Building Scalable Cisco Internetworks (BSCI)

Border Gateway Protocol (BGP)

http://www.INE.com

What Is BGP?

- Border Gateway Protocol Version 4
- Standards based
  - RFC 4271 “A Border Gateway Protocol 4 (BGP-4)”
- Exterior Gateway Protocol (EGP)
  - Used for inter-domain routing between Autonomous Systems
- Path vector routing
  - Uses multiple “attributes” for routing decision
- Classless
  - Supports VLSM and summarization

Inter-AS Routing and ASNs

- Autonomous System (AS)
  - *"…a set of routers under a single technical administration, using an interior gateway protocol (IGP) and common metrics to determine how to route packets within the AS, and using an inter-AS routing protocol to determine how to route packets to other ASes."* (RFC 4271)
- Like IP address space, Autonomous System Numbers (ASNs) allocated by Internet Assigned Numbers Authority (IANA)
  - http://www.iana.org/numbers/
- BGP ASNs originally 2-byte field
  - Values 0-65535
- RFC 4893 defines 4-byte ASNs
  - 0.0 – 65535.65535 notation
  - 0.[0-65535] denote original 2-byte ASNs

Why Use BGP?

- Scalability
  - IGPs can scale to thousands of routes
  - BGP can scale to hundreds of thousands of routes
  - Current Global (Internet) BGP table ~ 300,000 routes
- Stability
  - Internet routing table *never* converges
  - BGP stable enough to handle routing and decision making at the same time
- Enforce routing policy
  - IGP uses link cost for routing decision
    - Effective traffic engineering nearly impossible with IGP
  - BGP uses attributes of the route itself
    - Traffic engineering feasible and simple to implement

Who Needs BGP?

- Transit networks
  - SPs that sell access or transit bandwidth to customers
  - Need full routing table to make accurate decisions
  - Should not use default routing
- Multihomed networks
  - Enterprise networks with two or more connections to ISPs
  - Allows control of inbound and outbound routing policy

Example Transit Network

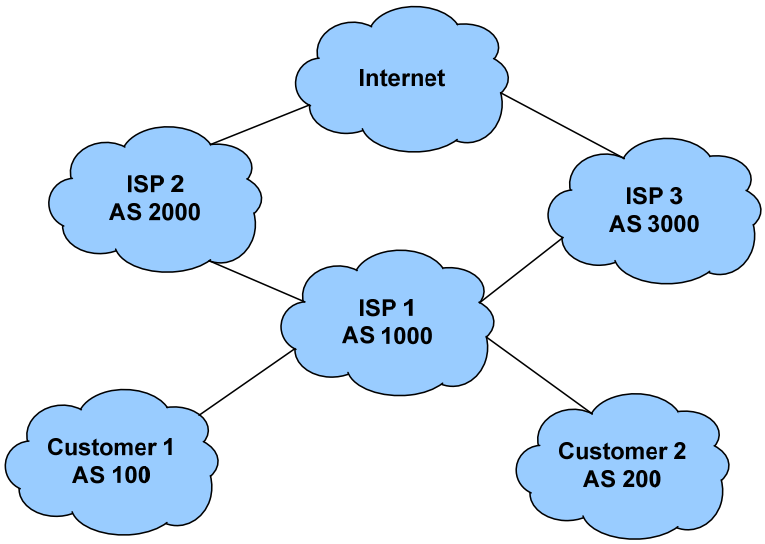

Example Multihomed Network

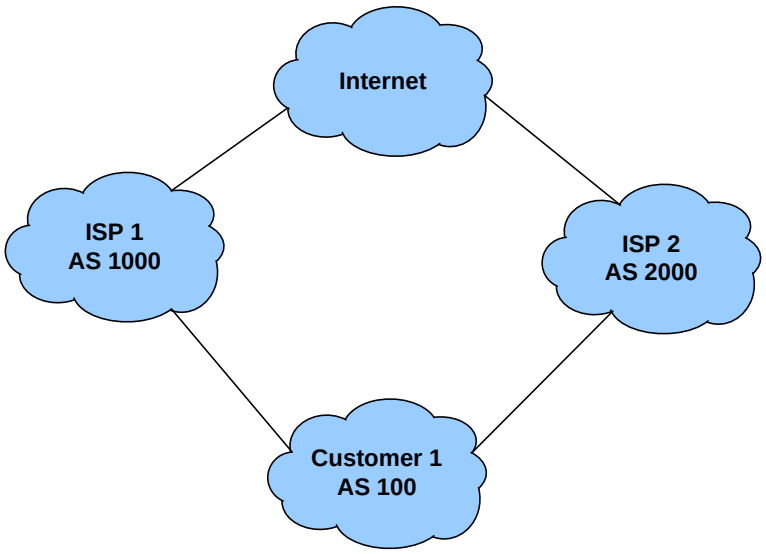

When *not* To Use BGP

- Single ISP connectivity
  - Default routing sufficient
- Limited memory and/or CPU resources
  - Global table needs ~ 1GB of memory just for storage
- Don't "own" your IPv4 addresses
  - ISP advertises "their" address space on your behalf
  - Red tape involved with getting PI address space and BGP ASN

BGP Data Structure

- Like EIGRP/OSPF/IS-IS, BGP uses a three table data structure
- Neighbor table
  - List of active adjacencies called “peerings”
- BGP table
  - All prefixes learned from all peers
- IP Routing table
  - The “best” routes from the BGP table actually used for routing

How BGP Works

- Establish BGP peerings to build neighbor table
- Exchange updates to build BGP table
- Choose BGP bestpaths to build routing table

Example Global BGP Neighbor Table

```
route-views.oregon-ix.net>show ip bgp summary
BGP router identifier 128.223.51.103, local AS number 6447
BGP table version is 14808442, main routing table version 14808442
311034 network entries using 41056488 bytes of memory
9577818 path entries using 498046536 bytes of memory
1570690/56881 BGP path/bestpath attribute entries using 232462120 bytes of memory
1359127 BGP AS-PATH entries using 36934358 bytes of memory
20032 BGP community entries using 1333024 bytes of memory
29 BGP extended community entries using 1406 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 809833932 total bytes of memory
Dampening enabled. 4650 history paths, 13012 dampened paths
BGP activity 533254/216360 prefixes, 53516863/43920540 paths, scan interval 60 secs

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
4.69.184.193    4  3356 2713078   73873 14808442    0    0 4w4d       287263
12.0.1.63       4  7018 7838671   44558 14808442    0    0 4w4d       288299
64.71.255.61    4   812 4340316   65169 14808442    0    0 2w2d       288960
64.125.0.137    4  6461       0       0        0    0    0 never    Active
65.106.7.139    4  2828 3331717   73873 14808442    0    0 1w3d       288575
66.59.190.221   4  6539 2199436   73908 14808442    0    0 7w4d       289238
66.185.128.48   4  1668 2439262   73872 14808442    0    0 3w4d       288084
89.149.178.10   4  3257 3433546     456 14808442    0    0 3w0d       288743
114.31.199.1    4  4826 3203817   73009 14808442    0    0 3w4d       290129
128.223.253.8   4  3582 3102763  145590 14808442    0    0 4d23h      289837
129.250.0.11    4  2914 6481745  145499 14808442    0    0 4w2d       289026
129.250.0.171   4  2914 6505329  145526 14808442    0    0 14:59:04   289026
134.222.87.1    4   286 4507669     452 14808442    0    0 5w3d       289366
144.228.241.81  4  1239       0       0        0    0    0 never    Active
154.11.11.113   4   852 3183907   73875 14808442    0    0 3w4d       272660
154.11.98.225   4   852 3483572   73875 14808442    0    0 3w4d       272659
157.130.10.233  4   701 3958967  145494 14808442    0    0 2w1d       288026
164.128.32.11   4  3303 1396623   44575 14808442    0    0 7w4d       113860
192.203.116.253 4 22388  306757   44577 14808445    0    0 2w1d        12109
193.0.0.56      4  3333 9431113  145515 14808445    0    0 2w0d       292853
193.251.245.6   4  5511       0       0        0    0    0 never    Active
194.85.4.55     4  3277 5820626   73573 14808445    0    0 13:53:44   292248
194.85.40.15    4  3267 4453220   73832 14808445    0    0 07:02:15   292385
195.66.232.239  4  5459 2386544   44575 14808445    0    0 7w4d       200134
195.219.96.239  4  6453 4064203   44474 14808445    0    0 3w3d       288282
<output omitted>
```


