messages, showing the number of BGP messages that have been sent to each peer.

Example 1-9 Verifying route origination on R1: “show ip bgp summary”

```
vyatta@R1:~$ show ip bgp summary
BGP router identifier 10.0.0.11, local AS number 100
RIB entries 25, using 1600 bytes of memory
Peers 4, using 10080 bytes of memory
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.0.0.22	4	100	50	55	0	0	0	00:48:02	0
10.0.0.33	4	100	50	55	0	0	0	00:48:01	0
10.0.0.44	4	100	54	58	0	0	0	00:16:30	7
88.88.88.2	4	200	4	5	0	0	0	00:11:01	5

```
Total number of neighbors 4
vyatta@R1:~$
```

Additionally you may notice that the **MsgRcvd** and **MsgSent** fields for peer 88.88.88.2 shows “4” and “5” respectively. This shows that R1 has received four BGP messages from 88.88.88.2 and sent it five, which are associated with the connection to the peer (the **show ip bgp neighbors** command can be used to provide additional detail). Also, you can see that router R1 has received seven route prefixes (**PfxRcd** column) from 10.0.0.44 and five route prefixes from 88.88.88.2. The prefixes received can be seen using the **show ip bgp** command as follows.

R1: show ip bgp

[Example 1-10](#) shows the output of the **show ip bgp** command for router R1 at this stage of the configuration.

Example 1-10 Verifying eBGP on R1: “show ip bgp”

```
vyatta@R1:~$ show ip bgp
BGP table version is 0, local router ID is 10.0.0.11
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 2.0.0.0/24       88.88.88.2             0           0 200 i
*> 2.1.0.0/24       88.88.88.2             0           0 200 i
*> 2.2.0.0/24       88.88.88.2             0           0 200 i
*>i3.0.0.0/24       99.99.99.2             0          100    0 300 i
*>i3.1.0.0/24       99.99.99.2             0          100    0 300 i
*>i3.2.0.0/24       99.99.99.2             0          100    0 300 i
*> 12.0.0.0         88.88.88.2             0           0 200 i
*>i13.0.0.0/24     99.99.99.2             0          100    0 300 i
*> 88.88.88.0/30    88.88.88.2             0           0 200 i
*>i99.99.99.0/30    99.99.99.2             0          100    0 300 i
*>i172.16.128.0/24 99.99.99.2             0          100    0 300 i
*>i192.168.2.0     99.99.99.2             0          100    0 300 i
```

```
Total number of prefixes 12
vyatta@R1:~$
```

From this output we can see that R1 knows about twelve prefixes within BGP, five from AS200 and seven from AS300.

There are two symbols that are displayed at the beginning of each of the prefixes shown in the output of **show ip bgp**. The first symbol is the status code for a valid route, which is an asterisk (“*”). Essentially all routes shown in the BGP table should be preceded by this symbol. The second symbol is the greater than character (“>”), which indicates which path is the best available path as determined by the BGP best path selection process. The **show ip bgp** command shows only the best path to each peer.

R1: show ip route bgp

[Example 1-11](#) shows the output of the **show ip route bgp** command for router R1 at this stage of the configuration.

Example 1-11 Verifying eBGP on R1: “show ip route bgp”

```
vyatta@R1:~$ show ip route bgp
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       I - ISIS, B - BGP, > - selected route, * - FIB route

B>* 2.0.0.0/24 [20/0] via 88.88.88.2, eth0, 00:06:28
B>* 2.1.0.0/24 [20/0] via 88.88.88.2, eth0, 00:06:28
B>* 2.2.0.0/24 [20/0] via 88.88.88.2, eth0, 00:06:28
B>* 3.0.0.0/24 [200/0] via 99.99.99.2, eth1 (recursive via 172.16.0.4), 00:06:56
B>* 3.1.0.0/24 [200/0] via 99.99.99.2, eth1 (recursive via 172.16.0.4), 00:06:56
B>* 3.2.0.0/24 [200/0] via 99.99.99.2, eth1 (recursive via 172.16.0.4), 00:06:56
B>* 12.0.0.0/8 [20/0] via 88.88.88.2, eth0, 00:06:28
B>* 13.0.0.0/24 [200/0] via 99.99.99.2, eth1 (recursive via 172.16.0.4), 00:06:56
B  88.88.88.0/30 [20/0] via 88.88.88.2 inactive, 00:06:28
B  99.99.99.0/30 [200/0] via 99.99.99.2 inactive, 00:06:56
B>* 172.16.128.0/24 [200/0] via 99.99.99.2, eth1 (recursive via 172.16.0.4), 00:06:56
B>* 192.168.2.0/24 [200/0] via 99.99.99.2, eth1 (recursive via 172.16.0.4), 00:06:56
vyatta@R1:~$
```

The **show ip route bgp** command displays the routes in the RIB that were learned via BGP. This is different from the output of **show ip bgp**, which shows all paths learned via BGP regardless of whether it is the best BGP path and whether the BGP candidate to the RIB for the prefix is the best route (for example, it has the lowest admin cost).

The output for the same operational BGP commands run on router R4 yields similar results.

R4: show ip bgp summary

[Example 1-12](#) shows the output of the `show ip bgp summary` command for router R4 at this stage of the configuration.

Example 1-12 Verifying eBGP on R4: “show ip bgp summary”

```
vyatta@R4:~$ show ip bgp summary
BGP router identifier 10.0.0.44, local AS number 100
RIB entries 23, using 1472 bytes of memory
Peers 4, using 10080 bytes of memory

Neighbor      V    AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.0.11     4   100   511    512     0    0    0 00:13:01    5
10.0.0.22     4   100   495    507     0    0    0 08:12:22    0
10.0.0.33     4   100   492    511     0    0    0 08:01:00    0
99.99.99.2    4   300    11     12     0    0    0 00:08:03    7

Total number of neighbors 4
vyatta@R4:~$
```

R4: show ip bgp

[Example 1-13](#) shows the output of the `show ip bgp` command for router R4 at this stage of the configuration.

Example 1-13 Verifying eBGP on R4: “show ip bgp ”

```
vyatta@R4:~$ show ip bgp
BGP table version is 0, local router ID is 10.0.0.44
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

Network        Next Hop        Metric LocPrf Weight Path
*>i2.0.0.0/24   88.88.88.2      0     100     0 200 i
*>i2.1.0.0/24   88.88.88.2      0     100     0 200 i
*>i2.2.0.0/24   88.88.88.2      0     100     0 200 i
*> 3.0.0.0/24   99.99.99.2      0                      0 300 i
*> 3.1.0.0/24   99.99.99.2      0                      0 300 i
```

```
*> 3.2.0.0/24      99.99.99.2      0      0 300 i
*>i12.0.0.0        88.88.88.2      0     100  0 200 i
*> 13.0.0.0/24    99.99.99.2      0      0 300 i
*>i88.88.88.0/30  88.88.88.2      0     100  0 200 i
*> 99.99.99.0/30  99.99.99.2      0      0 300 i
*> 172.16.128.0/24 99.99.99.2      0      0 300 i
*> 192.168.2.0    99.99.99.2      0      0 300 i
```

```
Total number of prefixes 12
vyatta@R4:~$
```

Router R4's BGP table contains the paths it learned from its eBGP peer, as well as the paths it learned from its iBGP neighbor R1.

Inbound Route Filtering

A common requirement for BGP configurations is to filter inbound routing announcements from a BGP peer. On the Vyatta system this is accomplished using routing policies that are then applied to the BGP process as “import” policies. In this instance we use prefix lists in conjunction with route maps to accomplish this.

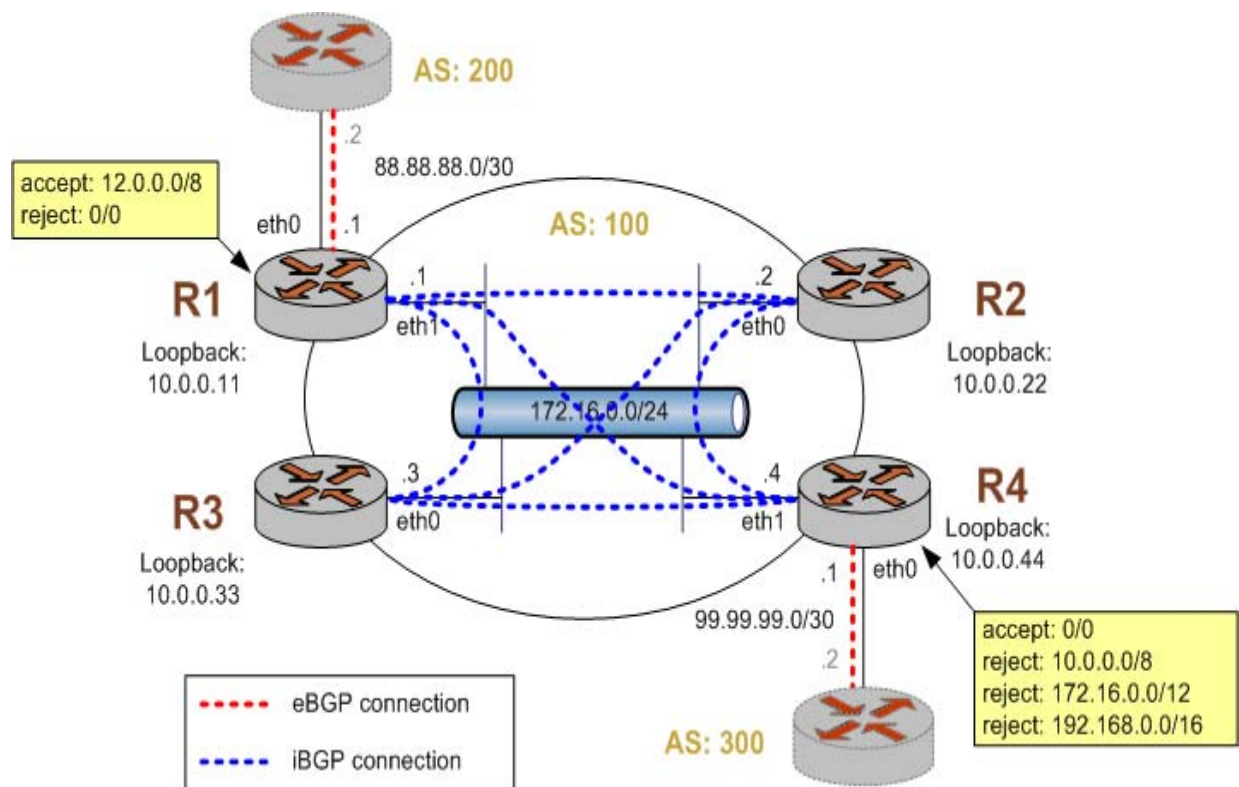
[Example 1-14](#) creates the following inbound filtering policies:

- R1 should only accept network 12.0.0.0/8 from its eBGP peer, and reject everything else.
- R4 should allow all Internet routes, but reject all RFC 1918 networks from its eBGP peer.

This import policy is shown in [Figure 1-9](#).

NOTE We assume that the routers in AS100 have been configured for iBGP and eBGP as shown and that the routers in AS200 and AS300 are configured appropriately as eBGP peers.

Figure 1-9 Filtering inbound routes



To create this inbound route filter, perform the following steps in configuration mode.

Example 1-14 Creating an import policy

Router	Step	Command(s)
R1	Create a list of prefixes to allow. In this case we just have one - 12.0.0.0/8.	vyatta@R1# set policy prefix-list ALLOW-PREFIXES rule 1 action permit vyatta@R1# set policy prefix-list ALLOW-PREFIXES rule 1 prefix 12.0.0.0/8
R1	Create a route map rule to permit all prefixes in our list.	vyatta@R1# set policy route-map eBGP-IMPORT rule 10 action permit vyatta@R1# set policy route-map eBGP-IMPORT rule 10 match ip address prefix-list ALLOW-PREFIXES
R1	Create a route map rule to deny all other prefixes.	vyatta@R1# set policy route-map eBGP-IMPORT rule 20 action deny

Example 1-14 Creating an import policy

R1	Assign the route map policy created as the import route map policy for AS 200.	<pre>vyatta@R1# set protocols bgp 100 neighbor 88.88.88.2 route-map import eBGP-IMPORT</pre>
R1	Commit the configuration.	<pre>vyatta@R1# commit</pre>
R1	Reset the BGP session to the peer so that the new policies are enabled.	<pre>vyatta@R1# run reset ip bgp 88.88.88.2</pre>
R1	Display the policy configuration.	<pre>vyatta@R1# show policy prefix-list ALLOW-PREFIXES { rule 1 { action permit prefix 12.0.0.0/8 } } route-map eBGP-IMPORT { rule 10 { action permit match { ip { address { prefix-list ALLOW-PREFIXES } } } } rule 20 { action deny } } vyatta@R1#</pre>
R1	Display the BGP configuration for eBGP neighbor 88.88.88.2.	<pre>vyatta@R1# show protocols bgp 100 neighbor 88.88.88.2 remote-as 200 route-map { import eBGP-IMPORT } vyatta@R1#</pre>

Example 1-14 Creating an import policy

R4	Create a rule to match any prefix from 10.0.0.0/8 to 32.	<pre>vyatta@R4# set policy prefix-list RFC1918PREFIXES rule 1 action permit vyatta@R4# set policy prefix-list RFC1918PREFIXES rule 1 le 32 vyatta@R4# set policy prefix-list RFC1918PREFIXES rule 1 prefix 10.0.0.0/8</pre>
R4	Create a rule to match any prefix from 172.16.0.0/12 to 32.	<pre>vyatta@R4# set policy prefix-list RFC1918PREFIXES rule 2 action permit vyatta@R4# set policy prefix-list RFC1918PREFIXES rule 2 le 32 vyatta@R4# set policy prefix-list RFC1918PREFIXES rule 2 prefix 172.16.0.0/12</pre>
R4	Create a rule to match any prefix from 192.168.0.0/16 to 32.	<pre>vyatta@R4# set policy prefix-list RFC1918PREFIXES rule 3 action permit vyatta@R4# set policy prefix-list RFC1918PREFIXES rule 3 le 32 vyatta@R4# set policy prefix-list RFC1918PREFIXES rule 3 prefix 192.168.0.0/16</pre>
R4	Create a route map rule to deny all prefixes in our list.	<pre>vyatta@R4# set policy route-map eBGP-IMPORT rule 10 action deny vyatta@R4# set policy route-map eBGP-IMPORT rule 10 match ip address prefix-list RFC1918PREFIXES</pre>
R4	Create a route map rule to permit all other prefixes.	<pre>vyatta@R4# set policy route-map eBGP-IMPORT rule 20 action permit</pre>
R4	Assign the route map policy created as the import route map policy for AS 300.	<pre>vyatta@R4# set protocols bgp 100 neighbor 99.99.99.2 route-map import eBGP-IMPORT</pre>
R4	Commit the configuration.	<pre>vyatta@R4# commit</pre>
R4	Reset the BGP session to the peer so that the new policies are enabled.	<pre>vyatta@R1# run reset ip bgp 99.99.99.2</pre>

Example 1-14 Creating an import policy

R4	Display the policy configuration.	<pre>vyatta@R4# show policy prefix-list RFC1918PREFIXES { rule 1 { action permit le 32 prefix 10.0.0.0/8 } rule 2 { action permit le 32 prefix 172.16.0.0/12 } rule 3 { action permit le 32 prefix 192.168.0.0/16 } } route-map eBGP-IMPORT { rule 10 { action deny match { ip { address { prefix-list RFC1918PREFIXES } } } } rule 20 { action permit } } vyatta@R4#</pre>
R4	Display the BGP configuration for eBGP neighbor 99.99.99.2.	<pre>vyatta@R4# show protocols bgp 100 neighbor 99.99.99.2 remote-as 300 route-map { import eBGP-IMPORT } vyatta@R4#</pre>

Verifying the Inbound Filter

The following commands can be used to verify the inbound filter configuration.

R1: show ip bgp

[Example 1-15](#) shows R1's BGP table *before* the import filter is applied.

Example 1-15 R1 inbound BGP routes before import filtering

```
vyatta@R1:~$ show ip bgp
BGP table version is 0, local router ID is 10.0.0.11
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 2.0.0.0/24       88.88.88.2         0           0 200 i
*> 2.1.0.0/24       88.88.88.2         0           0 200 i
*> 2.2.0.0/24       88.88.88.2         0           0 200 i
*>i3.0.0.0/24       99.99.99.2         0          100    0 300 i
*>i3.1.0.0/24       99.99.99.2         0          100    0 300 i
*>i3.2.0.0/24       99.99.99.2         0          100    0 300 i
*> 12.0.0.0         88.88.88.2         0           0 200 i
*>i13.0.0.0/24     99.99.99.2         0          100    0 300 i
*> 88.88.88.0/30    88.88.88.2         0           0 200 i
*>i99.99.99.0/30    99.99.99.2         0          100    0 300 i
*> 172.16.0.0/24    0.0.0.0            1           32768 i
* i                 10.0.0.44          1          100    0 i
*>i172.16.128.0/24  99.99.99.2         0          100    0 300 i
*>i192.168.2.0     99.99.99.2         0          100    0 300 i

Total number of prefixes 13
vyatta@R1:~$
```

R1: show ip bgp

[Example 1-16](#) shows R1's BGP table *after* the import filter is applied.

Example 1-16 R1 inbound BGP routes after import filtering

```
vyatta@R1:~$ show ip bgp
BGP table version is 0, local router ID is 10.0.0.11
```

Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
 r RIB-failure, S Stale, R Removed
 Origin codes: i - IGP, e - EGP, ? - incomplete

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i3.0.0.0/24	99.99.99.2	0	100	0 300	i
*>i3.1.0.0/24	99.99.99.2	0	100	0 300	i
*>i3.2.0.0/24	99.99.99.2	0	100	0 300	i
*> 12.0.0.0	88.88.88.2	0		0 200	i
*>i13.0.0.0/24	99.99.99.2	0	100	0 300	i
*>i99.99.99.0/30	99.99.99.2	0	100	0 300	i
*> 172.16.0.0/24	0.0.0.0	1		32768	i
* i	10.0.0.44	1	100	0	i
*>i172.16.128.0/24	99.99.99.2	0	100	0 300	i
*>i192.168.2.0	99.99.99.2	0	100	0 300	i

Total number of prefixes 9
 vyatta@R1:~\$

Note that only 12.0.0.0 from 88.88.88.2 is still in the table.

R4: show ip bgp

Example 1-17 shows R4's BGP table *before* the import filter is applied.

Example 1-17 R4 inbound BGP routes before import filtering

```
vyatta@R4:~$ show ip bgp
BGP table version is 0, local router ID is 10.0.0.44
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 3.0.0.0/24	99.99.99.2	0		0 300	i
*> 3.1.0.0/24	99.99.99.2	0		0 300	i
*> 3.2.0.0/24	99.99.99.2	0		0 300	i
*>i12.0.0.0	88.88.88.2	0	100	0 200	i
*> 13.0.0.0/24	99.99.99.2	0		0 300	i
*> 99.99.99.0/30	99.99.99.2	0		0 300	i
* i172.16.0.0/24	10.0.0.11	1	100	0	i
*>	0.0.0.0	1		32768	i
*> 172.16.128.0/24	99.99.99.2	0		0 300	i
*> 192.168.2.0	99.99.99.2	0		0 300	i

Total number of prefixes 9


```
vyatta@R4:~$
```

R4: show ip bgp

The output below shows R4's BGP table *after* the import filter is applied.

Example 1-18 R4 inbound BGP routes after import filtering

```
vyatta@R4:~$ show ip bgp
BGP table version is 0, local router ID is 10.0.0.44
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 3.0.0.0/24       99.99.99.2         0           0 300 i
*> 3.1.0.0/24       99.99.99.2         0           0 300 i
*> 3.2.0.0/24       99.99.99.2         0           0 300 i
*>i12.0.0.0         88.88.88.2         0      100     0 200 i
*> 13.0.0.0/24     99.99.99.2         0           0 300 i
*> 99.99.99.0/30   99.99.99.2         0           0 300 i
* i172.16.0.0/24   10.0.0.11          1      100     0 i
*>                  0.0.0.0            1           32768 i

Total number of prefixes 7
vyatta@R4:~$
```

Outbound Route Filtering

Filtering outbound prefixes is another common BGP configuration requirement. On the Vyatta system this is accomplished using routing policies that are then applied to the BGP process as “export” policies.

The example in this section assumes that AS100 does not want to be a transit AS for AS 200 or AS 300. This means that:

- eBGP routes from R1's eBGP peer (AS 200) should not be sent to R4's eBGP peer.
- Routes from R4's eBGP peer (AS 300) should not be sent to R1's eBGP peer.

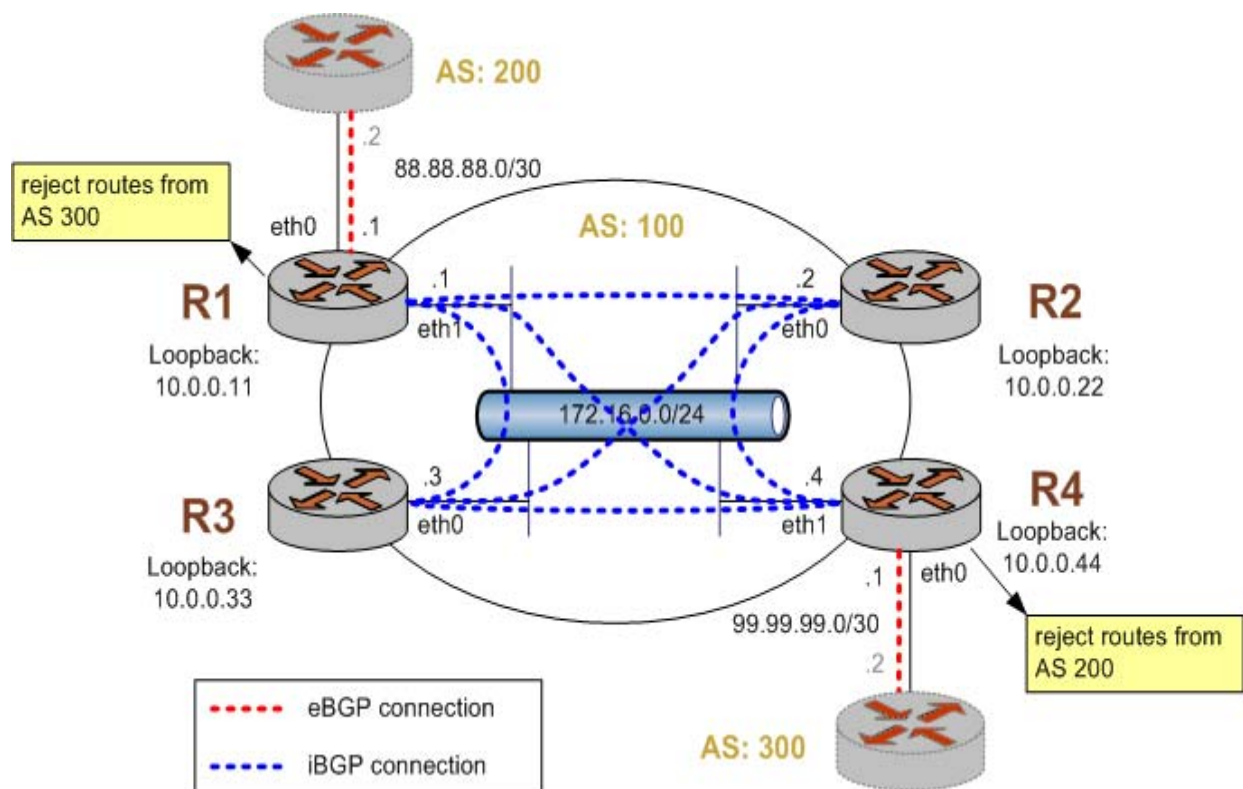
If we *did not* implement this filtering, AS 300 might send traffic destined for AS 200 to router R4, and this traffic would then be carried across the AS 100 network.

There are several ways that this routing policy could be implemented: two most common are basing the filter on the network prefix or basing it on the AS Path. In this example, we update the existing BGP export policy to add some additional restrictions that will prevent AS 100 from acting as a transit network for AS 200 and AS 300.

This export policy is shown in [Figure 1-10](#).

NOTE We assume that the routers in AS100 have been configured for iBGP and eBGP as shown and that the routers in AS200 and AS300 are configured appropriately as eBGP peers.

Figure 1-10 Filtering outbound routes



To create this export policy, perform the following steps in configuration mode.

Example 1-19 Creating an export policy

Router	Step	Command(s)
R1	Create a list of AS paths to deny. In this case we just have one - AS300.	vyatta@R1# set policy as-path-list AS300 rule 1 action permit vyatta@R1# set policy as-path-list AS300 rule 1 regex 300

Example 1-19 Creating an export policy

R1	Create a route map rule to deny all AS paths in our list.	vyatta@R1# set policy route-map eBGP-EXPORT rule 10 action deny vyatta@R1# set policy route-map eBGP-EXPORT rule 10 match as-path AS300
R1	Create a route map rule to permit all other prefixes.	vyatta@R1# set policy route-map eBGP-EXPORT rule 20 action permit
R1	Assign the route map policy created as the export route map policy for AS 200.	vyatta@R1# set protocols bgp 100 neighbor 88.88.88.2 route-map export eBGP-EXPORT
R1	Commit the configuration.	vyatta@R1# commit
R1	Reset the BGP session to the peer so that the new policies are enabled.	vyatta@R1# run reset ip bgp 88.88.88.2
R1	Display the policy configurations	vyatta@R1# show policy as-path-list AS300 rule 1 { action permit regex 300 } vyatta@R1# show policy route-map eBGP-EXPORT rule 10 { action deny match { as-path AS300 } } rule 20 { action permit } vyatta@R1#

Example 1-19 Creating an export policy

R1	Display the BGP configuration for eBGP neighbor 88.88.88.2.	<pre>vyatta@R1# show protocols bgp 100 neighbor 88.88.88.2 remote-as 200 route-map { export eBGP-EXPORT import eBGP-IMPORT } vyatta@R1#</pre>
R4	Create a list of AS paths to deny. In this case we just have one - AS200.	<pre>vyatta@R4# set policy as-path-list AS200 rule 1 action permit vyatta@R4# set policy as-path-list AS200 rule 1 regex 200</pre>
R4	Create a route map rule to deny all AS paths in our list.	<pre>vyatta@R4# set policy route-map eBGP-EXPORT rule 10 action deny vyatta@R4# set policy route-map eBGP-EXPORT rule 10 match as-path AS200</pre>
R4	Create a route map rule to permit all other prefixes.	<pre>vyatta@R4# set policy route-map eBGP-EXPORT rule 20 action permit</pre>
R4	Assign the route map policy created as the export route map policy for AS 300.	<pre>vyatta@R4# set protocols bgp 100 neighbor 99.99.99.2 route-map export eBGP-EXPORT</pre>
R4	Commit the configuration.	<pre>vyatta@R4# commit</pre>
R4	Reset the BGP session to the peer so that the new policies are enabled.	<pre>vyatta@R4# run reset ip bgp 99.99.99.2</pre>

Example 1-19 Creating an export policy

```

R4      Display the      vyatta@R4# show policy as-path-list AS200
policy  configurations
        .
        }
        vyatta@R4# show policy route-map eBGP-EXPORT
        rule 10 {
            action deny
            match {
                as-path AS200
            }
        }
        rule 20 {
            action permit
        }
        vyatta@R4#

R4      Display the      vyatta@R4# show protocols bgp 100 neighbor 99.99.99.2
BGP     BGP             remote-as 300
configuration
for eBGP route-map {
neighbor   export eBGP-EXPORT
99.99.99.2 import eBGP-IMPORT
        }
        vyatta@R4#

```

Verifying the Outbound Filter

The following commands can be used to verify the outbound filter configuration.

AS 200: show ip bgp

[Example 1-20](#) shows AS 200's BGP table *before* the export filter is applied.

Example 1-20 AS 200 outbound BGP routes before export filtering

```

vyatta@AS200:~$ show ip bgp
BGP table version is 0, local router ID is 10.0.11.11
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop           Metric LocPrf Weight Path
*> 2.0.0.0/24     0.0.0.0             0         32768 i
*> 2.1.0.0/24     0.0.0.0             0         32768 i

```

```

*> 2.2.0.0/24      0.0.0.0          0          32768 i
*> 3.0.0.0/24      88.88.88.1       0 100 300 i
*> 3.1.0.0/24      88.88.88.1       0 100 300 i
*> 3.2.0.0/24      88.88.88.1       0 100 300 i
*> 12.0.0.0        0.0.0.0          0          32768 i
*> 13.0.0.0/24     88.88.88.1       0 100 300 i
*> 88.88.88.0/30   0.0.0.0          0          32768 i
*> 99.99.99.0/30   88.88.88.1       0 100 300 i
*> 172.16.0.0/24   88.88.88.1       1          0 100 i

```

Total number of prefixes 11

vyatta@AS200:~\$

AS 200: show ip bgp

[Example 1-21](#) shows AS 200's BGP table *after* the export filter is applied.

Example 1-21 AS 200 outbound BGP routes after export filtering

```

vyatta@AS200:~$ show ip bgp
BGP table version is 0, local router ID is 10.0.11.11
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 2.0.0.0/24	0.0.0.0	0		32768	i
*> 2.1.0.0/24	0.0.0.0	0		32768	i
*> 2.2.0.0/24	0.0.0.0	0		32768	i
*> 12.0.0.0	0.0.0.0	0		32768	i
*> 88.88.88.0/30	0.0.0.0	0		32768	i
*> 172.16.0.0/24	88.88.88.1	1		0 100	i

Total number of prefixes 6

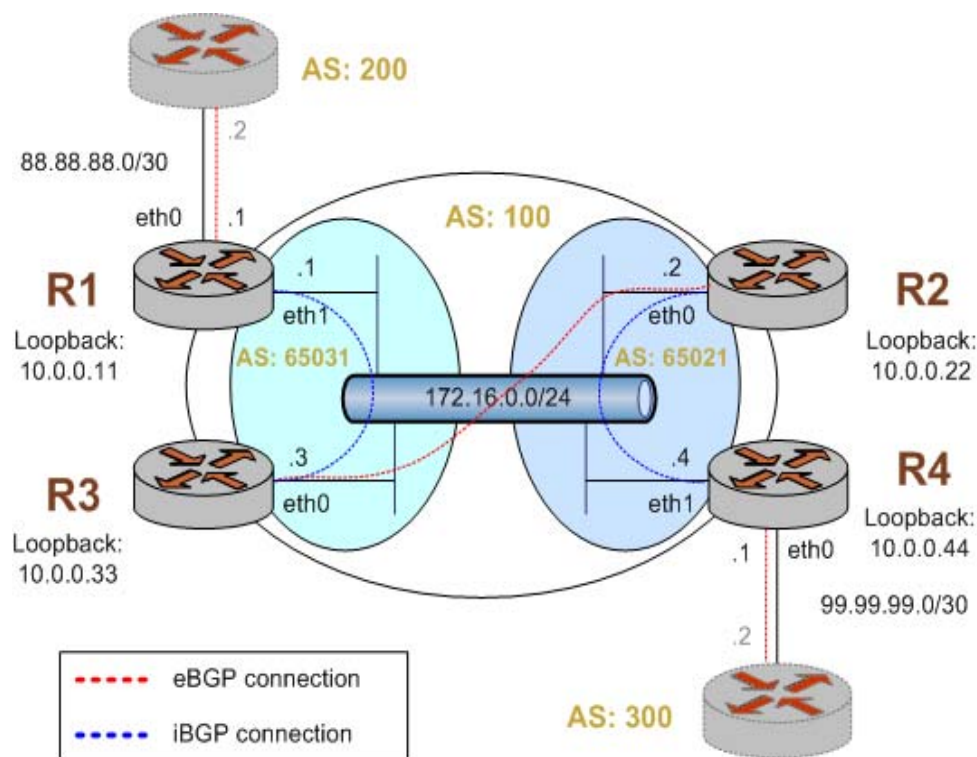
vyatta@AS200:~\$

Confederations

Confederations allow large Autonomous Systems to sub-divide the AS into sub-ASs. This helps solve the scalability issues associated with having to maintain a full mesh of iBGP connections between all iBGP routers in the AS. In the confederation example shown in [Figure 1-11](#), routers R1 and R3 are configured in one sub-AS (AS number 65031) and routers R2 and R4 are configured in a different sub-AS (AS number 65021).

NOTE We assume that the routers in AS200 and AS300 are configured appropriately as eBGP peers.

Figure 1-11 BGP confederation



This example assumes that the configurations in previous sections have been performed.

To create the confederation shown in [Figure 1-11](#), perform the following steps in configuration mode.

Example 1-22 Creating a BGP confederation

Router	Step	Command(s)
R1	Delete current BGP configuration.	<code>vyatta@R1# delete protocols bgp 100</code>
R1	To allow routes from AS200 to be injected into RIB on R3.	<code>vyatta@R1# set protocols bgp 65031 neighbor 10.0.0.33 nexthop-self</code>

Example 1-22 Creating a BGP confederation

R1	Set R3 in the same sub-AS as R1.	<pre>vyatta@R1# set protocols bgp 65031 neighbor 10.0.0.33 remote-as 65031</pre>
R1	Define the IP address on the local R1 router that is used to peer with the R3 router.	<pre>vyatta@R1# set protocols bgp 65031 neighbor 10.0.0.33 update-source 10.0.0.11</pre>
R1	Set the AS200 router in a different AS than R1.	<pre>vyatta@R1# set protocols bgp 65031 neighbor 88.88.88.2 remote-as 200</pre>
R1	Set the outbound filter.	<pre>vyatta@R1# set protocols bgp 65031 neighbor 88.88.88.2 route-map export eBGP-EXPORT</pre>
R1	Set the inbound filter.	<pre>vyatta@R1# set protocols bgp 65031 neighbor 88.88.88.2 route-map import eBGP-IMPORT</pre>
R1	Set the network to advertise.	<pre>vyatta@R1# set protocols bgp 65031 network 172.16.0.0/24</pre>
R1	Set the AS identifier for the confederation.	<pre>vyatta@R1# set protocols bgp 65031 parameters confederation identifier 100</pre>
R1	Set the peer for the sub-AS.	<pre>vyatta@R1# set protocols bgp 65031 parameters confederation peers 65021</pre>
R1	Set the router ID.	<pre>vyatta@R1# set protocols bgp 65031 parameters router-id 10.0.0.11</pre>
R1	Commit the configuration.	<pre>vyatta@R1# commit</pre>

Example 1-22 Creating a BGP confederation

R1	Display the BGP configuration.	<pre>vyatta@R1# show protocols bgp 65031 { neighbor 10.0.0.33 { nexthop-self remote-as 65031 update-source 10.0.0.11 } neighbor 88.88.88.2 { remote-as 200 route-map { export eBGP-EXPORT import eBGP-IMPORT } } network 172.16.0.0/24 { } parameters { confederation { identifier 100 peers 65021 } router-id 10.0.0.11 } } vyatta@R1#</pre>
R2	Delete current BGP configuration.	<pre>vyatta@R2# delete protocols bgp 100</pre>
R2	Set R3 in a different sub-AS than R2.	<pre>vyatta@R2# set protocols bgp 65021 neighbor 10.0.0.33 remote-as 65031</pre>
R2	Define the IP address on the local R2 router that is used to peer with the R3 router.	<pre>vyatta@R2# set protocols bgp 65021 neighbor 10.0.0.33 update-source 10.0.0.22</pre>
R2	Set R4 in the same sub-AS as R2.	<pre>vyatta@R2# set protocols bgp 65021 neighbor 10.0.0.44 remote-as 65021</pre>

Example 1-22 Creating a BGP confederation

R2	Define the IP address on the local R2 router that is used to peer with the R4 router.	<pre>vyatta@R2# set protocols bgp 65021 neighbor 10.0.0.44 update-source 10.0.0.22</pre>
R2	Set the network to advertise.	<pre>vyatta@R2# set protocols bgp 65021 network 172.16.0.0/24</pre>
R2	Set the AS identifier for the confederation.	<pre>vyatta@R2# set protocols bgp 65021 parameters confederation identifier 100</pre>
R2	Set the peer for the sub-AS.	<pre>vyatta@R2# set protocols bgp 65021 parameters confederation peers 65031</pre>
R2	Set the router ID.	<pre>vyatta@R2# set protocols bgp 65021 parameters router-id 10.0.0.22</pre>
R2	Commit the configuration.	<pre>vyatta@R2# commit</pre>
R2	Display the BGP configuration.	<pre>vyatta@R2# show protocols bgp 65021 { neighbor 10.0.0.33 { remote-as 65031 update-source 10.0.0.22 } neighbor 10.0.0.44 { remote-as 65021 update-source 10.0.0.22 } network 172.16.0.0/24 { } parameters { confederation { identifier 100 peers 65031 } router-id 10.0.0.22 } } vyatta@R2#</pre>

Example 1-22 Creating a BGP confederation

R3	Delete current BGP configuration.	<pre>vyatta@R3# delete protocols bgp 100</pre>
R3	Set R1 in the same sub-AS as R3.	<pre>vyatta@R3# set protocols bgp 65031 neighbor 10.0.0.11 remote-as 65031</pre>
R3	Define the IP address on the local R3 router that is used to peer with the R1 router.	<pre>vyatta@R3# set protocols bgp 65031 neighbor 10.0.0.11 update-source 10.0.0.33</pre>
R3	Set the R2 router in a different AS than R3.	<pre>vyatta@R3# set protocols bgp 65031 neighbor 10.0.0.22 remote-as 65021</pre>
R3	Define the IP address on the local R3 router that is used to peer with the R2 router.	<pre>vyatta@R3# set protocols bgp 65031 neighbor 10.0.0.22 update-source 10.0.0.33</pre>
R3	Set the network to advertise.	<pre>vyatta@R3# set protocols bgp 65031 network 172.16.0.0/24</pre>
R3	Set the AS identifier for the confederation.	<pre>vyatta@R3# set protocols bgp 65031 parameters confederation identifier 100</pre>
R3	Set the peer for the sub-AS.	<pre>vyatta@R3# set protocols bgp 65031 parameters confederation peers 65021</pre>
R3	Set the router ID.	<pre>vyatta@R3# set protocols bgp 65031 parameters router-id 10.0.0.33</pre>
R3	Commit the configuration.	<pre>vyatta@R3# commit</pre>

Example 1-22 Creating a BGP confederation

R3	Display the BGP configuration.	<pre>vyatta@R3# show protocols bgp 65031 { neighbor 10.0.0.11 { remote-as 65031 update-source 10.0.0.33 } neighbor 10.0.0.22 { remote-as 65021 update-source 10.0.0.33 } network 172.16.0.0/24 { } parameters { confederation { identifier 100 peers 65021 } router-id 10.0.0.33 } } vyatta@R3#</pre>
R4	Delete current BGP configuration.	<pre>vyatta@R4# delete protocols bgp 100</pre>
R4	To allow routes from AS300 to be injected into RIB on R2.	<pre>vyatta@R4# set protocols bgp 65021 neighbor 10.0.0.22 nexthop-self</pre>
R4	Set R2 in the same sub-AS as R4.	<pre>vyatta@R4# set protocols bgp 65021 neighbor 10.0.0.22 remote-as 65021</pre>
R4	Define the IP address on the local R4 router that is used to peer with the R2 router.	<pre>vyatta@R4# set protocols bgp 65021 neighbor 10.0.0.22 update-source 10.0.0.44</pre>

Example 1-22 Creating a BGP confederation

R4	Set the AS300 router in a different AS than R4.	vyatta@R4# set protocols bgp 65021 neighbor 99.99.99.2 remote-as 300
R4	Set the outbound filter.	vyatta@R4# set protocols bgp 65021 neighbor 99.99.99.2 route-map export eBGP-EXPORT
R4	Set the inbound filter.	vyatta@R4# set protocols bgp 65021 neighbor 99.99.99.2 route-map import eBGP-IMPORT
R4	Set the network to advertise.	vyatta@R4# set protocols bgp 65021 network 172.16.0.0/24
R4	Set the AS identifier for the confederation.	vyatta@R4# set protocols bgp 65021 parameters confederation identifier 100
R4	Set the peer for the sub-AS.	vyatta@R4# set protocols bgp 65021 parameters confederation peers 65031
R4	Set the router ID.	vyatta@R4# set protocols bgp 65021 parameters router-id 10.0.0.44
R4	Commit the configuration.	vyatta@R4# commit

Example 1-22 Creating a BGP confederation

