

## 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.

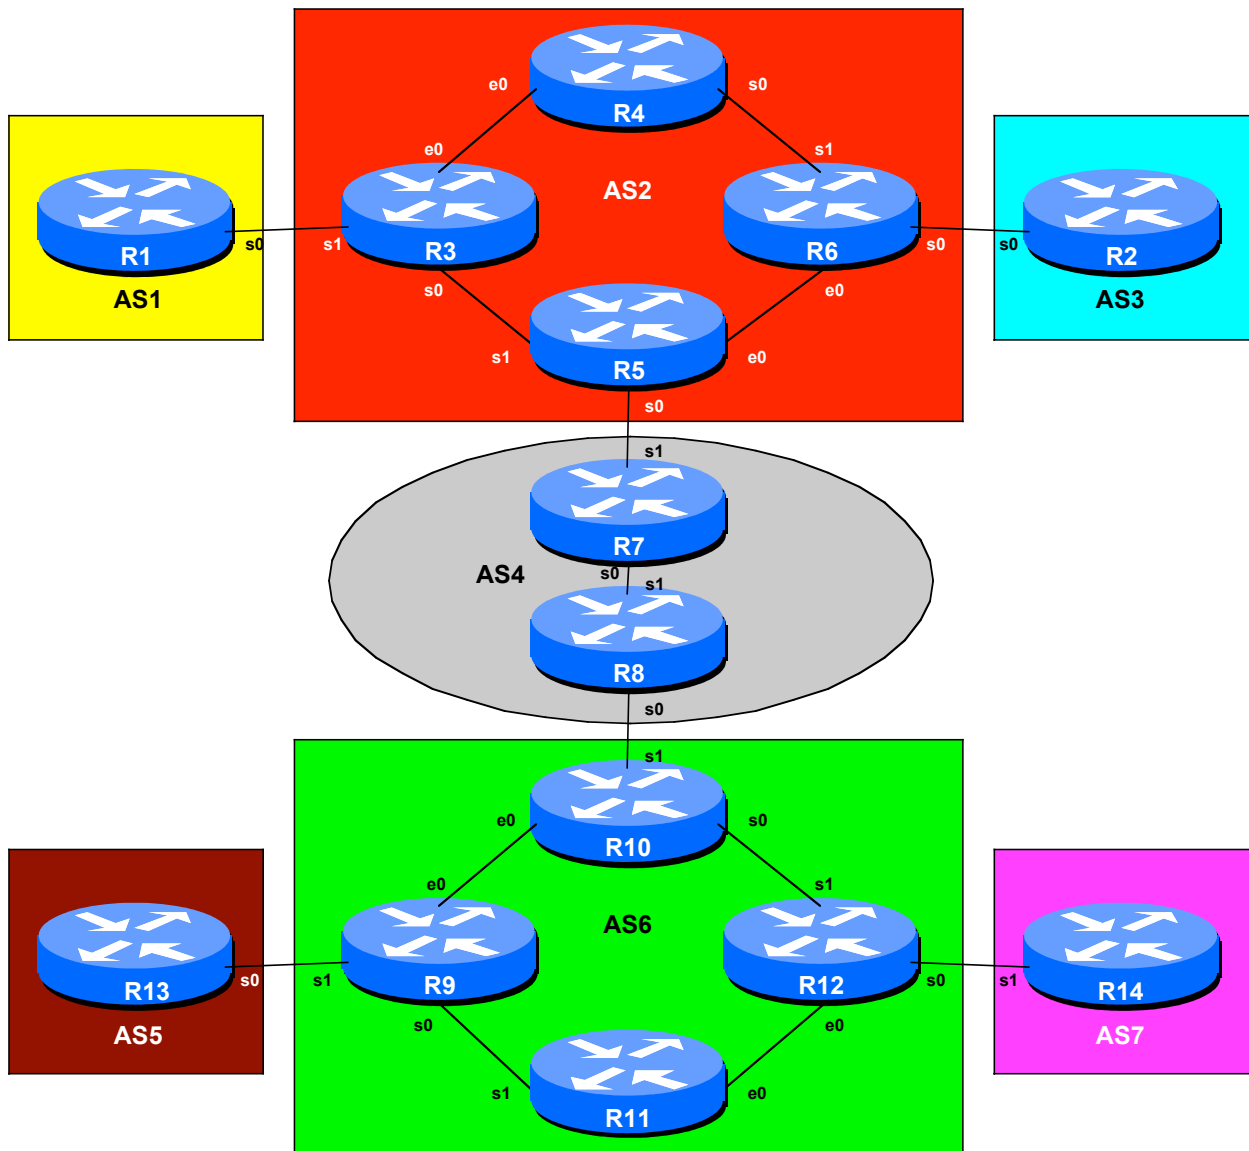


Figure 1 – ISP Lab Transit Configuration

## 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:

- AS1 and AS3 use AS2 for transit to reach each other
- AS2 and AS6 use AS4 for transit to reach each other
- AS5 and AS7 use AS6 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, AS2 should not announce either AS1 or AS3 to AS4, 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.

AS1	120.10.0.0/19	AS5	121.35.0.0/19
AS2	120.19.0.0/19	AS6	121.99.0.0/19
AS3	120.73.0.0/19	AS7	122.11.0.0/19
AS4	121.19.0.0/19		

3. **Routing Protocols.** OSPF (area 0 only) and iBGP should now be configured for each of the routers in AS 2, AS4 and AS6. 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.*

## 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.

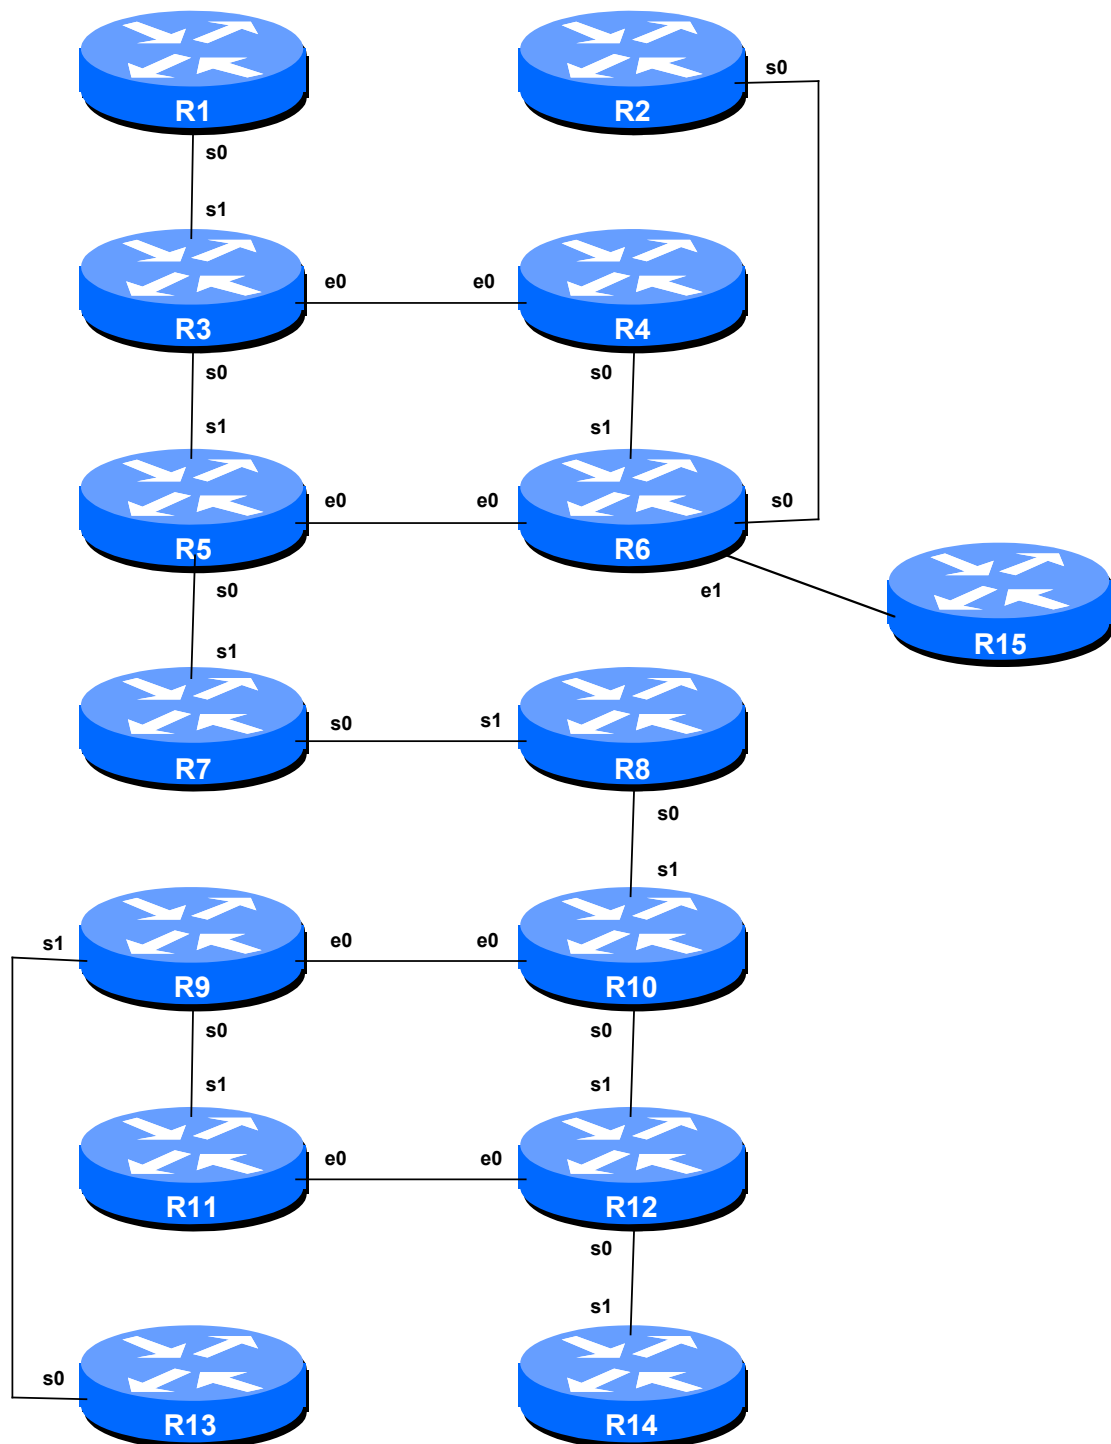


Figure 2 – Transit Lab Physical Layout

4. **Configuring edge ASes.** The router teams in ASes 1, 3, 5 and 7 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 AS1 might be:

```
ip prefix-list myblock permit 120.10.0.0/19
ip prefix-list default permit 0.0.0.0/0
!
router bgp 1
 network 120.10.0.0 mask 255.255.224.0
 neighbor <router3> remote-as 2
 neighbor <router3> description My Transit Provider to AS3
 neighbor <router3> prefix-list myblock out
 neighbor <router3> prefix-list default in
!
ip route 120.10.0.0 255.255.224.0 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 AS2 and AS6.** Next the routers in AS2 and AS6 which are providing the connection to the edge ASes need to be configured. Specifically these are Routers 3, 6, 9 and 12. 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 Router3 might be:

```
ip prefix-list Customer1 permit 120.10.0.0/19
ip prefix-list default permit 0.0.0.0/0
!
router bgp 2
 neighbor <router1> remote-as 1
 neighbor <router1> description Link to BGP Customer 1
 neighbor <router1> default-originate
 neighbor <router1> prefix-list Customer1 in
 neighbor <router1> prefix-list default out
!
```

The configurations for the other routers will be similar.

- 6. Configure the transit of AS2 and AS6 across AS4.** The next step is to configure the routes in AS2, AS4 and AS6 such that AS2 transits AS4 to reach AS6 – and vice-versa, of course. Remember that AS2 should not announce either AS1 or AS3 to AS4, and that AS6 should not announce either AS5 or AS7 to AS4. So Routers 5 and 10 should only announce their own ASes to AS4. A sample configuration for Router5 might be:

```
ip prefix-list myblock permit 120.19.0.0/19
ip prefix-list default permit 0.0.0.0/0
!
router bgp 2
 network 120.19.0.0 mask 255.255.224.0
 neighbor <router7> remote-as 4
 neighbor <router7> description Link to Transit Provider AS4
 neighbor <router7> prefix-list myblock out
 neighbor <router7> prefix-list default in
!
ip route 120.19.0.0 255.255.224.0 null 0
```

Notice that only the locally originated prefix is announced by AS2 to AS4. Later on in this module we will change the configuration so that AS2 and AS6 provide transit to other ASes too.

7. **Configure the AS4 routers.** Finally we need to configure the routers in AS4. 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.*

## 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. In this example AS2, AS4 and AS6 are transit providers as previously. However,

- AS2 should provide local transit to AS1 and global transit to AS3
- AS6 should provide local transit to AS5 and global transit to AS7

So, this means that:

- AS1 should only be able to see AS2 and AS3
- AS5 should only be able to see AS6 and AS7
- AS3 should see all ASes apart from AS5
- AS7 should see all ASes apart from AS1

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 4 and 11. 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 1, 3, 5 and 7 should now configure their routers to peer with their transit providers. AS3 and AS7 receive global transit so only need to receive the default route from their transit provider. AS1 and AS5 only receive prefixes local to AS2 and AS6 respectively. (Note that AS2 and AS6 could announce the default route, but a misconfiguration on AS2 and AS6's peering with AS4 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 AS1 might be:

```
ip prefix-list myblock permit 120.10.0.0/19
!
router bgp 1
```

```
network 120.10.0.0 mask 255.255.224.0
neighbor <router3> remote-as 2
neighbor <router3> description My Transit Provider to AS3
neighbor <router3> prefix-list myblock out
neighbor <router3> prefix-list rfc1918-sua in
!
ip route 120.10.0.0 255.255.224.0 null0
```

The rfc1918-sua 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 AS2 and AS6.** Next the routers in AS2 and AS6 which are providing the connection to the edge ASes need to be configured. Specifically these are Routers 3, 6, 9 and 12. The configuration is slightly modified from the previous section, noting especially that Routers 3 and 9 should only provide local transit services to their eBGP connected ASes. Routers 6 and 12 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 Router3 might be:

```
ip prefix-list Customer1 permit 120.10.0.0/19
!
ip as-path access-list 15 permit ^$
ip as-path access-list 15 permit ^3$
!
router bgp 2
neighbor <router1> remote-as 1
neighbor <router1> description Link to BGP Customer 1
neighbor <router1> prefix-list Customer1 in
neighbor <router1> prefix-list rfc1918-sua out
neighbor <router1> filter-list 15 out
!
```

**12. Configure the transit of AS2 and AS6 across AS4.** The next step is to configure the routers in AS2, AS4 and AS6 such that AS2 transits AS4 to reach AS6 – and vice-versa, of course. This time AS2 and AS6 announce everything apart from their local AS customer (AS1 and AS5 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 ^3$
!
router bgp 2
network 120.19.0.0 mask 255.255.224.0
neighbor <router7> remote-as 4
neighbor <router7> description Link to Transit Provider AS4
neighbor <router7> prefix-list rfc1918-sua out
neighbor <router7> prefix-list rfc1918-sua in
neighbor <router7> filter-list 15 out
!
ip route 120.19.0.0 255.255.224.0 null0
```

Notice the application of the rfc1918-sua 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 AS4 routers.** Finally we need to configure the routers in AS4. 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 ^2_$
!
!
router bgp 4
 network 121.19.0.0 mask 255.255.224.0
 neighbor <router10> remote-as 6
 neighbor <router10> description Link to Customer AS6
 neighbor <router10> prefix-list rfc1918-sua out
 neighbor <router10> prefix-list rfc1918-sua in
 neighbor <router10> filter-list 15 out
!
ip route 121.19.0.0 255.255.224.0 null0
```

Notice the as-path filter this time – basically anything heard from AS2 is allowed to transit to AS6. AS4 has no knowledge of AS2’s filtering policies, so doesn’t know that AS1 only has local connection within AS2 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:

- AS1 should only be able to see AS2 and AS3
- AS5 should only be able to see AS6 and AS7
- AS3 should see all ASes apart from AS5
- AS7 should see all ASes apart from AS1

***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.

## ***CONFIGURATION NOTES***

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