BGP Path Attributes

Attributes Categories & Types:

1. **Well-known Mandatory:** Recognized and Included in all BGP Update messages. Such as **ORIGIN , AS_PATH , NEXT_HOP**
2. **Well-known Discretionary:** Recognized and May or May not include in BGP Update messages, such as **LOCAL_PREF , ATOMIC_AGGREGATE** (atomic aggregate indicates that the prefix is an aggregate and not to an actual destination , not used now days)
3. **Optional Transitive:** Even if Not Supported it Still need to accept and Send in Update Message. Such as **AGGREGATOR , COMMUNITY**
4. **Optional Non-transitive:** Can be ignored and not advertise to peers. Such as **MULTI_EXIT_DISC , ORIGINATOR_ID , CLUSTER_LIST**

BGP Best path selection in order:

w=weight
L=local preference
L=locally injected routes
A=AS path
O=Origin code
M=MED
N=ebgp network over ibgp
i= igp cost to next-hop
L=lowest router ID
L=lowest Neighbor IP
http://www.cisco.com/c/en/us/support/docs/ip/border-gateway-protocol-bgp/13753-25.html

S. No.	Attribute	Description	Preference
1	Weight	Administrative Weight	Highest
2	Local Preference	Communicated between peers within an AS	Highest
3	Self-Originated	Prefer path originated locally	next hop 0.0.0.0
4	AS Path	AS hop required to reach destination	Shortes
5	Origin Code	IGP (i) < EGP (e) < Redistributed / Incomplete (?)	Lowest (IGP)
6	MED	Used by externally peers to enter an AS	Lowest
7	BGP AD	Prefer eBGP routes over iBGP	eBG
8	IGP Cost	IGP metric to reach next hop or neighbor	Lowest
9	eBGP Peering	Compared the age of the route	Oldest
10	Router ID	Compares router ID of neighbor	Lowest
11	Neighbor IP	Compares IP address of the neghbor	Lowest

Attributes Best Path Selection: lets explain each attribute and understand best path selection in details):

To find best path before start algorithm next hop must be reachable
Let say you type sh ip bgp 1.100.100.100 and inaccessible shown that mean next hop is not reachable

Weight

1- Weight only work in cisco routers (**Cisco proprietary**), **highest preferred**
Valid locally only (has meaning for router you typed it in ONLY)
Default weight for received prefixes is 0
Default weight of locally originated prefixes is 32768

R1 has two exit point, R1 will prefer exit point with highest number on another

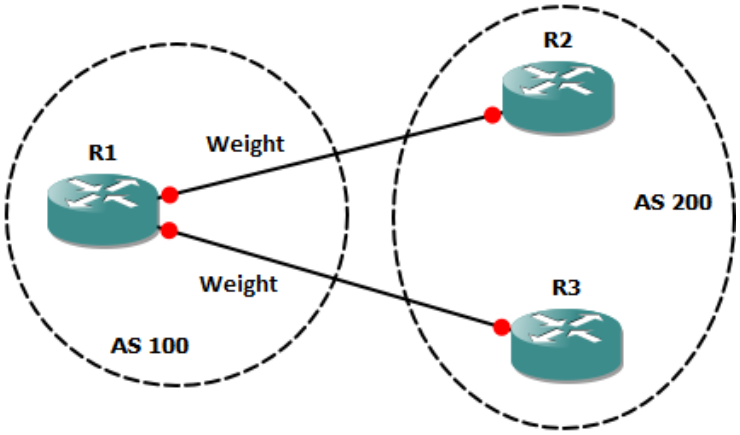

Highest Weight will prefered as exit point to AS 200

onfigured per neighbor for all prefixes

```
router bgp 1
neighbor 192.168.13.3 weight 500
```

Or configured per neighbor for specific prefixes by route-map

```
route-map SETWEIGHT permit 10
match ip address 1
set weight 400
exit
route-map SETWEIGHT permit 20
set weight 0
exit
access-list 1 permit 22.22.22.0 0.0.0.255

router bgp 1
neighbor 192.168.13.3 route-map SETWEIGHT in
```

