Module 18 – Multihoming Case Study

Objective: Starting from a single homed ISP, the objective is to build a network where the local ISP is multihomed to their upstream and to a local peer.

Prerequisites: Modules 11, 12 and 14

The following will be the common topology used.

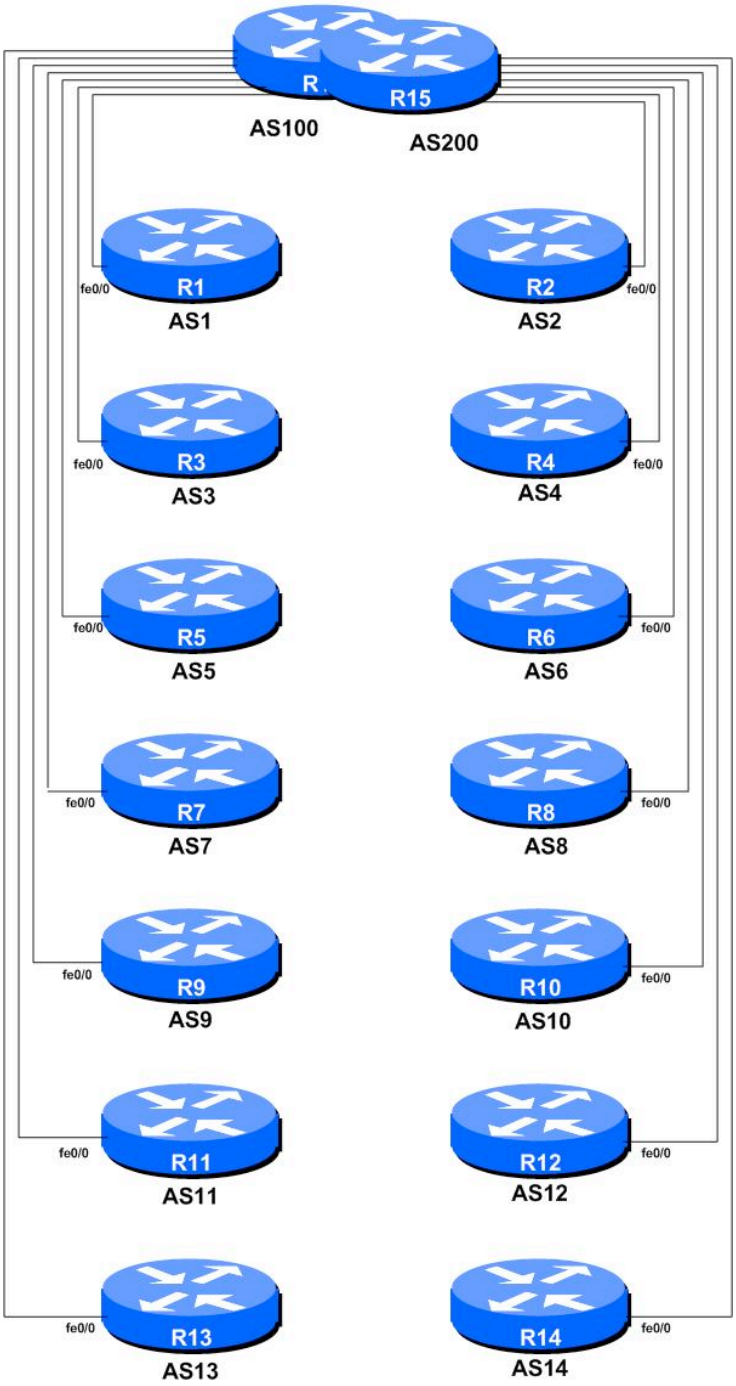

Figure 1 – ISP Lab Basic Configuration

Lab Notes

The purpose of this module is to construct a network building from a very basic singlehomed scenario to end with a situation where the ISP is multihomed between their upstream and another local ISP.

This module demonstrates the good principles of configuring eBGP to not provide transit to local peers, and configuring iBGP to ensure that functions with the correct next-hop values. The further aim of the module is to show how straightforward it is to transition from a single homed connection to an upstream provider, to having a true multihomed configuration for the local AS.

