Module 12 – Multihoming to the Same ISP

Objective: To investigate various methods for multihoming onto the same upstream's backbone

Prerequisites: Module 11 and Multihoming Presentation

The following will be the common topology used.

Figure 1 – ISP Lab Configuration

CISCO SYSTEMS

Lab Notes

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

- Enterprise customer requires more than one connection to the service provider to provide resiliency, and/or loadsharing.
- Enterprise customer has multiple sites which require connection to the Internet.
- Start-up ISP requires more than a single link to the Internet, but has little requirement to connect to more than one upstream ISP

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.

The accepted way to multihome to a single upstream ISP on the Internet today is to use a private AS number. The IANA defines the range 64512 to 65534 as being private ASes, in the same way that RFC1918 defines private address space. These ASes should **<u>never</u>** be visible on the Internet.

Finally, 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 more detailed policy**

There are rarely any exceptions to this. Prefix lists are very efficient access-lists and they make the 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. **Physical Layout.** 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. **IMPORTANT:** Each router team should ensure that their router has the basic configuration as covered in the first steps of Module 11. Figure 2 shows the physical layout of the lab. Note that you will have to move routers around to ensure you have sufficient ports. Also, try and minimise the disruption as you recable the lab.

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

2

**Figure 2 – Multihoming Lab Physical Layout**

2. **Basic Configuration.** Before starting to configure anything for this module, each router team should clean up the old configuration on the router they are using. The easiest way is probably to do a "write erase" and start afresh. The alternative is to remove any interface ip addressing, OSPF and all BGP configuration. The latter does need care – remember that all unused configuration must at all times be removed from the router.

3. **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.

AS65534	**220.10.0.0/19**	**AS110**	**221.19.0.0/19**
AS65533	**220.19.0.0/19**	**AS111**	**221.35.0.0/19**
AS109	**220.73.0.0/19**		

When constructing an addressing plan, don't forget to use a small block for loopback interfaces and another small block for point to point links. Also, agree between yourselves and your neighbouring ASes which addresses will be used for the point to point links between you. Remember, it is Internet convention that addresses from the upstream ISP's address block are used for point to point link addresses to their customers.

Suggestion:

220.10.0.0/19 network block

220.10.0.0/24	220.10.1.0		220.10.31.255

Loopbacks   Infrastructure addresses   Customer assigned addresses

4. **Routing Protocols.** OSPF (area 0 only) and iBGP should now be configured between the routers in each AS. 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.*

Scenario One – Basic Configuration, no redundancy

The first scenario is not found in common practice but serves as a good introduction to the concepts of BGP multihoming as found on the Internet. And it will also serve as a reminder of some of the configuration concepts covered Modules 3 and 4 in the Basic ISP Workshop.

There are two links between the upstream ISP and their customer. Both links require to be used equally for traffic. The best way of achieving this is to announce half the address space on one link, and the other half of the address space on the other link. (The alternative is to simply announce the whole address block on each link. However, this will probably not achieve any useful degree of loadsharing, and the BGP path selection will simply fall back to using the originating router identifier.)

5. **Enable eBGP between AS109, AS110 and AS111.** AS109 has a connection to AS110, and AS110 has a connection to AS111. The eBGP sessions between the relevant routers in those ASes should now be configured. All router teams in these ASes should ensure that they are able to see all the prefixes of AS109, AS110 and AS111 (basically there will be entries for 220.73.0.0/19, 221.19.0.0/19 and 221.35.0.0/19 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!

6. **Prepare to enable eBGP between AS109 and AS65534.** Both AS65534 and AS109 are running iBGP within their ASes. To announce AS65534's prefix to AS109 we will take the /19 address block and divide it into two. The aim is to achieve relative even utilisation of the links between AS65534 and AS109, and common practice is to subdivide the address space. AS109 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.

