# 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.

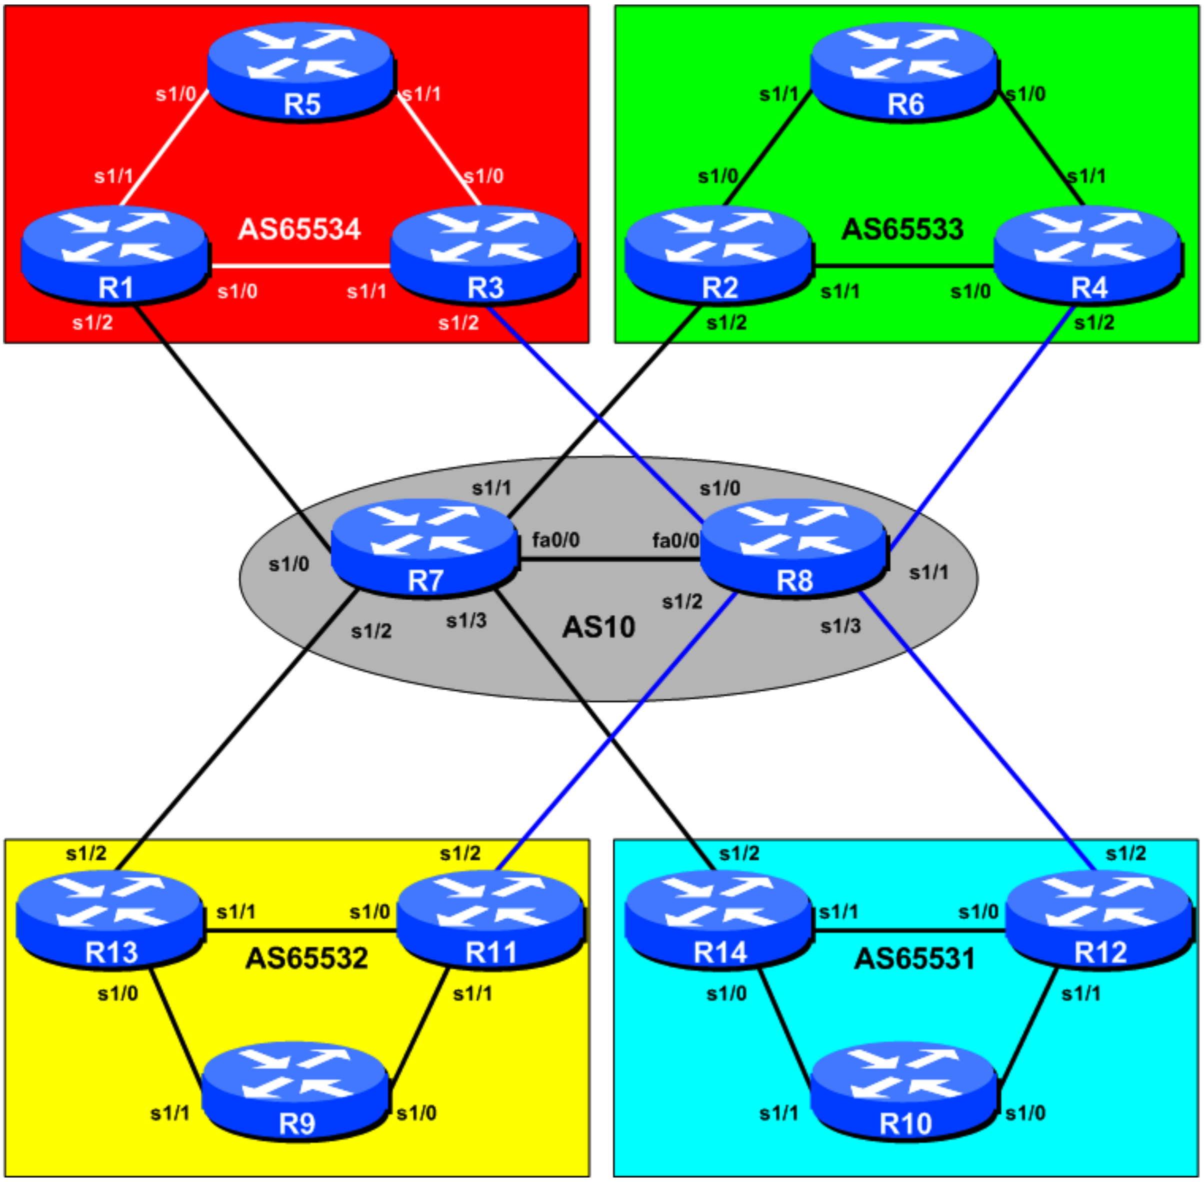


**Figure 1 – ISP Lab Configuration**

## 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 **never** 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 Router7 and Router8 have Ethernet interface names such as “fa0/0.10” etc. This is a .1q vlan configuration, and the lab instructors will help the teams operating Routers 7 and 8 with this.

   **IMPORTANT:** Each router team should ensure that their router has the basic configuration as covered in the first steps of Module 11.

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/ISIS 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** | **10.10.0.0/20** | **AS65531** | **10.40.0.0/20** |
| **AS65533** | **10.20.0.0/20** | **AS10** | **10.50.0.0/20** |
| **AS65532** | **10.30.0.0/20** | | |

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:

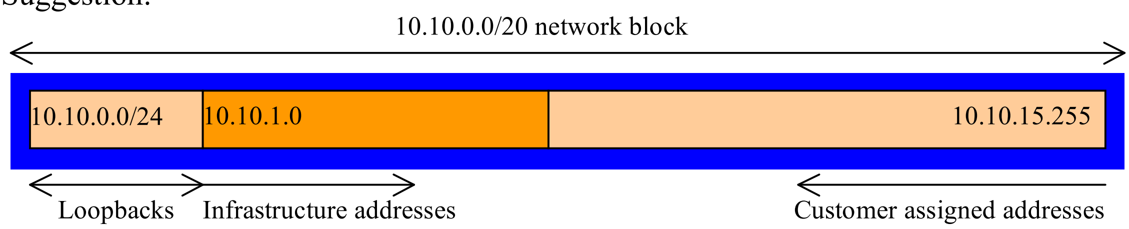


4. **Routing Protocols.** The chosen IGP (OSPF using area 0 only, or ISIS using L2 only) and iBGP should now be configured between the routers in each AS. Any interfaces which are not going to run OSPF/ISIS *MUST* be marked as passive in the configuration. And don't forget to use BGP peer groups for 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/ISIS prefixes in the routing table), please let the instructor know.*

## Scenario One – Primary link and backup link

This first scenario is more commonly employed where the customer has two circuits to their upstream: a large circuit which is used for all the inbound and outbound traffic, 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.)

