BGP – Lab

Peering

MYANMAR NETWORK OPERATORS GROUP

Author: Thein Myint Khine

Version: 0.0

Last Update: September 11, 2024

Peering?

Peering:

is a voluntary interconnection of administratively separate Internet networks for the purpose of exchanging traffic between the "down-stream" users of each network.

Neither party pays the other in association with the exchange of traffic.

Peering

Public peering :

Interconnection utilizing a multi-party shared switch fabric such as an Ethernet switch.

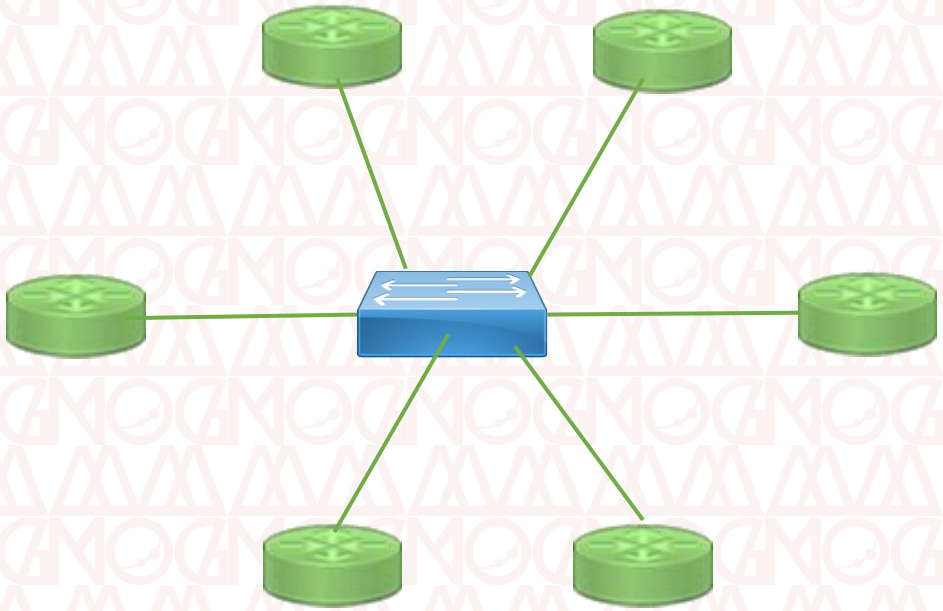

Private peering :

Interconnection utilizing a point-to-point link between two parties.

Peering

Multilateral Peering :

Hub and Spoke Topology, bgp peering session via Route Server, accepting all peering prefixes.

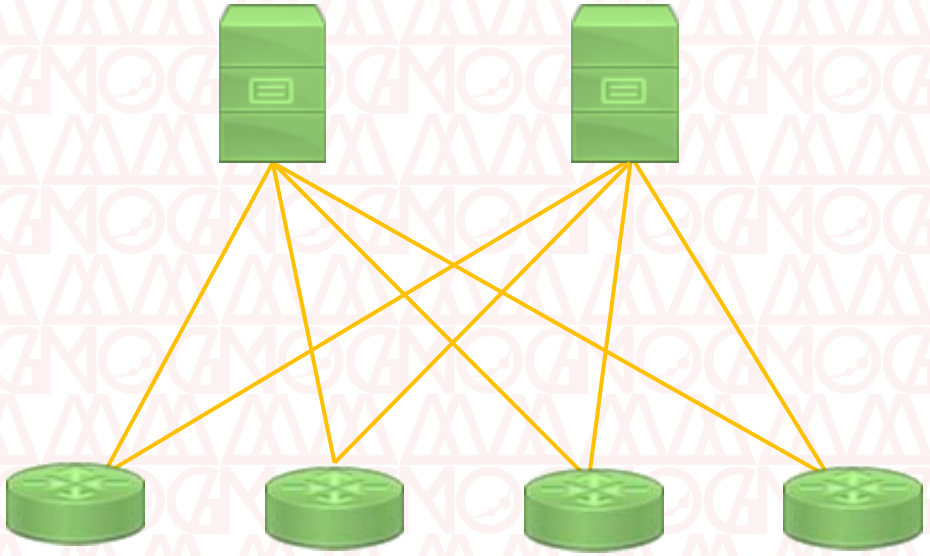

Bilateral peering :

Partial mesh Topology, direct bgp meshing sessions between peers.