Lab Exercise

The following list is typical for what needs to be done to bring up the lab configuration:

1. **Remove any configuration from the previous Modules.** Tidy up from the previous Modules so that the router has no configuration on it. Either do a "write erase" and then reboot the router (strongly recommended) or carefully remove the IP addressing, OSPF and BGP configuration from previous modules (not really recommended as classroom teams invariably leave some old configuration on the router, causing confusion further on in the module).

2. **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 Ethernet cables. Remember what was covered in Module 11! The workshop instructors will configure router15 to provide the necessary network configuration (they will reconfigure one of the Ethernet switches to use .1q VLAN trunking, thereby making Router15 appear as though it has 14 Ethernet ports on it).

3. **Addressing Plan.** These address ranges should be used throughout this module. You are welcome to use your own range within your AS if you really desire, just so long as you consult with the teams in other ASes to ensure there is no overlap, and that you let the workshop instructors know that you've used different address blocks. In the every day Internet, such address assignment is carried out by the Regional Internet Registries. AS100 and AS200 are the upstream providers used in this module. A /16 network block has been assigned to those providers – you will find out what this block is later on in this module.

AS1	**120.10.0.0/19**	**AS8**	**121.19.32.0/19**
AS2	**120.10.32.0/19**	**AS9**	**121.35.0.0/19**
AS3	**120.19.0.0/19**	**AS10**	**121.35.32.0/19**
AS4	**120.19.32.0/19**	**AS11**	**121.99.0.0/19**
AS5	**120.73.0.0/19**	**AS12**	**121.99.32.0/19**
AS6	**120.73.32.0/19**	**AS13**	**121.158.0.0/19**
AS7	**121.19.0.0/19**	**AS14**	**121.158.32.0/19**

4. **Addressing the router.** Each router team should come up with an appropriate addressing plan for their router. Remember that a Loopback interface will be required (for BGP router-id), as well as addresses for future point to point links. The lab instructors will have assigned point-to-point link addresses from Router15 to each of the classroom routers. The individual address blocks for these are as follows – Router15 uses the low address of the /30, the classroom router uses the high address:

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

Router1	**192.168.250.0/30**	**Router2**	**192.168.251.0/30**
Router3	**192.168.250.4/30**	**Router4**	**192.168.251.4/30**
Router5	**192.168.250.8/30**	**Router6**	**192.168.251.8/30**
Router7	**192.168.250.12/30**	**Router8**	**192.168.251.12/30**
Router9	**192.168.250.16/30**	**Router10**	**192.168.251.16/30**
Router11	**192.168.250.20/30**	**Router12**	**192.168.251.20/30**
Router13	**192.168.250.24/30**	**Router14**	**192.168.251.24/30**

5. **Static Routes.** With the basic configuration in place on the router, set up a static route so that you can ping the upstream (AS100 or AS200) router. Also talk to the lab assistant who is configuring the upstream router to ensure that they will put in a static route which allows them to see your router. Note that the odd numbered routers are connecting to AS100, the even numbered routers are connecting to AS200.

6. **Summary.** Once all the configuration has been completed, each router team should be able to ping every router in the classroom. This is quite a typical scenario, where a transit ISP is providing connectivity to several customer networks, with single-homing only.

Checkpoint #1: *call lab assistant to verify the connectivity. Save the configuration as it is on the router either on the worksheet on the end of this hand out, or own your own laptop, or on the classroom tftp server if it is available.*

STOP AND WAIT HERE

Scenario One – eBGP with upstream

7. **Set up eBGP with upstream ISP.** For the first exercise we are simply going to configure eBGP with our upstream ISP, AS100 or AS200.This typifies the situation where an ISP is only connected to a single upstream, and is about to consider the migration to a multihoming situation. An extract of a configuration might be:

```
!
router bgp 10
 network 121.35.32.0 mask 255.255.224.0
 neighbor x.x.x.x remote-as 100
 neighbor x.x.x.x description Peering with Upstream AS200
 neighbor x.x.x.x password cisco
!
ip route 121.35.32.0 255.255.224.0 null0
```

Note that the upstreams AS100 or AS200 will only originate a default route. They will not announce any routes apart from the default to you as you don't need any other routing information to see the whole Internet. Also note that we are not using any prefix filtering here. Prefix filtering will be covered later on in this module.

