Module 14 – Transit

Objective: To investigate methods for providing transit services.

Prerequisites: Modules 12 and 13, and the Transit Presentation
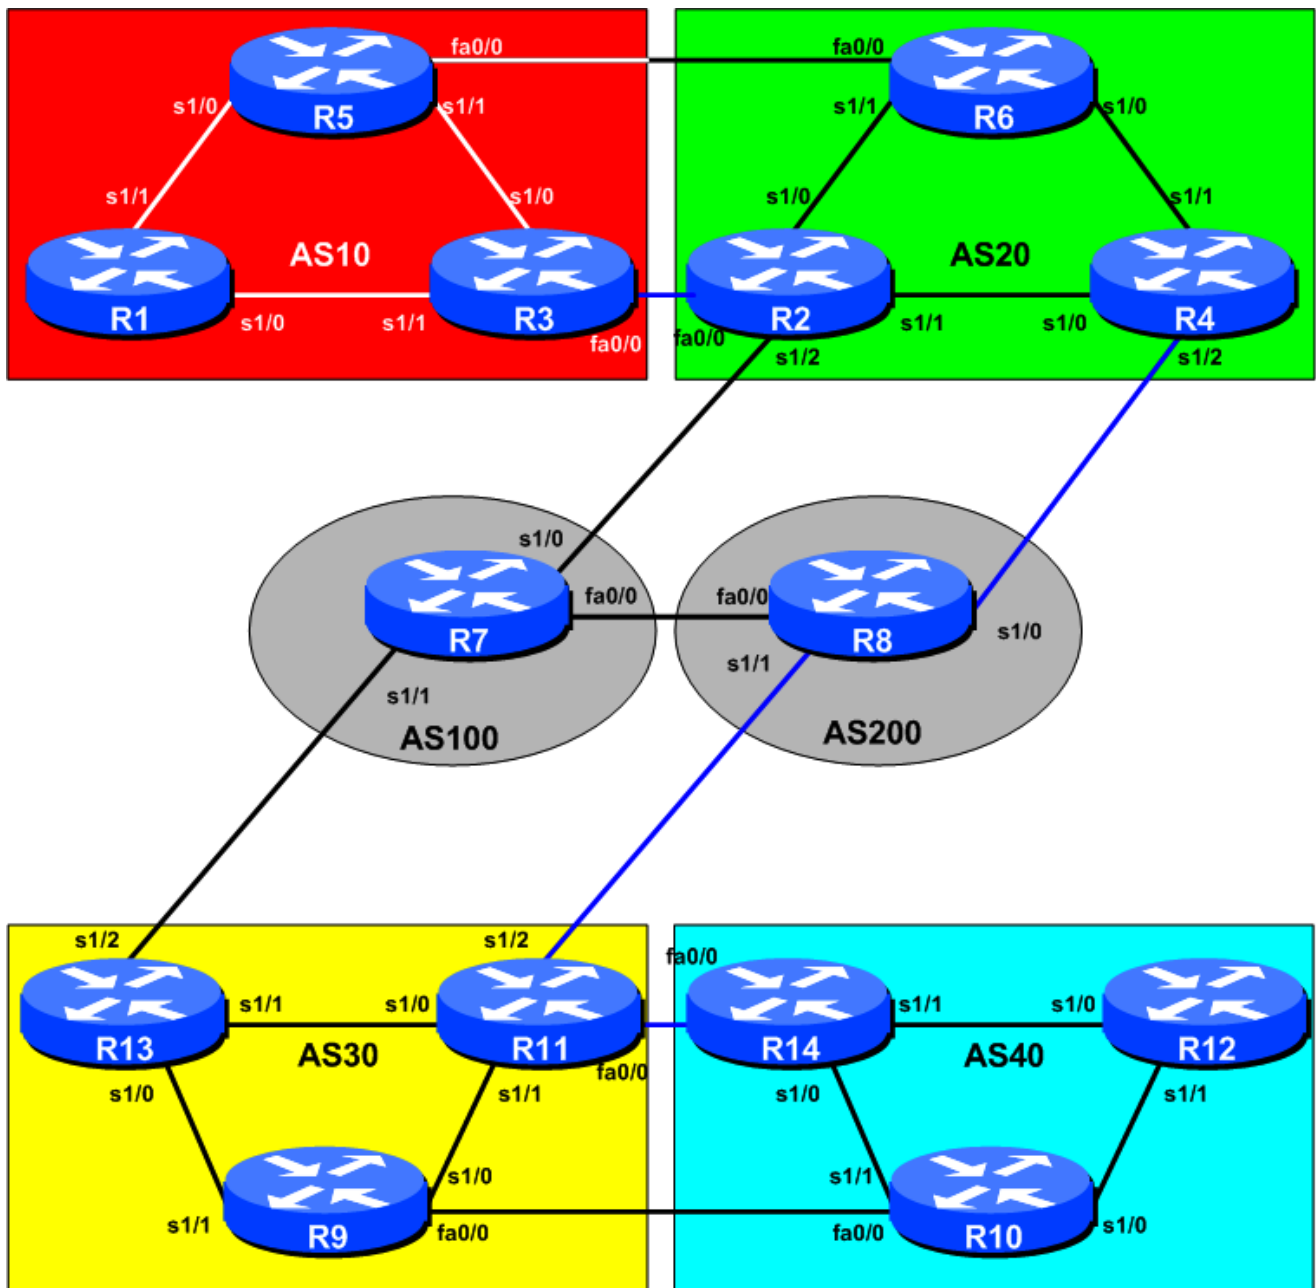
The following will be the common topology used.

Figure 1 – ISP Lab Transit Configuration

Lab Notes

The purpose of this module is to demonstrate how to configure backbones to provide transit services. Transit in its basic form means carrying packets with source and destination addresses which are not in the local AS. Not being in the local AS implies that these packets have originated in another AS, are being carried across the local AS, and transmitted to another connected AS.

The lab covers some of the configuration possibilities. AS10 and AS40 are edge ASNs and use the other ASes in the network to get transit between each other. AS20 and AS30 also get transit from AS100 and AS200 to reach each other – and provide transit to AS10 and AS40 respectively.

Lab Exercise

1. **Basic Configuration.** Each router team should configure their router to fit into the network topology depicted in Figure 1. Check all connections.

2. **Addressing Plan.** These address ranges, which we have retained from Module 13, should be used throughout this module – it means we don't have to renumber or reconfigure the OSPF or iBGP. (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.)

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.** Check that OSPF/ISIS is still working as per the previous module – also check that iBGP is still functioning as expected. Any interfaces which should not be running an IGP *MUST* be marked as passive in the configuration.

<u>Checkpoint #1:</u> 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 – Simple transit example

The first example is that of very simple transit. The example only uses prefix lists to configure filtering, and only provides transit from one AS, across an intermediate AS, to a remote AS. The example is typically that of a backbone ISP (NSP) providing connectivity services for access or content providers.

4. **Configuring edge ASes.** The router teams in ASes 10 and 40 should now configure their routers to peer with their transit providers. As in previous modules, they will only receive the default route from their transit provider. Also, they will only announce their own network block to their transit provider. This is very important to remember – ASes on the edge of the "Internet" should always attempt to minimise the complexity of their configuration. An example for Router5 in AS10 might be:

    ```
    ip prefix-list myblock permit 10.10.0.0/20
    ```

```
    ip prefix-list default permit 0.0.0.0/0
    !
    router bgp 10
     network 10.10.0.0 mask 255.255.240.0
     neighbor <router6> remote-as 20
     neighbor <router6> description My Transit Provider to AS20
     neighbor <router6> prefix-list myblock out
     neighbor <router6> prefix-list default in
    !
    ip route 10.10.0.0 255.255.240.0 null0
```

This should now be a very familiar configuration as it has been used several times in previous Modules. It is the best way of configuring an edge AS to connect to any Internet service provider.

5. **Configuring border routers in AS20 and AS30.** Next the routers in AS2 and AS6 which are providing the connection to the edge ASes need to be configured. Specifically these are Routers 2, 6, 9 and 11. The configuration again is similar to what has been repeatedly used in previous modules. Basically you want to announce default route only to your customer, and only accept the prefix block which they are entitled to send to you. A sample configuration for Router2 might be:

```
    ip prefix-list Customer1 permit 10.10.0.0/20
    ip prefix-list default permit 0.0.0.0/0
    !
    router bgp 20
     neighbor <router3> remote-as 10
     neighbor <router3> description Link to BGP Customer 1
     neighbor <router3> default-originate
     neighbor <router3> prefix-list Customer1 in
     neighbor <router3> prefix-list default out
    !
```

The configurations for the other routers will be similar.

6. **Configure the transit of AS20 and AS30 across AS100 and AS200.** The next step is to configure the routers in AS20, AS30, AS100 and AS200 such that AS20 transits AS100 and AS200 to reach AS30 – and vice-versa, of course. Remember that AS20 should not announce AS10 to AS100/200 and AS30 should not announce AS40 to AS100/200. So Routers 2, 4, 11 and 13 should only announce their own ASes to AS100 and AS200. A sample configuration for Router2 might be:

```
    ip prefix-list myblock permit 10.20.0.0/20
    ip prefix-list default permit 0.0.0.0/0
    !
    router bgp 20
     network 10.20.0.0 mask 255.255.240.0
     neighbor <router7> remote-as 100
     neighbor <router7> description Link to Transit Provider AS100
     neighbor <router7> prefix-list myblock out
     neighbor <router7> prefix-list default in
    !
    ip route 10.20.0.0 255.255.240.0 null 0
```

Notice that only the locally originated prefix is announced by AS20 to AS100. Later on in this module we will change the configuration so that AS20 and AS30 provide transit to other ASes too.

7. **Configure the AS100 and AS200 routers.** Finally we need to configure the routers in AS100 and AS200. The configuration is very similar to what was seen in Step 5 – only announce the default to the customer, only accept what the customer is supposed to announce to you.

8. **Connectivity Test.** Check connectivity throughout the lab network. Each router team should be able to see across their transit ASes to the routers in the opposite customer AS. When you are satisfied that BGP is working correctly, try running traceroutes to ensure that the proper paths are being followed. Note that the connectivity will be such that not every AS can see the entire network.

__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 – More complex transit

This situation above was a simple example of how to configure transit simply using prefix-list filters. However, this doesn't really represent reality in the majority of Internet cases. This next example introduces transit beyond the backbone AS we saw above. But now we want each AS to provide transit to the other ASNs, rather than the limited connectivity we saw in the first scenario.

9. **Reconfigure the network.** All routers should remove any prefix-list filtering and eBGP peering which has been configured. It is important that this is done so that the differences between this and the previous scenario can be understood.

10. **Configuring edge ASes.** The router teams in ASes 10 and 40 should now configure their routers to peer with their transit providers. AS10 and AS40 again receive global transit so only need to receive the default route from their transit provider. But this time we will send them the "full BGP table", hence required to do different filtering. An example for Router3 in AS10 might be:

```
ip prefix-list myblock permit 10.10.0.0/20
!
router bgp 10
 network 10.10.0.0 mask 255.255.240.0
 neighbor <router2> remote-as 20
 neighbor <router2> description My Transit Provider to the World
 neighbor <router2> prefix-list myblock out
 neighbor <router2> prefix-list bogons in
!
ip route 10.10.0.0 255.255.240.0 null0
```

The bogons prefix list was discussed during the presentation. It is good practice to have this prefix list configured on all eBGP peerings, in addition to any other filtering required. Refer to the BGP presentation material for the configuration of the list.

11. **Configuring border routers in AS20 and AS30.** Next the routers in AS20 and AS30 which are providing the connection to the edge ASes need to be configured. Specifically these are Routers 2, 6, 9 and 11. The configuration is slightly modified from the previous section, with the routers providing a full BGP feed to their customers. A sample configuration for Router2 might be:

```
ip prefix-list Customer1 permit 10.10.0.0/20
!
router bgp 20
 neighbor <router3> remote-as 10
 neighbor <router3> description Link to BGP Customer 1
 neighbor <router3> prefix-list Customer1 in
 neighbor <router3> prefix-list bogons out
!
```

12. **Configure the transit of AS20 and AS30 across AS100 and AS200.** The next step is to configure the routers in AS20 and AS30 such that AS20 transits AS100/200 to reach AS30 – and vice-versa, of course. AS20 and AS30 announce their customer address space as well as their own address space to AS100 and AS200. A sample configuration for Router 4 might be:

```
ip as-path access-list 5 permit ^$
ip as-path access-list 5 permit ^10$
!
router bgp 20
 network 10.20.0.0 mask 255.255.240.0
 neighbor <router8> remote-as 200
 neighbor <router8> description Link to Transit Provider AS200
 neighbor <router8> prefix-list bogons out
 neighbor <router8> filter-list 5 out
 neighbor <router8> prefix-list bogons in
!
ip route 10.20.0.0 255.255.240.0 null0
```

Notice the application of the bogons prefix list mentioned previously. Get into the habit of using it! Also there is no inbound filter-list. Some ISPs construct those too, using entries in the Internet Routing Registry (IRR)

13. **Configure the AS100 and AS200 routers.** Finally we need to configure the routers in AS100 and AS200. The configuration is very similar to the above – but use as-path filters to filter what your peers are sending you. An example for Router 8 is included:

```
ip as-path access-list 15 permit ^$
ip as-path access-list 15 permit ^20$
ip as-path access-list 15 permit ^20_10$
!
!
router bgp 200
 network 10.200.0.0 mask 255.255.240.0
 neighbor <router11> remote-as 30
 neighbor <router11> description Link to Customer AS30
 neighbor <router11> prefix-list bogons out
 neighbor <router11> prefix-list bogons in
 neighbor <router11> filter-list 15 out
!
ip route 10.200.0.0 255.255.240.0 null0
```

Notice the as-path filter this time – basically anything heard from AS20 is allowed to transit to AS30. Also notice the prefix-lists!

14. **Connectivity Test.** Check connectivity throughout the lab network. Each router team should be able to see across their transit ASes to the routers in the opposite customer AS. When you are

satisfied that BGP is working correctly, try running traceroutes to ensure that the proper paths are being followed.

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.*

15. **Summary.** This module has given examples of configurations used by Internet Transit Providers (NSPs). They have concentrated on using prefix-lists and as-path filters – more sophisticated configurations are possible by using communities. These examples 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.