BGP – Lab
Network Security

MYANMAR NETWORK OPERATORS GROUP

Author: Thein Myint Khine

Version: 0.0

Last Update: September 11, 2024

RTBH at an IXP

MMIX Supports Remote Trigger, black-hole filtering (RTBH)

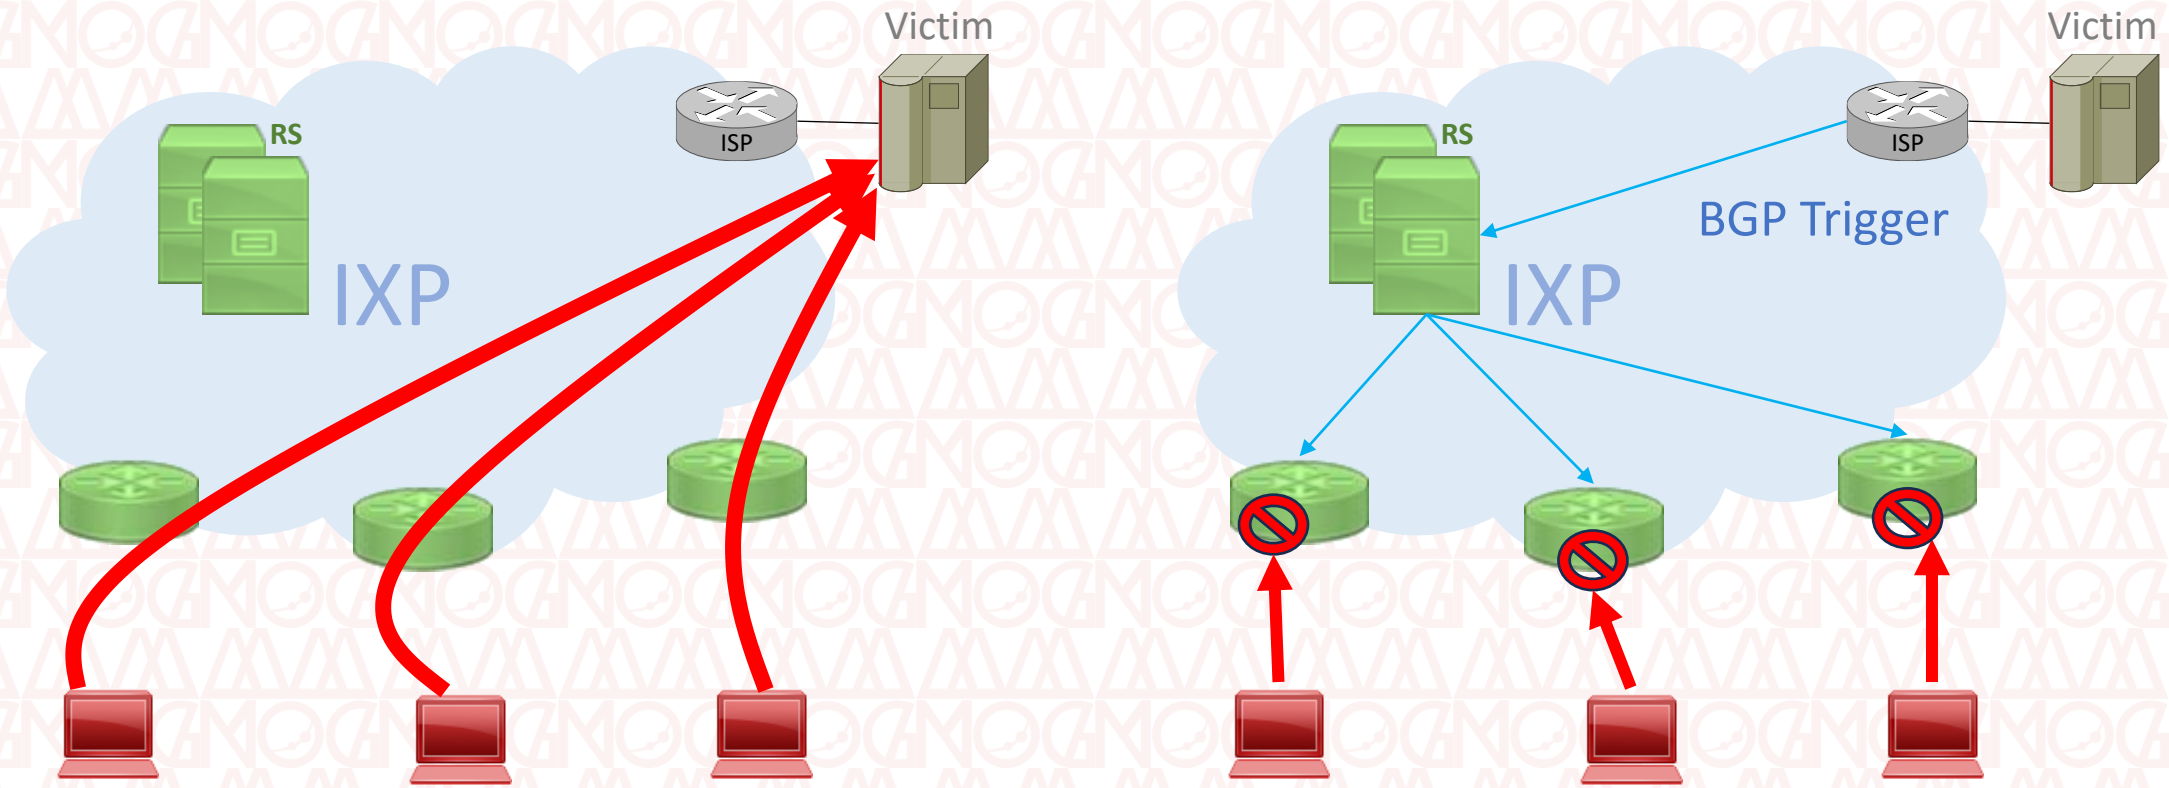

LAB - BGP - Network Security

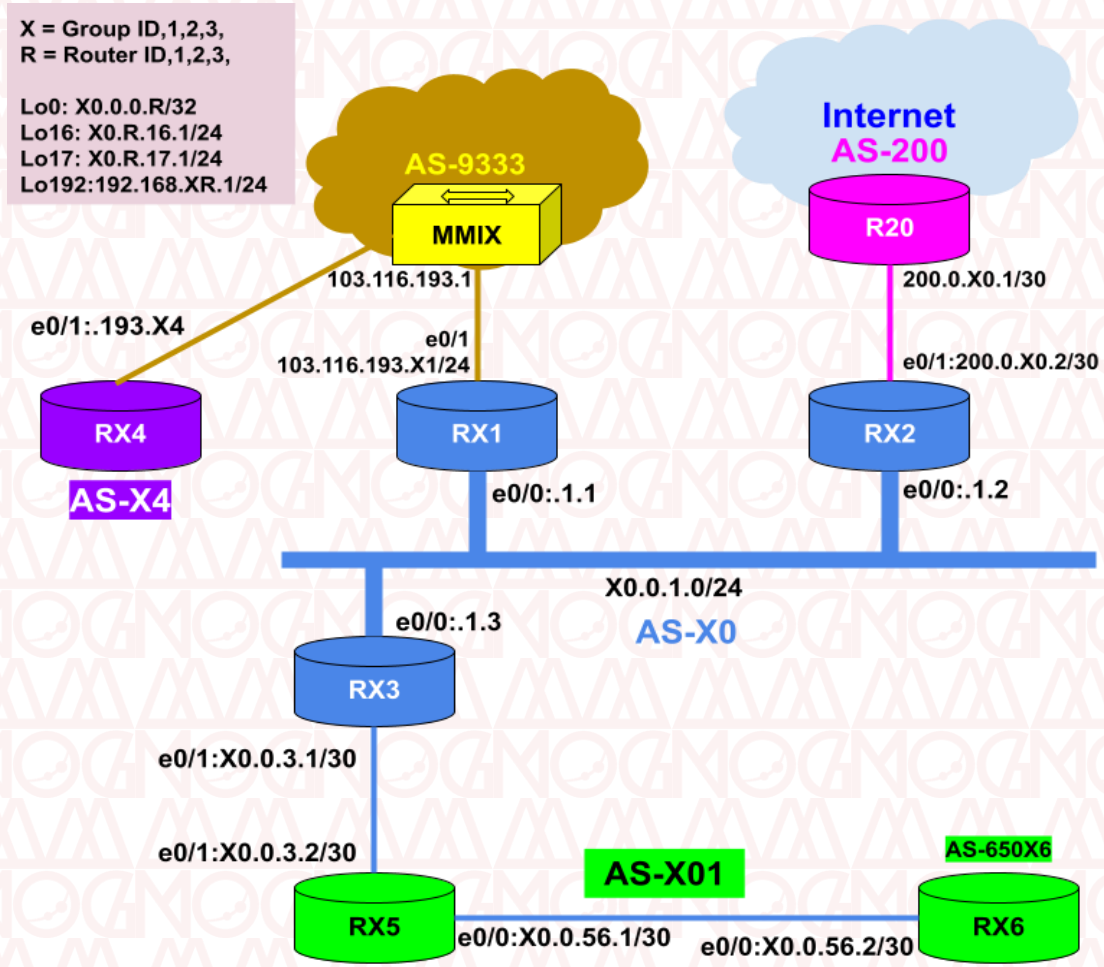

RTBH - Remotely Triggered Black Hole

MMIX is supporting RTBH features with the following parameters.

IXP: Mandalay
BGP Community: 9333:66
Next hop address: 103.116.193.66
trigger IP size: /32

Exercise

Configure to meet the following requirements.

1. Check configuration and ping test link IP addresses. Check also OSPF routes, BGP sessions and BGP routes.
2. Activate RTBH at Peer Router Rx1.

```
!## If there is no inbound filtering on peering session with MMIX,
! just need to configure a static null route and stop retransmission

Rx1(config)#
ip route 103.116.193.66 255.255.255.255 null 0
interface Null0
 no ip unreachable
```

```
!## If there is existing filtering, modify route-map as follow:-
!Create Prefix, community list and inbound route-map
ip prefix-list PRF-MMIX-HOST permit 0.0.0.0/0 ge 32
ip community-list standard CM-MMIX-RTBH permit 9333:66
!
route-map RM-MMIX-IN permit 10
 match ip address prefix-list PRF-MMIX-HOST
 match community CM-MMIX-RTBH
route-map RM-MMIX-IN permit 15
 ! existing Setting
!Bind to BGP session
router bgp X0
 neighbor 103.116.193.1 route-map RM-MMIX-IN in
```

