


BGP Lab7 - Configuring BGP Weight

?Lab Objectives?

1. When the local export router has a number of outward self-government system link, it needs to apply the weight attributes which can decide the data flow to outflow of the local self-government system from which exit link of the local route.
2. Weight attribute is the Cisco private property.

?Lab Topology?



?Lab Steps?

1. Configure the router's IP address, and use the Ping command to confirm the connect's interoperability of each router.
2. To configure the BGP protocols for each router and use show ip bgp to confirm that the protocols will work properly. We suggest to configure the BGP with the order as R1?R2?R3?R4?so to facilitate to carry out the lab. The main reason for emphasizing the configuration orders is because that when R1 reaches the two paths of network 150.150.1.0/24 in this lab, the BGP protocols will take for the two paths as the equivalent, but BGP is not used to sharing the load lab, so, BGP will prefer to choose from the R1 to reach the goal.

The reason is that BGP choosing route with rules: 1, The first received route is prior to the later one. 2, Prefer to choose the BGP's route with lower Router-ID.

3. Check the routing table of R1

R1#show ip route

Gateway of last resort is not set

C 192.168.4.0/24 is directly connected, FastEthernet0/0

C 192.168.1.0/24 is directly connected, Serial1/1

B 192.168.2.0/24 [20/0] via 192.168.1.2, 00:16:39

B 192.168.3.0/24 [20/0] via 192.168.4.1, 00:16:25

150.150.0.0/24 is subnetted, 1 subnets

B 150.150.1.0 [20/0] via 192.168.1.2, 00:16:39

4. Check BGP data table of R1:

```
R1#show ip bgp
BGP table version is 7, local router ID is 192.168.4.2
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
* 150.150.1.0/24 192.168.4.1 0 64515 64514 i
*> 192.168.1.2 0 64513 64514 i
* 192.168.1.0 192.168.1.2 0 0 64513 i
*> 0.0.0.0 0 32768 i
* 192.168.2.0 192.168.4.1 0 64515 64514 i
*> 192.168.1.2 0 0 64513 i
*> 192.168.3.0 192.168.4.1 0 0 64515 i
* 192.168.1.2 0 64513 64514 i
* 192.168.4.0 192.168.4.1 0 0 64515 i
*> 0.0.0.0 0 32768 i
```

5. Check neighbors table information of R1:

```
R1#show ip bgp neighbors
BGP neighbor is 192.168.1.2, remote AS 64513, external link
BGP version 4, remote router ID 192.168.2.1
BGP state = Established, up for 02:10:56
```

???

```
BGP neighbor is 192.168.4.1, remote AS 64515, external link
BGP version 4, remote router ID 192.168.4.1
BGP state = Established, up for 02:09:04
```

6. As can be seen through the topology, the best route for self-government system 64512 to reach network 150.150.1.0/24 is going through self-government system 64514 rather than 64513. Therefore, BGP chosen a second-best route once again.

7. To configure the weight property on R1 for solving the problem on choosing second-best route. The configuration is shown as below:

```
R1(config)#router bgp 64512
R1(config-router)#neighbor 192.168.4.1 weight 100
R1(config-router)#neighbor 192.168.1.2 weight 50
R1(config-router)#exit
R1#clear ip bgp * soft
```

8. In order to confirm the effect that the different weight values make on the route, to check R1 routing tables and BGP databases:

```
R1#show ip bgp
```

BGP table version is 13, local router ID is 192.168.4.2
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
*> 150.150.1.0/24 192.168.4.1 100 64515 64514 i
* 192.168.1.2 50 64513 64514 i
* 192.168.1.0 192.168.1.2 0 50 64513 i
*> 0.0.0.0 0 32768 i
*> 192.168.2.0 192.168.4.1 100 64515 64514 i
* 192.168.1.2 0 50 64513 i
*> 192.168.3.0 192.168.4.1 0 100 64515 i
* 192.168.1.2 50 64513 64514 i
* 192.168.4.0 192.168.4.1 0 100 64515 i
*> 0.0.0.0 0 32768 i
```

