

## Module 13 – Multihoming to Different ISPs

**Objective:** To investigate various methods for multihoming onto two different upstream ISPs.

**Prerequisites:** Module 12 and Multihoming Presentation

The following will be the common topology used.




Figure 1 – ISP Lab Multihoming Configuration

## Lab Notes

The purpose of this module is to demonstrate multihoming in the situation where the customer AS has one connection to more than one upstream service provider. There are at least two situations where this is applicable:

- Enterprise or Service Provider customer requires more than one connection to the Internet to provide resiliency, and/or loadsharing.
- Enterprise or Service Provider customer requires more than one Internet connection to give their network service provider redundancy.

It is important that you review the multihoming presentation before you start with this module. Only configuration examples will be given – it will be left to the workshop participant to use the presentation notes to help them configure their routers correctly.

To ensure an understandable and easy to follow configuration, as well as good practice, a few assumptions about configuring BGP will be made. These are:

- **Use prefix-lists to filter prefixes**
- **Use as-path access-lists to filter ASes**
- **Use route-maps to implement policy**

There are rarely any exceptions to this. Prefix lists are very efficient access-lists and make implementation of prefix filtering on AS borders (and elsewhere) very easy. Please review the BGP presentation materials if there is any uncertainty as to how prefix lists work.

## Lab Exercise

1. **Basic Configuration.** Each router team should configure their router to fit into the network layout depicted in Figure 1. Check all connections. Note that most links are using serial cables. Remember what was covered in Module 11!
2. **Addressing Plan.** These address ranges should be used throughout this module. You are welcome to use your own range within an AS if you desire, just so long as you consult with the teams in other ASes to ensure there is no overlap. In the every day Internet, such address assignment is carried out by the Regional Internet Registry. AS16384 is the transit provider used in this module, and as such represents the Internet at large. A /16 network block has been assigned to that provider.

|     |               |
|-----|---------------|
| AS1 | 120.10.0.0/19 |
| AS2 | 120.19.0.0/19 |
| AS3 | 120.73.0.0/19 |
| AS4 | 121.19.0.0/19 |

|         |               |
|---------|---------------|
| AS5     | 121.35.0.0/19 |
| AS6     | 121.99.0.0/19 |
| AS16384 | 122.11.0.0/16 |

3. **Routing Protocols.** OSPF (area 0 only) and iBGP should now be configured between the routers for each of AS 1, AS6 and AS16384. The other 4 ASes do not require OSPF or iBGP as they only have a single router in them. Any interfaces which should not be running OSPF *MUST* be marked as passive in the configuration. And don't forget to use BGP peer groups for iBGP peers.




Figure 2 – Multihoming Lab Physical Layout

**Checkpoint #1:** When you have properly configured your router, and the other routers in the AS are reachable (i.e. you can ping the other routers, and see BGP and OSPF prefixes in the routing table), please let the instructor know.

## **Scenario One – Primary link and backup link**

This first scenario is more commonly employed where the customer has a large circuit to their upstream and an inexpensive circuit they use almost exclusively for backup purposes to another ISP. (The case maybe that the primary ISP has good connection to the Internet, and the backup ISP is used only as a last resort – for technical, commercial, or political reasons.)

In this case the whole address block is announced out of both links. However, the announcement going out the backup link has its AS path length increased so that it is at a lower priority. Likewise, the incoming default route announcement from the ISP is “weighted” using local-preference. (**Hint:** remember the purpose of changing the AS Path length? If in doubt, review the BGP presentation material.)

