

Module 8 – Multihoming Strategies Lab

Objective: Introduction to routing policy, the manipulation of BGP attributes to control traffic flow in a multihomed network.

Prerequisite: Module 6 and 7

Topology :

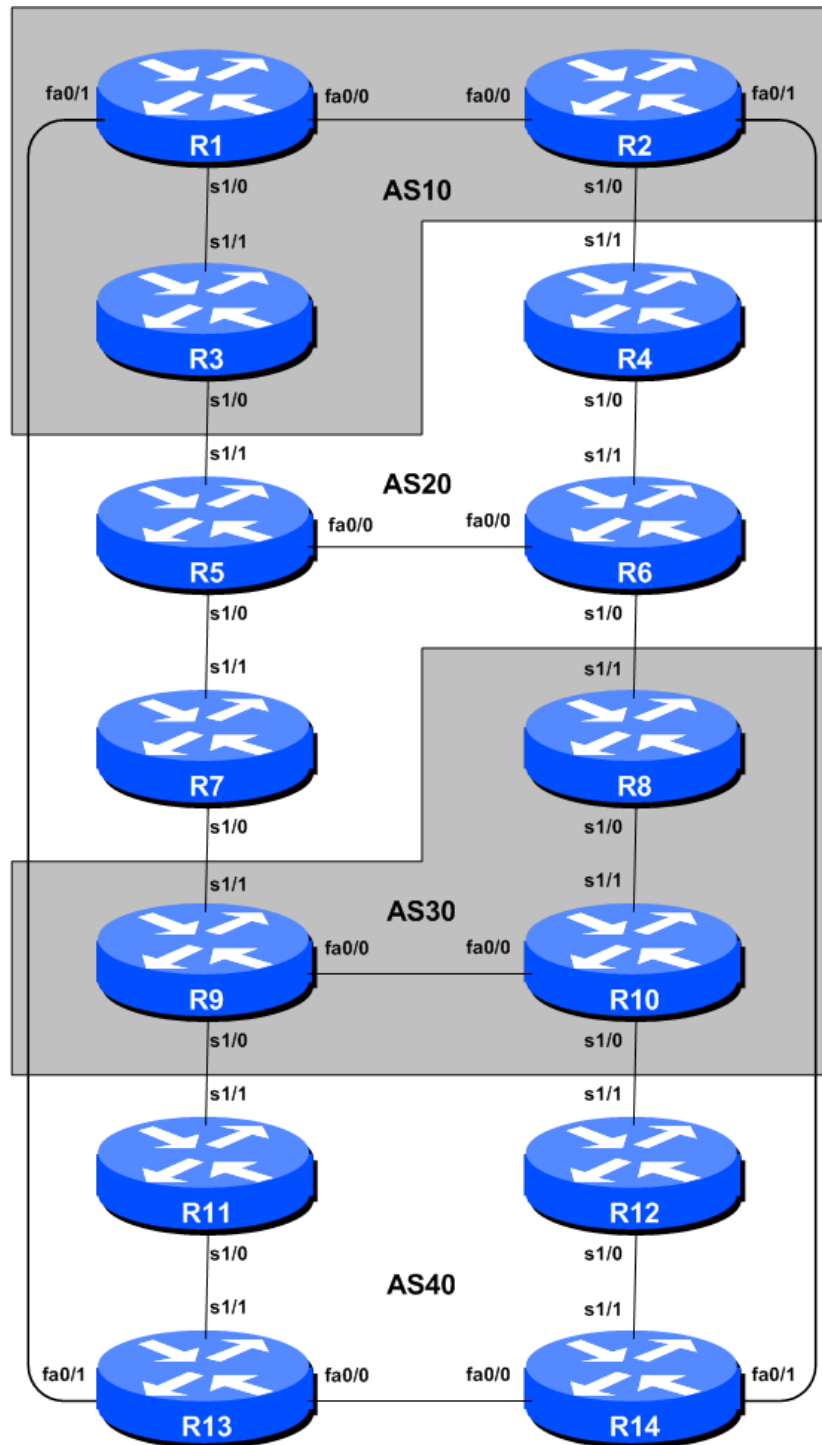


Figure 1 – BGP AS Numbers

Lab Notes

This module demonstrates how an AS can use Local Preference to control outbound routing paths, and how to use AS path prepend, MEDs (multi-exit discriminators or metrics) or communities to determine inbound routing paths. All four are very powerful tools for ISPs to control how their external peering links are used. Refer to the BGP presentations and documentation for more information about the BGP path selection process and the default values for, and priorities of, the “local_pref” and “metric” attributes.

