Basic BGP Lab

Introduction

The purpose of this exercise is to:

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

Pre-requisites

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

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

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

Address Space Allocation

Refer to the IP Address Plan document for information about the IP address plan for the network infrastructure for these labs.

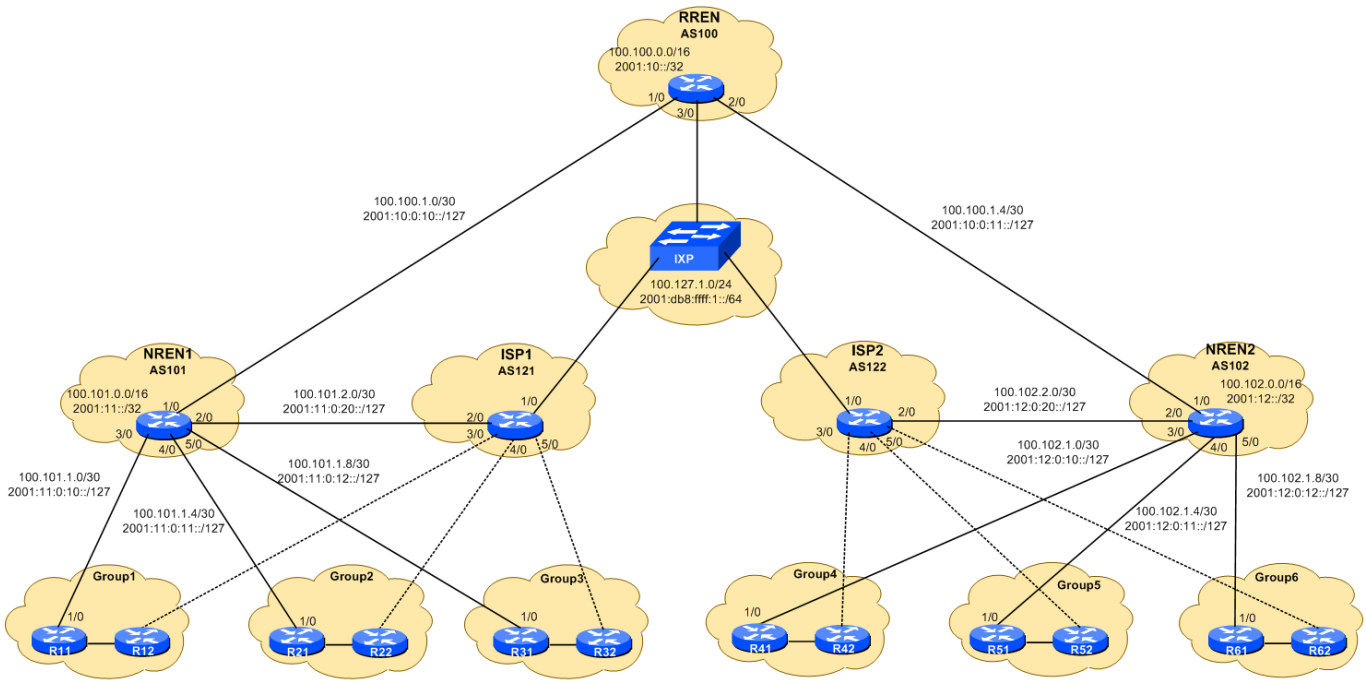

iBGP Configuration

Enable the BGP process

Before we set up iBGP, we need to do some basic preparation on the router. The Cisco IOS defaults are not optimised, so before we bring up BGP sessions, we should set the parameters that we require.

On Cisco routers, the default distance for eBGP is 20, the default distance for iBGP is 200, and the default distance for IS-IS is 115. This means that there is a potential for a prefix learned by eBGP to override the identical prefix carried by IS-IS. 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 neighbour state changes
- Configure deterministic calculations of MEDs
- 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 RX1 and RX2:

```
router bgp X0
 bgp log-neighbor-changes
 bgp deterministic-med
 no bgp default ipv4-unicast
 !
 address-family ipv4
  distance bgp 200 200 200
 address-family ipv6
  distance bgp 200 200 200
```

