Module 14 – Loadsharing using Communities

Objective: To investigate loadsharing using BGP Communities.

Prerequisites: Module 13 and Multihoming Presentation

The following will be the common topology used.

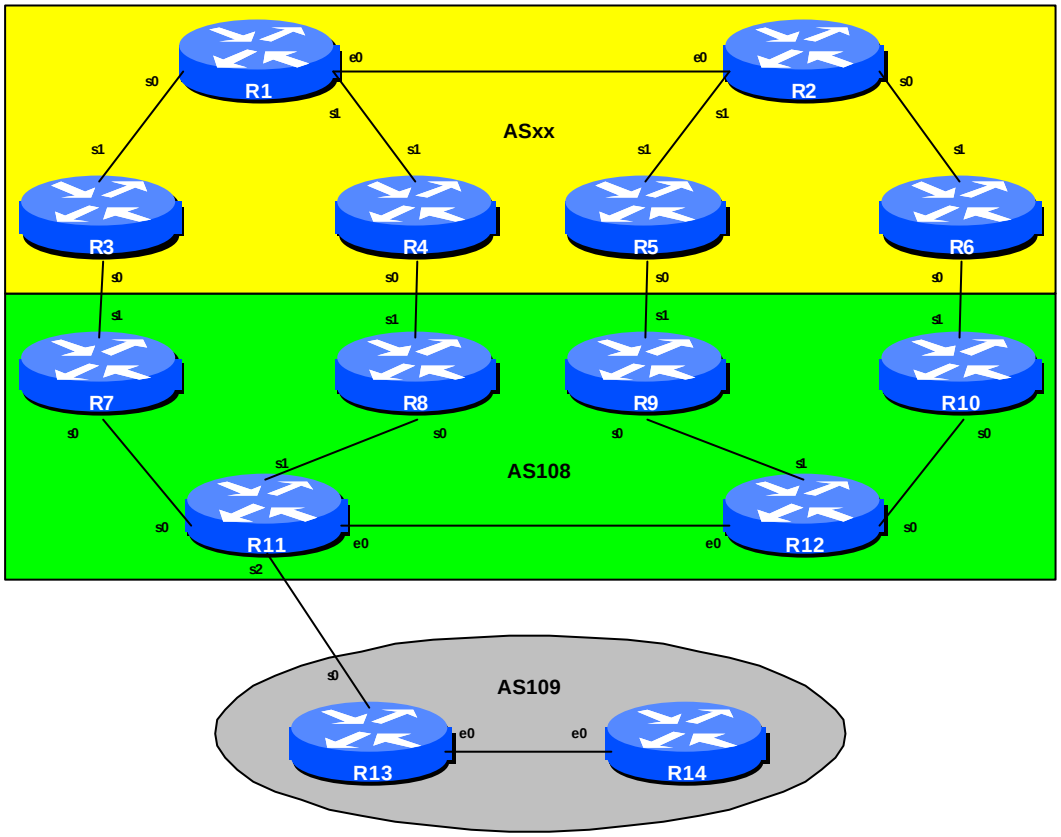

Figure 1 – ISP Lab Multihoming Configuration

Lab Notes

The purpose of this module is to demonstrate how to use communities for loadsharing. The example is one with four links between two ISP backbones. The autonomous system with ASxx is the end or origin AS. It receives transit through AS108 to reach AS109.

Two situations will be investigated. The first will involve a private AS, the other will involve a public AS. Conceptually there is little difference between the two cases – the former requires the private AS to be stripped out of any announcements to the "Internet", i.e. to AS109.