Local -Preference

2-Local -Preference highest Preferred

Default is 100

Exchanged between all IBGP peers to decide which exit point will be preferred

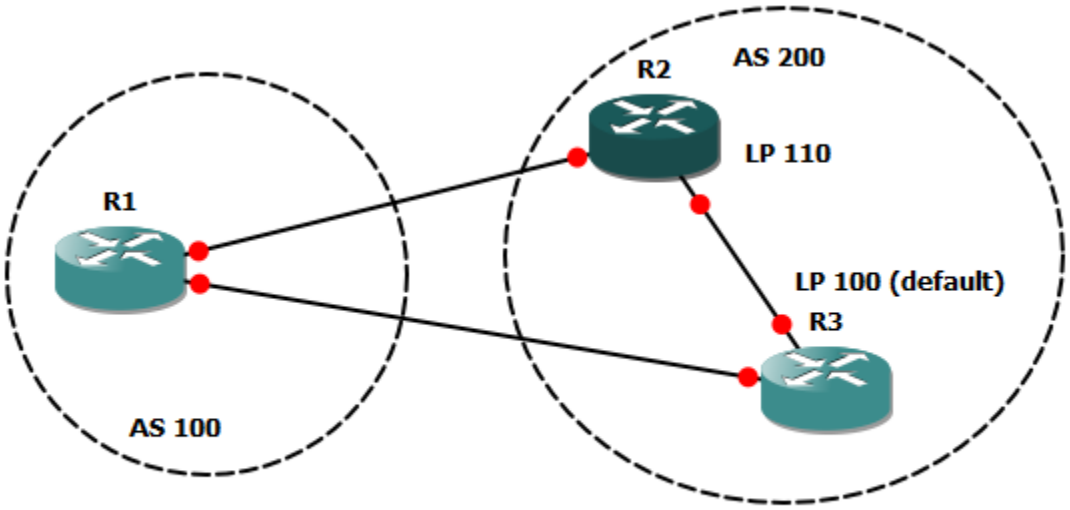

Highest Local Preferences will prefered as exit point to AS 100

Configured to be exchanged to all IBGP neighbors

router bgp 2

bgp default local-preference 600

Or configured to be exchanged with specific IBGP neighbor

route-map LOCALPREF permit 10

set local-preference 700

router bgp 2

neighbor 192.168.13.1 route-map LOCALPREF in

If you don't use a match statement in a route-map then all prefixes is matched by default.

Prefer Locally Originated prefix

3-Prefer Locally Originated prefix (with next hop 0.0.0.0) over same prefix we learn from neighbor

Notice : Path originated locally via network or redistribute command preferred over prefix originated by aggregate-address command

AS_PATH

4- Prefer Shortest AS_PATH

To skip this step we can configure **bgp bestpath-as ignore** So to not consider the autonomous system (AS) path during best path route selection .

we can repeat some AS numbers by prepend to manipulate AS Path , will talk about this later

```
route-map PREPEND permit 10
set as-path prepend 1 1 1 1 1
exit
router bgp 1
neighbor 192.168.12.2 route-map PREPEND out
```

Sending this to your remote neighbor it should be **outbound**!

show ip bgp

```
   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 1 1 1 1 1
```

Also we can remove Private AS from AS-Path

neighbor 10.0.0.1 remove-private-as

remove private AS from the path in outbound routing updates

ORIGIN Type

5-Prefer Lowest ORIGIN Type

IGP >EGP>Incomplete

i (IGP) prefix learned internal to the AS by IBGP or added to bgp with network command
e (EGP) not used now days
? (Incomplete) prefix leaned through redistribution

To Manipulate Origin Type:

```
route-map yasser permit 10
match as-path 10
set origin igp
exit
router bgp 100
redistribute eigrp 51 route-map yasser
```

MED

6-MED (Metric) Lowest Preferred

Exchanged between EBGP Peers to tell the external neighbor the suggest exit point to come to our AS

Default is 0

In R3 med is only comparing when same coming prefix came many times from different neighbors from the same AS

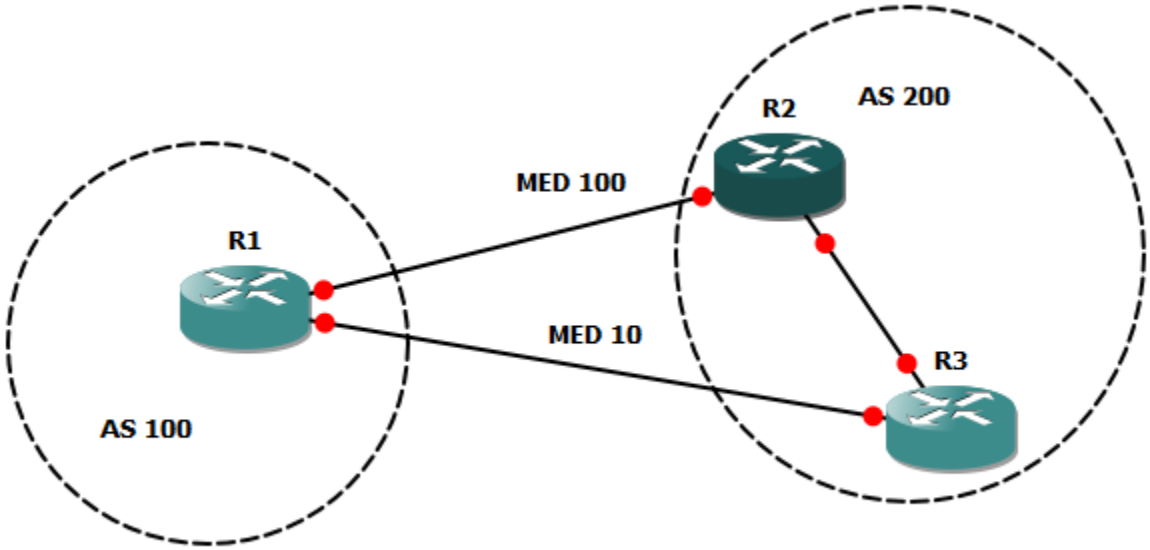

Lowest MED (Metric) will force EBGP nieghbor to prefered as exit point to AS 200

R2

```
route-map MED permit 10
set metric 100
exit
router bgp 1
neighbor 192.168.12.2 route-map MED out
```

