Multiprotocol BGP (MP-BGP)

- BGP was originally designed to support IPv4 routing (before IPv6 existed).
- It has now been extended to MultiProtocol BGP (MP-BGP) which supports other protocols such as IPv6 and MPLS VPN routing.

Multiprotocol BGP (MP-BGP)

- MP-BGP is an extension to standard BGP, not a separate routing protocol version (such as OSPFv3 or RIPng).
- Only one instance of BGP can be configured on a router.
- Multiprotocol BGP (MP-BGP) supports IPv4, IPv6, unicast and multicast, MPLS VPN etc. through the use of AFI Address Family Identifiers (eg IPv4 or IPv6) and SAFI Subsequent Address Family Identifiers (eg unicast or multicast).

Address Families

```
R1(config)# router bgp 65001
R1(config-router)# address-family ipv4 unicast        (AFI 1, SAFI 1)
R1(config-router)# address-family ipv4 multicast      (AFI 1, SAFI 2)
R1(config-router)# address-family ipv6 unicast        (AFI 2, SAFI 1)
R1(config-router)# address-family ipv6 multicast      (AFI 2, SAFI 2)
```

The default is unicast. For example if you enter:

```
R1(config)# router bgp 65001
R1(config-router)# address-family ipv4
```

Then the router enters ipv4 unicast configuration mode

Address Families

- There is a separate configuration and BGP table for each protocol (address family and sub-address family pair)
- Capabilities are exchanged during the initial open message negotiation
- The AFI and SAFI is included in route advertisements
- Routing policy can be different for each protocol

Address Families

- All protocols can operate over a single BGP session
- Best practice is to use IPv4 sessions to exchange IPv4 routes, and IPv6 sessions to exchange IPv6 routes
  - This avoids possible Next Hop address issues
  - If you are using IPv4 sessions to exchange both IPv4 and IPv6 routes, if there is a problem with your IPv4 network that will affect IPv6 routing also

Neighbor Address Family Activation

```
router bgp 65001
neighbor 2001:DB8:0:1::1 remote-as 65002
neighbor 203.0.113.1 remote-as 65002
!
 address-family ipv4 (used for IPv4 route exchange)
  neighbor 203.0.113.1 activate
  no neighbor 2001:DB8:0:1::1 activate
!
 address-family ipv6 (used for IPv6 route exchange)
  neighbor 2001:DB8:0:1::1 activate
  no neighbor 203.0.113.1 activate
```

Neighbor Address Family Activation

- To prevent new neighbors with IPv4 addresses from being automatically enabled for the IPv4 address family, use

  ```
  no bgp default ipv4-unicast
  ```

- Best practice is to manually activate/deactivate neighbors

Router ID

- Only one instance of BGP can be configured on a router.
- The BGP router ID is 32-bit value
- By default it is set to the highest IPv4 address of a loopback interface
- If no loopback interface is configured then it is set to the highest IPv4 address of a physical interface
- If the router has only IPv6 addresses (no IPv4) then you must manually set the BGP Router ID

```
R1(config)# router bgp 65011
R1(config-router)# bgp router-id 192.168.0.1
```

Neighbor Addresses

- Neighbor addresses must be configured in BGP, unlike IGPs which discover neighbors through multicast
- Best practice is to use Global Unicast addresses for IPv6 BGP sessions
- Link Local IPv6 addresses can change if stateless address autoconfiguration is used

Example 1: BGP IPv4 Only Configuration

- The default address family if one is not entered is IPv4 unicast


```
R1(config)# router bgp 65011
R1(config-router)# network 192.0.2.0
R1(config-router)# neighbor 192.168.0.2 remote-as 65011
R1(config-router)# neighbor 192.168.0.2 update-source Loopback0
R1(config-router)# neighbor 192.168.0.2 next-hop-self
R1(config-router)# neighbor 192.168.0.2 send-community
R1(config-router)# neighbor 198.51.100.1 remote-as 65001
R1(config-router)# neighbor 198.51.100.1 route-map NO_TRANSIT in
R1(config-router)# exit
R1(config)# route-map NO_TRANSIT permit 10
R1(config-route-map)# set community no-export
```

