

Advanced Routing Workshop

Basic BGP Lab

February 19, 2013

Contents

1	Introduction	3
2	Pre-requisites	3
3	Address Space Allocation	3
3.1	Regional REN (RREN)	3
3.2	National RENs (NRENs)	3
4	iBGP Configuration	4
4.1	Enable the BGP process	4
4.2	Configure iBGP neighbors	5
4.3	Advertise your network	5
5	Multihoming - eBGP Configuration	6
5.1	Connect to the NREN	6
5.2	BGP-peer with the NREN and the ISP	7
5.3	Filter what you send and receive	9
5.3.1	Remove static routes	10
6	Traffic Exchange (Peering)	11
6.1	Connect to your neighbor AS	12
7	Appendix A - RREN Configuration	15

8	Appendix B - NREN1 Configuration Example	17
9	Appendix C - ISP1 Configuration Example	20

1 Introduction

The purpose of this exercise is to:

- Understand the routing implications of connecting to multiple external routing domains
- Learn to configure basic eBGP to exchange routing information with multiple external peers and iBGP to carry that information inside your network.

2 Pre-requisites

This exercise builds upon the configurations implemented in the OSPF + Static routing lab. You must be able to:

- Ping your neighbor router in the same AS using its loopback address (both IPv4 and IPv6!).
- Ping your neighbor routers in other ASs using their point-to-point link addresses.

Note: Actually, if everyone configured their OSPF and static routes properly in the previous exercise, you should be able to ping every other router using their loopback address.

3 Address Space Allocation

3.1 Regional REN (RREN)

We only need one:

RREN	IPv4	IPv6	ASN
1	10.100.0.0/16	fd00:100::/32	100

3.2 National RENs (NRENs)

NREN	IPv4	IPv6	ASN
1	10.101.0.0/16	fd00:101::/32	101
2	10.102.0.0/16	fd00:102::/32	102

... and so on.

4 iBGP Configuration

4.1 Enable the BGP process

Before we set up iBGP, we need to do some basic preparation on the router. The IOS defaults are not optimized, so before we bring up BGP sessions, we should set the parameters that we require. The default distance for eBGP is 20, the default distance for iBGP is 200, and the default distance for OSPF is 110. This means that there is a potential for a prefix learned by eBGP to override the identical prefix carried by OSPF. To protect against accidents, the eBGP distance is set to 200 also.

The command to do this is the *distance bgp* subcommand:

```
distance bgp <external-routes> <internal-routes> <local-routes>
```

We also want to:

- Enable logging of BGP neighbor state changes
- Disable the requirement that a route must be present in the IGP table before it can be advertised by BGP (synchronization).
- Disable auto-summarization of routes to classful network boundaries
- Disable the automatic exchange of IPv4 unicast routes on every peering session.

This must be done in all future BGP configurations of this workshop:

On both R11 and R12:

```
router bgp 10
  distance bgp 200 200 200
  bgp log-neighbor-changes
  no synchronization
  no auto-summary
  no bgp default ipv4-unicast
```

4.2 Configure iBGP neighbors

Again, make sure that you can ping the neighbor router using its loopback address, otherwise the BGP session will not come up!

On R11:

```
router bgp 10
 address-family ipv4
   neighbor 10.10.255.2 remote-as 10
   neighbor 10.10.255.2 update-source loopback 0
   neighbor 10.10.255.2 description iBGP to R12
   neighbor 10.10.255.2 password N$RC
 address-family ipv6
   neighbor fd00:10:ff::2 remote-as 10
   neighbor fd00:10:ff::2 update-source loopback 0
   neighbor fd00:10:ff::2 description iBGP to R12
   neighbor fd00:10:ff::2 password N$RC
```

On R12:

```
router bgp 10
 address-family ipv4
   neighbor 10.10.255.1 remote-as 10
   neighbor 10.10.255.1 update-source loopback 0
   neighbor 10.10.255.1 description iBGP to R11
   neighbor 10.10.255.1 password N$RC
 address-family ipv6
   neighbor fd00:10:ff::1 remote-as 10
   neighbor fd00:10:ff::1 update-source loopback 0
   neighbor fd00:10:ff::1 description iBGP to R11
   neighbor fd00:10:ff::1 password N$RC
```

Check that the BGP sessions are up on both sides.

