Interdomain Multicast Lab: MSDP/MBGP

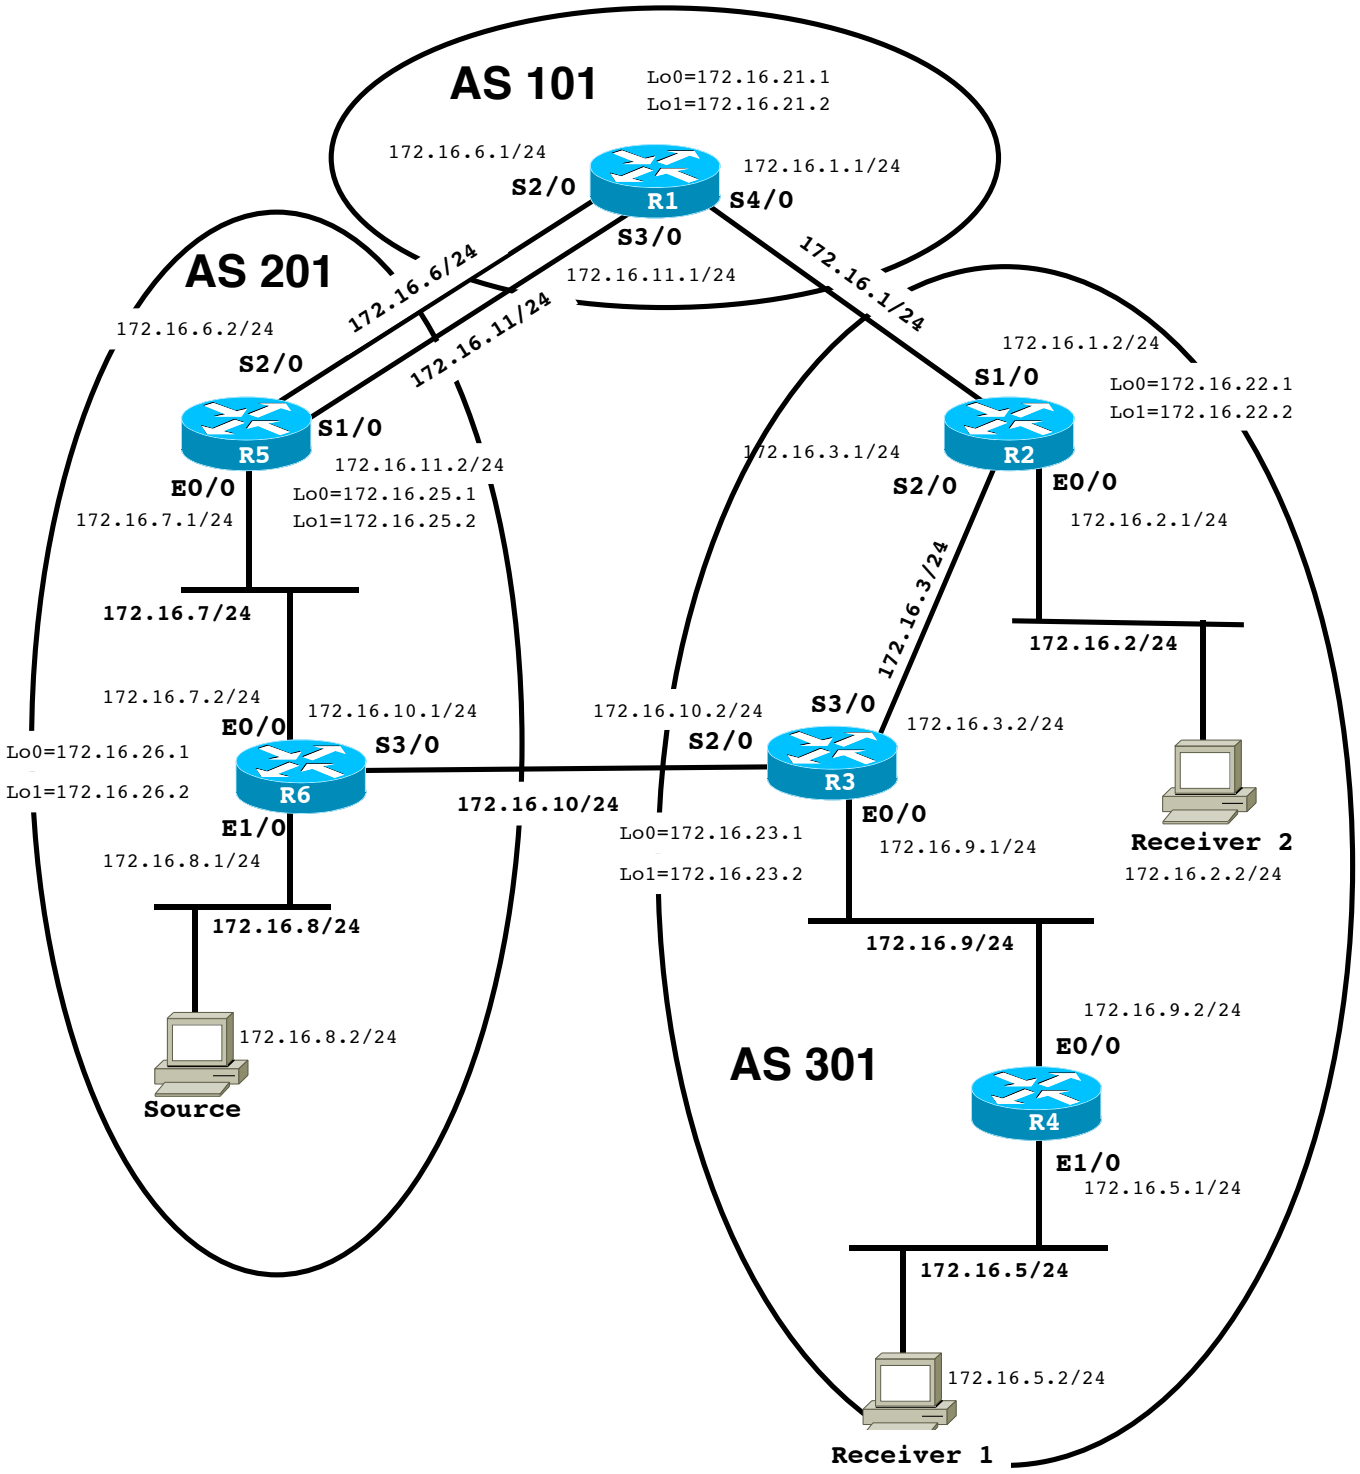

Figure 1 – MBGP/MSDP Lab Configuration

1. Lab Objectives

- Understand the operation of MBGP and MSDP protocol mechanics
    1. MBGP/MSDP control traffic
    2. Determine the state of the network from show commands
- Observe formation and maintenance of peering relationships
- Observe the source discovery and state creation

Figure 2 - MBGP/MSDP Logical Diagram

Figure 2 represents a logical diagram of the MBGP/MSDP lab. There are number of differences with the previous labs that are necessary to point out:

- The network is now divided into three Autonomous Systems. Each AS is a separate PIM domain with its own RP.
- Routers R1, R2, R3, R5 and R6 are either an eMBGP or iMBGP peer to their physical neighbors.
- R1, R2, R3 and R5 are MSDP peers. Notice that R3 is running MSDP even though it is not an RP.
- There are two physical links between R5 and R1. One of these links will be configured as a unicast only MBGP peer – the other a multicast only MBGP peer.

2. Pre-configuration Steps

Perform the following configuration steps in preparation for the exercises in this lab.

1. Ensure that all the interfaces of your router shown in the topology are configured for PIM Sparse mode.
2. Configure the second serial interface between R1 and R5 as shown in Figure 1.
3. Configure Loopback0 and Loopback1 interfaces on routers R1-R3, R5-R6.