Don't forget the required BGP configuration, so items like disabling synchronization and auto-summarisation, enabling logging, and setting the BGP distances appropriately. If you have forgotten, refer to your configuration in Module 11 and subsequent modules. Note that the odd numbered routers are connecting to AS100, the even numbered routers are connecting to AS200.

8. **Remove Static Routes.** Once the BGP peering is up and running, remove the static default route pointing towards your upstream provider. Also talk to the lab assistant who is monitoring the upstream provider router (Router 15) so that they will remove the static route pointing to your network. This step represents the step when you migrate from using static routing with your upstream to using BGP. Once you see your prefix in the Internet Routing Table, and are seeing the default from the upstream by BGP, you then know that you can safely remove the static routes which were set up.

9. **Status of BGP peering.** Use the BGP show commands to find out what is in your BGP table. What do you see? You should see a default route from your upstream, your own address block and the address blocks from the other ASNs transited by AS100 and AS200. If this is not the case, either try and work out what is wrong, or ask the lab assistants to help you debug the problem.

10. **Status of connectivity.** Use the "show ip route" commands to find out how routing is on the system. Also use trace and ping to check the routing to other ASNs in the classroom. You should find that AS100 and AS200 are providing transit to all the other ASNs in the network.

Checkpoint #2: *call lab assistant to verify the connectivity. Save the configuration as it is on the router either on the worksheet on the end of this hand out, or own your own laptop, or on the classroom tftp server if it is available.*

STOP AND WAIT HERE

Scenario Two – Adding in local peer.

11. **eBGP with local peer.** We now need to add configuration to the network such that ASx sets up a local eBGP session with ASx+1 – in other words, AS1 needs to set up eBGP with AS2, AS3 needs to set up eBGP with AS4, and so on. This situation is depicted in Figure 2. The connections between neighbour ASNs are formed using serial cables, connecting the serial interfaces on each router. The lab instructors will have already connected the cables between the routers before you start this exercise.

12. **Serial Connections.** Verify the serial connections to the neighbouring router. Establish whether you have the DTE or DCE cable connected to your router (if the latter, you will need to provide clock for the "circuit"). Agree on the IP addresses for the point to point link (remember, you will use a /30).

13. **Configuration**. Leaving the existing eBGP configuration with your upstream on the router, now add in the necessary eBGP configuration to your local peer. An example configuration might be:

```
!
router bgp 7
 neighbor y.y.y.y remote-as 8
 neighbor y.y.y.y description Peering with Local Peer AS8
 neighbor y.y.y.y password cisco
!
ip route 121.19.0.0 255.255.224.0 null0
```


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

Note that we still have no prefix filtering here. The purpose is to demonstrate how eBGP can go really badly wrong if no filters are used. This configuration must **NEVER** be used on the public Internet without proper filters.

