

The BGP synchronization rule states that if an AS provides transit service to another AS, BGP should not advertise a route until all of the routers within the AS have learned about the route via an IGP.

Consider the following scenario:

AS 5 (on the right hand side) is advertising 150.5.0.0/16.

Problem: AS100 (left) and AS1 (top left) are not seeing any path to 150.5.0.0/16

Analysis: The core of the problem is twofold:

- Routers in AS24 (bottom) and AS3 (top right) have synchronization enabled.

- Router SP4-1 is not redistributing 150.5.0.0/16 into its IGP

Lets start by looking at the synchronization scene.

Router SP4-1 learns about 150.5.0.0/16, as does SP3-2 via EBGP and both install the route into their routing table.

```
SP4-1#show ip route bgp
B 150.5.0.0/16 [20/1] via 4.0.45.2, 00:30:54
```

```
SP3-2#show ip route bgp
B 150.5.0.0/16 [20/0] via 3.3.35.2, 00:31:26
```

SP4-1 and SP3-2 distributes this route to other iBGP routers in their respective ASs, but these routers will NOT install the route into their routing table because synchronization is enabled. (synchronization was enabled by default on Cisco IOS before 12.8T).

We can see this clearly on router SP3-1, where the show ip bgp command shows the 150.5.0.0 network in the bgp table, but not the routing table.

```
SP3-1#show ip bgp
BGP table version is 1, local router ID is 3.3.10.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
Network Next Hop Metric LocPrf Weight Path
*
i150.5.0.0 3.3.35.2 0 100 0 5 i
```

Key concept: Since routers SP3-1 and SP2-1 don't install the route in their routing table, neither will advertise the route to the neighbouring ASs (AS1 and AS100).

The situation on routers SP2-1 and SP2-2 though is slightly different to SP3-1 because SP2-2 will not only be receiving iBGP updates from SP4-1 advertising 150.5.0.0, but also eBGP updates from SP3-2 - remember SP3-2 has 150.5.0.0/16 in its routing table too.

```
SP2-2#show ip bgp
BGP table version is 2, local router ID is 222.2.10.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete
```

```
Network Next Hop Metric
LocPrf Weight Path
* > 150.5.0.0 222.2.23.6 0 3 5 i
```

```
4.0.45.2 1 100 0 5
```

```
SP2-2#show ip route bgp
B 150.5.0.0/16 [20/0] via 222.2.23.6, 01:27:21
```

The output of the show ip bgp command on SP2-2 shows how it knows about two paths to 150.5.0.0/16 - the iBGP path via 4.0.52.2 learned from SP4-1, and the eBGP path via 222.2.23.6 learned from SP3-2. Note however that although both are marked as valid (indicated by *), the eBGP path is installed as the best path (indicated by >).

The result is virtually identical for SP2-1. This might seem contrary to what you would expect - after all, the path via SP3-2 will traverse two ASs, while the path via SP4-1 will just pop across to the AS next door.

Synchronization

Explained

However, if we look closely, we can see the whole reason for why we would want to use synchronization in AS24.

If SP2-2 needed to reach 150.5.0.0, its internal path would be via SP4-2. However, SP4-2 does not run bgp, so has no knowledge of the path to 150.5.0.0.

```
SP4-2#show ip route 150.5.0.0  
% Network not in table
```

So in this scenario, it is far better to leave synchronisation enabled, and have the path to 150.5.0.0 via AS3 than via SP4-2 which is a black hole as far as a route to 150.5.0.0 is concerned.

What happens if we disable synchronization?

If synchronization is disabled on AS24 SP2-1 and SP2-2 will allow the iBGP path to be installed into the routing table. Here's the same output from SP2-2 after synchronization has been disabled on AS24.

```
SP2-2#show ip bgp  
BGP table version is 3,  
local router ID is 222.2.10.2  
Status codes: s suppressed, d damped, h history, * valid, >  
best, i - internal,  
r RIB-failure, S Stale  
Origin codes: i - IGP, e - EGP, ? -  
incomplete  
Network          Next Hop          Metric LocPrf Weight Path  
150.5.0.0        222.2.23.6        0 3 5 i          4.0.45.2         1 100  
0 5
```

```
SP2-2#show ip route bgp  
B 150.5.0.0/16 [200/1] via  
4.0.45.2, 00:15:27
```

