

Module 18 – 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.

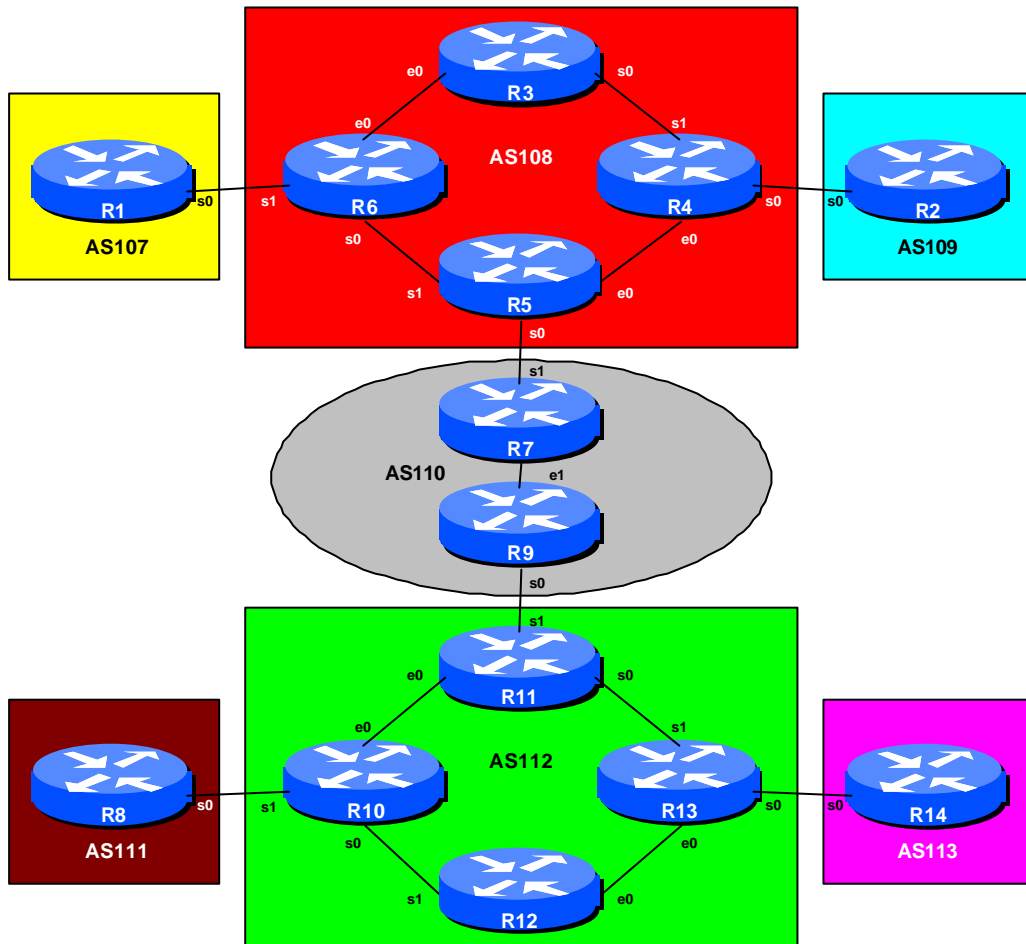


Figure 1 – ISP Lab Transit Configuration

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 example covered attempts to look at some of the configuration possibilities. The aim is as follows:

- AS107 and AS109 use AS108 for transit to reach each other
- AS108 and AS112 use AS110 for transit to reach each other
- AS111 and AS113 use AS112 for transit to reach each other

Note that the transit ASes should only provide transit as described above. They should not provide transit to any other of the connected ASes. For example, AS108 should not announce either AS107 or AS109 to AS110, etc.

Lab Exercise

1. **Basic Configuration.** Each router team should configure their router to fit into the network topology depicted in Figure 1 and physical layout depicted in Figure 2. 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.

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	AS113	222.11.0.0/19
AS110	221.19.0.0/19		

3. **Routing Protocols.** OSPF (area 0 only) and iBGP should now be configured for each of the routers in AS 108, AS110 and AS112. 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.




Figure 2 – Lab Physical Layout

Scenario One – Simple transit example

Monday, April 30, 2001

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 107, 109, 111 and 113 should now configure their routes 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 Router1 in AS107 might be:

```
ip prefix-list myblock permit 220.10.0.0/19
ip prefix-list default permit 0.0.0.0/0
!
router bgp 107
 network 220.10.0.0 mask 255.255.224.0
 neighbor <router6> remote-as 108
 neighbor <router6> description My Transit Provider to AS109
 neighbor <router6> prefix-list myblock out
 neighbor <router6> prefix-list default in
!
ip route 220.10.0.0 255.255.224.0 null0 250
```

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 AS108 and AS112.** Next the routers in AS108 and AS112 which are providing the connection to the edge ASes need to be configured. Specifically these are Routers 4, 6, 10 and 13. 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 Router6 might be:

```
ip prefix-list Customer107 permit 220.10.0.0/19
ip prefix-list default permit 0.0.0.0/0
!
router bgp 108
 neighbor <router1> remote-as 107
 neighbor <router1> description Link to BGP Customer 107
 neighbor <router1> default-originate
```

```
neighbor <router1> prefix-list Customer107 in
neighbor <router1> prefix-list default out
!
```

The configurations for the other routers will be similar.

- 6. Configure the transit of AS108 and AS112 across AS110.** The next step is to configure the routes in AS108, AS110 and AS112 such that AS108 transits AS110 to reach AS112 – and vice-versa, of course. Remember that AS108 should not announce either AS107 or AS109 to AS110, and that AS112 should not announce either AS111 or AS113 to AS110. So basically Routers 5 and 11 should only announce their own ASes to AS110. A sample configuration for Router 5 might be:

```
ip prefix-list myblock permit 220.19.0.0/19
ip prefix-list default permit 0.0.0.0/0
!
router bgp 108
network 220.19.0.0 mask 255.255.224.0
neighbor <router7> remote-as 110
neighbor <router7> description Link to Transit Provider AS110
neighbor <router7> prefix-list myblock out
neighbor <router7> prefix-list default in
!
ip route 220.19.0.0 255.255.224.0 null 0 250
```

Notice that only the locally originated prefix is announced by AS108 to AS110. Later on in this module we will change the configuration so that AS108 and AS112 provide transit to other ASes too.

- 7. Configure the AS110 routers.** Finally we need to configure the routers in AS110. 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.

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 – 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. In this example AS108, AS110 and AS112 are transit providers as previously. However,

- AS108 should provide local transit to AS107 and global transit to AS109
- AS112 should provide local transit to AS111 and global transit to AS113

So, this means that:

- AS107 should only be able to see AS108 and AS109
- AS111 should only be able to see AS112 and AS113
- AS109 should see all ASes apart from AS111
- AS113 should see all ASes apart from AS107

9. Reconfigure the network. All routers should remove any prefix-list filtering and eBGP peering which has been configured. That should be the whole classroom apart from Routers 3 and 12. 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 107, 109, 111 and 113 should now configure their routes to peer with their transit providers. AS109 and AS113 receive global transit so only need to receive the default route from their transit provider. AS107 and AS111 only receive prefixes local to AS108 and AS112 respectively. (Note that AS108 and AS112 could announce the default route, but a misconfiguration on AS108 and AS112's peering with AS110 might result in these two AS's receiving more connectivity than they expected/paid for.) As before, the four ASes will only announce their own network block to their transit provider. An example for Router1 in AS107 might be:

```
ip prefix-list myblock permit 220.10.0.0/19
!
router bgp 107
 network 220.10.0.0 mask 255.255.224.0
 neighbor <router6> remote-as 108
 neighbor <router6> description My Transit Provider to AS109
 neighbor <router6> prefix-list myblock out
 neighbor <router6> prefix-list rfc1918-dsua in
!
ip route 220.10.0.0 255.255.224.0 null0 250
```

The rfc1918-dsua 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 AS108 and AS112.** Next the routers in AS108 and AS112 which are providing the connection to the edge ASes need to be configured. Specifically these are Routers 4, 6, 10 and 13. The configuration is slightly modified from the previous section, noting especially that Routers 6 and 10 should only provide local transit services to their eBGP connected ASes. Routers 4 and 13 can announce the default route only to the customer, and only accept the prefix block which they are entitled to send to (as previously). A sample configuration for Router6 might be:

```
ip prefix-list Customer107 permit 220.10.0.0/19
!
ip as-path access-list 15 permit ^$
ip as-path access-list 15 permit ^109$
!
router bgp 108
 neighbor <router1> remote-as 109
 neighbor <router1> description Link to BGP Customer 107
 neighbor <router1> prefix-list Customer107 in
 neighbor <router1> prefix-list rfc1918-dsua out
 neighbor <router1> filter-list 15 out
!
```

- 12. Configure the transit of AS108 and AS112 across AS110.** The next step is to configure the routers in AS108, AS110 and AS112 such that AS108 transits AS110 to reach AS112 – and vice-versa, of course. This time AS108 and AS112 announce everything apart from their local AS customer (AS107 and AS111 respectively). This configuration is achieved using as-path filters (although prefix lists could be used). A sample configuration for Router 5 might be:

```
ip as-path access-list 15 permit ^$
ip as-path access-list 15 permit ^109$
!
router bgp 108
 network 220.19.0.0 mask 255.255.224.0
 neighbor <router7> remote-as 110
 neighbor <router7> description Link to Transit Provider AS110
 neighbor <router7> prefix-list rfc1918-dsua out
 neighbor <router7> prefix-list rfc1918-dsua in
 neighbor <router7> filter-list 15 out
!
ip route 220.19.0.0 255.255.224.0 null0 250
```

Monday, April 30, 2001

Notice the application of the rfc1918-dsua 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 AS110 routers. Finally we need to configure the routers in AS110. 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 9 is included:

```
ip as-path access-list 15 permit ^$
ip as-path access-list 15 permit ^108_$
!
!
router bgp 110
 network 221.19.0.0 mask 255.255.224.0
 neighbor <router11> remote-as 112
 neighbor <router11> description Link to Customer AS112
 neighbor <router11> prefix-list rfc1918-dsua out
 neighbor <router11> prefix-list rfc1918-dsua in
 neighbor <router11> filter-list 15 out
!
ip route 221.19.0.0 255.255.224.0 null0 250
```

Notice the as-path filter this time – basically anything heard from AS108 is allowed to transit to AS112. AS110 has no knowledge of AS108’s filtering policies, so doesn’t know that AS107 only has local connection within AS108 and it’s customer base. 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. As a reminder:

- AS107 should only be able to see AS108 and AS109
- AS111 should only be able to see AS112 and AS113
- AS109 should see all ASes apart from AS111
- AS113 should see all ASes apart from AS107

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.

Monday, April 30, 2001

CONFIGURATION NOTES

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