



BGP Configuration Guide, 17.2.0

Contents

About This Guide.....	21
BGP Configuration.....	22
BGP overview.....	22
iBGP and eBGP.....	22
BGP ID selection process.....	23
BGP path selection process.....	24
Scalability of BGP.....	25
Route flapping and flap damping.....	28
AS paths.....	28
BGP communities.....	28
Peer groups.....	29
BGP multipath support.....	30
IPv4 and IPv6 support.....	30
VPNv4 support.....	31
Supported standards.....	33
Configuring BGP.....	33
Basic iBGP configuration.....	33
Verifying the iBGP configuration.....	41



- Basic eBGP configuration..... 42
- Verifying the eBGP configuration..... 44
- Originating a route to eBGP neighbors..... 45
- Verifying the route origination..... 47
- Inbound route filtering..... 49
- Outbound route filtering..... 54
- Confederations..... 60
- Verifying the confederation..... 67
- Route reflectors..... 70
- Verifying the route reflector..... 77
- Route redirection..... 80
- Configuring IPv6 routing with BGP..... 80
 - Configure IPv6 BGP peers..... 80
 - Advertise connected networks..... 81
 - Confirm BGP neighbors..... 82
 - Confirm the peer session..... 82
 - Confirm advertised routes..... 82
 - Advertise IPv4 prefix for IPv6 enabled BGP peers..... 84
- Configuring symmetric routing with VRRP and BGP..... 92
- Advertising IPv6 and IPv4 prefixes over IPv4 and IPv6 eBGP peers..... 94



Setting BGP MED based on the IGP metric..... 97

Global and Router-Specific Commands..... 100

- monitor protocol bgp disable all..... 100
- monitor protocol bgp disable <id> bfd..... 100
- monitor protocol bgp disable <id> dampening..... 100
- monitor protocol bgp disable events..... 100
- monitor protocol bgp disable filters..... 101
- monitor protocol bgp disable fsm..... 101
- monitor protocol bgp disable keepalives..... 101
- monitor protocol bgp disable <id> msdp..... 101
- monitor protocol bgp disable <id> nsm..... 102
- monitor protocol bgp disable rib..... 102
- monitor protocol bgp disable updates..... 102
- monitor protocol bgp enable <id> all..... 102
- monitor protocol bgp enable <id> bfd..... 103
- monitor protocol bgp enable <id> dampening..... 103
- monitor protocol bgp enable events..... 103
- monitor protocol bgp enable filters..... 103
- monitor protocol bgp enable fsm..... 103
- monitor protocol bgp enable keepalives..... 104
- monitor protocol bgp enable <id> msdp..... 104
- monitor protocol bgp enable <id> nsm..... 104



monitor protocol bgp enable rib..... 104

monitor protocol bgp enable updates..... 105

protocols bgp..... 105

protocols bgp address-family ipv4-unicast..... 106

protocols bgp address-family ipv4-unicast aggregate-address..... 106

protocols bgp address-family ipv4-unicast network..... 107

protocols bgp address-family ipv4-unicast parameters dampening..... 108

protocols bgp address-family ipv4-unicast parameters distance global..... 109

protocols bgp address-family ipv4-unicast parameters network-synchronization..... 110

protocols bgp address-family ipv4-unicast parameters synchronization..... 111

protocols bgp address-family ipv6-unicast..... 111

protocols bgp address-family ipv6-unicast aggregate-address..... 112

protocols bgp address-family ipv6-unicast network..... 113

protocols bgp address-family ipv6-unicast network..... 113

protocols bgp address-family ipv6-unicast parameters dampening..... 114

protocols bgp address-family ipv6-unicast parameters distance global..... 115

protocols bgp address-family ipv6-unicast parameters network-synchronization..... 116

protocols bgp address-family ipv6-unicast parameters synchronization..... 117

protocols bgp asn address-family vpv4-unicast..... 118

protocols bgp asn log..... 118

protocols bgp asn log update..... 119



protocols bgp parameters always-compare-med..... 120

protocols bgp parameters bestpath as-path..... 120

protocols bgp parameters bestpath compare-routerid..... 121

protocols bgp parameters bestpath igp-metric-ignore..... 122

protocols bgp parameters bestpath med..... 122

protocols bgp <asn> parameters maximum-as-limit <number>..... 123

protocols bgp parameters med-out-delay..... 124

protocols bgp parameters default local-pref..... 124

protocols bgp parameters deterministic-med..... 125

protocols bgp parameters enforce-first-as..... 126

protocols bgp parameters extended-asn-capability..... 127

protocols bgp parameters graceful-restart..... 127

protocols bgp parameters log-neighbor-changes..... 128

protocols bgp parameters maximum-paths ebgp..... 129

protocols bgp parameters maximum-paths ibgp..... 130

protocols bgp parameters no-client-to-client-reflection..... 131

protocols bgp parameters no-fast-external-failover..... 132

protocols bgp parameters no-rtm..... 132

protocols bgp parameters router-id..... 133

protocols bgp parameters scan-time..... 133

protocols bgp timers..... 134

reset ip bgp..... 135



reset ip bgp ipv4 unicast.....	135
reset ip bgp ipv6 unicast.....	136
reset ip bgp all.....	136
reset ip bgp all ipv4 unicast.....	136
reset ip bgp all ipv4 unicast soft.....	136
reset ip bgp all ipv6 unicast.....	137
reset ip bgp all ipv6 unicast soft.....	137
reset ip bgp all soft.....	137
reset ip bgp ipv4 unicast dampening.....	138
reset ip bgp ipv6 unicast dampening.....	138
show ip bgp.....	138
show ip bgp attribute-info.....	138
show ip bgp cidr-only.....	139
show ip bgp cluster-ids.....	139
show ip bgp community.....	139
show ip bgp community-info.....	139
show ip bgp community-list.....	139
show ip bgp dampening dampened-paths.....	140
show ip bgp filter-list.....	140
show ip bgp ipv4 unicast.....	140
show ip bgp ipv4 unicast cidr-only.....	141



show ip bgp ipv4 unicast community..... 141

show ip bgp ipv4 unicast community-list..... 141

show ip bgp ipv4 unicast filter-list..... 141

show ip bgp ipv4 unicast paths..... 142

show ip bgp ipv4 unicast prefix-list..... 142

show ip bgp ipv4 unicast regexp..... 142

show ip bgp ipv4 unicast route-map..... 142

show ip bgp ipv4 unicast summary..... 142

show ip bgp ipv6 unicast..... 143

show ip bgp ipv6 unicast community..... 143

show ip bgp ipv6 unicast community-list..... 143

show ip bgp ipv6 unicast dampening..... 144

show ip bgp ipv6 unicast filter-list..... 144

show ip bgp ipv6 unicast prefix-list..... 144

show ip bgp ipv6 unicast regexp..... 144

show ip bgp paths..... 145

show ip bgp prefix-list..... 145

show ip bgp regexp..... 145

show ip bgp route-map..... 145

show ip bgp scan..... 146

show ip route bgp..... 146

show monitoring protocols bgp..... 146



- Route Reflection..... 148
 - Overview..... 148
 - BGP multiple cluster IDs..... 148
 - Configuring multiple cluster IDs..... 149
- Route Reflection Commands..... 153
 - protocols bgp neighbor address-family ipv4-unicast route-reflector-client..... 153
 - protocols bgp neighbor address-family ipv6-unicast route-reflector-client..... 154
 - protocols bgp neighbor cluster-id..... 155
 - protocols bgp parameters cluster-id..... 155
 - protocols bgp parameters no-client-to-client-reflection all..... 156
 - protocols bgp parameters no-client-to-client-reflection cluster-id any..... 157
 - protocols bgp parameters no-client-to-client-reflection cluster-id..... 157
 - protocols bgp peer-group address-family ipv4-unicast route-reflector-client..... 158
 - protocols bgp peer-group address-family ipv6-unicast route-reflector-client..... 159
 - protocols bgp peer-group cluster-id..... 159
- Confederations Commands..... 161
 - protocols bgp parameters confederation identifier..... 161
 - protocols bgp parameters confederation peers..... 161
- Synchronization..... 163
 - Basic synchronization configuration..... 163
 - Verifying the configuration..... 167



Graceful restart..... 170

Neighbors Commands..... 171

 protocols bgp neighbor..... 171

 protocols bgp neighbor address-family ipv4-unicast..... 171

 protocols bgp neighbor address-family ipv4-unicast allowas-in number..... 172

 protocols bgp neighbor address-family ipv4-unicast attribute-unchanged..... 173

 protocols bgp neighbor address-family ipv4-unicast capability graceful-restart..... 174

 protocols bgp neighbor address-family ipv4-unicast capability orf prefix-list..... 175

 protocols bgp neighbor address-family ipv4-unicast default-originate route-map..... 176

 protocols bgp neighbor address-family ipv4-unicast disable-send-community..... 177

 protocols bgp neighbor address-family ipv4-unicast distribute-list export..... 178

 protocols bgp neighbor address-family ipv4-unicast distribute-list import..... 179

 protocols bgp neighbor address-family ipv4-unicast filter-list export..... 180

 protocols bgp neighbor address-family ipv4-unicast filter-list import..... 180

 protocols bgp neighbor address-family ipv4-unicast maximum-prefix..... 181

 protocols bgp neighbor address-family ipv4-unicast nexthop-self..... 182

 protocols bgp neighbor address-family ipv4-unicast peer-group..... 183

 protocols bgp neighbor address-family ipv4-unicast prefix-list export..... 184

 protocols bgp neighbor address-family ipv4-unicast prefix-list import..... 184

 protocols bgp neighbor address-family ipv4-unicast remove-private-as..... 185

 protocols bgp neighbor address-family ipv4-unicast route-map export..... 186

 protocols bgp neighbor address-family ipv4-unicast route-map import..... 187



protocols bgp neighbor address-family ipv4-unicast soft-reconfiguration inbound..... 188

protocols bgp neighbor address-family ipv4-unicast unsuppress-map..... 189

protocols bgp neighbor address-family ipv4-unicast weight..... 189

protocols bgp neighbor address-family ipv6-unicast..... 190

protocols bgp neighbor address-family ipv6-unicast allowas-in..... 191

protocols bgp neighbor address-family ipv6-unicast attribute-unchanged..... 192

protocols bgp neighbor address-family ipv6-unicast capability graceful-restart..... 193

protocols bgp neighbor address-family ipv6-unicast capability orf..... 193

protocols bgp neighbor address-family ipv6-unicast default-originate..... 194

protocols bgp neighbor address-family ipv6-unicast disable-send-community..... 195

protocols bgp neighbor address-family ipv6-unicast distribute-list export..... 196

protocols bgp neighbor address-family ipv6-unicast distribute-list import..... 197

protocols bgp neighbor address-family ipv6-unicast filter-list export..... 198

protocols bgp neighbor address-family ipv6-unicast filter-list import..... 199

protocols bgp neighbor address-family ipv6-unicast maximum-prefix..... 199

protocols bgp neighbor address-family ipv6-unicast nexthop-self..... 201

protocols bgp neighbor address-family ipv6-unicast peer-group..... 201

protocols bgp neighbor address-family ipv6-unicast prefix-list export..... 202

protocols bgp neighbor address-family ipv6-unicast prefix-list import..... 203

protocols bgp neighbor address-family ipv6-unicast remove-private-as..... 204

protocols bgp neighbor address-family ipv6-unicast route-map export..... 204



protocols bgp neighbor address-family ipv6-unicast route-map import..... 205

protocols bgp neighbor address-family ipv6-unicast route-reflector-client..... 206

protocols bgp neighbor address-family ipv6-unicast soft-reconfiguration inbound..... 207

protocols bgp neighbor address-family ipv6-unicast unsuppress-map..... 207

protocols bgp neighbor address-family ipv6-unicast weight..... 208

protocols bgp asn neighbor id address-family vpnv4-unicast..... 209

protocols bgp asn neighbor id address-family vpnv4-unicast capability graceful-restart..... 209

protocols bgp asn neighbor id address-family vpnv4-unicast distribute-list..... 210

protocols bgp asn neighbor id address-family vpnv4-unicast distribute-list export access-list-name..... 211

protocols bgp asn neighbor id address-family vpnv4-unicast distribute-list export as-path-list-name..... 212

protocols bgp asn neighbor id address-family vpnv4-unicast filter-list import as-path-list-name..... 213

protocols bgp asn neighbor id address-family vpnv4-unicast maximum-prefix..... 213

protocols bgp asn neighbor id address-family vpnv4-unicast prefix-list export prefix-list-name..... 214

protocols bgp asn neighbor id address-family vpnv4-unicast prefix-list..... 215

protocols bgp asn neighbor id address-family vpnv4-unicast route-map..... 216

protocols bgp asn neighbor id address-family vpnv4-unicast route-map import map-name..... 217

protocols bgp asn neighbor id address-family vpnv4-unicast route-reflector-client..... 218

protocols bgp neighbor advertisement-interval..... 218

protocols bgp neighbor capability dynamic..... 219

protocols bgp neighbor description..... 220

protocols bgp neighbor disable-capability-negotiation..... 220

protocols bgp neighbor ebgp-multihop..... 221



protocols bgp neighbor enforce-multihop..... 222

protocols bgp neighbor local-as..... 222

protocols bgp asn neighbor <id> log..... 223

protocols bgp neighbor med-out igp..... 224

protocols bgp neighbor med-out igp delay- updates..... 225

protocols bgp neighbor med-out minimum-igp..... 225

protocols bgp neighbor override-capability..... 226

protocols bgp neighbor passive..... 227

protocols bgp neighbor password..... 227

protocols bgp neighbor port..... 228

protocols bgp neighbor remote-as..... 228

protocols bgp neighbor shutdown..... 229

protocols bgp neighbor strict-capability-match..... 230

protocols bgp neighbor timers..... 230

protocols bgp neighbor ttl-security hops..... 231

protocols bgp neighbor update-source..... 232

protocols bgp neighbor vrrp-failover vrrp-group med..... 233

protocols bgp neighbor vrrp-failover vrrp-group prepend-as..... 234

protocols bgp neighbor vrrp-failover vrrp-group route-map..... 235

reset ip bgp external..... 236

reset ip bgp external ipv4 unicast..... 237



reset ip bgp external ipv6 unicast..... 237

show ip bgp ipv4 unicast neighbors..... 238

show ip bgp ipv4 unicast neighbors advertised-routes..... 238

show ip bgp ipv4 unicast neighbors received prefix-filter..... 238

show ip bgp ipv4 unicast neighbors received-routes..... 239

show ip bgp ipv4 unicast neighbors routes..... 239

show ip bgp ipv6 unicast neighbors..... 239

show ip bgp ipv6 unicast neighbors advertised-routes..... 239

show ip bgp ipv6 unicast neighbors received-routes..... 240

show ip bgp ipv6 unicast neighbors routes..... 240

show ip bgp neighbors..... 240

show ip bgp neighbors advertised-routes..... 240

show ip bgp neighbors received prefix-filter..... 240

show ip bgp neighbors received-routes..... 241

show ip bgp neighbors routes..... 241

show ip bgp vpnv4 unicast..... 241

show ip bgp vpnv4 unicast prefix..... 241

show ip bgp vpnv4 unicast dampening..... 242

show ip bgp vpnv4 unicast filter-list..... 242

show ip bgp vpnv4 unicast neighbors..... 242

show ip bgp vpnv4 unicast neighbors id advertised-routes..... 242

show ip bgp vpnv4 unicast neighbors id received-routes..... 243



- show ip bgp vpnv4 unicast neighbors id routes..... 243
- show ip bgp vpnv4 unicast neighbors prefix-list..... 243
- show ip bgp vpnv4 unicast neighbors regex..... 243
- show ip bgp vpnv4 unicast neighbors summary..... 243
- Peer Groups Commands..... 245
 - protocols bgp peer-group..... 245
 - protocols bgp peer-group address-family..... 245
 - protocols bgp peer-group address-family ipv4-unicast..... 246
 - protocols bgp peer-group address-family ipv4-unicast allowas-in..... 247
 - protocols bgp peer-group address-family ipv4-unicast attribute-unchanged..... 248
 - protocols bgp peer-group address-family ipv4-unicast capability graceful-restart..... 249
 - protocols bgp peer-group address-family ipv4-unicast capability orf..... 250
 - protocols bgp peer-group address-family ipv4-unicast default-originate..... 251
 - protocols bgp peer-group address-family ipv4-unicast disable-send-community..... 252
 - protocols bgp peer-group address-family ipv4-unicast distribute-list import..... 253
 - protocols bgp peer-group address-family ipv4-unicast filter-list export..... 254
 - protocols bgp peer-group address-family ipv4-unicast filter-list import..... 254
 - protocols bgp peer-group address-family ipv4-unicast maximum-prefix..... 255
 - protocols bgp peer-group address-family ipv4-unicast nexthop-self..... 256
 - protocols bgp peer-group address-family ipv4-unicast prefix-list export..... 257
 - protocols bgp peer-group address-family ipv4-unicast prefix-list import..... 258



protocols bgp peer-group address-family ipv4-unicast remove-private-as..... 258

protocols bgp peer-group address-family ipv4-unicast route-map export..... 259

protocols bgp peer-group address-family ipv4-unicast route-map import..... 260

protocols bgp peer-group address-family ipv4-unicast soft-reconfiguration inbound..... 261

protocols bgp peer-group address-family ipv4-unicast unsuppress-map..... 262

protocols bgp peer-group address-family ipv4-unicast weight..... 263

protocols bgp peer-group address-family ipv6-unicast..... 263

protocols bgp peer-group address-family ipv6-unicast allowas-in..... 264

protocols bgp peer-group address-family ipv6-unicast attribute-unchanged..... 265

protocols bgp peer-group address-family ipv6-unicast capability graceful-restart..... 266

protocols bgp peer-group address-family ipv6-unicast capability orf..... 267

protocols bgp peer-group address-family ipv6-unicast capability orf prefix-list..... 268

protocols bgp peer-group address-family ipv6-unicast default-originate..... 269

protocols bgp peer-group address-family ipv6-unicast disable-send-community..... 270

protocols bgp peer-group address-family ipv6-unicast distribute-list export..... 271

protocols bgp peer-group address-family ipv6-unicast distribute-list import..... 272

protocols bgp peer-group address-family ipv6-unicast filter-list export..... 273

protocols bgp peer-group address-family ipv6-unicast filter-list import..... 273

protocols bgp peer-group address-family ipv6-unicast maximum-prefix..... 274

protocols bgp peer-group address-family ipv6-unicast nexthop-self..... 275

protocols bgp peer-group address-family ipv6-unicast prefix-list export..... 276

protocols bgp peer-group address-family ipv6-unicast prefix-list import..... 277



protocols bgp peer-group address-family ipv6-unicast remove-private-as..... 278

protocols bgp peer-group address-family ipv6-unicast route-map export..... 278

protocols bgp peer-group address-family ipv6-unicast route-map import..... 279

protocols bgp peer-group address-family ipv6-unicast soft-reconfiguration inbound..... 280

protocols bgp peer-group address-family ipv6-unicast unsuppress-map..... 281

protocols bgp peer-group address-family ipv6-unicast weight..... 282

protocols bgp peer-group advertisement-interval..... 283

protocols bgp peer-group capability dynamic..... 283

protocols bgp peer-group capability graceful-restart..... 284

protocols bgp peer-group description..... 285

protocols bgp peer-group disable-capability-negotiation..... 285

protocols bgp peer-group ebgp-multihop..... 286

protocols bgp peer-group local-as..... 287

protocols bgp peer-group med-out igp..... 287

protocols bgp peer-group med-out igp delay-updates..... 288

protocols bgp peer-group med-out minimum-igp..... 288

protocols bgp peer-group override-capability..... 289

protocols bgp peer-group passive..... 290

protocols bgp peer-group password..... 290

protocols bgp peer-group remote-as..... 291

protocols bgp peer-group shutdown..... 292



- protocols bgp peer-group strict-capability-match..... 292
- protocols bgp peer-group ttl-security hops..... 293
- protocols bgp peer-group update-source..... 294
- reset ip bgp peer-group..... 294
- reset ip bgp peer-group ipv4 unicast..... 295
- reset ip bgp peer-group ipv6 unicast..... 295
- Route Redistribution Commands..... 297
 - protocols bgp address-family ipv4-unicast redistribute connected..... 297
 - protocols bgp address-family ipv4-unicast redistribute kernel..... 298
 - protocols bgp address-family ipv4-unicast redistribute ospf..... 299
 - protocols bgp address-family ipv4-unicast redistribute rip..... 299
 - protocols bgp address-family ipv4-unicast redistribute static..... 300
 - protocols bgp address-family ipv6-unicast redistribute connected..... 301
 - protocols bgp address-family ipv6-unicast redistribute kernel..... 302
 - protocols bgp address-family ipv6-unicast redistribute ospfv3..... 303
 - protocols bgp address-family ipv6-unicast redistribute ripng..... 304
 - protocols bgp address-family ipv6-unicast redistribute static..... 305
- Route Server Commands..... 307
 - protocols bgp neighbor address-family ipv4-unicast route-server-client..... 307
 - protocols bgp neighbor address-family ipv6-unicast route-server-client..... 307
 - protocols bgp peer-group address-family ipv4-unicast route-server-client..... 308
 - protocols bgp peer-group address-family ipv6-unicast route-server-client..... 309



VRF Support for BGP..... 310

 VRF support for BGP..... 310

 Configuring BGP on a routing instance..... 310

 Configuring inter-VRF BGP on routing instances..... 312

 Command support for VRF routing instances..... 314

List of Acronyms..... 318

Copyright Statement

© 2017 AT&T Intellectual Property. All rights reserved. AT&T and Globe logo are registered trademarks of AT&T Intellectual Property. All other marks are the property of their respective owners.

The training materials and other content provided herein for assistance in training on the Vyatta vRouter may have references to Brocade as the Vyatta vRouter was formerly a Brocade product prior to AT&T's acquisition of Vyatta. Brocade remains a separate company and is not affiliated to AT&T.



About This Guide

This guide describes how to configure BGP on AT&T products that run on the AT&T Vyatta Network Operating System (referred to as a virtual router, vRouter, or router in the guide).



BGP Configuration

BGP overview

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 from one AS to another.

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. A BGP peer that is configured in a different AS than another peer 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 BGP table of the router. (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 AS 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 AS, 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 through AS 2, then from AS 2 through AS 3, and finally from AS 3 through 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. 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 by using `protocols bgp <asn> parameters router-id <id>` (page 133). If explicitly assigned, then this is the BGP ID that should be used. If it is not assigned explicitly, then the system choose one automatically. If a loopback address is configured as is not 127.0.0.1, it be used. If the loopback address is not used, then the largest IP address assigned to an interface in the system configuration 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.

Table 1: Best practice for setting the BGP ID

Router	Step	Commands
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>



Router	Step	Commands
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, 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 does not consider paths that contain a NEXT_HOP value that is not reachable via an entry in the RIB. For all valid paths, the AT&T Vyatta vRouter uses a BGP path selection process based on decision process described in RFC 4271, section 9.1 and RFC 5004. 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 the paths learned via eBGP over paths learned via iBGP.
- IGP METRIC. Prefer the 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 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 `show ip bgp` ([page 138](#)) 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 AS, which can have hundreds of iBGP routers. To scale well in a large AS, 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-AS 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 through 65534. Within a sub-AS, all the standard iBGP rules, including full-mesh peering, apply. The connections between confederation sub-AS use eBGP peering. One or more eBGP connections can be made between each sub-AS. The sub-AS are grouped as a confederation, which advertises as a single AS to external peers.

The following figure 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: iBGP full mesh



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

**Figure 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.

The following figure shows again the full mesh of iBGP connections in even a moderately sized AS.




Figure 3: iBGP full mesh



The following figure shows how introducing route reflection into the AS dramatically reduces the number of iBGP connections required within the AS.

Figure 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 AS. Each AS number represents an AS (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 AS 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
- user-defined or private communities

The AT&T Vyatta vRouter 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 AT&T Vyatta vRouter 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.

BGP community list is a user defined BGP communities attribute list. BGP community list can be used for matching or manipulating BGP communities attribute in updates.

The AT&T Vyatta vRouter also supports the following user-defined community lists:

- Standard
- Extended

Furthermore, standard and extended community lists support standard and expanded lists. A standard community list ranges from 1 through 99 and an expanded community list ranges from 100 through 199.

Standard Community List

- Standard list specifies standard community attributes.
- Expanded list specifies regular expression.

Extended Community List

- Standard list specifies route-target and site-of-origin.
- Expanded list specifies regular expression.

Note: Community lists do not support these features: named community list and appending of a community list on a received prefix with new communities.

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 BGP peer. The peer group configuration applies to all members of the peer group.

Note: A peer group can have a maximum of 32 BGP peers.

A neighbor can be associated with a single peer group although it can selectively inherit the properties of the given peer group per address family.

Depending on the IP address family type of the peer group, peers are added to a peer group by using `protocols bgp <asn> neighbor <id> address-family ipv4-unicast peer-group <group-name>` (page 183) or `protocols bgp <asn> neighbor <id> address-family ipv6-unicast peer-group <group-name>` (page 201).

If the neighbor is associated under an IPv4 address family and an IPv6 address family by using `protocols bgp <asn> neighbor <id> address-family ipv4-unicast peer-group <group-name>` (page 183) and `protocols bgp <asn> neighbor <id> address-family ipv6-unicast peer-group <group-name>` (page 201), then all of the following applies:



- Global or address-family agnostic configurations of the peer group are applied to the neighbor.
- IPv4 address-family specific configurations of the peer group are applied to the neighbor.
- IPv6 address-family specific configurations of the peer group are applied to the neighbor.

If the neighbor is associated under an IPv4 address family and not an IPv6 address family (using `protocols bgp <asn> neighbor <id> address-family ipv4-unicast peer-group <group-name>` ([page 183](#))), then the following applies:

- Global or address-family agnostic configurations of the peer group are applied to the neighbor and IPv4 address-family specific configurations of the peer group are applied to the neighbor.
- The IPv6 address-family specific configurations of the peer group are not applied to the neighbor.

If the neighbor is associated under an IPv6 address family and not an IPv4 address family (using `protocols bgp <asn> neighbor <id> address-family ipv6-unicast peer-group <group-name>` ([page 201](#))), then the following applies:

- Global or address-family agnostic configurations of the peer group are applied to the neighbor and IPv6 address-family specific configurations of the peer group are applied to the neighbor.
- The IPv4 address-family specific configurations of the peer group are not applied to the neighbor.

BGP multipath support

In general, BGP selects a single best path to a given destination. BGP multipath support allows 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> parameters maximum-paths ebgp <max-paths>` ([page 129](#)) and `protocols protocols bgp <asn> parameters maximum-paths ibgp <max-paths>` ([page 130](#)).

IPv4 and IPv6 support

The 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 you have enabled either of the following:

- IPv4 unicast address family for a neighbor (using `protocols bgp <asn> neighbor <id> address-family ipv4-unicast` ([page 171](#)))
- IPv4 unicast address family for the peer group to which the neighbor belongs (using `protocols bgp <asn> peer-group <name> address-family ipv4-unicast` ([page 246](#)))

IPv6 routes can be exchanged after you have enabled either of the following:

- IPv6 unicast address family for a neighbor (using `protocols bgp <asn> neighbor <id> address-family ipv6-unicast` ([page 190](#)))
- IPv6 unicast address family for the peer group to which the neighbor belongs (using `protocols bgp <asn> peer-group <name> address-family ipv6-unicast` ([page 263](#)))

The top-level address family configuration sets family-specific parameters while the neighbor-level family configuration sets support for the given family over that neighbor connection.



The following example defines IPv4 address family parameters, but does not enable any BGP neighbor connection to carry IPv4 address-family routing data.

```
bgp 7777 {
  address-family {
    ipv4-unicast {
      .....
    }
  }
}
```

The following example defines two BGP neighbors, one with an IPv4 address and one with an IPv6 address, that can carry both IPv4 and IPv6 routing information. Note that IPv4 address family support is implicitly defined.

```
bgp 7777 {
  neighbor 3ffe::c5:5 {
    address-family {
      ipv6-unicast {
      }
    }
    remote-as 7777
  }
  neighbor 10.1.1.1 {
    address-family {
      ipv6-unicast {
      }
    }
    remote-as 7777
  }
}
```

VPNv4 support


Route Reflector (RR) in BGP improves the scalability of iBGP sessions between iBGP peers and reduces the connectivity overload between BGP peers. The vRouter provides support for RR for VPNv4 address families in addition to IPv4 and IPv6 address families.

Note the following regarding RR support for VPNv4:

- RR cannot be used in the forwarding path for VPNv4.
- For best performance, AT&T recommends not configuring VRF on the VPNv4 RR.

The following diagram illustrates an RR configuration for reflecting the VPNv4 address family. The vRouter can be used as RR1 or RR2. PE1 and PE2 are third party devices that have full Layer 3 VPN support. The RR is not part of the forwarding path.

All BGP peerings within RR1 and RR2 are in the same AS (100 in the example). The AS is illustrated by the large blue circle in the diagram (RR applies to iBGP peerings only to avoid the full mesh between iBGP peers).

**Figure 5: Sample RR configuration for the VPNv4 address family**

The following commands show how to configure the example in the diagram. MPLS configuration is not necessary on RR1 or RR2, because traffic is sent between PE1 and PE2 over an MPLS enabled network.

```
vyatta@vyatta:~$ set interfaces dataplane dp0s4 address 192.168.1.4/24
vyatta@vyatta:~$ set interfaces dataplane dp0s5 address 192.168.12.4/24
vyatta@vyatta:~$ set interfaces loopback lo address 10.1.1.1/32
vyatta@vyatta:~$ set protocols bgp 100 neighbor 192.168.1.1 address-family vpnv4-unicast route-
reflector-client
vyatta@vyatta:~$ set protocols bgp 100 neighbor 192.168.1.1 remote-as 100
vyatta@vyatta:~$ set protocols bgp 100 neighbor 192.168.12.2 address-family vpnv4-unicast route-
reflector-client
vyatta@vyatta:~$ set protocols bgp 100 neighbor 192.168.12.2 remote-as 100
vyatta@vyatta:~$ set protocols bgp 100 parameters router-id 10.1.1.1
vyatta@vyatta:~$ set protocols ospf area 0 network 10.1.1.1/32
vyatta@vyatta:~$ set protocols ospf area 0 network 192.168.1.0/24
vyatta@vyatta:~$ set protocols ospf area 0 network 192.168.12.0/24
```

The following command displays summary information for the configuration.

```
vyatta@vyatta:~$ show ip bgp vpnv4 unicast summary
BGP router identifier 10.1.1.1, local AS number 100
BGP table version is 1
1 BGP GLOBAL AS-PATH entries
0 BGP community entries

Neighbor      V      AS MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  Statd
192.1268.1.1  4      100    11     8        1     0     0      00:02:45
3
192.1268.12.2 4      100    27    51        1     0     0      00:10:18  1

Total number of neighbors 2

Total number of Established sessions 2
```

The following example displays detailed information for the configuration.

```
vyatta@vyatta:~$ show ip bgp vpnv4 unicast
BGP table version is 1, local router ID is 10.1.1.1
```




```
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
               S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
Route Distinguisher: 1:1
*>i10.4.4.0/24      192.168.1.1          0       100       0       ?
*>i10.7.9.0/24      192.168.12.2         0       100       0       ?

BGP table version is 1, local router ID is 10.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
               S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Route Distinguisher: 2:2
*>i10.8.2.0/24      192.168.1.1          0       100       0       ?

BGP table version is 1, local router ID is 10.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
               S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

Route Distinguisher: 3:3
*>i10.8.3.0/24      192.168.1.1          0       100       0       ?
```

Supported standards

The AT&T Vyatta vRouter implementation of BGP complies with the following standards:

- RFC 1997: *BGP Communities Attribute*
- RFC 2545: *Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing*
- RFC 2796: *BGP Route Reflection - An Alternative to Full Mesh IBGP*
- RFC 2842: *Capabilities Advertisement with BGP-4*
- RFC 3065: *Autonomous System Confederations for BGP*
- RFC 4271: *A Border Gateway Protocol (BGP-4)*
- RFC 4273: *Definitions of Managed Objects for BGP-4*
- RFC 4760: *Multiprotocol Extensions for BGP-4*

Configuring BGP

This section presents sample configurations for BGP.

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 routing table of the router. This requires some form of Internal Gateway Protocol (IGP).

The following figure shows the BGP connections after you have completed the iBGP configuration.



Figure 6: Basic iBGP configuration



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

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

Table 2: Basic iBGP configuration

Router	Step	Commands
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 address-family ipv4- unicast</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 address-family ipv4- unicast</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>



Router	Step	Commands
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 address-family ipv4- unicast</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>
R1	Display the configuration.	<pre>vyatta@R1# show protocols bgp 100 { neighbor 10.0.0.22 { address-family { ipv4-unicast { update-source 10.0.0.11 } } } neighbor 10.0.0.33 { address-family { ipv4-unicast { update-source 10.0.0.11 } } } neighbor 10.0.0.44 { address-family { ipv4-unicast { update-source 10.0.0.11 } } } }</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 address-family ipv4-unicast</pre>
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>



Router	Step	Commands
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 address-family ipv4- unicast</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 address-family ipv4- unicast</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	Activate the neighbor for the IPv4 unicast address family to enable the exchange IPv4 unicast routes.	<pre>vyatta@R2# set protocols bgp 100 neighbor 10.0.0.44 address-family ipv4-unicast</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>



Router	Step	Commands
R2	Display the configuration.	<pre>vyatta@R2# show protocols bgp 100 { neighbor 10.0.0.11 { address-family { ipv4-unicast { update-source 10.0.0.22 } } } neighbor 10.0.0.33 { address-family { ipv4-unicast { update-source 10.0.0.22 } } } neighbor 10.0.0.44 { address-family { ipv4-unicast { update-source 10.0.0.22 } } } neighbor 10.0.0.44 { address-family ipv4-unicast } 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	Advertise 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	Advertise 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>



Router	Step	Commands
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 address-family ipv4- unicast</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 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 address-family ipv4-unicast</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 address-family ipv4- unicast</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>



Router	Step	Commands
R3	Display the configuration.	<pre>vyatta@R3# show protocols bgp 100 { neighbor 10.0.0.11 { address-family { ipv4-unicast { update-source 10.0.0.33 } } } neighbor 10.0.0.22 { address-family { ipv4-unicast { update-source 10.0.0.33 } } } neighbor 10.0.0.44 { address-family { ipv4-unicast { update-source 10.0.0.33 } } } neighbor 10.0.0.44 { address-family ipv4-unicast } 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	Advertise 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	Advertise 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	Advertise 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>



Router	Step	Commands
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>
R4	Configure the dp0p1p1 interface to be a passive interface (i.e. do not advertise our internal network routes over the external network).	<pre>vyatta@R4# set protocols ospf passive-interface dp0p1p1</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 address-family ipv4- unicast</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 address-family ipv4- unicast</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 address-family ipv4- unicast</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>



Router	Step	Commands
R4	Display the configuration.	<pre>vyatta@R4# show protocols bgp 100 { neighbor 10.0.0.11 { address-family { ipv4-unicast { update-source 10.0.0.44 } } } neighbor 10.0.0.22 { address-family { ipv4-unicast { update-source 10.0.0.44 } } } neighbor 10.0.0.33 { address-family { ipv4-unicast { update-source 10.0.0.44 } } } neighbor 10.0.0.44 { address-family ipv4-unicast } 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 dp0p1p1 }</pre>

Verifying the iBGP configuration

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

R1: show ip bgp summary

The following example shows the output of the `show ip bgp summary` command for router R1 at this stage of the configuration.

```
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 be refused.

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 the following figure.



Figure 7: Basic eBGP configuration



This example assumes the following:

- The configuration in the preceding example 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.

Table 3: Basic eBGP configuration

Router	Step	Commands
R1	Create an eBGP peer for R1. The peer is an eBGP peer because it resides in a different AS than this router.	<pre>vyatta@R1# set protocols bgp 100 neighbor 88.88.88.2 remote-as 200</pre>
R1	Activate the neighbor for the IPv4 unicast address family to enable the exchange IPv4 unicast routes.	<pre>vyatta@R1# set protocols bgp 100 neighbor 88.88.88.2 address-family ipv4-unicast</pre>
R1	Commit the configuration.	<pre>vyatta@R1# commit</pre>



Router	Step	Commands
R1	Display the configuration.	<pre>vyatta@R1# show protocols protocols { bgp 100 { neighbor 88.88.88.2 { address-family { ipv4-unicast } } } remote-as 200 }</pre>
R4	Create an eBGP peer for R4. The peer is an eBGP peer because it resides in a different AS than this router.	<pre>vyatta@R4# set protocols bgp 100 neighbor 99.99.99.2 remote-as 300</pre>
R4	Activate the neighbor for the IPv4 unicast address family to enable the exchange IPv4 unicast routes.	<pre>vyatta@R4# set protocols bgp 100 neighbor 99.99.99.2 address-family ipv4-unicast</pre>
R4	Commit the configuration.	<pre>vyatta@R4# commit</pre>
R4	Commit the configuration.	<pre>vyatta@R4# show protocols protocols { bgp 100 { neighbor 99.99.99.2 { address-family { ipv4-unicast } } } remote-as 300 }</pre>

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

The following example shows the output of the `show ip bgp summary` command for router R1 at this stage of the configuration.