First, the good news

Now that routers SP2-1 and SP3-1 have the route to 150.5.0.0/16 in their routing tables, they will now advertise the route to their neighbouring ASs - AS1 and AS100.

```
SP1-2#show ip bgp  
BGP  
table version is 3, local router ID is 1.0.10.2  
Status codes: s suppressed, d damped, h history, *  
valid, > best, i - internal,  
r RIB-failure, S Stale  
Origin codes: i - IGP, e - EGP, ? -  
incomplete  
Network          Next Hop          Metric LocPrf Weight Path  
150.5.0.0  
1.0.13.2          0 100 0 3 5 i *      1.2.2.2          0 100 24 5 i  
SP1-2#show ip route bgp  
B 150.5.0.0/16 [200/0] via 1.0.13.2, 00:07:43
```

Problem solved. Or is it?

Now the bad news

Note that back on router SP2-1, show ip bgp shows the best path as being the iBGP route - via 4.0.45.2. The problem is, for SP3-2 to get to 4.0.52.2, it will want to go via router SP4-1. And as I explained earlier, SP4-1 doesn't know how to forward any packets that are headed to 150.5.0.0. A traceroute from SP2-1 shows how the trace only gets as far as SP4-2.

```
SP2-1#traceroute 150.5.0.1  
Type escape sequence to abort.  
Tracing the route to 150.5.0.1  
 1 222.2.0.2 44 msec 12 msec 100 msec  
 2 222.2.24.2 80 msec 188 msec 20 msec  
 3 222.2.24.2 !H * !H
```

Removing the Black Hole

Essentially the most logical way to remove the black hole is for router SP4-2 to learn the route to 150.5.0.0. You could add it as a static route, or you could redistribute the bgp update from SP5-2 into the IGP. When this is done, the network converges.

```
SP2-1#traceroute 150.5.0.1  
Type escape sequence to abort. Tracing the  
route to 150.5.0.1  
 1 222.2.0.2 88 msec 60 msec 24 msec  
 2 222.2.24.2 112 msec 88 msec  
 68 msec  
 3 4.0.0.1 76 msec 168 msec 184 msec  
 4 4.0.45.2 44 msec 196 msec 24 msec  
 5 150.5.0.1 336 msec * 64 msec [Note: this assumes 150.5.0.1 has a route back to SP2-1]
```

When to enable or disable synchronization

From the example above, it seems that you would never want to enable synchronization, and indeed this is generally the case. However, for AS24, or you could create black holes if your BGP routes are not being redistributed into your IGP.

<http://www.cisco.com/en/US/docs/internetnetworking/case/studies/icsbgp4.html#wp19382> Cisco gives the following advice on when it is permissible to disable synchronization:

You can disable synchronization if one of the following conditions is

true:

- Your AS does not pass traffic from one AS to another AS.
- All the transit routers in your AS run BGP.

Following this advice, we can see that we can avoid the black hole in AS24 when the bgp routes are NOT being redistributed.

So lets finally examine what happens when:

- bgp is NOT being redistributed, and
- synchronization is enabled for AS24

Firstly, lets recall that above when we had 1) bgp NOT being distributed AND 2) synchronization disabled, we had a bit-pit at router SP4-2 for traffic leaving SP2-1.

```
SP2-1#traceroute 150.5.0.1
 1 150.5.0.1 0 msec 0 msec 0 msec
 2 222.2.24.2 80 msec 188 msec 20 msec
 3 222.2.24.2 !H * !H
```

Now lets turn off the redistribution AND enable synchronization:

```
SP2-1#traceroute 150.5.0.1
 1 150.5.0.1 0 msec 0 msec 0 msec
 2 222.2.23.2 16 msec 32 msec 64 msec
 3 3.3.35.2 152 msec * 120 msec
```

[Note: 3.3.35.2 is the same router as 150.5.0.1, so the traceroute succeeded]

Here the black hole is avoided, and the traffic sent to 150.0.5.1 via AS3

So the best practice for this example is to keep synchronization enabled for AS24, but to captialise on the best paths, we redistribute bgp routes into the IGP as well.