Before starting this module, retain the topology and configurations as used in Module 6. This requires the removal of **all** the filtering and community configurations examined in Module 7.

Recommendation: Remember, if any configuration on a router is not in use, **it should be removed**. Surplus configuration usually gives rise to delayed error detection and debugging of configurations in cases of routing problems or other network failures.

Lab Exercise

1. **Tidy up from Module 7.** If the previous module completed was Module 7, the router configuration needs to be tidied up substantially before this module is attempted. The following steps show exactly what is required.

Example: Router R1

```
Router1#conf t
Router1(config)#router bgp 10
!
! First remove BGP neighbour route-map statement
!
Router1(config-router)#no neighbor 10.10.15.14 route-map infiltrer in
!
! Now remove community-tag from network statement
!
Router1(config-router)#no network 10.10.0.0 mask 255.255.255.192 route-map
community-tag
Router1(config-router)#network 10.10.0.0 mask 255.255.255.192
!
! Now remove route-maps
!
Router1(config)#no route-map community-tag
Router1(config)#no route-map infiltrer
!
! Now remove community list
!
Router1(config)#no ip community-list 1
!
! That's the configuration nice and tidy, the way it should be.
!
Router1(config)#end
!
! Now clear the bgp peering so that the old policy is removed
!
Router1#clear ip bgp 40 in
Router1#
```

Checkpoint #1: Call your lab instructor and display the following:

- i) Output of a “show ip route” and “show ip bgp”*
- ii) Outputs of the ‘ping’ and ‘trace’ to various destinations within the network*
- iii) Outputs of the ‘ping’ and ‘trace’ after the primary link fails.*

2. Aim of the Module:

The aim of the module is to demonstrate how to achieve a particular traffic flow using four different methods. These four methods involve applying policies for the modification of outbound traffic flow, and three ways of applying policies to modify inbound traffic flow. The reader should review the BGP presentation given prior to this module as a reminder on how to influence path choice using BGP policies.

The diagram following (Figure 2) displays the desired traffic flows between particular routers and ASes. Eight traffic flows are to be implemented. The arrows in the figure show the flows that will be configured. Each arrow originates from a border router in an AS, and terminates on one of the routers in another AS. This signifies the traffic flows desired for the links between the two systems. Each of the following steps has a description on how to implement the traffic flow represented by each arrow. If at any time there is any doubt as to the configuration required, ask the lab instructors.

3. **Assigning Address space within each ASN.** To make the different policies in this module work, we will use the customer prefixes we assigned to each router in Module 6. This way we will be able to target our traffic engineering based on what individual routers announce to the rest of the lab network. This is similar to the real life scenario where an ISP needs to do traffic engineering to meet particular operational requirements. As a reminder, the assignments for each ASN made in Module 6 are given in the following table:

AS10		AS30	
Router1	10.10.0.0/26	Router8	10.30.0.0/26
Router2	10.10.0.64/26	Router9	10.30.0.64/26
Router3	10.10.0.128/26	Router10	10.30.0.128/26
AS20		AS40	
Router4	10.20.0.0/26	Router11	10.40.0.0/26
Router5	10.20.0.64/26	Router12	10.40.0.64/26
Router6	10.20.0.128/26	Router13	10.40.0.128/26
Router7	10.20.0.192/26	Router14	10.40.0.192/26

As we will be doing traceroutes through the network to check that the policies have worked, it’s also a good idea to configure another loopback interface with one address out of the /26 address block. Recall that we routed the block to the Null0 interface in lieu of having a real customer router to connect to. Setting up another Loopback as a /32 from this /26 block will serve our purpose well.