Example Global BGP Table

```
route-views.oregon-ix.net>show ip bgp
BGP table version is 14808445, local router ID is 128.223.51.103
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
*> 1.0.0.0          207.46.32.34                           0 8075 8069 4538 i
*  3.0.0.0          194.85.40.15                           0 3267 6453 9304 80 i
*                   194.85.4.55                            0 3277 3267 6453 9304 80 i
*                   129.250.0.171          368             0 2914 9304 80 i
*                   128.223.253.8                          0 3582 4600 11537 15412 9304 80 i
*                   164.128.32.11                          0 3303 2914 9304 80 i
*                   207.172.6.1              0             0 6079 2914 9304 80 i
*                   217.75.96.60             0             0 16150 15412 9304 80 i
*                   65.106.7.139             3             0 2828 2914 9304 80 i
*                   196.7.106.245            0             0 2905 701 2914 9304 80 i
*                   193.0.0.56                             0 3333 3356 15412 9304 80 i
*                   208.51.134.254        2500             0 3549 2914 9304 80 i
*                   157.130.10.233                         0 701 2914 9304 80 i
*                   64.71.255.61                           0 812 6453 9304 80 i
*                   203.181.248.168                        0 7660 4635 9304 80 i
*                   89.149.178.10           10             0 3257 6453 9304 80 i
*                   207.172.6.20             0             0 6079 2914 9304 80 i
*                   207.45.223.244                         0 6453 9304 80 i
*                   195.219.96.239                         0 6453 9304 80 i
*                   154.11.11.113            0             0 852 15412 9304 80 i
*                   66.185.128.48          511             0 1668 6453 9304 80 i
*                   114.31.199.1                           0 4826 3356 15412 9304 80 i
*                   154.11.98.225            0             0 852 15412 9304 80 i
*                   207.46.32.34                           0 8075 15412 9304 80 i
*>                  129.250.0.11           293             0 2914 9304 80 i
*                   203.62.252.186           0             0 1221 4637 9304 80 i
*                   4.69.184.193             0             0 3356 15412 9304 80 i
*                   12.0.1.63                              0 7018 2914 9304 80 i
*                   206.24.210.102                         0 3561 2914 9304 80 i
*                   202.232.0.2                            0 2497 2914 9304 80 i
*                   134.222.87.1                           0 286 15412 9304 80 i
*                   202.249.2.86                           0 7500 2497 2914 9304 80 i
*                   66.59.190.221                          0 6539 15412 9304 80 i
*                   216.218.252.164                        0 6939 15412 9304 80 i
*                   195.66.232.239                         0 5459 15412 9304 80 i
<output omitted>
```


Example BGP Routing Table

```
route-views.oregon-ix.net>show ip route
Codes: 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

Gateway of last resort is 128.223.51.1 to network 0.0.0.0

B    216.221.5.0/24 [20/0] via 206.24.210.102, 1d12h
B    216.187.99.0/24 [20/0] via 4.69.184.193, 4w4d
B    210.51.225.0/24 [20/0] via 4.69.184.193, 1w4d
B    210.34.240.0/24 [20/0] via 207.46.32.34, 2w5d
B    209.136.89.0/24 [20/0] via 216.218.252.164, 1d09h
B    209.34.243.0/24 [20/0] via 157.130.10.233, 2w0d
B    205.204.1.0/24 [20/0] via 157.130.10.233, 2w1d
B    204.255.51.0/24 [20/0] via 157.130.10.233, 2d18h
B    204.238.34.0/24 [20/0] via 157.130.10.233, 2w1d
B    204.221.17.0/24 [20/0] via 134.222.87.1, 1d09h
B    204.17.221.0/24 [20/0] via 4.69.184.193, 4w0d
B    203.255.52.0/24 [20/0] via 207.46.32.34, 2w4d
     203.238.37.0/24 is variably subnetted, 3 subnets, 2 masks
B       203.238.37.0/24 [20/0] via 216.218.252.164, 3w4d
B       203.238.37.96/27 [20/0] via 203.62.252.186, 2w3d
B       203.238.37.128/27 [20/0] via 203.62.252.186, 2w3d
B    203.170.97.0/24 [20/0] via 216.218.252.164, 2w5d
B    203.34.233.0/24 [20/0] via 203.62.252.186, 4w4d
B    203.17.218.0/24 [20/0] via 164.128.32.11, 4w4d
B    202.153.83.0/24 [20/0] via 195.66.232.239, 3w1d
B    202.119.189.0/24 [20/0] via 207.46.32.34, 2w5d
B    202.85.159.0/24 [20/0] via 203.181.248.168, 14:17:04
B    198.17.215.0/24 [20/3] via 65.106.7.139, 1w3d
B    194.153.91.0/24 [20/0] via 4.69.184.193, 2d12h
B    194.136.74.0/24 [20/0] via 4.69.184.193, 2w6d
B    194.68.134.0/24 [20/0] via 216.218.252.164, 4w3d
B    194.0.194.0/24 [20/0] via 216.218.252.164, 1d12h
B    193.102.167.0/24 [20/0] via 4.69.184.193, 4w1d
B    193.17.208.0/24 [20/0] via 216.218.252.164, 2w3d
B    192.68.132.0/24 [20/0] via 66.59.190.221, 7w0d
     170.170.0.0/16 is variably subnetted, 3 subnets, 3 masks
<output omitted>
```


Establishing BGP Peerings

- Like IGP, first step in BGP is to find neighbors to exchange information with
- Unlike IGP…
  - BGP does not have its own transport
    - Uses TCP port 179
  - BGP neighbors are not discovered
    - Manually configured via **`neighbor`** statement
  - BGP neighbors do not have to be connected
    - IGP is always on a link-by-link basis
    - BGP is a logical peering over TCP
    - Implies that BGP always needs IGP underneath
  - BGP has different types of neighbors
    - External BGP vs. Internal BGP

BGP Packet Formats

- Peering establishment and maintenance uses four types of packets
  - OPEN
  - KEEPALIVE
  - UPDATE
  - NOTIFICATION

BGP OPEN Message

- Used to negotiate parameters for peering
- Includes…
  - BGP version
    - Should be 4
  - Local ASN
  - Local Router-ID
  - Hold time
    - Negotiated to lowest requested value
  - Options
    - AKA “capabilities”

BGP KEEPALIVE Message

- Used for dead neighbor detection
- If hold time = 0, keepalives disabled

BGP UPDATE Message

- Used to advertise or withdraw a prefix
- Includes
  - Withdrawn routes
    - List of routes that should be discarded
  - NLRI
    - Route being advertised
  - Path vector attributes
    - Attributes of route being advertised
    - Used for bestpath selection

BGP NOTIFICATION Message

- Used to convey error messages
- After notification sent, BGP session closed
- Examples
  - Unsupported Version Number
  - Unsupported Optional Parameter
  - Unacceptable Hold Time
  - Hold Timer Expired

BGP Peering State Machine

- BGP state machine tracks peering establishment
- Idle
  - Waiting to start 3-way handshake
- Connect
  - Waiting to complete 3-way handshake
- Active
  - 3-way handshake failed, try again
- Open sent
  - 3-way handshake complete, OPEN message sent
- Open confirm
  - OPEN message received, parameters agreed upon
- Established
  - Peering complete

BGP Peering State Machine Debug

```
R1#debug ip bgp
BGP debugging is on for address family: IPv4 Unicast
R1#config t
Enter configuration commands, one per line.  End with CNTL/Z.
R1(config)#router bgp 1
R1(config-router)#neighbor 10.1.146.6 remote-as 1
R1(config-router)#end
%SYS-5-CONFIG_I: Configured from console by console
R1#
BGP: 10.1.146.6 went from Idle to Connect
BGP: 10.1.146.6 rcv message type 1, length (excl. header) 26
BGP: 10.1.146.6 rcv OPEN, version 4, holdtime 180 seconds
BGP: 10.1.146.6 went from Connect to OpenSent
BGP: 10.1.146.6 sending OPEN, version 4, my as: 1, holdtime 180 seconds
BGP: 10.1.146.6 rcv OPEN w/ OPTION parameter len: 16
BGP: 10.1.146.6 rcvd OPEN w/ optional parameter type 2 (Capability) len 6
BGP: 10.1.146.6 OPEN has CAPABILITY code: 1, length 4
BGP: 10.1.146.6 OPEN has MP_EXT CAP for afi/safi: 1/1
BGP: 10.1.146.6 rcvd OPEN w/ optional parameter type 2 (Capability) len 2
BGP: 10.1.146.6 OPEN has CAPABILITY code: 128, length 0
BGP: 10.1.146.6 OPEN has ROUTE-REFRESH capability(old) for all address-families
BGP: 10.1.146.6 rcvd OPEN w/ optional parameter type 2 (Capability) len 2
BGP: 10.1.146.6 OPEN has CAPABILITY code: 2, length 0
BGP: 10.1.146.6 OPEN has ROUTE-REFRESH capability(new) for all address-families
BGP: 10.1.146.6 rcvd OPEN w/ remote AS 1
BGP: 10.1.146.6 went from OpenSent to OpenConfirm
BGP: 10.1.146.6 send message type 1, length (incl. header) 45
BGP: 10.1.146.6 went from OpenConfirm to Established
%BGP-5-ADJCHANGE: neighbor 10.1.146.6 Up
```