```

R4      Display the      vyatta@R4# show protocols bgp
      BGP                65021 {
      configuration.     neighbor 10.0.0.22 {
                          nexthop-self
                          remote-as 65021
                          update-source 10.0.0.44
                        }
                        neighbor 99.99.99.2 {
                          remote-as 300
                          route-map {
                            export eBGP-EXPORT
                            import eBGP-IMPORT
                          }
                        }
                        network 172.16.0.0/24 {
                        }
                        parameters {
                          confederation {
                            identifier 100
                            peers 65031
                          }
                          router-id 10.0.0.44
                        }
                      }
      vyatta@R4#
  
```

Verifying the Confederation

The following commands can be used to verify the confederation configuration.

R1: show ip bgp summary

Example 1-23 shows the output of the `show ip bgp summary` command for router R1 at this stage of the configuration.

Example 1-23 Verifying confederations on R1: “show ip bgp summary”

```

vyatta@R1:~$ show ip bgp summary
BGP router identifier 10.0.0.11, local AS number 65031
RIB entries 13, using 832 bytes of memory
Peers 2, using 5040 bytes of memory

Neighbor      V    AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.0.33     4 65031   1159   1167     0    0  0 00:00:23      6
  
```

```
88.88.88.2      4    200    1150    1159      0    0    0 00:00:06      1
```

```
Total number of neighbors 2
vyatta@R1:~$
```

R1: show ip bgp

[Example 1-24](#) shows the output of the `show ip bgp` command for router R1 at this stage of the configuration.

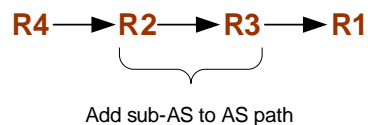
Example 1-24 Verifying confederations on R1: “show ip bgp”

```
vyatta@R1:~$ show ip bgp
BGP table version is 0, local router ID is 10.0.0.11
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*>i3.0.0.0/24      10.0.0.44          0     100     0 (65021) 300 i
*>i3.1.0.0/24      10.0.0.44          0     100     0 (65021) 300 i
*>i3.2.0.0/24      10.0.0.44          0     100     0 (65021) 300 i
*> 12.0.0.0         88.88.88.2         0           0 200 i
*>i13.0.0.0/24     10.0.0.44          0     100     0 (65021) 300 i
*>i99.99.99.0/30   10.0.0.44          0     100     0 (65021) 300 i
* i172.16.0.0/24  10.0.0.33          1     100     0 i
*>                  0.0.0.0            1           32768 i

Total number of prefixes 7
vyatta@R1:~$
```

Note that the routes learned from router R4 (Next Hop 10.0.0.44) include the confederation sub-AS in the AS Path. All confederation sub-ASs will be shown inside brackets (). This information is not transmitted outside of the true AS (AS 100).



R2: show ip bgp summary

[Example 1-25](#) shows the output of the `show ip bgp summary` command for router R2 at this stage of the configuration.

Example 1-25 Verifying confederations on R2: “show ip bgp summary”

```
vyatta@R2:~$ show ip bgp summary
BGP router identifier 10.0.0.22, local AS number 65021
RIB entries 13, using 832 bytes of memory
Peers 2, using 5040 bytes of memory

Neighbor      V    AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.0.33     4 65031   1165   1163      0    0    0 00:23:14      2
10.0.0.44     4 65021   1159   1167      0    0    0 00:23:32      6

Total number of neighbors 2
vyatta@R2:~$
```

R2: show ip bgp

[Example 1-26](#) shows the output of the `show ip bgp` command for router R2 at this stage of the configuration.

Example 1-26 Verifying confederations on R2: “show ip bgp”

```
vyatta@R2:~$ show ip bgp
BGP table version is 0, local router ID is 10.0.0.22
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*>i3.0.0.0/24       10.0.0.44          0     100     0 300 i
*>i3.1.0.0/24       10.0.0.44          0     100     0 300 i
*>i3.2.0.0/24       10.0.0.44          0     100     0 300 i
*> 12.0.0.0         10.0.0.11          0     100     0 (65031) 200 i
*>i13.0.0.0/24     10.0.0.44          0     100     0 300 i
*>i99.99.99.0/30    10.0.0.44          0     100     0 300 i
* 172.16.0.0/24    10.0.0.33          1     100     0 (65031) i
* i                10.0.0.44          1     100     0 i
*>                 0.0.0.0            1           32768 i

Total number of prefixes 7
```



```
vyatta@R2:~$
```

R3: show ip bgp summary

Example 1-27 shows the output of the `show bgp peers` command for router R3 at this stage of the configuration.

Example 1-27 Verifying confederations on R3: “show bgp peers”

```
vyatta@R3:~$ show ip bgp summary
BGP router identifier 10.0.0.33, local AS number 65031
RIB entries 13, using 832 bytes of memory
Peers 2, using 5040 bytes of memory

Neighbor      V    AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.0.11     4 65031   1172   1180     0    0   0 00:20:04    2
10.0.0.22     4 65021   1161   1177     0    0   0 00:27:51    6

Total number of neighbors 2
vyatta@R3:~$
```

R3: show ip bgp

Example 1-28 shows the output of the `show ip bgp` command for router R3 at this stage of the configuration.

Example 1-28 Verifying confederations on R3: “show ip bgp”

```
vyatta@R3:~$ show ip bgp
BGP table version is 0, local router ID is 10.0.0.33
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop           Metric LocPrf Weight Path
*> 3.0.0.0/24      10.0.0.44           0     100     0 (65021) 300 i
*> 3.1.0.0/24      10.0.0.44           0     100     0 (65021) 300 i
*> 3.2.0.0/24      10.0.0.44           0     100     0 (65021) 300 i
*>i12.0.0.0        10.0.0.11           0     100     0 200 i
*> 13.0.0.0/24    10.0.0.44           0     100     0 (65021) 300 i
*> 99.99.99.0/30  10.0.0.44           0     100     0 (65021) 300 i
* i172.16.0.0/24  10.0.0.11           1     100     0 i
```

```
*          10.0.0.22          1    100    0 (65021) i
*>         0.0.0.0           1          32768 i
```

```
Total number of prefixes 7
vyatta@R3:~$
```

R4: show ip bgp summary

[Example 1-29](#) shows the output of the `show ip bgp summary` command for router R4 at this stage of the configuration.

Example 1-29 Verifying confederations on R4: “show ip bgp summary”

```
vyatta@R4:~$ show ip bgp summary
BGP router identifier 10.0.0.44, local AS number 65021
RIB entries 13, using 832 bytes of memory
Peers 2, using 5040 bytes of memory

Neighbor      V    AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
10.0.0.22     4  65021   1165   1168      0     0    0 00:32:56      2
99.99.99.2    4   300   1155   1162      0     0    0 00:33:30      5

Total number of neighbors 2
vyatta@R4:~$
```

R4: show ip bgp

[Example 1-30](#) shows the output of the `show ip bgp` command for router R4 at this stage of the configuration.

Example 1-30 Verifying confederations on R4: “show ip bgp”

```
vyatta@R4:~$ show ip bgp
BGP table version is 0, local router ID is 10.0.0.44
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

Network        Next Hop        Metric LocPrf Weight Path
*> 3.0.0.0/24   99.99.99.2      0           0 300 i
*> 3.1.0.0/24   99.99.99.2      0           0 300 i
*> 3.2.0.0/24   99.99.99.2      0           0 300 i
```

```

*>i12.0.0.0      10.0.0.11      0    100    0 (65031) 200 i
*> 13.0.0.0/24   99.99.99.2     0          0 300 i
*> 99.99.99.0/30 99.99.99.2     0          0 300 i
* i172.16.0.0/24 10.0.0.22     1    100    0 i
*>                0.0.0.0       1          32768 i

```

```

Total number of prefixes 7
vyatta@R4:~$

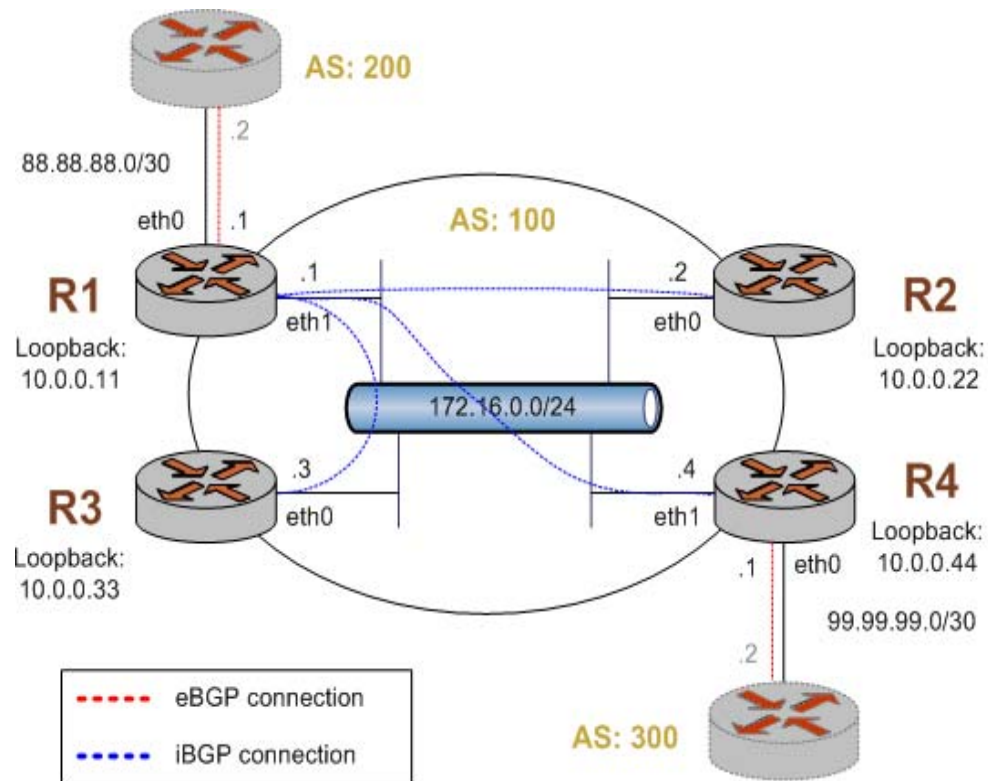
```

Route Reflectors

Router reflectors are another technology designed to help BGP scale to large Autonomous Systems. In a route reflector configuration there is at least one route reflector server and one or more route reflector clients. In the example shown in [Figure 1-12](#), router R1 is the route reflector server and router R2, R3, and R4 are the route reflector clients.

NOTE We assume that the routers in AS200 and AS300 are configured appropriately as eBGP peers.

Figure 1-12 BGP route reflector



This example assumes that the configurations in previous sections have been performed, and that interfaces and OSPF have been configured. If you are starting from a clean base system you need not delete previous configuration.

Example 1-31 Creating route reflectors

Router	Step	Command(s)
R1	Delete current BGP configuration.	vyatta@R1# delete protocols bgp
R1	To allow routes from AS200 to be injected into RIB on R2.	vyatta@R1# set protocols bgp 100 neighbor 10.0.0.22 nexthop-self
R1	Set R2 in the same AS as R1.	vyatta@R1# set protocols bgp 100 neighbor 10.0.0.22 remote-as 100
R1	Set R2 as a route reflector client.	vyatta@R1# set protocols bgp 100 neighbor 10.0.0.22 route-reflector-client
R1	Define the IP address on the local R1 router that is used to peer with the R2 router.	vyatta@R1# set protocols bgp 100 neighbor 10.0.0.22 update-source 10.0.0.11
R1	To allow routes from AS200 to be injected into RIB on R3.	vyatta@R1# set protocols bgp 100 neighbor 10.0.0.33 nexthop-self
R1	Set R3 in the same AS as R1.	vyatta@R1# set protocols bgp 100 neighbor 10.0.0.33 remote-as 100
R1	Set R3 as a route reflector client.	vyatta@R1# set protocols bgp 100 neighbor 10.0.0.33 route-reflector-client
R1	Define the IP address on the local R1 router that is used to peer with the R3 router.	vyatta@R1# set protocols bgp 100 neighbor 10.0.0.33 update-source 10.0.0.11
R1	To allow routes from AS200 to be injected into RIB on R4.	vyatta@R1# set protocols bgp 100 neighbor 10.0.0.44 nexthop-self
R1	Set R4 in the same AS as R1.	vyatta@R1# set protocols bgp 100 neighbor 10.0.0.44 remote-as 100

Example 1-31 Creating route reflectors

R1	Set R4 as a route reflector client.	<pre>vyatta@R1# set protocols bgp 100 neighbor 10.0.0.44 route-reflector-client</pre>
R1	Define the IP address on the local R1 router that is used to peer with the R4 router.	<pre>vyatta@R1# set protocols bgp 100 neighbor 10.0.0.44 update-source 10.0.0.11</pre>
R1	Set the AS200 router in a different AS than R1.	<pre>vyatta@R1# set protocols bgp 100 neighbor 88.88.88.2 remote-as 200</pre>
R1	Set the outbound filter.	<pre>vyatta@R1# set protocols bgp 100 neighbor 88.88.88.2 route-map export eBGP-EXPORT</pre>
R1	Set the inbound filter.	<pre>vyatta@R1# set protocols bgp 100 neighbor 88.88.88.2 route-map import eBGP-IMPORT</pre>
R1	Set the network to advertise.	<pre>vyatta@R1# set protocols bgp 100 network 172.16.0.0/24</pre>
R1	Set the router ID.	<pre>vyatta@R1# set protocols bgp 100 parameters router-id 10.0.0.11</pre>
R1	Commit the configuration.	<pre>vyatta@R1# commit</pre>

Example 1-31 Creating route reflectors

R1	Display the BGP configuration.	<pre>vyatta@R1# show protocols bgp 100 { neighbor 10.0.0.22 { nexthop-self remote-as 100 route-reflector-client update-source 10.0.0.11 } neighbor 10.0.0.33 { nexthop-self remote-as 100 route-reflector-client update-source 10.0.0.11 } neighbor 10.0.0.44 { nexthop-self remote-as 100 route-reflector-client update-source 10.0.0.11 } neighbor 88.88.88.2 { remote-as 200 route-map { export eBGP-EXPORT import eBGP-IMPORT } } network 172.16.0.0/24 { } parameters { router-id 10.0.0.11 } } vyatta@R1#</pre>
<hr/>		
R2	Delete current BGP configuration.	<pre>vyatta@R2# delete protocols bgp</pre>
<hr/>		
R2	Set R1 in the same AS as R2.	<pre>vyatta@R2# set protocols bgp 100 neighbor 10.0.0.11 remote-as 100</pre>

Example 1-31 Creating route reflectors

R2	Define the IP address on the local R2 router that is used to peer with the R1 router.	vyatta@R2# set protocols bgp 100 neighbor 10.0.0.11 update-source 10.0.0.22
R2	Set the network to advertise.	vyatta@R2# set protocols bgp 100 network 172.16.0.0/24
R2	Set the router ID.	vyatta@R2# set protocols bgp 100 parameters router-id 10.0.0.22
R2	Commit the configuration.	vyatta@R2# commit
R2	Display the BGP configuration.	vyatta@R2# show protocols bgp 100 { neighbor 10.0.0.11 { remote-as 100 update-source 10.0.0.22 } network 172.16.0.0/24 { } parameters { router-id 10.0.0.22 } } vyatta@R2#
R3	Delete current BGP configuration.	vyatta@R3# delete protocols bgp
R3	Set R1 in the same AS as R3.	vyatta@R3# set protocols bgp 100 neighbor 10.0.0.11 remote-as 100
R3	Define the IP address on the local R3 router that is used to peer with the R1 router.	vyatta@R3# set protocols bgp 100 neighbor 10.0.0.11 update-source 10.0.0.33
R3	Set the network to advertise.	vyatta@R3# set protocols bgp 100 network 172.16.0.0/24
R3	Set the router ID.	vyatta@R3# set protocols bgp 100 parameters router-id 10.0.0.33
R3	Commit the configuration.	vyatta@R3# commit

Example 1-31 Creating route reflectors

R3	Display the BGP configuration.	<pre>vyatta@R3# show protocols bgp 100 { neighbor 10.0.0.11 { remote-as 100 update-source 10.0.0.33 } network 172.16.0.0/24 { } parameters { router-id 10.0.0.33 } } vyatta@R3#</pre>
R4	Delete current BGP configuration.	<pre>vyatta@R4# delete protocols bgp</pre>
R4	To allow routes from AS300 to be injected into RIB on R1.	<pre>vyatta@R4# set protocols bgp 100 neighbor 10.0.0.11 next-hop-self</pre>
R4	Set R1 in the same AS as R4.	<pre>vyatta@R4# set protocols bgp 100 neighbor 10.0.0.11 remote-as 100</pre>
R4	Define the IP address on the local R4 router that is used to peer with the R1 router.	<pre>vyatta@R4# set protocols bgp 100 neighbor 10.0.0.11 update-source 10.0.0.44</pre>
R4	Set the AS300 router in a different AS than R4.	<pre>vyatta@R4# set protocols bgp 100 neighbor 99.99.99.2 remote-as 300</pre>
R4	Set the outbound filter.	<pre>vyatta@R4# set protocols bgp 100 neighbor 99.99.99.2 route-map export eBGP-EXPORT</pre>
R4	Set the inbound filter.	<pre>vyatta@R4# set protocols bgp 100 neighbor 99.99.99.2 route-map import eBGP-IMPORT</pre>
R4	Set the network to advertise.	<pre>vyatta@R4# set protocols bgp 100 network 172.16.0.0/24</pre>
R4	Set the router ID.	<pre>vyatta@R4# set protocols bgp 100 parameters router-id 10.0.0.44</pre>

Example 1-31 Creating route reflectors

R4	Commit the configuration.	vyatta@R4# commit
R4	Display the BGP configuration.	<pre> vyatta@R4# show protocols bgp 100 { neighbor 10.0.0.11 { nexthop-self remote-as 100 update-source 10.0.0.44 } neighbor 99.99.99.2 { remote-as 300 route-map { export eBGP-EXPORT import eBGP-IMPORT } } network 172.16.0.0/24 { } parameters { router-id 10.0.0.44 } } vyatta@R4# </pre>

Verifying the Route Reflector

The following commands can be used to verify the route reflector configuration.

R1: show ip bgp summary

[Example 1-32](#) shows the output of the `show ip bgp summary` command for router R1 at this stage of the configuration.

Example 1-32 Verifying route reflector on R1: “show ip bgp summary”

```

vyatta@R1:~$ show ip bgp summary
BGP router identifier 10.0.0.11, local AS number 100
RIB entries 13, using 832 bytes of memory
Peers 4, using 10080 bytes of memory

Neighbor      V    AS MsgRcvd MsgSent   TblVer  InQ  OutQ Up/Down  State/PfxRcd
10.0.0.22     4   100     46     66       0    0    0 00:26:12    1
10.0.0.33     4   100     48     66       0    0    0 00:25:50    1

```

```

10.0.0.44      4   100    40    54    0    0    0 00:25:05    6
88.88.88.2    4   200    33    36    0    0    0 00:27:14    1

```

```

Total number of neighbors 4
vyatta@R1:~$

```

R1: show ip bgp

[Example 1-33](#) shows the output of the `show ip bgp` command for router R1 at this stage of the configuration.

Example 1-33 Verifying route reflector on R1: “show ip bgp”

```

vyatta@R1:~$ show ip bgp
BGP table version is 0, local router ID is 10.0.0.11
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

```

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i3.0.0.0/24	10.0.0.44	0	100	0	300 i
*>i3.1.0.0/24	10.0.0.44	0	100	0	300 i
*>i3.2.0.0/24	10.0.0.44	0	100	0	300 i
*> 12.0.0.0	88.88.88.2	0		0	200 i
*>i13.0.0.0/24	10.0.0.44	0	100	0	300 i
*>i99.99.99.0/30	10.0.0.44	0	100	0	300 i
* i172.16.0.0/24	10.0.0.44	1	100	0	i
* i	10.0.0.33	1	100	0	i
* i	10.0.0.22	1	100	0	i
*>	0.0.0.0	1		32768	i

```

Total number of prefixes 7
vyatta@R1:~$

```

R2: show ip bgp summary

[Example 1-34](#) shows the output of the `show ip bgp summary` command for router R2 at this stage of the configuration.

Example 1-34 Verifying route reflector on R2: “show ip bgp summary”

```

vyatta@R2:~$ show ip bgp summary

```

```
BGP router identifier 10.0.0.22, local AS number 100
RIB entries 13, using 832 bytes of memory
Peers 1, using 2520 bytes of memory
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.0.0.11	4	100	61	55	0	0	0	00:29:21	7

```
Total number of neighbors 1
vyatta@R2:~$
```

R2: show ip bgp

[Example 1-35](#) shows the output of the `show ip bgp` command for router R2 at this stage of the configuration.

Example 1-35 Verifying route reflector on R2: “show ip bgp”

```
vyatta@R2:~$ show ip bgp
BGP table version is 0, local router ID is 10.0.0.22
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i3.0.0.0/24	10.0.0.44	0	100	0	300 i
*>i3.1.0.0/24	10.0.0.44	0	100	0	300 i
*>i3.2.0.0/24	10.0.0.44	0	100	0	300 i
*>i12.0.0.0	10.0.0.11	0	100	0	200 i
*>i13.0.0.0/24	10.0.0.44	0	100	0	300 i
*>i99.99.99.0/30	10.0.0.44	0	100	0	300 i
* i172.16.0.0/24	10.0.0.11	1	100	0	i
*>	0.0.0.0	1		32768	i

```
Total number of prefixes 7
vyatta@R2:~$
```

R3: show ip bgp summary

[Example 1-36](#) shows the output of the `show ip bgp summary` command for router R3 at this stage of the configuration.

Example 1-36 Verifying route reflector on R3: “show ip bgp summary”

```
vyatta@R3:~$ show ip bgp summary
BGP router identifier 10.0.0.33, local AS number 100
RIB entries 13, using 832 bytes of memory
Peers 1, using 2520 bytes of memory

Neighbor      V    AS MsgRcvd MsgSent   TblVer  InQ  OutQ Up/Down  State/PfxRcd
10.0.0.11     4   100     60     54       0    0    0 00:32:15      7

Total number of neighbors 1
vyatta@R3:~$
```

R3: show ip bgp

[Example 1-37](#) shows the output of the `show ip bgp` command for router R3 at this stage of the configuration.

Example 1-37 Verifying route reflector on R3: “show ip bgp”

```
vyatta@R3:~$ show ip bgp
BGP table version is 0, local router ID is 10.0.0.33
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop           Metric LocPrf Weight Path
*>i3.0.0.0/24      10.0.0.44           0     100     0 300 i
*>i3.1.0.0/24      10.0.0.44           0     100     0 300 i
*>i3.2.0.0/24      10.0.0.44           0     100     0 300 i
*>i12.0.0.0        10.0.0.11           0     100     0 200 i
*>i13.0.0.0/24    10.0.0.44           0     100     0 300 i
*>i99.99.99.0/30  10.0.0.44           0     100     0 300 i
* i172.16.0.0/24  10.0.0.11           1     100     0 i
*>                 0.0.0.0             1           32768 i

Total number of prefixes 7
vyatta@R3:~$
```

R4: show ip bgp summary

[Example 1-38](#) shows the output of the `show ip bgp summary` command for router R4 at this stage of the configuration.

Example 1-38 Verifying route reflector on R4: “show ip bgp summary”

```
vyatta@R4:~$ show ip bgp summary
BGP router identifier 10.0.0.44, local AS number 100
RIB entries 13, using 832 bytes of memory
Peers 2, using 5040 bytes of memory

Neighbor      V    AS MsgRcvd MsgSent  TblVer  InQ OutQ Up/Down  State/PfxRcd
10.0.0.11     4   100    51     54      0    0  0 00:34:25    2
99.99.99.2    4   300    44     48      0    0  0 00:34:19    5

Total number of neighbors 2
vyatta@R4:~$
```

R4: show ip bgp

[Example 1-39](#) shows the output of the `show ip bgp` command for router R4 at this stage of the configuration.

Example 1-39 Verifying route reflector on R4: “show ip bgp”

```
vyatta@R4:~$ show ip bgp
BGP table version is 0, local router ID is 10.0.0.44
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
               r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 3.0.0.0/24       99.99.99.2        0             0 300 i
*> 3.1.0.0/24       99.99.99.2        0             0 300 i
*> 3.2.0.0/24       99.99.99.2        0             0 300 i
*>i12.0.0.0         10.0.0.11         0      100     0 200 i
*> 13.0.0.0/24     99.99.99.2        0             0 300 i
*> 99.99.99.0/30   99.99.99.2        0             0 300 i
* i172.16.0.0/24   10.0.0.11         1      100     0 i
*>                  0.0.0.0           1             32768 i

Total number of prefixes 7
vyatta@R4:~$
```

Route Redirection

Route redirection in BGP is performed by means of routing policies. For more information about routing policies, please see the *Vyatta Routing Policies Reference Guide*.

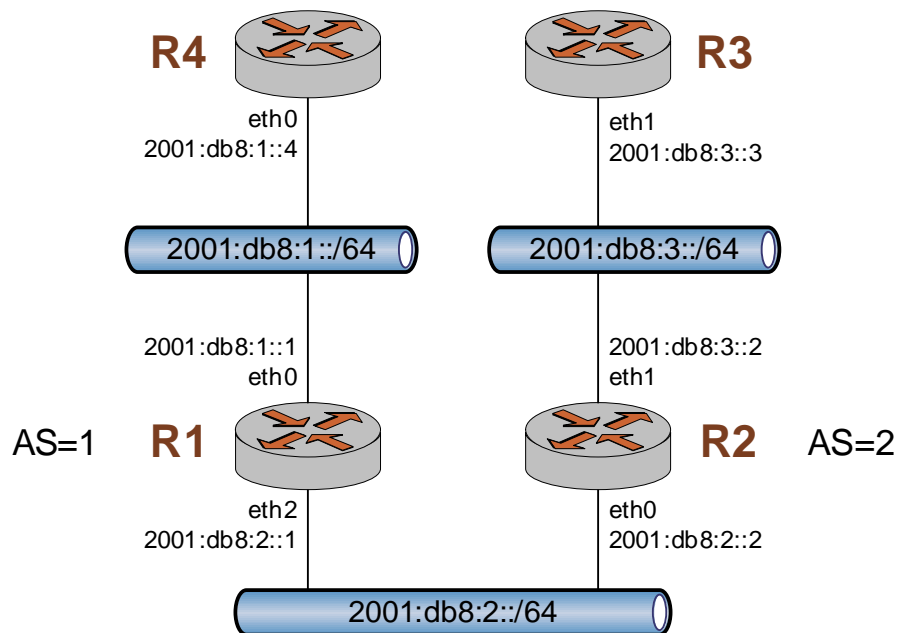
Configuring IPv6 Routing with BGP

This section presents the following topics:

- [Enable Forwarding on R1 and R2](#)
- [Configure IPv6 BGP Peer](#)
- [Advertise Connected Networks](#)
- [Confirm Advertised Routes](#)

This section presents sample configurations for dynamic IPv6 routing with BGP. The configuration examples are based on the reference diagram in [Figure 1-13](#). In this example we will show configuration of the nodes using dynamic IPv6 routing with BGP to enable R3 and R4 to communicate via R1 and R2.

Figure 1-13 Dynamic IPv6 routing example - BGP



Enable Forwarding on R1 and R2

For R1 to be able to pass data between interfaces eth0 and eth2, and for R2 to be able to pass data between interfaces eth0 and eth1, they must be configured to enable forwarding. To enable forwarding on R1, perform the following steps in configuration mode.

Example 1-40 Enable forwarding on R1

Step	Command
Enable forwarding on R1.	<code>vyatta@R1# delete system ipv6 disable-forwarding</code>
Commit the change.	<code>vyatta@R1# commit</code>

To enable forwarding on R2, perform the following steps in configuration mode.

Example 1-41 Enable forwarding on R2

Step	Command
Enable forwarding on R2.	<code>vyatta@R2# delete system ipv6 disable-forwarding</code>
Commit the change.	<code>vyatta@R2# commit</code>

Configure IPv6 BGP Peer

R1 and R2 must be configured to determine how to access remote Autonomous Systems (AS). To configure R1 to see AS=2, perform the following steps in configuration mode.

Example 1-42 Configure R1 to see AS=2

Step	Command
Configure a BGP peer on R1.	<code>vyatta@R1# set protocols bgp 1 neighbor 2001:db8:2::2 remote-as 2</code>
Commit the change.	<code>vyatta@R1# commit</code>

Similarly, to configure R2 to see AS=1, perform the following steps in configuration mode.

Example 1-43 Configure R2 to see AS=1

Step	Command
Configure a BGP peer on R2.	<code>vyatta@R2# set protocols bgp 2 neighbor 2001:db8:2::1 remote-as 1</code>

Example 1-43 Configure R2 to see AS=1

```
Commit the change.          vyatta@R2# commit
```

To confirm that the peer session is "Established", perform the following step in operational mode.

Example 1-44 Confirm peer session is established

Step	Command
Display the status of the BGP neighbor.	<pre>vyatta@R1:~\$ show ip bgp neighbors 2001:db8:2::2 BGP neighbor is 2001:db8:2::2, remote AS 2, local AS 1, external link BGP version 4, remote router ID 172.16.139.160 BGP state = Established, up for 00:01:24 Last read 00:00:24, hold time is 180, keepalive interval is 60 seconds Neighbor capabilities: 4 Byte AS: advertised and received Route refresh: advertised and received(old & new) Address family IPv6 Unicast: advertised and received Message statistics: Inq depth is 0 Outq depth is 0 Sent Rcvd Opens: 2 0 Notifications: 0 0 Updates: 0 0 Keepalives: 3 2 Route Refresh: 0 0 Capability: 0 0 Total: 5 2 Minimum time between advertisement runs is 30 seconds For address family: IPv6 Unicast Community attribute sent to this neighbor(both) 0 accepted prefixes Connections established 1; dropped 0 Last reset never Local host: 2001:db8:2::1, Local port: 179 Foreign host: 2001:db8:2::2, Foreign port: 55711 Nexthop: 172.16.117.128 Nexthop global: 2001:db8:2::1 Nexthop local: fe80::20c:29ff:fed6:8180 BGP connection: shared network Read thread: on Write thread: off</pre>

Advertise Connected Networks

The **redistribute** command is then used to advertise the connected networks. To advertise connected networks on R1, perform the following steps in configuration mode.

Example 1-45 Advertise connected networks on R1

Step	Command
Advertise connected networks via bgp.	vyatta@R1# set protocols bgp 1 address-family ipv6-unicast redistribute connected
Commit the change.	vyatta@R1# commit

Confirm Advertised Routes

To see which routes are being advertised by R1 and which routes have been learned from peers, perform the following steps in operational mode.

Example 1-46 Confirm routes advertised and learned by R1

Step	Command
Display routes advertised by R1.	<pre>vyatta@R1:~\$ show ipv6 bgp neighbors 2001:db8:2::2 advertised-routes BGP table version is 0, local router ID is 172.16.117.128 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, r RIB-failure, S Stale, R Removed Origin codes: i - IGP, e - EGP, ? - incomplete Network Next Hop Metric LocPrf Weight Path *> 2001:db8:1::/64 2001:db8:2::1 1 32768 ? *> 2001:db8:2::/64 2001:db8:2::1 1 32768 ? Total number of prefixes 2</pre>

Example 1-46 Confirm routes advertised and learned by R1

```

Display routes learned by R1.      vyatta@R1:~$ show ipv6 bgp neighbors 2001:db8:2::2 routes
BGP table version is 0, local router ID is 172.16.117.128
Status codes: s suppressed, d damped, h history, * valid, >
best, i - internal,
                r RIB-failure, S Stale, R Removed
Origin codes: i - IGP, e - EGP, ? - incomplete

      Network          Next Hop          Metric LocPrf Weight Path
* 2001:db8:2::/64    2001:db8:2::2          1         0 2 ?
*> 2001:db8:3::/64  2001:db8:2::2          1         0 2 ?