Cisco Systems Inc
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

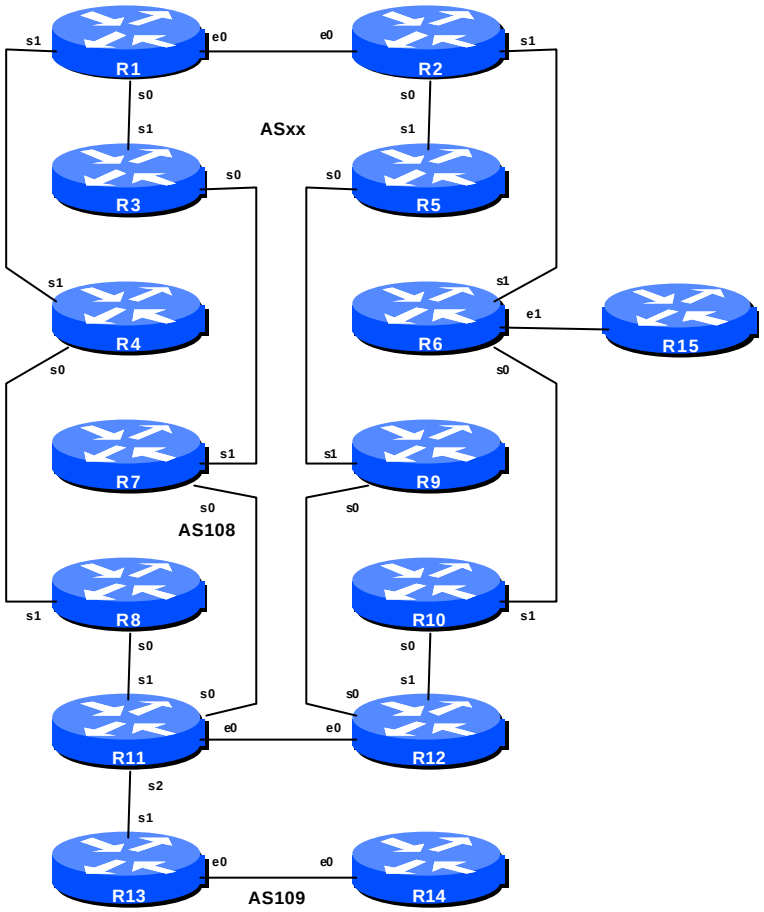

Figure 2 – ISP Lab Physical Layout

Lab Exercise

1. **Basic Configuration.** Each router team should configure their router to fit into the network layout depicted in Figure 2. **Notice that Router11 requires 3 serial ports** (the 3620s in the ISP Workshop kit have 4 serial ports). 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.

ASxx	**220.10.0.0/19**	**AS109**	**220.73.0.0/19**
AS108	**220.19.0.0/19**		

3. **Routing Protocols.** OSPF (area 0 only) should now be configured between the routers for each of ASxx, AS108 and AS109. Any interfaces which should not be running OSPF *MUST* be marked as passive in the configuration.

4. **Communities**. As communities are being used, don't forget to enable the new format style (which isn't yet a default in IOS) – recall this was covered in Module 11.

```
ip bgp-community new-format
```

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 OSPF prefixes in the routing table), please let the instructor know.*

Scenario One – Private AS multihoming using communities

The first scenario examines the use of communities for multihoming when the end-site has a private AS. Community usage is an alternative to using prefix-lists and the necessary filtering – recall that the "no-export" community can be used to automatically filter prefixes on eBGP peerings.

The AS to be used for the private AS is 65534. The way to implement this in AS65534 is to split the /19 into 4 pieces, one for each link. The aggregate is announced out of each link, as well as one of the subprefixes. The subprefix is tagged with the community "no-export" – this ensures that it will not be advertised to any another AS.

5. **Configure iBGP in each of the 3 ASes.** Previous modules should have made the workshop participants quite fluent in configuring iBGP by now. Don't forget to use peer-groups!

6. **Enable eBGP between AS108 and AS109.** AS108 and AS109 should now enable the BGP peering between themselves. You should see the other ASes prefixes appearing in the BGP table. 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.

7. **Prepare to enable eBGP between AS65534 and AS108.** AS65534 should currently be running iBGP within its own network. To announce AS65534's prefix to AS108 and AS109 we will take the /19 address block and divide it into four. The aim is to achieve relative even utilisation of the links between AS65534 and AS108, and common practice is to subdivide the address space. AS108 will not announce any prefixes to AS65534 – it will simply announce a default route. There is no need for any more routing information to be injected into the customer site.