BGP Peering Types

- External BGP (EBGP) Peers
  - Neighbors outside my Autonomous System
- Internal BGP (iBGP) Peers
  - Neighbors inside my Autonomous System
- Update and path selection rules change depending on what type of peer a route is being sent to/received from

EBGP Peerings

- Peers in different ASes
- Usually directly connected neighbors
  - e.g. DS3 Frame Relay link to ISP
- Can be "multihop", but TTL defaults to 1
  - **`neighbor [address] ebgp-multihop [ttl]`**
- Uses AS-Path attribute for loop prevention
  - If I receive an update from an EBGP peer with my own ASN in the AS-Path, discard it

iBGP Peerings

- Peers in the same AS
- Many times *not* directly connected
  - Implies IGP needed to provide TCP transport
- Loop prevention via route suppression
  - Routes learned from an iBGP peer cannot be advertised on to another iBGP peer
- Implies that all routers running BGP within the AS must peer with each other
  - i.e. "iBGP full mesh"
  - *n**(*n*-1)/2 peerings

iBGP Full Mesh Example

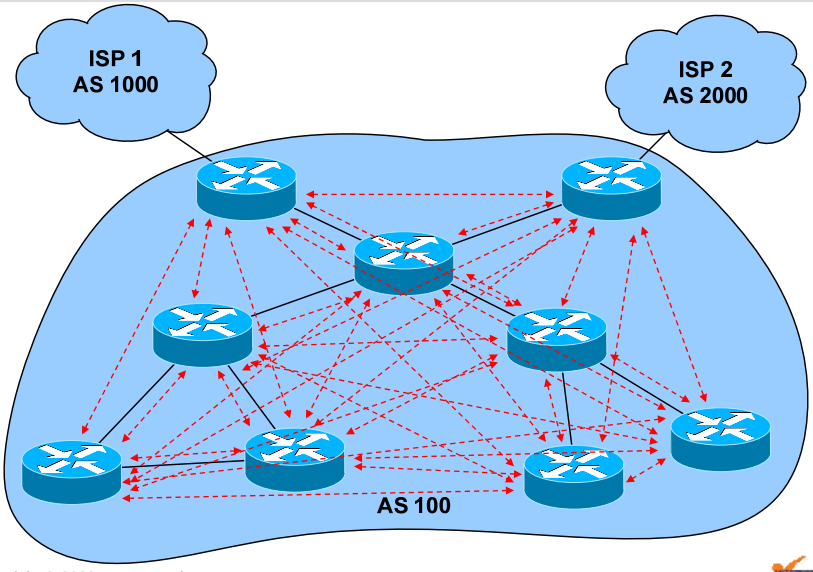

iBGP Full Mesh Scalability

- *n**(*n*-1)/2 doesn't scale
  - 10 routers, 45 peerings
  - 100 routers, 4950 peerings
  - 1000 routers, 499,500 peerings
- Can be fixed with two exceptions
  - Route Reflectors
    - Same logic as DR/DIS
  - Confederation
    - Split the AS into smaller Sub-ASes

BGP Route Reflectors

- Eliminates need for full mesh
  - Only need peering(s) to the RR(s)
- Like OSPF DR & IS-IS DIS, minimizes prefix replication
  - Send one update to the RR
  - RR sends the update to its “clients”
- Loop prevention through Cluster-ID
  - If I am a RR and I receive a route with my own Cluster-ID, discard it

Route Reflector Example

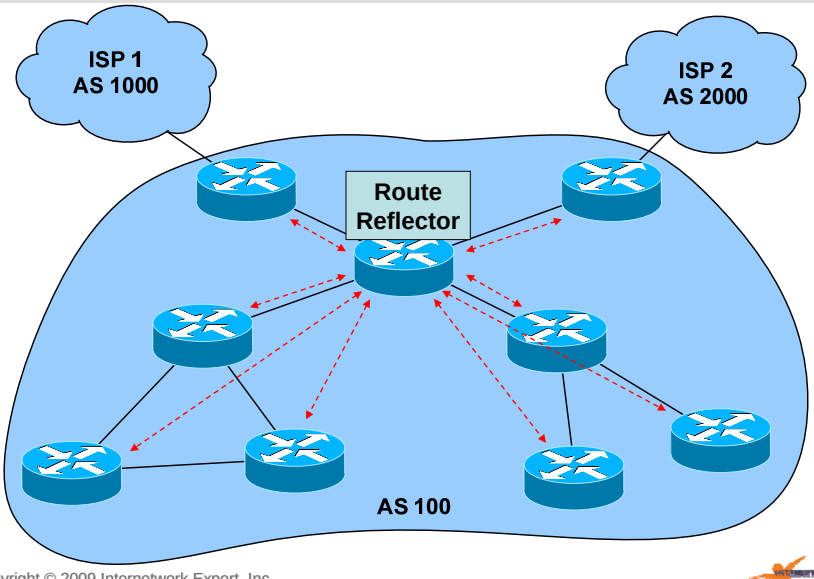

BGP Confederation

- Reduces full mesh iBGP need by splitting AS into smaller Sub-ASes
  - Inside Sub-AS full mesh or RR need remains
  - Between Sub-AS acts like EBGP
- Devices outside the confederation do not know about the internal structure
  - Sub-AS numbers are stripped from advertisements to true EBGP peers
- Typically uses ASNs in private range (64512 – 65535)

BGP Confederation Example

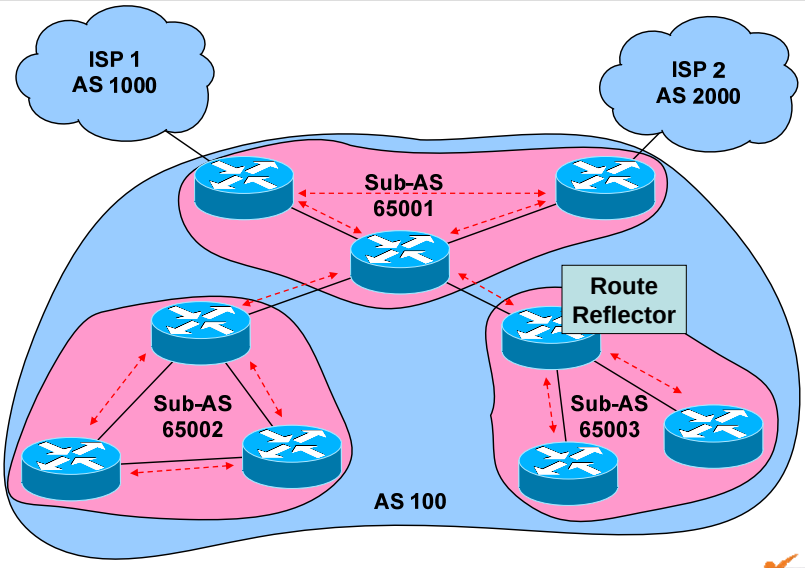

BGP Peering Redundancy

- BGP peering is based on TCP reachability to peer address
- If peer address is unreachable, peering goes down
  - e.g. if IP address of Serial link is used for peering and Serial link is down, peer goes down
- Using Loopback addresses for peerings allows rerouting around link failures and adds redundancy
  - e.g. as long as any link is up, Loopback can be reached
- Defined as **`update-source`** for TCP session

BGP Loopback Redundancy Example

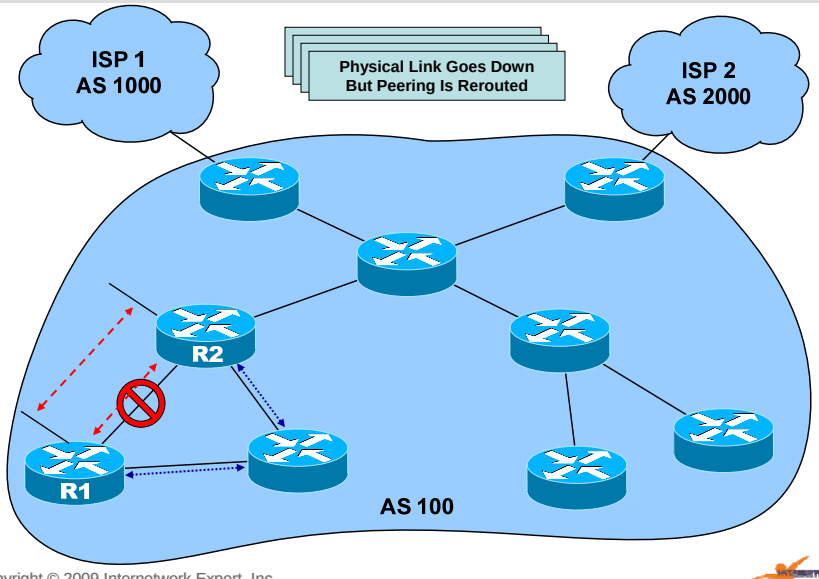

BGP Peer Groups

- Typically many peers share the same update policy
  - e.g. a route reflector's clients
- BGP Peer Groups reduce configuration and processing overhead by applying a template to the peers
- Peer group is assigned parameters such as…
  - **`remote-as`**
  - **`route-reflector-client`**
  - **`route-map`**
- Neighbor is specified as a member of the group
  - Peers in a group must be either all iBGP or all EBGP