   Note: some customers request/demand a full Internet routing table to be announced to their network – if after education they still insist, give it to them, but creative ISPs often charge for this service. Plus the customer needs to be aware they will need a router with at least 64Mbytes of useable memory given today's size of the Internet routing table.

7. **Create AS65534 prefix lists.** First, create the prefix lists on the routers in AS65534. For example, Router1 will announce the first subblock, Router3 will announce the second subblock. Both will accept the default route. Example for Router1:

   ```
   ip prefix-list subblock1 permit 220.10.0.0/20
   ip prefix-list default permit 0.0.0.0/0
   ```

8. **Create AS109 prefix-lists.** The routers in AS109 are the customer aggregation routers and should only accept those prefixes which the customer is entitled to announce. So prefix lists need to be installed on Routers 2 and 4 to do this. The example given is for Router 4. Notice how a range of addresses has been listed in the permit statement – this is so that the customer can make changes without the upstream having to reconfigure their filters:

```
ip prefix-list Customer permit 220.10.0.0/19 le 20
ip prefix-list default permit 0.0.0.0/0
```

9. **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 subblock2 permit 220.10.16.0/20
ip prefix-list default permit 0.0.0.0/0
!
router bgp 65534
 neighbor <router4> remote-as 109
 neighbor <router4> description Peering with Router 4 in AS109
 neighbor <router4> soft-reconfiguration in
 neighbor <router4> prefix-list subblock2 out
 neighbor <router4> prefix-list default in
!
```

10. **Configure eBGP in AS109.** AS109 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 2:

```
router bgp 109
 neighbor <router1> remote-as 65534
 neighbor <router1> description Dualhomed Customer
 neighbor <router1> default-originate
 neighbor <router1> soft-reconfiguration in
 neighbor <router1> prefix-list Customer in
 neighbor <router1> prefix-list default out
!
```

11. **Strip the private AS from external announcements from AS109.** Without further configuration changes in AS109, the private AS65534 will be announced by AS109 routers to other ASes. To stop private ASes from being announced, AS109 will need to configure the BGP command *remove-private-AS*. This needs to be done on all border routes in AS109 – in this case it is Router6.

```
ip prefix-list mynets permit 220.10.0.0/19 le 20
ip prefix-list mynets permit 220.73.0.0/19
!
```

```
router bgp 109
 neighbor <router8> remote-as 110
 neighbor <router8> description Peering with AS110
 neighbor <router8> soft-reconfiguration in
 neighbor <router8> remove-private-AS
 neighbor <router8> prefix-list mynets out
!
```

12. **AS111 and AS65533.** The same types of configuration concepts are also required on AS111 and AS65533. AS65533 is a multihomed customer of AS111. 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.

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. Notice that you should **not** see any private ASes in the BGP table of AS110.*

Scenario Two – Primary link and backup link

The second scenario is more commonly employed, especially where the customer has a large circuit to their upstream, and an inexpensive circuit they use almost exclusively for backup purposes.

In this case, the whole address block is announced out of both links. However, the announcement going out the backup link is "weighted" using MEDs 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 MEDs and local-preference? If in doubt, review the BGP presentation material.)

13. **Cleanup the private ASes.** Remove the configuration which split the address blocks in the two private ASes into two pieces and inserted them into the BGP table.

14. **Configure the main link.** Configure the main link between the private AS and the ISP. For AS65534, the link between Router1 and Router2 in AS109 is the main link – the link between Router3 and Router5 is the backup. For AS65533, the main link is between Router14 and Router13 in AS111. Example configuration for Router1:

```
ip prefix-list myblock permit 220.10.0.0/19
ip prefix-list default permit 0.0.0.0/0
!
router bgp 65534
 network 220.10.0.0 mask 255.255.224.0
 neighbor <router2> remote-as 109
 neighbor <router2> description Link to Router2 in AS109
```

```
  neighbor <router2> soft-reconfiguration in
  neighbor <router2> prefix-list myblock out
  neighbor <router2> prefix-list default in
 !
 ip route 220.10.0.0 255.255.224.0 null0
```

15. **Configure the backup link.** Configure the backup link between the private AS and the ISP. Set the metric on outbound announcements to 20, and set local preference on inbound announcements to 80. Remember that lowest metric 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 221.19.0.0/19
ip prefix-list default permit 0.0.0.0/0
!
route-map outfilter permit 10
 match ip address prefix-list myblock
 set metric 20
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 65533
 network 221.19.0.0 mask 255.255.224.0
 neighbor <router11> remote-as 111
 neighbor <router11> description Link to Router11 in AS111
 neighbor <router11> soft-reconfiguration in
 neighbor <router11> prefix-list myblock out
 neighbor <router11> prefix-list default in
 neighbor <router11> route-map outfilter out
 neighbor <router11> route-map infilter in
!
ip route 221.19.0.0 255.255.224.0 null0
```

16. **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 #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*

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

8

other router teams to resolve those. Notice that you still should __not__ see any private ASes in the BGP table of AS110.

Scenario Three – Loadsharing

The third scenario is the most commonly deployed. Most dualhomed 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.

In this case, 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.

17. **Clean up the private ASes.** 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.

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

```
ip prefix-list subblock1 permit 220.10.0.0/19
ip prefix-list subblock1 permit 220.10.0.0/20
!
ip prefix-list default permit 0.0.0.0/0
!
router bgp 65534
 network 220.10.0.0 mask 255.255.224.0
 network 220.10.0.0 mask 255.255.240.0
!
ip route 220.10.0.0 255.255.224.0 null0
ip route 220.10.0.0 255.255.240.0 null0
```

19. **Configure BGP in the private ASes.** Configure BGP on the border routers in the private ASes so that the prefix and one sub prefix is announced to the direct peer. For example, Router1 could announce *subblock1* as above, whereas Router3 could announce an equivalent *subblock2*. For example:

```
router bgp 65533
 neighbor <router11> remote-as 111
 neighbor <router11> description Link to Router11 in AS111
 neighbor <router11> soft-reconfiguration in
 neighbor <router11> prefix-list subblock1 out
```

```
    neighbor <router11> prefix-list default in
    !
```

20. **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. Notice that you still should __not__ see any private ASes in the BGP table of AS110.

Scenario Four – Scaling to support multiple dualhomed customers

The final scenario shows how to scale the third scenario described above. ISPs will offer multiple location connections as a service, so it is important to consider how to scale the configuration of the ISP's aggregation routers.

The customer configuration is unchanged from the previous step – both the customer address block and its subprefixes are announced to the upstream. However, the customers can all use the same private ASN – the ASN information is not transited by the ISP, the customer simply point default at the upstream, so BGP loop detection is not an issue. This device is used to great effect by many ISPs for their multihoming customers – not only does it ensure that the ISP doesn't have multiple private ASes in their backbone, it also ensures that their configuration tools have less complexity to deal with. Simplicity is always the design goal in successful ISP operations.

Note: RFC2270 describes this type of multihoming in more detail.

21. **Reconfigure the network.** Routers in AS111 and AS65533 should be reconfigured to become customers of AS109. Please refer to Figure 3 for connection details. As previously, each router team will need to set up OSPF and iBGP within their own AS. So, for example, in the Green (middle) network, Routers 9, 11 and 13 will need to set up OSPF and iBGP within their own network.

22. **Configure the address blocks and subblocks within each private AS.** The address blocks to use are as follows:

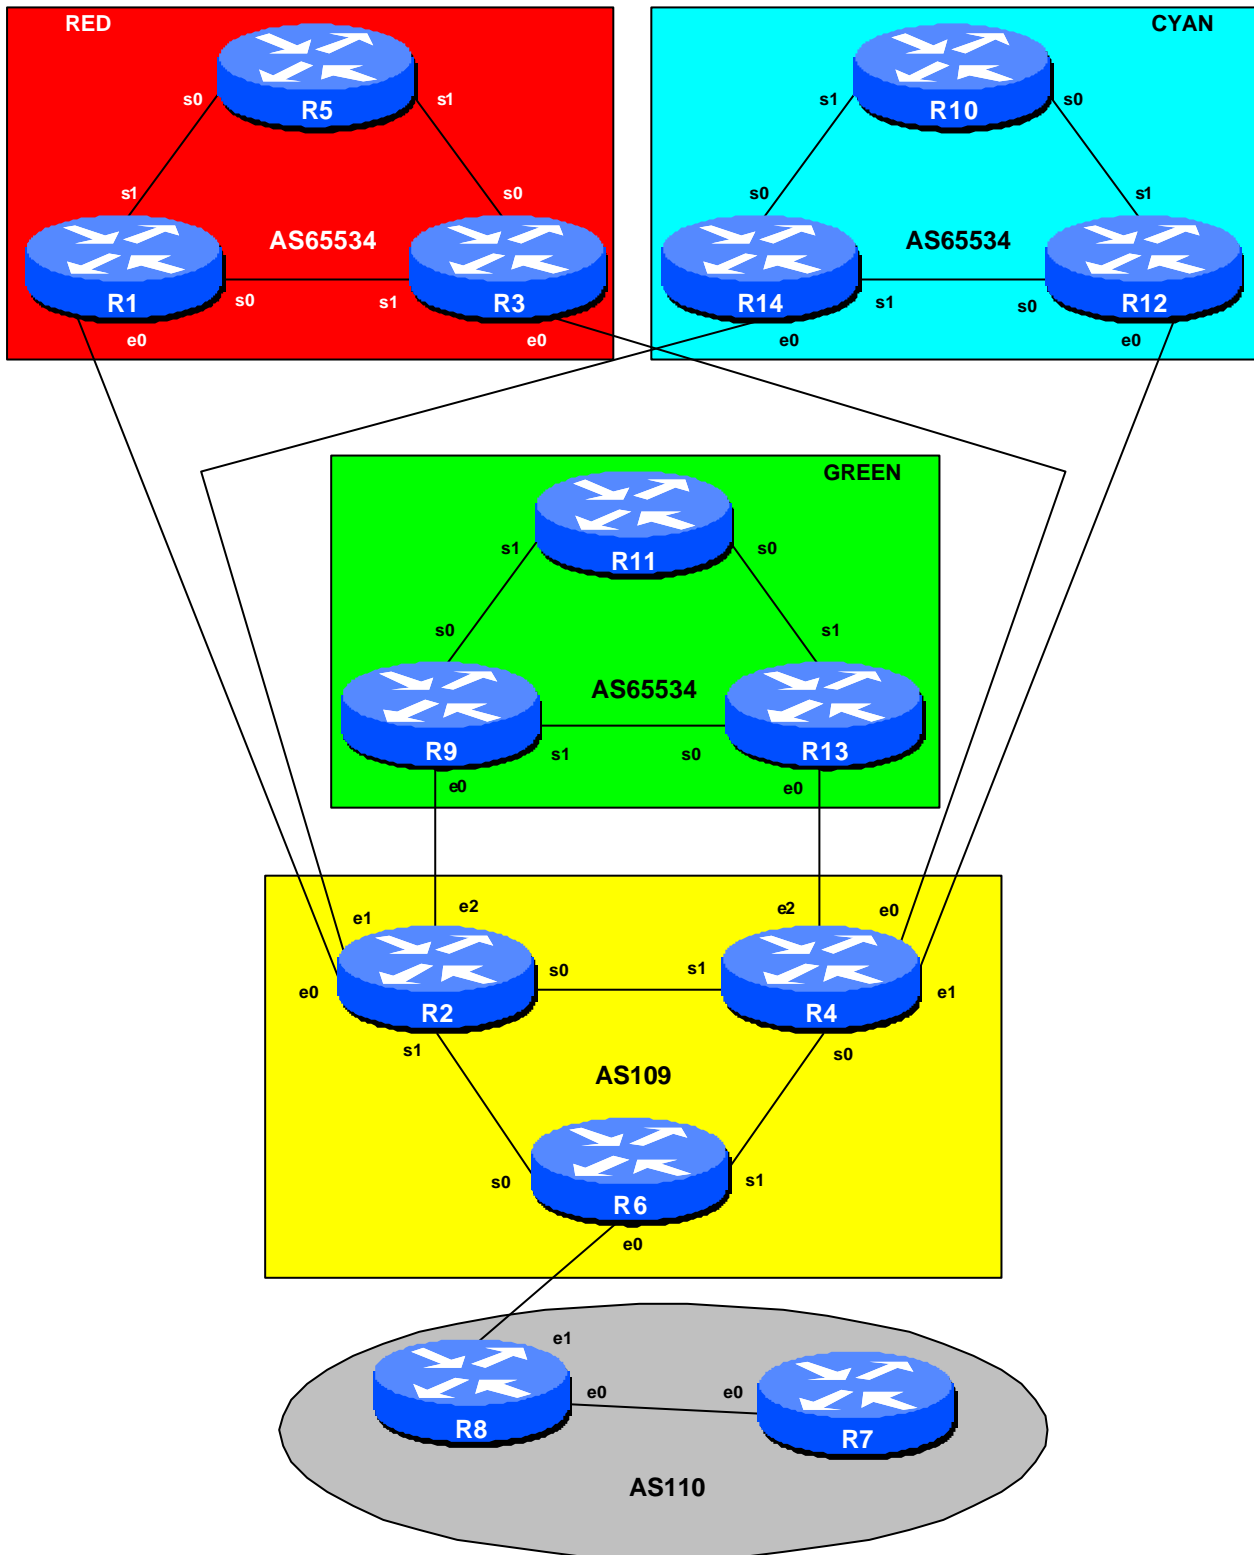|			
AS110	**221.19.0.0/19**	**AS65534 (R9,11,13)**	**221.35.0.0/19**
AS109	**220.73.0.0/19**	**AS65534 (R10,12,14)**	**220.19.0.0/19**
AS65534 (R1,3,5)	**220.10.0.0/19**		

CISCO SYSTEMS

Sunday, January 12, 2003

Figure 3 – Multiple Dualhomed Customers

CISCO SYSTEMS

Figure 4 – RFC2270 Lab Physical Layout

23. Configure eBGP between each AS65534 customer and AS109. Following the configuration hints in the previous section, each router team in AS65534 should configure their border routers to peer eBGP with AS109. Hint – the configuration should look something like:

```
ip prefix-list subblock1 permit y.y.0.0/19
ip prefix-list subblock1 permit y.y.0.0/20
!
ip prefix-list default permit 0.0.0.0/0
!
router bgp 65534
 network y.y.0.0 mask 255.255.224.0
 network y.y.0.0 mask 255.255.240.0
 neighbor x.x.x.x remote-as 111
 neighbor x.x.x.x description Link to RouterX in Aszzz
 neighbor x.x.x.x soft-reconfiguration in
 neighbor x.x.x.x prefix-list subblock1 out
 neighbor x.x.x.x prefix-list default in
!
ip route y.y.0.0 255.255.224.0 null 0
ip route y.y.0.0 255.255.240.0 null 0
```

24. Configure eBGP on AS109 border routers. Scalable eBGP configuration on Routers 2 and 4 is required. If AS109 has multiple BGP customers, it ensures that the growth of the AS109 network is not hindered by having to handcraft a configuration for every new customer.

The first first step is to use peer-groups for this. All the customers have the same outbound configuration, basically announce a default route. Remember that inbound policy can still be modified per peergroup neighbour – peergroups must have uniform **outbound** policy.

```
router bgp 109
 neighbor bgp-customers peer-group
 neighbor bgp-customers remote-as 65534
 neighbor bgp-customers default-originate
 neighbor bgp-customers soft-reconfiguration in
 neighbor bgp-customers prefix-list default out
!
```

After creating the peer-group, it can be applied to every BGP customer connecting to the router. Don't forget to create a prefix-list to filter the customer's inbound announcements. This is still required on a per customer basis.

```
ip prefix-list default permit 0.0.0.0/0
ip prefix-list RedCustomer permit 220.10.0.0/19 le 20
```

```
  ip prefix-list GreenCustomer permit 221.35.0.0/19 le 20
  ip prefix-list CyanCustomer permit 220.19.0.0/19 le 20
  !
  router bgp 109
   neighbor x.x.x.x peer-group bgp-customers
   neighbor x.x.x.x description Red AS customer
   neighbor x.x.x.x prefix-list RedCustomer in
   neighbor x.x.x.x peer-group bgp-customers
   neighbor x.x.x.x description Green AS customer
   neighbor x.x.x.x prefix-list GreenCustomer in
   neighbor x.x.x.x peer-group bgp-customers
   neighbor x.x.x.x description Cyan AS customer
   neighbor x.x.x.x prefix-list CyanCustomer in
  !
```

25. **Configuring AS109 border router to AS110.** The configuration of the AS109 border router connecting to AS110 (Router6) should be little changed from previous examples. It still requires the configuration to remove the private AS. And notice that it should only be allowing the customer blocks through, not the subprefixes of the customer blocks. As a reminder, the configuration of Router6 should look something like:

```
  ip prefix-list mynets permit 220.10.0.0/19
  ip prefix-list mynets permit 220.19.0.0/19
  ip prefix-list mynets permit 220.73.0.0/19
  ip prefix-list mynets permit 221.35.0.0/19
  !
  router bgp 109
   neighbor x.x.x.x remote-as 110
   neighbor x.x.x.x description Peering with AS110
   neighbor x.x.x.x soft-reconfiguration in
   neighbor x.x.x.x remove-private-AS
   neighbor x.x.x.x prefix-list mynets out
```

If the prefix-list is omitted, AS109 will leak the sub-prefixes of its multihomed customers to AS110. As there is no need to leak these sub-prefixes, this is frowned upon as bad practise on the Internet today.

26. **Announcing Prefixes – Important Note.** In the above step, the outbound prefix-list "mynets" listed all the prefixes which needed to be announced to the Internet. This case is correct **ONLY** when the customer address blocks do not come from the upstream ISP – in other words, the customer has so called provider independent (PI) space. This situation is becoming more rare on the Internet, and more often customers will have provider aggregatable (PA) space – in other words, the customer's address space comes from the upstream ISP address block. In that situation, the upstream ISP **MUST NOT** announce the sub-prefixes – he should only announce his own block.

Cisco Systems

For example, consider the situation where an ISP has the 221.19.0.0/19 address block. All his multihomed customers come out of this address block. His router configuration should be:

```
ip prefix-list myblock permit 220.10.0.0/19
!
router bgp 109
 network 220.10.0.0 mask 255.255.224.0
 neighbor x.x.x.x remote-as 110
 neighbor x.x.x.x description Peering with AS110
 neighbor x.x.x.x soft-reconfiguration in
 neighbor x.x.x.x remove-private-AS
 neighbor x.x.x.x prefix-list myblock out
```

Notice how the "remove-private-AS" directive has been retained – it is an extra precaution in case of accidents with the prefix-list.

27. **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.

28. **Summary.** This module has covered the major situations where a customer requires to multihomed onto the service provider backbone. It has demonstrated how to implement this multihoming using prefix-lists, MEDs and local-preference where appropriate. It has also demonstrated the *remove-private-AS* BGP command, which ensures that private ASes are stripped out of any announcements to the wider Internet.

CONFIGURATION NOTES

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

Cisco Systems