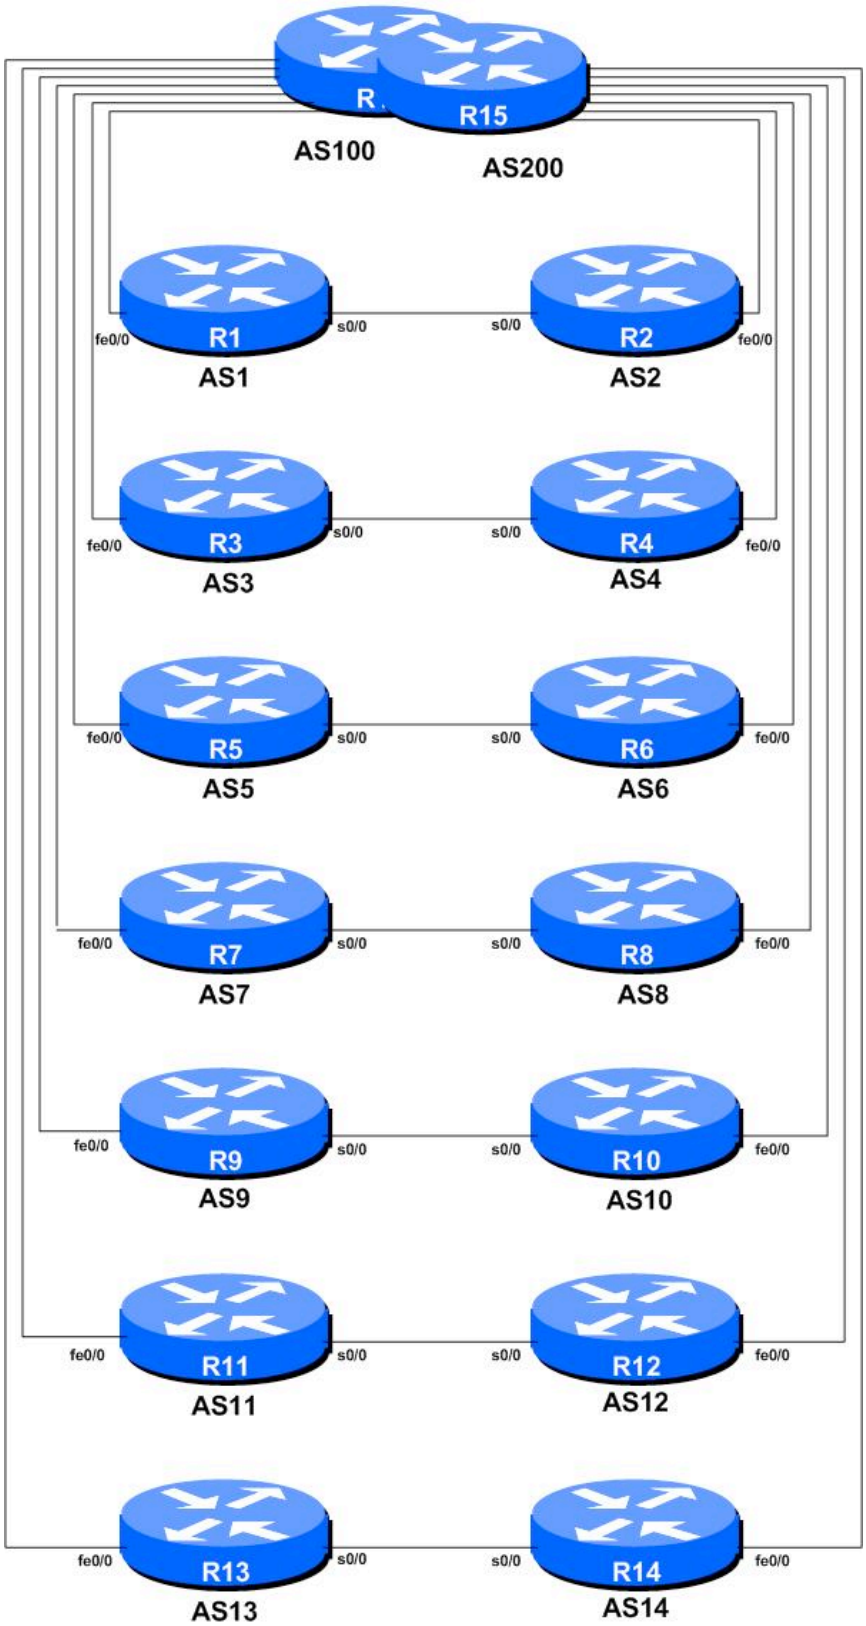

Figure 2 - Peering with local peer

14. **Status of BGP peering.** Use the BGP show commands to find out what is in your BGP table. What do you see? You should see a default route from your upstream, your own address block, the address block announced to you by your neighbouring AS, as well as the address blocks from all the other ASNs in the network. If this is not the case, either try and work out what is wrong, or ask the lab assistants to help you debug the problem.

15. **Status of connectivity.** Use the "show ip route" commands to find out how routing is on the system. Also use trace and ping to check the routing to other ASNs in the classroom. How do you reach the other ASNs in the classroom? Why?

Checkpoint #3: *call lab assistant to verify the connectivity. Save the configuration as it is on the router either on the worksheet on the end of this hand out, or own your own laptop, or on the classroom tftp server if it is available.*

STOP AND WAIT HERE

Scenario Three – Filtering using AS PATH filters.

16. **Filtering BGP sessions using AS-PATH.** So far we have used no filtering whatsoever on the BGP sessions. And as you saw, the BGP table ended up in a significant mess, with surplus information from the upstream ISPs AS100 and AS200, and with your local peering AS providing you with free transit to other ASNs in the network. The latter situation is commercially a disaster for any ISP, the former simply unnecessarily wastes lots of router memory.