5. **Prepare to enable eBGP between AS10 and AS65534.** Both AS65534 and AS10 should now be running iBGP within their ASes. To announce AS65534's prefix to AS1 we will take the /20 address block and announce it on both BGP peerings between the ASNs. AS10 will not announce any specific 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 128Mbytes of free memory given today's size of the Internet routing table.

6. **Prepare to enable eBGP between the other private ASes and AS10.** The process for enabling eBGP in this case is the same as for between AS10 and AS65534.

7. **Create AS65534 prefix lists.** First, create the prefix lists on the routers in AS65534. Both Router1 and Router3 will announce the /20. Both will accept the default route. Example for Router1:

```
ip prefix-list myblock permit 10.10.0.0/20
ip prefix-list default permit 0.0.0.0/0
```

8. **Create AS10 prefix-lists.** The routers in AS10 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 7 and 8 to do this. The example given is for Router 7:

```
ip prefix-list Customer permit 10.10.0.0/20
ip prefix-list default permit 0.0.0.0/0
```

9. **Configure the main link.** Configure the main link between the private AS and the ISP. For AS65534, the link between Router1 and Router7 in AS10 is the main link – the link between Router3 and Router8 is the backup (marked in blue). Example configuration for Router1:

```
router bgp 65534
 network 10.10.0.0 mask 255.255.240.0
 neighbor <router7> remote-as 10
 neighbor <router7> description Link to Router8 in AS10
 neighbor <router7> prefix-list myblock out
 neighbor <router7> prefix-list default in
!
ip route 10.10.0.0 255.255.240.0 null0
```

Configure the other ASNs' primary link – they are marked in black on the diagram in Figure 1.