BGP Peer Group Example

```
router bgp 1
 neighbor IBGP_PEER_GROUP peer-group
 neighbor IBGP_PEER_GROUP remote-as 1
 neighbor IBGP_PEER_GROUP update-source Loopback0
 neighbor IBGP_PEER_GROUP route-reflector-client
 neighbor IBGP_PEER_GROUP next-hop-self
 neighbor 1.2.3.4 peer-group IBGP_PEER_GROUP
 neighbor 5.6.7.8 peer-group IBGP_PEER_GROUP
 neighbor 9.10.11.12 peer-group IBGP_PEER_GROUP
 neighbor 13.14.15.16 peer-group IBGP_PEER_GROUP
```


BGP Authentication

- Like IGP authentication, BGP peer authentication protects control plane against attacks and misconfigurations
  - Without authentication, BGP susceptible to TCP RST attacks
    - Interesting read: “Slipping in the Window: TCP Reset attacks”
- Uses MD5 as defined in RFC 2385
  - “Protection of BGP Sessions via the TCP MD5 Signature Option”
- Simply configured as **`neighbor [`*`address`*`] password [`*`password`*`]`**

Basic BGP Configuration

- Enable global BGP process
  - **`router bgp [`*`ASN`*`]`**
- Establish BGP peers
  - **`neighbor [`*`address`*`] remote-as [`*`remote ASN`*`]`**

Basic BGP Verification

- Verify BGP peerings
  - **`show ip bgp summary`**
- Verify BGP table
  - **`show ip bgp`**
- Verify BGP table detail
  - **`show ip bgp [network] [mask]`**
- Verify BGP routing table
  - **`show ip route [bgp]`**

Misc. BGP Configuration

- Modify peering source address
  - **`neighbor [address] update-source [interface]`**
- Enabling BGP authentication
  - **`neighbor [address] password [password]`**
- Configuring BGP peer group
  - **`neighbor [Peer-Group-Name] peer-group`**
  - **`neighbor [Peer-Group-Name] [attrbiutes]`**
  - **`neighbor [address] peer-group [Peer-Group-Name]`**

Misc. BGP Configuration (cont.)

- Enabling Route Reflection
  - **`neighbor [address] route-reflector-client`**
- Enabling Confederation
  - Enable global BGP process
    - **`router bgp [Sub-ASN]`**
  - Define global ASN
    - **`bgp confederation-id [ASN]`**
  - Define other Sub-ASes
    - **`bgp confederation peers [`$Sub\text{-}ASN_1$`] [`$Sub\text{-}ASN_2$`] [`$Sub\text{-}ASN_n$`]`**

Building the BGP Table

- Once peerings are established, UPDATE messages are exchanged to advertise NLRI and build the BGP table
- Routes local to the AS can be originated either by process level **`network [network] mask [mask]`** statement or redistribution
- Unlike IGP, networks do not have to be directly connected to be advertised, they only have to be in the routing table
  - e.g. prefixes in local routing table learned via OSPF can be advertised with BGP **`network`** statement

BGP NLRI Aggregation

- BGP aggregation, like IGP summarization, is used to reduce resource requirements needed to process the BGP table
  - Configured as **`aggregate-address [`*`network`*`] [`*`mask`*`] [summary-only|as-set|route-map|...]`**
- Can be applied at any point in the network
  - No hierarchy like OSPF/IS-IS
- Does not automatically stop subnet advertisements
  - **`summary-only`** argument
- Can be used for longest match routing traffic engineering

BGP Path Vector Attributes

- UPDATE includes path vector attributes for a route
  - Next-hop
  - AS-Path
  - Origin
  - Local preference
  - Multi-Exit Discriminator (MED)
  - Atomic aggregate
  - Aggregator

BGP Attribute Types

- Attributes fall into different categories…
  - Well-known vs. optional
    - Well-known must be implemented
    - Optional may or may not be implemented
  - Mandatory vs. discretionary
    - Mandatory must be present in update
    - Discretionary may or may not be present
  - Transitive vs. non-transitive
    - Transitive passes between EBGP and iBGP neighbors
    - Non-transitive passes only between iBGP neighbors
- Valid combinations are…
  - Well-known mandatory
  - Well-known discretionary
  - Optional transitive
  - Optional non-transitive

BGP Next-Hop

- Well-known mandatory attribute
- If UPDATE comes from EBGP peer
  - Next-hop is the IP address they use to peer with you
    - i.e. their **`update-source`**
- If UPDATE comes from iBGP peer
  - Next-hop is the IP address used to peer with the EBGP neighbor they learned it from
    - i.e. the next-hop is unmodified
  - Implies that iBGP neighbors must have an IGP route to the links between EBGP neighbors
  - Behavior can be changed with **`neighbor [address] next-hop-self`**

BGP AS-Path

- Well-known mandatory attribute
- Defines which Autonomous Systems the route has passed through
- When sending an UPDATE to an EBGP neighbor, the local ASN is “prepended” to the route
  - Example path “100 1000 2000”
    - Originated in 2000
    - Passed through 1000
    - Learned from 100
- Shorter AS-Path length is preferred
  - e.g. “100 1000 2000” (3) vs “100 1000 2000 3000” (4)

BGP Origin

- Well-known mandatory attribute
- Possible values…
  - 0 – IGP
  - 1 – EGP
  - 2 – Incomplete
- Defines how prefix was advertised into BGP
  - IGP – interior to the AS
  - EGP – the actual protocol “EGP” (deprecated)
  - Incomplete – some other means
    - e.g. redistribution
- Lower origin code is preferred

BGP Local Preference

- Well-known discretionary attribute
- 4 byte field
  - Value of 0 – 4,294,967,295
- Only exchanged in iBGP updates
- Higher local preference is preferred

BGP Multi-Exit Discriminator

- AKA MED or simply “metric”
- optional non-transitive attribute
- 4 byte field
  - Value of 0 – 4,294,967,295
- Used to choose (discriminate) between multiple exit points out of the AS
  - Many exceptions to MED comparison
  - Rarely used in practice
- Lower MED is preferred

BGP Atomic Aggregate and Aggregator

- Atomic Aggregate
  - Well-known discretionary attribute
- Aggregator
  - Optional transitive attribute
- Both used when BGP prefixes are summarized (aggregated) together
- "Aggregate" prefix has
  - Atomic Aggregate = TRUE
  - Aggregator = BGP Router-ID who performed summarization

BGP Bestpath Selection

- Once updates are exchanged, path selection begins
- Bestpath selection algorithm compares path vector attributes and elects one route as "best" for each prefix
  - Denoted by "**>**" in the **`show ip bgp`** output
- Like RIPv2 & EIGRP, only best route is sent to the routing table and to other peers

Bestpath Selection Example

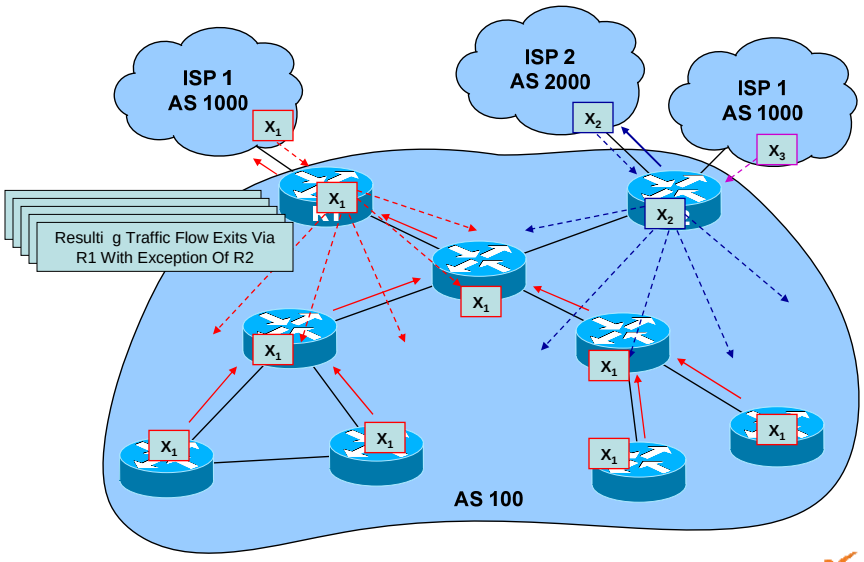

BGP Bestpath Selection Order

- Algorithm not standardized, Cisco IOS selection order is…
  - Weight (highest)
    - Locally significant Cisco proprietary attribute
  - Local Preference (highest)
  - Locally originated routes
  - AS-Path (shortest)
  - Origin (lowest)
  - MED (lowest)
  - EBGP learned routes over iBGP learned routes
  - Smallest IGP metric to next-hop value
- Algorithm runs top down until a deciding match occurs
- Other tie-breaking checks occur if no bestpath
  - Oldest route, lowest Router-ID, lowest interface IP address, etc.
- See BGP Best Path Selection Algorithm on cisco.com for details