```
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.

Originating a route to eBGP neighbors


One of the main requirements for BGP configurations is to originate network prefixes to BGP peers. On the AT&T Vyatta vRouter 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 the following figure.

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 8: Originating a route to eBGP neighbors





This example assumes that the configurations in previous sections have been performed. For more information about configuring an eBGP connection, refer to [Basic eBGP configuration \(page 81\)](#) and for more information about configuring an iBGP connection, refer to [Basic iBGP configuration \(page 81\)](#).

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

Table 4: Originating routes to eBGP neighbors

Router	Step	Commands
R1	Advertise the local network to BGP.	<pre>vyatta@R1# set protocols bgp 100 address-family ipv4-unicast network 172.16.0.0/24</pre>
R1	Commit the configuration.	<pre>vyatta@R1# commit</pre>
R1	Verify the configuration.	<pre>vyatta@R1# show protocols bgp bgp 100 { neighbor 172.16.0.2 { address-family { ipv4-unicast { } } remote-as 100 update-source 172.16.0.1 } neighbor 172.16.0.3 { address-family { ipv4-unicast { } } remote-as 100 update-source 172.16.0.1 } neighbor 172.16.0.4 { address-family { ipv4-unicast { } } remote-as 100 update-source 172.16.0.1 } } }</pre>
V1: AS-200	Advertise the remote -as AS-200 to BGP.	<pre>vyatta@V1# set protocols bgp 200 neighbor 88.88.88.1 remote-as 100 vyatta@V1#set protocols bgp 200 neighbor 88.88.88.1 address-family ipv4-unicast</pre>
V1	Commit the configuration	<pre>vyatta@V1# commit</pre>



Router	Step	Commands
V1	Verify the RA configuration.	<pre>vyatta@V1# show protocols bgp bgp 200 { address-family { ipv4-unicast { network 12.0.0.0/8 { } } } neighbor 88.88.88.1 { address-family { ipv4-unicast { } } ebgp-multihop 1 } } [edit]</pre>
V2: AS-300	Advertise remote-as AS-300 to BGP.	<pre>vyatta@V2# set protocols bgp 300 neighbor 99.99.99.1 remote-as 100 vyatta@V2# set protocols bgp 300 neighbor 99.99.99.1 address-family ipv4-unicast</pre>
V2	Commit the configuration	<pre>vyatta@V2## commit</pre>
V2	Verify the configuration	<pre>vyatta@V2# show protocols bgp bgp 300 { address-family { ipv4-unicast { } } neighbor 99.99.99.1 { address-family { ipv4-unicast { } } remote-as 100 } } }</pre>

Verifying the route origination

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

R1: show ip bgp summary

The following example 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.

```
vyatta@R1:~$ show ip bgp summary
```



```
BGP router identifier 10.0.0.11, local AS number 100
BGP table version is 35
3 BGP AS-PATH entries
0 BGP community entries

IPv4 Unicast address family:
Neighbor      V    AS  MsgRcvd  MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
10.0.0.22     4   100    6067    6072     35    0    0  2d02h08m    0
10.0.0.33     4   100    6064    6071     35    0    0  2d02h08m    0
10.0.0.44     4   100    6066    6069     35    0    0  1d03h52m    4
88.88.88.2    4   200    6026    6030     35    0    0  00:24:05    1
Total number of neighbors 4
```

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

The following example shows the output of the `show ip bgp` command for router R1 at this stage of the configuration.

```
vyatta@R1:~$ show ip bgp
BGP table version is 36, local router ID is 10.0.0.11
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
               S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop           Metric LocPrf Weight Path
*>i10.0.0.0         99.99.99.2         0     100      0 300 i
*> 12.0.0.0         88.88.88.2         0           0 200 i
*>i172.16.0.0/12    99.99.99.2         0     100      0 300 i
*>i172.16.0.0/24    10.0.0.44          0     100      0 i
*>i192.168.0.0/16  99.99.99.2         0     100      0 300 i

Total number of prefixes 5
```

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

The following example shows the output of the `show ip route bgp` command for router R1 at this stage of the configuration.

```
vyatta@R1:~$ show ip route bgp
B      10.0.0.0/8 [200/0] via 99.99.99.2, 00:12:04
B      *> 12.0.0.0/8 [20/0] via 88.88.88.2, dp0p161p1, 00:17:51
B      *> 172.16.0.0/12 [200/0] via 99.99.99.2 (recursive via 172.16.0.4 ), 00:12:04
B      172.16.0.0/24 [200/0] via 10.0.0.44, 00:18:37
B      *> 192.168.0.0/16 [200/0] via 99.99.99.2 (recursive via 172.16.0.4 ), 00:12:04
```

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

The following example shows the output of the `show ip bgp summary` command for router R4 at this stage of the configuration.

```
vyatta@R4:~$ show ip bgp summary
BGP router identifier 10.0.0.44, local AS number 100
BGP table version is 39
3 BGP AS-PATH entries
0 BGP community entries

IPv4 Unicast address family:
Neighbor      V    AS MsgRcvd MsgSent   TblVer   InQ  OutQ  Up/Down   State/PfxRcd
10.0.0.11     4   100   6076   6076     39    0    0 1d03h57m    1
10.0.0.22     4   100   6051   6057     39    0    0 1d03h57m    0
10.0.0.33     4   100   6043   6048     39    0    0 1d03h57m    0
99.99.99.2    4   300  35755  44056     39    0    0 01:10:11    3
Total number of neighbors 4
vyatta@R4:~$
```

R4: show ip bgp

The following example shows the output of the `show ip bgp` command for router R4 at this stage of the configuration.

```
vyatta@R4:~$ show ip bgp
BGP table version is 70, local router ID is 10.0.0.44
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
                S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop           Metric LocPrf Weight Path
*> 10.0.0.0        99.99.99.2          0           0 300 i
*>i12.0.0.0        88.88.88.2          0    100     0 200 i
*> 172.16.0.0/12   99.99.99.2          0           0 300 i
*> 172.16.0.0/24   0.0.0.0             100  32768 i
*> 192.168.0.0/16  99.99.99.2          0           0 300 i

Total number of prefixes 5
```

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 AT&T Vyatta vRouter 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.

The following table 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 the following figure.

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 9: Filtering inbound routes



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

Table 5: Creating an import policy

Router	Step	Commands
R1	Create a list of prefixes to allow. In this case we just have one - 12.0.0.0/8.	<pre>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</pre>
R1	Create a route map rule to permit all prefixes in our list.	<pre>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</pre>
R1	Create a route map rule to deny all other prefixes.	<pre>vyatta@R1# set policy route- map eBGP-IMPORT rule 20 action deny</pre>
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 address-family ipv4-unicast 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>



Router	Step	Commands
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{ address-family { ipv4-unicast { route-map { import eBGP-IMPORT } } ipv6-unicast { } } ebgp-multihop 1 remote-as 200 } vyatta@R1#</pre>
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	Commit the configuration.	<pre>vyatta@R4# commit</pre>



Router	Step	Commands
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	Commit the configuration.	<pre>vyatta@R4# commit</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@R4# run reset ip bgp 99.99.99.2</pre>
R4	Display the policy configuration.	<pre>vyatta@R4# show policy route { prefix-list RFC1918PREFIXES { rule 1 { action permit le 32 prefix 10.0.0.0/8 } } route-map eBGP-IMPORT { rule 10 { action deny match { ip { address { prefix-list RFC1918PREFIXES } } } } rule 20 { action permit } } } vyatta@R4#</pre>



Router	Step	Commands
R4	Display the BGP configuration for eBGP neighbor 99.99.99.2.	<pre>vyatta@R4# show protocols bgp 100 neighbor 99.99.99.2 address-family { ipv4-unicast { route-map { import } } eBGP-IMPORT } ipv6-unicast { } } ebgp-multihop 1 remote-as 300 } vyatta@R4#</pre>

Verifying the inbound filter

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

R1: show ip bgp

The following output shows R1's BGP table before the import filter is applied.

```
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

R1: show ip bgp

The following output shows R1's BGP table after the import filter is applied. Note that only 12.0.0.0 from 88.88.88.2 is still in the table.

```
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
*> 12.0.0.0	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
*>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

R4: show ip bgp

The following output shows R4's BGP table before the import filter is applied.

```
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.

```
vyatta@R4:~$ show ip bgp
BGP table version is 2, local router ID is 10.0.0.44
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
               S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*>i12.0.0.0	88.88.88.2	0	100		0 200 i

Total number of prefixes 1

Outbound route filtering

Filtering outbound prefixes is another common BGP configuration requirement. On the AT&T Vyatta vRouter 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 the following figure.

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 10: Filtering outbound routes



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

Table 6: Creating an export policy

Router	Step	Commands
R1	Create a list of AS paths to deny. In this case we just have one - AS300.	<pre>vyatta@R1# set policy route as-path-list AS300 rule 1 action permit vyatta@R1# set policy route as-path-list AS300 rule 1 regex 300</pre>
R1	Create a route map rule to deny all AS paths in our list.	<pre>vyatta@R1# set policy route route-map eBGP-EXPORT rule 10 action deny vyatta@R1# set policy route route-map eBGP-EXPORT rule 10 match as-path AS300</pre>
R1	Create a route map rule to permit all other prefixes.	<pre>vyatta@R1# set policy route route-map eBGP-EXPORT rule 20 action permit</pre>



Router	Step	Commands
R1	Assign the route map policy created as the export and import route map policy for AS 200.	<pre>vyatta@R1# set protocols bgp 100 neighbor 88.8.88.2 remote-as 200 vyatta@R1# set protocols bgp 100 neighbor 88.88.88.2 address-family ipv4-unicast route-map export eBGP- EXPORT vyatta@R1# set protocols bgp 100 neighbor 88.88.88.2 address-family ipv4-unicast route-map import eBGP- IMPORT vyatta@R1# set protocols bgp 100 neighbor 88.88.88.2 ebgp-multihop 1</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 configurations.	<pre>vyatta@R1# show policy route { as-path-list AS300 { rule 1 { action permit regex 300 } } route-map eBGP-EXPORT { rule 10 { action deny match { as-path AS300 } } rule 20 { action permit } } }</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 address-family { ipv4-unicast { route-map { export eBGP- EXPORT import eBGP- IMPORT } } } ebgp-multihop 1 remote-as 200</pre>



Router	Step	Commands
R4	Create a list of AS paths to deny. In this case we just have one - AS200.	<pre>vyatta@R4# set policy route route-map eBGP-EXPORT rule 20 action permit vyatta@R4# set policy route as-path-list AS200 rule 1 regex 200 vyatta@R4# commit</pre>
R4	Create a route map rule to deny all AS paths in our list.	<pre>vyatta@R4# set policy route route-map eBGP-EXPORT rule 10 action deny vyatta@R4# set policy route 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 route-map eBGP-EXPORT rule 20 action permit vyatta@R4# commit</pre>
R4	Assign the route map policy created as the export route map policy for AS 300.	<pre>vyatta@R4# set protocol bgp 100 neigh 99.99.99.2 address-family ipv4-unicast 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>



Router	Step	Commands
R4	Display the policy configurations.	<pre>vyatta@R4# show policy route { as-path-list AS200 { rule 1 { action permit regex 200 } } prefix-list RFC1918PREFIXES { rule 1 { action permit le 32 prefix 10.0.0.0/8 } } route-map eBGP-EXPORT { rule 10 { action deny match { as-path AS200 } } rule 20 { action permit } } route-map eBGP-IMPORT { rule 10 { action deny match { ip { address { prefix-list RFC1918PREFIXES } } } rule 20 { action permit } } }</pre>



Router	Step	Commands
R4	Display the BGP configuration for eBGP neighbor 99.99.99.2.	<pre>vyatta@R4# show protocols bgp 100 neighbor 99.99.99.2 address-family { ipv4-unicast { route-map { import eBGP-IMPORT } soft-reconfiguration { inbound } } ipv6-unicast { } } ebgp-multihop 1 remote-as 300</pre>

Verifying the outbound filter

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

AS 200: show ip bgp

The following example shows AS 200's BGP table before the export filter is applied.

```
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

The following example shows AS 200's BGP table after the export filter is applied.

```
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 AS to sub-divide the AS into sub-AS. 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 the following figure, 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 11: BGP confederation



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

To create the confederation shown in the preceding figure, perform the following steps in configuration mode.

Table 7: Creating a BGP confederation

Router	Step	Commands
R1	Delete current BGP configuration.	vyatta@R1# delete protocols bgp 100



Router	Step	Commands
R1	To allow routes from AS200 to be injected into RIB on R3.	<pre>vyatta@R1# set protocols bgp 65031 neighbor 10.0.0.33 nexthop-self</pre>
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 address-family ipv4-unicast 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>



Router	Step	Commands
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 } } address-family ipv4- unicast { 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>
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>



Router	Step	Commands
R2	Activate the neighbor for the IPv4 unicast address family to enable the exchange IPv4 unicast routes.	<pre>vyatta@R2# set protocols bgp 100 neighbor 10.0.0.44 address-family ipv4-unicast</pre>
R2	Set the network to advertise.	<pre>vyatta@R2# set protocols bgp 65021 address-family ipv4-unicast 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>



Router	Step	Commands
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 } neighbor 10.0.0.44 { address-family ipv4-unicast } address-family ipv4- unicast { network 172.16.0.0/24 } parameters { confederation { identifier 100 peers 65031 } router-id 10.0.0.22 } } vyatta@R2#</pre>
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	Activate the neighbor for the IPv4 unicast address family to enable the exchange IPv4 unicast routes.	<pre>vyatta@R3# set protocols bgp 100 neighbor 10.0.0.44 address-family ipv4-unicast</pre>



Router	Step	Commands
R3	Set the network to advertise.	<pre>vyatta@R3# set protocols bgp 65031 address-family ipv4-unicast 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>
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 } neighbor 10.0.0.44 { address-family ipv4-unicast } address-family ipv4- unicast { 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>



Router	Step	Commands
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>
R4	Set the AS300 router in a different AS than R4.	<pre>vyatta@R4# set protocols bgp 65021 neighbor 99.99.99.2 remote-as 300</pre>
R4	Set the outbound filter.	<pre>vyatta@R4# set protocols bgp 65021 neighbor 99.99.99.2 route-map export eBGP- EXPORT</pre>
R4	Set the inbound filter.	<pre>vyatta@R4# set protocols bgp 65021 neighbor 99.99.99.2 route-map import eBGP- IMPORT</pre>
R4	Activate the neighbor for the IPv4 unicast address family to enable the exchange IPv4 unicast routes.	<pre>vyatta@R4# set protocols bgp 100 neighbor 10.0.0.44 address-family ipv4-unicast</pre>
R4	Set the network to advertise.	<pre>vyatta@R4# set protocols bgp 65021 network 172.16.0.0/24</pre>
R4	Set the AS identifier for the confederation.	<pre>vyatta@R4# set protocols bgp 65021 parameters confederation identifier 100</pre>
R4	Set the peer for the sub-AS.	<pre>vyatta@R4# set protocols bgp 65021 parameters confederation peers 65031</pre>
R4	Set the router ID.	<pre>vyatta@R4# set protocols bgp 65021 parameters router-id 10.0.0.44</pre>
R4	Commit the configuration.	<pre>vyatta@R4# commit</pre>



Router	Step	Commands
R4	Display the BGP configuration.	<pre>vyatta@R4# show protocols bgp 65021 { 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 } } neighbor 10.0.0.44 { address-family ipv4-unicast } address-family ipv4- unicast { network 172.16.0.0/24 } parameters { confederation { identifier 100 peers 65031 } router-id 10.0.0.44 } } vyatta@R4#</pre>

Verifying the confederation

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

R1: show ip bgp summary

The following example shows the output of the `show ip bgp summary` command for router R1 at this stage of the configuration.

```
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
```



R1: show ip bgp


The following example shows the output of the `show ip bgp` command for router R1 at this stage of the configuration.

```
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
```

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-AS be shown inside brackets (). This information is not transmitted outside of the true AS (AS 100).



R2: show ip bgp summary

The following example shows the output of the `show ip bgp summary` command for router R2 at this stage of the configuration.

```
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
```

R2: show ip bgp

The following example shows the output of the `show ip bgp` command for router R2 at this stage of the configuration.

```
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

R3: show ip bgp summary

The following example shows the output of the `show bgp peers` command for router R3 at this stage of the configuration.

```
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

R3: show ip bgp

The following example shows the output of the `show ip bgp` command for router R3 at this stage of the configuration.

```
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

R4: show ip bgp summary

The following example shows the output of the `show ip bgp summary` command for router R4 at this stage of the configuration.

```
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



R4: show ip bgp

The following example shows the output of the show ip bgp command for router R4 at this stage of the configuration.

```

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


```

Route reflectors

Router reflectors are another technology designed to help BGP scale to a large AS. In a route reflector configuration there is at least one route reflector server and one or more route reflector clients. In example shown in the following figure, 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 13: 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.

Table 8: Creating route reflectors

Router	Step	Commands
R1	Delete current BGP configuration.	<pre>vyatta@R1# delete protocols bgp</pre>
R1	To allow routes from AS200 to be injected into RIB on R2.	<pre>vyatta@R1# set protocols bgp 100 neighbor 10.0.0.22 nexthop-self</pre>
R1	Set R2 in the same AS as R1.	<pre>vyatta@R1# set protocols bgp 100 neighbor 10.0.0.22 remote-as 100</pre>
R1	Set R2 as a route reflector client.	<pre>vyatta@R1# set protocols bgp 100 neighbor 10.0.0.22 address-family ipv4-unicast route-reflector-client</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	To allow routes from AS200 to be injected into RIB on R3.	<pre>vyatta@R1# set protocols bgp 100 neighbor 10.0.0.33 nexthop-self</pre>
R1	Set R3 in the same AS as R1.	<pre>vyatta@R1# set protocols bgp 100 neighbor 10.0.0.33 remote-as 100</pre>
R1	Set R3 as a route reflector client.	<pre>vyatta@R1# set protocols bgp 100 neighbor 10.0.0.33 address-family ipv4-unicast route-reflector-client</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	To allow routes from AS200 to be injected into RIB on R4.	<pre>vyatta@R1# set protocols bgp 100 neighbor 10.0.0.44 nexthop-self</pre>
R1	Set R4 in the same AS as R1.	<pre>vyatta@R1# set protocols bgp 100 neighbor 10.0.0.44 remote-as 100</pre>



Router	Step	Commands
R1	Activate the neighbor for the IPv4 unicast address family to enable the exchange IPv4 unicast routes.	<pre>vyatta@R1# set protocols bgp 100 neighbor 10.0.0.44 address-family ipv4-unicast</pre>
R1	Set R4 as a route reflector client.	<pre>vyatta@R1# set protocols bgp 100 neighbor 10.0.0.44 address-family ipv4-unicast 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 address-family ipv4-unicast 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>



Router	Step	Commands
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 10.0.0.44 { address-family ipv4-unicast } neighbor 88.88.88.2 { remote-as 200 route-map { export eBGP- EXPORT import eBGP- IMPORT } } address-family ipv4- unicast { network 172.16.0.0/24 } parameters { router-id 10.0.0.11 } } vyatta@R1#</pre>
R2	Delete current BGP configuration.	<pre>vyatta@R2# delete protocols bgp</pre>
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>



Router	Step	Commands
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	Activate the neighbor for the IPv4 unicast address family to enable the exchange IPv4 unicast routes.	<pre>vyatta@R2# set protocols bgp 100 neighbor 10.0.0.44 address-family ipv4-unicast</pre>
R2	Set the network to advertise.	<pre>vyatta@R2# set protocols bgp 100 address-family ipv4-unicast network 172.16.0.0/24</pre>
R2	Set the router ID.	<pre>vyatta@R2# set protocols bgp 100 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 100 { neighbor 10.0.0.11 { remote-as 100 update-source 10.0.0.22 } neighbor 10.0.0.44 { address-family ipv4-unicast } address-family ipv4- unicast { network 172.16.0.0/24 } parameters { router-id 10.0.0.22 } } vyatta@R2#</pre>
R3	Delete current BGP configuration.	<pre>vyatta@R3# delete protocols bgp</pre>
R3	Set R1 in the same AS as R3.	<pre>vyatta@R3# set protocols bgp 100 neighbor 10.0.0.11 remote-as 100</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 100 neighbor 10.0.0.11 update-source 10.0.0.33</pre>



Router	Step	Commands
R3	Activate the neighbor for the IPv4 unicast address family to enable the exchange IPv4 unicast routes.	<pre>vyatta@R3# set protocols bgp 100 neighbor 10.0.0.44 address-family ipv4-unicast</pre>
R3	Set the network to advertise.	<pre>vyatta@R3# set protocols bgp 100 address-family ipv4-unicast network 172.16.0.0/24</pre>
R3	Set the router ID.	<pre>vyatta@R3# set protocols bgp 100 parameters router-id 10.0.0.33</pre>
R3	Commit the configuration.	<pre>vyatta@R3# commit</pre>
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 } neighbor 10.0.0.44 { address-family ipv4-unicast } address-family ipv4- unicast { 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 nexthop-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>



Router	Step	Commands
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	Activate the neighbor for the IPv4 unicast address family to enable the exchange IPv4 unicast routes.	<pre>vyatta@R4# set protocols bgp 100 neighbor 10.0.0.44 address-family ipv4-unicast</pre>
R4	Set the network to advertise.	<pre>vyatta@R4# set protocols bgp 100 address-family ipv4-unicast 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>
R4	Commit the configuration.	<pre>vyatta@R4# commit</pre>



Router	Step	Commands
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 } } address-family ipv4- unicast { 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

The following example shows the output of the `show ip bgp summary` command for router R1 at this stage of the configuration.

```
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
```

R1: show ip bgp

The following example shows the output of the `show ip bgp` command for router R1 at this stage of the configuration.

```
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

R2: show ip bgp summary

The following example shows the output of the `show ip bgp summary` command for router R2 at this stage of the configuration.

```
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
```

R2: show ip bgp

The following example shows the output of the `show ip bgp` command for router R2 at this stage of the configuration.

```
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
```

R3: show ip bgp summary

The following example shows the output of the `show ip bgp summary` command for router R3 at this stage of the configuration.

```
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

The following example shows the output of the `show ip bgp` command for router R3 at this stage of the configuration.

```
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
```

R4: show ip bgp summary

The following example shows the output of the `show ip bgp summary` command for router R4 at this stage of the configuration.

```
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
```

R4: show ip bgp

The following example shows the output of the `show ip bgp` command for router R4 at this stage of the configuration.

```
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, see the AT&T Vyatta Network Operating System Routing Policies Configuration Guide.

Configuring IPv6 routing with BGP

This section presents sample configurations for dynamic IPv6 routing with BGP.

The configuration examples that follow are based on using internal BGP (iBGP) and external BGP (eBGP) to connect and achieve communication. In this example we show configuration of the nodes using dynamic IPv6 routing with iBGP and eBGP to enable R1 and R3 to communicate through R2. To achieve this, you must configure R1 and R2 as iBGP peers, then configure R2 and R3 as eBGP peers.

For information on iBGP and eBGP, see [iBGP and eBGP \(page 22\)](#).

For information on configuring neighbors, see [Neighbor Commands \(page 171\)](#).

Configure IPv6 BGP peers

To configure an iBGP peer between R1 and R2, perform the following steps in configuration mode.

Table 9: Configuring R1 and R2 as iBGP peers

Router	Step	Commands
R1	Configure an iBGP peer.	<pre>vyatta@R1# set protocols bgp 1 neighbor 2001:db8:2::2 remote-as 1</pre>
R1	Configure the address-family.	<pre>vyatta@R1# set protocols bgp 1 neighbor 2001:db8:2::2 address-family ipv6-unicast</pre>
R1	Commit the changes.	<pre>vyatta@R1# commit</pre>
R2	Configure a iBGP peer.	<pre>vyatta@R2# set protocols bgp 1neighbor 2001:db8:2::1 remote-as 1</pre>
R2	Configure the address-family.	<pre>vyatta@R2# set protocols bgp 1neighbor 2001:db8:2::1 address-family ipv6-unicast</pre>
R2	Commit the change.	<pre>vyatta@R2# commit</pre>

To configure eBGP between R2 and R3, perform the following steps in configuration mode on R2.

**Table 10: Configuring R2 and R3 as eBGP peers**

Router	Step	Commands
R2	Configure an eBGP peer.	<pre>vyatta@R2# set protocols bgp 2 neighbor 2001:db8:2::2 remote-as 3</pre>
R2	Configure a BGP multihop peer.	<pre>vyatta@R2# set protocols bgp 2 neighbor 2001:db8:2::1 ebgp-multihop '2'</pre>
R2	Configure the address-family.	<pre>vyatta@R2# set protocols bgp 2 neighbor 2001:db8:2::2 address-family ipv6-unicast</pre>
R2	Commit the changes.	<pre>vyatta@R2# commit</pre>
R3	Configure an eBGP peer.	<pre>vyatta@R3# set protocols bgp 3 neighbor 2001:db8:2::1 remote-as 2</pre>
R3	Configure a BGP multihop peer.	<pre>vyatta@R3# set protocols bgp 3 neighbor 2001:db8:2::1 ebgp-multihop '2'</pre>
R3	Configure the address-family.	<pre>vyatta@R3# set protocols bgp 3 neighbor 2001:db8:2::1 address-family ipv6-unicast</pre>
R3	Commit the changed.	<pre>vyatta@R3# commit</pre>

Advertise connected networks

To advertise connected networks on R1, perform the following steps in configuration mode.

Table 11: Advertising connected networks on r1

Step	Commands
Advertise connected networks through BGP.	<pre>vyatta@R1# set protocols bgp 1 address-family ipv6-unicast redistribute connected</pre>
Commit the change.	<pre>vyatta@R1# commit</pre>

Confirm address-family

To confirm the address-family, use the following command in operational mode:

```
vyatta@R1# show protocols bgp
bgp 1 {
  neighbor 2001:db8:2::2 {
```



```
address-family {
  ipv6-unicast {
  }
}
remote-as 2
}
}
[edit]
```

Confirm BGP neighbors

To confirm BGP neighbors, use the following command in configuration mode.

```
vyatta@R1# run show ip bgp ipv6 unicast neighbors 20:1:1::1 received-routes
BGP table version is 6, local router ID is 180.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete
```

Network	Next Hop	Metric	LocPrf	Weight	Path
*> 21:1:1::/64	20:1:1::1(fe80::20c:29ff:fe5a:b4da)			0	100 i

Total number of prefixes 1

Confirm the peer session

To confirm the peer session, perform the following step in configuration mode (only R1 is shown below).

```
vyatta@R1# run show ip bgp ipv6 summary
BGP router identifier 180.2.2.2, local AS number 200
BGP table version is 6
1 BGP AS-PATH entries
0 BGP community entries
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
20:1:1::1	4	100	452	458	6	0	0	00:04:00	1

Confirm advertised routes

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

**Table 12: Confirm routes advertised and learned by r1**

Step	Commands
Display routes advertised by R1.	<pre>vyatta@R1# show ip bgp ipv6 unicast 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>
Display routes learned by R1.	<pre>vyatta@R1:~\$ show ip bgp ipv6 unicast 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</pre>
Display the routing table.	<pre>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, dp0p1p1 C>* 2001:db8:2::/64 is directly connected, dp0p1p3 B>* 2001:db8:3::/64 [20/1] via fe80::20c:29ff:fe4e:fc6, dp0p1p3, 00:22:47 C * fe80::/64 is directly connected, dp0p1p3 C>* fe80::/64 is directly connected, dp0p1p1</pre>




Step	Commands
Commit the change.	vyatta@R1# commit

Advertise IPv4 prefix for IPv6 enabled BGP peers

IPv6 enabled BGP peers that are configured with `address-family ipv4-unicast` should have a route-map to define the next hop for the ipv4 prefix. This example defines a route-map for IPv6 enabled BGP peers. The route-map enables the device with a physical IPv6 address to advertise an IPv4 prefix with proper ipv4 address as the next hop.

Figure 14: Advertise IPv4 prefix for IPv6 enabled BGP peers



In this example Router 2 is connected to Router 4 and has IPv4 and IPv6 interfaces. The physical direct connection link is configured with IPv4 30.1.1/24 and IPv6 30:1:1::/64 addresses. Router 2 and Router 4 are configured as BGP peers by using ipv6 30:1:1::/64 address.

The network should look as shown in the figure.

Table 13: Advertise IPv4 prefix for IPv6 enabled BGP peers

Router	Step	Commands
R2	Display IPv4 prefix advertised by the BGP peer.	<pre>vyatta@R2# run show ip bgp ipv4 unicast BGP table version is 5, local router ID is 170.2.2.2 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled S Stale Origin codes: i - IGP, e - EGP, ? - incomplete Network Next Hop Metric LocPrf Weight Path *> 20.1.1.0/24 30.1.1.2 32768 ? Total number of prefixes 1</pre>



Router	Step	Commands
R2	Display the configuration.	<pre>vyatta@R2# show protocol bgp bgp 100 { address-family { ipv4-unicast } redistribute { connected { route-map map1 } } neighbor 30:1:1::1 { address-family { ipv4-unicast ipv6-unicast } remote-as 200 } }</pre>



Router	Step	Commands
R2	Display the policy configuration.	<pre>vyatta@R2# show policy policy { route { prefix-list } p1 { rule } 10 { action permit prefix 20.1.1.0/24 } } route-map } map1 { rule } 10 { action permit match { ip { address } } prefix-list p1 } } } set { ip-next-hop 30.1.1.2 } } }</pre>
R2	Create a rule to match prefix 20.1.1.0/24.	<pre>vyatta@R2# set policy route prefix-list p1 rule 10 action permit vyatta@R2# set policy route prefix-list p1 rule 10 prefix 20.1.1.0/24</pre>



Router	Step	Commands
R2	Create a route-map rule to allow all IP address in the prefix list p1.	<pre>vyatta@R2# set policy route route-map map1 rule 10 action 'permit' vyatta@R2# set policy route route-map map1 rule 10 match ip address prefix- list 'p1'</pre>
R2	Create a route-map rule to configure the ip-next hop as '30.1.1.2'.	<pre>vyatta@R2# set policy route route-map map1 rule 10 set ip-next-hop '30.1.1.2'</pre>
R2	Redistribute ipv4 unicast directly connected route into BGP.	<pre>vyatta@R2# set protocols bgp 100 address-family ipv4-unicast redistribute connected route-map map1</pre>
R2	Configure IPv4 and IPv6 unicast address family between BGP peers.	<pre>vyatta@R2# set protocols bgp 100 neighbor 30:1:1::1 address-family ipv4-unicast vyatta@R2# set protocols bgp 100 neighbor 30:1:1::1 address-family ipv6-unicast vyatta@R2# set protocols bgp 100 neighbor 30:1:1::1 remote-as 200</pre>



Router	Step	Commands
R2	Display interfaces to verify the configuration.	<pre>vyatta@R2# show interfaces Codes: S - State, L - Link, u - Up, D - Down, A - Admin Down Interface IP Address S/L Description ----- ----- dp0p160p1 10.18.177.130/21 u/u dp0p161p1 20.1.1.2/24 u/u 20:1:1::2/64 dp0p162p1 - u/u dp0p192p1 - u/u dp0p193p1 - u/u dp0p193p1.10 110.1.1.1/24 A/D dp0p194p1 30.1.1.2/24 u/u 30:1:1::2/64 dp0p224p1 - u/u dp0p225p1 - u/u dp0p256p1 - u/u dp0p257p1 - u/u lo1 170.2.2.2/24 u/u lo99 - u/u</pre>



Router	Step	Commands
R2	Display advertised BGP IPv4 unicast routes for a BGP neighbor.	<pre>vyatta@R2# run show ip bgp ipv4 unicast neigh 30:1:1::1 advertised-routes BGP table version is 25, local router ID is 170.2.2.2 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal Origin codes: i - IGP, e - EGP, ? - incomplete Network Next Hop Metric LocPrf Weight Path *> 20.1.1.0/24 30.1.1.2 32768 ? Total number of prefixes 1</pre>



Router	Step	Commands
R4	Display IPv4 unicast routes for R4.	<pre>vyatta@R4# show ip bgp ipv4 unicast BGP table version is 7, local router ID is 170.4.4.4 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled S Stale Origin codes: i - IGP, e - EGP, ? - incomplete Network Next Hop Metric LocPrf Weight Path *> 20.1.1.0/24 30.1.1.2 0 100 ? Total number of prefixes 1 vyatta@R4# sh ip bgp ipv4 unicast neighbors 30:1:1::2 received-routes BGP table version is 7, local router ID is 170.4.4.4 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal Origin codes: i - IGP, e - EGP, ? - incomplete Network Next Hop Metric LocPrf Weight Path *> 20.1.1.0/24 30.1.1.2 0 100 ? Total number of prefixes 1</pre>



Router	Step	Commands
R4	Display interfaces for R4.	<pre>vyatta@R4# show interfaces Codes: S - State, L - Link, u - Up, D - Down, A - Admin Down Interface IP Address S/L Description ----- ----- dp0p160p1 10.18.177.132/21 u/u dp0p161p1 30.1.1.1/24 u/u 30:1:1::1/64 dp0p192p1 - u/u dp0p193p1 - u/u dp0p224p1 - u/u dp0p225p1 - u/u dp0p256p1 - A/D dp0p257p1 - A/D lo1 170.4.4.4/32 u/u lo2 - u/u</pre>
R4	Display the configuration on R4.	<pre>vyatta@R4# show protocol bgp bgp 200 { neighbor 30:1:1::2 { address- family { ipv4-unicast { soft-reconfiguration { inbound } } ipv6-unicast } remote-as 100 } }</pre>



Router	Step	Commands
R4	Configure the IPv6 neighbor.	<pre>vyatta@R4# set protocols bgp 200 neighbor 30:1:1::2 address-family ipv4-unicast soft-reconfiguration inbound</pre>
R4	Display the configuration.	<pre>vyatta@R4# show ip bgp summary BGP router identifier 170.4.4.4, local AS number 200 BGP table version is 7 2 BGP AS-PATH entries 0 BGP community entries IPv4 Unicast address family: Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down Statd 30:1:1::2 4 100 52 52 7 0 0 00:25:00 1 Total number of neighbors 1 IPv6 Unicast address family: Neighbor V AS MsgRcvd MsgSent TblVer InQ OutQ Up/Down Statd 30:1:1::2 4 100 52 52 1 0 0 00:25:00 0 Total number of neighbors 1</pre>

Configuring symmetric routing with VRRP and BGP

In this topology, VRRP (Virtual Router Redundancy Protocol) runs on the local internal LAN network. BGP runs between the ISP and the VRRP routers. The northbound traffic from the local internal LAN network routers does not use the same path as does the southbound (incoming) traffic from the ISP network routers. This mode of routing is referred to as asymmetric routing.