10. **Configure the backup link.** Configure the backup links between the private ASNs and the ISP. Set the metric on outbound announcements to 20, and set local preference on inbound announcements to 90. 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 Router4:

```
ip prefix-list myblock permit 10.20.0.0/20
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 90
route-map infilter permit 20
!
router bgp 65533
 network 10.20.0.0 mask 255.255.240.0
 neighbor <router11> remote-as 10
 neighbor <router11> description Link to Router8 in AS10
 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 100.10.0.0 255.255.240.0 null0
```

11. **Configure eBGP with private ASNs in AS10.** AS10 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 7:

```
router bgp 10
 neighbor <router1> remote-as 65534
 neighbor <router1> description Dualhomed Customer
 neighbor <router1> default-originate
 neighbor <router1> prefix-list Customer in
 neighbor <router1> prefix-list default out
!
```

12. **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. For example:

```
Router8>trace 10.10.0.5

Type escape sequence to abort.
Tracing the route to 10.10.0.5

  1 fa0-1.router7  4 msec 4 msec 0 msec
  2 fa0-0.router1  0 msec 4 msec 4 msec
  3 ser0-0.router5 [AS 65534] 4 msec *  4 msec
```

With the cable between Router 1 and Router 7 disconnected, the trace becomes:

```
Router8>trace 10.10.0.5

Type escape sequence to abort.
Tracing the route to 10.10.0.5

  1 fa0-0.router3  [AS 1] 4 msec 5 msec 4 msec
  2 ser0-1.router5 [AS 1] 4 msec *  4 msec
Router8>
```

***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 – Loadsharing (Method One)*

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. The result of this is that traffic for the first /21 comes in one path, and traffic for the second /21 comes in the other path. If either path fails, the advertisement of the /21 address block (aggregate) ensures continued connectivity.

13. **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.

14. **Configure the address block and subprefixes in the private ASes.** Modify the router configuration so that the /20 address block and two /21 subprefixes are present in the BGP table. Also set up prefix lists to cater for these blocks. For example:

```
ip prefix-list path1 permit 10.10.0.0/20
ip prefix-list path1 permit 10.10.0.0/21
!
ip prefix-list path2 permit 10.10.0.0/20
ip prefix-list path2 permit 10.10.8.0/21
!
ip prefix-list default permit 0.0.0.0/0
!
router bgp 65534
 network 10.10.0.0 mask 255.255.240.0
 network 10.10.0.0 mask 255.255.248.0
!
ip route 10.10.0.0 255.255.240.0 null0
ip route 10.10.8.0 255.255.248.0 null0
```

15. **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 *path1* as above, whereas Router3 could announce an equivalent *path2*. For example, for AS65533:

```
router bgp 65533
 neighbor <router7> remote-as 10
 neighbor <router7> description Link to Router7 in AS10
 neighbor <router7> prefix-list path1 out
 neighbor <router7> prefix-list default in
!
```

16. **Check filters in AS10.** Ensure that Routers7 and 8 have appropriate filters to allow the /20 and /21 prefixes in from the private AS customers. Here is an example:

```
router bgp 10
 neighbor <router1> remote-as 65534
 neighbor <router1> description Link to Router1 in AS65534
 neighbor <router1> prefix-list Customer in
...
!
ip prefix-list Customer permit 10.10.0.0/20 le 21
!
```

