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.

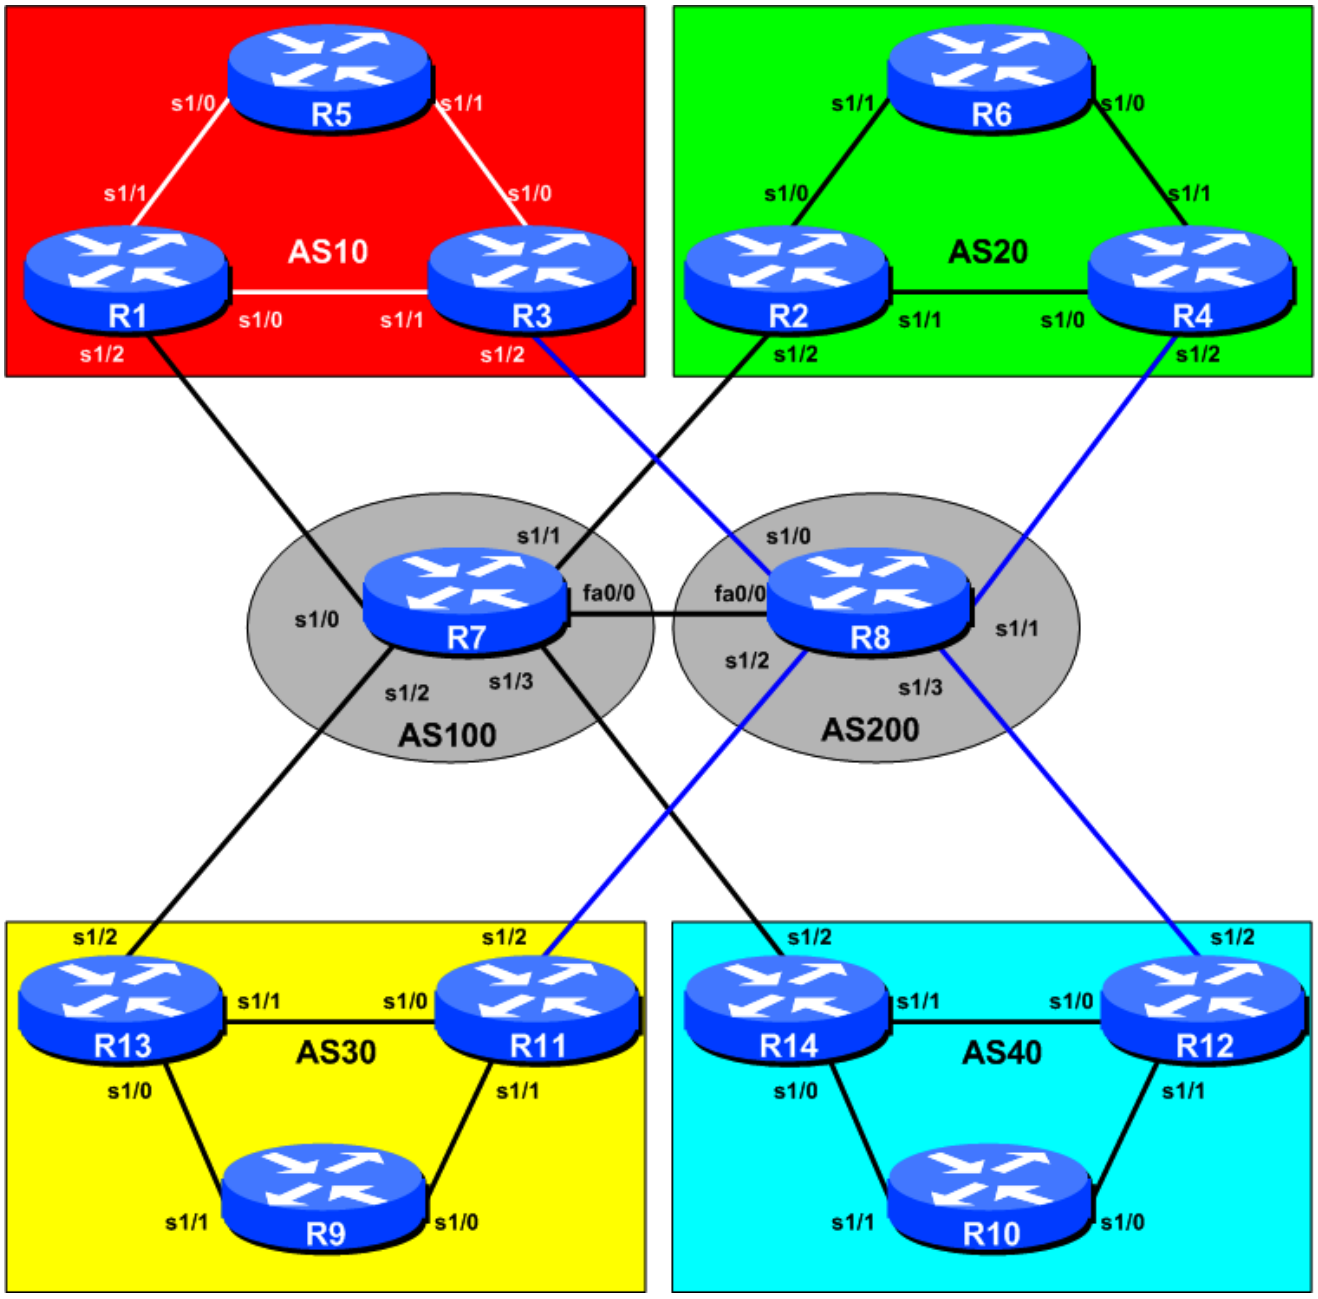

Figure 1 – ISP Lab Multihoming Configuration

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. 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.

AS10	**10.10.0.0/20**	**AS40**	**10.40.0.0/20**
AS20	**10.20.0.0/20**	**AS100**	**10.100.0.0/20**
AS30	**10.30.0.0/20**	**AS200**	**10.200.0.0/20**

3. **Routing Protocols.** The IGP (either OSPF or ISIS) and iBGP should now be configured between the routers for each AS. Any interfaces which will not be running the IGP *MUST* be marked as passive in the configuration. And don't forget to use BGP peer-groups for iBGP peers.

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.** AS100 and AS200 should now enable their eBGP link to each other. All router teams in these ASes must ensure that they can see all the prefixes of AS100 and AS200. 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 AS10 and its two upstreams.** AS10 should currently be running iBGP within its own network. To announce AS10's prefix to AS100 and AS200 we will take the /20 address block and announce it on both peerings between the ASNs. AS100 and AS200 will not announce any prefixes to AS10 – they will simply announce a default route. There is no need for any more specific routing information to be injected into the customer site at this stage in the lab.

6. **Prepare to enable eBGP between the other ASes and each of their two upstreams.** The configuration steps to enable BGP between AS20, AS30, AS40 and their two upstreams are the same as those being described for AS10.

7. **Create AS10 prefix lists.** First, create the prefix lists on the routers in AS10. Both Router1 and Router3 will announce the aggregate. Both will accept the default route. Example for Router1:

```
ip prefix-list myblock permit 10.10.0.0/20
ip prefix-list default permit 0.0.0.0/0
```

