


BGP Exercises – iBGP, eBGP

In the following exercises we will set up basic iBGP and eBGP sessions.. Each AS will peer with an Internet Service Provider (ISP), and later with other groups via an Internet Exchange Point (IXP). Please use the physical and logical topology diagrams to guide you.

This lab is a continuation of the Layer 2 and OSPF labs, so we assume that OSPF has already been configured internally, and that there is full connectivity among workstations and routers.



iBGP

1. Verify that we have good connectivity to each of the loopback interfaces.

```
ping 10.X0.254.2          ping ipv6 FEC0:X0:0:FE::2
ping 10.X0.254.3          ping ipv6 FEC0:X0:0:FE::3
ping 10.X0.64.2           ping ipv6 FEC0:X0:2:64::2
ping 10.X0.64.3           ping ipv6 FEC0:X0:2:64::3
```

IPv4:

```
show ip ospf              : shows general OSPF information
show ip ospf interface    : shows the status of OSPF in an interface
show ip ospf neighbor     : shows OSPF neighbor list.
show ip ospf database     : shows OSPF topology DB
```

IPv6:

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

2. Log the BGP neighbor changes and set other global parameters.

```
router bgp X0
  bgp log-neighbor-changes
  no synchronization
  no auto-summary
  bgp deterministic-med
  distance bgp 200 200 200
```

3. Configure iBGP neighbors.

RX1:

```
router bgp X0
  neighbor 10.X0.254.2 remote-as X0
  neighbor 10.X0.254.2 update-source loopback 0
  neighbor 10.X0.254.2 description iBGP to RX2
  neighbor 10.X0.254.2 password NSRC
  neighbor 10.X0.254.3 remote-as X0
  neighbor 10.X0.254.3 update-source loopback 0
  neighbor 10.X0.254.3 description iBGP to RX3
  neighbor 10.X0.254.3 password NSRC
  neighbor FEC0:X0:0:FE::2 remote-as X0
  neighbor FEC0:X0:0:FE::2 update-source loopback 0
  neighbor FEC0:X0:0:FE::2 description iBGP to RX2
  neighbor FEC0:X0:0:FE::2 password NSRC
  neighbor FEC0:X0:0:FE::3 remote-as X0
  neighbor FEC0:X0:0:FE::3 update-source loopback 0
  neighbor FEC0:X0:0:FE::3 description iBGP to RX3
```

```

neighbor FEC0:X0:0:FE::3 password NSRC
address-family ipv4
no neighbor FEC0:X0:0:FE::2 activate
no neighbor FEC0:X0:0:FE::3 activate
address-family ipv6
neighbor FEC0:X0:0:FE::2 activate
neighbor FEC0:X0:0:FE::3 activate

```

4. STOP -- Checkpoint 1.

IPv4:

```

show ip bgp summary      : shows BGP information and neighbors
show ip bgp              : shows a list of learned BGP paths
show ip route            : shows all installed routes

```

IPv6:

```

show bgp ipv6 unicast summary
show bgp ipv6
show ipv6 route

```

5. Create “customer” networks to advertise via iBGP

On the internal routers (RX2 and RX3) only:

```

router bgp X0
network 10.X0.64.0 mask 255.255.192.0
address-family ipv6
network FEC0:X0:2::/48

```

6. STOP -- Checkpoint 2.

On RX2 and RX3:

```

show ip bgp neighbors 10.X0.254.1 advertised-routes
show bgp ipv6 unicast neighbors FEC0:X0:0:FE::1 advertised-routes

```

On the border router:

```

show ip bgp
show bgp ipv6 unicast

```

Why isn't the prefix being announced?

7. Create a static route for your prefix.

On the internal routers (RX2 and RX3) only:

```

ip route 10.X0.64.0 255.255.192.0 Null0 250
ipv6 route FEC0:X0:2::/48 Null0 250


```

8. STOP -- Checkpoint

```
ping 10.X0.254.2
ping 10.X0.254.3
ping 10.X0.1.1
ping 10.X0.1.2
ping 10.X0.64.1
ping 10.X0.65.1
ping FEC0:X0:0:FE::2
ping FEC0:X0:0:FE::3
ping FEC0:X0:2:64::2
ping FEC0:X0:2:64::3
ping FEC0:X0:2:65::2
ping FEC0:X0:2:65::3
```

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

Logical AS Topology 1



1. Configure your interface to the ISP.

RX1:

```
interface FastEthernet0/0
description Link to ISP
ip address 10.254.X0.2 255.255.255.252
ipv6 enable
ipv6 address FEC0:FE:0:X0::2/64
no ip redirects
no ip directed-broadcast
no ip proxy-arp
no shutdown
```

ISP:

All customer links are aggregated at a switch managed by the ISP, and from there, all connections are trunked into one physical router interface. For that reason, we use 802.1Q tagging at the ISP router. The VLAN ID corresponds to the AS number:

```
interface FastEthernet0/0.X0
description Link to asX0
encapsulation dot1Q X0
ip address 10.254.X0.1 255.255.255.252
ipv6 address FEC0:FE:0:X0::1/64
ipv6 enable
no ip redirects
no ip directed-broadcast
no ip proxy-arp
no shutdown
!
```

