

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

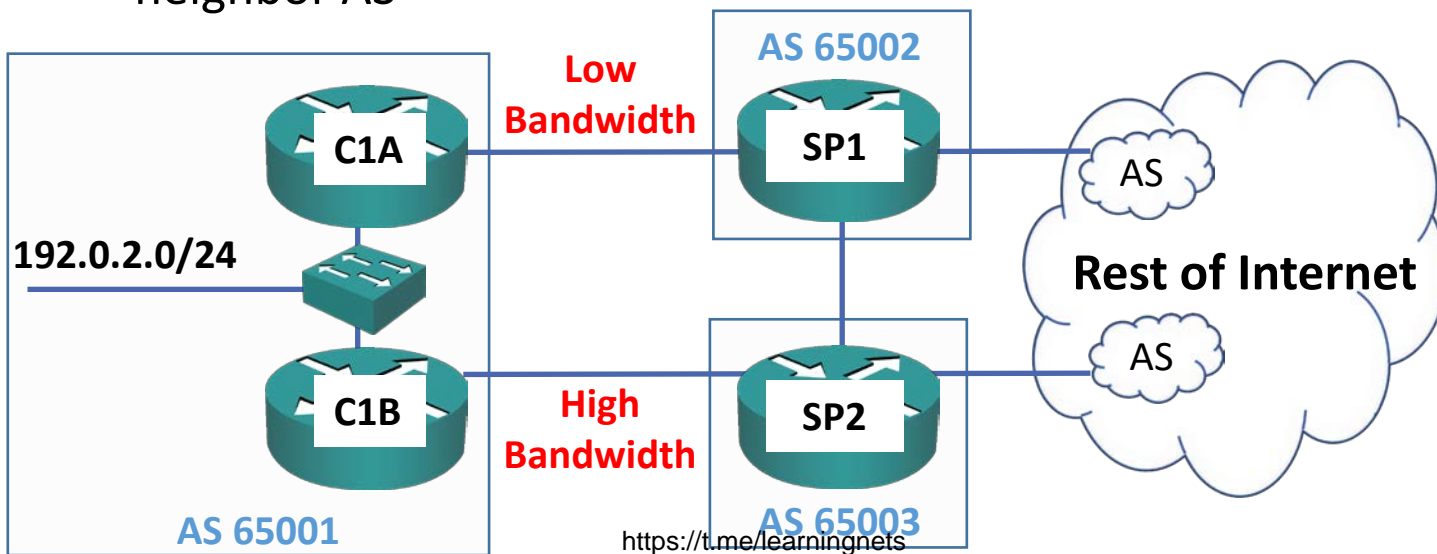
Use Longest Prefix Match.

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.

Influencing Inbound Path Selection

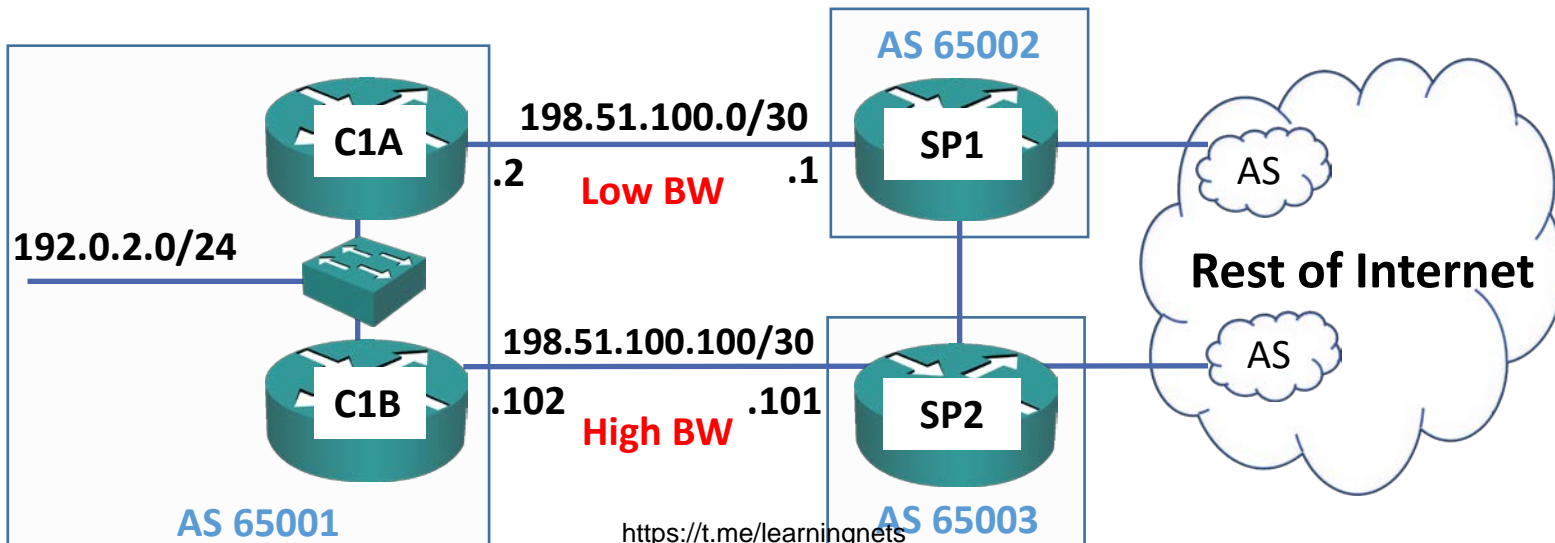
- In the example below, AS 65001 would prefer all neighbor ASs to send traffic to its 192.0.2.0/24 network over the high bandwidth link
- By default, AS 65002 will send traffic over its direct link to C1A (rather than via SP2) because it has a shorter AS Path
- Internet Service Providers use fixed policies and configuration templates
- They typically avoid custom configurations per customer (such as configuring Local Preference) because it would add considerable complexity to their administration and troubleshooting
- AS Path Prepending provides a way to influence incoming traffic with no changes required in the neighbor AS