8. **Create AS100 and AS200 prefix-lists.** The routers in AS100 and AS200 should only accept those prefixes which the customer is entitled to announce. So a prefix list needs to be installed on both Router7 and Router8 to do this. For example:

```
ip prefix-list Customer permit 10.10.0.0/20
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 AS10, the link between Router1 and Router7 in AS100 is the main link (coloured in black) – the link between Router3 and Router8 in AS200 is the backup (coloured in blue). An example configuration for Router1 might be:

```
router bgp 10
 network 10.10.0.0 mask 255.255.240.0
```

```
 neighbor <router7> remote-as 100
 neighbor <router7> description Link to Router7 in AS100
 neighbor <router7> prefix-list myblock out
 neighbor <router7> prefix-list default in
!
ip route 10.10.0.0 255.255.240.0 null0
```

10. **Configure the backup link.** Configure the backup link between the customer 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 Router12:

```
ip prefix-list myblock permit 10.40.0.0/20
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 40 40 40
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 40
 network 10.40.0.0 mask 255.255.240.0
 neighbor <router8> remote-as 200
 neighbor <router8> description Link to Router8 in AS200
 neighbor <router8> prefix-list myblock out
 neighbor <router8> prefix-list default in
 neighbor <router8> route-map outfilter out
 neighbor <router8> route-map infilter in
!
ip route 10.40.0.0 255.255.240.0 null0
```

11. **Configure primary and back up links for the other 3 ASNs.** The teams in the other three customer ASes should use the above steps as a guidline to configure their primary and backup paths with AS100 and AS200. Refer to the diagram in Figure 1. The black coloured links are the primary paths, the blue coloured links are the backup paths.

12. **Configure eBGP in AS100 and AS200 with AS10.** AS100 and AS200 are going to originate the default route in the peering with AS10. The BGP command `default-originate` is used to do this. Example configuration for Router 7:

```
router bgp 100
 neighbor <router3> remote-as 10
 neighbor <router3> description Multihomed Customer
 neighbor <router3> default-originate
 neighbor <router3> prefix-list AS10 in
 neighbor <router3> prefix-list default out
!
```

Once the teams operating Router7 and Router8 have done the configuration for AS10, they should also complete the configuration for AS20, AS30 and AS40 customers.

13. **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 /21 comes in one path, and traffic for the second /21 comes in the other path. If either path fails, the advertisement of the /20 address block (aggregate) ensures continued connectivity.

14. **Clean up the configuration of AS10, AS20, AS30 and AS40.** 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.

15. **Configure the address block and subprefixes in the customer ASes.** Modify the router configuration so that the /20 address block and the two /21 subprefixes are present in the BGP table. Also set up prefix lists to cater for these blocks. For example:

```
ip prefix-list aggregate permit 10.10.0.0/20
!
ip prefix-list path1 permit 10.10.0.0/20
ip prefix-list path1 permit 10.10.0.0/21
!
ip prefix-list path2 permit 10.10.0.0/20
ip prefix-list path2 permit 10.10.8.0/21
!
ip prefix-list default permit 0.0.0.0/0
!
router bgp 10
 network 10.10.0.0 mask 255.255.240.0
 network 10.10.0.0 mask 255.255.248.0
!
ip route 10.10.0.0 255.255.240.0 null0
ip route 10.10.0.0 255.255.248.0 null0
```

16. **Configure BGP in the customer ASes.** For AS10, the link between Router1 and Router7 in AS100 is the first link – the link between Router3 and Router8 in AS200 is the second (and is the

one which should also announce the subprefix). 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* and *path1* as above, whereas Router3 could announce *aggregate* and *path2*. For example on Router12:

```
!
router bgp 40
 network 10.40.0.0 mask 255.255.240.0
 network 10.40.8.0 mask 255.255.248.0
 neighbor <router8> remote-as 200
 neighbor <router8> description Link to Router8 in AS200
 neighbor <router8> prefix-list path2 out
 neighbor <router8> prefix-list default in
!
ip route 10.40.8.0 255.255.248.0 null0
```

17. **Connectivity targets.** So that connectivity via each /21 can be tested, Routers 5, 6, 9 and 10 should set up a second loopback interface with an IP address from their ASN's respective second /21 block. Note that the /32 address should be announced by the IGP so that other routers in the ASN know how to get to the destination. The following configuration snippet shows a possible configuration for Router 10 using OSPF:

```
interface loopback 1
 ip address 10.30.15.10 255.255.255.255
 ip ospf 65533 area 0                    ! ONLY for IOS 12.4 and later
!
router ospf 65533
 network 10.30.15.10 0.0.0.0 area 0     ! ONLY for IOS <12.4
!
```

The following configuration snippet shows a possible configuration for Router 10 using ISIS:

```
interface loopback 1
 ip address 10.30.15.10 255.255.255.255
!
router isis workshop
 passive-interface loopback1
!
```

18. **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 /21 is taken out of the /20 address block and announced on one link between the customer and the ISP – and the /20 is announced with a longer AS-PATH on this link also. The other link sees just the standard announcement of the /20.

19. **Clean up the configuration of AS10, AS20, AS30 and AS40.** 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.

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

```
ip prefix-list aggregate permit 10.10.0.0/20
!
ip prefix-list subblocks permit 10.10.0.0/20 le 21
!
ip prefix-list default permit 0.0.0.0/0
!
router bgp 10
 network 10.10.0.0 mask 255.255.240.0
 network 10.10.0.0 mask 255.255.248.0
!
ip route 10.10.0.0 255.255.240.0 null0
ip route 10.10.0.0 255.255.248.0 null0
```

21. **Configure BGP in the customer ASes.** For AS10, the link between Router1 and Router7 in AS100 is the first link – the link between Router3 and Router8 in AS200 is the second (and is the one which should also announce the subprefix). 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 10 10 10
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 10
 network 10.10.0.0 mask 255.255.240.0
 neighbor <router13> remote-as 200
 neighbor <router13> description Link to Router8 in AS200
 neighbor <router13> prefix-list subblocks out
```

```
 neighbor <router13> prefix-list default in
 neighbor <router13> route-map outfilter out
 neighbor <router13> route-map infilter in
!
```

22. **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.*

23. **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.

24. **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.