

Consider only routes with no AS loops and a valid next hop.


Use Longest Prefix Match.

Where multiple routes are available to identical network and prefix:

- Prefer highest weight (local to router).
- Prefer highest local preference (global within AS).
- Prefer route originated by the local router ('network' command or redistribution).
- Prefer shortest AS path.
- Prefer lowest origin code: IGP ('network') < EGP (legacy) < incomplete (redistributed).
- Prefer lowest MED Multi Exit Discriminator (exchanged between autonomous systems).
- Prefer EBGP path over IBGP path.
- Prefer the path through the closest IGP neighbor.
- Prefer oldest route for EBGP paths.
- Prefer the path with the lowest neighbor BGP router ID.
- Prefer the path with the lowest neighbor IP address.

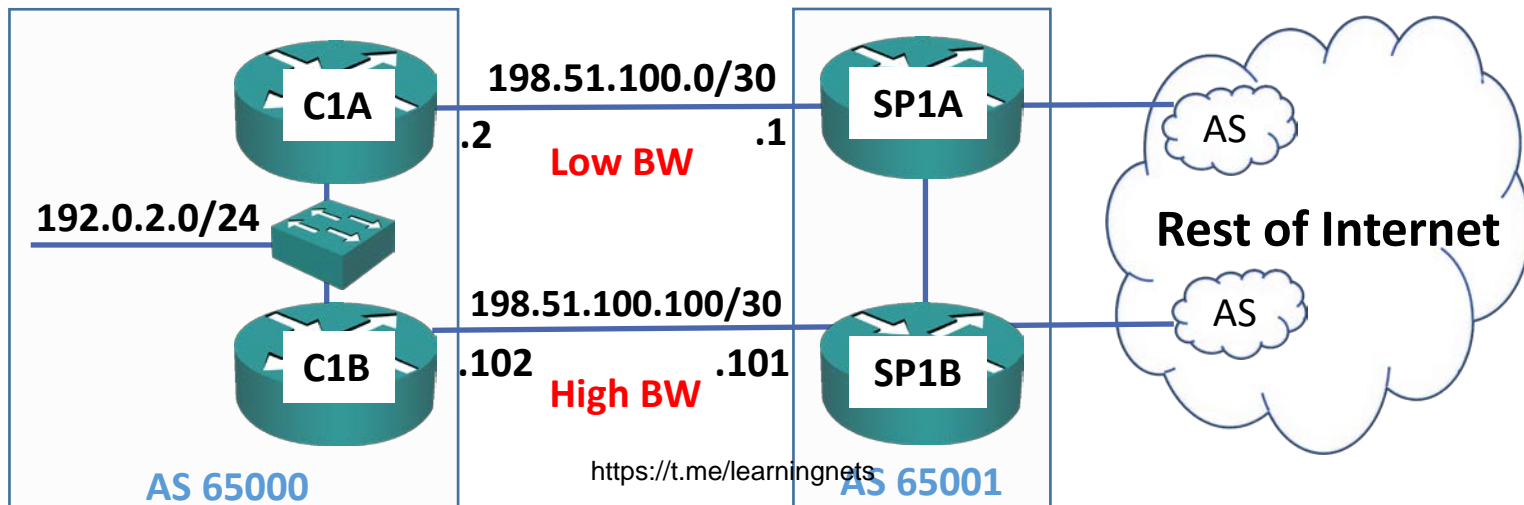
Influencing Inbound Path Selection

- In the example below, AS 65000 would prefer AS 65001 to send traffic to its 192.0.2.0/24 network over the high bandwidth link. AS 65001 could achieve this by setting a higher Local Preference for the link
- Internet Service Providers use fixed policies and configuration templates
- They typically avoid custom configurations per customer because it would add considerable complexity to their administration and troubleshooting
- MED provides a way to influence incoming traffic with no changes required in the neighbor AS