Total number of prefixes 2

```

```

Display the routing table.        vyatta@R1:~$ show ipv6 route
Codes: K - kernel route, C - connected, S - static, R - RIPng,
O - OSPFv3,
      I - ISIS, B - BGP, * - FIB route.

C>* ::1/128 is directly connected, lo
C>* 2001:db8:1::/64 is directly connected, eth0
C>* 2001:db8:2::/64 is directly connected, eth2
B>* 2001:db8:3::/64 [20/1] via fe80::20c:29ff:fe4e:fc6,
eth2, 00:22:47
C * fe80::/64 is directly connected, eth2
C>* fe80::/64 is directly connected, eth0

```

Chapter 2: Global and Router-Specific Configuration

This chapter describes commands for global and router-specific configuration for BGP.

This chapter presents the following topics:

- [Global and Router-Specific Commands](#)

Global and Router-Specific Commands

This chapter contains the following commands.

Configuration Commands	
Global BGP Configuration	
<code>protocols bgp <asn></code>	Creates a BGP instance on the router and locates it within an Autonomous System (AS).
<code>protocols bgp <asn> aggregate-address <ipv4net></code>	Specifies a block of IP addresses to aggregate.
<code>protocols bgp <asn> maximum-paths ebgp <max-paths></code>	Specifies the maximum number of eBGP paths to a destination.
<code>protocols bgp <asn> maximum-paths ibgp <max-paths></code>	Specifies the maximum number of iBGP paths to a destination.
<code>protocols bgp <asn> network <ipv4net></code>	Specifies a network to be advertised by the BGP routing process.
<code>protocols bgp <asn> timers</code>	Sets BGP timers globally for the local router.
Global BGP Configuration - IPv6	
<code>protocols bgp <asn> address-family ipv6-unicast</code>	Defines parameters for controlling IPv6 unicast routes.
<code>protocols bgp <asn> address-family ipv6-unicast aggregate-address <ipv6net></code>	Specifies a block of IPv6 addresses to aggregate.
<code>protocols bgp <asn> address-family ipv6-unicast network <ipv6net></code>	Specifies an IPv6 network to be advertised by the BGP routing process.
Router-Specific BGP Configuration	
<code>protocols bgp <asn> parameters always-compare-med</code>	Directs the router to compare the MED for paths from neighbors in different autonomous systems.
<code>protocols bgp <asn> parameters bestpath as-path</code>	Directs the router to compare the AS paths during best path selection.
<code>protocols bgp <asn> parameters bestpath compare-routerid</code>	Directs the router to compare identical routes received from different external peers during best path selection.
<code>protocols bgp <asn> parameters bestpath med</code>	Directs the router to compare the Multi Exit Discriminator (MED) among paths learned from confederation peers during best path selection.

<code>protocols bgp <asn> parameters dampening</code>	Enables or disables route dampening and sets route dampening values.
<code>protocols bgp <asn> parameters default</code>	Sets default routing behaviors for the system.
<code>protocols bgp <asn> parameters deterministic-med</code>	Enables or disables enforcing of deterministic MED.
<code>protocols bgp <asn> parameters distance global</code>	Sets administrative distance for all BGP routes.
<code>protocols bgp <asn> parameters distance prefix <ipv4net> distance <dist></code>	Sets administrative distance for BGP routes for a specific prefix.
<code>protocols bgp <asn> parameters disable-network-import-check</code>	Disables IGP route check for network statements.
<code>protocols bgp <asn> parameters enforce-first-as</code>	Enables or disables forcing eBGP peers to list AS number at the beginning of the AS_PATH attribute in incoming updates.
<code>protocols bgp <asn> parameters graceful-restart</code>	Enables or disables graceful restart of the BGP process.
<code>protocols bgp <asn> parameters log-neighbor-changes</code>	Enables or disables logging of the status of BGP neighbors.
<code>protocols bgp <asn> parameters no-fast-external-failover</code>	Enables or disables automatic resetting of BGP sessions for failed links.
<code>protocols bgp <asn> parameters router-id <id></code>	Sets a fixed BGP router ID for the router, overriding the automatic ID selection process.
<code>protocols bgp <asn> parameters scan-time <seconds></code>	Sets the scanning interval for the router.
Operational Commands	
<code>monitor protocol bgp disable</code>	Disables debug message generation related to the acquisition of the BGP router ID and the sending and receiving of BGP messages.
<code>monitor protocol bgp disable all</code>	Disables all BGP debugging.
<code>monitor protocol bgp disable events</code>	Disables debug message generation related to BGP events.
<code>monitor protocol bgp disable filters</code>	Disables debug message generation related to BGP filters.
<code>monitor protocol bgp disable fsm</code>	Disables debug message generation related to the BGP Finite State Machine (FSM).
<code>monitor protocol bgp disable keepalives</code>	Disable debugging information related to sending and receiving BGP keep-alive messages.
<code>monitor protocol bgp disable rib</code>	Disable debugging information related to the Routing Information Base (RIB).

<code>monitor protocol bgp disable updates</code>	Disable debugging information related to BGP routing updates.
<code>monitor protocol bgp enable</code>	Enables debug message generation related to the acquisition of the BGP router ID and the sending and receiving of BGP messages.
<code>monitor protocol bgp enable events</code>	Enables debug message generation related to BGP events.
<code>monitor protocol bgp enable filters</code>	Enables debug message generation related to BGP filters.
<code>monitor protocol bgp enable fsm</code>	Enables debug message generation related to the BGP Finite State Machine (FSM).
<code>monitor protocol bgp enable keepalives</code>	Display debugging information related to sending and receiving BGP keep-alive messages.
<code>monitor protocol bgp enable rib</code>	Display debugging information related to the Routing Information Base (RIB).
<code>monitor protocol bgp enable updates</code>	Display debugging information related to BGP routing updates.
<code>reset ip bgp <address></code>	Resets a BGP connection.
<code>reset ip bgp <address> ipv4 unicast</code>	Resets an IPv4 unicast BGP connection.
<code>reset ip bgp all</code>	Resets all BGP peering sessions.
<code>reset ip bgp dampening</code>	Resets BGP route dampening information and unsuppresses suppressed routes.
<code>show ip bgp</code>	Displays BGP routes.
<code>show ip bgp attribute-info</code>	Displays BGP attribute information.
<code>show ip bgp cidr-only</code>	Displays BGP routes with CIDR network masks.
<code>show ip bgp community-info</code>	Displays BGP community information.
<code>show ip bgp community <community></code>	Displays BGP routes belonging to the specified BGP community.
<code>show ip bgp community-list <list-name></code>	Displays BGP routes permitted by the specified community list.
<code>show ip bgp dampened-paths</code>	Displays BGP routes that are currently dampened.
<code>show ip bgp filter-list <as-path-list-name></code>	Displays routes matching a list of autonomous system paths.
<code>show ip bgp flap-statistics</code>	Displays route flap statistics for BGP routes.

<code>show ip bgp flap-statistics cidr-only</code>	Displays only route flap statistics for BGP routes with CIDR network masks.
<code>show ip bgp flap-statistics filter-list <as-path-list-name></code>	Displays route flap statistics for BGP routes matching a list of AS paths.
<code>show ip bgp flap-statistics prefix-list <list-name></code>	Displays route flap statistics for BGP routes matching a prefix list.
<code>show ip bgp flap-statistics regexp <expr></code>	Displays route flap statistics for BGP routes matching an AS path regular expression.
<code>show ip bgp flap-statistics route-map <map-name></code>	Displays route flap statistics for BGP routes matching a route map.
<code>show ip bgp ipv4 unicast</code>	Displays information for IPv4 unicast BGP routes.
<code>show ip bgp ipv4 unicast cidr-only</code>	Displays IPv4 unicast BGP routes with CIDR network masks.
<code>show ip bgp ipv4 unicast community <community></code>	Displays IPv4 unicast BGP routes belonging to the specified community.
<code>show ip bgp ipv4 unicast community-list <list-name></code>	Displays IPv4 unicast BGP routes permitted by the specified community list.
<code>show ip bgp ipv4 unicast filter-list <as-path-list-name></code>	Displays IPv4 unicast BGP routes permitted by the specified AS path list.
<code>show ip bgp ipv4 unicast paths</code>	Displays BGP IPv4 unicast path information.
<code>show ip bgp ipv4 unicast prefix-list <list-name></code>	Displays IPv4 unicast BGP routes matching a prefix list.
<code>show ip bgp ipv4 unicast regexp <regexp></code>	Displays IPv4 unicast BGP routes matching an AS path regular expression.
<code>show ip bgp ipv4 unicast route-map <map-name></code>	Displays IPv4 unicast BGP routes matching a route map.
<code>show ip bgp ipv4 unicast statistics</code>	Displays statistics for BGP IPv4 unicast routes.
<code>show ip bgp ipv4 unicast summary</code>	Displays a summary of BGP IPv4 unicast information.
<code>show ip bgp memory</code>	Displays memory usage for BGP.
<code>show ip bgp paths</code>	Displays all BGP paths.
<code>show ip bgp prefix-list <list-name></code>	Displays BGP routes matching a prefix list.
<code>show ip bgp regexp <regexp></code>	Displays routes matching an AS path regular expression.
<code>show ip bgp route-map <map-name></code>	Displays routes matching a route map.

<code>show ip bgp scan</code>	Displays BGP scan status.
<code>show ip route bgp</code>	Displays BGP routes.
<code>show ipv6 bgp</code>	Displays BGP routes.
<code>show ipv6 bgp community <community></code>	Displays BGP routes belonging to the specified BGP community.
<code>show ipv6 bgp community-list <list-name></code>	Displays BGP routes permitted by the specified community list.
<code>show ipv6 bgp filter-list <as-path-list-name></code>	Displays routes matching a list of autonomous system paths.
<code>show ipv6 bgp prefix-list <list-name></code>	Displays BGP routes matching a prefix list.
<code>show ipv6 bgp regexp <regexp></code>	Displays routes matching an AS path regular expression.
<code>show monitoring protocols bgp</code>	Displays BGP protocol debugging flags.

monitor protocol bgp disable

Disables debug message generation related to the acquisition of the BGP router ID and the sending and receiving of BGP messages.

Syntax

```
monitor protocol bgp disable
```

Command Mode

Operational mode.

Parameters

None.

Default

None

Usage Guidelines

Use this command to disable debugging related to the acquisition of the BGP router ID and the sending and receiving of BGP messages.

monitor protocol bgp disable all

Disables all BGP debugging.

Syntax

```
monitor protocol bgp disable all
```

Command Mode

Operational mode.

Parameters

None.

Default

None.

Usage Guidelines

Use this command to disable all BGP debug message generation.

monitor protocol bgp disable events

Disables debug message generation related to BGP events.

Syntax

```
monitor protocol bgp disable events
```

Command Mode

Operational mode.

Parameters

None.

Default

None

Usage Guidelines

Use this command to disable generation of trace-level messages related to BGP events.

monitor protocol bgp disable filters

Disables debug message generation related to BGP filters.

Syntax

```
monitor protocol bgp disable filters
```

Command Mode

Operational mode.

Parameters

None.

Default

None

Usage Guidelines

Use this command to disable generation of trace-level messages related to BGP filters.

monitor protocol bgp disable fsm

Disables debug message generation related to the BGP Finite State Machine (FSM).

Syntax

```
monitor protocol bgp disable fsm
```

Command Mode

Operational mode.

Parameters

None.

Default

None

Usage Guidelines

Use this command to disable generation of trace-level messages related to the BGP finite state machine (FSM).

A BGP router uses an FSM consisting of six states, as defined by RFC 1771. The FSM describes how and when the BGP router should make decisions about its operations with other BGP neighbors.

monitor protocol bgp disable keepalives

Disable debugging information related to sending and receiving BGP keep-alive messages.

Syntax

```
monitor protocol bgp disable keepalives
```

Command Mode

Operational mode.

Parameters

None.

Default

None

Usage Guidelines

Use this command to disable generation of trace-level messages related to sending and receiving BGP keep-alive messages.

monitor protocol bgp disable rib

Disable debugging information related to the Routing Information Base (RIB).

Syntax

```
monitor protocol bgp disable rib
```

Command Mode

Operational mode.

Parameters

None.

Default

None.

Usage Guidelines

Use this command to disable generation of trace-level messages related to the Routing Information Base.

monitor protocol bgp disable updates

Disable debugging information related to BGP routing updates.

Syntax

```
monitor protocol bgp disable updates
```

Command Mode

Operational mode.

Parameters

None.

Default

None.

Usage Guidelines

Use this command to disable generation of trace-level messages related to BGP routing updates.

monitor protocol bgp enable

Enables debug message generation related to the acquisition of the BGP router ID and the sending and receiving of BGP messages.

Syntax

```
monitor protocol bgp enable
```

Command Mode

Operational mode.

Parameters

None.

Default

None

Usage Guidelines

Use this command to enable generation of trace-level messages related to the acquisition of the BGP router ID and the sending and receiving of BGP messages.

monitor protocol bgp enable events

Enables debug message generation related to BGP events.

Syntax

```
monitor protocol bgp enable events
```

Command Mode

Operational mode.

Parameters

None.

Default

None

Usage Guidelines

Use this command to enable generation of trace-level messages related to BGP events.

monitor protocol bgp enable filters

Enables debug message generation related to BGP filters.

Syntax

```
monitor protocol bgp enable filters
```

Command Mode

Operational mode.

Parameters

None.

Default

None

Usage Guidelines

Use this command to enable generation of trace-level messages related to BGP filters.

monitor protocol bgp enable fsm

Enables debug message generation related to the BGP Finite State Machine (FSM).

Syntax

```
monitor protocol bgp enable fsm
```

Command Mode

Operational mode.

Parameters

None.

Default

None

Usage Guidelines

Use this command to enable generation of trace-level messages related to the BGP finite state machine (FSM).

A BGP router uses an FSM consisting of six states, as defined by RFC 1771. The FSM describes how and when the BGP router should make decisions about its operations with other BGP neighbors.

monitor protocol bgp enable keepalives

Display debugging information related to sending and receiving BGP keep-alive messages.

Syntax

```
monitor protocol bgp enable keepalives
```

Command Mode

Operational mode.

Parameters

None.

Default

None

Usage Guidelines

Use this command to enable generation of trace-level messages related to sending and receiving BGP keep-alive messages.

monitor protocol bgp enable rib

Display debugging information related to the Routing Information Base (RIB).

Syntax

```
monitor protocol bgp enable rib
```

Command Mode

Operational mode.

Parameters

None.

Default

None.

Usage Guidelines

Use this command to enable generation of trace-level messages related to the Routing Information Base.

monitor protocol bgp enable updates

Display debugging information related to BGP routing updates.

Syntax

```
monitor protocol bgp enable updates [in | out]
```

Command Mode

Operational mode.

Parameters

in	Optional. Debug information is generated only for inbound routing updates.
out	Optional. Debug information is generated only for outbound routing updates.

Default

Debug messages are generated for both inbound and outbound routing messages.

Usage Guidelines

Use this command to enable generation of trace-level messages related to BGP routing updates.

protocols bgp <asn>

Creates a BGP instance on the router and locates it within an Autonomous System (AS).

Syntax

```
set protocols bgp asn
delete protocols bgp asn
show protocols bgp [asn]
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
    bgp asn {
    }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. Any peers of this router must be configured to know this AS number—if there is a mismatch, a peering will not be established. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
------------	--

Default

None.

Usage Guidelines

Use the set form of this command to enable a BGP instance on the router, and set its Autonomous System. All subsequent BGP configuration falls under this configuration node.

Note that you cannot create another BGP instance (that is, you cannot specify a second or further AS) unless you declare multiple BGP instances using the **multiple-instance** command.

Use the **delete** form of this command to disable BGP on the router, removing all BGP configuration.

Use the **show** form of this command to view all BGP configuration.

protocols bgp <asn> address-family ipv6-unicast

Defines parameters for controlling IPv6 unicast routes.

Syntax

```
set protocols bgp asn address-family ipv6-unicast
delete protocols bgp asn address-family ipv6-unicast
show protocols bgp asn address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    address-family {
      ipv6-unicast {}
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides.
------------	---

Usage Guidelines

- Use this command to define parameters for IPv6 unicast routes.
- Use the **set** form of this command to specify the configuration node.
- Use the **delete** form of this command to delete the configuration.
- Use the **show** form of this command to view the configuration.

protocols bgp <asn> address-family ipv6-unicast aggregate-address <ipv6net>

Specifies a block of IPv6 addresses to aggregate.

Syntax

```
set protocols bgp asn address-family ipv6-unicast aggregate-address ipv6net  
[summary-only]
```

```
delete protocols bgp asn address-family ipv6-unicast aggregate-address ipv6net
```

```
show protocols bgp asn address-family ipv6-unicast aggregate-address [ipv6net]
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        address-family {  
            ipv6-unicast {  
                aggregate-address ipv6net {  
                    summary-only  
                }  
            }  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides.
<i>ipv6net</i>	Mandatory. The IPv6 network from which routes are to be aggregated. The format is ipv6-address/prefix.
summary-only	Specifies that aggregated routes are summarized. These routes will not be announced.

Usage Guidelines

Use this command to specify a block of IPv6 addresses to aggregate. The command applies to IPv6 unicast routes only.

Use the **set** form of this command to specify a contiguous block of IPv6 addresses to aggregate.

Use the **delete** form of this command to delete an aggregate address.

Use the **show** form of this command to view aggregate address configuration settings.

protocols bgp <asn> address-family ipv6-unicast network <ipv6net>

Specifies an IPv6 network to be advertised by the BGP routing process.

Syntax

```
set protocols bgp asn address-family ipv6-unicast network ipv6net [path-limit limit | route-map map-name]
```

```
delete protocols bgp asn address-family ipv6-unicast network ipv6net [path-limit | route-map]
```

```
show protocols bgp asn address-family ipv6-unicast network ipv6net
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    address-family {
      ipv6-unicast {
        network ipv6net {
          path-limit limit
          route-map map-name
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>ipv6net</i>	Mandatory. Multi-node. An IPv6 network in the format <i>ipv6-address/prefix</i> . You can advertise to multiple networks by creating multiple network configuration nodes.

<i>limit</i>	AS path hop count limit. The range of values is 0 to 255.
<i>map-name</i>	The configured route map to be used when advertising the network.

Default

None.

Usage Guidelines

Use this command to advertise networks to BGP neighbors. The command applies to IPv6 unicast routes only. Note that **path-limit** and **route-map** options are mutually exclusive.

Use the **set** form of this command to specify a network to be announced via BGP.

Use the **delete** form of this command to remove a network from the list of networks to be announced by BGP.

Use the **show** form of this command to view BGP network advertising configuration settings.

protocols bgp <asn> aggregate-address <ipv4net>

Specifies a block of IP addresses to aggregate.

Syntax

```
set protocols bgp asn aggregate-address ipv4net [as-set | summary-only]
delete protocols bgp asn aggregate-address ipv4net
show protocols bgp asn aggregate-address [ipv4net]
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    aggregate-address ipv4net {
      as-set
      summary-only
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides.
<i>ipv4net</i>	Mandatory. The network from which routes are to be aggregated. The format is ip-address/prefix.
as-set	Specifies that the routes resulting from aggregation include the AS set.
summary-only	Specifies that aggregated routes are summarized. These routes will not be announced.

Usage Guidelines

Use the set form of this command to specify a contiguous block of IP addresses to aggregate.

Use the **delete** form of this command to delete an aggregate address.

Use the **show** form of this command to view aggregate address configuration settings.

protocols bgp <asn> maximum-paths ebgp <max-paths>

Specifies the maximum number of eBGP paths to a destination.

Syntax

```
set protocols bgp asn maximum-paths ebgp max-paths
delete protocols bgp asn maximum-paths ebgp
show protocols bgp asn maximum-paths ebgp
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    maximum-paths {
      ebgp max-paths
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>max-paths</i>	Mandatory. The maximum number of eBGP paths to a destination that will be inserted in the routing table. The range of values is 1 to 255, where 1 indicates that a single path can be inserted into the routing table. This has the same effect as not configuring the maximum-paths feature.

Default

A single eBGP path can be inserted into the routing table.

Usage Guidelines

Use this command to specify the maximum number of eBGP paths to a destination to allow in the routing table. Normally, the single best path is determined and then added to the routing table. This command allows more than one path to be added to the routing table.

To be considered as one of a set of multipaths to a destination, candidate paths must have identical weight, LOCAL PREFERENCE, AS PATH LENGTH, ORIGIN, MULTI_EXIT_DISC (MED), and PEER AS.

Once the paths are selected, traffic load is shared across them on a per session basis. That is, each new session is routed across the path that has the fewest active sessions. To achieve this, the route is advertised with the nexthop address being the interface address rather than the address of the best path. Once traffic is routed to the interface, it can select the path based on current traffic loads.

The attribute set for the advertised route is based on the attribute set from the best path, with the following attributes are aggregated from the multipaths, including the best path: AS_PATH, ORIGIN, COMMUNITIES, and EXTENDED COMMUNITIES.

If the deterministic MED option is configured (see [“protocols bgp <asn> parameters deterministic-med” on page 123](#)), a set of multipaths is derived for each peer AS. The set of paths for the peer AS associated with the selected best path are added to the routing table.

If the always-compare MED option is configured (see [“protocols bgp <asn> parameters always-compare-med” on page 111](#)), the MED is compared across all paths from all ASs. The set of multipaths can be distributed across ASs.

NOTE *Since each path requires additional memory and processing resources, the number of **maximum-paths** should be kept to a minimum. This is especially true for large routing tables.*

NOTE *The multipath feature is only supported in IPv4 unicast environments.*

Use the **set** form of this command to specify the maximum number of paths to a destination.

Use the **delete** form of this command to remove the configuration and allow only a single best path to a destination in the routing table.

Use the **show** form of this command to view eBGP maximum path configuration settings.

protocols bgp <asn> maximum-paths ibgp <max-paths>

Specifies the maximum number of iBGP paths to a destination.

Syntax

```
set protocols bgp asn maximum-paths ibgp max-paths
delete protocols bgp asn maximum-paths ibgp
show protocols bgp asn maximum-paths ibgp
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    maximum-paths {
      ibgp max-paths
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>max-paths</i>	Mandatory. The maximum number of iBGP paths to a destination that will be inserted in the routing table. The range of values is 1 to 255, where 1 indicates that a single path can be inserted into the routing table. This has the same effect as not configuring the maximum-paths feature.

Default

A single iBGP path can be inserted into the routing table.

Usage Guidelines

Use this command to specify the maximum number of iBGP paths to a destination to allow in the routing table. Normally, the single best path is determined and then added to the routing table. This command allows more than one path to be added to the routing table.

To be considered as one of a set of multipaths to a destination, candidate paths must have identical weight, LOCAL PREFERENCE, AS PATH LENGTH, ORIGIN, MULTI_EXIT_DISC (MED), and AS_PATH.

Once the paths are selected, traffic load is shared across them on a per session basis. That is, each new session is routed across the path that has the fewest active sessions. To achieve this, the route is advertised with the nexthop address being the interface address rather than the address of the best path. Once traffic is routed to the interface, it can select the path based on current traffic loads.

The attribute set for the advertised route is based on the attribute set from the best path, with the following attributes are aggregated from the multipaths, including the best path: AS_PATH, ORIGIN, COMMUNITIES, and EXTENDED COMMUNITIES.

If the deterministic MED option is configured (see [“protocols bgp <asn> parameters deterministic-med” on page 123](#)), a set of multipaths is derived for each peer AS. The set of paths for the peer AS associated with the selected best path are added to the routing table.

If the always-compare MED option is configured (see [“protocols bgp <asn> parameters always-compare-med” on page 111](#)), the MED is compared across all paths from all ASs. The set of multipaths can be distributed across ASs.

NOTE *Since each path requires additional memory and processing resources, the number of **maximum-paths** should be kept to a minimum. This is especially true for large routing tables.*

NOTE *The multipath feature is only supported in IPv4 unicast environments.*

Use the **set** form of this command to specify the maximum number of paths to a destination.

Use the **delete** form of this command to remove the configuration and allow only a single best path to a destination in the routing table.

Use the **show** form of this command to view iBGP maximum path configuration settings.

protocols bgp <asn> network <ipv4net>

Specifies a network to be advertised by the BGP routing process.

Syntax

```
set protocols bgp asn network ipv4net [backdoor | route-map map-name]  
delete protocols bgp asn network ipv4net [backdoor | route-map]  
show protocols bgp asn network
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        network ipv4net {  
            backdoor  
            route-map map-name  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>ipv4net</i>	Mandatory. Multi-node. An IPv4 network in the format <i>ip-address/prefix</i> . You can advertise to multiple networks by creating multiple network configuration nodes.
backdoor	Optional. Indicates that this network is reachable by a backdoor route. A backdoor network is considered to be like a local network, but is not advertised.
<i>map-name</i>	Optional. Specifies a configured route map to be used when advertising the network.

Default

None.

Usage Guidelines

Use this command to advertise networks to BGP neighbors.

Use the **set** form of this command to specify a network to be announced via BGP.

Use the **delete** form of this command to remove a network from the list of networks to be announced by BGP.

Use the **show** form of this command to view BGP network advertising configuration settings.

protocols bgp <asn> parameters always-compare-med

Directs the router to compare the MED for paths from neighbors in different autonomous systems.

Syntax

```
set protocols bgp asn parameters always-compare-med
delete protocols bgp asn parameters always-compare-med
show protocols bgp asn parameters
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    parameters {
      always-compare-med
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
------------	---

Default

The MED is not compared.

Usage Guidelines

Use the **set** form of this command to have the router compare the Multi Exit Discriminator (MED) for paths from neighbors in different autonomous systems. The MED is compared only if the AS path for the compared routes is identical. Use the **delete** form of this command to disable MED comparison.

Use the **show** form of this command to view global BGP configuration settings.

protocols bgp <asn> parameters bestpath as-path

Directs the router to compare the AS paths during best path selection.

Syntax

```
set protocols bgp asn parameters bestpath as-path [confed | ignore]
delete protocols bgp asn parameters bestpath as-path
show protocols bgp asn parameters bestpath
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    parameters {
      bestpath {
        as-path {
          confed
          ignore
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
confed	Optional. Directs the router to compare the AS paths within a confederation during best path selection.
ignore	Optional. Directs the router to ignore AS_PATH during best path selection.

Default

By default, when making the best-path selection the router does not compare AS_PATHs within a confederation and does not ignore the AS_PATH.

Usage Guidelines

Use the **set** form of this command to direct the router to compare the AS paths during best path selection.

Use the **delete** form of this command to restore the default best-path selection behavior.

Use the **show** form of this command to view BGP best path selection configuration settings.

protocols bgp <asn> parameters bestpath compare-routerid

Directs the router to compare identical routes received from different external peers during best path selection.

Syntax

```
set protocols bgp asn parameters bestpath compare-routerid
delete protocols bgp asn parameters bestpath compare-routerid
show protocols bgp asn parameters bestpath
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    parameters {
      bestpath {
        compare-routerid
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
------------	---

Default

By default, when making the best-path selection the router does not compare identical routes received from different external peers.

Usage Guidelines

Use the **set** form of this command to direct the router to compare identical routes received from different external peers during best path selection, and select as the best path the route with the lowest router ID.

Use the **delete** form of this command to restore the default best-path selection behavior.

Use the **show** form of this command to view BGP best path selection configuration settings.

protocols bgp <asn> parameters bestpath med

Directs the router to compare the Multi Exit Discriminator (MED) among paths learned from confederation peers during best path selection.

Syntax

```
set protocols bgp asn parameters bestpath med [confed | missing-as-worst]
delete protocols bgp asn parameters bestpath med [confed | missing-as-worst]
show protocols bgp asn parameters bestpath
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    parameters {
      bestpath {
        med {
          confed
          missing-as-worst
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>confed</i>	Optional. Compare the MED among confederation paths
<i>missing-as-worst</i>	Optional. Treat a missing MED as the least preferred one.

Default

By default, when making the best-path selection the router does not consider the MED.

Usage Guidelines

Use the **set** form of this command to direct the router to compare the Multi Exit Discriminator (MED) among paths learned from confederation peers during best-path selection.

Use the **delete** form of this command to restore the default best-path selection behavior.

Use the **show** form of this command to view BGP best path selection configuration settings.

protocols bgp <asn> parameters dampening

Enables or disables route dampening and sets route dampening values.

Syntax

```
set protocols bgp asn parameters dampening [half-life minutes | re-use penalty |  
start-suppress-time penalty | max-suppress-time minutes]
```

```
delete protocols bgp asn parameters dampening [half-life | re-use |  
start-suppress-time | max-suppress-time]
```

```
show protocols bgp asn parameters dampening
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        parameters {  
            dampening {  
                half-life minutes  
                max-suppress-time minutes  
                re-use penalty  
                start-suppress-time penalty  
            }  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>half-life minutes</i>	Optional. The time period, in minutes, after which the penalty assigned to a route because of flapping is reduced by half. The range is 1 to 45. The default is 15.

max-suppress-time <i>minutes</i>	Optional. The maximum time, in minutes, that a route may be suppressed. The range is 1 to 20000. The default is four times the half-life period.
re-use <i>penalty</i>	Optional. The reuse threshold. If the penalty for a flapping route is reduced below this point, the route is to be brought back into use. The range is 1 to 20000. The default is 750.
start-suppress-time <i>penalty</i>	Optional. The route suppression threshold. If the accumulated penalty for a flapping route reaches this limit, the route is suppressed. The range is 1 to 20000. The default is 2000.

Default

Route dampening is disabled.

Usage Guidelines

Use the **set** form of this command to configure route dampening. When used with no options, this command enables route dampening at the default values.

Use the **delete** form of this command to disable route dampening or reset route dampening parameters to default values.

Use the **show** form of this command to view BGP route dampening configuration settings.

protocols bgp <asn> parameters default

Sets default routing behaviors for the system.

Syntax

```
set protocols bgp asn parameters default [local-pref pref | no-ipv4-unicast]
delete protocols bgp asn parameters default [local-pref | no-ipv4-unicast]
show protocols bgp asn parameters default
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    parameters {
      default {
        local-pref pref
        no-ipv4-unicast
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
local-pref <i>pref</i>	Optional. Specifies the degree of preference that iBGP peers are to give local routes during BGP best path selection. The higher the value, the more the route is to be preferred. The range is 0 to 4294967295. The default is 100.
no-ipv4-unicast	Optional. Disables the IPv4 unicast address family as the default for peering session establishment. By default, IPv4 address family prefixes are automatically exchanged.

Default

The value of the local-pref attribute is 100. IPv4 unicast is the default address family.

Usage Guidelines

Use the **set** form of this command to override default local route preferences and automatic address family exchanges.

Use the **delete** form of this command to reset the route preferences and address family exchanges to the default.

Use the **show** form of this command to view BGP default routing configuration settings.

protocols bgp <asn> parameters deterministic-med

Enables or disables enforcing of deterministic MED.

Syntax

```
set protocols bgp asn parameters deterministic-med
delete protocols bgp asn parameters deterministic-med
show protocols bgp asn parameters
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    parameters {
      deterministic-med
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
------------	---

Default

Deterministic MED is not enforced.

Usage Guidelines

Use the **set** form of this command to enforce the deterministic comparison of the Multi Exit Discriminator (MED) value between all paths received from within the same autonomous system.

Use the **delete** form of this command to disable required MED comparison.

Use the **show** form of this command to view global BGP configuration settings.

protocols bgp <asn> parameters distance global

Sets administrative distance for all BGP routes.