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

4.3 Advertise your network

1. Use the ‘network’ command to tell BGP which prefixes you want to announce.

On R11 and R12:

```
router bgp 10
 address-family ipv4
   network 10.10.0.0 mask 255.255.0.0
 address-family ipv6
   network fd00:10::/32
```

Get the list of learned paths:

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

Q. Do you see any paths?

2. Create a static route for the prefix being announced on each router:

On R11 and R12:

```
ip route 10.10.0.0 255.255.0.0 null0
ipv6 route fd00:10::/32 null0
```

These are called a “pull up routes”

Get the list of learned paths again. You should see both your prefix and the neighbor's.

Q. Why are these routes needed?

5 Multihoming - eBGP Configuration

5.1 Connect to the NREN

1. Configure your RX1 router to connect to the NREN with a point-to-point link.

NRENs: Use configuration in Appendix.

On R11:

```
interface GigabitEthernet1/0
 description P2P Link to NREN1
 ip address 10.101.254.2 255.255.255.252
 no ip directed-broadcast
 no ip redirects
```

```
no ip proxy-arp
ipv6 address fd00:101:fe::1/127
ipv6 nd ra suppress
no shutdown
```

Make sure that it's up and that you can ping the other side:

```
ping 10.101.254.1
ping fd00:101:fe::0
```

Do some traceroutes to other networks again:

```
R11# traceroute 10.20.255.1
R11# traceroute 10.30.255.1
```

Has anything changed since the last exercise?

Notice that before we had only one connection to the Internet - via the ISP. Now we have two. But we are still using a default route pointing to the ISP only!

We could add another default route pointing to the NREN, but that would not give us much flexibility in terms of traffic policies. Keep going.

5.2 BGP-peer with the NREN and the ISP

1. Configure eBGP sessions to the ISP and the NREN

On R11:

```
router bgp 10
address-family ipv4
neighbor 10.101.254.1 remote-as 101
neighbor 10.101.254.1 description eBGP to NREN1
neighbor 10.101.254.1 password N$RC
address-family ipv6
neighbor fd00:101:fe:: remote-as 101
neighbor fd00:101:fe:: description eBGP to NREN1
neighbor fd00:101:fe:: password N$RC
```

Q. Notice that with eBGP we no longer use the loopback address as the endpoint of the BGP session, as we did with iBGP

On R12:


```

router bgp 10
  address-family ipv4
    neighbor 10.201.254.1 remote-as 201
    neighbor 10.201.254.1 description eBGP to ISP1
    neighbor 10.201.254.1 password N$RC
  address-family ipv6
    neighbor fd00:201:fe:: remote-as 201
    neighbor fd00:201:fe:: description eBGP to ISP1
    neighbor fd00:201:fe:: password N$RC

```

Check that the BGP sessions are up on both routers:

```

show ip bgp summary
show bgp ipv6 unicast summary

```

Once those are up, check if you are learning any prefixes:

```

R11# show ip bgp nei 10.101.254.1 routes
R11# sh bgp ipv6 uni neigh fd00:101:fe:: routes

```

Inject the point-to-point subnets that connect to your upstreams into your IGP (OSPF), to make sure that the external next hops are in the routing tables of each iBGP peer.

On R11 and R12:

```

interface GigabitEthernet1/0
  ip ospf 10 area 0
  ipv6 ospf 10 area 0

```

Verify what you are advertising to the NREN:

```

R11# show ip bgp nei 10.101.254.1 advertised-routes
R11# sh bgp ipv6 uni neigh fd00:101:fe:: advertised

```

... and to the ISP:

```

R12# show ip bgp neighbor 10.201.254.1 advertised-routes
R12# sh bgp ipv6 uni neigh fd00:201:fe:: advertised

```

Are you perhaps announcing other prefixes that don't originate in your AS? If so, can you remember what serious negative implications this could have? Please stop and think about this. Ask the instructors if you need clarification.

