
Networkforyou
Welcome
To
Network for you
BGP
Networkforyou

BGP Border Gateway Protocol:

- BGP stand for Border Gateway Protocol.
- BGP Protocols uses TCP Port Number 179.
- **BGP is use to advertise internal server on internet.**
- **BGP is only which can manipulate incoming and outgoing traffic.**
- BGP is also a Dynamic Routing Protocol.
- BGP is Exterior Gateway Protocol (EGP).
- BGP is called Protocol of the Internet.
- **BGP is loop free protocol.**
- BGP is called an application layer protocol.
- BGP is also called Policy-based routing protocol.
- BGP is an AS-by-AS dynamic routing protocol.
- BGP is open standard routing protocol.
- BGP is mainly used for Scalability and Reliability.
- BGP uses the concepts of Autonomous Systems (AS).
- Routers running BGP are known as BGP speakers.
- BGP neighbors are known as the BGP peers.
- BGP supports FLSM, VLSM, and CIDR.
- BGP support auto and manual summarization.
- BGP Updates are incremental and triggered.
- BGP Updates are sent as unicast to manually defined neighbors.
- BGP has two flavors eBGP and iBGP.
- External BGP: If peers are in the different AS called External BGP (EBGP).
- Internal BGP: If peers are in the same AS called internal BGP (iBGP).
- BGP Administrative Distance is 20 for External Updates
- BGP Administrative Distance is 200 for Internal Updates.
- Neighbor relationships for eBGP and iBGP are slightly different.
- IBGP the neighbors no need to be connected directly.

- EBGP the neighbors need to be connected directly.
- BGP has no auto discovery mechanism peers to be set manually.
- BGP Protocol is used to exchange Internet routes.

AS (Autonomous System):

- An AS is a collection of networks under a single administrative domain
- In the world of BGP, each routing domain is known as an Autonomous System, or AS.
- Autonomous System Number
- What is an ASN or AS? An Autonomous System Number (AS number or just ASN) is a special number assigned by IANA used primarily with Border Gateway Protocol which uniquely identifies a network under a single technical administration that has a unique routing policy, or is multi-homed to the public internet.
- Like IP addresses, Autonomous System Numbers have to be unique on the Internet.
- The Internet is nothing more but a bunch of AS those are connected to each other.
- Within an Autonomous System, use an IGP like OSPF, RIP, ISIS or EIGRP.
- Between different Autonomous Systems, use an External Gateway Protocol.
- BGP uses the Autonomous System (AS) number for its loop prevention mechanism.
- Autonomous System numbers are 16-Bit or 2-Octed Autonomous System numbers.
- There is total 65535 numbers of Autonomous System (AS) to choose from.
- An extension has been created that supports 32-Bit or 4-Octed AS numbers.
- It means total about 4294967296 Autonomous System (AS) numbers to choose from.

There is range of public and private AS numbers as given below:

Number	Bits	Description	Reference
0	16	Reserved for RPKI unallocated space invalidation[12]	RFC6483, RFC7607
1 - 23455	16	Public ASNs	
23456	16	Reserved for AS Pool Transition	RFC6793
23457 - 64495	16	Public ASNs	
64496 - 64511	16	Reserved for use in documentation/sample code	RFC5398
64512 - 65534	16	Reserved for private use	RFC1930, RFC6996
65535	16	Reserved	RFC7300
65536 - 65551	32	Reserved for use in documentation and sample code	RFC4893, RFC5398
65552 - 131071	32	Reserved	
131072 - 4199999999	32	Public 32-bit ASNs	
4200000000 - 4294967294	32	Reserved for private use	RFC6996
4294967295	32	Reserved	RFC7300

BGP Flavors:

There are two flavors of BGP Internal BGP and External BGP.

IBGP (Interior Border Gateway Protocol):

- If the peers are in the same AS called Internal BGP (iBGP).
- Internal BGP (IBGP) is between same Autonomous System Number.
- By default, Internal BGP (IBGP) peers are set with TTL value = 255
- Internal BGP (IBGP) routes have Administrative Distance of 200.
- Internal BGP (IBGP) peers do not need to be directly connected.

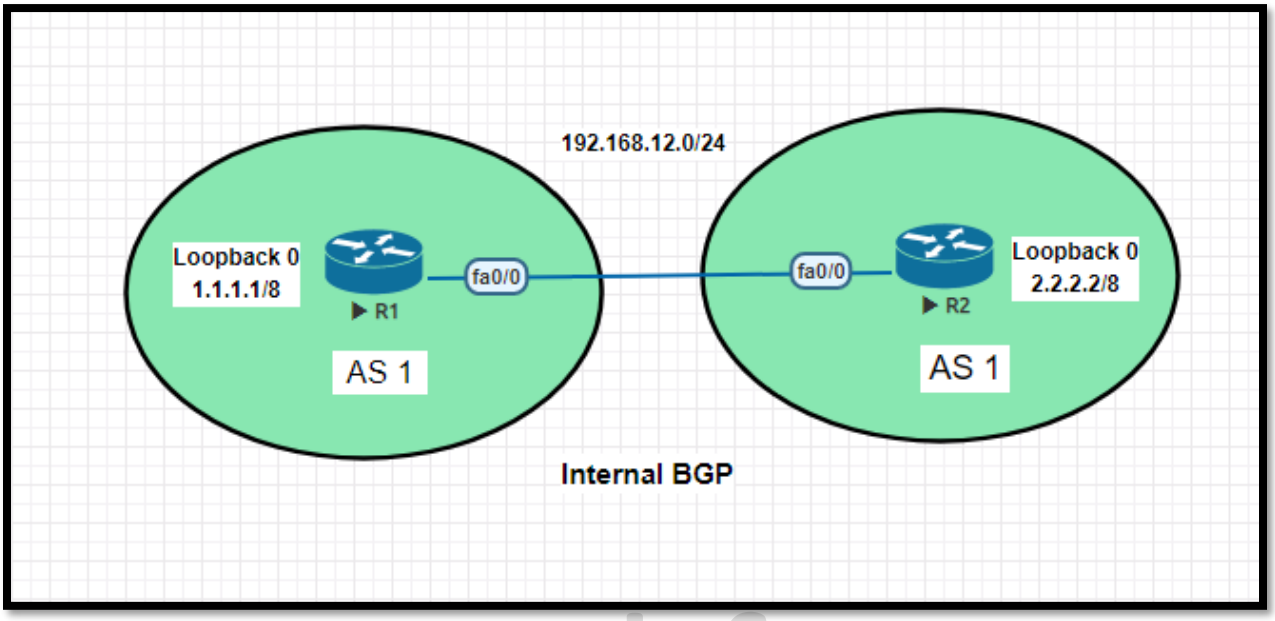

R1 Configuration	R2 Configuration
en config t hostname R1 int f0/0 ip add 192.168.12.1 255.255.255.0 no sh int lo 0 ip add 1.1.1.1 255.0.0.0 router bgp 1 neighbor 192.168.12.2 remote-as 1 network 1.0.0.0 mask 255.0.0.0	en config t hostname R2 int f0/0 ip add 192.168.12.2 255.255.255.0 no sh int lo 0 ip add 2.2.2.2 255.0.0.0 router bgp 1 neighbor 192.168.12.1 remote-as 1 network 2.0.0.0 mask 255.0.0.0

```
R1#sh ip route  bgp
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

B     2.0.0.0/8 [200/0] via 192.168.12.2, 00:02:46
```

```
R1#sh ip bgp neighbors
BGP neighbor is 192.168.12.2,  remote AS 1, internal link
  BGP version 4, remote router ID 2.2.2.2
  BGP state = Established, up for 00:09:20
  Last read 00:00:18, last write 00:00:25, hold time is 180, keepalive interval
is 60 seconds
  Neighbor sessions:
    1 active, is not multisession capable (disabled)
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
    Enhanced Refresh Capability: advertised and received
```

```
  Address tracking is enabled, the RIB does have a route to 192.168.12.2
  Connections established 3; dropped 2
  Last reset 00:02:23, due to Peer closed the session of session 1
  Transport(tcp) path-mtu-discovery is enabled
  Graceful-Restart is disabled
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled
Mininum incoming TTL 0, Outgoing TTL 255
Local host: 192.168.12.1, Local port: 179
Foreign host: 192.168.12.2, Foreign port: 41872
Connection tableid (VRF): 0
```

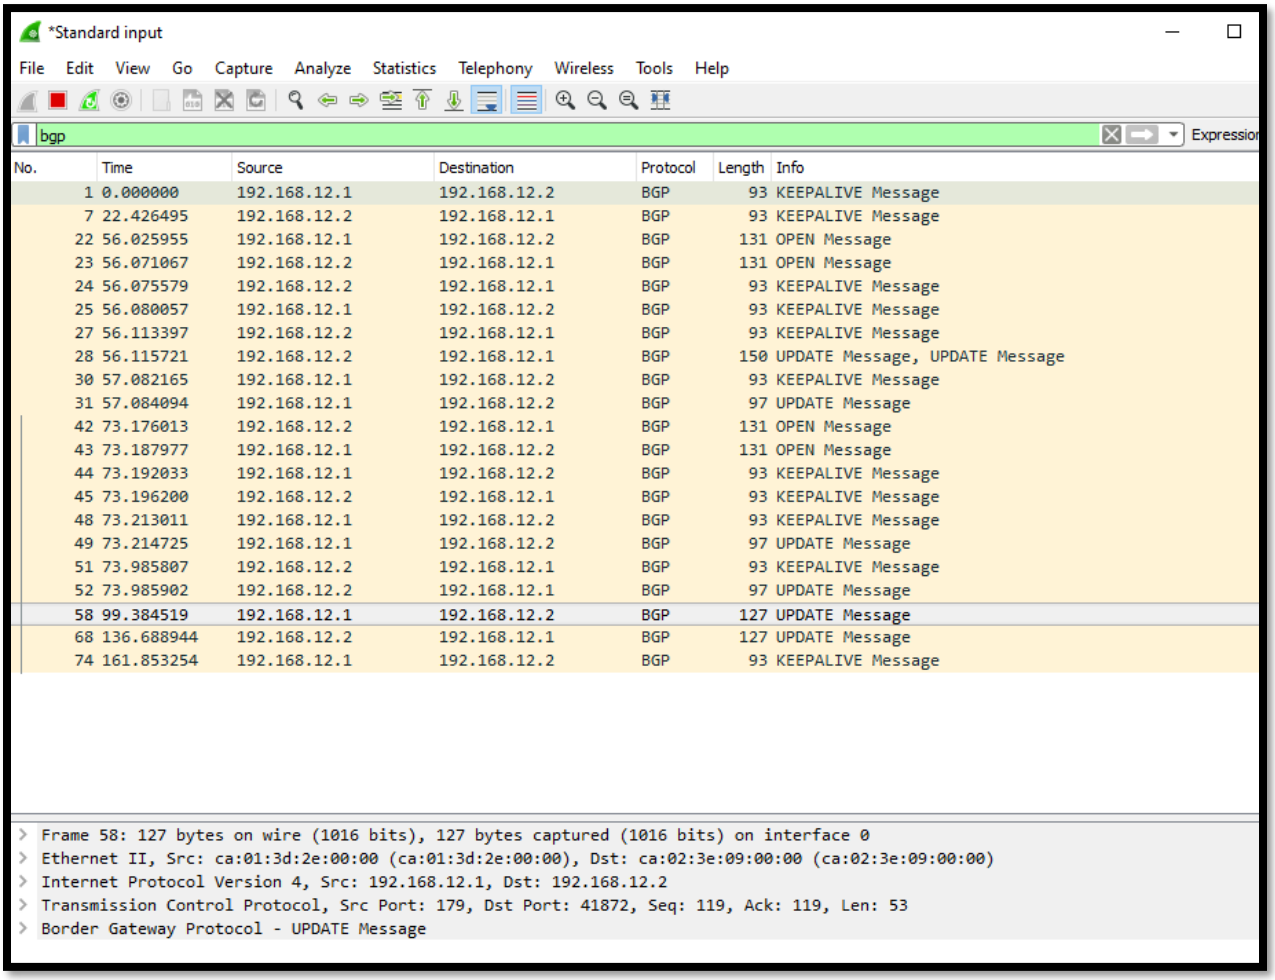

EBGP (Exterior Border Gateway Protocol):

- EBGP peers are set with TTL = 1, means neighbors directly connected.
- External BGP (EBGP) routes have Administrative Distance of 20.
- External BGP (EBGP) the neighbors need to be connected directly.

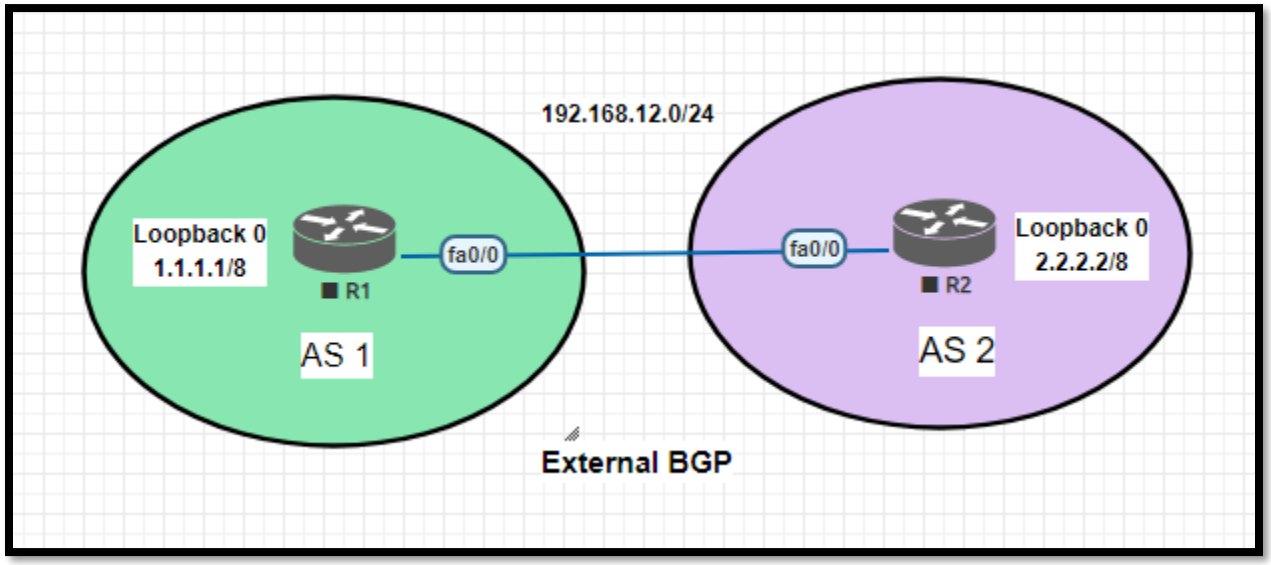

R1 Configuration	R2 Configuration
en config t hostname R1 int f0/0 ip add 192.168.12.1 255.255.255.0 no sh int lo 0 ip add 1.1.1.1 255.0.0.0 router bgp 1 neighbor 192.168.12.2 remote-as 2 network 1.0.0.0 mask 255.0.0.0	en config t hostname R2 int f0/0 ip add 192.168.12.2 255.255.255.0 no sh int lo 0 ip add 2.2.2.2 255.0.0.0 router bgp 2 neighbor 192.168.12.1 remote-as 1 network 2.0.0.0 mask 255.0.0.0

```
R1#sh ip route bgp
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

B     2.0.0.0/8 [20/0] via 192.168.12.2, 00:00:04
```

BGP Tables:

Like most routing protocols, BGP also uses three tables:

Neighbor Table:

- This contains the list of all configured BGP neighbors.
- 'show ip bgp summary' command is use to check neighbor table.


```
R1#sh ip bgp summary
BGP router identifier 1.1.1.1, local AS number 1
BGP table version is 3, main routing table version 3
2 network entries using 288 bytes of memory
2 path entries using 160 bytes of memory
2/2 BGP path/bestpath attribute entries using 272 bytes of memory
1 BGP AS-PATH entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 744 total bytes of memory
BGP activity 2/0 prefixes, 2/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.12.2    4            2      51      55        3    0    0 00:47:31        1
```

BGP router identifier	The IP address representing this router
Local AS number	The local router's Autonomous System Number
BGP table version	Local BGP table increases when the BGP table changes
Main routing table version	Last version of BGP database in the main routing table
Neighbor	The IP address, used in the neighbor statement
V (Version)	The version of BGP this router is running
AS (Autonomous System)	The listed neighbor's Autonomous System Number
MsgRcvd (Message Received)	The number of BGP messages received from neighbor
MsgSent (Message Sent)	The number of BGP messages sent to this neighbor
TblVer (Table Version)	Last version of the BGP table that was sent to neighbor
InQ (In Queue)	In Queue input messages in Queue
OutQ (Out Queue)	Out Queue Output messages in Queue
Up/Down	Time since BGP session was established
State	The current state of the BGP session: active, idle etc
PfxRcd (Prefix Received)	Number of BGP network entries received from this neighbor

Forwarding Table:

- This table contains a list of networks along their path and attributes which are known by BGP.
- 'show ip bgp' command would display the information.

```
R1#show ip bgp
BGP table version is 3, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  1.0.0.0          0.0.0.0                  0         32768 i
 *>  2.0.0.0          192.168.12.2             0             0 2 i
```

BGP table version	Local BGP table increases when the BGP table changes
Local router ID	The IP address representing this router
Network	Learn network with subnet masks
*	This is a valid route and that BGP is able to use it
>	This entry has been selected as the best path
Next Hop	0.0.0.0 means that this network originated on this router R1 learn about this network from 192.168.12.2
Metric	BGP attributes that are used to select the best path
LocPrf	BGP attributes that are used to select the best path
Weight	BGP attributes that are used to select the best path

Path	A sequence of Autonomous Systems in the path from Left to Right
Path i	Network was advertised using the network command
Path 2	AS path 2
Path ?	Redistributed Networks

Routing Table:

- This table lists the best path to the destination networks and also the next hop for each network.
- show ip route command is use to check routing table.

```
R1#sh ip route bgp
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       + - replicated route, % - next hop override

Gateway of last resort is not set

B    2.0.0.0/8 [20/0] via 192.168.12.2, 00:00:57
```

B	This route was learned through BGP
2.0.0.0/8	Destination learn network and 8 is subnet mask
20	20 is the Administrative Distance of eBGP protocol
192.168.12.2	Next Hop IP Address where to send the traffic
00:00:57	Time since the route was learnt

BGP Neighbor States:

- Like OSPF or EIGRP, BGP establishes a neighbor adjacency with other BGP routers.
- To create a neighbor relationship, BGP router first tries to establish TCP connection.
- When the connection is established, the cisco router will send a BGP open message.
- BGP open message is similar to hello message like in OSPF or EIGRP routing protocol.
- Then the parameters are matched, & routers become Border Gateway Protocol peers.
- BGP protocol does not use broadcast or multicast to “Discover” other BGP neighbors.
- Border Gateway Protocol (BGP) neighbors have to be configured manually or statically.
- Border Gateway Protocol (BGP) uses TCP port number 179 for the connection to peer.

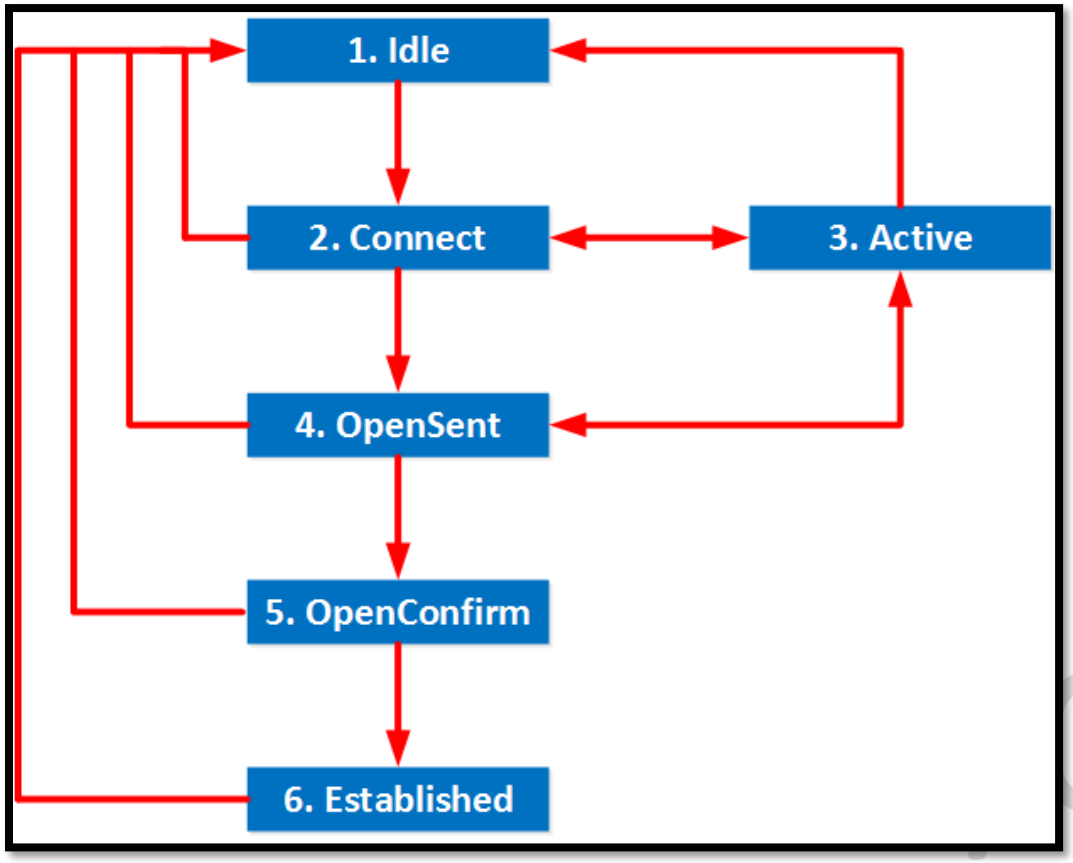

Idle:

- This is the first state where BGP waits for a “start event”.
- The start event occurs when someone configures a new BGP neighbor or when we reset an established BGP peering.
- After the start event, BGP will initialize some resources, resets a ConnectRetry timer and initiates a TCP connection to the remote BGP neighbor.
- It will also start listening for a connection in case the remote BGP neighbor tries to establish a connection.
- When successful, BGP moves to the Connect state. When it fails, it will remain in the Idle state.

Or in Simple words we can say.

- Device is in IDLE state until we configure a BGP, it will refuse all incoming BGP connections.
- After we configure BGP on it, it will initiate TCP Connection with BGP Configure peer and start listening TCP connection from its peer.
- It changes it state from idle to connect.
- If error occur at any state BGP session is terminated and returned into Idle state.
- There are some reasons why router does not progress from the idle state.
- TCP Port 179 is not open.
- AS number configure incorrectly.
- Peer address configure incorrectly.

Connect:

- The BGP process is waiting for the TCP connect to be established.
- BGP is waiting for the TCP three-way handshake to complete.
- If successful, it will continue to the OpenSent State.
- If fails, it will continue to the Active State.
- If BGP reset is, send it will move back to the Idle State.

Or Other simple method:

- In this state device competes the three-way handshaking with the peer.
- After the successful TCP connection established with the peer, device sends an open message to a peer.
- The moment it send an open message to peer, it jumps to a OpenSent state.
- If any error occur while TCP connection establishment , it jumps to an active state.