8. **Create AS65534 prefix lists.** First, create the prefix lists on the routers in AS65534. For example, Router3 will announce the first subblock, Router4 will announce the second subblock, etc. All of Routers3 to 6 will accept the default route. Example for Router 3:

```
ip prefix-list subblocks1 permit 220.10.0.0/19
ip prefix-list subblocks1 permit 220.10.0.0/21
!
ip prefix-list firstblock permit 220.10.0.0/21
!
ip prefix-list default permit 0.0.0.0/0
```

9. **Create AS108 prefix-lists.** The routers in AS108 should only accept those prefixes which the customer is entitled to announce. So a prefix list needs to be installed on Routers 7 to 10 to do this. Notice in the example how a range of addresses has been listed in the permit statement – this is so that the customer can make changes to which block is being announced where without the upstream having to reconfigure their filters:

```
ip prefix-list Customer permit 220.10.0.0/19 le 21
ip prefix-list default permit 0.0.0.0/0
```

10. **Configure eBGP in AS65534.** With the prefix lists configured it is now possible to set up eBGP. It is good practice to configure the filters first, then configure BGP, not the other way around. This helps prevent accidents. Example configuration for Router 3:

```
ip prefix-list subblocks1 permit 220.10.0.0/19
ip prefix-list subblocks1 permit 220.10.0.0/21
!
ip prefix-list firstblock permit 220.10.0.0/21
ip prefix-list default permit 0.0.0.0/0
!
route-map Router7-out permit 10
 match ip address prefix-list firstblock
```

```
 set community no-export
route-map Router7-out permit 20
!
router bgp 65534
 neighbor <router7> remote-as 108
 neighbor <router7> send-community
 neighbor <router7> description Peering with Router 7 in AS108
 neighbor <router7> prefix-list subblocks1 out
 neighbor <router7> prefix-list default in
 neighbor <router7> route-map Router7-out out
!
```

Routers 4 to 6 will have similar configurations but using the other subprefix blocks of 221.10.0.0/19 address space.

11. **Configure eBGP in AS108.** AS108 is going to originate the default route in the peering with AS65534. The BGP command "default-originate" is used to do this. Example configuration for Router 7:

```
router bgp 108
 neighbor <router3> remote-as 65534
 neighbor <router3> description Multihomed Customer
 neighbor <router3> default-originate
 neighbor <router3> prefix-list Customer in
 neighbor <router3> prefix-list default out
!
```

12. **AS108 eBGP configuration with AS109.** Without further configuration changes in AS108, AS65534 will be announced by AS108 to AS109 routers to other ASes. To prevent this, the "remove-private-AS" BGP command is required. The example below is for Router11 in AS108:

```
ip prefix-list mynets permit 220.10.0.0/19 le 21
ip prefix-list mynets permit 220.19.0.0/19
!
router bgp 108
 neighbor <router13> remote-as 109
 neighbor <router13> remove-private-AS
 neighbor <router13> description Peering with AS109 – The Internet
 neighbor <router13> prefix-list mynets out
```

Notice the mynets prefix-list! The full range of prefixes which could be announced by AS65534 has been permitted. This is in case AS65534 wishes to leak subprefixes further out into the Internet by not tagging them with "no-export" community.

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 check the path being followed. Also check that backup via one of the alternative paths still functions (do this by disconnecting the cable between the two routers on the primary path) – you will see that a 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.*

Scenario Two – Public AS multihoming using communities

The first scenario examines the use of communities for multihoming when the end-site has a private AS. Community usage is an alternative to using prefix-lists and the necessary filtering – recall that the "no-export" community can be used to automatically filter prefixes on eBGP peerings.

The AS to be used for the public AS is 107. The way to implement this in AS107 is to split the /19 into 4 pieces, one for each link. The aggregate is announced out of each link, as well as one of the subprefixes. The subprefix is tagged with the community "no-export" – this ensures that it will not be advertised to any another AS.