5.3 Filter what you send and receive

1. Create prefix lists for your inbound/outbound filters.

On R11:

```
ip prefix-list out-peer permit 10.10.0.0/16 le 32
ip prefix-list nren-in-peer deny 10.10.0.0/16 le 32
ip prefix-list nren-in-peer permit 0.0.0.0/0 le 32
ipv6 prefix-list ipv6-out-peer permit fd00:10::/32 le 128
ipv6 prefix-list ipv6-nren-in-peer deny fd00:10::/32 le 128
ipv6 prefix-list ipv6-nren-in-peer permit ::/0 le 128
```

On R12:

```
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 fd00:10::/32 le 128
ipv6 prefix-list ipv6-isp-in-peer deny fd00:10::/32 le 128
ipv6 prefix-list ipv6-isp-in-peer permit ::/0 le 128
```

2. Now create inbound/outbound filters using those prefix lists

R11:

```
router bgp 10
 address-family ipv4
   neighbor 10.101.254.1 prefix-list nren-in-peer in
   neighbor 10.101.254.1 prefix-list out-peer out
 address-family ipv6
   neighbor fd00:101:fe:: prefix-list ipv6-nren-in-peer in
   neighbor fd00:101:fe:: prefix-list ipv6-out-peer out
```

R12:

```
router bgp 10
 address-family ipv4
   neighbor 10.201.254.1 prefix-list isp-in-peer in
   neighbor 10.201.254.1 prefix-list out-peer out
 address-family ipv6
   neighbor fd00:201:fe:: prefix-list ipv6-isp-in-peer in
   neighbor fd00:201:fe:: prefix-list ipv6-out-peer out
```

Use the *BGP refresh* capability to resend the information to the peer:

```
R11#clear ip bgp 10.101.254.1 out
R11#clear bgp ipv6 unicast fd00:101:fe:: out
```

```
R12#clear ip bgp 10.201.254.1 out
R12#clear bgp ipv6 unicast fd00:201:fe:: out
```

You should now be advertising only your own address space. Check with the ISP and NREN administrators to make sure that they are receiving your prefix.

5.3.1 Remove static routes

1. The ISPs remove their static routes towards their customers.

Now your ISP has learned a route to reach your network, correct? The ISPs can now safely remove the static routes pointing to you and the other customers:

ISP1:

```
no ip route 10.10.0.0 255.255.0.0 10.201.254.2
no ip route 10.20.0.0 255.255.0.0 10.201.254.6
no ip route 10.30.0.0 255.255.0.0 10.201.254.10
!
no ipv6 route fd00:10::/32 fd00:201:fe::1
no ipv6 route fd00:20::/32 fd00:201:fe::3
no ipv6 route fd00:30::/32 fd00:201:fe::5
```

ISP2:

```
no ip route 10.40.0.0 255.255.0.0 10.202.254.2
no ip route 10.50.0.0 255.255.0.0 10.202.254.6
no ip route 10.60.0.0 255.255.0.0 10.202.254.10
!
no ipv6 route fd00:40::/32 fd00:202:fe::1
no ipv6 route fd00:50::/32 fd00:202:fe::3
no ipv6 route fd00:60::/32 fd00:202:fe::5
```

2. Remove your static default routes

In the previous exercise, we created default routes on both routers. But thanks to BGP, we should now be receiving routes from our NREN and our ISP.

Let's check first (do this on both routers):

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

You should be learning routes advertised by other groups, and also from the NRENs and the ISPs.

Remove your static default routes:

R11:

```
no ip route 0.0.0.0 0.0.0.0 10.10.254.2
no ipv6 route ::/0 fd00:10:fe::1
```

R12:

```
no ip route 0.0.0.0 0.0.0.0 10.201.254.1
no ipv6 route ::/0 fd00:201:fe::
```

You should be able to ping any other router now. If you can't, wait for other groups to finish, or ask the instructors.

Use traceroute to verify the paths that packets are following towards various destinations:

```
R11# traceroute 10.100.255.1
R11# traceroute 10.30.255.2
...
```

Repeat the same tests from the other router in your AS and compare. Use the diagram to help you visualize it.