Verification – show ip bgp (Before Policy)

```
SP1#sh ip bgp
```

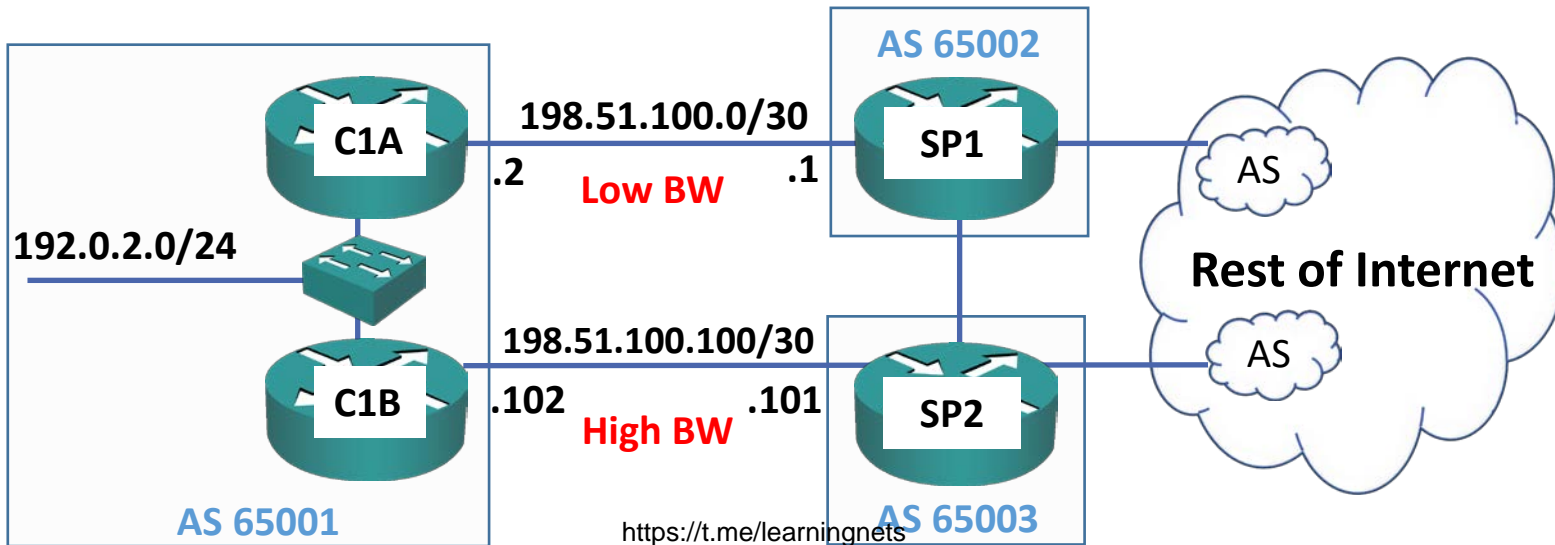
```
      Network                Next Hop                Metric LocPrf  Weight  Path
*      192.0.2.0              203.0.113.2 (SP2)          0         0 65003 65001 i
*>      198.51.100.2 (C1A)    198.51.100.2 (C1A)  0         0 65001 i
```



Verification – show ip bgp (Before Policy)


```
SP2#sh ip bgp
```

```
      Network                Next Hop                Metric LocPrf Weight Path
* >   192.0.2.0              198.51.100.102 (C1B)      0         0 65001 i
*     203.0.113.1           203.0.113.1   (SP1)      0         0 65002 65001 i
```



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 65002 can reach 192.0.2.0/24 via either path with the same AS Path of 65001