R3

```
route-map MED permit 10
set metric 10
exit
router bgp 1
neighbor 192.168.23.2 route-map MED out
```

- MED can be used to advertise to your neighbors how they should **enter your AS**.
- MED is exchanged **between** autonomous systems.
- MED is propagated to all routers within the neighbor AS but not passed along any other autonomous systems.

Enabling the **bgp deterministic-med** command ensures the comparison of the MED variable when choosing routes advertised by different peers in the same autonomous system. Enabling the **bgp always-compare-med** command ensures the comparison of the MED for paths from neighbors in different autonomous systems.
Note: The **bgp deterministic-med** and **bgp always-compare-med** commands are not enabled by default. Also, the two commands are separate; enabling one does not automatically enable the other.
Note: **bgp-always-compare-med** used to compare med and typed in the neighbor will do this compare

7- Prefer EBGP over IBGP path

8- Prefer Lowest IGP Metric to NEXT_HOP

9-Prefer Oldest Received Path
when two paths leaned by EBGP the oldest Preferred , this prevent route flapping.
We can skip this step and go to step 10 using:
bgp best path comapre-router id

10-Prefer Path From Lowest Router ID
BGP got his router id with same methods we use with ospf (router-id command , virtual , physical)
To configure a fixed router ID for the local BGP routing process, use the **bgp router-id** command in router or address family configuration mode.

Route Reflector Specific Case
(Prefer the path directly peered client over a path that was reflected by RR)

12- Prefer Lowest Neighbor ID

Finally there are two questions important to answer and related to what we just learned :
What does a next hop of 0.0.0.0 mean in the show ip bgp command output?
A network in the BGP table with a next hop address of 0.0.0.0 means that the network is locally originated via redistribution of Interior Gateway Protocol (IGP) into BGP, or via a network or aggregate command in the BGP configuration.

```
R5#sh ip bgp
BGP table version is 25, local router ID is 5.5.5.5
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
*> 1.1.1.1/32       25.25.25.2              11             0 123 ?
*> 2.2.2.2/32       25.25.25.2               0             0 123 ?
*> 3.3.3.3/32       25.25.25.2              11             0 123 ?
*> 4.4.4.4/32       45.45.45.4          409600         32768 ?
* i                  45.45.45.4               0    100      0 ?
*> 5.5.5.5/32       0.0.0.0                  0         32768 ?
* i                  45.45.45.4          409600    100      0 ?
```

IF you see `r RIB-Failure`

```
   Network          Next Hop            Metric LocPrf Weight Path
r i56.56.56.0/24    46.46.46.6               0    100      0 600 ?
r>                  56.56.56.6               0             0 600 ?
r i56.56.56.5/32    46.46.46.6               0    100      0 600 ?
r>                  56.56.56.6               0             0 600 ?
```

This normal Route with better administrative distance already present in IGP. For example, if a static or directly connected route already exists in IP Routing table. Type sh ip route in R5:

```
C       56.56.56.0/24 is directly connected, Serial0/1
C       56.56.56.6/32 is directly connected, Serial0/1
```

A rib failure usually means that the route has entered the routing table via another source.
And because IBGP AD is 200 and EIGRP AD is 90 (Internal) and 170 (External). So the BGP route cannot be inserted into the routing table because the EIGRP route is preferred and in the route table.
run "sh ip bgp rib-failure to see more info

What is the difference between when a route is injected in BGP via redistribute command or a network command?

When you use the redistribution of IGP into BGP to advertise the route, then there is no need to specify the network statement for all the subnets individually. Also when the route is obtained from any other routing protocols into BGP table by redistribution, the Origin attribute is Incomplete (?) and when you specify the network command then it is Internal/IGP (i). During the route selection, the lowest origin code is preferred (IGP<EGP<Incomplete).

HowTo display what u advertise to neighbor?
sh ip bgp nei x.x.x.x advertised-routes

how To display what you received from neighbor ?
Sh ip bgp nei x.x.x.x routes or **sh ip bgp nei x.x.x.x received-routes**

Good Luck
CCIE & CCSI: Yasser Auda
https://www.facebook.com/YasserRamzyAuda
https://learningnetwork.cisco.com/people/yasserramzy/content
https://www.youtube.com/user/yasserramzyauda