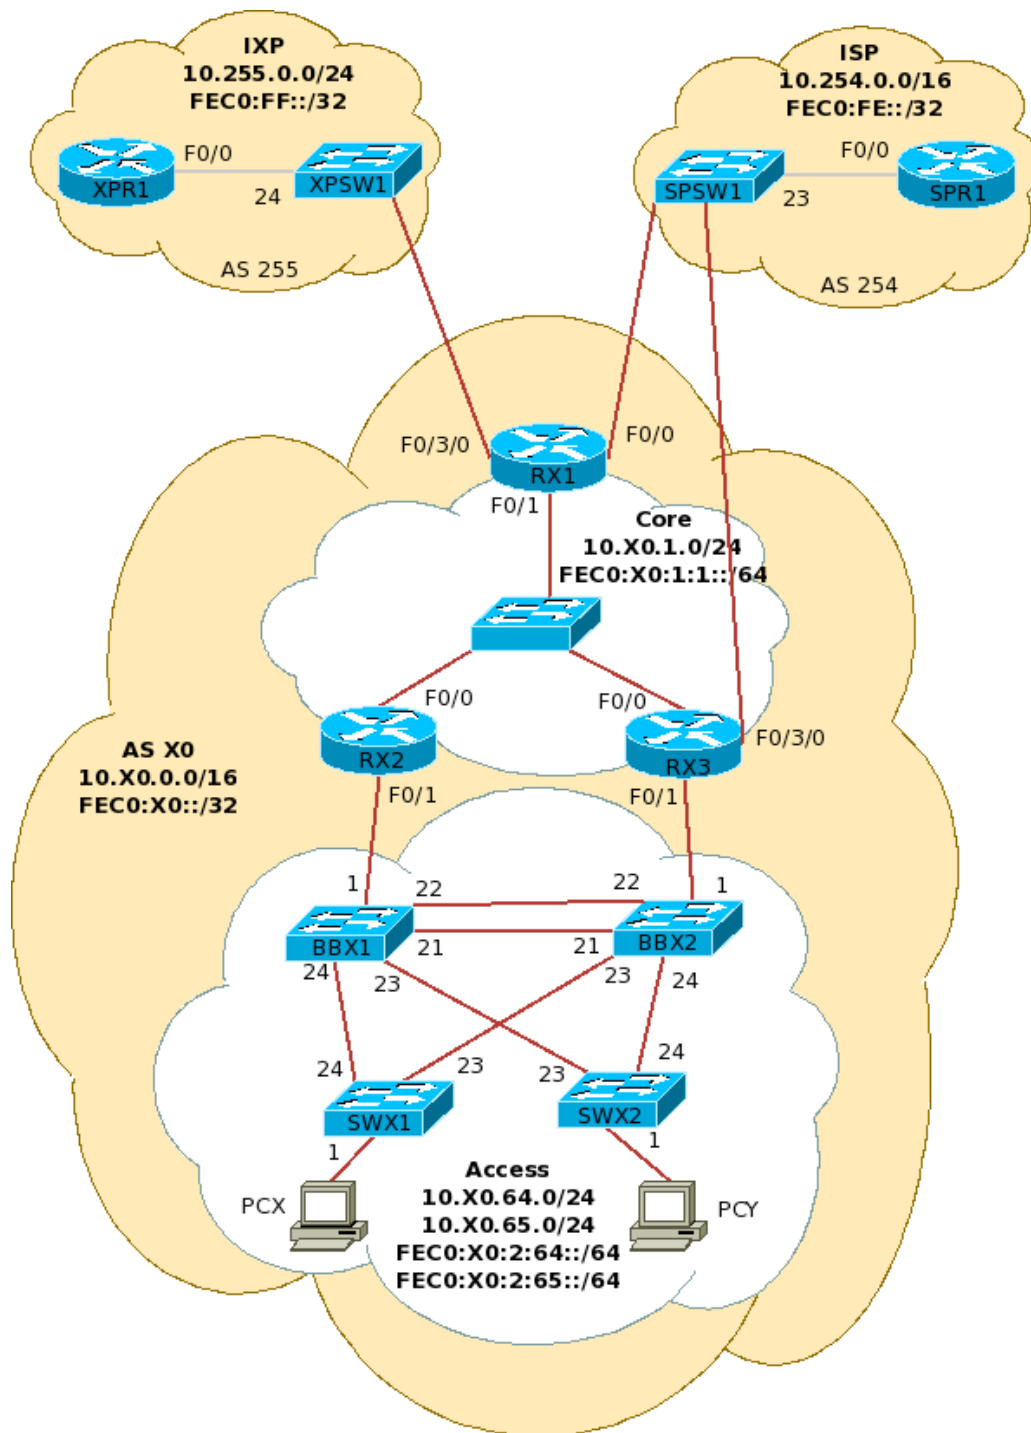


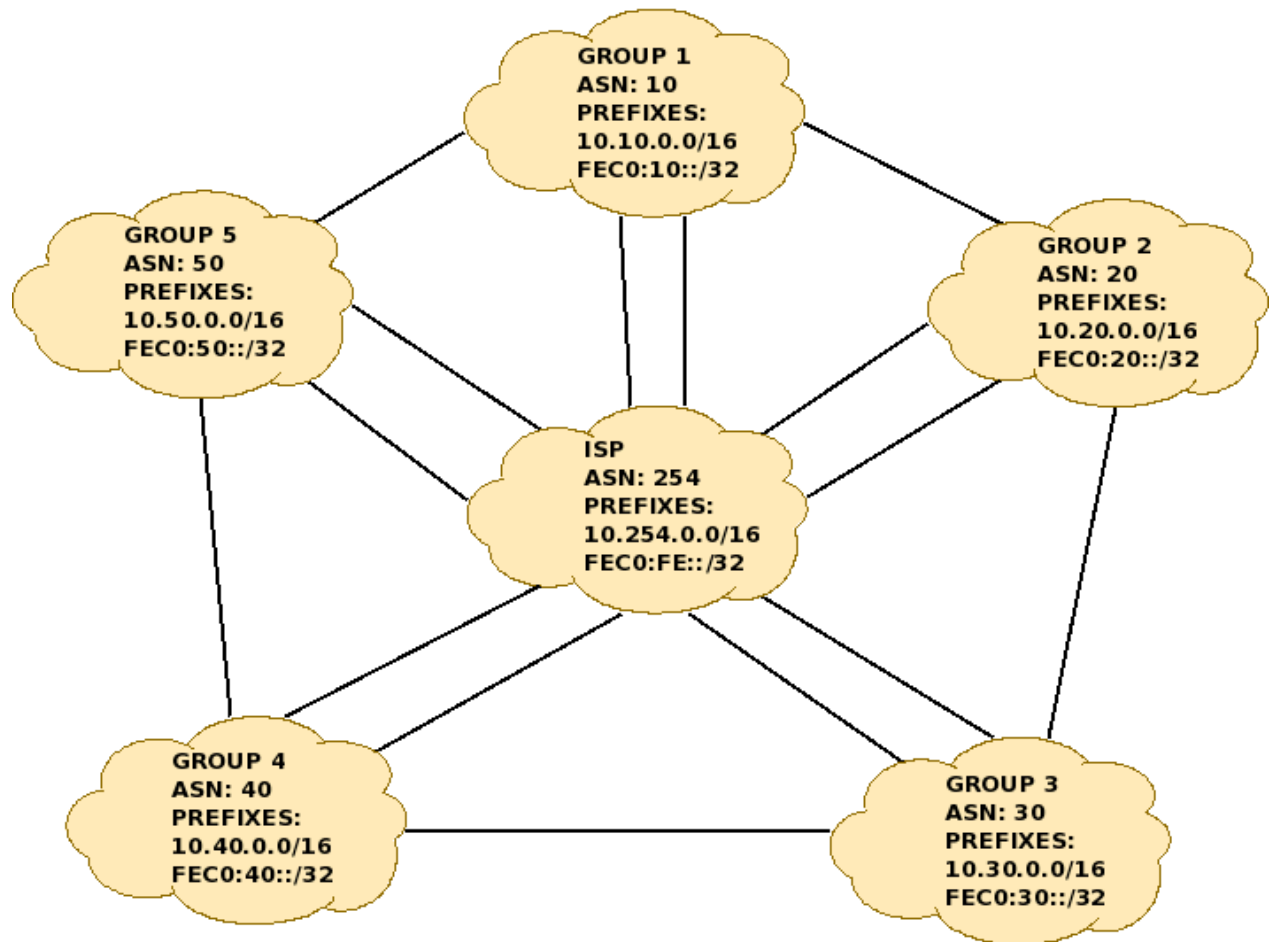
BGP Exercises – Implementing Policy

In these exercises we will use several BGP attributes to achieve desired traffic policies. Use the physical and logical network diagrams to help you understand the configurations and the traffic flows.



