Ejercicios de BGP– iBGP, eBGP

En la siguiente práctica estableceremos sesiones básicas de iBGP y eBGP. Cada AS intercambiará rutas con un proveedor de servicios de Internet (ISP), y luego con los AS de otros grupos por vía de un punto de intercambio de Internet (IXP). Por favor utilice los diagramas de topología para guiarse.

Esta práctica es una continuación de los ejercicios de capa 2 y de OSPF, así que asumiremos que OSPF ya ha sido configurado internamente, y que hay conectividad completa entre enrutadores y estaciones de trabajo.

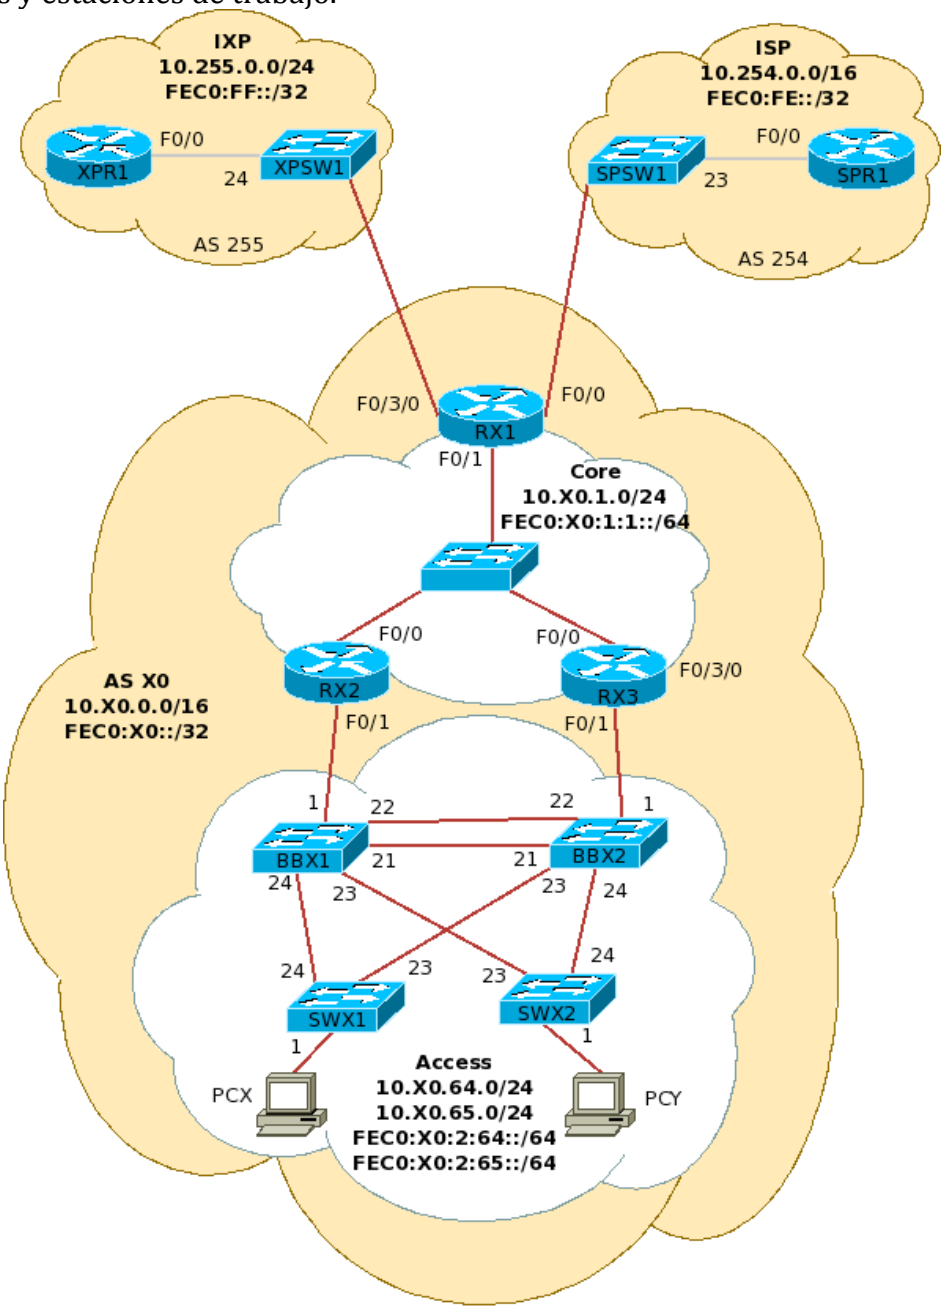

iBGP

1. **Verificar que tenemos buena conectividad hacia cada una de las interfaces loopback.**

```
ping 10.10.254.2
ping 10.10.254.3
ping 10.10.64.2
ping 10.10.64.3
```

```
ping ipv6 FEC0:10:0:FE::2
ping ipv6 FEC0:10:0:FE::3
ping ipv6 FEC0:10:2:64::2
ping ipv6 FEC0:10:2:64::3
```

