BGP Exercise

Campus Network Design Workshop

Contents

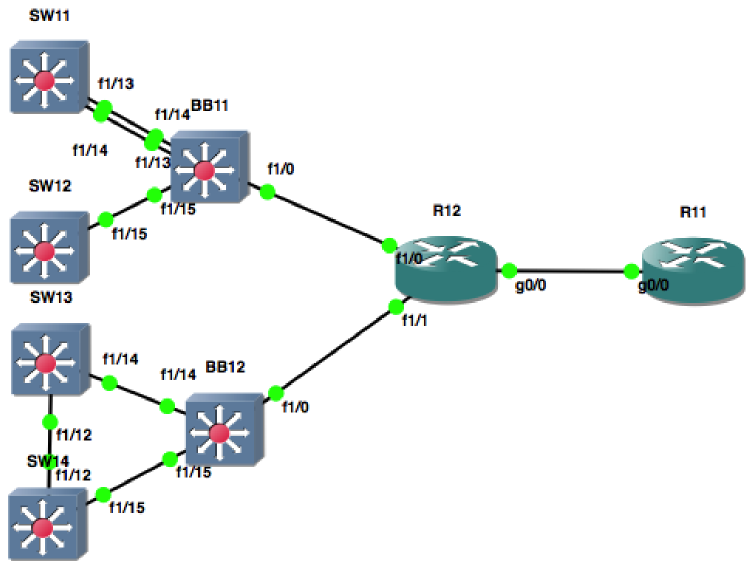

Figure 1: Physical Topology

1 Introduction

The purpose of this exercise is to learn how to configure BGP on a Cisco router so that it can exchange network reachability information with an external peer.

We've already configured two sets of switches with VLANS to represent Engineering and Computer Science (ECS) and the Library (LIB) on a campus and connected those to the campus Core (or backbone) routers.

All participants will work within a group as a team. Each group has two routers and six switches to work with. There is a certain dependency between the labs as the exercises progress. Make sure to maintain your configuration unless otherwise instructed. All exercises will use a common IP addressing scheme and network topology. As you go through the exercises all the examples are given from the point of view of R11, the border router in group 1.

Make sure to take the examples and adapt them to your own router, network topology and addressing scheme.

1.1 Router types used in the lab

Cisco 7206 VXR

1.2 Address Space Allocation

Your address space was allocated in the Layer-2 and OSPF labs. We've added details of your AS number:

Group	IPv4 Block	IPv6 Block	AS number
1	10.110.0.0/16	fd00:110::/32	65001
2	10.120.0.0/16	fd00:120::/32	65002
3	10.130.0.0/16	fd00:130::/32	65003
4	10.140.0.0/16	fd00:140::/32	65004
5	10.150.0.0/16	fd00:150::/32	65005
5	10.160.0.0/16	fd00:160::/32	65006

Your upstream provider will be NSRC and their AS number is 65000. They have a BGP router that you will peer with and the IP addresses are shown below. They have allocated IP addresses for your router's external interface as follows:

Group	IPv4 Address	IPv6 Address
1	10.10.0.101/24	fd00:0:1:1::101/64
2	10.10.0.102/24	fd00:0:1:1::102/64
3	10.10.0.103/24	fd00:0:1:1::103/64
4	10.10.0.104/24	fd00:0:1:1::104/64
5	10.10.0.105/24	fd00:0:1:1::105/64
6	10.10.0.106/24	fd00:0:1:1::106/64
NSRC	10.10.0.254/24	fd00:0:1:1::254/64

1.2.1 Using private address space

We are using private address space for the exercises - if you are planning to peer using BGP with your upstream provider(s) you will need to have IPv4 and IPv6 public address space and an AS number. You can get these from your local Regional Internet Registry.

Region	RIR
Africa	AfriNIC
Europe	RIPE
Asia Pacific	APNIC
Latin America	LACNIC
North America	ARIN

2 Exercises

2.1 Basic Router Configuration

Your border router should be configured as though you've completed the Layer-2 lab and the OSPF lab. For example, Appendix A shows the confguration for R11

2.2 Configuring BGP

2.2.1 Configure a BGP process and establish a peering session with the NSRC router.

Our first thing to do is create a BGP router process. This is similar to creating an OSPF process which we did in the last exercise. We use the AS number allocated above. We'll also use the IPv4 loopback address we configured earlier as a unique identifier for BGP.

The last three lines define a new neighbor (note the US spelling). We specify their AS number, a description and a password for the session. (Don't use a password like this on a live network!)