Example 1: BGP IPv4 Only Configuration


```
R1#sh run | begin bgp
!
router bgp 65011
 bgp log-neighbor-changes
 network 192.0.2.0
 neighbor 192.168.0.2 remote-as 65011
 neighbor 192.168.0.2 update-source Loopback0
 neighbor 192.168.0.2 next-hop-self
 neighbor 192.168.0.2 send-community
 neighbor 198.51.100.1 remote-as 65001
 neighbor 198.51.100.1 route-map NO_TRANSIT in
!
route-map NO_TRANSIT permit 10
 set community no-export
```

Example 2: BGP IPv6 Only Configuration


```
R1(config)# router bgp 65011
R1(config-router)# bgp router-id 192.168.0.1
R1(config-router)# neighbor 2001:DB8:0:3::2 remote-as 65011
R1(config-router)# neighbor 2001:DB8:0:4::1 remote-as 65001

R1(config-router)# address-family ipv6
R1(config-router-af)# network 2001:DB8:0:1::/64
R1(config-router-af)# neighbor 2001:DB8:0:3::2 activate
R1(config-router-af)# neighbor 2001:DB8:0:3::2 update-source Loopback0
R1(config-router-af)# neighbor 2001:DB8:0:3::2 send-community
R1(config-router-af)# neighbor 2001:DB8:0:3::2 next-hop-self
R1(config-router-af)# neighbor 2001:DB8:0:4::1 activate
R1(config-router-af)# neighbor 2001:DB8:0:4::1 route-map NO_TRANSIT in

R1(config)# route-map NO_TRANSIT permit 10
R1(config-route-map)# set community no-export
```

Example 2: BGP IPv6 Only Configuration


```
R1#sh run | begin bgp
!
router bgp 65011
 bgp router-id 192.168.0.1
 bgp log-neighbor-changes
 neighbor 2001:DB8:0:3::2 remote-as 65011
 neighbor 2001:DB8:0:3::2 update-source Loopback0
 neighbor 2001:DB8:0:4::1 remote-as 65001
 !
 address-family ipv4
  no neighbor 2001:DB8:0:3::2 activate
  no neighbor 2001:DB8:0:4::1 activate
 exit-address-family
 !
 address-family ipv6
  network 2001:DB8:0:1::/64
  neighbor 2001:DB8:0:3::2 activate
  neighbor 2001:DB8:0:3::2 send-community
  neighbor 2001:DB8:0:3::2 next-hop-self
  neighbor 2001:DB8:0:4::1 activate
  neighbor 2001:DB8:0:4::1 route-map NO_TRANSIT in
 exit-address-family
!
route-map NO_TRANSIT permit 10
 set community no-export
```

Example 3: Merged Configuration

- IPv4 and IPv6, router starts with IPv4 BGP configuration only


```
interface Loopback0
 ip address 192.168.0.1 255.255.255.255
 ipv6 address 2001:DB8:0:100::1/128
!
interface GigabitEthernet0/0
 ip address 198.51.100.2 255.255.255.252
 ipv6 address 2001:DB8:0:4::2/64
 no shutdown
!
interface GigabitEthernet0/1
 ip address 192.0.2.1 255.255.255.0
 ipv6 address 2001:DB8:0:1::1/64
 no shutdown
```

Example 3: Merged Configuration

- Router starts with IPv4 configuration only:


```
R1#sh run | begin bgp
!
router bgp 65011
 bgp log-neighbor-changes
 network 192.0.2.0
 neighbor 192.168.0.2 remote-as 65011
 neighbor 192.168.0.2 update-source Loopback0
 neighbor 192.168.0.2 next-hop-self
 neighbor 192.168.0.2 send-community
 neighbor 198.51.100.1 remote-as 65001
 neighbor 198.51.100.1 route-map NO_TRANSIT in
!
route-map NO_TRANSIT permit 10
 set community no-export
```

Example 3: Merged Configuration

- IPv6 configuration is added (no policy):

FLACKBOX
www.flackbox.com

```
R1(config)#router bgp 65011
R1(config-router)# neighbor 2001:DB8:0:3::2 remote-as 65011
R1(config-router)# neighbor 2001:DB8:0:4::1 remote-as 65001

R1(config-router)# address-family ipv6
R1(config-router-af)# network 2001:DB8:0:1::/64
R1(config-router-af)# neighbor 2001:DB8:0:3::2 activate
R1(config-router-af)# neighbor 2001:DB8:0:4::1 activate
```