Active:

- BGP will try another TCP three-way handshake to establish a connection with the remote BGP neighbor.
- If it is successful, it will move to the OpenSent state.
- If the ConnectRetry timer expires then we move back to the Connect state.
- BGP will also keep listening for incoming connections in case the remote BGP neighbor tries to establish a connection.
- Other events can cause the router to go back to the Idle state (resetting BGP for example).

Or Other simple method:

- If the router was unable to establish a successful TCP session, it ends up in Active state.
- It tries to restart another TCP session with the peer, if it is successful then it sends an open message to a peer and jumps open sent state.
- If it is unsuccessful againt, it jumps back to an idle state.

Reasons to toggle between active and idle state is as given below.

- TCP Port 179 is not open.
- Wrong BGP Configuration.
- Network connectivity issues link is flapping.

OpenSent:

- In this state BGP will be waiting for an Open message from the remote BGP neighbor.
- The Open message will be checked for errors, if something is wrong (incorrect version numbers, wrong AS number, etc.) then BGP will respond with a Notification message and jumps back to the Idle state.

- This is also the moment where BGP decides whether we use EBGP or IBGP (since we check the AS number).
- If everything is OK then BGP starts sending keepalive messages and resets its keepalive timer. At this moment, the hold time is negotiated (lowest value is picked) between the two BGP routers.
- In case the TCP session fails, BGP will jump back to the Active state. When any other errors occur (expiration of hold timer), BGP will send a notification message with the error code and jumps back to the Idle state.
- In case someone resets the BGP process, we also jump back to the Idle state.

Or Other simple method:

- Router sends an open message and receives an open message from its Peer.
- Onces it receives an open message it checks it validity like (BGP Version number, AS number, neighbor statement).
- If every thing is ok then sends a keepalives message and jumps to an openconfirm state else sends a notification message.

OpenConfirm:

- BGP waits for a keepalive message from the remote BGP neighbor.
- When we receive the keepalive, we can move to the established state and the neighbor adjacency will be completed.
- When this occurs, it will reset the hold timer. If we receive a notification message from the remote BGP neighbor then we fall back to the Idle state.
- BGP will keep sending keepalive messages.

Established:

- The BGP neighbor adjacency is complete and the BGP routers will send update packets to exchange routing information.

- Every time we receive a keepalive or update message, the hold timer will be resetted.
- In case we receive a notification message we will jump back to the Idle state.

Lab:

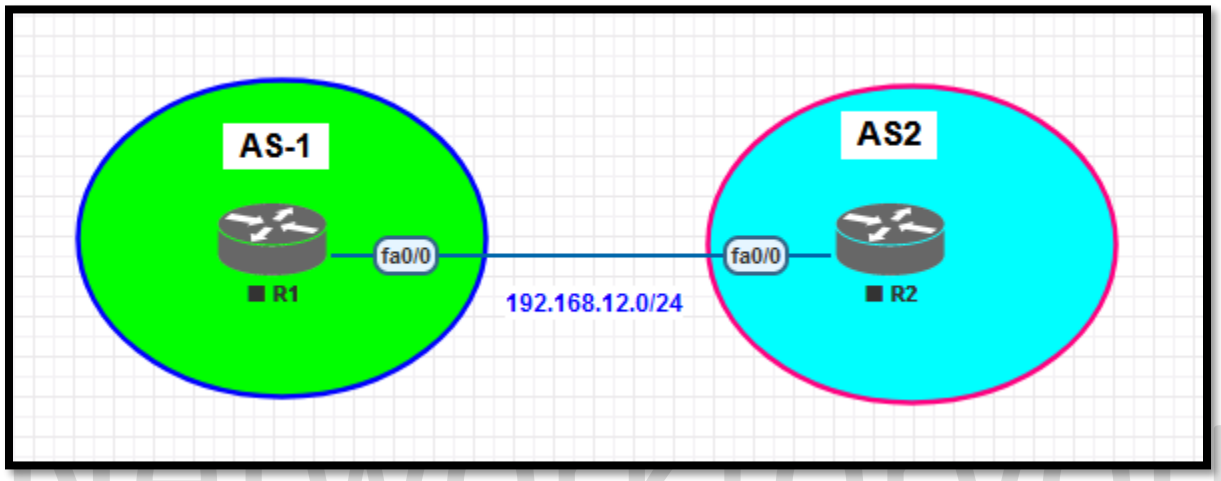

R1 Configuration	R2 Configuration
En Config t Hostname R1 Int f0/0 Ip add 192.168.12.1 255.255.255.0 No sh Router bgp 1 Neighbor 192.168.12.2 remote-as 2	En Config t Hostname R2 Int f0/0 Ip add 192.168.12.2 255.255.255.0 No sh Router bgp 2 Neighbor 192.168.12.1 remote-as 1
Show ip bgp summary	To check status

```
*Apr 10 20:11:25.243: BGP: ses global 192.168.12.2 (0x672FFD30:0) act Reset (Active open failed).
*Apr 10 20:11:25.243: BGP: 192.168.12.2 active went from Active to Idle
*Apr 10 20:11:25.243: BGP: nbr global 192.168.12.2 Active open failed - open timer running
*Apr 10 20:11:25.243: BGP: nbr global 192.168.12.2 Active open failed - open timer running
*Apr 10 20:11:37.515: BGP: 192.168.12.2 active went from Idle to Active
*Apr 10 20:11:37.515: BGP: 192.168.12.2 open active, local address 192.168.12.1
*Apr 10 20:11:37.543: BGP: ses global 192.168.12.2 (0x687F905C:0) act Adding topology IPv4 Unicast:base
*Apr 10 20:11:37.543: BGP: ses global 192.168.12.2 (0x687F905C:0) act Send OPEN
*Apr 10 20:11:37.543: BGP: 192.168.12.2 active went from Active to OpenSent
*Apr 10 20:11:37.543: BGP: 192.168.12.2 active sending OPEN, version 4, my as: 1, holdtime 180 seconds, ID C0A80C01
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active rcv message type 1, length (excl. header) 39
*Apr 10 20:11:37.551: BGP: ses global 192.168.12.2 (0x687F905C:0) act Receive OPEN
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active rcv OPEN, version 4, holdtime 180 seconds
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active rcv OPEN w/ OPTION parameter len: 29
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active rcvd OPEN w/ optional parameter type 2 (Capability) len 6
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has CAPABILITY code: 1, length 4
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has MP_EXT CAP for afi/safi: 1/1
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active rcvd OPEN w/ optional parameter type 2 (Capability) len 2
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has CAPABILITY code: 128, length 0
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has ROUTE-REFRESH capability(old) for all address-families
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active rcvd OPEN w/ optional parameter type 2 (Capability) len 2
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has CAPABILITY code: 2, length 0
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has ROUTE-REFRESH capability(new) for all address-families
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active rcvd OPEN w/ optional parameter type 2 (Capability) len 3
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has CAPABILITY code: 131, length 1
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has MULTISESSION capability, without grouping
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active rcvd OPEN w/ optional parameter type 2 (Capability) len 6
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has CAPABILITY code: 65, length 4
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has 4-byte ASN CAP for: 2
*Apr 10 20:11:37.551: BGP: nbr global 192.168.12.2 neighbor does not have IPv4 MDT topology activated
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active rcvd OPEN w/ remote AS 2, 4-byte remote AS 2
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active went from OpenSent to OpenConfirm
*Apr 10 20:11:37.551: BGP: 192.168.12.2 active went from OpenConfirm to Established
*Apr 10 20:11:37.551: BGP: ses global 192.168.12.2 (0x687F905C:1) act Assigned ID
*Apr 10 20:11:37.555: BGP: ses global 192.168.12.2 (0x687F905C:1) Up
```

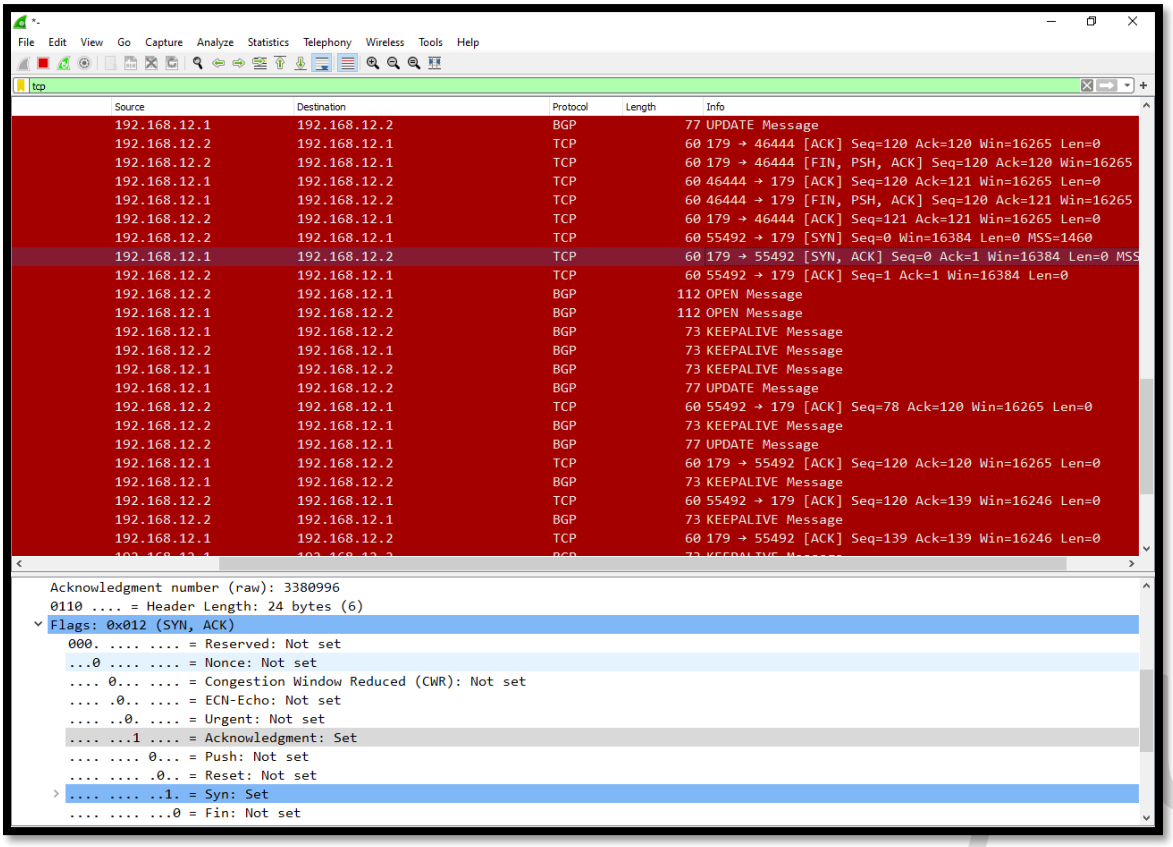

R1# debug ip bgp all

*Apr 10 20:11:25.243: BGP: nbr global 192.168.12.2 Active open failed - open timer running

*Apr 10 20:11:37.515: BGP: 192.168.12.2 active went from Idle to Active

*Apr 10 20:11:37.515: BGP: 192.168.12.2 open active, local address 192.168.12.1

*Apr 10 20:11:37.543: BGP: ses global 192.168.12.2 (0x687F905C:0) act Adding topology IPv4 Unicast:base

*Apr 10 20:11:37.543: BGP: ses global 192.168.12.2 (0x687F905C:0) act Send OPEN

*Apr 10 20:11:37.543: BGP: 192.168.12.2 active went from Active to OpenSent

*Apr 10 20:11:37.543: BGP: 192.168.12.2 active sending OPEN, version 4, my as: 1, holdtime 180 seconds, ID C0A80C01

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active rcv message type 1, length (excl. header) 39

*Apr 10 20:11:37.551: BGP: ses global 192.168.12.2 (0x687F905C:0) act Receive OPEN

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active rcv OPEN, version 4, holdtime 180 seconds

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active rcv OPEN w/ OPTION parameter len: 29

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active rcvd OPEN w/ optional parameter type 2 (Capability) len 6

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has CAPABILITY code: 1, length 4

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has MP_EXT CAP for afi/safi: 1/1

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active rcvd OPEN w/ optional parameter type 2 (Capability) len 2

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has CAPABILITY code: 128, length 0

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has ROUTE-REFRESH capability(old) for all address-families

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active rcvd OPEN w/ optional parameter type 2 (Capability) len 2

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has CAPABILITY code: 2, length 0

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has ROUTE-REFRESH capability(new) for all address-families

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active rcvd OPEN w/ optional parameter type 2 (Capability) len 3

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has CAPABILITY code: 131, length 1

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has MULTISESSION capability, without grouping

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active rcvd OPEN w/ optional parameter type 2 (Capability) len 6

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has CAPABILITY code: 65, length 4

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active OPEN has 4-byte ASN CAP for: 2

*Apr 10 20:11:37.551: BGP: nbr global 192.168.12.2 neighbor does not have IPv4 MDT topology activated

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active rcvd OPEN w/ remote AS 2, 4-byte remote AS 2

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active went from **OpenSent to OpenConfirm**

*Apr 10 20:11:37.551: BGP: 192.168.12.2 active went from **OpenConfirm to Established**

*Apr 10 20:11:37.551: BGP: ses global 192.168.12.2 (0x687F905C:1) act Assigned ID

*Apr 10 20:11:37.555: BGP: ses global 192.168.12.2 (0x687F905C:1) Up

*Apr 10 20:11:37.555: %BGP-5-ADJCHANGE: **neighbor 192.168.12.2 Up**

BGP Message Types:

BGP uses four message types in its operations:

1. Open:

- Open message is sent after BGP neighbor is configured.
- Exchanges BGP values and capabilities.
- Open message is sent to establish or form peering with that neighbor.
- Contains information such as **Version, AS Number, Router ID** and **the Hold-Time**.

2. Update:

- Routing information between peers is transferred using Update message.
- Update message include, new routes, withdrawn routes and the path attributes.

3. Keepalive:

- Keepalive message is similar to Hellos message in other routing protocols.
- BGP uses keepalives, which help in keeping the peering session active.
- Sent periodically to maintain neighbor relationship.
- The BGP peers exchange Keepalive messages every 60 seconds.
- The BGP peers Hold Timer is 180 seconds by default.

4. Notification:

- In the event of a problem which causes the router to end the BGP peering session.
- BGP neighbor send a notification message and the connection is closed.
- Typically resets the neighbor relationship.
- Such ad Bad peer AS, bad BGP identifier, unacceptable hold time etc.

Lab for BGP Messages Types:

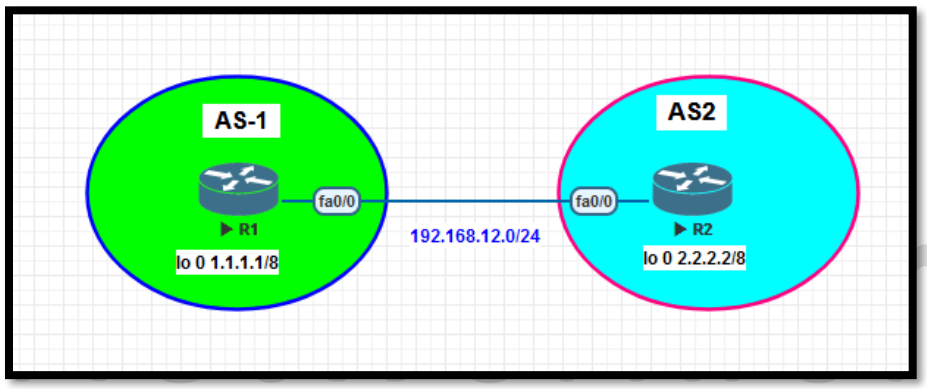

R1 Configuration	R2 Configuration
en config t hostname R1 int f0/0 ip add 192.168.12.1 255.255.255.0 no sh int lo 0 ip add 1.1.1.1 255.0.0.0 router bgp 1	en config t hostname R2 int f0/0 ip add 192.168.12.2 255.255.255.0 no sh int lo 0 ip add 2.2.2.2 255.0.0.0 router bgp 2

neighbor 192.168.12.2 remote-as 2 network 1.0.0.0 mask 255.0.0.0	neighbor 192.168.12.1 remote-as 1 network 2.0.0.0 mask 255.0.0.0

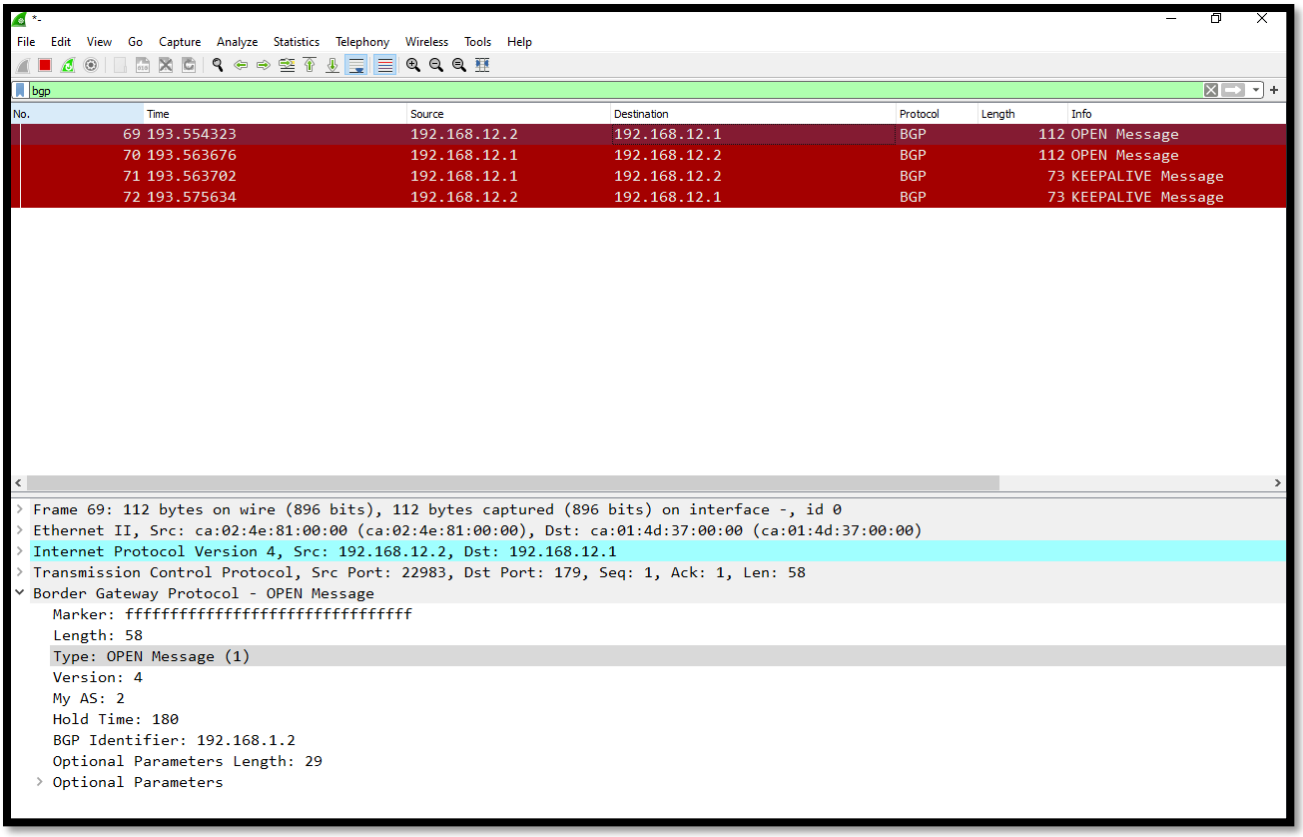

R2(config-if)# sh

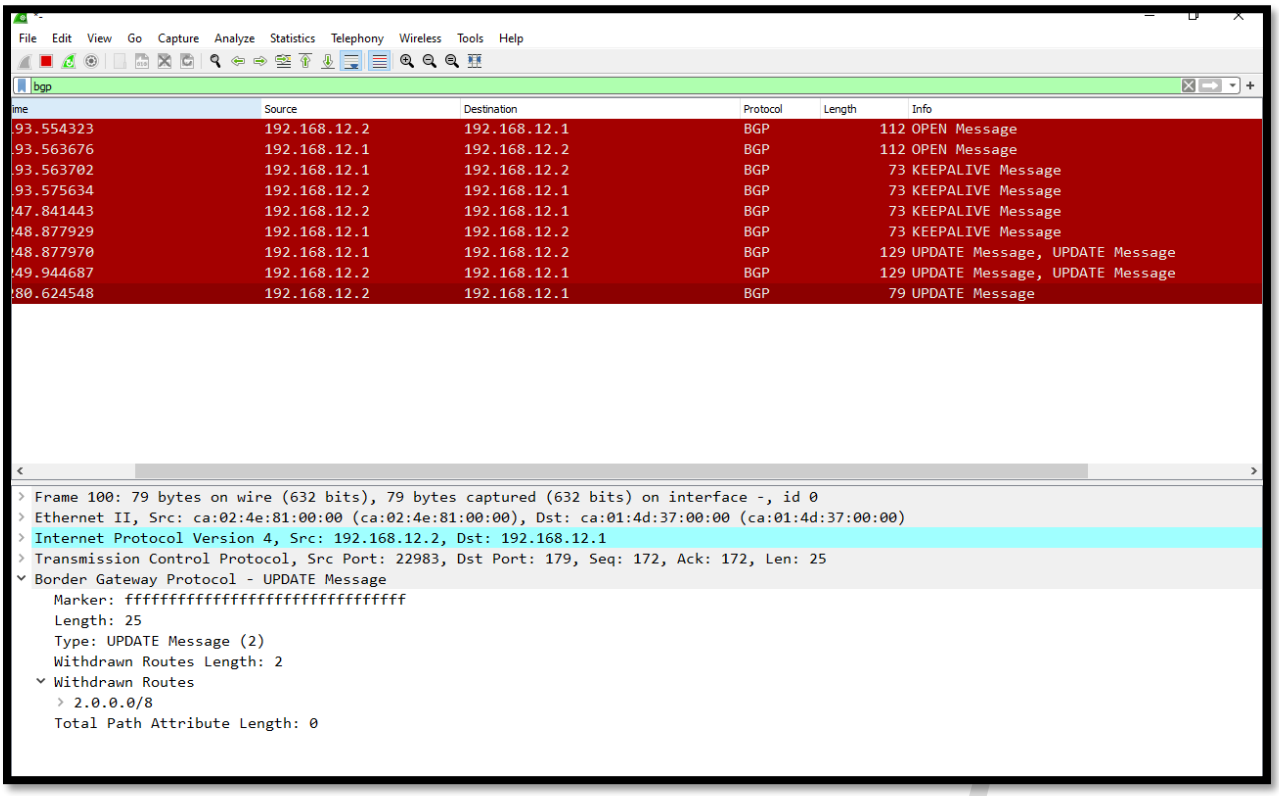

R2(config-if)# no sh

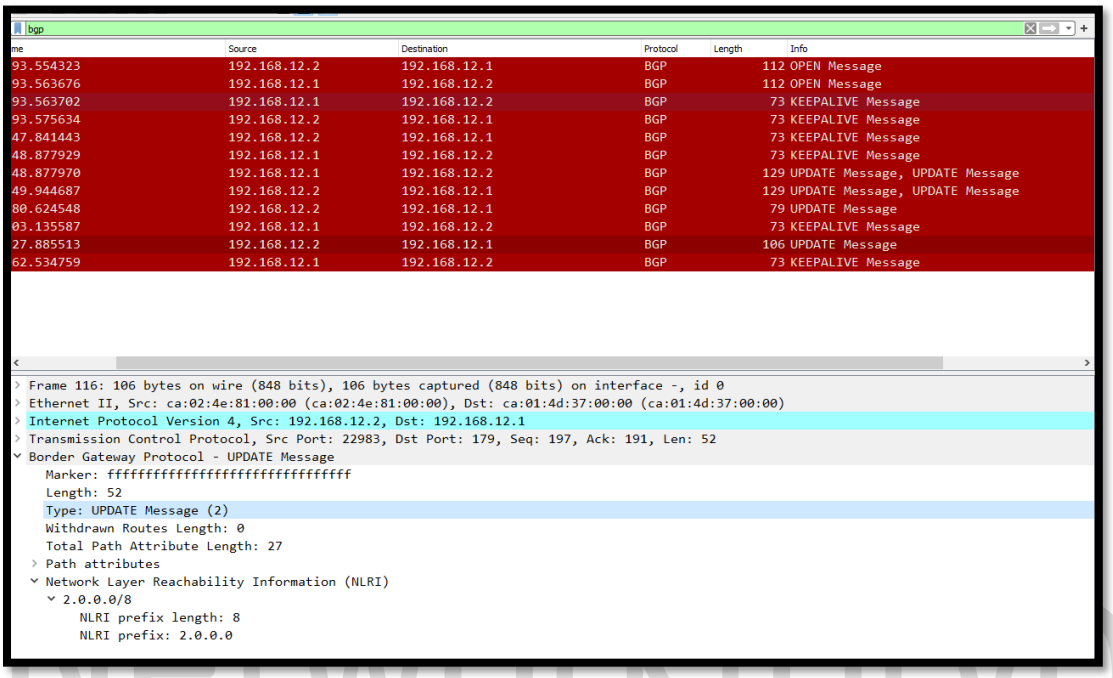


```
R2(config-if)#no router bgp 2
R2(config)#router
*Apr 10 20:56:33.631: %BGP-5-ADJCHANGE: neighbor 192.168.12.1 Down Neighbor deleted
*Apr 10 20:56:33.631: %BGP_SESSION-5-ADJCHANGE: neighbor 192.168.12.1 IPv4 Unicast topology base removed from session  Neighb
or deletedbgp
% Incomplete command.

R2(config)#
R2(config)#
R2(config)#router bgp 3
R2(config-router)#neighbor 192.168.12.1 remote-as 1
R2(config-router)#router bgp 3
*Apr 10 20:57:01.175: %BGP-3-NOTIFICATION: received from neighbor 192.168.12.1 passive 2/2 (peer in wrong AS) 2 bytes 00
*Apr 10 20:57:01.179: %BGP_SESSION-5-ADJCHANGE: neighbor 192.168.12.1 IPv4 Unicast topology base removed from session  BGP No
tification received
*Apr 10 20:57:06.839: %BGP-3-NOTIFICATION: received from neighbor 192.168.12.1 active 2/2 (peer in wrong AS) 2 bytes 0003
*Apr 10 20:57:06.839: %BGP_SESSION-5-ADJCHANGE: neighbor 192.168.12.1 IPv4 Unicast topology base removed from session  BGP No
tification received
```