R11:

```
router bgp 65001
bgp router-id 10.110.10.1
 neighbor 10.10.0.1 remote-as 65000
 neighbor 10.10.0.1 description NSRC
 neighbor 10.10.0.1 password nsrc
```

At this point we should be able to exit configuration mode and check if our BGP session is up using the command:

```
R11#show ip bgp summary
BGP router identifier 10.10.254.1, local AS number 65001
BGP table version is 4, main routing table version 4
1 network entries using 136 bytes of memory
1 path entries using 56 bytes of memory
1/1 BGP path/bestpath attribute entries using 128 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 320 total bytes of memory
BGP activity 1/0 prefixes, 2/1 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.10.0.1       4        65000      18      20        4    0    0 00:13:38        1
```

The last column should show **at least** one prefix being received from NSRC.

2.2.2 Check the routes we are receiving from NSRC

We can check this more closely:

```
R11#sh ip bgp neighbors 10.10.0.1 routes
BGP table version is 6, local router ID is 10.10.254.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, x best-external, f RT-Filt
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 10.10.0.0/16     10.10.0.1                0             0 65000 i

Total number of prefixes 1
```

If you see more prefixes they will be those of others in your class who are a little ahead of you.

2.2.3 Configure the routes we are going to send to NSRC

```
R11#sh ip bgp neighbors 10.10.0.1 advertised-routes

Total number of prefixes 0
```

We need to tell the BGP process which routes it should send and we do this using a simple static declaration. We don't need to pass on all our subnet details to the outside world so we use the 'aggregate-address' command to limit things to our campus block.

```
router bgp 65001
network 10.110.0.0 mask 255.255.0.0
aggregate-address 10.110.0.0 255.255.0.0
```

Now we can check what we're sending again:

```
R11#sh ip bgp neighbors 10.10.0.1 advertised-routes
BGP table version is 7, local router ID is 10.10.254.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, x best-external, f RT-Filt
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
   Network          Next Hop            Metric LocPrf Weight Path
*> 10.110.0.0/16    0.0.0.0                  0         32768 i

Total number of prefixes 1
```

2.2.4 What happens if OSPF stops advertising our subnets from R12?

Our edge router will stop advertising out network! We can prevent that by adding a static route for our network block that discards traffic. The OSPF process will override this static route. Don't worry - if the OSPF process isn't working your campus will be off line anyway.

```
R11(config)# ip route 10.110.0.0 255.255.0.0 Null 0  200
```

2.3 Testing the connection

Finally we should be able to test a connection to your Campus network from outside. Try ping these addresses from your laptop or workstation:

```
ping 10.1X0.0.2
ping 10.1X0.64.2
ping 10.1X0.65.2
ping 10.1X0.254.2
ping 10.1X0.74.2
ping 10.1X0.75.2
ping 10.1X0.255.2
```

- Can you reach your switches?
- Can you telnet to them from outside the campus?
- Is this a good idea?

3 Further BGP work

We have only configured very basic IPv4 peering. IPv6 is a little more involved but the principles are just the same.

We have not done anything to create or enforce any policy around the routes we send and receive. This is almost always done but is outside of the scope of this workshop.

You should not use a configuration as basic as this in a live network.

4 Appendix A

4.1 Base configuration

```
hostname R11
!
!
aaa new-model
aaa authentication login default local
aaa authentication enable default enable
aaa session-id common
!
ip cef
!
no ip domain-lookup

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
ipv6 unicast-routing
!
interface loopback 0
 ip address 10.110.10.1 255.255.255.255
 ipv6 address fd00:10:a::1/128
!
interface GigabitEthernet0/0
 ip address 10.110.1.1 255.255.255.0
 description Link to Core
 ipv6 address fd00:0:1:1::1/64
 no ip redirects
 no ip directed-broadcast
 no ip proxy-arp
 no shutdown
!
interface GigabitEthernet1/0
 ip address 10.10.0.101 255.255.255.0
```

```
 description Link to Internet (NSRC Backbone)
 ipv6 address fd00:0:1:1::101/64
 no ip redirects
 no ip directed-broadcast
 no ip proxy-arp
 no shutdown
!
router ospf 10
 log-adjacency-changes
 passive-interface default
 area 0 authentication message-digest
 no passive-interface GigabitEthernet0/0
 auto-cost reference-bandwidth 1000
!
ipv6 router ospf 10
 log-adjacency-changes
 passive-interface default
 no passive-interface GigabitEthernet0/0
 area 0 authentication ipsec
 spi 256 md5 0123456789ABCDEF0123456789ABCDEF
 auto-cost reference-bandwidth 1000
!
interface Loopback0
 ip ospf 10 area 0
 ipv6 ospf 10 area 0
!
interface GigabitEthernet0/0
 ip ospf 10 area 0
 ip ospf authentication-key nsrc
 ipv6 ospf 10 area 0
```