Note:

If you are upgrading the AT&T Vyatta vRouter to the 4.0 release from any previous release, ensure that unique VRRP group IDs are used across interfaces to avoid conflicting configurations among those interfaces. Furthermore, by employing only the following command, use unique VRRP group IDs for the VRRP groups that are tracked by BGP.

```
protocols bgp asn neighbor id interface interface-name vrrp-failover vrrp-group vrrp-group-id
```

. VRRP groups that are not tracked do not require unique IDs within the vRouter.

This example explains how to set up symmetric routing using BGP and VRRP.

**Figure 15: Configuring symmetric routing for VRRP and BGP**

The following components of the topology are shown in the preceding figure.

- Vyatta01 and Vyatta02 are the two routers in the local internal LAN network that are running VRRP. Vyatta01 is the master router and Vyatta02 is the backup router.
- CE01 and CE02 are the ISP gateway routers that are running BGP.

Both the VRRP master and backup routers, Vyatta01 and Vyatta02, are connected to an ISP gateway router, by using BGP. This topology provides two possible outgoing and incoming paths for the traffic.

For example, the HTTP traffic passes from CE01 to the master router, Vyatta01. To make routing symmetric, the data sent from Vyatta01 must pass through CE01.

When VRRP failover occurs, the flow of northbound traffic switches to the backup router, Vyatta02 which becomes the new master router.

After the switchover, the southbound traffic coming from the ISP network to the LAN network should pass traffic through Vyatta02 to achieve symmetric routing.

BGP recognizes the multiple paths to send traffic. BGP determines the best path to send traffic by using the path selection algorithm. Modify the BGP path selection process in ISP routers by influencing the following parameters:

- Multi-exit discriminator (MED)
- AS PATH Length
- BGP route map for VRRP failover

To enable BGP updates to ISP routers for influencing the BGP attributes perform following steps in configuration mode.


**Table 14: Configuring VRRP settings**

Router	Step	Commands
Vyatta01 and Vyatta02	Configure BGP as a client to notify VRRP state changes.	<pre>vyatta@R1#set interfaces dataplane interface-name vrrp vrrp-group vrrp group- id notify bgp</pre>
Vyatta01 and Vyatta02	Set MED configuration for the neighbor.	<pre>vyatta@R1# set protocols bgp asn neighbor id interface interface-name vrrp- failover vrrp-group vrrp- group-id med med-value</pre>
Vyatta01 and Vyatta02	Set prepend-as configuration for the neighbor.	<pre>vyatta@R1# set protocols bgp asn neighbor id interface interface-name vrrp- failover vrrp-group vrrp- group-id prepend-as as-path- string</pre>
Vyatta01 and Vyatta02	Set route-map configuration for the neighbor.	<pre>vyatta@R1#set protocols bgp asn neighbor id interface interface-name vrrp- failover vrrp-group vrrp- group-id route-map route- map-name</pre>

Advertising IPv6 and IPv4 prefixes over IPv4 and IPv6 eBGP peers

The BGP speaker can now carry multiple network layer protocols like IPv4 and IPv6. The following configuration examples show how to advertise an IPv6 prefix over an IPv4 connection and an IPv4 prefix over an IPv6 connection.

Advertising an IPv6 prefix over an IPv4 eBGP peer

Figure 16: Advertising an IPv6 prefix over an IPv4 eBGP peer

The following components of the topology are shown in the preceding figure.

- R1 router with a router ID of 1.1.1 and R2 router with a router ID of 2.2.2.2 are two routers.
- The IPv4 unicast address of R1 router is 198.51.100.2.

**Table 15: Advertising an IPv6 prefix over an IPv4 eBGP peer**

Router	Step	Commands
R1	Activate the neighbor for the IPv6 unicast address family to enable the exchange of IPv6 unicast routes.	<pre>vyatta@R1# set protocols bgp 100 address-family ipv6-unicast network 2001:db8::/64</pre>
R1	Set the network to advertise.	<pre>vyatta@R1# set protocols bgp 100 neighbor 198.51.100.2 address-family ipv4-unicast vyatta@R1# set protocols bgp 100 neighbor 198.51.100.2 address-family ipv6-unicast</pre>
R1	Set R2 in the same AS as R1.	<pre>vyatta@R1# set protocols bgp 100 neighbor 198.51.100.2 remote-as 200</pre>
R1	Show the advertised routes.	<pre>vyatta@R1:~\$ show ip bgp neighbors 198.51.100.2 advertised-routes BGP table version is 2, local router ID is 1.1.1.1 Status codes: s suppressed, d damped, h history, * valid, > best, i - internal Origin codes: i - IGP, e - EGP, ? - incomplete IPv6 Unicast address family: Network Next Hop Metric LocPrf Weight Path *> 2001:db8::/64 :ffff:c633:6401 32768 i Total number of prefixes 1</pre>

From the preceding configuration example, you observe the following:


- The 2001:db8::/64 prefix is configured with a network command.
- The BGP IPv4 transport session with the IPv6 unicast address family is configured.
- The 2001:db8::/64 prefix is advertised with the next hop that is assigned to the IPv4-mapped IPv6 address (::ffff:198.51.100.1) from router 1 to router 2.

Note: The forwarding next hop must be an active route. You must ensure that the forwarding next hop is an active route because the IPv4 mapped address and the IPv6-address is derived. Therefore AT&T recommends that you use a route map to set the IPv6 next hop (for example, 2001:cafe::1/64). When a route-map is used, prefix 2001:db8::/64 is advertised with nexthop 2001:cafe::1/64.

Advertising an IPv4 prefix over an IPv6 eBGP peer



Figure 17: Advertising an IPv4 prefix over an IPv6 eBGP peer



The following components of the topology are shown in the preceding figure.

- R1 router with a router ID of 1.1.1.1 and R2 router with a router ID of 2.2.2.2 are two routers.
- The IPv6 unicast address of R1 router is 203.0.113.0/24.

Table 16: Advertising an IPv4 prefix over an IPv6 eBGP peer

Router	Step	Command
R2	Activate the neighbor for the IPv6 unicast address family to enable the exchange of IPv6 unicast routes.	<pre>vyatta@R1# set protocols bgp 100 address-family ipv4-unicast network 203.0.113.0/24</pre>
R2	Set the network to advertise.	<pre>vyatta@R1# set protocols bgp 100 neighbor 2001:cafe::2 address-family ipv4-unicast vyatta@R1# set protocols bgp 100 neighbor 2001:cafe::2 address-family ipv6-unicast</pre>
R2	Set R2 in the same AS as R1.	<pre>vyatta@R1# set protocols bgp 100 neighbor 2001:cafe::2 remote-as 200</pre>
R2	Show the advertised routes.	<pre>vyatta@R1:~\$ show ip bgp neighbors 2001:cafe::2 advertised-routes BGP table version is 1, local router ID is 1.1.1.1 Status codes: s suppressed, d damped, h history, valid, > best, i - internal Origin codes: i - IGP, e - EGP, ? - incomplete IPv4 Unicast address family: Network Next Hop Metric LocPrf Weight Path *> 203.0.113.0 1.1.1.1 32768 i Total number of prefixes 1</pre>

From the preceding configuration example, you observe the following:

- The 203.0.113.0/24 prefix is configured with a network command.
- The BGP IPv6 transport session with the IPv4 unicast address family is configured.



- The 203.0.113.0/24 prefix is advertised with the next hop that is assigned to the router ID 1.1.1.1 from router 1 to router 2.

Note: AT&T recommends that you use a route map to set the IPv4 next hop (for example, 203.0.113.1). When a route map is used the 198.51.100.0/24 prefix is advertised with the 203.0.113.1 next hop.

Setting BGP MED based on the IGP metric

A multi-exit discriminator (MED) is a path attribute that the device considers during BGP best-path selection. The value of a MED informs the external neighbors about the preferred path into an autonomous system (AS) that has multiple entry points. During BGP best-path selection, the device prefers a lower MED value as compared with a higher value.


You can send a MED value which is based on the IGP metric. When you configure a MED value, based on the IGP metric for a neighbor, the MED value is copied from the IGP metric and sent to the neighbor. Copying the MED value from the IGP metric allows eBGP peers from neighboring AS to send traffic to local AS by using the shortest path. The shortest path is based on the IGP administrative domain.

However, if a route-map for a particular neighbor is already configured with `set-metric`, then the `set-metric` command takes priority and the `med-out` configuration is ignored. For more information about setting a MED value based on the IGP metric, see the following commands.

- `protocols bgp <asn> neighbor <id> med-out igp` (page 224)
- `protocols bgp <asn> neighbor <id> med-out igp delay-updates` (page 225)
- `protocols bgp <asn> neighbor <id> med-out minimum-igp` (page 225)
- `protocols bgp <asn> peer-group <name> med-out igp` (page 287)
- `protocols bgp <asn> peer-group <name> med-out igp delay-updates` (page 288)
- `protocols bgp <asn> peer-group <name> med-out minimum-igp` (page 288).

The following example illustrates the configuration of a MED value that is based on IGP cost. In the example, R1, R2, R3 are IBGP peers. R1 is connected to R2 and R3 by using OSPF protocol. They are connected in the same AS. R4 is a neighbor of AS. R4 is an eBGP neighbor for R2 and R3. The IGP metric for R2 and R1 is 30 and the IGP metric for R1 and R3 is 20.

Figure 18: Configuration of MED value that is based on IGP metric





Router	Step	Commands
R4	Display the state of BGP routing table before setting the MED value.	<pre>vyatta@R4# show ip bgp 172.11.0.0/24 BGP routing table entry for 172.11.0.0/24 Paths: (2 available, best #1, table Default-IP- Routing-Table) Advertised to non peer- group peers: 10.34.0.3 100 10.24.0.2 from 10.24.0.2 (1.1.1.2) Origin IGP, metric 0, localpref 100, valid, external, best Last update: Tue Feb 24 15:06:20 2015 100 10.34.0.3 from 10.34.0.3 (1.1.1.3) Origin IGP, metric 0, localpref 100, valid, external Last update: Tue Feb 24 15:05:50 2015</pre>
R2	Set the MED value that is based on IGP metric for R2.	<pre>vyatta@R2# set protocols bgp 100 neighbor 10.24.0.4 med- out igp</pre>
R3	Set the MED value that is based on IGP metric for R3.	<pre>vyatta@R3# set protocols bgp 100 neighbor 10.24.0.4 med- out igp</pre>
R1	Commit the configuration.	<pre>vyatta@R1# commit</pre>



Router	Step	Commands
R4	<p>Display the state of BGP routing table after setting the MED value.</p> <p>Note: In the output, best determines the best path recommended by the device during BGP best-path selection.</p>	<pre>vyatta@R4# show ip bgp 172.11.0.0/24 BGP routing table entry for 172.11.0.0/24 Paths: (2 available, best #1, table Default-IP- Routing-Table) Advertised to non peer- group peers: 10.24.0.2 100 10.34.0.3 from 10.34.0.3 (1.1.1.3) Origin IGP, metric 20, localpref 100, valid, external, best Last update: Tue Feb 24 16:34:17 2015 100 10.24.0.2 from 10.24.0.2 (1.1.1.2) Origin IGP, metric 30, localpref 100, valid, external Last update: Tue Feb 24 16:33:41 2015</pre>



Global and Router-Specific Commands

monitor protocol bgp disable <id> all

Disables all BGP debugging for the specified neighbor.

Syntax:

```
monitor protocol bgp disable [ id ] all
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to disable all BGP debug message for the specified neighbor.

monitor protocol bgp disable <id> bfd

Disables all BGP BFD debugging for the specified neighbor.

Syntax:

```
monitor protocol bgp disable [ id ] bfd
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to disable all BGP BFD debug message generation for the specified neighbor.

monitor protocol bgp disable <id> dampening

Disables all BGP dampening debugging for the specified neighbor.

Syntax:

```
monitor protocol bgp disable [ id ] dampening
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to disable all BGP dampening debug message generation for the specified neighbor.

monitor protocol bgp disable <id> events

Disables debug message generation related to BGP events for the specified neighbor.

Syntax:

```
monitor protocol bgp disable [ id ] events
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode



Use this command to disable generation of trace-level messages related to BGP events for the specified neighbor.

monitor protocol bgp disable <id> filters

Disables debug message generation related to BGP filters for the specified neighbor.

Syntax:

```
monitor protocol bgp disable [ id ] filters
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to disable generation of trace-level messages related to BGP filters for the specified neighbor.

monitor protocol bgp disable <id> fsm

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

Syntax:

```
monitor protocol bgp disable [ id ] fsm
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

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

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 <id> keepalives

Disable debugging information related to sending and receiving BGP keep-alive messages for the specified neighbor.

Syntax:

```
monitor protocol bgp disable [ id ] keepalives
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

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

monitor protocol bgp disable <id> msdp

Disables debug message generation related to the BGP multicast source discovery protocol (MSDP) for the specified neighbor.

Syntax:

```
monitor protocol bgp disable [ id ] msdp
```

id

The IPv4 or IPv6 address of a BGP neighbor.



Operational mode

Use this command to disable generation of trace-level messages related to the BGP multicast source discovery protocol (MSDP) for the specified neighbor.

monitor protocol bgp disable <id> nsm

Disables debug message generation related to the BGP network and security manager (NSM) for the specified neighbor.

Syntax:

```
monitor protocol bgp disable [ id ] nsm
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to disable generation of trace-level messages related to the BGP network and security manager (NSM) for the specified neighbor.

monitor protocol bgp disable <id> rib

Disables debugging information related to the Routing Information Base (RIB) for the specified neighbor.

Syntax:

```
monitor protocol bgp disable [ id ] rib
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to disable generation of trace-level messages related to the RIB for the specified neighbor.

monitor protocol bgp disable <id> updates

Disables debugging information related to BGP routing updates for the specified neighbor.

Syntax:

```
monitor protocol bgp disable [ id ] updates
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

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

monitor protocol bgp enable <id> all

Enables all BGP debugging for the specified neighbor.

Syntax:

```
monitor protocol bgp enable [ id ] all
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to enable all BGP debug message generation for the specified neighbor.



monitor protocol bgp enable <id> bfd

Enables all BGP BFD debugging for the specified neighbor.

Syntax:

```
monitor protocol bgp enable [ id ] bfd
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to enable BGP BFD debug message generation for the specified neighbor.

monitor protocol bgp enable <id> dampening

Enables all BGP dampening debugging for the specified neighbor.

Syntax:

```
monitor protocol bgp enable [ id ] dampening
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to enable BGP dampening debug message generation for the specified neighbor.

monitor protocol bgp enable <id> events

Enables debug message generation related to BGP events for the specified neighbor.

Syntax:

```
monitor protocol bgp enable [ id ] events
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to enable generation of trace-level messages related to BGP events for the specified neighbor.

monitor protocol bgp enable <id> filters

Enables debug message generation related to BGP filters for the specified neighbor.

Syntax:

```
monitor protocol bgp enable [ id ] filters
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to enable the generation of trace-level messages related to BGP filters for the specified neighbor.

monitor protocol bgp enable <id> fsm

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

Syntax:



```
monitor protocol bgp enable [ id ] fsm
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

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

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 <id> keepalives

Display debugging information related to sending and receiving BGP keep-alive messages for the specified neighbor.

Syntax:

```
monitor protocol bgp enable [ id ] keepalives
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

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

monitor protocol bgp enable <id> msdp

Enables debug message generation related to the BGP multicast source discovery protocol (MSDP) for the specified neighbor.

Syntax:

```
monitor protocol bgp enable [ id ] msdp
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to enable generation of trace-level messages related to the BGP multicast source discovery protocol (MSDP) for the specified neighbor.

monitor protocol bgp enable <id> nsm

Enables debug message generation related to the BGP network and security manager (NSM) for the specified neighbor.

Syntax:

```
monitor protocol bgp enable [ id ] nsm
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to enable generation of trace-level messages related to the BGP network and security manager (NSM) for the specified neighbor.

monitor protocol bgp enable <id> rib

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

**Syntax:**

```
monitor protocol bgp enable [ id ] rib
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to enable the generation of trace-level messages related to the RIB for the specified neighbor.

monitor protocol bgp enable <id> updates

Display debugging information related to BGP routing updates for the specified neighbor.

Syntax:

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

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

id

The IPv4 or IPv6 address of a BGP neighbor.

in

Debug information is generated only for inbound routing updates.

out

Debug information is generated only for outbound routing updates.

Operational mode

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

protocols bgp <asn>

Creates a BGP instance on the router and locates it within an AS.

Syntax:

```
set protocols bgp asn
```

Syntax:

```
delete protocols bgp asn
```

Syntax:

```
show protocols bgp [ asn ]
```

asn

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 is not established.

The number range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

Configuration mode

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

Use the `set` form of this command to enable a BGP instance on the router, and set its AS. 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 ipv4-unicast

Defines parameters for controlling IPv4 unicast routes.

Syntax:

```
set protocols bgp asn address-family ipv4-unicast
```

Syntax:

```
delete protocols bgp asn address-family ipv4-unicast
```

Syntax:

```
show protocols bgp asn address-family ipv4-unicast
```

asn

The number of the AS in which this router resides.

Configuration mode

```
protocols {
  bgp asn {
    address-family {
      ipv4-unicast {
      }
    }
  }
}
```

Use the `set` form of this command to define the parameters for IPv4 unicast routes.

Use the `delete` form of this command to delete the parameters for IPv4 unicast routes.

Use the `show` form of this command to view the parameters for IPv4 unicast routes.

protocols bgp <asn> address-family ipv4-unicast aggregate-address <ipv4net>

Specifies a block of IP addresses to aggregate.

Syntax:

```
set protocols bgp asn address-family ipv4-unicast aggregate-address ipv4net { as-set | summary-only }
```

Syntax:

```
delete protocols bgp asn address-family ipv4-unicast aggregate-address ipv4net [ as-set | summary-only ]
```

Syntax:

```
show protocols bgp asn address-family ipv4-unicast aggregate-address ipv4net
```

asn

The number of the AS in which this router resides.

ipv4net

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 are not be announced.

Configuration mode

```
protocols {
  bgp asn {
    address-family {
      ipv4-unicast {
        aggregate-address ipv4net {
          as-set
          summary-only
        }
      }
    }
  }
}
```

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> address-family ipv4-unicast network <ipv4net> <action>

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

Syntax:

```
set protocols bgp asn address-family ipv4-unicast network ipv4net { backdoor | route-map map-name }
```

Syntax:

```
delete protocols bgp asn address-family ipv4-unicast network ipv4net { backdoor | route-map [ map-name ] }
```

Syntax:

```
show protocols bgp asn address-family ipv4-unicast network
```

asn

The number for the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

ipv4net

Multi-node. An IPv4 network in the format *ip-address/prefix*.

You can advertise to multiple networks by creating multiple network configuration nodes.

backdoor

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.

map-name

Specifies a configured route map to be used when advertising the network.

Configuration mode

```
protocols {
  bgp asn {
    address-family {
      ipv4-unicast {
        network ipv4net {
          backdoor
        }
      }
    }
  }
}
```



```
route-map map-name
}
}
}
}
```

Use the `set` form of this command to advertise networks to BGP neighbors.

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.

Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

protocols bgp <asn> address-family ipv4-unicast parameters dampening <time>

Enables or disables route dampening and sets route dampening values on the system.

Syntax:

```
set protocols bgp asn address-family ipv4-unicast parameters dampening { half-life minutes | max-suppress-time minutes | re-use penalty | start-suppress-time penalty }
```

Syntax:

```
delete protocols bgp asn address-family ipv4-unicast parameters dampening [ half-life | max-suppress-time | re-use | start-suppress-time ]
```

Syntax:

```
show protocols bgp asn address-family ipv4-unicast parameters dampening [ half-life | re-use | start-suppress-time | max-suppress-time ]
```

Route dampening is disabled.

asn

The number of the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

half-life minutes

The time period, in minutes, after which the penalty assigned to a route because of flapping is reduced by half. The numbers range from 1 through 45. The default is 15.

max-suppress-time minutes

The maximum time, in minutes, that a route may be suppressed. The numbers range from 1 through 20000. The default is four times the half-life period.

re-use penalty

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 numbers range from 1 through 20000. The default is 750.

start-suppress-time penalty

The route suppression threshold. If the accumulated penalty for a flapping route reaches this limit, the route is suppressed. The numbers range from 1 through 20000. The default is 2000.

Configuration mode



```
protocols {
  bgp asn {
    address-family {
      ipv4-unicast {
        parameters {
          dampening {
            half-life minutes
            max-suppress-time minutes
            re-use penalty
            start-suppress-time penalty
          }
        }
      }
    }
  }
}
```

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> address-family ipv4-unicast parameters distance global <distance>

Sets the global administrative distance for IPv4 BGP routes.

Syntax:

```
set protocols bgp asn address-family ipv4-unicast parameters distance global { external distance | internal distance | local distance }
```

Syntax:

```
delete protocols bgp asn address-family ipv4-unicast parameters distance global { external | internal | local }
```

Syntax:

```
show protocols bgp asn address-family ipv4-unicast parameters distance global [ external | internal | local ]
```

asn

The number of the AS in which this router resides. The number ranges from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for a private AS.

external *distance*

Specifies the externally learned (eBGP) route distance. The number ranges from 1 through 255.

internal *distance*

Specifies the internally learned (iBGP) route distance. The number ranges from 1 through 255.

local *distance*

Specifies the locally generated route distance. The number ranges from 1 through 255.

Configuration mode

```
protocols {
  bgp asn {
    address-family {
      ipv6-unicast {
        parameters {
          distance {
```



```
global {
  external distance
  internal distance
  local distance
}
}
```

Use the `set` form of this command to set the global administrative distance for IPv4 BGP routes.

Use the `delete` form of this command to delete the global administrative for IPv4 BGP routes.

Use the `show` form of this command to display the global administrative for IPv4 BGP routes.

protocols bgp <asn> address-family ipv4-unicast parameters network-synchronization

Enables BGP network synchronization for an IPv4 unicast address family.

Syntax:

```
set protocols bgp asn address-family ipv4-unicast parameters network-synchronization
```

Syntax:

```
delete protocols bgp asn address-family ipv4-unicast parameters network-synchronization
```

Syntax:

```
show protocols bgp asn address-family ipv4-unicast parameters network-synchronization
```

BGP network synchronization is disabled by default.

asn

The number of the AS in which this router resides. The number ranged from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for a private AS.

Configuration mode

```
protocols {
  bgp asn {
    address-family {
      ipv4-unicast {
        parameters {
          network-synchronization
        }
      }
    }
  }
}
```

BGP network synchronization applies to a static bgp route that is configured using either `protocols bgp <asn> address-family ipv4-unicast network <ipv4net> <action>` ([page 107](#)) or `protocols bgp <asn> address-family ipv6-unicast network <ipv6net>` ([page 113](#)).

Use the `set` form of this command to enable BGP network synchronization for an IPv4 unicast address family.

Use the `delete` form of this command to delete BGP network synchronization for an IPv4 unicast address family.

Use the `show` form of this command to display BGP network synchronization for an IPv4 unicast address family.



protocols bgp <asn> address-family ipv4-unicast parameters synchronization

Enables BGP synchronization for an IPv4 unicast address family.

Syntax:

```
set protocols bgp asn address-family ipv4-unicast parameters synchronization
```

Syntax:

```
delete protocols bgp asn address-family ipv4-unicast parameters synchronization
```

Syntax:

```
show protocols bgp asn address-family ipv4-unicast parameters synchronization
```

BGP synchronization is disabled by default.

asn

The number of the AS in which this router resides. The number ranged from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for a private AS.

Configuration mode

```
protocols {
  bgp asn {
    address-family {
      ipv4-unicast {
        parameters {
          synchronization
        }
      }
    }
  }
}
```

BGP synchronization applies to dynamic BGP routes that are learned from sources other than a static configuration.

Use the `set` form of this command to enable BGP synchronization for an IPv4 unicast address family.

Use the `delete` form of this command to delete BGP synchronization for an IPv4 unicast address family.

Use the `show` form of this command to display BGP synchronization for an IPv4 unicast address family.

protocols bgp <asn> address-family ipv6-unicast

Defines parameters for controlling IPv6 unicast routes.

Syntax:

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

Syntax:

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

Syntax:

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

asn

The number of the AS in which this router resides.

Configuration mode



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

Use the `set` form of this command to define the parameters for IPv6 unicast routes.

Use the `delete` form of this command to delete the parameters for IPv6 unicast routes.

Use the `show` form of this command to view the parameters for IPv6 unicast routes.

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

Specifies a contiguous block of IPv6 addresses to aggregate.

Syntax:

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

Syntax:

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

Syntax:

```
show protocols bgp asn address-family ipv6-unicast aggregate-address
```

asn

The number of the AS in which this router resides.

ipv6net

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 are not announced.

Configuration mode

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

This 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 [ route-map map-name ]
```

Syntax:

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

Syntax:

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

asn

The number for the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

ipv6net

Multi-node. An IPv6 network in the format *ipv6-address/prefix*.

You can advertise to multiple networks by creating multiple network configuration nodes.

map-name

The configured route map to be used when advertising the network.

Configuration mode

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

The command applies to IPv6 unicast routes only.

Use the `set` form of this command to specify a network to advertise networks to BGP neighbors.

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

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

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

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

Syntax:

```
set protocols bgp asn address-family ipv6-unicast network ipv6net { backdoor | route-map map-name }
```

Syntax:

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

Syntax:

**Syntax:**

```
delete protocols bgp asn address-family ipv6-unicast network ipv6net { backdoor | route-map [ map-name ] }
```

Syntax:

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

asn

The number for the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

ipv6net

Multi-node. An IPv6 network in the format *ip-address/prefix*.

You can advertise to multiple networks by creating multiple network configuration nodes.

backdoor

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.

map-name

Specifies a configured route map to be used when advertising the network.

Configuration mode

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

Use the `set` form of this command to advertise networks to BGP neighbors.

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> address-family ipv6-unicast parameters dampening <time>

Enables or disables route dampening and sets route dampening values on the system.

Syntax:

```
set protocols bgp asn address-family ipv6-unicast parameters dampening { half-life minutes | max-suppress-time minutes | re-use penalty | start-suppress-time penalty }
```

Syntax:

```
delete protocols bgp asn address-family ipv6-unicast parameters dampening [ half-life | max-suppress-time | re-use | start-suppress-time ]
```

Syntax:

```
show protocols bgp asn address-family ipv6-unicast parameters dampening [ half-life | re-use | start-suppress-time | max-suppress-time ]
```

Route dampening is disabled.

**asn**

The number of the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

half-life *minutes*

The time period, in minutes, after which the penalty assigned to a route because of flapping is reduced by half. The numbers range from 1 through 45. The default is 15.

max-suppress-time *minutes*

The maximum time, in minutes, that a route may be suppressed. The numbers range from 1 through 20000. The default is four times the half-life period.

re-use *penalty*

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 numbers range from 1 through 20000. The default is 750.

start-suppress-time *penalty*

The route suppression threshold. If the accumulated penalty for a flapping route reaches this limit, the route is suppressed. The numbers range from 1 through 20000. The default is 2000.

Configuration mode

```
protocols {
  bgp asn {
    address-family {
      ipv6-unicast {
        parameters {
          dampening {
            half-life minutes
            max-suppress-time minutes
            re-use penalty
            start-suppress-time penalty
          }
        }
      }
    }
  }
}
```

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> address-family ipv6-unicast parameters distance global <distance>

Sets the global administrative distance of BGP routes for an IPv6 unicast address family.

Syntax:

```
set protocols bgp asn address-family ipv6-unicast parameters distance global { external distance |  
internal distance | local distance }
```

Syntax:

```
delete protocols bgp asn address-family ipv6-unicast parameters distance global { external |  
internal | local }
```

Syntax:

```
show protocols bgp asn address-family ipv6-unicast parameters distance global [ external | internal  
| local ]
```



The default values for the external, internal, and local distances are mentioned in the Parameters table.

asn

The number of the AS in which this router resides. The number ranges from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for a private AS.

external distance

Specifies the externally learned (eBGP) route distance. The number ranges from 1 through 255. Default is 20.

internal distance

Specifies the internally learned (iBGP) route distance. The number ranges from 1 through 255. Default is 200.

local distance

The locally generated route distance. The number ranges from 1 through 255. Default is 200.

Configuration mode

```
protocols {
  bgp asn {
    address-family {
      ipv6-unicast {
        parameters {
          distance {
            global {
              external distance
              internal distance
              local distance
            }
          }
        }
      }
    }
  }
}
```

Use the `set` form of this command to set the global administrative distance of BGP routes for an IPv6 unicast address family.

Use the `delete` form of this command to delete the global administrative distance of BGP routes for an IPv6 unicast address family.

Use the `show` form of this command to display the global administrative distance of BGP routes for an IPv6 unicast address family.

protocols bgp <asn> address-family ipv6-unicast parameters network-synchronization

Enables BGP network synchronization for an IPv6 unicast address family.

Syntax:

```
set protocols bgp asn address-family ipv6-unicast parameters network-synchronization
```

Syntax:

```
delete protocols bgp asn address-family ipv6-unicast parameters network-synchronization
```

Syntax:

```
show protocols bgp asn address-family ipv6-unicast parameters network-synchronization
```

BGP network synchronization is disabled by default.

asn



The number of the AS in which this router resides. The number ranges from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for a private AS.

Configuration mode

```
protocols {
  bgp asn {
    address-family {
      ipv6-unicast {
        parameters {
          network-synchronization {
          }
        }
      }
    }
  }
}
```

When BGP network synchronization is enabled, the system checks its own routing table to see whether it has learned an iBGP route from IGP before it advertises the route to eBGP peers.

BGP network synchronization applies to a static bgp route that is configured using either `protocols bgp <asn> address-family ipv4-unicast network <ipv4net> <action>` ([page 107](#)) or `protocols bgp <asn> address-family ipv6-unicast network <ipv6net>` ([page 113](#)).

Use the `set` form of this command to enable BGP network synchronization for an IPv6 unicast address family.

Use the `delete` form of this command to disable BGP network synchronization for an IPv6 unicast address family.

Use the `show` form of this command to display the BGP network synchronization configuration for an IPv6 unicast address family.

protocols bgp <asn> address-family ipv6-unicast parameters synchronization

Enables BGP synchronization for an IPv6 unicast address family.

Syntax:

```
set protocols bgp asn address-family ipv6-unicast parameters synchronization
```

Syntax:

```
delete protocols bgp asn address-family ipv6-unicast parameters synchronization
```

Syntax:

```
show protocols bgp asn address-family ipv6-unicast parameters synchronization
```

BGP synchronization is disabled by default.

asn

The number of the AS in which this router resides. The number ranged from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for a private AS.

Configuration mode

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



```
}  
}
```

BGP synchronization applies to dynamic BGP routes that are learned from sources other than a static configuration.

Use the `set` form of this command to enable BGP synchronization for an IPv6 unicast address family.

Use the `delete` form of this command to delete BGP synchronization for an IPv6 unicast address family.

Use the `show` form of this command to display BGP synchronization for an IPv6 unicast address family.

protocols bgp <asn> address-family vpnv4-unicast

Defines parameters for controlling VPNv4 unicast routes.

Syntax:

```
set protocols bgp asn address-family vpnv4-unicast
```

Syntax:

```
delete protocols bgp asn address-family vpnv4-unicast
```

Syntax:

```
show protocols bgp asn address-family vpnv4-unicast
```

asn

The number of the AS in which this router resides.

Configuration mode

```
protocols {  
  bgp asn {  
    address-family {  
      vpnv4-unicast {  
      }  
    }  
  }  
}
```

Use the `set` form of this command to define the parameters for VPNv4 unicast.

Use the `delete` form of this command to delete the parameters for VPNv4 unicast.

Use the `show` form of this command to view the parameters for VPNv4 unicast.

protocols bgp *asn* log

Enables BGP logs.

Syntax:

```
set protocols bgp asn log { all | events | filters | fsm | keepalive | nsm | update }
```

Syntax:

```
delete protocols bgp asn log { all | events | filters | fsm | keepalive | nsm | update }
```

Syntax:

```
show protocols bgp asn log { all | events | filters | fsm | keepalive | nsm | update }
```

None

asn

The number of the AS in which this router resides. The number ranges from 1 through 4294967295.

all



- events** Enables all BGP logs.
- filters** Enables only BGP event logs.
- filters** Enables only BGP filter logs.
- fsm** Enables only BGP finite-state machine logs.
- keepalive** Enables only BGP keepalive logs.
- nsm** Enables only BGP network services module (NSM) logs.
- update** Enables only BGP update logs.

Configuration mode

```
protocols {
  bgp asn {
    log {
      all
      events
      filters
      fsm
      keepalive
      nsm
      update
    }
  }
}
```

Use the `set` form of this command to enable BGP logs.

Use the `delete` form of this command to remove BGP logs.

Use the `show` form of this command to view BGP logs.

protocols bgp asn log update

Enables BGP update logs.

Syntax:

```
set protocols bgp asn log { all | in | out }
```

Syntax:

```
delete protocols bgp asn log { all | in | out }
```

Syntax:

```
show protocols bgp asn log { all | in | out }
```

None

asn

The number of the AS in which this router resides. The number ranges from 1 through 4294967295.

all

Enables all BGP logs.

in

Enables BGP update-in logs.

out

Enables BGP update-out logs.

Configuration mode



```
protocols {
  bgp asn {
    log update {
      all
      in
      out
    }
  }
}
```

Use the `set` form of this command to enable BGP update logs.

Use the `delete` form of this command to remove BGP update logs.

Use the `show` form of this command to view BGP update logs.

protocols bgp <asn> parameters always-compare-med

Directs the router to compare the multi-exit discriminator (MED) for paths from neighbors in different AS.

Syntax:

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

Syntax:

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

Syntax:

```
show protocols bgp asn parameters
```

The MED is not compared.

asn

The number of the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

Configuration mode

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

Use the `set` form of this command to have the router compare the MED for paths from neighbors in different AS.

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 <action>

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

Syntax:

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

Syntax:

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

Syntax:



```
show protocols bgp asn parameters bestpath
```

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.

asn

The number of the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

confed

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

ignore

Directs the router to ignore AS_PATH during best path selection.

Configuration mode

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

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
```

Syntax:

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

Syntax:

```
show protocols bgp asn parameters bestpath
```

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

asn

The number of the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

Configuration mode

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



```
}  
}
```

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 igp-metric-ignore

Ignores Interior Gateway Protocol (IGP) metrics during BGP best-path selection.