Influencing Inbound Path Selection

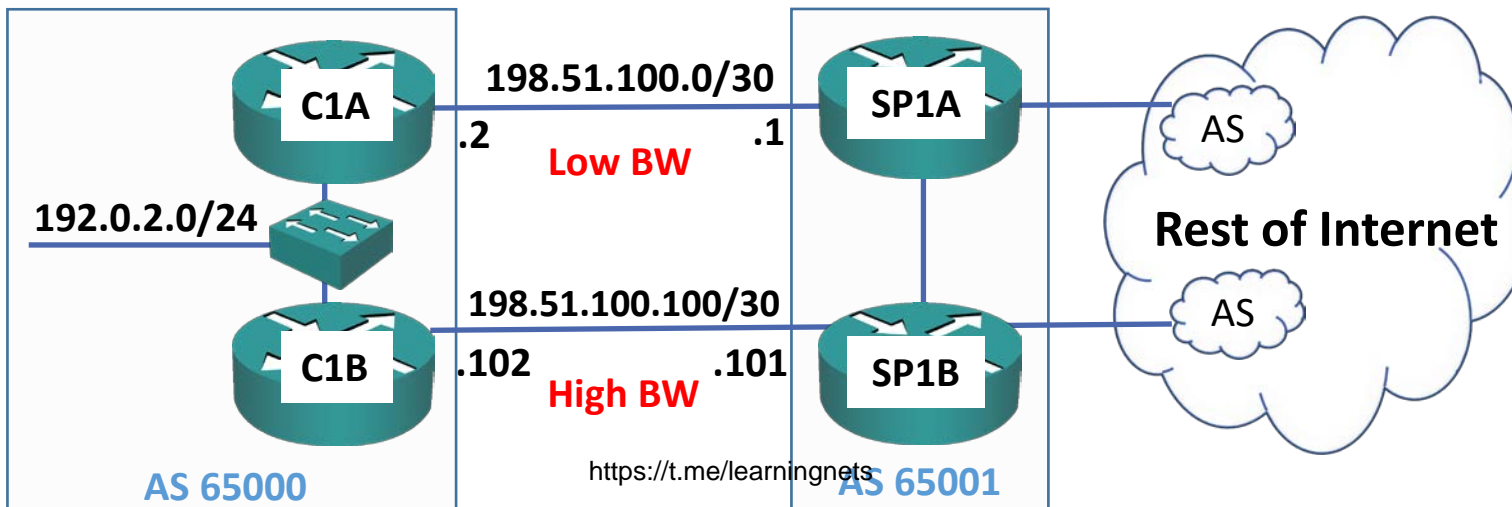
- With all other factors being equal, SP1A and SP1B will send traffic to 192.0.2.0/24 over their direct eBGP links to AS 65000:
 - Prefer highest weight (local to router).
 - Prefer highest local preference (global within AS).
 - Prefer route originated by the local router ('network' command or redistribution).
 - Prefer shortest AS path.
 - Prefer lowest origin code: IGP ('network') < EGP (legacy) < incomplete (redistributed).
 - Prefer lowest MED (exchanged between autonomous systems).
 - Prefer EBGP path over IBGP path.



Verification – show ip bgp (Before Policy)

```
SP1A#sh ip bgp
```


```
      Network                Next Hop                Metric LocPrf Weight Path
*i 192.0.2.0                 192.168.0.2 (SP1B)      0      100      0 65000 i
*>                          198.51.100.2 (C1A)      0      100      0 65000 i
```



Verification – show ip bgp (Before Policy)

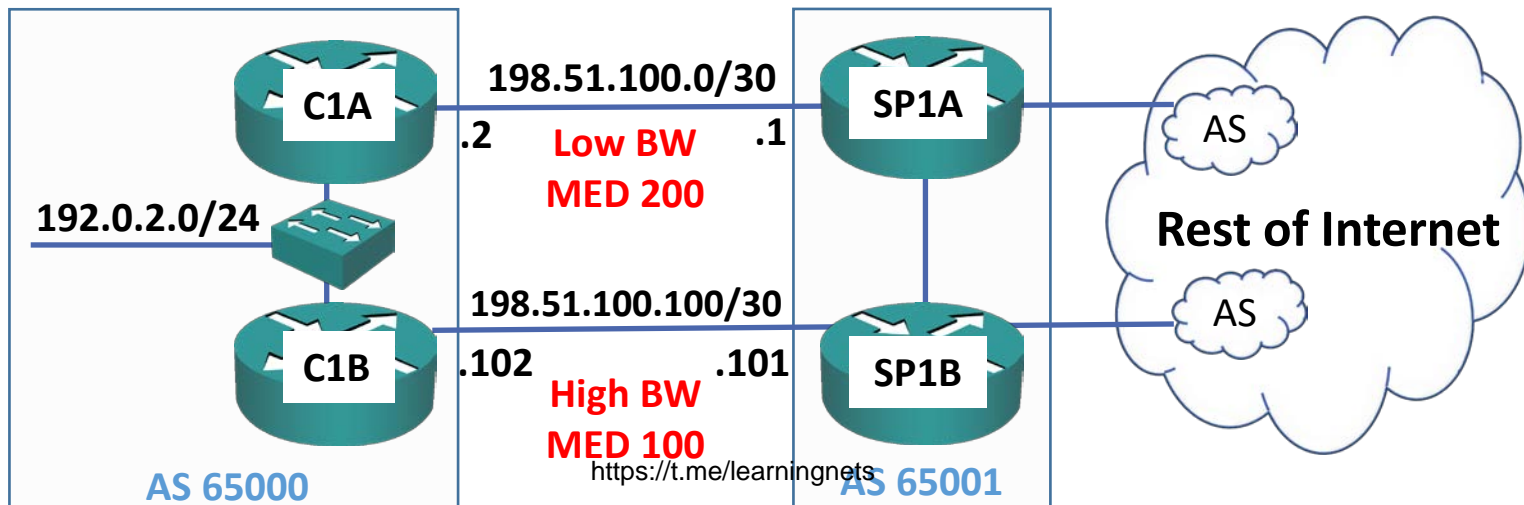
```
SP1B#sh ip bgp
```

```
      Network                Next Hop                Metric LocPrf Weight Path
*i  192.0.2.0                192.168.0.1 (SP1A)      0      100      0 65000 i
*> 198.51.100.101 (C1B)    0      100      0 65000 i
```