Syntax

```
set protocols bgp asn parameters distance global {external dist | internal dist | local dist}
```

```
delete protocols bgp asn parameters distance global {external | internal | local}
```

```
show protocols bgp asn parameters distance global {external | internal | local}
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    parameters {
      distance {
        global {
          external dist
          internal dist
          local dist
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>external dist</i>	Mandatory. The externally learned (eBGP) route distance. The range of values is 1 to 255.
<i>internal dist</i>	Mandatory. The internally learned (iBGP) route distance. The range of values is 1 to 255.

local <i>dist</i>	Mandatory. The locally generated route distance. The range of values is 1 to 255.
--------------------------	---

Default

None.

Usage Guidelines

This command sets the administrative distance for BGP routes. All three distances (external, internal, and local) must be defined.

Use the **set** form of this command to set the administrative distance for BGP routes.

Use the **delete** form of this command to remove distance configuration settings.

Use the **show** form of this command to view distance configuration settings.

protocols bgp <asn> parameters distance prefix <ipv4net> distance <dist>

Sets administrative distance for BGP routes for a specific prefix.

Syntax

```
set protocols bgp asn parameters distance prefix ipv4net distance dist
delete protocols bgp asn parameters distance prefix ipv4net
show protocols bgp asn parameters distance prefix ipv4net
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    parameters {
      distance {
        prefix ipv4net {
          distance dist
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>ipv4net</i>	Mandatory. Multi-node. An IPv4 network in the format <i>ip-address/prefix</i> . You can specify multiple prefixes by creating multiple prefix configuration nodes.
distance <i>dist</i>	The route distance for the specified prefix. The range of values is 1 to 255.

Default

None.

Usage Guidelines

This command sets the administrative distance for a specific prefix.

Use the **set** form of this command to set the administrative distance for the specified prefix.

Use the **delete** form of this command to remove distance configuration settings.

Use the **show** form of this command to view distance configuration settings.

protocols bgp <asn> parameters disable-network-import-check

Disables IGP route check for network statements.

Syntax

```
set protocols bgp asn parameters disable-network-import-check
delete protocols bgp asn parameters disable-network-import-check
show protocols bgp asn parameters
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    parameters {
      disable-network-import-check
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
------------	---

Default

IGP route check for network statements is enabled.

Usage Guidelines

This command will cause BGP to advertise a network regardless of the Interior Gateway Protocol (IGP) in use.

Use the **set** form of this command to disable IGP route checks for network statements.

Use the **delete** form of this command to enable IGP route checks for network statements.

Use the **show** form of this command to view global BGP configuration settings.

protocols bgp <asn> parameters enforce-first-as

Enables or disables forcing eBGP peers to list AS number at the beginning of the AS_PATH attribute in incoming updates.

Syntax

```
set protocols bgp asn parameters enforce-first-as
delete protocols bgp asn parameters enforce-first-as
show protocols bgp asn parameters
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    parameters {
      enforce-first-as
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
------------	---

Default

Disabled.

Usage Guidelines

Use the set form of this command to direct the router to enforce listing of an eBGP peer's AS number at the beginning of the AS_PATH.

When this feature is enabled, the router will deny an update received from an external BGP (eBGP) peer unless the AS number is listed at the beginning of the AS_PATH in the incoming update. This prevents “spoof” situations where a misconfigured or unauthorized peer misdirecting traffic by advertising a route as if it were sourced from another autonomous system.

Use the **delete** form of this command to disable this behavior.

Use the **show** form of this command to view global BGP configuration settings.

protocols bgp <asn> parameters graceful-restart

Enables or disables graceful restart of the BGP process.

Syntax

```
set protocols bgp asn parameters graceful-restart [stalepath-time seconds]  
delete protocols bgp asn parameters graceful-restart  
show protocols bgp asn parameters
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    parameters {  
      graceful-restart {  
        stalepath-time seconds  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
stalepath-time <i>seconds</i>	Optional. Sets the maximum interval, in seconds, to retain stale paths for a restarting peer. If this interval is exceeded, all stale paths are removed. The range is 1 to 3600. The default is 360. The default value for this attribute is optimal for most deployments, and changing this value can have negative effects on network behavior. We recommend that only experienced network operators change this values.

Default

By default, paths for restarting peers are retained for at most 360 seconds.

Usage Guidelines

Use the **set** form of this command to direct the router to restart gracefully when it is reset.

Use the **delete** form of this command to disable graceful restart.

Use the **show** form of this command to view global BGP configuration settings.

protocols bgp <asn> parameters log-neighbor-changes

Enables or disables logging of the status of BGP neighbors.

Syntax

```
set protocols bgp asn parameters log-neighbor-changes
delete protocols bgp asn parameters log-neighbor-changes
show protocols bgp asn parameters
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    parameters {
      log-neighbor-changes
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
------------	---

Default

Disabled.

Usage Guidelines

Use the **set** form of this command to direct the router to log changes in adjacencies for BGP neighbors.

This feature helps detect network problems, by recording when BGP neighbors come up or go down, and when they reset. Log messages are sent to the main log file.

Use the **delete** form of this command to disable logging of neighbor status changes.

Use the **show** form of this command to view global BGP configuration settings.

Even when this feature is disabled, the system tracks neighbor resets; these resets can be seen in the output of the **show ip bgp neighbors** command.

protocols bgp <asn> parameters no-fast-external-failover

Enables or disables automatic resetting of BGP sessions for failed links.

Syntax

```
set protocols bgp asn parameters no-fast-external-failover
delete protocols bgp asn parameters no-fast-external-failover
show protocols bgp asn parameters
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    parameters {
      no-fast-external-failover
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
------------	---

Default

Disabled. By default, sessions are automatically reset.

Usage Guidelines

Use the **set** form of this command to disable fast external failover.

When fast external failover is enabled, then BGP sessions for directly adjacent external peers are immediately reset if the link fails.

Use the **delete** form of this command to restore fast external failover.

Use the **show** form of this command to view global BGP configuration settings.

protocols bgp <asn> parameters router-id <id>

Sets a fixed BGP router ID for the router, overriding the automatic ID selection process.

Syntax

```
set protocols bgp asn parameters router-id id
delete protocols bgp asn parameters router-id id
show protocols bgp asn parameters
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
    bgp asn {
        parameters {
            router-id id
        }
    }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The ID to be used by the router as the BGP router ID.

Default

When router ID is not explicitly set, then the router ID is automatically set to the IP address of the loopback address, if configured. If the loopback address is not configured, the router ID is set to the first IP address configured on a physical interface.

Usage Guidelines

Use the **set** form of this command to configure a fixed router ID for the local BGP routing process. This ID will override the router ID automatic set by the system.

Use the **delete** form of this command to remove the fixed router ID and restore the automatically selected ID.

Use the **show** form of this command to view global BGP configuration settings.

protocols bgp <asn> parameters scan-time <seconds>

Sets the scanning interval for the router.

Syntax

```
set protocols bgp asn parameters scan-time seconds
delete protocols bgp asn parameters scan-time seconds
show protocols bgp asn parameters
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    parameters {
      scan-time seconds
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>seconds</i>	Mandatory. The interval, in seconds, at which the router scans for BGP routing information. The range is 5 to 60. The default is 15.

Default

Scans take place every 15 seconds.

Usage Guidelines

Use the **set** form of this command to set the interval at which the router scans for BGP routing information.

Use the **delete** form of this command to reset the scanning interval to the default.

Use the **show** form of this command to view global BGP configuration settings.

protocols bgp <asn> timers

Sets BGP timers globally for the local router.

Syntax

```
set protocols bgp asn timers [keepalive seconds | holdtime seconds]  
delete protocols bgp asn timers [keepalive | holdtime]  
show protocols bgp asn timers [keepalive | holdtime]
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        timers {  
            keepalive seconds  
            holdtime seconds  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>keepalive seconds</i>	Optional. The frequency, in seconds, with which the local router sends keep-alive messages to neighbors. The range is 1 to 65535. The default is 60.
<i>holdtime seconds</i>	Optional. The maximum interval, in seconds, after which if the local router has not received a keep-alive message from the neighbor, a neighbor is declared dead. The range is 0 and 4 to 65535, where 0 disables the holdtime timer. The default is 180.

Default

The default for the keep-alive timer is 60 seconds. The default for the holdtime timer is 180 seconds.

Usage Guidelines

Use the **set** form of this command to set global BGP timers for monitoring the health of remote peers. These timers will be applied to all remote peers unless a neighbor has timers explicitly configured for it. Timers explicitly specified for a neighbor override the timers set globally.

Use the **delete** form of this command to restore global BGP timers to default values.

Use the **show** form of this command to view global BGP timer configuration settings.

reset ip bgp <address>

Resets a BGP connection.

Syntax

```
reset ip bgp {ipv4|ipv6} [in [prefix-filter] | out | rsclient | soft [in | out]]
```

Command Mode

Operational mode.

Parameters

<i>ipv4</i>	Resets the connection for the IPv4 BGP peer at the specified address.
<i>ipv6</i>	Resets the connection for the IPv6 BGP peer at the specified address.
in	Optional. Resets the inbound session only.
prefix-filter	Optional. Resets the BGP outbound route filter (ORF). This keyword is ignored unless ORF capabilities have been enabled on the local system or received from the sending BGP peer. In this case, a normal inbound soft reset is performed.
out	Optional. Resets the outbound session only.
rsclient	Optional. Resets only connections in the route server client Routing Information Base (RIB).
soft	Optional. Uses soft reconfiguration for the reset.
in	Optional. Resets with soft reconfiguration only inbound sessions.
out	Optional. Resets with soft reconfiguration only outbound sessions.

Default

Both inbound and outbound sessions are reset.

Usage Guidelines

Use this command to reset a BGP connection to a peer.

reset ip bgp <address> ipv4 unicast

Resets an IPv4 unicast BGP connection.

Syntax

```
reset ip bgp {ipv4|ipv6} ipv4 unicast {in [prefix-filter] | out | soft [in | out]}
```

Command Mode

Operational mode.

Parameters

<i>ipv4</i>	Resets the connection for the IPv4 BGP peer at the specified address.
<i>ipv6</i>	Resets the connection for the IPv6 BGP peer at the specified address.
in	Optional. Resets inbound sessions only.
prefix-filter	Optional. Resets the BGP outbound route filter (ORF). The <i>prefix-filter</i> keyword is ignored unless ORF capabilities have been enabled on the local system or received from the sending BGP peer. In this case, a normal inbound soft reset is performed.
out	Optional. Resets outbound sessions only.
soft	Optional. Uses soft reconfiguration for the reset.
in	Optional. Resets with soft reconfiguration only inbound sessions.
out	Optional. Resets with soft reconfiguration only outbound sessions.

Default

When used without the **soft** option, reset connections are dropped, both inbound and outbound.

Usage Guidelines

Use this command to reset an inbound BGP IPv4 unicast session for a given IP address. This forces BGP updates to be generated and new BGP policies to be applied.

Unless the **soft** option is used, all connections are dropped (a “hard reset”): TCP connections are terminated and all routes received from the neighbor are removed from the BGP routing table. Then the connection with the neighbor is reestablished.

If the **soft** option is used, routes from the neighbor are marked as stale but are not immediately removed from the BGP table. Stale routes that are not received from the neighbor when the connection is reestablished are removed from the BGP table at that point.

reset ip bgp all

Resets all BGP peering sessions.

Syntax

```
reset ip bgp all
```

Command Mode

Operational mode.

Parameters

None.

Default

None.

Usage Guidelines

Use this command to reset all BGP peering sessions.

reset ip bgp dampening

Resets BGP route dampening information and unsuppresses suppressed routes.

Syntax

```
reset ip bgp dampening [ipv4 [ipv4-mask] | ipv4net]
```

Command Mode

Operational mode.

Parameters

<i>ipv4</i>	Optional. Resets route dampening information for the IPv4 peer at the specified address.
<i>ipv4-mask</i>	Optional. An IPv4 network mask.
<i>ipv4net</i>	Optional. Resets route dampening information for all peers on the specified network. The format is <i>ip-address/prefix</i> .

Default

When used with no option, this command resets route dampening information and unsuppresses routes for all BGP peers.

Usage Guidelines

Use this command on a router running BGP to reset information related to route dampening and unsuppress routes that are currently suppressed.

show ip bgp

Displays BGP routes.

Syntax

```
show ip bgp [ipv4 | ipv4net [longer-prefixes] | summary]
```

Command Mode

Operational mode.

Parameters

<i>ipv4</i>	Optional. Displays routes for the neighbor at the specified IPv4 address.
<i>ipv4net</i>	Optional. Displays routes for the specified IPv4 network.
longer-prefixes	Optional. Displays any routes more specific than the one specified.
summary	Optional. Shows summary BGP route information.

Default

Displays all BGP routes by default.

Usage Guidelines

Use this command to display the BGP routing table.

show ip bgp attribute-info

Displays BGP attribute information.

Syntax

```
show ip bgp attribute-info
```

Command Mode

Operational mode.

Parameters

None.

Default

None.

Usage Guidelines

Use this command to display BGP attribute information.

show ip bgp cidr-only

Displays BGP routes with CIDR network masks.

Syntax

```
show ip bgp cidr-only
```

Command Mode

Operational mode.

Parameters

None.

Default

None.

Usage Guidelines

Use this command to display only routes with non-natural network masks; that is, Classless Inter Domain Routing network masks.

show ip bgp community-info

Displays BGP community information.

Syntax

```
show ip bgp community-info
```

Command Mode

Operational mode.

Parameters

None.

Default

None.

Usage Guidelines

Use this command to display BGP community information.

show ip bgp community <community>

Displays BGP routes belonging to the specified BGP community.

Syntax

```
show ip bgp community community [exact-match]
```

Command Mode

Operational mode.

Parameters

<i>community</i>	Mandatory. A BGP community identifier in the form AA:NN (where AA and NN are in the range of 0-65535), one of the well-known BGP communities local-AS , no-export , or no-advertise , or a space-separated list of up to four community identifiers.
------------------	---

exact-match	Optional. Displays only routes that have an exact match.
--------------------	--

Default

None.

Usage Guidelines

Use this command to display the BGP routes belonging to up to four BGP communities.

show ip bgp community-list <list-name>

Displays BGP routes permitted by the specified community list.

Syntax

```
show ip bgp community-list list-name [exact-match]
```

Command Mode

Operational mode.

Parameters

<i>list-name</i>	Mandatory. A preconfigured list of BGP communities.
------------------	---

exact-match	Optional. Displays only route that have an exact match.
--------------------	---

Default

None.

Usage Guidelines

Use this command to display the BGP routes permitted by the specified community list.

show ip bgp dampened-paths

Displays BGP routes that are currently dampened.

Syntax

```
show ip bgp dampened-paths
```

Command Mode

Operational mode.

Parameters

None.

Default

None.

Usage Guidelines

Use this command to display the BGP routes that are currently dampened.

show ip bgp filter-list <as-path-list-name>

Displays routes matching a list of autonomous system paths.

Syntax

```
show ip bgp filter-list as-path-list-name
```

Command Mode

Operational mode.

Parameters

<i>as-path-list-name</i>	Mandatory. Name of an autonomous system path list.
--------------------------	--

Default

None.

Usage Guidelines

Use this command to filter displayed routes according to preconfigured list of autonomous system paths.

BGP filter lists are defined using the **policy** command. For information about creating filter lists, see the *Vyatta Routing Policies Reference Guide*.

show ip bgp flap-statistics

Displays route flap statistics for BGP routes.

Syntax

```
show ip bgp flap-statistics [ipv4 | ipv4net] [longer-prefixes]
```

Command Mode

Operational mode.

Parameters

<i>ipv4</i>	Optional. Displays flap statistics for routes matching the specified IPv4 address.
<i>ipv4net</i>	Optional. Displays flap statistics for routes matching the specified IPv4 network.
longer-prefixes	Optional. Displays any routes more specific than the one specified.

Default

Displays route flap statistics for all BGP routes.

Usage Guidelines

Use this command to display statistics for flapping BGP routes.

show ip bgp flap-statistics cidr-only

Displays only route flap statistics for BGP routes with CIDR network masks.

Syntax

```
show ip bgp flap-statistics cidr-only
```

Command Mode

Operational mode.

Parameters

None.

Default

None.

Usage Guidelines

Use this command to display route flap statistics for BGP routes with non-natural network masks; that is, Classless Inter Domain Routing (CIDR) network masks.

show ip bgp flap-statistics filter-list <as-path-list-name>

Displays route flap statistics for BGP routes matching a list of AS paths.

Syntax

```
show ip bgp flap-statistics filter-list as-path-list-name
```

Command Mode

Operational mode.

Parameters

<i>as-path-list-name</i>	Mandatory. Name of an autonomous system path list.
--------------------------	--

Default

None.

Usage Guidelines

Use this command to display route flap statistics for BGP routes matching a preconfigured access list of autonomous system (AS) paths.

AS path filter lists are configured using the **policy** command. For information about creating path filter lists, see the *Vyatta Routing Policies Reference Guide*.

show ip bgp flap-statistics prefix-list <list-name>

Displays route flap statistics for BGP routes matching a prefix list.

Syntax

```
show ip bgp flap-statistics prefix-list list-name
```

Command Mode

Operational mode.

Parameters

<i>list-name</i>	Mandatory. The name of a defined prefix list.
------------------	---

Default

None.

Usage Guidelines

Use this command to display BGP statistics for BGP routes matching a list of network prefixes.

Prefix lists are configured using using the **policy** command. For information about creating filter lists, see the *Vyatta Routing Policies Reference Guide*.

show ip bgp flap-statistics regexp <expr>

Displays route flap statistics for BGP routes matching an AS path regular expression.

Syntax

```
show ip bgp flap-statistics regexp expr
```

Command Mode

Operational mode.

Parameters

<i>expr</i>	Mandatory. A POSIX-style regular expression representing a set of AS paths.
-------------	---

Default

Displays flap statistics for all BGP routes by default.

Usage Guidelines

Use this command to display route flap statistics for BGP routes matching a regular expression representing a set of autonomous system (AS) paths.

show ip bgp flap-statistics route-map <map-name>

Displays route flap statistics for BGP routes matching a route map.

Syntax

```
show ip bgp flap-statistics route-map map-name
```

Command Mode

Operational mode.

Parameters

<i>map-name</i>	Optional. The name of a defined route map.
-----------------	--

Default

Displays flap statistics for all BGP routes by default.

Usage Guidelines

Use this command to display route flap statistics for BGP routes matching a preconfigured route map.

Route maps are configured using the **policy** command. For information about creating route maps, see the *Vyatta Routing Policies Reference Guide*.

show ip bgp ipv4 unicast

Displays information for IPv4 unicast BGP routes.

Syntax

```
show ip bgp ipv4 unicast [ipv4 | ipv4net [longer-prefixes] | summary]
```

Command Mode

Operational

Parameters

<i>ipv4</i>	Optional. Show BGP information for the specified address.
<i>ipv4net</i>	Optional. Show BGP information for the specified network.
longer-prefixes	Optional. Displays any routes more specific than the one specified.
summary	Optional. Displays summary IPv4 unicast route information.

Default

Displays all IPv4 unicast BGP routes by default.

Usage Guidelines

Use this command to display IPv4 unicast routes in the BGP routing table.

show ip bgp ipv4 unicast cidr-only

Displays IPv4 unicast BGP routes with CIDR network masks.

Syntax

```
show ip bgp ipv4 unicast cidr-only
```

Command Mode

Operational

Parameters

None.

Default

None.

Usage Guidelines

Use this command to display only BGP IPv4 unicast routes with non-natural network masks; that is, Classless Inter Domain Routing (CIDR) network masks.

show ip bgp ipv4 unicast community <community>

Displays IPv4 unicast BGP routes belonging to the specified community.

Syntax

```
show ip bgp ipv4 unicast community community [exact-match]
```

Command Mode

Operational mode.

Parameters

<i>community</i>	Mandatory. A BGP community identifier in the form AA:NN (where AA and NN are in the range of 0-65535), one of the well-known BGP communities local-AS , no-export , or no-advertise , or a space-separated list of up to four community identifiers.
------------------	---

exact-match	Optional. Displays only routes that have an exact match.
--------------------	--

Default

None.

Usage Guidelines

Use this command to display BGP IPv4 unicast routes belonging to up to four BGP communities.

show ip bgp ipv4 unicast community-list <list-name>

Displays IPv4 unicast BGP routes permitted by the specified community list.

Syntax

```
show ip bgp ipv4 unicast community-list list-name [exact-match]
```

Command Mode

Operational mode.

Parameters

<i>list-name</i>	Mandatory. A preconfigured list of BGP communities.
exact-match	Optional. Displays routes that have an exact match.

Default

None.

Usage Guidelines

Use this command to display BGP IPv4 unicast routes permitted by the specified community list.

show ip bgp ipv4 unicast filter-list <as-path-list-name>

Displays IPv4 unicast BGP routes permitted by the specified AS path list.

Syntax

```
show ip bgp ipv4 unicast filter-list as-path-list-name
```

Command Mode

Operational mode.

Parameters

<i>as-path-list-name</i>	Mandatory. Name of an autonomous system path list.
--------------------------	--

Default

None.

Usage Guidelines

Use this command to display BGP IPv4 unicast routes permitted by the specified AS path list.

show ip bgp ipv4 unicast paths

Displays BGP IPv4 unicast path information.

Syntax

```
show ip bgp ipv4 unicast paths
```

Command Mode

Operational mode.

Parameters

None.

Default

None.

Usage Guidelines

Use this command to display BGP IPv4 unicast path information.

show ip bgp ipv4 unicast prefix-list <list-name>

Displays IPv4 unicast BGP routes matching a prefix list.

Syntax

```
show ip bgp ipv4 unicast prefix-list list-name
```

Command Mode

Operational mode.

Parameters

<i>list-name</i>	Mandatory. Name of a defined prefix list.
------------------	---

Default

None.

Usage Guidelines

Use this command to display BGP IPv4 unicast routes matching a preconfigured prefix list.

Prefix lists are configured using the **policy** command. For information about creating prefix lists, see the *Vyatta Routing Policies Reference Guide*.

show ip bgp ipv4 unicast regexp <regexp>

Displays IPv4 unicast BGP routes matching an AS path regular expression.

Syntax

```
show ip bgp ipv4 unicast regexp regexp
```

Command Mode

Operational mode.

Parameters

<i>regexp</i>	Mandatory. A POSIX-style regular expression representing a set of AS paths.
---------------	---

Default

None.

Usage Guidelines

Use this command to display BGP IPv4 unicast routes matching the specified AS path regular expression.

show ip bgp ipv4 unicast route-map <map-name>

Displays IPv4 unicast BGP routes matching a route map.

Syntax

```
show ip bgp ipv4 unicast route-map map-name
```

Command Mode

Operational mode.

Parameters

<i>map-name</i>	Mandatory. Name of a defined route map.
-----------------	---

Default

None.

Usage Guidelines

Use this command to display IPv4 unicast BGP routes matching a preconfigured route map.

Route maps are defined using the **policy** command. For information about creating route maps, see the *Vyatta Routing Policies Reference Guide*.

show ip bgp ipv4 unicast statistics

Displays statistics for BGP IPv4 unicast routes.

Syntax

```
show ip bgp ipv4 unicast statistics
```

Command Mode

Operational mode.

Parameters

None.

Default

None.

Usage Guidelines

Use this command to display BGP IPv4 unicast statistics.

show ip bgp ipv4 unicast summary

Displays a summary of BGP IPv4 unicast information.

Syntax

```
show ip bgp ipv4 unicast summary
```

Command Mode

Operational mode.

Parameters

None.

Default

None.

Usage Guidelines

Use this command to display a summary of BGP IPv4 unicast information.

show ip bgp memory

Displays memory usage for BGP.

Syntax

```
show ip bgp memory
```

Command Mode

Operational mode.

Parameters

None.

Default

None.

Usage Guidelines

Use this command to display the amount of memory being used for BGP, including the RIB, cache entries, attributes, AS-PATH entries, and hashes.

show ip bgp paths

Displays all BGP paths.

Syntax

```
show ip bgp paths
```

Command Mode

Operational mode.

Parameters

None.

Default

None.

Usage Guidelines

Use this command to display all BGP paths.

show ip bgp prefix-list <list-name>

Displays BGP routes matching a prefix list.

Syntax

```
show ip bgp prefix-list list-name
```

Command Mode

Operational mode.

Parameters

<i>list-name</i>	Mandatory. Name of a defined prefix list.
------------------	---

Default

None.

Usage Guidelines

Use this command to display routes that match a preconfigured prefix list.

Prefix lists are configured using the **policy** command. For information about creating prefix lists, see the *Vyatta Routing Policies Reference Guide*.

show ip bgp regexp <regexp>

Displays routes matching an AS path regular expression.

Syntax

```
show ip bgp regexp regexp
```

Command Mode

Operational mode.

Parameters

<i>regexp</i>	Mandatory. A POSIX-style regular expression representing a set of AS paths.
---------------	---

Default

None.

Usage Guidelines

Use this command to display routes matching a regular expression representing an autonomous system (AS) path list.

show ip bgp route-map <map-name>

Displays routes matching a route map.

Syntax

```
show ip bgp route-map map-name
```

Command Mode

Operational mode.

Parameters

<i>map-name</i>	Optional. The name of a defined route map.
-----------------	--

Default

None.

Usage Guidelines

Use this command to display routes matching a preconfigured route map.

Route maps are defined using the **policy** command. For information about creating route maps, see the *Vyatta Routing Policies Reference Guide*.

show ip bgp scan

Displays BGP scan status.

Syntax

```
show ip bgp scan
```

Command Mode

Operational mode.

Parameters

None.

Default

None.

Usage Guidelines

Use this command to display BGP scan status.

show ip route bgp

Displays BGP routes.

Syntax

```
show ip route bgp
```

Command Mode

Operational mode.

Parameters

None.

Default

None.

Usage Guidelines

Use this command to display BGP routes.

Examples

[Example 2-1](#) shows BGP routes.

Example 2-1 “show ip route bgp”: Displaying BGP routes

```
vyatta@vyatta:~$ show ip route bgp
Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
       I - ISIS, B - BGP, > - selected route, * - FIB route

B   10.1.0.0/24 [20/75] via 10.1.0.4 inactive, 1d00h46m
B>* 10.100.100.4/32 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h46m
B>* 10.104.104.4/32 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h46m
B>* 172.16.0.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
B>* 172.20.0.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
B>* 172.20.1.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
B>* 172.20.2.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
B>* 172.20.3.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
B>* 172.20.4.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
B>* 172.20.5.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
```

```
B>* 172.20.6.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
B>* 172.20.7.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
B>* 172.20.8.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
B>* 172.20.9.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
B>* 172.21.0.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
B>* 172.22.0.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
B>* 172.23.0.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
B>* 172.24.0.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
B>* 172.25.0.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
B>* 172.26.0.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
B>* 172.27.0.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
B>* 172.28.0.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
B>* 172.29.0.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 1d00h36m
B>* 192.168.3.0/24 [20/75] via 10.1.0.4 (recursive via 10.3.0.1), 05:15:56
vyatta@vyatta:~$
```

show ipv6 bgp

Displays BGP routes.

Syntax

```
show ipv6 bgp [ipv6 | ipv6net [longer-prefixes] | summary]
```

Command Mode

Operational mode.

Parameters

<i>ipv6</i>	Optional. Displays routes for the neighbor at the specified IPv6 address.
<i>ipv6net</i>	Optional. Displays routes for the specified IPv6 network.
<i>longer-prefixes</i>	Optional. Displays any routes more specific than the one specified.
<i>summary</i>	Optional. Shows summary BGP route information.

Default

Displays all BGP routes by default.

Usage Guidelines

Use this command to display the BGP routing table.

show ipv6 bgp community <community>

Displays BGP routes belonging to the specified BGP community.

Syntax

```
show ipv6 bgp community community [exact-match]
```

Command Mode

Operational mode.

Parameters

<i>community</i>	Mandatory. A BGP community identifier in the form AA:NN (where AA and NN are in the range of 0-65535), one of the well-known BGP communities local-AS , no-export , or no-advertise , or a space-separated list of up to four community identifiers.
exact-match	Optional. Displays only routes that have an exact match.

Default

None.

Usage Guidelines

Use this command to display the BGP routes belonging to up to four BGP communities.

show ipv6 bgp community-list <list-name>

Displays BGP routes permitted by the specified community list.

Syntax

```
show ipv6 bgp community-list list-name [exact-match]
```

Command Mode

Operational mode.

Parameters

<i>list-name</i>	Mandatory. A preconfigured list of BGP communities.
------------------	---

exact-match	Optional. Displays only route that have an exact match.
--------------------	---

Default

None.

Usage Guidelines

Use this command to display the BGP routes permitted by the specified community list.

show ipv6 bgp filter-list <as-path-list-name>

Displays routes matching a list of autonomous system paths.

Syntax

```
show ipv6 bgp filter-list as-path-list-name
```

Command Mode

Operational mode.

Parameters

<i>as-path-list-name</i>	Mandatory. Name of an autonomous system path list.
--------------------------	--

Default

None.

Usage Guidelines

Use this command to filter displayed routes according to preconfigured list of autonomous system paths.

BGP filter lists are defined using the the **policy as-path-list** command.

show ipv6 bgp prefix-list <list-name>

Displays BGP routes matching a prefix list.

Syntax

```
show ipv6 bgp prefix-list list-name
```

Command Mode

Operational mode.

Parameters

<i>list-name</i>	Mandatory. Name of a defined prefix list.
------------------	---

Default

None.

Usage Guidelines

Use this command to display routes that match a preconfigured prefix list.

Prefix lists are configured using the **policy prefix-list** command.

show ipv6 bgp regexp <regexp>

Displays routes matching an AS path regular expression.

Syntax

```
show ipv6 bgp regexp regexp
```

Command Mode

Operational mode.

Parameters

<i>regexp</i>	Mandatory. A POSIX-style regular expression representing a set of AS paths.
---------------	---

Default

None.

Usage Guidelines

Use this command to display routes matching a regular expression representing an autonomous system (AS) path list.

show monitoring protocols bgp

Displays BGP protocol debugging flags.

Syntax

```
show monitoring protocols bgp
```

Command Mode

Operational mode.

Parameters

None.

Default

None.

Usage Guidelines

Use this command to display BGP protocol debugging flags.

Chapter 3: Route Reflection

This chapter describes commands for for BGP route reflection.

This chapter presents the following topics:

- [Route Reflection Commands](#)

Route Reflection Commands

This chapter contains the following commands.

Configuration Commands

<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast route-reflector-client</code>	Specify this neighbor as a route reflector client.
<code>protocols bgp <asn> neighbor <id> route-reflector-client</code>	Defines the local router as a BGP route reflector, with the neighbor as a route reflector client.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast route-reflector-client</code>	Specify this peer group as a route reflector client.
<code>protocols bgp <asn> peer-group <group-name> route-reflector-client</code>	Specify this peer group as a route reflector client.
<code>protocols bgp <asn> parameters cluster-id <id></code>	Sets the cluster ID for a BGP route reflection cluster.
<code>protocols bgp <asn> parameters no-client-to-client-reflection</code>	Enables or disables route reflection from a BGP route reflector to clients.

Operational Commands

None.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast route-reflector-client

Specify this neighbor as a route reflector client.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast route-reflector-client
```

```
delete protocols bgp asn neighbor id address-family ipv6-unicast route-reflector-client
```

```
show protocols bgp asn neighbor id address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          route-reflector-client
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Default

Disabled.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to configure the BGP neighbor as a route reflector client.

Use the **delete** form of this command to remove the configuration.

Use the **show** form of this command to view the configuration.

protocols bgp <asn> neighbor <id> route-reflector-client

Defines the local router as a BGP route reflector, with the neighbor as a route reflector client.

Syntax

```
set protocols bgp asn neighbor id route-reflector-client
delete protocols bgp asn neighbor id route-reflector-client
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      route-reflector-client
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Default

There is no route reflector in the autonomous system.

Usage Guidelines

Use the set form of this command to make the local router a BGP route reflector, and to designate the specified neighbor as a route reflector client.

Another technology designed to help ASs with large numbers of iBGP peers is route reflection. In a standard BGP implementation, all iBGP peers must be fully meshed. Because of this requirement, when an iBGP peer learns a route from another iBGP peer, the receiving router does not forward the route to any of its iBGP peers, since these routers should have learned the route directly from the announcing router.

In a route reflector environment the iBGP peers are no longer fully meshed. Instead, each iBGP peer has an iBGP connection to one or more route reflectors (RRs). Routers configured with a connection to an RR server are referred to as RR clients. Only the RR server is configured to be aware that the RR client is part of an RR configuration; from the RR client's point of view, it is configured normally, and does not have any awareness that it is part of a RR configuration.

In route reflection, internal peers of an RR server are categorized into two types:

- **Client peers.** The RR server and its client peers form a cluster. Within a cluster, client peers need not be fully meshed, but must have an iBGP connection to at least one RR in the cluster.
- **Non-client peers.** Non-client peers, including the RR server, must be fully meshed.

An RR environment is unlike a regular environment, where iBGP peers never forward a route update to other iBGP peers (which is the reason why each iBGP peer must peer with all other peers). When an RR server receives an iBGP update from an RR client, these route updates can also be sent to all other RR clients. When an RR server receives a route update from a peer, it selects the best path based on its path selection rule. After the best path is selected, the RR server chooses its action depending on the type of the peer from which it learned the best path.

- If the route was learned from a client peer, the RR reflects the route to both client and non-client peers. All iBGP updates from client peers are reflected to all other client peers in the cluster. This is done regardless of whether the update was the best path for the RR itself.
- If the route was learned from a non-client iBGP peer, it is reflected out to all RR client peers.
- If the route was learned from an eBGP peer, the route is reflected to all RR clients and all non-clients.

Use the **delete** form of this command to remove the neighbor as a route reflector client.

Use the **show** form of this command to view BGP neighbor configuration settings.

Usage Guidelines

Use the **set** form of this command to disable route reflection from a BGP route reflector to clients. When client-to-client route reflection is enabled, the configured route reflector reflects routes from one client to other clients.

Use the **delete** form of this command to enable client-to-client route reflection.

Use the **show** form of this command to view global BGP configuration settings.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast route-reflector-client

Specify this peer group as a route reflector client.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast  
route-reflector-client
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
route-reflector-client
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      address-family {  
        ipv6-unicast {  
          route-reflector-client  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

Default

Disabled.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to configure the BGP peer group as a route reflector client.

Use the **delete** form of this command to remove the configuration.

Use the **show** form of this command to view the configuration.

protocols bgp <asn> peer-group <group-name> route-reflector-client

Specify this peer group as a route reflector client.

Syntax

```
set protocols bgp asn peer-group group-name route-reflector-client
delete protocols bgp asn peer-group group-name route-reflector-client
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      route-reflector-client
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

Default

Disabled.

Usage Guidelines

Use the set form of this command to configure the BGP peer group as a route reflector client.

Use the **delete** form of this command to remove the configuration.

Use the **show** form of this command to view the configuration.

protocols bgp <asn> parameters cluster-id <id>

Sets the cluster ID for a BGP route reflection cluster.