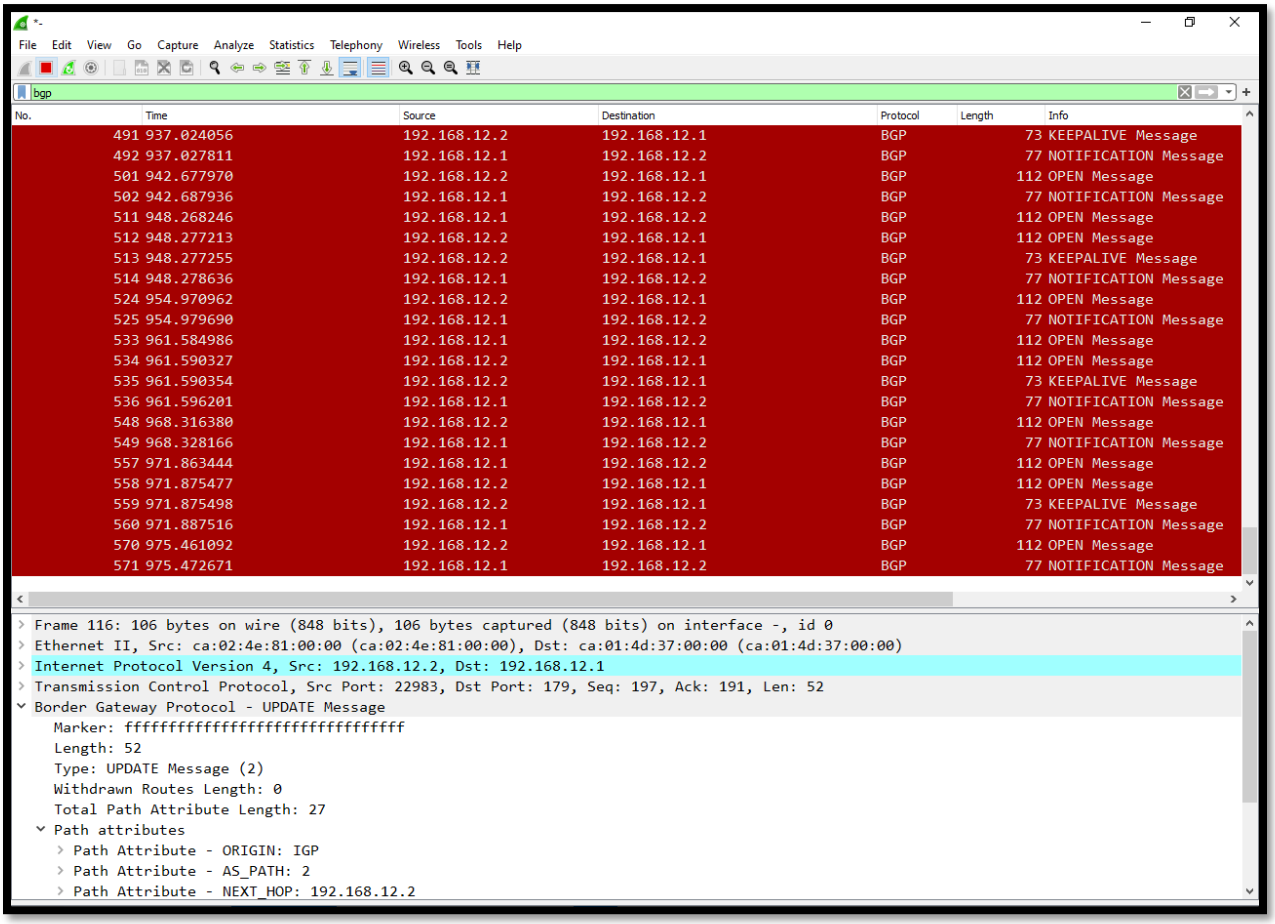
File Edit View Go Capture Analyze Statistics Telephony Wireless Tools Help
bgp
No. Time Source Destination Protocol Length Info
491 937.024056 192.168.12.2 192.168.12.1 BGP 73 KEEPALIVE Message
492 937.027811 192.168.12.1 192.168.12.2 BGP 77 NOTIFICATION Message
501 942.677970 192.168.12.2 192.168.12.1 BGP 112 OPEN Message
502 942.687936 192.168.12.1 192.168.12.2 BGP 77 NOTIFICATION Message
511 948.268246 192.168.12.1 192.168.12.2 BGP 112 OPEN Message
512 948.277213 192.168.12.2 192.168.12.1 BGP 112 OPEN Message
513 948.277255 192.168.12.2 192.168.12.1 BGP 73 KEEPALIVE Message
514 948.278636 192.168.12.1 192.168.12.2 BGP 77 NOTIFICATION Message
524 954.970962 192.168.12.2 192.168.12.1 BGP 112 OPEN Message
525 954.979690 192.168.12.1 192.168.12.2 BGP 77 NOTIFICATION Message
533 961.584986 192.168.12.1 192.168.12.2 BGP 112 OPEN Message
534 961.590327 192.168.12.2 192.168.12.1 BGP 112 OPEN Message
535 961.590354 192.168.12.2 192.168.12.1 BGP 73 KEEPALIVE Message
536 961.596201 192.168.12.1 192.168.12.2 BGP 77 NOTIFICATION Message
548 968.316380 192.168.12.2 192.168.12.1 BGP 112 OPEN Message
549 968.328166 192.168.12.1 192.168.12.2 BGP 77 NOTIFICATION Message
557 971.863444 192.168.12.1 192.168.12.2 BGP 112 OPEN Message
558 971.875477 192.168.12.2 192.168.12.1 BGP 112 OPEN Message
559 971.875498 192.168.12.2 192.168.12.1 BGP 73 KEEPALIVE Message
560 971.887516 192.168.12.1 192.168.12.2 BGP 77 NOTIFICATION Message
570 975.461092 192.168.12.2 192.168.12.1 BGP 112 OPEN Message
571 975.472671 192.168.12.1 192.168.12.2 BGP 77 NOTIFICATION Message
Frame 116: 106 bytes on wire (848 bits), 106 bytes captured (848 bits) on interface -, id 0
Ethernet II, Src: ca:02:4e:81:00:00 (ca:02:4e:81:00:00), Dst: ca:01:4d:37:00:00 (ca:01:4d:37:00:00)
Internet Protocol Version 4, Src: 192.168.12.2, Dst: 192.168.12.1
Transmission Control Protocol, Src Port: 22983, Dst Port: 179, Seq: 197, Ack: 191, Len: 52
Border Gateway Protocol - UPDATE Message
Marker: ffffffffffffffffffffffffffffffff
Length: 52
Type: UPDATE Message (2)
Withdrawn Routes Length: 0
Total Path Attribute Length: 27
Path attributes
Path Attribute - ORIGIN: IGP
Path Attribute - AS_PATH: 2
Path Attribute - NEXT_HOP: 192.168.12.2

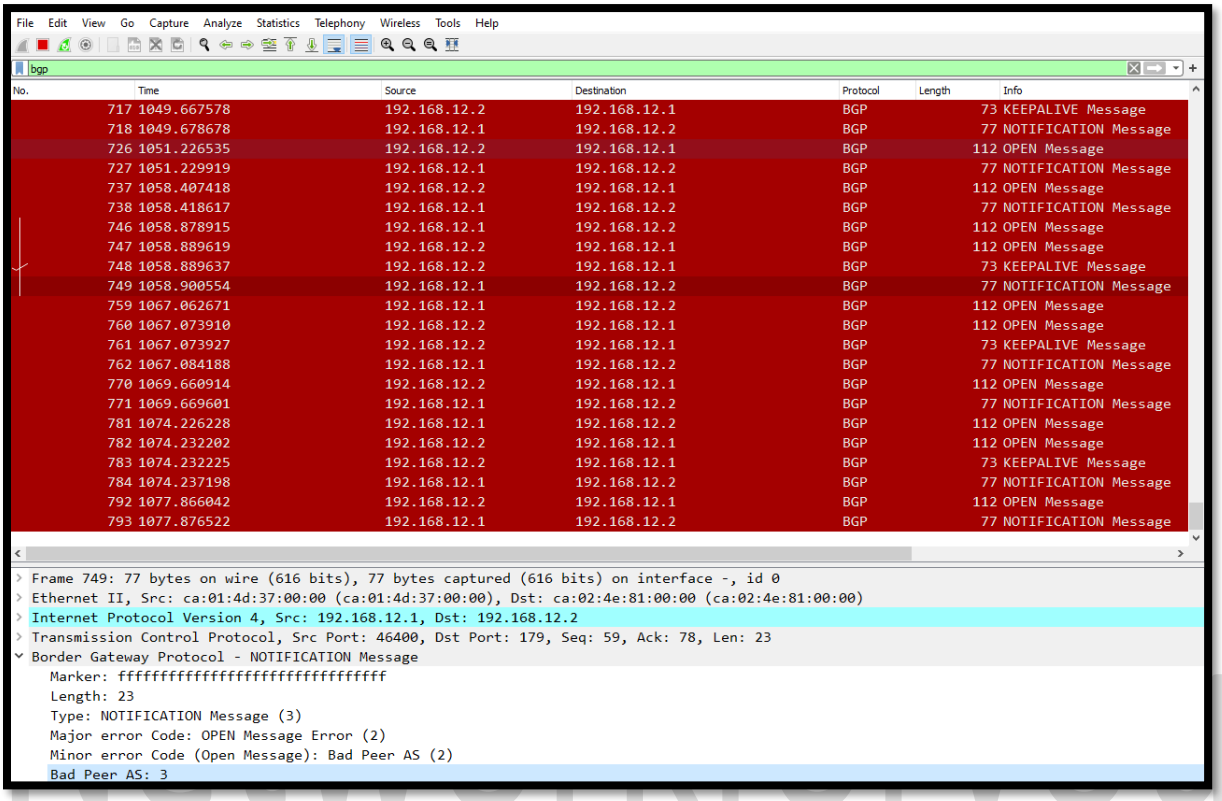

BGP Active and Passive:

- By default, neighbor with lowest BGP RID will establish connection.
- **Active having a higher random port > 1023**.
- **Passive having TCP Port 179.**
- **Active is called Client and Passive is called Server**.
- This behavior can be modified.

Lab for Active and Passive:

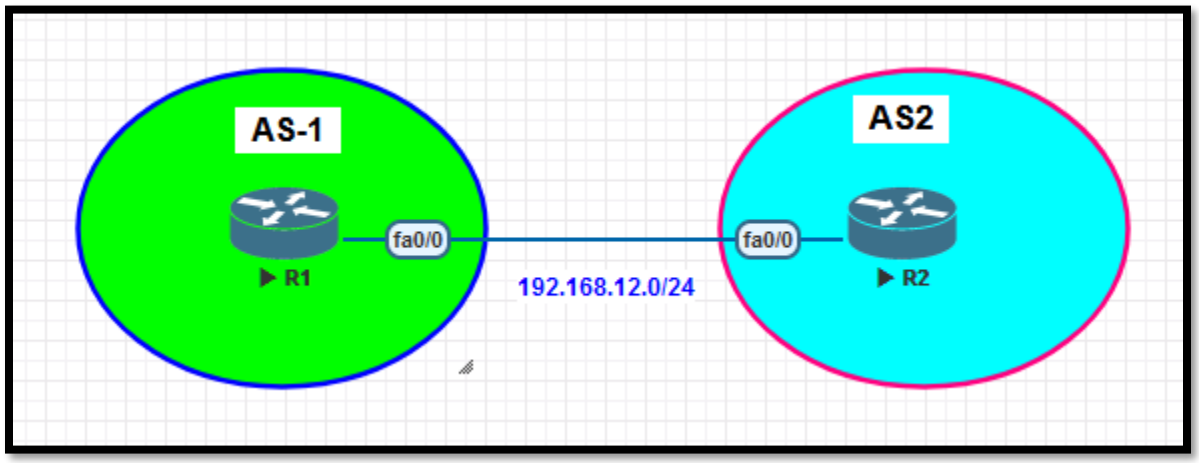

R1 Configuration	R2 Configuration
en config t hostname R1 int f0/0 ip add 192.168.12.1 255.255.255.0 no sh router bgp 1 neighbor 192.168.12.2 remote-as 2	en config t hostname R2 int f0/0 ip add 192.168.12.2 255.255.255.0 no sh router bgp 2 neighbor 192.168.12.1 remote-as 1

```
R1#sh ip bgp neighbors | include host
Local host: 192.168.12.1, Local port: 179
Foreign host: 192.168.12.2, Foreign port: 36145
R1#
```

```
R2#sh ip bgp neighbors | include host
Local host: 192.168.12.2, Local port: 36145
Foreign host: 192.168.12.1, Foreign port: 179
R2#
```

After clear BGP this time R1 become active and R2 become Passive.

```
R1#clear ip bgp *
R1#sh
*Apr 10 21:12:23.643: %BGP-5-ADJCHANGE: neighbor 192.168.12.2
 Down User reset
*Apr 10 21:12:23.643: %BGP_SESSION-5-ADJCHANGE: neighbor 192.
168.12.2 IPv4 Unicast topology base removed from session  Use
r reset
*Apr 10 21:12:24.095: %BGP-5-ADJCHANGE: neighbor 192.168.12.2
 Up ip bgp neg
R1#sh ip bgp neighbors | include host
Local host: 192.168.12.1, Local port: 30947
Foreign host: 192.168.12.2, Foreign port: 179
R1#
```

```
R2#sh ip bgp neighbors | include host
Local host: 192.168.12.2, Local port: 36145
Foreign host: 192.168.12.1, Foreign port: 179
R2#
*Apr 10 21:12:24.167: %BGP-5-ADJCHANGE: neighbor 192.168.12
.1 Down Peer closed the session
*Apr 10 21:12:24.167: %BGP_SESSION-5-ADJCHANGE: neighbor 19
2.168.12.1 IPv4 Unicast topology base removed from session
 Peer closed the session
*Apr 10 21:12:24.619: %BGP-5-ADJCHANGE: neighbor 192.168.12
.1 Up
R2#sh ip bgp neighbors | include host
Local host: 192.168.12.2, Local port: 179
Foreign host: 192.168.12.1, Foreign port: 30947
R2#
```

Manually Active and Passive	
R1(config)#router bgp 1 R1(config-router)#neighbor 192.168.12.2 transport connection-mode passive	R2(config)#router bgp 2 R2(config-router)#neighbor 192.168.12.1 transport connection-mode active

BGP Timers:

- Basic BGP times are Keepalive and Hold-down timer intervals.
- By default, Keepalive timer is 60 seconds.
- By default, hold-down timer is 3 x Keepalive or 180 seconds.
- Once peers is UP, router starts a hold-down timer counting from 0 second up.
- Every Keepalive message the neighbor peer resets this timer back to 0 seconds.
- Failing to receive 3 keepalives in a row will make the hold-down timer reach 180 seconds.
- Neighbor is considered down and routes from this neighbor are flushed.
- To verify current timers, issue the “show ip bgp neighbor” command.
- BGP timers can be changed, both the defaults and on a per-neighbor basis.
- In the OPEN message, BGP routers exchange the hold time they want to use.
- Values of 1 or 2 are illegal, minimum working value for the hold time is 3 seconds.

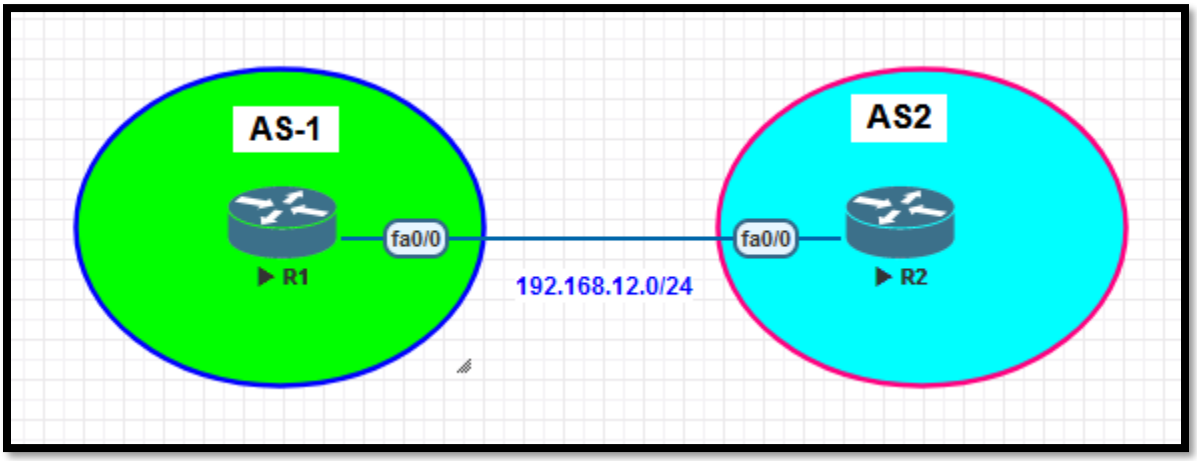

R1 Configuration	R2 Configuration
en config t hostname R1 int f0/0 ip add 192.168.12.1 255.255.255.0 no sh router bgp 1 neighbor 192.168.12.2 remote-as 2	en config t hostname R2 int f0/0 ip add 192.168.12.2 255.255.255.0 no sh router bgp 2 neighbor 192.168.12.1 remote-as 1

```
R1#sh ip bgp neighbors
BGP neighbor is 192.168.12.2,  remote AS 2, external link
  BGP version 4, remote router ID 192.168.1.2
  BGP state = Established, up for 00:05:42
  Last read 00:00:29, last write 00:00:30, hold time is 180, keepalive interval is 60 seconds
```

```
R2#sh ip bgp neighbors
BGP neighbor is 192.168.12.1,  remote AS 1, external link
  BGP version 4, remote router ID 192.168.12.1
  BGP state = Established, up for 00:08:46
  Last read 00:00:01, last write 00:00:48, hold time is 180, keepalive interval is 60 seconds
  Neighbor sessions:
```

If we want to change timer we can do this way.
R2(config)#router bgp 2
R2(config-router)#timers bgp 60 180
R2#clear ip bgp *

```
R1#sh ip bgp neighbors
BGP neighbor is 192.168.12.2,  remote AS 2, external link
  BGP version 4, remote router ID 192.168.12.2
  BGP state = Established, up for 00:00:04
  Last read 00:00:04, last write 00:00:04, hold time is 180, keepalive interval is 60 seconds
```


BGP Next-Hope-Self:

- **By default, feature of BGP while forwarding routes from EBGP to IBGP router next hop does not change.**
- Sometimes this can cause reachability issues.
- The router will make sure that the next-hop attribute reflects its IP address.
- BGP Next-Hop-Self solves reachability problems.

Let check with Lab.