Manipulating BGP Bestpath Selection

- Vector attributes can be manually modified to define different routing policy for different routes
  - E.g. control inbound/outbound traffic flow on a per-prefix basis
- Attributes typically modified are…
  - Weight
  - Local-Preference
  - AS-Path
  - MED
- Inbound routing policy affects outbound traffic
  - Change weight or local-pref in to affect traffic out
- Outbound routing policy affects incoming traffic
  - Change AS-Path or MED to affect traffic in

Manipulating Weight Example

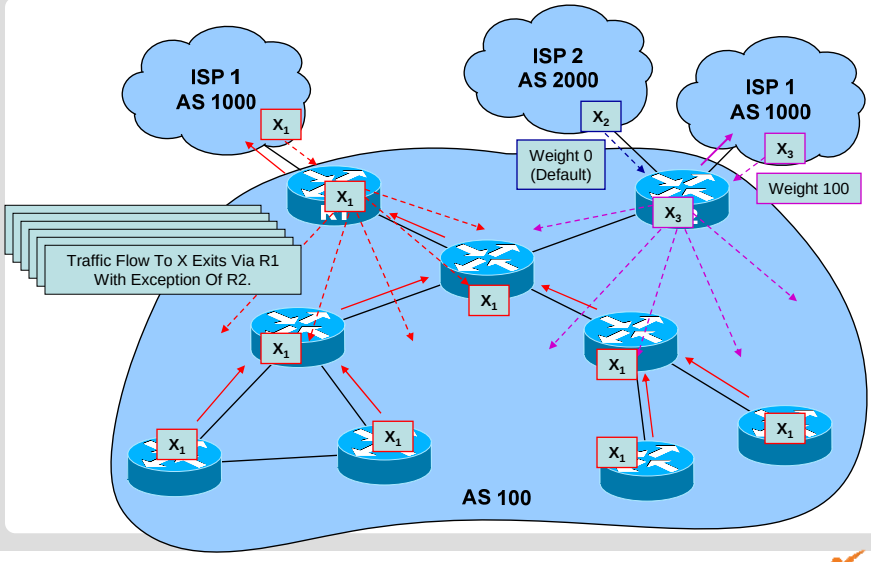

Manipulating Local Preference Example

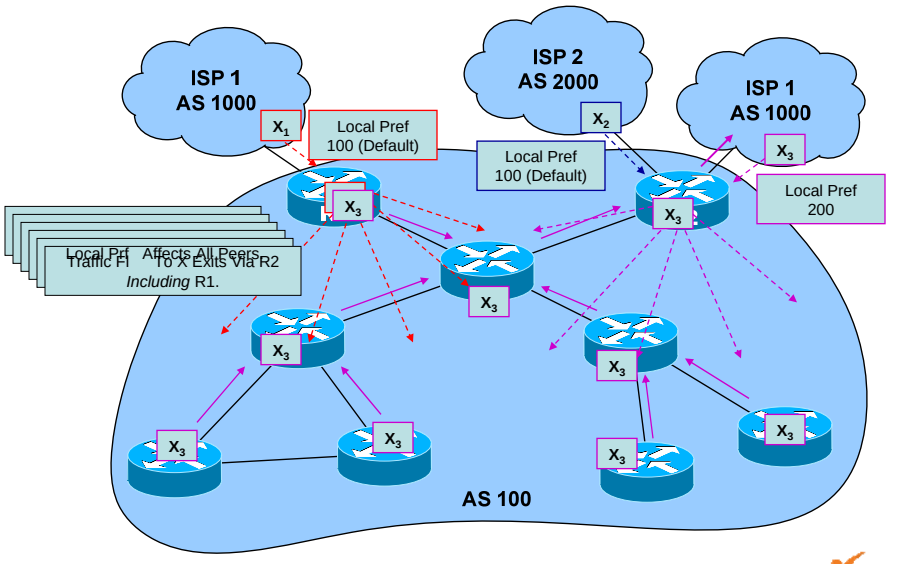

Manipulating AS-Path Example

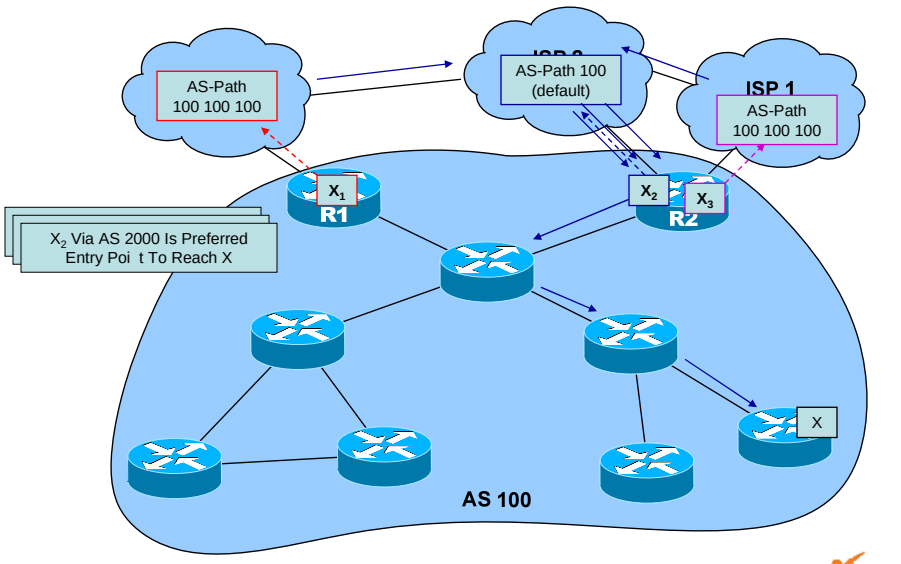

Manipulating MED Example

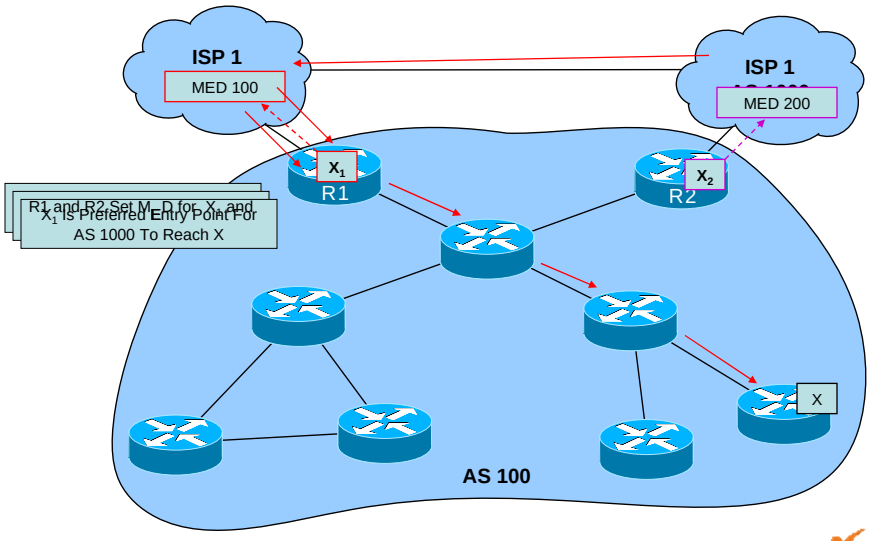

BGP Configuration Topology

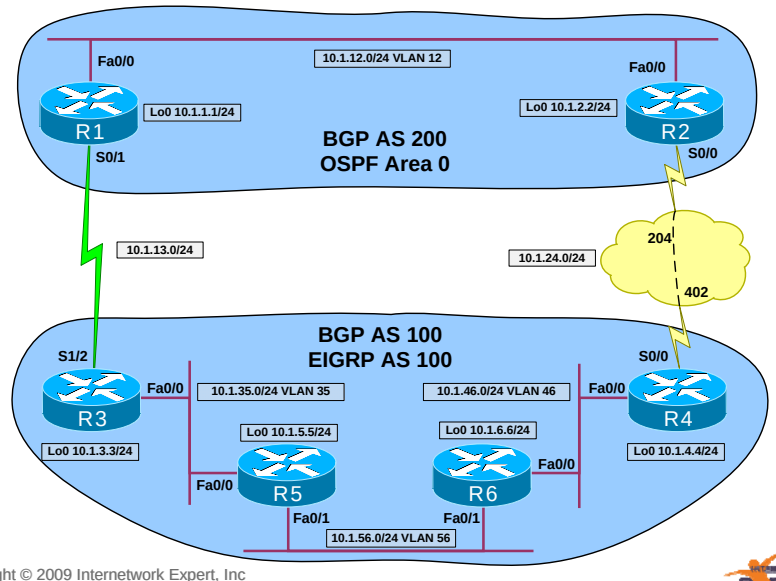

Basic BGP Peering Configuration

```
R1#
router bgp 200
 neighbor 10.1.12.2 remote-as 200
 neighbor 10.1.13.3 remote-as 100

R2#
router bgp 200
 neighbor 10.1.12.1 remote-as 200
 neighbor 10.1.24.4 remote-as 100

R3#
router bgp 100
 neighbor 10.1.4.4 remote-as 100
 neighbor 10.1.5.5 remote-as 100
 neighbor 10.1.6.6 remote-as 100
 neighbor 10.1.4.4 update-source Loopback0
 neighbor 10.1.5.5 update-source Loopback0
 neighbor 10.1.6.6 update-source Loopback0
 neighbor 10.1.13.1 remote-as 200
```

```
R4#
router bgp 100
 neighbor 10.1.3.3 remote-as 100
 neighbor 10.1.5.5 remote-as 100
 neighbor 10.1.6.6 remote-as 100
 neighbor 10.1.3.3 update-source Loopback0
 neighbor 10.1.5.5 update-source Loopback0
 neighbor 10.1.6.6 update-source Loopback0
 neighbor 10.1.24.2 remote-as 200

R5#
router bgp 100
 neighbor 10.1.3.3 remote-as 100
 neighbor 10.1.4.4 remote-as 100
 neighbor 10.1.6.6 remote-as 100
 neighbor 10.1.3.3 update-source Loopback0
 neighbor 10.1.4.4 update-source Loopback0
 neighbor 10.1.6.6 update-source Loopback0

R6#
router bgp 100
 neighbor 10.1.3.3 remote-as 100
 neighbor 10.1.4.4 remote-as 100
 neighbor 10.1.5.5 remote-as 100
 neighbor 10.1.3.3 update-source Loopback0
 neighbor 10.1.4.4 update-source Loopback0
 neighbor 10.1.5.5 update-source Loopback0
```


BGP Peering Verification

```
R1#show ip bgp summary
BGP router identifier 10.1.1.1, local AS number 200
BGP table version is 15, main routing table version 15

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.12.2       4   200      37      37       15    0    0 00:26:26        0
10.1.13.3       4   100      38      33       15    0    0 00:24:59        0

R2#show ip bgp summary
BGP router identifier 10.1.2.2, local AS number 200
BGP table version is 15, main routing table version 15

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.12.1       4   200      37      37       15    0    0 00:26:27        0
10.1.24.4       4   100      37      32       15    0    0 00:24:58        0

R3#show ip bgp summary
BGP router identifier 10.1.3.3, local AS number 100
BGP table version is 1, main routing table version 1

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.4.4        4   100      27      27        1    0    0 00:24:45        0
10.1.5.5        4   100      28      27        1    0    0 00:24:59        0
10.1.6.6        4   100      28      27        1    0    0 00:24:59        0
10.1.13.1       4   200      29      28        1    0    0 00:25:02        0
```


BGP Peering Verification (cont.)

```
R4#show ip bgp summary
BGP router identifier 10.1.4.4, local AS number 100
BGP table version is 1, main routing table version 1

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.3.3        4   100      27      27        1    0    0 00:24:47        0
10.1.5.5        4   100      29      28        1    0    0 00:25:02        0
10.1.6.6        4   100      28      27        1    0    0 00:24:54        0
10.1.24.2       4   200      29      28        1    0    0 00:25:01        0

R5#show ip bgp summary
BGP router identifier 10.1.5.5, local AS number 100
BGP table version is 34, main routing table version 34

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.3.3        4   100      55      46       34    0    0 00:25:01        0
10.1.4.4        4   100      56      46       34    0    0 00:25:03        0
10.1.6.6        4   100      43      43       34    0    0 00:40:12        0

R6#show ip bgp summary
BGP router identifier 10.1.6.6, local AS number 100
BGP table version is 34, main routing table version 34

Neighbor        V    AS MsgRcvd MsgSent   TblVer  InQ OutQ Up/Down  State/PfxRcd
10.1.3.3        4   100      55      46       34    0    0 00:25:03        0
10.1.4.4        4   100      55      45       34    0    0 00:24:57        0