Syntax:

```
set protocols bgp asn parameters bestpath igp-metric-ignore
```

Syntax:

```
delete protocols bgp asn parameters bestpath igp-metric-ignore
```

By default, the router considers IGP metrics during best-path selection.

asn

Number of the AS in which this router resides. Number ranges from 1 through 4294967294. Subrange of 64512 through 65534 is reserved for private AS.

Configuration mode

```
protocols {  
  bgp asn {  
    parameters {  
      bestpath {  
        igp-metric-ignore  
      }  
    }  
  }  
}
```

Use the `set` form of this command to disable IGP metrics during BGP best-path selection.

Use the `delete` form of this command to accept IGP metrics during the behavior of default best-path selection.

protocols bgp <asn> parameters bestpath med <action>

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 }
```

Syntax:

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

Syntax:

```
show protocols bgp asn parameters bestpath
```

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

asn

The number of the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

confed

Compare the MED among confederation paths.

**missing-as-worst**

Treat a missing MED as the least preferred one.

Configuration mode

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

Use the `set` form of this command to direct the router to compare the 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 maximum-as-limit <number>

Allows BGP routes with the AS_PATH or AS4_PATH attribute containing ASes lesser than the specified number. If the AS_PATH or AS4_PATH attribute contains ASes greater than the configured value, the update messages containing the AS_PATH or AS4_PATH attribute is ignored and the session is retained.

Syntax:

```
set protocols bgp asn parameters maximum-as-limit number
```

Syntax:

```
delete protocols bgp asn parameters maximum-as-limit number
```

None.

asn

The number for the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

number

Number specifying the ASes in the AS_PATH or AS4_PATH attribute. The value ranges from 1 through 254.

Configuration mode.

```
protocols {
  bgp asn {
    parameters {
      maximum-as-limit number
    }
  }
}
```



Use the `set` form of this command to allow BGP routes with `AS_PATH` or `AS4_PATH` attribute containing ASes lesser than the specified number.

Use the `delete` form of this command to delete BGP routes with `AS_PATH` or `AS4_PATH` attribute containing ASes lesser than the specified number.

Note:

When you configure this command and do a soft reset in the peer, the BGP table should flush out the prefixes that have an as-path length that is more than the configured `maximum-as-limit` value.

When soft-reconfiguration is enabled irrespective of the configured `maximum-as-limit` value, all the received prefixes should be stored in Adj-Rib-In, and then `maximum-as-limit` should be applied and prefixes with an as-path length more than the configured value should be flushed out of the BGP table. For more information on soft-reconfiguration, see [protocols bgp <asn> neighbor <id> address-family ipv6-unicast soft-reconfiguration inbound \(page 207\)](#).

protocols bgp <asn> parameters med-out-delay <minutes>

Specifies the time interval between the attempts to set a MED value that is based on the IGP metric.

Syntax:

```
set protocols bgp asn parameters med-out-delay minutes
```

Syntax:

```
delete protocols bgp asn parameters med-out-delay
```

The default time interval for a MED update is 10 minutes.

asn

Number for the AS in which this router resides. Number ranges from 1 through 4294967294. Subrange of 64512 through 65534, is reserved for private AS.

minutes

The time interval in minutes, that a MED is updated. The interval ranges from 10 through 600 minutes.

Configuration mode

```
protocols {
  bgp asn {
    parameters {
      med-out-delay minutes {
      }
    }
  }
}
```

Use the `set` form of this command to specify the time interval between the attempts to set a MED value.

Use the `delete` form of this command to disable the periodic MED update and hence the MED is updated every time the IGP metric changes.

protocols bgp <asn> parameters default local-pref <pref>

Sets default routing behaviors on the system.

Syntax:

```
set protocols bgp asn parameters default local-pref pref
```

Syntax:

```
delete protocols bgp asn parameters default local-pref
```

**Syntax:**

```
show protocols bgp asn parameters default
```

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

asn

The number of the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

local-pref *pref*

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 numbers range from 0 through 4294967295. The default is 100.

Configuration mode

```
protocols {
  bgp asn {
    parameters {
      default {
        local-pref pref
      }
    }
  }
}
```

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 Multi-Exit Discriminator (MED).

Syntax:

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

Syntax:

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

Syntax:

```
show protocols bgp asn parameters
```

Deterministic MED is not enforced.

asn

The number of the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

Configuration mode

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



Use the `set` form of this command to enforce the deterministic comparison of the 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 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
```

Syntax:

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

Syntax:

```
show protocols bgp asn parameters
```

Disabled.

asn

The number of the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

Configuration mode

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

When this feature is enabled, the router denies updates 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 is misdirecting traffic by advertising a route as if it were sourced from another autonomous system.

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.

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

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

Note:

When you configure this command and the number of prefixes received in an update message exceeds configured `enforce-first-as` limit, the session gets reset.

When soft-reconfiguration is enabled and the number of prefixes received exceeds `enforce-first-as` value, there will not be any corresponding unprocessed entries stored in Adj-RIB-In because the peer-session resets when the number of prefixes received exceeds `enforce-first-as` value. For more information on soft-reconfiguration, see [protocols bgp <asn> neighbor <id> address-family ipv6-unicast soft-reconfiguration inbound \(page 207\)](#).



protocols bgp <asn> parameters extended-asn-capability

Enables or disables extended ASN capability.

Syntax:

```
set protocols bgp asn parameters extended-asn-capability
```

Syntax:

```
delete protocols bgp asn parameters extended-asn-capability
```

Syntax:

```
show protocols bgp asn parameters extended-asn-capability
```

Disabled.

asn

The number of the AS in which this router resides. The number ranges from 1 through 4294967294. The subrange 64512 through 65534 is reserved for private AS.

parameters

BGP parameters.

Configuration mode

```
protocols {
  bgp asn {
    parameters {
      extended-asn-capability
    }
  }
}
```

Use the `set` form of this command to enable extended ASN capability.

Use the `delete` form of this command to disable extended ASN capability.

Use the `show` form of this command to view extended ASN capability configuration settings.

protocols bgp <asn> parameters graceful-restart <action>

Enables or disables graceful restart of the BGP process.

Syntax:

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

Syntax:

```
delete protocols bgp asn parameters graceful-restart [ graceful-reset | restart-time seconds | stalepath-time seconds ]
```

Syntax:

```
show protocols bgp asn parameters graceful-restart
```

Paths for restarting peers are retained for at most 360 seconds.

asn

The number of the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange 64512 through 65534 is reserved for private AS.

graceful-reset

Enables graceful reset capability.

restart-time *seconds*



Enables maximum time needed for neighbor to restart. The default is 360 seconds.

stalepath-time seconds

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 number range from 1 through 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 value.

Configuration mode

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

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
```

Syntax:

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

Syntax:

```
show protocols bgp asn parameters
```

Disabled.

asn

The number of the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

Configuration mode

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

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



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 maximum-paths ebgp <max-paths>

Specifies the maximum number of eBGP paths to a destination.

Syntax:

```
set protocols bgp asn parameters maximum-paths ebgp max-paths
```

Syntax:

```
delete protocols bgp asn parameters maximum-paths ebgp [ max-paths ]
```

Syntax:

```
show protocols bgp asn parameters maximum-paths ebgp
```

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

asn

The number for the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

max-paths

The maximum number of eBGP paths to a destination to be inserted in the routing table. The numbers range from 2 through 64, 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.

Configuration mode

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

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 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 \(page 125\)](#)), 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 \(page 120\)](#)), the MED is compared across all paths from all AS. The set of multipaths can be distributed from one AS to another.

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> parameters maximum-paths ibgp <max-paths>

Specifies the maximum number of iBGP paths to a destination.

Syntax:

```
set protocols bgp asn parameters maximum-paths ibgp max-paths
```

Syntax:

```
delete protocols bgp asn parameters maximum-paths ibgp [ max-paths ]
```

Syntax:

```
show protocols bgp asn parameters maximum-paths ibgp
```

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

asn

The number for the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

max-paths

The maximum number of iBGP paths to a destination to be inserted in the routing table. The numbers range from 2 through 64, 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.

Configuration mode

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

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 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 \(page 125\)](#)), 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 \(page 120\)](#)), the MED is compared across all paths from all AS. The set of multipaths can be distributed across AS.

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> parameters no-client-to-client-reflection

Disables client-to-client route reflection.

Syntax:

```
set protocols bgp asn parameters no-client-to-client-reflection
```

Syntax:

```
delete protocols bgp asn parameters no-client-to-client-reflection
```

Syntax:

```
show protocols bgp asn parameters
```

By default, client-to-client route reflection is disabled.

asn

The number of the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

Configuration mode

```
protocols {
  bgp asn {
    parameters {
      no-client-to-client-reflection
    }
  }
}
```

Use the **set** form of this command to disable client-to-client route reflection.

Use the **delete** form of this command to restore default setting (no client-to-client route reflection).

Use the **show** form of this command to view client-to-client route reflection.



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
```

Syntax:

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

Syntax:

```
show protocols bgp asn parameters
```

Disabled. By default, sessions are automatically reset.

asn

The number of the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

Configuration mode

```
protocols {
  bgp asn {
    parameters {
      no-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 `set` form of this command to disable fast external failover.

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 no-rtm

Enables the installing of BGP routes into the main routing table.

Syntax:

```
set protocols bgp asn parameters no-rtm
```

Syntax:

```
delete protocols bgp asn parameters no-rtm
```

Syntax:

```
show protocols bgp asn parameters
```

Disabled.

asn

The number of the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private autonomous systems.

Configuration mode

```
protocols {
  bgp asn {
```



```
    parameters {
      no-rtm
    }
  }
}
```

This command enables better support of iBGP route reflector for iBGP scaling. When configured, learned eBGP and iBGP routes are advertised to the peers but not installed into the main routing table. As a result, these routes cannot be forwarded by this router. The `no-rtm` command should be configured first before BGP peering. The `no-rtm` command takes effect immediately once configured.

Previous routes should already be installed to the main routing table prior to issuing this command. Therefore, configure, commit, and save the `no-rtm` command, then reboot the system to ensure a consistent state with BGP route installations.

Use the `set` form of this command to enable the installing of BGP routes into the main routing table.

Use the `delete` form of this command to disable the installing of BGP routes into the main routing table.

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
```

Syntax:

```
delete protocols bgp asn parameters router-id [ id ]
```

Syntax:

```
show protocols bgp asn parameters
```

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.

asn

The number of the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

id

The ID to be used by the router as the BGP router ID.

Configuration mode

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

Use the `set` form of this command to configure a fixed router ID for the local BGP routing process. This ID overrides 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
```

Syntax:

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

Syntax:

```
show protocols bgp asn parameters scan-time
```

Scans take place every 15 seconds.

asn

The number of the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

seconds

The interval, in seconds, at which the router scans for BGP routing information. The numbers range from 5 through 60. The default is 15.

Configuration mode

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

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 }
```

Syntax:

```
delete protocols bgp asn timers [ keepalive | holdtime ]
```

Syntax:

```
show protocols bgp asn timers [ keepalive | holdtime ]
```

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

asn

The number for the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

keepalive* *seconds

The frequency, in seconds, with which the local router sends keep-alive messages to neighbors. The numbers range from 1 through 65535. The default is 60.

holdtime* *seconds

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 through 65535, where 0 disables the holdtime timer. The default is 180.

Configuration mode



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

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 } [ soft [ in [ prefix-filter ] | out ] ]
```

Both inbound and outbound sessions are reset.

ipv4

Resets the connection for the IPv4 BGP neighbor at the specified address.

ipv6

Resets the connection for the IPv6 BGP neighbor at the specified address.

soft

Uses soft reconfiguration for the reset.

in

Resets with soft reconfiguration for only inbound sessions.

prefix-filter

Resets the outbound route filter (ORF) filter.

out

Resets with soft reconfiguration only outbound sessions.

Operational mode

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

reset ip bgp <address> ipv4 unicast <action>

Resets the BGP neighbor for the specified IPv4 unicast address family.

Syntax:

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

ipv4

Resets the connection for the IPv4 BGP neighbor at the specified address.

ipv6

Resets the connection for the IPv6 BGP neighbor at the specified address.

soft

Uses soft reconfiguration for the reset.

in

Resets with soft reconfiguration for only inbound sessions.

prefix-filter

Resets the existing ORF prefix-list for specified IPv4 unicast neighbor.

out



Resets with soft reconfiguration for only outbound sessions.

Operational mode

Use this command to reset the BGP neighbor for the specified IPv4 unicast address family.

reset ip bgp <address> ipv6 unicast <action>

Resets the BGP neighbor for the specified IPv6 unicast address family.

Syntax:

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

ipv4

Resets the connection for the IPv4 BGP neighbor at the specified address.

ipv6

Resets the connection for the IPv6 BGP neighbor at the specified address.

soft

Uses soft reconfiguration for the reset.

in

Resets with soft reconfiguration only inbound sessions.

prefix-filter

Resets the outbound route filter (ORF) filter prefix-list for the specified IPv6 unicast neighbor.

out

Resets with soft reconfiguration only outbound sessions.

Operational mode

Use this command to reset the BGP neighbor for the specified IPv6 unicast address family.

reset ip bgp all

Resets all BGP peering sessions.

Syntax:

```
reset ip bgp all
```

Operational mode

Use this command to reset all BGP peering sessions.

reset ip bgp all ipv4 unicast

Resets all IPv4 BGP unicast peering sessions.

Syntax:

```
reset ip bgp all ipv4 unicast
```

Operational mode

Use this command to reset all IPv4 BGP unicast peering sessions.

reset ip bgp all ipv4 unicast soft

Resets all IPv4 BGP unicast peering sessions without clearing the BGP session information.

Syntax:

```
reset ip bgp all ipv4 unicast soft [ in [ prefix-filter ] | out ]
```

in

Performs a reset and restore of BGP routes on inbound traffic. When soft-reconfiguration inbound is not configured, the system sends a route refresh notice to its peers.

prefix-filter



- out** Stores the received prefix and checks the policy rules to add the proper prefix to inbound traffic.
- out** Performs a reset and restore of BGP routes on outbound traffic and sends a route update to its peers. When soft-reconfiguration outbound is configured, the system checks the BGP RIB, policy rule, and sends out the proper routes in an update notice to its peers.

Operational mode

Use this command to reset all IPv4 BGP unicast peering sessions without clearing the BGP session information.

reset ip bgp all ipv6 unicast

Resets all IPv6 BGP unicast peering sessions.

Syntax:

```
reset ip bgp all ipv6 unicast
```

Operational mode

Use this command to reset all IPv6 BGP unicast peering sessions.

reset ip bgp all ipv6 unicast soft

Resets all IPv6 BGP unicast peering sessions without clearing the BGP session information.

Syntax:

```
reset ip bgp all ipv6 unicast soft [ in [ prefix-filter ] | out ]
```

in

Performs a reset and restore of BGP routes on inbound traffic. When soft-reconfiguration inbound is not configured, the system sends a route refresh notice to its peers.

prefix-filter

Stores the received prefix and checks the policy rules to add the proper prefix to inbound traffic.

out

Performs a reset and restore of BGP routes on outbound traffic and sends a route update to its peers. When soft-reconfiguration outbound is configured, the system checks the BGP RIB, policy rule, and sends out the proper routes in an update notice to its peers.

Operational mode

Use this command to reset all IPv6 BGP unicast peering sessions without clearing the BGP session information.

reset ip bgp all soft

Resets the BGP routing tables and removes the TCP session between two devices, but does not clear the BGP establishment of a peering session.

Syntax:

```
reset ip bgp all soft [ in [ prefix-filter ] | out ]
```

in

Performs a reset and restore of BGP routes on inbound traffic. When soft-reconfiguration inbound is not configured, the system sends a route refresh notice to its peers.

prefix-filter

Stores the received prefix and checks the policy rules to add the proper prefix to inbound traffic.

out

Performs a reset and restore of BGP routes on outbound traffic and sends a route update to its peers. When soft-reconfiguration outbound is configured, the system checks the BGP RIB, policy rule, and sends out the proper routes in an update notice to its peers.

Operational mode



Use this command to reset the BGP routing tables and removes the TCP session between two devices, but does not clear the BGP establishment of a peering session.

reset ip bgp ipv4 unicast dampening

Resets BGP route flap dampening information for IPv4 unicast.

Syntax:

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

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

ipv4

Resets route flap dampening information for the IPv4 peer at the specified address.

ipv4-mask

An IPv4 network mask.

ipv4net

Resets route flap dampening information for all peers on the specified network. The format is *ip-address/prefix*.

Operational mode

Use this command to reset BGP route flap dampening information for IPv4 unicast.

reset ip bgp ipv6 unicast dampening

Resets BGP route flap dampening information for IPv6 unicast.

Syntax:

```
reset ip bgp ipv6 unicast dampening [ ipv6 [ ipv6-mask ] | ipv6net ]
```

ipv6

Resets route flap dampening information for the IPv6 peer at the specified address.

ipv6-mask

An IPv6 network mask.

ipv6net

Resets route flap dampening information for all peers on the specified network. The format is *ip-address/prefix*.

Operational mode

Use this command to reset BGP route flap dampening information for IPv6 unicast.

show ip bgp

Displays BGP routes.

Syntax:

```
show ip bgp [ summary ]
```

Displays all BGP routes by default.

summary

Shows summary BGP route information.

Operational mode

Use this command to display the BGP routing table.

show ip bgp attribute-info

Displays BGP attribute information.

**Syntax:**

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

Operational mode

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
```

Operational mode

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

show ip bgp cluster-ids

Displays the BGP cluster IDs.

Syntax:

```
show ip bgp cluster-ids
```

Operational mode

Use this command to display BGP cluster IDs.

show ip bgp community <community>

Displays BGP routes belonging to the specified BGP community.

Syntax:

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

community

A BGP community identifier in the form *AA:NN* (where *AA* and *NN* are in the range of 0 through 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

Displays only routes that have an exact match.

Operational mode

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

show ip bgp community-info

Displays BGP community information.

Syntax:

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

Operational mode

Use this command to display BGP community information.

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 ]
```

list-name

A preconfigured list of BGP communities.

exact-match

Displays only routes that have an exact match.

Operational mode

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

show ip bgp dampening dampened-paths <action>

Displays IPv4 BGP routes that are currently dampened.

Syntax:

```
show ip bgp dampening { dampened-paths | flap-statistics }
```

dampened-paths

Display dampened BGP paths.

flap-statistics

Display statistics for BGP route flaps.

Operational mode

Use this command to display the IPv4 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
```

as-path-list-name

Name of an autonomous system path list.

Operational mode

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

BGP filter lists are defined using the `policy` command. For information about creating filter lists, see the AT&T Vyatta Network Operating System Routing Policies Configuration Guide.

show ip bgp ipv4 unicast

Displays information for IPv4 unicast BGP routes.

Syntax:

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

Displays all IPv4 unicast BGP routes by default.

ipv4

Show BGP information for the specified address.

ipv4net

Show BGP information for the specified network.

longer-prefixes

Displays any routes more specific than the one specified.

summary

Displays summary IPv4 unicast route information.

Operational mode



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 Classless Inter Domain Routing (CIDR) network masks.

Syntax:

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

Operational mode

Use this command to display only BGP IPv4 unicast routes with non-natural network masks; that is, 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 ]
```

community

A BGP community identifier in the form *AA:NN* (where *AA* and *NN* are in the range of 0 through 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

Displays only routes that have an exact match.

Operational mode

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 ]
```

list-name

A preconfigured list of BGP communities.

exact-match

Displays routes that have an exact match.

Operational mode

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
```

as-path-list-name

Name of an autonomous system path list.

Operational mode

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
```

Operational mode

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
```

list-name

Name of a defined prefix list.

Operational mode

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 AT&T Vyatta Network Operating System Routing Policies Configuration Guide.

show ip bgp ipv4 unicast regexp <expression>

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

Syntax:

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

expression

A POSIX-style regular expression representing a set of AS paths.

Operational mode

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
```

map-name

Name of a defined route map.

Operational mode

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 AT&T Vyatta Network Operating System Routing Policies Configuration Guide.

show ip bgp ipv4 unicast summary

Displays a summary of BGP IPv4 unicast information.

Syntax:



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

Operational mode

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

show ip bgp ipv6 unicast

Displays information for IPv6 unicast BGP routes.

Syntax:

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

Displays all IPv6 unicast BGP routes by default.

ipv6

Displays routes for the neighbor at the specified IPv6 address.

ipv6net

Displays routes for the specified IPv6 network.

longer-prefixes

Displays any routes more specific than the one specified.

summary

Shows summary BGP route information for BGP neighbor.

Operational mode

Use this command to display IPv6 unicast routes in the BGP routing table.

show ip bgp ipv6 unicast community <community>

Displays BGP routes belonging to the specified BGP community.

Syntax:

```
show ip bgp ipv6 unicast community [ AA : NN | exact-match | local-AS | no-advertise | no-export ]
```

AA : NN

A BGP community identifier in the form *AA:NN* (where *AA* and *NN* are in the range of 0 through 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

Displays only routes that have an exact match.

local-AS

Show BGP information for specified community number.

no-advertise

Show routes not advertised to any peer.

no-export

Show routes not exported to next AS.

Operational mode

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

show ip bgp ipv6 unicast community-list <list-name>

Displays BGP routes permitted by the specified community list.

Syntax:

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

list-name

A preconfigured list of BGP communities.

exact-match

Displays routes that have an exact match.



Operational mode

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

show ip bgp ipv6 unicast dampening <action>

Displays detailed information about dampening in IPv6 environments.

Syntax:

```
show ip bgp ipv6 unicast dampening { dampened-paths | flap-statistics }
```

dampened-paths

Display paths suppressed due to dampening.

flap-statistics

Display flap statistics of routes.

Operational mode

Use this command to display detailed information about dampening in IPv6 environments. Enable BGP IPv6 dampening to maintain dampened-path information in memory.

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

Displays routes matching a list of autonomous system paths.

Syntax:

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

as-path-list-name

Name of an autonomous system path list.

Operational mode

Use this command to filter displayed routes according to a preconfigured list of autonomous system paths. BGP filter lists are defined using the `policy as-path-list` command.

show ip bgp ipv6 unicast prefix-list <list-name>

Displays IPv6 unicast BGP routes matching a prefix list.

Syntax:

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

list-name

Name of a defined prefix list.

Operational mode

Use this command to display BGP IPv6 unicast routes matching a preconfigured prefix list.

Prefix lists are configured using the `policy` command. For information about creating prefix lists, see the AT&T Vyatta Network Operating System Routing Policies Configuration Guide.

show ip bgp ipv6 unicast regexp <expression>

Displays routes matching an AS path regular expression.

Syntax:

```
show ip bgp ipv6 unicast regexp expression
```

expression

A POSIX-style regular expression representing a set of AS paths.

**Operational mode**

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

show ip bgp paths

Displays all BGP paths.

Syntax:

```
show ip bgp paths
```

Operational mode

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
```

list-name

Name of a defined prefix list.

Operational mode

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 AT&T Vyatta Network Operating System Routing Policies Configuration Guide.

show ip bgp regexp <expression>

Displays routes matching an AS path regular expression.

Syntax:

```
show ip bgp regexp expression
```

expression

A POSIX-style regular expression representing a set of AS paths.

Operational mode

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

show ip bgp route-map <map-name>

Displays routes matching a route map.

Syntax:

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

map-name

The name of a defined route map.

Operational mode

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 AT&T Vyatta Network Operating System Routing Policies Configuration Guide.



show ip bgp scan

Displays BGP scan status.

Syntax:

```
show ip bgp scan
```

Operational mode

Use this command to display BGP scan status.

show ip route bgp

Displays BGP routes.

Syntax:

```
show ip route bgp
```

Operational mode

Use this command to display BGP routes.

The following example shows 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 monitoring protocols bgp

Displays BGP protocol debugging flags.

Syntax:



```
show monitoring protocols bgp
```

Operational mode

Use this command to display BGP protocol debugging flags.



Route Reflection

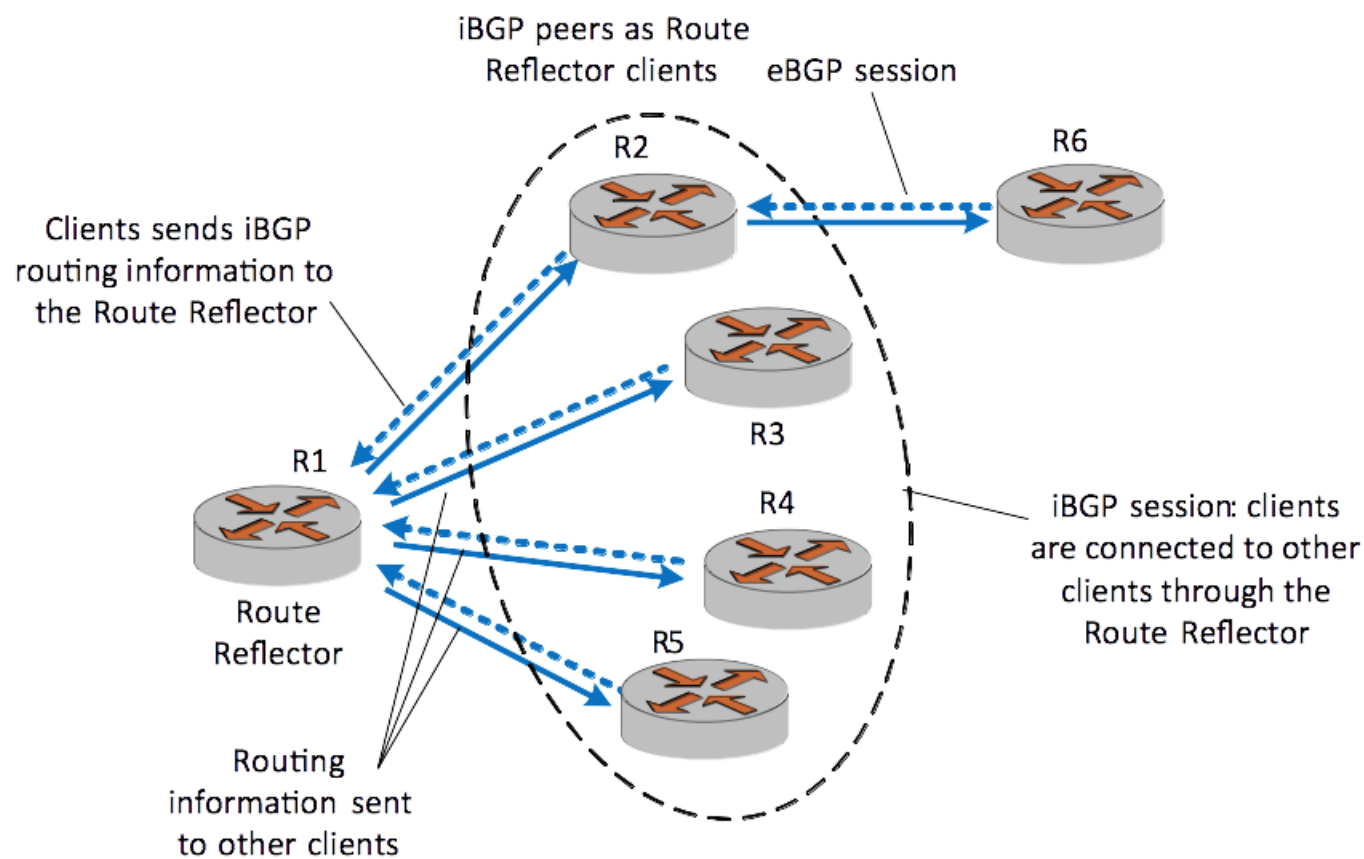
Overview

BGP Virtual Route Reflectors are typically deployed by installing hardware capable of handling high speed data plane forwarding and is deployed without accompanying hardware. Virtual Route Reflector functionality in the software enables routing with fewer routing pairs and eliminates the need for fully meshed Asynchronous Systems (AS).

Multiple peer routers can connect through iBGP with the route reflector which serves as a central point-of-contact server, enabling access to clients of the route reflector, that is, other routers participating in a full mesh. One or more router is configured as route reflector in a BGP topology to provide redundancy.

Route reflector requires BGP to be enabled. The BGP software component is not part on the data plane architecture.

Figure 19: Route reflector overview



BGP multiple cluster IDs

The BGP multiple cluster IDs let you to configure a route reflector such that the route reflector associates a cluster ID with each peer to unite groups of clients into different clusters. This association also allows you to disable the reflection of routes between route reflection clients in the same cluster, that is, intra-cluster route reflection.



A route reflector and its clients form a cluster. If the route reflector is connected to a group of clients that are fully meshed, you can assign a cluster ID to each peer of the route reflector peer or assign a cluster ID to a peer group. You can also disable client-to-client route reflection for the cluster that is connected to the route reflector. On disabling the intra-cluster client-to-client route reflection, the loop prevention mechanism is modified automatically to consider multiple cluster IDs.

You can reduce the number of updates in the network by disabling the intra-cluster client-to-client route reflection which in-turn increases the convergence speed.

Configuring multiple cluster IDs


In this section, you configure multiple cluster IDs for a network that contains routers and route reflectors. In the example, there are five routers one of which is a route reflector (RR), and four routers are route reflector clients (R1, R2, R3, and R4).

Clients R1 and R2 are fully meshed and the two routers can share prefixes with each other without involving the RR.

There are two logical clusters, cluster 1 and cluster 2. Clients R1 and R2 are in cluster 1 and have 1 as the cluster ID that is configured on RR. Clients R3 and R4 are in cluster 2 and have 2 as the cluster ID. In the given example, RR reflects routes from R3 to R1 and from R1 to R4 but blocks route reflection between R1 and R2 because R1 and R2 are fully meshed.

You can now disable client-to-client route reflection inside a specific cluster by using `protocols bgp <asn> parameters no-client-to-client-reflection cluster-id <cluster-id>` (page 157). Therefore, RR does not reflect routes between clients from the cluster for which the client-to-client route reflection is disabled. All other types of reflection, such as client-to-non client routes, from one cluster to another continue to work. In this example the client-to-client route reflection is disabled for cluster 1.

Figure 20: BGP multiple cluster-ids



To configure multiple cluster IDs, perform the following steps in configuration mode.

**Table 17: Configuring multiple cluster IDs**

Router	Step	Command(s)
RR	Configure cluster ID as 2 for routers R3 and R4.	<pre>vyatta@RR# set protocols bgp 100 neighbor 3.3.3.3 address-family ipv4-unicast 'route-reflector-client' vyatta@RR# set protocols bgp 100 neighbor 3.3.3.3 cluster-id 2 vyatta@RR# set protocols bgp 100 neighbor 3.3.3.3 remote-as '100' vyatta@RR# set protocols bgp 100 neighbor 4.4.4.4 address-family ipv4-unicast 'route-reflector-client' vyatta@RR# set protocols bgp 100 neighbor 4.4.4.4 cluster-id 2 vyatta@RR# set protocols bgp 100 neighbor 4.4.4.4 remote-as '100'</pre>
RR	Commit the configuration.	<pre>vyatta@RR# commit</pre>
RR	Configure cluster ID as 1 for routers R1 and R2.	<pre>vyatta@RR# set protocols bgp 100 neighbor 1.1.1.1 address-family ipv4-unicast 'route-reflector-client' vyatta@RR# set protocols bgp 100 neighbor 1.1.1.1 cluster-id 1 vyatta@RR# set protocols bgp 100 neighbor 1.1.1.1 remote-as '100' vyatta@RR# set protocols bgp 100 neighbor 2.2.2.2 address-family ipv4-unicast 'route-reflector-client' vyatta@RR# set protocols bgp 100 neighbor 2.2.2.2 cluster-id 1 vyatta@RR# set protocols bgp 100 neighbor 2.2.2.2 remote-as '100'</pre>
R1	Configure R1	<pre>vyatta@R1# set protocols bgp 100 address-family ipv4- unicast network '1.1.1.1/32' vyatta@R1# set protocols bgp 100 neighbor 5.5.5.5 address-family 'ipv4- unicast' vyatta@R1# set protocols bgp 100 neighbor 5.5.5.5 remote-as '100'</pre>



Router	Step	Command(s)
R1	Commit the configuration.	<pre>vyatta@R1# commit</pre>
R2	Configure R2	<pre>vyatta@R2# set protocols bgp 100 address-family ipv4- unicast network '2.2.2.2/32' vyatta@R2# set protocols bgp 100 neighbor 5.5.5.5 address-family 'ipv4- unicast' vyatta@R2# set protocols bgp 100 neighbor 5.5.5.5 remote-as '100</pre>
R2	Commit the configuration	<pre>vyatta@R2# commit</pre>
R3	Configure R3	<pre>vyatta@R3# set protocols bgp 100 address-family ipv4- unicast network '3.3.3.3/32' vyatta@R3# set protocols bgp 100 neighbor 5.5.5.5 address-family 'ipv4- unicast' vyatta@R3# set protocols bgp 100 neighbor 5.5.5.5 remote-as '100</pre>
R3	Commit the configuration	<pre>vyatta@R3# commit</pre>
R4	Configure R4	<pre>vyatta@R4# set protocols bgp 100 address-family ipv4- unicast network '4.4.4.4/32' vyatta@R4# set protocols bgp 100 neighbor 5.5.5.5 address-family 'ipv4- unicast' vyatta@R4# set protocols bgp 100 neighbor 5.5.5.5 remote-as '100</pre>
R4	Commit the configuration	<pre>vyatta@R4# commit</pre>
RR	Disable intra-cluster client-to-client route reflection	<pre>vyatta@RR# set protocols bgp 100 parameters no- client-to-client-reflection cluster-id 1</pre>



Router	Step	Command(s)
RR	Verify the cluster IDs configured on the route-reflector.	<pre>vyatta@RR# show ip bgp cluster-ids Global cluster-id: 1.1.1.1 BGP client-to-client reflection: Configured Used all (inter-cluster and intra-cluster): Enabled intra-cluster: Enabled Enabled List of Cluster-Ids: Cluster-Id Neighbors C2C Reflection: Configured Used 0.0.0.1 2 Enabled Enabled</pre>
RR	Verify the neighbors of the route reflector.	<pre>vyatta@RR# show ip bgp neighbors 1.1.1.1 BGP neighbor is 1.1.1.1, remote AS 100, local AS 100, internal link <...> Minimum time between advertisement runs is 5 seconds Configured with cluster-id 0.0.0.1 For address family: IPv4 Unicast <...></pre>



Route Reflection Commands

protocols bgp <asn> neighbor <id> address-family ipv4-unicast route-reflector-client

Specifies this neighbor as a route reflector client for IPv4 routes.

Syntax:

```
set protocols bgp asn neighbor id address-family ipv4-unicast route-reflector-client
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast route-reflector-client
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast
```

Disabled.

asn

The number for the AS in which this router resides. The numbers ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          route-reflector-client
        }
      }
    }
  }
}
```

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.