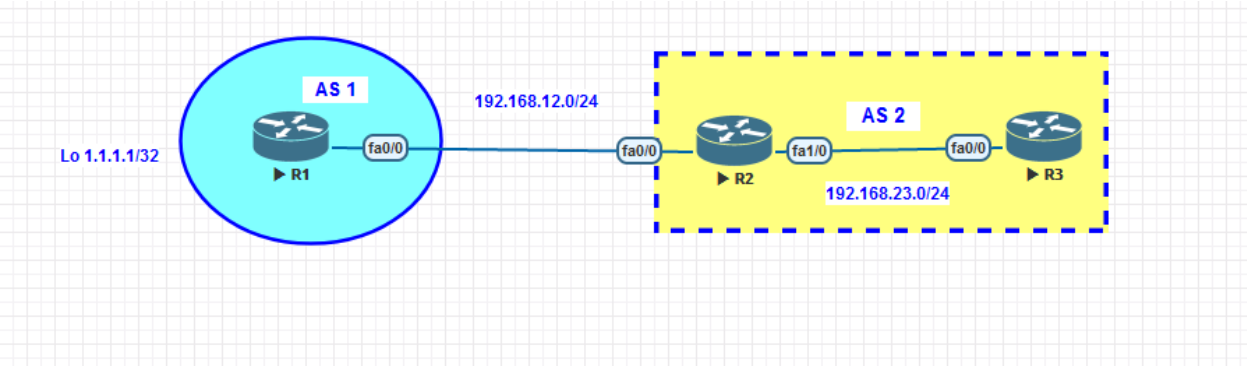

R1 Configuration	R2 Configuration	R3 Configuration
en config t hostname R1 int f0/0 ip add 192.168.12.1 255.255.255.0 no sh int lo 0 ip add 1.1.1.1 255.255.255.255 router bgp 1 neighbor 192.168.12.2 remote-as 2 network 1.1.1.1 mask 255.255.255.255	en config t hostname R2 int f0/0 ip add 192.168.12.2 255.255.255.0 no sh int f1/0 ip add 192.168.23.1 255.255.255.0 no sh router bgp 2 neighbor 192.168.12.1 remote-as 1 neighbor 192.168.23.2 remote-as 2	en config t hostname R3 int f0/0 ip add 192.168.23.2 255.255.255.0 no sh router bgp 2 neighbor 192.168.23.1 remote-as 2

R2 has installed 1.1.1.1 in its BGP table and it is a valid route, the next hop is 192.168.12.1

```
R2#sh ip bgp
BGP table version is 2, local router ID is 192.168.23.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.1.1/32       192.168.12.1             0             0 1 i
```

R3 learns the prefix but unable to install it in the routing table.

```
R3#sh ip bgp
BGP table version is 2, local router ID is 192.168.23.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
* i1.1.1.1/32       192.168.12.1             0    100      0 1 i
```

Solution is to advertise the network and put “Next-Hop-Self” command.

```
R2(config)#router bgp 2
R2(config-router)# neighbor 192.168.23.2 remote-as 2
R2(config-router)#neighbor 192.168.23.2 next-hop-self
R2(config-router)#network 192.168.23.0 mask 255.255.255.0
```


Now after **Next-Hop-Self** R3 learn the prefix and install in routing table.

```
R3#sh ip bgp
BGP table version is 5, local router ID is 192.168.23.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.1.1.1/32       192.168.23.1             0    100      0 1 i
r>i192.168.23.0     192.168.23.1             0    100      0 i
```

BGP Multihop & Update-Source:

- EBGP routers use a TTL value of one for their BGP packets.
- BGP neighbor away more than one hop TTL decrement to 0 & discarded.
- The solution is to use to Multihop command.
- Use the ebgp-multihop command to increase the TTL value.
- Multihop command does not apply to Internal BGP.
- If source EBGP from the loopback interfaces, then require Multihop.

- If source EBGP from the loopback interfaces also require update-source loopback.
- Using a loopback interface as update source, BGP session will not go down.
- Update source can be configured per neighbor or per peer-group.
- Static route is required to ensure that the loopback is reachable from both ends.
- Multihop enables the peers to pass through the other routers to form peer relationships.

Lab time:

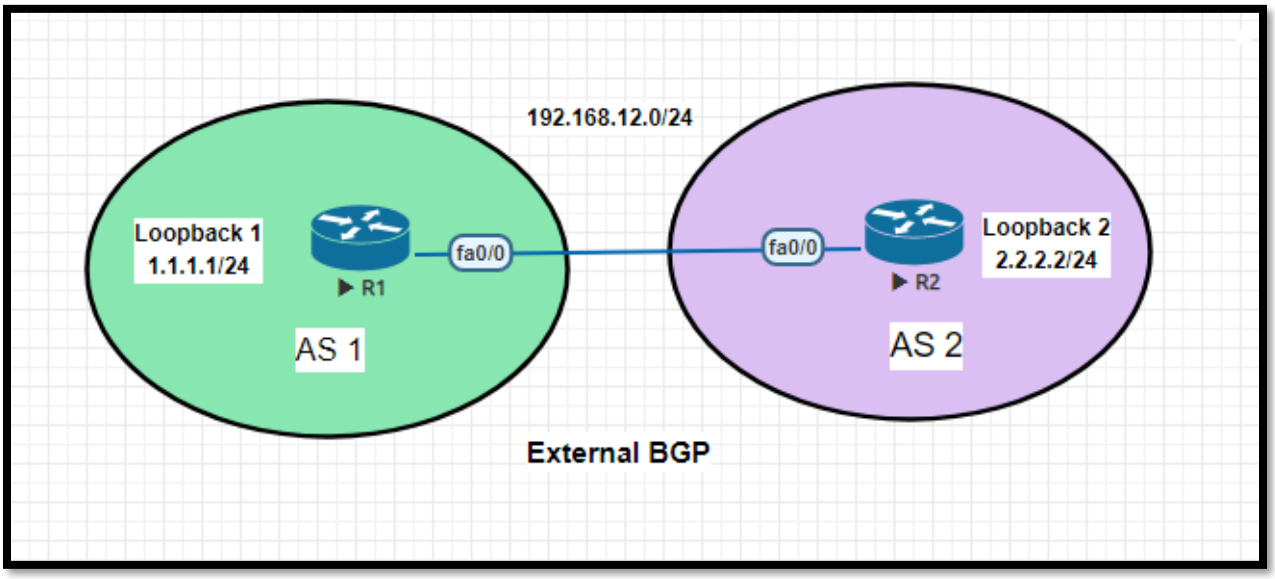

R1 Configuration	R2 Configuration
en config t hostname R1	en config t hostname R2

int f0/0 ip add 192.168.12.1 255.255.255.0 no sh int lo 1 ip add 1.1.1.1 255.255.255.0 router bgp 1 neighbor 2.2.2.2 remote-as 2 network 1.1.1.0 mask 255.255.255.0 ip route 2.2.2.0 255.255.255.0 192.168.12.2	int f0/0 ip add 192.168.12.2 255.255.255.0 no sh int lo 2 ip add 2.2.2.2 255.255.255.0 router bgp 2 neighbor 1.1.1.1 remote-as 1 network 2.2.2.0 mask 255.255.255.0 ip route 1.1.1.0 255.255.255.0 192.168.12.1
neighbor 2.2.2.2 update-source loopback 1 **neighbor 2.2.2.2 ebgp-multihop 2**	**neighbor 1.1.1.1 update-source loopback 2** **neighbor 1.1.1.1 ebgp-multihop 2**


```
R1#sh ip bgp neighbors
BGP neighbor is 2.2.2.2,  remote AS 2, external link
  BGP version 4, remote router ID 2.2.2.2
  BGP state = Established, up for 00:24:59
  Last read 00:00:04, last write 00:00:21, hold time is 180, keepalive interval is 60 seconds
  Neighbor sessions:
    1 active, is not multisession capable (disabled)
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
    Enhanced Refresh Capability: advertised and received
    Multisession Capability:
    Stateful switchover support enabled: NO for session 1
```

```
  Address tracking is enabled, the RIB does have a route to 2.2.2.2
  Connections established 1; dropped 0
  Last reset never
  External BGP neighbor may be up to 2 hops away.
  Transport(tcp) path-mtu-discovery is enabled
  Graceful-Restart is disabled
Connection state is ESTAB, I/O status: 1, unread input bytes: 0
Connection is ECN Disabled
Mininum incoming TTL 0, Outgoing TTL 2
Local host: 1.1.1.1, Local port: 58246
Foreign host: 2.2.2.2, Foreign port: 179
Connection tableid (VRF): 0
```

BGP Peer Group:

- When you configure BGP on a router it's possible that some of the BGP neighbors **share the exact same configuration**.
- This can be annoying since you have to type in the **exact same commands for each of these neighbors.**
- Also, when BGP prepares updates it does this separately for each neighbor.
- This means that it has to use **CPU resources to prepare the update** for each neighbor.
- To **simplify the configuration of BGP and to reduce the number of updates BGP has to create, we can use peer groups.**
- We **can add neighbors to a peer group and then apply all our configurations to the peer group.**
- BGP will prepare the updates for the peer group **which requires less CPU resources than preparing them for each neighbor separately.**

Lab Time:

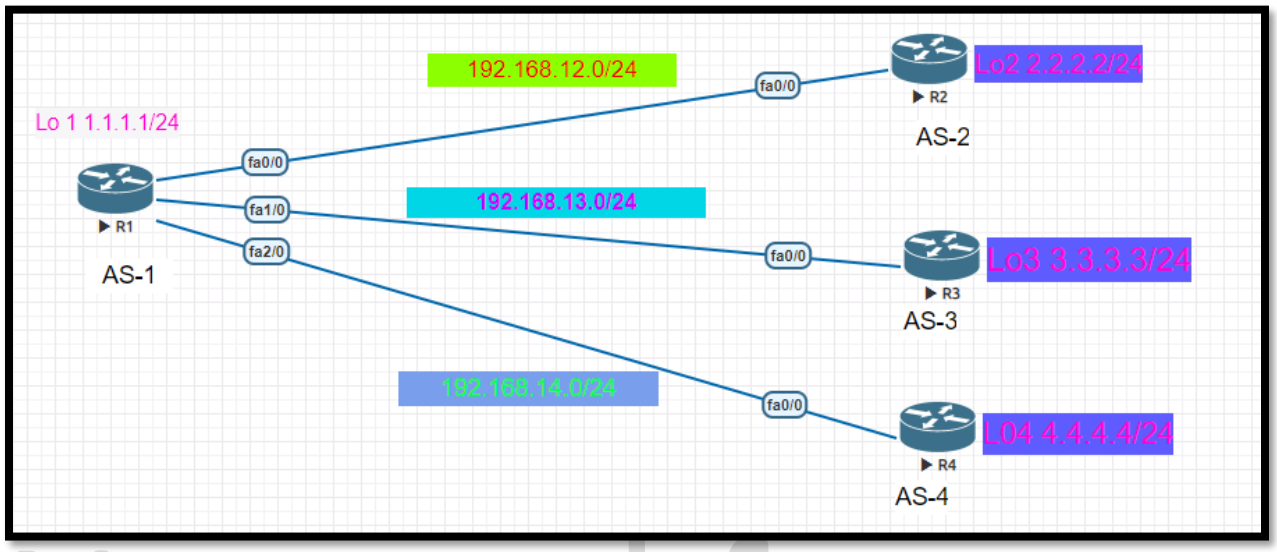

R1 Configuration	R2 Configuration	R3 Configuration
en config t hostname R1 int f0/0 ip add 192.168.12.1 255.255.255.0 no sh int f1/0 ip add 192.168.13.1 255.255.255.0 no sh int f2/0	en config t hostname R2 int f0/0 ip add 192.168.12.2 255.255.255.0 no sh int lo 2 ip add 2.2.2.2 255.255.255.0 ip route 1.1.1.0 255.255.255.0	en config t hostname 3 int f0/0 ip add 192.168.13.2 255.255.255.0 no sh int lo 3 ip add 3.3.3.3 255.255.255.0 ip route 1.1.1.0 255.255.255.0

ip add 192.168.14.1 255.255.255.0 no sh int lo 1 ip add 1.1.1.1 255.255.255.0 ip route 2.2.2.0 255.255.255.0 192.168.12.2 ip route 3.3.3.0 255.255.255.0 192.168.13.2 ip route 4.4.4.0 255.255.255.0 192.168.14.2	192.168.12.1 router bgp 2 neighbor 1.1.1.1 remote-as 1 neighbor 1.1.1.1 update-source loopback 2 neighbor 1.1.1.1 ebgp-multihop 2 network 2.2.2.0 mask 255.255.255.0	192.168.13.1 router bgp 3 neighbor 1.1.1.1 remote-as 1 neighbor 1.1.1.1 update-source loopback 3 neighbor 1.1.1.1 ebgp-multihop 2 network 3.3.3.0 mask 255.255.255.0
R4 Configuration		
en config t hostname R4 int f0/0 ip add 192.168.14.2 255.255.255.0 no sh int lo 4 ip add 4.4.4.4 255.255.255.0 ip route 1.1.1.0 255.255.255.0		

192.168.14.1 router bgp 4 neighbor 1.1.1.1 remote-as 1 neighbor 1.1.1.1 update-source loopback 4 neighbor 1.1.1.1 ebgp-multihop 2 network 4.4.4.0 mask 255.255.255.0		

R1 Configuration without Peer Group
R1(config)#router bgp 1 R1(config-router)#neighbor 2.2.2.2 remote-as 2 R1(config-router)#neighbor 3.3.3.3 remote-as 3 R1(config-router)#neighbor 4.4.4.4 remote-as 4 R1(config-router)#network 1.1.1.0 mask 255.255.255.0 R1(config-router)#neighbor 2.2.2.2 update-source loopback 1 R1(config-router)#neighbor 3.3.3.3 update-source loopback 1 R1(config-router)#neighbor 4.4.4.4 update-source loopback 1 R1(config-router)#neighbor 2.2.2.2 ebgp-multihop 2 R1(config-router)#neighbor 3.3.3.3 ebgp-multihop 2 R1(config-router)#neighbor 4.4.4.4 ebgp-multihop 2

R1 Configuration with Peer Group
R1(config)#router bgp 1 R1(config-router)#neighbor 2.2.2.2 remote-as 2 R1(config-router)#neighbor 3.3.3.3 remote-as 3 R1(config-router)#neighbor 4.4.4.4 remote-as 4 R1(config-router)#neighbor gp1 peer-group R1(config-router)#neighbor 2.2.2.2 peer-group gp1 R1(config-router)#neighbor 3.3.3.3 peer-group gp1 R1(config-router)#neighbor 4.4.4.4 peer-group gp1 R1(config-router)#neighbor gp1 update-source loopback 1 R1(config-router)#neighbor gp1 ebgp-multihop 2

BGP Attributes:

- BGP is a very flexible and extensible protocol.
- BGP path selection is done through the best path algorithm.
- Best path uses various attributes assigned to each route.
- BGP attributes are similar to metrics in OSPF and EIGRP.
- BGP selects the best path based on a list of attributes.
- BGP use attributes to decide the best route.
- BGP has 0 to 14 attributes for Cisco.
- BGP has 1 to 14 attributes for other venders.
- BGP does not use metrics but use set of attributes.

BGP has four main types of attributes.

- Well-Known Mandatory.
- Well-Known Discretionary.
- Optional Transitive.
- Optional Non-Transitive.

Well-Known Mandatory:

- As the name suggests it is mandatory and must.
- **These attribute must appear in every Update message**.
- Must be recognized & supported by all BGP speakers.
- If these attributes are missing a Notification, error is generated.
- If these attributes are missing the session will be closed.
- Well-Known mandatory attributes are **AS Path, Next Hop Address, & Origin**.

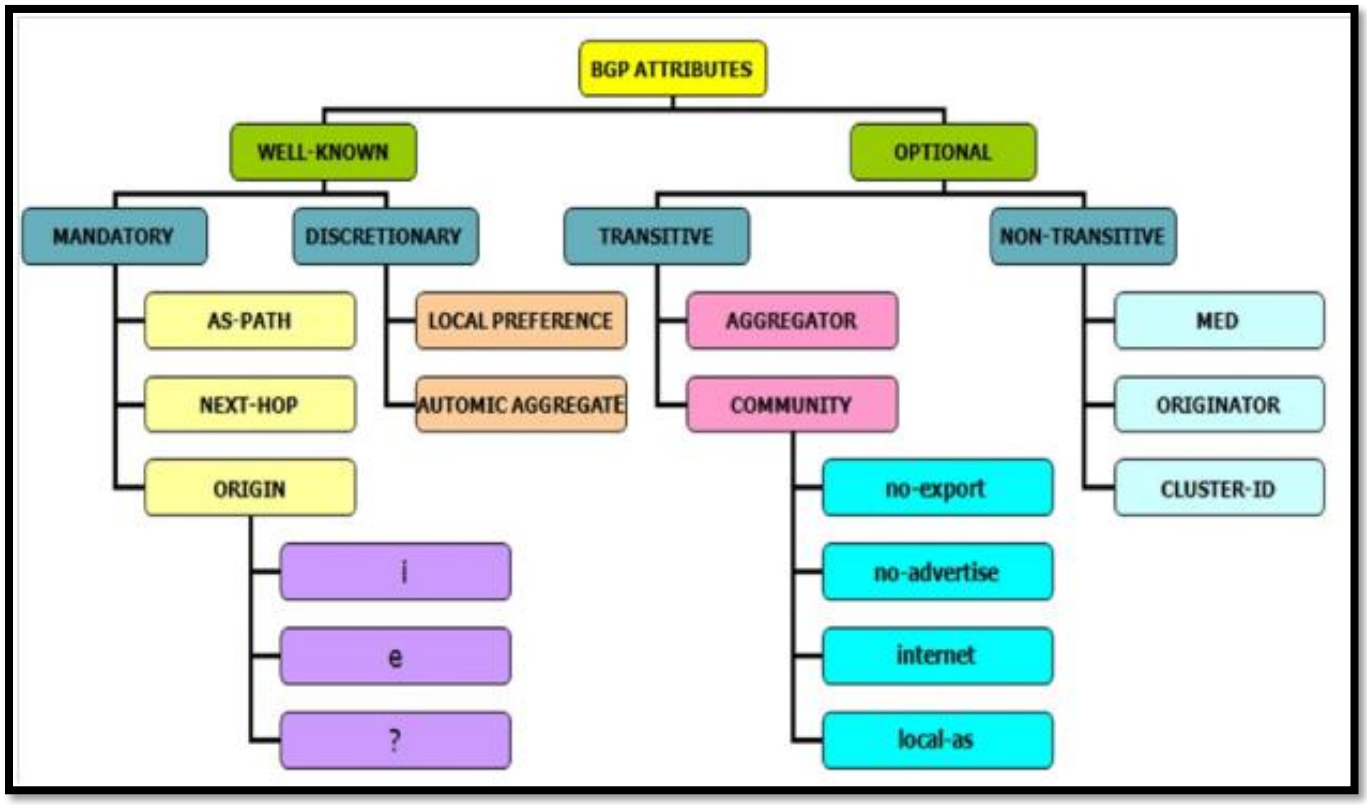

Well-Known Discretionary:

- Must be recognized & supported by all BGP speakers.
- May or may not appear in every BGP Update message.
- Does not have to be included in every BGP update message.

- Well-Known Discretionary attributes are **Local Preference & Atomic Aggregate.**

Optional Transitive:

- May or may not be supported by all BGP speakers.
- Will be passed on if not recognized by the receiver.
- The attribute should be accepted and passed along to other peers.
- Optional Transitive attributes are **Aggregator and Community**.
- Transitive, these attributes are across AS boundaries.

Optional Non-Transitive:

- May or may not be supported by all BGP speakers.
- Not required to pass on, may be safely ignored.
- The attribute should be ignored and not passed on to other peers.
- Optional Non-Transitive attributes are **MED, Originator ID and Cluster List.**
- Non-transitive, these attributes are restricted to the same AS.

Type Code	Attribute Name	Category
1	Origin	Well-Known Mandatory
2	AS Path	Well-Known Mandatory
3	Next Hop	Well-Known Mandatory
4	Multi Exit Desc (MED)	Optional Non-Transitive
5	Local Pref	Well-Known Discretionary
6	Atomic Aggregate	Well-Known Discretionary
7	Aggregator	Optional Transitive
8	Community	Optional Transitive
9	Originator ID	Optional Non-Transitive
10	Cluster List	Optional Non-Transitive

BGP Best Path Selection:

- BGP (Border Gateway Protocol) routers usually receive multiple paths to the same destination.
- Like how our IGPs (RIP, EIGRP, OSPF) work, we need to select the best path to each destination.
- IGPs select the path with the lowest metric.

For example:

- RIP selects the path with the lowest hop count.
- OSPF selects the path with the lowest cost.
- EIGRP selects the path with the highest bandwidth and lowest delay (unless you change the K values).
- BGP however, selects the best path based on a list of attributes.
- BGP sends update packet to a peer with path attributes associated with prefix.
- BGP selects the best path based on a list of attributes.
- BGP attributes are similar to metrics in OSPF and EIGRP.
- BGP use path attributes to pick the best route to a destination.
- BGP store multiple paths to a destination in BGP table.
- BGP only install one best route in the routing table.
- BGP path algorithm decides best path to install in the IP routing table.
- BGP path algorithm decides best path to use for traffic forwarding.
- BGP goes through the following steps to select the best path route.

Priority	Attribute	Preference
1	Weight	Highest
2	Local Preference	Highest
3	Originate Local	Local
4	AS Path Shortest	Shortest
5	Origin Code Lowest	Lowest

6	MED	Lowest
7	EBGP Path Over IBGP Path Prefer	Prefer eBGP
8	Shortest IGP Path to BGP Next Hop	Lowest IGP Metric
9	Oldest Path	Received First
10	Router ID	Lowest
11	Neighbor IP Address	Lowest Neighbor IP

Weight:

- Weight is the first BGP attribute in the list.
- Weight is Cisco-Proprietary value.
- Weight is only local on the router.
- Weight is not exchanged between BGP routers.
- Weight is never advertised to other routers.
- The path with the highest weight is preferred.
- Weight for a route originated on the local router is 32768.
- Weight is zero for all other routes.

BGP Weight Attribute Lab:

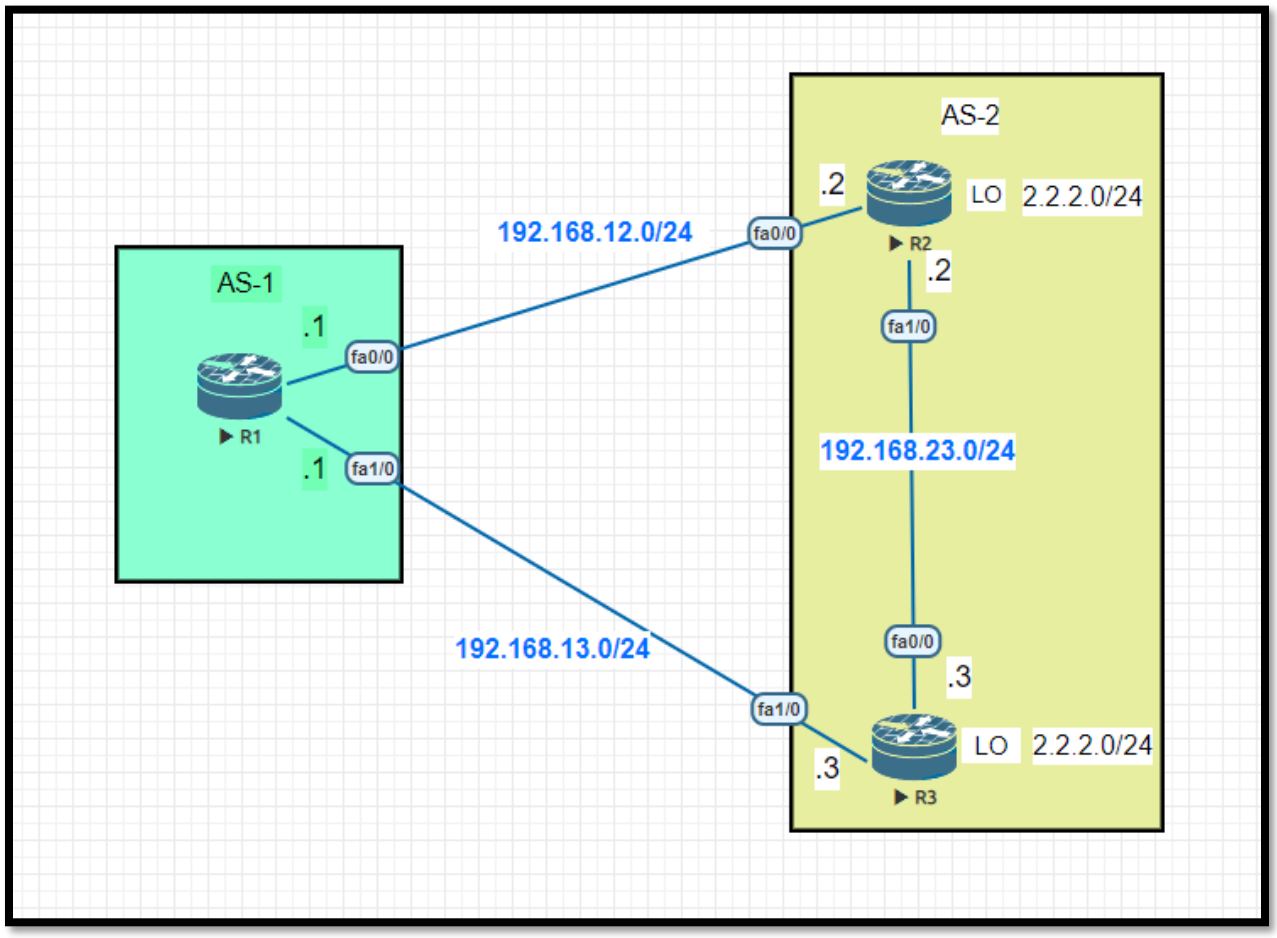

- Above we have a simple scenario with two autonomous systems.
- R2 and R3 both have network 2.2.2.0/24 configured on their loopback0 interface and we will advertise that in BGP.