Each team should choose **Loopback1** and give it the first address out of the block. For example, on Router8:

```
interface loopback 1
description Routing Policy Target Address
ip address 10.30.0.1 255.255.255.255
!
```

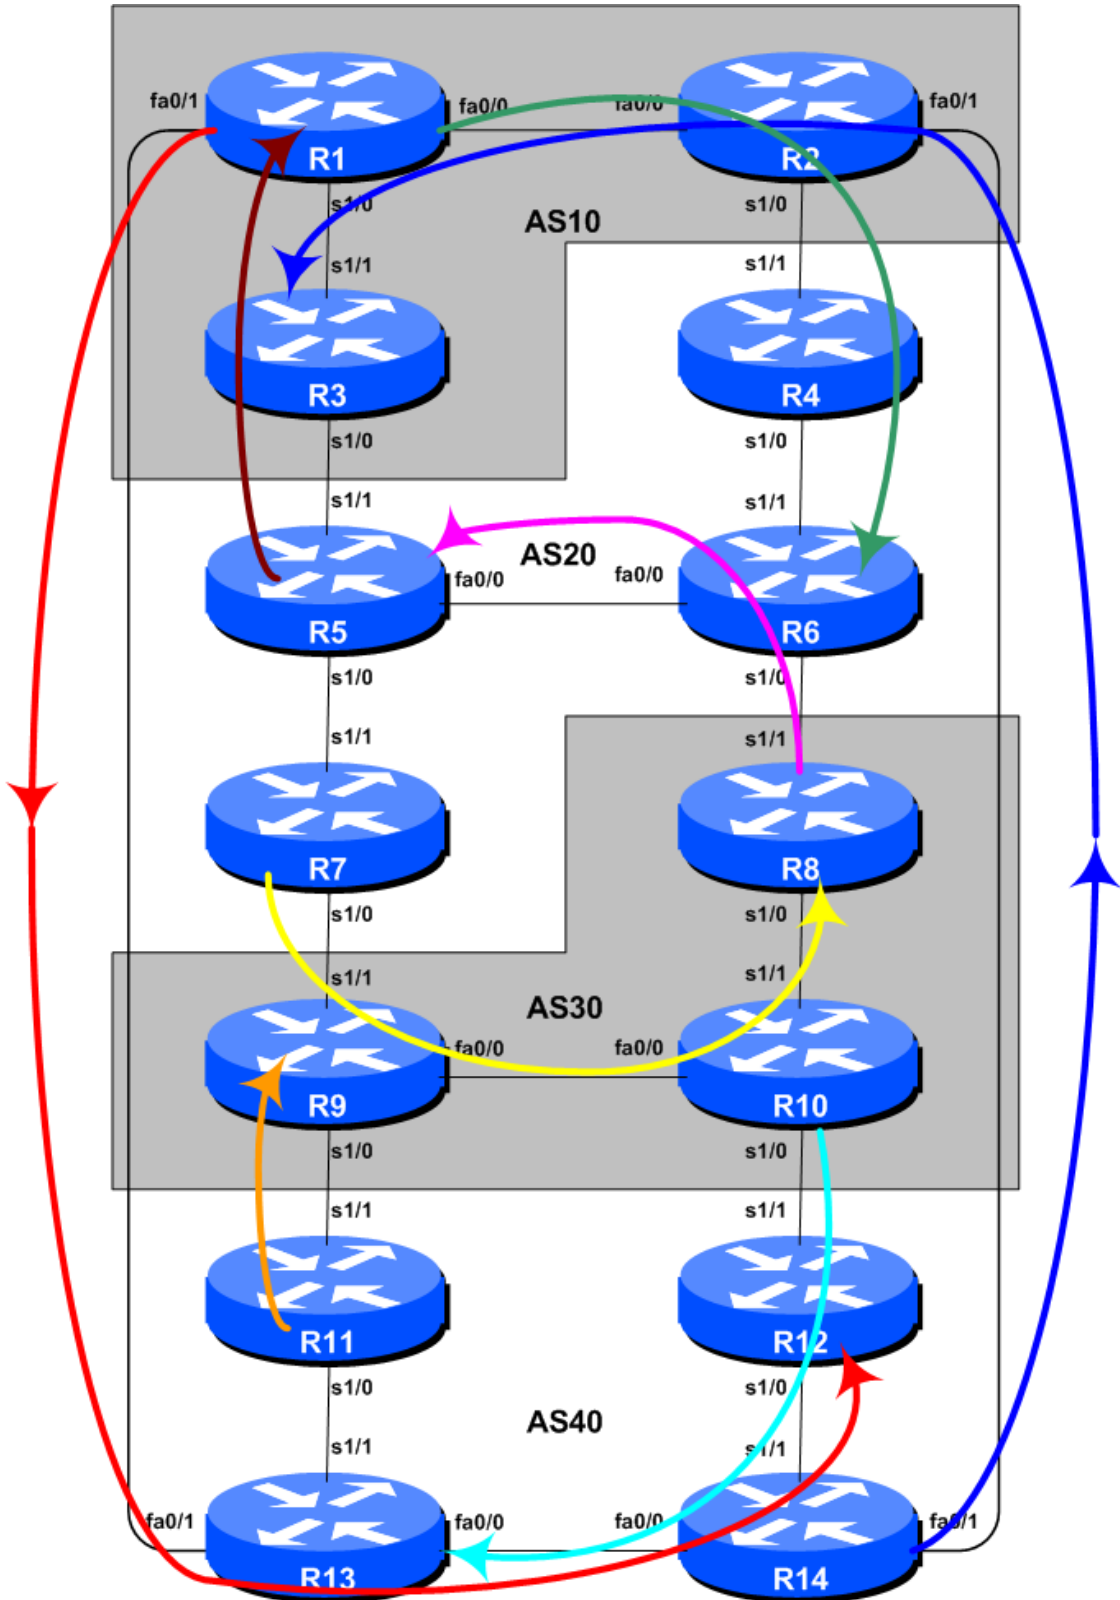


Figure 2 – Preferred Paths

4. Implement the following policies to influence outbound traffic flows using LOCAL_PREF:

This scenario demonstrates how to influence outbound traffic flows for each AS. Local preference is applied to inbound routing announcements. The address blocks defined in the previous step are used as the target: they are matched on the **eBGP peerings** and local preference is set high. While we set local preference high on the preferred outbound path, it is important that backup paths should still function – for that reason, we also set local preference low on the other exits from the ASN. Here are the primary paths/routing that we want to achieve:

AS 10:

- All traffic **TO** 10.20.0.128/26 (R6) must exit AS 10 via Router R2 only.
- All traffic **TO** 10.40.0.64/26 (R12) must exit AS 10 via Router R1 only.

AS 20:

- All traffic **TO** 10.10.0.0/26 (R1) must exit AS 20 via Router R5 only.
- All traffic **TO** 10.30.0.0/26 (R8) must exit AS 20 via Router R7 only.

AS 30:

- All traffic **TO** 10.20.0.64/26 (R5) must exit AS 30 via Router R8 only.
- All traffic **TO** 10.40.0.128/26 (R13) must exit AS 30 via Router R10 only.

AS 40:

- All traffic **TO** 10.10.0.128/26 (R3) must exit AS 40 via Router R14 only.
- All traffic **TO** 10.30.0.64/26 (R9) must exit AS 40 via Router R11 only.

Note that we are only trying to define outgoing traffic flow. The return path has no policies implemented, and the router's normal path decision process applies.