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

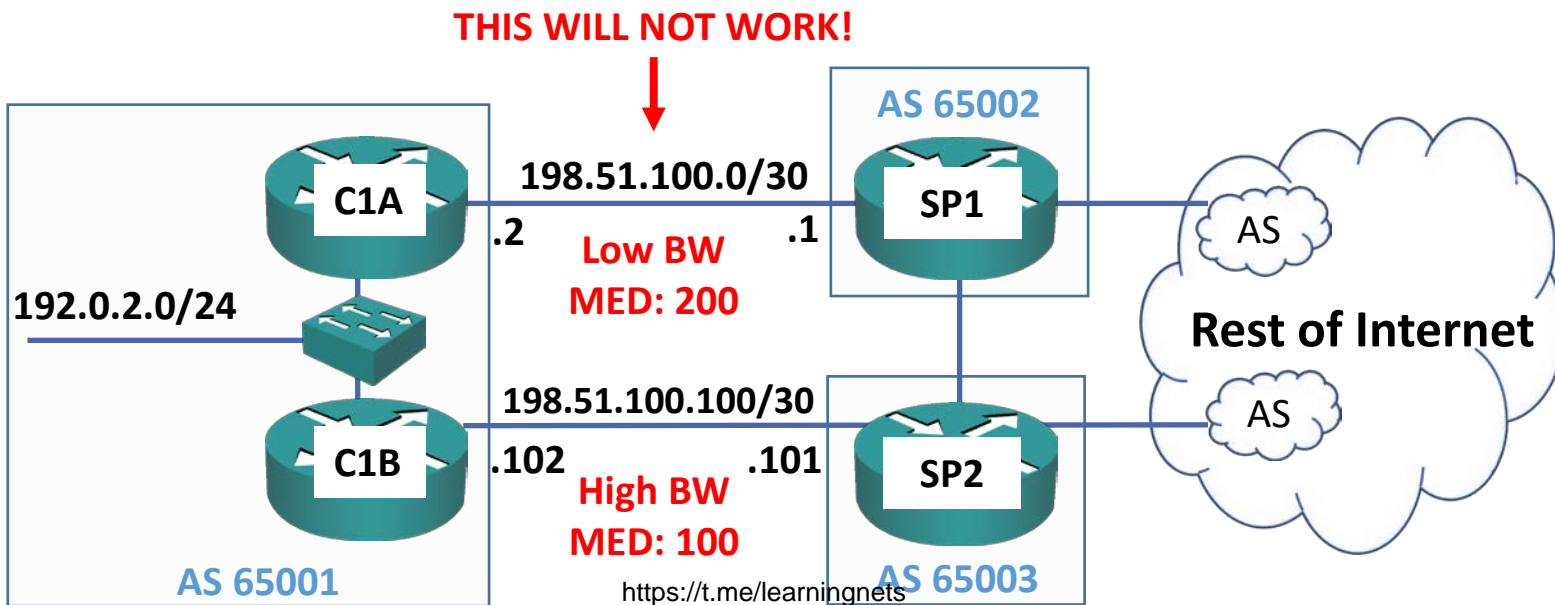
Use Longest Prefix Match.

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.