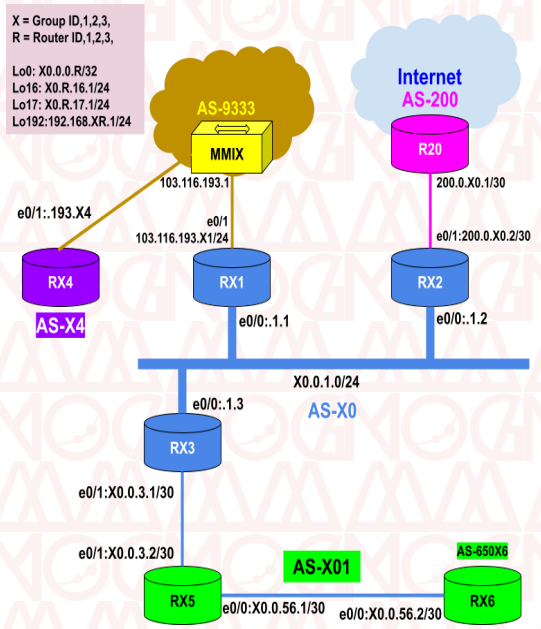

Q3: Setup RTBH Trigger

3. Considering Rx3, IP X0.3.16.1 is under attack by some MMIX networks, trigger RTBH using the features supported by MMIX.

```
! Check reachability to victim IP X0.3.16.1 from Rx4
R14#ping 10.3.16.1 source lo16
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 10.3.16.1, timeout is 2 seconds:
Packet sent with a source address of 14.4.16.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

! At Rx1, setup Trigger
....
```

```
!! Option 1
! create prefix list and null route
ip prefix-list PRF-RTBH seq 5 permit 10.3.16.1/32
ip route 10.3.16.1 255.255.255.255 Null0
! create outbound route-map
route-map RM-MMIX-OUT permit 10
 match ip address prefix-list PRF-RTBH
 set community 9333:66 additive
route-map RM-MMIX-OUT permit 65535
! bind to BGP and set other setting
router bgp 10
 network 10.3.16.1 mask 255.255.255.255
 neighbor 103.116.193.1 send-community
 neighbor 103.116.193.1 route-map RM-MMIX-OUT out
```

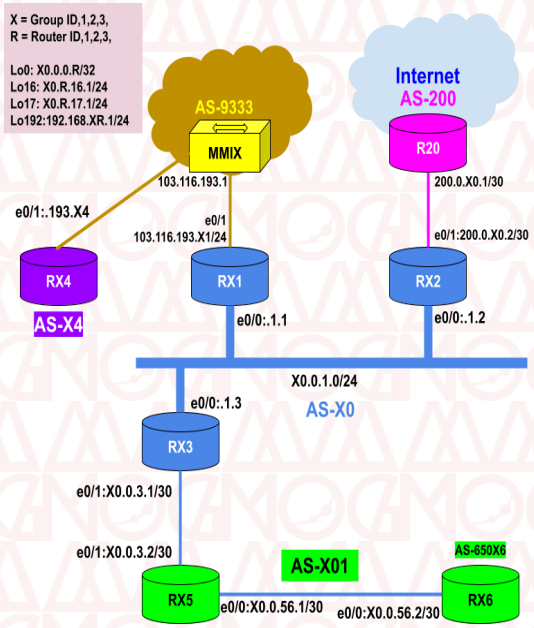

Q3: Setup RTBH Trigger

```
!! Option 2
ip route 10.3.16.1 255.255.255.255 Null0 tag 66
route-map RM-RTBH permit 10
 match tag 66
 set local-preference 200
 set community 9333:66
!
router bgp 10
 redistribute static route-map RM-RTBH
 neighbor 103.116.193.1 send-community
```

