Module 15 – Multihoming to the Same ISP using RFC1998

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

Prerequisites: Module 12 and Advanced Communities Presentation

The following will be the common topology used.

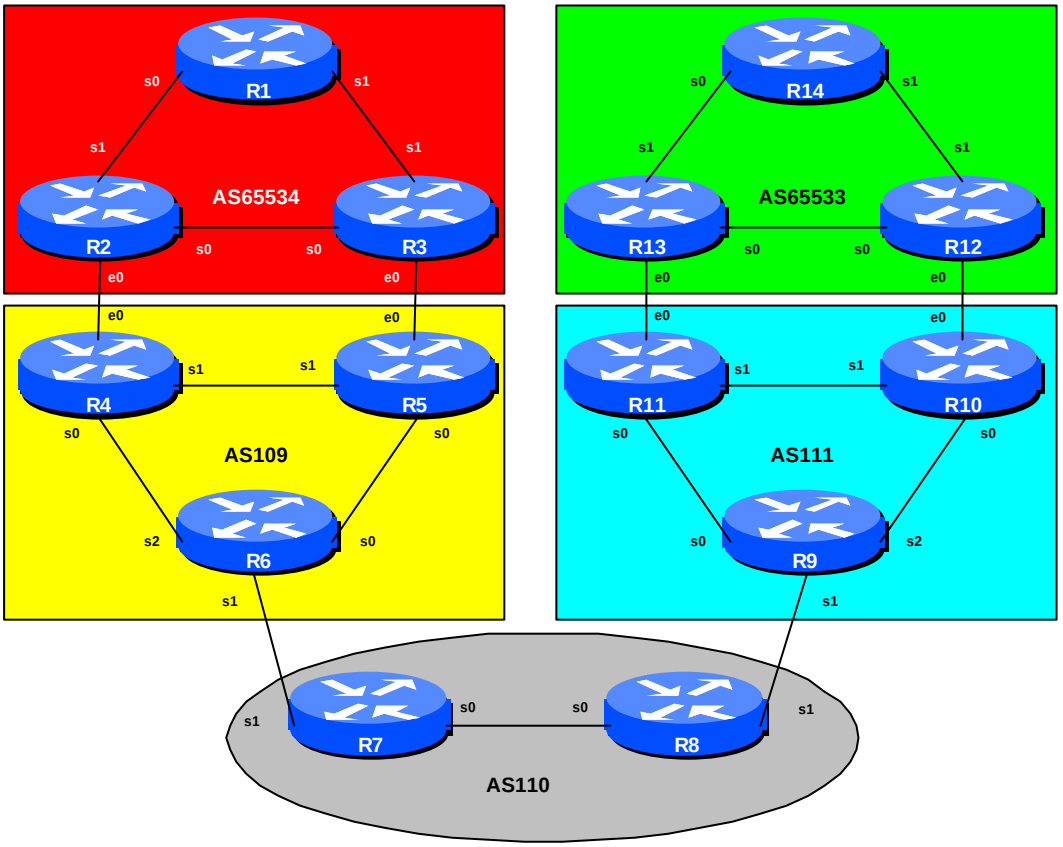

Figure 1 – ISP Lab Basic Configuration

Lab Notes

The purpose of this module is to demonstrate multihoming in the situation where the customer AS has more than one connection to the upstream service provider using RFC1998.

RFC1998 is an informational RFC which defines certain communities to have particular meanings in an ISP environment. Some ISPs follow the recommendations of RFC1998 – the situation where the ISP does support the RFC gives greater flexibility to the customer, and reduces the burden on the ISP when configuration changes are required. Some ISPs have taken the basics of RFC1998 and greatly enhanced them – these will be encountered in Module 17.

It is important that you review the RFC1998 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 situation being examined is one where one of the links is used and the other is kept purely for backup. This can be extended into a more complex situation, but that is left as an exercise for the reader.