R1:

```
interface Loopback0
 ip address 172.16.21.1 255.255.255.255
```

```
interface Loopback1
 ip address 172.16.21.2 255.255.255.255
```

R2:

```
interface Loopback0
 ip address 172.16.22.1 255.255.255.255

interface Loopback1
 ip address 172.16.22.2 255.255.255.255
```

R3:

```
interface Loopback0
 ip address 172.16.23.1 255.255.255.255

interface Loopback1
 ip address 172.16.23.2 255.255.255.255
```

R5:

```
interface Loopback0
 ip address 172.16.25.1 255.255.255.255

interface Loopback1
 ip address 172.16.25.2 255.255.255.255
```

R6:

```
interface Loopback0
 ip address 172.16.26.1 255.255.255.255

interface Loopback1
 ip address 172.16.26.2 255.255.255.255
```

4. Configure a static RP on every router for each Area

R1: `ip pim rp-address 172.16.21.2`

R2, R3, R4: `ip pim rp-address 172.16.22.2`

R5, R6: `ip pim rp-address 172.16.25.2`

• Ensure you can successfully ping and RPF check to the RP from your router.

4. MBGP Exercises

Perform the following pre-configuration steps:

1. Configure the eigrp interfaces between ASs to be passive and define some key static routes for injection into the BGP RIB's:

   R1:

```
router eigrp g
 passive-interface Serial2/0
 passive-interface Serial3/0
 passive-interface Serial4/0
 network 172.16.0.0
 no auto-summary
!
ip route 172.16.21.0 255.255.255.0 Null0
```

   R2:

```
router eigrp g
 passive-interface Serial1/0
 network 172.16.0.0
 no auto-summary
!
ip route 172.16.22.0 255.255.255.0 Null0
```

   R3:

```
router eigrp g
 redistribute static
 passive-interface Serial2/0
 network 172.16.0.0
 no auto-summary
```

```
!
ip route 0.0.0.0 0.0.0.0 Null0
ip route 172.16.23.0 255.255.255.0 Null0
```

R4:

```
router eigrp g
network 172.16.0.0
 no auto-summary
```

R5:

```
router eigrp g
 passive-interface Serial1/0
 passive-interface Serial2/0
 network 172.16.0.0
 no auto-summary
```

R6:

```
router eigrp g
 passive-interface Serial3/0
 network 172.16.0.0
 no auto-summary
```

2. Configure BGP between each of the Areas with the address families as follows:

R1:

```
router bgp 101
 no synchronization
 network 172.16.21.0 mask 255.255.255.0 nlri unicast multicast
 neighbor 172.16.1.2 remote-as 301 nlri unicast multicast
 neighbor 172.16.6.2 remote-as 201
 neighbor 172.16.11.2 remote-as 201 nlri multicast
 no auto-summary
```

R2:

```
router bgp 301
 no synchronization
 network 172.16.2.0 mask 255.255.255.0 nlri unicast multicast
 network 172.16.3.0 mask 255.255.255.0 nlri unicast multicast
 network 172.16.22.0 mask 255.255.255.0 nlri unicast multicast
 neighbor 172.16.1.1 remote-as 101 nlri unicast multicast
 neighbor 172.16.23.2 remote-as 301 nlri unicast multicast
 neighbor 172.16.23.2 update-source Loopback1
 neighbor 172.16.23.2 next-hop-self
 no auto-summary
```

R3:

```
router bgp 301
 no synchronization
 network 172.16.3.0 mask 255.255.255.0 nlri unicast multicast
 network 172.16.5.0 mask 255.255.255.0 nlri unicast multicast
 network 172.16.9.0 mask 255.255.255.0 nlri unicast multicast
 network 172.16.23.0 mask 255.255.255.0 nlri unicast multicast
 neighbor 172.16.10.1 remote-as 201 nlri unicast multicast
 neighbor 172.16.22.2 remote-as 301 nlri unicast multicast
 neighbor 172.16.22.2 update-source Loopback1
 neighbor 172.16.22.2 next-hop-self
 no auto-summary
```

R5:

```
router bgp 201
 no synchronization
 network 172.16.7.0 mask 255.255.255.0 nlri unicast multicast
 network 172.16.8.0 mask 255.255.255.0 nlri unicast multicast
 neighbor 172.16.6.1 remote-as 101
 neighbor 172.16.11.1 remote-as 101 nlri multicast
 neighbor 172.16.26.2 remote-as 201 nlri unicast multicast
```

```
 neighbor 172.16.26.2 update-source Loopback1
 neighbor 172.16.26.2 next-hop-self
 no auto-summary
```

R6:

```
router bgp 201
 no synchronization
 network 172.16.7.0 mask 255.255.255.0 nlri unicast multicast
 network 172.16.8.0 mask 255.255.255.0 nlri unicast multicast
 neighbor 172.16.10.2 remote-as 301 nlri unicast multicast
 neighbor 172.16.25.2 remote-as 201 nlri unicast multicast
 neighbor 172.16.25.2 update-source Loopback1
 neighbor 172.16.25.2 next-hop-self
 no auto-summary
```

4.1 MBGP Peers

Use the "*show ip bgp summary"*, "*show ip mbgp summary"* and "*show ip bgp neighbors*" command to answer the following question:

Questions :

1) Are all of the peering sessions up?