Example 3: Merged Configuration

```
R1#sh run | begin bgp
router bgp 65011
 bgp log-neighbor-changes
 neighbor 2001:DB8:0:3::2 remote-as 65011
 neighbor 2001:DB8:0:4::1 remote-as 65001
 neighbor 192.168.0.2 remote-as 65011
 neighbor 192.168.0.2 update-source Loopback0
 neighbor 198.51.100.1 remote-as 65001
 !
 address-family ipv4
  network 192.0.2.0
  no neighbor 2001:DB8:0:3::2 activate
  no neighbor 2001:DB8:0:4::1 activate
  neighbor 192.168.0.2 activate
  neighbor 192.168.0.2 send-community
  neighbor 192.168.0.2 next-hop-self
  neighbor 198.51.100.1 activate
  neighbor 198.51.100.1 route-map NO_TRANSIT in
 exit-address-family
 !
 address-family ipv6
  network 2001:DB8:0:1::/64
  neighbor 2001:DB8:0:3::2 activate
  neighbor 2001:DB8:0:4::1 activate
 exit-address-family
```

Example 3: Merged Configuration

- IPv6 configuration is added (with policy):

FLACKBOX
www.flackbox.com

```
R1(config)#router bgp 65011
R1(config-router)# neighbor 2001:DB8:0:3::2 remote-as 65011
R1(config-router)# neighbor 2001:DB8:0:4::1 remote-as 65001

R1(config-router)# address-family ipv6
R1(config-router-af)# network 2001:DB8:0:1::/64
R1(config-router-af)# neighbor 2001:DB8:0:3::2 activate
R1(config-router-af)# neighbor 2001:DB8:0:3::2 update-source Loopback0
R1(config-router-af)# neighbor 2001:DB8:0:3::2 send-community
R1(config-router-af)# neighbor 2001:DB8:0:3::2 next-hop-self
R1(config-router-af)# neighbor 2001:DB8:0:4::1 activate
R1(config-router-af)# neighbor 2001:DB8:0:4::1 route-map NO_TRANSIT in
```

Example 3: Merged Configuration


```
R1#sh run | begin bgp
router bgp 65011
 bgp log-neighbor-changes
 neighbor 2001:DB8:0:3::2 remote-as 65011
 neighbor 2001:DB8:0:3::2 update-source Loopback0
 neighbor 2001:DB8:0:4::1 remote-as 65001
 neighbor 192.168.0.2 remote-as 65011
 neighbor 192.168.0.2 update-source Loopback0
 neighbor 198.51.100.1 remote-as 65001
 !
 address-family ipv4
  network 192.0.2.0
  no neighbor 2001:DB8:0:3::2 activate
  no neighbor 2001:DB8:0:4::1 activate
  neighbor 192.168.0.2 activate
  neighbor 192.168.0.2 send-community
  neighbor 192.168.0.2 next-hop-self
  neighbor 198.51.100.1 activate
  neighbor 198.51.100.1 route-map NO_TRANSIT in
 exit-address-family
 !
 address-family ipv6
  network 2001:DB8:0:1::/64
  neighbor 2001:DB8:0:3::2 activate
  neighbor 2001:DB8:0:3::2 send-community
  neighbor 2001:DB8:0:3::2 next-hop-self
  neighbor 2001:DB8:0:4::1 activate
  neighbor 2001:DB8:0:4::1 route-map NO_TRANSIT in
 exit-address-family
!
route-map NO_TRANSIT permit 10
 set community no-export
```

Default IPv4 Address Family

- The default address family if one is not entered is IPv4 unicast
- Ensure you are in the correct address family when entering commands

```
R1(config)#router bgp 65011
R1(config-router)#neighbor 2001:DB8:0:4::1 route-map NO_TRANSIT in

R1#sh run | section bgp
router bgp 65011
neighbor 2001:DB8:0:4::1 remote-as 65001
 !
 address-family ipv4
  neighbor 2001:DB8:0:4::1 activate
  neighbor 2001:DB8:0:4::1 route-map NO_TRANSIT in
 exit-address-family
 !
 address-family ipv6
  neighbor 2001:DB8:0:4::1 activate
 exit-address-family
```