6 Traffic Exchange (Peering)

Direct traffic exchanges are usually established at no charge between two autonomous systems that want to save costs. The savings are achieved by not having to carry that traffic over expensive transit links via commercial providers. Also, these direct exchanges have the added benefit of reducing latency because there are fewer hops.

Usually traffic exchanges occur at public exchange points, also known as IXPs. The simplest kind of exchange point is a Layer-2 switch. In this exercise, we will simply configure direct links between routers, which from the point of view of BGP is equivalent to connecting through a switch.




Figure 2: Peering Topology

6.1 Connect to your neighbor AS

1. Configure a point-to-point link to your neighbor AS as shown in the diagram. You will have to agree with your peer on which address space to use. **Make sure to pick a point-to-point subnet that is not already used!**

The instructor will draw a map of the network at the front of the class and will ask you to document the subnet that was used for the peering session, so everybody can use that information when troubleshooting.

For example, on R12:

```
interface GigabitEthernet3/0
description Link to R21
ip address 10.10.254.5 255.255.255.252
no ip directed-broadcast
no ip redirects
no ip proxy-arp
ipv6 address fd00:10:fe::2/127
ipv6 nd ra suppress
no shutdown
```

Don't forget to inject that subnet into OSPF.

R12:

```
interface GigabitEthernet3/0
 ip ospf 10 area 0
 ipv6 ospf 10 area 0
```

Q. Remember why this is needed? If not, please ask.

2. Configure prefix lists for your inbound filters

On R12:

```
ip prefix-list AS20-in-peer permit 10.20.0.0/16 le 32
ipv6 prefix-list ipv6-AS20-in-peer permit fd00:20::/32 le 128
```

The equivalent needs to be done in R21.

3. Configure prefix lists for your outbound filters

You should have these from a previous step. You can verify like this:

```
R12#show ip prefix-list out-peer
R12#show ipv6 prefix-list ipv6-out-peer
```

4. Now create the BGP sessions and apply those inbound/outbound filters:

On R12:

```
router bgp 10
 address-family ipv4
  neighbor 10.10.254.6 remote-as 20
  neighbor 10.10.254.6 description eBGP to AS20
  neighbor 10.10.254.6 password N$RC
  neighbor 10.10.254.6 prefix-list out-peer out
  neighbor 10.10.254.6 prefix-list AS20-in-peer in
 address-family ipv6
  neighbor fd00:10:fe::3 remote-as 20
  neighbor fd00:10:fe::3 description eBGP to AS20
  neighbor fd00:10:fe::3 password N$RC
  neighbor fd00:10:fe::3 prefix-list ipv6-out-peer out
  neighbor fd00:10:fe::3 prefix-list ipv6-AS20-in-peer in
```

The equivalent needs to be done in R21.

Verify that the sessions are up:

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

..and that you are learning the prefix directly from the neighbor:

```
R12#show ip bgp neighbor 10.10.254.6 routes
R12#show bgp ipv6 unicast neighbors fd00:10:fe::3 routes
```

5. Do some traceroutes towards your peer and make sure that the path is direct.

Remember to save your configurations.

You are done! You have configured BGP in a multihomed environment and BGP is selecting the paths based on default values.