17. **Configuration**. Leaving the existing eBGP configuration with AS100/200 and your local peer on the router, now add in the necessary configuration so that you are using AS-PATH filters on your eBGP peerings. AS100/200 should be filtered so that only the prefixes it originates (the default route) is permitted in. The peer AS should also be filtered in a similar fashion. An example configuration might be:

```
!
router bgp 7
 neighbor x.x.x.x remote-as 100
 neighbor x.x.x.x filter-list 1 out
 neighbor x.x.x.x filter-list 2 in
 neighbor y.y.y.y remote-as 8
 neighbor y.y.y.y filter-list 1 out
 neighbor y.y.y.y filter-list 3 in
!
ip as-path access-list 1 permit ^$
ip as-path access-list 2 permit _100$
ip as-path access-list 3 permit _8$
!
```

Note that the neighbouring AS will only announce the prefixes they originate. Also note that the upstream will only send you a default route – a default route implies the whole Internet so there is really no need to get any further information from the upstream. The inbound AS-PATH filter is really important. You simply cannot assume that your eBGP peers will do prefix filtering as you

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

are doing, so it is vitally important that you have inbound (and outbound) route filters on your eBGP sessions.

18. **Status of BGP peering.** Use the BGP show commands to find out what is in your BGP table. What do you see? You should see a default route from your upstream, your own address block, and the address block announced to you by your neighbouring AS. If this is not the case, either try and work out what is wrong, or ask the lab assistants to help you debug the problem.

19. **Status of connectivity.** Use the "show ip route" commands to find out how routing is on the system. Also use trace and ping to check the routing to other ASNs in the classroom. You should find that AS100 & AS200 are providing transit to all the other ASNs apart from the local ASN you are peering with.

Checkpoint #4: *call lab assistant to verify the connectivity. Save the configuration as it is on the router either on the worksheet on the end of this hand out, or own your own laptop, or on the classroom tftp server if it is available.*

STOP AND WAIT HERE

Scenario Four – Filtering using prefix filters.

20. **Filtering BGP sessions using prefix filters.** Remove the AS-PATH filters and their application in the eBGP peering we used in the previous scenario. We are now going to look at a more strict filtering mechanism, using prefix-lists. While these are higher maintenance, they are generally preferred, as they don't imply a trust of what the peer AS is introducing into the routing system.

21. **Configuration**. Leaving the existing eBGP configuration with AS100 and AS200 as appropriate and with your local peer on the router, now add in the necessary configuration so that you are using prefix filters on your eBGP peerings. An example configuration might be:

```
!
router bgp 7
 neighbor x.x.x.x remote-as 100
 neighbor x.x.x.x prefix-list ourblock out
 neighbor x.x.x.x prefix-list default in
 neighbor y.y.y.y remote-as 8
 neighbor y.y.y.y prefix-list ourblock out
 neighbor y.y.y.y prefix-list AS8block in
!
ip prefix-list default permit 0.0.0.0/0
ip prefix-list ourblock permit 121.19.0.0/19
ip prefix-list AS8block permit 121.19.32.0/19
!
```

Note that the neighbouring AS will only announce the prefixes they originate. Also note that the upstream will only send you a default route – a default route implies the whole Internet so there is really no need to get any further information from the upstream. The inbound prefix filter is really important. You simply cannot assume that your eBGP peers will do prefix filtering as you are doing, so it is vitally important that you have inbound (and outbound) route filters on your eBGP sessions.

22. **Status of BGP peering.** Use the BGP show commands to find out what is in your BGP table. What do you see? You should see a default route from your upstream, your own address block, and the address block announced to you by your neighbouring AS. If this is not the case, either try and work out what is wrong, or ask the lab assistants to help you debug the problem.

23. **Status of connectivity.** Use the "show ip route" commands to find out how routing is on the system. Also use trace and ping to check the routing to other ASNs in the classroom. You should find that AS100 and AS200 are providing transit to all the other ASNs apart from the local ASN you are peering with.

24. **AS-PATH and prefix-lists.** Some ISPs merge the configurations we examined in Scenarios 3 and 4 so that they have prefix-lists and AS-PATH filters applied to each eBGP session. The reason for doing this is for backup – if the prefix-list becomes broken for some reason, the AS-PATH filter still provides some suitable filtering to protect the local ASNs routing system.

