# Module 17 – Advanced Community Usage (RFC1998++)

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

**Prerequisites: Modules 15 and 16, and Advanced Communities Presentation**

The following will be the common topology used.

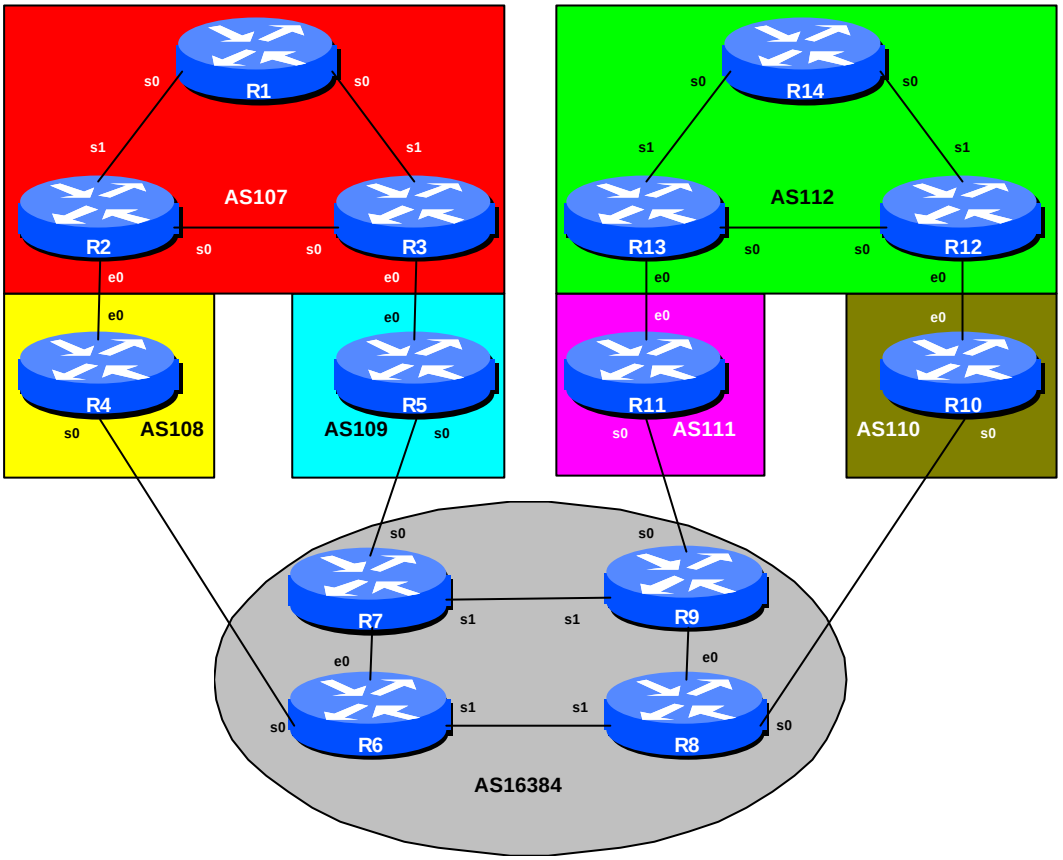


**Figure 1 – ISP Lab Multihoming Configuration**

## Lab Notes

The purpose of this module is to demonstrate complex multihoming using some example extensions to RFC1998 which are in use in many service provider backbones today.

The example covered here is a simple example of loadsharing between two upstream ISPs who have no interconnect. They do support the advanced community usage, which allows the customer to control loadsharing in very fine detail.

## 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.

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.

| | | | |
|---|---|---|---|
| **AS107** | **220.10.0.0/19** | **AS111** | **221.35.0.0/19** |
| **AS108** | **220.19.0.0/19** | **AS112** | **221.99.0.0/19** |
| **AS109** | **220.73.0.0/19** | **AS16384** | **222.11.0.0/16** |
| **AS110** | **221.19.0.0/19** | | |

3. **Routing Protocols.** OSPF (area 0 only) and iBGP should now be configured between the routers for each of AS 107, AS112 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.

***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.*

4. **Community definitions.** The following community definitions will be used in this module:

| | |
|---|---|
| **ASx:140** | **set local pref high on upstreams** |
| **ASx:130** | **set local pref low on upstreams** |
| **ASx:120** | **more preferred route – opposite of ASx:80** |
| **ASx:110** | **more preferred route if dual homed on ASx** |
| **ASx:100** | **preferred route** |
| **ASx:90** | **backup route if dualhomed on ASx** |
| **ASx:80** | **main link is to another ISP with the same AS path length** |
| **ASx:70** | **main link is to another ISP** |
| **ASx:60** | **As ASx:90 but add 2 times upstream AS** |
| **ASx:50** | **don't announce to any upstream** |
| **ASx:40** | **set local pref high on upstreams** |
| **ASx:30** | **set local pref low on upstreams** |

Refer to the accompanying presentation on the meaning of these community values. Note that these extra communities lets the end-site control how his networks are loadshared between his upstream ISPs and their upstreams – they move the control one step further away in the Internet, and allow the customer to fine tune their connectivity even more than is possible with existing BGP methods. The point is that the configuration is entirely in the customer's hand – he can control the loadsharing he desires simply by sending communities to his upstream ISP.