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

Specifies this neighbor as a route reflector client for IPv6 routes.

Syntax:

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

Syntax:

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

Syntax:

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

Disabled.

asn

Mandatory. The number for the AS in which this router resides. The ranges from 1 through 4294967295.

id

Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

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

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> cluster-id <cluster-id>

Defines a cluster id for a neighbor.

Syntax:

```
set protocols bgp asn neighbor id cluster-id cluster-id
```

Syntax:

```
delete protocols bgp asn neighbor id cluster-id cluster-id
```

Syntax:

```
show protocols bgp asn neighbor id cluster-id cluster-id
```

asn

The number of the AS in which this router resides. The number ranges from 1 through 4294967295.

id

Multi-node. The IPv4 or IPv6 address of a BGP neighbor.

You can define multiple BGP neighbors by creating multiple neighbor configuration nodes.

cluster-id

Neighbor cluster ID. When you change a cluster ID for a neighbor, BGP automatically does an inbound soft refresh and an outbound soft refresh for all iBGP peers.

The cluster ID can be in dotted decimal format such as 192.168.7.4 or decimal format such as 23, with a maximum of 4 bytes.

The router modifies a cluster ID that is configured in decimal format such as 23 to dotted decimal format such as 0.0.0.23 when it appears in a configuration file.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      cluster-id cluster-id {
      }
    }
  }
}
```

Use the `set` form of this command to define a cluster ID for a neighbor.

Use the `delete` form of this command to remove a cluster ID for a neighbor.

Use the `show` form of this command to view a cluster ID for a neighbor.

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
```

Syntax:

```
delete protocols bgp asn parameters cluster-id id
```

Syntax:

```
show protocols bgp asn parameters
```

asn

The number of the AS in which this router resides. The number ranges from 1 through 4294967295.

**id**

A network address uniquely identifying the route reflection cluster.

Configuration mode

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

Use the `set` form of this command to specify the route reflection cluster-ID for an internal route reflection cluster.

When a cluster has only one route reflector, 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 all

Enables or disables route reflection from a BGP route reflector to clients.

Syntax:

```
set protocols bgp asn parameters no-client-to-client-reflection all
```

Syntax:

```
delete protocols bgp asn parameters no-client-to-client-reflection all
```

Syntax:

```
show protocols bgp asn parameters
```

By default, client-to-client-reflection is enabled.

asn

The number of the AS in which this router resides. The numbers ranges from 1 through 4294967295.

Configuration mode

```
protocols {
  bgp asn {
    parameters {
      no-client-to-client-reflection
    }
  }
}
```

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> parameters no-client-to-client-reflection cluster-id any

Enables or disables route reflection for all the intra cluster client-to-client route reflection.

Syntax:

```
set protocols bgp asn parameters no-client-to-client-reflection cluster-id any
```

Syntax:

```
delete protocols bgp asn parameters no-client-to-client-reflection cluster-id any
```

Syntax:

```
show protocols bgp asn parameters
```

asn

The number of the AS in which this router resides. The number ranges from 1 through 4294967295.

any

All clusters

Configuration mode

```
protocols {
  bgp asn {
    parameters no-client-to-client-reflection cluster-id any {
    }
  }
}
```

Use the `set` form of this command to disable intra-cluster client-to-client route reflection from a BGP route reflector to all clusters.

Use the `delete` form of this command to enable intra-cluster client-to-client route reflection from a BGP route reflector to all clusters.

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

protocols bgp <asn> parameters no-client-to-client-reflection cluster-id <cluster-id>

Enables or disables route reflection for a cluster.

Syntax:

```
set protocols bgp asn parameters no-client-to-client-reflection cluster-id cluster-id
```

Syntax:

```
delete protocols bgp asn parameters no-client-to-client-reflection cluster-id cluster-id
```

Syntax:

```
show protocols bgp asn parameters
```

asn

The number of the AS in which this router resides. The number ranges from 1 through 4294967295.

cluster-id

Neighbor cluster-id.

The cluster ID can be in dotted decimal format such as 192.168.7.4 or decimal format such as 23, with a maximum of 4 bytes.



The router modifies a cluster ID that is configured in decimal format such as 23 to dotted decimal format such as 0.0.0.23 when it appears in a configuration file.

Configuration mode

```
protocols {
  bgp asn {
    parameters no-client-to-client-reflection cluster-id cluster-id {
    }
  }
}
```

Use the `set` form of this command to disable intra-cluster client-to-client route reflection from a BGP route reflector to the specified cluster. When the software changes reflection state for the given cluster ID BGP sends an outbound soft refresh to all clients.

Use the `delete` form of this command to enable intra-cluster client-to-client route reflection from a BGP route reflector to the specified cluster.

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

protocols bgp <asn> peer-group <group-name> address-family ipv4-unicast route-reflector-client

Specify this peer group as a route reflector client.

Syntax:

```
set protocols bgp asn peer-group group-name address-family ipv4-unicast route-reflector-client
```

Syntax:

```
delete protocols bgp asn peer-group group-name address-family ipv4-unicast route-reflector-client
```

Syntax:

```
show protocols bgp asn peer-group group-name address-family ipv4-unicast
```

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

group-name

The peer group name.

Configuration mode

```
protocols {
  bgp asn {
    peer-group group-name {
      address-family {
        ipv4-unicast {
          route-reflector-client
        }
      }
    }
  }
}
```

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> 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
```

Syntax:

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

Syntax:

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

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

group-name

The peer group name.

Configuration mode

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

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> cluster-id <cluster-id>

Defines a cluster ID for a peer group.

Syntax:

```
set protocols bgp asn peer-group name cluster-id cluster-id
```

Syntax:

```
delete protocols bgp asn peer-group name cluster-id cluster-id
```

Syntax:

```
show protocols bgp asn peer-group name cluster-id cluster-id
```

asn

The number of the AS in which this router resides. The number ranges from 1 through 4294967295.

**name**

Multi-node. The peer-group name.

cluster-id

Neighbor cluster ID. When you change a cluster ID for a neighbor, BGP automatically does an inbound soft refresh and an outbound soft refresh for all iBGP peers.

The cluster ID can be in dotted decimal format such as 192.168.7.4 or decimal format such as 23, with a maximum of 4 bytes.

The vrouter modifies a cluster ID that is configured in decimal format such as 23 to dotted decimal format such as 0.0.0.23 when it appears in a configuration file.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name{
      cluster-id cluster-id {
    }
  }
}
```

Use the `set` form of this command to define a cluster ID for a peer group.

Use the `delete` form of this command to remove the cluster ID for a peer group.

Use the `show` form of this command to view the cluster ID for a peer group.



Confederation Commands

protocols bgp <asn> parameters confederation identifier <asn>

Defines a BGP confederation.

Syntax:

```
set protocols bgp asn parameters confederation identifier asn
```

Syntax:

```
delete protocols bgp asn parameters confederation identifier asn
```

Syntax:

```
show protocols bgp asn parameters confederation
```

bgp *asn*

The number of the AS in which this router resides. The number ranges from 1 through 4294967295.

identifier *asn*

The AS number of the BGP confederation. The numbers range from 1 through 4294967295.

Configuration mode

```
protocols {
  bgp asn {
    parameters {
      confederation {
        identifier asn
      }
    }
  }
}
```

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 AS that make up a BGP confederation.

Syntax:

```
set protocols bgp asn parameters confederation peers asn
```

Syntax:

```
delete protocols bgp asn parameters confederation peers asn
```

Syntax:

```
show protocols bgp asn parameters confederation
```

bgp *asn*

The number of the AS in which this router resides. The number ranges from 1 through 4294967295.

peers *asn*



The subautonomous systems that make up BGP confederation. The numbers range from 1 through 4294967295. Multiple AS can be specified in a space-separated list.

Configuration mode

```
protocols {
  bgp asn {
    parameters {
      confederation {
        peers asn
      }
    }
  }
}
```

Use the `set` form of this command to list the subautonomous systems that to 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.



Synchronization


Basic synchronization configuration

To enable traffic to be routed through multiple asynchronous systems (AS), IGP must propagate iBGP routes within its AS before they advertise them to peers outside its AS. All routers in an AS must learn of a route before BGP advertises the route outside of the AS. A BGP router must have synchronization disabled (the default setting) to validate these routes and install them into its routing table.

In this section, you configure BGP synchronization on the routers labeled BGP17, BGP18, BGP19, BGP20, and BGP21 in the reference network diagram. The routers are advertising their routes on networks 10.0.40.0/24 and 10.0.50.0/24.

This section presents a sample configuration for BGP synchronization. The configuration example is based on the reference diagram in the following figure.

Figure 21: BGP synchronization topology



BGP synchronization must have AS200 configured as a transit between AS100 and AS300, and AS200 must have all routers meshed through iBGP.

BGP17 advertises a route to BGP18 through eBGP and the route is forwarded using iBGP to BGP20 through BGP19. For BGP20 to forward the route outside its AS to BGP 21, BGP19 must add the network in its routing table or the packet be dropped. BGP synchronization enables BGP19 to learn the network and to notify BGP20, before allowing BGP20 to forward the route to BGP21.

To create a basic BGP synchronization configuration, perform the steps in the following table:

**Table 18: Basic BGP synchronization configuration**

Router	Step	Commands
BGP17	Configure interfaces.	<pre>vyatta@BGP17# set interfaces dataplane dp0p160p1 address 10.1.34.133/24 vyatta@BGP17# set interfaces dataplane dp0p192p1 address 20.1.1.2/24 vyatta@BGP17# set interfaces dataplane dp0p224p1 address 100.1.1.1/24 vyatta@BGP17# set interfaces loopback lo</pre>
BGP17	Establish BGP, by configuring BGP neighbors.	<pre>vyatta@BGP17# set protocols bgp 100 neighbor 20.1.1.1 ebgp-multihop 1 vyatta@BGP17# set protocols bgp 100 neighbor 20.1.1.1 remote-as 200</pre>
BGP17	Advertise the route to establish the BGP route as the next hop and reachable. The eBGP route becomes an iBGP route.	<pre>vyatta@BGP17# set protocols bgp 100 network 100.1.1.0/24 vyatta@BGP17# set protocols ospf area 0.0.0.0 network 20.1.1.0/24</pre>
BGP18	Configure interfaces.	<pre>vyatta@BGP18# set interfaces dataplane dp0p160p1 address 10.1.34.134/24 vyatta@BGP18# set interfaces dataplane dp0p192p1 address 20.1.1.1/24 vyatta@BGP18# set interfaces dataplane dp0p224p1 address 30.1.1.1/24 vyatta@BGP18# set interfaces loopback lo address 175.18.18.18/32</pre>
BGP18	Establish BGP, by configuring BGP neighbors.	<pre>vyatta@BGP18# set protocols bgp 200 neighbor 20.1.1.2 ebgp-multihop 1 vyatta@BGP18# set protocols bgp 200 neighbor 20.1.1.2 remote-as 100 vyatta@BGP18# set protocols bgp 200 neighbor 175.20.20.20 remote-as 200 vyatta@BGP18# set protocols bgp 200 neighbor 175.20.20.20 update-source 175.18.18.18</pre>



Router	Step	Commands
BGP18	Advertise the route to establish the BGP route as the next hop and reachable. The eBGP route becomes an iBGP route.	<pre>vyatta@BGP18# set protocols bgp 200 parameters vyatta@BGP18# set protocols ospf area 0.0.0.0 network 20.1.1.0/24 vyatta@BGP18# set protocols ospf area 0.0.0.0 network 30.1.1.0/24 vyatta@BGP18# set protocols ospf area 0.0.0.0 network 175.18.18.18/32 vyatta@BGP18# set protocols ospf redistribute bgp</pre>
BGP19	Configure interfaces.	<pre>vyatta@BGP19# set interfaces dataplane dp0p160p1 address 10.1.34.135/24 vyatta@BGP19# set interfaces dataplane dp0p192p1 address 30.1.1.2/24 vyatta@BGP19# set interfaces dataplane dp0p224p1 address 40.1.1.2/24 vyatta@BGP19# set interfaces loopback lo</pre>
BGP19	Advertise the route to establish the BGP route as the next hop and reachable. The eBGP route becomes an iBGP route.	<pre>vyatta@BGP19# set protocols ospf area 0.0.0.0 network 20.1.1.0/24 vyatta@BGP19# set protocols ospf area 0.0.0.0 network 30.1.1.0/24 vyatta@BGP19# set protocols ospf area 0.0.0.0 network 40.1.1.0/24</pre>
BGP20	Configure Interfaces	<pre>vyatta@BGP20# set interfaces dataplane dp0p160p1 address 10.1.34.136/24 vyatta@BGP20# set interfaces dataplane dp0p192p1 address 40.1.1.1/24 vyatta@BGP20# set interfaces dataplane dp0p224p1 address 50.1.1.1/24 vyatta@BGP20# set interfaces loopback lo address 175.20.20.20/32</pre>



Router	Step	Commands
BGP20	Establish BGP, by configuring BGP neighbors.	<pre>vyatta@BGP20# set protocols bgp 200 neighbor 50.1.1.2 ebgp-multihop 1 vyatta@BGP20# set protocols bgp 200 neighbor 50.1.1.2 remote-as 300 vyatta@BGP20# set protocols bgp 200 neighbor 175.18.18.18 remote-as 200 vyatta@BGP20# set protocols bgp 200 neighbor 175.18.18.18 soft- reconfiguration inbound vyatta@BGP20# set protocols bgp 200 neighbor 175.18.18.18 update-source 175.20.20.20</pre>
BGP20	Configure synchronization and advertise the route to establish the BGP route as the next hop and reachable. The eBGP route becomes an iBGP route.	<pre>vyatta@BGP20# set protocols bgp 200 parameters synchronization vyatta@BGP20# set protocols ospf area 0.0.0.0 network 40.1.1.0/24 vyatta@BGP20# set protocols ospf area 0.0.0.0 network 50.1.1.0/24 vyatta@BGP20# set protocols ospf area 0.0.0.0 network 175.20.20.20/32</pre>
BGP21	Configure Interfaces	<pre>vyatta@BGP21# set interfaces dataplane dp0p160p1 address 10.1.34.137/24 vyatta@BGP21#set interfaces dataplane dp0p192p1 address 50.1.1.2/24 vyatta@BGP21# set interfaces loopback lo address 175.21.21.21/32</pre>
BGP21	Establish BGP, by configuring BGP neighbors.	<pre>vyatta@BGP21# set protocols bgp 300 neighbor 50.1.1.1 ebgp-multihop 1 vyatta@BGP21# set protocols bgp 300 neighbor 50.1.1.1 remote-as 200 vyatta@BGP21# set protocols bgp 300 neighbor 50.1.1.1 soft-reconfiguration inbound</pre>



Router	Step	Commands
BGP21	Advertise the route to establish the BGP route as the next hop and reachable. The eBGP route becomes an iBGP route.	<pre>vyatta@BGP21#set protocols ospf area 0.0.0.0 network 50.1.1.0/24 vyatta@BGP21# set protocols ospf area 0.0.0.0 network 175.21.21.21/32</pre>

Verifying the configuration

To display routes and their status for BGP 17, enter the command in the following example:

```
vyatta@BGP17# show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
> - selected route, * - FIB route, p - stale info

S *> 10.0.0.0/8 [1/0] via 10.1.34.1, dp0p160p1
C *> 10.1.34.0/24 is directly connected, dp0p160p1
O 20.1.1.0/24 [110/10] is directly connected, dp0p192p1, 15:39:00
C *> 20.1.1.0/24 is directly connected, dp0p192p1
O *> 30.1.1.0/24 [110/20] via 20.1.1.1, dp0p192p1, 00:26:54
O *> 40.1.1.0/24 [110/30] via 20.1.1.1, dp0p192p1, 00:26:54
O *> 50.1.1.0/24 [110/40] via 20.1.1.1, dp0p192p1, 00:26:54
C *> 100.1.1.0/24 is directly connected, dp0p224p1
C *> 127.0.0.0/8 is directly connected, lo
O *> 175.18.18.18/32 [110/20] via 20.1.1.1, dp0p192p1, 00:26:54
O *> 175.20.20.20/32 [110/40] via 20.1.1.1, dp0p192p1, 00:26:54
O *> 175.21.21.21/32 [110/50] via 20.1.1.1, dp0p192p1, 00:26:54
```

To verify the network, next hop routes, and their status in the BGP synchronization configuration, enter the command in the following example:

```
vyatta@BGP17# show ip bgp neigh 20.1.1.1 advertised-routes
BGP table version is 2, local router ID is 100.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 100.1.1.0/24     20.1.1.2           100   32768  i

Total number of prefixes 1
```

To display the routes and their status for BGP 18, enter the command in the following example:

```
vyatta@BGP18# show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
> - selected route, * - FIB route, p - stale info
S *> 10.0.0.0/8 [1/0] via 10.1.34.1, dp0p160p1
C *> 10.1.34.0/24 is directly connected, dp0p160p1
O 20.1.1.0/24 [110/10] is directly connected, dp0p192p1, 1d00h18m
C *> 20.1.1.0/24 is directly connected, dp0p192p1
O 30.1.1.0/24 [110/10] is directly connected, dp0p224p1, 21:20:09
C *> 30.1.1.0/24 is directly connected, dp0p224p1
```



```

O *> 40.1.1.0/24 [110/20] via 30.1.1.2, dp0p224p1, 16:28:46
O *> 50.1.1.0/24 [110/30] via 30.1.1.2, dp0p224p1, 00:50:46
B *> 100.1.1.0/24 [20/0] via 20.1.1.2, dp0p192p1, 00:28:08
C *> 127.0.0.0/8 is directly connected, lo
C *> 175.18.18.18/32 is directly connected, lo
O *> 175.20.20.20/32 [110/30] via 30.1.1.2, dp0p224p1, 00:40:19
O *> 175.21.21.21/32 [110/40] via 30.1.1.2, dp0p224p1, 00:44:16

```

To display advertised routes for BGP 18, enter the command in the following example:

```

vyatta@BGP18# show ip bgp neigh 175.20.20.20 advertised-routes

BGP table version is 3, local router ID is 175.18.18.18
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*>i100.1.1.0/24    20.1.1.2                100      0 100 i

Total number of prefixes 1

```

To display ospf updates for BGP 18, enter the command in the following example:

```

vyatta@BGP18# show ip ospf data

OSPF Router with ID (20.1.1.1) (Process ID 0)

          Router Link States (Area 0.0.0.0)

Link ID      ADV Router    Age Seq#      CkSum Link count
20.1.1.1    20.1.1.1     1338 0x80000180 0xc944 3
20.1.1.2    20.1.1.2     1740 0x80000172 0x4e29 1
30.1.1.2    30.1.1.2     1584 0x80000030 0x6ab7 2
40.1.1.1    40.1.1.1     656 0x8000002a 0x2bbe 3
50.1.1.2    50.1.1.2     906 0x80000007 0x86db 2

          Net Link States (Area 0.0.0.0)

Link ID      ADV Router    Age Seq#      CkSum
20.1.1.2    20.1.1.2     1736 0x80000007 0x2bc9
30.1.1.2    30.1.1.2     964 0x80000027 0x5462
40.1.1.1    40.1.1.1     1036 0x80000002 0x8a2b
50.1.1.2    50.1.1.2     1466 0x80000002 0x4e3d

          AS External Link States

Link ID      ADV Router    Age Seq#      CkSum Route          Tag
100.1.1.0    20.1.1.1     1337 0x80000001 0x0537 E2 100.1.1.0/24 0

```

To display routes and their status for BGP 19, enter command in the following example:

```

vyatta@BGP19# run show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       > - selected route, * - FIB route, p - stale info

S   *> 10.0.0.0/8 [1/0] via 10.1.34.1, dp0p160p1
C   *> 10.1.34.0/24 is directly connected, dp0p160p1
O   *> 20.1.1.0/24 [110/20] via 30.1.1.1, dp0p192p1, 15:18:49
O   30.1.1.0/24 [110/10] is directly connected, dp0p192p1, 20:58:07
C   *> 30.1.1.0/24 is directly connected, dp0p192p1
O   40.1.1.0/24 [110/10] is directly connected, dp0p224p1, 16:08:04
C   *> 40.1.1.0/24 is directly connected, dp0p224p1

```




```

O   *> 50.1.1.0/24 [110/20] via 40.1.1.1, dp0p224p1, 00:30:04
O E2 *> 100.1.1.0/24 [110/1] via 30.1.1.1, dp0p192p1, 00:00:05
C   *> 127.0.0.0/8 is directly connected, lo
O   *> 175.18.18.18/32 [110/20] via 30.1.1.1, dp0p192p1, 00:25:26
O   *> 175.20.20.20/32 [110/20] via 40.1.1.1, dp0p224p1, 00:19:36
O   *> 175.21.21.21/32 [110/30] via 40.1.1.1, dp0p224p1, 00:23:34

```

To display advertised routes for BGP 20, enter the command in the following example:

```

vyatta@BGP20# show ip bgp neigh 50.1.1.2 advertised-routes

BGP table version is 6, local router ID is 175.20.20.20
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 100.1.1.0/24     50.1.1.1              100      0 100 i

Total number of prefixes 1

```

To display routes and their status for BGP 20, enter the command in the following example:

```

vyatta@BGP20# run show ip route
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       > - selected route, * - FIB route, p - stale info

S   *> 10.0.0.0/8 [1/0] via 10.1.34.1, dp0p160p1
C   *> 10.1.34.0/24 is directly connected, dp0p160p1
O   *> 20.1.1.0/24 [110/30] via 40.1.1.2, dp0p192p1, 01:00:23
O   *> 30.1.1.0/24 [110/20] via 40.1.1.2, dp0p192p1, 01:00:23
O   40.1.1.0/24 [110/10] is directly connected, dp0p192p1, 16:34:02
C   *> 40.1.1.0/24 is directly connected, dp0p192p1
O   50.1.1.0/24 [110/10] is directly connected, dp0p224p1, 00:56:19
C   *> 50.1.1.0/24 is directly connected, dp0p224p1
O E2 *> 100.1.1.0/24 [110/1] via 40.1.1.2, dp0p192p1, 00:26:20
B   100.1.1.0/24 [200/0] via 20.1.1.2 (recursive via 40.1.1.2 ), 00:26:44
C   *> 127.0.0.0/8 is directly connected, lo
O   *> 175.18.18.18/32 [110/30] via 40.1.1.2, dp0p192p1, 00:51:41
C   *> 175.20.20.20/32 is directly connected, lo
O   *> 175.21.21.21/32 [110/20] via 50.1.1.2, dp0p224p1, 00:49:48

```

To verify the network, next hop routes, and their status in the BGP synchronization configuration, enter the command in the following example:

```

vyatta@BGP20# show ip bgp
BGP table version is 6, local router ID is 175.20.20.20
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal, l - labeled
              S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*>i100.1.1.0/24     20.1.1.2              0    100      0 100 i

Total number of prefixes 1

```

To display received routes for BGP 21, enter the command in the following example:

```

vyatta@BGP21# show ip bgp neigh 50.1.1.1 received-routes
BGP table version is 8, local router ID is 175.21.21.21
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal
Origin codes: i - IGP, e - EGP, ? - incomplete

```



```
Network      Next Hop      Metric LocPrf Weight Path
*> 100.1.1.0/24  50.1.1.1          0 200 100 i
```

```
Total number of prefixes 1
```

To display routes and their status for BGP 21, enter the command in the following example:

```
vyatta@BGP21# show ip routes
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
       O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
       > - selected route, * - FIB route, p - stale info

S   *> 10.0.0.0/8 [1/0] via 10.1.34.1, dp0p160p1
C   *> 10.1.34.0/24 is directly connected, dp0p160p1
O   *> 20.1.1.0/24 [110/40] via 50.1.1.1, dp0p192p1, 00:55:30
O   *> 30.1.1.0/24 [110/30] via 50.1.1.1, dp0p192p1, 00:55:30
O   *> 40.1.1.0/24 [110/20] via 50.1.1.1, dp0p192p1, 00:55:30
C   50.1.1.0/24 [110/10] is directly connected, dp0p192p1, 00:56:54
C   *> 50.1.1.0/24 is directly connected, dp0p192p1
O E2 100.1.1.0/24 [110/1] via 50.1.1.1, dp0p192p1, 00:28:38
B   *> 100.1.1.0/24 [20/0] via 50.1.1.1, dp0p192p1, 00:28:47
C   *> 127.0.0.0/8 is directly connected, lo
O   *> 175.18.18.18/32 [110/40] via 50.1.1.1, dp0p192p1, 00:54:00
O   *> 175.20.20.20/32 [110/20] via 50.1.1.1, dp0p192p1, 00:48:10
C   *> 175.21.21.21/32 is directly connected, lo
```

Graceful restart

To configure graceful restart if it has not been configured, enter the following command:

- `protocols bgp <asn> parameters graceful-restart <action>` ([page 127](#))



Neighbor Commands

protocols bgp <asn> neighbor <id>

Defines a BGP neighbor.

Syntax:

```
set protocols bgp asn neighbor id
```

Syntax:

```
delete protocols bgp asn neighbor id
```

Syntax:

```
show protocols bgp asn neighbor id
```

asn

The number for the autonomous system (AS) in which this router resides. The number ranges from 1 through 4294967295.

id

Multi-node. The IPv4 or IPv6 address of a BGP neighbor.

You can define multiple BGP neighbors by creating multiple neighbor configuration nodes.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
    }
  }
}
```

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 ipv4-unicast

Enables the IPv4 unicast address-family for a neighbor and defines parameters for controlling IPv4 unicast routes when subsequent configuration parameters and options are present.

Syntax:

```
set protocols bgp asn neighbor id address-family ipv4-unicast
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast
```

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

**id**

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
        }
      }
    }
  }
}
```

Use the `set` form of this command to enable the IPv4 unicast address-family for a neighbor and to define parameters for controlling IPv4 unicast routes when subsequent configuration parameters and options are present.

Use the `delete` form of this command to disable the IPv4 unicast address-family for a neighbor and to remove the configured parameters for controlling IPv4 unicast routes.

Use the `show` form of this command to view the configuration settings for the neighbor IPv4 unicast address-family.

protocols bgp <asn> neighbor <id> address-family ipv4-unicast allowas-in number <number>

Allows or disallows receiving BGP advertisements containing the AS path of the local router.

Syntax:

```
set protocols bgp asn neighbor id address-family ipv4-unicast allowas-in number number
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast allowas-in number number
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast
```

Readvertisement is disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

number

Specifies the number of times to accept an AS path containing the readvertisement of the ASN of the local router. The numbers range from 1 through 10. The default is 3.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          allowas-in {
            number number
          }
        }
      }
    }
  }
}
```



```
    }  
  }  
}
```

Use the `set` form of this command to allow the router to accept BGP AS paths advertising the AS number of the router.

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 AS path of the router.

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

Note:

When you configure this command and do a soft reset in the peer, the BGP table should flush out the prefixes that have an as-path length that is more than the configured `allows-in` value.

When soft-reconfiguration is enabled irrespective of the configured `allows-in` value, all the received prefixes should be stored in Adj-Rib-In, and then `allows-in` should be applied and prefixes with an as-path length more than the configured value should be flushed out of the BGP table. For more information on soft-reconfiguration, see [protocols bgp <asn> neighbor <id> address-family ipv4-unicast soft-reconfiguration inbound \(page 188\)](#).

protocols bgp <asn> neighbor <id> address-family ipv4-unicast attribute-unchanged <action>

Allows the router to send updates to a neighbor with unchanged attributes.

Syntax:

```
set protocols bgp asn neighbor id address-family ipv4-unicast attribute-unchanged [ as-path | med | next-hop ]
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast attribute-unchanged [ as-path | med | next-hop ]
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast attribute-unchanged
```

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

as-path

Propagates the route update with unchanged AS_PATH attribute.

med

Propagates the route update with unchanged Multi Exit Discriminator (MED).

next-hop

Propagates the route update with unchanged next hop.

Configuration mode

```
protocols {  
  bgp asn {  
    neighbor id {  
      address-family {  
        ipv4-unicast {  
          attribute-unchanged {
```



```
as-path
med
next-hop
}
}
}
}
}
```

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 ipv4-unicast capability graceful-restart

Enables or disables the graceful restart capability for the specified BGP neighbor.

Syntax:

```
set protocols bgp asn neighbor id address-family ipv4-unicast capability graceful-restart [ disable ]
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast graceful-restart [ disable ]
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast graceful-restart
```

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

disable

Disables the graceful restart capability.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          graceful-restart {
            disable
          }
        }
      }
    }
  }
}
```

Use the `set` form of this command to enable the graceful restart capability for the specified BGP neighbor.

Use the `delete` form of this command to disable the graceful restart capability for the specified BGP neighbor.

Use the `show` form of this command to view the BGP graceful restart configuration settings.



protocols bgp <asn> neighbor <id> address-family ipv4-unicast capability orf prefix-list <state>

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 ipv4-unicast capability orf prefix-list [ receive | send ]
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast capability orf prefix-list [ receive | send ]
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast capability orf prefix-list
```

By default, the session is brought up with minimal capability on both sides.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

prefix-list

Advertise prefix-list ORF capability to this neighbor.

receive

Advertise the ability to receive the ORF from this neighbor.

send

Advertise the capability to send the ORF to this neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          capability {
            orf {
              prefix-list {
                receive
                send
              }
            }
          }
        }
      }
    }
  }
}
```

In general, BGP-4 requires that if a BGP speaker receives an OPEN message with an unrecognized optional parameter, 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> address-family ipv4-unicast default-originate route-map <map-name>

Enables or disables forwarding of the default route to a BGP neighbor.

Syntax:

```
set protocols bgp asn neighbor id address-family ipv4-unicast default-originate route-map map-name
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast default-originate route-map map-name
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast default-originate
```

The default route is not distributed.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

map-name

Specifies a configured route map to be used when advertising the default route.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          default-originate {
            route-map map-name
          }
        }
      }
    }
  }
}
```

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>` ([page 222](#)).

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.

Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.



The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

protocols bgp <asn> neighbor <id> address-family ipv4-unicast disable-send-community <state>

Disables sending of community attributes to the specified neighbor.

Syntax:

```
set protocols bgp asn neighbor id address-family ipv4-unicast disable-send-community [ extended | standard ]
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast disable-send-community [ extended | standard ]
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast disable-send-community
```

Communities are sent by default.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

extended

Disable sending extended community attributes.

standard

Disable sending standard community attributes.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          disable-send-community {
            extended
            standard
          }
        }
      }
    }
  }
}
```

Use the `set` form of this command to disable sending community attributes between BGP neighbors.

Use the `delete` form of this command to restore sending community attributes between BGP neighbors.

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



protocols bgp <asn> neighbor <id> address-family ipv4-unicast distribute-list export <acl-number>

Applies an access list to filter outbound routing updates to this neighbor.

Syntax:

```
set protocols bgp asn neighbor id address-family ipv4-unicast distribute-list export acl-number
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast distribute-list export acl-number
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast distribute-list
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

acl-number

The number of a standard or extended access list. The range for a standard access list is 1 through 99. The range for an extended access list is 100 through 199.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          distribute-list {
            export acl-number
          }
        }
      }
    }
  }
}
```

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.

To create access lists, use the `policy route access-list rule` commands described in AT&T Vyatta Network Operating System Routing Policies Configuration Guide

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.

In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.



protocols bgp <asn> neighbor <id> address-family ipv4-unicast distribute-list import <acl-number>

Applies an access list to filter inbound routing updates from this neighbor.

Syntax:

```
set protocols bgp asn neighbor id address-family ipv4-unicast distribute-list import acl-number
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast distribute-list acl-number
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast distribute-list
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

acl-number

The number of a standard or extended access list. The range for a standard access list is 1 through 99. The range for an extended access list is 100 through 199.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          distribute-list {
            import acl-number
          }
        }
      }
    }
  }
}
```

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.

To create access lists, use the `policy route access-list rule` commands described in AT&T Vyatta Network Operating System Routing Policies Configuration Guide

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 can be applied to the specified direction.

In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.



protocols bgp <asn> neighbor <id> address-family ipv4-unicast 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 address-family ipv4-unicast filter-list export as-path-list-name
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast filter-list export as-path-list-name
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast filter-list
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

as-path-list-name

Name of an AS path list.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          filter-list {
            export as-path-list-name
          }
        }
      }
    }
  }
}
```

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 ipv4-unicast 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 address-family ipv4-unicast filter-list import as-path-list-name
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast filter-list import as-path-list-name
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast filter-list import
```

asn



The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

as-path-list-name

Name of an AS path list.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          filter-list {
            import as-path-list-name
          }
        }
      }
    }
  }
}
```

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 ipv4-unicast maximum-prefix <max-number>

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 ipv4-unicast maximum-prefix max-number [ threshold threshold-value [ warning-only ] warning-only ]
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast maximum-prefix [ threshold threshold-value [ warning-only ] warning-only ]
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast maximum-prefix
```

No maximum to the number of prefixes accepted.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

max-number

The maximum number of prefixes to accept from this neighbor before the neighbor is taken down.

threshold-value

The percentage of the maximum value of prefixes at which to start warning.

warning-only

Does not reset the connection once the maximum prefixes have been received. It only continues to warn.

Configuration mode

```
protocols {
```



```
bgp asn {
  peer-group name {
    address-family {
      ipv4-unicast {
        maximum-prefix max-number {
          threshold threshold-value {
            warning-only
          }
          warning-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.

Note: You can either configure `maximum-prefix max-number [warning-only]` or `maximum-prefix max-number threshold threshold-value [warning-only]`. The two values do not coexist in a configuration.

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.

Note:

When you configure this command and the number of prefixes received in an update message exceeds configured `maximum-prefix` limit, the session gets reset.

When soft-reconfiguration is enabled and the number of prefixes received exceeds `maximum-prefix` value, there will not be any corresponding unprocessed entries stored in Adj-RIB-In because the peer-session resets when the number of prefixes received exceeds `maximum-prefix` value. For more information on soft-reconfiguration, see [protocols bgp <asn> neighbor <id> address-family ipv4-unicast soft-reconfiguration inbound \(page 188\)](#).

protocols bgp <asn> neighbor <id> address-family ipv4-unicast nexthop-self

Sets the local router as the next hop for this neighbor.

Syntax:

```
set protocols bgp asn neighbor id address-family ipv4-unicast nexthop-self
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast nexthop-self
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast
```

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
```



```
bgp asn {
  neighbor id {
    address-family {
      ipv4-unicast {
        nexthop-self
      }
    }
  }
}
```

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> address-family ipv4-unicast peer-group <group-name>

Assigns a neighbor as a member of the specified peer group.