IPv4:

```
show ip ospf            : Muestra información general de OSPF
show ip ospf interface  : Muestra el estatus de OSPF en cada interfaz
show ip ospf neighbor   : Muestra la lista de vecinos OSPF.
show ip ospf database   : Muestra la base de datos topológica de OSPF
```

IPv6:

```
show ipv6 ospf
show ipv6 ospf interface
show ipv6 ospf neighbor
show ipv6 ospf database
```

2. **Establezca los parámetros globales.**

```
router bgp 10
 bgp log-neighbor-changes
 no synchronization
 no auto-summary
 bgp deterministic-med
 distance bgp 200 200 200
```

3. **Configure los vecinos iBGP.**

R11:

```
router bgp 10
 neighbor 10.10.254.2 remote-as 10
 neighbor 10.10.254.2 update-source loopback 0
 neighbor 10.10.254.2 description iBGP to R12
 neighbor 10.10.254.2 password NSRC
 neighbor 10.10.254.3 remote-as 10
 neighbor 10.10.254.3 update-source loopback 0
 neighbor 10.10.254.3 description iBGP to R13
 neighbor 10.10.254.3 password NSRC
 neighbor FEC0:10:0:FE::2 remote-as 10
 neighbor FEC0:10:0:FE::2 update-source loopback 0
 neighbor FEC0:10:0:FE::2 description iBGP to R12
 neighbor FEC0:10:0:FE::2 password NSRC
 neighbor FEC0:10:0:FE::3 remote-as 10
 neighbor FEC0:10:0:FE::3 update-source loopback 0
```

```
 neighbor FEC0:10:0:FE::3 description iBGP to R13
 neighbor FEC0:10:0:FE::3 password NSRC
address-family ipv4
 no neighbor FEC0:10:0:FE::2 activate
 no neighbor FEC0:10:0:FE::3 activate
address-family ipv6
 neighbor FEC0:10:0:FE::2 activate
 neighbor FEC0:10:0:FE::3 activate
```

4. **PARE -- Checkpoint 1.**

IPv4:

```
show ip bgp summary     : Muestra información de BGP y de los vecinos
show ip bgp             : Muestra la lista de las rutas aprendidas
show ip route           : Muestra todas las rutas instaladas en la FIB
```

IPv6:

```
show bgp ipv6 unicast summary
show bgp ipv6
show ipv6 route
```

5. **Crear redes de "clientes" para anunciar via iBGP**

En los enrutadores interiores (RX2 y RX3) solamente:

```
router bgp 10
 network 10.10.64.0 mask 255.255.192.0
 address-family ipv6
  network FEC0:10:2::/48
```

6. **PARE -- Checkpoint 2.**

En R12 y R13:

```
show ip bgp neighbors 10.10.254.1 advertised-routes
show bgp ipv6 unicast neighbors FEC0:10:0:FE::1 advertised-routes
```

En el enrutador de frontera:

```
show ip bgp
show bgp ipv6 unicast
```

Por qué no está siendo anunciado el prefijo?

7. **Crear una ruta estática para su prefijo.**

En los enrutadores interiores (RX2 y RX3) solamente:

```
ip route 10.10.64.0 255.255.192.0 Null0 250
ipv6 route FEC0:10:2::/48 Null0 250
```

8. STOP -- Checkpoint

```
ping 10.10.254.2
ping 10.10.254.3
ping 10.10.1.1
ping 10.10.1.2
ping 10.10.64.1
ping 10.10.65.1
ping FEC0:10:0:FE::2
ping FEC0:10:0:FE::3
ping FEC0:10:2:64::2
ping FEC0:10:2:64::3
ping FEC0:10:2:65::2
ping FEC0:10:2:65::3
```

```
write memory
show running-config
show startup-config
```