14. **Remove the all BGP configuration involving AS65534.** Remove the BGP configuration relating to AS65534 from the previous module. The iBGP configuration of AS108 and AS109 can remain. The router teams in AS107 should reconfigure their iBGP mesh to use AS107. The router teams in AS108 should remove all configuration related to AS65534.

15. **Prepare to enable eBGP between AS107 and AS108.** AS107 should now be running iBGP within its own network. To announce AS107's prefix to AS108 and AS109 we will take the /19 address block and divide it into four. The aim is to achieve relatively even utilisation of the links between AS107 and AS108, and common practice is to subdivide the address space. AS108 will not announce any prefixes to AS107 – it will simply announce a default route. There is no need for any more routing information to be injected into the customer site.

16. **Create AS107 prefix lists.** First, create the prefix lists on the routers in AS107. For example, Router3 will announce the first subblock, Router4 will announce the second subblock, etc. All of Routers3 to 6 will accept the default route. Example for Router 3:

```
ip prefix-list subblocks1 permit 220.10.0.0/19
```

```
ip prefix-list subblocks1 permit 220.10.0.0/21
!
ip prefix-list firstblock permit 220.10.0.0/21
!
ip prefix-list default permit 0.0.0.0/0
```

17. **Create AS108 prefix-lists.** The routers in AS108 should only accept those prefixes which the customer is entitled to announce. So a prefix list needs to be installed on Routers 7 to 10 to do this. Notice in the example how a range of addresses has been listed in the permit statement – this is so that the customer can make changes to which block is being announced where without the upstream having to reconfigure their filters:

```
ip prefix-list Customer permit 220.10.0.0/19 le 21
ip prefix-list default permit 0.0.0.0/0
```

18. **Configure eBGP in AS107.** With the prefix lists configured it is now possible to set up eBGP. It is good practice to configure the filters first, then configure BGP, not the other way around. This helps prevent accidents. Example configuration for Router 3:

```
ip prefix-list subblocks1 permit 220.10.0.0/19
ip prefix-list subblocks1 permit 220.10.0.0/21
!
ip prefix-list firstblock permit 220.10.0.0/21
ip prefix-list default permit 0.0.0.0/0
!
route-map Router7-out permit 10
 match ip address prefix-list firstblock
 set community no-export
route-map Router7-out permit 20
!
router bgp 107
 neighbor <router3> remote-as 108
 neighbor <router3> send-community
 neighbor <router3> description Peering with Router 7 in AS108
 neighbor <router3> prefix-list subblocks1 out
 neighbor <router3> prefix-list default in
 neighbor <router3> route-map Router7-out out
!
```

Routers 4 to 6 will have similar configurations but using the other subprefix blocks of 221.10.0.0/19 address space.

Cisco Systems Inc
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

19. **Configure eBGP in AS108.** AS108 is going to originate the default route in the peering with AS107. The BGP command "default-originate" is used to do this. Example configuration for Router 7:

```
router bgp 108
 neighbor <router3> remote-as 107
 neighbor <router3> description Multihomed Customer
 neighbor <router3> default-originate
 neighbor <router3> prefix-list Customer in
 neighbor <router3> prefix-list default out
!
```

20. **AS108 eBGP configuration with AS109.** Without further configuration changes in AS108, everything learnt by AS108 will be announced to AS109 routers and to other ASes. The previous example used prefix-lists to implement the filtering. This time as-path filters will be used. The example below is for Router11 in AS108:

```
ip as-path access-list 1 permit ^107$
ip as-path access-list 1 permit ^$
!
router bgp 108
 neighbor <router13> remote-as 109
 neighbor <router13> description Peering with AS109 – The Internet
 neighbor <router13> filter-list 1 out
```

    The filter list allows AS107 to leak subprefixes further out into the Internet by not tagging them with "no-export" community.

21. **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 one of the alternative paths still functions (do this by disconnecting the cable between the two routers on the primary path) – you will see that a 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.*

22. **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 communities.

Cisco Systems Inc
170 West Tasman Drive.
San Jose, CA 95134-1706
Phone: +1 408 526-4000
Fax: +1 408 536-4100

CONFIGURATION NOTES

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