Each AS will now have a second connection to the ISP via RX3.

Logical AS Topology 3



The address space for the secondary ISP links is

- 10.254.1X0.0/30
- FEC0:FE:0:1X0::/64

where X = your ASN

1. Configure your second interface to the ISP.

R13:

```
interface FastEthernet0/3/0
description Link to ISP
ip address 10.254.110.2 255.255.255.252
ipv6 enable
ipv6 address FEC0:FE:0:110::2/64
no ip redirects
no ip directed-broadcast
no ip proxy-arp
no shutdown
```

ISP:

```
interface FastEthernet0/0.110
description Link to AS10
encapsulation dot1Q 110
ip address 10.254.110.1 255.255.255.252
ipv6 address FEC0:FE:0:110::1/64
ipv6 enable
no ip redirects
no ip directed-broadcast
no ip proxy-arp
no shutdown
!
```

2. STOP -- Checkpoint.

Make sure that you can ping the remote side before continuing.

```
ping 10.254.110.2
ping FEC0:FE:0:110:2
```

3. Configure the second eBGP session to the ISP.

R13:

```
router bgp 10
neighbor 10.254.110.1 remote-as 254
neighbor 10.254.110.1 description eBGP to ISP
neighbor 10.254.110.1 password NSRC
neighbor FEC0:FE:0:110::1 remote-as 254
neighbor FEC0:FE:0:110::1 description eBGP to ISP
neighbor FEC0:FE:0:110::1 password NSRC
aggregate-address 10.10.0.0 255.255.0.0 summary-only
!
address-family ipv4
no neighbor FEC0:FE:0:110::1 activate
!
address-family ipv6
neighbor FEC0:FE:0:110::1 activate
aggregate-address FEC0:10::/32 summary-only
```

ISP:

```
router bgp 254
 neighbor 10.254.110.2 remote-as 10
 neighbor 10.254.110.2 description eBGP to AS10
 neighbor 10.254.110.2 password NSRC
 neighbor FEC0:FE:0:110::2 remote-as 10
 neighbor FEC0:FE:0:110::2 description eBGP a AS10
 neighbor FEC0:FE:0:110::2 password NSRC
!
 address-family ipv4
  no neighbor FEC0:FE:0:110::2 activate
!
 address-family ipv6
  neighbor FEC0:FE:0:110::2 activate
```

4. Add the point-to-point network to your IGP

R13:

```
router ospf 100
 passive-interface FastEthernet0/3/0
!
ipv6 router ospf 100
 passive-interface FastEthernet0/3/0
exit
!
interface FastEthernet0/3/0
 ip ospf 100 area 0
 ipv6 ospf 100 area 0
```

Explain what happens if you don't do this.

5. STOP -- Checkpoint.

```
show ip bgp summary
show ip bgp neighbors 10.254.110.1
show ip bgp neighbors 10.254.110.1 advertised-routes
show ip bgp neighbors 10.254.110.1 routes
show ip bgp
show bgp ipv6 unicast summary
show bgp ipv6 unicast neighbors FEC0:FE:0:110::1
show bgp ipv6 unicast neighbors FEC0:FE:0:110::1 advertised-routes
show bgp ipv6 unicast neighbors FEC0:FE:0:110::1 routes
show bgp ipv6 unicast
```

6. Create prefix-lists and input and output policies.

Note: You may have already created these prefix-lists in previous exercises, but you still need to associate them with the new peerings below:

