

Internal BGP Lab

Introduction

The purpose of this exercise is to:

- Set up Internal BGP (iBGP) within the each Group autonomous system to carry routing information within the AS.

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 all other routers in the lab 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 BX1, BX2 and CX1:

```
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

As you will remember from the presentation, iBGP needs to be fully meshed. We have three routers in our Group, so we need to set up iBGP between all three - which means that each router will have two iBGP neighbours.

First, make sure that you can ping the other routers using their loopback addresses. If you cannot ping them, there is no possibility that the BGP session will come up!

BX1 peering with BX2:

```
router bgp X0
  address-family ipv4
    neighbor 100.68.X.2 remote-as X0
    neighbor 100.68.X.2 update-source loopback 0
    neighbor 100.68.X.2 description iBGP with BX2
    neighbor 100.68.X.2 password NSRC-BGP
    neighbor 100.68.X.2 next-hop-self
    neighbor 100.68.X.2 send-community
    neighbor 100.68.X.2 activate
  !
  address-family ipv6
    neighbor 2001:db8:X::2 remote-as X0
    neighbor 2001:db8:X::2 update-source loopback 0
    neighbor 2001:db8:X::2 description iBGP with BX2
    neighbor 2001:db8:X::2 password NSRC-BGP
    neighbor 2001:db8:X::2 next-hop-self
    neighbor 2001:db8:X::2 send-community
    neighbor 2001:db8:X::2 activate
```

You need a similar configuration for BX1 peering with CX1. What needs to change from the above configuration example?

BX2 peering with BX1:

```
router bgp X0
 address-family ipv4
  neighbor 100.68.X.1 remote-as X0
  neighbor 100.68.X.1 update-source loopback 0
  neighbor 100.68.X.1 description iBGP with BX1
  neighbor 100.68.X.1 password NSRC-BGP
  neighbor 100.68.X.1 next-hop-self
  neighbor 100.68.X.1 send-community
  neighbor 100.68.X.1 activate
!
 address-family ipv6
  neighbor 2001:db8:X::1 remote-as X0
  neighbor 2001:db8:X::1 update-source loopback 0
  neighbor 2001:db8:X::1 description iBGP with BX1
  neighbor 2001:db8:X::1 password NSRC-BGP
  neighbor 2001:db8:X::1 next-hop-self
  neighbor 2001:db8:X::1 send-community
  neighbor 2001:db8:X::1 activate
!
```

You need a similar configuration for BX2 peering with CX1. What needs to change from the above configuration example?

Also, what do you think the configuration for CX1 would look like? Is it similar? What is different? The instructors will discuss this in class during the exercise.

There is a Cisco IOS configuration scaling feature called a *peer-group*. Many network operators using Cisco IOS routers use *peer-group* to scale their BGP configurations.

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. We will do this on all routers in our group (although strictly it is sufficient to do this on CX1 given that if the Core Router goes down, the campus has no connectivity to the Internet at all).

On BX1, BX2 and CX1:

```
router bgp X0
  address-family ipv4
    network 100.68.X.0 mask 255.255.255.0
  address-family ipv6
    network 2001:db8:X::/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 BX1, BX2 and CX1:

```
ip route 100.68.X.0 255.255.255.0 null0
ipv6 route 2001:db8:X::/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?

From:

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

Permanent link:

<https://workshops.nsrc.org/dokuwiki/2016/nsrc-asren-bgp/2-lab-bgp-ibgp>

Last update: **2016/09/07 08:37**