The example configurations below should be used as guidance for the router configuration for each router team. Note that for each scenario, one router in each AS will have to set high local preference and the remaining routers in the AS will set low local preference. The motivation for doing this is redundancy of configuration. If, for example, Router 7 lost its local-preference configuration due to some operator error, the low local preference set on the other three routers will ensure that the traffic policies required will still be implemented. It's quite common for many ISPs to have more than one configuration to implement a particular policy – a primary configuration, and an “inverse” secondary configuration on other routers which could be impacted.

Given there are two scenarios for each AS, each router in the AS will have to create a route-map covering both scenarios. Note that only one route-map can be used per eBGP session. The example below shows what Router 7's configuration could look like.

Sample configurations for Router 7 for the AS 20 scenario above (using LOCAL_PREF).

```
ip prefix-list R1-prefix permit 10.10.0.0/26
ip prefix-list R8-prefix permit 10.30.0.0/26
!
route-map as30-localpref permit 10
  match ip address prefix-list R1-prefix
  set local-preference 50
!
```

```
route-map as30-localpref permit 20
  match ip address prefix-list R8-prefix
  set local-preference 200
!
route-map as30-localpref permit 30
!
router bgp 20
  neighbor 10.30.15.5 remote-as 30
  neighbor 10.30.15.5 route-map as30-localpref in
!
```

Once the policy has been implemented, try doing a traceroute (`trace` command, and remember to specify the source interface – why?) through the network, to the destination you have targeted. Does the traceroute follow the arrows in the diagram? If not, why not? Maybe check the BGP table and if your colleagues in your AS have completed their steps in the exercise as well.

