Experiment 1d - Setup iBGP

DE-CIX Academy

Version 1.4 DENOG

1 Purpose

iBGP is simply BGP within one Autonomous System. It is used to redistribute prefixes received from other providers and to redistribute your own or your customers prefixes.

2 Network Setup

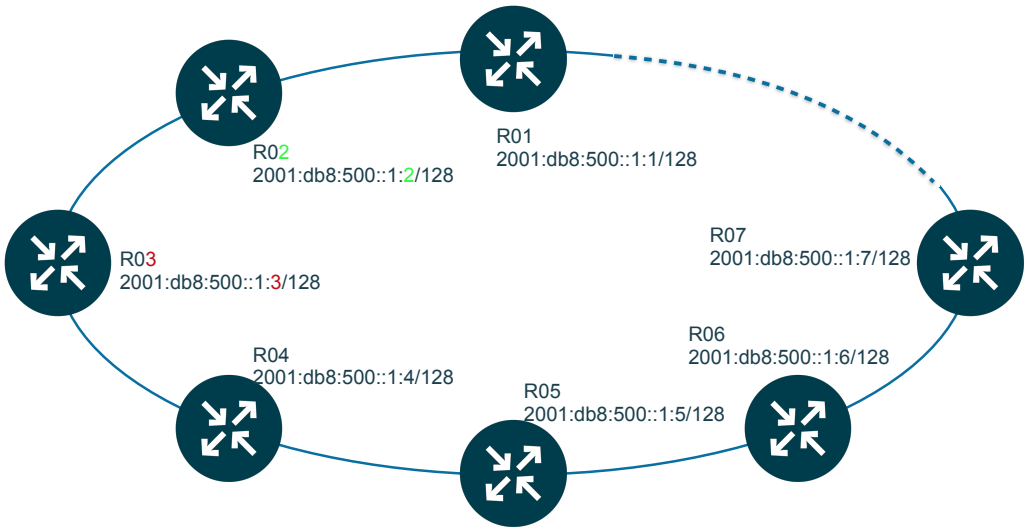


Figure 1: Network Setup

Figure 1 shows the network topology for this experiment. All devices are connected in a ring like structure, each device has two neighbors. Loopback addresses are distributed via an IGP (OSPF or IS-IS).

3 Slides

Setup iBGP for IPv4

- We use *peer-groups* for configuration
- All *common* config statements go into the peer-group
- All individual statements to into the peer entry

```
router bgp 64500
 no bgp default ipv4-unicast

 neighbor internal-v4 peer-group
 neighbor internal-v4 remote-as 64500
 neighbor internal-v4 update-source loopback0
 address-family ipv4 unicast
   neighbor internal-v4 activate
 exit-address-family

 neighbor 172.16.1.1 peer-group internal-v4
 ...
```

- Configure only the first neighbor - others will be added automatically
- We change the default so IPv4 is not automatically activated

Show commands:

- show bgp ipv4 summary
- show bgp ipv4 neighbors
- show bgp ipv4 neighbor 172.16.1.yy

Setup iBGP for IPv6

- We again use a *peer-group*
- You need separate peer-groups for IPv4 and IPv6
- All *common* config statements go into the peer-group
- All individual statements to into the peer entry

```
router bgp 64500
 ! commented out as we already did that last slide
 ! no bgp default ipv4-unicast

 neighbor internal-v6 peer-group
 neighbor internal-v6 remote-as 64500
 neighbor internal-v6 update-source loopback0
 address-family ipv6 unicast
  neighbor internal-v6 activate
 exit-address-family

 neighbor 2001:db8:500::1:1 peer-group internal-v6
 ...
```

- See the *address-family* entry?
- This is about what is being *transported* via BGP
- We have to activate IPv6
- and deactivate (*no bgp...*) IPv4 (we simply change the default)

Show commands:

- show bgp ipv6 summary
- show bgp ipv6 neighbors
- show bgp ipv6 neighbor 2001:db8:500::1:yy

4 Setup iBGP

4.1 Tasks:

- Define iBGP peer groups for IPv4 and IPv6
- Configure all parameters needed in the peer group
- Set up iBGP sessions to all other routers on IPv4 and IPv6 using the peer groups
- Check if the sessions are up

4.2 To set up iBGP you need to:

- start BGP using your AS number
- define one iBGP peer group for IPv4 and one for IPv6
- configure all iBGP neighbors

Information you need:

- Your AS number: **64500**

The IP addresses of your iBGP neighbors you can find out by doing a *show ip route ospf* or *show ip route isis* (depending on the IGP in use) and look for /32 prefixes for IPv4. For IPv6 do a *show ipv6 route ospf6* or *show ipv6 route isis* and look for /128 prefixes.

