

BGP Lab Exercise 3

1. This exercise will use AS Path filters on the BGP sessions. These will ensure that we only announce our prefixes to our neighbours. Create an AS Path filter which permits only prefixes originated by your peer AS to enter your network

```
RouterA(config)# ip as-path access-list 1 permit ^2$
! permit prefixes originated by AS2
```

2. Create an AS Path filter which permits only prefixes originated by your AS to leave your network

```
RouterA(config)# ip as-path access-list 2 permit ^$
! permit prefixes originated by local AS
```

3. Enable soft reconfiguration for the BGP session. This is only so that we can analyse the effect the applying the filter on our BGP session later on in this exercise. The use of soft reconfiguration in general ISP operations is not recommended apart from during session debugging (as it requires extra memory and has been superseded by a BGP feature called "route refresh").

```
RouterA(config)# router bgp 1
RouterA(config-router)# address-family ipv4
RouterA(config-router-af)# neighbor 196.200.220.xxx soft-reconfiguration in
```

```
RouterA(config-router)# address-family ipv6
RouterA(config-router-af)# neighbor 2001:4348:220:xx::x soft-reconfig in
```

4. Now apply filter-list 1 inbound on the eBGP session you have with your peer AS.

```
RouterA(config-router)# address-family ipv4
RouterA(config-router-af)# neighbor 196.200.220.xxx filter-list 1 in
! apply as-path filter 1 inbound
```

And apply filter-list 2 to outbound on the eBGP session you have with both AS100 and your peer AS.

```
RouterA(config-router-af)#neighbor 196.200.220.xxx filter-list 2 out
RouterA(config-router-af)#neighbor 196.200.222.1 filter-list 2 out
! apply as-path filter 2 outbound
```

Do the same for the IPv6 peerings. Remember you want to make the changes to the IPv6 address family and to the IPv6 neighbours. The filter-list can be used for both IPv4 and IPv6 peerings.

5. Now that you have applied the policy, you need to tell the BGP

session to refresh the prefixes it has received from the neighbour.
To do this, we use the route refresh capability of BGP.

```
routerA#clear ip bgp 2 in      ! refresh the peering with AS2
routerA#clear ip bgp 2 out    ! refresh the peering with AS2
routerA#clear ip bgp 100 out  ! refresh the peering with AS100
```

6. Do the same for IPv6.

```
routerA#clear bgp ipv6 unicast 2 in  ! refresh the peering with AS2
routerA#clear bgp ipv6 unicast 2 out ! refresh the peering with AS2
routerA#clear bgp ipv6 unicast 100 out ! refresh the peering with AS100
```

7. Check the BGP process using "sh ip bgp" and "sh bgp ipv6" to see what prefixes you are now receiving from your neighbouring AS. Use the more specific commands "sh ip bgp neighbor 196.200.220.xxx routes" and "sh ip bgp neighbor 196.200.220.xxx received-routes". What difference do you see? Write your answer here:

What about for IPv6? Write the results of the IPv6 version of the above commands here: