

BGP Configuration Guide 2 – Cisco Routers

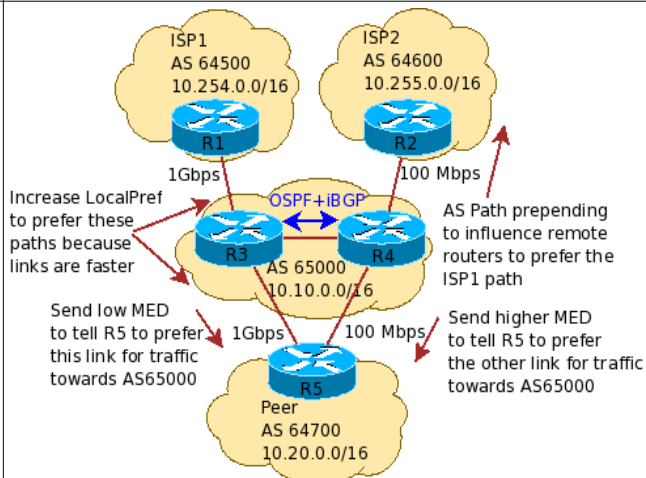


Configuration Example

```
hostname R3
!
interface Loopback 0
 ip address 10.10.254.3 255.255.255.255 ! Announce all these
!                                     ! subnets in IGP
interface GigabitEthernet0/0
 ip address 10.254.100.2 255.255.255.252 ! to ensure that BGP
!                                     ! paths always have
description link to ISP1                ! a next hop in the
!                                     ! forwarding table
interface GigabitEthernet0/1
 ip address 10.20.100.2 255.255.255.252 ! E.g. Use OSPF's
!                                     ! redistribute
description link to Peer                ! connected subnets
!
router BGP 65000
 bgp log-neighbor-changes
 no synchronization
 no auto-summary
 distance bgp 200 200 200
 bgp default local-preference 100
 neighbor 10.254.100.1 remote-as 64500 ! eBGP with ISP1
 neighbor 10.254.100.1 password N$RC
 neighbor 10.254.100.1 prefix-list from-ME out
 neighbor 10.254.100.1 route-map SET-LP-AS64500 in
 neighbor 10.20.100.1 remote-as 64700 ! eBGP with Peer
 neighbor 10.20.100.1 password N$RC
 neighbor 10.20.100.1 route-map SET-LP-AS64700 in
 neighbor 10.20.100.1 route-map SET-MED out
 neighbor 10.10.254.4 remote-as 65000 ! iBGP peer R4
 neighbor 10.10.254.4 password N$RC
 neighbor 10.10.254.4 update-source Loopback0
 network 10.10.0.0 mask 255.255.0.0
!
ip route 10.10.0.0 255.255.0.0 null0 250 ! Need for network cmd
!
ip prefix-list from-ME permit 10.10.0.0/16
!
ip prefix-list from-AS64500 deny 10.10.0.0/16 le 32
ip prefix-list from-AS64500 permit 0.0.0.0/0 le 24
!
ip prefix-list from-AS64700 deny 10.10.0.0/16 le 32
ip prefix-list from-AS64700 permit 10.20.0.0/16
!
route-map SET-LP-AS64500 permit 10 ! Higher LocalPref for
 match ip address prefix-list from-AS64500 ! prefixes coming via
 set local-preference 150 ! ISP1 and denies else
!
route-map SET-LP-AS64700 permit 10 ! Higher LocalPref
 match ip address prefix-list from-AS64700 ! and denies prefixes
 set local-preference-peer 150 ! not from peer
!
route-map SET-MED permit 10 ! Sets higher MED
 match ip address prefix-list from-ME ! and allows only my
 set metric 10 ! prefix out
```

```
hostname R5
!
interface Loopback 0
 ip address 10.20.254.5 255.255.255.255 ! Announce in IGP
!
interface FastEthernet0/0
 ip address 10.20.100.1 255.255.255.252 ! Announce in IGP
description link to R3
!
interface FastEthernet0/1
 ip address 10.20.200.1 255.255.255.252 ! Announce in IGP
description link to R4
!
router BGP 64700
 bgp log-neighbor-changes
 no synchronization
 no auto-summary
 distance bgp 200 200 200
 neighbor AS65000 peer-group ! Define peer group
 neighbor AS65000 remote-as 65000
 neighbor AS65000 password N$RC
 neighbor AS65000 prefix-list from-AS65000 in
 neighbor AS65000 filter-list 5 in
 neighbor AS65000 prefix-list from-ME out
 neighbor 10.20.100.1 peer-group AS65000 ! Assign neighbors
 neighbor 10.20.200.1 peer-group AS65000 ! to peer group
 network 10.20.0.0 mask 255.255.0.0
!
ip route 10.20.0.0 255.255.0.0 null0 250
!
ip prefix-list from-AS65000 deny 10.20.0.0/16 le 32
ip prefix-list from-AS65000 permit 10.10.0.0/16
ip prefix-list from-ME permit 10.20.0.0/16
! Only allow prefixes directly from AS65000 (AS_PATH length=1)
ip as-path access-list 5 permit ^65000$
```

```
hostname R4
!
interface Loopback 0
 ip address 10.10.254.4 255.255.255.255 ! Announce in IGP
!
interface FastEthernet0/0
 ip address 10.255.100.2 255.255.255.252 ! Announce in IGP
description link to ISP2
!
interface FastEthernet0/1
 ip address 10.20.200.2 255.255.255.252 ! Announce in IGP
description link to Peer
!
router BGP 65000
 bgp log-neighbor-changes
 no synchronization
 no auto-summary
 distance bgp 200 200 200
 bgp default local-preference 100
 neighbor 10.255.100.1 remote-as 64600 ! eBGP with ISP2
 neighbor 10.255.100.1 password N$RC
 neighbor 10.255.100.1 prefix-list from-AS64600 in
 neighbor 10.255.100.1 prefix-list from-ME out
 neighbor 10.255.100.1 route-map SET-PREPEND out
 neighbor 10.20.200.1 remote-as 64700 ! eBGP with Peer
 neighbor 10.20.200.1 password N$RC
 neighbor 10.20.200.1 prefix-list from-AS64700 in
 neighbor 10.20.200.1 prefix-list from-ME out
 neighbor 10.20.200.1 route-map SET-MED out
 neighbor 10.10.254.3 remote-as 65000 ! iBGP peer R3
 neighbor 10.10.254.3 password N$RC
 neighbor 10.10.254.3 update-source Loopback0
 network 10.10.0.0 mask 255.255.0.0
!
ip route 10.10.0.0 255.255.0.0 null0 250 ! Need for network cmd
!
ip prefix-list from-ME permit 10.10.0.0/16
!
ip prefix-list from-AS64600 deny 10.10.0.0/16 le 32
ip prefix-list from-AS64600 permit 0.0.0.0/0 le 24
!
ip prefix-list from-AS64700 deny 10.10.0.0/16 le 32
ip prefix-list from-AS64700 permit 10.20.0.0
!
route-map SET-PREPEND permit 10 ! Prepends my AS twice
 match ip address prefix-list from-ME ! in AS_PATH to infl.
 set as-path prepend 65000 65000 ! inbound traffic
route-map SET-PREPEND permit 20
!
route-map SET-MED permit 10 ! Sets lower MED and
 match ip address prefix-list from-ME ! and allows only my
 set metric 50 ! prefix out
```