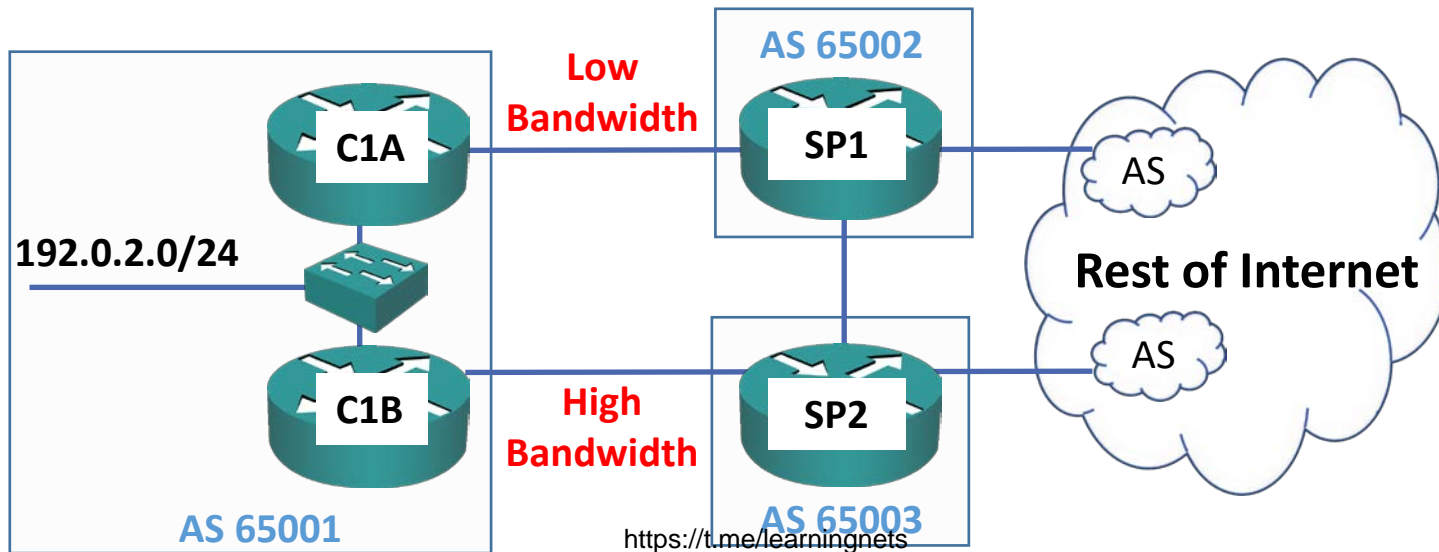
Influencing Inbound Path Selection

- MED can only be used to influence incoming traffic **from a single AS**
- AS 65002 can reach 192.0.2.0/24 via the low bandwidth path with an AS Path of 65001, or via the high bandwidth path with an AS Path of 65003_65001
- 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



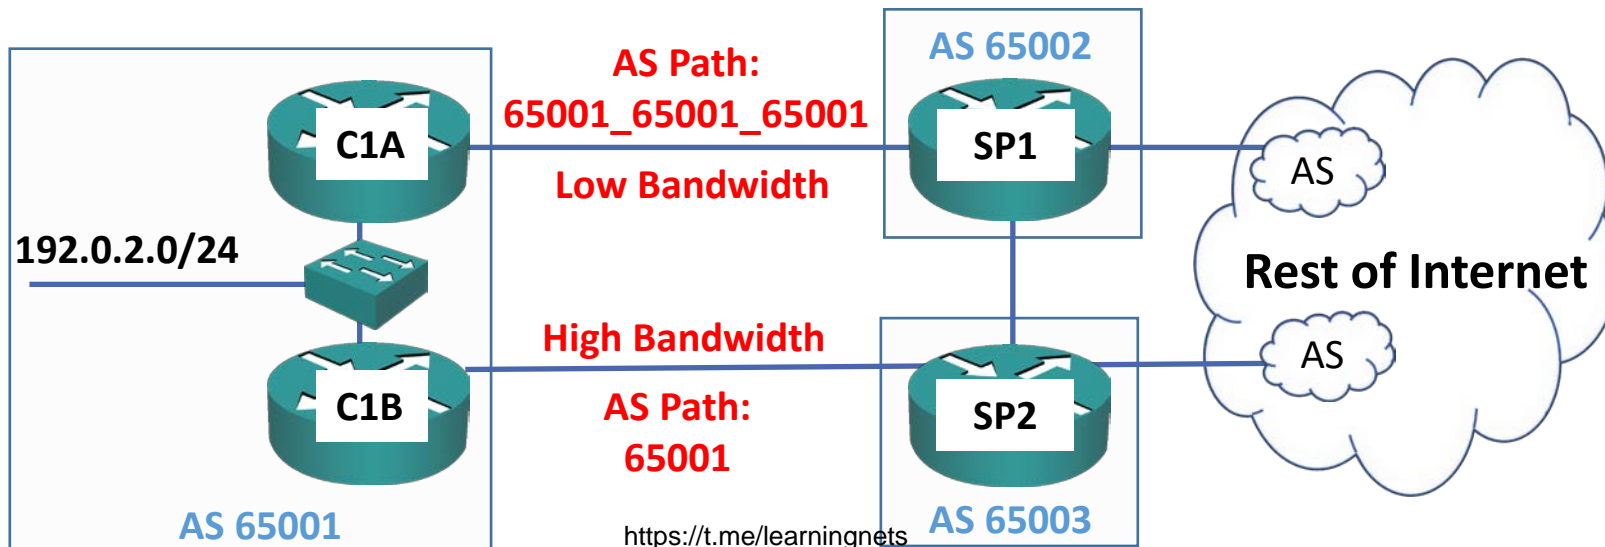
AS Path Prepending

- AS Path Prepending can be used to influence incoming path selection from multiple neighbor ASs
- It works by extending an AS Path with multiple duplicate copies of the AS number to make that path longer and less preferred