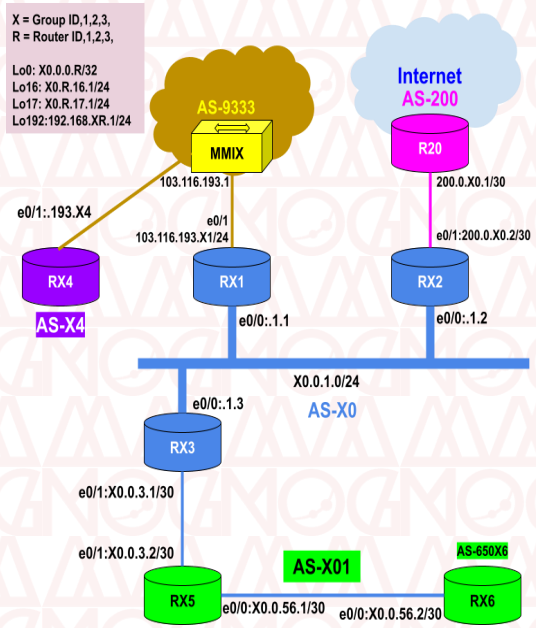

Q3: Setup RTBH Trigger

```
! check result at Rx4 and test reachability
R16#sh ip bgp reg _10$
     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.3.16.0/24     103.116.193.11           0             0 10 i
 *>  10.3.16.1/32     103.116.193.66           0             0 10 i
 *>  10.3.17.0/24     103.116.193.11           0             0 10 i
 *>  10.4.16.0/24     103.116.193.11           0             0 10 i
R14#sh ip bgp 10.3.16.1
BGP routing table entry for 10.3.16.1/32, version 68
Paths: (1 available, best #1, table default)
  Not advertised to any peer
  Refresh Epoch 1
  10
    103.116.193.66 from 103.116.193.1 (103.116.193.1)
      Origin IGP, metric 0, localpref 100, valid, external, best
      Community: 9333:66
      rx pathid: 0, tx pathid: 0x0
```

Q4: Inbound Filter

4. For ASX0 Blue ISP, create inbound filters for
    a. Private IP addresses
    b. Longer prefix length
    c. Default gateway

```
! Check existing routes.
! You will notice private IPs, longer prefixes and default gateway
R12#sh ip bgp
BGP table version is 104, local router ID is 10.0.0.2
     Network          Next Hop            Metric LocPrf Weight Path
 *>  0.0.0.0          200.0.10.1                             0 200 i
 r>i 10.0.0.1/32      10.0.0.1                 0    100      0 ?
 *>  10.0.0.2/32      0.0.0.0                  0         32768 ?
 *   10.0.2.0/30      10.0.2.2                 0             0 65015 ?
 *>                   0.0.0.0                  0         32768 ?
 *>i 10.0.3.0/30      10.0.0.3                 0    100      0 ?
 *>i 10.4.17.0/24     10.0.0.4                 0    100      0 i
 *>  15.0.0.5/32      10.0.2.2                 0             0 65015 ?
 *>  15.5.16.0/24     10.0.2.2                 0             0 65015 i
 *   16.6.16.0/24     200.0.10.1                             0 200 40 16 i
```

Q4: Inbound Filter

```
! Check routing tables before changes
R12#sh ip bgp sum
Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  St/PfxRc
10.0.0.1        4    10      24      20       89    0    0 00:05:30        27
10.0.0.3        4    10      17      19       89    0    0 00:05:30        15
200.0.10.1      4   200     103      80       89    0    0 00:05:30        67
! you will notice private IPs, longer prefixes and default gateway
! Create Filter list.  This is the list just for this Lab only.
ip prefix-list BOGON deny 192.168.0.0/16 le 32     <- private IP
ip prefix-list BOGON deny 0.0.0.0/0 ge 25 le 32    <- Longer prefixes
ip prefix-list BOGON permit 0.0.0.0/0 le 32        <- default gateway
! bind to bgp sessions
router bgp 10
 neighbor 200.0.10.1 prefix-list BOGON in
 neighbor 200.0.10.1 prefix-list BOGON out
! check received and advertised routes
! will notice there are no longer private IPs and longer prefixes.
```

5. Filter Private ASNs: At Rx2, do not accept Private ASN via AS-200.

```
! Check routing tables before changes
! Wil notice private AS numbers at the AS paths.
! At Rx2, Create Filter list.  This is the list just for this Lab only.
ip as-path access-list 1 deny _650[0-9][0-9]_
ip as-path access-list 1 permit .*
! bind to bgp sessions
router bgp 10
 neighbor 200.0.10.1 filter-list 1 in
! check received routes
! will notice there are no longer private ASN from Global.
```