Lab Exercise

Scenario One – Introduction to Community Based Multihoming

1. **Basic Configuration.** Each router team should configure their router to fit into the network layout depicted in Figure 2. **Notice that Router6 and Router9 require 3 serial ports** (the 3620s in the ISP Workshop kit have 4 serial ports). 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. 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**		

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

3. **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.

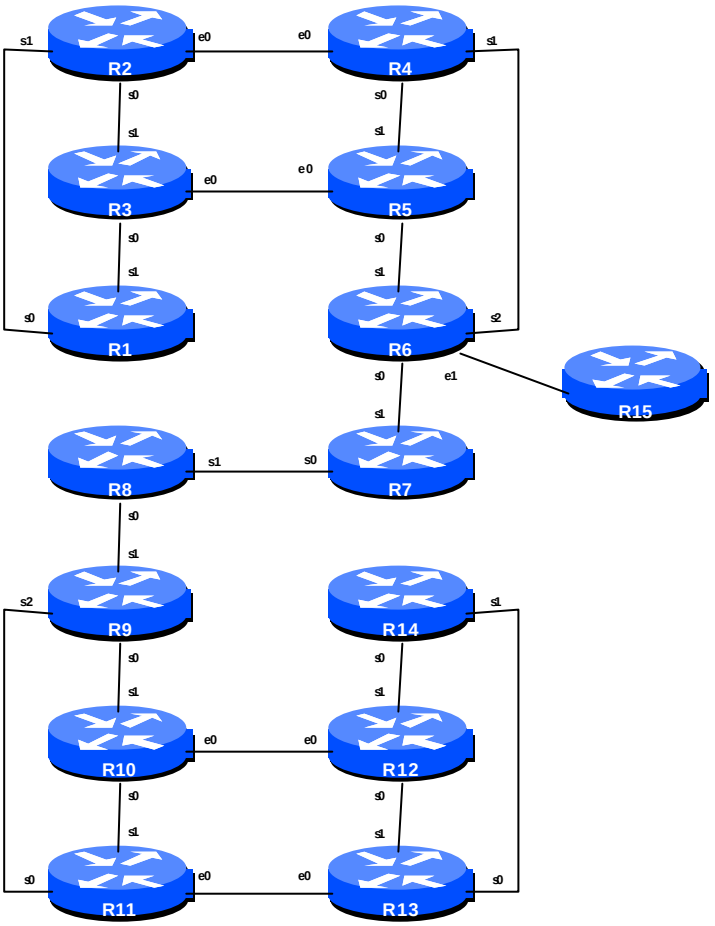

Figure 2 – Multihoming Lab Physical Layout

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.*

4. **RFC1998 definitions.** RFC1998 basically defines particular communities to have special meaning. These are summarised as follows:

ASx:100	**preferred route**
ASx:90	**backup route if dualhomed on ASx**
ASx:80	**main link is to another ISP with the same AS path length**
ASx:70	**main link is to another ISP**

If a customer wants to change how the loadsharing between him and his upstreams is functioning, he sends the appropriate community to achieve this. For example, if he is dualhomed onto on upstream and desires all his traffic to come in a particular link, he would send the community **ASx:100** out that link, and the community **ASx:90** out the other link. If at some stage in the future he wants to change this around, he starts sending **ASx:100** out the second link, and **ASx:90** out the first link. The point is that the configuration is entirely in the customer's hand – he can control the loadsharing he desires simply by sending communities to his upstream ISP.