Configure iBGP neighbours

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

On RX1:

```
router bgp X0
 address-family ipv4
  neighbor 100.68.X0.2 remote-as X0
  neighbor 100.68.X0.2 update-source loopback 0
  neighbor 100.68.X0.2 description iBGP to RX2
```

```
  neighbor 100.68.X0.2 password NSRC-BGP
  neighbor 100.68.X0.2 next-hop-self
  neighbor 100.68.X0.2 send-community
  neighbor 100.68.X0.2 activate
!
 address-family ipv6
  neighbor 2001:db8:X0::2 remote-as X0
  neighbor 2001:db8:X0::2 update-source loopback 0
  neighbor 2001:db8:X0::2 description iBGP to RX2
  neighbor 2001:db8:X0::2 password NSRC-BGP
  neighbor 2001:db8:X0::2 next-hop-self
  neighbor 2001:db8:X0::2 send-community
  neighbor 2001:db8:X0::2 activate
```

On RX2:

```
router bgp X0
 address-family ipv4
  neighbor 100.68.X0.1 remote-as X0
  neighbor 100.68.X0.1 update-source loopback 0
  neighbor 100.68.X0.1 description iBGP to RX1
  neighbor 100.68.X0.1 password NSRC-BGP
  neighbor 100.68.X0.1 next-hop-self
  neighbor 100.68.X0.1 send-community
  neighbor 100.68.X0.1 activate
!
 address-family ipv6
  neighbor 2001:db8:X0::1 remote-as X0
  neighbor 2001:db8:X0::1 update-source loopback 0
  neighbor 2001:db8:X0::1 description iBGP to RX1
  neighbor 2001:db8:X0::1 password NSRC-BGP
  neighbor 2001:db8:X0::1 next-hop-self
  neighbor 2001:db8:X0::1 send-community
  neighbor 2001:db8:X0::1 activate
```

Once the configuration has been entered, check the configuration as displayed by the router:

```
show run | begin router bgp
```

Notice how the router has “rearranged” the BGP configuration, separating the generic from the address-family specific configuration.

Check that the BGP sessions are up on both sides.

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

Explaining some of the above commands we used for the BGP configuration:

update-source specifies the interface which should be used as the source of all BGP packets originated by the router. The default is the outgoing interface.

next-hop-self tells iBGP to use the source address of the BGP update message as the value of the next-hop attribute sent to the iBGP peer, rather than the default value which is the IP address of the router that we heard the BGP update from.

send-community tells BGP to include the BGP community attribute when sending BGP updates to neighbouring BGP speakers. Cisco IOS does not include the community attribute by default. It is important to send BGP communities to all iBGP neighbours, but be very careful about sending communities to external BGP neighbours, as we will see in the following BGP Policy Lab.

activate tells the router to activate this BGP peering inside this address family. At the time of writing, Cisco IOS activates IPv4 peers automatically inside IPv4 address families, but does not activate IPv6 peers inside IPv6 address families. Safest just to include the configuration in all templates.

Advertise your network

Use the *network* command to tell BGP which prefixes you want to announce.

On RX1 and RX2:

```
router bgp X0
 address-family ipv4
  network 100.68.X0.0 mask 255.255.255.0
 address-family ipv6
  network 2001:db8:X0::/48
```

Get the list of learned paths:

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

Do you see any paths? Why not?

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

On RX1 and RX2:

```
ip route 100.68.X0.0 255.255.255.0 null0
ipv6 route 2001:db8:X0::/48 null0
```

These are called a “pull up routes”

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

Q. Why are these routes needed?

STOP - Checkpoint One

All groups must finish this part before continuing. **Please do NOT continue until the instructor says so**.

Multihoming - eBGP Configuration

Connect to the NREN

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

NRENs: Use the configuration in the Appendix.

On R11:

```
interface GigabitEthernet1/0
 description P2P Link to NREN1
 ip address 100.101.1.2 255.255.255.252
 no ip directed-broadcast
 no ip redirects
 no ip proxy-arp
 ipv6 address 2001:11:0:10::1/127
 ipv6 nd prefix default no-advertise
 ipv6 nd ra suppress
 no shutdown
```

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

```
R11# ping 100.101.1.1
R11# ping 2001:11:0:10::0
```

Do some traceroutes to other networks again:

```
R11# traceroute 100.68.20.1
R11# traceroute 100.68.30.1
```

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.