Checkpoint #5: *call lab assistant to verify the connectivity. Save the configuration as it is on the router either on the worksheet on the end of this hand out, or own your own laptop, or on the classroom tftp server if it is available.*

STOP AND WAIT HERE

Scenario Five – Providing backup for the local peer.

25. **Background.** The next exercise in the BGP labs is to modify the policy configuration on the local routers so that the local peers provide backup for each other. The network layout is still as depicted in Figure 2.

26. **Configuration.** The technique for providing local backup is a modification of the previous step. The router in the local AS will accept the prefix announced by its local peer, and announce this prefix to the upstream with a longer AS-PATH. In this example we will use a 4 times pre-pend. The router in the local AS will also announce the default route to the local peer, again with a pre-pend – we will use a 2 times pre-pend in this case.

27. **Reconfiguring eBGP with the upstream.** Because we want to send our local peer's address block to our upstream, we need to modify our filters to allow their address block out to our upstream. But, unlike Scenario 2, we need to apply outbound policy on their address block so that they don't use us as the primary path into their network. To do this, we use a route-map with an as-path prepend option. An example configuration might be:

```
!
router bgp 7
 neighbor x.x.x.x remote-as 100
 neighbor x.x.x.x prefix-list AS7out out
 neighbor x.x.x.x prefix-list default in
 neighbor x.x.x.x route-map AS8-backup out
!
ip prefix-list default permit 0.0.0.0/0
ip prefix-list AS7out permit 121.19.0.0/19
ip prefix-list AS7out permit 121.19.32.0/19
```


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

```
ip prefix-list AS8block permit 121.19.32.0/19
!
route-map AS8-backup permit 10
 match ip address prefix-list AS8block
 set as-path prepend 7 7 7 7
route-map AS8-backup permit 20
!
```

Note that the local peer's address block has been added to our outbound prefix-list to our upstream. A separate prefix-list has been set up for their address block for use in the pre-pending route-map. There is no way of concatenating prefix-lists in IOS at the moment.

28. **Reconfiguring eBGP with our local peer.** We also want to send the default route we hear from our upstream to our local peer. It is very important that we transit the default route, rather than originating a default route ourselves. If we do the latter, a failure in our upstream link will mean that we end up blackholing our local peers traffic as well as our own. An example configuration might be:

```
!
router bgp 7
 neighbor y.y.y.y remote-as 8
 neighbor y.y.y.y prefix-list AS8out out
 neighbor y.y.y.y prefix-list AS8in in
 neighbor y.y.y.y route-map default-prepend out
 neighbor y.y.y.y route-map set-lp-low in
!
ip prefix-list default permit 0.0.0.0/0
ip prefix-list AS8out permit 121.19.0.0/19
ip prefix-list AS8out permit 0.0.0.0/0
ip prefix-list AS8in permit 121.19.32.0/19
ip prefix-list AS8in permit 0.0.0.0/0
!
route-map default-prepend permit 10
 match ip address prefix-list default
 set as-path prepend 7 7
route-map default-prepend permit 20
!
route-map set-lp-low permit 10
 match ip address prefix-list default
 set local-preference 50
route-map set-lp-low permit 20
!
```

Note that the default route has been added to our outbound prefix-list to our local peer. Once you have applied the new configuration, don't forget to use route refresh to update the routing information heard from and sent to your BGP peers.

29. **Status of BGP peering.** Use the BGP show commands to find out what is in your BGP table. What do you see? You should still see a default route from your upstream, a default route from your local peer which is prepended by two of their ASN, your own address block, and the address block announced to you by your neighbouring AS. If this is not the case, either try and work out what is wrong, or ask the lab assistants to help you debug the problem.

30. **Testing the backup.** The lab assistants will now demonstrate what is seen on the two upstream routers in AS100 and AS200. You should see your path learned directly from your router, as well

as a path with a 4 times prepend learned via your local peer. If this is the case, try disconnecting the link between you and your upstream, and see if the back up works. If not, try and work out what is wrong, or ask the lab assistants for help.

Checkpoint #6: *call lab assistant to verify the connectivity. Save the configuration as it is on the router either on the worksheet on the end of this hand out, or own your own laptop, or on the classroom tftp server if it is available.*

STOP AND WAIT HERE

Scenario Six – Adding iBGP.