2. Configure eBGP session to the ISP.

RX1:

```
router bgp X0
neighbor 10.254.X0.1 remote-as 254
neighbor 10.254.X0.1 description eBGP to ISP
neighbor 10.254.X0.1 password NSRC
neighbor FEC0:FE:0:X0::1 remote-as 254
neighbor FEC0:FE:0:X0::1 description eBGP to ISP
neighbor FEC0:FE:0:X0::1 password NSRC
address-family ipv4
no neighbor FEC0:FE:0:X0::1 activate
address-family ipv6
neighbor FEC0:FE:0:X0::1 activate
```

ISP:

```
router bgp 254
bgp log-neighbor-changes
neighbor 10.254.X0.2 remote-as X0
neighbor 10.254.X0.2 description eBGP to asX0
neighbor 10.254.X0.2 password NSRC
neighbor FEC0:FE:0:X0::2 remote-as X0
neighbor FEC0:FE:0:X0::2 description eBGP to asX0
neighbor FEC0:FE:0:X0::2 password NSRC
address-family ipv4
no neighbor FEC0:FE:0:X0::2 activate
exit-address-family
!
address-family ipv6
neighbor FEC0:FE:0:X0::2 activate
```

3. STOP -- Checkpoint.

```

show ip bgp summary
show ip bgp neighbors 10.254.X0.1
show ip bgp neighbors 10.254.X0.1 advertised-routes
show ip bgp neighbors 10.254.X0.1 routes
show ip bgp
show bgp ipv6 unicast summary
show bgp ipv6 unicast neighbors FEC0:FE:0:X0::1
show bgp ipv6 unicast neighbors FEC0:FE:0:X0::1 advertised-routes
show bgp ipv6 unicast neighbors FEC0:FE:0:X0::1 routes
show bgp ipv6 unicast

```

Can you see any of the ISP prefixes in your BGP rib? Why not?

4. Add the point-to-point network to your IGP, but make sure that you don not establish OSPF adjacencies with your eBGP peer!

```

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. Aggregate CIDR blocks.

```

router bgp X0
  aggregate-address 10.X0.0.0 255.255.0.0
  address-family ipv6
    aggregate-address FEC0:X0::/32

```

6. STOP -- Checkpoint.

```

show ip bgp neighbors 10.254.X0.1 advertised-routes
show ip bgp neighbors 10.254.X0.1 routes
show bgp ipv6 unicast neighbors FEC0:FE:0:X0::1 advertised-routes
show bgp ipv6 unicast neighbors FEC0:FE:0:X0::1 routes

```

7. Advertise only a summary aggregate.

```

router bgp X0
  aggregate-address 10.X0.0.0 255.255.0.0 summary-only
  address-family ipv6
    R11(config-rtr-af)#aggregate-address FEC0:X0::/32 summary-only

```

Another option would be:

```

router bgp X0
  no aggregate-address 10.X0.0.0 255.255.0.0 summary-only
  network 10.X0.0.0 mask 255.255.0.0

```

```

address-family ipv6
  no aggregate-address FEC0:X0::/32 summary-only
  network FEC0:X0::/32
exit
exit
ip route 10.X0.0.0 255.255.0.0 Null0 250
ipv6 route FEC0:X0::/32 Null0 250

```

8. Create prefix lists for inbound/outbound policies.

RX1:

```

ip prefix-list out-peer permit 10.X0.0.0/16 le 32
ip prefix-list isp-in-peer deny 10.X0.0.0/16 le 32
ip prefix-list isp-in-peer permit 0.0.0.0/0 le 32
ipv6 prefix-list ipv6-out-peer permit FEC0:X0::/32 le 128
ipv6 prefix-list ipv6-isp-in-peer deny FEC0:X0::/32 le 128
ipv6 prefix-list ipv6-isp-in-peer permit ::/0 le 128

```

ISP:

```

ip prefix-list asX0-in-peer permit 10.X0.0.0/16 le 32
ipv6 prefix-list ipv6-asX0-in-peer permit FEC0:X0::/32 le 128

```

9. Create input and output policies.

RX1:

```

router bgp X0
  neighbor 10.254.X0.1 prefix-list out-peer out
  neighbor 10.254.X0.1 prefix-list isp-in-peer in
  neighbor FEC0:FE:0:X0::1 prefix-list ipv6-out-peer out
  neighbor FEC0:FE:0:X0::1 prefix-list ipv6-isp-in-peer in

```

ISP:

```

router bgp 254
  neighbor 10.254.X0.2 prefix-list asX0-in-peer in
  neighbor FEC0:FE:0:X0::2 prefix-list ipv6-asX0-in-peer in

```

10. STOP -- Checkpoint.

```

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


```


Exchange Points

We will now introduce a layer-2 exchange point (IXP). A layer-2 IXP is basically a switch which allows ASs to establish private peerings where they can exchange traffic, avoiding unnecessarily traversing costly ISP circuits. In this lab, each AS will peer with two other ASs, in addition to their ISP peering.

One key element of a private peering arrangement is to use prefix filters to avoid BGP announcements that could cause your AS to become a transit point (passing traffic which neither originates nor terminates in your network).

Logical AS Topology 2



1. Configure the interface that connects to the 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. Make sure to carry the IXP network into your IGP.

R11:

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

3. Configure your eBGP peering sessions for every neighbor you will have at the IXP.

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. STOP – 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. Create prefix lists for your outbound policies.

```
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. Create prefix lists for your inbound policies.

```
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. Create input and output policies.

```
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. STOP – 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
```