MED Multi Exit Discriminator


- MED (the 'metric') is used to advertise to EBGP neighbors how to exit their AS to reach networks owned by this AS
- MED influences inbound traffic to an AS. Weight and Local Preference influence outbound traffic from the AS
- Paths with a lower MED value are preferred (MED is also known as 'Metric')
- In the example below, AS 65000 advertises to AS 65001 a preferred MED value for the higher bandwidth link to its 192.0.2.0/24 network



MED Multi Exit Discriminator




- If the originating router is directly connected to the network then default MED is 0
- If originating router reaches network through static route or IGP then the default MED value is copied from the metric in the router's routing table
- The receiving router uses a default value of 0 if no MED is received
- When using MED, ensure you set it on all non-preferred eBGP routers facing the neighbor AS



MED Multi Exit Discriminator

- MED is only passed to direct neighbor ASs, they do not propagate it to further remote ASs
- In the example below, AS 65001 will advertise the 192.0.2.0/24 network to further downstream ASs in the rest of the Internet, but will not include the MED received from AS 65000 in the Update



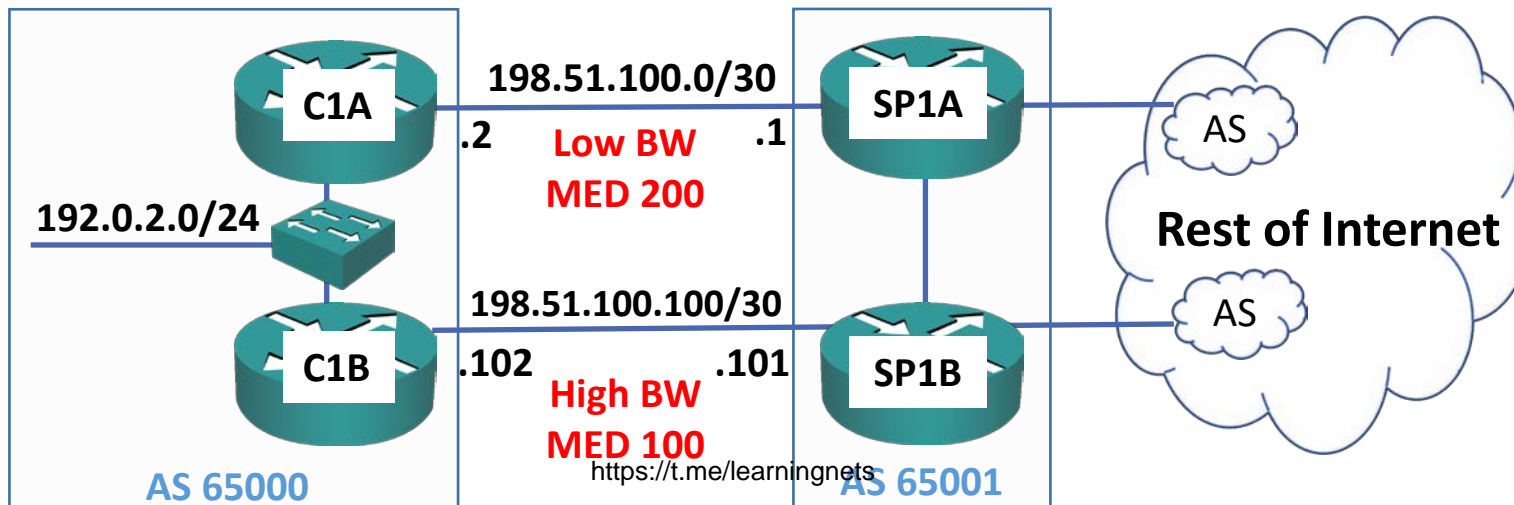
MED – Primary/Backup Links



```
C1A(config)#route-map DEMO permit 10
C1A(config-route-map)#set metric 200
C1A(config)#router bgp 65000
C1A(config-router)#neighbor 198.51.100.1 route-map DEMO out
```



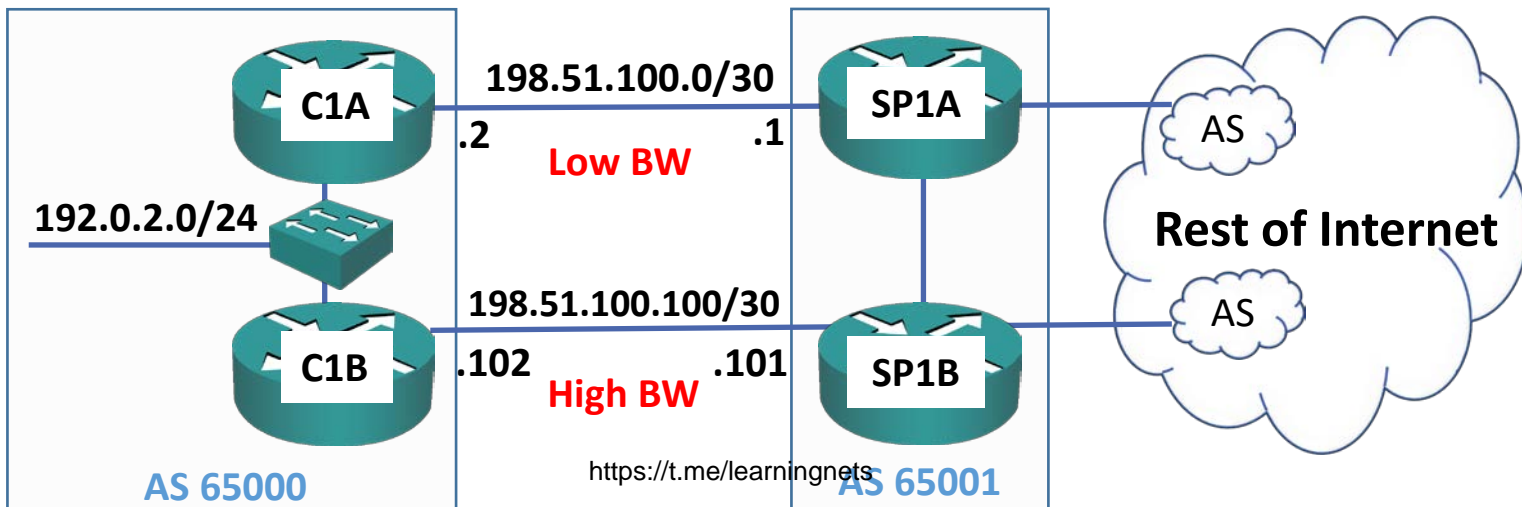
```
C1B(config)#route-map DEMO permit 10
C1B(config-route-map)#set metric 100
C1B(config)#router bgp 65000
C1B(config-router)#neighbor 198.51.100.101 route-map DEMO out
```



Verification – show ip bgp (Before Policy)

```
SP1A#sh ip bgp
```

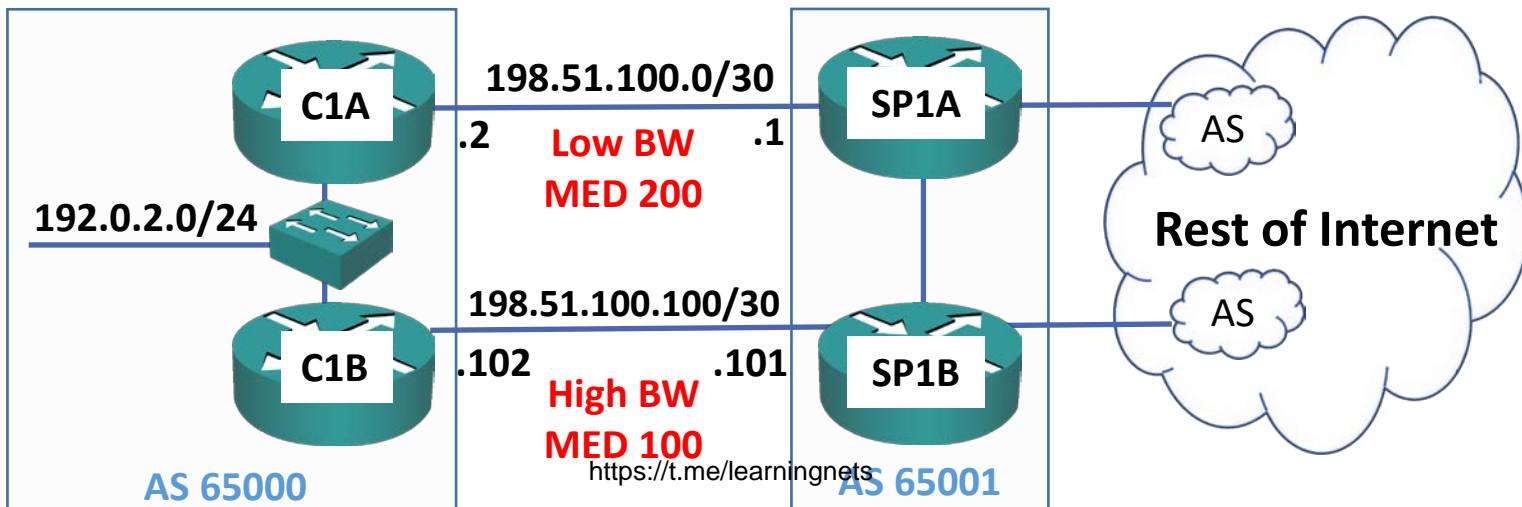
```
      Network                Next Hop                Metric LocPrf Weight Path
*i 192.0.2.0                 192.168.0.2 (SP1B)      0      100      0 65000 i
*>                          198.51.100.2 (C1A)      0      100      0 65000 i
```



Verification – show ip bgp (After Policy)

```
SP1A#sh ip bgp
```

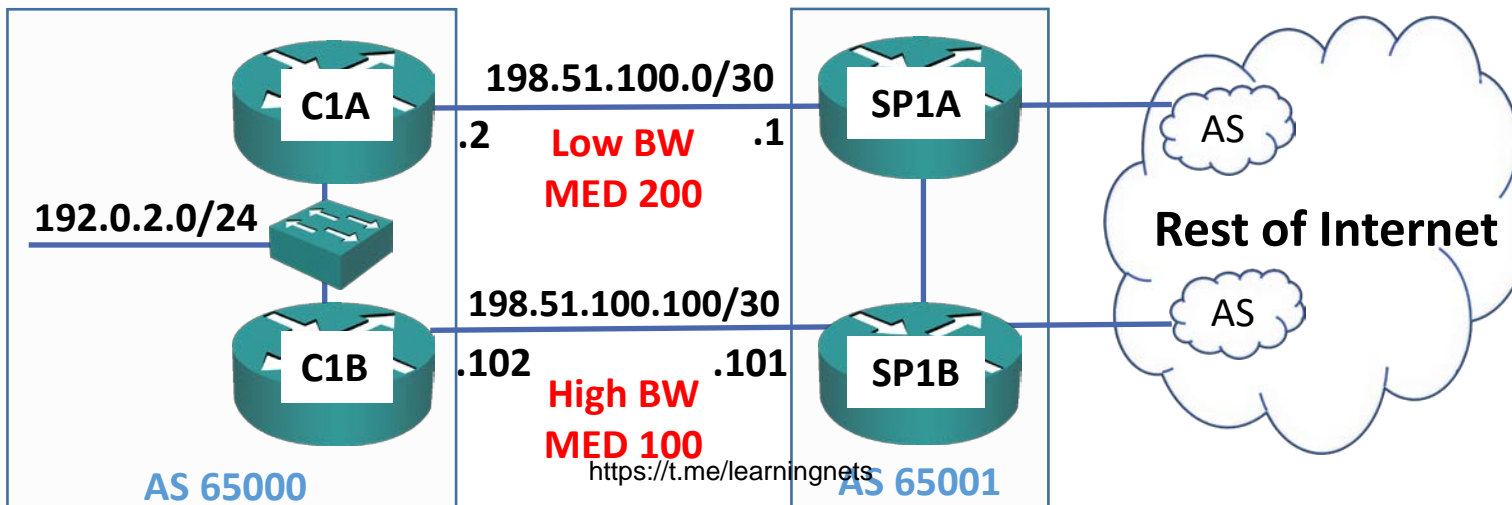
```
      Network                Next Hop                Metric LocPrf Weight Path
* >i 192.0.2.0                192.168.0.2 (SP1B)      100   100     0 65000 i
*                                198.51.100.2 (C1A)      200   100     0 65000 i
```



Verification – show ip bgp (Before Policy)

```
SP1B#sh ip bgp
```

```
      Network                Next Hop                Metric LocPrf Weight Path
 *i  192.0.2.0                192.168.0.1 (SP1A)      0      100      0 65000 i
 *>  198.51.100.101 (C1B)      0      100      0 65000 i
```

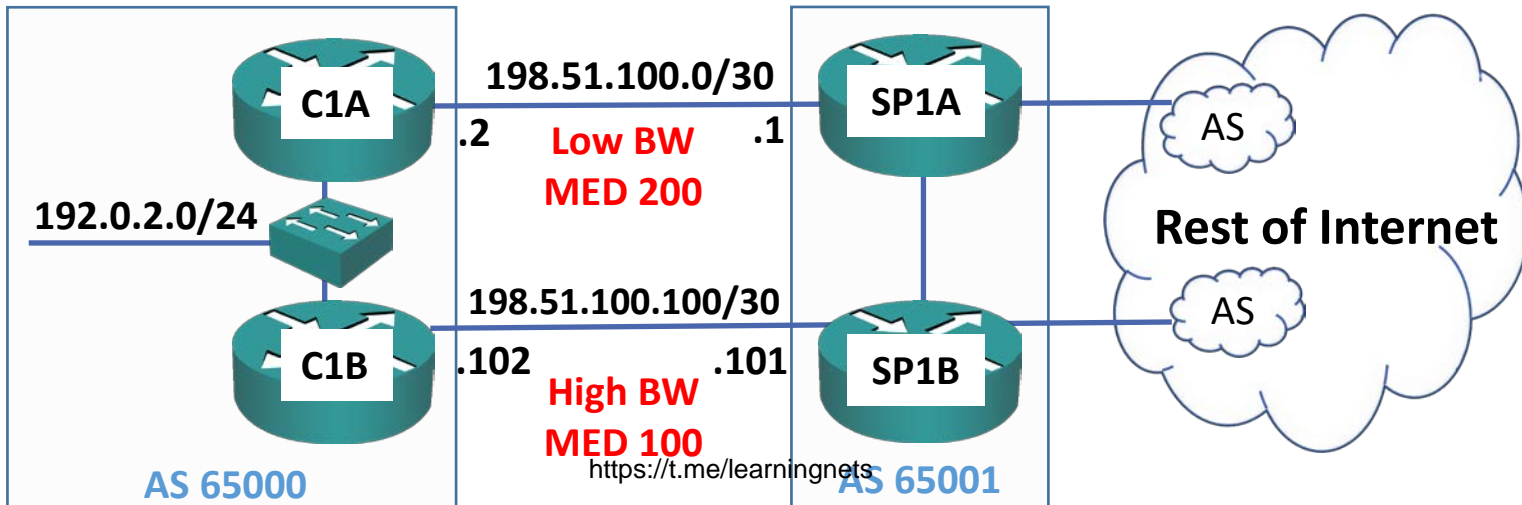


Verification – show ip bgp (After Policy)

```
SP1B#sh ip bgp
```

```
Network                Next Hop                Metric LocPrf Weight Path
* > 192.0.2.0           198.51.100.102 (C1B)  100                0 65000 i
```

- SP1A's preferred path to the 192.0.2.0/24 network is via SP1B
- Loop prevention stops the path from being advertised from SP1A back to SP1B



Verification – show ip route



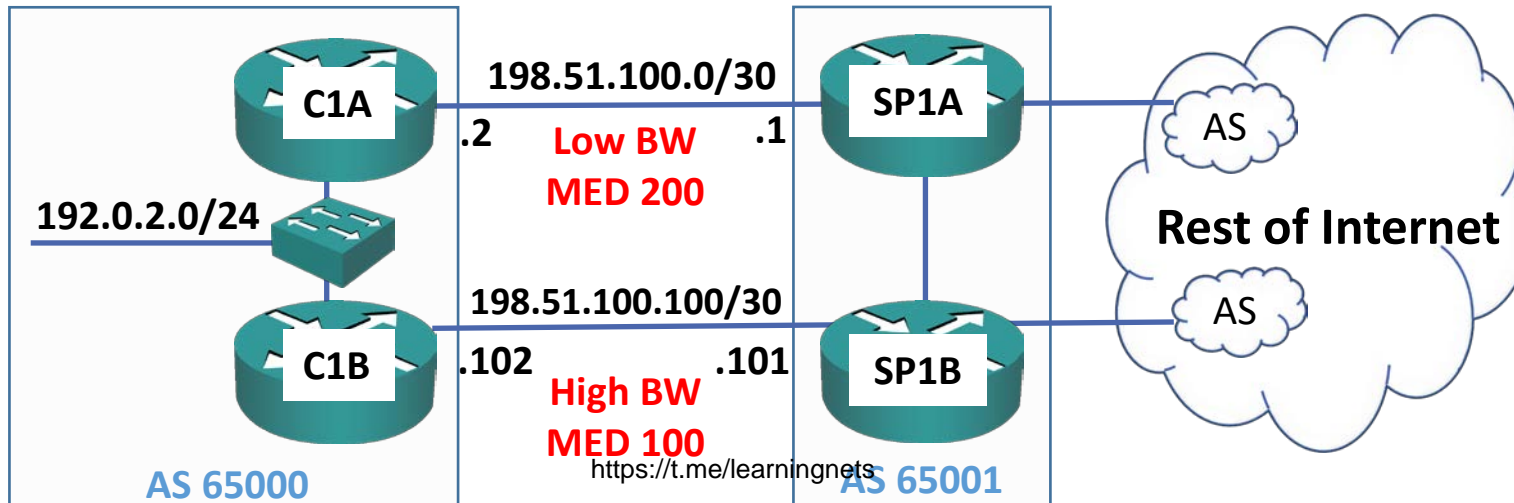
```
SP1A#show ip route bgp
```

```
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, p - overrides from PfR
```




```
Gateway of last resort is not set
```

```
B 192.0.2.0 [200/100] via 192.168.0.2, 00:25:07
```



MED – Load Balancing – C1A

```
C1A(config)#ip prefix-list TO_192.0.2.0/24 permit 192.0.2.0/24
C1A(config)#route-map DEMO permit 10
C1A(config-route-map)#match ip address prefix-list TO_192.0.2.0/24
C1A(config-route-map)#set metric 100
C1A(config)#route-map DEMO permit 20
C1A(config-route-map)#set metric 200
C1A(config)#router bgp 65000
C1A(config-router)#neighbor 198.51.100.1 route-map DEMO out
```




Traffic from AS 65001 to the 192.0.2.0/24 network goes via the top link to C1A

Traffic to the 203.0.113.0/24 network goes via the bottom link to C1B

MED – Load Balancing – C1B

```
C1B(config)#ip prefix-list TO_203.0.113.0/24 permit 203.0.113.0/24
C1B(config)#route-map DEMO permit 10
C1B(config-route-map)#match ip address prefix-list TO_203.0.113.0/24
C1B(config-route-map)#set metric 100
C1B(config)#route-map DEMO permit 20
C1B(config-route-map)#set metric 200
C1B(config)#router bgp 65000
C1B(config-router)#neighbor 198.51.100.101 route-map DEMO out
```