31. **Background.** The final exercise in the BGP labs is to simulate a situation which is commonly found in many service provider networks around the world. The situation is where an ISP multihomes between two upstream ISPs, and uses two local routers for redundancy. This necessitates using eBGP towards the upstreams, and iBGP between the two local routers. The network layout is depicted in Figure 3.

32. **Reconfiguring BGP between local ASNs**. The eBGP session between ASx and ASx+1 should be removed now. We will be converting this to iBGP. For example, the router in AS1 should remove the eBGP peering with the router in AS2, etc.

33. **Changing ASN**. The routers in AS1, AS3, AS5, AS7, AS9, AS11 and AS13 should now reconfigure their BGP session so that they are now members of AS2, AS4, AS6, AS8, AS10, AS12, and AS14 respectively. There is no simple way of doing this on the router, apart from simply deleting the entire "router bgp" configuration and re-entering the configuration with the new AS number. Note that these routers should still peer with their upstream AS100.

34. **Implementing iBGP**. We are now going to set up iBGP between neighbouring routers. The important point to remember here is that iBGP runs between the loopback interfaces of the router. For the loopback to be functional, OSPF needs to run on the router first. OSPF is used so that the two routers in the network can work out how to reach each other.

35. **Setting up OSPF**. Make sure that the loopback interface is configured on the router. Now set up OSPF – there is only one active interface, the interface pointing towards the neighbouring router in the AS. Also make sure that the loopback interface address is configured in OSPF. Do not put the external peering interfaces in OSPF (so those towards AS100 and AS200 respectively) as we will be using the BGP next-hop-self concept to ensure reachability of external networks. A configuration example might be:

```
router ospf 41
 log-adjacency-changes
 passive-interface default
 no passive-interface serial 0/0
 area 0 authentication message-digest
 network x.x.x.x 0.0.0.3 area 0     ! p2p link to neighbour
 network 1.1.1.1 0.0.0.0 area 0     ! loopback interface address
!
```


