

BGP Lab

Introduction


The purpose of this exercise is to learn how to configure BGP between the Border and Core routers.

In the last routing lab we configured OSPF between the Border and Core routers. The introduction of BGP allows simpler announcement of address space of the local autonomous network, and allows future possibility of the local network connecting to two different network operators (for example a commercial ISP and the REN).

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

Lab Layout

As a reminder, the following diagram shows the layout of the devices and all the links for each campus:



Our campus network consists of two routers, **bdr1-campusX** and **core1-campusX** as well as six switches.

The following table shows the connections between each device in the campus:

Device	Interface	Remote Device	Remote Interface
dist1-bY-campusX	FastEthernet1/12	edge1-bY-campusX	FastEthernet1/14
	FastEthernet1/13	edge1-bY-campusX	FastEthernet1/15
	FastEthernet1/14	edge2-bY-campusX	FastEthernet1/15
core1-campusX	FastEthernet1/0	bdr1-campusX	FastEthernet0/1
	FastEthernet1/1	dist1-b1-campusX	FastEthernet1/15
	FastEthernet1/2	dist1-b2-campusX	FastEthernet1/15
	FastEthernet1/15	pc1-campusX	

Replace **Y** with your building number and **X** with your campus number.

Accessing the Lab

The Workshop Instructors will let you know what the lab environment is. It will either be run on a Virtual Platform, or on real physical switches provided in the Training Room.

Refer to the **correct** document below for information about logging into the devices that have been assigned to you:

VIRTUAL ENVIRONMENT: [Lab Access Instructions - Virtualised Platform](#)

PHYSICAL HARDWARE: [Lab Access Instructions - Physical Hardware](#)

Summary Commands

During all exercises, verify the output of the following commands:

```
show ip bgp sum           : show status of the IPv4 BGP peering
show bgp ipv6 unicast sum : show status of the IPv6 BGP peering
show ip bgp               : show the IPv4 BGP table
show bgp ipv6 unicast     : show the IPv6 BGP table
```

Dynamic Routing with BGP

We will now configure a new BGP routing process for both IPv4 and IPv6. As with OSPF, BGP automatically creates the router-id it needs from the IPv4 address configured on the Loopback interface. However, we will bypass this automatic process and instead manually configure the router-id but still using the IPv4 address on the Loopback interface.

On the Core router:

```
router bgp X0
  no bgp default ipv4-unicast
  bgp deterministic-med
  bgp router-id 100.68.X.242
```

```
!  
address-family ipv4  
  neighbor 100.68.X.241 remote-as X0  
  neighbor 100.68.X.241 update-source loopback 0  
  neighbor 100.68.X.241 description iBGP with Border Router  
  network 100.68.X.0 mask 255.255.255.0  
  network 172.2X.0.0 mask 255.255.0.0  
  distance bgp 200 200 200  
!  
address-family ipv6  
  neighbor 2001:DB8:X:2::241 remote-as X0  
  neighbor 2001:DB8:X:2::241 update-source loopback0  
  neighbor 2001:DB8:X:2::241 description iBGP with Border Router  
  network 2001:DB8:X::/48  
  distance bgp 200 200 200  
!  
ip route 100.68.X.0 255.255.255.0 Null0  
ip route 172.2X.0.0 255.255.0.0 Null0  
ipv6 route 2001:DB8:X::/48 Null0
```

Here is an explanation of each command used above - some were not mentioned in the presentation.

- “bgp default ipv4-unicast” means that IOS BGP must not assume that all neighbours will exchange IPv4 prefixes. IPv6 neighbours do not normally exchange IPv4 prefixes!
- “bgp deterministic-med” is a best practice which fixes BGP’s path selection in the instance there are multiple paths with multiple MEDs to the same destination
- “distance bgp 200 200 200” means that we will set eBGP’s administrative distance to be the same as that of iBGP - another best practice to avoid accidents.

On the Border router:

```
router bgp X0  
  no bgp default ipv4-unicast  
  bgp deterministic-med  
  bgp router-id 100.68.X.241  
!  
address-family ipv4  
  neighbor 100.68.X.242 remote-as X0  
  neighbor 100.68.X.242 update-source loopback 0  
  neighbor 100.68.X.242 description iBGP with Core Router  
  distance bgp 200 200 200  
!  
address-family ipv6  
  neighbor 2001:DB8:X:2::242 remote-as X0  
  neighbor 2001:DB8:X:2::242 update-source loopback0  
  neighbor 2001:DB8:X:2::242 description iBGP with Core Router  
  network 2001:DB8:X::/48  
  distance bgp 200 200 200  
!
```

STOP - Checkpoint.

```
show ip bgp summary      : show neighbours
show ip route            : show routes in routing table
show ip bgp              : shows general BGP information
```

The status commands for IPv6 are very similar: simply replace “ip bgp” in the above show commands with “bgp ipv6”.

Question: How many routes do you see in BGP for both IPv4 and IPv6?