BGP-peer with the NREN and the ISP

Configure eBGP sessions to the ISP and the NREN

On R11:

```
router bgp 10
```

```
 address-family ipv4
  neighbor 100.101.1.1 remote-as 101
  neighbor 100.101.1.1 description eBGP with NREN1
  neighbor 100.101.1.1 password NSRC-BGP
  neighbor 100.101.1.1 activate
!
 address-family ipv6
  neighbor 2001:11:0:10:: remote-as 101
  neighbor 2001:11:0:10:: description eBGP with NREN1
  neighbor 2001:11:0:10:: password NSRC-BGP
  neighbor 2001:11:0:10:: activate
```

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

On R12:

```
router bgp 10
 address-family ipv4
  neighbor 100.121.1.1 remote-as 121
  neighbor 100.121.1.1 description eBGP to ISP1
  neighbor 100.121.1.1 password NSRC-BGP
  neighbor 100.121.1.1 activate
!
 address-family ipv6
  neighbor 2001:18:0:10:: remote-as 121
  neighbor 2001:18:0:10:: description eBGP to ISP1
  neighbor 2001:18:0:10:: password NSRC-BGP
  neighbor 2001:18:0:10:: activate
```

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 neighbor 100.101.1.1 routes
R11# show bgp ipv6 uni neighbor 2001:11:0:10:: routes
```

Verify what you are advertising to the NREN:

```
R11# show ip bgp neighbor 100.101.1.1 advertised-routes
R11# show bgp ipv6 uni neighbor 2001:11:0:10:: advertised-routes
```

… and to the ISP:

```
R12# show ip bgp neighbor 100.121.1.1 advertised-routes
R12# sh bgp ipv6 uni neigh 2001:18:0:10:: 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 instructor if you need clarification.

Filter what you send and receive

Create prefix lists for your outbound filters. Outbound filters for RX1 and RX2 are the same:

```
ip prefix-list ASX0-out permit 100.68.X0.0/24
ipv6 prefix-list ASX0-v6-out permit 2001:db8:X0::/48
```

As you can see, we are only allowing our aggregate out to the Internet. Sending smaller prefixes (if we had any) serves no useful purpose at all.

Now create prefix lists for your inbound filters. Notice the descriptive names for the prefix lists.

On RX1:

```
ip prefix-list nren-in deny 100.68.X0.0/24 le 32
ip prefix-list nren-in permit 0.0.0.0/0 le 32
!
ipv6 prefix-list nren-v6-in deny 2001:db8:X0::/48 le 128
ipv6 prefix-list nren-v6-in permit ::/0 le 128
```

On RX2:

```
ip prefix-list isp-in deny 100.68.X0.0/24 le 32
ip prefix-list isp-in permit 0.0.0.0/0 le 32
!
ipv6 prefix-list isp-v6-in deny 2001:db8:X0::/48 le 128
ipv6 prefix-list isp-v6-in permit ::/0 le 128
```

Notice how we are matching the aggregate as well as all possible subnets of our address blocks. This protects against some other network announcing any of our prefixes (whatever subnet size) to us.

Now apply these prefix-lists to the BGP sessions with the ISP and the NREN.

On R11:

```
router bgp 10
 address-family ipv4
  neighbor 100.101.1.1 prefix-list nren-in in
  neighbor 100.101.1.1 prefix-list AS10-out out
!
 address-family ipv6
  neighbor 2001:11:0:10:: prefix-list nren-v6-in in
  neighbor 2001:11:0:10:: prefix-list AS10-v6-out out
```

On R12:

```
router bgp 10
 address-family ipv4
  neighbor 100.121.1.1 prefix-list isp-in in
  neighbor 100.121.1.1 prefix-list AS10-out out
!
 address-family ipv6
  neighbor 2001:18:0:10:: prefix-list isp-v6-in in
  neighbor 2001:18:0:10:: prefix-list AS10-v6-out out
```

Use the *BGP route refresh* capability to resend the information to the peer. Use the AS number of the peer rather than the IP address (much less typing for IPv6) in the route refresh command:

```
R11# clear ip bgp 101 out
R11# clear bgp ipv6 unicast 101 out
```

```
R12# clear ip bgp 121 out
R12# clear bgp ipv6 unicast 121 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.

Remove static routes

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 100.68.10.0 255.255.255.0 100.121.1.2
no ip route 100.68.20.0 255.255.255.0 100.121.1.6
no ip route 100.68.30.0 255.255.255.0 100.121.1.10
!
no ipv6 route 2001:db8:10::/48 2001:18:0:10::1
no ipv6 route 2001:db8:20::/48 2001:18:0:11::1
no ipv6 route 2001:db8:30::/48 2001:18:0:12::1
```