The upstream ISP who supports RFC1998 has a generic route map which he uses on all his customer BGP peerings. This scales, and is much easier than handcrafting a configuration per customer. A sample configuration might be:

```
ip community-list 70 permit 109:70
ip community-list 80 permit 109:80
ip community-list 90 permit 109:90
ip community-list 100 permit 109:100
!
route-map customer-in permit 10
 match community 70
 set local-preference 70
route-map customer-in permit 20
 match community 80
 set local-preference 80
route-map customer-in permit 30
 match community 90
 set local-preference 90
route-map customer-in permit 40
 match community 100
```

```
 set local-preference 100
route-map customer-in permit 50
 ! fall through – nothing to be changed
!
router bgp 109
 neighbor x.x.x.x remote-as 110
 neighbor x.x.x.x route-map customer-in in
!
```

5. **Configure the main link.** Configure the main link between the private AS and the ISP. For AS65534, the link between Router2 and Router4 in AS109 is the main link – the link between Router3 and Router5 is the backup. For AS65533, the main link is between Router 13 and Router 11 in AS111. Example configuration for Router2:

```
ip prefix-list myblock permit 220.10.0.0/19
ip prefix-list default permit 0.0.0.0/0
!
route-map outfilter permit 10
 match ip address prefix-list myblock
 set community 109:100
route-map outfilter permit 20
!
route-map infilter permit 10
 match ip address prefix-list default
 set local-preference 100
route-map infilter permit 20
!
router bgp 65534
 network 220.10.0.0 mask 255.255.224.0
 neighbor <router4> remote-as 109
 neighbor <router4> description Link to Router4 in AS109
 neighbor <router4> prefix-list myblock out
 neighbor <router4> prefix-list default in
 neighbor <router4> route-map outfilter out
 neighbor <router4> route-map infilter in
!
ip route 220.10.0.0 255.255.224.0 null0 250
```