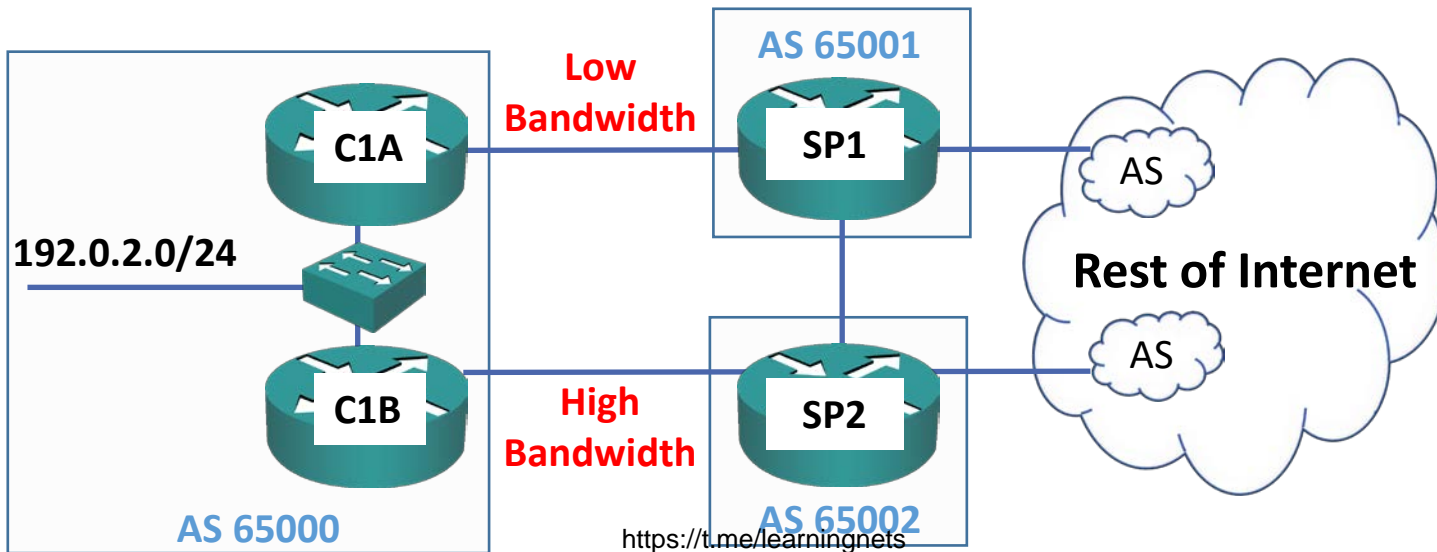


Traffic from AS 65001 to the 192.0.2.0/24 network goes via the top link to C1A

Traffic to the 203.0.113.0/24 network goes via the bottom link to C1B

Influencing Inbound Path Selection

- MED can only be used to influence incoming traffic **from a single AS**
- AS 65001 can reach 192.0.2.0/24 via the low bandwidth path with an AS Path of 65000, or via the high bandwidth path with an AS Path of 65002_65000
- The low bandwidth path will be preferred because of the shorter AS Path
- MED values have no effect because AS Path length is considered before MED in BGP path selection



Consider only routes with no AS loops and a valid next hop.

Use Longest Prefix Match.

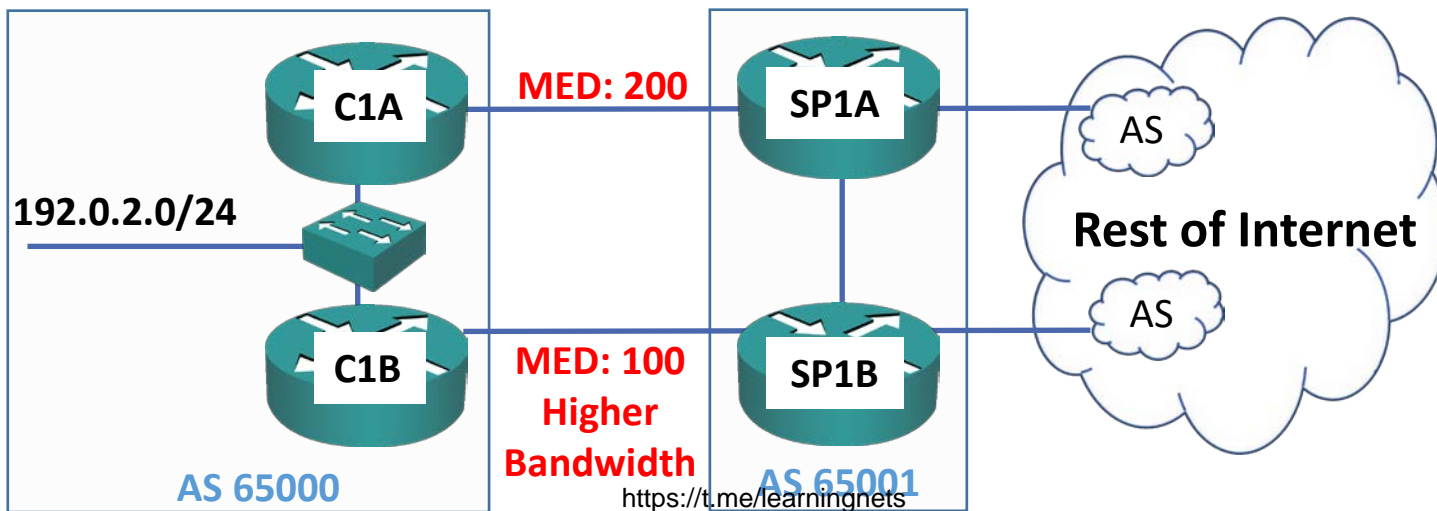
Where multiple routes are available to identical network and prefix:

- Prefer highest weight (local to router).
- Prefer highest local preference (global within AS).
- Prefer route originated by the local router ('network' command or redistribution).
- Prefer shortest AS path.
- Prefer lowest origin code: IGP ('network') < EGP (legacy) < incomplete (redistributed).
- Prefer lowest MED (exchanged between autonomous systems).
- Prefer EBGP path over IBGP path.
- Prefer the path through the closest IGP neighbor.
- Prefer oldest route for EBGP paths.
- Prefer the path with the lowest neighbor BGP router ID.
- Prefer the path with the lowest neighbor IP address.

MED Multi Exit Discriminator



- MED can be used to influence incoming traffic from a single AS because the AS Path length is the same over all links



AS Path Prepending

- AS Path Prepending can be used to influence incoming path selection from multiple neighbor ASs