ISP2:

```
no ip route 100.68.40.0 255.255.255.0 100.122.1.2
no ip route 100.68.50.0 255.255.255.0 100.122.1.6
no ip route 100.68.60.0 255.255.255.0 100.122.1.10
!
no ipv6 route 2001:db8:40::/48 2001:19:0:10::1
no ipv6 route 2001:db8:50::/48 2001:19:0:11::1
no ipv6 route 2001:db8:60::/48 2001:19:0:12::1
```

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 from routers RX2:

On R12:

```
no ip route 0.0.0.0 0.0.0.0 100.121.1.1
no ipv6 route ::/0 2001:18:0:10::
```

You should be able to ping any other router now. If you can't, wait for other groups to finish. If other groups are finished, work them to work out what might be wrong. If you cannot see what is wrong, ask the instructors.

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

```
R11# traceroute 100.100.0.1
R11# traceroute 100.68.30.2
...
```

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

STOP - Checkpoint Two

All groups must finish this part before continuing. **Please do NOT continue until the instructor says so**.

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 is basically the same thing as connecting through a switch.

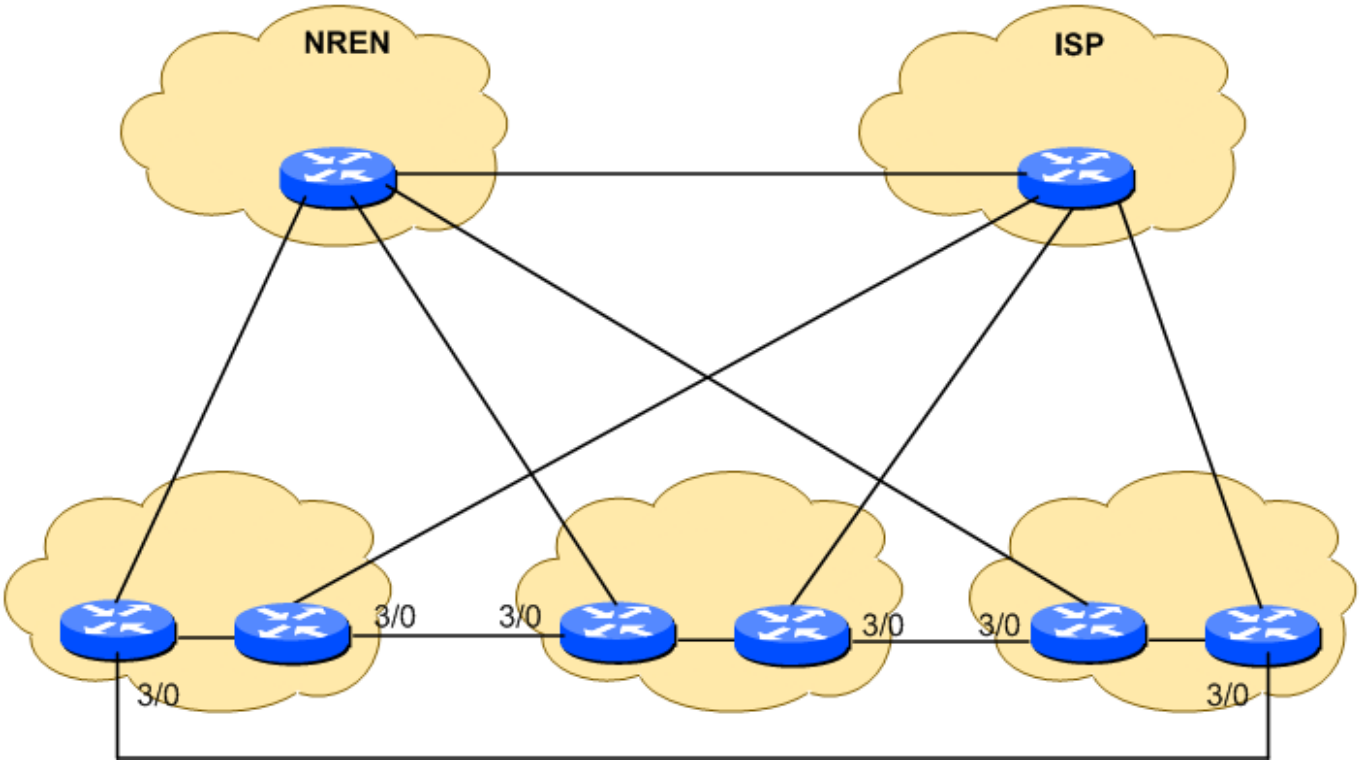

Connect to your neighbour AS

Configure a point-to-point link to your neighbour 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 100.68.10.21 255.255.255.252
 no ip directed-broadcast
 no ip redirects
 no ip proxy-arp
 ipv6 address 2001:db8:10:11::/127
 ipv6 nd prefix default no-advertise
 ipv6 nd ra suppress
 no shutdown
```

Configure prefix lists for your inbound filters

On R12:

```
ip prefix-list AS20-in permit 100.68.20.0/24
ipv6 prefix-list AS20-v6-in permit 2001:db8:20::/48
```

The equivalent needs to be done in R21.

Prefix lists for outbound filters should still exist from a previous step. You can verify like this:

```
R12#show ip prefix-list AS10-out
R12#show ipv6 prefix-list AS10-v6-out
```

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

On R12:

```