eBGP

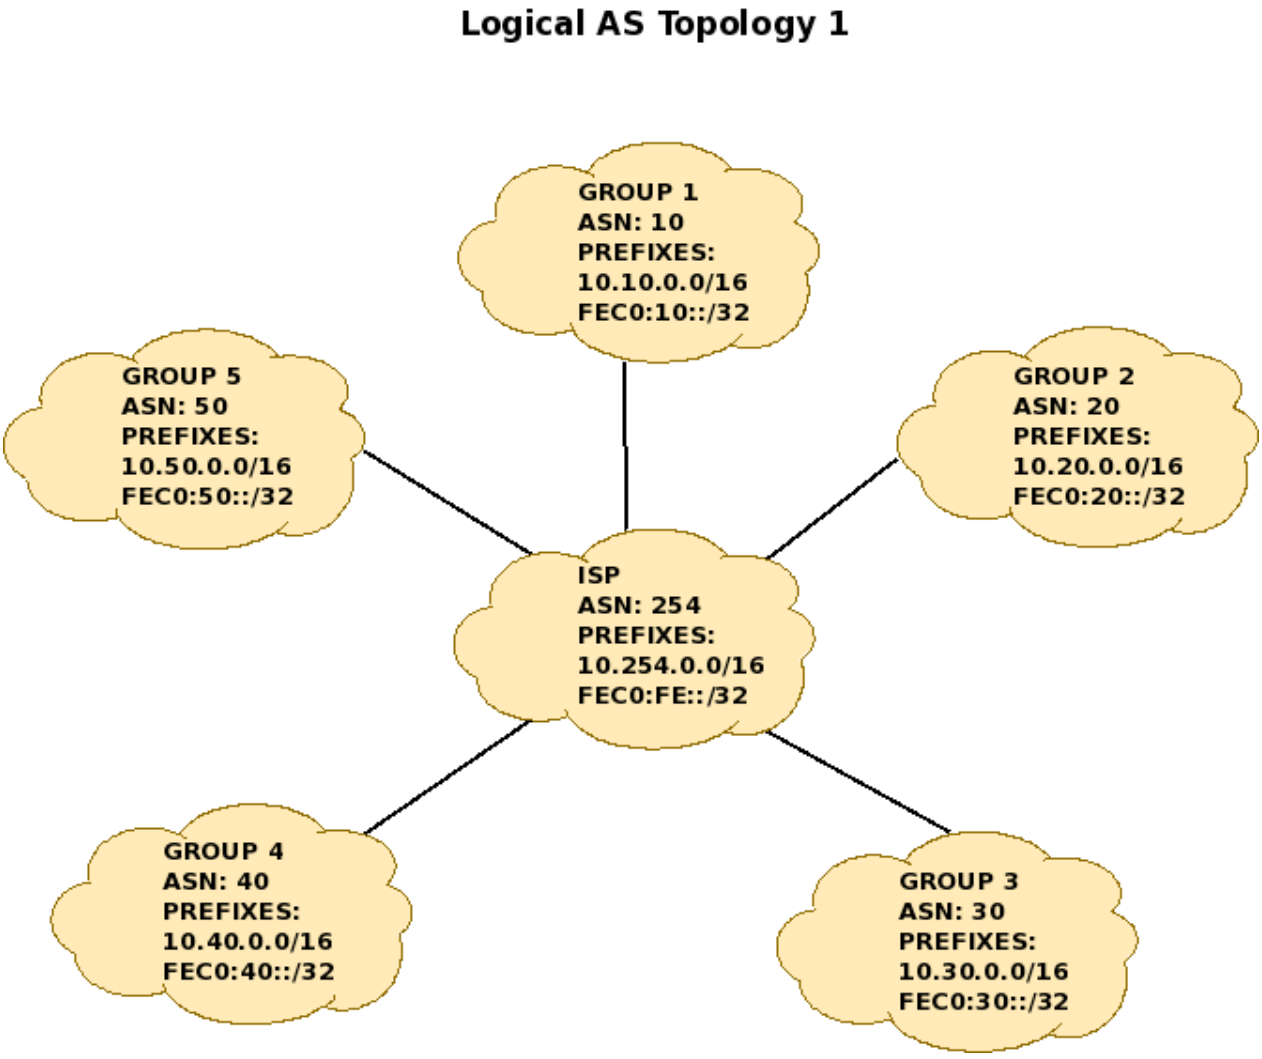

1. Configure su interfaz con el ISP.

R11:

```
interface FastEthernet0/0
 description Link to ISP
 ip address 10.254.10.2 255.255.255.252
 ipv6 enable
 ipv6 address FEC0:FE:0:10::2/64
 no ip redirects
 no ip directed-broadcast
 no ip proxy-arp
 no shutdown
```