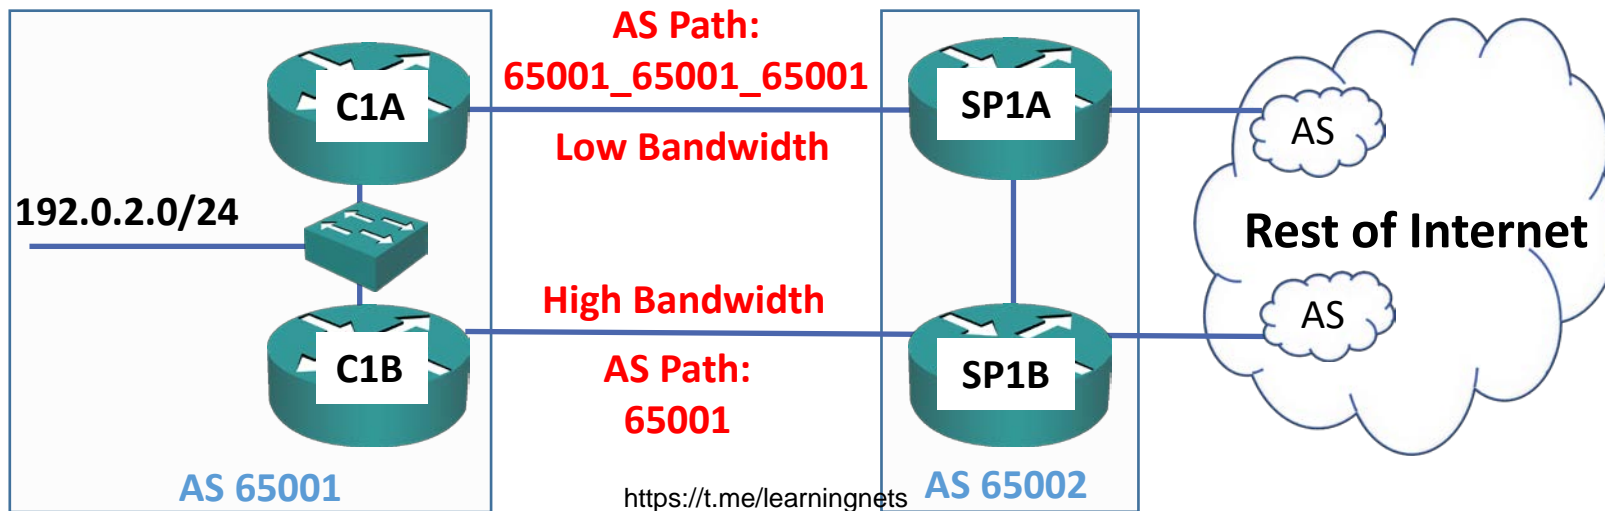
How AS Path Prepending Works

- The Update from C1B to SP2 for the 192.0.2.0/24 network is unchanged, with an AS Path of 65001
- C1A uses AS Path prepending to advertise the network to SP1 with an AS Path of 65001_65001_65001 (rather than the default 65001)
- SP1 also receives the standard advertisement from SP2 for the 192.0.2.0/24 network with an AS Path of 65003_65001
- SP1 sees the shorter AS Path via SP2 and prefers it



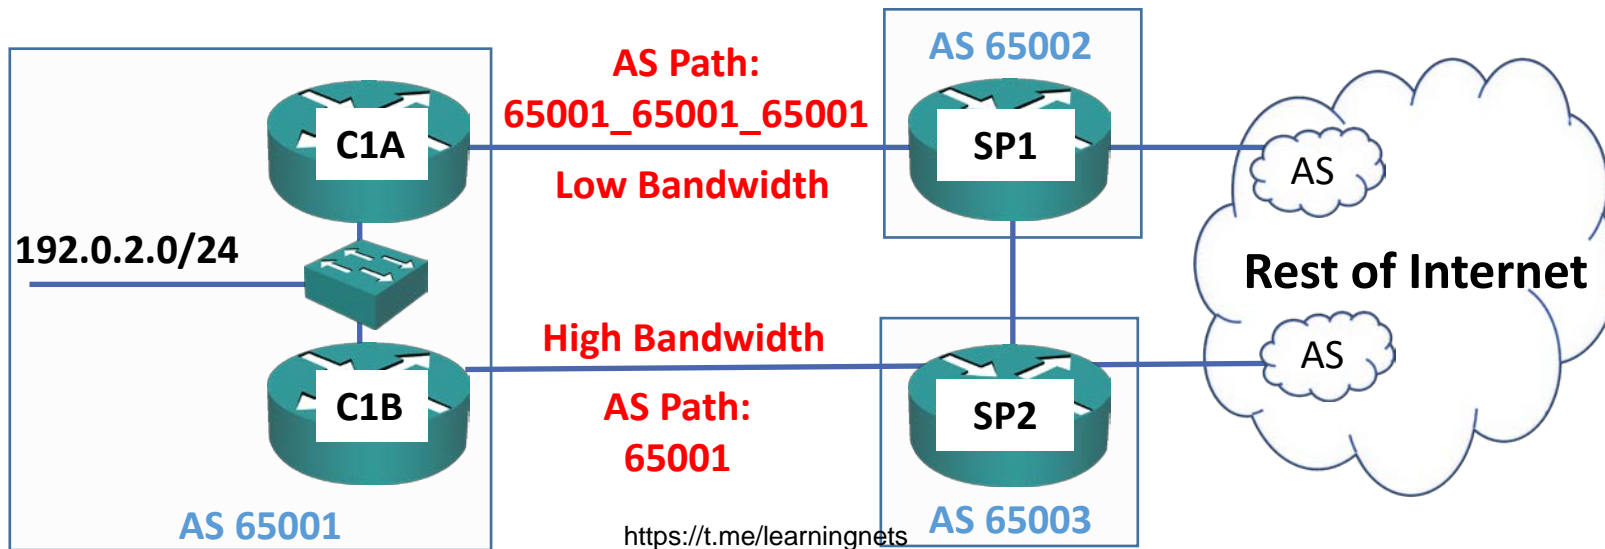
AS Path Prepending

- AS Path Prepending can influence incoming path selection across multiple paths from the same or from different ASs