4. **Enable eBGP between the transit ASes.** AS2 and AS3 should now enable their eBGP link to AS16384. AS4 and AS5 should do likewise. All router teams in these ASes must ensure that they can see all the prefixes of AS2, AS3, AS4, AS5 and AS16384. If they are not there, work with your team members to ensure they appear. Don’t forget the static pull-up route when injecting prefixes into BGP! Also, at this stage there is no need to install prefix filters between these ASes – if you would like to, don’t forget that you need to allow through the network blocks of the ASes you are providing transit to.
5. **Prepare to enable eBGP between AS1 and its two upstreams.** AS1 should currently be running iBGP within its own network. To announce AS1’s prefix to AS2 and AS3 we will take the /19 address block and announce it on both peerings between the ASNs. AS2 and AS3 will not announce any prefixes to AS1 – they will simply announce a default route. There is no need for any more specific routing information to be injected into the customer site.
6. **Prepare to enable eBGP between AS6 and its two upstreams.** The configuration steps to enable BGP between AS6 and its two upstreams are the same as those being described for AS1.
7. **Create AS1 prefix lists.** First, create the prefix lists on the routers in AS1. Both Router1 and Router3 will announce the aggregate. Both will accept the default route. Example for Router1:

```
ip prefix-list myblock permit 120.10.0.0/19
ip prefix-list default permit 0.0.0.0/0
```

8. **Create AS2 and AS3 prefix-lists.** The routers in AS2 and AS3 should only accept those prefixes which the customer is entitled to announce. So a prefix list needs to be installed on both Router2 and Router4 to do this. For example:

```
ip prefix-list Customer permit 120.10.0.0/19
ip prefix-list default permit 0.0.0.0/0
```

9. **Configure the main link.** Configure the main link between the customer AS and the ISP. For AS1, the link between Router1 and Router2 in AS2 is the main link – the link between Router3 and Router4 in AS3 is the backup. For AS6, the main link is between Router 12 and Router 11 in AS4. Example configuration for Router1:

```
router bgp 1
 network 120.10.0.0 mask 255.255.224.0
 neighbor <router2> remote-as 2
```

```

neighbor <router2> description Link to Router2 in AS2
neighbor <router2> prefix-list myblock out
neighbor <router2> prefix-list default in
!
ip route 120.10.0.0 255.255.224.0 null0

```

- 10. Configure the backup link.** Configure the backup link between the private AS and the ISP. Increase the AS Path Length on outbound announcements to 3, and set local preference on inbound announcements to 80. Remember that the shortest AS Path Length and highest local-preference win during the BGP path selection process. To do this, use a route-map on the peering – you will require an inbound and outbound route-map. Example configuration for Router14:

```

ip prefix-list myblock permit 121.99.0.0/19
ip prefix-list default permit 0.0.0.0/0
!
route-map outfilter permit 10
match ip address prefix-list myblock
set as-path prepend 6 6 6
route-map outfilter permit 20
!
route-map infilter permit 10
match ip address prefix-list default
set local-preference 80
route-map infilter permit 20
!
router bgp 6
network 121.99.0.0 mask 255.255.224.0
neighbor <router13> remote-as 5
neighbor <router13> description Link to Router13 in AS5
neighbor <router13> prefix-list myblock out
neighbor <router13> prefix-list default in
neighbor <router13> route-map outfilter out
neighbor <router13> route-map infilter in
!
ip route 121.99.0.0 255.255.224.0 null0

```

- 11. Configure eBGP in AS2 and AS3 with AS1.** AS2 and AS3 are going to originate the default route in the peering with AS1. The BGP command “default-originate” is used to do this. Example configuration for Router 4:

```

router bgp 3
neighbor <router3> remote-as 1
neighbor <router3> description Multihomed Customer
neighbor <router3> default-originate
neighbor <router3> prefix-list Customer in
neighbor <router3> prefix-list default out
!

```

- 12. AS2 and AS3 eBGP configuration with AS16384.** Without further configuration changes in AS3, AS1 will be announced by AS2 and AS3 routers to other ASes. It is always good practice to only announce the prefixes you are entitled to announce to the Internet – so now is the time to add prefix filters to the peering with AS16384. The example below is for AS3:

```

ip prefix-list mynets permit 120.10.0.0/19
ip prefix-list mynets permit 120.73.0.0/19
!