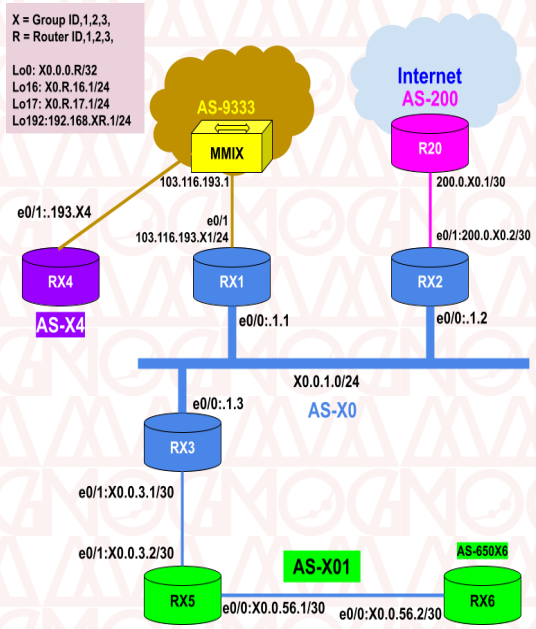

6. Outbound Filter: Just advertised prefixes of your owned and downstream.
    - At Rx1, advertise prefixes of AS-X0 and AS-X5 to AS-9333 (MMIX)

```
! check advertised routes before configuration.
! you will notice other prefixes including MMIX Peering LAN
! Create a Prefix list and bind to BGP.
R11(config)#
ip prefix-list PRF-MINE permit 10.0.0.0/8 ge 23 le 24
ip prefix-list PRF-MINE permit 11.0.0.0/8 ge 23 le 24
router bgp 10
 neighbor 103.116.193.1 prefix-list PRF-MINE out
! check advertised routes. will notice only owned
! prefixes are announced.
```

    - At Rx5, advertise owned prefixes to AS-X0.

```
! At Rx5, configure the same like Rx1
```

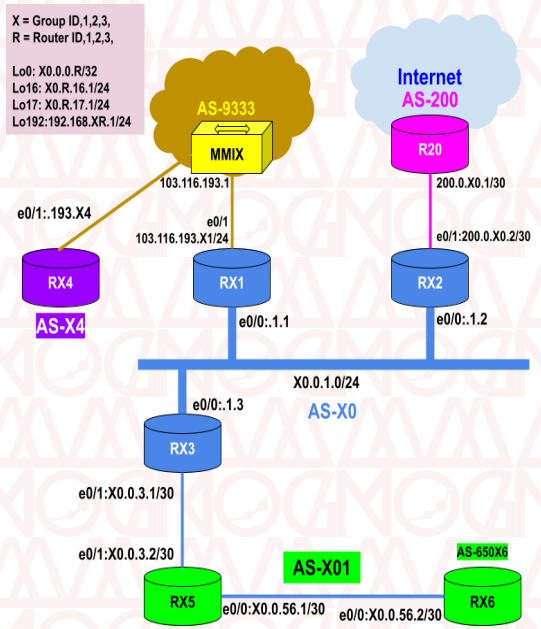

- At Rx2, advertise prefixes of AS-X0 and AS-X5 to AS-200 (Internet)

```
! At Rx2, check routes before changes.
! create as-path list and bind to BGP
R12(config)#
ip as-path access-list 2 permit ^$
ip as-path access-list 2 permit ^101$
!
router bgp 10
 neighbor 200.0.10.1 filter-list 2 out
! Here is the result
R12#sh ip bgp nei 200.0.10.1 adver
     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.0.1.0/24      0.0.0.0                  0         32768 ?
 *>i 10.3.16.0/24     10.0.0.3                 0    100      0 i
 *>i 10.3.17.0/24     10.0.0.3                 0    100      0 i
 *>i 11.5.16.0/24     10.0.0.3                 0    100      0 101 i
 *>i 11.5.17.0/24     10.0.0.3                 0    100      0 101 I
 *>i 103.116.193.0/24 10.0.0.1                 0    100      0 ?
```