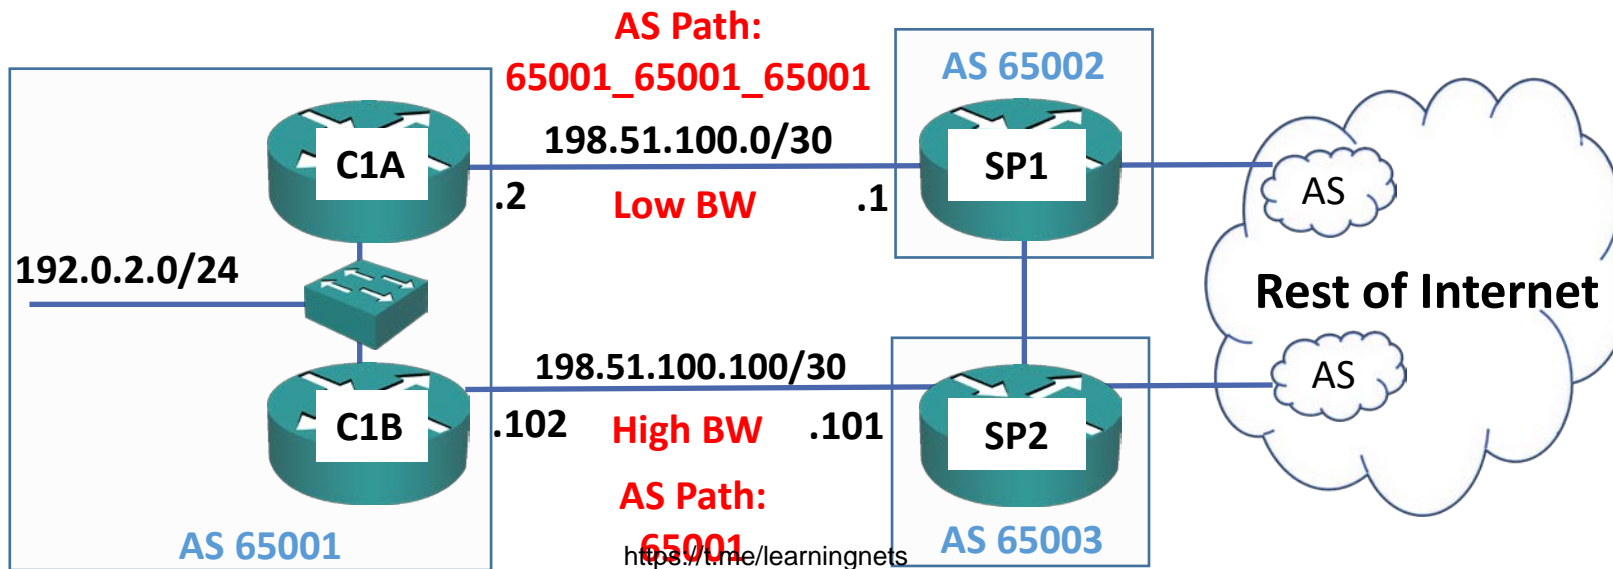
AS Path Prepending Propagation

- The extended AS path is passed to direct neighbor ASs, and they also propagate it to further remote ASs
- In the example below, AS 65002 will advertise the 192.0.2.0/24 network to further downstream ASs with the longer AS Path



AS Path – Primary/Backup Links


```
C1A(config)#route-map DEMO permit 10
C1A(config-route-map)#set as-path prepend 65001 65001
C1A(config)#router bgp 65001
C1A(config-router)#neighbor 198.51.100.1 route-map DEMO out
```