7 Appendix A - RREN Configuration

```
hostname RREN
aaa new-model
aaa authentication login default local
aaa authentication enable default enable
username nsrc secret nsrc
enable secret nsrc
service password-encryption
line vty 0 4
  transport preferred none
line console 0
  transport preferred none
no logging console
logging buffered 8192 debugging
no ip domain-lookup
ip subnet-zero
ip classless
no ip source-route
ipv6 unicast-routing
!
interface Loopback0
  ip address 10.100.255.1 255.255.255.255
  ipv6 address fd00:100:ff::1/128
!
interface GigabitEthernet1/0
  description P2P Link to RREN1
  ip address 10.100.254.1 255.255.255.252
  no ip directed-broadcast
  no ip redirects
  no ip proxy-arp
  ipv6 address fd00:100:fe::/127
  ipv6 nd ra suppress
  no shutdown
!
interface GigabitEthernet2/0
  description P2P Link to RREN2
  ip address 10.100.254.5 255.255.255.252
  no ip directed-broadcast
  no ip redirects
  no ip proxy-arp
  ipv6 address fd00:100:fe::2/127
  ipv6 nd ra suppress
  no shutdown
!
```



```

interface GigabitEthernet3/0
  description Link to IXP
  ip address 10.251.1.3 255.255.255.0
  no ip directed-broadcast
  no ip redirects
  no ip proxy-arp
  ipv6 address fd00:251:1::3/64
  ipv6 nd ra supress
  no shutdown
!
router bgp 100
  bgp log-neighbor-changes
  no synchronization
  no auto-summary
  no bgp default ipv4-unicast
  distance bgp 200 200 200
  address-family ipv4
    network 10.100.0.0 mask 255.255.0.0
    neighbor 10.100.254.2 remote-as 101
    neighbor 10.100.254.2 description eBGP to AS101
    neighbor 10.100.254.2 password N$RC
    neighbor 10.100.254.6 remote-as 102
    neighbor 10.100.254.6 description eBGP to AS102
    neighbor 10.100.254.6 password N$RC
    neighbor 10.251.1.1 remote-as 201
    neighbor 10.251.1.1 description eBGP to AS201
    neighbor 10.251.1.1 password N$RC
    neighbor 10.251.1.2 remote-as 202
    neighbor 10.251.1.2 description eBGP to AS202
    neighbor 10.251.1.2 password N$RC address-family ipv6
  network fd00:100::/32
    neighbor fd00:100:fe::1 remote-as 101
    neighbor fd00:100:fe::1 description eBGP to AS101
    neighbor fd00:100:fe::1 password N$RC
    neighbor fd00:100:fe::3 remote-as 102
    neighbor fd00:100:fe::3 description eBGP to AS102
    neighbor fd00:100:fe::3 password N$RC
    neighbor fd00:251:1::1 remote-as 201
    neighbor fd00:251:1::1 description eBGP to AS201
    neighbor fd00:251:1::1 password N$RC
    neighbor fd00:251:1::2 remote-as 202
    neighbor fd00:251:1::2 description eBGP to AS202
    neighbor fd00:251:1::2 password N$RC
  !
ip route 10.100.0.0 255.255.0.0 null0
ipv6 route fd00:100::/32 null0

```

8 Appendix B - NREN1 Configuration Example

```
hostname NREN1
aaa new-model
aaa authentication login default local
aaa authentication enable default enable
username nsrc secret nsrc
enable secret nsrc
service password-encryption
line vty 0 4
  transport preferred none
line console 0
  transport preferred none
no logging console
logging buffered 8192 debugging
no ip domain-lookup
ip subnet-zero
ip classless
no ip source-route
ipv6 unicast-routing
!
interface Loopback0
  ip address 10.101.255.1 255.255.255.255
  ipv6 address fd00:101:ff::1/128
!
interface GigabitEthernet1/0
  description P2P Link to RREN
  ip address 10.100.254.2 255.255.255.252
  no ip directed-broadcast
  no ip redirects
  no ip proxy-arp
  ipv6 address fd00:100:fe::1/127
  ipv6 nd ra suppress
  no shutdown
!
interface GigabitEthernet2/0
  description P2P Link to ISP1
  ip address 10.101.254.13 255.255.255.252
  no ip directed-broadcast
  no ip redirects
  no ip proxy-arp
  ipv6 address fd00:101:fe::6/127
  ipv6 nd ra suppress
  no shutdown
!
```

```

interface GigabitEthernet3/0
  description P2P Link to R11
  ip address 10.101.254.1 255.255.255.252
  no ip directed-broadcast
  no ip redirects
  no ip proxy-arp
  ipv6 address fd00:101:fe::0/127
  ipv6 nd ra suppress
  no shutdown
!
ip prefix-list AS10-in-peer permit 10.10.0.0/16 le 32
ipv6 prefix-list ipv6-AS10-in-peer permit fd00:10::/32 le 128
!
router bgp 101
  bgp log-neighbor-changes
  no synchronization
  no auto-summary
  no bgp default ipv4-unicast
  distance bgp 200 200 200
  address-family ipv4
    network 10.101.0.0 mask 255.255.0.0
    neighbor 10.101.254.2 remote-as 10
    neighbor 10.101.254.2 description eBGP to AS10
    neighbor 10.101.254.2 password N$RC
    neighbor 10.101.254.2 prefix-list AS10-in-peer in
    neighbor 10.101.254.14 remote-as 201
    neighbor 10.101.254.14 description eBGP to AS201
    neighbor 10.101.254.14 password N$RC
    neighbor 10.100.254.1 remote-as 100
    neighbor 10.100.254.1 description eBGP to AS100
    neighbor 10.100.254.1 password N$RC
  address-family ipv6
    network fd00:101::/32
    neighbor fd00:101:fe::1 remote-as 10
    neighbor fd00:101:fe::1 description eBGP to AS10
    neighbor fd00:101:fe::1 password N$RC
    neighbor fd00:101:fe::1 prefix-list ipv6-AS10-in-peer in
    neighbor fd00:101:fe::7 remote-as 201
    neighbor fd00:101:fe::7 description eBGP to AS201
    neighbor fd00:101:fe::7 password N$RC
    neighbor fd00:100:fe:: remote-as 100
    neighbor fd00:100:fe:: description eBGP to AS100
    neighbor fd00:100:fe:: password N$RC
!

ip route 10.101.0.0 255.255.0.0 null0