ISP:

Todos los enlaces de clientes son agregados en un switch gestionado por el ISP, y desde allí todas las conexiones son transportadas en una troncal que termina en una interfaz física del enrutador. Por esa razón utilizaremos etiquetado 802.1Q en el enrutador del ISP. La ID de la VLAN corresponderá con el número de AS:

```
interface FastEthernet0/0.10
 description Link to as10
 encapsulation dot1Q 10
 ip address 10.254.10.1 255.255.255.252
 ipv6 address FEC0:FE:0:10::1/64
 ipv6 enable
 no ip redirects
 no ip directed-broadcast
 no ip proxy-arp
 no shutdown
!
```

2. Configure la sesión eBGP con el ISP.

R11:

```
router bgp 10
 neighbor 10.254.10.1 remote-as 254
 neighbor 10.254.10.1 description eBGP to ISP
 neighbor 10.254.10.1 password NSRC
 neighbor FEC0:FE:0:10::1 remote-as 254
 neighbor FEC0:FE:0:10::1 description eBGP to ISP
 neighbor FEC0:FE:0:10::1 password NSRC
 address-family ipv4
  no neighbor FEC0:FE:0:10::1 activate
 address-family ipv6
  neighbor FEC0:FE:0:10::1 activate
```

ISP:

```
router bgp 254
 bgp log-neighbor-changes
 neighbor 10.254.10.2 remote-as 10
 neighbor 10.254.10.2 description eBGP to as10
 neighbor 10.254.10.2 password NSRC
 neighbor FEC0:FE:0:10::2 remote-as 10
 neighbor FEC0:FE:0:10::2 description eBGP to as10
 neighbor FEC0:FE:0:10::2 password NSRC
 address-family ipv4
  no neighbor FEC0:FE:0:10::2 activate
exit-address-family
 !
 address-family ipv6
  neighbor FEC0:FE:0:10::2 activate
```

3. PARE -- Checkpoint.

```
show ip bgp summary
show ip bgp neighbors 10.254.10.1
show ip bgp neighbors 10.254.10.1 advertised-routes
show ip bgp neighbors 10.254.10.1 routes
show ip bgp
show bgp ipv6 unicast summary
show bgp ipv6 unicast neighbors FEC0:FE:0:10::1
show bgp ipv6 unicast neighbors FEC0:FE:0:10::1 advertised-routes
show bgp ipv6 unicast neighbors FEC0:FE:0:10::1 routes
show bgp ipv6 unicast
```

¿Puede observar alguno de los prefijos del ISP en su RIB de BGP? ¿Por qué?

4. **Agregue la red punto-a-punto en su IGP, pero asegúrese de no establecer una adyacencia OSPF con su vecino eBGP!**

```
router ospf 100
 passive-interface FastEthernet0/0
!
ipv6 router ospf 100
 passive-interface FastEthernet0/0
```

```
interface FastEthernet0/0
 ip ospf 100 area 0
 ipv6 ospf 100 area 0
```

5. **Haga agregación de los bloques CIDR.**

```
router bgp 10
 aggregate-address 10.10.0.0 255.255.0.0
 address-family ipv6
  aggregate-address FEC0:10::/32
```

6. **PARE -- Checkpoint.**

```
show ip bgp neighbors 10.254.10.1 advertised-routes
show ip bgp neighbors 10.254.10.1 routes
show bgp ipv6 unicast neighbors FEC0:FE:0:10::1 advertised-routes
show bgp ipv6 unicast neighbors FEC0:FE:0:10::1 routes
```

7. **Anuncie sólo el bloque agregado, no sus componentes.**

```
router bgp 10
 aggregate-address 10.10.0.0 255.255.0.0 summary-only
 address-family ipv6
  aggregate-address FEC0:10::/32 summary-only
```

Otra opción sería:

```
router bgp 10
 no aggregate-address 10.10.0.0 255.255.0.0 summary-only
```

```
 network 10.10.0.0 mask 255.255.0.0
 address-family ipv6
  no aggregate-address FEC0:10::/32 summary-only
  network FEC0:10::/32
 exit
exit
ip route 10.10.0.0 255.255.0.0 Null0 250
ipv6 route FEC0:10::/32 Null0 250
```

8. Cree listas de prefijos para políticas de entrada/salida.

R11:

```
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
```

ISP:

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

9. Cree las políticas de entrada/salida.

R11:

```
router bgp 10
 neighbor 10.254.10.1 prefix-list out-peer out
 neighbor 10.254.10.1 prefix-list isp-in-peer in
 neighbor FEC0:FE:0:10::1 prefix-list ipv6-out-peer out
 neighbor FEC0:FE:0:10::1 prefix-list ipv6-isp-in-peer in
```

ISP:

```
router bgp 254
 neighbor 10.254.10.2 prefix-list as10-in-peer in
 neighbor FEC0:FE:0:10::2 prefix-list ipv6-as10-in-peer in
```

10. PARE -- Checkpoint.

```
show ip bgp summary
show ip bgp
ping <ip_address>
traceroute <ip_address>
write memory
show running-config
show startup-config
```

Puntos de Intercambio

Ahora introduciremos un punto de intercambio de capa-2 (IXP). Un IXP de capa-2 es básicamente un switch que permite a los AS establecer sesiones de intercambio de tráfico (peering), evitando atravesar innecesariamente circuitos del ISP que cuestan dinero. En esta práctica, cada AS hará peering con otros dos AS, además de su sesión con el ISP.

Un elemento clave en el establecimiento de un peering es usar filtros de prefijos para evitar anuncios BGP que pudieran causar que su AS se convierta en un punto de tránsito (pasando tráfico que ni origina ni termina en su red).