R13:

```
ip prefix-list out-peer permit 10.10.0.0/16 le 32
ip prefix-list isp-in-peer deny 10.10.0.0/16 le 32
```

```

ip prefix-list isp-in-peer permit 0.0.0.0/0 le 32
ipv6 prefix-list ipv6-out-peer permit FEC0:10::/32 le 128
ipv6 prefix-list ipv6-isp-in-peer deny FEC0:10::/32 le 128
ipv6 prefix-list ipv6-isp-in-peer permit ::/0 le 128

router bgp 10
 neighbor 10.254.110.1 prefix-list out-peer out
 neighbor 10.254.110.1 prefix-list isp-in-peer in
!
 address-family ipv6
  neighbor FEC0:FE:0:110::1 prefix-list ipv6-out-peer out
  neighbor FEC0:FE:0:110::1 prefix-list ipv6-isp-in-peer in

```

ISP:

```

ip prefix-list as10-in permit 10.10.0.0/16 le 32
ipv6 prefix-list ipv6-as10-in permit FEC0:10::/32 le 128

router bgp 254
 neighbor 10.254.110.2 prefix-list as10-in in
!
 address-family ipv6
  neighbor FEC0:FE:0:110::2 prefix-list ipv6-as10-in in

```

7. STOP -- Checkpoint.

Perform a soft reset of your BGP sessions:

```

R11# clear ip bgp * in
R11# clear ip bgp * out

R13# clear ip bgp * in
R13# clear ip bgp * out

show ip bgp summary
show ip bgp
show ip route
traceroute <ip_address>

```

You should now see that traffic leaving your AS towards non-peer ASs will traverse the ISP via one link or the other. For example, traffic from AS10 to AS30 and AS40 will either traverse R11 or R13. This will depend on your target, the point in your network from where you originate your traceroutes, and other factors like router-ids (lowest wins).

Local Preference (LOCAL_PREF)

Here are our traffic policy goals. We want traffic:

- From anywhere in AS10 to go to AS30 leaving out of R11
- From anywhere in AS10 to go to AS40 leaving out of R13
- From anywhere in AS20 to go to AS40 leaving out of R21
- From anywhere in AS20 to go to AS50 leaving out of R23
- From anywhere in AS30 to go to AS50 leaving out of R31
- From anywhere in AS30 to go to AS10 leaving out of R33
- From anywhere in AS40 to go to AS10 leaving out of R41
- From anywhere in AS40 to go to AS20 leaving out of R43
- From anywhere in AS50 to go to AS20 leaving out of R51
- From anywhere in AS50 to go to AS30 leaving out of R53

8. Use the Local Preference attribute to prefer an exit point depending on the destination AS. Use AS10 configurations below as an example.

R11:

```
ip prefix-list as30-prefix permit 10.30.0.0/16
ip prefix-list as40-prefix permit 10.40.0.0/16
!
route-map set-local-pref permit 10
  match ip address prefix-list as30-prefix
  set local-preference 150
route-map set-local-pref permit 20
  match ip address prefix-list as40-prefix
  set local-preference 50
route-map set-local-pref permit 30
!
ipv6 prefix-list ipv6-as30-prefix permit FEC0:30::/32
ipv6 prefix-list ipv6-as40-prefix permit FEC0:40::/32
!
route-map ipv6-set-local-pref permit 10
  match ipv6 address prefix-list ipv6-as30-prefix
  set local-preference 150
route-map ipv6-set-local-pref permit 20
  match ipv6 address prefix-list ipv6-as40-prefix
  set local-preference 50
route-map ipv6-set-local-pref permit 50
!
router bgp 10
  neighbor 10.254.10.1 route-map set-local-pref in
  address-family ipv6
    neighbor FEC0:FE:0:10::1 route-map ipv6-set-local-pref in
  !
```

R13:

```
ip prefix-list as30-prefix permit 10.30.0.0/16
ip prefix-list as40-prefix permit 10.40.0.0/16
!
route-map set-local-pref permit 10
  match ip address prefix-list as30-prefix
  set local-preference 50
route-map set-local-pref permit 20
  match ip address prefix-list as40-prefix
  set local-preference 150
route-map set-local-pref permit 30
!
ipv6 prefix-list ipv6-as30-prefix permit FEC0:30::/32
ipv6 prefix-list ipv6-as40-prefix permit FEC0:40::/32
!
route-map ipv6-set-local-pref permit 10
  match ipv6 address prefix-list ipv6-as30-prefix
  set local-preference 50
route-map ipv6-set-local-pref permit 20
  match ipv6 address prefix-list ipv6-as40-prefix
  set local-preference 150
route-map ipv6-set-local-pref permit 50
!

router bgp 10
  neighbor 10.254.110.1 route-map set-local-pref in
  address-family ipv6
    neighbor FEC0:FE:0:110:1 route-map ipv6-set-local-pref in
  !
```

Notice that the default LOCAL_PREF value is 100. Here we increase the preference on one side and lower it on the opposite side. Can you think of why?