6. **Configure the backup link.** Configure the backup link between the private AS and the ISP. Send community 109:90 on outbound announcements, and set local preference on inbound announcements to 90. 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 community 109:90
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 221.19.0.0 mask 255.255.224.0
 neighbor <router10> remote-as 111
 neighbor <router10> description Link to Router10 in AS111
 neighbor <router10> prefix-list myblock out
 neighbor <router10> prefix-list default in
 neighbor <router10> route-map outfilter out
 neighbor <router10> route-map infilter in
!
ip route 221.19.0.0 255.255.224.0 null0 250
```

7. **Configure the routers in AS109 and AS111 to peer with the private ASes.** Using the definitions given in step 4 above, configure the routers in AS109 and AS111 to peer with the private ASes. Don't forget the prefix-list to filter inbound and outbound announcements. And remember to originate the default route for the private ASes.

8. **Configure the routers in AS109 and AS111 to peer with those in AS110.** This configuration is very similar to that covered in Module 12. Remember to remove the private AS using the remove-private-AS BGP command. Don't forget prefix-list filtering inbound and outbound, etc.

9. **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 #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 still should* ***not*** *see any private ASes in the BGP table of AS110.*

10. **Swap primary and backup paths.** Alter the configuration on the private AS border routers so that the primary path and backup paths are swapped. In other words, you now want to configure the path between Router3 and Router5 to be the main link between AS65534 and AS109. And similarly for AS65533 and AS111. **Hint:** All you have to do is change the "set community" command in the outbound route-map and clear the bgp session. And you will similarly need to change the local-preference setting for the inbound route-map.

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

Scenario Two – Scaling to support multiple dualhomed customers

The second scenario shows how to scale the 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.

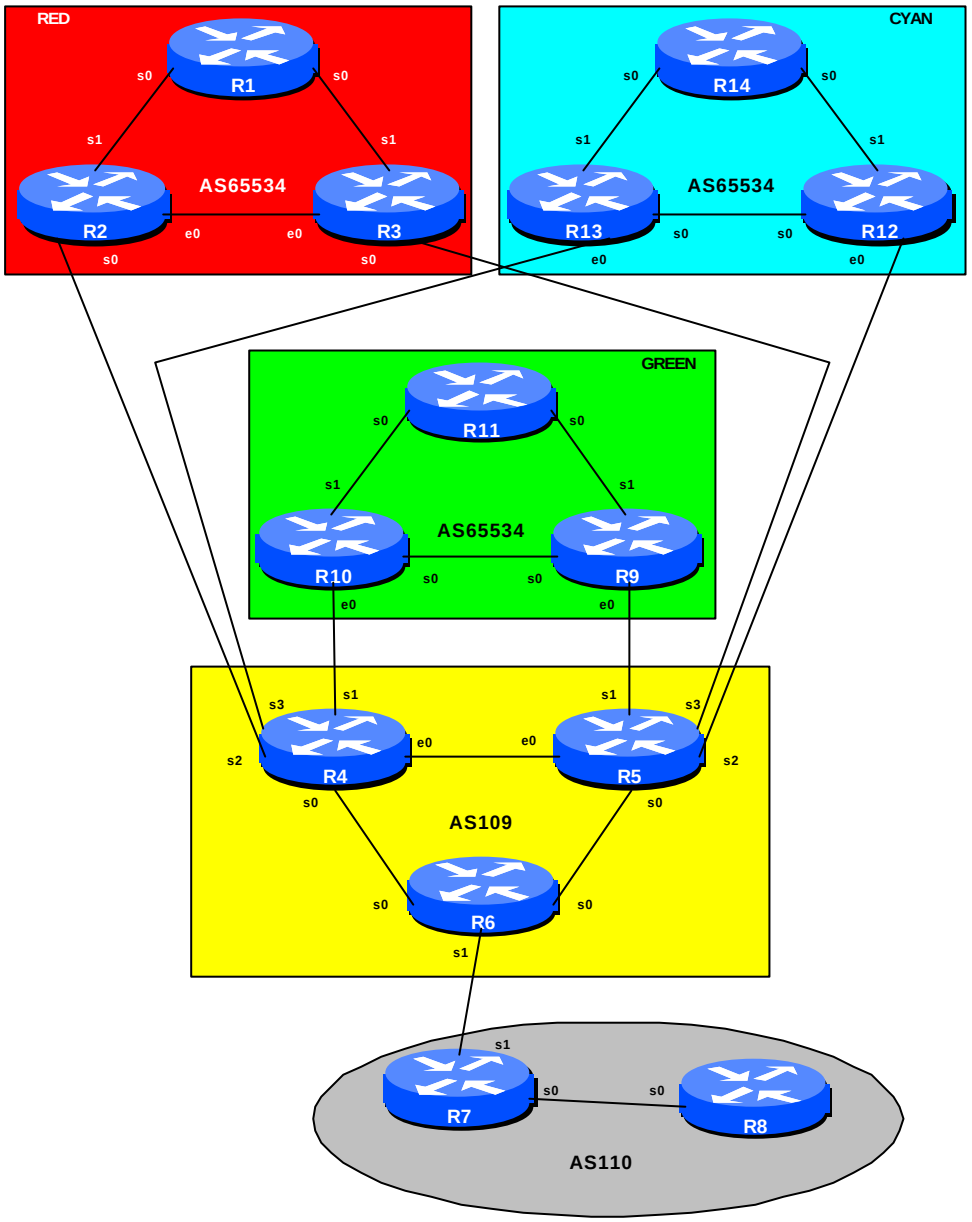

Figure 3 – Multiple Dualhomed Customers

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

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.

12. **Reconfigure the network.** Routers in AS110, AS111 and AS65533 (i.e. Router9 to Router14) should be reconfigured to become customers of AS109. Please refer to Figure 3 for AS topology details and to Figure 4 for the physical layout. 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 to 11 will need to set up OSPF and iBGP within their own network.

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

AS110		**AS65534 (R1-R3)**	**220.10.0.0/19**
	221.19.0.0/19	**AS65534 (R9-R11)**	**221.35.0.0/19**
AS109		**AS65534 (R12-R14)**	**220.19.0.0/19**
	220.73.0.0/19		

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