```

```
ipv6 route fd00:101::/32 null0
```

9 Appendix C - ISP1 Configuration Example

Note: *This is in addition to what was configured in the previous exercise.*

```
interface GigabitEthernet2/0
  description P2P Link to NREN1
  ip address 10.101.254.14 255.255.255.252
  no ip directed-broadcast
  no ip redirects
  no ip proxy-arp
  ipv6 address fd00:101:fe::7/127
  ipv6 nd ra suppress
  no shutdown
!
ip prefix-list AS10-in-peer permit 10.10.0.0/16 le 32
ipv6 prefix-list ipv6-AS10-in-peer permit fd00:10::/32 le 128
!

router bgp 201
  bgp log-neighbor-changes
  no synchronization
  no auto-summary
  no bgp default ipv4-unicast
  bgp deterministic-med
  distance bgp 200 200 200
  address-family ipv4
    network 10.201.0.0 mask 255.255.0.0
    neighbor 10.201.254.2 remote-as 10
    neighbor 10.201.254.2 description eBGP to AS10
    neighbor 10.201.254.2 password N$RC
    neighbor 10.201.254.2 prefix-list AS10-in-peer in
    neighbor 10.101.254.13 remote-as 101
    neighbor 10.101.254.13 description eBGP to AS101
    neighbor 10.101.254.13 password N$RC
    neighbor 10.251.1.2 remote-as 202
    neighbor 10.251.1.2 description eBGP to AS202
    neighbor 10.251.1.2 password N$RC
    neighbor 10.251.1.3 remote-as 100
    neighbor 10.251.1.3 description eBGP to AS100
    neighbor 10.251.1.3 password N$RC
  address-family ipv6
    network fd00:201::/32
    neighbor fd00:201:fe::1 remote-as 10
    neighbor fd00:201:fe::1 description eBGP to AS10
    neighbor fd00:201:fe::1 password N$RC
```

```
neighbor fd00:201:fe::1 prefix-list AS10-in-peer in
neighbor fd00:101:fe::6 remote-as 101
neighbor fd00:101:fe::6 description eBGP to AS101
neighbor fd00:101:fe::6 password N$RC
neighbor fd00:251:1::2 remote-as 202
neighbor fd00:251:1::2 description eBGP to AS202
neighbor fd00:251:1::2 password N$RC
neighbor fd00:251:1::3 remote-as 100
neighbor fd00:251:1::3 description eBGP to AS100
neighbor fd00:251:1::3 password N$RC
!

ip route 10.201.0.0 255.255.0.0 null0
ipv6 route fd00:201::/32 null0
```