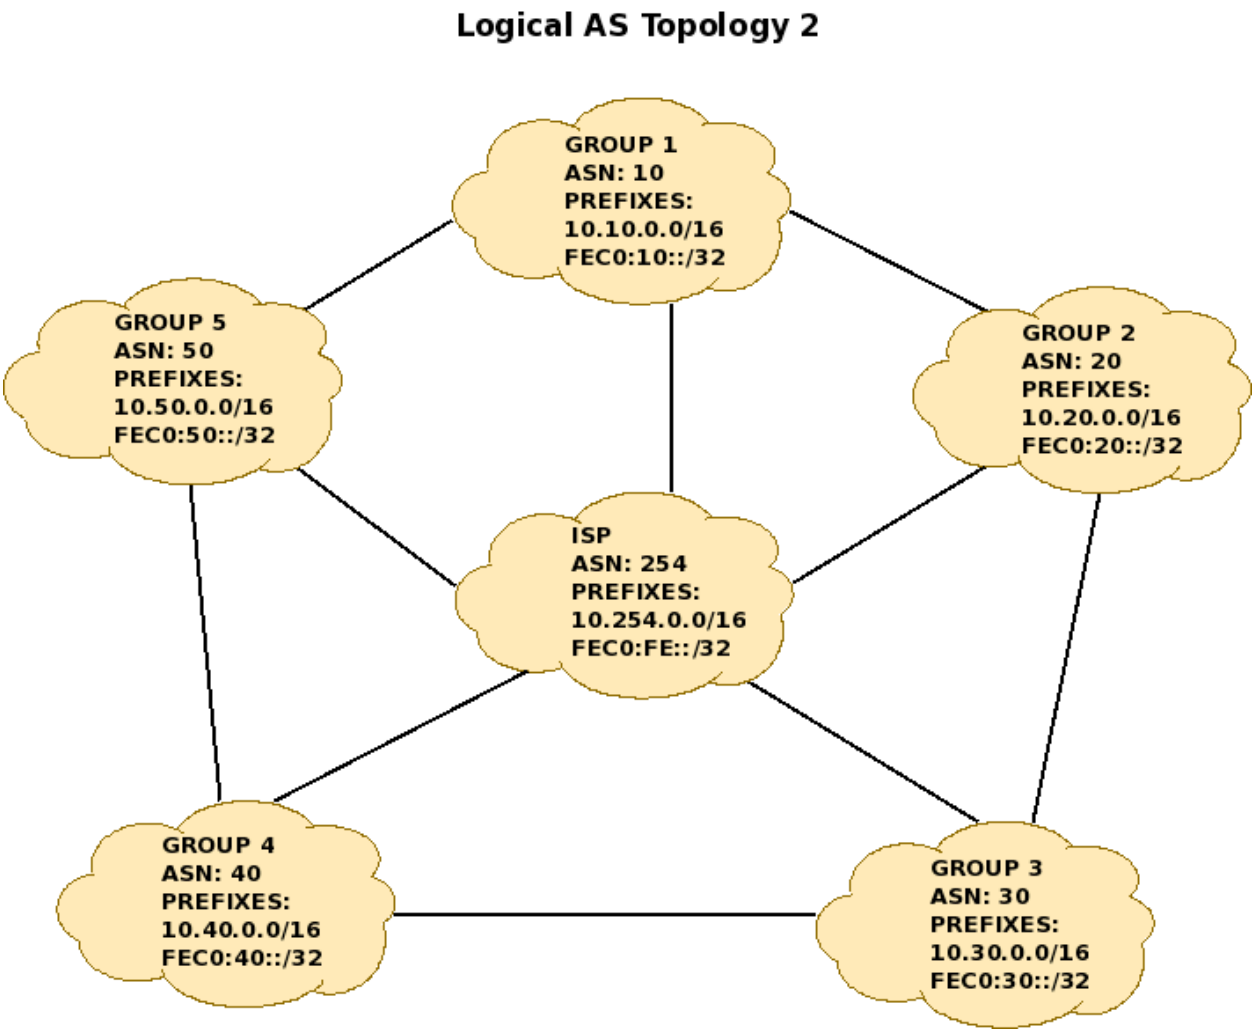

1. Configure la interfaz que conecta con el IXP.

R11:

```
interface FastEthernet0/3/0
 description Link to IXP
 ip address 10.255.0.10 255.255.255.0
 ipv6 enable
 ipv6 address FEC0:FF::10/64
 no ip redirects
 no ip directed-broadcast
 no ip proxy-arp
 no shutdown
```

2. Asegúrese de incluir la red del IXP en su IGP.

R11:

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

interface FastEthernet0/3/0
 ip ospf 100 area 0
 ipv6 ospf 100 area 0
```

3. Configure sus sesiones eBGP por cada vecino en el IXP. Recuerde mirar el diagrama.

R11:

```
router bgp 10
 neighbor 10.255.0.20 remote-as 20
 neighbor 10.255.0.20 description IXP to AS20
 neighbor 10.255.0.50 remote-as 50
 neighbor 10.255.0.50 description IXP to AS50
 neighbor FEC0:FF::20 remote-as 20
 neighbor FEC0:FF::20 description IXP to AS20
 neighbor FEC0:FF::50 remote-as 50
 neighbor FEC0:FF::50 description IXP to AS50
 address-family ipv4
  no neighbor FEC0:FE::20 activate
  no neighbor FEC0:FE::50 activate
 address-family ipv6
  neighbor FEC0:FF::20 activate
  neighbor FEC0:FF::50 activate
```

4. PARE – Checkpoint.

```
show ip bgp summary
show ip bgp neighbors 10.255.0.20
show ip bgp neighbors 10.255.0.20 advertised-routes
show ip bgp neighbors 10.255.0.20 routes
show ip bgp
```

```
show bgp ipv6 unicast summary
show bgp ipv6 unicast neighbors FEC0:FF::20
show bgp ipv6 unicast neighbors FEC0:FF::20 advertised-routes
show bgp ipv6 unicast neighbors FEC0:FF::20 routes
show bgp ipv6 unicast
```

5. **Cree listas de prefijos para sus políticas de salida.**

```
ip prefix-list out-peer permit 10.10.0.0/16 le 32
ipv6 prefix-list ipv6-out-peer permit FEC0:10::/32 le 128
```

6. **Y para sus políticas de entrada.**

```
ip prefix-list as20-in-peer permit 10.20.0.0/16 le 32
ip prefix-list as50-in-peer permit 10.50.0.0/16 le 32
!
ipv6 prefix-list ipv6-as20-in-peer permit FEC0:20::/32 le 128
ipv6 prefix-list ipv6-as50-in-peer permit FEC0:50::/32 le 128
```

7. **Cree las políticas de entrada/salida.**

```
router bgp 10
 neighbor 10.255.0.20 prefix-list out-peer out
 neighbor 10.255.0.20 prefix-list as20-in-peer in
 neighbor 10.255.0.50 prefix-list out-peer out
 neighbor 10.255.0.50 prefix-list as50-in-peer in
 neighbor FEC0:FF::20 prefix-list ipv6-out-peer out
 neighbor FEC0:FF::20 prefix-list ipv6-as20-in-peer in
 neighbor FEC0:FF::50 prefix-list ipv6-out-peer out
 neighbor FEC0:FF::50 prefix-list ipv6-as50-in-peer in
```

8. **PARE – Checkpoint.**

```
show ip bgp summary
show ip bgp
ping <ip_address>
traceroute <ip_address>
show bgp ipv6 unicast summary
show bgp ipv6 unicast
ping ipv6 <ipv6_address>
traceroute ipv6 <ipv6_address>
```

```
write memory
show running-config
```