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

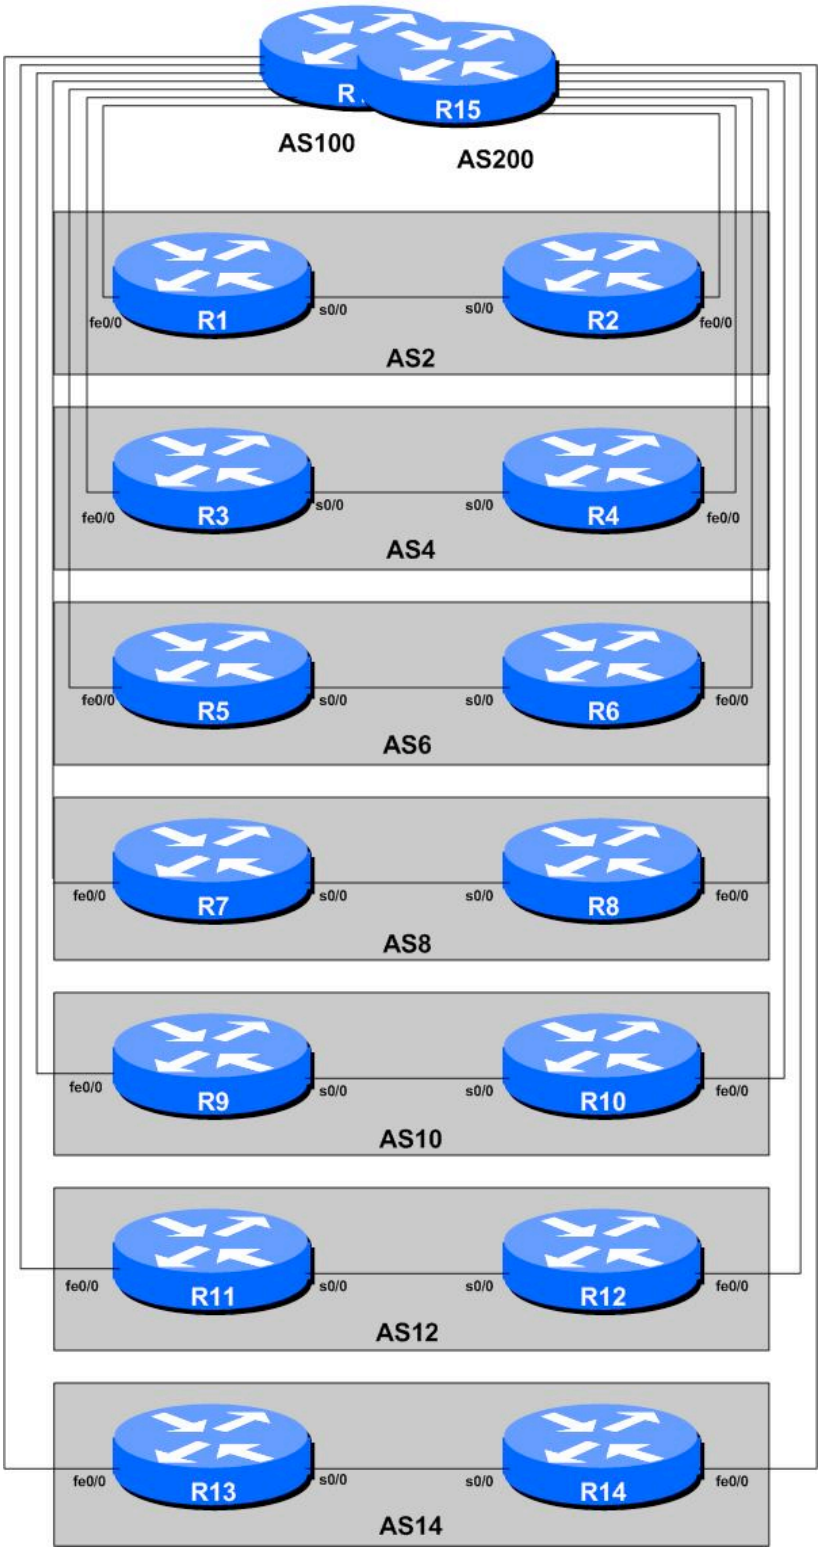

Figure 3 - iBGP and eBGP

36. **Setting up iBGP**. Once you can ping the loopback interface of the neighbouring router in your AS, set up iBGP between your and the neighbouring router. Because we are not carrying the external facing addresses in OSPF, we need to use the BGP next-hop-self concept to ensure reachability of external networks. An example configuration might be:

```
!
router bgp 8
 network 121.19.0.0 mask 255.255.224.0
 network 121.19.32.0 mask 255.255.224.0
 neighbor ibgp-peers peer-group
 neighbor ibgp-peers remote-as 8
 neighbor ibgp-peers password cisco
 neigbhor ibgp-peers send-community
 neighbor ibgp-peers next-hop-self
 neighbor ibgp-peers update-source loopback 0
 neighbor x.x.x.x remote-as 100
 neighbor x.x.x.x prefix-list ourblock out
 neighbor x.x.x.x prefix-list default in
 neighbor x.x.x.x password cisco
 neighbor y.y.y.y peer-group ibgp-peers
!
ip prefix-list default permit 0.0.0.0/0
ip prefix-list ourblock permit 121.19.0.0/19
ip prefix-list ourblock permit 121.19.32.0/19
!
```

37. **Status of BGP peering.** Use the BGP show commands to find out what is in your BGP table. What do you see? You should see a default route from your upstream and your own address blocks. If this is not the case, either try and work out what is wrong, or ask the lab assistants to help you debug the problem.

38. **Status of connectivity.** Use the "show ip route" commands to find out how routing is on the system. Also use trace and ping to check the routing to other ASNs in the classroom. You should find that either AS100 or AS200 is providing transit to all the other ASNs in the network.

39. **Extra tasks**. If you have time, notice that you can aggregate the two /19 address blocks used in your ASN. So change the configuration on both router so that you are announcing a /18 rather than two /19s. Change the filters and the BGP configuration to suit. Do you still have the same connectivity? If not, investigate why not. Do you need to tell any BGP neighbours about your changes?

Checkpoint #7: *call the lab assistant to verify the connectivity.*

40. **Summary.** This module has given an example of the kind of migration strategy and steps which are necessary to move from a single statically routed connection to an upstream, to using a BGP multihomed connection with two routers. Each router team has been encouraged to save their configuration steps covered in each scenario – as these configuration steps are exactly those required in a real live situation on the Internet.

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.