MP−EBGP Configuration Example

Document ID: 113555

Contents

Introduction
Prerequisites
Requirements
Components Used
Conventions
Configure
Network Diagram
Configurations
Verify
Related Information

Introduction

This document provides information how to configure multiprotocol extended Border Gateway Protocol (MP−EBGP) in Cisco IOS routers. MP−BGP is an extended BGP that allows BGP to carry routing information for multiple network layer protocols IPv6, VPNv4, and others. MP−BGP allows you to have a unicast routing topology different from a multicast routing topology, which helps to control the network and resources.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

The configurations in this document are based on the Cisco 3700 series router that runs Cisco IOS® Software Release 12.4 (15)T 13.

Conventions

Refer to Cisco Technical Tips Conventions for more information on document conventions.

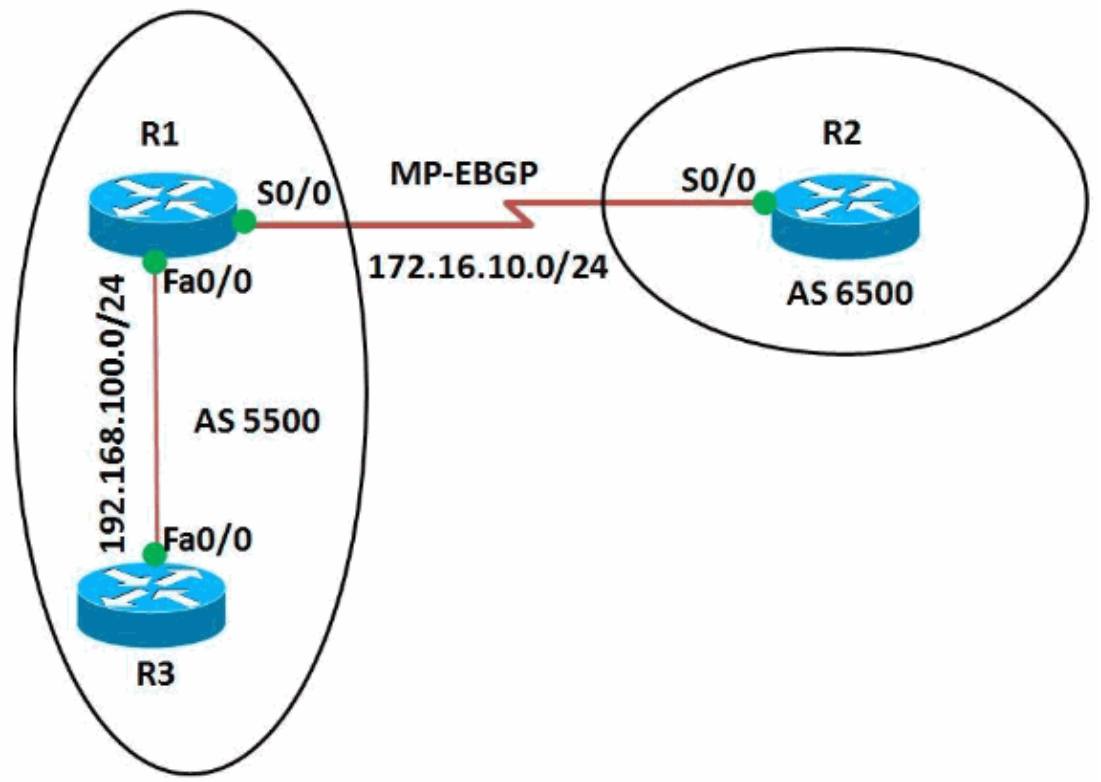
Configure

In this example, the R1 and R3 routers are configured to be in AS 5500 forming iBGP. The R2 router is configured to be in AS 6500. The R1 and R2 routers communicate with each other using MP−EBGP. All the routers are configured with loopback addresses.

Note: Use the Command Lookup Tool (registered customers only) to find more information on the commands used in this document.

Network Diagram

This document uses this network setup:



Configurations

This document uses these configurations:

- Router R1
- Router R2
- Router R3

Configuration on Router R1

```
R1#show run
Building configuration...
!
version 12.4
!
hostname R1
!
ip cef
!
!
interface Loopback0
 ip address 10.10.10.10 255.255.255.0
!
interface FastEthernet0/0
 ip address 192.168.100.10 255.255.255.0
 duplex auto
 speed auto
!
interface Serial0/0
 ip address 172.16.10.1 255.255.255.0
```

```
 mpls ip
 clock rate 2000000
!
router bgp 5500
 no synchronization
 bgp router-id 10.10.10.10
 bgp log-neighbor-changes
 network 192.168.100.0
 redistribute connected
 neighbor 172.16.10.2 remote-as 6500
 neighbor 172.16.10.2 soft-reconfiguration inbound
 neighbor 192.168.100.11 remote-as 5500
 no auto-summary
 !
 address-family vpnv4
  neighbor 172.16.10.2 activate
  neighbor 172.16.10.2 send-community both

!--- Sends the community attribute to a BGP neighbor.

 exit-address-family
!
!
end
```

Configuration on Router R2

```
R2#show run
Building configuration...
!
version 12.4
!
hostname R2
!
ip cef
!
ip vrf WAN
 rd 2020:1
 route-target export 2020:1
 route-target import 2020:1
!
!
interface Loopback0
 ip vrf forwarding WAN

!--- Associates a VRF instance with an interface or subinterface.

 ip address 20.20.20.20 255.255.255.255
!
interface Serial0/0
 ip vrf forwarding WAN
 ip address 172.16.10.2 255.255.255.0
 mpls ip
 clock rate 2000000
!
router bgp 6500
 no synchronization
 bgp router-id 20.20.20.20
 bgp log-neighbor-changes
 neighbor 172.16.10.1 remote-as 5500
 no auto-summary
 !
 !
 address-family vpnv4
  neighbor 172.16.10.1 activate
```

```
  neighbor 172.16.10.1 send-community both
 exit-address-family
 !
 address-family ipv4 vrf WAN
  redistribute connected
  redistribute static
  neighbor 172.16.10.1 remote-as 5500
  neighbor 172.16.10.1 activate
  no synchronization
 exit-address-family
!
!
!
end
```