Checkpoint #2: Call your lab instructor and display the following:

i) Each router in an AS will be asked to do a ‘trace’ to selected destinations, the trace must show packets exiting the AS as specified in the exercise given above.

ii) Explain your configuration used to achieve the desired result to the instructor. Display the output of “show ip bgp”, and “show ip bgp x.x.x.x” for the networks with local preference set to 200. Show the output of a trace according to instructions.

iii) Wait until the lab instructor gives the goahead to move onto the next step.

STOP AND WAIT HERE

- 5. Remove configuration from the previous step.** Before moving on to the next step it is important that the configuration from the previous step is removed. This involves removing the route-maps, prefix-lists and the per-neighbour configuration to set local preference. All router teams should do this, and then do a soft reset of their eBGP peerings.
- 6. Implement the following policies to influence inbound traffic flows using MEDs.** This step introduces one of three methods of influencing inbound policies. Here MEDs are used, while the next steps will introduce the use of BGP communities and AS path prepends. As for the previous step, read the instructions careful, and discuss in your team, and in your AS, how you are going to implement the following.

The example in this step achieves exactly the same traffic flow between neighbouring ASes as in the previous step for the networks in question – remember that local preference is used by an AS to influence outbound traffic paths, whereas MEDs are used to influence inbound traffic paths. Refer to Figure 2 for a picture of traffic flow...

AS 10:

- All traffic **TO** 10.10.0.0/26 (R1) from anywhere in AS 20 must enter AS 10 via the R5 – R3 link. (**Hint:** this means that R2 must announce 10.10.0.0/26 to AS 20 with a higher metric than the equivalent announcement from R3.)

- All traffic **TO** 10.10.0.128/26 (R3) from anywhere in AS 40 must enter AS 10 via the R14 – R2 link. (**Hint:** this means that R1 must announce 10.10.0.128/26 to AS 40 with a higher metric than the equivalent announcement from R2.)

AS 20:

- All traffic **TO** 10.20.0.64/26 (R5) from anywhere in AS 30 must enter AS 20 via the R8 – R6 link. (**Hint:** this means that R7 must announce 10.20.0.64/26 to AS 30 with a higher metric than the equivalent announcement from R6.)
- All traffic **TO** 10.20.0.128/26 (R6) from anywhere in AS 10 must enter AS 20 via the R2 – R4 link. (**Hint:** this means that R5 must announce 10.20.0.128/26 to AS 10 with a higher metric than the equivalent announcement from R4.)

AS 30:

- All traffic **TO** 10.30.0.0/26 (R8) from anywhere in AS 20 must enter AS 30 via the R7 – R9 link. (**Hint:** this means that R8 must announce 10.30.0.0/26 to AS 20 with a higher metric than the equivalent announcement from R9.)
- All traffic **TO** 10.30.0.64/26 (R9) from anywhere in AS 40 must enter AS 30 via the R11 – R9 link. (**Hint:** this means that R10 must announce 10.30.0.64/26 to AS 40 with a higher metric than the equivalent announcement from R9.)

AS 40:

- All traffic **TO** 10.40.0.64/26 (R12) from anywhere in AS 10 must enter AS 40 via the R1 – R13 link. (**Hint:** this means that R14 must announce 10.40.0.64/26 to AS 10 with a higher metric than the equivalent announcement from R13.)
- All traffic **TO** 10.40.0.128/26 (R13) from anywhere in AS 30 must enter AS 40 via the R10 – R12 link. (**Hint:** this means that R11 must announce 10.40.0.128/26 to AS 30 with a higher metric than the equivalent announcement from R12.)

Suggestion: Choose an MED of 10 for the “low MED” configuration and an MED of 200 for “high MED” configuration. But most important of all, confer with your colleagues in your AS to make sure that your “low MED” value is lower than their chosen “high MED” value.