10.1.5.5        4   100      43      43       34    0    0 00:40:13        0
```


Originating NLRI Configuration

```
R1#
router bgp 200
 network 10.1.1.0 mask 255.255.255.0
 network 10.1.2.0 mask 255.255.255.0
 network 10.1.12.0 mask 255.255.255.0

R2#
router bgp 200
 network 10.1.1.0 mask 255.255.255.0
 network 10.1.2.0 mask 255.255.255.0
 network 10.1.12.0 mask 255.255.255.0

R3#
router bgp 100
 network 10.1.3.0 mask 255.255.255.0
 network 10.1.4.0 mask 255.255.255.0
 network 10.1.5.0 mask 255.255.255.0
 network 10.1.6.0 mask 255.255.255.0
 network 10.1.35.0 mask 255.255.255.0
 network 10.1.46.0 mask 255.255.255.0
 network 10.1.56.0 mask 255.255.255.0

R4#
router bgp 100
 network 10.1.3.0 mask 255.255.255.0
 network 10.1.4.0 mask 255.255.255.0
 network 10.1.5.0 mask 255.255.255.0
 network 10.1.6.0 mask 255.255.255.0
 network 10.1.35.0 mask 255.255.255.0
 network 10.1.46.0 mask 255.255.255.0
 network 10.1.56.0 mask 255.255.255.0
```


BGP Table Verification

```
R1#show ip bgp
BGP table version is 28, local router ID is 10.1.1.1
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
*> 10.1.1.0/24      0.0.0.0                  0         32768 i
*>i10.1.2.0/24      10.1.12.2                0    100      0 i
*> 10.1.3.0/24      10.1.13.3                0             0 100 i
*>i10.1.4.0/24      10.1.24.4                0    100      0 100 i
*                   10.1.13.3           161280             0 100 i
*> 10.1.5.0/24      10.1.13.3           156160             0 100 i
*  10.1.6.0/24      10.1.13.3           158720             0 100 i
*>i                 10.1.24.4           156160    100      0 100 i
* i10.1.12.0/24     10.1.12.2                0    100      0 i
*>                  0.0.0.0                  0         32768 i
*> 10.1.35.0/24     10.1.13.3                0             0 100 i
*  10.1.46.0/24     10.1.13.3            33280             0 100 i
*>i                 10.1.24.4                0    100      0 100 i
*> 10.1.56.0/24     10.1.13.3            30720             0 100 i
* i                 10.1.24.4            30720    100      0 100 i
```


BGP Table Verification (cont.)

```
R2#show ip bgp
BGP table version is 30, local router ID is 10.1.2.2
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
*>i10.1.1.0/24      10.1.12.1                0    100      0 i
*> 10.1.2.0/24      0.0.0.0                  0         32768 i
*  10.1.3.0/24      10.1.24.4           161280             0 100 i
*>i                 10.1.13.3                0    100      0 100 i
*> 10.1.4.0/24      10.1.24.4                0             0 100 i
*>i10.1.5.0/24      10.1.13.3           156160    100      0 100 i
*                   10.1.24.4           158720             0 100 i
*> 10.1.6.0/24      10.1.24.4           156160             0 100 i
*> 10.1.12.0/24     0.0.0.0                  0         32768 i
* i                 10.1.12.1                0    100      0 i
*>i10.1.35.0/24     10.1.13.3                0    100      0 100 i
*                   10.1.24.4            33280             0 100 i
*> 10.1.46.0/24     10.1.24.4                0             0 100 i
* i10.1.56.0/24     10.1.13.3            30720    100      0 100 i
*>                  10.1.24.4            30720             0 100 i
```


BGP Table Verification (cont.)

```
R3#show ip bgp
BGP table version is 11, local router ID is 10.1.3.3
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
* i10.1.1.0/24      10.1.24.2                0    100      0 200 i
*>                  10.1.13.1                0             0 200 i
* i10.1.2.0/24      10.1.24.2                0    100      0 200 i
*>                  10.1.13.1                             0 200 i
* i10.1.3.0/24      10.1.46.6           161280    100      0 i
*>                  0.0.0.0                  0         32768 i
* i10.1.4.0/24      10.1.4.4                 0    100      0 i
*>                  10.1.35.5           161280         32768 i
* i10.1.5.0/24      10.1.46.6           158720    100      0 i
*>                  10.1.35.5           156160         32768 i
* i10.1.6.0/24      10.1.46.6           156160    100      0 i
*>                  10.1.35.5           158720         32768 i
* i10.1.12.0/24     10.1.24.2                0    100      0 200 i
*>                  10.1.13.1                0             0 200 i
* i10.1.35.0/24     10.1.46.6            33280    100      0 i
*>                  0.0.0.0                  0         32768 i
* i10.1.46.0/24     10.1.4.4                 0    100      0 i
*>                  10.1.35.5            33280         32768 i
* i10.1.56.0/24     10.1.46.6            30720    100      0 i
*>                  10.1.35.5            30720         32768 I
```