router bgp 10
 address-family ipv4
  neighbor 100.68.10.22 remote-as 20
  neighbor 100.68.10.22 description eBGP to AS20
  neighbor 100.68.10.22 password NSRC-BGP
  neighbor 100.68.10.22 prefix-list AS10-out out
  neighbor 100.68.10.22 prefix-list AS20-in in
  neighbor 100.68.10.22 activate
!
 address-family ipv6
  neighbor 2001:db8:10:11::1 remote-as 20
  neighbor 2001:db8:10:11::1 description eBGP to AS20
  neighbor 2001:db8:10:11::1 password NSRC-BGP
  neighbor 2001:db8:10:11::1 prefix-list AS10-v6-out out
  neighbor 2001:db8:10:11::1 prefix-list AS20-v6-in in
  neighbor 2001:db8:10:11::1 activate
```

The equivalent needs to be done for R21.

Verify that the sessions are up:

```
R12# show ip bgp summary
R12# show ipv6 bgp unicast summary
```

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

```
R12# show ip bgp neighbor 100.68.10.22 routes
R12# show bgp ipv6 unicast neighbor 2001:db8:10:11::1 routes
```

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.

Appendix A - RREN Configuration

```
hostname RREN
aaa new-model
aaa authentication login default local
aaa authentication enable default enable
username nsrc secret nsrc-PW
enable secret nsrc-EN
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
ipv6 unicast-routing
ipv6 cef
no ip source-route
no ipv6 source-route
!
interface Loopback0
 ip address 100.100.0.1 255.255.255.255
 ipv6 address 2001:10::1/128
!
interface GigabitEthernet1/0
 description P2P Link to NREN1
 ip address 100.100.1.1 255.255.255.252
 no ip directed-broadcast
 no ip redirects
 no ip proxy-arp
 ipv6 address 2001:10:0:10::/127
 ipv6 nd prefix default no-advertise
 ipv6 nd ra suppress
 no shutdown
!
interface GigabitEthernet2/0
 description P2P Link to NREN2
 ip address 100.100.1.5 255.255.255.252
 no ip directed-broadcast
 no ip redirects
 no ip proxy-arp
 ipv6 address 2001:10:0:11::/127
 ipv6 nd prefix default no-advertise
 ipv6 nd ra suppress
 no shutdown
!
interface GigabitEthernet3/0
```