2) Are the outputs of the "*show ip bgp summary"* and "*show ip mbgp summary"* commands the same. If not, why?

3) What type of NLRI exchange has been negotiated between the peers; unicast, multicast or both?

R1:	**(R2):**	**R2:**	**(R1):**	**R3:**	**(R2):**
	(R5):		**(R3):**		**(R6):**
R5:	**(R1):**	**R6:**	**(R1):**		
	(R6):		**(R5):**		

4.2 BGP Unicast NLRI Exchange

Examine the output of "*show ip bgp", "show ip bgp <network-address>"* commands to answer the following questions:

Questions :

1) Which BGP peer advertised the best unicast path to network 172.16.21.0/24 to your router?

R1: **R2:** **R3:**

R5: **R6:**

2) What is the BGP next-hop address to network 172.16.21.0/24 of this path?

R1: **R2:** **R3:**

R5: **R6:**

3) What is the AS-path to 172.16.21.0/24 from your router?

R1: **R2:** **R3:**

R5: **R6:**

4.3 BGP Multicast NLRI Exchange

Examine the output of "*show ip mbgp", "show ip mbgp <network-address>"* commands to answer the following questions:

Questions :

1) Which MBGP peer advertised the best multicast RPF path back to network 172.16.21.0/24 to your router (sh ip rpf)?

R1: **R2:** **R3:**

R5: **R6:**

2) What is the next hop address to network 172.16.21.0/24 of this path?

R1: **R2:** **R3:**

R5: **R6:**

3) What is the AS-path to 172.16.21.0/24 from your router?

R1: **R2:** **R3:**

R5: **R6:**

5. MSDP Exercises

Perform the following pre-configuration steps:

1. Configure MSDP peers between R1, R2 and R5 as follows:

R1:

```
ip msdp peer 172.16.22.2 connect-source Loopback1 remote-as 301
ip msdp peer 172.16.25.2 connect-source Loopback1 remote-as 201
ip msdp cache-sa-state
```

R2:

```
ip msdp peer 172.16.21.2 connect-source Loopback1 remote-as 101
ip msdp peer 172.16.23.1 connect-source Loopback0 remote-as 201
ip msdp cache-sa-state
```

R3:

```
ip msdp peer 172.16.22.1 connect-source Loopback0 remote-as 301
ip msdp peer 172.16.25.2 connect-source Loopback1 remote-as 201
ip msdp cache-sa-state
```

R5:

```
ip msdp peer 172.16.21.2 connect-source Loopback1 remote-as 101
ip msdp peer 172.16.23.2 connect-source Loopback1 remote-as 301
ip msdp cache-sa-state
```

5.1 MSDP Peer Connections

Examine the output of "*show ip msdp summary*", "*show ip msdp peer", "show ip route"* and "*show ip bgp"* commands on R1, R2 and R5 and answer the following question:

Questions :

1) What is the connection status of each peer session?

R1: (R2):	**R2: (R1):**	**R5: (R1):**	**R3: (R2):**
(R5):	**(R3):**	**(R3):**	**(R5):**

2) What is the "Connection Source" of each peer session?

R1: (R2):	**R2: (R1):**	**R5: (R1):**	**R3: (R2):**
(R5):	**(R3):**	**(R3):**	**(R5):**

3) Why is the status of the peering sessions on R5 "Listen" and not "Connect" or "Down" as they are on R1 and R2?

4) Why are the peering sessions to R5 not coming up?

Add the following commands to R5 and recheck the connections with the "*show ip msdp summary"* command:

```
!
ip route 172.16.25.0 255.255.255.0 Null0
!
router bgp 201
 network 172.16.25.0 mask 255.255.255.0 nlri unicast multicast
```

5) Are all MSDP sessions up now?