17. **Connectivity targets.** So that connectivity via each /21 can be tested, Routers 5, 6, 9 and 10 should set up a second loopback interface with an IP address from their ASN's respective second /21 block. Note that the /32 address should be announced by the IGP so that other routers in the ASN know how to get to the destination. The following configuration snippet shows a possible configuration for Router 10 using OSPF:

```
interface loopback 1
 ip address 10.30.15.10 255.255.255.255
 ip ospf 65533 area 0                   ! ONLY for IOS 12.4 and later
!
router ospf 65533
 network 10.30.15.10 0.0.0.0 area 0    ! ONLY for IOS <12.4
!
```

The following configuration snippet shows a possible configuration for Router 10 using ISIS:

```
interface loopback 1
 ip address 10.30.15.10 255.255.255.255
!
router isis workshop
 passive-interface loopback1
!
```

18. **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 Routers 5, 6, 9 and 10. Check the path being followed for the first and second /21s for each ASN. For example:

```
Router8>trace 10.10.0.5

Type escape sequence to abort.
Tracing the route to 120.10.0.5

  1 fa0-0.router7  4 msec 4 msec 0 msec
  2 fa0-0.router1  4 msec 4 msec 4 msec
  3 ser0-0.router5 [AS 65534] 4 msec *  4 msec
Router8>trace 10.10.15.5

Type escape sequence to abort.
Tracing the route to 10.10.15.5

  1 fa0-0.router3  4 msec 5 msec 4 msec
  2 ser0-1.router5 [AS 65534] 4 msec *  4 msec
Router8>
```

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 #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.*

## STOP AND WAIT HERE

## *Scenario Three – Loadsharing (Method Two)*

The third scenario is a variation on the second scenario and provides another example.

As before the whole address block is announced out of both links. In fact, one of the key features of multihoming, and providing redundancy, is that the ISP's address blocks are always announced out of each external link. The key to loadbalancing is how those external announcements are made. In this example, one /21 is taken out of the /20 address block and announced on one link between the customer and the ISP as well as the /20. The other link sees just the standard announcement of the /20.

19. **Clean up the private ASes.** Remove the configuration which subdivided the address space for the previous example. Remember it is always very important to remove any configuration which isn't being used from the router.

20. **Configure the address block and subprefixes in the private ASes.** Modify the configuration of the routers in the private ASNs so that the /20 address block and one /21 subprefix are present in the BGP table. Also set up prefix lists to cater for these blocks. For example:

```
ip prefix-list aggregate permit 10.10.0.0/20
!
ip prefix-list subblock permit 10.10.0.0/20
ip prefix-list subblock permit 10.10.0.0/21
!
ip prefix-list default permit 0.0.0.0/0
!
router bgp 65534
 network 120.10.0.0 mask 255.255.240.0
 network 120.10.0.0 mask 255.255.248.0
!
ip route 120.10.0.0 255.255.240.0 null0
ip route 120.10.0.0 255.255.248.0 null0
```

21. **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 on one link, and just the aggregate is announced on the other link. The router teams in each private ASN should discuss amongst themselves who will leak the /21 prefix. For example, Router1 could announce *subblock* as above, whereas Router3 could announce just the aggregate. A configuration example for AS65533 might look like:

```
router bgp 65533
 neighbor <router11> remote-as 10
 neighbor <router11> description Link to Router8 in AS10
 neighbor <router11> prefix-list subblock out
 neighbor <router11> prefix-list default in
!
```

22. **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.

***<u>Checkpoint #4:</u>*** *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 Four – RFC2270*

The final scenario describes RFC2270 style multihoming and represents an efficiency saving over the configuration concepts described in the previous three steps.

23. **Routers in Private ASNs.** ISPs who deploy RFC2270 style multihoming use only **ONE** ASN for their multihoming customers. Each private AS should now replace the private ASN in their BGP sessions with AS64512 (as in Figure 2). Do not change your BGP peers, just change the ASN referred to. For example, all the routers in AS65534 will renumber their BGP to use AS64512. This means removing BGP totally, and then adding it back in (there is no way to easily migrate BGP ASNs in IOS). The best way of doing this is to copy the BGP configuration on to another terminal window, edit the 65534 to replace it with 64512, do `no router bgp 65534` on the router, and the copy the modified BGP configuration back onto the router. For example:

    **Before Modification:**

    ```
    router bgp 65534
     neighbor ibgp peer-group
     neighbor ibgp remote-as 65534
     neighbor ibgp update-source loopback0
     neighbor ibgp send-community
     neighbor ibgp next-hop-self
     neighbor ibgp password cisco
     neighbor 10.10.0.3 peer-group ibgp
     neighbor 10.10.0.3 description R3
     neighbor 10.10.0.5 peer-group ibgp
     neighbor 10.10.0.5 description R5
    !
    ```

    **After Modification:**

    ```
    router bgp 64512
     neighbor ibgp peer-group
     neighbor ibgp remote-as 64512
     neighbor ibgp update-source loopback0
     neighbor ibgp send-community
     neighbor ibgp next-hop-self
     neighbor ibgp password cisco
     neighbor 10.10.0.3 peer-group ibgp
     neighbor 10.10.0.3 description R3
     neighbor 10.10.0.5 peer-group ibgp
     neighbor 10.10.0.5 description R5
    !
    ```

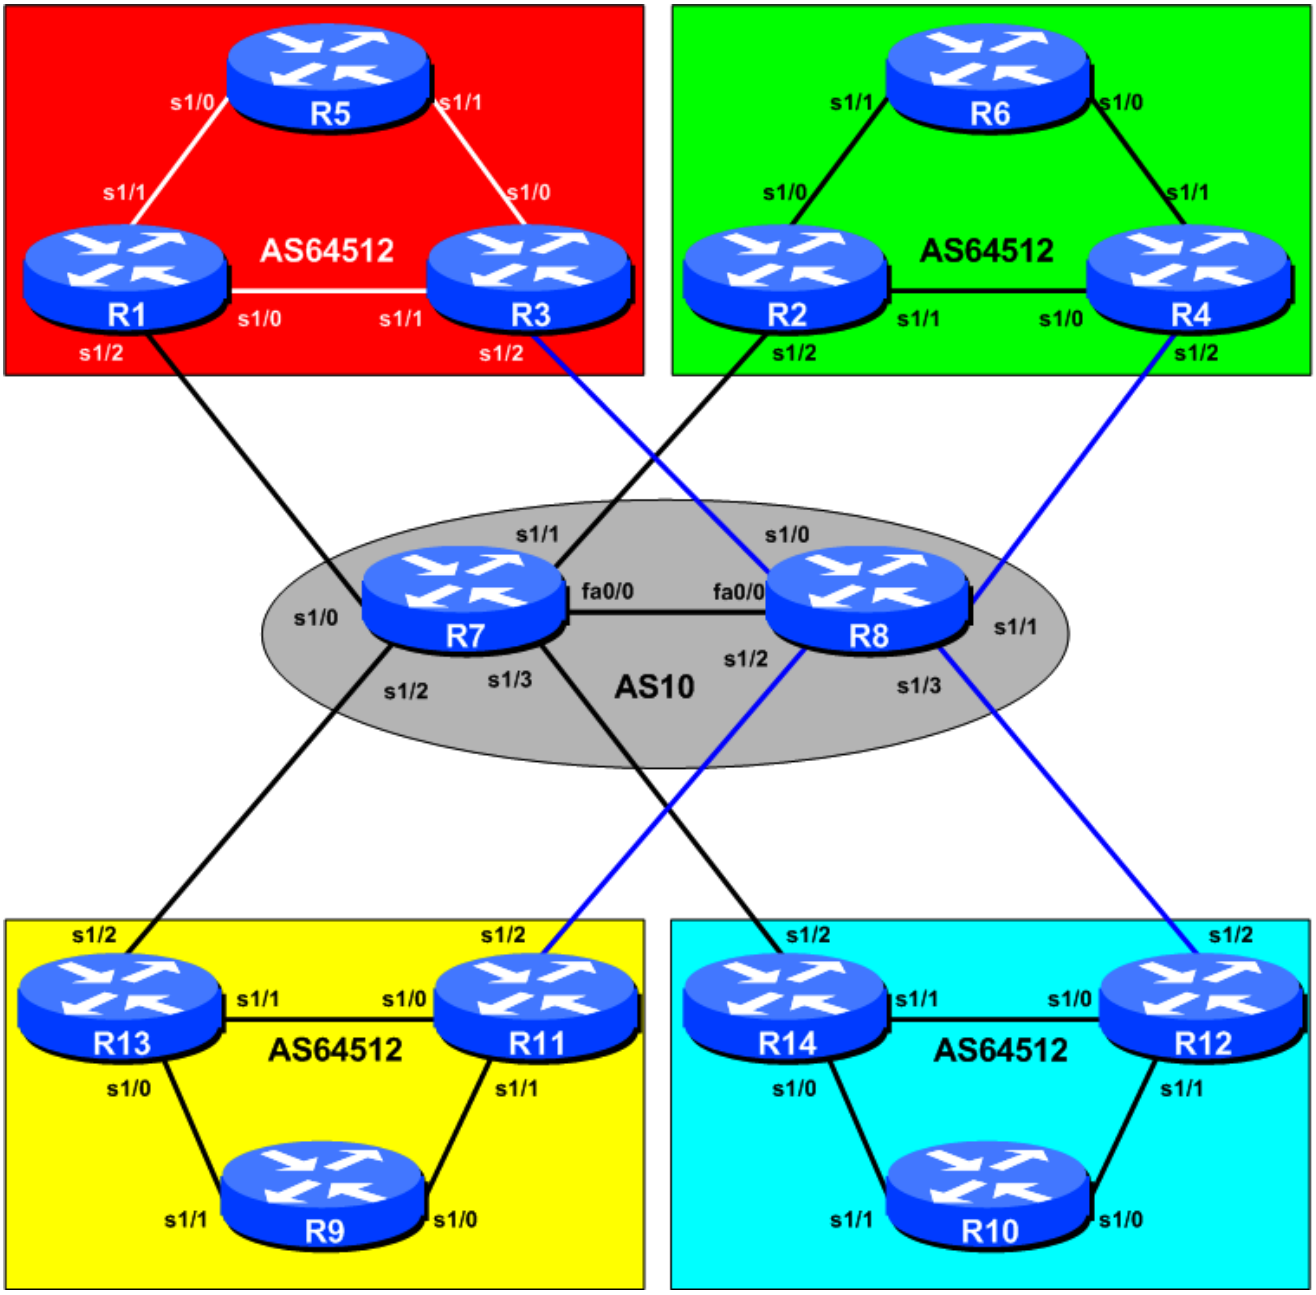


**Figure 2 - RFC2270 configuration**

**24. Routers in AS10.** The teams operating Router7 and Router8 should make their BGP configuration more efficient. Because each neighbour is now in the same ASN, each team can replace the individual BGP configurations for each neighbouring ASN in the previous steps with a generic peer-group based configuration. Create a peer-group for the multihoming private ASNs – for example:

```
router bgp 10
 neighbor rfc2270 peer-group
 neighbor rfc2270 remote-as 64512
 neighbor rfc2270 default-originate
 neighbor rfc2270 send-community
 neighbor rfc2270 prefix-list default out
 neighbor rfc2270 password cisco
```

And then apply the peer-group to each eBGP customer, for example:

```
router bgp 10
 neighbor <router1> peer-group rfc2270
 neighbor <router1> prefix-list r1 in
 neighbor <router2> peer-group rfc2270
 neighbor <router2> prefix-list r2 in
 neighbor <router13> peer-group rfc2270
 neighbor <router13> prefix-list r13 in
 neighbor <router14> peer-group rfc2270
 neighbor <router14> prefix-list r14 in
!
```

25. **Final check.** This step introduced a more efficient way of configuring multiple multihomed customers, as described in RFC2270. While there is no requirement to configure customers in this style, the transit ISP has a more efficient and therefore scalable configuration. Check that the routing and connectivity is exactly as it was at the end of the previous scenario.

***Checkpoint #5:*** *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.*

26. **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.