Syntax

```
set protocols bgp asn parameters cluster-id id
delete protocols bgp asn parameters cluster-id id
show protocols bgp asn parameters
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    parameters {
      cluster-id id
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. A network address uniquely identifying the route reflection cluster.

Default

None.

Usage Guidelines

Use the set form of this command to specify the route reflection cluster ID for an internal route reflection cluster.

When there is only one route reflector in a cluster, the cluster uses the router ID of the route reflector as an ID. If more than one route reflector is deployed in a cluster (for example, to provide redundancy), you must provide an ID for the cluster. In this case, the cluster ID is assigned to every route reflector in the cluster.

Use the **delete** form of this command to remove a cluster ID.

Use the **show** form of this command to view global BGP configuration settings.

protocols bgp <asn> parameters no-client-to-client-reflection

Enables or disables route reflection from a BGP route reflector to clients.

Syntax

```
set protocols bgp asn parameters no-client-to-client-reflection
delete protocols bgp asn parameters no-client-to-client-reflection
show protocols bgp asn parameters
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    parameters {
      no-client-to-client-reflection
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
------------	---

Default

By default client-to-client-reflection is enabled.

Usage Guidelines

Use the **set** form of this command to disable route reflection from a BGP route reflector to clients. When client-to-client route reflection is enabled, the configured route reflector reflects routes from one client to other clients.

Use the **delete** form of this command to enable client-to-client route reflection.

Use the **show** form of this command to view global BGP configuration settings.

Chapter 4: Confederations

This chapter describes commands for BGP confederations.

This chapter presents the following topics:

- [Confederation Commands](#)

Confederation Commands

This chapter contains the following commands.

Configuration Commands

`protocols bgp <asn> parameters confederation
identifier <asn>` Defines a BGP confederation.

`protocols bgp <asn> parameters confederation peers
<asn>` Defines the autonomous systems that make up a BGP confederation.

Operational Commands

None.

protocols bgp <asn> parameters confederation identifier <asn>

Defines a BGP confederation.

Syntax

```
set protocols bgp asn parameters confederation identifier asn
delete protocols bgp asn parameters confederation identifier asn
show protocols bgp asn parameters confederation
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    parameters {
      confederation {
        identifier asn
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
identifier <i>asn</i>	Mandatory. The AS number of the BGP confederation. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.

Default

None.

Usage Guidelines

Use the **set** form of this command to create a BGP confederation.

Use the **delete** form of this command to remove the BGP confederation.

Use the **show** form of this command to view BGP confederation configuration settings.

protocols bgp <asn> parameters confederation peers <asn>

Defines the autonomous systems that make up a BGP confederation.

Syntax

```
set protocols bgp asn parameters confederation peers asn [asn... asn]  
delete protocols bgp asn parameters confederation peers asn [asn... asn]  
show protocols bgp asn parameters confederation
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    parameters {  
      confederation {  
        peers asn  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides.
<i>peers asn</i>	Mandatory. The subautonomous systems that will make up the BGP confederation. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems. Multiple ASs can be specified in a space-separated list.

Default

None.

Usage Guidelines

Use the **set** form of this command to list the subautonomous systems that will be the members of a BGP confederation. To a peer outside the confederation, the confederation appears as a single autonomous system.

Use the **delete** form of this command to remove an AS from a confederation.

Use the **show** form of this command to view BGP confederation configuration settings.

Chapter 5: Neighbors

This chapter describes commands for BGP neighbors .

This chapter presents the following topics:

- [Neighbor Commands](#)

Neighbor Commands

This chapter contains the following commands.

Configuration Commands	
<code>protocols bgp <asn> neighbor <id></code>	Defines a BGP neighbor.
<code>protocols bgp <asn> neighbor <id> advertisement-interval <seconds></code>	Sets the minimum interval for BGP route updates.
<code>protocols bgp <asn> neighbor <id> allowas-in</code>	Allows or disallows receiving BGP advertisements containing the AS path of the local router.
<code>protocols bgp <asn> neighbor <id> attribute-unchanged</code>	Allows the router to send updates to a neighbor with unchanged attributes.
<code>protocols bgp <asn> neighbor <id> capability dynamic</code>	Advertises support for dynamic update of BGP capabilities advertised and received from this neighbor.
<code>protocols bgp <asn> neighbor <id> capability orf</code>	Enables or disables forwarding of the default route to a BGP neighbor.
<code>protocols bgp <asn> neighbor <id> default-originate</code>	Provides a brief description for a BGP neighbor.
<code>protocols bgp <asn> neighbor <id> description <desc></code>	Provides a brief description for a BGP neighbor.
<code>protocols bgp <asn> neighbor <id> disable-capability-negotiation</code>	Disables BGP capability negotiation.
<code>protocols bgp <asn> neighbor <id> disable-connected-check</code>	Disables direct connection verification for single-hop eBGP peers.
<code>protocols bgp <asn> neighbor <id> disable-send-community</code>	Disables sending of community attributes to the specified neighbor.
<code>protocols bgp <asn> neighbor <id> distribute-list export <acl-num></code>	Applies an access list to filter outbound routing updates to this neighbor.
<code>protocols bgp <asn> neighbor <id> distribute-list import <acl-num></code>	Applies an access list to filter inbound routing updates from this neighbor.
<code>protocols bgp <asn> neighbor <id> ebgp-multihop <ttl></code>	Allows eBGP neighbors not on directly connected networks.
<code>protocols bgp <asn> neighbor <id> filter-list export <as-path-list-name></code>	Applies an AS path list to routing updates to this neighbor.
<code>protocols bgp <asn> neighbor <id> filter-list import <as-path-list-name></code>	Applies an AS path list to routing updates from this neighbor.

<code>protocols bgp <asn> neighbor <id> local-as <asn></code>	Defines a local autonomous system number for eBGP peerings.
<code>protocols bgp <asn> neighbor <id> maximum-prefix <max-num></code>	Sets the maximum number of prefixes to accept from this neighbor before that neighbor is taken down.
<code>protocols bgp <asn> neighbor <id> nexthop-self</code>	Sets the local router as the next hop for this neighbor.
<code>protocols bgp <asn> neighbor <id> override-capability</code>	Overrides capability negotiation to allow a peering session to be established with a neighbor that does not support capability negotiation.
<code>protocols bgp <asn> neighbor <id> passive</code>	Directs the router not to initiate connections with this neighbor.
<code>protocols bgp <asn> neighbor <id> password <pwd></code>	Defines a BGP MD5 password.
<code>protocols bgp <asn> neighbor <id> peer-group <group-name></code>	Assigns a neighbor as a member of the specified peer group.
<code>protocols bgp <asn> neighbor <id> port <port-num></code>	Specifies the port on which the neighbor is listening for BGP signals.
<code>protocols bgp <asn> neighbor <id> prefix-list export <list-name></code>	Applies a prefix list to filter updates to this neighbor.
<code>protocols bgp <asn> neighbor <id> prefix-list import <list-name></code>	Applies a prefix list to filter updates from this neighbor.
<code>protocols bgp <asn> neighbor <id> remote-as <asn></code>	Specifies the autonomous system number of the neighbor.
<code>protocols bgp <asn> neighbor <id> remove-private-as</code>	Directs the router to remove private AS numbers from updates sent to this neighbor.
<code>protocols bgp <asn> neighbor <id> route-map export <map-name></code>	Applies a route map to filter updates to this neighbor.
<code>protocols bgp <asn> neighbor <id> route-map import <map-name></code>	Applies a route map to filter updates to or from this neighbor.
<code>protocols bgp <asn> neighbor <id> shutdown</code>	Administratively shuts down a BGP neighbor.
<code>protocols bgp <asn> neighbor <id> soft-reconfiguration inbound</code>	Directs the router to store received routing updates.
<code>protocols bgp <asn> neighbor <id> strict-capability-match</code>	Directs the router to strictly match the capabilities of the neighbor.
<code>protocols bgp <asn> neighbor <id> timers</code>	Sets BGP timers for this neighbor.
<code>protocols bgp <asn> neighbor <id> ttl-security hops <hops></code>	Sets Time to Live (TTL) security hop count.

<code>protocols bgp <asn> neighbor <id> unsuppress-map <map-name></code>	Directs the router to selectively advertise routes suppressed by aggregating addresses, based on a route map.
<code>protocols bgp <asn> neighbor <id> update-source <source></code>	Specifies the source IP address or interface for routing updates.
<code>protocols bgp <asn> neighbor <id> weight <weight></code>	Defines a default weight for routes from this neighbor.
Neighbor Commands - IPv6 Routes	
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast</code>	Defines parameters for controlling IPv6 unicast routes for a peering session.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast allowas-in</code>	Allows or disallows receiving BGP advertisements containing the AS path of the local router.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast attribute-unchanged</code>	Allows the router to send updates to a neighbor with unchanged attributes.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast capability dynamic</code>	Advertises support for dynamic update of BGP capabilities advertised and received from this neighbor.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast capability orf</code>	Advertises support for Outbound Route Filtering (ORF) for updating BGP capabilities advertised and received from this neighbor.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast default-originate</code>	Enables or disables forwarding of the default route to a BGP neighbor.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast disable-send-community</code>	Disables sending of community attributes to the specified neighbor.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast distribute-list export <access-list6-name></code>	Applies an access list to filter outbound routing updates to this neighbor.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast distribute-list import <access-list6-name></code>	Applies an access list to filter inbound routing updates from this neighbor.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast filter-list export <as-path-list-name></code>	Applies a filter list to routing updates to this neighbor.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast filter-list import <as-path-list-name></code>	Applies a filter list to routing updates from this neighbor.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast maximum-prefix <max-num></code>	Sets the maximum number of prefixes to accept from this neighbor before that neighbor is taken down.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast nexthop-local unchanged</code>	Specifies that the IPv6 link-local address is not changed when the prefix is announced to the peer.

<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast nexthop-self</code>	Sets the local router as the next hop for this neighbor.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast peer-group <group-name></code>	Specifies the IPv6 peer group for this peer.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast prefix-list export <prefix-list6-name></code>	Applies a prefix list to filter updates to this neighbor.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast prefix-list import <prefix-list6-name></code>	Applies a prefix list to filter updates from this neighbor.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast remove-private-as</code>	Directs the router to remove private AS numbers from updates sent to this neighbor.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast route-map export <map-name></code>	Applies a route map to filter updates to this neighbor.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast route-map import <map-name></code>	Applies a route map to filter updates to or from this neighbor.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast soft-reconfiguration inbound</code>	Directs the router to store received routing updates.
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast unsuppress-map <map-name></code>	Directs the router to selectively advertise routes suppressed by aggregating addresses, based on a route map.
Operational Commands	
<code>reset ip bgp external</code>	Resets peering sessions for eBGP neighbors.
<code>reset ip bgp external ipv4 unicast</code>	Resets IPv4 unicast peering sessions for eBGP neighbors.
<code>show ip bgp ipv4 unicast neighbors</code>	Displays IPv4 unicast route information for BGP neighbors.
<code>show ip bgp ipv4 unicast neighbors <id> advertised-routes</code>	Displays advertised BGP IPv4 unicast routes for a BGP neighbor.
<code>show ip bgp ipv4 unicast neighbors <id> prefix-counts</code>	Displays IPv4 unicast prefix-counts for a BGP neighbor.
<code>show ip bgp ipv4 unicast neighbors <id> received prefix-filter</code>	Displays the IPv4 unicast prefix-lists received from a BGP neighbor.
<code>show ip bgp ipv4 unicast neighbors <id> received-routes</code>	Displays the IPv4 unicast routes received from a BGP neighbor.
<code>show ip bgp ipv4 unicast neighbors <id> routes</code>	Displays IPv4 unicast received and accepted routes from a BGP neighbor.
<code>show ip bgp neighbors</code>	Displays BGP neighbor information.

<code>show ip bgp neighbors <id> advertised-routes</code>	Displays advertised routes for a BGP neighbor.
<code>show ip bgp neighbors <id> dampened-routes</code>	Displays dampened routes to a BGP neighbor.
<code>show ip bgp neighbors <id> flap-statistics</code>	Displays route flap statistics for routes learned from a BGP neighbor.
<code>show ip bgp neighbors <id> prefix-counts</code>	Displays prefix counts for a BGP neighbor.
<code>show ip bgp neighbors <id> received prefix-filter</code>	Displays prefix lists received from a BGP neighbor.
<code>show ip bgp neighbors <id> received-routes</code>	Displays routes received from a BGP neighbor.
<code>show ip bgp neighbors <id> routes</code>	Displays all received and accepted routes from a BGP neighbor.
<code>show ipv6 bgp neighbors</code>	Displays BGP neighbor information.
<code>show ipv6 bgp neighbors <ipv6> advertised-routes</code>	Displays advertised routes for a BGP neighbor.
<code>show ipv6 bgp neighbors <ipv6> received-routes</code>	Displays routes received from a BGP neighbor.
<code>show ipv6 bgp neighbors <ipv6> routes</code>	Displays all received and accepted routes from a BGP neighbor.

Some command for working with BGP neighbors are described in other chapters.

Related Commands Documented Elsewhere

<code>protocols bgp <asn> parameters log-neighbor-changes</code>	See “Chapter 2: Global and Router-Specific Configuration.”
<code>protocols bgp <asn> neighbor <id> route-reflector-client</code>	See “Chapter 3: Route Reflection.”
<code>protocols bgp <asn> parameters confederation peers <asn></code>	See “Chapter 4: Confederations.”

protocols bgp <asn> neighbor <id>

Defines a BGP neighbor.

Syntax

```
set protocols bgp asn neighbor id
delete protocols bgp asn neighbor id
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {}
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. Multi-node. The IPv4 or IPv6 address of a BGP neighbor. You can define multiple BGP neighbors by creating multiple neighbor configuration nodes.

Default

None.

Usage Guidelines

Use the **set** form of this command to define a BGP neighbor.

Use the **delete** form of this command to remove a BGP neighbor.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast

Defines parameters for controlling IPv6 unicast routes for a peering session.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast
delete protocols bgp asn neighbor id address-family ipv6-unicast
show protocols bgp asn neighbor id address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {}
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Usage Guidelines

Use this command to define parameters for IPv6 unicast routes for a peering session. Creating this configuration node activates the peer to receive IPv6 prefixes.

Use the **set** form of this command to specify the configuration node.

Use the **delete** form of this command to delete the **ipv6-unicast** configuration.

Use the **show** form of this command to view the **ipv6-unicast** configuration.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast allowas-in

Allows or disallows receiving BGP advertisements containing the AS path of the local router.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast allowas-in [number num]
```

```
delete protocols bgp asn neighbor id address-family ipv6-unicast allowas-in
```

```
show protocols bgp asn neighbor id address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          allowas-in {
            number num
          }
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

number <i>num</i>	Optional. Specifies the number of times to accept an AS path containing the readvertisement of the local router's ASN. The range is 1 to 10. The default is 3.
-----------------------------	--

Default

Readvertisement is disabled.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to allow the router to accept BGP AS paths advertising the router's own AS number.

This situation could indicate a routing loop, and by default such an update is dropped. However, you can use this command to direct the router to accept updates readvertising its AS number for a specified number of times.

Use the **delete** form of this command to disallow readvertisement of the router's AS path.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast attribute-unchanged

Allows the router to send updates to a neighbor with unchanged attributes.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast attribute-unchanged  
[as-path | med | next-hop]
```

```
delete protocols bgp asn neighbor id address-family ipv6-unicast  
attribute-unchanged [as-path | med | next-hop]
```

```
show protocols bgp asn neighbor id address-family ipv6-unicast attribute-unchanged
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    neighbor id {  
      address-family {  
        ipv6-unicast {  
          attribute-unchanged {  
            as-path  
            med  
            next-hop  
          }  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
------------	--

<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
as-path	Optional. Propagates the route update with unchanged AS_PATH attribute.
med	Optional. Propagates the route update with unchanged Multi Exit Discriminator.
next-hop	Optional. Propagates the route update with unchanged next hop.

Default

Disabled.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to direct the router to propagate routing updates without modifying the BGP AS_PATH, Multi Exit Discriminator (MED), or next-hop attribute.

Use the **delete** form of this command to restore normal modification of BGP attributes.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast capability dynamic

Advertises support for dynamic update of BGP capabilities advertised and received from this neighbor.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast capability dynamic
delete protocols bgp asn neighbor id address-family ipv6-unicast capability dynamic
show protocols bgp asn neighbor id address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          capability {
            dynamic
          }
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Default

By default, the session is brought up with minimal capability on both sides.

Usage Guidelines

In general, BGP requires that if a BGP speaker receives an OPEN message with any unrecognized optional parameters, the speaker must terminate BGP peering. This makes it difficult to introduce new capabilities into the protocol. The Capabilities parameter allows graceful negotiation of BGP capabilities without requiring a peer to terminate peering.

This command specifies use of the BGP Dynamic Capability feature, which allows dynamic update of capabilities over an established BGP session.

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to specify dynamic update of BGP capabilities to be used in capability negotiation.

Use the **delete** form of this command to remove the dynamic update capability.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast capability orf

Advertises support for Outbound Route Filtering (ORF) for updating BGP capabilities advertised and received from this neighbor.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast capability orf  
[prefix-list [receive | send]]
```

```
delete protocols bgp asn neighbor id address-family ipv6-unicast capability orf
```

```
show protocols bgp asn neighbor id address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    neighbor id {  
      address-family {  
        ipv6-unicast {  
          capability {  
            orf {  
              prefix-list {  
                receive  
                send  
              }  
            }  
          }  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
prefix-list	Optional. Advertise prefix-list ORF capability to this neighbor.
receive	Optional. Advertise the ability to receive the ORF from this neighbor.
send	Optional. Advertise the capability to send the ORF to this neighbor.

Default

By default, the session is brought up with minimal capability on both sides.

Usage Guidelines

This command specifies use of BGP Outbound Route Filtering (ORF) to send and receive capabilities. Using ORF minimizes the number of BGP updates that are sent between peer routers.

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to specify BGP ORF capabilities to be used in capability negotiation.

Use the **delete** form of this command to remove the ORF capability.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast default-originate

Enables or disables forwarding of the default route to a BGP neighbor.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast default-originate  
[route-map map-name]
```

```
delete protocols bgp asn neighbor id address-family ipv6-unicast default-originate  
[route-map map-name]
```

```
show protocols bgp asn neighbor id address-family ipv6-unicast default-originate
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    neighbor id {  
      address-family {  
        ipv6-unicast {  
          default-originate {  
            route-map map-name  
          }  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

<i>map-name</i>	Optional. Specifies a configured route map to be used when advertising the default route.
-----------------	---

Default

The default route is not distributed.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to enable the router to advertise the default route `::/0` to this BGP neighbor. This route can then be used by the neighbor to reach the local router if no other routes are available.

The route `::/0` need not be explicitly configured on the local router.

If a route map is specified, the default route is advertised if two conditions are satisfied:

- The route map includes a **match ip** address clause.
- A route exists that exactly matches the IP access list.

Route maps are configured using `protocols bgp <asn> neighbor <id> local-as <asn> command`.

Use the **delete** form of this command to disable forwarding of the default route or to delete a route map.

Use the **show** form of this command to view BGP neighbor default route origination settings.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast disable-send-community

Disables sending of community attributes to the specified neighbor.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast
  disable-send-community [extended | standard]
delete protocols bgp asn neighbor id address-family ipv6-unicast
  disable-send-community
show protocols bgp asn neighbor id address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          disable-send-community {
            extended
            standard
          }
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
extended	Optional. Disable sending extended community attributes.

standard	Optional. Disable sending standard community attributes.
-----------------	--

Default

Communities are sent by default.

Usage Guidelines

Use this command to enable and disable sending community attributes between BGP neighbors.

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to disable sending community attributes.

Use the **delete** form of this command to restore sending community attributes.

Use the **show** form of this command to view BGP neighbor configuration settings

protocols bgp <asn> neighbor <id> address-family ipv6-unicast distribute-list export <access-list6-name>

Applies an access list to filter outbound routing updates to this neighbor.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast distribute-list export access-list6-name
```

```
delete protocols bgp asn neighbor id address-family ipv6-unicast distribute-list export
```

```
show protocols bgp asn neighbor id address-family ipv6-unicast distribute-list export
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    neighbor id {  
      address-family {  
        ipv6-unicast {  
          distribute-list {  
            export access-list6-name  
          }  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

<i>access-list6-name</i>	Optional. The name of the IPv6 access-list.
--------------------------	---

Default

None.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to apply an access list to filter outbound routing updates to a BGP neighbor.

Use the **delete** form of this command to disable outbound distribute list filtering.

Use the **show** form of this command to view BGP neighbor distribute list configuration settings.

NOTE *A neighbor distribute list cannot be used together with a neighbor prefix list in the same direction. These two lists are mutually exclusive, and only one list may be applied to a given direction.*

protocols bgp <asn> neighbor <id> address-family ipv6-unicast distribute-list import <access-list6-name>

Applies an access list to filter inbound routing updates from this neighbor.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast distribute-list import access-list6-name
```

```
delete protocols bgp asn neighbor id address-family ipv6-unicast distribute-list import
```

```
show protocols bgp asn neighbor id address-family ipv6-unicast distribute-list import
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          distribute-list {
            import access-list6-name
          }
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

<i>access-list6-name</i>	Optional. The name of the IPv6 access-list.
--------------------------	---

Default

None.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to apply an access list to filter inbound routing updates from a BGP neighbor.

Use the **delete** form of this command to disable inbound distribute list filtering.

Use the **show** form of this command to view BGP neighbor distribute list configuration settings.

NOTE *A neighbor distribute list cannot be used together with a neighbor prefix list in the same direction. These two lists are mutually exclusive, and only one list may be applied to the specified direction.*

protocols bgp <asn> neighbor <id> address-family ipv6-unicast filter-list export <as-path-list-name>

Applies a filter list to routing updates to this neighbor.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast filter-list export as-path-list-name
```

```
delete protocols bgp asn neighbor id address-family ipv6-unicast filter-list export as-path-list-name
```

```
show protocols bgp asn neighbor id address-family ipv6-unicast filter-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    neighbor id {  
      address-family {  
        ipv6-unicast {  
          filter-list {  
            export as-path-list-name  
          }  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>as-path-list-name</i>	Mandatory. Name of an autonomous system path list.

Default

None.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to apply an access list for filtering outbound routing updates.

Use the **delete** form of this command to disable outbound routing update filtering.

Use the **show** form of this command to view BGP neighbor filter list configuration settings.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast filter-list import <as-path-list-name>

Applies a filter list to routing updates from this neighbor.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast filter-list import
as-path-list-name
```

```
delete protocols bgp asn neighbor id address-family ipv6-unicast filter-list import
as-path-list-name
```

```
show protocols bgp asn neighbor id address-family ipv6-unicast filter-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          filter-list {
            import as-path-list-name
          }
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>as-path-list-name</i>	Mandatory. Name of an autonomous system path list.

Default

None.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to apply an access list for filtering inbound routing updates.

Use the **delete** form of this command to disable inbound routing update filtering.

Use the **show** form of this command to view BGP neighbor filter list configuration settings.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast maximum-prefix <max-num>

Sets the maximum number of prefixes to accept from this neighbor before that neighbor is taken down.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast maximum-prefix max-num
```

```
delete protocols bgp asn neighbor id address-family ipv6-unicast maximum-prefix
```

```
show protocols bgp asn neighbor id address-family ipv6-unicast maximum-prefix
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    neighbor id {  
      address-family {  
        ipv6-unicast {  
          maximum-prefix max-num  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>max-num</i>	Mandatory. The maximum number of prefixes to accept from this neighbor before the neighbor is taken down.

Default

No maximum to the number of prefixes accepted.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to set the maximum number of prefixes to accept from this neighbor before the neighbor is taken down. This helps alleviate situations where the router receives more routes than it can handle.

Use the **delete** form of this command to remove the maximum prefix setting.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast nexthop-local unchanged

Specifies that the IPv6 link-local address is not changed when the prefix is announced to the peer.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast nexthop-local unchanged
```

```
delete protocols bgp asn neighbor id address-family ipv6-unicast nexthop-local
```

```
show protocols bgp asn neighbor id address-family ipv6-unicast nexthop-local
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          nexthop-local {
            unchanged
          }
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Default

The IPv6 link-local address is changed when the prefix is announced to the peer.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to specify that the IPv6 link-local address is not changed when the prefix is announced to the peer.

Use the **delete** form of this command to return the system to its default behavior.

Use the **show** form of this command to view BGP neighbor **ipv6-unicast nexthop-local** configuration settings.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast nexthop-self

Sets the local router as the next hop for this neighbor.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast nexthop-self
delete protocols bgp asn neighbor id address-family ipv6-unicast nexthop-self
show protocols bgp asn neighbor id address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          nexthop-self
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Default

Disabled.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to set this router as the next hop for this neighbor. This disables the next-hop calculation for this neighbor.

Use the **delete** form of this command to restore next-hop calculation for the neighbor.

Use the **show** form of this command to view BGP neighbor **ipv6-unicast** configuration settings.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast peer-group <group-name>

Specifies the IPv6 peer group for this peer.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast peer-group group-name
```

```
delete protocols bgp asn neighbor id address-family ipv6-unicast peer-group group-name
```

```
show protocols bgp asn neighbor id address-family ipv6-unicast peer-group
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    neighbor id {  
      address-family {  
        ipv6-unicast {  
          peer-group group-name  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>group-name</i>	Mandatory. Name of an IPv6 peer group.

Default

None.

Usage Guidelines

This command specifies the IPv6 peer group for the peer.

Use the **set** form of this command to specify the IPv6 peer group for the peer.

Use the **delete** form of this command to remove the IPv6 peer group for the peer.

Use the **show** form of this command to view the IPv6 peer group for the peer.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast prefix-list export <prefix-list6-name>

Applies a prefix list to filter updates to this neighbor.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast prefix-list export
prefix-list6-name
```

```
delete protocols bgp asn neighbor id address-family ipv6-unicast prefix-list export
prefix-list6-name
```

```
show protocols bgp asn neighbor id address-family ipv6-unicast prefix-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          prefix-list {
            export prefix-list6-name
          }
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>prefix-list6-name</i>	Mandatory. Name of a configured IPv6 prefix list.

Default

None.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to restrict distribution of outbound BGP neighbor information by filtering with a prefix list.

Use the **delete** form of this command to remove an outbound prefix list filter.

Use the **show** form of this command to view BGP neighbor prefix list export configuration settings.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast prefix-list import <prefix-list6-name>

Applies a prefix list to filter updates from this neighbor.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast prefix-list import
prefix-list6-name
```

```
delete protocols bgp asn neighbor id address-family ipv6-unicast prefix-list import
prefix-list6-name
```

```
show protocols bgp asn neighbor id address-family ipv6-unicast prefix-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          prefix-list {
            import prefix-list6-name
          }
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>prefix-list6-name</i>	Mandatory. Name of a configured IPv6 prefix list.

Default

None.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to restrict distribution of inbound BGP neighbor information by filtering with a prefix list.

Use the **delete** form of this command to remove an inbound prefix list filter.

Use the **show** form of this command to view BGP neighbor prefix list import configuration settings.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast remove-private-as

Directs the router to remove private AS numbers from updates sent to this neighbor.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast remove-private-as
delete protocols bgp asn neighbor id address-family ipv6-unicast remove-private-as
show protocols bgp asn neighbor id address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          remove-private-as
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Default

Private AS numbers are included in outgoing updates.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to direct the router to exclude private autonomous system (AS) numbers from updates to eBGP peers. When this feature is enabled, the router omits private AS numbers from the AS_PATH attribute. The range of private AS numbers is 64512 to 65534.

Note that it is a configuration error to include both private and public AS numbers in an AS path. If the router detects this error, it does not remove private AS numbers.

This command may be used in confederations provided that the private AS numbers are appended after the confederation portion of the AS path.

This command applies only to eBGP peers; it cannot be used with iBGP peers.

Use the **delete** form of this command to restore the default behavior.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast route-map export <map-name>

Applies a route map to filter updates to this neighbor.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast route-map export map-name
```

```
delete protocols bgp asn neighbor id address-family ipv6-unicast route-map export map-name
```

```
show protocols bgp asn neighbor id address-family ipv6-unicast route-map export map-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    neighbor id {  
      address-family {  
        ipv6-unicast {  
          route-map {  
            export map-name  
          }  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>map-name</i>	Mandatory. Name of a configured route map.

Default

None.

Usage Guidelines

Use the **set** form of this command to restrict distribution of outbound BGP neighbor information by filtering with a route map.

Use the **delete** form of this command to remove an outbound route map filter.

Use the **show** form of this command to view BGP neighbor route map configuration settings.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast route-map import <map-name>

Applies a route map to filter updates to or from this neighbor.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast route-map import map-name
```

```
delete protocols bgp asn neighbor id address-family ipv6-unicast route-map import map-name
```

```
show protocols bgp asn neighbor id address-family ipv6-unicast route-map import map-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          route-map {
            import map-name
          }
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>map-name</i>	Mandatory. Name of a configured route map.

Default

None.

Usage Guidelines

Use the **set** form of this command to restrict routing updates received from this BGP neighbor by filtering with a route map.

Use the **delete** form of this command to remove a route map filter.

Use the **show** form of this command to view BGP neighbor route map configuration settings.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast soft-reconfiguration inbound

Directs the router to store received routing updates.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast soft-reconfiguration inbound
```

```
delete protocols bgp asn neighbor id address-family ipv6-unicast soft-reconfiguration inbound
```

```
show protocols bgp asn neighbor id address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          soft-reconfiguration {
            inbound
          }
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Default

Disabled.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to enable soft reconfiguration.

When this command is issued, the router begins storing routing updates, which can be used subsequently for inbound soft reconfiguration. Outbound BGP soft reconfiguration can be performed without enabling inbound soft reconfiguration.

Use the **delete** form of this command to disable soft reconfiguration.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast unsuppress-map <map-name>

Directs the router to selectively advertise routes suppressed by aggregating addresses, based on a route map.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast unsuppress-map map-name
```

```
delete protocols bgp asn neighbor id address-family ipv6-unicast unsuppress-map map-name
```

```
show protocols bgp asn neighbor id address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    neighbor id {  
      address-family {  
        ipv6-unicast {  
          unsuppress-map map-name  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>map-name</i>	Mandatory. The name of a configured route map.

Default

Routes suppressed by address aggregation are not advertised.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to selectively advertise routes suppressed by aggregating addresses.

Use the **delete** form of this command to restore the default behavior.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> advertisement-interval <seconds>

Sets the minimum interval for BGP route updates.

Syntax

```
set protocols bgp asn neighbor id advertisement-interval seconds  
delete protocols bgp asn neighbor id advertisement-interval  
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        neighbor id {  
            advertisement-interval seconds  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>seconds</i>	Mandatory. The minimum interval, in seconds, between BGP routing updates to this neighbor. The range is 0 to 600. The default is 30 for eBGP peers and 5 for iBGP peers.

Default

The default advertisement interval is 30 seconds for eBGP peers and 5 seconds for iBGP peers.

Usage Guidelines

Use the **set** form of this command to set the minimum interval between BGP routing advertisements to a BGP neighbor.

Use the **delete** form of this command to restore the advertisement interval to the default.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> allowas-in

Allows or disallows receiving BGP advertisements containing the AS path of the local router.

Syntax

```
set protocols bgp asn neighbor id allowas-in [number num]
```

```
delete protocols bgp asn neighbor id allowas-in
```

```
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      allowas-in {
        number num
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
number <i>num</i>	Optional. Specifies the number of times to accept an AS path containing the readvertisement of the local router's ASN. The range is 1 to 10. The default is 3.

Default

Readvertisement is disabled.

Usage Guidelines

Use the **set** form of this command to allow the router to accept BGP AS paths advertising the router's own AS number.

This situation could indicate a routing loop, and by default such an update is dropped. However, you can use this command to direct the router to accept updates readvertising its AS number for a specified number of times.

Use the **delete** form of this command to disallow readvertisement of the router's AS path.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> attribute-unchanged

Allows the router to send updates to a neighbor with unchanged attributes.

Syntax

```
set protocols bgp asn neighbor id attribute-unchanged [as-path | med | next-hop]
delete protocols bgp asn neighbor id attribute-unchanged [as-path | med | next-hop]
show protocols bgp asn neighbor id attribute-unchanged
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      attribute-unchanged {
        as-path
        med
        next-hop
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>as-path</i>	Optional. Propagates the route update with unchanged AS_PATH attribute.
<i>med</i>	Optional. Propagates the route update with unchanged Multi Exit Discriminator.

next-hop	Optional. Propagates the route update with unchanged next hop.
-----------------	--

Default

Disabled.

Usage Guidelines

Use the **set** form of this command to direct the router to propagate routing updates without modifying the BGP AS_PATH, Multi Exit Discriminator (MED), or next-hop attribute.

Use the **delete** form of this command to restore normal modification of BGP attributes.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> capability dynamic

Advertises support for dynamic update of BGP capabilities advertised and received from this neighbor.

Syntax

```
set protocols bgp asn neighbor id capability dynamic
delete protocols bgp asn neighbor id capability dynamic
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      capability {
        dynamic
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Default

By default, the session is brought up with minimal capability on both sides.

Usage Guidelines

In general, BGP-4 requires that if a BGP speaker receives an OPEN message with an unrecognized optional parameters, the speaker must terminate BGP peering. This makes it difficult to introduce new capabilities into the protocol. The Capabilities parameter allows graceful negotiation of BGP capabilities without requiring a peer to terminate peering.

This command specifies use of the BGP Dynamic Capability feature, which allows dynamic update of capabilities over an established BGP session.

Use the **set** form of this command to specify dynamic update of BGP capabilities to be used in capability negotiation.

Use the **delete** form of this command to remove the dynamic update capability.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> capability orf

Advertises support for Outbound Route Filtering (ORF) for updating BGP capabilities advertised and received from this neighbor.

Syntax

```
set protocols bgp asn neighbor id capability orf [prefix-list [receive | send]]
```

```
delete protocols bgp asn neighbor id capability orf
```

```
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      capability {
        orf {
          prefix-list {
            receive
            send
          }
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
prefix-list	Optional. Advertise prefix-list ORF capability to this neighbor.
receive	Optional. Advertise the ability to receive the ORF from this neighbor.

send	Optional. Advertise the capability to send the ORF to this neighbor.
-------------	--

Default

By default, the session is brought up with minimal capability on both sides.

Usage Guidelines

In general, BGP-4 requires that if a BGP speaker receives an OPEN message with an unrecognized optional parameters, the speaker must terminate BGP peering. This makes it difficult to introduce new capabilities into the protocol. The Capabilities parameter allows graceful negotiation of BGP capabilities without requiring a peer to terminate peering.