Configuring BGP with the NREN

The final task for this lab is to set up BGP between the Border router and the NREN router. To do this we are setting up eBGP on the border router, peering on the point to point link address of the link to the NREN.

As was mentioned in the presentation, BGP is very chatty. In the absence of any filters, it will send all prefixes on the router, and will accept all prefixes it hears from a neighbour. Before we connect any BGP speaking router to the public Internet we must prepare filters for it.

```
ip prefix-list NREN-in permit 0.0.0.0/0
!
ip prefix-list CampusX-out permit 100.68.X.0/24
!
ipv6 prefix-list NREN-v6in permit ::/0
!
ipv6 prefix-list CampusX-v6out permit 2001:DB8:X::/48
```

These filters in order, do the following:

- Only allow IPv4 default route in from the NREN
- Only allow the Campus public IPv4 address block out to the NREN
- Only allow IPv6 default route in from the NREN
- Only allow the Campus IPv6 address block out to the NREN

Now the filters are prepared we can set up BGP with the NREN. Here is a configuration example for the Border router - modify this example to suit your Campus. The instructors will have already configured BGP on the NREN router to talk to the campus border router.

```
router bgp X0
!
address-family ipv4
  neighbor 100.68.y.z remote-as 65534
  neighbor 100.68.y.z description eBGP with NREN
  neighbor 100.68.y.z prefix-list NREN-in in
  neighbor 100.68.y.z prefix-list CampusX-out out
!
```

```
address-family ipv6
  neighbor 2001:DB8:0:X:: remote-as 65534
  neighbor 2001:DB8:0:X:: description eBGP with NREN
  neighbor 2001:DB8:0:X:: prefix-list NREN-v6in in
  neighbor 2001:DB8:0:X:: prefix-list CampusX-v6out out
!
```

Check the BGP table. Do you see the default route learned from the NREN?

If you do, now is a good time to remove the static default currently pointing to the NREN - we set this up in an earlier exercise.

```
no ip route 0.0.0.0 0.0.0.0 <nren point to point IPv4 address>
no ipv6 route ::/0 <nren point to point IPv6 address>
```

Now look at the BGP table again. Do you see any difference from previously?

Also, if you now look at the Routing Table and compare with the BGP table, what do you also notice.

Here is an example from a previous Campus Network Design Workshop. How does this BGP table and Routing table compare with the entries you see?

Routing table:

```
bdr1-campus1#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static
route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       + - replicated route, % - next hop override
```

Gateway of last resort is 100.68.0.1 to network 0.0.0.0

```
B*    0.0.0.0/0 [200/0] via 100.68.0.1, 00:11:54
      100.0.0.0/8 is variably subnetted, 8 subnets, 4 masks
C      100.68.0.0/30 is directly connected, FastEthernet0/0
L      100.68.0.2/32 is directly connected, FastEthernet0/0
B      100.68.1.0/24 [200/0] via 100.68.1.242, 00:11:55
C      100.68.1.0/28 is directly connected, FastEthernet0/1
L      100.68.1.1/32 is directly connected, FastEthernet0/1
O      100.68.1.128/28 [110/2] via 100.68.1.2, 00:12:14, FastEthernet0/1
C      100.68.1.241/32 is directly connected, Loopback0
O      100.68.1.242/32 [110/2] via 100.68.1.2, 00:12:14, FastEthernet0/1
      172.21.0.0/16 is variably subnetted, 7 subnets, 2 masks
B      172.21.0.0/16 [200/0] via 100.68.1.242, 00:11:55
```

```

0      172.21.10.0/24 [110/2] via 100.68.1.2, 00:12:14, FastEthernet0/1
0      172.21.11.0/24 [110/2] via 100.68.1.2, 00:12:14, FastEthernet0/1
0      172.21.12.0/24 [110/2] via 100.68.1.2, 00:12:14, FastEthernet0/1
0      172.21.20.0/24 [110/2] via 100.68.1.2, 00:12:14, FastEthernet0/1
0      172.21.21.0/24 [110/2] via 100.68.1.2, 00:12:14, FastEthernet0/1
0      172.21.22.0/24 [110/2] via 100.68.1.2, 00:12:14, FastEthernet0/1

```

and BGP table:

```

bdr1-campus1#sh ip bgp
BGP table version is 4, local router ID is 100.68.1.241
Status codes: s suppressed, d damped, h history, * valid, > best, i -
internal,
                r RIB-failure, S Stale, m multipath, b backup-path, f RT-
Filter,
                x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	0.0.0.0	100.68.0.1			0	65534 i
*>i	100.68.1.0/24	100.68.1.242	0	100	0	i
*>i	172.21.0.0	100.68.1.242	0	100	0	i

From:

<https://workshops.nsrc.org/dokuwiki/> - Workshops

Permanent link:

<https://workshops.nsrc.org/dokuwiki/2016/nsrc-tein-bdren/bgp>

Last update: **2016/10/13 00:11**