5.2 MSDP SA RPF Checking

1. Activate the multicast Source so that it begins sending traffic to group (*,224.1.1.1).
    - Copy the config in file "pagent_SSM_config" and paste it into the cli of the "Source". The file can be viewed from the "Status" tab of WebIOU.
    - At the "Source" cli, enter *"tgn"* to access the traffic generator and "*start"* to enable the stream.
2. Use the *"show ip mroute 244.1.1.1"* and observe which flags are set for the entry.
3. Use the *"show ip msdp peer <address> advertised-SAs"* command to answer the following questions:

Questions :

1) What source is being advertised to the router's MSDP peers:

R1: (172.16.22.2 - R2):
(172.16.25.2 - R5):

R2: (172.16.21.2 – R1):
(172.16.23.1 – R3):

R3: (172.16.22.1 – R2):
(172.16.25.2 – R5):

R5: (172.16.21.2 – R1):
(172.16.23.2 – R3):

2) Why aren't there any sources being advertised by R2?

Now use the "*show ip msdp sa-cache"* on routers R1, R2, R3 and R5 to list the contents of the router's SA cache.

3) Does the source appear in the router's SA cache? Why?

R1: **R2:** **R3:** **R5:**

Enable "*debug ip msdp routes" "debug ip msdp peer"* on routers R1, R2, R3 and R5. Observe the debug output for a few minutes.

4) From which MSDP peers is each router receiving a copy of the SA message and which messages are passing the SA RPF check?

	Rcvd [Y/N]	**RPF Check[Pass/Fail]**
R1: (172.16.22.2 - R2):		
(172.16.25.2 - R5):		
R2: (172.16.21.2 – R1):		
(172.16.23.1 – R3):		
R3: (172.16.22.1 – R2):		
(172.16.25.2 – R5):		
R5: (172.16.21.2 – R1):		
(172.16.23.2 – R3):		

5) Why aren't any SA messages passing the RPF check on router R2?

Correct the configuration of R2 and R3 using the following configuration changes:

R2:

```
no ip msdp peer 172.16.23.1 connect-source Loopback0 remote-as
201
ip msdp peer 172.16.23.2 connect-source Loopback1 remote-as 201
```

R3:

```
no ip msdp peer 172.16.22.1 connect-source Loopback0 remote-as
101
```

```
ip msdp peer 172.16.22.2 connect-source Loopback1 remote-as 101
```

Observe the debug msdp messages on router R2 to verify that the SA messages arriving from R3 are passing the RPF check. Use the "*show ip msdp sa-cache"* command to verify that R2's SA cache now contains an entry for the source.

5.3 State with Receiver 1 – no source

1. Insure that your routers are configured as shown in Figure-1 and that no multicast sources or receivers are active – type "stop" on the source and wait for the mroute and msdp cache to timeout. This will take several minutes. (Note: You may ignore the 224.0.1.40 group. This is the "RP-Announce" group and is automatically joined by the routers in order to receive Auto- RP information.)

2. Add the following configuration command to router R2:

```
router bgp 301
 neighbor 172.16.1.1 weight 400
```

3. Clear the bgp peer connections on router R2 using the *"clear ip bgp *"* command. (Use the *"show ip bgp"* and *"show ip mbgp"* commands to verify that the weight is now 400 on all updates received from 172.16.1.1.

4. Using the "*show ip msdp sa-cache"* command, make sure that the SA caches in routers R1, R5 and R6 are clear. *"clear ip mroute" "clear ip msdp sa"* ..to speed things along.

5. Using the "*show ip mroute summary"* command, verify that no ip multicast state exists for multicast group `224.1.1.1`.

6. Turn on *"debug ip pim"* and *"debug ip mroute"* on R2 and R4. Activate multicast Receiver 1 shown in the configuration diagram so that it "joins" the `224.1.1.1` group – on eth 0/0 type *"ip igmp join 224.1.1.1"*

On each of the routers, examine the multicast state for group `224.1.1.1` using the *"show ip mroute 224.1.1.1"* command and answer the following questions.

Questions :

1) What is the current status of the Flags associated with the (*,G) and (S,G) entries for group `224.1.1.1` ?