This command specifies use of BGP Outbound Route Filtering (ORF) to send and receive capabilities. Using ORF minimizes the number of BGP updates that are sent between peer routers.

Use the **set** form of this command to specify BGP ORF capabilities to be used in capability negotiation.

Use the **delete** form of this command to remove the ORF capability.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> default-originate

Enables or disables forwarding of the default route to a BGP neighbor.

Syntax

```
set protocols bgp asn neighbor id default-originate [route-map map-name]  
delete protocols bgp asn neighbor id default-originate [route-map map-name]  
show protocols bgp asn neighbor id default-originate
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    neighbor id {  
      default-originate {  
        route-map map-name  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>map-name</i>	Optional. Specifies a configured route map to be used when advertising the default route.

Default

The default route is not distributed.

Usage Guidelines

Use the **set** form of this command to enable the router to advertise the default route 0.0.0.0 to this BGP neighbor. This route can then be used by the neighbor to reach the local router if no other routes are available.

If a route map is specified, the default route is advertised if two conditions are satisfied:

- The route map includes a **match ip** address clause.
- A route exists that exactly matches the IP access list.

Route maps are configured using `protocols bgp <asn> neighbor <id> local-as <asn> command`.

Use the **delete** form of this command to disable forwarding of the default route or to delete a route map.

Use the **show** form of this command to view BGP neighbor default route origination settings.

protocols bgp <asn> neighbor <id> description <desc>

Provides a brief description for a BGP neighbor.

Syntax

```
set protocols bgp asn neighbor id description desc  
delete protocols bgp asn neighbor id description  
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        neighbor id {  
            description desc  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>desc</i>	Mandatory. A description of up to 80 characters for the neighbor. If spaces are included, the description must be enclosed in quotes.

Default

None.

Usage Guidelines

Use the set form of this command to enter a description for this BGP neighbor.

Use the **delete** form of this command to delete the neighbor's description.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> disable-capability-negotiation

Disables BGP capability negotiation.

Syntax

```
set protocols bgp asn neighbor id disable-capability-negotiation
delete protocols bgp asn neighbor id disable-capability-negotiation
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      disable-capability-negotiation
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Default

Capability negotiation is performed.

Usage Guidelines

Use the set form of this command to disable BGP capability negotiation.

Use the **delete** form of this command to delete this attribute and restore BGP capability negotiation.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> disable-connected-check

Disables direct connection verification for single-hop eBGP peers.

Syntax

```
set protocols bgp asn neighbor id disable-connected-check
delete protocols bgp asn neighbor id disable-connected-check
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      disable-connected-check
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Default

Connection verification is performed.

Usage Guidelines

Use the set form of this command to disable connection verification.

Connection verification is used to establish an eBGP peering session with a single-hop peer. A single-hop peer is defined using the **neighbor ebgp-multihop** command and specifying a time-to-live (TTL) value of 1. Such a peer is reachable by a single hop, but is configured on a loopback interface or is otherwise configured with a non-directly connected IP address.

For these peers, the BGP process normally checks to determine whether the eBGP peer is directly connected to the same network as the local router. If not, the peering session is not established.

If connection verification is disabled, source updating must be enabled using `protocols bgp <asn> neighbor <id> update-source <source> command` so that the BGP process can use the loopback interface for the peering session.

Use the **delete** form of this command to restore connection verification for eBGP peering sessions.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> disable-send-community

Disables sending of community attributes to the specified neighbor.

Syntax

```
set protocols bgp asn neighbor id disable-send-community [extended | standard]
delete protocols bgp asn neighbor id disable-send-community
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      disable-send-community {
        extended
        standard
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
extended	Optional. Disable sending extended community attributes.
standard	Optional. Disable sending standard community attributes.

Default

Communities are sent by default.

Usage Guidelines

Use this command to enable and disable sending community attributes between BGP neighbors.

Use the **set** form of this command to disable sending community attributes.

Use the **delete** form of this command to restore sending community attributes.

Use the **show** form of this command to view BGP neighbor configuration settings

protocols bgp <asn> neighbor <id> distribute-list export <acl-num>

Applies an access list to filter outbound routing updates to this neighbor.

Syntax

```
set protocols bgp asn neighbor id distribute-list export acl-num
delete protocols bgp asn neighbor id distribute-list
show protocols bgp asn neighbor id distribute-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      distribute-list {
        export acl-num
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>acl-num</i>	Optional. The number of a standard or extended access list. The range for a standard access list is 1 to 99. The range for an extended access list is 100 to 199.

Default

None.

Usage Guidelines

Use the **set** form of this command to apply an access list to filter outbound routing updates to a BGP neighbor.

Use the **delete** form of this command to disable outbound distribute list filtering.

Use the **show** form of this command to view BGP neighbor distribute list configuration settings.

NOTE *A neighbor distribute list cannot be used together with a neighbor prefix list in the same direction. These two lists are mutually exclusive, and only one list may be applied to a given direction.*

protocols bgp <asn> neighbor <id> distribute-list import <acl-num>

Applies an access list to filter inbound routing updates from this neighbor.

Syntax

```
set protocols bgp asn neighbor id distribute-list import acl-num
delete protocols bgp asn neighbor id distribute-list
show protocols bgp asn neighbor id distribute-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      distribute-list {
        import acl-num
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>acl-num</i>	Optional. The number of a standard or extended access list. The range for a standard access list is 1 to 99. The range for an extended access list is 100 to 199.

Default

None.

Usage Guidelines

Use the **set** form of this command to apply an access list to filter inbound routing updates from a BGP neighbor.

Use the **delete** form of this command to disable inbound distribute list filtering.

Use the **show** form of this command to view BGP neighbor distribute list configuration settings.

NOTE *A neighbor distribute list cannot be used together with a neighbor prefix list in the same direction. These two lists are mutually exclusive, and only one list may be applied to the specified direction.*

protocols bgp <asn> neighbor <id> ebgp-multihop <ttl>

Allows eBGP neighbors not on directly connected networks.

Syntax

```
set protocols bgp asn neighbor id ebgp-multihop ttl
delete protocols bgp asn neighbor id ebgp-multihop
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      ebgp-multihop ttl
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>ttl</i>	Mandatory. The time-to-live, or maximum number of hops allowed. The range is 1 to 255.

Default

Only directly connected neighbors are allowed.

Usage Guidelines

Use this command to configure eBGP neighbors that are not on directly connected networks.

For eBGP peers, there are three possible configurations:

- 1** By default, outgoing TTL is set to 1 and no checking is done on TTL of incoming packets.
- 2** If **ebgp-multihop** is configured, the TTL for outgoing packets is set to that value and no checking is done on TTL of incoming packets.
- 3** If **ttl-security hops** is configured, the TTL for outgoing packets is set to 255 and the TTL of incoming packets is compared against the value set for **ttl-security hops**.

Note that **ttl-security hops** can not be configured at the same time as **ebgp-multihop**.

Packets will always be dropped when TTL reaches zero.

Use the **set** form of this command to allow connections to eBGP peers residing on networks that are not directly connected.

Use the **delete** form of this command to restrict connections to directly connected peers.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> filter-list export <as-path-list-name>

Applies an AS path list to routing updates to this neighbor.

Syntax

```
set protocols bgp asn neighbor id filter-list export as-path-list-name
delete protocols bgp asn neighbor id filter-list export as-path-list-name
show protocols bgp asn neighbor id filter-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      filter-list {
        export as-path-list-name
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>as-path-list-name</i>	Mandatory. Name of an autonomous system path list.

Default

None.

Usage Guidelines

Use the **set** form of this command to apply an access list for filtering outbound routing updates.

Use the **delete** form of this command to disable outbound routing update filtering.

Use the **show** form of this command to view BGP neighbor filter list configuration settings.

protocols bgp <asn> neighbor <id> filter-list import <as-path-list-name>

Applies an AS path list to routing updates from this neighbor.

Syntax

```
set protocols bgp asn neighbor id filter-list import as-path-list-name
delete protocols bgp asn neighbor id filter-list import as-path-list-name
show protocols bgp asn neighbor id filter-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      filter-list {
        import as-path-list-name
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>as-path-list-name</i>	Mandatory. Name of an autonomous system path list.

Default

None.

Usage Guidelines

Use the **set** form of this command to apply an access list for filtering inbound routing updates.

Use the **delete** form of this command to disable inbound routing update filtering.

Use the **show** form of this command to view BGP neighbor filter list configuration settings.

protocols bgp <asn> neighbor <id> local-as <asn>

Defines a local autonomous system number for eBGP peerings.

Syntax

```
set protocols bgp asn neighbor id local-as asn [no-prepend]
delete protocols bgp asn neighbor id local-as asn [no-prepend]
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      local-as asn {
        no-prepend
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>asn</i>	Optional. Valid autonomous system number. This may not be the number of the autonomous system to which the neighbor belongs. The range is 1 to 4294967294.
no-prepend	Optional. Directs the router to not prepend the local autonomous system number to routes received from an external peer.

Default

None.

Usage Guidelines

Use the **set** form of this command to set a local autonomous system number for eBGP peer groupings. This number is used by all peers in the group for peering. It cannot be applied to individual peers in the group.

A local autonomous system number can only be applied to a true eBGP peer; it cannot be applied to peers in different subautonomous systems within a confederation.

The **no-prepend** keyword can cause routing loops and should be used with care. It should be used only to change the autonomous system number in a BGP network. After the network transition has completed, this setting should be deleted.

Use the **delete** form of this command to remove a local autonomous system number, or to remove the **no-prepend** keyword.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> maximum-prefix <max-num>

Sets the maximum number of prefixes to accept from this neighbor before that neighbor is taken down.

Syntax

```
set protocols bgp asn neighbor id maximum-prefix max-num
delete protocols bgp asn neighbor id maximum-prefix max-num
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      maximum-prefix max-num
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>max-num</i>	Mandatory. The maximum number of prefixes to accept from this neighbor before the neighbor is taken down.

Default

No maximum to the number of prefixes accepted.

Usage Guidelines

Use the **set** form of this command to set the maximum number of prefixes to accept from this neighbor before the neighbor is taken down. This helps alleviate situations where the router receives more routes than it can handle.

Use the **delete** form of this command to remove the maximum prefix setting.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> nexthop-self

Sets the local router as the next hop for this neighbor.

Syntax

```
set protocols bgp asn neighbor id nexthop-self
delete protocols bgp asn neighbor id nexthop-self
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      nexthop-self
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Default

Disabled.

Usage Guidelines

Use the **set** form of this command to set this router as the next hop for this neighbor. This disables the next-hop calculation for this neighbor.

Use the **delete** form of this command to restore next-hop calculation for the neighbor.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> override-capability

Overrides capability negotiation to allow a peering session to be established with a neighbor that does not support capability negotiation.

Syntax

```
set protocols bgp asn neighbor id override-capability
delete protocols bgp asn neighbor id override-capability
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      override-capability
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Default

A session cannot be established if the neighbor does not support capability negotiation.

Usage Guidelines

Use the set form of this command to enable a peering session to be established with a BGP neighbor that does not support capability negotiation.

Normally, if a BGP peer does not support capability negotiation, a peering session cannot be established and the connection is terminated. Setting this value overrides this process and allows the session to be established.

Use the **delete** form of this command to restore the default behavior.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> passive

Directs the router not to initiate connections with this neighbor.

Syntax

```
set protocols bgp asn neighbor id passive
delete protocols bgp asn neighbor id passive
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      passive
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Default

The router both accepts inbound connections and initiates outbound connections.

Usage Guidelines

Use the **set** form of this command to configure the local router such that accepts inbound connections from the neighbor, but does not initiate outbound connections by sending a BGP OPEN message.

Use the **delete** form of this command to restore the default behavior.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> password <pwd>

Defines a BGP MD5 password.

Syntax

```
set protocols bgp asn neighbor id password pwd
delete protocols bgp asn neighbor id password pwd
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      password pwd
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>pwd</i>	Mandatory. A password to be used to generate an MD5 digest.

Default

None.

Usage Guidelines

Use this command to specify a password-like string for generating an MD5 authentication string. This string is used for verification by the BGP neighbor. If the MD5 digest generated by the neighbor does not match the MD5 digest generated by the local router, any communication is silently discarded by the neighbor.

Use the **set** form of this command to specify or modify the BGP MD5 password.

Use the **delete** form of this command to remove the BGP MD5 password.

Use the **show** form of this command to view BGP MD5 password configuration.

protocols bgp <asn> neighbor <id> peer-group <group-name>

Assigns a neighbor as a member of the specified peer group.

Syntax

```
set protocols bgp asn neighbor id peer-group group-name
delete protocols bgp asn neighbor id peer-group group-name
show protocols bgp asn neighbor id peer-group
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      peer-group group-name
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>group-name</i>	Mandatory. A string identifying a BGP peer group.

Default

By default, peer group members inherit all configured peer group settings.

Usage Guidelines

Use the **set** form of this command to assign this BGP neighbor to a peer-group.

Use the **delete** form of this command to remove a neighbor from a peer group.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> port <port-num>

Specifies the port on which the neighbor is listening for BGP signals.

Syntax

```
set protocols bgp asn neighbor id port port-num
delete protocols bgp asn neighbor id port
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      port port-num
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>port-num</i>	Mandatory. The port on which the BGP neighbor will be listening for BGP messages. The range is 1 to 65535. The default is 179.

Default

By default, the router uses the well-known port for BGP, which is 179.

Usage Guidelines

Use the set form of this command to specify the port number to which BGP signals will be sent.

Use the **delete** form of this command to restore the port to the default.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> prefix-list export <list-name>

Applies a prefix list to filter updates to this neighbor.

Syntax

```
set protocols bgp asn neighbor id prefix-list export list-name
delete protocols bgp asn neighbor id prefix-list export list-name
show protocols bgp asn neighbor id prefix-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      prefix-list {
        export list-name
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>list-name</i>	Mandatory. Name of a configured prefix list.

Default

None.

Usage Guidelines

Use the **set** form of this command to restrict distribution of outbound BGP neighbor information by filtering with a prefix list.

Use the **delete** form of this command to remove an outbound prefix list filter.

Use the **show** form of this command to view BGP neighbor prefix list configuration settings.

protocols bgp <asn> neighbor <id> prefix-list import <list-name>

Applies a prefix list to filter updates from this neighbor.

Syntax

```
set protocols bgp asn neighbor id prefix-list import list-name
delete protocols bgp asn neighbor id prefix-list import list-name
show protocols bgp asn neighbor id prefix-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      prefix-list {
        import list-name
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>list-name</i>	Mandatory. Name of a configured prefix list.

Default

None.

Usage Guidelines

Use the **set** form of this command to restrict distribution of inbound BGP neighbor information by filtering with a prefix list.

Use the **delete** form of this command to remove an inbound prefix list filter.

Use the **show** form of this command to view BGP neighbor prefix list configuration settings.

protocols bgp <asn> neighbor <id> remote-as <asn>

Specifies the autonomous system number of the neighbor.

Syntax

```
set protocols bgp asn neighbor id remote-as asn
delete protocols bgp asn neighbor id remote-as
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      remote-as asn
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Default

None.

Usage Guidelines

Use the set form of this command to specify the autonomous system (AS) number of a BGP neighbor.

If the AS number of the neighbor is the same as that of the local router, the neighbor is an internal BGP (iBGP) peer. If it is different the neighbor is an external BGP (eBGP) peer.

Use the **delete** form of this command to remove AS number settings for the neighbor.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> remove-private-as

Directs the router to remove private AS numbers from updates sent to this neighbor.

Syntax

```
set protocols bgp asn neighbor id remove-private-as
delete protocols bgp asn neighbor id remove-private-as
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      remove-private-as
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Default

Private AS numbers are included in outgoing updates.

Usage Guidelines

Use the set form of this command to direct the router to exclude private autonomous system (AS) numbers from updates to eBGP peers. When this feature is enabled, the router omits private AS numbers from the AS_PATH attribute. The range of private AS numbers is 64512 to 65534.

Note that it is a configuration error to include both private and public AS numbers in an AS path. If the router detects this error, it does not remove private AS numbers.

This command may be used in confederations provided that the private AS numbers are appended after the confederation portion of the AS path.

This command applies only to eBGP peers; it cannot be used with iBGP peers.

Use the **delete** form of this command to restore the default behavior.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> route-map export <map-name>

Applies a route map to filter updates to this neighbor.

Syntax

```
set protocols bgp asn neighbor id route-map export map-name
delete protocols bgp asn neighbor id route-map export map-name
show protocols bgp asn neighbor id route-map export map-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      route-map {
        export map-name
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>map-name</i>	Mandatory. Name of a configured route map.

Default

None.

Usage Guidelines

Use the **set** form of this command to restrict distribution of outbound BGP neighbor information by filtering with a route map.

Use the **delete** form of this command to remove an outbound route map filter.

Use the **show** form of this command to view BGP neighbor route map configuration settings.

protocols bgp <asn> neighbor <id> route-map import <map-name>

Applies a route map to filter updates to or from this neighbor.

Syntax

```
set protocols bgp asn neighbor id route-map import map-name
delete protocols bgp asn neighbor id route-map import map-name
show protocols bgp asn neighbor id route-map import map-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      route-map {
        import map-name
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>map-name</i>	Mandatory. Name of a configured route map.

Default

None.

Usage Guidelines

Use the **set** form of this command to restrict routing updates received from this BGP neighbor by filtering with a route map.

Use the **delete** form of this command to remove a route map filter.

Use the **show** form of this command to view BGP neighbor route map configuration settings.

protocols bgp <asn> neighbor <id> shutdown

Administratively shuts down a BGP neighbor.

Syntax

```
set protocols bgp asn neighbor id shutdown
delete protocols bgp asn neighbor id shutdown
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      shutdown
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Default

Disabled.

Usage Guidelines

Use the set form of this command to administratively shut down a BGP neighbor. Shutting down a BGP speaker administratively terminates any active sessions for the neighbor and removes any associated routing information. If the neighbor is a peer group, this could terminate a large number of sessions.

To determine whether a BGP neighbor has been administratively shut down, you can view BGP summary information. A neighbor with a status of Idle with a comment of Admin have been administratively shut down.

Use the **delete** form of this command to administratively reenable a BGP neighbor.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> soft-reconfiguration inbound

Directs the router to store received routing updates.

Syntax

```
set protocols bgp asn neighbor id soft-reconfiguration inbound
delete protocols bgp asn neighbor id soft-reconfiguration inbound
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      soft-reconfiguration {
        inbound
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Default

Disabled.

Usage Guidelines

Use the set form of this command to enable soft reconfiguration.

When this command is issued, the router begins storing routing updates, which can be used subsequently for inbound soft reconfiguration. Outbound BGP soft reconfiguration can be performed without enabling inbound soft reconfiguration.

Use the **delete** form of this command to disable soft reconfiguration.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> strict-capability-match

Directs the router to strictly match the capabilities of the neighbor.

Syntax

```
set protocols bgp asn neighbor id strict-capability-match
delete protocols bgp asn neighbor id strict-capability-match
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      strict-capability-match
    }
  }
}
```

Parameters

<i>asn</i>	<i>Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.</i>
<i>id</i>	<i>Mandatory. The IPv4 or IPv6 address of a BGP neighbor.</i>

Default

Disabled.

Usage Guidelines

Use the set form of this command to direct the router strictly compare the capabilities of the local router and the remote peer. If the capabilities are different, the session is terminated.

Use the **delete** form of this command to disable strict capability matching.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> timers

Sets BGP timers for this neighbor.

Syntax

```
set protocols bgp asn neighbor id timers [connect seconds | keepalive seconds | holdtime seconds]
```

```
delete protocols bgp asn neighbor id timers [connect | keepalive | holdtime]
```

```
show protocols bgp asn neighbor id timers
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      timers {
        connect seconds
        keepalive seconds
        holdtime seconds
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

connect <i>seconds</i>	Optional. The amount of time, in seconds, that the system waits for the transport protocol connection to complete. If this timer expires, the state remains in Connect state, the timer is reset, and the system tries to initiate a new transport connection. The range is 0 to 65535. The default is 120. During the connect period, the remote BGP peer can establish a connection to the local system.
keepalive <i>seconds</i>	Optional. The frequency, in seconds, with which the local router sends keep-alive messages to this neighbor. The range is 1 to 65535. The default is 60.
holdtime <i>seconds</i>	Optional. The maximum interval, in seconds, after which if the local router has not received a keep-alive message from this neighbor, the neighbor is declared dead. The range is 0 and 4 to 65535, where 0 disables the holdtime timer. The default is 180.

Default

The default for the connect timer is 120. The default for the keep-alive timer is 60 seconds. The default for the holdtime timer is 180 seconds.

Usage Guidelines

Use the **set** form of this command to set timers for monitoring the health of the remote peer.

- If this command is issued and values are specified, the values configured here override global timers set for the local router.
- If this command is issued but values are not specified, the defaults for this command apply.
- If this command is not issued, or if the configuration statement is deleted, timers set globally for the router using [protocols bgp <asn> timers command](#) apply to this neighbor.

Use the **delete** form of this command to remove explicitly configured timers for a neighbor. In this case, timers set globally for the router apply to this neighbor.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> ttl-security hops <hops>

Sets Time to Live (TTL) security hop count.

Syntax

```
set protocols bgp asn neighbor id ttl-security hops hops
delete protocols bgp asn neighbor id ttl-security hops
show protocols bgp asn neighbor id ttl-security hops
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      ttl-security {
        hops hops
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>hops</i>	Optional. The maximum number of hops away from the local peer a peering session connection will be accepted. The range is 1 to 254.

Default

The default for the hop count is 1 for eBGP connections.

Usage Guidelines

Use this command to set the Time to Live (TTL) security for the local peer.

This option provides an additional level of security for eBGP peering sessions. It helps prevent BGP hijacking attempts by determining whether or not the BGP peer is on a directly connected network using the BGP packet TTL value.

For eBGP peers, there are three possible configurations:

- 1** By default, outgoing TTL is set to 1 and no checking is done on TTL of incoming packets.
- 2** If **ebgp-multihop** is configured, the TTL for outgoing packets is set to that value and no checking is done on TTL of incoming packets.
- 3** If **ttl-security hops** is configured, the TTL for outgoing packets is set to 255 and the TTL of incoming packets is compared against the value set for **ttl-security hops**.

Note that **ttl-security hops** can not be configured at the same time as **ebgp-multihop**.

Packets will always be dropped when TTL reaches zero.

Use the **set** form of this command to specify the hop count.

Use the **delete** form of this command to return the system to its default.

Use the **show** form of this command to view the parameter configuration.

protocols bgp <asn> neighbor <id> unsuppress-map <map-name>

Directs the router to selectively advertise routes suppressed by aggregating addresses, based on a route map.

Syntax

```
set protocols bgp asn neighbor id unsuppress-map map-name
delete protocols bgp asn neighbor id unsuppress-map map-name
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      unsuppress-map map-name
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>map-name</i>	Mandatory. The name of a configured route map.

Default

Routes suppressed by address aggregation are not advertised.

Usage Guidelines

Use the **set** form of this command to selectively advertise routes suppressed by aggregating addresses.

Use the **delete** form of this command to restore the default behavior.

Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> neighbor <id> update-source <source>

Specifies the source IP address or interface for routing updates.

Syntax

```
set protocols bgp asn neighbor id update-source source
delete protocols bgp asn neighbor id update-source
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      update-source source
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>source</i>	Mandatory. The IPv4 address of the router to receive routing updates from or the interface to receive routing updates from..

Default

None.

Usage Guidelines

Use the **set** form of this command to direct the system to use a specific source for routing updates.

Use the **delete** form of this command to remove the routing update source.

Use the **show** form of this command to view the configuration settings.

protocols bgp <asn> neighbor <id> weight <weight>

Defines a default weight for routes from this neighbor.

Syntax

```
set protocols bgp asn neighbor id weight weight
delete protocols bgp asn neighbor id weight
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      weight weight
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
<i>weight</i>	Mandatory. The weight to be assigned to routes from this neighbor. The range is 0 to 65535.

Default

Routes learned from a BGP neighbor have a weight of 0. Routes sourced by the local router have a weight of 32768.

Usage Guidelines

Use the **set** form of this command to set the default weights for routes learned from a BGP neighbor.

Use the **delete** form of this command to restore route weighting to the default. Existing route weights are not changed.

Use the **show** form of this command to view BGP neighbor configuration settings.

reset ip bgp external

Resets peering sessions for eBGP neighbors.

Syntax

```
reset ip bgp external [in [prefix-filter] | out | soft [in | out]]
```

Command Mode

Operational mode.

Parameters

in	Optional. Resets inbound sessions only.
out	Optional. Resets outbound sessions only.
prefix-filter	Optional. Resets the BGP outbound route filter (ORF). The prefix-filter keyword is ignored unless ORF capabilities have been enabled on the local system or received from the sending BGP peer. In this case, a normal inbound soft reset is performed.
soft	Optional. Sessions are not reset. Resend updates that have changed due to an export policy change or recalculate any updates received based on import policy changes.
in	Optional. Re-read import policies and recalculates information in the BGP table based on import policy changes. This requires soft-reconfiguration inbound to be configured on the neighbor.
out	Optional. Re-read export policies and resend updates that have changed due to an export policy change.

Default

When used without the **soft** option, reset connections are dropped, both inbound and outbound.

Usage Guidelines

Use this command on a router running BGP to reset sessions for external BGP (eBGP) neighbors. This forces BGP updates to be generated and new BGP policies to be applied.

Unless the **soft** option is used, all connections are dropped (a “hard reset”): TCP connections are terminated and all routes received from the neighbor are removed from the BGP routing table. Then the connection with the neighbor is reestablished.

If the **soft** option is used, sessions are not reset. Changes to export policies are resent to peers. Changes to import policies causes recalculation of information in the BGP table (given that **soft-reconfiguration inbound** is configured on the peer).

reset ip bgp external ipv4 unicast

Resets IPv4 unicast peering sessions for eBGP neighbors.

Syntax

```
reset ip bgp external ipv4 unicast [in [prefix-filter] | out | soft [in | out]]
```

Command Mode

Operational mode.

Parameters

in	Optional. Resets inbound sessions only.
prefix-filter	Optional. Reset the BGP outbound route filter (ORF). The prefix-filter keyword is ignored unless ORF capabilities have been enabled on the local system or received from the sending BGP peer. In this case, a normal inbound soft reset is performed.
out	Optional. Resets outbound sessions only.
soft	Optional. Sessions are not reset. Resend updates that have changed due to an export policy change or recalculate any updates received based on import policy changes.
in	Optional. Re-read import policies and recalculates information in the BGP table based on import policy changes. This requires soft-reconfiguration inbound to be configured on the neighbor.
out	Optional. Re-read export policies and resend updates that have changed due to an export policy change.

Default

When used without the **soft** option, reset connections are dropped, both inbound and outbound.

Usage Guidelines

Use this command to reset IPv4 unicast peering sessions for eBGP neighbors. This forces BGP updates to be generated and new BGP policies to be applied.

Unless the **soft** option is used, all connections are dropped (a “hard reset”): TCP connections are terminated and all routes received from the neighbor are removed from the BGP routing table. Then the connection with the neighbor is reestablished.

If the **soft** option is used, sessions are not reset. Changes to export policies are resent to peers. Changes to import policies causes recalculation of information in the BGP table (given that **soft-reconfiguration inbound** is configured on the peer).

show ip bgp ipv4 unicast neighbors

Displays IPv4 unicast route information for BGP neighbors.

Syntax

```
show ip bgp ipv4 unicast neighbors [ipv4|ipv6]
```

Command Mode

Operational mode.

Parameters

<i>ipv4</i>	Displays detailed IPv4 unicast route information for the specified IPv4 BGP neighbor.
<i>ipv6</i>	Displays detailed IPv4 unicast route information for the specified IPv6 BGP neighbor.

Default

IPv4 unicast route information is shown for all BGP neighbors.

Usage Guidelines

Use this command to display IPv4 unicast route information for BGP neighbors.

show ip bgp ipv4 unicast neighbors <id> advertised-routes

Displays advertised BGP IPv4 unicast routes for a BGP neighbor.

Syntax

```
show ip bgp ipv4 unicast neighbors id advertised-routes
```

Command Mode

Operational mode.

Parameters

<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
-----------	--

Default

None.

Usage Guidelines

Use this command to display advertised BGP IPv4 unicast routes for a BGP neighbor.

show ip bgp ipv4 unicast neighbors <id> prefix-counts

Displays IPv4 unicast prefix-counts for a BGP neighbor.

Syntax

```
show ip bgp ipv4 unicast neighbors id prefix-counts
```

Command Mode

Operational mode.

Parameters

<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
-----------	--

Default

None.

Usage Guidelines

Use this command to display IPv4 unicast prefix counts for a BGP neighbor.

show ip bgp ipv4 unicast neighbors <id> received prefix-filter

Displays the IPv4 unicast prefix-lists received from a BGP neighbor.

Syntax

```
show ip bgp ipv4 unicast neighbors id } received prefix-filter
```

Command Mode

Operational mode.

Parameters

<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
-----------	--

Default

None.

Usage Guidelines

Use this command to display BGP IPv4 unicast prefix-lists received from a BGP neighbor.

show ip bgp ipv4 unicast neighbors <id> received-routes

Displays the IPv4 unicast routes received from a BGP neighbor.

Syntax

```
show ip bgp ipv4 unicast neighbors id received-routes
```

Command Mode

Operational mode.

Parameters

<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
-----------	--

Default

None.

Usage Guidelines

Use this command to display IPv4 unicast routes (both accepted and rejected) received from a BGP neighbor.

show ip bgp ipv4 unicast neighbors <id> routes

Displays IPv4 unicast received and accepted routes from a BGP neighbor.

Syntax

```
show ip bgp ipv4 unicast neighbors id routes
```

Command Mode

Operational mode.

Parameters

<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.
-----------	--

Default

None.

Usage Guidelines

Use this command to display IPv4 unicast received and accepted routes from a BGP neighbor.

show ip bgp neighbors

Displays BGP neighbor information.

Syntax

```
show ip bgp neighbors [ipv4|ipv6]
```

Command Mode

Operational mode.

Parameters

<i>ipv4</i>	Displays detailed information for the specified IPv4 BGP neighbor.
<i>ipv6</i>	Displays detailed information for the specified IPv6 BGP neighbor.

Default

Information for all neighbors is displayed.

Usage Guidelines

Use this command to display BGP neighbor information.

show ip bgp neighbors <id> advertised-routes

Displays advertised routes for a BGP neighbor.

Syntax

```
show ip bgp neighbors id advertised-routes
```

Command Mode

Operational mode.

Parameters

<i>id</i>	The IPv4 or IPv6 address of a BGP neighbor.
-----------	---

Default

None.

Usage Guidelines

Use this command to display advertised routes for a BGP neighbor.

show ip bgp neighbors <id> dampened-routes

Displays dampened routes to a BGP neighbor.

Syntax

```
show ip bgp neighbors id dampened-routes
```

Command Mode

Operational mode.

Parameters

<i>id</i>	The IPv4 or IPv6 address of a BGP neighbor.
-----------	---

Default

None.

Usage Guidelines

Use this command to display routes that have been dampened (suppressed) to a BGP neighbor due to route flapping.

show ip bgp neighbors <id> flap-statistics

Displays route flap statistics for routes learned from a BGP neighbor.

Syntax

```
show ip bgp neighbors id flap-statistics
```

Command Mode

Operational mode.

Parameters

<i>id</i>	The IPv4 or IPv6 address of a BGP neighbor.
-----------	---

Default

None.

Usage Guidelines

Use this command to display route flap statistics for routes learned from a BGP neighbor.

show ip bgp neighbors <id> prefix-counts

Displays prefix counts for a BGP neighbor.

Syntax

```
show ip bgp neighbors id prefix-counts
```

Command Mode

Operational mode.

Parameters

<i>id</i>	The IPv4 or IPv6 address of a BGP neighbor.
-----------	---

Default

None.

Usage Guidelines

Use this command to display prefix counts for a BGP neighbor.

show ip bgp neighbors <id> received prefix-filter

Displays prefix lists received from a BGP neighbor.

Syntax

```
show ip bgp neighbors id received prefix-filter
```

Command Mode

Operational mode.

Parameters

<i>id</i>	The IPv4 or IPv6 address of a BGP neighbor.
-----------	---

Default

None.

Usage Guidelines

Use this command to display prefix lists received from a BGP neighbor.

show ip bgp neighbors <id> received-routes

Displays routes received from a BGP neighbor.

Syntax

```
show ip bgp neighbors id received-routes
```

Command Mode

Operational mode.

Parameters

<i>id</i>	The IPv4 or IPv6 address of a BGP neighbor.
-----------	---

Default

None.

Usage Guidelines

Use this command to display routes (both accepted and rejected) received from a BGP neighbor.

show ip bgp neighbors <id> routes

Displays all received and accepted routes from a BGP neighbor.

Syntax

```
show ip bgp neighbors id routes
```

Command Mode

Operational mode.

Parameters

<i>id</i>	The IPv4 or IPv6 address of a BGP neighbor.
-----------	---

Default

None.

Usage Guidelines

Use this command to display received and accepted routes from a BGP neighbor.

show ipv6 bgp neighbors

Displays BGP neighbor information.

Syntax

```
show ipv6 bgp neighbors [ipv6]
```

Command Mode

Operational mode.

Parameters

<i>ipv6</i>	Optional. The IPv6 address of a BGP neighbor.
-------------	---

Default

Information is shown for all BGP neighbors.

Usage Guidelines

Use this command to display BGP neighbor information.

show ipv6 bgp neighbors <ipv6> advertised-routes

Displays advertised routes for a BGP neighbor.

Syntax

```
show ipv6 bgp neighbors ipv6 advertised-routes
```

Command Mode

Operational mode.

Parameters

<i>ipv6</i>	Mandatory. The IPv6 address of a BGP neighbor.
-------------	--

Default

None.

Usage Guidelines

Use this command to display advertised routes for a BGP neighbor.

show ipv6 bgp neighbors <ipv6> received-routes

Displays routes received from a BGP neighbor.

Syntax

```
show ipv6 bgp neighbors ipv6 received-routes
```

Command Mode

Operational mode.

Parameters

<i>ipv6</i>	Mandatory. The IPv6 address of a BGP neighbor.
-------------	--

Default

None.

Usage Guidelines

Use this command to display routes (both accepted and rejected) received from a BGP neighbor.

show ipv6 bgp neighbors <ipv6> routes

Displays all received and accepted routes from a BGP neighbor.

Syntax