4.3 To configure iBGP in config mode you need:

router bgp <asnumber> to start BGP with <*asnumber*> as your AS number

neighbor <name> peer-group to start configuring a peergroup named <*name*>

neighbor <name> remote-as <asnumber> to set a common remote AS number for all peer group members. If configuring iBGP <*asnumber*> is **your** AS number.

neighbor <name> update-source <interface> make BGP use the IP address of <*interface*> as source IP when setting up connections. Use your loopback interface name (*loopback0*) here.

neighbor <name> next-hop-self makes iBGP to set this routers IP (of *loopback0*) for next-hop address when distributing prefixes received via eBGP

neighbor <name> next-hop-self all makes iBGP to set this routers IP (of *loopback0*) for next-hop address when distributing prefixes received via eBGP *or iBGP*

neighbor <name> send-community both to forward BGP communities. We have not covered them yet, but you will need this later on.

neighbor <ip address> peer-group <name> sets up a BGP session to *<ip address>* as a member of peer-group <*name*>. That means all configs from <*name*> are copied.

address-family ipv6 switches into IPv6 context. You need to activate explicitly distributing IPv6 prefixes in this context.

address-family ipv4 switches into IPv4 context. You do not need to activate the distribution of IPv4 prefixes via IPv4 - this is activated by default. But you need to *de*activate the distribution of IPv6 addresses.

neighbor <ipv6 address> activate in address-family context activates distribution of IPv6 prefixes in this context.

neighbor <ipv4 address> activate in address-family context activates distribution of IPv4 prefixes. Not needed in IPv4 context.

no neighbor <ipv6 address> activate in address-family context prevents distribution of IPv6 prefixes.

no neighbor <ipv4 address> activate in address-family context prevents distribution of IPv4 prefixes.

It is recommended to distribute prefixes only in their own context. So you usually configure:

```
...
address-family ipv4 unicast
  neighbor <ipv4 address> activate
  ... (for all IPv4 neighbors)
  no neighbor <ipv6 address> activate
  ... (for all IPv6 neighbors)
exit-address-family
address-family ipv6 unicast
  no neighbor <ipv4 address> activate
  ... (for all IPv4 neighbors)
  neighbor <ipv6 address> activate
  ... (for all IPv6 neighbors)
exit-address-family
```

In FRRouting you can configure this also for a whole peer group and you can change the default so IPv4 is *not* activated automatically. See solution below.

4.4 Commands to check if iBGP is running:

show ip bgp summary shows you all configured neighbor and if sessions to them are up or not.

show ip bgp neighbors show you detailed information about all BGP neighbors (long!)

show ip bgp neigbors <ip address> same as above, but only for one neigbor. Preferred.

5 Solution

5.1 Solution for FRRouting

5.1.1 IPv4

```
router bgp 64500
  no bgp default ipv4-unicast
  neighbor internal-v4 peer-group
  neighbor internal-v4 remote-as 64500
  neighbor internal-v4 update-source loopback0
  address-family ipv4 unicast
    neighbor internal-v4 activate
  exit-address-family
  neighbor 172.16.1.YY peer-group internal-v4
  ...
```

Where *YY* are the IPv4 addresses of all other routers in the network (except your own).

5.1.2 IPv6

```
router bgp 64500
  no bgp default ipv4-unicast
  neighbor internal-v6 peer-group
  neighbor internal-v6 remote-as 64500
  neighbor internal-v6 update-source loopback0
  address-family ipv6 unicast
    neighbor internal-v6 activate
  exit-address-family
  neighbor 2001:db8:500::1:Y peer-group internal-v6
  ...
```

Where *Y* are the IPv6 addresses of all other routers in the network (except your own).

5.2 Solution for Juniper

```
routing-options {
  autonomous-system 64500;
}
protocols {
  bgp {
    local-as 64500;
    peer-as 64500;

    group internal-v4 {
      type internal;
      local-address 172.16.1.X;
      peer-as 64500;
      family inet {
        unicast;
      }
      neighbor 172.16.1.YY;
      ...
    }
    group internal-v6 {
      type internal;
      local-address 2001:db8:500::1:X;
      peer-as 64500;
      family inet6 {
        unicast;
      }
      neighbor 2001:db8:500::1:Y;
      ...
    }
  }
}
```

5.3 Solution for Mikrotik

```
/routing bgp instance
set default as=64500 router-id=172.16.1.XX
/routing bgp peer
add name=Ry-v4 remote-address=172.16.1.yy remote-as=64500\
  update-source=loopback0 address-families=ip
...
add name=Ry-v6 remote-address=2001:db8:500::1:Y remote-as=64500\
  update-source=loopback0 address-families=ipv6
...
```