```
 description Link to IXP
 ip address 100.127.1.3 255.255.255.0
 no ip directed-broadcast
 no ip redirects
 no ip proxy-arp
 ipv6 address 2001:db8:ffff:1::3/64
 ipv6 nd prefix default no-advertise
 ipv6 nd ra suppress
 no shutdown
!
router bgp 100
 bgp log-neighbor-changes
 bgp deterministic-med
 no bgp default ipv4-unicast
 address-family ipv4
  distance bgp 200 200 200
  network 100.100.0.0 mask 255.255.0.0
  neighbor 100.100.1.2 remote-as 101
  neighbor 100.100.1.2 description eBGP to NREN1 (AS101)
  neighbor 100.100.1.2 password NSRC-BGP
  neighbor 100.100.1.2 activate
  neighbor 100.100.1.6 remote-as 102
  neighbor 100.100.1.6 description eBGP to NREN2 (AS102)
  neighbor 100.100.1.6 password NSRC-BGP
  neighbor 100.100.1.6 activate
  neighbor 100.127.1.1 remote-as 121
  neighbor 100.127.1.1 description eBGP to ISP1 (AS121)
  neighbor 100.127.1.1 password NSRC-BGP
  neighbor 100.127.1.1 activate
  neighbor 100.127.1.2 remote-as 122
  neighbor 100.127.1.2 description eBGP to ISP2 (AS122)
  neighbor 100.127.1.2 password NSRC-BGP
  neighbor 100.127.1.2 activate
 !
 address-family ipv6
  distance bgp 200 200 200
  network 2001:10::/32
  neighbor 2001:10:0:10::1 remote-as 101
  neighbor 2001:10:0:10::1 description eBGP to NREN1 (AS101)
  neighbor 2001:10:0:10::1 password NSRC-BGP
  neighbor 2001:10:0:10::1 activate
  neighbor 2001:10:0:11::1 remote-as 102
  neighbor 2001:10:0:11::1 description eBGP to NREN2 (AS102)
  neighbor 2001:10:0:11::1 password NSRC-BGP
  neighbor 2001:10:0:11::1 activate
  neighbor 2001:db8:ffff:1::1 remote-as 121
  neighbor 2001:db8:ffff:1::1 description eBGP to ISP1 (AS121)
  neighbor 2001:db8:ffff:1::1 password NSRC-BGP
  neighbor 2001:db8:ffff:1::1 activate
  neighbor 2001:db8:ffff:1::2 remote-as 122
  neighbor 2001:db8:ffff:1::2 description eBGP to ISP2 (AS122)
```

```
  neighbor 2001:db8:ffff:1::2 password NSRC-BGP
  neighbor 2001:db8:ffff:1::2 activate
!
ip route 100.100.0.0 255.255.0.0 null0
ipv6 route 2001:10::/32 null0
```

Appendix B - NREN1 Sample Configuration

The sample configuration for NREN1 is below. NREN2's configuration will be very similar, so modify accordingly.

```
hostname NREN1
aaa new-model
aaa authentication login default local
aaa authentication enable default enable
username nsrc secret nsrc-PW
enable secret nsrc-EN
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
ipv6 unicast-routing
ipv6 cef
no ip source-route
no ipv6 source-route
!
interface Loopback0
 ip address 100.101.0.1 255.255.255.255
 ipv6 address 2001:11::1/128
!
interface GigabitEthernet1/0
 description P2P Link to RREN
 ip address 100.100.1.2 255.255.255.252
 no ip directed-broadcast
 no ip redirects
 no ip proxy-arp
 ipv6 address 2001:10:0:10::1/127
 ipv6 nd prefix default no-advertise
 ipv6 nd ra suppress
 no shutdown
!
interface GigabitEthernet2/0
 description P2P Link to ISP1
 ip address 100.101.2.1 255.255.255.252
 no ip directed-broadcast
```