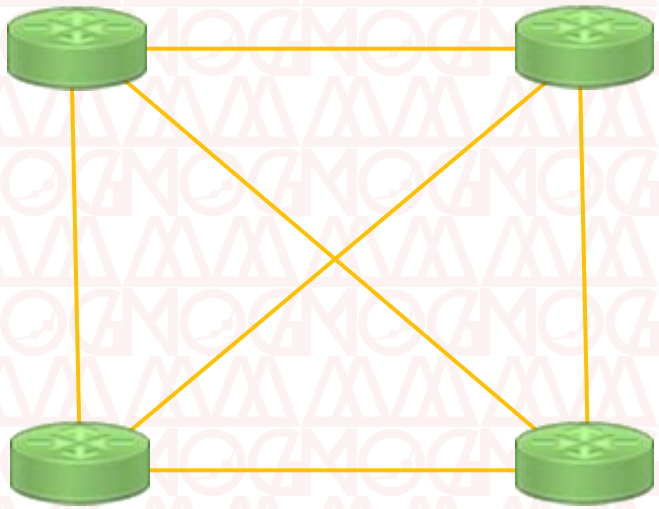

Client Connectivity's

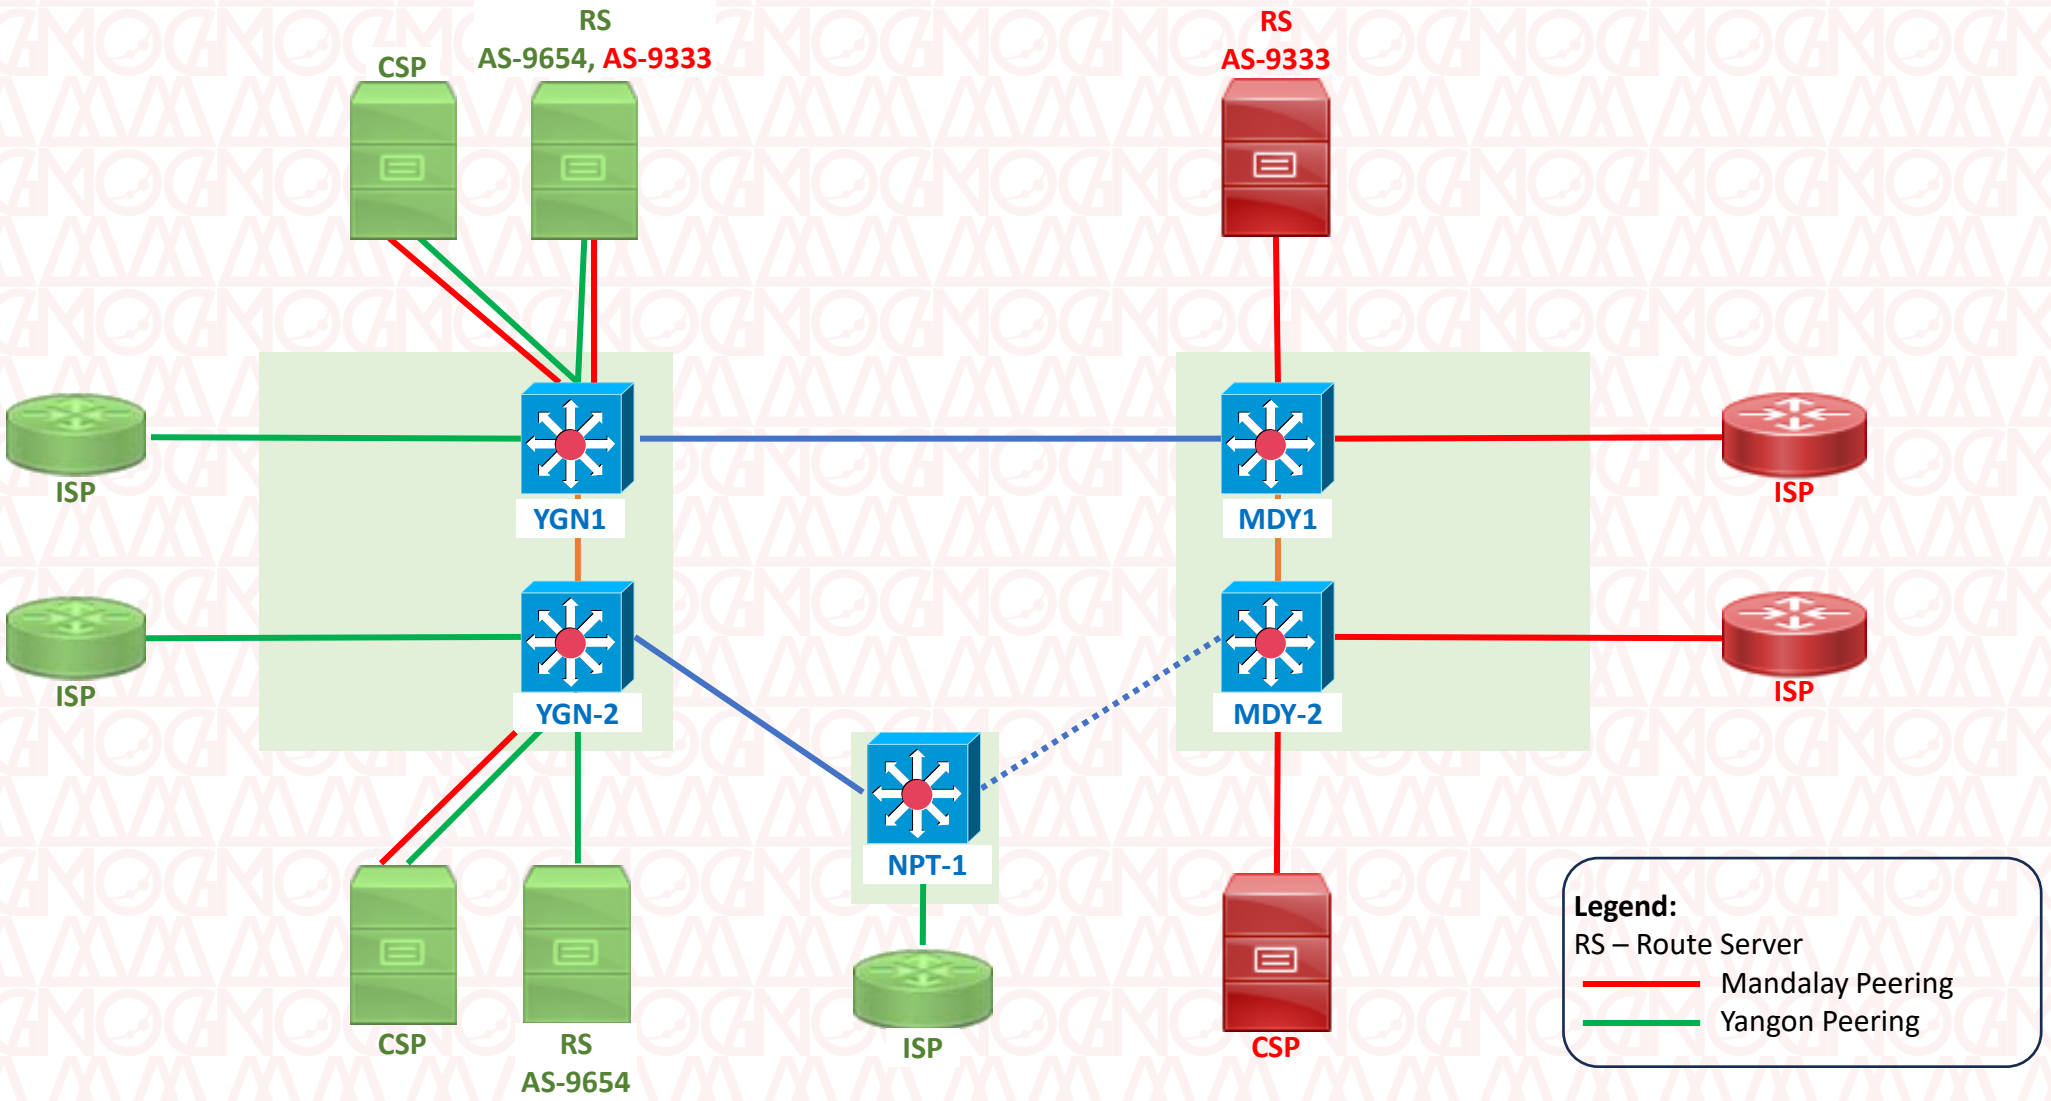

Peering IP & Route Servers for an IXP

peer IP :

It is not Point to Point link like /30 but Peering LAN address shall be /24, mor or less. Same concept for IPv6, it shall be /64.