Let see the configuration:

R1 Configuration:	R2 Configuration:	R3 Configuration:
en config t hostname R1	en config t hostname R2	en config t hostname R3

int f0/0 ip add 192.168.12.1 255.255.255.0 no sh int f1/0 ip add 192.168.13.1 255.255.255.0 no sh router bgp 1 bgp router-id 1.1.1.1 neighbor 192.168.12.2 remote-as 2 neighbor 192.168.13.3 remote-as 2	int f0/0 ip add 192.168.12.2 255.255.255.0 no sh int f1/0 ip add 192.168.23.2 255.255.255.0 no sh int lo 0 ip add 2.2.2.2 255.255.255.0 router bgp 2 bgp router-id 2.2.2.2 neighbor 192.168.12.1 remote-as 1 neighbor 192.168.23.3 remote-as 2 network 2.2.2.0 mask 255.255.255.0	int f0/0 ip add 192.168.23.3 255.255.255.0 no sh int f1/0 ip add 192.168.13.3 255.255.255.0 no sh int lo 0 ip add 2.2.2.2 255.255.255.0 router bgp 2 bgp router-id 3.3.3.3 neighbor 192.168.13.1 remote-as 1 neighbor 192.168.23.2 remote-as 2 network 2.2.2.0 mask 255.255.255.0

Let’s take a detailed look at R1:

```
R1#show ip bgp
BGP table version is 2, local router ID is 192.168.13.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  2.2.2.0/24       192.168.13.3             0             0 2 i
*>                  192.168.12.2             0             0 2 i
```

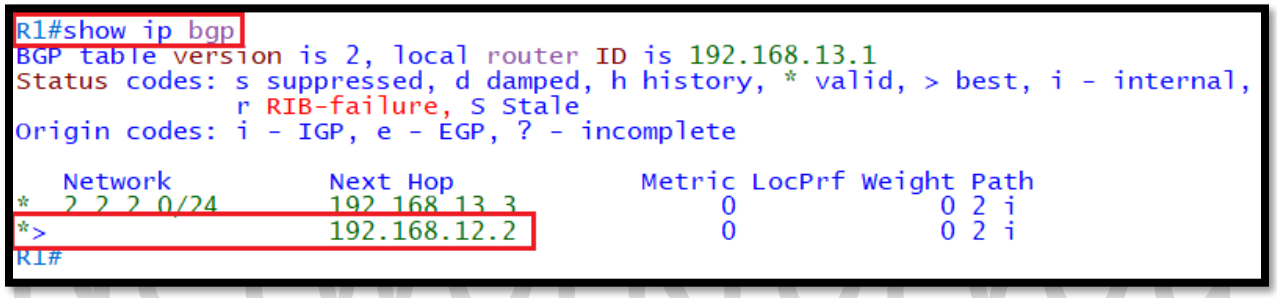

Router R1 decided to use 192.168.12.2 as the next hop. All the BGP attributes are the same so it came to the router ID to select a winner. Now let's change this behavior using the weight attribute.

```
R1(config)#router bgp 1
R1(config-router)#neighbor 192.168.13.3 weight 500
```

You can configure weight per neighbor using the weight command.

All prefixes from this neighbor will have a weight of 500.

```
R1#clear ip bgp *
```

```
R1#sh ip bgp
BGP table version is 2, local router ID is 192.168.13.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.2.2.0/24       192.168.13.3             0             500 2 i
*                   192.168.12.2             0               0 2 i
R1#
```

Now we can see that 192.168.13.3 has been selected as the next hop because the weight is now 500.

What if we want to set the weight to 500 for just a couple of prefixes from AS 2?

```
! adding lo 1 with new ip

R2(config)#interface loopback 1
R2(config-if)#ip address 22.22.22.22 255.255.255.0
R2(config)#router bgp 2
R2(config-router)#network 22.22.22.0 mask 255.255.255.0
```

```
! adding lo 1 with new ip

R3(config)#interface loopback 1
R3(config-if)#ip address 22.22.22.22 255.255.255.0
R3(config)#router bgp 2
R3(config-router)#network 22.22.22.0 mask 255.255.255.0
```

Now we create a new loopback interface on R2 and R3 we will advertise network 22.22.22.0/24 in BGP.

Let see what R1 see now.

```
R1#sh ip bgp
BGP table version is 5, local router ID is 192.168.13.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 2.2.2.0/24       192.168.13.3             0             500 2 i
*                   192.168.12.2             0               0 2 i
*> 22.22.22.0/24    192.168.13.3             0             500 2 i
*                   192.168.12.2             0               0 2 i
```

As we can see above router R1 willl use 192.168.13.3 as the next hop for both prefixes. What if we want to change the weight for just 1 prfix?

We will use Route-maps.

```
! Remove weight
R1(config)#router bgp 1
R1(config-router)#no neighbor 192.168.13.3 weight 500
```



```
!Route-Map
R1(config)#access-list 1 permit 22.22.22.0 0.0.0.255
R1(config)#route-map setweight permit 10
R1(config-route-map)#match ip address 1
R1(config-route-map)#set weight 500
R1(config-route-map)#exit
R1(config)#route-map setweight permit 20
R1(config-route-map)#set weight 0
R1(config-route-map)#exit
```

Here's the route-map that we will use. If the prefixes match access-list 1 we will set the weight to 500.

```
R1(config) router bgp 1
R1(config-router)#neighbor 192.168.13.3 route-map setweight in
```

To complete the configuration, we have to apply it to our neighbor in AS 2.

```
R1#clear ip bgp *
```

```
R1#sh ip bgp
BGP table version is 3, local router ID is 192.168.13.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  2.2.2.0/24       192.168.13.3             0             0 2 i
*>                  192.168.12.2             0             0 2 i
*> 22.22.22.0/24    192.168.13.3             0           500 2 i
*                   192.168.12.2             0             0 2 i
```

Path Selection:

- When BGP has multiple paths to a destination they are stored in the BGP table.
- All paths are in the BGP table but only one gets installed in the routing table.
- Which path do we select? We start at the top of the list with BGP attributes and work our way to the bottom:
- We start with weight because it's at the top of the BGP attributes list. We now have two options:
- If one path has a better weight then we select this path as the best path.
- If the weight is equal, we move down to the next attribute.
- The next attribute is local preference. Once again, we have two options:
- If one path has a better local preference then we select this path as the best path.
- If the local preference is equal, we move down to the next attribute.
- We work our way down this attribute list until we have a tiebreaker to select the best path. If all paths have the same BGP attributes, then we end up with the neighbor IP address.

Weight:

- Prefer the path with the highest weight.
- This is a value that is local to the router and it's Cisco proprietary.
- The default value is 0 for all routes that are not originated by the local router.

Local Preference:

- Local Preference is the second BGP attribute.
- The local preference is used within an autonomous system and exchanged between iBGP routers.
- We prefer the path with the highest local preference.
- The default value is 100.

Originate:

- Originate is the third BGP attribute.
- Prefer the path that the local router originated.
- In the BGP table, Local router originated see next hop 0.0.0.0.
- Routes with weight set to "32768" is considered as local routes.
- Path in the BGP table through Network Command, Aggregation, or Redistribution.
- BGP router will prefer routes that it installed into BGP itself to another router installed.

AS path length:

- AS Path is the fourth BGP attribute.
- Prefer the path with the shortest AS path length.
- For example, AS path 1 2 3 is preferred over AS path 1 2 3 4 5.

Origin code:

- Prefer the lowest origin code.
- There are three origin codes:
- IGP,EGP and INCOMPLETE
- IGP is lower than EGP and EGP is lower than Incomplete.
- IGP (shows up as i) use the network command for BGP.
- EGP (shows up as e) is an old routing protocol no more.
- Incomplete (shows up as ?) means redistributed something into BGP.
- IGP is lower than EGP and EGP is lower than INCOMPLETE.

MED:

- MED (Multi-Exit Discriminator) is the sixth BGP attribute.
- Multi-Exit Discriminator (MED) is optional non-transitive attribute.
- The lowest MED is the preferred path.
- The MED is exchanged between Autonomous Systems.
- MED is used to advertise the neighbors how to enter the AS.
- MED is propagated to all routers within the neighbor AS.
- MED is not passed along any other Autonomous Systems.
- MED can influence routers in the same AS but not on different AS.

eBGP path over iBGP path:

- Prefer eBGP (external BGP) over iBGP (internal BGP) paths.

Shortest IGP path to BGP next hop:

- Prefer the path within the autonomous system with the lowest IGP metric to the BGP next hop.

Oldest Path:

- Prefer the path that we received first, in other words, the oldest path.

Router ID:

- Prefer the path with the lowest BGP neighbor router ID.
- The router ID is based on the highest IP address.
- If you have a loopback interface, then the IP address on the loopback will be used.
- The router ID can also be manually configured.

Neighbor IP address:

- Prefer the path with the lowest neighbor IP address.
- If you have two eBGP routers and two links in between then the router ID will be the same.
- In this case, the neighbor IP address is the tiebreaker.

BGP Local Preference Lab:

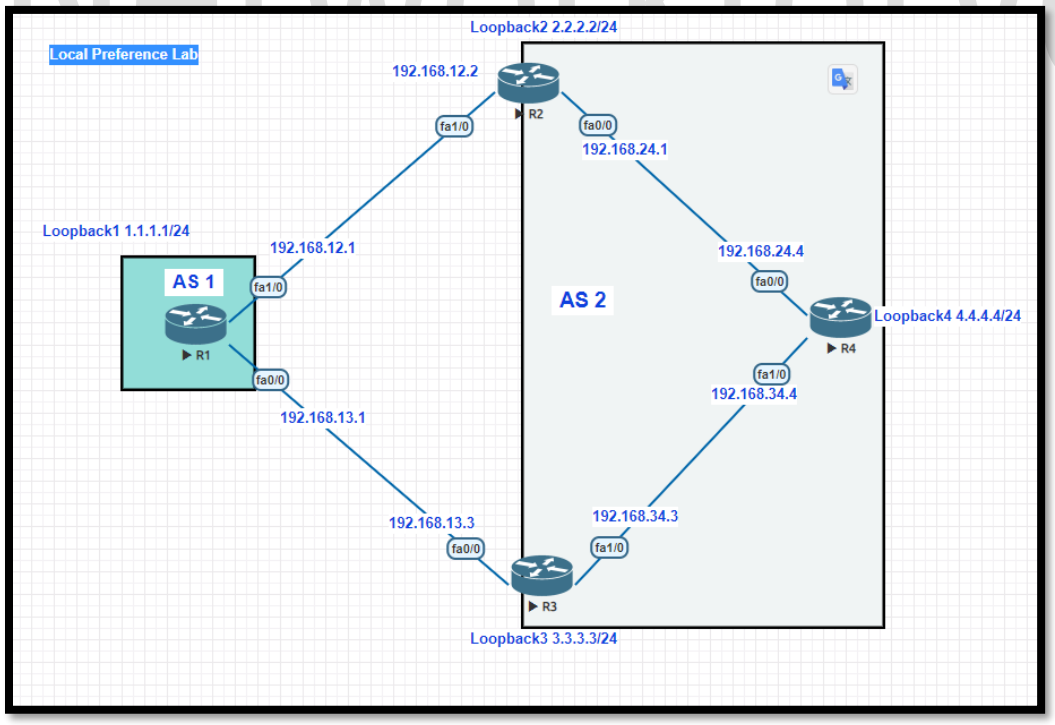

R1 Configuration	R2 Configuration
en config t hostname R1 int f1/0 ip add 192.168.12.1 255.255.255.0 no sh int f0/0 ip add 192.168.13.1 255.255.255.0 no sh int lo 1 ip add 1.1.1.1 255.255.255.0 router bgp 1 neighbor 192.168.12.2 remote-as 2 neighbor 192.168.13.3 remote-as 2 network 1.1.1.0 mask 255.255.255.0	en config t hostname R2 int f1/0 ip add 192.168.12.2 255.255.255.0 no sh int f0/0 ip add 192.168.24.1 255.255.255.0 no sh int lo 2 ip add 2.2.2.2 255.255.255.0 router ospf 1 int f0/0 ip ospf 1 area 0 int lo 2 ip ospf 1 area 0 router bgp 2 neighbor 192.168.12.1 remote-as 1 neighbor 3.3.3.3 remote-as 2 neighbor 3.3.3.3 update-source loopback2 neighbor 4.4.4.4 remote-as 2

	neighbor 4.4.4.4 update-source loopback2 neighbor 4.4.4.4 next-hop-self
R3 Configuration	**R4 Configuration**
en config t hostname R3 int f1/0 ip add 192.168.34.3 255.255.255.0 no sh int f0/0 ip add 192.168.13.3 255.255.255.0 no sh int lo 3 ip add 3.3.3.3 255.255.255.0 router ospf 1 int f1/0 ip ospf 1 area 0 int lo 3 ip ospf 1 area 0 router bgp 2 neighbor 192.168.13.1 remote-as 1 neighbor 2.2.2.2 remote-as 2 neighbor 2.2.2.2 update-source loopback3	en config t hostname R4 int f1/0 ip add 192.168.34.4 255.255.255.0 no sh int f0/0 ip add 192.168.24.4 255.255.255.0 no sh int lo 4 ip add 4.4.4.4 255.255.255.0 router ospf 1 int f1/0 ip ospf 1 area 0 int f0/0 ip ospf 1 area 0 int lo 4 ip ospf 1 area 0 router bgp 2 neighbor 2.2.2.2 remote-as 2

neighbor 4.4.4.4 remote-as 2 neighbor 4.4.4.4 update-source loopback3 neighbor 4.4.4.4 next-hop-self	neighbor 2.2.2.2 update-source loopback 4 neighbor 3.3.3.3 remote-as 2 neighbor 3.3.3.3 update-source loopback 4

We need to run some internet routing protocol between R2,R3 and R4 we run IGP (OSPF) to have connectivity.

In Router R4 in BGP summary it is showing 1.1.1.1 is coming from 192.168.12.1 (before we put next-hop-self) but it is not connected it should show 192.168.24.2(so for that purpose we need to use next-hop-self)

After putting next-hop-self command it is showing that 1.1.1.1 is coming from 3.3.3.3 and 2.2.2.2

```
R4#sh ip bgp
BGP table version is 2, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
* i1.1.1.0/24      3.3.3.3                  0    100      0 1 i
*>i               2.2.2.2                  0    100      0 1 i
```

Now we will increase local preference in R3 as 600 so that best router should change from R2 (2.2.2.2)to R3 (3.3.3.3).

So go to R3 and perform below mention commands.

```
R3(config)#router bgp 2
R3(config-router)#bgp default local-preference 600
```

Then check in R4

```
R4#sh ip bgp
BGP table version is 3, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.1.1.0/24       3.3.3.3                  0    600      0 1 i
* i                 2.2.2.2                  0    100      0 1 i
```

So if you see in R4 best router is showing with R3 (3.3.3.3).

Another way to configure local preference, first clear the configured

command.

```
R3(config)#router bgp 2
R3(config-router)#no bgp default local-preference 600
```

Now we will use route-map to do this.

```
R3(config)#route-map abc permit 10
R3(config-route-map)#set local-preference 700
R3(config)#router bgp 2
R3(config-router)#neighbor 192.168.13.1 route-map abc in
```

Now let see sh ip bgp in R4 router.

Bgp is slow protocol so if you are not getting then do clear ip bgp * in R4 and R3 and check but do not do this in production network. If you want to do hard reset for BGP try to get approval in production network.

```
R4#sh ip bgp
BGP table version is 5, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.1.1.0/24       3.3.3.3                  0    700      0 1 i
* i                 2.2.2.2                  0    100      0 1 i
```

But here we will have one issues like if any other network is coming from R1 that also automatically have local Preference 700 so to overcome with that we can do create ACL and apply in Route-map.

But before that remove above apply router map and neig.

```
R3(config)#no route-map abc permit 10
R3(config)#router bgp 2
R3(config-router)#no ne
R3(config-router)#no neigh 192.168.13.1 route-map abc in
```

Route map apply for that:

Let create one more loopback for R1

```
R1(config)#int lo 0
R1(config-if)#ip add 11.11.11.11 255.255.255.0
R1(config-if)#router bgp 1
R1(config-router)#network 11.11.11.0 mask 255.255.255.0
```

Now let see sh ip bgp:

```
R4#sh ip bgp
BGP table version is 14, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*>i1.1.1.0/24       2.2.2.2                  0    100      0 1 i
* i                 3.3.3.3                  0    100      0 1 i
*>i11.11.11.0/24    2.2.2.2                  0    100      0 1 i
* i                 3.3.3.3                  0    100      0 1 i
```



```
!Route-Map
R3(config)#access-list 1 permit 11.11.11.0 0.0.0.255
R3(config)#route-map abc permit 10
R3(config-route-map)#match ip address 1
R3(config-route-map)#set local-preference 500
R3(config-route-map)#exit
R3(config)#route-map abc permit 20
R3(config-route-map)#set local-preference 100
R3(config-route-map)#exit
```

Here's the route-map that we will use. If the prefixes match access-list 1 we will set the Local preference 500.

Then apply to that neighbor 192.168.12.1

(Here in example router map name we take abc u can take any name)

```
R3(config)#router bgp 2
R3(config-router)#neighbor 192.168.13.1 route-map abc in
```

Bgp is slow protocol so if you are not getting then do clear ip bgp * in R4 and R3 and check but do not do this in production network. If you want to do hard reset for BGP try to get approval in production network.

```
R4#sh ip bgp
BGP table version is 4, local router ID is 4.4.4.4
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
* i1.1.1.0/24      3.3.3.3                  0    100      0 1 i
*>i                2.2.2.2                  0    100      0 1 i
*>i11.11.11.0/24   3.3.3.3                  0    500      0 1 i
* i                2.2.2.2                  0    100      0 1 i
```

So if you see now we have two prefix from R1 our R4 is going with different Router like for 1.1.1.1 is going with R2 (2.2.2.2) and 11.11.11 is going with R3 (3.3.3.3) so with this way we can play with BGP.

AS-Path Prepending Lab:

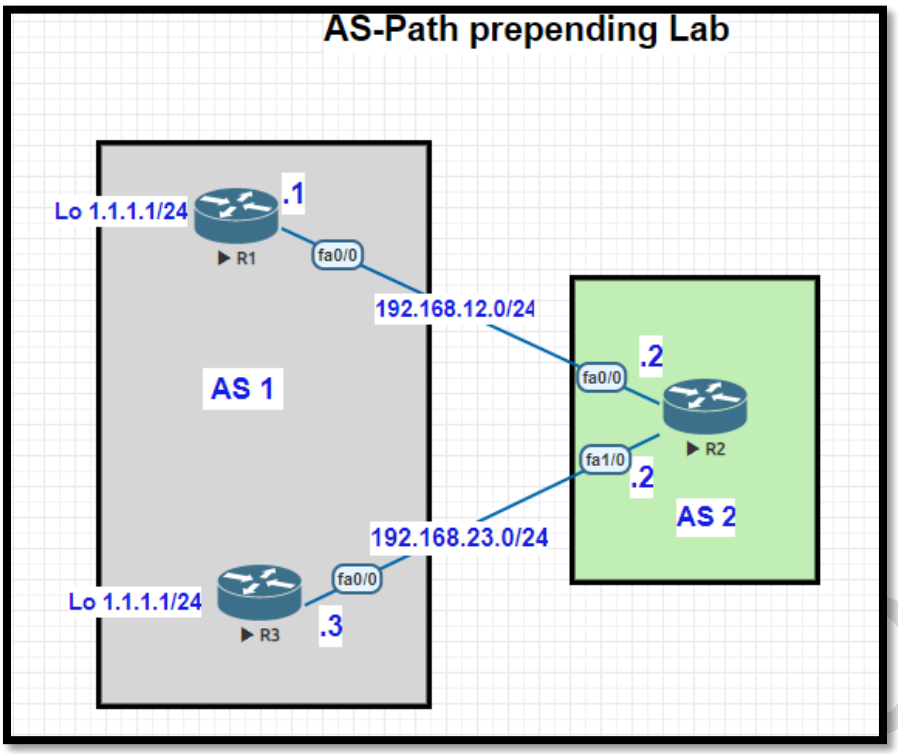

R1 Configuration	R2 Configuration	R3 Configuration
en config t hostname R1 int f0/0 ip add 192.168.12.1 255.255.255.0 no sh int lo 0 ip add 1.1.1.1 255.255.255.0	en config t hostname R2 int f0/0 ip add 192.168.12.2 255.255.255.0 no sh int f1/0 ip add 192.168.23.2 255.255.255.0	en config t hostname R3 int f0/0 ip add 192.168.23.3 255.255.255.0 no sh int lo 0 ip add 1.1.1.1 255.255.255.0

router bgp 1 neighbor 192.168.12.2 remote-as 2 network 1.1.1.0 mask 255.255.255.0	no sh router bgp 2 neighbor 192.168.12.1 remote-as 1 neighbor 192.168.23.3 remote-as 1	 router bgp 1 neighbor 192.168.23.2 remote-as 2 network 1.1.1.0 mask 255.255.255.0

```
R2#sh ip bgp
BGP table version is 2, local router ID is 192.168.23.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  1.1.1.0/24       192.168.23.3             0             0 1 i
*>                  192.168.12.1             0             0 1 i
R2#
```



```
!Route-Map
R1(config)#access-list 1 permit 1.1.1.0 0.0.0.255
R1(config)#route-map abc permit 10
R1(config-route-map)#match ip address 1
R1(config-route-map)#set as-path prepend 11111
R1(config-route-map)#exit
R1(config)#route-map abc permit 20
R1(config-route-map)#exit
```

Here's the route-map that we will use. If the prefixes match access-list 1 we will set the as-path prepend 11111.