Sample configuration for Router 6 for the AS 20 scenario above (using MED).

```
ip prefix-list R5-prefix permit 10.20.0.64/26
!
route-map as30-med permit 10
  match ip address prefix-list R5-prefix
  set metric 10
!
route-map as30-med permit 20
!
router bgp 20
  neighbor 10.20.15.18 remote-as 30
  neighbor 10.20.15.18 route-map as30-med out
!
```

Checkpoint #3: Call your lab instructor and display the following:

ij Each router in an AS will be asked to do a ‘traceroute’ to selected destinations, the trace must show packets exiting the AS as specified in the exercise given above.

ii) Explain your configuration used to achieve the desired result to the instructor. Display the output of “show ip bgp”, and “show ip bgp x.x.x.x” for the networks with MED set to 50. Show the output of a trace according to instructions.

STOP AND WAIT HERE

- 7. Remove configuration from the previous step.** Before moving on to the next step it is important that the configuration from the previous step is removed. This involves removing the route-maps, prefix-lists and the per-neighbour configuration to set MEDs. All router teams should do this, and then do a route refresh of their eBGP peerings.
- 8. Implementing policies using BGP communities.** This section describes how to use BGP communities to influence inbound policies. Rather than using MEDs as we did in the previous scenario (Step 6), we will signal to our neighbouring ASN by using BGP communities. To prepare for this we need to establish which communities will implement the policies replacing the MEDs. Setting low MED means that the path would be preferred over one with high MED. This could be replicated in the peer AS by that AS setting high local preference on the path that would have heard the low MED, and by setting low local preference on the path that would have heard the high MED.
- 9. Choose communities to set high and low priority for outbound traffic.** As a recommendation, configure community <localASN>:150 to set high priority for outbound traffic, and <localASN>:50 to set low priority for outbound traffic. This means that each AS will have to configure a community matching route-map to set the appropriate local-preference value.

Example configuration for Router in AS20:

```
ip community-list 1 permit 20:150
ip community-list 2 permit 20:50
```

- 10. Set up the route-map configuration to implement the community policy.** We now create a route map which will set local-preference for each of the communities. If the eBGP neighbour sends us the community attached to a prefix, we will apply local preference to the prefix depending on the community value attached.

Example configuration:

```
route-map customer-comm permit 10
  match community 1
  set local-preference 150
!
route-map customer-comm permit 20
  match community 2
  set local-preference 50
!
route-map customer-comm permit 30
!
```


- 11. Apply the route-map to eBGP neighbours.** With the route-map configured, we now apply it to all our eBGP neighbours. When they send prefixes with the appropriate communities set, we will now set the local preference.

Example configuration for Router 8:

```
router bgp 30
  neighbor 10.20.15.17 remote-as 20
  neighbor 10.20.15.17 description eBGP peering with R6
  neighbor 10.20.15.17 route-map customer-comm in
!
```

While the whole concept looks more complex at first glance, it actually scales a lot better as the service provider is able to standardise their policy configuration using communities all across their access network.

- 12. Implement the following policies to influence inbound traffic flows using BGP communities.** As for the previous step, please read the instructions carefully, and discuss within your team, and in your AS, how you are going to implement the following.

The example in this step achieves exactly the same traffic flow between neighbouring ASes as in the previous step for the networks in question. Refer to Figure 2 for a picture of traffic flow...

AS 10:

- All traffic **TO** 10.10.0.0/26 (R1) from anywhere in AS 20 must enter AS 10 via the R5 – R3 link. (**Hint:** this means that R2 must announce 10.10.0.0/26 to AS 20 with AS20's low priority community whereas the equivalent announcement from R3 needs AS20's high priority community.)
- All traffic **TO** 10.10.0.128/26 (R3) from anywhere in AS 40 must enter AS 10 via the R14 – R2 link. (**Hint:** this means that R1 must announce 10.10.0.128/26 to AS 40 with AS40's low priority community whereas the equivalent announcement from R2 needs AS40's high priority community.)

AS 20:

- All traffic **TO** 10.20.0.64/26 (R5) from anywhere in AS 30 must enter AS 20 via the R8 – R6 link. (**Hint:** this means that R7 must announce 10.20.0.64/26 to AS 30 with AS30's low priority community whereas the equivalent announcement from R6 needs AS30's high priority community.)
- All traffic **TO** 10.20.0.128/26 (R6) from anywhere in AS 10 must enter AS 20 via the R2 – R4 link. (**Hint:** this means that R5 must announce 10.20.0.128/26 to AS 10 with AS10's low priority community whereas the equivalent announcement from R4 needs AS10's high priority community.)