```
ip prefix-list myblock permit y.y.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 community 109:90
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 65534
 network y.y.0.0 mask 255.255.224.0
 neighbor x.x.x.x remote-as 109
 neighbor x.x.x.x description Link to RouterX in AS109
 neighbor x.x.x.x prefix-list myblock out
 neighbor x.x.x.x prefix-list default in
```

```
 neighbor x.x.x.x route-map outfilter out
 neighbor x.x.x.x route-map infilter in
!
ip route y.y.0.0 255.255.224.0 null0 250
```

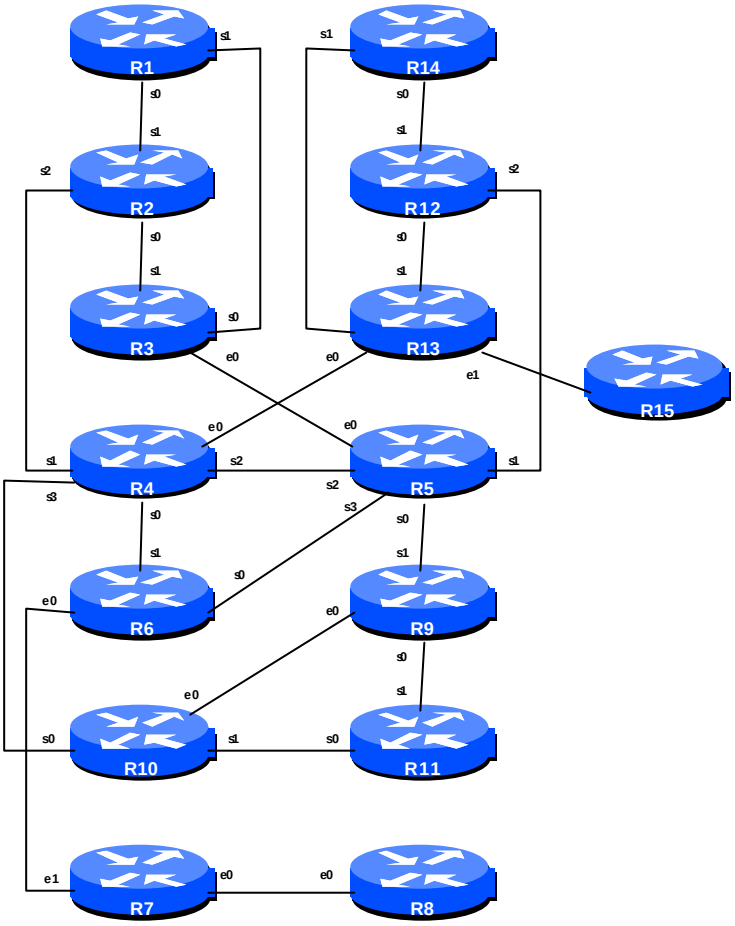

Figure 4 – Lab Physical Layout

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

15. **Configure eBGP on AS109 border routers.** Scalable eBGP configuration on Routers 4 and 5 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 only. The "customer-in" route-map is as used in the initial part of this module, and implements the community based policy matching.

```
router bgp 109
 neighbor bgp-customers peer-group
 neighbor bgp-customers remote-as 65534
 neighbor bgp-customers default-originate
 neighbor bgp-customers prefix-list default out
 neighbor bgp-customers route-map customer-in in
!
```

    After creating the peer-group, it can then be applied uniformly 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
!
```

16. **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 remove-private-AS
 neighbor x.x.x.x prefix-list mynets out
```

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

17. **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.

18. **Summary.** This module has covered the major situations where a customer requires to be multihomed onto the service provider backbone. It has demonstrated how to implement this multihoming using prefix-lists and communities based on RFC1918 where appropriate. It has also demonstrated the bgp remove-private-AS command, which ensures that private ASes are stripped out of any announcements to the wider Internet.

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

CONFIGURATION NOTES

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