Syntax:

```
set protocols bgp asn neighbor id address-family ipv4-unicast peer-group group-name
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast peer-group group-name
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast peer-group
```

By default, peer group members inherit all configured peer group settings.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

group-name

A string identifying a BGP peer group.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          peer-group group-name
        }
      }
    }
  }
}
```

Refer to [Peer groups \(page 29\)](#) for information about peer groups.

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> address-family ipv4-unicast prefix-list export <list-name>

Applies a prefix list to filter updates to this neighbor.

Syntax:

```
set protocols bgp asn neighbor id address-family ipv4-unicast prefix-list export list-name
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast prefix-list export list-name
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast prefix-list
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

list-name

Name of a configured prefix list.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          prefix-list {
            export list-name
          }
        }
      }
    }
  }
}
```

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> address-family ipv4-unicast prefix-list import <list-name>

Applies a prefix list to filter updates from this neighbor.

Syntax:

```
set protocols bgp asn neighbor id address-family ipv4-unicast prefix-list import list-name
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast prefix-list import list-name
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast prefix-list
```


**asn**

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

list-name

Name of a configured prefix list.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          prefix-list {
            import list-name
          }
        }
      }
    }
  }
}
```

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> address-family ipv4-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 ipv4-unicast remove-private-as
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast remove-private-as
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast
```

Private AS numbers are included in outgoing updates.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          remove-private-as
        }
      }
    }
  }
}
```



```
}  
}  
}
```

Use the `set` form of this command to direct the router to exclude private AS numbers from updates to eBGP peers. When this feature is enabled, the router omits private AS numbers from the `AS_PATH` attribute. The private AS numbers range from 64512 through 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 ipv4-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 ipv4-unicast route-map export map-name
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast route-map export map-name
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast route-map export
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

map-name

Name of a configured route map.

Configuration mode

```
protocols {  
  bgp asn {  
    neighbor id {  
      address-family {  
        ipv4-unicast {  
          route-map {  
            export map-name  
          }  
        }  
      }  
    }  
  }  
}
```

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.



Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

protocols bgp <asn> neighbor <id> address-family ipv4-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 ipv4-unicast route-map import map-name
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast route-map import map-name
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast route-map import map-name
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

map-name

Name of a configured route map.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          route-map {
            import map-name
          }
        }
      }
    }
  }
}
```

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.

Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.



The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

protocols bgp <asn> neighbor <id> address-family ipv4-unicast soft-reconfiguration inbound

Directs the router to store received routing updates.

Syntax:

```
set protocols bgp asn neighbor id address-family ipv4-unicast soft-reconfiguration inbound
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast soft-reconfiguration inbound
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast soft-reconfiguration
```

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          soft-reconfiguration {
            inbound
          }
        }
      }
    }
  }
}
```

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. This soft reconfiguration command ensures maximum utilization of memory because you request for a route-refresh from a peer and update the routing table only when there is a change in the network.

The `soft-reconfiguration` command helps conserve memory utilization because you request for a route-refresh from a peer and update the routing table only when there is a change in the network.

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 ipv4-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 ipv4-unicast unsuppress-map map-name
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast unsuppress-map map-name
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast unsuppress-map
```

Routes suppressed by address aggregation are not advertised.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

map-name

The name of a configured route map.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          unsuppress-map map-name
        }
      }
    }
  }
}
```

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> address-family ipv4-unicast weight <weight>

Defines a default weight for routes from this neighbor.

Syntax:

```
set protocols bgp asn neighbor id address-family ipv4-unicast weight weight
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast weight weight
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast weight
```



Routes learned from a BGP neighbor have a weight of 0. Routes sourced by the local router have a weight of 32768.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

weight

The weight to be assigned to routes from this neighbor. The numbers range from 0 through 65535.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          weight weight
        }
      }
    }
  }
}
```

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 routes weighting to the default. Existing route weights are not changed.

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

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

Enables the IPv6 unicast address-family for a neighbor and defines parameters for controlling IPv6 unicast routes when subsequent configuration parameters and options are present.

Syntax:

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

Syntax:

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

Syntax:

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

asn

The number of the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

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



```
}
```

Use the `set` form of this command to enable the IPv6 unicast address-family for a neighbor and to define parameters for controlling IPv6 unicast routes when subsequent configuration parameters and options are present.

Use the `delete` form of this command to disable the IPv6 unicast address-family for a neighbor and to remove the configured parameters for controlling IPv6 unicast routes.

Use the `show` form of this command to view the configuration settings for the neighbor IPv6 unicast address-family.

protocols bgp <asn> neighbor <id> address-family ipv6-unicast allowas-in number <number>

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 number ]
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv6-unicast allowas-in
```

Syntax:

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

Readvertisement is disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

number number

Specifies the number of times to accept an AS path containing the readvertisement of the ASN of the local router. The numbers range from 1 through 10. The default is 3.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          allowas-in {
            number number
          }
        }
      }
    }
  }
}
```

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 AS number of the router. 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 AS path of the router.

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

**Note:**

When you configure this command and do a soft reset in the peer, the BGP table should flush out the prefixes that have an as-path length that is more than the configured `allows-in` value.

When soft-reconfiguration is enabled irrespective of the configured `allows-in` value, all the received prefixes should be stored in Adj-Rib-In, and then `allows-in` should be applied and prefixes with an as-path length more than the configured value should be flushed out of the BGP table. For more information on soft-reconfiguration, see [protocols bgp <asn> neighbor <id> address-family ipv6-unicast soft-reconfiguration inbound \(page 207\)](#).

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 ]
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv6-unicast attribute-unchanged [ as-path | med | next-hop ]
```

Syntax:

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

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

as-path

Propagates the route update with unchanged AS_PATH attribute.

med

Propagates the route update with unchanged Multi Exit Discriminator (MED).

next-hop

Propagates the route update with unchanged next hop.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          attribute-unchanged {
            as-path
            med
            next-hop
          }
        }
      }
    }
  }
}
```

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, 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 graceful-restart

Enables or disables the graceful restart capability for the specified BGP neighbor.

Syntax:

```
set protocols bgp asn neighbor id address-family ipv6-unicast capability graceful-restart [ disable ]
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv6-unicast graceful-restart [ disable ]
```

Syntax:

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

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

disable

Disables the graceful restart capability.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          graceful-restart {
            disable
          }
        }
      }
    }
  }
}
```

Use the `set` form of this command to enable the graceful restart capability for the specified BGP neighbor.

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

Use the `show` form of this command to view the BGP graceful restart 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 ] ]
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv6-unicast capability orf
```

Syntax:

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

By default, the session is brought up with minimal capability on both sides.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

prefix-list

Advertise prefix-list ORF capability to this neighbor.

receive

Advertise the ability to receive the ORF from this neighbor.

send

Advertise the capability to send the ORF to this neighbor.

Configuration mode

```
protocols {  
  bgp asn {  
    neighbor id {  
      address-family {  
        ipv6-unicast {  
          capability {  
            orf {  
              prefix-list {  
                receive  
                send  
              }  
            }  
          }  
        }  
      }  
    }  
  }  
}
```

This command specifies use of BGP 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 ]
```

**Syntax:**

```
delete protocols bgp asn neighbor id address-family ipv6-unicast default-originate [ route-map map-name ]
```

Syntax:

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

The default route is not distributed.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

map-name

Specifies a configured route map to be used when advertising the default route.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          default-originate {
            route-map map-name
          }
        }
      }
    }
  }
}
```

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>` ([page 222](#)).

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 <state>

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 ]
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv6-unicast disable-send-community
```

Syntax:



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

Communities are sent by default.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

extended

Disable sending extended community attributes.

standard

Disable sending standard community attributes.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          disable-send-community {
            extended
            standard
          }
        }
      }
    }
  }
}
```

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 the specified neighbor.

Syntax:

```
set protocols bgp asn neighbor id address-family ipv6-unicast distribute-list export access-list6-name
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv6-unicast distribute-list export access-list6-name
```

Syntax:

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

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

access-list6-name

The name of the IPv6 access-list.



Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          distribute-list {
            export access-list6-name
          }
        }
      }
    }
  }
}
```

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
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv6-unicast distribute-list import access-list6-name
```

Syntax:

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

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

access-list6-name

The name of the IPv6 access-list.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          distribute-list {
            import access-list6-name
          }
        }
      }
    }
  }
}
```



```
    }  
  }  
}
```

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
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv6-unicast filter-list export as-path-list-name
```

Syntax:

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

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

as-path-list-name

Name of an AS path list.

Configuration mode

```
protocols {  
  bgp asn {  
    neighbor id {  
      address-family {  
        ipv6-unicast {  
          filter-list {  
            export as-path-list-name  
          }  
        }  
      }  
    }  
  }  
}
```

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 access 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
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv6-unicast filter-list import as-path-list-name
```

Syntax:

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

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

as-path-list-name

Name of an AS path list.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          filter-list {
            import as-path-list-name
          }
        }
      }
    }
  }
}
```

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-number>

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-number [ threshold threshold-value [ warning-only ] warning-only ]
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv6-unicast maximum-prefix max-number [ threshold threshold-value [ warning-only ] warning-only ]
```

Syntax:



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

No maximum to the number of prefixes accepted.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

max-number

The maximum number of prefixes to accept from this neighbor before the neighbor is taken down.

threshold-value

The percentage of the maximum value of prefixes at which to start warning.

warning-only

Does not reset the connection once the maximum prefixes have been received. It only continues to warn.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv6-unicast {
          maximum-prefix max-number {
            threshold threshold-value {
              warning-only
            }
            warning-only
          }
        }
      }
    }
  }
}
```

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.

Note: You can either configure `maximum-prefix max-number [warning-only]` or `maximum-prefix max-number threshold threshold-value [warning-only]`. The two values do not coexist in a configuration.

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.

Note:

When you configure this command and the number of prefixes received in an update message exceeds configured `maximum-prefix` limit, the session gets reset.

When soft-reconfiguration is enabled and the number of prefixes received exceeds `maximum-prefix` value, there will not be any corresponding unprocessed entries stored in Adj-RIB-In because the peer-session resets when the number of prefixes received exceeds `maximum-prefix` value. For more information on soft-reconfiguration, see [protocols bgp <asn> neighbor <id> address-family ipv6-unicast soft-reconfiguration inbound \(page 207\)](#).



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
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv6-unicast nexthop-self
```

Syntax:

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

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          nexthop-self
        }
      }
    }
  }
}
```

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
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv6-unicast peer-group group-name
```

Syntax:

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

asn



The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

group-name

Name of an IPv6 peer group for the peer to join.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          peer-group group-name
        }
      }
    }
  }
}
```

This command specifies the IPv6 peer group for the peer. Refer to [Peer groups \(page 29\)](#) for information about peer groups.

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
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv6-unicast prefix-list export prefix-list6-name
```

Syntax:

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

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

prefix-list6-name

Name of a configured IPv6 prefix list.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          prefix-list {
            export prefix-list6-name
          }
        }
      }
    }
  }
}
```



```
    }  
  }  
}
```

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
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv6-unicast prefix-list import prefix-list6-name
```

Syntax:

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

asn

The number for the AS in which this router resides. The numbers range from 1 through 4294967294. The subrange, 64512 through 65534, is reserved for private AS.

id

The IPv4 or IPv6 address of a BGP neighbor.

prefix-list6-name

Name of a configured IPv6 prefix list.

Configuration mode

```
protocols {  
  bgp asn {  
    neighbor id {  
      address-family {  
        ipv6-unicast {  
          prefix-list {  
            import prefix-list6-name  
          }  
        }  
      }  
    }  
  }  
}
```

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
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv6-unicast remove-private-as
```

Syntax:

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

Private AS numbers are included in outgoing updates.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          remove-private-as
        }
      }
    }
  }
}
```

This command applies to IPv6 unicast routes only.

Use the `set` form of this command to direct the router to exclude private 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 through 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
```

**Syntax:**

```
delete protocols bgp asn neighbor id address-family ipv6-unicast route-map export map-name
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv6-unicast route-map export map-name
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

map-name

Name of a configured route map.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          route-map {
            export map-name
          }
        }
      }
    }
  }
}
```

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
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv6-unicast route-map import map-name
```

Syntax:

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

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

map-name

Name of a configured route map.

Configuration mode

```
protocols {
```



```
bgp asn {
  neighbor id {
    address-family {
      ipv6-unicast {
        route-map {
          import map-name
        }
      }
    }
  }
}
```

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 route-reflector-client

Configures the specified neighbor as a route reflector client.

Syntax:

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

Syntax:

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

Syntax:

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

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

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

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

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

Use the `show` form of this command to view the route reflector client configuration for the specified neighbor.



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
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv6-unicast soft-reconfiguration inbound
```

Syntax:

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

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          soft-reconfiguration {
            inbound
          }
        }
      }
    }
  }
}
```

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.

The `soft-reconfiguration` command helps conserve memory utilization because you request for a route-refresh from a peer and update the routing table only when there is a change in the network.

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
```

Syntax:



```
delete protocols bgp asn neighbor id address-family ipv6-unicast unsuppress-map map-name
```

Syntax:

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

Routes suppressed by address aggregation are not advertised.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

map-name

The name of a configured route map.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          unsuppress-map map-name
        }
      }
    }
  }
}
```

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> address-family ipv6-unicast weight <weight>

Defines a default weight for routes from this neighbor.

Syntax:

```
set protocols bgp asn neighbor id address-family ipv6-unicast weight weight
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv6-unicast weight weight
```

Syntax:

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

Routes learned from a BGP neighbor have a weight of 0. Routes sourced by the local router have a weight of 32768.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

weight

The weight to be assigned to routes from this neighbor. The numbers range from 0 through 65535.

Configuration mode



```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv6-unicast {
          weight weight
        }
      }
    }
  }
}
```

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 routes weighting to the default. Existing route weights are not changed.

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

protocols bgp <asn> neighbor <id> address-family vpnv4-unicast

Enables the VPNv4 unicast address family for a neighbor and defines parameters for controlling VPNv4 unicast routes when subsequent configuration parameters are present.

Syntax:

```
set protocols bgp asn neighbor id address-family vpnv4-unicast
```

Syntax:

```
delete protocols bgp asn neighbor id address-family vpnv4-unicast
```

Syntax:

```
show protocols bgp asn neighbor id address-family vpnv4-unicast
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        vpnv4-unicast
      }
    }
  }
}
```

Use the `set` form of this command to enable the VPNv4 unicast address family for a specified neighbor.

Use the `delete` form of this command to disable the VPNv4 unicast address family for a specified neighbor.

Use the `show` form of this command to view the VPNv4 unicast address family configuration.

protocols bgp <asn> neighbor <id> address-family vpnv4-unicast capability graceful-restart

Advertises the graceful-restart capability to the specified neighbor.

**Syntax:**

```
set protocols bgp asn neighbor id address-family vpnv4-unicast capability graceful-restart
```

Syntax:

```
delete protocols bgp asn neighbor id address-family vpnv4-unicast capability graceful-restart
```

Syntax:

```
show protocols bgp asn neighbor id address-family vpnv4-unicast capability graceful-restart
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        vpnv4-unicast {
          capability {
            graceful-restart
          }
        }
      }
    }
  }
}
```

Use the `set` form of this command to advertise the graceful-restart capability to the specified neighbor.

Use the `delete` form of this command to disable the graceful-restart capability for the specified neighbor.

Use the `show` form of this command to view the graceful restart configuration.

protocols bgp <asn> neighbor <id> address-family vpnv4-unicast distribute-list export <access-list-name>

Applies an access list to filter outbound routing updates to the specified neighbor.

Syntax:

```
set protocols bgp asn neighbor id address-family vpnv4-unicast distribute-list export access-list-name
```

Syntax:

```
delete protocols bgp asn neighbor id address-family vpnv4-unicast distribute-list export access-list-name
```

Syntax:

```
show protocols bgp asn neighbor id address-family vpnv4-unicast distribute-list export
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

access-list-name

The name of the access-list.



Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        vpnv4-unicast {
          distribute-list {
            export access-list-name
          }
        }
      }
    }
  }
}
```

Use the `set` form of this command to apply an access list to filter outbound routing VPNv4 messages 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 vpnv4-unicast distribute-list import <access-list-name>

Applies an access list to filter inbound routing updates from this neighbor.

Syntax:

```
set protocols bgp asn neighbor id address-family vpnv4-unicast distribute-list import access-list-name
```

Syntax:

```
delete protocols bgp asn neighbor id address-family vpnv4-unicast distribute-list import access-list-name
```

Syntax:

```
show protocols bgp asn neighbor id address-family vpnv4-unicast distribute-list import
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

access-list-name

The name of the access-list.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        vpnv4-unicast {
          distribute-list {
            import access-list-name
          }
        }
      }
    }
  }
}
```



```
    }  
  }  
}
```

Use the `set` form of this command to apply an access list to filter inbound routing VPNv4 messages 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 vpnv4-unicast 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 address-family vpnv4-unicast filter-list export as-path-list-name
```

Syntax:

```
delete protocols bgp asn neighbor id address-family vpnv4-unicast filter-list export as-path-list-name
```

Syntax:

```
show protocols bgp asn neighbor id address-family vpnv4-unicast filter-list
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

as-path-list-name

Name of an AS path list.

Configuration mode

```
protocols {  
  bgp asn {  
    neighbor id {  
      address-family {  
        vpnv4-unicast {  
          filter-list {  
            export as-path-list-name  
          }  
        }  
      }  
    }  
  }  
}
```

Use the `set` form of this command to apply an access list for filtering outbound VPNv4 messages.

Use the `delete` form of this command to disable outbound filtering of VPNv4 messages.

Use the `show` form of this command to view BGP neighbor filter list configuration settings.



protocols bgp <asn> neighbor <id> address-family vpnv4-unicast 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 address-family vpnv4-unicast filter-list import as-path-list-name
```

Syntax:

```
delete protocols bgp asn neighbor id address-family vpnv4-unicast filter-list import as-path-list-name
```

Syntax:

```
show protocols bgp asn neighbor id address-family vpnv4-unicast filter-list import
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

as-path-list-name

Name of an AS path list.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        vpnv4-unicast {
          filter-list {
            import as-path-list-name
          }
        }
      }
    }
  }
}
```

Use the set form of this command to apply an access list for filtering inbound VPNv4 messages.

Use the delete form of this command to disable inbound VPNv4 filtering.

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

protocols bgp <asn> neighbor <id> address-family vpnv4-unicast maximum-prefix <max-number>

Sets the maximum number of prefixes to accept from this VPNv4 neighbor before that neighbor is taken down.

Syntax:

```
set protocols bgp asn neighbor id address-family vpnv4-unicast maximum-prefix max-number [
threshold threshold-value [ warning-only ] warning-only ]
```

Syntax:

```
delete protocols bgp asn neighbor id address-family vpnv4-unicast maximum-prefix [ threshold
threshold-value [ warning-only ] warning-only ]
```

Syntax:



```
show protocols bgp asn neighbor id address-family vpnv4-unicast maximum-prefix
```

No maximum to the number of prefixes accepted.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

max-number

The maximum number of prefixes to accept from this neighbor before the neighbor is taken down.

threshold-value

The percentage of the maximum value of prefixes at which to start warning.

warning-only

Does not reset the connection once the maximum prefixes have been received. It only continues to warn.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        vpnv4-unicast {
          maximum-prefix max-number {
            threshold threshold-value {
              warning-only
            }
            warning-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.

Note: You can either configure `maximum-prefix max-number [warning-only]` or `maximum-prefix max-number threshold threshold-value [warning-only]`. The two values do not coexist in a configuration.

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.

Note:

When you configure this command and the number of prefixes received in an update message exceeds configured `maximum-prefix` limit, the session gets reset.

protocols bgp <asn> neighbor <id> address-family vpnv4-unicast prefix-list export <list-name>

Applies a prefix list to filter VPNv4 messages to this neighbor.

Syntax:

```
set protocols bgp asn neighbor id address-family vpnv4-unicast prefix-list export list-name
```

Syntax:



```
delete protocols bgp asn neighbor id address-family vpnv4-unicast prefix-list export list-name
```

Syntax:

```
show protocols bgp asn neighbor id address-family vpnv4-unicast prefix-list
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

list-name

Name of a configured prefix list.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        vpnv4-unicast {
          prefix-list {
            export list-name
          }
        }
      }
    }
  }
}
```

Use the `set` form of this command to restrict distribution of outbound BGP neighbor information by filtering VPNv4 messages 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> address-family vpnv4-unicast prefix-list import <list-name>

Applies a prefix list to filter VPNv4 messages from this neighbor.

Syntax:

```
set protocols bgp asn neighbor id address-family vpnv4-unicast prefix-list import list-name
```

Syntax:

```
delete protocols bgp asn neighbor id address-family vpnv4-unicast prefix-list import list-name
```

Syntax:

```
show protocols bgp asn neighbor id address-family vpnv4-unicast prefix-list
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

list-name

Name of a configured prefix list.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
```



```
address-family {
  vpnv4-unicast {
    prefix-list {
      import list-name
    }
  }
}
```

Use the `set` form of this command to restrict distribution of inbound BGP neighbor information by filtering VPNv4 messages 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> address-family vpnv4-unicast route-map export <map-name>

Applies a route map to filter VPNv4 messages to this neighbor.

Syntax:

```
set protocols bgp asn neighbor id address-family vpnv4-unicast route-map export map-name
```

Syntax:

```
delete protocols bgp asn neighbor id address-family vpnv4-unicast route-map export map-name
```

Syntax:

```
show protocols bgp asn neighbor id address-family vpnv4-unicast route-map export
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

map-name

Name of a configured route map.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        vpnv4-unicast {
          route-map {
            export map-name
          }
        }
      }
    }
  }
}
```

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.



Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

protocols bgp <asn> neighbor <id> address-family vpnv4-unicast route-map import <map-name>

Applies a route map to filter VPNv4 messages from this neighbor.

Syntax:

```
set protocols bgp asn neighbor id address-family vpnv4-unicast route-map import map-name
```

Syntax:

```
delete protocols bgp asn neighbor id address-family vpnv4-unicast route-map import map-name
```

Syntax:

```
show protocols bgp asn neighbor id address-family vpnv4-unicast route-map import map-name
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

map-name

Name of a configured route map.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        vpnv4-unicast {
          route-map {
            import map-name
          }
        }
      }
    }
  }
}
```

Use the `set` form of this command to restrict VPNv4 messages 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.

Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.



The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

protocols bgp <asn> neighbor <id> address-family vpnv4-unicast route-reflector-client

Configures the specified neighbor as a VPNv4 route reflector client.

Syntax:

```
set protocols bgp asn neighbor id address-family vpnv4-unicast route-reflector-client
```

Syntax:

```
delete protocols bgp asn neighbor id address-family vpnv4-unicast route-reflector-client
```

Syntax:

```
show protocols bgp asn neighbor id address-family vpnv4-unicast route-reflector-client
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        vpnv4-unicast {
          route-reflector-client
        }
      }
    }
  }
}
```

Use the `set` form of this command to configure the specified neighbor as a VPNv4 route reflector client.

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

Use the `show` form of this command to view the route reflector client configuration for the specified neighbor.

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
```

Syntax:

```
delete protocols bgp asn neighbor id advertisement-interval [ seconds ]
```

**Syntax:**

```
show protocols bgp asn neighbor id advertisement-interval
```

The default advertisement interval is 30 seconds for eBGP peers and 5 seconds for iBGP peers.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

seconds

The minimum interval, in seconds, between BGP routing updates to this neighbor. The numbers range from 0 through 600. The default is 30 for eBGP peers and 5 for iBGP peers.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      advertisement-interval seconds
    }
  }
}
```

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*> 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
```

Syntax:

```
delete protocols bgp asn neighbor id capability dynamic
```

Syntax:

```
show protocols bgp asn neighbor id
```

By default, the session is brought up with minimal capability on both sides.

asn

Mandatory. The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

Mandatory. The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      capability {
        dynamic
      }
    }
  }
}
```



```
}
```

In general, BGP-4 requires that if a BGP speaker receives an OPEN message with an unrecognized optional parameter, 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> description <description>

Provides a brief description for a BGP neighbor.

Syntax:

```
set protocols bgp asn neighbor id description description
```

Syntax:

```
delete protocols bgp asn neighbor id description [ description ]
```

Syntax:

```
show protocols bgp asn neighbor id description
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

description

A description of up to 80 characters for the neighbor. If spaces are included, the description must be enclosed in quotes.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      description description
    }
  }
}
```

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
```

Syntax:

```
delete protocols bgp asn neighbor id disable-capability-negotiation
```

Syntax:

```
show protocols bgp asn neighbor id
```

Capability negotiation is performed.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      disable-capability-negotiation
    }
  }
}
```

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> ebgp-multihop <ttl>

Allows the configuration of eBGP neighbors not on directly connected networks.

Syntax:

```
set protocols bgp asn neighbor id ebgp-multihop ttl
```

Syntax:

```
delete protocols bgp asn neighbor id ebgp-multihop [ ttl ]
```

Syntax:

```
show protocols bgp asn neighbor id
```

Only directly connected neighbors are allowed.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

ttl

The time-to-live, or maximum number of hops allowed. The numbers range from 1 through 255.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      ebgp-multihop ttl
    }
  }
}
```



```
}
```

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 **t1-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 **t1-security hops**.

Note that **t1-security hops** can not be configured at the same time as **ebgp-multihop**.

Packets are 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> enforce-multihop

Enforces multihop for eBGP neighbors.

Syntax:

```
set protocols bgp asn neighbor id enforce-multihop
```

Syntax:

```
delete protocols bgp asn neighbor id enforce-multihop
```

Syntax:

```
show protocols bgp asn neighbor id enforce-multihop
```

Multihop for eBGP neighbors is not configured.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      enforce-multihop
    }
  }
}
```

Use the **set** form of this command to enforce multihop for eBGP neighbors residing on networks that are not directly connected.

Use the **delete** form of this command to not enforce multihop for eBGP neighbors.

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

protocols bgp <asn> neighbor <id> local-as <asn>

Defines a local AS number for eBGP peerings.

Syntax:



```
set protocols bgp asn neighbor id local-as asn
```

Syntax:

```
delete protocols bgp asn neighbor id local-as asn
```

Syntax:

```
show protocols bgp asn neighbor id local-as
```

asn

Mandatory if an ASN other than the parent ASN is to be sent to the neighbor. The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      local-as asn
    }
  }
}
```

Use the `set` form of this command to set a local AS 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 AS number can only be applied to a true eBGP peer; it cannot be applied to peers in different subautonomous systems within a confederation.

Use the `delete` form of this command to remove a local AS number.

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

protocols bgp *asn* neighbor <*id*> log

Enables BGP neighbor logs.

Syntax:

```
set protocols bgp asn neighbor id log { all | events | filters | fsm | keepalive | nsm | updates [ all | in | out ] }
```

Syntax:

```
delete protocols bgp asn neighbor id log { all | events | filters | fsm | keepalive | nsm | updates [ all | in | out ] }
```

Syntax:

```
show protocols bgp asn neighbor id log { all | events | filters | fsm | keepalive | nsm | updates [ all | in | out ] }
```

asn

The number of the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

all

Enables all BGP neighbor logs.

events

Enables only BGP neighbor event logs.

filters

Enables only BGP neighbor filter logs.

fsm



- enable** Enables only BGP neighbor finite-state machine logs.
- keepalive** Enables only BGP neighbor keepalive logs.
- nsm** Enables only BGP neighbor network services module (NSM) logs.
- all** Enables all BGP neighbor logs.
- in** Enables BGP neighbor update-in logs.
- out** Enables BGP neighbor update-out logs.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      log {
        all
        events
        filters
        fsm
        keepalive
        nsm
        update {
          all
          in
          out
        }
      }
    }
  }
}
```

Use the `set` form of this command to enable BGP neighbor logs.

Use the `delete` form of this command to remove BGP neighbor logs.

Use the `show` form of this command to view BGP neighbor logs.

protocols bgp <asn> neighbor <id> med-out igp

Sets a multi-exit discriminator (MED) value that is based on changes in the Interior Gateway Protocol (IGP) metric.

Syntax:

```
set protocols bgp asn neighbor id med-out igp
```

Syntax:

```
delete protocols bgp asn neighbor id med-out igp
```

asn

Mandatory if an ASN other than the parent ASN is to be sent to the neighbor. The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
```




```
        med-out igp
    }
}
```

Use the `set` form of this command to send a MED update that is based on the changes in IGP metric. Use the `delete` form of this command to disable MED update according to changes in the IGP metric.

protocols bgp <asn> neighbor <id> med-out igp delay-updates

Delays the update of the multi-exit Discriminator (MED) to a peer.

Syntax:

```
set protocols bgp asn neighbor id med-out igp delay-updates
```

Syntax:

```
delete protocols bgp asn neighbor id med-out igp delay-updates
```

asn

Mandatory if an ASN other than the parent ASN is to be sent to the neighbor. The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      med-out igp delay-updates
    }
  }
}
```

Use the `set` form of this command to delay the MED updates to a peer. Enter the command only if the device is already configured to assign the MED based on the changes in the IGP metric. The time delay is based on the default timer that is configured in [protocols bgp <asn> parameters med-out-delay <minutes>](#) ([page 124](#)).

Note: The device delays the MED updates when the IGP metric changes and the new metric is greater than the old IGP metric.

Use the `delete` form of this command to update the MED value as soon as the IGP metric changes.

protocols bgp <asn> neighbor <id> med-out minimum-igp

Sets the value of the multi-exit discriminator (MED) based on the smaller IGP metric.

Syntax:

```
set protocols bgp asn neighbor id med-out minimum-igp
```

Syntax:

```
delete protocols bgp asn neighbor id med-out minimum-igp
```

asn

Mandatory if an ASN other than the parent ASN is to be sent to the neighbor. The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.



Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      med-out minimum-igp
    }
  }
}
```

Use the `set` form of this command to set the MED value that is based on the minimum value of IGP metric. The minimum IGP metric is the smallest IGP metric received by BGP from the RIB.

Note: The `protocols bgp <asn> neighbor <id> med-out igp` and `protocols bgp <asn> neighbor <id> med-out minimum-igp` commands are mutually exclusive.

Use the `delete` form of this command to disable MED updates that is based on the IGP metric.

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
```

Syntax:

```
delete protocols bgp asn neighbor id override-capability
```

Syntax:

```
show protocols bgp asn neighbor id
```

A session cannot be established if the neighbor does not support capability negotiation.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          override-capability
        }
      }
    }
  }
}
```

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
```

Syntax:

```
delete protocols bgp asn neighbor id passive
```

Syntax:

```
show protocols bgp asn neighbor id
```

The router both accepts inbound connections and initiates outbound connections.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      passive
    }
  }
}
```

Use the `set` form of this command to configure the local router such that it 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 <password>

Defines a BGP MD5 password.

Syntax:

```
set protocols bgp asn neighbor id password password
```

Syntax:

```
delete protocols bgp asn neighbor id password
```

Syntax:

```
show protocols bgp asn neighbor id password
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

password

A password to be used to generate an MD5 digest.

Configuration mode

```
protocols {
```



```
    bgp asn {
      neighbor id {
        password password
      }
    }
  }
```

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> port <port-number>

Specifies the source port of a BGP neighbor.

Syntax:

```
set protocols bgp asn neighbor id port port-number
```

Syntax:

```
delete protocols bgp asn neighbor id port port-number
```

Syntax:

```
show protocols bgp asn neighbor id port
```

A randomly-selected, unused port is chosen as the source port of the BGP neighbor.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

port-number

Mandatory. Specifies the source port of a BGP neighbor. The numbers range from 1 through 65535. The default is a randomly-selected, unused port.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      port port-number
    }
  }
}
```

Use the `set` form of this command to specify the source port of a BGP neighbor.

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> remote-as <asn>

Specifies the AS number of the neighbor.

Syntax:

```
set protocols bgp asn neighbor id remote-as asn
```

**Syntax:**

```
delete protocols bgp asn neighbor id remote-as asn
```

Syntax:

```
show protocols bgp asn neighbor id remote-as
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      remote-as asn
    }
  }
}
```

Use the `set` form of this command to specify the 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> shutdown

Administratively shuts down a BGP neighbor.

Syntax:

```
set protocols bgp asn neighbor id shutdown
```

Syntax:

```
delete protocols bgp asn neighbor id shutdown
```

Syntax:

```
show protocols bgp asn neighbor id
```

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      shutdown
    }
  }
}
```

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 has been administratively shut down.

Use the `delete` form of this command to administratively re-enable a BGP neighbor.

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
```

Syntax:

```
delete protocols bgp asn neighbor id strict-capability-match
```

Syntax:

```
show protocols bgp asn neighbor id
```

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      strict-capability-match
    }
  }
}
```

Use the `set` form of this command to direct the router to 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 <action>

Sets BGP timers for this neighbor.

Syntax:

```
set protocols bgp asn neighbor id timers { connect | keepalive seconds | holdtime seconds }
```

Syntax:

```
delete protocols bgp asn neighbor id timers [ connect | keepalive | holdtime ]
```

Syntax:

```
show protocols bgp asn neighbor id timers
```

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.

**asn**

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

connect seconds

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 numbers range from 0 through 65535. The default is 120.

During the connect period, the remote BGP peer can establish a connection to the local system.

connect

The BGP connect time for this BGP neighbor. The numbers are 0 and 1 through 65535, where 0 disables the connect timer.

holdtime seconds

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 through 65535, where 0 disables the holdtime timer. The default is 180.

keepalive seconds

The frequency, in seconds, with which the local router sends keep-alive messages to this neighbor. The numbers range from 1 through 65535. The default is 60.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      timers {
        connect seconds
        keepalive seconds
        holdtime seconds
      }
    }
  }
}
```

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` (page 134) 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
```

Syntax:

```
delete protocols bgp asn neighbor id ttl-security hops hops
```

Syntax:

```
show protocols bgp asn neighbor id ttl-security hops
```



The default for the hop count is 1 for eBGP connections.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

hops

The maximum number of hops away from the local peer a peering session connection be accepted. The numbers range from 1 through 254.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      ttl-security {
        hops hops
      }
    }
  }
}
```

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 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> update-source <source>

Specifies the source IP address or interface for routing updates.

Syntax:

```
set protocols bgp asn neighbor id update-source source
```

Syntax:

```
delete protocols bgp asn neighbor id update-source source
```

Syntax:

```
show protocols bgp asn neighbor id
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

source



The IPv4 address of the router to receive routing updates from or the interface to receive routing updates from.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      update-source source
    }
  }
}
```

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> interface <interface-name> vrrp-failover vrrp-group <vrrp-group-id> med <med-value>

Dynamically changes the BGP best path taken for incoming traffic in accordance with the path taken by outgoing traffic, if MED is used as the control mechanism.

Syntax:

```
set protocols bgp asn neighbor id interface interface-name vrrp-failover vrrp-group vrrp-group-id med med-value
```

Syntax:

```
delete protocols bgp asn neighbor id interface interface-name vrrp-failover vrrp-group vrrp-group-id med med-value
```

Syntax:

```
show protocols bgp asn neighbor id interface interface-name vrrp-failover vrrp-group vrrp-group-id
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

interface-name

The name of a BGP neighbor interface.

vrrp-group-id

The ID of a VRRP group instance.

med-value

The MED to use on a backup path.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      interface interface-name {
        vrrp-failover
        vrrp-group vrrp-group-id
        med med-value
      }
    }
  }
}
```



```
}
```

Use the `set` form of this command to dynamically change the BGP best path taken by incoming traffic in accordance with the path taken by outgoing traffic, if MED is used as the control mechanism.

Use the `delete` form of this command to remove the user-configured MED attribute in the BGP route update.

Use the `show` form of this command to view the configured MED value.

The update notification contains Multi-Exit Discriminator (MED) changes according to configuration to influence the change in BGP best paths for incoming traffic.