The route-map is an extended version of the one covered in Modules 15 and 16. It is left as an exercise to the reader to complete the route-map. Use the worksheet at the rear of this handout if required.

5. **Configure the main link.** Configure the main link between the customer AS and the ISP. For AS107, the link between Router2 and Router4 in AS108 is the main link – the link between Router3 and Router5 in AS109 is the backup. For AS112, the main link is between Router 12 and Router 10 in AS110. Example configuration for Router2:

```
ip prefix-list myblock permit 220.10.0.0/19
ip prefix-list default permit 0.0.0.0/0
!
route-map outfilter permit 10
 match ip address prefix-list myblock
 set community 109:100
route-map outfilter permit 20
!
route-map infilter permit 10
 match ip address prefix-list default
 set local-preference 100
route-map infilter permit 20
```

```
!
router bgp 107
 network 220.10.0.0 mask 255.255.224.0
 neighbor <router4> remote-as 108
 neighbor <router4> description Link to Router4 in AS108
 neighbor <router4> prefix-list myblock out
 neighbor <router4> prefix-list default in
 neighbor <router4> route-map outfilter out
 neighbor <router4> route-map infilter in
!
ip route 220.10.0.0 255.255.224.0 null0 250
```

6. **Configure the backup link.** Configure the link from AS107 to AS108. Set local preference on inbound announcements to 80 and send community 108:50 on outbound announcements. To do this, use a route-map on the peering – you will require an inbound and outbound route-map. Example configuration for Router3:

```
ip prefix-list myblock permit 220.10.0.0/19
ip prefix-list default permit 0.0.0.0/0
!
!
route-map outfilter permit 10
 match ip address prefix-list myblock
 set community 108:50
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 107
 network 220.10.0.0 mask 255.255.224.0
 neighbor <router5> remote-as 109
 neighbor <router5> description Link to Router5 in AS109
 neighbor <router5> prefix-list myblock out
 neighbor <router5> prefix-list default in
 neighbor <router5> route-map outfilter out
 neighbor <router5> route-map infilter in
!
ip route 220.10.0.0 255.255.224.0 null 0 250
```

7. **Configure the routers in AS108, AS109, AS110 and AS111 to peer with the customer ASes.** Using the definitions given in step 4 above, construct the appropriate route-map and configure

the routers in AS108 through AS111 to peer with the customer ASes (107 and 112). Don't forget the prefix-list to filter inbound and outbound announcements. And remember to originate the default route for the private ASes.

8. **Configure the routers in AS108, AS109, AS110 and AS111 to peer with AS16384.** This configuration is very similar to that covered in Module 13. Don't forget prefix-list filtering inbound and outbound, etc. Remember that the customer ASes (107 and 112) have sent a community ASx:60 which these transit ASes need to act upon. Remember that ASx:60 means add twice the upstream AS to any outbound announcement – in other words, an AS-PATH prepend. Example configuration on Router5 of AS109 in it's peering with AS16384:

```
ip prefix-list all-blocks permit 220.10.0.0/19
ip prefix-list all-blocks permit 220.73.0.0/19
ip prefix-list default permit 0.0.0.0/0
!
ip community-list 1 permit 108:60
!
route-map outfilter permit 10
 match ip community 1
 set as-path prepend 109 109
route-map outfilter permit 20
!
router bgp 109
 neighbor x.x.x.x remote-as 16384
 neighbor x.x.x.x description Link to Router7 in AS16384
 neighbor x.x.x.x prefix-list myblock out
 neighbor x.x.x.x prefix-list default in
 neighbor x.x.x.x route-map outfilter out
!
```

All transit routers in AS108 to AS111 should configure this, even though their customer ASes may not be sending the community. If the customer decides to change the loadsharing, he simply needs to send a different community – the upstream ISP should not need to change their configuration.

9. **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.*

10. **Swap primary and backup paths.** Alter the configuration on the border routers of AS107 and AS112 so that the primary path and backup paths are swapped. In other words, you now want to configure the path between Router3 and Router5 to be the main link. And similarly for AS112. **Hint:** All you have to do is change the "set community" command in the outbound route-map and clear the bgp session. And you will similarly need to change the local-preference setting for the inbound route-map.

11. **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.*

12. **Summary.** This module has given a brief example of how advanced community usage can greatly simplify ISP configuration and multihoming. Clearly there are more complex examples, which are left to the reader to consider. If there is time at the end of the workshop, ask the Instructor to test out some other scenarios.

    **Caution: Allowing a customer, or a customer of a customer, to modify your routing policy at their will through community manipulation could be potentially harmful to your network. Only allow what you want to be allowed, and ensure that your charging model is appropriate for the service on offer. While there are great benefits to be achieved through automation, due care and attention is required at all times so that the benefits aren't overridden by potential harm to the network.**

## *CONFIGURATION NOTES*

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