Then apply to that neighbor 192.168.12.1

(Here in example router map name we take abc u can take any name)

```
R1(config)#router bgp 1
R1(config-router)#neighbor 192.168.12.2 route-map abc out
```

Bgp is slow protocol so if you are not getting then do clear ip bgp * in R2 and R1 and check but do not do this in production network. If you want to do hard reset for BGP try to get approval in production network.


```
R2#sh ip bgp
BGP table version is 2, local router ID is 192.168.23.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.1.0/24       192.168.23.3             0             0 1 i
*                   192.168.12.1             0             0 1 111111 i
```

BGP Origin Lab:

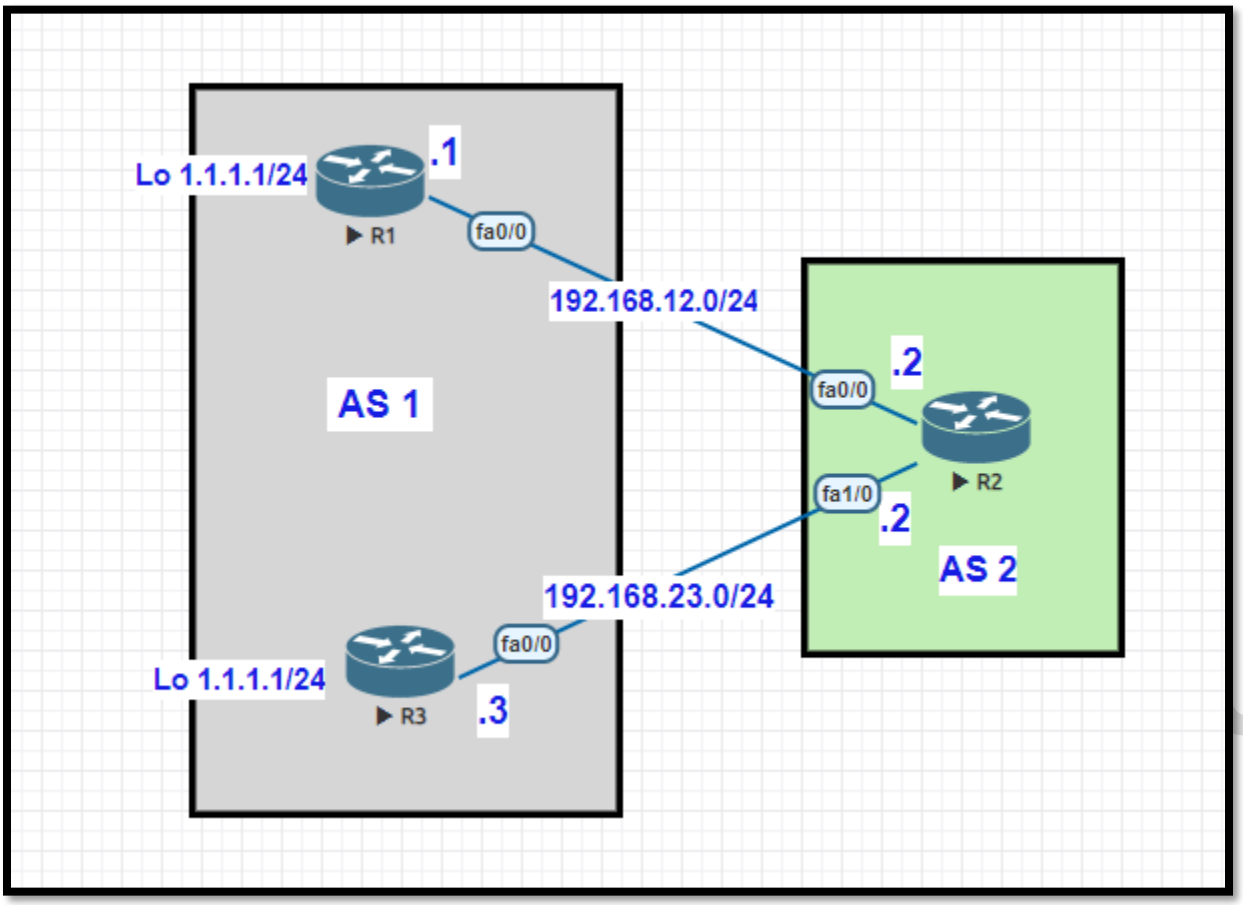

Before redistribute connected:

```
R2#sh ip bgp
BGP table version is 5, local router ID is 192.168.23.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.1.0/24       192.168.12.1             0             0 1 i
```

R1 Configuration	R2 Configuration	R3 Configuration
en config t hostname R1 int f0/0 ip add 192.168.12.1 255.255.255.0 no sh int lo 0 ip add 1.1.1.1 255.255.255.0 router bgp 1 neighbor 192.168.12.2 remote-as 2 network 1.1.1.0 mask 255.255.255.0	en config t hostname R2 int f0/0 ip add 192.168.12.2 255.255.255.0 no sh int f1/0 ip add 192.168.23.2 255.255.255.0 no sh router bgp 2 neighbor 192.168.12.1 remote-as 1 neighbor 192.168.23.3 remote-as 1	en config t hostname R3 int f0/0 ip add 192.168.23.3 255.255.255.0 no sh int lo 0 ip add 1.1.1.1 255.255.255.0 router bgp 1 neighbor 192.168.23.2 remote-as 2 redistribute connected

After redistribute connected

```
R2#sh ip bgp
BGP table version is 7, local router ID is 192.168.23.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  1.1.1.0/24       192.168.23.3             0             0 1 ?
*>                  192.168.12.1             0             0 1 i
```

BGP MED Attribute Lab:

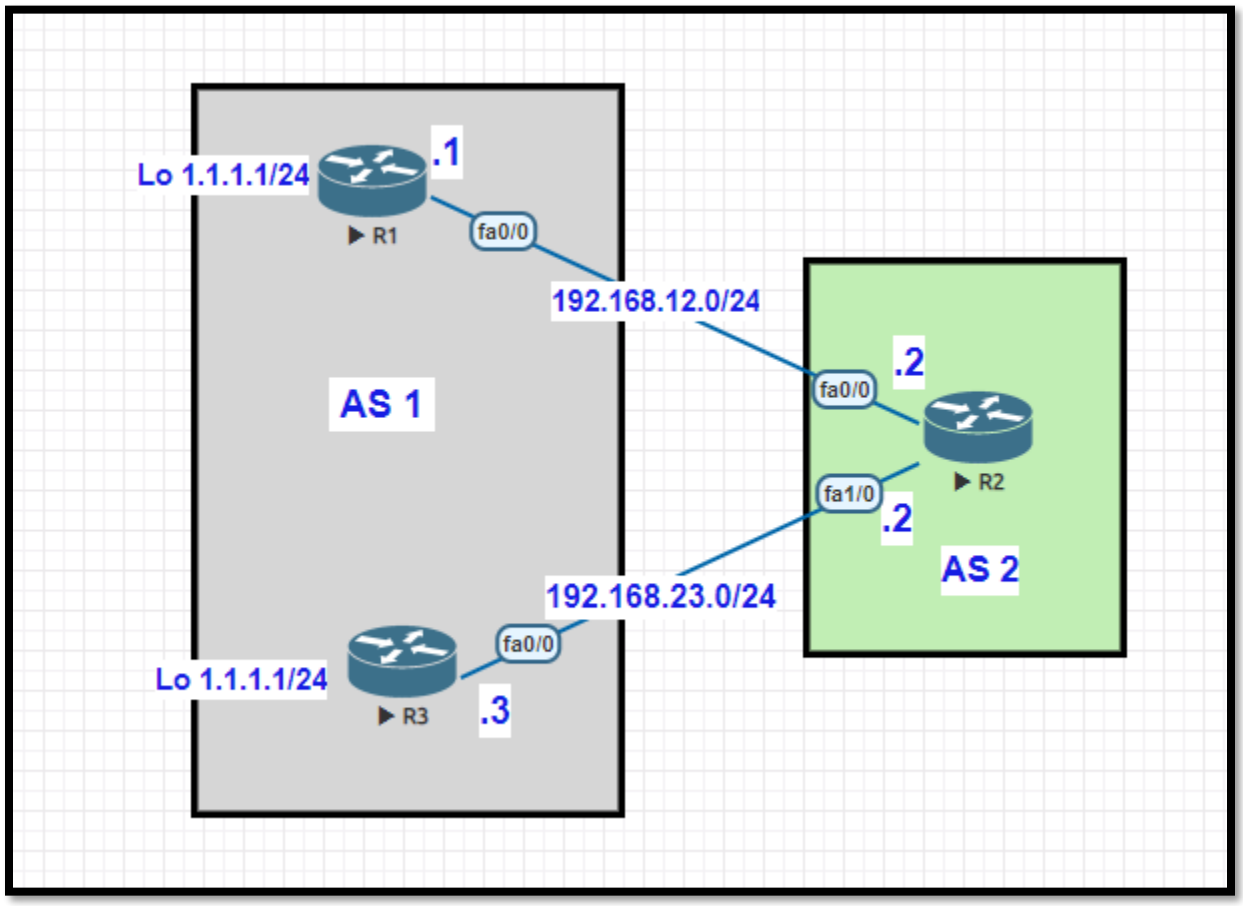

R1 Configuration	R2 Configuration	R3 Configuration
en config t hostname R1 int f0/0 ip add 192.168.12.1 255.255.255.0	en config t hostname R2 int f0/0 ip add 192.168.12.2 255.255.255.0	en config t hostname R3 int f0/0 ip add 192.168.23.3 255.255.255.0

no sh int lo 0 ip add 1.1.1.1 255.255.255.0 router bgp 1 neighbor 192.168.12.2 remote-as 2 network 1.1.1.0 mask 255.255.255.0	no sh int f1/0 ip add 192.168.23.2 255.255.255.0 no sh router bgp 2 neighbor 192.168.12.1 remote-as 1 neighbor 192.168.23.3 remote-as 1	no sh int lo 0 ip add 1.1.1.1 255.255.255.0 router bgp 1 neighbor 192.168.23.2 remote-as 2 network 1.1.1.0 mask 255.255.255.0

Without MED value:

```
R2#sh ip bgp
BGP table version is 2, local router ID is 192.168.23.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*  1.1.1.0/24       192.168.23.3             0             0 1 i
*>                  192.168.12.1             0             0 1 i
```

R1:

route-map MED permit 10
set metric 700
exit
router bgp 1
neighbor 192.168.12.2 route-map MED out

After MED value increase of R1

```
R2#sh ip bgp
BGP table version is 3, local router ID is 192.168.23.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop            Metric LocPrf Weight Path
*> 1.1.1.0/24       192.168.23.3             0             0 1 i
*                   192.168.12.1           700             0 1 i
```

BGP Neighbor Relationship:

- BGP neighbors are routers forming TCP connection for exchanging BGP updates.
- Border gateway protocol (BGP) neighbors also called BGP peers or BGP speakers.
- There are two types of BGP neighbor relationship internal BGP and External BGP.
- BGP first forms neighbor relationship with Border gateway protocol Speakers (Peer).
- Than its learns information from its peer, BGP place that information in its BGP table.
- Local Router's Autonomous System Number must match neighboring Router's ASN.
- Ensure that Border gateway protocol (BGP) peers have IP reachability to each other.
- Configure basic External Border gateway protocol (eBGP) on each Cisco Router.
- Transmission Control Protocol (TCP) three 3-Way Handshake must be completed.
- The Border Gateway Protocol Peering must be completed & BGP Update Exchange.
- Source IP address of incoming TCP connection must be from configured BGP peer.

- Peer's advertisement of his Border gateway protocol AS# must be what we expect.
- If Border gateway protocol authentication is used, same password must configured.
- The BGP peers must have unique Border Gateway Protocol (BGP) Router-IDs.
- The BGP Peers must use the same Border Gateway Protocol (BGP) version only.

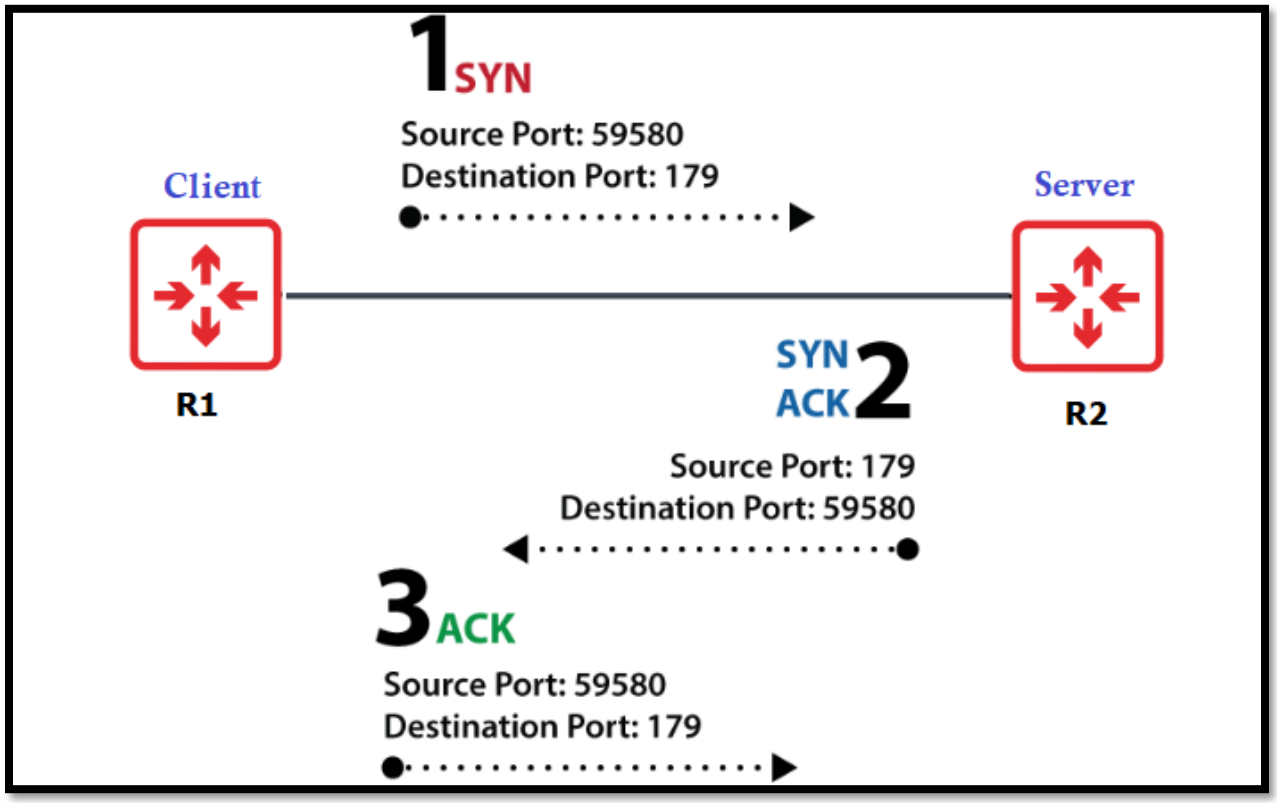

Overview of Internal BGP (iBGP) and External BGP (eBGP):

Router BGP 1 **Neighbor 2.2.2.2 remote-as 2**	**eBGP Peering**
Router BGP 1 **Neighbor 2.2.2.2 remote-as 1**	**iBGP Peering**

Let us check the BGP Peering Overview in below image (Lab):

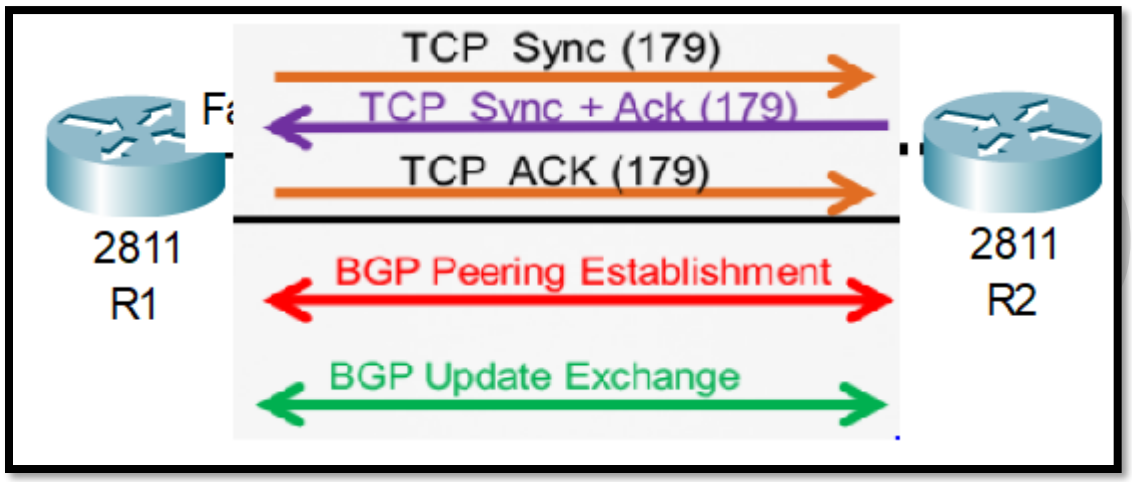

R1 Configuration	R2 Configuration
en config t hostname R1 int e0/0 ip add 192.168.12.1 255.255.255.0 no sh	en config t hostname R2 int e0/0 ip add 192.168.12.2 255.255.255.0 no sh

int lo 0 ip add 1.1.1.1 255.0.0.0 router bgp 1 neighbor 192.168.12.2 remote-as 2 network 1.0.0.0 mask 255.0.0.0 neighbor 192.168.12.2 password cisco	int lo 0 ip add 2.2.2.2 255.0.0.0 router bgp 2 neighbor 192.168.12.1 remote-as 1 network 2.0.0.0 mask 255.0.0.0 neighbor 192.168.12.1 password cisco

Keep the Same Router ID in R1
Router bgp 1 Bgp router-id 2.2.2.2

Mismatch Authentication in R1
Router bgp 1 Neighbor 192.168.12.2 password abc

BGP Authentication:

- Easiest ways to reduce security risks on BGP network is to use BGP peer authentication.
- The BGP supports authentication mechanism using the Message Digest 5 (MD5) algorithm.
- When authentication is enabled, TCP segment belonging to BGP exchanged between peers.
- TCP segment belonging to BGP is verified and accepted only if authentication is successful.
- For authentication to be successful, both peers must be configured with same password.
- If the authentication fails, BGP Protocol neighbor relationship is not being established.
- If the authentication fails, BGP neighbor relationship goes down and not be established.

Lab Time:

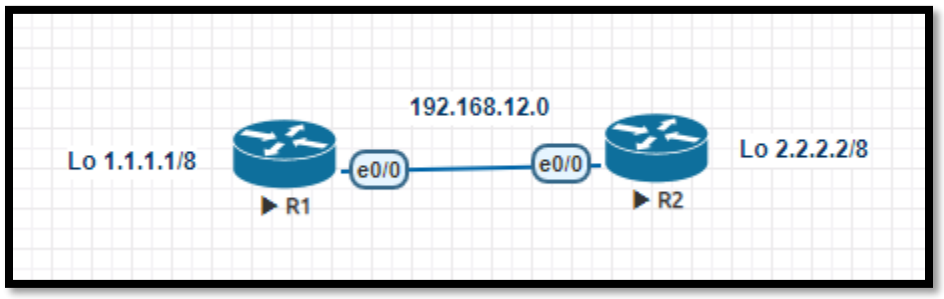

R1 Configuration	R2 Configuration
en config t hostname R1 int e0/0	en config t hostname R2 int e0/0

ip add 192.168.12.1 255.255.255.0 no sh int lo 0 ip add 1.1.1.1 255.0.0.0 router bgp 1 neighbor 192.168.12.2 remote-as 2 network 1.0.0.0 mask 255.0.0.0 neighbor 192.168.12.2 password cisco	ip add 192.168.12.2 255.255.255.0 no sh int lo 0 ip add 2.2.2.2 255.0.0.0 router bgp 2 neighbor 192.168.12.1 remote-as 1 network 2.0.0.0 mask 255.0.0.0 neighbor 192.168.12.1 password cisco

BGP Synchronization:

- BGP will not advertise something that it learns from IBGP neighbor to EBGP neighbor.
- If the prefix can't be validated in its IGP and by default, BGP synchronization is disabled.
- The BGP Synchronization mostly applies to interactions between two (2) separate AS.
- No synchronization command tells Routers that don't want them to "synchronize" iBGP.
- A Route learned via BGP will not be used nor advertised to an external neighbor BGP.
- And unless that same prefix is learned via an Interior Gateway Protocol (IGP) as well.
- A BGP border Router will not propagate a BGP-learned prefix to an eBGP neighbor.
- Unless that same IP prefix has been learned via an Interior Gateway Protocol (IGP).
- The BGP synchronization rule refers to prefix synchronization between BGP and IGP.
- If it is enabled, a BGP speakers will not advertise routes learned from an iBGP peer.
- Unless the destination described in the route is also reachable through the local IGP.
- **When synchronization is enabled then IBGP whatever routes learn from IBGP neighbor will not give to EBGP.**

Lab Time Synchronization:

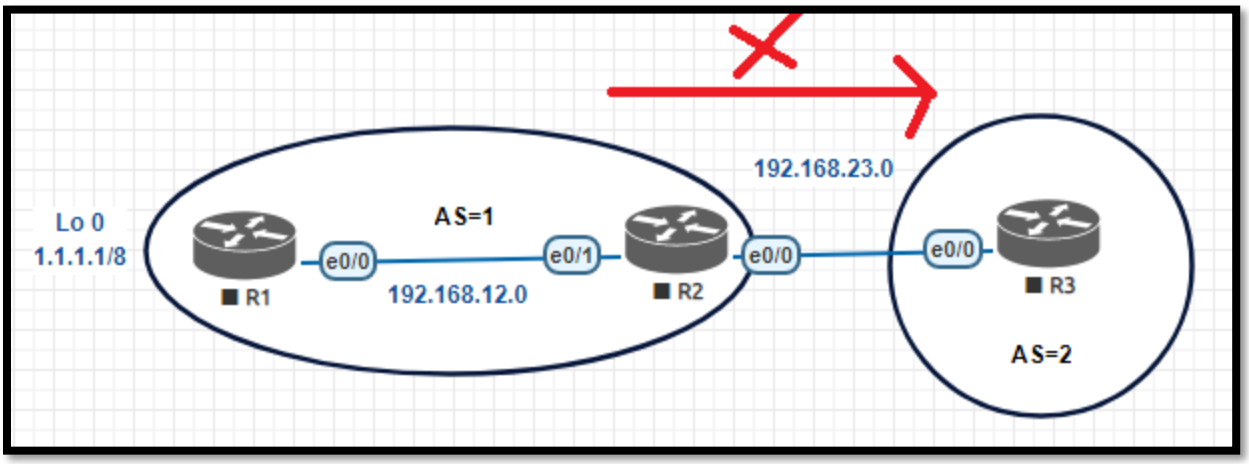

R1 Configuration	R2 Configuration	R3 Configuration
en config t hostname R1 int e0/0 ip add 192.168.12.1 255.255.255.0 no sh int lo 0 ip add 1.1.1.1 255.0.0.0 router bgp 1 neighbor 192.168.12.2 remote-as 1 network 1.0.0.0 mask 255.0.0.0	en config t hostname R2 int e0/1 ip add 192.168.12.2 255.255.255.0 no sh int e0/0 ip add 192.168.23.2 255.255.255.0 no sh router bgp 1 neighbor 192.168.12.1 remote-as 1 neighbor 192.168.23.3 remote-as 2	en config t hostname R3 int e0/0 ip add 192.168.23.3 255.255.255.0 no sh router bgp 2 neighbor 192.168.23.2 remote-as 1

In R1 BGP configuration BGP Synchronization is disable by default.

```
R1#sh running-config | section bgp
router bgp 1
 bgp log-neighbor-changes
 network 1.0.0.0
 neighbor 192.168.12.2 remote-as 1
```

In R2 BGP configuration BGP Synchronization is disable by default.

```
R2#sh running-config | sec bgp
router bgp 1
 bgp log-neighbor-changes
 neighbor 192.168.12.1 remote-as 1
 neighbor 192.168.23.3 remote-as 2
```

R3 BGP configuration BGP Synchronization is disable by default.

```
R3#sh running-config | section  bgp
router bgp 2
 bgp log-neighbor-changes
 neighbor 192.168.23.2 remote-as 1
```

BGP Synchronization is disable by default on R1 and R2 that's why R3 get R1 Routes 1.0.0.0

```
R3#sh ip route
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

B     1.0.0.0/8 [20/0] via 192.168.23.2, 00:08:26
      192.168.23.0/24 is variably subnetted, 2 subnets, 2 masks
C        192.168.23.0/24 is directly connected, Ethernet0/0
L        192.168.23.3/32 is directly connected, Ethernet0/0
```

In R2 also getting R1 route in Routing table and BGP table as well.

```
R2#sh ip route bgp
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

B    1.0.0.0/8 [200/0] via 192.168.12.1, 00:09:54
```

```
R2#sh ip bgp
BGP table version is 2, local router ID is 192.168.23.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>i 1.0.0.0          192.168.12.1             0    100      0 i
```

Let enable BGP Synchronization:

R1 Configuration	R2 Configuration	R3 Configuration
Router bgp 1 synchronization	Router bgp 1 synchronization	Router bgp 2 synchronization

R3 is not receiving anymore R1 1.0.0.0 network in Routing table nor in BGP table.

```
R3#sh ip route bgp
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set
```

R2 is still receiving R1 route 1.0.0.0 in BGP table but not in Routing Table anyhow, also R2 stop to advertise R1 routes to R3 Router.

```
R2#sh ip bgp
BGP table version is 1, local router ID is 192.168.23.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 * i 1.0.0.0          192.168.12.1             0    100      0 i
R2#
```

```
R2#sh ip route bgp
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set
```

iBGP the split-horizon rules states that an iBGP router will not advertise iBGP routes to other iBGP neighbors.

iBGP requires full mesh between all BGP Speaking routers ie.