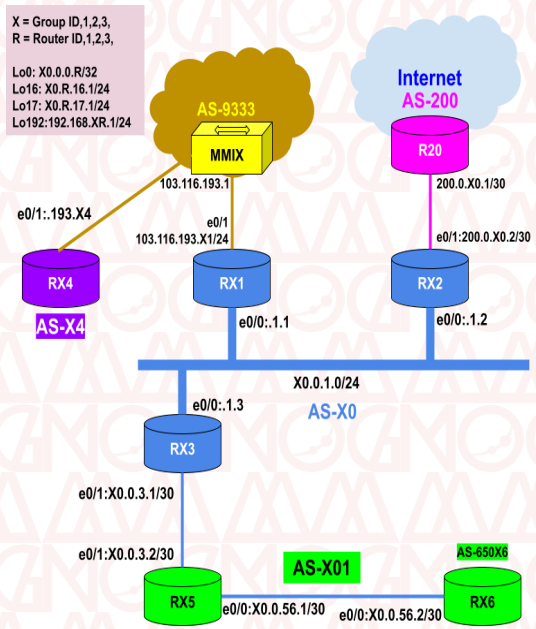

```
! You will still notice not your owned IP, MMIX Peering LAN
! Big mistake and fix it.
Rx1(config)#
router bgp 10
 no redistribute connected
! You will also notice Rx6 prefixes are not advertised.
R12#sh ip bgp reg _101_
     Network          Next Hop            Metric LocPrf Weight Path
 *>i 11.0.0.5/32      10.0.0.3                 0    100      0 101 ?
 *>i 11.0.0.6/32      10.0.0.3                 0    100      0 101 65016 ?
 *>i 11.0.56.0/30     10.0.0.3                 0    100      0 101 ?
 *>i 11.5.16.0/24     10.0.0.3                 0    100      0 101 i
 *>i 11.5.17.0/24     10.0.0.3                 0    100      0 101 i
 *>i 11.6.16.0/24     10.0.0.3                 0    100      0 101 65016 i
 *>i 11.6.17.0/24     10.0.0.3                 0    100      0 101 65016 i
 *>i 192.168.15.0     10.0.0.3                 0    100      0 101 ?
 *>i 192.168.16.0     10.0.0.3                 0    100      0 101 65016 ?
```

! Private AS….

7. Remove Private AS from AS-X01

```
! From Rx3, check routes before changes
R13#sh ip bgp regexp ^101_
…
! will notice private ASN
! remove private AS from ASX5.
R15#
router bgp 101
 neighbor 10.0.3.1 remove-private-as
! check received routes from Rx3
! check advertised routes from Rx2 to R20
R12#sh ip bgp nei 200.0.10.1 advertised-routes
     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.0.1.0/24      0.0.0.0                  0         32768 ?
 *>i 10.3.16.0/24     10.0.0.3                 0    100      0 i
 *>i 10.3.17.0/24     10.0.0.3                 0    100      0 i
 *>i 11.5.16.0/24     10.0.0.3                 0    100      0 101 i
 *>i 11.5.17.0/24     10.0.0.3                 0    100      0 101 i
 *>i 11.6.16.0/24     10.0.0.3                 0    100      0 101 i
 *>i 11.6.17.0/24     10.0.0.3                 0    100      0 101 i
```