BGP Attributes		
Attribute	Description	Type
Origin	How the route was originated (IGP, EGP, Incomplete)	Well-known Mandatory
AS Path	List of ASs traversed by the route advertisement	Well-known Mandatory
Next Hop	The next router to send the packet to for a given route	Well-known Mandatory
Local Preference	Metric to influence internal selection of paths for outbound traffic	Well-known Discretionary
Atomic Aggregate	Includes ASs not shown in the path because of route aggregation	Well-known Discretionary
Aggregator	ID and AS of router in the path that is aggregating prefixes	Optional Transitive
Community	A label assigned to a prefix or group of prefixes	Optional Transitive
Multiple Exit Discriminator (MED)	Metric sent to neighbor to influence their path selection for traffic destined to us	Optional Non-Transitive
Originator ID	Identification for a route reflector	Optional Non-Transitive
Cluster List	List of cluster IDs	Optional Non-Transitive
Cluster ID	Originating Cluster	Optional Non-Transitive
Weight	Preference local to router	Cisco proprietary

BGP Selection Process	
Order	Description
1	Do not consider path if there is no route to next hop (Internally, AS should run an IGP to announce loopbacks)
2	Highest Weight (Only Cisco)
3	Highest Local Preference (global within AS)
4	Prefer locally-originated route
5	Shortest AS Path
6	Lowest Origin Code IGP < EGP < Incomplete
7	Lowest Multiple Exit Discriminator (MED). Default is 0
8	Prefer eBGP over iBGP path
9	Path with Lowest IGP metric to next hop
10	For eBGP paths: <ul style="list-style-type: none"> • If multipath is enabled, install N parallel routes in forwarding table • If Router ID is not the same, select oldest route • if Router ID is the same, go to next step
11	Lowest Router ID (originator ID for reflected routes)
12	Shortest Cluster List (Client must be aware of Route Reflector attributes)
13	Lowest neighbor address

Prefix List Examples	
ip prefix-list mylist permit 10.10.0.0/16 le 32	(less than or equal /32) Allows all prefixes within 10.10.0.0/16, including 10.10.0.0/16
ip prefix-list mylist permit 10.20.0.0/16 le 24	Allows prefixes within 10.20.0.0/16 except /25, /26, /27, /28, /29, /30, /31 and /32
ip prefix-list mylist deny 10.20.0.0/16 ge 25	(greater than or equal /25) Denies prefixes within 10.20.0.0/16 of size /25, /26, /27, /28, /29, /30, /31 and /32 (same result as list above)

Regular Expressions for AS Path Matching			
.	Match anything	_65000_64500_	Via AS 65000 and AS 64500
+	Match at least one character	_(65000)_+	Any sequence of this same AS (prepending)
^\$	Match routes local to this AS	^[0-9]+\$	Match AS path of length 1 (neighbor ASs)
65000\$	Originated by AS 65000	^[0-9]*[0-9]+\$	Match AS path length 1 or 2
65000	Via AS 65000	_(100 200)_	Any path through either AS 100 or AS 200