Experiment 1e - Distribute prefixes

DE-CIX Academy

Version 1.2

1 Introduction

BGP is all about prefixes. To get prefixes *into* BGP there are two methods, in this experiment we will use the simplest method adding prefixes using a *network* statement.

2 Network Setup

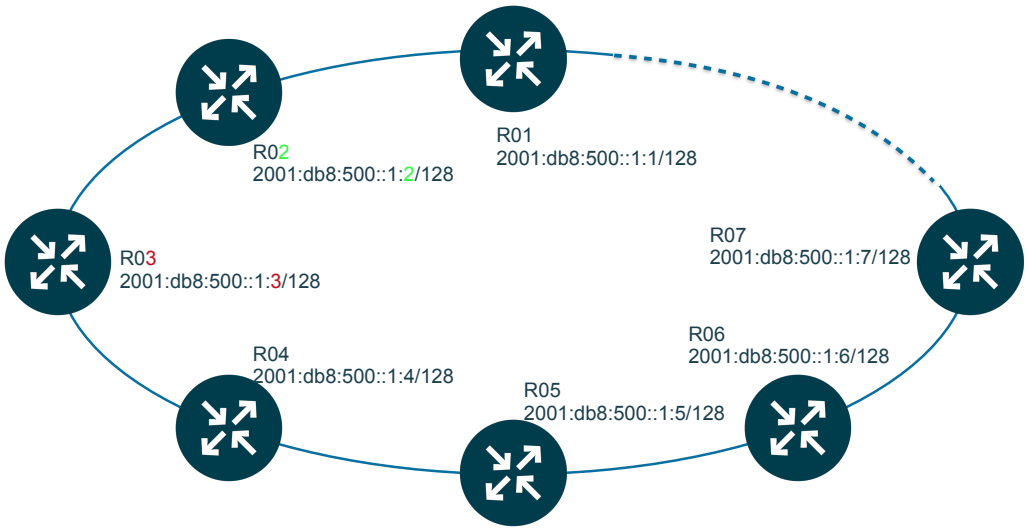


Figure 1: Network Setup

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).

iBGP has already been set up between the loopback addresses of the routers.

3 Distribute prefixes

3.1 Tasks:

- Put in static routes for one IPv4 and one IPv6 prefix
- Distribute these prefixes in BGP

3.2 To distribute prefixes you need to:

- have the prefix to be distributed in your routing table
- tell BGP to distribute that prefix

Information you need:

- A couple of prefixes to distribute
- Have a look at the last part of your loopback address. If your loopback address is for example 172.16.1.**X**/32, then you can use
  - 10.**X**.0.0/16
  - 2001:db8:**X**::/48

  to redistribute

3.3 To configure routes and redistribution in config mode:

ipv6 route <prefix> <destination>

ip route <prefix> <destination> This creates static routes in your routing table for prefix <*prefix*>. <*destination*> is where packets to this prefix are sent, if you want just a static route so you can put the prefix into BGP you can use either the *Null0* or your *loopback0* interface as destination.

address-family ipv4 unicast or

address-family ipv6 unicast in BGP context selects an adress family for the following *network* config commands

network <prefix> in BGP address-family context tells your BGP router to put <*prefix*> into BGP **if there is a corresponding route to it in your routing table**

network <prefix> route-map <name> same as above, but uses route-map <*name*> to modify attributes before putting it into BGP. We will cover this later.

3.4 Commands to check if your prefix is announced:

show ipv6 bgp

show ip bgp lists all prefixes in the BGP table

show ipv6 bgp <*prefix*>

show ip bgp <*prefix*> gives you detailed information about <*prefix*> including to which neigbors it is announced.

4 Distribution of prefixes learned via eBGP

We did not configure eBGP yet, but it is important to learn how prefixes received via eBGP are distributed in iBGP. By default, all prefixes leared via eBPG are distributed *unchanged*, that means all attributes and also the next-hop IP address stay as they are received.

Especially for the next-hop address this can be a problem, if you do not distributed all interface addresses via your IGP. Thankfully there is a solution for that. You can configure BGP to replace the next-hop address with the IP address of the distributing iBGP router.

4.1 To configure next-hop address behaviour:

This is configured in *address-family* context.

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*

5 Solution

5.1 IPv4

```
ip route 10.something.somethingelse.0/24 Null0
router bgp 64500
  network 10.something.somethingelse.0/24
```

Choose any valid prefix for this. Ask your neighbors if the see the prefix in the BGP and routing tables.

5.2 IPv6

```
ipv6 route 2001:db8:something::/48 Null0
router bgp 64500
  address-family IPv6 unicast
    network 2001:db8:something::/48
```

Here the *network* statement has to be within the *address-family* context. The only reason you can define an IPv4 statement without the context is backward compatibility.

6 Slides

Distribute an IPv4 prefix

- any prefix you want to distribute must be in the routing table first
- so you do need a static route entry for it

```
ip route 10.11.xx.0/24 null0

router bgp 64500
 address-family ipv4 unicast
  network 10.11.xx.0/24
```

- For *xx* use your router id (01, 02, ...)
- Instead of 10.11.xx.0/24 you can also announce any other prefix
- A *route target* of “null0” means “we have no target for this prefix but want it in our routing table anyway”

Show commands:

- show ip route static
- show bgp ipv4
- show bgp ipv4 10.11.xx.0/24
- show ip route bgp

Distribute an IPv6 prefix

- You also need a static route for your IPv6 prefix
- Be aware of the address-family statement

```
ipv6 route 2001:db8:xx::/48 Null0

router bgp 64500
 address-family ipv6 unicast
  network 2001:db8:xx::/48
```

- For *xx* use your id (01, 02, ...)
- You can also use any other IPv6 prefix for this
- A *route target* of “null0” means “we have no target for this prefix but want it in our routing table anyway”

Show commands:

- show ipv6 route static
- show bgp ipv6
- show bgp ipv6 2001:db8:xx::/48
- show ipv6 route bgp

Change the next-hop address

- Check your BGP table - show bgp ipv4
- The prefix *r01* distributes is *inaccessible* by the other routers
- because the next hop it has is not distributed
- *r01* needs to set itself as the next-hop for iBGP

```
router bgp 64500
 address-family ipv4 unicast
  neighbor internal-v4 next-hop-self
 exit-address-family
 address-family ipv6 unicast
  neighbor internal-v6 next-hop-self
 exit-address-family
```

Show commands:

- show bgp ipv4
- show bgp ipv6