- Large number of TCP sessions.
- Duplicate routing traffic.
- High processing
- High memory usage.
- Delay Stability

Solutions:

1. Route reflector modify iBGP split horizon rules
2. BGP Confederation modify iBGP AS path processing.

BGP Route Reflector (RR):

- By default, BGP speakers does not advertise iBGP-learned prefixes to iBGP peers.
- Due to the BGP split horizon rule all iBGP peers within an AS must be fully meshed.
- Any routes learn from iBGP neighbor never be advertised to any other iBGP neighbor.
- A route reflector is BGP router that is allowed to break the iBGP loop avoidance rule.
- Route reflectors can advertise updates received from iBGP peer to another iBGP peer.
- Route reflectors are one method to get rid of full-mesh of IBGP peers in your network.
- A route reflector is a BGP router that is allowed to ignore the IBGP loop avoidance rule.

- Route Reflector allowed to advertise updates received from an IBGP peer to another IBGP peer in AS.
- Route Reflector allows all IBGP speakers within autonomous network to learn about available routes.
- Route Reflector allow all IBGP speakers to learn about route without introducing loops.
- Route reflectors allows BGP speaker to advertise IBGP learn routes to certain IBGP peers.

Lab time Route reflectors:

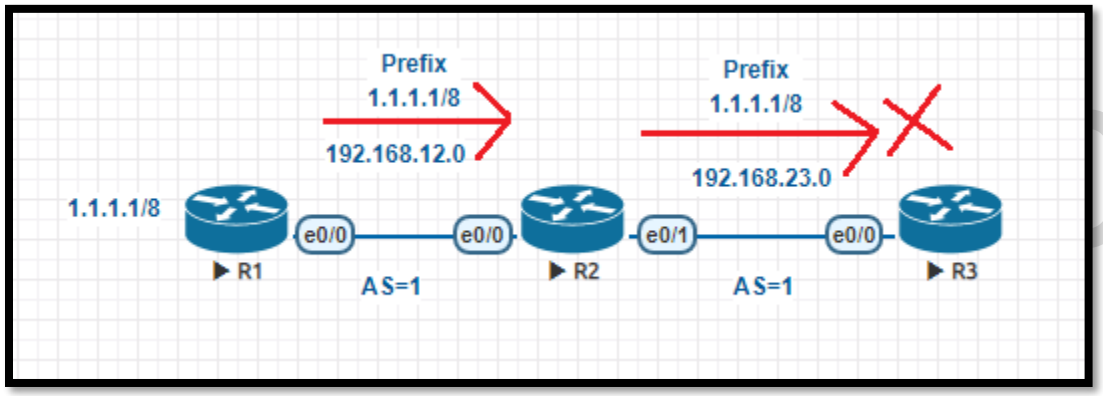

R1 Configuration:	R2 Configuration:
en config t hostname R1 int e0/0 ip add 192.168.12.1 255.255.255.0 no sh	en config t hostname R2 int e0/0 ip add 192.168.12.2 255.255.255.0 no sh

int lo 0 ip add 1.1.1.1 255.0.0.0 router bgp 1 neighbor 192.168.12.2 remote-as 1 network 1.0.0.0 mask 255.0.0.0	int e0/1 ip add 192.168.23.2 255.255.255.0 no sh router bgp 1 neighbor 192.168.12.1 remote-as 1 neighbor 192.168.23.3 remote-as 1
R3 Configuration:	Output we can check
en config t hostname R3 int e0/0 ip add 192.168.23.3 255.255.255.0 no sh router bgp 1 neighbor 192.168.23.3 remote-as 1	Sh ip route bgp

R2 Receiving R1 Route 1.0.0.0 in BGP Table also in the Routing table.

```
R2#sh ip bgp
BGP table version is 2, local router ID is 192.168.23.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>i 1.0.0.0          192.168.12.1             0    100      0 i
```

```
R2#sh ip route bgp
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set

B     1.0.0.0/8 [200/0] via 192.168.12.1, 00:00:20
```

R3 is not receiving Route 1.0.0.0 which R2 Learn from R1 in BGP Table not in Routing table.

```
R3#sh ip route bgp
Codes: L - local, C - connected, S - static, R - RIP, M - mobile, B - BGP
       D - EIGRP, EX - EIGRP external, O - OSPF, IA - OSPF inter area
       N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
       E1 - OSPF external type 1, E2 - OSPF external type 2
       i - IS-IS, su - IS-IS summary, L1 - IS-IS level-1, L2 - IS-IS level-2
       ia - IS-IS inter area, * - candidate default, U - per-user static route
       o - ODR, P - periodic downloaded static route, H - NHRP, l - LISP
       a - application route
       + - replicated route, % - next hop override

Gateway of last resort is not set
```

```
R3#sh ip bgp summary
BGP router identifier 192.168.23.3, local AS number 1
BGP table version is 1, main routing table version 1

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.23.2    4            1      11      11        1    0    0 00:06:39        0
```

After enable Route Reflector now R3 receiving R1 network 1.0.0.0 in BGP Table.

Enable Route Reflector on R2
router bgp 1 neighbor 192.168.23.3 route-reflector-client

```
R3#sh ip bgp sum
BGP router identifier 192.168.23.3, local AS number 1
BGP table version is 1, main routing table version 1
1 network entries using 140 bytes of memory
1 path entries using 80 bytes of memory
1/0 BGP path/bestpath attribute entries using 144 bytes of memory
1 BGP rrinfo entries using 24 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 388 total bytes of memory
BGP activity 1/0 prefixes, 1/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.23.2    4            1      14      13        1    0    0 00:06:35        1
```

```
R3#sh ip bgp
BGP table version is 1, local router ID is 192.168.23.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 * i 1.0.0.0          192.168.12.1             0    100      0 i
```

For reachability need to configure any IGP protocol on all Three Routers to work.

Configure IGP Protocol in R1:
router eigrp 1 network 0.0.0.0
Configure IGP Protocol in R2:
router eigrp 1 network 0.0.0.0

Configure IGP Protocol in R3:
router eigrp 1 network 0.0.0.0

```
R3#sh ip bgp
BGP table version is 2, local router ID is 192.168.23.3
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 r>i 1.0.0.0          192.168.12.1             0    100      0 i
```

BGP Confederation:

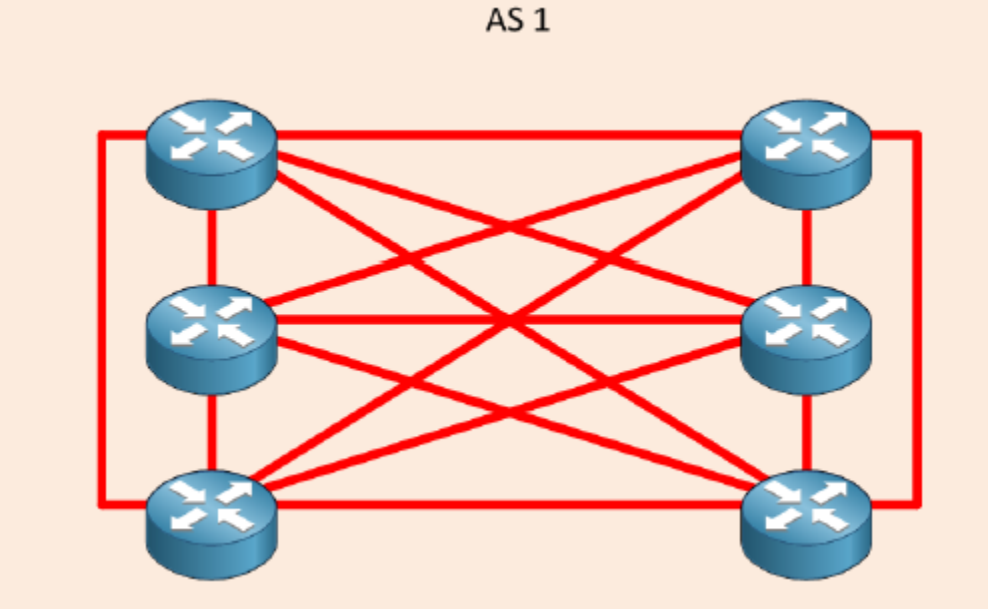

- Above we have AS 1 with 6 routers running IBGP.
- The number of IBGP peering's can be calculated with the full mesh formula:

N(N-1)/2

So in our case that's:

6 * (6-1 = 5) / 2 = 15 IBGP peering's.

- A BGP confederation divides our AS into sub-ASes to reduce the number of required IBGP peerings.
- Within a sub-AS we still require full-mesh IBGP but between these sub-ASes we use something that looks like EBGP but behaves like IBGP called confederation BGP.
- Here's an example of what a BGP confederation could look like:

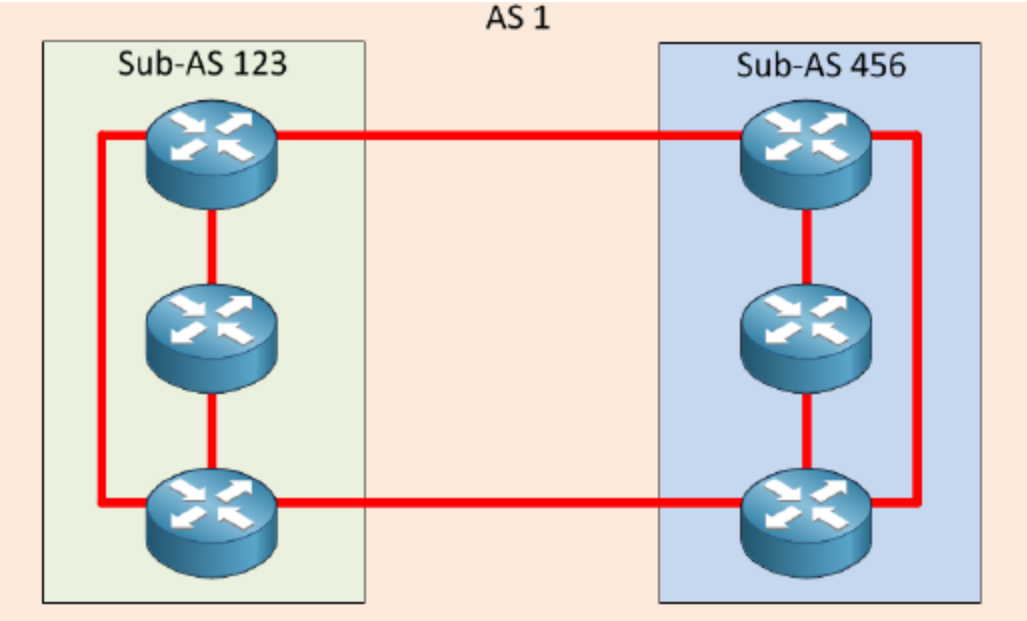

- By dividing our main AS into two sub-ASes we reduced the number of IBGP peerings from 15 to 8.
- Within the sub-AS we still have the full-mesh IBGP requirement. Between sub-ASes it's just like EBGP, it's up to you how many peerings you want.
- The outside world will never see your sub-AS numbers, they will only see the main AS number.

- Since the sub-AS numbers are not seen outside of your network you will often see private AS numbers used for the sub-ASes (64512 – 65535) but you can pick any number you like.
- **Or in other words we can say BGP Confederations is feature used to split an autonomous system into smaller autonomous systems.**
- Confederations are usable only for huge autonomous systems where you can afford to split them into several sub-Ases.
- Note that each sub-AS in a confederation needs to have its internal IBGP peer either fully meshed, or use route reflection internally.
- As we can see, the confederations are not much of an advantage for small ASes having a few BGP routers.

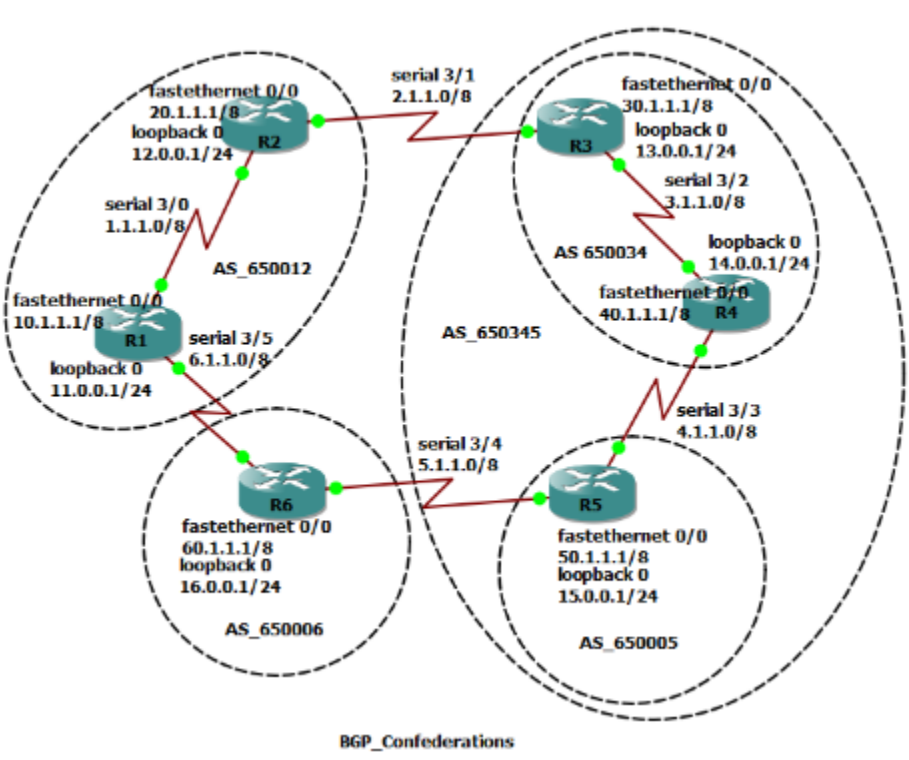

Lab Time: BGP Confederation

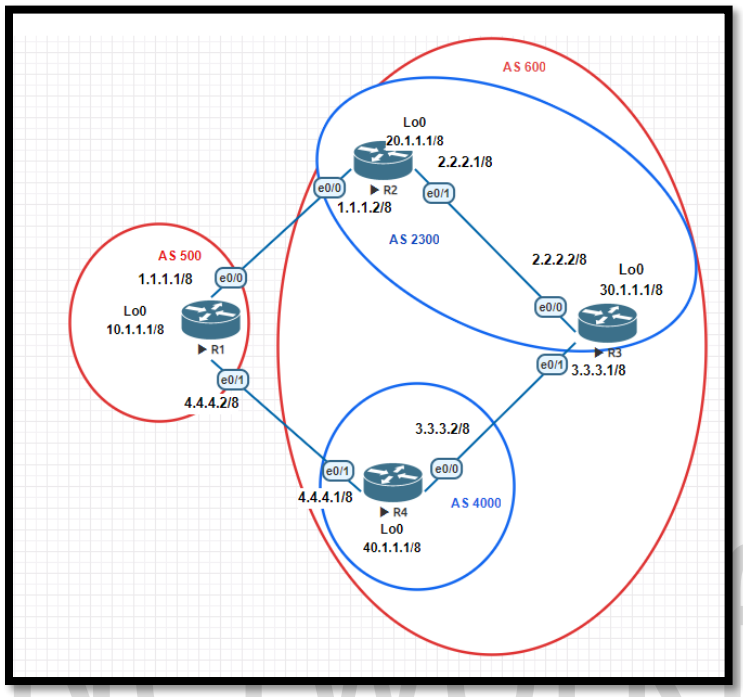

R1 Configuration	R2 Configuration
en config t hostname R1 int e0/0 ip add 1.1.1.1 255.0.0.0 no sh int e0/1 ip add 4.4.4.2 255.0.0.0 no sh int lo 0 ip add 10.1.1.1 255.0.0.0	en config t hostname R2 int e0/0 ip add 1.1.1.2 255.0.0.0 no sh int e0/1 ip add 2.2.2.1 255.0.0.0 no sh int lo 0 ip add 20.1.1.1 255.0.0.0

router bgp 500 neighbor 1.1.1.2 remote-as 600 neighbor 4.4.4.1 remote-as 600 network 10.0.0.0 mask 255.0.0.0	router bgp 2300 bgp confederation identifier 600 neighbor 1.1.1.1 remote-as 500 neighbor 2.2.2.2 remote-as 2300 network 20.0.0.0 mask 255.0.0.0
R3 Configuration	**R4 Configuration**
en config t hostname R3 int e0/0 ip add 2.2.2.2 255.0.0.0 no sh int e0/1 ip add 3.3.3.1 255.0.0.0 no sh int lo 0 ip add 30.1.1.1 255.0.0.0 router bgp 2300 bgp confederation identifier 600 bgp confederation peers 4000 neighbor 2.2.2.1 remote-as 500 neighbor 3.3.3.2 remote-as 4000 network 30.0.0.0 mask 255.0.0.0	en config t hostname R4 int e0/1 ip add 4.4.4.1 255.0.0.0 no sh int e0/0 ip add 3.3.3.2 255.0.0.0 no sh int lo 0 ip add 40.1.1.1 255.0.0.0 router bgp 4000 bgp confederation identifier 600 bgp confederation peers 2300 neighbor 3.3.3.1 remote-as 2300 neighbor 4.4.4.2 remote-as 500 network 40.0.0.0 mask 255.0.0.0

```
R3#sh ip bgp
BGP table version is 4, local router ID is 30.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 * i 10.0.0.0         1.1.1.1                  0    100      0 500 i
 *                    4.4.4.2                  0    100      0 (4000) 500 i
 *>i 20.0.0.0         2.2.2.1                  0    100      0 i
 *>  30.0.0.0         0.0.0.0                  0         32768 i
 *>  40.0.0.0         3.3.3.2                  0    100      0 (4000) i
```

```
R4#sh ip bgp
BGP table version is 4, local router ID is 40.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  10.0.0.0         4.4.4.2                  0             0 500 i
 *   20.0.0.0         2.2.2.1                  0    100      0 (2300) i
 *>  30.0.0.0         3.3.3.1                  0    100      0 (2300) i
 *>  40.0.0.0         0.0.0.0                  0         32768 i
```

```
R3#sh ip bgp sum
BGP router identifier 30.1.1.1, local AS number 2300
BGP table version is 4, main routing table version 4
4 network entries using 560 bytes of memory
5 path entries using 400 bytes of memory
5/3 BGP path/bestpath attribute entries using 720 bytes of memory
3 BGP AS-PATH entries using 72 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 1752 total bytes of memory
BGP activity 4/0 prefixes, 5/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
2.2.2.1         4         2300      13      13        4    0    0 00:06:50        2
3.3.3.2         4         4000      15      14        4    0    0 00:09:51        2
```

```
R4#sh ip bgp sum
BGP router identifier 40.1.1.1, local AS number 4000
BGP table version is 4, main routing table version 4
4 network entries using 560 bytes of memory
4 path entries using 320 bytes of memory
3/3 BGP path/bestpath attribute entries using 432 bytes of memory
2 BGP AS-PATH entries using 48 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 1360 total bytes of memory
BGP activity 4/0 prefixes, 4/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
3.3.3.1         4         2300      16      18        4    0    0 00:12:05        2
4.4.4.2         4          500      19      17        4    0    0 00:12:04        1
```

BGP Route Refresh:

Here are the 3 options that we have to refresh our BGP table when our policy changes:

- Hard reset
- Soft reconfiguration
- Route refresh capability

Hard rest:

- The hard reset is the simplest method (clear ip bgp * command).
- It kills the TCP session with your BGP neighbor which forces it to restart and as a result you'll receive all prefixes from your neighbor again.
- It works but will interrupt your network, not a good idea.

Soft reconfiguration:

- The soft reconfiguration will store everything that you receive from a BGP neighbor in a separate table before applying the policy.
- This works but it's not very efficient. Our router will store an entire table for each BGP neighbor with the unmodified prefixes, you'll need extra memory.

Route refresh capability:

- Route refresh capability is the most preferred method. When you change your BGP policy you just send a message to your BGP neighbor and it will resend you all its prefixes, there will be no disruption at all.

BGP Route Refresh Capability Lab time:

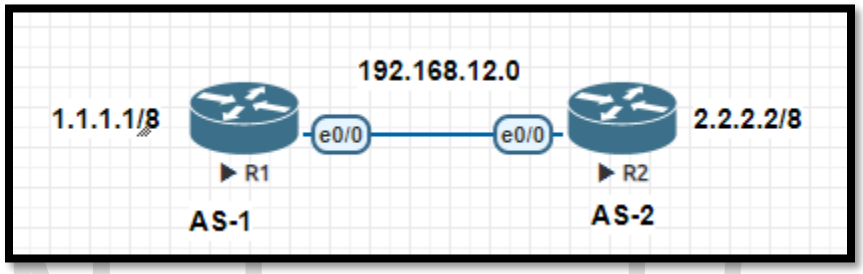

R1 Configuration	R2 Configuration
en config t hostname R1 int e0/0 ip add 192.168.12.1 255.255.255.0 no sh int lo 0 ip add 1.1.1.1 255.0.0.0 router bgp 1 neighbor 192.168.12.2 remote-as 2 network 1.0.0.0 mask 255.0.0.0	en config t hostname R2 int e0/0 ip add 192.168.12.2 255.255.255.0 no sh int lo 0 ip add 2.2.2.2 255.0.0.0 router bgp 2 neighbor 192.168.12.1 remote-as 1 network 2.0.0.0 mask 255.0.0.0

Show Commends
Sh ip bgp Sh ip bgp neighbors 192.168.12.2

```
R2#sh ip bgp
BGP table version is 3, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  1.0.0.0          192.168.12.1             0             0 1 i
 *>  2.0.0.0          0.0.0.0                  0         32768 i
```

Changes in R2
route-map abc permit 10 set weight 500 exit router bgp 2 neighbor 192.168.12.1 route-map abc in

Route-map is applied but nothing has changed yet need to use route refresh method to fix.

```
R2#sh ip bgp
BGP table version is 3, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  1.0.0.0          192.168.12.1             0             0 1 i
 *>  2.0.0.0          0.0.0.0                  0         32768 i
```

Let's apply Border Gateway Protocol (BGP) Hard Reset Option on R2.