```
show ipv6 bgp neighbors ipv6 routes
```

Command Mode

Operational mode.

Parameters

<i>ipv6</i>	Mandatory. The IPv6 address of a BGP neighbor.
-------------	--

Default

None.

Usage Guidelines

Use this command to display received and accepted routes from a BGP neighbor.

Chapter 6: Peer Groups

This chapter describes commands for BGP peer groups.

This chapter presents the following topics:

- [Peer Group Commands](#)

Peer Group Commands

This chapter contains the following commands.

Configuration Commands	
<code>protocols bgp <asn> peer-group <group-name></code>	Defines a BGP peer group.
<code>protocols bgp <asn> peer-group <group-name> allowas-in</code>	Allows or disallows receiving BGP advertisements containing the AS path of the local router.
<code>protocols bgp <asn> peer-group <group-name> attribute-unchanged</code>	Allows the router to send updates to a peer group with unchanged attributes.
<code>protocols bgp <asn> peer-group <group-name> capability dynamic</code>	Advertises support for dynamic update of BGP capabilities advertised and received from this peer group.
<code>protocols bgp <asn> peer-group <group-name> capability orf</code>	Enables or disables forwarding of the default route to a BGP peer group.
<code>protocols bgp <asn> peer-group <group-name> default-originate</code>	Provides a brief description for a BGP peer group.
<code>protocols bgp <asn> peer-group <group-name> description <desc></code>	Provides a brief description for a BGP peer group.
<code>protocols bgp <asn> peer-group <group-name> disable-capability-negotiation</code>	Disables BGP capability negotiation.
<code>protocols bgp <asn> peer-group <group-name> disable-connected-check</code>	Disables direct connection verification for single-hop eBGP peers.
<code>protocols bgp <asn> peer-group <group-name> disable-send-community</code>	Disables sending of community attributes to the specified peer group.
<code>protocols bgp <asn> peer-group <group-name> distribute-list export <acl-num></code>	Applies an access list to filter outbound routing updates to this peer group.
<code>protocols bgp <asn> peer-group <group-name> distribute-list import <acl-num></code>	Applies an access list to filter inbound routing updates from this peer group.
<code>protocols bgp <asn> peer-group <group-name> ebgp-multihop <tth></code>	Allows eBGP neighbors not on directly connected networks.
<code>protocols bgp <asn> peer-group <group-name> filter-list export <as-path-list-name></code>	Applies an AS path list to routing updates to this peer group.
<code>protocols bgp <asn> peer-group <group-name> filter-list import <as-path-list-name></code>	Applies an AS path list to routing updates from this peer group.

<code>protocols bgp <asn> peer-group <group-name> local-as <asn></code>	Defines a local autonomous system number for eBGP peerings.
<code>protocols bgp <asn> peer-group <group-name> maximum-prefix <max-num></code>	Sets the maximum number of prefixes to accept from this peer group before that peer group is taken down.
<code>protocols bgp <asn> peer-group <group-name> nexthop-self</code>	Sets the local router as the next hop for this peer group.
<code>protocols bgp <asn> peer-group <group-name> override-capability</code>	Overrides capability negotiation to allow a peering session to be established with a peer group that does not support capability negotiation.
<code>protocols bgp <asn> peer-group <group-name> passive</code>	Directs the router not to initiate connections with this peer group.
<code>protocols bgp <asn> peer-group <group-name> password <pwd></code>	Defines a BGP MD5 password.
<code>protocols bgp <asn> peer-group <group-name> prefix-list export <list-name></code>	Applies a prefix list to filter updates to this peer group.
<code>protocols bgp <asn> peer-group <group-name> prefix-list import <list-name></code>	Applies a prefix list to filter updates from this peer group.
<code>protocols bgp <asn> peer-group <group-name> remote-as <asn></code>	Specifies the autonomous system number of the peer group.
<code>protocols bgp <asn> peer-group <group-name> remove-private-as</code>	Directs the router to remove private AS numbers from updates sent to this peer group.
<code>protocols bgp <asn> peer-group <group-name> route-map export <map-name></code>	Applies a route map to filter updates to this peer group.
<code>protocols bgp <asn> peer-group <group-name> route-map import <map-name></code>	Applies a route map to filter updates to or from this peer group.
<code>protocols bgp <asn> peer-group <group-name> shutdown</code>	Administratively shuts down a BGP peer group.
<code>protocols bgp <asn> peer-group <group-name> soft-reconfiguration inbound</code>	Directs the router to store received routing updates.
<code>protocols bgp <asn> peer-group <group-name> ttl-security hops <hops></code>	Sets Time to Live (TTL) security hop count for this peer group.
<code>protocols bgp <asn> peer-group <group-name> unsuppress-map <map-name></code>	Directs the router to selectively advertise routes suppressed by aggregating addresses, based on a route map.
<code>protocols bgp <asn> peer-group <group-name> update-source <source></code>	Specifies the source IP address or interface for routing updates.

<code>protocols bgp <asn> peer-group <group-name> weight <weight></code>	Defines a default weight for routes from this peer group.
Peer Group Commands - IPv6 Routes	
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast</code>	Defines parameters for controlling IPv6 unicast routes for a peering session.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast allowas-in</code>	Allows or disallows receiving BGP advertisements containing the AS path of the local router.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast attribute-unchanged</code>	Allows the router to send updates to a peer group with unchanged attributes.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast capability dynamic</code>	Advertises support for dynamic update of BGP capabilities advertised and received from this peer group.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast capability orf</code>	Advertises support for Outbound Route Filtering (ORF) for updating BGP capabilities advertised and received from this peer group.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast default-originate</code>	Enables or disables forwarding of the default route to a BGP peer group.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast disable-send-community</code>	Disables sending of community attributes to the specified peer group.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast distribute-list export <access-list6-name></code>	Applies an access list to filter outbound routing updates to this peer group.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast distribute-list import <access-list6-name></code>	Applies an access list to filter inbound routing updates from this peer group.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast filter-list export <as-path-list-name></code>	Applies an AS path list to routing updates to this peer group.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast filter-list import <as-path-list-name></code>	Applies an AS path list to routing updates from this peer group.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast maximum-prefix <max-num></code>	Sets the maximum number of prefixes to accept from this peer group before that peer group is taken down.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast nexthop-local unchanged</code>	Specifies that the IPv6 link-local address is not changed when the prefix is announced to the peer.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast nexthop-self</code>	Sets the local router as the next hop for this peer group.

<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast prefix-list export <prefix-list6-name></code>	Applies a prefix list to filter updates to this peer group.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast prefix-list import <prefix-list6-name></code>	Applies a prefix list to filter updates from this peer group.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast remove-private-as</code>	Directs the router to remove private AS numbers from updates sent to this peer group.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast route-map export <map-name></code>	Applies a route map to filter updates to this peer group.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast route-map import <map-name></code>	Applies a route map to filter updates to or from this peer group.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast soft-reconfiguration inbound</code>	Directs the router to store received routing updates.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast unsuppress-map <map-name></code>	Directs the router to selectively advertise routes suppressed by aggregating addresses, based on a route map.
Operational Commands	
<code>reset ip bgp peer-group <group-name></code>	Resets sessions for all members of a peer group.
<code>reset ip bgp peer-group <group-name> ipv4 unicast</code>	Resets IPv4 unicast sessions for all members of a peer group.

protocols bgp <asn> peer-group <group-name>

Defines a BGP peer group.

Syntax

```
set protocols bgp asn peer-group group-name
delete protocols bgp asn peer-group group-name
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {}
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. Multi-node. The peer group name.

Default

None.

Usage Guidelines

Configuring a peer group simplifies configuration for neighbors with similar update policies. Once a peer group is created, it can be configured using the same commands for configuring individual BGP neighbors. Each member of the peer group inherits the peer group's update policies unless overridden by explicit configuration of the individual peer. This also makes update calculation more efficient.

All members of a peer group must share identical outbound routing policies—that is, they must have identified distribute lists, filter lists, prefix lists, and route maps applied. They need not have identical settings for **default-originate**, as this is always processed on a per-peer basis. Inbound update routing policies may vary per peer group member.

Use the **set** form of this command to define a BGP peer group.

Use the **delete** form of this command to remove a BGP peer group.

Use the **show** form of this command to view BGP peer group origination settings.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast

Defines parameters for controlling IPv6 unicast routes for a peering session.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast
delete protocols bgp asn peer-group group-name address-family ipv6-unicast
show protocols bgp asn peer-group group-name address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      address-family {
        ipv6-unicast {}
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number of the AS in which this router resides.
<i>group-name</i>	Mandatory. The peer group name.

Usage Guidelines

Use this command to define parameters for IPv6 unicast routes for a peering session.

Use the **set** form of this command to specify the configuration node.

Use the **delete** form of this command to delete the **ipv6-unicast** configuration.

Use the **show** form of this command to view the **ipv6-unicast** configuration.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast allowas-in

Allows or disallows receiving BGP advertisements containing the AS path of the local router.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast allowas-in  
[number num]
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
allowas-in
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      address-family {  
        ipv6-unicast {  
          allowas-in {  
            number num  
          }  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

number <i>num</i>	Optional. Specifies the number of times to accept an AS path containing the readvertisement of the local router's ASN. The range is 1 to 10. The default is 3.
-----------------------------	--

Default

Readvertisement is disabled.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to allow the router to accept BGP AS paths advertising the router's own AS number.

This situation could indicate a routing loop, and by default such an update is dropped. However, you can use this command to direct the router to accept updates readvertising its AS number for a specified number of times.

Use the **delete** form of this command to disallow readvertisement of the router's AS path.

Use the **show** form of this command to view configuration settings.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast attribute-unchanged

Allows the router to send updates to a peer group with unchanged attributes.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast  
attribute-unchanged [as-path | med | next-hop]
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
attribute-unchanged [as-path | med | next-hop]
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast  
attribute-unchanged
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      address-family {  
        ipv6-unicast {  
          attribute-unchanged {  
            as-path  
            med  
            next-hop  
          }  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
------------	--

<i>group-name</i>	Mandatory. The peer group name.
as-path	Optional. Propagates the route update with unchanged AS_PATH attribute.
med	Optional. Propagates the route update with unchanged Multi Exit Discriminator.
next-hop	Optional. Propagates the route update with unchanged next hop.

Default

Disabled.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to direct the router to propagate routing updates without modifying the BGP AS_PATH, Multi Exit Discriminator (MED), or next-hop attribute.

Use the **delete** form of this command to restore normal modification of BGP attributes.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast capability dynamic

Advertises support for dynamic update of BGP capabilities advertised and received from this peer group.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast capability dynamic
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast capability dynamic
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        peer-group group-name {  
            address-family {  
                ipv6-unicast {  
                    capability {  
                        dynamic  
                    }  
                }  
            }  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

Default

By default, the session is brought up with minimal capability on both sides.

Usage Guidelines

In general, BGP requires that if a BGP speaker receives an OPEN message with any unrecognized optional parameters, the speaker must terminate BGP peering. This makes it difficult to introduce new capabilities into the protocol. The Capabilities parameter allows graceful negotiation of BGP capabilities without requiring a peer to terminate peering.

This command specifies use of the BGP Dynamic Capability feature, which allows dynamic update of capabilities over an established BGP session.

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to specify dynamic update of BGP capabilities to be used in capability negotiation.

Use the **delete** form of this command to remove the dynamic update capability.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast capability orf

Advertises support for Outbound Route Filtering (ORF) for updating BGP capabilities advertised and received from this peer group.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast capability  
orf [prefix-list [receive | send]]
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
capability orf
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      address-family {  
        ipv6-unicast {  
          capability {  
            orf {  
              prefix-list {  
                receive  
                send  
              }  
            }  
          }  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
prefix-list	Optional. Advertise prefix-list ORF capability to this peer group.
receive	Optional. Advertise the ability to receive the ORF from this peer group.
send	Optional. Advertise the capability to send the ORF to this peer group.

Default

By default, the session is brought up with minimal capability on both sides.

Usage Guidelines

This command specifies use of BGP Outbound Route Filtering (ORF) to send and receive capabilities. Using ORF minimizes the number of BGP updates that are sent between peer routers.

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to specify BGP ORF capabilities to be used in capability negotiation.

Use the **delete** form of this command to remove the ORF capability.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast default-originate

Enables or disables forwarding of the default route to a BGP peer group.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast  
default-originate [route-map map-name]
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
default-originate [route-map map-name]
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast  
default-originate
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      address-family {  
        ipv6-unicast {  
          default-originate {  
            route-map map-name  
          }  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

<i>map-name</i>	Optional. Specifies a configured route map to be used when advertising the default route.
-----------------	---

Default

The default route is not distributed.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to enable the router to advertise the default route `::/0` to this BGP peer group. This route can then be used by the peer group to reach the local router if no other routes are available.

The route `::/0` need not be explicitly configured on the local router.

If a route map is specified, the default route is advertised if two conditions are satisfied:

- The route map includes a **match ip** address clause.
- A route exists that exactly matches the IP access list.

Route maps are configured using [protocols bgp <asn> peer-group <group-name> local-as <asn> command](#).

Use the **delete** form of this command to disable forwarding of the default route or to delete a route map.

Use the **show** form of this command to view BGP peer group default route origination settings.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast disable-send-community

Disables sending of community attributes to the specified peer group.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast  
disable-send-community [extended | standard]
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
disable-send-community
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      address-family {  
        ipv6-unicast {  
          disable-send-community {  
            extended  
            standard  
          }  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>extended</i>	Optional. Disable sending extended community attributes.

standard	Optional. Disable sending standard community attributes.
-----------------	--

Default

Communities are sent by default.

Usage Guidelines

Use this command to enable and disable sending community attributes between BGP neighbors.

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to disable sending community attributes.

Use the **delete** form of this command to restore sending community attributes.

Use the **show** form of this command to view BGP peer group configuration settings

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast distribute-list export <access-list6-name>

Applies an access list to filter outbound routing updates to this peer group.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast  
distribute-list export access-list6-name
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
distribute-list export
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast  
distribute-list export
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      address-family {  
        ipv6-unicast {  
          distribute-list {  
            export access-list6-name  
          }  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
------------	--

<i>group-name</i>	Mandatory. The peer group name.
<i>access-list6-name</i>	Optional. The name of the IPv6 access-list.

Default

None.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to apply an access list to filter outbound routing updates to a BGP peer group.

Use the **delete** form of this command to disable outbound distribute list filtering.

Use the **show** form of this command to view BGP peer group distribute list configuration settings.

NOTE *A peer group distribute list cannot be used together with a peer group prefix list in the same direction. These two lists are mutually exclusive, and only one list may be applied to a given direction.*

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast distribute-list import <access-list6-name>

Applies an access list to filter inbound routing updates from this peer group.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast  
distribute-list import access-list6-name
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
distribute-list import
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast  
distribute-list import
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      address-family {  
        ipv6-unicast {  
          distribute-list {  
            import access-list6-name  
          }  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
------------	--

<i>group-name</i>	Mandatory. The peer group name.
<i>access-list6-name</i>	Optional. The name of the IPv6 access-list.

Default

None.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to apply an access list to filter inbound routing updates from a BGP peer group.

Use the **delete** form of this command to disable inbound distribute list filtering.

Use the **show** form of this command to view BGP peer group distribute list configuration settings.

NOTE *A peer group distribute list cannot be used together with a peer group prefix list in the same direction. These two lists are mutually exclusive, and only one list may be applied to the specified direction.*

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast filter-list export <as-path-list-name>

Applies an AS path list to routing updates to this peer group.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast filter-list  
export as-path-list-name
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
filter-list export as-path-list-name
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast  
filter-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      address-family {  
        ipv6-unicast {  
          filter-list {  
            export as-path-list-name  
          }  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
------------	--

<i>group-name</i>	Mandatory. The peer group name.
<i>as-path-list-name</i>	Mandatory. Name of an autonomous system path list.

Default

None.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to apply an access list for filtering outbound routing updates.

Use the **delete** form of this command to disable outbound routing update filtering.

Use the **show** form of this command to view BGP peer group filter list configuration settings.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast filter-list import <as-path-list-name>

Applies an AS path list to routing updates from this peer group.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast filter-list  
import as-path-list-name
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
filter-list import as-path-list-name
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast  
filter-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        peer-group group-name {  
            address-family {  
                ipv6-unicast {  
                    filter-list {  
                        import as-path-list-name  
                    }  
                }  
            }  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
------------	--

<i>group-name</i>	Mandatory. The peer group name.
<i>as-path-list-name</i>	Mandatory. Name of an autonomous system path list.

Default

None.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to apply an access list for filtering inbound routing updates.

Use the **delete** form of this command to disable inbound routing update filtering.

Use the **show** form of this command to view BGP peer group filter list configuration settings.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast maximum-prefix <max-num>

Sets the maximum number of prefixes to accept from this peer group before that peer group is taken down.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast  
maximum-prefix max-num
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
maximum-prefix max-num
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      address-family {  
        ipv6-unicast {  
          maximum-prefix max-num  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>max-num</i>	Mandatory. The maximum number of prefixes to accept from this peer group before the peer group is taken down.

Default

No maximum to the number of prefixes accepted.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to set the maximum number of prefixes to accept from this peer group before the peer group is taken down. This helps alleviate situations where the router receives more routes than it can handle.

Use the **delete** form of this command to remove the maximum prefix setting.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast nexthop-local unchanged

Specifies that the IPv6 link-local address is not changed when the prefix is announced to the peer.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast  
nexthop-local unchanged
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
nexthop-local
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast  
nexthop-local
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      address-family {  
        ipv6-unicast {  
          nexthop-local {  
            unchanged  
          }  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

Default

The IPv6 link-local address is changed when the prefix is announced to the peer.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to specify that the IPv6 link-local address is not changed when the prefix is announced to a peer.

Use the **delete** form of this command to return the system to its default behavior.

Use the **show** form of this command to view BGP peer group **ipv6-unicast nexthop-local** configuration settings.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast nexthop-self

Sets the local router as the next hop for this peer group.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast  
nexthop-self
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
nexthop-self
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      address-family {  
        ipv6-unicast {  
          nexthop-self  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

Default

Disabled.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to set this router as the next hop for this peer group. This disables the next-hop calculation for this peer group.

Use the **delete** form of this command to restore next-hop calculation for the peer group.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast prefix-list export <prefix-list6-name>

Applies a prefix list to filter updates to this peer group.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast prefix-list  
export prefix-list6-name
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
prefix-list export prefix-list6-name
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast  
prefix-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        peer-group group-name {  
            address-family {  
                ipv6-unicast {  
                    prefix-list {  
                        export prefix-list6-name  
                    }  
                }  
            }  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
------------	--

<i>group-name</i>	Mandatory. The peer group name.
<i>prefix-list6-name</i>	Mandatory. Name of a configured IPv6 prefix list.

Default

None.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to restrict distribution of outbound BGP peer group information by filtering with a prefix list.

Use the **delete** form of this command to remove an outbound prefix list filter.

Use the **show** form of this command to view BGP peer group prefix list configuration settings.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast prefix-list import <prefix-list6-name>

Applies a prefix list to filter updates from this peer group.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast prefix-list  
import prefix-list6-name
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
prefix-list import prefix-list6-name
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast  
prefix-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      address-family {  
        ipv6-unicast {  
          prefix-list {  
            import prefix-list6-name  
          }  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
------------	--

<i>group-name</i>	Mandatory. The peer group name.
<i>prefix-list6-name</i>	Mandatory. Name of a configured IPv6 prefix list.

Default

None.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to restrict distribution of inbound BGP peer group information by filtering with a prefix list.

Use the **delete** form of this command to remove an inbound prefix list filter.

Use the **show** form of this command to view BGP peer group prefix list configuration settings.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast remove-private-as

Directs the router to remove private AS numbers from updates sent to this peer group.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast  
remove-private-as
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
remove-private-as
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      address-family {  
        ipv6-unicast {  
          remove-private-as  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

Default

Private AS numbers are included in outgoing updates.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to direct the router to exclude private autonomous system (AS) numbers from updates to eBGP peers. When this feature is enabled, the router omits private AS numbers from the AS_PATH attribute. The range of private AS numbers is 64512 to 65534.

Note that it is a configuration error to include both private and public AS numbers in an AS path. If the router detects this error, it does not remove private AS numbers.

This command may be used in confederations provided that the private AS numbers are appended after the confederation portion of the AS path.

This command applies only to eBGP peers; it cannot be used with iBGP peers.

Use the **delete** form of this command to restore the default behavior.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast route-map export <map-name>

Applies a route map to filter updates to this peer group.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast route-map  
export map-name
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
route-map export map-name
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast  
route-map export map-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        peer-group group-name {  
            address-family {  
                ipv6-unicast {  
                    route-map {  
                        export map-name  
                    }  
                }  
            }  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
------------	--

<i>group-name</i>	Mandatory. The peer group name.
-------------------	---------------------------------

<i>map-name</i>	Mandatory. Name of a configured route map.
-----------------	--

Default

None.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to restrict distribution of outbound BGP peer group information by filtering with a route map.

Use the **delete** form of this command to remove an outbound route map filter.

Use the **show** form of this command to view BGP peer group route map configuration settings.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast route-map import <map-name>

Applies a route map to filter updates to or from this peer group.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast route-map  
import map-name
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
route-map import map-name
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast  
route-map import map-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      address-family {  
        ipv6-unicast {  
          route-map {  
            import map-name  
          }  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
------------	--

<i>group-name</i>	Mandatory. The peer group name.
-------------------	---------------------------------

<i>map-name</i>	Mandatory. Name of a configured route map.
-----------------	--

Default

None.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to restrict routing updates received from this BGP neighbor by filtering with a route map.

Use the **delete** form of this command to remove a route map filter.

Use the **show** form of this command to view BGP peer group route map configuration settings.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast soft-reconfiguration inbound

Directs the router to store received routing updates.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast  
soft-reconfiguration inbound
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
soft-reconfiguration inbound
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      address-family {  
        ipv6-unicast {  
          soft-reconfiguration {  
            inbound  
          }  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

Default

Disabled.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to enable soft reconfiguration.

When this command is issued, the router begins storing routing updates, which can be used subsequently for inbound soft reconfiguration. Outbound BGP soft reconfiguration can be performed without enabling inbound soft reconfiguration.

Use the **delete** form of this command to disable soft reconfiguration.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast unsuppress-map <map-name>

Directs the router to selectively advertise routes suppressed by aggregating addresses, based on a route map.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast  
unsuppress-map map-name
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
unsuppress-map map-name
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      address-family {  
        ipv6-unicast {  
          unsuppress-map map-name  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>map-name</i>	Mandatory. The name of a configured route map.

Default

Routes suppressed by address aggregation are not advertised.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to selectively advertise routes suppressed by aggregating addresses.

Use the **delete** form of this command to restore the default behavior.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> allowas-in

Allows or disallows receiving BGP advertisements containing the AS path of the local router.

Syntax

```
set protocols bgp asn peer-group group-name allowas-in [number num]
```

```
delete protocols bgp asn peer-group group-name allowas-in
```

```
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        peer-group group-name {  
            allowas-in {  
                number num  
            }  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
number <i>num</i>	Optional. Specifies the number of times to accept an AS path containing the readvertisement of the local router's ASN. The range is 1 to 10. The default is 3.

Default

Readvertisement is disabled.

Usage Guidelines

Use the **set** form of this command to allow the router to accept BGP AS paths advertising the router's own AS number.

This situation could indicate a routing loop, and by default such an update is dropped. However, you can use this command to direct the router to accept updates readvertising its AS number for a specified number of times.

Use the **delete** form of this command to disallow readvertisement of the router's AS path.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> attribute-unchanged

Allows the router to send updates to a peer group with unchanged attributes.

Syntax

```
set protocols bgp asn peer-group group-name attribute-unchanged [as-path | med |  
next-hop]
```

```
delete protocols bgp asn peer-group group-name attribute-unchanged [as-path | med  
| next-hop]
```

```
show protocols bgp asn peer-group group-name attribute-unchanged
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      attribute-unchanged {  
        as-path  
        med  
        next-hop  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>as-path</i>	Optional. Propagates the route update with unchanged AS_PATH attribute.

med	Optional. Propagates the route update with unchanged Multi Exit Discriminator.
next-hop	Optional. Propagates the route update with unchanged next hop.

Default

Disabled.

Usage Guidelines

Use the **set** form of this command to direct the router to propagate routing updates without modifying the BGP AS_PATH, Multi Exit Discriminator (MED), or next-hop attribute.

Use the **delete** form of this command to restore normal modification of BGP attributes.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> capability dynamic

Advertises support for dynamic update of BGP capabilities advertised and received from this peer group.

Syntax

```
set protocols bgp asn peer-group group-name capability dynamic
delete protocols bgp asn peer-group group-name capability dynamic
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      capability {
        dynamic
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

Default

By default, the session is brought up with minimal capability on both sides.

Usage Guidelines

In general, BGP-4 requires that if a BGP speaker receives an OPEN message with an unrecognized optional parameters, the speaker must terminate BGP peering. This makes it difficult to introduce new capabilities into the protocol. The Capabilities parameter allows graceful negotiation of BGP capabilities without requiring a peer to terminate peering.

This command specifies use of the BGP Dynamic Capability feature, which allows dynamic update of capabilities over an established BGP session.

Use the **set** form of this command to specify dynamic update of BGP capabilities to be used in capability negotiation.

Use the **delete** form of this command to remove the dynamic update capability.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> capability orf

Advertises support for Outbound Route Filtering (ORF) for updating BGP capabilities advertised and received from this peer group.

Syntax

```
set protocols bgp asn peer-group group-name capability orf [prefix-list [receive | send]]
```

```
delete protocols bgp asn peer-group group-name capability orf
```

```
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      capability {  
        orf {  
          prefix-list {  
            receive  
            send  
          }  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>prefix-list</i>	Optional. Advertise prefix-list ORF capability to this peer group.

receive	Optional. Advertise the ability to receive the ORF from this peer group.
send	Optional. Advertise the capability to send the ORF to this peer group.

Default

By default, the session is brought up with minimal capability on both sides.

Usage Guidelines

In general, BGP-4 requires that if a BGP speaker receives an OPEN message with an unrecognized optional parameters, the speaker must terminate BGP peering. This makes it difficult to introduce new capabilities into the protocol. The Capabilities parameter allows graceful negotiation of BGP capabilities without requiring a peer to terminate peering.

This command specifies use of BGP Outbound Route Filtering (ORF) to send and receive capabilities. Using ORF minimizes the number of BGP updates that are sent between peer routers.

Use the **set** form of this command to specify BGP ORF capabilities to be used in capability negotiation.

Use the **delete** form of this command to remove the ORF capability.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> default-originate

Enables or disables forwarding of the default route to a BGP peer group.

Syntax

```
set protocols bgp asn peer-group group-name default-originate [route-map  
map-name]
```

```
delete protocols bgp asn peer-group group-name default-originate [route-map  
map-name]
```

```
show protocols bgp asn peer-group group-name default-originate
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        peer-group group-name {  
            default-originate {  
                route-map map-name  
            }  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>map-name</i>	Optional. Specifies a configured route map to be used when advertising the default route.

Default

The default route is not distributed.

Usage Guidelines

Use the **set** form of this command to enable the router to advertise the default route 0.0.0.0 to this BGP peer group. This route can then be used by the peer group to reach the local router if no other routes are available.

The route 0.0.0.0 need not be explicitly configured on the local router.

If a route map is specified, the default route is advertised if two conditions are satisfied:

- The route map includes a **match ip** address clause.
- A route exists that exactly matches the IP access list.

Route maps are configured using [protocols bgp <asn> peer-group <group-name> local-as <asn> command](#).

Use the **delete** form of this command to disable forwarding of the default route or to delete a route map.

Use the **show** form of this command to view BGP peer group default route configuration settings.

protocols bgp <asn> peer-group <group-name> description <desc>

Provides a brief description for a BGP peer group.

Syntax

```
set protocols bgp asn peer-group group-name description desc  
delete protocols bgp asn peer-group group-name description  
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        peer-group group-name {  
            description desc  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>desc</i>	Mandatory. A description of up to 80 characters for the peer group. If spaces are included, the description must be enclosed in quotes.

Default

None.

Usage Guidelines

Use the **set** form of this command to enter a description for this BGP peer group.

Use the **delete** form of this command to delete the peer group's description.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> disable-capability-negotiation

Disables BGP capability negotiation.

Syntax

```
set protocols bgp asn peer-group group-name disable-capability-negotiation
delete protocols bgp asn peer-group group-name disable-capability-negotiation
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      disable-capability-negotiation
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

Default

Capability negotiation is performed.

Usage Guidelines

Use the set form of this command to disable BGP capability negotiation.

Use the **delete** form of this command to delete this attribute and restore BGP capability negotiation.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> disable-connected-check

Disables direct connection verification for single-hop eBGP peers.

Syntax

```
set protocols bgp asn peer-group group-name disable-connected-check
delete protocols bgp asn peer-group group-name disable-connected-check
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      disable-connected-check
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

Default

Connection verification is performed.

Usage Guidelines

Use the set form of this command to disable connection verification.

Connection verification is used to establish an eBGP peering session with a single-hop peer. A single-hop peer is defined using the `peer group ebgp-multihop` command and specifying a time-to-live (TTL) value of 1. Such a peer is reachable by a single hop, but is configured on a loopback interface or is otherwise configured with a non-directly connected IP address.

For these peers, the BGP process normally checks to determine whether the eBGP peer is directly connected to the same network as the local router. If not, the peering session is not established.

If connection verification is disabled, source updating must be enabled using `protocols bgp <asn> peer-group <group-name> update-source <source>` command so that the BGP process can use the loopback interface for the peering session.

Use the `delete` form of this command to restore connection verification for eBGP peering sessions.

Use the `show` form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> disable-send-community

Disables sending of community attributes to the specified peer group.

Syntax

```
set protocols bgp asn peer-group group-name disable-send-community [extended | standard]
```

```
delete protocols bgp asn peer-group group-name disable-send-community
```

```
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        peer-group group-name {  
            disable-send-community {  
                extended  
                standard  
            }  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>extended</i>	Optional. Disable sending extended community attributes.
<i>standard</i>	Optional. Disable sending standard community attributes.

Default

Communities are sent by default.

Usage Guidelines

Use this command to enable and disable sending community attributes between BGP peer groups.

Use the **set** form of this command to disable sending community attributes.

Use the **delete** form of this command to restore sending community attributes.

Use the **show** form of this command to view BGP peer group configuration settings

protocols bgp <asn> peer-group <group-name> distribute-list export <acl-num>

Applies an access list to filter outbound routing updates to this peer group.

Syntax

```
set protocols bgp asn peer-group group-name distribute-list export acl-num
delete protocols bgp asn peer-group group-name distribute-list
show protocols bgp asn peer-group group-name distribute-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      distribute-list {
        export acl-num
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>acl-num</i>	Optional. The number of a standard or extended access list. The range for a standard access list is 1 to 99. The range for an extended access list is 100 to 199.

Default

None.

Usage Guidelines

Use the **set** form of this command to apply an access list to filter outbound routing updates to a BGP peer group.

Use the **delete** form of this command to disable outbound distribute list filtering.

Use the **show** form of this command to view BGP peer group distribute list configuration settings.

NOTE *A peer group distribute list cannot be used together with a peer group prefix list in the same direction. These two lists are mutually exclusive, and only one list may be applied to a given direction.*

protocols bgp <asn> peer-group <group-name> distribute-list import <acl-num>

Applies an access list to filter inbound routing updates from this peer group.

Syntax

```
set protocols bgp asn peer-group group-name distribute-list import acl-num
delete protocols bgp asn peer-group group-name distribute-list
show protocols bgp asn peer-group group-name distribute-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      distribute-list {
        import acl-num
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>acl-num</i>	Optional. The number of a standard or extended access list. The range for a standard access list is 1 to 99. The range for an extended access list is 100 to 199.

Default

None.

Usage Guidelines

Use the **set** form of this command to apply an access list to filter inbound routing updates from a BGP peer group.

Use the **delete** form of this command to disable inbound distribute list filtering.

Use the **show** form of this command to view BGP peer group distribute list configuration settings.

NOTE *A peer group distribute list cannot be used together with a peer group prefix list in the same direction. These two lists are mutually exclusive, and only one list may be applied to the specified direction.*

protocols bgp <asn> peer-group <group-name> ebgp-multihop <ttl>

Allows eBGP neighbors not on directly connected networks.

Syntax

```
set protocols bgp asn peer-group group-name ebgp-multihop ttl
delete protocols bgp asn peer-group group-name ebgp-multihop
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      ebgp-multihop ttl
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>ttl</i>	Mandatory. The time-to-live, or maximum number of hops allowed. The range is 1 to 255.

Default

Only directly connected neighbors are allowed.

Usage Guidelines

Use the **set** form of this command to allow connections to eBGP peers residing on networks that are not directly connected.

Use the **delete** form of this command to restrict connections to directly connected peers.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> filter-list export <as-path-list-name>

Applies an AS path list to routing updates to this peer group.

Syntax

```
set protocols bgp asn peer-group group-name filter-list export as-path-list-name
delete protocols bgp asn peer-group group-name filter-list export as-path-list-name
show protocols bgp asn peer-group group-name filter-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      filter-list {
        export as-path-list-name
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>as-path-list-name</i>	Mandatory. Name of an autonomous system path list.

Default

None.

Usage Guidelines

Use the **set** form of this command to apply an access list for filtering outbound routing updates.

Use the **delete** form of this command to disable outbound routing update filtering.

Use the **show** form of this command to view BGP peer group filter list configuration settings.

protocols bgp <asn> peer-group <group-name> filter-list import <as-path-list-name>

Applies an AS path list to routing updates from this peer group.

Syntax

```
set protocols bgp asn peer-group group-name filter-list import as-path-list-name
delete protocols bgp asn peer-group group-name filter-list import as-path-list-name
show protocols bgp asn peer-group group-name filter-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      filter-list {
        import as-path-list-name
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>as-path-list-name</i>	Mandatory. Name of an autonomous system path list.

Default

None.

Usage Guidelines

Use the **set** form of this command to apply an access list for filtering inbound routing updates.

Use the **delete** form of this command to disable inbound routing update filtering.

Use the **show** form of this command to view BGP peer group filter list configuration settings.

protocols bgp <asn> peer-group <group-name> local-as <asn>

Defines a local autonomous system number for eBGP peerings.