Route Servers :

Normally, 2 Route Servers are used for redundancy.

Rules :

Do not advertise Peering LAN prefix to other Peers and Internet.

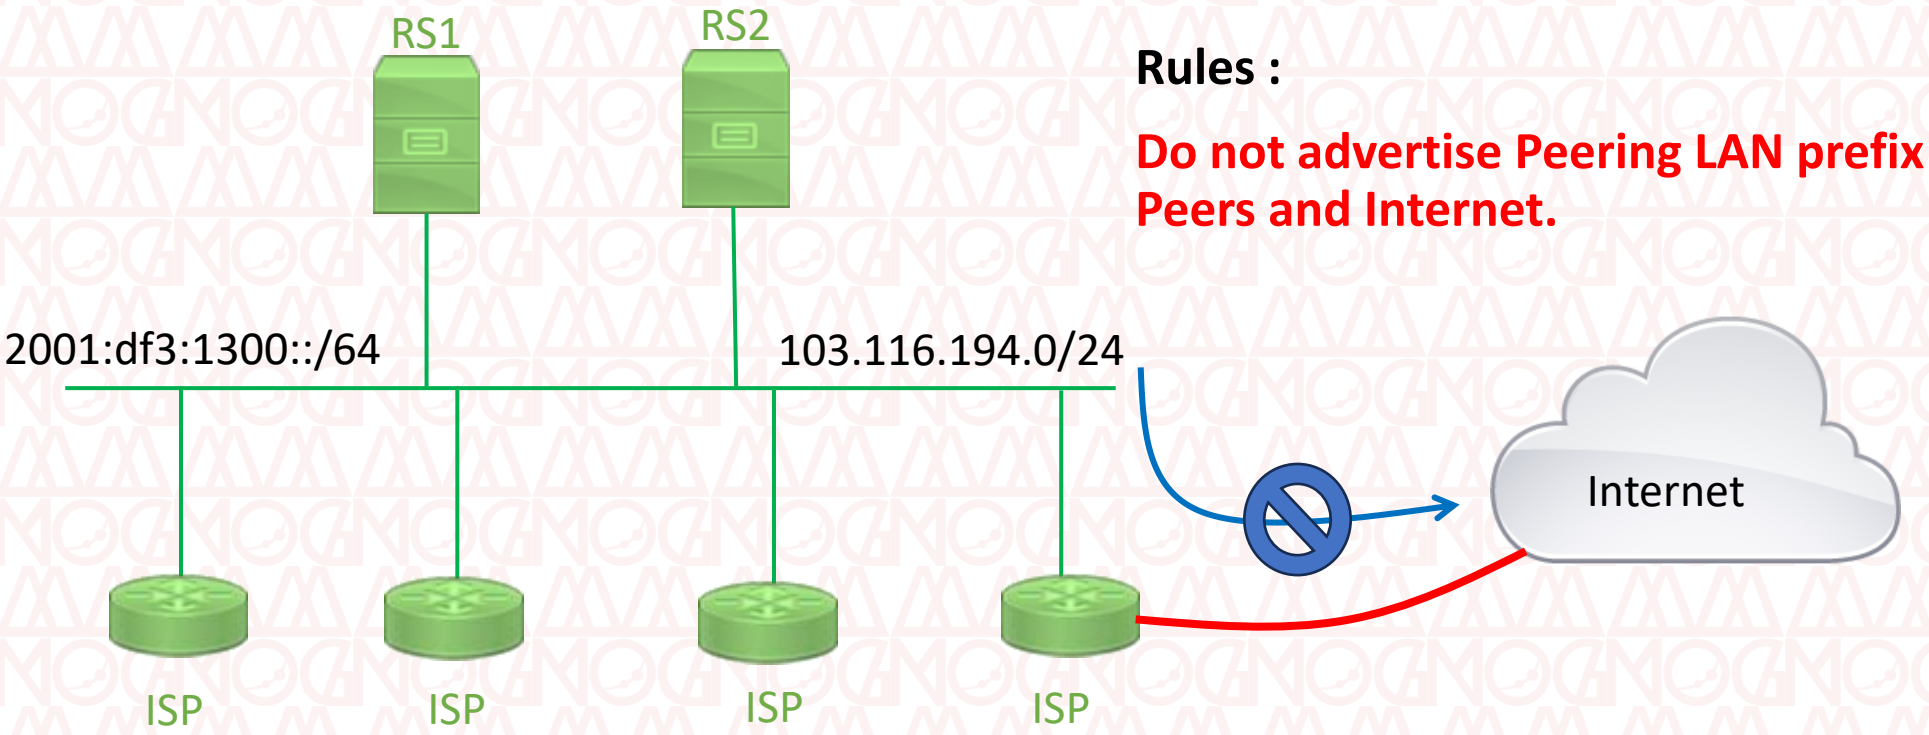

BGP Communities (MMIX)

Most IXPs are supporting BGP communities for route control.