Master VRRP router will always send route updates with default BGP attribute values, not the one configured under "bgp neighbor vrrp-failover."

The MED that is used on the active path or the path that includes the VRRP elected master will advertise the default MED value.

The following configuration is required to enable the notification function on vRouters:

```
set interface dataplane downlink-interface vrrp vrrp-group vrrp-group-instance-id notify bgp
```

Note:

If you are upgrading the AT&T Vyatta vRouter to the 4.0 release from any previous release, ensure that unique VRRP group IDs are used across interfaces to avoid conflicting configurations among those interfaces. Furthermore, by employing only the following command, use unique VRRP group IDs for the VRRP groups that are tracked by BGP.

```
protocols bgp asn neighbor id interface interface-name vrrp-failover vrrp-group vrrp-group-id
```

. VRRP groups that are not tracked do not require unique IDs within the vRouter.

Note: Local route maps and policies, if any, override the MED.

protocols bgp <asn> neighbor <id> interface <interface-name> vrrp-failover vrrp-group <vrrp-group-id> prepend-as <as-path-string>

Dynamically changes the BGP best path taken by incoming traffic according to the path taken by outgoing traffic, if prepend-as string is used as the control mechanism.

Syntax:

```
set protocols bgp asn neighbor id interface interface-name vrrp-failover vrrp-group vrrp-group-id  
prepend-as as-path-string
```

Syntax:

```
delete protocols bgp asn neighbor id interface interface-name vrrp-failover vrrp-group vrrp-group-id  
prepend-as
```

Syntax:

```
show protocols bgp asn neighbor id interface interface-name vrrp-failover vrrp-group vrrp-group-id
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

interface-name

The name of a BGP neighbor interface.

vrrp-group-id

The ID of the VRRP group instance.

as-path-string



prepend-as path string to be used on backup path.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      interface interface-name {
        vrrp-failover
          vrrp-group vrrp-group-id {
            prepend-as as-path-string
          }
        }
      }
    }
  }
}
```

Use the `set` form of this command to direct the system to dynamically change the BGP best path taken by incoming traffic in accordance with the path taken by outgoing traffic, if `prepend-as` string is used as the control mechanism.

Use the `delete` form of this command to remove the user-configured `prepend-as` attribute in the BGP route update.

Use the `show` form of this command to view the BGP configured `prepend-as` value.

The update notification contains `prepend-as-path` changes as per configuration to influence the change in BGP best paths for incoming traffic.

The `as-path-string` used on the active path or the path including the VRRP elected master is `NONE`, that is, no AS-path prepended.

The following configuration is required to enable the notification function on vRouters:

```
set interface dataplane downlink-interface vrrp vrrp-group vrrp-group-id notify bgp
```

Note:

If you are upgrading the AT&T Vyatta vRouter to the 4.0 release from any previous release, ensure that unique VRRP group IDs are used across interfaces to avoid conflicting configurations among those interfaces. Furthermore, by employing only the following command, use unique VRRP group IDs for the VRRP groups that are tracked by BGP.

```
protocols bgp asn neighbor id interface interface-name vrrp-failover vrrp-group vrrp-group-id
```

. VRRP groups that are not tracked do not require unique IDs within the vRouter.

Note: Local route maps and policies, if any, override the `prepend-as`.

protocols bgp <asn> neighbor <id> interface <interface-name> vrrp-failover vrrp-group <group-id> route-map <route-map-name>

Sets route-map configuration in the AT&T Vyatta vRouter for the neighbor.

Syntax:

```
set protocols bgp asn neighbor id interface interface-name vrrp-failover vrrp-group group-id route-map route-map-name
```

Syntax:

```
delete protocols bgp asn neighbor id interface interface-name vrrp-failover vrrp-group group-id route-map route-map-name
```

**Syntax:****asn**

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

interface-name

The name of a BGP neighbor interface.

group-id

The ID of a VRRP group instance.

route-map-name

The route map name used for a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      interface interface-name {
        vrrp-failover
        vrrp-group group-id {
          route map route-map-name
        }
      }
    }
  }
}
```

Use the `set` form of this command to set route map configuration to influence the path.

Use the `delete` form of this command to remove the route map configuration.

Use the `show` form of this command to view the BGP configured route-map.

Note:

If you are upgrading the AT&T Vyatta vRouter to the 4.0 release from any previous release, ensure that unique VRRP group IDs are used across interfaces to avoid conflicting configurations among those interfaces. Furthermore, by employing only the following command, use unique VRRP group IDs for the VRRP groups that are tracked by BGP.

```
protocols bgp asn neighbor id interface interface-name vrrp-failover vrrp-group vrrp-group-id
```

. VRRP groups that are not tracked do not require unique IDs within the vRouter.

reset ip bgp external

Resets peering sessions for eBGP neighbors.

Syntax:

```
reset ip bgp external [ soft [ in [ prefix-filter ] | out ] ]
```

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

soft

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

Resets BGP external peer information via soft reconfig of inbound update

prefix-filter

Resets the outbound route filter (ORF) filter.

out



Resets BGP external peer information via soft reconfig of outbound update

Operational mode

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 re-established.

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 <state>

Resets IPv4 unicast peering sessions for eBGP neighbors.

Syntax:

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

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

soft

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

Re-read import policies and recalculate information in the BGP table based on import policy changes. This requires **soft-reconfiguration inbound** to be configured on the neighbor.

prefix-filter

Resets the outbound route filter (ORF) filter.

out

Re-read export policies and resend updates that have changed due to an export policy change.

Operational mode

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 re-established.

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 ipv6 unicast <state>

Resets IPv6 unicast peering sessions for eBGP neighbors.

Syntax:

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

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

soft

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

Re-read import policies and recalculate information in the BGP table based on import policy changes. This requires **soft-reconfiguration inbound** to be configured on the neighbor.

prefix-filter

Resets the outbound route filter (ORF) filter.

**out**

Re-read export policies and resend updates that have changed due to an export policy change.

Operational mode

Use this command to reset IPv6 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 re-established.

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 address
```

IPv4 unicast route information is shown for all BGP neighbors.

address

Displays BGP information for the specified IPv4 BGP neighbor. Can be in the following formats: *X::X::X::X* or *x.x.x.x*.

Operational mode

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
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to display advertised BGP IPv4 unicast routes 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
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

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 when soft-reconfiguration inbound is enabled for the peer.

Syntax:

```
show ip bgp ipv4 unicast neighbors id received-routes
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

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
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to display IPv4 unicast received and accepted routes from a BGP neighbor.

show ip bgp ipv6 unicast neighbors

Displays BGP neighbor information on TCP and BGP neighbor connections.

Syntax:

```
show ip bgp ipv6 unicast neighbors [ ipv6-address ]
```

Information is shown for all BGP neighbors.

ipv6-address

The IPv6 address of a BGP neighbor.

Operational mode

Use this command to display BGP neighbor information.

show ip bgp ipv6 unicast neighbors <ipv6> advertised-routes

Displays advertised routes for a BGP neighbor.

Syntax:

```
show ip bgp ipv6 unicast neighbors ipv6 advertised-routes
```

ipv6

The IPv6 address of a BGP neighbor.

Operational mode

Use this command to display advertised routes for a BGP neighbor.



show ip bgp ipv6 unicast neighbors <ipv6> received-routes

Displays routes received from a BGP neighbor when soft-reconfiguration inbound is enabled for the peer.

Syntax:

```
show ip bgp ipv6 unicast neighbors ipv6 received-routes
```

ipv6

The IPv6 address of a BGP neighbor.

Operational mode

Use this command to display routes (both accepted and rejected) received from a BGP neighbor.

show ip bgp ipv6 unicast neighbors <ipv6> routes

Displays all received and accepted routes from a BGP neighbor.

Syntax:

```
show ip bgp ipv6 unicast neighbors ipv6 routes
```

ipv6

The IPv6 address of a BGP neighbor.

Operational mode

Use this command to display received and accepted routes from a BGP neighbor.

show ip bgp neighbors

Displays BGP neighbor information.

Syntax:

```
show ip bgp neighbors [ address ]
```

Information for all neighbors is displayed.

address

Displays detailed IPv4 or IPv6 BGP information for a BGP neighbor.

Operational mode

Use this command to display BGP neighbor information.

Note: This command does not display peer's graceful-restart capability.

show ip bgp neighbors <id> advertised-routes

Displays advertised routes for a BGP neighbor.

Syntax:

```
show ip bgp neighbors id advertised-routes
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to display advertised routes 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
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

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 when soft-reconfiguration inbound is enabled for the peer.

Syntax:

```
show ip bgp neighbors id received-routes
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

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
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to display received and accepted routes from a BGP neighbor.

show ip bgp vpnv4 unicast

Displays VPNv4 unicast route information.

Syntax:

```
show ip bgp vpnv4 unicast
```

Operational mode

Use this command to display VPNv4 unicast route information for BGP neighbors.

The following example shows command output.

Network	Next Hop	Metric	LocPrf	Weight	Path
Route Distinguisher:	1:1				
*>i100.100.100.0/24	60.60.60.1	0	100	0	?
*>i110.110.110.0/24	60.60.60.1	0	100	0	?
*>i120.120.120.0/24	60.60.60.1	0	100	0	?
*>i130.130.130.0/24	60.60.60.1	0	100	0	?

show ip bgp vpnv4 unicast <prefix>

Displays VPNv4 unicast route information for the specified address or prefix.

**Syntax:**

```
show ip bgp vpnv4 unicast prefix
```

***prefix* [longer-prefixes]**

An IP address (x.x.x.x) or prefix (x.x.x.x/n).

Operational mode

Use this command to display VPNv4 unicast route information for the specified address or prefix. Use the optional **longer-prefixes** keyword to show more specific route information.

show ip bgp vpnv4 unicast dampening

Displays dampened BGP VPNv4 unicast routes or flap statistics.

Syntax:

```
show ip bgp vpnv4 unicast dampening { dampened-paths | flap-statistics }
```

Operational mode

Use this command to display dampened BGP VPNv4 routes or flap statistics.

show ip bgp vpnv4 unicast filter-list

Displays the BGP VPNv4 unicast routes that conform to the specified filter list.

Syntax:

```
show ip bgp vpnv4 unicast filter-list filter-list
```

filter-list

An access control list.

Operational mode

Use this command to display the routes that conform to the specified filter list.

show ip bgp vpnv4 unicast neighbors

Displays information about BGP VPNv4 unicast neighbor connections.

Syntax:

```
show ip bgp vpnv4 unicast neighbors
```

Operational mode

Use this command to display information about BGP VPNv4 neighbor connections.

show ip bgp vpnv4 unicast neighbors <id> advertised-routes

Displays advertised BGP VPNv4 unicast routes for a BGP neighbor.

Syntax:

```
show ip bgp vpnv4 unicast neighbors id advertised-routes
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to display advertised BGP vpnv4 unicast routes for a BGP neighbor.



show ip bgp vpnv4 unicast neighbors <id> received-routes

Displays the BGP VPNv4 unicast routes received from a BGP neighbor.

Syntax:

```
show ip bgp vpnv4 unicast neighbors id received-routes
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to display VPNv4 unicast routes (both accepted and rejected) received from a BGP neighbor.

show ip bgp vpnv4 unicast neighbors <id> routes

Displays BGP VPNv4 unicast received and accepted routes from a BGP neighbor.

Syntax:

```
show ip bgp vpnv4 unicast neighbors id routes
```

id

The IPv4 or IPv6 address of a BGP neighbor.

Operational mode

Use this command to display VPNv4 unicast received and accepted routes from a BGP neighbor.

show ip bgp vpnv4 unicast prefix-list

Displays the BGP VPNv4 unicast routes matching the specified prefix list.

Syntax:

```
show ip bgp vpnv4 unicast prefix-list
```

Syntax:

```
show ip bgp vpnv4 unicast prefix-list list
```

Operational mode

Use this command to display the VPNv4 routes matching the specified prefix list.

show ip bgp vpnv4 unicast regexp

Displays the BGP VPNv4 unicast routes that match the specified AS path regular expression.

Syntax:

```
show ip bgp vpnv4 unicast regexp expression
```

expression

A POSIX-style regular expression representing a set of AS paths.

Operational mode

Use this command to display the BGP VPNv4 unicast routes that match the specified AS path regular expression.

show ip bgp vpnv4 unicast summary

Displays a summary of BGP VPNv4 unicast neighbor status.

Syntax:



```
show ip bgp vpnv4 unicast summary
```

Operational mode

Use this command to display a summary of BGP VPNv4 unicast neighbor status.



Peer Group Commands

protocols bgp <asn> peer-group <name>

Defines a BGP peer group.

Syntax:

```
set protocols bgp asn peer-group name
```

Syntax:

```
delete protocols bgp asn peer-group name
```

Syntax:

```
show protocols bgp asn peer-group name
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

Multi-node. The peer group name.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
    }
  }
}
```

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 <name> address-family

Defines an address family for a BGP peer group.

Syntax:

```
set protocols bgp asn peer-group name address-family
```

Syntax:

```
delete protocols bgp asn peer-group name address-family
```

Syntax:



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

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

Multi-node. The peer group name.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
      }
    }
  }
}
```

Use the `set` form of this command to define an address family for a BGP peer group.

Use the `delete` form of this command to delete an address family from a BGP peer group.

Use the `show` form of this command to display an address family for a BGP peer group.

protocols bgp <asn> peer-group <name> address-family ipv4-unicast

Enables the IPv4 unicast address-family for a peer group and defines parameters for controlling IPv4 unicast routes when subsequent configuration parameters and options are present.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv4-unicast
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv4-unicast
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv4-unicast {
        }
      }
    }
  }
}
```

Use the `set` form of this command to enable the IPv4 unicast address-family for a peer group and to define parameters for controlling IPv4 unicast routes when subsequent configuration parameters and options are present.



Use the `delete` form of this command to disable the IPv4 unicast address-family for a peer group and to remove the configured parameters for controlling IPv4 unicast routes.

Use the `show` form of this command to view the configuration settings for the IPv4 unicast address-family of a peer group.

protocols bgp <asn> peer-group <name> address-family ipv4-unicast allows-in number <number>

Allows or disallows receiving BGP advertisements containing the AS path of the local router.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv4-unicast allows-in number number
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv4-unicast allows-in [ number number ]
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast
```

Readvertisement is disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

number *number*

The number of times to accept an AS path containing the readvertisement of the ASN of the local router. The numbers range from 1 through 10. The default is 3.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv4-unicast {
          allows-in {
            number number
          }
        }
      }
    }
  }
}
```

Use the `set` form of this command to allow the router to accept BGP AS paths advertising the AS number of the router.

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 AS path of the router.

Use the `show` form of this command to view BGP peer group configuration settings.

Note:

When you configure this command and do a soft reset in the peer, the BGP table should flush out the prefixes that have an as-path length that is more than the configured `allows-in` value.

When soft-reconfiguration is enabled irrespective of the configured `allows-in` value, all the received prefixes should be stored in Adj-Rib-In, and then `allows-in` should be applied and prefixes with an as-path length more than the configured value should be flushed out of the BGP table. For more information



on soft-reconfiguration, see [protocols bgp <asn> neighbor <id> address-family ipv4-unicast soft-reconfiguration inbound \(page 188\)](#).

protocols bgp <asn> peer-group <name> address-family ipv4-unicast attribute-unchanged <action>

Allows the router to send updates to a peer group with unchanged attributes.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv4-unicast attribute-unchanged { as-path | med | next-hop }
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv4-unicast attribute-unchanged [ as-path | med | next-hop ]
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast attribute-unchanged
```

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

as-path

Propagates the route update with unchanged AS_PATH attribute.

med

Propagates the route update with unchanged Multi Exit Discriminator (MED).

next-hop

Propagates the route update with unchanged next hop.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv4-unicast {
          attribute-unchanged {
            as-path
            med
            next-hop
          }
        }
      }
    }
  }
}
```

Use the `set` form of this command to direct the router to propagate routing updates without modifying the BGP AS_PATH, 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 <name> address-family ipv4-unicast capability graceful-restart

Enables or disables the graceful restart capability for the specified BGP peer group of an IPv4 unicast address family.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv4-unicast capability graceful-restart [ disable ]
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv4-unicast graceful-restart [ disable ]
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast graceful-restart
```

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

disable

Disables the graceful restart capability.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv4-unicast {
          capability
            graceful-restart {
              disable
            }
        }
      }
    }
  }
}
```

When a neighbor is a member of a peer group, graceful restart cannot be configured on the neighbor level. A graceful restart configured at the neighbor level or the peer-group level, will override the global-level graceful restart configuration.

Use the `set` form of this command to enable the graceful restart capability for the specified BGP peer group of an IPv4 unicast address family.

Use the `delete` form of this command to disable the graceful restart capability for the specified BGP peer group of an IPv4 unicast address family.

Use the `show` form of this command to display the graceful restart capability for the specified BGP peer group of an IPv4 unicast address family.



protocols bgp <asn> peer-group <name> address-family ipv4-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 name address-family ipv4-unicast capability orf { prefix-list [ receive | send ] }
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv4-unicast capability orf [ prefix-list [ receive | send ] ]
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast
```

By default, the session is brought up with minimal capability on both sides.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

prefix-list

Advertise prefix-list ORF capability to this peer group.

receive

Advertise the ability to receive the ORF from this peer group.

send

Advertise the capability to send the ORF to this peer group.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv4-unicast {
          capability {
            orf {
              prefix-list {
                receive
                send
              }
            }
          }
        }
      }
    }
  }
}
```

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 <name> address-family ipv4-unicast default-originate

Enables or disables forwarding of the default route to a BGP peer group.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv4-unicast default-originate [ route-map map-name ]
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv4-unicast default-originate [ route-map [ map-name ] ]
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast default-originate
```

The default route is not distributed.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

map-name

Specifies a configured route map to be used when advertising the default route.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv4-unicast {
          default-originate {
            route-map map-name
          }
        }
      }
    }
  }
}
```

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 <name> local-as <asn>` ([page 287](#)).

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.



Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

protocols bgp <asn> peer-group <name> address-family ipv4-unicast disable-send-community <action>

Disables sending of community attributes to the specified peer group.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv4-unicast disable-send-community { extended | standard }
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv4-unicast disable-send-community [ extended | standard ]
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast
```

Communities are sent by default.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

extended

Disable sending extended community attributes.

standard

Disable sending standard community attributes.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv4-unicast {
          disable-send-community {
            extended
            standard
          }
        }
      }
    }
  }
}
```

Use this command to enable or disable the sending of community attributes between BGP peer groups.

Use the `set` form of this command to disable the sending of community attributes between peer groups.

Use the `delete` form of this command to restore sending community attributes between peer groups.



Use the `show` form of this command to view BGP peer group configuration settings between peer groups.

protocols bgp <asn> peer-group <name> address-family ipv4-unicast distribute-list import <acl-number>

Applies an access list to filter inbound routing updates from this peer group.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv4-unicast distribute-list import acl-number
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv4-unicast distribute-list acl-number
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast distribute-list
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

acl-number

The number of a standard or extended access list. The number range is 1 through 99 for a standard access list, and 100 through 199 for an extended access list.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv4-unicast {
          distribute-list {
            import acl-number
          }
        }
      }
    }
  }
}
```

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.

Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.



Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

protocols bgp <asn> peer-group <name> address-family ipv4-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 *name* address-family ipv4-unicast filter-list export *as-path-list-name*

Syntax:

delete protocols bgp *asn* peer-group *name* address-family ipv4-unicast filter-list export [*as-path-list-name*]

Syntax:

show protocols bgp *asn* peer-group *name* address-family ipv4-unicast filter-list export

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

as-path-list-name

Name of an autonomous system path list.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv4-unicast {
          filter-list {
            export as-path-list-name
          }
        }
      }
    }
  }
}
```

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 <name> address-family ipv4-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 *name* address-family ipv4-unicast filter-list import *as-path-list-name*

**Syntax:**

```
delete protocols bgp asn peer-group name address-family ipv4-unicast filter-list import [ as-path-list-name ]
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast filter-list import
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

as-path-list-name

Name of an autonomous system path list.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv4-unicast {
          filter-list {
            import as-path-list-name
          }
        }
      }
    }
  }
}
```

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 <name> address-family ipv4-unicast maximum-prefix <max-number>

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 name address-family ipv4-unicast maximum-prefix max-number [ threshold threshold-value [ warning-only ] warning-only ]
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv4-unicast maximum-prefix [ max-number ] [ threshold threshold-value [ warning-only ] warning-only ]
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast maximum-prefix
```

No maximum to the number of prefixes accepted.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

max-number

The maximum number of prefixes to accept from this peer group before the peer group is taken down.

**threshold-value**

The percentage of the maximum value of prefixes at which to start warning.

warning-only

Does not reset the connection once the maximum prefixes have been received. It only continues to warn.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv4-unicast {
          maximum-prefix max-number {
            threshold threshold-value {
              warning-only
            }
          }
          warning-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.

Note: You can either configure `maximum-prefix max-number [warning-only]` or `maximum-prefix max-number threshold threshold-value [warning-only]`. The two values do not coexist in a configuration.

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.

Note:

When you configure this command and the number of prefixes received in an update message exceeds configured `maximum-prefix` limit, the session gets reset.

When soft-reconfiguration is enabled and the number of prefixes received exceeds `maximum-prefix` value, there will not be any corresponding unprocessed entries stored in Adj-RIB-In because the peer-session resets when the number of prefixes received exceeds `maximum-prefix` value. For more information on soft-reconfiguration, see `protocols bgp <asn> neighbor <id> address-family ipv4-unicast soft-reconfiguration inbound` ([page 188](#)).

protocols bgp <asn> peer-group <name> address-family ipv4-unicast nexthop-self

Sets the local router as the next hop for this peer group.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv4-unicast nexthop-self
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv4-unicast nexthop-self
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast
```




Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv4-unicast {
          nexthop-self
        }
      }
    }
  }
}
```

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 <name> address-family ipv4-unicast prefix-list export <list-name>

Applies a prefix list to filter updates to this peer group.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv4-unicast prefix-list export list-name
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv4-unicast prefix-list export list-name
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast prefix-list
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

list-name

Name of a configured prefix list.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv4-unicast {
          prefix-list {
            export list-name
          }
        }
      }
    }
  }
}
```



```
    }  
  }  
}
```

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 <name> address-family ipv4-unicast prefix-list import <list-name>

Applies a prefix list to filter updates from this peer group.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv4-unicast prefix-list import list-name
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv4-unicast prefix-list import list-name
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast prefix-list
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

list-name

Name of a configured prefix list.

Configuration mode

```
protocols {  
  bgp asn {  
    peer-group name {  
      address-family {  
        ipv4-unicast {  
          prefix-list {  
            import list-name  
          }  
        }  
      }  
    }  
  }  
}
```

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 <name> address-family ipv4-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 name address-family ipv4-unicast remove-private-as
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv4-unicast remove-private-as
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast
```

Private AS numbers are included in outgoing updates.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv4-unicast {
          remove-private-as
        }
      }
    }
  }
}
```

Use the `set` form of this command to direct the router to exclude private 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 through 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 <name> address-family ipv4-unicast route-map export <map-name>

Applies a route map to filter updates to this peer group.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv4-unicast route-map export map-name
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv4-unicast route-map export map-name
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name



The peer group name.

map-name

Name of a configured route map.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv4-unicast {
          route-map {
            export map-name
          }
        }
      }
    }
  }
}
```

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.

Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

protocols bgp <asn> peer-group <name> address-family ipv4-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 name address-family ipv4-unicast route-map import map-name
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv4-unicast route-map import map-name
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast route-map
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

map-name

Name of a configured route map.

Configuration mode



```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv4-unicast {
          route-map {
            import map-name
          }
        }
      }
    }
  }
}
```

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.

Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

protocols bgp <asn> peer-group <name> address-family ipv4-unicast soft-reconfiguration inbound

Directs the router to store received routing updates.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv4-unicast soft-reconfiguration inbound
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv4-unicast soft-reconfiguration inbound
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast
```

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
```



```
        ipv4-unicast {
            soft-reconfiguration {
                inbound
            }
        }
    }
}
```

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.

The `soft-reconfiguration` command helps conserve memory utilization because you request for a route-refresh from a peer and update the routing table only when there is a change in the network.

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 <name> address-family ipv4-unicast unsuppress-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 name address-family ipv4-unicast unsuppress-map map-name
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv4-unicast unsuppress-map [ map-name ]
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast
```

Routes suppressed by address aggregation are not advertised.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

map-name

The name of a configured route map.

Configuration mode

```
protocols {
    bgp asn {
        peer-group name {
            address-family {
                ipv4-unicast {
                    peer-group name {
                        unsuppress-map map-name
                    }
                }
            }
        }
    }
}
```

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 <name> address-family ipv4-unicast weight <weight>

Defines a default weight for routes from this peer group.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv4-unicast weight weight
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv4-unicast weight [ weight ]
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast
```

Routes learned from a BGP neighbor have a weight of 0. Routes sourced by the local router have a weight of 32768.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

weight

The weight to be assigned to routes from this peer group. The numbers range from 0 through 65535.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv4-unicast {
          weight weight
        }
      }
    }
  }
}
```

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.

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

Enables the IPv6 unicast address-family for a peer group and defines parameters for controlling IPv6 unicast routes when subsequent configuration parameters and options are present.

Syntax:

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

Syntax:



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

Syntax:

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

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

Configuration mode

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

Use the set form of this command to enable the IPv6 unicast address-family for a peer group and to define parameters for controlling IPv6 unicast routes when subsequent configuration parameters and options are present.

Use the delete form of this command to disable the IPv6 unicast address-family for a peer group and to remove the configured parameters for controlling IPv6 unicast routes.

Use the show form of this command to view the configuration settings for the IPv6 unicast address-family of a peer group.

protocols bgp <asn> peer-group <name> address-family ipv6-unicast allows-in number <number>

Allows or disallows receiving BGP advertisements containing the AS path of the local router.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv6-unicast allows-in [ number number ]
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv6-unicast allows-in
```

Syntax:

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

Readvertisement is disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

number *number*

Specifies the number of times to accept an AS path containing the readvertisement of the ASN of the local router. The numbers range from 1 through 10. The default is 3.

Configuration mode

```
protocols {
  bgp asn {
```




```
peer-group name {
  address-family {
    ipv6-unicast {
      allowas-in {
        number number
      }
    }
  }
}
```

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 AS number of the router.

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 AS path of the router.

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

Note:

When you configure this command and do a soft reset in the peer, the BGP table should flush out the prefixes that have an as-path length that is more than the configured `allowas-in` value.

When soft-reconfiguration is enabled irrespective of the configured `allowas-in` value, all the received prefixes should be stored in Adj-Rib-In, and then `allowas-in` should be applied and prefixes with an as-path length more than the configured value should be flushed out of the BGP table. For more information on soft-reconfiguration, see [protocols bgp <asn> neighbor <id> address-family ipv6-unicast soft-reconfiguration inbound \(page 207\)](#).

protocols bgp <asn> peer-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 name address-family ipv6-unicast attribute-unchanged [ as-path | med | next-hop ]
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv6-unicast attribute-unchanged [ as-path | med | next-hop ]
```

Syntax:

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

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

as-path

Propagates the route update with unchanged AS_PATH attribute.

med

Propagates the route update with unchanged Multi Exit Discriminator (MED).

next-hop



Propagates the route update with unchanged next hop.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv6-unicast {
          attribute-unchanged {
            as-path
            med
            next-hop
          }
        }
      }
    }
  }
}
```

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, 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 <name> address-family ipv6-unicast capability graceful-restart

Enables or disables the graceful restart capability for the specified BGP peer group of an IPv6 unicast address family.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv6-unicast capability graceful-restart [ disable ]
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv6-unicast graceful-restart [ disable ]
```

Syntax:

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

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

disable

Disables the graceful restart capability.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv6-unicast {
          capability
        }
      }
    }
  }
}
```



```
        graceful-restart {
            disable
        }
    }
}
}
```

When a neighbor is a member of a peer group, graceful restart cannot be configured on the neighbor level. A graceful restart configured at the neighbor level or the peer-group level, will override the global-level graceful restart configuration.

Use the `set` form of this command to enable the graceful restart capability for the specified BGP peer group of an IPv6 unicast address family.

Use the `delete` form of this command to disable the graceful restart capability for the specified BGP peer group of an IPv6 unicast address family.

Use the `show` form of this command to display the graceful restart capability for the specified BGP peer group of an IPv6 unicast address family.

protocols bgp <asn> peer-group <name> address-family ipv6-unicast capability orf

Advertises support for Outbound Route Filtering (ORF) for updating BGP capabilities for updating BGP capabilities advertised and received from this peer group.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv6-unicast capability orf { prefix-list [ receive | send ] }
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv6-unicast capability orf [ prefix-list [ receive | send ] ]
```

Syntax:

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

By default, the session is brought up with minimal capability on both sides.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

prefix-list

Advertise prefix-list ORF capability to this peer group.

receive

Advertise the ability to receive the ORF from this peer group.

send

Advertise the capability to send the ORF to this peer group.

Configuration mode

```
protocols {
    bgp asn {
        peer-group name {
            address-family {
                ipv6-unicast
                capability {
                    orf {
                        prefix-list {
```



```

    receive
    send
}
}
}
}
}
}
}
}
}
}
}
}
}

```

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 <name> address-family ipv6-unicast capability orf prefix-list <action>

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 name address-family ipv6-unicast capability orf prefix-list {
receive | send }
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv6-unicast capability orf prefix-list [
receive | send ]
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv6-unicast capability orf prefix-list
```

By default, the session is brought up with minimal capability on both sides.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

prefix-list

Advertise prefix-list ORF capability to this peer group.

receive

Advertise the ability to receive the ORF from this peer group.

send

Advertise the capability to send the ORF to this peer group.

Configuration mode

```

protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv6-unicast {
          capability {
            orf {
              prefix-list {
                receive
                send
              }
            }
          }
        }
      }
    }
  }
}

```



```
}
  }
}
  }
}
  }
}
```

This command applies to IPv6 unicast routes only.

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 <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 name address-family ipv6-unicast default-originate route-map map-name
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv6-unicast default-originate route-map [map-name]
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv6-unicast default-originate route-map
```

The default route is not distributed.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

map-name

Specifies a configured route map to be used when advertising the default route.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv6-unicast {
          default-originate {
            route-map map-name
          }
        }
      }
    }
  }
}
```

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 <name> local-as <asn>` ([page 287](#)).

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.

Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

`protocols bgp <asn> peer-group <name> address-family ipv6-unicast disable-send-community`

Disables sending of community attributes to a specified peer group.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv6-unicast disable-send-community [ extended | standard ]
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv6-unicast disable-send-community [ extended | standard ]
```

Syntax:

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

Communities are sent by default.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

extended

Disable sending extended community attributes.

standard

Disable sending standard community attributes.

Configuration mode

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



```

        disable-send-community {
            extended
            standard
        }
    }
}

```

This command applies to IPv6 unicast routes only.

Use the `set` form of this command to enable and disable sending community attributes between BGP neighbors.

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 <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 name address-family ipv6-unicast distribute-list export access-list6-name
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv6-unicast distribute-list export [ access-list6-name ]
```

Syntax:

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

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

access-list6-name

The name of the IPv6 access-list.

Configuration mode

```

protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv6-unicast {
          distribute-list {
            export access-list6-name
          }
        }
      }
    }
  }
}

```

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 <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 name address-family ipv6-unicast distribute-list import access-list6-name`

Syntax:

`delete protocols bgp asn peer-group name address-family ipv6-unicast distribute-list import [access-list6-name]`

Syntax:

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

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

access-list6-name

The name of the IPv6 access-list.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv6-unicast {
          distribute-list {
            import access-list6-name
          }
        }
      }
    }
  }
}
```

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 <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 name address-family ipv6-unicast filter-list export as-path-list-name
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv6-unicast filter-list export [ as-path-list-name ]
```

Syntax:

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

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

as-path-list-name

Name of an autonomous system path list.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv6-unicast {
          filter-list {
            export as-path-list-name
          }
        }
      }
    }
  }
}
```

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 <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 name address-family ipv6-unicast filter-list import as-path-list-name
```

Syntax:



```
delete protocols bgp asn peer-group name address-family ipv6-unicast filter-list import [ as-path-list-name ]
```

Syntax:

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

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

as-path-list-name

Name of an autonomous system path list.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv6-unicast {
          filter-list {
            import as-path-list-name
          }
        }
      }
    }
  }
}
```

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 <name> address-family ipv6-unicast maximum-prefix <max-number>

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 name address-family ipv6-unicast maximum-prefix max-number [ threshold threshold-value [ warning-only ] warning-only ]
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv6-unicast maximum-prefix max-number [ threshold threshold-value [ warning-only ] warning-only ]
```

Syntax:

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

No maximum to the number of prefixes accepted.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

max-number

The maximum number of prefixes to accept from this peer group before the peer group is taken down.

threshold-value



The percentage of the maximum value of prefixes at which to start warning.

warning-only

Does not reset the connection once the maximum prefixes have been received. It only continues to warn.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv6-unicast {
          maximum-prefix max-number {
            threshold threshold-value {
              warning-only
            }
          }
        }
      }
    }
  }
}
```

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.

Note: You can either configure `maximum-prefix max-number [warning-only]` or `maximum-prefix max-number threshold threshold-value [warning-only]`. The two values do not coexist in a configuration.

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.

Note:

When you configure this command and the number of prefixes received in an update message exceeds configured `maximum-prefix` limit, the session gets reset.

When soft-reconfiguration is enabled and the number of prefixes received exceeds `maximum-prefix` value, there will not be any corresponding unprocessed entries stored in Adj-RIB-In because the peer-session resets when the number of prefixes received exceeds `maximum-prefix` value. For more information on soft-reconfiguration, see [protocols bgp <asn> neighbor <id> address-family ipv6-unicast soft-reconfiguration inbound \(page 207\)](#).

protocols bgp <asn> peer-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 name address-family ipv6-unicast nexthop-self
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv6-unicast nexthop-self
```

Syntax:

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



Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv6-unicast {
          nexthop-self
        }
      }
    }
  }
}
```

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 <name> address-family ipv6-unicast prefix-list export <name>

Applies a prefix list to filter updates to this peer group.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv6-unicast prefix-list export prefix-list6-name
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv6-unicast prefix-list export prefix-list6-name
```

Syntax:

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

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

prefix-list6-name

Name of a configured IPv6 prefix list.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv6-unicast {
          prefix-list {
            export prefix-list6-name
          }
        }
      }
    }
  }
}
```



```
    }
  }
}
```

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 <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 name address-family ipv6-unicast prefix-list import prefix-list6-name
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv6-unicast prefix-list import [ prefix-list6-name ]
```

Syntax:

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

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

prefix-list6-name

Name of a configured IPv6 prefix list.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv6-unicast {
          prefix-list {
            import prefix-list6-name
          }
        }
      }
    }
  }
}
```

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 <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 name address-family ipv6-unicast remove-private-as
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv6-unicast remove-private-as
```

Syntax:

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

Private AS numbers are included in outgoing updates.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv6-unicast {
          remove-private-as
        }
      }
    }
  }
}
```

This command applies to IPv6 unicast routes only.

Use the `set` form of this command to direct the router to exclude private AS numbers from updates to eBGP peers. When this feature is enabled, the router omits private AS numbers from the AS_PATH attribute. The private AS numbers range from 64512 through 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 <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 name address-family ipv6-unicast route-map export map-name
```

**Syntax:**

```
delete protocols bgp asn peer-group name address-family ipv6-unicast route-map export map-name
```

Syntax:

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

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

map-name

Name of a configured route map.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv6-unicast {
          route-map {
            export map-name
          }
        }
      }
    }
  }
}
```

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.

Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

protocols bgp <asn> peer-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 name address-family ipv6-unicast route-map import map-name
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv6-unicast route-map import [ map-name ]
```

**Syntax:**

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

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

map-name

Name of a configured route map.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv6-unicast {
          route-map {
            import map-name
          }
        }
      }
    }
  }
}
```

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.

Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

protocols bgp <asn> peer-group <name> address-family ipv6-unicast soft-reconfiguration inbound

Directs the router to store received routing updates.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv6-unicast soft-reconfiguration inbound
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv6-unicast soft-reconfiguration inbound
```

Syntax:

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




Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv6-unicast {
          soft-reconfiguration {
            inbound
          }
        }
      }
    }
  }
}
```

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.

The `soft-reconfiguration` command helps conserve memory utilization because you request for a route-refresh from a peer and update the routing table only when there is a change in the network.

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 <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 name address-family ipv6-unicast unsuppress-map map-name
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv6-unicast unsuppress-map map-name
```

Syntax:

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

Routes suppressed by address aggregation are not advertised.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

map-name

The name of a configured route map.

Configuration mode



```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv6-unicast {
          unsuppress-map map-name
        }
      }
    }
  }
}
```

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 <name> address-family ipv6-unicast weight <weight>

Defines a default weight for routes from this peer group.

Syntax:

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

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv6-unicast weight [ weight ]
```

Syntax:

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

Routes learned from a BGP neighbor have a weight of 0. Routes sourced by the local router have a weight of 32768.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

weight

The weight to be assigned to routes from this peer group. The numbers range from 0 through 65535.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv6-unicast {
          weight weight
        }
      }
    }
  }
}
```

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.



protocols bgp <asn> peer-group <name> advertisement-interval <seconds>

Sets the minimum interval for sending BGP route updates.

Syntax:

```
set protocols bgp asn peer-group name advertisement-interval seconds
```

Syntax:

```
delete protocols bgp asn peer-group name advertisement-interval [ seconds ]
```

Syntax:

```
show protocols bgp asn peer-group name advertisement-interval
```

The default advertisement interval is 30 seconds for eBGP peers and 5 seconds for iBGP peers.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

seconds

The minimum interval, in seconds, between BGP routing updates to this neighbor. The numbers range from 0 through 600. The default is 30 for eBGP peers and 5 for iBGP peers.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      advertisement-interval seconds
    }
  }
}
```

Use the `set` form of this command to set the minimum interval between BGP routing advertisements to a BGP peer group.

Use the `delete` form of this command to restore the advertisement interval to the default.

Use the `show` form of this command to view the configuration settings for a BGP peer group.

protocols bgp <asn> peer-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 name capability dynamic
```

Syntax:

```
delete protocols bgp asn peer-group name capability dynamic
```

Syntax:

```
show protocols bgp asn peer-group name
```

By default, the session is brought up with minimal capability on both sides.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name



The peer group name.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      capability {
        dynamic
      }
    }
  }
}
```

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 <name> capability graceful-restart

Advertises support for dynamic update of the graceful restart of BGP capabilities advertised and received from this peer group.

Syntax:

```
set protocols bgp asn peer-group name capability graceful-restart [ disable ]
```

Syntax:

```
delete protocols bgp asn peer-group name capability graceful-restart [ disable ]
```

Syntax:

```
show protocols bgp asn peer-group name graceful-restart
```

The session is brought up with minimal capability on both sides.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      capability {
        graceful-restart
        graceful-restart disable
      }
    }
  }
}
```



```
}
```

Use the `set` form of this command to advertise support for dynamic update of a graceful restart of BGP capabilities advertised and received from this peer group.

Use the `delete` form of this command to remove support for dynamic update of a graceful restart of BGP capabilities advertised and received from this peer group.

Use the `show` form of this command to view support for dynamic update of a graceful restart of BGP capabilities advertised and received from this peer group.

protocols bgp <asn> peer-group <name> description <description>

Provides a brief description for a BGP peer group.

Syntax:

```
set protocols bgp asn peer-group name description description
```

Syntax:

```
delete protocols bgp asn peer-group name description
```

Syntax:

```
show protocols bgp asn peer-group name
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

description

A description of up to 80 characters for the peer group. If spaces are included, the description must be enclosed in quotes.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      description description
    }
  }
}
```

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 <name> disable-capability-negotiation

Disables BGP capability negotiation.

Syntax:

```
set protocols bgp asn peer-group name disable-capability-negotiation
```

Syntax:

```
delete protocols bgp asn peer-group name disable-capability-negotiation
```

**Syntax:**

```
show protocols bgp asn peer-group name
```

Capability negotiation is performed.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      disable-capability-negotiation
    }
  }
}
```

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 <name> ebgp-multihop <ttl>

Allows eBGP neighbors not on directly connected networks.

Syntax:

```
set protocols bgp asn peer-group name ebgp-multihop ttl
```

Syntax:

```
delete protocols bgp asn peer-group name ebgp-multihop [ ttl ]
```

Syntax:

```
show protocols bgp asn peer-group name ebgp-multihop
```

Only directly connected neighbors are allowed.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

ttl

The time-to-live, or maximum number of hops allowed. The numbers range from 1 through 255.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      ebgp-multihop ttl
    }
  }
}
```

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 <name> local-as <asn>

Defines a local autonomous system number for eBGP peerings.

Syntax:

```
set protocols bgp asn peer-group name local-as asn
```

Syntax:

```
delete protocols bgp asn peer-group name local-as [ asn ]
```

Syntax:

```
show protocols bgp asn peer-group name local-as
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

asn

Valid autonomous system number. This may not be the number of the autonomous system to which the peer group belongs. The numbers range from 1 through 4294967294.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      local-as asn
    }
  }
}
```

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.

Use the `delete` form of this command to remove a local autonomous system number.

Use the `show` form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <name> med-out igp

Sets the value of a multi-exit discriminator (MED) that is based on the Interior Gateway Protocol (IGP) metric to all peers in a peer group.

Syntax:

```
set protocols bgp asn peer-group name med-out igp
```

Syntax:

```
delete protocols bgp asn peer-group name med-out igp
```

asn

Mandatory if an ASN other than the parent ASN is to be sent to the neighbor. The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The name of the peer group.



Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      med-out igp
    }
  }
}
```

Use the `set` form of this command to set the value of a multi-exit discriminator (MED) that is based on the Interior Gateway Protocol (IGP) metric to all peers in a peer group.

Use the `delete` form of this command to remove the MED.

protocols bgp <asn> peer-group <name> med-out igp delay-updates

Delays the update of the multi-exit discriminator (MED) to all peers in the peer-group.

Syntax:

```
set protocols bgp asn peer-group name med-out igp delay-updates
```

Syntax:

```
delete protocols bgp asn peer-group name med-out igp delay-updates
```

10 minutes.

asn

Mandatory if an ASN other than the parent ASN is to be sent to the neighbor. The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      med-out igp delay-updates
    }
  }
}
```

Use the `set` form of this command to delay MED updates to all peers in the peer group. Enter this command if the device is already configured to assign the MED that is based on the IGP metric. The time delay is based on the default timer that is configured in `protocols bgp <asn> parameters med-out-delay <minutes>` ([page 124](#)).

Note: The device delays the MED updates if the IGP metric changes and the new metric is greater than the old IGP metric.

Use the `delete` form of this command to update the MED value as soon as the IGP metric changes.

protocols bgp <asn> peer-group <name> med-out minimum-igp

Sets the value of the multi-exit discriminator (MED) that is based on the smaller Interior Gateway Protocol (IGP) metric.

Syntax:



```
set protocols bgp asn peer-group name med-out minimum-igp
```

Syntax:

```
delete protocols bgp asn peer-group name med-out minimum-igp
```

asn

Mandatory if an ASN other than the parent ASN is to be sent to the neighbor. The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      med-out minimum-igp
    }
  }
}
```

Use the `set` form of this command to set the MED value based on the smaller IGP metric. The minimum value of IGP metric is the smallest IGP metric received by BGP from RIB.

Note: This command and `protocols bgp <asn> peer-group <name> med-out igp (page 287)` are mutually exclusive.

Use the `delete` form of this command to disable MED updates that are based on the IGP metric.

protocols bgp <asn> peer-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 name override-capability
```

Syntax:

```
delete protocols bgp asn peer-group name override-capability
```

Syntax:

```
show protocols bgp asn peer-group name
```

A session cannot be established if the peer group does not support capability negotiation.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      override-capability
    }
  }
}
```



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 <name> passive

Directs the router not to initiate connections with this peer group.

Syntax:

```
set protocols bgp asn peer-group name passive
```

Syntax:

```
delete protocols bgp asn peer-group name passive
```

Syntax:

```
show protocols bgp asn peer-group name
```

The router both accepts inbound connections and initiates outbound connections.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      passive
    }
  }
}
```

Use the `set` form of this command to configure the local router such that it 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 <name> password <password>

Defines a BGP MD5 password.

Syntax:

```
set protocols bgp asn peer-group name password password
```

Syntax:

```
delete protocols bgp asn peer-group name password
```

Syntax:

```
show protocols bgp asn peer-group name password
```

**asn**

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

pwd

A password to be used to generate an MD5 digest.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      password pwd
    }
  }
}
```

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 <name> remote-as <asn>

Specifies the AS number of the peer group.

Syntax:

```
set protocols bgp asn peer-group name remote-as asn
```

Syntax:

```
delete protocols bgp asn peer-group name remote-as [ asn ]
```

Syntax:

```
show protocols bgp asn peer-group name remote-as
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      remote-as asn
    }
  }
}
```

Use the `set` form of this command to specify the 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 <name> shutdown

Administratively shuts down a BGP peer group.

Syntax:

```
set protocols bgp asn peer-group name shutdown
```

Syntax:

```
delete protocols bgp asn peer-group name shutdown
```

Syntax:

```
show protocols bgp asn peer-group name
```

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      shutdown
    }
  }
}
```

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 has been administratively shut down.

Use the delete form of this command to administratively re-enable a BGP peer group.

Use the show form of this command to view BGP peer group configuration settings.

protocols bgp <asn> peer-group <name> strict-capability-match

Directs the router to strictly match the capabilities of the BGP peer group.

Syntax:

```
set protocols bgp asn peer-group name strict-capability-match
```

Syntax:

```
delete protocols bgp asn peer-group name strict-capability-match
```

Syntax:

```
show protocols bgp asn peer-group name
```

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id



The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      strict-capability-match
    }
  }
}
```

Use the `set` form of this command to direct the router to strictly compare the capabilities of the local router and the BGP 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 the configuration settings of a BGP remove peer.

protocols bgp <asn> peer-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 name ttl-security hops hops
```

Syntax:

```
delete protocols bgp asn peer-group name ttl-security hops [ hops ]
```

Syntax:

```
show protocols bgp asn peer-group name ttl-security hops
```

The default for the hop count is 254 for IBGP and 1 for EBGP connections.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

hops

The maximum number of hops away from the local peer a peering session connection be accepted. The numbers range from 1 through 254.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      ttl-security {
        hops hops
      }
    }
  }
}
```

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 <name> update-source <source>

Specifies the source IP address or interface for routing updates.

Syntax:

```
set protocols bgp asn peer-group name update-source source
```

Syntax:

```
delete protocols bgp asn peer-group name update-source [ source ]
```

Syntax:

```
show protocols bgp asn peer-group name
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

source

The IPv4 address of the router to receive routing updates from or the interface to receive routing updates from.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      update-source source
    }
  }
}
```

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.

reset ip bgp peer-group <name>

Resets sessions for all members of a peer group.

Syntax:

```
reset ip bgp peer-group name [ soft [ in [ prefix-filter ] | out ] ]
```

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

name

The name of a defined BGP peer group.

soft

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

Re-read import policies and recalculate information in the BGP table based on import policy changes. This requires `soft-reconfiguration inbound` to be configured on the neighbor.

prefix-filter

Resets the existing ORF prefix list to trigger a soft reconfiguration.

out



Re-read export policies and resend updates that have changed due to an export policy change.

Operational mode

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 re-established.

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 <name> ipv4 unicast

Resets IPv4 unicast sessions for all members of a peer group.

Syntax:

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

name

The name of a defined BGP peer group.

soft

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

Re-read import policies and recalculate information in the BGP table based on import policy changes. This requires **soft-reconfiguration inbound** to be configured on the neighbor.

prefix-filter

Resets the outbound route filter (ORF) filter.

out

Re-read export policies and resend updates that have changed due to an export policy change.

Operational mode

Use this command to reset inbound IPv4 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 re-established.

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 <name> ipv6 unicast

Resets IPv6 unicast sessions for all members of a peer group.

Syntax:

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

name

The name of a defined BGP peer group.

soft

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

Re-read import policies and recalculate information in the BGP table based on import policy changes. This requires **soft-reconfiguration inbound** to be configured on the neighbor.

prefix-filter



Resets the outbound route filter (ORF) filter.

out

Re-read export policies and resend updates that have changed due to an export policy change.

Operational mode

Use this command to reset inbound IPv6 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 re-established.

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.



Route Redistribution Commands

protocols bgp <asn> address-family ipv4-unicast redistribute connected

Redistributes directly connected routes into BGP.

Syntax:

```
set protocols bgp asn address-family ipv4-unicast redistribute connected [ route-map map-name ]
```

Syntax:

```
delete protocols bgp asn address-family ipv4-unicast redistribute connected [ route-map [ map-name ] ]
```

Syntax:

```
show protocols bgp asn address-family ipv4-unicast redistribute connected
```

When this command has not been set, directly connected routes are not distributed into BGP.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

map-name

The name of a configured route map to be used for redistributing connected routes.

Configuration mode

```
protocols {
  bgp asn {
    address-family {
      ipv4-unicast {
        redistribute {
          connected {
            route-map map-name
          }
        }
      }
    }
  }
}
```

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.

Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.



Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

protocols bgp <asn> address-family ipv4-unicast redistribute kernel

Redistributes kernel routes into BGP.

Syntax:

```
set protocols bgp asn address-family ipv4-unicast redistribute kernel [ route-map map-name ]
```

Syntax:

```
delete protocols bgp asn address-family ipv4-unicast redistribute kernel [ route-map [ map-name ] ]
```

Syntax:

```
show protocols bgp asn address-family ipv4-unicast redistribute kernel
```

When this command has not been set, kernel routes are not distributed into BGP.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

map-name

The name of a configured route map to be used for redistributing kernel routes.

Configuration mode

```
protocols {
  bgp asn {
    address-family {
      ipv4-unicast {
        redistribute {
          kernel {
            route-map map-name
          }
        }
      }
    }
  }
}
```

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.

Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.



protocols bgp <asn> address-family ipv4-unicast redistribute ospf

Redistributes routes learned from Open Shortest Path First (OSPF) into BGP.

Syntax:

```
set protocols bgp asn address-family ipv4-unicast redistribute ospf [ route-map map-name ]
```

Syntax:

```
delete protocols bgp asn address-family ipv4-unicast redistribute ospf [ route-map [ map-name ] ]
```

Syntax:

```
show protocols bgp asn address-family ipv4-unicast redistribute ospf
```

When this command has not been set, routes learned from OSPF are not distributed into BGP.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

map-name

The name of a configured route map to be used for redistributing OSPF routes.

Configuration mode

```
protocols {
  bgp asn {
    address-family {
      ipv4-unicast {
        redistribute {
          ospf {
            route-map map-name
          }
        }
      }
    }
  }
}
```

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.

Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

protocols bgp <asn> address-family ipv4-unicast redistribute rip

Redistributes routes learned from Routing Information Protocol (RIP) into BGP.

**Syntax:**

```
set protocols bgp asn address-family ipv4-unicast redistribute rip [ route-map map-name ]
```

Syntax:

```
delete protocols bgp asn address-family ipv4-unicast redistribute rip [ route-map [ map-name ] ]
```

Syntax:

```
show protocols bgp asn address-family ipv4-unicast redistribute rip
```

When this command has not been set, routes learned from RIP are not distributed into BGP.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

map-name

The name of a configured route map to be used for redistributing RIP routes.

Configuration mode

```
protocols {
  bgp asn {
    address-family {
      ipv4-unicast {
        redistribute {
          rip {
            route-map map-name
          }
        }
      }
    }
  }
}
```

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.

Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

protocols bgp <asn> address-family ipv4-unicast redistribute static

Redistributes static routes into BGP.

Syntax:

```
set protocols bgp asn address-family ipv4-unicast redistribute static [ route-map map-name ]
```

Syntax:



```
delete protocols bgp asn address-family ipv4-unicast redistribute static [ route-map [ map-name ] ]
```

Syntax:

```
show protocols bgp asn address-family ipv4-unicast redistribute static
```

When this command has not been set, static routes are not distributed into BGP.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

map-name

The name of a configured route map to be used for redistributing static routes.

Configuration mode

```
protocols {
  bgp asn {
    address-family {
      ipv4-unicast {
        redistribute {
          static {
            route-map map-name
          }
        }
      }
    }
  }
}
```

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.

Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

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 [ route-map map-name ]
```

Syntax:

```
delete protocols bgp asn address-family ipv6-unicast redistribute connected [ route-map [ map-name ] ]
```

Syntax:



```
show protocols bgp asn address-family ipv6-unicast redistribute connected
```

When this command has not been set, directly connected routes are not distributed into BGP.

asn

The number for the autonomous system (AS) in which this router resides. The number ranges from 1 through 4294967295.

map-name

The name of a configured route map to be used for redistributing connected routes.

Configuration mode

```
protocols {
  bgp asn {
    address-family {
      ipv6-unicast {
        redistribute {
          connected {
            route-map map-name
          }
        }
      }
    }
  }
}
```

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.

Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

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 [ route-map map-name ]
```

Syntax:

```
delete protocols bgp asn address-family ipv6-unicast redistribute kernel [ route-map [ map-name ] ]
```

Syntax:

```
show protocols bgp asn address-family ipv6-unicast redistribute kernel
```

When this command has not been set, kernel routes are not distributed into BGP.

**asn**

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

map-name

The name of a configured route map to be used for redistributing kernel routes.

Configuration mode

```
protocols {
  bgp asn {
    address-family {
      ipv6-unicast {
        redistribute {
          kernel {
            route-map map-name
          }
        }
      }
    }
  }
}
```

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.

Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

protocols bgp <asn> address-family ipv6-unicast redistribute ospfv3

Redistributes IPv6 unicast routes learned from Open Shortest Path First Version 3 (OSPFv3) into BGP.

Syntax:

```
set protocols bgp asn address-family ipv6-unicast redistribute ospfv3 [ route-map map-name ]
```

Syntax:

```
delete protocols bgp asn address-family ipv6-unicast redistribute ospfv3 [ route-map [ map-name ] ]
```

Syntax:

```
show protocols bgp asn address-family ipv6-unicast redistribute ospfv3
```

When this command has not been set, routes learned from OSPFv3 are not distributed into BGP.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

map-name



The name of a configured route map to be used for redistributing OSPFv3 routes.

Configuration mode

```
protocols {
  bgp asn {
    address-family {
      ipv6-unicast {
        redistribute {
          ospfv3 {
            route-map map-name
          }
        }
      }
    }
  }
}
```

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 [ route-map map-name ]
```

Syntax:

```
delete protocols bgp asn address-family ipv6-unicast redistribute ripng [ route-map [ map-name ] ]
```

Syntax:

```
show protocols bgp asn address-family ipv6-unicast redistribute ripng
```

When this command has not been set, routes learned from RIPng are not distributed into BGP.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

map-name

The name of a configured route map to be used for redistributing RIPng routes.

Configuration mode

```
protocols {
  bgp asn {
    address-family {
      ipv6-unicast {
        redistribute {
          ripng {
            route-map map-name
          }
        }
      }
    }
  }
}
```




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.

Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.

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 [ route-map map-name ]
```

Syntax:

```
delete protocols bgp asn address-family ipv6-unicast redistribute static [ route-map [ map-name ] ]
```

Syntax:

```
show protocols bgp asn address-family ipv6-unicast redistribute static
```

When this command has not been set, static routes are not distributed into BGP.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

map-name

The name of a configured route map to be used for redistributing static routes.

Configuration mode

```
protocols {
  bgp asn {
    address-family {
      ipv6-unicast {
        redistribute {
          static {
            route-map map-name
          }
        }
      }
    }
  }
}
```

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.

Note: In the case of the IPv4 address family, if extended IPv4 ACLs are used in a route map, the route must be specified in the source field of the ACL and the net mask of the route must be specified in the destination field of the ACL.

The inverse mask for the source field is applied to the prefix before the comparison with the source. The inverse mask for the destination field is applied to the net mask of the prefix before the comparison with the net mask as specified in the destination field.

Effectively, the prefix must match the source with application of inverse mask of the source field. And the net mask for the prefix must match the destination with application of inverse mask of the destination field.



Route Server Commands

protocols bgp <asn> neighbor <id> address-family ipv4-unicast route-server-client

Specifies this neighbor as a route server client.

Syntax:

```
set protocols bgp asn neighbor id address-family ipv4-unicast route-server-client
```

Syntax:

```
delete protocols bgp asn neighbor id address-family ipv4-unicast route-server-client
```

Syntax:

```
show protocols bgp asn neighbor id address-family ipv4-unicast route-server-client
```

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

```
protocols {
  bgp asn {
    neighbor id {
      address-family {
        ipv4-unicast {
          route-server-client
        }
      }
    }
  }
}
```

Use the `set` form of this command to configure the specified neighbor as a route server client.

Use the `delete` form of this command to remove the specified neighbor as a route server client.

Use the `show` form of this command to view the route server client configuration for the specified neighbor.

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
```

Syntax:

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

Syntax:

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



Disabled.

asn

The number for the autonomous system (AS) in which this router resides. The number ranges from 1 through 4294967295.

id

The IPv4 or IPv6 address of a BGP neighbor.

Configuration mode

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

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> peer-group <name> address-family ipv4-unicast route-server-client

Specify this peer group as a route server client.

Syntax:

```
set protocols bgp asn peer-group name address-family ipv4-unicast route-server-client
```

Syntax:

```
delete protocols bgp asn peer-group name address-family ipv4-unicast route-server-client
```

Syntax:

```
show protocols bgp asn peer-group name address-family ipv4-unicast
```

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295

name

The peer group name.

Configuration mode

```
protocols {
  bgp asn {
    peer-group name {
      address-family {
        ipv4-unicast {
          route-server-client
        }
      }
    }
  }
}
```



```
}
```

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 <name> address-family ipv6-unicast route-server-client

Specify this peer group as a route server client.

Syntax:

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

Syntax:

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

Syntax:

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

Disabled.

asn

The number for the AS in which this router resides. The number ranges from 1 through 4294967295.

name

The peer group name.

Configuration mode

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

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.



VRF Support for BGP

VRF support for BGP

The implementation of VRF on the AT&T Vyatta vRouter supports BGP.

If you configure BGP on the vRouter without specifying a VRF instance, the router uses the default routing table. BGP supports route-leaking between VRF's by using route target. To configure BGP for a particular VRF routing instance, specify the instance in the command syntax.

Note: You must configure a unique route distinguisher, before you configure BGP with a VRF routing instance. You cannot further modify the route distinguisher configuration if BGP is already configured within a routing instance. Use the following command to configure a route distinguisher:

```
vyatta@R1# set routing routing-instance VRF-NAME route-distinguisher RDVALUE
```

BGP monitor commands log messages from the default VRF, unless the name of the VRF is specified in the configuration of the routing instance.

The following BGP features are supported only on the default routing instance:


- Extended ASN capability
- Graceful restart
- BGP scan timer
- Cluster ID
- Confederation identifier
- Maximum AS limit

Configuring BGP on a routing instance

This example shows how to assign BGP to the RED routing instance on the IPv4 address family.

In this example the R1 vRouter and R2 vRouter consists of RED and BLUE routing instances. The R1 and R2 vRouters are connected with the RED routing instance.

Figure 22: Configuring BGP on a routing instance



Note: Ensure that you have configured the BGP AS number in the default mode prior to configuring BGP on a routing instance.

**Table 19: Configuring BGP on a routing instance**

Step	Command
The configuration before reassigning routing instances.	<pre>vyatta@R1# show protocols protocols { bgp 2 { parameters { router-id 5.5.5.5 } } }</pre>
Configure a unique route distinguisher before you configure BGP on a routing instance.	<pre>vyatta@R1# set routing routing-instance RED route-distinguisher 2:1</pre>
Add the AS number for the address family.	<pre>vyatta@R1# set routing routing-instance RED protocols bgp 2 neighbor 12.12.12.1 remote-as 1</pre>
Enable the IPv4 unicast address family for a neighbor on the RED routing instance.	<pre>vyatta@R1# set routing routing-instance RED protocols bgp 2 neighbor 12.12.12.1 address-family ipv4-unicast</pre>
Redistribute the IPv4 unicast static route into BGP.	<pre>vyatta@R1# set routing RED protocols bgp 2 address-family ipv4-unicast redistribute static</pre>
Configure a router ID for BGP within RED routing instance.	<pre>vyatta@R1# set routing routing-instance RED protocols bgp 2 parameters router-id 66.66.66.66</pre>
Enable all BGP logs for the RED routing instance.	<pre>vyatta@R1# set routing routing-instance RED protocols bgp 2 log all</pre>
Commit the configuration.	<pre>vyatta@R1# commit</pre>



Step	Command
View the candidate configuration.	<pre>vyatta@R1# show routing routing { routing-instance RED { protocols { bgp 2 { log { all } neighbor 12.12.12.1 { address-family { ipv4-unicast } redistribute-static } remote-as 1 } parameters { router-id 66.66.66.66 } } route-distinguisher 2:1 } }</pre>
View the route table for the RED routing instance.	<pre>vyatta@R1:~\$ show ip bgp routing-instance RED BGP table version is 1, local router ID is 12.12.12.12 Status codes: s suppressed, d damped, h history, * valid, > best, i -internal, l -labeled S Stale Origin codes: i -IGP, e -EGP, ? - incomplete Network Next Hop Metric LocPrf Weight Path *> 100.1.1.1/32 12.12.12.1 0 0 1 ? *> 100.2.2.2/32 12.12.12.1 0 0 1 ? *> 100.11.11.11/32 12.12.12.1 0 0 1 i *> 200.1.1.1/32 12.12.12.1 0 32768 ? *> 200.2.2.2/32 12.12.12.1 0 32768 ? Total number of prefixes 5</pre>

Configuring inter-VRF BGP on routing instances

This example shows how to configure inter-vrf BGP between the RED and BLUE routing instances.

You can leak BGP routes from one VRF to another by ensuring that the import and export route targets match.



This example shows how to configure a route target for inter-VRF BGP route leaking. You can leak routes from one VRF to another on a BGP router that consists of two routing instances namely, RED and BLUE, by using a route target. Ensure that the import and export route targets match with each other to configure inter-VRF route leaking.

Figure 23: Configuring inter-VRF BGP on routing instances



For more information about configuring BGP on routing instances, refer to [Configuring BGP on routing instances \(page 310\)](#).

Table 20: Configuring inter-VRF BGP on routing instances

Step	Command
Configure a route distinguisher for the RED routing instance.	<pre>vyatta@R1# set routing routing-instance RED route-distinguisher 2:1</pre>
Configure a route target to export routes from the RED routing instance.	<pre>vyatta@R1# set routing routing-instance RED route-target 2:1 type export</pre>
Configure a route distinguisher for the BLUE routing instance.	<pre>vyatta@R1# set routing routing-instance BLUE route-distinguisher 2:2</pre>
Configure a route target to export routes to the BLUE routing instance.	<pre>vyatta@R1# set routing routing-instance BLUE route-target 2:1 type import</pre>
Add the RED routing instance to a network	<pre>vyatta@R1# set routing routing-instance RED protocols bgp 100 address-family ipv4-unicast network '100.11.11.11/32</pre>
Redistribute the IPv4 unicast static route into BGP.	<pre>vyatta@R1# set routing routing-instance RED protocols bgp 2 address-family ipv4-unicast redistribute static</pre>
Set a router for BGP with the RED routing instance.	<pre>vyatta@R1# set routing routing-instance RED protocols bgp 2 parameters router-id 66.66.66.66</pre>
Enable all BGP logs for the RED routing instance.	<pre>vyatta@R1# set routing routing-instance RED protocols bgp 2 log all</pre>



Step	Command
Commit the configuration,	<pre>vyatta@R1# commit</pre>
View the route table for the BLUE routing instance.	<pre>vyatta@R1# run show ip bgp routing-instance BLUE BGP table version is 1, local router ID is 0.0.0.0 Status codes: s suppressed, d damped, h history, * valid, > best, i -internal, l - labeled S Stale Origin codes: i -IGP, e -EGP, ? -incomplete Network Next Hop Metric LocPrf Weight Path *> 100.1.1.1/32 12.12.12.10 0 0 1 ? *> 100.2.2.2/32 12.12.12.10 0 0 1 ? *> 100.11.11.11/32 12.12.12.1 0 0 1 i *> 200.1.1.1/32 12.12.12.10 0 32768 ? *> 200.2.2.2/32 12.12.12.10 0 32768 ? Total number of prefixes 5</pre>

Command support for VRF routing instances

VRF allows an AT&T Vyatta vRouter to support multiple routing tables, one for each VRF routing instance. Some commands in this guide support VRF and can be applied to particular routing instances.

Use the guidelines in this section to determine correct syntax when adding VRF routing instances to commands. For more information about VRF, refer to AT&T Vyatta Network Operating System Basic Routing Configuration Guide. This guide includes an overview of VRF, VRF configuration examples, information about VRF-specific features, and a list of commands that support VRF routing instances.

Adding a VRF routing instance to a Configuration mode command

For most Configuration mode commands, specify the VRF routing instance at the beginning of a command. Add the appropriate VRF keywords and variable to follow the initial action (**set**, **show**, or **delete**) and before the other keywords and variables in the command.

Example: Configuration mode example: syslog

The following command configures the syslog logging level for the specified syslog host. The command does not include a VRF routing instance, so the command applies to the default routing instance.

```
vyatta@R1# set system syslog host 10.10.10.1 facility all level debug
vyatta@R1# show system syslog
syslog {
  host 10.10.10.1 {
    facility all {
      level debug
    }
  }
}
```



The following example shows the same command with the VRF routing instance (GREEN) added. Notice that **routing routing-instance GREEN** has been inserted between the basic action (**set** in the example) and the rest of the command. Most Configuration mode commands follow this convention.

```
vyatta@R1# set routing routing-instance GREEN system syslog host 10.10.10.1 facility all
level debug
vyatta@R1# show routing
routing {
  routing-instance GREEN {
    system {
      syslog {
        host 11.12.13.2:514 {
          facility all {
            level debug
          }
        }
      }
    }
  }
}
```

Example: Configuration mode example: SNMP

Some features, such as SNMP, are not available on a per-routing instance basis but can be bound to a specific routing instance. For these features, the command syntax is an exception to the convention of specifying the routing instance at the beginning of Configuration mode commands.

The following example shows how to configure the SNMPv1 or SNMPv2c community and context for the RED and BLUE routing instances. The first two commands specify the RED routing instance as the context for community A and BLUE routing instance as the context for community B. The subsequent commands complete the configuration.

For more information about configuring SNMP, refer to AT&T Vyatta Network Operating System Remote Management Configuration Guide.

```
vyatta@R1# set service snmp community commA context RED
vyatta@R1# set service snmp community commB context BLUE
vyatta@R1# set service snmp view all oid 1
vyatta@R1# set service snmp community commA view all
vyatta@R1# set service snmp community commB view all
vyatta@R1# show service snmp community
community commA {
  context RED
  view all
}
community commB {
  context BLUE
  view all
}
[edit]
vyatta@vyatta#
```

Adding a VRF routing instance to an Operational mode command

The syntax for adding a VRF routing instance to an Operational mode command varies according to the type of command parameters:

- If the command does not have optional parameters, specify the routing instance at the end of the command.



- If the command has optional parameters, specify the routing instance after the required parameters and before the optional parameters.

Example: Operational mode examples without optional parameters

The following command displays dynamic DNS information for the default routing instance.

```
vyatta@vyatta:~$ show dns dynamic status
```

The following command displays the same information for the specified routing instance (GREEN). The command does not have any optional parameters, so the routing instance is specified at the end of the command.

```
vyatta@vyatta:~$ show dns dynamic status routing-instance GREEN
```

Example: Operational mode example with optional parameters

The following command obtains multicast path information for the specified host (10.33.2.5). A routing instance is not specified, so the command applies to the default routing instance.

```
vyatta@vyatta:~$ mtrace 10.33.2.5 detail
```

The following command obtains multicast path information for the specified host (10.33.2.5) and routing instance (GREEN). Notice that the routing instance is specified before the optional **detail** keyword.

```
vyatta@vyatta:~$ mtrace 10.33.2.5 routing-instance GREEN detail
```

Example: Operational mode example output: SNMP

The following SNMP **show** commands display output for routing instances.

```
vyatta@vyatta:~$ show snmp routing-instance
Routing Instance SNMP Agent is Listening on for Incoming Requests:
Routing-Instance      RDID
-----
RED                   5
```

```
vyatta@vyatta:~$ show snmp community-mapping
SNMPv1/v2c Community/Context Mapping:
Community             Context
-----
commA                 'RED'
commB                 'BLUE'
deva                  'default'
```

```
vyatta@vyatta:~$ show snmp trap-target
SNMPv1/v2c Trap-targets:
Trap-target           Port   Routing-Instance Community
-----
1.1.1.1               ----  'RED'           'test'
```

```
vyatta@vyatta:~$ show snmp v3 trap-target
SNMPv3 Trap-targets:
```



Trap-target Instance	User	Port	Protocol	Auth	Priv	Type	EngineID	Routing-
2.2.2.2	'test'	'162'	'udp'	'md5'		'infor'		'BLUE'



List of Acronyms

Acronym	Description
ACL	access control list
ADSL	Asymmetric Digital Subscriber Line
AH	Authentication Header
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
DMVPN	dynamic multipoint VPN
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



Acronym	Description
IGMP	Internet Group Management Protocol
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
ISAKMP	Internet Security Association and Key Management Protocol
ISM	Internet Standard Multicast
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
mGRE	multipoint GRE
MIB	Management Information Base
MLD	Multicast Listener Discovery
MLPPP	multilink PPP
MRRU	maximum received reconstructed unit
MTU	maximum transmission unit
NAT	Network Address Translation
NBMA	Non-Broadcast Multi-Access
ND	Neighbor Discovery
NHRP	Next Hop Resolution Protocol
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
PIM	Protocol Independent Multicast
PIM-DM	PIM Dense Mode
PIM-SM	PIM Sparse Mode
PKI	Public Key Infrastructure
PPP	Point-to-Point Protocol
PPPoA	PPP over ATM



Acronym	Description
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
RP	Rendezvous Point
RPF	Reverse Path Forwarding
RSA	Rivest, Shamir, and Adleman
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
SPT	Shortest Path Tree
SSH	Secure Shell
SSID	Service Set Identifier
SSM	Source-Specific Multicast
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