Verification – show ip bgp (Before Policy)

```
SP1#sh ip bgp
```

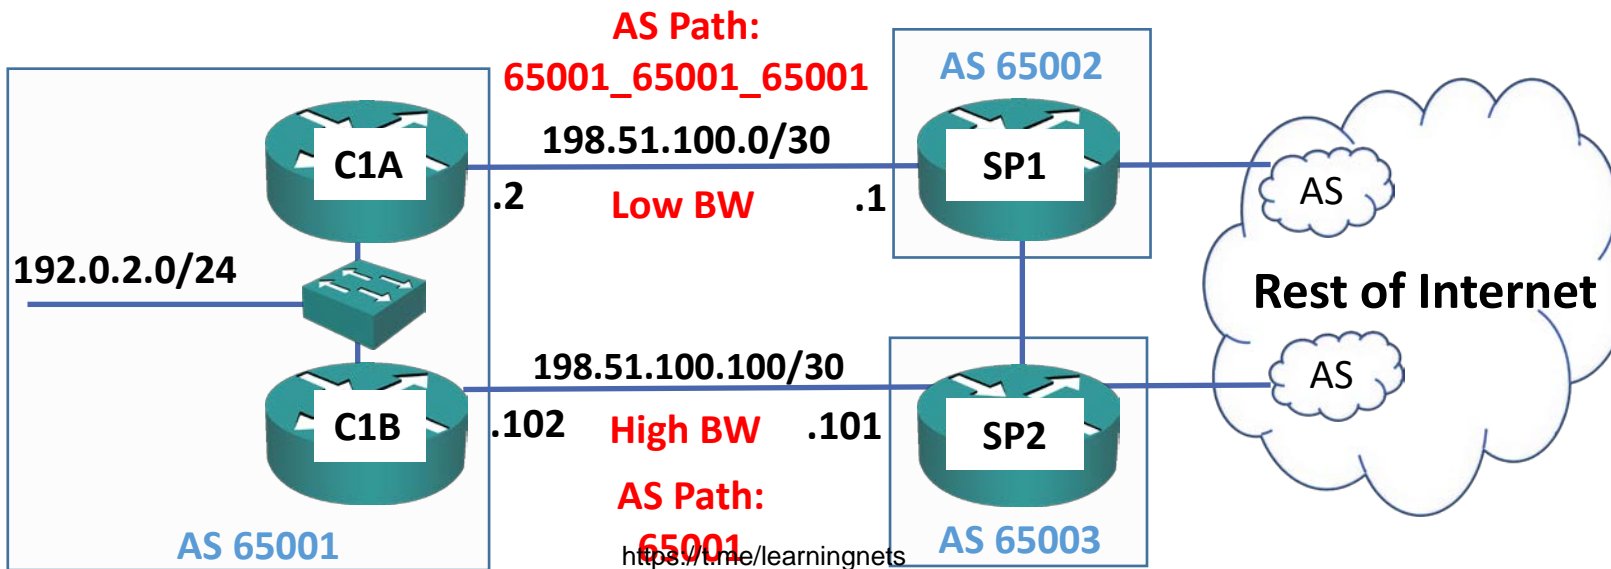
```
      Network                Next Hop                Metric LocPrf Weight Path
*      192.0.2.0              203.0.113.2 (SP2)          0         0 65003 65001 i
*>      198.51.100.2 (C1A)    198.51.100.2 (C1A)      0         0 65001 i
```



Verification – show ip bgp (After Policy)

```
SP1#sh ip bgp
```


Network	Next Hop	Metric	LocPrf	Weight	Path
*> 192.0.2.0	203.0.113.2 (SP2)	0	0	65003	65001 i
*	198.51.100.2 (C1A)	0	0	65001	65001 65001 i



Verification – show ip bgp (Before Policy)

```
SP2#sh ip bgp
```

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	192.0.2.0	198.51.100.102 (C1B)	0		0	65001 i
*		203.0.113.1 (SP1)	0		0	65002 65001 i