AS 30:

- All traffic **TO** 10.30.0.0/26 (R8) from anywhere in AS 20 must enter AS 30 via the R7 – R9 link. (**Hint:** this means that R8 must announce 10.30.0.0/26 to AS 20 with AS20's low priority community whereas the equivalent announcement from R9 needs AS20's high priority community.)
- All traffic **TO** 10.30.0.64/26 (R9) from anywhere in AS 40 must enter AS 30 via the R11 – R9 link. (**Hint:** this means that R10 must announce 10.30.0.64/26 to AS 40 with AS40's low

priority community whereas the equivalent announcement from R9 needs AS40's high priority community.)

AS 40:

- All traffic **TO** 10.40.0.64/26 (R12) from anywhere in AS 10 must enter AS 40 via the R1 – R13 link. (**Hint:** this means that R14 must announce 10.40.0.64/26 to AS 10 with AS10's low priority community whereas the equivalent announcement from R13 needs AS10's high priority community.)
- All traffic **TO** 10.40.0.128/26 (R13) from anywhere in AS 30 must enter AS 40 via the R10 – R12 link. (**Hint:** this means that R11 must announce 10.40.0.128/26 to AS 30 with AS30's low priority community whereas the equivalent announcement from R12 needs AS30's high priority community.)

Look at the following examples to see what needs to be done for each router team.

Sample configuration for Router8 for the AS 20 scenario above (using Community).

```
ip prefix-list R8-prefix permit 10.30.0.0/26
!
route-map set-comm permit 10
  match ip address prefix-list R8-prefix
  set community 30:50
!
route-map set-comm permit 20
!
router bgp 30
  neighbor 10.20.15.17 remote-as 20
  neighbor 10.20.15.17 description eBGP peering with R6
  neighbor 10.20.15.17 route-map set-comm out
  neighbor 10.20.15.17 route-map customer-comm in
!
```

Checkpoint #4: Call your lab instructor and display the following:

i] Each router in an AS will be asked to do a 'traceroute' to selected destinations, the trace must show packets exiting the AS as specified in the exercise given above.

ii] Explain your configuration used to achieve the desired result to the instructor. Display the output of "show ip bgp", and "show ip bgp x.x.x.x" for the networks with configured communities. Show the output of a trace according to instructions.

STOP AND WAIT HERE

13. Remove the configuration used for the previous step. Before moving on to the next step it is important that the configuration from the previous step is removed. This involves removing the route-maps, prefix-lists and the per-neighbour configuration to set communities. All router teams should do this, and then do a soft reset of their eBGP peerings.

14. Implement the following policies to influence inbound traffic flows using the AS path prepend method. This step introduces the second of two methods of influencing inbound policies.

As for the previous step, read the instructions carefully, and discuss within your team, and within your AS, how you are going to implement the following.

The example in this step achieves exactly the same traffic flow between neighbouring ASes as in the previous step for the networks in question. Refer to Figure 2 for a picture of traffic flow...

AS 10:

- All traffic **TO** 10.10.0.128/26 (R3) from anywhere in the lab topology must enter AS 10 via the R14 – R2 link. (**Hint:** this means that R1 and R3 must announce 10.10.0.128/26 with a longer AS path than the other networks in AS 10. R2 needs to announce 10.10.0.128/26 with a longer AS path in its peering with R4.)
- All traffic **TO** 10.10.0.0/26 (R1) from anywhere in the lab topology must enter AS 10 via the R5 – R3 link. (**Hint:** this means that R1 and R2 must announce 10.10.0.0/26 with a longer AS path than the other networks in AS 10.)

AS 20:

- All traffic **TO** 10.20.0.64/26 (R5) from anywhere in the lab topology must enter AS 20 via the R8 – R6 link. (**Hint:** this means that R4, R5 and R7 must announce 10.20.0.64/26 with a longer AS path than the other networks in AS 20.)
- All traffic **TO** 10.20.0.128/26 (R6) from anywhere in the lab topology must enter AS 20 via the R2 – R4 link. (**Hint:** this means that R5, R6 and R7 must announce 10.20.0.128/26 with a longer AS path than the other networks in AS 20.)