```

```
router bgp 3
 neighbor <router8> remote-as 16384
 neighbor <router8> description Peering with AS16384 - The Internet
 neighbor <router8> prefix-list mynets out
!
```

- 13. AS16384 eBGP configuration.** The configuration for the routers in AS16384 is left as an exercise. You can choose to filter your customer ASes if you wish – the prefix filters will be based on those configured by your “customers” in the previous step. You will announce all prefixes you have to your customers, so in this case there is no need for an outbound prefix filter (although you can construct one if you really wish to do so).
- 14. AS6 and its upstreams of AS4 and AS5.** The same types of configuration concepts are also required on AS4, AS5 and AS6. AS6 is a multihomed customer of AS4 and AS5. The teams looking after the routers in these two ASes should use the above configuration examples as hints to set up their own peering sessions.
- 15. Connectivity Test.** Check connectivity throughout the lab network. Each router team should be able to see all other routers in the room. When you are satisfied that BGP is working correctly, try running traceroutes to ensure that the primary paths are being followed. When you are satisfied this is the case, check that the backup functions (do this by disconnecting the cable between the two routers on the primary path) – you will see that the backup path is now used.

***Checkpoint #2:*** *Once the BGP configuration has been completed, check the routing table and ensure that you have complete reachability over the entire network. If there are any problems, work with the other router teams to resolve those.*

## STOP AND WAIT HERE

### Scenario Two – Loadsharing

Most multihomed sites want to implement some kind of loadsharing on the circuits they have to their upstream provider. The example here discusses only two circuits, but the techniques work equally well for a greater number.

To do this, the whole address block is announced out of both links. Also, the address block is split into two pieces, with one subprefix being announced out of one link, and the other being announced out of the other link. The result of this is that traffic for the first /20 comes in one path, and traffic for the second /20 comes in the other path. If either path fails, the advertisement of the /19 address block (aggregate) ensures continued connectivity.

- 16. Clean up the configuration of AS1 and AS6.** Remove the configuration which set the weighting for the previous example – specifically the route-maps. They must be removed from the BGP configuration, and from the main configuration.
- 17. Configure the address block and subprefixes in AS1 and AS6.** Modify the router configuration so that the /19 address block and the two /20 subprefix is present in the BGP table. Also set up prefix lists to cater for these blocks. For example:

```

ip prefix-list aggregate permit 120.10.0.0/19
!
ip prefix-list subblock1 permit 120.10.0.0/19
ip prefix-list subblock1 permit 120.10.0.0/20
!
ip prefix-list subblock2 permit 120.10.0.0/19
ip prefix-list subblock2 permit 120.10.16.0/20
!
ip prefix-list default permit 0.0.0.0/0
!
router bgp 1
 network 120.10.0.0 mask 255.255.224.0
 network 120.10.0.0 mask 255.255.240.0
!
ip route 120.10.0.0 255.255.224.0 null0
ip route 120.10.0.0 255.255.240.0 null0

```

- 18. Configure BGP in AS1 and AS6.** For AS1, the link between Router1 and Router2 in AS2 is the first link – the link between Router3 and Router4 in AS3 is the second (and is the one which should also announce the subprefix). For AS6, the first link is between Router 12 and Router 11 in AS4. Configure BGP on the border routers in the customer ASes so that the prefix and one subprefix is announced to the direct peer as described earlier. For example, Router1 could announce *aggregate* and *subblock1* as above, whereas Router3 could announce *aggregate* and *subblock2*. For example on Router3:

```

!
router bgp 1
 network 120.19.0.0 mask 255.255.224.0
 network 120.19.16.0 mask 255.255.240.0
 neighbor <router4> remote-as 3
 neighbor <router4> description Link to Router4 in AS3
 neighbor <router4> prefix-list subblock2 out
 neighbor <router4> prefix-list default in
!
ip route 120.19.16.0 255.255.240.0 null0

```

- 19. Connectivity test.** Check connectivity throughout the lab network. Each router team should be able to see all other routers in the room. When you are satisfied that BGP is working correctly, try running traceroutes to check the path being followed. Also check that backup via the alternative path still functions (do this by disconnecting the cable between the two routers on the primary path) – you will see that the backup path is now used.