Yangon IXP	
0:9654	Block Announcement of prefixes to all peers
0:(peer-as)	Block Announcement of prefixes to certain peer
9654:(peer-as)	Advertise to certain peer
40027:40000	Advertise to Netflix
9654:20940	Advertise to Akamai
9654:54994	Advertise to Wangsu
9654:11344	Advertise to GGC

Mandalay IXP	
0:9333	Block Announcement of prefixes to all peers
0:(peer-as)	Block Announcement of prefixes to certain peer
9333:(peer-as)	Advertise to Certain peer
40027:40000	Advertise to Netflix
9333:20940	Advertise to Akamai
9333:54994	Advertise to Wangsu

More Learning

"**IXP Peering Tutorial**" of NOG6 – Yangon

On the MMNOG portal
At the "MMNOG" Youtube Channel.

LAB-IXP Peering

X = Group ID,1,2,3,
R = Router ID,1,2,3,

Lo0: X0.0.0.R/32
Lo16: X0.R.16.1/24
Lo17: X0.R.17.1/24

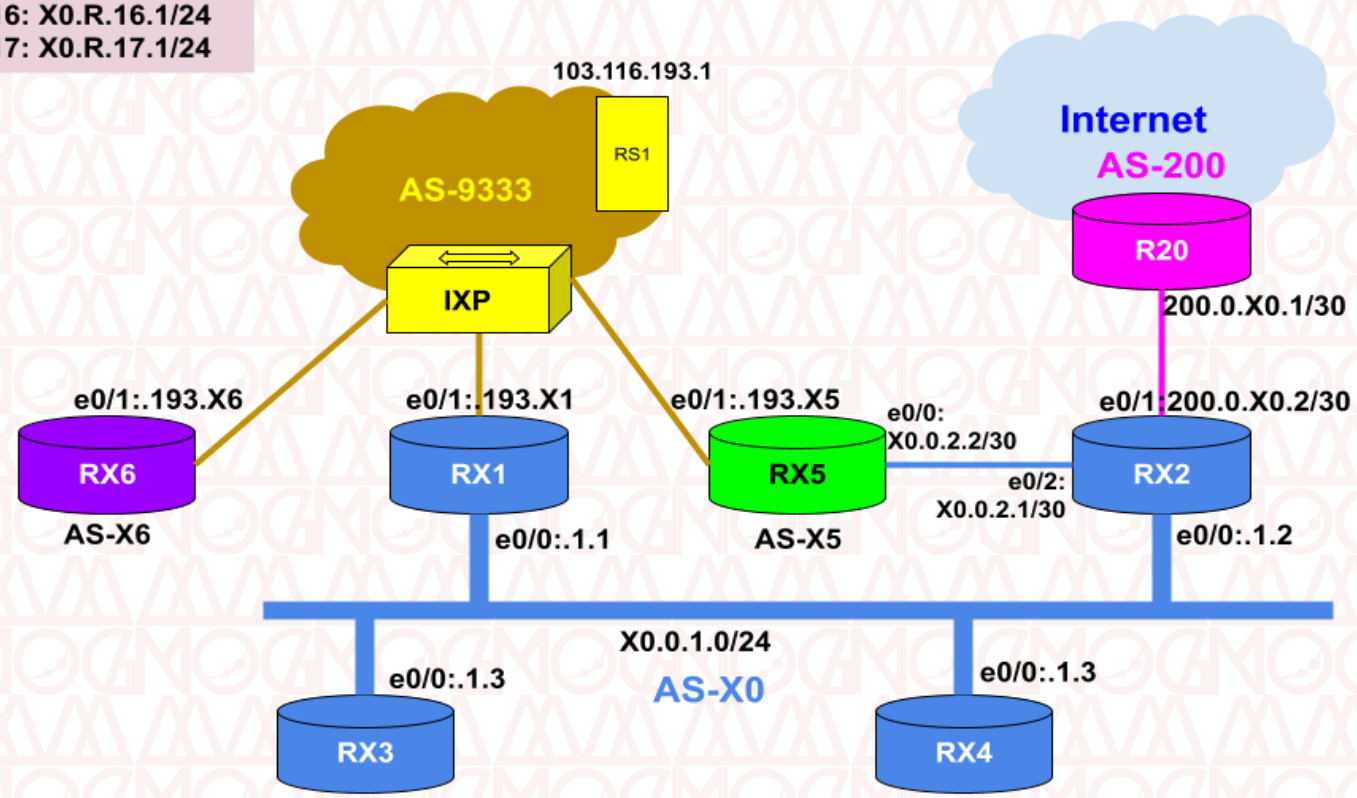

Exercise

Check the preconfigure setting.

1. Check configuration and ping test link IP addresses. Check also OSPF routes, BGP sessions and BGP routes.
2. At RX1 and RX6, setup bgp peer sessions with RS1 of MMIX.

```
! at Rx1 & Rx6 setup BGP session with MMIX's Route Server 1
router bgp 10
 no bgp enforce-first-as
 neighbor 103.116.193.1 remote-as 9333
!Check routing table
!will notice direct peering routes without MMIX RS's AS number.
!will notice downstream of peers.
 *>  30.4.16.0/24     103.116.193.31          0             0 30 i
 *>  30.4.17.0/24     103.116.193.31          0             0 30 i
 *>  35.5.16.0/24     103.116.193.31          0             0 30 35 i
 *>  35.5.17.0/24     103.116.193.31          0             0 30 35 i
```

3. At RX5, setup bgp peer sessions with RS1 of MMIX.

```
! At Rx5, setup bgp session with RS1 of MMIX
! will notice no more downstream as downstream are directly peering.
! MMIX is supporting the following BGP communities
```

0:9333 - Block announcement of prefixes to all Peers.
0:<peer-as> - Block announcement of prefixes to certain peer.
9333:9333 - Advertise to all peers.
9333:<peer-as> - Advertise to certain peer.

4. At Rx2 setup bgp session with its upstream AS200.

