Routing Hints, Top Tips, and FAQs

Here are some of the Routing hints and documentation that seems to be woefully lacking or just misleading on the greater wider Internet.

Propagating default using IS-IS rather than IBGP

It is often desirable on multihoming end-sites to propagate the default route learned from external peers using IS-IS rather than using IBGP itself.

The typical set up would be an organisation having two border routers (BR1 and BR2) and two core routers (CR1 and CR2). And these would be linked together in the typical cross-fashion, so:

- BR1 connects to CR1 and CR2
- BR2 connects to CR2 and CR1
- CR1 connects to CR2

There isn't any need for BR1 and BR2 to interconnect, usually.

How do we inject the default route into IS-IS if it is heard from an EBGP neighbour? The following configuration example uses Cisco IOS syntax but can be easily adapted for your own favourite vendor (unless your favourite vendor has an easy mechanism to do this).

Set EBGP distance

Before we start, we need to make sure that the default distance for EBGP learned routes is set to something lower than those learned by IS-IS, otherwise the following will not work. Yes, Cisco's defaults are for EBGP routes to be distance 20, IBGP learned routes to be distance 200, and locally originated routes to be also distance 200. But most operators change the EBGP distance to be the same as for IBGP distance - and the following trick will not work if the EBGP distance is greater than that of IS-IS.

Default Originate in IS-IS

The way to make this work is to use the default-information originate command in IS-IS. However, without any options applied to it, this command will unconditionally originate the default route. Even if the router has no path to the default. Which can be somewhat problematic if the default route learned from the EBGP speaker disappears.

To remedy this we need to use the route-map option for the default-information originate command. The route-map we are going to add needs to look for the default route appearing in the router's global RIB and, if it exists, inject it into the IS-IS RIB so that it is then propagated to IS-IS speaking routers in the network. But left like this, the default route seen in the RIB could be created by the IS-IS process on a neighbouring internal router. So we need to add another condition to make sure that the default we are seeing in the RIB is learned from the EBGP neighbour. Our resulting

configuration is below:

```
ip access-list standard BGP-NH
  remark External BGP speaker
  permit 10.10.10.1
!
ip prefix-list DEFAULT permit 0.0.0.0/0
!
route-map DEFAULT-ORIG permit 5
  match ip address prefix-list DEFAULT
  match ip next-hop BGP-NH
!
```

We can now apply this route-map to the Border Router IS-IS process:

```
router isis ISP
...configuration...
default-information originate route-map DEFAULT-ORIG
!
```

If we would like to prefer one border router over the other, then we would set a metric in the route-map. For example, if we want to prefer BR1 over BR2 as the exit point, then BR1 route-map would look like this:

```
route-map DEFAULT-ORIG permit 5
match ip address prefix-list DEFAULT
match ip next-hop BGP-NH
!
```

and BR2's route-map would look like this:

```
route-map DEFAULT-ORIG permit 5
match ip address prefix-list DEFAULT
match ip next-hop BGP-NH
set metric 10
!
```

The other routers in the network will see the default being propagated by IS-IS from the BR1 router. If BR1 loses its externally learned default, the other routers will then see the default being propagated by IS-IS from the BR2 router, with a metric increased by 10.

Removing default from IBGP

Now that we are propagating the default from both routers using IS-IS, we no longer need to carry the default route in our IBGP. So we augment the existing inbound policy applied to the EBGP peers with a statement which looks for the default route and tags it with the no-advertise community.

The following configuration example shows how to do this:

```
route-map EBGP-in permit 5
description Do not propagate default route
match ip address prefix-list default
set community no-advertise
!
route-map EBGP-in permit 10
description Other policy
...other policy configuration...
!
router bgp 64512
address-family ipv4
neighbor 10.10.10.1 route-map EBGP-in in
distance bgp 20 200 200
!
```

The no-advertise community will make sure that the matched route will not be announced to any other BGP speaker.

Conclusion

With the configuration applied to both Border Routers now, the rest of the routers will see a default route via IS-IS only. There will be no default route in BGP apart from at the Border Routers. Each Border Router will see the default route being learned from the attached external neighbour. And the best path to the default route will be determined by the IS-IS metric attached to the default being propagated by IS-IS.

Complete Configuration

For BR1 (making BR1 the main default gateway):

```
router isis ISP
...

default-information originate route-map DEFAULT-ORIG
!
router bgp 64512
...

address-family ipv4
...

neighbor 10.10.10.1 route-map EBGP-in in
distance bgp 20 200 200
!
ip access-list standard BGP-NH
remark External BGP speaker
permit 10.10.10.1
!
ip prefix-list DEFAULT permit 0.0.0.0/0
!
route-map DEFAULT-ORIG permit 5
```

```
match ip address prefix-list DEFAULT
match ip next-hop BGP-NH
!
route-map EBGP-in permit 5
description Do not propagate default route
match ip address prefix-list default
set community no-advertise
!
route-map EBGP-in permit 10
description Other policy
...
!
```

Back to Home page

From:

https://mail.bgp4all.com/pfs/ - Philip Smith's Internet Development Site

Permanent link:

https://mail.bgp4all.com/pfs/hints/routing?rev=1600773062

Last update: 2020/09/22 11:11