9. STOP -- Checkpoint.

Perform a soft reset of your BGP sessions and verify paths.

On all AS10 routers:

```
clear ip bgp * in
clear ip bgp * out

show ip bgp 10.30.0.0
show bgp ipv6 unicast FEC0:30::/32
traceroute 10.30.254.1
traceroute 10.30.254.2
traceroute 10.30.254.3
traceroute FEC0:30:...
```

Repeat for AS40.

Multiple Exit Discriminator (MULTI_EXIT_DISC)

- 10. Remove the configurations from the previous step. This is the order in which it should be done (and from now on, in every case):**

```
router bgp 10
  no neighbor 10.254.10.1 route-map set-local-pref in
  address-family ipv6
    no neighbor FEC0:FE:0:10::1 route-map ipv6-set-local-pref in
  exit
no route-map set-local-pref
no ip prefix-list as30-prefix permit 10.30.0.0/16
end
```

- 11. perform a soft reset on all your BGP peerings**

```
clear ip bgp * in
clear ip bgp * out
```

The goal now is to instruct the ISP to send traffic into our AS always via RX1.

- 12. Ask your non-peers to do a traceroute towards different points in your AS. Keep note of the path that traffic is taking to reach your different points in AS from these other networks.**
- 13. Configure your border routers to set MEDs when advertising your prefix to the ISP:**

R11:

```
ip prefix-list as10-prefix permit 10.10.0.0/16
!
route-map set-med permit 10
  match ip address prefix-list as10-prefix
  set metric 10
route-map set-med permit 20
!
ipv6 prefix-list ipv6-as10-prefix permit FEC0:10::/32
!
route-map ipv6-set-med permit 10
  match ipv6 address prefix-list ipv6-as10-prefix
  set metric 10
route-map ipv6-set-med permit 20
!
router bgp 10
  neighbor 10.254.10.1 route-map set-med out
  address-family ipv6
    neighbor FEC0::FE:0:10 route-map ipv6-set-med out
  !
```

R13:

```
ip prefix-list as10-prefix permit 10.10.0.0/16
```



```

!
route-map set-med permit 10
  match ip address prefix-list as10-prefix
  set metric 50
route-map set-med permit 20
!
ipv6 prefix-list ipv6-as10-prefix permit FEC0:10::/32
!
route-map ipv6-set-med permit 10
  match ipv6 address prefix-list ipv6-as10-prefix
  set metric 50
route-map ipv6-set-med permit 20
!
router bgp 10
  neighbor 10.254.110.1 route-map set-med out
  address-family ipv6
    neighbor FEC0::FE:0:110:1 route-map ipv6-set-med out
!

```

14. STOP -- Checkpoint.

ISP:

```

show ip bgp 10.10.0.0
show ip route 10.10.0.0
show bgp ipv6 unicast FEC0:10::
show ip route FEC0:10::

```

Ask your non-peer ASs to repeat the traceroutes towards you and verify that traffic is entering your AS via router RX1.

BGP Communities (COMMUNITY)

15. Remove the configurations from the previous step and perform a soft reset on all your BGP peerings

```
clear ip bgp * in
clear ip bgp * out
```

The goal now is the same as in the previous step, but in this case, we set different communities on our prefix as it is announced to the ISP by our two border routers. The ISP will then apply different local preference attributes depending on the value of the community. Obviously, this has to be coordinated with our ISP.