```
! setup BGP session, learn routing tables at Rx1, Rx2, Rx3, Rx4,and Rx5.
router bgp 10
 neighbor 200.0.10.1 remote-as 200
! learn routing tables again.  Will notice more routes
```

5. Higher preference for peering traffic.

```
! learn routing tables before changes
! set preference value into 110 for the IX prefixes
Rx1(config)#
route-map RM-MMIX-IN permit 10
 set local-preference 110
router bgp 10
 neighbor 103.116.193.1 route-map RM-MMIX-IN in
! learn bgp tables at Rx1,2,3,4
R11#sh ip bgp
     Network          Next Hop            Metric LocPrf Weight Path
 *>i 0.0.0.0          10.0.0.2                 0    100      0 200 i
 *>  10.1.16.0/24     0.0.0.0                  0         32768 i
 *>i 10.4.16.0/24     10.0.0.4                 0    100      0 i
 *>i 10.4.17.0/24     10.0.0.4                 0    100      0 i
 *>  15.5.16.0/24     103.116.193.15           0    110      0 15 i
 *>  16.6.17.0/24     103.116.193.16           0    110      0 16 i
```

```
R12#sh ip bgp regexp _15$
     Network          Next Hop            Metric LocPrf Weight Path
 *>i 15.5.16.0/24     10.0.0.1                 0    110      0 15 i
 *                    10.0.2.2                 0             0 15 i
 *>i 15.5.17.0/24     10.0.0.1                 0    110      0 15 i
 *                    10.0.2.2                 0             0 15 i
```

6. At Rx2 adjust the downstream path to Rx5 to go via private peer link.

```
! create route-map for weight and bind to inbound of R15
R12(config)#
route-map RM-15-IN permit 10
 set weight 100
router bgp 10
 neighbor 10.0.2.2 route-map RM-15-IN in
! check routing table after bgp update
R12#sh ip bgp regexp _15$
     Network          Next Hop            Metric LocPrf Weight Path
 * i 15.5.16.0/24     10.0.0.1                 0    110      0 15 i
 *>                   10.0.2.2                 0            100 15 i
 * i 15.5.17.0/24     10.0.0.1                 0    110      0 15 i
 *>                   10.0.2.2                 0            100 15 i
```