AS 30:

- All traffic **TO** 10.30.0.0/26 (R8) from anywhere in the lab topology must enter AS 30 via the R7 – R9 link. (**Hint:** this means that R8 and R10 must announce 10.30.0.0/26 with a longer AS path than the other networks in AS 30. R9 needs to announce 10.30.0.0/26 with a longer AS path in its peering with R11.)
- All traffic **TO** 10.30.0.64/26 (R9) from anywhere in the lab topology must enter AS 30 via the R11 – R9 link. (**Hint:** this means that R8 and R10 must announce 10.30.0.64/26 with a longer AS path than the other networks in AS 30. R9 needs to announce 10.30.0.64/26 with a longer AS path in its peering with R7.)

AS 40:

- All traffic **TO** 10.40.0.64/26 (R12) from anywhere in the lab topology must enter AS 40 via the R1 – R13 link. (**Hint:** this means that R11, R12 and R14 must announce 10.40.0.64/26 with a longer AS path than the other networks in AS 40.)
- All traffic **TO** 10.40.0.128/26 (R13) from anywhere in the lab topology must enter AS 40 via the R10 – R12 link. (**Hint:** this means that R11, R13 and R14 must announce 10.40.0.128/26 with a longer AS path than the other networks in AS 40.)

AS_PREPEND is commonly used by smaller ISPs who are multihoming to their upstream providers. It is convention on the Internet to add at least two ASes when using AS_PREPEND. More usually, three ASes are added, especially if the upstream ISPs have links to each other going through a third party.

Sample configuration for Router 9 for the AS 30 scenario above (using AS PATH prepend):

```
ip prefix-list R8-prefix permit 10.30.0.0/26
ip prefix-list R9-prefix permit 10.30.0.64/26
!
```

```
route-map set-path-as20 permit 10
  match ip address prefix-list R9-prefix
  set as-path prepend 30 30 30
!
route-map set-path-as20 permit 20
!
route-map set-path-as40 permit 10
  match ip address prefix-list R8-prefix
  set as-path prepend 30 30 30
!
route-map set-path-as40 permit 20
!
router bgp 30
  neighbor 10.20.15.6 remote-as 20
  neighbor 10.20.15.6 description eBGP peering with R7
  neighbor 10.20.15.6 route-map set-path-as20 out
  neighbor 10.40.15.1 remote-as 40
  neighbor 10.40.15.1 description eBGP peering with R11
  neighbor 10.40.15.1 route-map set-path-as40 out
!
```

Checkpoint #5: Call the lab instructor and display the following:

ij Each router in an AS will be asked to do a ‘traceroute’ to selected destinations, the trace must show packets exiting the AS as specified in the exercise given above.

ii Explain your configuration used to achieve the desired result to the instructor. Display the output of “show ip bgp”, and “show ip bgp x.x.x.x” for the networks with increased AS path length. Show the output of a trace according to instructions.

iii How has AS Path prepend changed the BGP table and the Routing decision. Can the decision be overridden using any other BGP configuration within an AS? **Answer:** review the BGP route selection rules from the slide in the presentation section.

15. Summary:

This module has demonstrated several ways of influencing inbound and outbound routing policy.

Q: What is the difference in the resulting effects using the methods in steps 6, 8 and 14?

A: AS PATH prepend affects routing announcements between two ASes, and is visible everywhere, even outside the two ASes which are making use of the prepend information. MEDs only apply between multiple peerings between the same AS. If the peer AS is announcing the local AS onwards, the metric set is that of the peer AS, not the local AS. Communities are generally only used in peerings between neighbouring ASNs.

Consult the BGP documentation for more information. There are many possible variations on the examples given in this module. Remember the following points:

- local preference is used for influencing policy within an AS
- MEDs are used to influence policy over multiple links between the local and an immediately neighbouring AS

- BGP Communities can also be used to influence policy over multiple links between the local and neighbouring AS
- The AS path prepend is used to influence external policy on a global scale (which includes the immediately neighbouring AS).

Review Questions

1. Which is the most effective way of influencing how traffic leaves your network?
2. How useful do you think MEDs are in the real live Internet? Consider the answer to question one before replying!