Ie. Clear ip bgp *

```
R2#clear ip bgp *
R2#
*Nov  2 16:13:52.783: %BGP-5-ADJCHANGE: neighbor 192.168.12.1 Down User reset
*Nov  2 16:13:52.783: %BGP_SESSION-5-ADJCHANGE: neighbor 192.168.12.1 IPv4 Unicast
*Nov  2 16:13:53.331: %BGP-5-ADJCHANGE: neighbor 192.168.12.1 Up
```

Let's look at the BGP table again this time the Weight is changed.

```
R2#sh ip bgp
BGP table version is 3, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  1.0.0.0          192.168.12.1             0             500 1 i
 *>  2.0.0.0          0.0.0.0                  0         32768 i
```

Changes in R2
route-map abc permit 10 set weight 700 exit router bgp 2 neighbor 192.168.12.1 route-map abc in neighbor 192.168.12.1 soft-reconfiguration inbound

When enable Soft Reconfiguration it stores everything in separate table which require memory.

```
R2#sh ip bgp neighbors 192.168.12.1 received-routes
BGP table version is 4, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *   1.0.0.0          192.168.12.1             0             0 1 i

Total number of prefixes 1
```

Let's apply BGP Soft Reconfiguration Option on R2.

```
R2#clear ip bgp 192.168.12.1 soft
```

Let's look at the BGP table again this time the Weight is changed

```
R2#sh ip bgp
BGP table version is 4, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  1.0.0.0          192.168.12.1             0             700 1 i
 *>  2.0.0.0          0.0.0.0                  0         32768 i
R2#
```

Changes in R2
route-map abc permit 10 set weight 800 exit router bgp 2 neighbor 192.168.12.1 route-map abc in no neighbor 192.168.12.1 soft-reconfiguration inbound
Show commends: sh ip bgp
Clear ip bgp 192.168.12.1 soft

After removing Soft Reconfiguration command there is no extra BGP table.

```
R2#sh ip bgp neighbors 192.168.12.1 received-routes
% Inbound soft reconfiguration not enabled on 192.168.12.1
R2#
```

Clear ip bgp 192.168.12.1 soft

```
R2#sh ip bgp
BGP table version is 5, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  1.0.0.0          192.168.12.1             0             800 1 i
 *>  2.0.0.0          0.0.0.0                  0         32768 i
R2#
```

BGP Address Families:

- Normal version of BGP only **supported IPv4 unicast prefixes**.
- BGP was intended for only routing IPv4 prefixes but added Multi-Protocol BGP capability.
- Modern BGP is very robust and is capable of so much more than **IPv4 and IPV6 routes.**
- **Added Multiprotocol BGP adding an extension called the address family identifier (AFI).**
- Nowadays we use **MP-BGP (Multiprotocol BGP) which supports different addresses etc.**
- Multiprotocol BGP support IPv4 unicast, IPv4 multicast, IPv6 unicast and IPv6 multicast.
- To allow these new addresses, MBGP has new features that the old BGP doesn't have.
- Address family identifier is shared among the BGP peers during the BGP OPEN message.
- In case of multiple address families there will be more than one optional parameter fields.
- This BGP address family is the most common use in the MPLS Virtual Private Networks.

- Every address family maintains a separate database and configuration for each protocol.
- IPv4 address family is used to identify routing sessions for protocols BGP that use IPV4.
- Unicast or multicast address prefixes can be specified within the IPv4 address family.
- The VRF instances can also be associated with IPv4 AFI configuration mode commands.
- IPv6 address family is used to identify routing sessions for protocols BGP that use IPv6.
- Unicast or multicast address prefixes can be specified within the IPv6 address family.

L3-ADVENTERPRISEK9-M-15.4-2T.bin

Lab time:

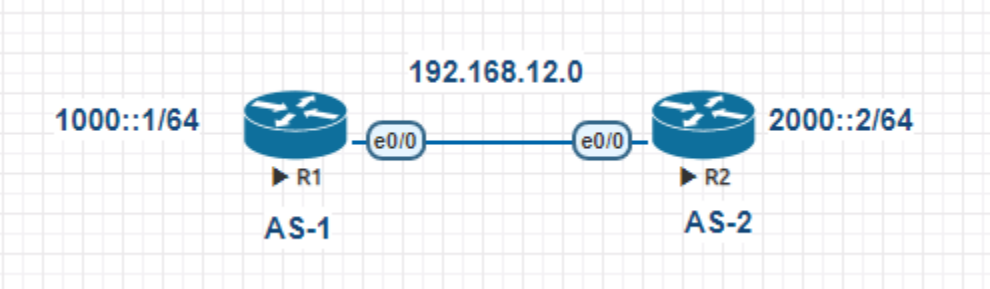

R1 Configuration	R2 Configuration
en config t hostname R1 interface Ethernet0/0 ip address 192.168.12.1 255.255.255.0	en config t hostname R2 interface Ethernet0/0 ip address 192.168.12.2 255.255.255.0

no shutdown exit interface loopback 0 ipv6 address 1000::1/64 no shutdown ipv6 unicast-routing router bgp 1 neighbor 192.168.12.2 remote-as 2 address-family ipv6 network 1000::0/64 neighbor 192.168.12.2 activate	no shutdown exit interface loopback 0 ipv6 address 2000::2/64 no shutdown ipv6 unicast-routing router bgp 2 neighbor 192.168.12.1 remote-as 1 address-family ipv6 network 2000::0/64 neighbor 192.168.12.1 activate
Show command:	
R1#show bgp ipv6 unicast R2#show bgp ipv6 unicast R1#show ip bgp summary R2#show ip bgp summary	


```
R1#sh ip bgp ipv6 unicast
BGP table version is 2, local router ID is 192.168.12.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  1000::/64       ::                       0         32768 i
 *   2000::/64       ::FFFF:192.168.12.2
                                              0             0 2 i
```

```
R2#sh ip bgp ipv6 unicast
BGP table version is 2, local router ID is 192.168.12.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *   1000::/64        ::FFFF:192.168.12.1
                                                0             0 1 i
 *>  2000::/64        ::                      0         32768 i
```

Lab 2:

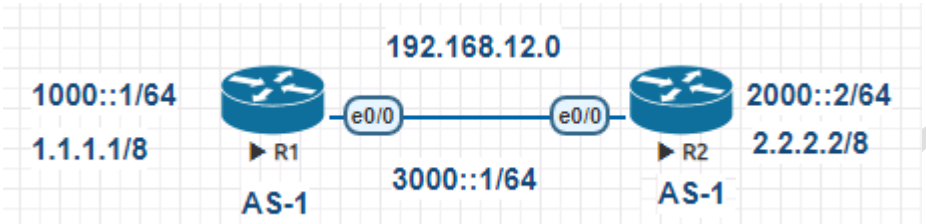

R1 Configuration:	R2 Configuration:
en	en
config t	config t
hostname R1	hostname R2
interface Ethernet0/0	interface Ethernet0/0
ip address 192.168.12.1 255.255.255.0	ip address 192.168.12.2 255.255.255.0
ipv6 address 3000::1/64	ipv6 address 3000::2/64
ipv6 enable	ipv6 enable
no shutdown	no shutdown
exit	exit
interface loopback 1	interface loopback 2
ip address 1.1.1.1 255.0.0.0	ip address 2.2.2.2 255.0.0.0
exit	exit
interface loopback 11	interface loopback 22

ipv6 address 1000::1/64 ipv6 enable ipv6 unicast-routing router bgp 1 neighbor 192.168.12.2 remote-as 1 neighbor 3000::2 remote-as 1 address-family ipv4 unicast neighbor 192.168.12.2 activate network 192.168.12.0 mask 255.255.255.0 network 1.0.0.0 mask 255.0.0.0 exit-address-family address-family ipv6 unicast neighbor 3000::2 activate network 3000::0/64 network 1000::0/64	ipv6 address 2000::2/64 ipv6 enable ipv6 unicast-routing router bgp 1 neighbor 192.168.12.1 remote-as 1 neighbor 3000::1 remote-as 1 address-family ipv4 unicast neighbor 192.168.12.1 activate network 192.168.12.0 mask 255.255.255.0 network 2.0.0.0 mask 255.0.0.0 exit-address-family address-family ipv6 unicast neighbor 3000::1 activate network 3000::0/64 network 2000::0/64


```
R1#show bgp ipv4 unicast
BGP table version is 4, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  1.0.0.0          0.0.0.0                  0         32768 i
 *>i 2.0.0.0          192.168.12.2             0    100      0 i
 * i 192.168.12.0     192.168.12.2             0    100      0 i
 *>                   0.0.0.0                  0         32768 i
```

```
R2#show bgp ipv4 unicast
BGP table version is 4, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>i 1.0.0.0          192.168.12.1             0    100      0 i
 *>  2.0.0.0          0.0.0.0                  0         32768 i
 * i 192.168.12.0     192.168.12.1             0    100      0 i
 *>                   0.0.0.0                  0         32768 i
```

```
R1#sh ip bgp summary
BGP router identifier 1.1.1.1, local AS number 1
BGP table version is 4, main routing table version 4
3 network entries using 420 bytes of memory
4 path entries using 320 bytes of memory
2/2 BGP path/bestpath attribute entries using 288 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 1028 total bytes of memory
BGP activity 6/0 prefixes, 8/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.12.2    4            1       8       8        4    0    0 00:04:00        2
```

```
R2#sh ip bgp summary
BGP router identifier 2.2.2.2, local AS number 1
BGP table version is 4, main routing table version 4
3 network entries using 420 bytes of memory
4 path entries using 320 bytes of memory
2/2 BGP path/bestpath attribute entries using 288 bytes of memory
0 BGP route-map cache entries using 0 bytes of memory
0 BGP filter-list cache entries using 0 bytes of memory
BGP using 1028 total bytes of memory
BGP activity 6/0 prefixes, 8/0 paths, scan interval 60 secs

Neighbor        V           AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
192.168.12.1    4            1       8       8        4    0    0 00:04:24        2
```

```
R1#sh bgp ipv6 unicast
BGP table version is 4, local router ID is 1.1.1.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  1000::/64       ::                       0         32768 i
 *>i 2000::/64       3000::2                  0    100      0 i
 * i 3000::/64       3000::2                  0    100      0 i
 *>                  ::                       0         32768 i
```

```
R2#sh bgp ipv6 unicast
BGP table version is 4, local router ID is 2.2.2.2
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>i 1000::/64       3000::1                  0    100      0 i
 *>  2000::/64       ::                       0         32768 i
 * i 3000::/64       3000::1                  0    100      0 i
 *>                  ::                       0         32768 i
```

Filtering:

- Route filters can be used when exchanging routing updates to with accept or deny update.
- In Dynamic, routing protocols to filter routes that are sent out from one router to another.
- It used to **manipulate traffic flows, reduce memory utilization, or to improve the security.**
- Filter routes that are received into router possibilities can be configured to filter routes.
- **These are Access Control Lists (ACLs), Distribute Lists, Prefix Lists and the Route Maps.**

ACLs:

- For route, filtering Access Control Lists (ACLs) are used in a different way not like normal.
- Access Control Lists are used to specify which routes to allow and which routes to filter.
- The Access Control Lists (ACLs) are used differently with the different types of filtering.
- In Filtering the ACLs are not denying or permitting traffic to/from a different network.
- Instead, the ACLs are either allowing or denying a route from being advertised or learned.

Distribute Lists:

- Distribute-List is used to control routing updates either coming to or leaving from router.
- It mainly used for route filtering means to permit or block routes from being advertised.
- **The Distribute-List itself does not block or permit any route but it will call an Access-List.**
- **The Distribute-List itself does not block or permit any route but it will call a Prefix-List.**
- Action will be taken based on whatever specified in that Access-List or in the Prefix-List.
- So, Distribute-List is mainly used to control the incoming or the outgoing traffic to router.
- The Distribute List can be configured in Cisco Router to control in two directions, in or out.
- The Distribute List command to configure will be run in Cisco Router configuration mode.

- With the BGP, access-lists, prefix-lists, and route-maps are all options of a distribute list.

1. First of all, in Cisco Router, define what routes you want to filter for in or out direction.

2. Create an Access Control Lists (ACL) in configuration mode to filter out that traffic.

3. Create a Distribute-List that references the Access Control Lists & defines the direction.

4. The last thing needs to verify that the route which want to filter has been removed.

Prefix List:

- A prefix list in similar to an Access List in that it will be used to match a route prefix.
- However, processing is much faster than an ACL and provides flexibility for filtering.
- Just like Access Control Lists, there is a default “deny any” at the end of a prefix-list.
- Prefix-list is mainly used to filter the Routes, it used in routing protocols only to filter.
- Main difference in Access-List and Prefix-List is that Access-List only matches the bits.
- Which specified by wildcard mask, but Prefix-List also match the subnet mask as well.
- Can specify range of Subnetmask, which need to be matched to be permitted or denied.

Route-Maps:

- The Route-Maps can be used in Distribute Lists for the purpose of route filtering in router.
- For the purposes of filtering routes in BGP, use a route map that references Access List.
- For the purposes of filtering routes in BGP, use a route map that references or prefix list.
- When looking for matches on routes to filter, Route-Maps use the concept of sequence No.
- In Route-Maps specify sequence number at the end of the route-map command in router.
- By default, the Sequence Number will be Ten 10, If do not specify the sequence number.
- Route-Maps can be used for many purposes on routers such as in PBR, Redistribution route.
- In Cisco router the Route-Map command can either have a permit clause or deny clause.

BGP Filtering Lab:

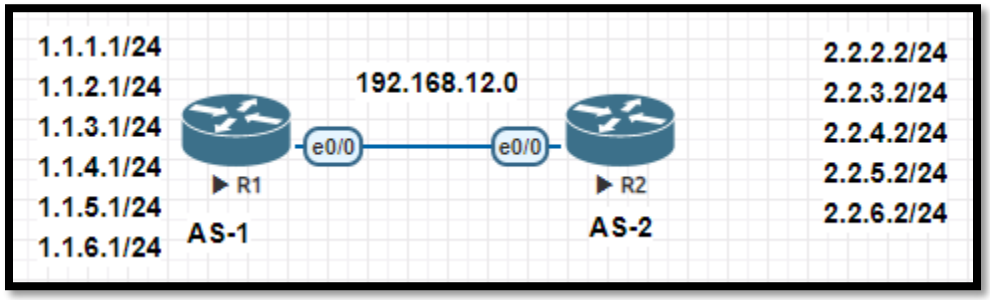

R1 Configuration:	R2 Configuration:
en config t	en config t

hostname R1 interface Ethernet0/0 ip address 192.168.12.1 255.255.255.0 no shutdown interface loopback 1 ip address 1.1.1.1 255.255.255.0 interface loopback 2 ip address 1.1.2.1 255.255.255.0 interface loopback 3 ip address 1.1.3.1 255.255.255.0 interface loopback 4 ip address 1.1.4.1 255.255.255.0 interface loopback 5 ip address 1.1.5.1 255.255.255.0 interface loopback 6 ip address 1.1.6.1 255.255.255.0 router bgp 1 neighbor 192.168.12.2 remote-as 2 network 1.1.1.0 mask 255.255.255.0 network 1.1.2.0 mask 255.255.255.0 network 1.1.3.0 mask 255.255.255.0 network 1.1.4.0 mask 255.255.255.0 network 1.1.5.0 mask 255.255.255.0 network 1.1.6.0 mask 255.255.255.0	hostname R2 interface Ethernet0/0 ip address 192.168.12.2 255.255.255.0 no shutdown interface loopback 2 ip address 2.2.2.2 255.255.255.0 interface loopback 3 ip address 2.2.3.2 255.255.255.0 interface loopback 4 ip address 2.2.4.2 255.255.255.0 interface loopback 5 ip address 2.2.5.2 255.255.255.0 interface loopback 6 ip address 2.2.6.2 255.255.255.0 router bgp 2 neighbor 192.168.12.1 remote-as 1 network 2.2.2.0 mask 255.255.255.0 network 2.2.3.0 mask 255.255.255.0 network 2.2.4.0 mask 255.255.255.0 network 2.2.5.0 mask 255.255.255.0 network 2.2.6.0 mask 255.255.255.0

```
R1#sh ip bgp
BGP table version is 12, local router ID is 1.1.6.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  1.1.1.0/24      0.0.0.0                  0         32768 i
 *>  1.1.2.0/24      0.0.0.0                  0         32768 i
 *>  1.1.3.0/24      0.0.0.0                  0         32768 i
 *>  1.1.4.0/24      0.0.0.0                  0         32768 i
 *>  1.1.5.0/24      0.0.0.0                  0         32768 i
 *>  1.1.6.0/24      0.0.0.0                  0         32768 i
 *>  2.2.2.0/24      192.168.12.2             0             0 2 i
 *>  2.2.3.0/24      192.168.12.2             0             0 2 i
 *>  2.2.4.0/24      192.168.12.2             0             0 2 i
 *>  2.2.5.0/24      192.168.12.2             0             0 2 i
 *>  2.2.6.0/24      192.168.12.2             0             0 2 i
```

Lab time using Access List:

Create ACL to Deny 2.2.2.0 Network
access-list 1 deny 2.2.2.0 0.0.0.255 access-list 1 permit any
Attached ACL to BGP
router bgp 1 distribute-list 1 in
We can check debug:
debug ip bgp updates clear ip bgp *

Debug message clear show that 2.2.2.0/24 network has been denied.

***Nov 3 19:07:32.996: BGP(0): 192.168.12.2 rcvd 2.2.2.0/24 -- DENIED due to: distribute/prefix-list;**

Now let us check BGP routing table there is no more 2.2.2.0/24 network.

```
R1#sh ip bgp
BGP table version is 11, local router ID is 1.1.6.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  1.1.1.0/24      0.0.0.0                  0         32768 i
 *>  1.1.2.0/24      0.0.0.0                  0         32768 i
 *>  1.1.3.0/24      0.0.0.0                  0         32768 i
 *>  1.1.4.0/24      0.0.0.0                  0         32768 i
 *>  1.1.5.0/24      0.0.0.0                  0         32768 i
 *>  1.1.6.0/24      0.0.0.0                  0         32768 i
 *>  2.2.3.0/24      192.168.12.2             0             0 2 i
 *>  2.2.4.0/24      192.168.12.2             0             0 2 i
 *>  2.2.5.0/24      192.168.12.2             0             0 2 i
 *>  2.2.6.0/24      192.168.12.2             0             0 2 i
R1#
```

Lab time using Prefix list:

```
R1#sh ip bgp
BGP table version is 12, local router ID is 1.1.6.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  1.1.1.0/24      0.0.0.0                  0         32768 i
 *>  1.1.2.0/24      0.0.0.0                  0         32768 i
 *>  1.1.3.0/24      0.0.0.0                  0         32768 i
 *>  1.1.4.0/24      0.0.0.0                  0         32768 i
 *>  1.1.5.0/24      0.0.0.0                  0         32768 i
 *>  1.1.6.0/24      0.0.0.0                  0         32768 i
 *>  2.2.2.0/24      192.168.12.2             0             0 2 i
 *>  2.2.3.0/24      192.168.12.2             0             0 2 i
 *>  2.2.4.0/24      192.168.12.2             0             0 2 i
 *>  2.2.5.0/24      192.168.12.2             0             0 2 i
 *>  2.2.6.0/24      192.168.12.2             0             0 2 i
R1#
```

Create Prefix List to deny 2.2.2.0/24 Network:
ip prefix-list abc deny 2.2.2.0/24 ip prefix-list abc permit 0.0.0.0/0 le 32

Attached ACL to BGP
router bgp 1 distribute-list prefix abc in
We can check debug:
debug ip bgp updates clear ip bgp *

Debug message clear show that 2.2.2.0/24 network has been denied.

***Nov 3 19:17:34.940: BGP(0): 192.168.12.2 rcvd 2.2.2.0/24 -- DENIED due to: distribute/prefix-list;**

Now let us check BGP routing table there is no more 2.2.2.0/24 network.

```
R1#sh ip bgp
BGP table version is 11, local router ID is 1.1.6.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  1.1.1.0/24      0.0.0.0                  0         32768 i
 *>  1.1.2.0/24      0.0.0.0                  0         32768 i
 *>  1.1.3.0/24      0.0.0.0                  0         32768 i
 *>  1.1.4.0/24      0.0.0.0                  0         32768 i
 *>  1.1.5.0/24      0.0.0.0                  0         32768 i
 *>  1.1.6.0/24      0.0.0.0                  0         32768 i
 *>  2.2.3.0/24      192.168.12.2             0             0 2 i
 *>  2.2.4.0/24      192.168.12.2             0             0 2 i
 *>  2.2.5.0/24      192.168.12.2             0             0 2 i
 *>  2.2.6.0/24      192.168.12.2             0             0 2 i
R1#
```

Lab time using Route-map:

```
R1#sh ip bgp
BGP table version is 12, local router ID is 1.1.6.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  1.1.1.0/24      0.0.0.0                  0         32768 i
 *>  1.1.2.0/24      0.0.0.0                  0         32768 i
 *>  1.1.3.0/24      0.0.0.0                  0         32768 i
 *>  1.1.4.0/24      0.0.0.0                  0         32768 i
 *>  1.1.5.0/24      0.0.0.0                  0         32768 i
 *>  1.1.6.0/24      0.0.0.0                  0         32768 i
 *>  2.2.2.0/24      192.168.12.2             0             0 2 i
 *>  2.2.3.0/24      192.168.12.2             0             0 2 i
 *>  2.2.4.0/24      192.168.12.2             0             0 2 i
 *>  2.2.5.0/24      192.168.12.2             0             0 2 i
 *>  2.2.6.0/24      192.168.12.2             0             0 2 i
R1#
```

Create ACL to Deny 2.2.2.0 Network:
access-list 1 permit 2.2.2.0 0.0.0.255
Attached ACL to Route Map:
route-map abc deny 10 match ip address 1 exit route-map abc permit 20 exit
Attached ACL to BGP:
router bgp 1 neighbor 192.168.12.2 route-map abc in
We can check debug:
debug ip bgp updates clear ip bgp *

Debug message clearly show that 2.2.2.0/24 network has been denied by route-map.

***Nov 3 19:31:15.329: BGP(0): 192.168.12.2 rcvd 2.2.2.0/24 -- DENIED due to: route-map;**

```
R1#sh ip bgp
BGP table version is 11, local router ID is 1.1.6.1
Status codes: s suppressed, d damped, h history, * valid, > best, i - internal,
              r RIB-failure, S Stale, m multipath, b backup-path, f RT-Filter,
              x best-external, a additional-path, c RIB-compressed,
Origin codes: i - IGP, e - EGP, ? - incomplete
RPKI validation codes: V valid, I invalid, N Not found

     Network          Next Hop            Metric LocPrf Weight Path
 *>  1.1.1.0/24      0.0.0.0                  0         32768 i
 *>  1.1.2.0/24      0.0.0.0                  0         32768 i
 *>  1.1.3.0/24      0.0.0.0                  0         32768 i
 *>  1.1.4.0/24      0.0.0.0                  0         32768 i
 *>  1.1.5.0/24      0.0.0.0                  0         32768 i
 *>  1.1.6.0/24      0.0.0.0                  0         32768 i
 *>  2.2.3.0/24      192.168.12.2             0             0 2 i
 *>  2.2.4.0/24      192.168.12.2             0             0 2 i
 *>  2.2.5.0/24      192.168.12.2             0             0 2 i
 *>  2.2.6.0/24      192.168.12.2             0             0 2 i
```