Syntax

```
set protocols bgp asn peer-group group-name local-as asn [no-prepend]
delete protocols bgp asn peer-group group-name local-as asn [no-prepend]
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      local-as asn {
        no-prepend
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>asn</i>	Optional. Valid autonomous system number. This may not be the number of the autonomous system to which the peer group belongs. The range is 1 to 4294967294.
no-prepend	Optional. Directs the router to not prepend the local autonomous system number to routes received from an external peer.

Default

None.

Usage Guidelines

Use the **set** form of this command to set a local autonomous system number for eBGP peer groupings. This number is used by all peers in the group for peering. It cannot be applied to individual peers in the group.

A local autonomous system number can only be applied to a true eBGP peer; it cannot be applied to peers in different subautonomous systems within a confederation.

The **no-prepend** keyword can cause routing loops and should be used with care. It should be used only to change the autonomous system number in a BGP network. After the network transition has completed, this setting should be deleted.

Use the **delete** form of this command to remove a local autonomous system number, or to remove the **no-prepend** keyword.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> maximum-prefix <max-num>

Sets the maximum number of prefixes to accept from this peer group before that peer group is taken down.

Syntax

```
set protocols bgp asn peer-group group-name maximum-prefix max-num  
delete protocols bgp asn peer-group group-name maximum-prefix max-num  
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        peer-group group-name {  
            maximum-prefix max-num  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>max-num</i>	Mandatory. The maximum number of prefixes to accept from this peer group before the peer group is taken down.

Default

No maximum to the number of prefixes accepted.

Usage Guidelines

Use the **set** form of this command to set the maximum number of prefixes to accept from this peer group before the peer group is taken down. This helps alleviate situations where the router receives more routes than it can handle.

Use the **delete** form of this command to remove the maximum prefix setting.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> nexthop-self

Sets the local router as the next hop for this peer group.

Syntax

```
set protocols bgp asn peer-group group-name nexthop-self
delete protocols bgp asn peer-group group-name nexthop-self
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      nexthop-self
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

Default

Disabled.

Usage Guidelines

Use the set form of this command to set this router as the next hop for this peer group. This disables the next-hop calculation for this peer group.

Use the **delete** form of this command to restore next-hop calculation for the peer group.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> override-capability

Overrides capability negotiation to allow a peering session to be established with a peer group that does not support capability negotiation.

Syntax

```
set protocols bgp asn peer-group group-name override-capability
delete protocols bgp asn peer-group group-name override-capability
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
    bgp asn {
        peer-group group-name {
            override-capability
        }
    }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

Default

A session cannot be established if the peer group does not support capability negotiation.

Usage Guidelines

Use the **set** form of this command to enable a peering session to be established with a BGP peer group that does not support capability negotiation.

Normally, if a BGP peer does not support capability negotiation, a peering session cannot be established and the connection is terminated. Setting this value overrides this process and allows the session to be established.

Use the **delete** form of this command to restore the default behavior.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> passive

Directs the router not to initiate connections with this peer group.

Syntax

```
set protocols bgp asn peer-group group-name passive
delete protocols bgp asn peer-group group-name passive
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      passive
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

Default

The router both accepts inbound connections and initiates outbound connections.

Usage Guidelines

Use the **set** form of this command to configure the local router such that accepts inbound connections from the peer group, but does not initiate outbound connections by sending a BGP OPEN message.

Use the **delete** form of this command to restore the default behavior.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> password <pwd>

Defines a BGP MD5 password.

Syntax

```
set protocols bgp asn peer-group group-name password pwd
delete protocols bgp asn peer-group group-name password pwd
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      password pwd
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>pwd</i>	Mandatory. A password to be used to generate an MD5 digest.

Default

None.

Usage Guidelines

Use this command to specify a password-like string for generating an MD5 authentication string. This string is used for verification by the BGP peer group. If the MD5 digest generated by the peer group does not match the MD5 digest generated by the local router, any communication is silently discarded by the peer group.

Use the **set** form of this command to specify or modify the BGP MD5 password.

Use the **delete** form of this command to remove the BGP MD5 password.

Use the **show** form of this command to view BGP MD5 password configuration.

protocols bgp <asn> peer-group <group-name> prefix-list export <list-name>

Applies a prefix list to filter updates to this peer group.

Syntax

```
set protocols bgp asn peer-group group-name prefix-list export list-name
delete protocols bgp asn peer-group group-name prefix-list export list-name
show protocols bgp asn peer-group group-name prefix-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      prefix-list {
        export list-name
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>list-name</i>	Mandatory. Name of a configured prefix list.

Default

None.

Usage Guidelines

Use the **set** form of this command to restrict distribution of outbound BGP peer group information by filtering with a prefix list.

Use the **delete** form of this command to remove an outbound prefix list filter.

Use the **show** form of this command to view BGP peer group prefix list configuration settings.

protocols bgp <asn> peer-group <group-name> prefix-list import <list-name>

Applies a prefix list to filter updates from this peer group.

Syntax

```
set protocols bgp asn peer-group group-name prefix-list import list-name
delete protocols bgp asn peer-group group-name prefix-list import list-name
show protocols bgp asn peer-group group-name prefix-list
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      prefix-list {
        import list-name
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>list-name</i>	Mandatory. Name of a configured prefix list.

Default

None.

Usage Guidelines

Use the **set** form of this command to restrict distribution of inbound BGP peer group information by filtering with a prefix list.

Use the **delete** form of this command to remove an inbound prefix list filter.

Use the **show** form of this command to view BGP peer group prefix list configuration settings.

protocols bgp <asn> peer-group <group-name> remote-as <asn>

Specifies the autonomous system number of the peer group.

Syntax

```
set protocols bgp asn peer-group group-name remote-as asn
delete protocols bgp asn peer-group group-name remote-as
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      remote-as asn
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

Default

None.

Usage Guidelines

Use the set form of this command to specify the autonomous system (AS) number of a BGP peer group.

If the AS number of the peer group is the same as that of the local router, the neighbor is an internal BGP (iBGP) peer. If it is different the neighbor is an external BGP (eBGP) peer.

Use the **delete** form of this command to remove AS number settings for the peer group.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> remove-private-as

Directs the router to remove private AS numbers from updates sent to this peer group.

Syntax

```
set protocols bgp asn peer-group group-name remove-private-as
delete protocols bgp asn peer-group group-name remove-private-as
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      remove-private-as
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

Default

Private AS numbers are included in outgoing updates.

Usage Guidelines

Use the **set** form of this command to direct the router to exclude private autonomous system (AS) numbers from updates to eBGP peers. When this feature is enabled, the router omits private AS numbers from the AS_PATH attribute. The range of private AS numbers is 64512 to 65534.

Note that it is a configuration error to include both private and public AS numbers in an AS path. If the router detects this error, it does not remove private AS numbers.

This command may be used in confederations provided that the private AS numbers are appended after the confederation portion of the AS path.

This command applies only to eBGP peers; it cannot be used with iBGP peers.

Use the **delete** form of this command to restore the default behavior.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> route-map export <map-name>

Applies a route map to filter updates to this peer group.

Syntax

```
set protocols bgp asn peer-group group-name route-map export map-name  
delete protocols bgp asn peer-group group-name route-map export map-name  
show protocols bgp asn peer-group group-name route-map export map-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        peer-group group-name {  
            route-map {  
                export map-name  
            }  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>map-name</i>	Mandatory. Name of a configured route map.

Default

None.

Usage Guidelines

Use the **set** form of this command to restrict distribution of outbound BGP peer group information by filtering with a route map.

Use the **delete** form of this command to remove an outbound route map filter.

Use the **show** form of this command to view BGP peer group route map configuration settings.

protocols bgp <asn> peer-group <group-name> route-map import <map-name>

Applies a route map to filter updates to or from this peer group.

Syntax

```
set protocols bgp asn peer-group group-name route-map import map-name  
delete protocols bgp asn peer-group group-name route-map import map-name  
show protocols bgp asn peer-group group-name route-map import map-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        peer-group group-name {  
            route-map {  
                import map-name  
            }  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>map-name</i>	Mandatory. Name of a configured route map.

Default

None.

Usage Guidelines

Use the **set** form of this command to restrict routing updates received from this BGP peer group by filtering with a route map.

Use the **delete** form of this command to remove a route map filter.

Use the **show** form of this command to view BGP peer group route map configuration settings.

protocols bgp <asn> peer-group <group-name> shutdown

Administratively shuts down a BGP peer group.

Syntax

```
set protocols bgp asn peer-group group-name shutdown
delete protocols bgp asn peer-group group-name shutdown
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      shutdown
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

Default

Disabled.

Usage Guidelines

Use the set form of this command to administratively shut down a BGP peer group.

Shutting down a BGP speaker administratively terminates any active sessions for the peer group and removes any associated routing information. This could terminate a large number of sessions.

To determine whether a BGP neighbor has been administratively shut down, you can view BGP summary information. A neighbor with a status of Idle with a comment of Admin have been administratively shut down.

Use the **delete** form of this command to administratively reenab a BGP peer group.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> soft-reconfiguration inbound

Directs the router to store received routing updates.

Syntax

```
set protocols bgp asn peer-group group-name soft-reconfiguration inbound
delete protocols bgp asn peer-group group-name soft-reconfiguration inbound
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      soft-reconfiguration {
        inbound
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

Default

Disabled.

Usage Guidelines

Use the set form of this command to enable soft reconfiguration.

When this command is issued, the router begins storing routing updates, which can be used subsequently for inbound soft reconfiguration. Outbound BGP soft reconfiguration can be performed without enabling inbound soft reconfiguration.

Use the **delete** form of this command to disable soft reconfiguration.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> ttl-security hops <hops>

Sets Time to Live (TTL) security hop count for this peer group.

Syntax

```
set protocols bgp asn peer-group group-name ttl-security hops hops  
delete protocols bgp asn peer-group group-name ttl-security hops  
show protocols bgp asn peer-group group-name ttl-security hops
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      ttl-security {  
        hops hops  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>hops</i>	Optional. The maximum number of hops away from the local peer a peering session connection will be accepted. The range is 1 to 254.

Default

The default for the hop count is 254 for IBGP and 1 for EBGP connections.

Usage Guidelines

Use this command to set the Time to Live (TTL) security for the local peer group.

TTL security is incompatible with EBGP Multihops.

Use the **set** form of this command to specify the hop count.

Use the **delete** form of this command to return the system to its default.

Use the **show** form of this command to view the parameter configuration.

protocols bgp <asn> peer-group <group-name> unsuppress-map <map-name>

Directs the router to selectively advertise routes suppressed by aggregating addresses, based on a route map.

Syntax

```
set protocols bgp asn peer-group group-name unsuppress-map map-name
delete protocols bgp asn peer-group group-name unsuppress-map map-name
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
    bgp asn {
        peer-group group-name {
            unsuppress-map map-name
        }
    }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>map-name</i>	Mandatory. The name of a configured route map.

Default

Routes suppressed by address aggregation are not advertised.

Usage Guidelines

Use the **set** form of this command to selectively advertise routes suppressed by aggregating addresses.

Use the **delete** form of this command to restore the default behavior.

Use the **show** form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <group-name> update-source <source>

Specifies the source IP address or interface for routing updates.

Syntax

```
set protocols bgp asn peer-group group-name update-source source  
delete protocols bgp asn peer-group group-name update-source  
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        peer-group group-name {  
            update-source source  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>source</i>	Mandatory. The IPv4 address of the router to receive routing updates from or the interface to receive routing updates from..

Default

None.

Usage Guidelines

Use the **set** form of this command to direct the system to use a specific source for routing updates.

Use the **delete** form of this command to remove the routing update source.

Use the **show** form of this command to view the configuration settings.

protocols bgp <asn> peer-group <group-name> weight <weight>

Defines a default weight for routes from this peer group.

Syntax

```
set protocols bgp asn peer-group group-name weight weight
delete protocols bgp asn peer-group group-name weight
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      weight weight
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.
<i>weight</i>	Mandatory. The weight to be assigned to routes from this peer group. The range is 0 to 65535.

Default

Routes learned from a BGP neighbor have a weight of 0. Routes sourced by the local router have a weight of 32768.

Usage Guidelines

Use the **set** form of this command to set the default weights for routes learned from a BGP peer group.

Use the **delete** form of this command to restore route weighting to the default. Existing route weights are not changed.

Use the **show** form of this command to view BGP peer group configuration settings.

reset ip bgp peer-group <group-name>

Resets sessions for all members of a peer group.

Syntax

```
reset ip bgp peer-group group-name [in [prefix-filter] | out | soft [in | out]]
```

Command Mode

Operational mode.

Parameters

<i>group-name</i>	Mandatory. The name of a defined BGP peer group.
in	Optional. Resets inbound sessions only.
out	Optional. Resets outbound sessions only.
prefix-filter	Optional. Resets the BGP outbound route filter (ORF). This keyword is ignored unless ORF capabilities have been enabled on the local system or received from the sending BGP peer. In this case, a normal inbound soft reset is performed.
soft	Optional. Sessions are not reset. Resend updates that have changed due to an export policy change or recalculate any updates received based on import policy changes.
in	Optional. Re-read import policies and recalculates information in the BGP table based on import policy changes. This requires soft-reconfiguration inbound to be configured on the neighbor.
out	Optional. Re-read export policies and resend updates that have changed due to an export policy change.

Default

When used without the **soft** option, reset connections are dropped, both inbound and outbound.

Usage Guidelines

Use this command to reset sessions for all members of a peer group. This forces BGP updates to be generated and new BGP policies to be applied.

Unless the **soft** option is used, all connections are dropped (a “hard reset”): TCP connections are terminated and all routes received from the neighbor are removed from the BGP routing table. Then the connection with the neighbor is reestablished.

If the **soft** option is used, routes from the neighbor are marked as stale but are not immediately removed from the BGP table. Stale routes that are not received from the neighbor when the connection is reestablished are removed from the BGP table at that point.

reset ip bgp peer-group <group-name> ipv4 unicast

Resets IPv4 unicast sessions for all members of a peer group.

Syntax

```
reset ip bgp peer-group group-name ipv4 unicast [in [prefix-filter] | out | soft [in | out]]
```

Command Mode

Operational mode.

Parameters

<i>group-name</i>	Mandatory. The name of a defined BGP peer group.
in	Optional. Resets inbound sessions only.
out	Optional. Resets outbound sessions only.
prefix-filter	Optional. Resets the BGP outbound route filter (ORF). This keyword is ignored unless ORF capabilities have been enabled on the local system or received from the sending BGP peer. In this case, a normal inbound soft reset is performed.
soft	Optional. Sessions are not reset. Resend updates that have changed due to an export policy change or recalculate any updates received based on import policy changes.
in	Optional. Re-read import policies and recalculates information in the BGP table based on import policy changes. This requires soft-reconfiguration inbound to be configured on the neighbor.
out	Optional. Re-read export policies and resend updates that have changed due to an export policy change.

Default

None.

Usage Guidelines

Use this command to reset inbound unicast sessions for all members of a peer group. This forces BGP updates to be generated and new BGP policies to be applied.

Unless the **soft** option is used, all connections are dropped (a “hard reset”): TCP connections are terminated and all routes received from the neighbor are removed from the BGP routing table. Then the connection with the neighbor is reestablished.

If the **soft** option is used, routes from the neighbor are marked as stale but are not immediately removed from the BGP table. Stale routes that are not received from the neighbor when the connection is reestablished are removed from the BGP table at that point.

Chapter 7: Route Redistribution

This chapter describes commands for BGP route redistribution.

This chapter presents the following topics:

- [Route Redistribution Commands](#)

Route Redistribution Commands

This chapter contains the following commands.

Configuration Commands	
Route Redistribution Configuration - IPv6	
<code>protocols bgp <asn> address-family ipv6-unicast redistribute connected</code>	Redistributes IPv6 unicast directly connected routes into BGP.
<code>protocols bgp <asn> address-family ipv6-unicast redistribute kernel</code>	Redistributes IPv6 unicast kernel routes into BGP.
<code>protocols bgp <asn> address-family ipv6-unicast redistribute ospfv3</code>	Redistributes IPv6 unicast routes learned from OSPFv3 into BGP.
<code>protocols bgp <asn> address-family ipv6-unicast redistribute ripng</code>	Redistributes IPv6 unicast routes learned from RIPng into BGP.
<code>protocols bgp <asn> address-family ipv6-unicast redistribute static</code>	Redistributes IPv6 unicast static routes into BGP.
Route Redistribution Configuration	
<code>protocols bgp <asn> redistribute connected</code>	Redistributes directly connected routes into BGP.
<code>protocols bgp <asn> redistribute kernel</code>	Redistributes kernel routes into BGP.
<code>protocols bgp <asn> redistribute ospf</code>	Redistributes routes learned from OSPF into BGP.
<code>protocols bgp <asn> redistribute rip</code>	Redistributes routes learned from RIP into BGP.
<code>protocols bgp <asn> redistribute static</code>	Redistributes static routes into BGP.
Operational Commands	
None.	

protocols bgp <asn> address-family ipv6-unicast redistribute connected

Redistributes IPv6 unicast directly connected routes into BGP.

Syntax

```
set protocols bgp asn address-family ipv6-unicast redistribute connected [metric | route-map map-name]
```

```
delete protocols bgp asn address-family ipv6-unicast redistribute connected [metric | route-map]
```

```
show protocols bgp asn address-family ipv6-unicast redistribute connected
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    address-family {
      ipv6-unicast {
        redistribute {
          connected {
            metric metric
            route-map map-name
          }
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>metric</i>	Optional. The metric to be applied to redistributed connected routes.

<i>map-name</i>	Optional. The name of a configured route map to be used for redistributing connected routes.
-----------------	--

Default

When this command has not been set, directly connected routes are not distributed into BGP.

Usage Guidelines

Use this command to redistribute directly connected routes into BGP. The command applies to IPv6 unicast routes only.

Use the **set** form of this command to direct the router to redistribute directly connected routes into BGP.

Use the **delete** form of this command to prevent redistribution of directly connected routes into BGP.

Use the **show** form of this command to view route redistribution configuration settings.

protocols bgp <asn> address-family ipv6-unicast redistribute kernel

Redistributes IPv6 unicast kernel routes into BGP.

Syntax

```
set protocols bgp asn address-family ipv6-unicast redistribute kernel [metric metric |  
route-map map-name]
```

```
delete protocols bgp asn address-family ipv6-unicast redistribute kernel [metric |  
route-map]
```

```
show protocols bgp asn address-family ipv6-unicast redistribute kernel
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        address-family {  
            ipv6-unicast {  
                redistribute {  
                    kernel {  
                        metric metric  
                        route-map map-name  
                    }  
                }  
            }  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>metric</i>	Optional. The metric to be applied to redistributed kernel routes.

<i>map-name</i>	Optional. The name of a configured route map to be used for redistributing kernel routes.
-----------------	---

Default

When this command has not been set, kernel routes are not distributed into BGP.

Usage Guidelines

Use this command to redistribute kernel routes into BGP. The command applies to IPv6 unicast routes only.

Use the **set** form of this command to direct the router to redistribute kernel routes into BGP.

Use the **delete** form of this command to prevent redistribution of kernel routes into BGP.

Use the **show** form of this command to view route redistribution configuration settings.

protocols bgp <asn> address-family ipv6-unicast redistribute ospfv3

Redistributes IPv6 unicast routes learned from OSPFv3 into BGP.

Syntax

```
set protocols bgp asn address-family ipv6-unicast redistribute ospfv3 [metric metric | route-map map-name]
```

```
delete protocols bgp asn address-family ipv6-unicast redistribute ospfv3 [metric | route-map]
```

```
show protocols bgp asn address-family ipv6-unicast redistribute ospfv3
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    address-family {
      ipv6-unicast {
        redistribute {
          ospfv3 {
            metric metric
            route-map map-name
          }
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>metric</i>	Optional. The metric to be applied to redistributed OSPFv3 routes.

<i>map-name</i>	Optional. The name of a configured route map to be used for redistributing OSPFv3 routes.
-----------------	---

Default

When this command has not been set, routes learned from OSPFv3 are not distributed into BGP.

Usage Guidelines

Use this command to redistribute Open Shortest Path First Version 3 (OSPFv3) routes into BGP. The command applies to IPv6 unicast routes only.

Use the **set** form of this command to direct the router to redistribute routes learned from OSPFv3 into BGP.

Use the **delete** form of this command to prevent redistribution of routes learned from OSPFv3 into BGP.

Use the **show** form of this command to view route redistribution configuration settings.

protocols bgp <asn> address-family ipv6-unicast redistribute ripng

Redistributes IPv6 unicast routes learned from RIPng into BGP.

Syntax

```
set protocols bgp asn address-family ipv6-unicast redistribute ripng [metric metric |  
route-map map-name]
```

```
delete protocols bgp asn address-family ipv6-unicast redistribute ripng [metric |  
route-map]
```

```
show protocols bgp asn address-family ipv6-unicast redistribute ripng
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        address-family {  
            ipv6-unicast {  
                redistribute {  
                    ripng {  
                        metric metric  
                        route-map map-name  
                    }  
                }  
            }  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>metric</i>	Optional. The metric to be applied to redistributed RIPng routes.

<i>map-name</i>	Optional. The name of a configured route map to be used for redistributing RIPng routes.
-----------------	--

Default

When this command has not been set, routes learned from RIPng are not distributed into BGP.

Usage Guidelines

Use this command to redistribute Routing Information Protocol - next generation (RIPng) routes into BGP. The command applies to IPv6 unicast routes only.

Use the **set** form of this command to direct the router to redistribute routes learned from RIPng into BGP.

Use the **delete** form of this command to prevent redistribution of routes learned from RIPng into BGP.

Use the **show** form of this command to view route redistribution configuration settings.

protocols bgp <asn> address-family ipv6-unicast redistribute static

Redistributes IPv6 unicast static routes into BGP.

Syntax

```
set protocols bgp asn address-family ipv6-unicast redistribute static [metric metric |  
route-map map-name]
```

```
delete protocols bgp asn address-family ipv6-unicast redistribute static [metric |  
route-map]
```

```
show protocols bgp asn address-family ipv6-unicast redistribute static
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        address-family {  
            ipv6-unicast {  
                redistribute {  
                    static {  
                        metric metric  
                        route-map map-name  
                    }  
                }  
            }  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>metric</i>	Optional. The metric to be applied to redistributed static routes.

<i>map-name</i>	Optional. The name of a configured route map to be used for redistributing static routes.
-----------------	---

Default

When this command has not been set static routes are not distributed into BGP.

Usage Guidelines

Use this command to redistribute static routes into BGP. The command applies to IPv6 unicast routes only.

Use the **set** form of this command to direct the router to redistribute static routes into BGP.

Use the **delete** form of this command to prevent redistribution of static routes into BGP.

Use the **show** form of this command to view route redistribution configuration settings.

protocols bgp <asn> redistribute connected

Redistributes directly connected routes into BGP.

Syntax

```
set protocols bgp asn redistribute connected [metric metric | route-map map-name]  
delete protocols bgp asn redistribute connected [metric | route-map]  
show protocols bgp asn redistribute
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        redistribute {  
            connected {  
                metric metric  
                route-map map-name  
            }  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>metric</i>	Optional. The metric to be applied to redistributed connected routes.
<i>map-name</i>	Optional. The name of a configured route map to be used for redistributing connected routes.

Default

When this command has not been set, directly connected routes are not distributed into BGP.

Usage Guidelines

Use this command to redistribute directly connected routes into BGP.

Use the **set** form of this command to direct the router to redistribute directly connected routes into BGP.

Use the **delete** form of this command to prevent redistribution of directly connected routes into BGP.

Use the **show** form of this command to view route redistribution configuration settings.

protocols bgp <asn> redistribute kernel

Redistributes kernel routes into BGP.

Syntax

```
set protocols bgp asn redistribute kernel [metric metric | route-map map-name]  
delete protocols bgp asn redistribute kernel [metric | route-map]  
show protocols bgp asn redistribute
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    redistribute {  
      kernel {  
        metric metric  
        route-map map-name  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>metric</i>	Optional. The metric to be applied to redistributed kernel routes.
<i>map-name</i>	Optional. The name of a configured route map to be used for redistributing kernel routes.

Default

When this command has not been set, kernel routes are not distributed into BGP.

Usage Guidelines

Use this command to redistribute kernel routes into BGP.

Use the **set** form of this command to direct the router to redistribute kernel routes into BGP.

Use the **delete** form of this command to prevent redistribution of kernel routes into BGP.

Use the **show** form of this command to view route redistribution configuration settings.

protocols bgp <asn> redistribute ospf

Redistributes routes learned from OSPF into BGP.

Syntax

```
set protocols bgp asn redistribute ospf [metric metric | route-map map-name]  
delete protocols bgp asn redistribute ospf [metric | route-map]  
show protocols bgp asn redistribute
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    redistribute {  
      ospf {  
        metric metric  
        route-map map-name  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>metric</i>	Optional. The metric to be applied to redistributed OSPF routes.
<i>map-name</i>	Optional. The name of a configured route map to be used for redistributing OSPF routes.

Default

When this command has not been set, routes learned from OSPF are not distributed into BGP.

Usage Guidelines

Use this command to redistribute Open Shortest Path First (OSPF) routes into BGP.

Use the **set** form of this command to direct the router to redistribute routes learned from OSPF into BGP.

Use the **delete** form of this command to prevent redistribution of routes learned from OSPF into BGP.

Use the **show** form of this command to view route redistribution configuration settings.

protocols bgp <asn> redistribute rip

Redistributes routes learned from RIP into BGP.

Syntax

```
set protocols bgp asn redistribute rip [metric metric | route-map map-name]  
delete protocols bgp asn redistribute rip [metric | route-map]  
show protocols bgp asn redistribute
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    redistribute {  
      rip {  
        metric metric  
        route-map map-name  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>metric</i>	Optional. The metric to be applied to redistributed RIP routes.
<i>map-name</i>	Optional. The name of a configured route map to be used for redistributing RIP routes.

Default

When this command has not been set, routes learned from RIP are not distributed into BGP.

Usage Guidelines

Use this command to redistribute Routing Information Protocol (RIP) routes into BGP.

Use the **set** form of this command to direct the router to redistribute routes learned from RIP into BGP.

Use the **delete** form of this command to prevent redistribution of routes learned from RIP into BGP.

Use the **show** form of this command to view route redistribution configuration settings.

protocols bgp <asn> redistribute static

Redistributes static routes into BGP.

Syntax

```
set protocols bgp asn redistribute static [metric metric | route-map map-name]  
delete protocols bgp asn redistribute static [metric | route-map]  
show protocols bgp asn redistribute
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
    bgp asn {  
        redistribute {  
            static {  
                metric metric  
                route-map map-name  
            }  
        }  
    }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>metric</i>	Optional. The metric to be applied to redistributed static routes.
<i>map-name</i>	Optional. The name of a configured route map to be used for redistributing static routes.

Default

When this command has not been set static routes are not distributed into BGP.

Usage Guidelines

Use this command to redistribute static routes into BGP.

Use the **set** form of this command to direct the router to redistribute static routes into BGP.

Use the **delete** form of this command to prevent redistribution of static routes into BGP.

Use the **show** form of this command to view route redistribution configuration settings.

Chapter 8: Route Server

This chapter describes commands for BGP route server.

This chapter presents the following topics:

- [Route Server Commands](#)

Route Server Commands

This chapter contains the following commands.

Configuration Commands	
<code>protocols bgp <asn> neighbor <id> address-family ipv6-unicast route-server-client</code>	Specify this neighbor as a route server client.
<code>protocols bgp <asn> neighbor <id> route-server-client</code>	Defines the local router as a BGP route server, with the neighbor as a route server client.
<code>protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast route-server-client</code>	Specify this peer group as a route server client.
<code>protocols bgp <asn> peer-group <group-name> route-server-client</code>	Specify this peer group as a route server client.
Operational Commands	
<code>show ip bgp ipv4 unicast rsclient summary</code>	Displays IPv4 unicast BGP route server client summary information.
<code>show ip bgp rsclient <address></code>	Displays BGP route server client information.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast route-server-client

Specify this neighbor as a route server client.

Syntax

```
set protocols bgp asn neighbor id address-family ipv6-unicast route-server-client
delete protocols bgp asn neighbor id address-family ipv6-unicast route-server-client
show protocols bgp asn neighbor id address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          route-server-client
        }
      }
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Default

Disabled.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to configure the BGP neighbor as a route server client.

Use the **delete** form of this command to remove the configuration.

Use the **show** form of this command to view the configuration.

protocols bgp <asn> neighbor <id> route-server-client

Defines the local router as a BGP route server, with the neighbor as a route server client.

Syntax

```
set protocols bgp asn neighbor id route-server-client
delete protocols bgp asn neighbor id route-server-client
show protocols bgp asn neighbor id
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    neighbor id {
      route-server-client
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>id</i>	Mandatory. The IP address of a BGP neighbor or the name of a BGP peer group.

Default

There is no route server in the autonomous system.

Usage Guidelines

Use the set form of this command to make the local router a BGP route server, and to designate the specified neighbor as a route server client.

Use the **delete** form of this command to remove the neighbor as a route server client.
Use the **show** form of this command to view BGP neighbor configuration settings.

protocols bgp <asn> peer-group <group-name> address-family ipv6-unicast route-server-client

Specify this peer group as a route server client.

Syntax

```
set protocols bgp asn peer-group group-name address-family ipv6-unicast  
route-server-client
```

```
delete protocols bgp asn peer-group group-name address-family ipv6-unicast  
route-server-client
```

```
show protocols bgp asn peer-group group-name address-family ipv6-unicast
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {  
  bgp asn {  
    peer-group group-name {  
      address-family {  
        ipv6-unicast {  
          route-server-client  
        }  
      }  
    }  
  }  
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

Default

Disabled.

Usage Guidelines

This command applies to IPv6 unicast routes only.

Use the **set** form of this command to configure the BGP peer group as a route server client.

Use the **delete** form of this command to remove the configuration.

Use the **show** form of this command to view the configuration.

protocols bgp <asn> peer-group <group-name> route-server-client

Specify this peer group as a route server client.

Syntax

```
set protocols bgp asn peer-group group-name route-server-client
delete protocols bgp asn peer-group group-name route-server-client
show protocols bgp asn peer-group group-name
```

Command Mode

Configuration mode.

Configuration Statement

```
protocols {
  bgp asn {
    peer-group group-name {
      route-server-client
    }
  }
}
```

Parameters

<i>asn</i>	Mandatory. The number for the AS in which this router resides. The range of values is 1 to 4294967294. The subrange 64512 to 65534 is reserved for private autonomous systems.
<i>group-name</i>	Mandatory. The peer group name.

Default

Disabled.

Usage Guidelines

Use the set form of this command to configure the BGP peer group as a route server client.

Use the **delete** form of this command to remove the configuration.

Use the **show** form of this command to view the configuration.

show ip bgp ipv4 unicast rsclient summary

Displays IPv4 unicast BGP route server client summary information.

Syntax

```
show ip bgp ipv4 unicast rsclient summary
```

Command Mode

Operational mode.

Parameters

None.

Default

None.

Usage Guidelines

Use this command to display IPv4 unicast BGP route server client summary information.

show ip bgp rsclient <address>

Displays BGP route server client information.

Syntax

```
show ip bgp rsclient address [ipv4 | ipv4net | summary]
```

Command Mode

Operational mode.

Parameters

<i>ipv4</i>	Optional. Displays routes for the route server client at the specified IPv4 address.
<i>ipv4net</i>	Optional. Displays routes for route server clients on the specified IPv4 network.
summary	Optional. Displays summary information for all route server clients.

Default

Detailed information is displayed for all route server clients.

Usage Guidelines

Use this command to display BGP route server client information.

Glossary of Acronyms

ACL	access control list
ADSL	Asymmetric Digital Subscriber Line
AMI	Amazon Machine Image
API	Application Programming Interface
AS	autonomous system
ARP	Address Resolution Protocol
AWS	Amazon Web Services
BGP	Border Gateway Protocol
BIOS	Basic Input Output System
BPDU	Bridge Protocol Data Unit
CA	certificate authority
CCMP	AES in counter mode with CBC-MAC
CHAP	Challenge Handshake Authentication Protocol
CLI	command-line interface
DDNS	dynamic DNS
DHCP	Dynamic Host Configuration Protocol
DHCPv6	Dynamic Host Configuration Protocol version 6

DLCI	data-link connection identifier
DMI	desktop management interface
DMZ	demilitarized zone
DN	distinguished name
DNS	Domain Name System
DSCP	Differentiated Services Code Point
DSL	Digital Subscriber Line
eBGP	external BGP
EBS	Amazon Elastic Block Storage
EC2	Amazon Elastic Compute Cloud
EGP	Exterior Gateway Protocol
ECMP	equal-cost multipath
ESP	Encapsulating Security Payload
FIB	Forwarding Information Base
FTP	File Transfer Protocol
GRE	Generic Routing Encapsulation
HDLC	High-Level Data Link Control
I/O	Input/Output
ICMP	Internet Control Message Protocol
IDS	Intrusion Detection System
IEEE	Institute of Electrical and Electronics Engineers
IGP	Interior Gateway Protocol
IPS	Intrusion Protection System
IKE	Internet Key Exchange
IP	Internet Protocol
IPOA	IP over ATM

IPsec	IP security
IPv4	IP Version 4
IPv6	IP Version 6
ISP	Internet Service Provider
KVM	Kernel-Based Virtual Machine
L2TP	Layer 2 Tunneling Protocol
LACP	Link Aggregation Control Protocol
LAN	local area network
LDAP	Lightweight Directory Access Protocol
LLDP	Link Layer Discovery Protocol
MAC	medium access control
MIB	Management Information Base
MLPPP	multilink PPP
MRRU	maximum received reconstructed unit
MTU	maximum transmission unit
NAT	Network Address Translation
ND	Neighbor Discovery
NIC	network interface card
NTP	Network Time Protocol
OSPF	Open Shortest Path First
OSPFv2	OSPF Version 2
OSPFv3	OSPF Version 3
PAM	Pluggable Authentication Module
PAP	Password Authentication Protocol
PAT	Port Address Translation
PCI	peripheral component interconnect

PKI	Public Key Infrastructure
PPP	Point-to-Point Protocol
PPPoA	PPP over ATM
PPPoE	PPP over Ethernet
PPTP	Point-to-Point Tunneling Protocol
PTMU	Path Maximum Transfer Unit
PVC	permanent virtual circuit
QoS	quality of service
RADIUS	Remote Authentication Dial-In User Service
RHEL	Red Hat Enterprise Linux
RIB	Routing Information Base
RIP	Routing Information Protocol
RIPng	RIP next generation
Rx	receive
S3	Amazon Simple Storage Service
SLAAC	Stateless Address Auto-Configuration
SNMP	Simple Network Management Protocol
SMTP	Simple Mail Transfer Protocol
SONET	Synchronous Optical Network
SSH	Secure Shell
SSID	Service Set Identifier
STP	Spanning Tree Protocol
TACACS+	Terminal Access Controller Access Control System Plus
TBF	Token Bucket Filter
TCP	Transmission Control Protocol
TKIP	Temporal Key Integrity Protocol

ToS	Type of Service
TSS	TCP Maximum Segment Size
Tx	transmit
UDP	User Datagram Protocol
VHD	virtual hard disk
vif	virtual interface
VLAN	virtual LAN
VPC	Amazon virtual private cloud
VPN	Virtual Private Network
VRRP	Virtual Router Redundancy Protocol
WAN	wide area network
WAP	wireless access point
WPA	Wired Protected Access