16. Configure your border routers using this example.

R11:

```
ip bgp-community new-format
!
ip prefix-list as10-prefix permit 10.10.0.0/16
!
route-map set-comm permit 10
  match ip address prefix-list as10-prefix
  set community 10:150
route-map set-comm permit 20
!
ipv6 prefix-list ipv6-as10-prefix permit FEC0:10::/32
!
route-map ipv6-set-comm permit 10
  match ipv6 address prefix-list ipv6-as10-prefix
  set community 10:150
route-map ipv6-set-comm permit 20
!
router bgp 10
  neighbor 10.254.10.1 send-community
  neighbor 10.254.10.1 route-map set-comm out
  address-family ipv6
    neighbor FEC0:FE:0:10:1 send-community
    neighbor FEC0:FE:0:10:1 route-map ipv6-set-comm out
!
```

R13:

```
ip bgp-community new-format
!
ip prefix-list as10-prefix permit 10.10.0.0/16
!
route-map set-comm permit 10
  match ip address prefix-list as10-prefix
  set community 10:50
route-map set-comm permit 20
!
ipv6 prefix-list ipv6-as10-prefix permit FEC0:10::/32
!
```

```

route-map ipv6-set-comm permit 10
  match ipv6 address prefix-list ipv6-as10-prefix
  set community 10:50
route-map ipv6-set-comm permit 20
!
router bgp 10
  neighbor 10.254.110.1 send-community
  neighbor 10.254.110.1 route-map set-comm out
  address-family ipv6
    neighbor FEC0:FE:0:110:1 send-community
    neighbor FEC0:FE:0:110:1 route-map ipv6-set-comm out
  !

```

ISP:

```

ip bgp-community new-format
!
ip community-list 11 permit 10:150
ip community-list 12 permit 10:50
ip community-list 21 permit 20:150
ip community-list 22 permit 20:50

route-map customer-comm permit 10
  match community 11
  set local-preference 150
!
route-map customer-comm permit 20
  match community 12
  set local-preference 50
!
route-map customer-comm permit 30
!
router bgp 254
  neighbor 10.254.10.2 route-map customer-comm in
  neighbor 10.254.110.2 route-map customer-comm in
  address-family ipv6
    neighbor FEC0:FE:0:10:2 route-map customer-comm in
    neighbor FEC0:FE:0:110:2 route-map customer-comm in
  !

```

17. STOP -- Checkpoint.

ISP:

```

show ip bgp 10.10.0.0
show ip route 10.10.0.0

```

AS Path Prepending (AS_PATH)

18. Remove the configurations from the previous step and perform a soft reset on all your BGP peerings

```
clear ip bgp * in
clear ip bgp * out
```

We will now use a technique called *AS path prepending*, which consists of adding extra “fake” hops to a path using our ASN multiple times. The goal will be to force one of our peers to send traffic to us via the ISP, instead of over our direct peering via the IXP. A possible reason for this policy could be, for example, if the link to our peer via the IXP was not reliable, or if it had insufficient capacity. The traffic coming from our peer would traverse the ISP by default, but we would still have the direct link as a backup in case the ISP path were to go down.

19. Using the following example, configure your RX1 border router to prepend the path announced to your peer to the right (think clockwise). For example, for AS10, the peer would be AS20.

R11:

```
ip prefix-list as10-prefix permit 10.10.0.0/16
!
route-map set-prepend permit 10
  match ip address prefix-list as10-prefix
  set as-path prepend 10 10
route-map set-prepend permit 20
!
ipv6 prefix-list ipv6-as10-prefix permit FEC0:10::/32
!
route-map ipv6-set-prepend permit 10
  match ipv6 address prefix-list ipv6-as10-prefix
  set as-path prepend 10 10
route-map ipv6-set-prepend permit 20
!

router bgp 10
  neighbor 10.255.0.20 route-map set-prepend out
  address-family ipv6
    neighbor FEC0:FF:0:20:1 route-map ipv6-set-prepend out
  !
```

Ask your peer to do a soft reset on their router (the one that peers at the IXP), and then verify that their path to you is now via the ISP:

```
R21#show ip route 10.10.0.0
R21#show ipv6 route FEC0:10::/32
R21#show ip bgp 10.10.0.0
R21#show bgp ipv6 unicast FEC0:10::/32
```

```
R21#traceroute 10.10...  
R21#traceroute FEC0:10...
```

Compare this with the paths towards you as seen from your other peer (AS50, in the example).