

Chapter 22

Multi-Protocol - Border Gateway Protocol (MPBGP)

22.1 MPBGP Overview

The Multi-Protocol Border Gateway Protocol (MPBGP) is a set of extensions to BGP-4 to make the protocol capable of carrying routing information for IPv6 and Multicast routes. These extensions are backward compatible, making it possible for BGP-4 routers to interoperate with MPBGP routers. BGP is an exterior routing (or inter-domain routing) protocol used for exchanging routing information between autonomous systems. (See “Chapter 14 Border Gateway Protocol (BGP)” on page 61 for more information about BGP.) MPBGP will support the use of multi-protocol extensions for BGP-4, IPv6 and Multicast.

For more details on the multiple RIBs and the multicast RIB, see “Chapter 9 Multiple Routing Information Bases (RIBs)” on page 33.

MPBGP adds two new attributes, Multiprotocol Reachable NLRI (MP_REACH_NLRI) and Multiprotocol Unreachable NLRI (MP_UNREACH_NLRI). MP_REACH_NLRI carries the set of reachable destinations with next hop information. MP_UNREACH_NLRI contains the set of unreachable destinations.

Capability Advertisement is used to determine whether the Multiprotocol Extensions can be used with a particular peer. If a peer is found not to support these extensions, the MPBGP router drops the peering session.

22.2 MPBGP Syntax

Note: You must specify the AS and routerid at the top of your configuration file in order for BGP to work.

```
mpbgp ( on | off )
{
  [ clusterid host-id ; ]
  [ defaultmetric metric ; ]
  [ discard-nonprefixed-confederations |
  ignore-nonprefixed-confederations ]
  [ open-on-accept ; ]
  [ preference mpbgppreference ; ]
  [ traceoptions trace_options ; ]
  group type
    ( external peeras autonomoussystem
    | internal peeras autonomoussystem
    | routing peeras autonomoussystem proto protocol )
```

```

[ ascount count ] # external only
[ comm community_values]
[ confed ]
[ gateway gateway_ip_address]
[ holdtime time ]
[ ignorefirstashop ] # external only
[ interface interface_list ] # routing only
[ keep ( all | none ) ]
[ keepalivesalways ]
[ localas autonomous_system ] # external only
[ localtcp local_address ]
[ localv4addr ipv4_address ]
[ localv6addr ipv6_address ]
[ logupdown ]
[ med ]
[ metricout metric ]
[ nexthopself ] # external only
[ noaggregatorid ]
[ nogendefault ]
[ nov4asloop ]
[ outdelay time ]
[ passive ]
[ preference grouppreference ]
[ preference2 grouppreference2 ]
[ recvbuffer buffer_size ]
[ reflector-client [ no-client-reflect ] ] # internal or
# routing only

[ remotev4addr ipv4_address ]
[ remotev6addr ipv6_address ]
[ routetopeer ]
[ sendbuffer buffer_size ]
[ setpref metric ] # internal or routing only
[ showwarnings ]
[ traceoptions trace_options ]
[ ttl t1 ] # routing only
[ caps { [v4u] | [v4m] | [v4um] | [v6u] | [v6m] | [v6um]
| [no-caps] } ]
#
# There can be zero or one "allow" clauses within
# a peer group.
#
{
[inet6] allow {
    all ;
    | host ipaddress ;
    | classful_network ;
    | network mask mask ;
    | network masklen masklennumber ;

```

```

        | network / masklength ;
    } ;
#
# There can be zero or more "peer" clauses within
# a peer group.
#
peer host
    [ ascount count ]
    [ gateway gateway ]
    [ holdtime time ]
    [ ignorefirstshop ]
    [ keep ( all | none ) ]
    [ keepalivesalways ]
    [ localtcp local_address ]
    [ localv4addr ipv4_address ]
    [ localv6addr ipv6_address ]
    [ logupdown ]
    [ med ]
    [ metricout metric ]
    [ nexthopself ]
    [ noagggregatorid ]
    [ nogendefault ]
    [ nov4asloop ]
    [ outdelay time ]
    [ passive ]
    [ preference peerpreference ]
    [ preference2 peerpreference2 ]
    [ rcvbuffer buffer_size ]
    [ remotev4addr ipv4_address ]
    [ remotev6addr ipv6_address ]
    [ routetopeer ]
    [ sendbuffer buffer_size ]
    [ showwarnings ]
    [ traceoptions trace_options ]
    [ ttl ttl ]
} ;
#
# There should be at least one "allow" or "peer"
# clause within a "group type" statement.
#
} ;
} ;

```

More detailed descriptions of these commands can be found on page 429 of the *Command Reference Guide*.

22.3 MPBGP Configurable Options

See the following sections for more information about specific MPBGP options.

22.3.1 Route Reflection

MPBGP supports route reflection for internal peer groups. When using route reflection, the rule that a router may not readvertise routes from internal peers to other internal peers is relaxed for some routers, called “route reflectors.” See “Route Reflection Overview and Examples” on page 71 for more information about route reflection.

22.3.2 Weighted Route Damping

The basic idea of weighted route damping is to treat routes that are being announced and withdrawn (flapping) at a rapid rate as unreachable. See “Weighted Route Damping Overview, Syntax, and Defaults” on page 73 for more information about weighted route damping.

22.3.3 Setpref/Local_Pref

The `setpref` option allows GateD to set the `Local_Pref` to reflect GateD’s own internal preference for the route, as given by the global protocol preference value. `Local_Pref` may be used by a BGP speaker to inform other BGP speakers in its own autonomous system of the originating speaker’s degree of preference for an advertised route. See “Setpref/Local_Pref Overview” on page 75 for more information about `setpref`.

22.3.4 Communities and Extended Communities

The communities attribute allows the administrator of a routing domain to tag groups of routes with a community tag. Using communities allows the administrator to limit the routes that can be imported or exported. See “Communities Overview and Examples” on page 77 for more information about communities.

22.3.5 Multi-Exit Discriminator

The Multi-Exit Discriminator, or MED, allows the administrator of a routing domain to choose between various exits from a peering autonomous system (AS). See “Multi-Exit Discriminator Overview and Examples” on page 79 for more information about Multi-Exit Discriminator.

22.3.6 Confederations

A BGP Confederation is a collection of ASs that present themselves as a single AS to peers outside of the confederation. See “Confederations” on page 80 for more information about Confederations.