```
R1#show ip route
```

Codes: C - connected, S - static, R - RIP, M - mobile, B - BGP
D - EIGRP, EX - EIGRP external, 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, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
ia - IS-IS inter area, * - candidate default, U - per-user static route
o - ODR, P - periodic downloaded static route

Gateway of last resort is not set

C 192.168.4.0/24 is directly connected, FastEthernet0/0
C 192.168.1.0/24 is directly connected, Serial1/1
B 192.168.2.0/24 [20/0] via 192.168.4.1, 00:01:01
B 192.168.3.0/24 [20/0] via 192.168.4.1, 03:12:01
150.150.0.0/24 is subnetted, 1 subnets
B 150.150.1.0 [20/0] via 192.168.4.1, 00:01:01

9. By the analysis of the BGP database, it shows that the weight values of the route sent from route R2 are all 50, while the values are all 100 from R4. As a result of the rules for BGP to choose the route: prefer to the route with higher weight value. Therefore, when the data flow flows out of 64512 self-government system to local self-government system, they all use R4 as the next hop. However, there will be another problem, namely, it is not wise that R1 uses R4 to reach the subnet 192.168.2.0/24 of R2, because the best route from R4 to reach network 192.168.2.0/24 is sent directly to the R2. The configuration method to use neighbor *.*.* weight *** is not the perfect configuration at all. In order to solve this problem, we recommend to configure with route map, and to set different weight values for the target networks.

10. To configure the weight value with route-map on R1, the configuration is show as below:

```
R1(config)#router bgp 64512
R1(config-router)#no neighbor 192.168.1.2 weight 50
R1(config-router)#no neighbor 192.168.4.1 weight 100
R1(config-router)#exit
R1(config)#
R1(config)#access-list 1 permit 150.150.1.0 0.0.0.255
R1(config)#
R1(config)#route-map set_weight permit 10
R1(config-route-map)#match ip address 1
R1(config-route-map)#set weight 100
R1(config-route-map)#exit
R1(config)#
R1(config)#
R1(config)#route-map set_weight permit 20
R1(config-route-map)#exit
R1(config)#
R1(config)#router bgp 64512
R1(config-router)#neighbor 192.168.4.1 route-map set_weight in
R1(config-router)#exit
R1(config)#exit
R1#clear ip bgp * soft
```

11. Check the routing table of R1:

```
R1#show ip route
```

Gateway of last resort is not set

```
C 192.168.4.0/24 is directly connected, FastEthernet0/0
C 192.168.1.0/24 is directly connected, Serial1/1
B 192.168.2.0/24 [20/0] via 192.168.1.2, 00:03:45
B 192.168.3.0/24 [20/0] via 192.168.4.1, 00:03:45
 150.150.0.0/24 is subnetted, 1 subnets
B 150.150.1.0 [20/0] via 192.168.4.1, 00:00:00
```

12. Check BGP database of R1:

```
R1#show ip bgp
```

BGP table version is 21, local router ID is 192.168.4.2

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

```
*> 150.150.1.0/24 192.168.4.1 100 64515 64514 i
* 192.168.1.2 0 64513 64514 i
* 192.168.1.0 192.168.1.2 0 0 64513 i
*> 0.0.0.0 0 32768 i
* 192.168.2.0 192.168.4.1 0 64515 64514 i
*> 192.168.1.2 0 0 64513 i
*> 192.168.3.0 192.168.4.1 0 0 64515 i
* 192.168.1.2 0 64513 64514 i
* 192.168.4.0 192.168.4.1 0 0 64515 i
*> 0.0.0.0 0 32768 i
```

13. To confirm the validity of the route by using commands ping and traceroute:

R1#ping 150.150.1.1

Type escape sequence to abort.

Sending 5, 100-byte ICMP Echos to 150.150.1.1, timeout is 2 seconds:

!!!!

Success rate is 100 percent (5/5), round-trip min/avg/max = 68/126/216 ms

R1#

R1#traceroute 150.150.1.1

Type escape sequence to abort.

Tracing the route to 150.150.1.1

1 192.168.4.1 16 msec 60 msec 64 msec

2 192.168.3.1 [AS 64515] 144 msec * 32 msec

14. Lab completed.

Hope to helpful for you!