```
 no ip redirects
 no ip proxy-arp
 ipv6 address 2001:11:0:20::/127
 ipv6 nd prefix default no-advertise
 ipv6 nd ra suppress
 no shutdown
!
! (repeat for Group 2 and Group 3 using Gig4/0 and 5/0)
interface GigabitEthernet3/0
 description P2P Link to R11
 ip address 100.101.1.1 255.255.255.252
 no ip directed-broadcast
 no ip redirects
 no ip proxy-arp
 ipv6 address 2001:11:0:10::/127
 ipv6 nd prefix default no-advertise
 ipv6 nd ra suppress
 no shutdown
!
! inbound filter for AS10 - repeat for AS20 and AS30
ip prefix-list AS10-in permit 100.68.10.0/24
ipv6 prefix-list AS10-v6-in permit 2001:db8:10::/48
!
router bgp 101
 bgp log-neighbor-changes
 bgp deterministic-med
 no bgp default ipv4-unicast
 address-family ipv4
  distance bgp 200 200 200
  network 100.101.0.0 mask 255.255.0.0
  neighbor 100.101.1.2 remote-as 10
  neighbor 100.101.1.2 description eBGP with AS10
  neighbor 100.101.1.2 password NSRC-BGP
  neighbor 100.101.1.2 prefix-list AS10-in in
  neighbor 100.101.1.2 activate
  (repeat for AS20 and AS30)
  neighbor 100.101.2.2 remote-as 121
  neighbor 100.101.2.2 description eBGP with ISP1 (AS121)
  neighbor 100.101.2.2 password NSRC-BGP
  neighbor 100.101.2.2 activate
  neighbor 100.100.1.1 remote-as 100
  neighbor 100.100.1.1 description eBGP with RREN (AS100)
  neighbor 100.100.1.1 password NSRC-BGP
  neighbor 100.100.1.1 activate
 !
 address-family ipv6
  distance bgp 200 200 200
  network 2001:11::/32
  neighbor 2001:11:0:10::1 remote-as 10
  neighbor 2001:11:0:10::1 description eBGP with AS10
  neighbor 2001:11:0:10::1 password NSRC-BGP
```

```
  neighbor 2001:11:0:10::1 prefix-list AS10-v6-in in
  neighbor 2001:11:0:10::1 activate
  (repeat for AS20 and AS30)
  neighbor 2001:11:0:20::1 remote-as 121
  neighbor 2001:11:0:20::1 description eBGP with ISP1 (AS121)
  neighbor 2001:11:0:20::1 password NSRC-BGP
  neighbor 2001:11:0:20::1 activate
  neighbor 2001:10:0:10:: remote-as 100
  neighbor 2001:10:0:10:: description eBGP with RREN (AS100)
  neighbor 2001:10:0:10:: password NSRC-BGP
  neighbor 2001:10:0:10:: activate
!
ip route 100.101.0.0 255.255.0.0 null0
ipv6 route 2001:11::/32 null0
```

Appendix C - ISP1 Sample Configuration

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

```
interface GigabitEthernet2/0
 description P2P Link to NREN1
 ip address 100.101.2.2 255.255.255.252
 no ip directed-broadcast
 no ip redirects
 no ip proxy-arp
 ipv6 address 2001:11:0:20::1/127
 ipv6 nd prefix default no-advertise
 ipv6 nd ra suppress
 no shutdown
!
! (filters for Group 1 - repeat for Group 2 and 3)
ip prefix-list AS10-in permit 100.68.10.0/24
ipv6 prefix-list AS10-v6-in permit 2001:db8:10::/48
!
router bgp 121
 bgp log-neighbor-changes
 no bgp default ipv4-unicast
 bgp deterministic-med
 address-family ipv4
  distance bgp 200 200 200
  network 100.121.0.0 mask 255.255.0.0
  neighbor 100.121.1.2 remote-as 10
  neighbor 100.121.1.2 description eBGP with AS10
  neighbor 100.121.1.2 password NSRC-BGP
  neighbor 100.121.1.2 prefix-list AS10-in in
  neighbor 100.121.1.2 activate
  (repeat for AS20 and AS30)
  neighbor 100.101.2.1 remote-as 101
  neighbor 100.101.2.1 description eBGP with NREN1 (AS101)
```