BGP Table Verification (cont.)

```
R4#show ip bgp
BGP table version is 26, local router ID is 10.1.4.4
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
*> 10.1.1.0/24      10.1.24.2                             0 200 i
* i                 10.1.13.1                0    100      0 200 i
* i10.1.2.0/24      10.1.13.1                0    100      0 200 i
*>                  10.1.24.2                0             0 200 i
*> 10.1.3.0/24      10.1.46.6           161280         32768 i
* i                 10.1.3.3                 0    100      0 i
*> 10.1.4.0/24      0.0.0.0                  0         32768 i
* i                 10.1.35.5           161280    100      0 i
*> 10.1.5.0/24      10.1.46.6           158720         32768 i
* i                 10.1.35.5           156160    100      0 i
*> 10.1.6.0/24      10.1.46.6           156160         32768 i
* i                 10.1.35.5           158720    100      0 i
* i10.1.12.0/24     10.1.13.1                0    100      0 200 i
*>                  10.1.24.2                0             0 200 i
*> 10.1.35.0/24     10.1.46.6            33280         32768 i
* i                 10.1.3.3                 0    100      0 i
*> 10.1.46.0/24     0.0.0.0                  0         32768 i
* i                 10.1.35.5            33280    100      0 i
*> 10.1.56.0/24     10.1.46.6            30720         32768 i
* i                 10.1.35.5            30720    100      0 i
```


BGP Table Verification (cont.)

```
R5#show ip bgp
BGP table version is 52, local router ID is 10.1.5.5
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
*>i10.1.1.0/24      10.1.24.2                0    100      0 200 i
* i                 10.1.13.1                0    100      0 200 i
* i10.1.2.0/24      10.1.13.1                0    100      0 200 i
*>i                 10.1.24.2                0    100      0 200 i
r i10.1.3.0/24      10.1.46.6           161280    100      0 i
r>i                 10.1.3.3                 0    100      0 i
r>i10.1.4.0/24      10.1.4.4                 0    100      0 i
r>i10.1.5.0/24      10.1.46.6           158720    100      0 i
r>i10.1.6.0/24      10.1.46.6           156160    100      0 i
* i10.1.12.0/24     10.1.13.1                0    100      0 200 i
*>i                 10.1.24.2                0    100      0 200 i
r i10.1.35.0/24     10.1.46.6            33280    100      0 i
r>i                 10.1.3.3                 0    100      0 i
r>i10.1.46.0/24     10.1.4.4                 0    100      0 i
r>i10.1.56.0/24     10.1.46.6            30720    100      0 i
```


BGP Table Verification (cont.)

```
R6#show ip bgp
BGP table version is 54, local router ID is 10.1.6.6
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
*>i10.1.1.0/24      10.1.24.2                0    100      0 200 i
* i                 10.1.13.1                0    100      0 200 i
* i10.1.2.0/24      10.1.13.1                0    100      0 200 i
*>i                 10.1.24.2                0    100      0 200 i
r>i10.1.3.0/24      10.1.3.3                 0    100      0 i
r>i10.1.4.0/24      10.1.4.4                 0    100      0 i
r i                 10.1.35.5           161280    100      0 i
r>i10.1.5.0/24      10.1.35.5           156160    100      0 i
r>i10.1.6.0/24      10.1.35.5           158720    100      0 i
* i10.1.12.0/24     10.1.13.1                0    100      0 200 i
*>i                 10.1.24.2                0    100      0 200 i
r>i10.1.35.0/24     10.1.3.3                 0    100      0 i
r>i10.1.46.0/24     10.1.4.4                 0    100      0 i
r i                 10.1.35.5            33280    100      0 i
r>i10.1.56.0/24     10.1.35.5            30720    100      0 I
```


BGP Table Verification Detail

```
R1#show ip bgp 10.1.56.0
BGP routing table entry for 10.1.56.0/24, version 25
Paths: (2 available, best #1, table Default-IP-Routing-
  Table)
  Advertised to update-groups:
     2
  100
    10.1.13.3 from 10.1.13.3 (10.1.3.3)
      Origin IGP, metric 30720, localpref 100, valid,
  external, best
  100
    10.1.24.4 (metric 20) from 10.1.12.2 (10.1.2.2)
      Origin IGP, metric 30720, localpref 100, valid,
  internal
```


Modifying BGP Next-Hop Configuration

```
R4#
router bgp 100
 neighbor 10.1.3.3 next-hop-self
 neighbor 10.1.5.5 next-hop-self
 neighbor 10.1.6.6 next-hop-self

R5#show ip bgp
BGP table version is 58, local router ID is 10.1.5.5
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
*>i10.1.1.0/24      10.1.4.4                 0    100      0 200 i
* i                 10.1.13.1                0    100      0 200 i
* i10.1.2.0/24      10.1.13.1                0    100      0 200 i
*>i                 10.1.4.4                 0    100      0 200 i
r i10.1.3.0/24      10.1.4.4            161280    100      0 i
r>i                 10.1.3.3                 0    100      0 i
r>i10.1.4.0/24      10.1.4.4                 0    100      0 i
r>i10.1.5.0/24      10.1.4.4            158720    100      0 i
r>i10.1.6.0/24      10.1.4.4            156160    100      0 i
* i10.1.12.0/24     10.1.13.1                0    100      0 200 i
*>i                 10.1.4.4                 0    100      0 200 i
r i10.1.35.0/24     10.1.4.4             33280    100      0 i
r>i                 10.1.3.3                 0    100      0 i
r>i10.1.46.0/24     10.1.4.4                 0    100      0 i
r>i10.1.56.0/24     10.1.4.4             30720    100      0 I
```


BGP Aggregation Configuration

```
R3#
router bgp 100
 aggregate-address 10.1.4.0 255.255.254.0 summary-only

R1#show ip bgp
BGP table version is 31, local router ID is 10.1.1.1
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
*> 10.1.1.0/24      0.0.0.0                  0         32768 i
*>i10.1.2.0/24      10.1.12.2                0    100      0 i
*> 10.1.3.0/24      10.1.13.3                0             0 100 i
*>i10.1.4.0/24      10.1.24.4                0    100      0 100 i
* i10.1.4.0/23      10.1.24.4                0    100      0 100 i
*>                  10.1.13.3                0             0 100 i
*>i10.1.5.0/24      10.1.24.4           158720    100      0 100 i
*  10.1.6.0/24      10.1.13.3           158720             0 100 i
*>i                 10.1.24.4           156160    100      0 100 i
* i10.1.12.0/24     10.1.12.2                0    100      0 i
*>                  0.0.0.0                  0         32768 i
*> 10.1.35.0/24     10.1.13.3                0             0 100 i
*  10.1.46.0/24     10.1.13.3            33280             0 100 i
*>i                 10.1.24.4                0    100      0 100 i
*> 10.1.56.0/24     10.1.13.3            30720             0 100 i
i                 10.1.24.4            30720    100      0 100 I
```


BGP Weight Configuration

```
R2#show ip bgp 10.1.3.0
BGP routing table entry for 10.1.3.0/24, version 28
Paths: (2 available, best #2, table Default-IP-Routing-Table)
  Advertised to update-groups:
     1
  100
    10.1.24.4 from 10.1.24.4 (10.1.4.4)
      Origin IGP, metric 161280, localpref 100, valid, external
  100
    10.1.13.3 (metric 20) from 10.1.12.1 (10.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, internal, best

R2#config t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)#ip prefix-list R3_LOOPBACK permit 10.1.3.0/24
R2(config)#route-map R4_INBOUND permit 10
R2(config-route-map)#match ip address prefix-list R3_LOOPBACK
R2(config-route-map)#set weight 100
R2(config-route-map)#route-map R4_INBOUND permit 100
R2(config-route-map)#router bgp 200
R2(config-router)#neighbor 10.1.24.4 route-map R4_INBOUND in
R2(config-router)#end
R2#clear ip bgp * in
R2#show ip bgp 10.1.3.0
BGP routing table entry for 10.1.3.0/24, version 34
Paths: (2 available, best #1, table Default-IP-Routing-Table)
Flag: 0x4940
  Advertised to update-groups:
     2
  100
    10.1.24.4 from 10.1.24.4 (10.1.4.4)
      Origin IGP, metric 161280, localpref 100, weight 100, valid, external, best
  100
    10.1.13.3 (metric 20) from 10.1.12.1 (10.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, internal
```