7. Control peer partner using BGP communities. Rx5 does not like to peer with Rx1. From RX5, do not announce owned prefixes to RX1 using BGP Community supported by IXP, do not accept prefixes of RX1.

```
! before changes, learn bgp tables at Rx5 and Rx1
R11#sh ip bgp regexp _15$
     Network          Next Hop            Metric LocPrf Weight Path
 *>  15.5.16.0/24     103.116.193.15           0    110      0 15 i
 *>  15.5.17.0/24     103.116.193.15           0    110      0 15 i
R15#sh ip bgp regexp _10$
     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.3.16.0/24     103.116.193.11           0    110      0 10 i
 *                    10.0.2.1                                0 10 i
 *>  10.3.17.0/24     103.116.193.11           0    110      0 10 i
 *                    10.0.2.1                                0 10 i
 *>  10.4.16.0/24     103.116.193.11           0    110      0 10 i
 *                    10.0.2.1                                0 10 i
 *>  10.4.17.0/24     103.116.193.11           0    110      0 10 i
 *                    10.0.2.1                                0 10 i
```

```
! Create outbound route-map at Rx5 not to announce to Rx1
ip bgp-community new-format
!
R15#route-map RM-MMIX-OUT permit 10
 set community 0:10
router bgp 15
 neighbor 103.116.193.1 send-community
 neighbor 103.116.193.1 route-map RM-MMIX-out out

! check routes update at Rx1
R11#sh ip bgp regexp _15$
     Network          Next Hop            Metric LocPrf Weight Path
 *>i 15.5.16.0/24     10.0.0.2                 0    100      0 15 i
 *>i 15.5.17.0/24     10.0.0.2                 0    100      0 15 i
```

```
! At Rx5, block prefixes of ASX0 not to accept via IXP
! at Rx5, check existing routes of Rx1, will notice routes via IXP
R15#sh ip bgp regexp _10$
     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.3.16.0/24     103.116.193.11           0    110      0 10 i
 *>  10.3.17.0/24     103.116.193.11           0    110      0 10 i
 *>  10.4.16.0/24     103.116.193.11           0    110      0 10 i
 *>  10.4.17.0/24     103.116.193.11           0    110      0 10 i
```

```
! At Rx5, check existing inbound route-map
R15(config)#
route-map RM-MMIX-IN permit 10
 set local-preference 110
! modify route-map
ip as-path access-list 10 permit ^10_
route-map RM-MMIX-IN deny 5
 match as-path 10
! check back route table after bgp update
R15#sh ip bgp regexp _10$
     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.3.16.0/24     10.0.2.1                               0 10 i
 *>  10.3.17.0/24     10.0.2.1                               0 10 i
 *>  10.4.16.0/24     10.0.2.1                               0 10 i
 *>  10.4.17.0/24     10.0.2.1                               0 10 i
! will notice route via local path, not via IX path
```

← lower seq number

! Do the same config at Rx1 not to community with Rx5 via IX

8. Setup Peering partner. Rx5 likes to peer with Rx6. At Rx5 advertise all prefixes using bgp community value support by MMIX, so Rx5 prefixes shall be announced to Rx6.

```
! At Rx5, tag bgp community 9333:x6 to IX peering session.
! So, Rx5 prefixes will arrive to Rx6.
! create route-map.  If there is an existing route-map, modify it.
route-map RM-MMIX-OUT permit 10
 set community 9333:16
router bgp 15
 neighbor 103.116.193.1 send-community
 neighbor 103.116.193.1 route-map RM-MMIX-OUT out
! check the routing table at Rx6.
! You will receive prefixes from neighbor Rx1.
R16#sh ip bgp
     Network          Next Hop            Metric LocPrf Weight Path
 *>  15.5.16.0/24     103.116.193.15           0             0 15 i
 *>  15.5.17.0/24     103.116.193.15           0             0 15 i
 *>  16.6.16.0/24     0.0.0.0                  0         32768 i
 *>  16.6.17.0/24     0.0.0.0                  0         32768 i
```

```
! check bgp communities
R16#sh ip bgp 15.5.16.0/24
BGP routing table entry for 15.5.16.0/24, version 84
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 1
  15
    103.116.193.15 from 103.116.193.1 (103.116.193.1)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: 0:10 9333:16
      rx pathid: 0, tx pathid: 0x0
```

Thank you

Q&A

Thein Myint Khine
theinmyintkhine@mm-ix.net

www.mmnog.net.mm

event@mm-ix.net