```
  neighbor 100.101.2.1 password NSRC-BGP
  neighbor 100.101.2.1 activate
  neighbor 100.127.1.2 remote-as 122
  neighbor 100.127.1.2 description eBGP with ISP2 (AS122)
  neighbor 100.127.1.2 password NSRC-BGP
  neighbor 100.127.1.2 activate
  neighbor 100.127.1.3 remote-as 100
  neighbor 100.127.1.3 description eBGP to RREN (AS100)
  neighbor 100.127.1.3 password NSRC-BGP
  neighbor 100.127.1.3 activate
 !
 address-family ipv6
  distance bgp 200 200 200
  network 2001:18::/32
  neighbor 2001:18:0:10::1 remote-as 10
  neighbor 2001:18:0:10::1 description eBGP with AS10
  neighbor 2001:18:0:10::1 password NSRC-BGP
  neighbor 2001:18:0:10::1 prefix-list AS10-v6-in in
  neighbor 2001:18:0:10::1 activate
  (repeat for AS20 and AS30)
  neighbor 2001:11:0:20:: remote-as 101
  neighbor 2001:11:0:20:: description eBGP with NREN1 (AS101)
  neighbor 2001:11:0:20:: password NSRC-BGP
  neighbor 2001:11:0:20:: activate
  neighbor 2001:db8:ffff:1::2 remote-as 122
  neighbor 2001:db8:ffff:1::2 description eBGP with ISP2 (AS122)
  neighbor 2001:db8:ffff:1::2 password NSRC-BGP
  neighbor 2001:db8:ffff:1::2 activate
  neighbor 2001:db8:ffff:1::3 remote-as 100
  neighbor 2001:db8:ffff:1::3 description eBGP with RREN (AS100)
  neighbor 2001:db8:ffff:1::3 password NSRC-BGP
  neighbor 2001:db8:ffff:1::3 activate
!
ip route 100.121.0.0 255.255.0.0 null0
ipv6 route 2001:18::/32 null0
```

Appendix D - Complete Lab Diagram & Address Plan

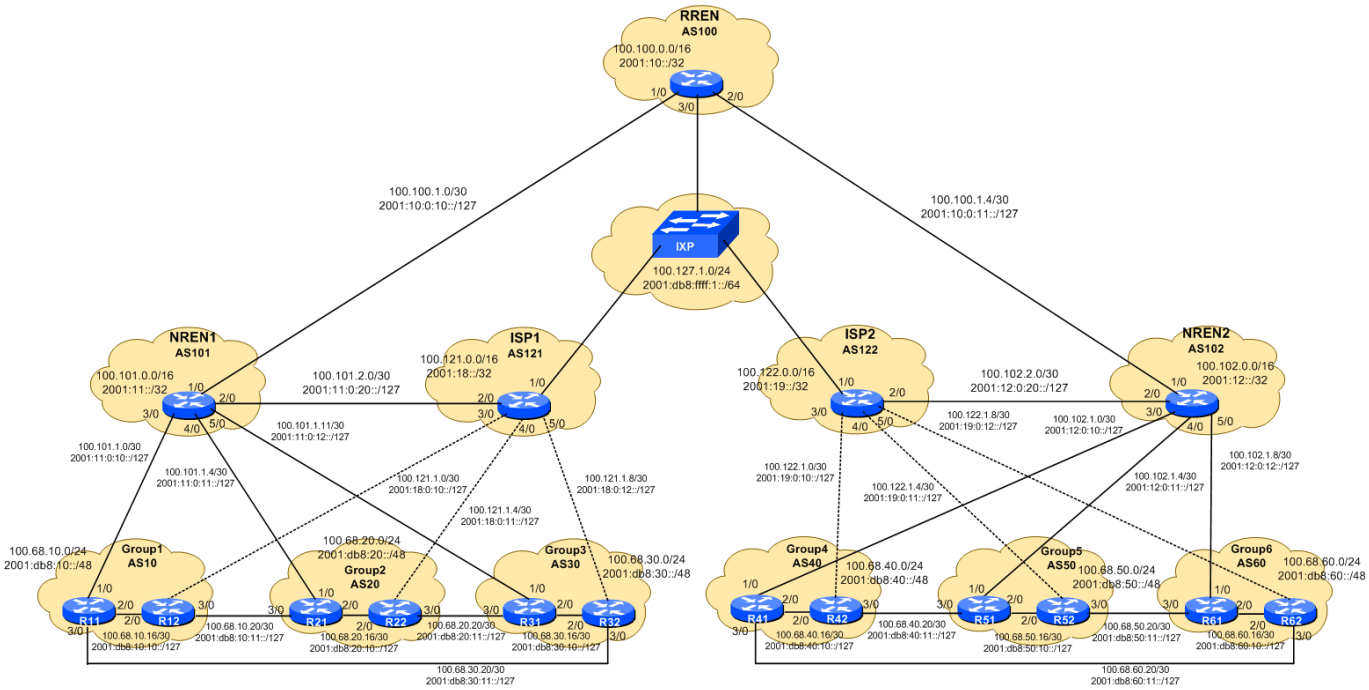

From:
https://wiki.lpnz.org/ - **Workshops**

Permanent link:
https://wiki.lpnz.org/doku.php?id=2015:ubuntunet-bgp:2-lab-bgp-basic

Last update: **2015/11/17 14:36**