**Checkpoint #3:** *Once the BGP configuration has been completed, check the routing table and ensure that you have complete reachability over the entire network. If there are any problems, work with the other router teams to resolve those.*

**STOP AND WAIT HERE**

## Scenario Three – More Controlled Loadsharing

The third scenario is a variation on the second scenario and provides another example.

As before the whole address block is announced out of both links. In fact, one of the key features of multihoming, and providing redundancy, is that the ISP's address blocks are always announced out of each external link. The key to loadbalancing is how those external announcements are made. In this example, one /20 is taken out of the /19 address block and announced on one link between the customer and the ISP – and the /19 is announced with a longer AS-PATH on this link also. The other link sees just the standard announcement of the /19.

**20. Clean up the configuration of AS1 and AS6.** Remove the configuration which subdivided the address space for the previous example. Remember it is always very important to remove any configuration which isn't being used from the router.

**21. Configure the address block and subprefixes in AS1 and AS6.** Modify the router configuration so that the /19 address block and one /20 subprefix is present in the BGP table. Also set up prefix lists to cater for these blocks. For example:

```
ip prefix-list aggregate permit 120.10.0.0/19
!
ip prefix-list subblocks permit 120.10.0.0/19 le 20
!
ip prefix-list default permit 0.0.0.0/0
!
router bgp 1
 network 120.10.0.0 mask 255.255.224.0
 network 120.10.0.0 mask 255.255.240.0
!
ip route 120.10.0.0 255.255.224.0 null0
ip route 120.10.0.0 255.255.240.0 null0
```

**22. Configure BGP in AS1 and AS6.** For AS1, the link between Router1 and Router2 in AS2 is the first link – the link between Router3 and Router4 in AS3 is the second (and is the one which should also announce the subprefix). For AS6, the first link is between Router 12 and Router 11 in AS4. Configure BGP on the border routers in the customer ASes so that the prefix and one sub prefix is announced to the direct peer as described earlier. For example, Router1 could announce *aggregate* as above, whereas Router3 could announce *aggregate* with a lengthened AS Path, and announce *subblock1* as is. For example on Router3:

```
route-map outfilter permit 10
 match ip address prefix-list aggregate
 set as-path prepend 1 1 1
route-map outfilter permit 20
!
route-map infilter permit 10
 match ip address prefix-list default
 set local-preference 80
route-map infilter permit 20
!
router bgp 1
 network 120.19.0.0 mask 255.255.224.0
 neighbor <router4> remote-as 3
 neighbor <router4> description Link to Router4 in AS3
```

```
neighbor <router4> prefix-list subblocks out
neighbor <router4> prefix-list default in
neighbor <router4> route-map outfilter out
neighbor <router4> route-map infilter in
!
```

**23. Connectivity test.** Check connectivity throughout the lab network. Each router team should be able to see all other routers in the room. When you are satisfied that BGP is working correctly, try running traceroutes to check the path being followed. Also check that backup via the alternative path still functions (do this by disconnecting the cable between the two routers on the primary path) – you will see that the backup path is now used.

**Checkpoint #4:** *Once the BGP configuration has been completed, check the routing table and ensure that you have complete reachability over the entire network. If there are any problems, work with the other router teams to resolve those.*

**24. Check the network paths.** Run traceroutes between your router and other routers in the classroom. Ensure that all routers are reachable. If any are not, work with the other router teams to establish what might be wrong.

**25. Summary.** This module has covered the major situations where a customer requires to multihomed onto more than one service provider backbone. It has demonstrated how to implement this multihoming using prefix-lists, AS Path Length modifications and local-preference where appropriate.

## ***CONFIGURATION NOTES***

Documentation is critical! You should record the configuration at each ***Checkpoint***, as well as the configuration at the end of the module.