Link Local Address Mapping

- In the IPv6 BGP table, next hop addresses pointing to Ethernet Global Unicast addresses are automatically mapped to that interface's Link Local address

```
SP1#show bgp ipv6 unicast 2001:DB8:0:1::/64
BGP routing table entry for 2001:DB8:0:1::/64, version 2
Paths: (2 available, best #2, table default)
Flag: 0x100
  Advertised to update-groups:
     1           2
  Refresh Epoch 2
  65002 65011
    2001:DB8:0:6::2 (FE80::5054:FF:FE1F:2349) from 2001:DB8:0:6::2 (192.168.0.6)
      Origin IGP, localpref 100, valid, external
      rx pathid: 0, tx pathid: 0
  Refresh Epoch 1
  65011
    2001:DB8:0:4::2 (FE80::5054:FF:FE0D:8BE3) from 2001:DB8:0:4::2 (192.168.0.1)
      Origin IGP, metric 0, localpref 100, valid, external, best
      rx pathid: 0, tx pathid: 0x0
```

Link Local Address Mapping

- Link Local Addresses are used in the routing table for Ethernet next hops

```
SP1#sh ipv6 route bgp
IPv6 Routing Table - default - 12 entries
Codes: C - Connected, L - Local, S - Static, U - Per-user Static route
       B - BGP, HA - Home Agent, MR - Mobile Router, R - RIP
       H - NHRP, I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       IS - ISIS summary, D - EIGRP, EX - EIGRP external, NM - NEMO
       ND - ND Default, NDp - ND Prefix, DCE - Destination, NDr - Redirect
       RL - RPL, O - OSPF Intra, OI - OSPF Inter, OE1 - OSPF ext 1
       OE2 - OSPF ext 2, ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
       la - LISP alt, lr - LISP site-registrations, ld - LISP dyn-eid
       lA - LISP away, a - Application
B    2001:DB8:0:1::/64 [20/0]
       via FE80::5054:FF:FE0D:8BE3, GigabitEthernet0/0
```

AFI-SAFI Commands

- You can use the new AFI-SAFI version of show commands
- Deprecated commands have no context-sensitive help in the CLI

Deprecated	Traditional	New AFI-SAFI Format
show bgp	show ip bgp	show bgp ipv4 unicast
		show bgp ipv6 unicast
show bgp summary	show ip bgp summary	show bgp ipv4 unicast summary
show bgp neighbors	show ip bgp neighbors	show bgp ipv4 unicast neighbors

Verification – show bgp ipv6 unicast summary

```
R1#show bgp ipv6 unicast summary
BGP router identifier 192.168.0.1, local AS number 65011
BGP table version is 1, main routing table version 1

Neighbor          V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2001:DB8:0:4::1   4        65001       7       6        1    0    0 00:02:14        0
2001:DB8:0:100::2 4        65011      18      18        1    0    0 00:13:13        0
```

Verification – show bgp ipv6 unicast neighbors

```
R1#show bgp ipv6 unicast neighbors
BGP neighbor is 2001:DB8:0:4::1,  remote AS 65001, external link
  Description: SP1
  BGP version 4, remote router ID 192.168.0.5
  BGP state = Established, up for 00:07:25
  Last read 00:00:14, last write 00:00:56, hold time is 180, keepalive interval is 60 seconds
  Neighbor sessions:
    1 active, is not multisession capable (disabled)
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv6 Unicast: advertised and received
    Enhanced Refresh Capability: advertised and received
    Multisession Capability:
    Stateful switchover support enabled: NO for session 1
  Message statistics:
    InQ depth is 0
    OutQ depth is 0

                         Sent       Rcvd
    Opens:                  1          1
    Notifications:          0          0
    Updates:                1          1
    Keepalives:             9          9
```

Verification – show bgp ipv6 unicast neighbors - Piped

```
R1#show bgp ipv6 unicast neighbors | include (BGP neighbor is|Description|BGP state)

BGP neighbor is 2001:DB8:0:4::1,  remote AS 65001, external link
 Description: SP1
  BGP state = Established, up for 00:09:43

BGP neighbor is 2001:DB8:0:100::2,  remote AS 65011, internal link
 Description: R2
  BGP state = Established, up for 00:20:42
```

Verification – show bgp ipv6 unicast

- The Next Hop is shown as :: when the BGP router is directly connected to the advertised network

```
R1#show bgp ipv6 unicast
Status codes: * valid, > best
! truncated

Network          Next Hop            Metric LocPrf Weight Path
*>   2001:DB8:0:0::/64
                      2001:DB8:0:4::1                        0 65001 65012 i
* i  2001:DB8:0:1::/64
                      2001:DB8:0:100::2
                                           0    100      0 i
 *>                   ::                    0         32768 i
```