R1: (*, G):	**R2: (*, G):**	**R3: (*, G):**
(S, G):	**(S,G):**	**(S,G):**
R4: (*, G):	**R5: (*, G):**	**R6: (*, G):**
(S, G):	**(S, G):**	**(S,G):**

2) Why is there no state on router R1?

3) Why does R2 have an OIF for the (S,G)?

4) Why is the (S,G) RPF out S1/0 on R2?

5) Why does R5 have a null OIF for the (S,G)?

5.4 State with Receiver and Source

5.4.1 State with Receiver 1 and Source – Shared Tree

1. Make sure that the following command is configured on R4 in the network:

   ip pim spt-threshold infinity

2. Activate the multicast Source so that it begins sending traffic to group `(*,224.1.1.1).`

3. On each of the routers, examine the multicast state for group `224.1.1.1` using the *"show ip mroute 224.1.1.1"* command and answer the following questions.

Questions :

1) What is the current status of the Flags associated with the (*,G) and (S,G) entries for group `224.1.1.1` ?

R1: (*, G):	**R2: (*, G):**	**R3: (*, G):**
(S, G):	**(S,G):**	**(S,G):**
R4: (*, G):	**R5: (*, G):**	**R6: (*, G):**
(S, G):	**(S, G):**	**(S,G):**

2) Examine the output of "show ip mdsp sa-cache"on each of the MSDP peers. Which RP is associated with the source entry?

5.4.2 State with Receiver 1 and Source – Source Tree

1. Stop the multicast Source and wait a brief time for the network to stabilize.

2. Make sure that the following command is configured on all routers in the network:

   ip pim spt-threshold 0

3. Activate the multicast Source so that it begins sending traffic to group `(*,224.1.1.1).`

4. On each of the routers, examine the multicast state for group `224.1.1.1` using the *"show ip mroute 224.1.1.1"* command and answer the following questions.

Questions :

1) What is the current status of the Flags associated with the (*,G) and (S,G) entries for group `224.1.1.1` ?

R1: (*, G):	**R2: (*, G):**	**R3: (*, G):**
(S, G):	**(S,G):**	**(S,G):**

R4: (*, G):	**R5: (*, G):**	**R6: (*, G):**
(S, G):	**(S, G):**	**(S,G):**

2) Examine the output of "*show ip mdsp sa-cache*" on each of the MSDP peers. Which RP is associated with the source entry?

5.5 Timing out of state

1. Stop the multicast Source from sending packets to group 224.1.1.1.

2. Use the *"show ip mroute 224.1.1.1"* command to answer the following questions:

Questions :

1) Are all the routes in prune state?

2) How long does it take for all the routes to disappear from AS 201?

3) Why does state remain in AS 301?

5.6 State with Receiver 2 and source

1. Activate multicast Receiver 2 shown in the configuration diagram so that it "joins" the `224.1.1.1` group.
2. Activate the multicast Source so that it begins sending traffic to group `(*,224.1.1.1).`
3. On each of the routers, examine the multicast state for group `224.1.1.1` using the *"show ip mroute 224.1.1.1"* command and answer the following questions.

Questions :

1) What is the current status of the Flags associated with the (*,G) and (S,G) entries for group `224.1.1.1` ?

R1: (*, G):	**R2: (*, G):**	**R3: (*, G):**
(S, G):	**(S,G):**	**(S,G):**
R4: (*, G):	**R5: (*, G):**	**R6: (*, G):**
(S, G):	**(S, G):**	**(S,G):**

5.7 State changes with Link Failure

1. Shut the serial interface between R6 and R3
2. Activate the multicast Source so that it begins sending traffic to group `(*,224.1.1.1).`
3. On each of the routers, examine the multicast state for group `224.1.1.1` using the *"show ip mroute 224.1.1.1"* command and answer the following questions.

Questions :

1) What is the current status of the Flags associated with the (*,G) and (S,G) entries for group `224.1.1.1` and why?

R1: (*, G):	**R2: (*, G):**	**R3: (*, G):**
(S, G):	**(S,G):**	**(S,G):**
R4: (*, G):	**R5: (*, G):**	**R6: (*, G):**
(S, G):	**(S, G):**	**(S,G):**