BGP Local Preference Configuration

```
R3#show ip bgp 10.1.1.0
BGP routing table entry for 10.1.1.0/24, version 2
Paths: (2 available, best #2, table Default-IP-Routing-Table)
  Advertised to update-groups:
     2
  200
    10.1.4.4 (metric 161280) from 10.1.4.4 (10.1.4.4)
      Origin IGP, metric 0, localpref 100, valid, internal
  200
    10.1.13.1 from 10.1.13.1 (10.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, external, best

R4#show ip bgp 10.1.1.0
BGP routing table entry for 10.1.1.0/24, version 3
Paths: (2 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
     2
  200
    10.1.24.2 from 10.1.24.2 (10.1.2.2)
      Origin IGP, localpref 100, valid, external, best
  200
    10.1.13.1 (metric 20519680) from 10.1.3.3 (10.1.3.3)
      Origin IGP, metric 0, localpref 100, valid, internal

R3#traceroute 10.1.1.1

Type escape sequence to abort.
Tracing the route to 10.1.1.1

  1 10.1.13.1 12 msec *  13 msec

R4#traceroute 10.1.1.1

Type escape sequence to abort.
Tracing the route to 10.1.1.1

  1 10.1.24.2 28 msec 28 msec 28 msec
  2 10.1.12.1 [AS 200] 28 msec *  28 msec
```


BGP Local Preference Configuration (cont.)

```
R4#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R4(config)#ip prefix-list R1_LOOPBACK permit 10.1.1.0/24
R4(config)#route-map R2_INBOUND permit 10
R4(config-route-map)#match ip address prefix-list R1_LOOPBACK
R4(config-route-map)#set local-pref
R4(config-route-map)#set local-preference 200
R4(config-route-map)#route-map R2_INBOUND permit 100
R4(config-route-map)#router bgp 100
R4(config-router)#neighbor 10.1.24.2 route-map R2_INBOUND in
R4(config-router)#end
R4#clear ip bgp * in
R4#show ip bgp 10.1.1.0
BGP routing table entry for 10.1.1.0/24, version 28
Paths: (1 available, best #1, table Default-IP-Routing-Table)
Flag: 0x800
  Advertised to update-groups:
     2
  200
    10.1.24.2 from 10.1.24.2 (10.1.2.2)
      Origin IGP, localpref 200, valid, external, best

R4#traceroute 10.1.1.1

Type escape sequence to abort.
Tracing the route to 10.1.1.1

  1 10.1.24.2 28 msec 28 msec 28 msec
  2 10.1.12.1 [AS 200] 32 msec *  28 msec
```

BGP Local Preference Configuration (cont.)

```
R3#show ip bgp 10.1.1.0
BGP routing table entry for 10.1.1.0/24, version 15
Paths: (2 available, best #1, table Default-IP-Routing-Table)
Flag: 0x940
  Advertised to update-groups:
     1
  200
    10.1.4.4 (metric 161280) from 10.1.4.4 (10.1.4.4)
      Origin IGP, metric 0, localpref 200, valid, internal, best
  200
    10.1.13.1 from 10.1.13.1 (10.1.1.1)
      Origin IGP, metric 0, localpref 100, valid, external

R3#traceroute 10.1.1.1

Type escape sequence to abort.
Tracing the route to 10.1.1.1

  1 10.1.35.5 4 msec 4 msec 0 msec
  2 10.1.56.6 4 msec 4 msec 0 msec
  3 10.1.46.4 4 msec 4 msec 4 msec
  4 10.1.24.2 20 msec 20 msec 20 msec
  5 10.1.12.1 [AS 200] 20 msec *  20 msec
```


BGP AS-Path Configuration

```
R3#show ip bgp 10.1.2.0
BGP routing table entry for 10.1.2.0/24, version 11
Paths: (2 available, best #2, table Default-IP-Routing-Table)
  Advertised to update-groups:
     2
  200
    10.1.4.4 (metric 161280) from 10.1.4.4 (10.1.4.4)
      Origin IGP, metric 0, localpref 100, valid, internal
  200
    10.1.13.1 from 10.1.13.1 (10.1.1.1)
      Origin IGP, localpref 100, valid, external, best

R4#show ip bgp 10.1.2.0
BGP routing table entry for 10.1.2.0/24, version 26
Paths: (2 available, best #2, table Default-IP-Routing-Table)
  Advertised to update-groups:
     2
  200
    10.1.13.1 (metric 20519680) from 10.1.3.3 (10.1.3.3)
      Origin IGP, metric 0, localpref 100, valid, internal
  200
    10.1.24.2 from 10.1.24.2 (10.1.2.2)
      Origin IGP, metric 0, localpref 100, valid, external, best

R2#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R2(config)#ip prefix-list R2_LOOPBACK permit 10.1.2.0/24
R2(config)#route-map R4_OUTBOUND permit 10
R2(config-route-map)#match ip address prefix-list R2_LOOPBACK
R2(config-route-map)#set as-path prepend 200 200 200
R2(config-route-map)#route-map R4_OUTBOUND permit 100
R2(config-route-map)#router bgp 200
R2(config-router)#neighbor 10.1.24.4 route-map R4_OUTBOUND out
R2(config-router)#end
R2#clear ip bgp * out
```

BGP AS-Path Configuration (cont.)

```
R3#show ip bgp 10.1.2.0
BGP routing table entry for 10.1.2.0/24, version 11
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
     2
  200
    10.1.13.1 from 10.1.13.1 (10.1.1.1)
      Origin IGP, localpref 100, valid, external, best

R4#show ip bgp 10.1.2.0
BGP routing table entry for 10.1.2.0/24, version 29
Paths: (2 available, best #1, table Default-IP-Routing-Table)
Flag: 0x940
  Advertised to update-groups:
     1
  200
    10.1.13.1 (metric 20519680) from 10.1.3.3 (10.1.3.3)
      Origin IGP, metric 0, localpref 100, valid, internal, best
  200 200 200 200
    10.1.24.2 from 10.1.24.2 (10.1.2.2)
      Origin IGP, metric 0, localpref 100, valid, external

R4#traceroute 10.1.2.2

Type escape sequence to abort.
Tracing the route to 10.1.2.2

  1 10.1.46.6 4 msec 4 msec 0 msec
  2 10.1.56.5 4 msec 4 msec 0 msec
  3 10.1.35.3 0 msec 4 msec 4 msec
  4 10.1.13.1 24 msec 24 msec 28 msec
  5 10.1.12.2 [AS 200] 24 msec *  24 msec
```


BGP MED Configuration

```
R1#show ip bgp 10.1.6.0
BGP routing table entry for 10.1.6.0/24, version 27
Paths: (2 available, best #2, table Default-IP-Routing-Table)
  Advertised to update-groups:
     1
  100
    10.1.13.3 from 10.1.13.3 (10.1.3.3)
      Origin IGP, metric 158720, localpref 100, valid, external
  100
    10.1.24.4 (metric 20) from 10.1.12.2 (10.1.2.2)
      Origin IGP, metric 156160, localpref 100, valid, internal, best

R2#show ip bgp 10.1.6.0
BGP routing table entry for 10.1.6.0/24, version 25
Paths: (1 available, best #1, table Default-IP-Routing-Table)
  Advertised to update-groups:
     2
  100
    10.1.24.4 from 10.1.24.4 (10.1.4.4)
      Origin IGP, metric 156160, localpref 100, valid, external, best

R3#conf t
Enter configuration commands, one per line.  End with CNTL/Z.
R3(config)#ip prefix-list R6_LOOPBACK permit 10.1.6.0/24
R3(config)#route-map R1_OUTBOUND permit 10
R3(config-route-map)#matc ip address prefix-list R6_LOOPBACK
R3(config-route-map)#match ip address prefix-list R6_LOOPBACK
R3(config-route-map)#set metric 100
R3(config-route-map)#route-map R1_OUTBOUND permit 100
R3(config-route-map)#router bgp 100
R3(config-router)#neighbor 10.1.13.1 route-map R1_OUTBOUND out
R3(config-router)#end
R3#clear ip bgp * out
```

BGP MED Configuration (cont.)

```
R1#show ip bgp 10.1.6.0
BGP routing table entry for 10.1.6.0/24, version 32
Paths: (1 available, best #1, table Default-IP-Routing-Table)
Flag: 0x4940
  Advertised to update-groups:
     2
  100
    10.1.13.3 from 10.1.13.3 (10.1.3.3)
      Origin IGP, metric 100, localpref 100, valid, external, best

R2#show ip bgp 10.1.6.0
BGP routing table entry for 10.1.6.0/24, version 35
Paths: (2 available, best #1, table Default-IP-Routing-Table)
Flag: 0x4940
  Advertised to update-groups:
     1
  100
    10.1.13.3 (metric 20) from 10.1.12.1 (10.1.1.1)
      Origin IGP, metric 100, localpref 100, valid, internal, best
  100
    10.1.24.4 from 10.1.24.4 (10.1.4.4)
      Origin IGP, metric 156160, localpref 100, valid, external
```


BGP Q&A