Configuration on Router R3

```
R3#show run
Building configuration...
!
version 12.4
!
hostname R3
!
ip cef
!
!
!
interface Loopback0
 ip address 11.11.11.11 255.255.255.255
!
interface FastEthernet0/0
 ip address 192.168.100.11 255.255.255.0
 duplex auto
 speed auto
!
router bgp 5500
 no synchronization
 bgp router-id 11.11.11.11
 bgp log-neighbor-changes
 neighbor 192.168.100.10 remote-as 5500
 no auto-summary
!
end
```

Verify

In order to display entries in the (BGP) routing table, use the **show ip bgp** command.

show ip bgp

In router R1

```
R1#show ip bgp 172.16.10.2
BGP routing table entry for 172.16.10.2/32, version 14
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
        1    2
  Local
    0.0.0.0 from 0.0.0.0 (10.10.10.10)
      Origin incomplete, metric 0, localpref 100, weight 32768, valid, sourced, best

!--- Displays the routing table entries for the host 172.16.10.2
```

```
R1#sh ip bgp 192.168.100.11
BGP routing table entry for 192.168.100.0/24, version 4
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
        1    2
  Local
    0.0.0.0 from 0.0.0.0 (10.10.10.10)
      Origin IGP, metric 0, localpref 100, weight 32768, valid, sourced, local, best

!--- Displays the entries for the host 192.168.100.11
```

In router R3

```
R3#sh ip bgp 192.168.100.10
BGP routing table entry for 192.168.100.0/24, version 4
Paths: (1 available, best #1, table Default-IP-Routing-Table, RIB-failure(17))
  Not advertised to any peer
  Local
    192.168.100.10 from 192.168.100.10 (10.10.10.10)
      Origin IGP, metric 0, localpref 100, valid, internal, best

!--- Displays the entries for the host 192.168.100.10
```

In router R2, use the **show ip bgp vpnv4** command to display (VPNv4) address information from the (BGP) table.

show ip bgp vpnv4

In router R2

```
R2#sh ip bgp vpnv4 vrf WAN
BGP table version is 24, local router ID is 20.20.20.20
Status codes: s suppressed, d damped, h history, * valid, > best, I - internal,
              r RIB-failure, S Stale
Origin codes: I - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
Route Distinguisher: 2020:1 (default for vrf WAN)
*> 10.10.10.0/24    172.16.10.1              0             0 5500 ?
*> 20.20.20.20/32   0.0.0.0                  0         32768 ?
*  172.16.10.0/24   172.16.10.1              0             0 5500 ?
*>                  0.0.0.0                  0         32768 ?
r> 172.16.10.2/32   172.16.10.1              0             0 5500 ?
*> 192.168.100.0    172.16.10.1              0             0 5500 I

!--- Displays prefixes associated with the (VRF) instance WAN.


R2#show ip bgp vpnv4 vrf WAN 172.16.10.1
BGP routing table entry for 2020:1:172.16.10.0/24, version 7
Paths: (2 available, best #2, table WAN)
  Advertised to update-groups:
        1
  5500
    172.16.10.1 from 172.16.10.1 (10.10.10.10)
      Origin incomplete, metric 0, localpref 100, valid, external
      Extended Community: RT:2020:1
      mpls labels in/out 18/nolabel
  Local
    0.0.0.0 from 0.0.0.0 (20.20.20.20)
```

```
      Origin incomplete, metric 0, localpref 100, weight 32768, valid, sourced, best
      Extended Community: RT:2020:1
      mpls labels in/out 18/aggregate(WAN)

!--- Displays prefixes associated with neighbor 172.16.10.1
```

MP−EBGP is established between the R1 and R2 routers. Use the **ping** command to verify the reachability from R1 to R2 and vice versa.

ping

In router R1

```
R1#ping 172.16.10.2

Type escape sequence to abort.
Sending 5, 100−byte ICMP Echos to 172.16.10.2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round−trip min/avg/max = 12/64/208 ms

R1#ping 192.168.100.11

Type escape sequence to abort.
Sending 5, 100−byte ICMP Echos to 192.168.100.11, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round−trip min/avg/max = 12/41/96 ms

!--- Router R1 can successfully ping the routers R2 and R3.
```

In router R2

```
R2#ping vrf WAN 172.16.10.1

Type escape sequence to abort.
Sending 5, 100−byte ICMP Echos to 172.16.10.1, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round−trip min/avg/max = 4/32/96 ms

R2#ping vrf WAN 192.168.100.11

Type escape sequence to abort.
Sending 5, 100−byte ICMP Echos to 192.168.100.11, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round−trip min/avg/max = 32/73/204 ms


!--- Router R2 can successfully reach router R1 and R3.
```

Related Information

- **Border Gateway Protocol (BGP)**
- **Multiprotocol BGP Extensions for IP Multicast Commands**
- **Technical Support & Documentation − Cisco Systems**

Contacts & Feedback | Help | Site Map
© 2014 − 2015 Cisco Systems, Inc. All rights reserved. Terms & Conditions | Privacy Statement | Cookie Policy | Trademarks of Cisco Systems, Inc.