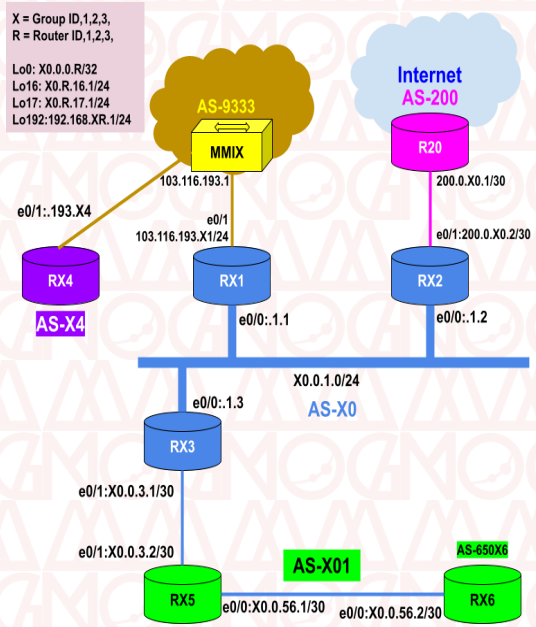

8. No default-GW at Peer Router. Filter default gateway at Rx1.

```
! create prefix list and bind to BGP
R11(config)#
ip prefix-list PRF-INSIDE-IN seq 5 deny 0.0.0.0/0
ip prefix-list PRF-INSIDE-IN seq 10 permit 0.0.0.0/0 le 32
!
router bgp 10
 Neighbor IBGP prefix-list PRF-INSIDE-IN in
! check the routing table.  You will notice no more default-GW.
```

9. Don’t receive owned and downstream prefixes from Global.
    - At Rx2, filter prefixes of AS-X0 and AS-X01 received from global.

```
! Create inbound route-map and bind to bgp
Rx2(config)#
ip prefix-list PRF-MINE permit 10.0.0.0/8 ge 23 le 24
ip prefix-list PRF-MINE permit 11.0.0.0/8 ge 23 le 24
!
route-map RM-AS200-IN deny 10
 match ip address prefix-list PRF-MINE
Route-map RM-AS200-IN permit 65535
!
router bgp 10
 neighbor 200.0.10.1 remote-map RM-AS200 in
```

10. Null route for unused IP addresses.

```
! Add null routes for private IP addresses at full route routers.
ip route 10.0.0.0 255.0.0.0 null 0
ip route 172.16.0.0 255.240.0.0 null 0
ip route 192.168.0.0 255.255.0.0 null 0

! Also null route for unused, owned prefixes.
ip route X0.0.0.0 255.0.0.0 null 0

! Don't play on downstream prefixes.
! They are not yours.
```

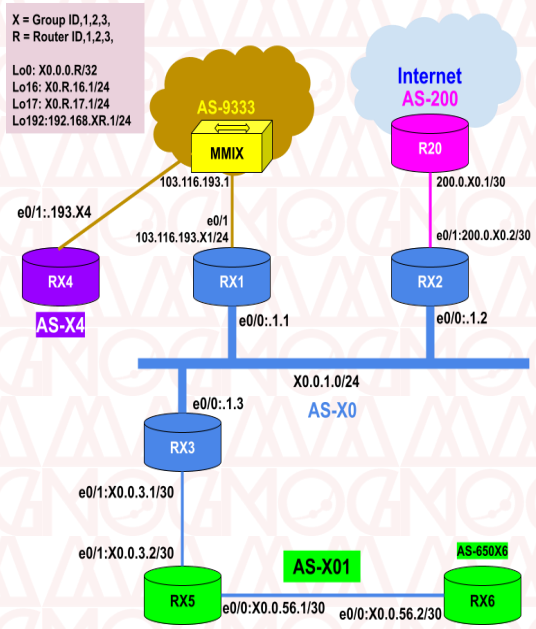

Learn More:

RPF – Reverse Path Forwarding

ROA – Route Origin Authorization

RPKI – Resource Public Key Infrastructure

RADB – The Internet Routing Registry

MANRS – Mutually Agreed Norms for Routing Security

Apnic – Asia Pacific Network Information Centre

Team CYMRU – Free Security Solutions and References

Thank you

Q&A

Thein Myint Khine
theinmyintkhine@mm-ix.net

www.mmnog.net.mm

event@mm-ix.net