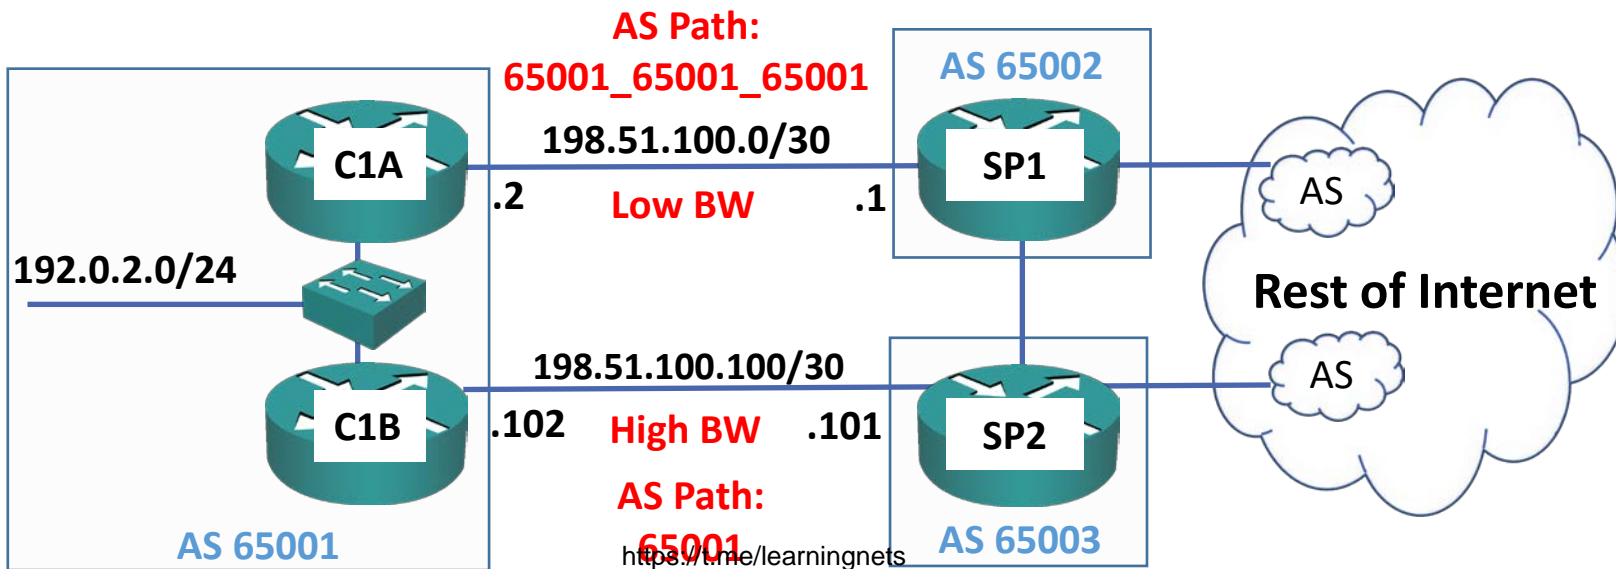


Verification – show ip bgp (After Policy)

```
SP2#sh ip bgp
```

```
Network          Next Hop          Metric LocPrf Weight Path
*> 192.0.2.0      198.51.100.102 (C1B)    0           0 65001 i
```

- The route via SP1 is no longer in the BGP table due to AS Path loop prevention
- (SP2 sees its own AS in SP1's preferred path)



Verification – show ip route



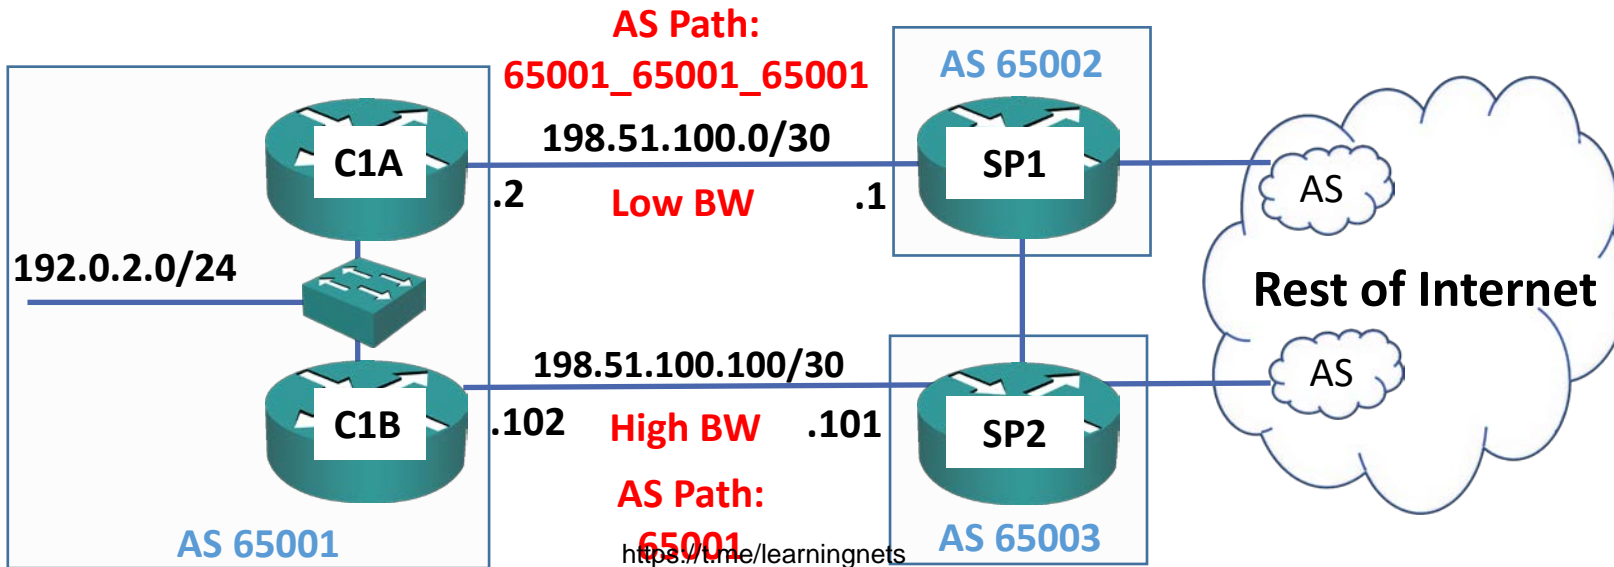
```
SP1#show ip route bgp
```

```
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 [20/0] via 203.0.113.2, 00:25:01
```



Verification – show ip route



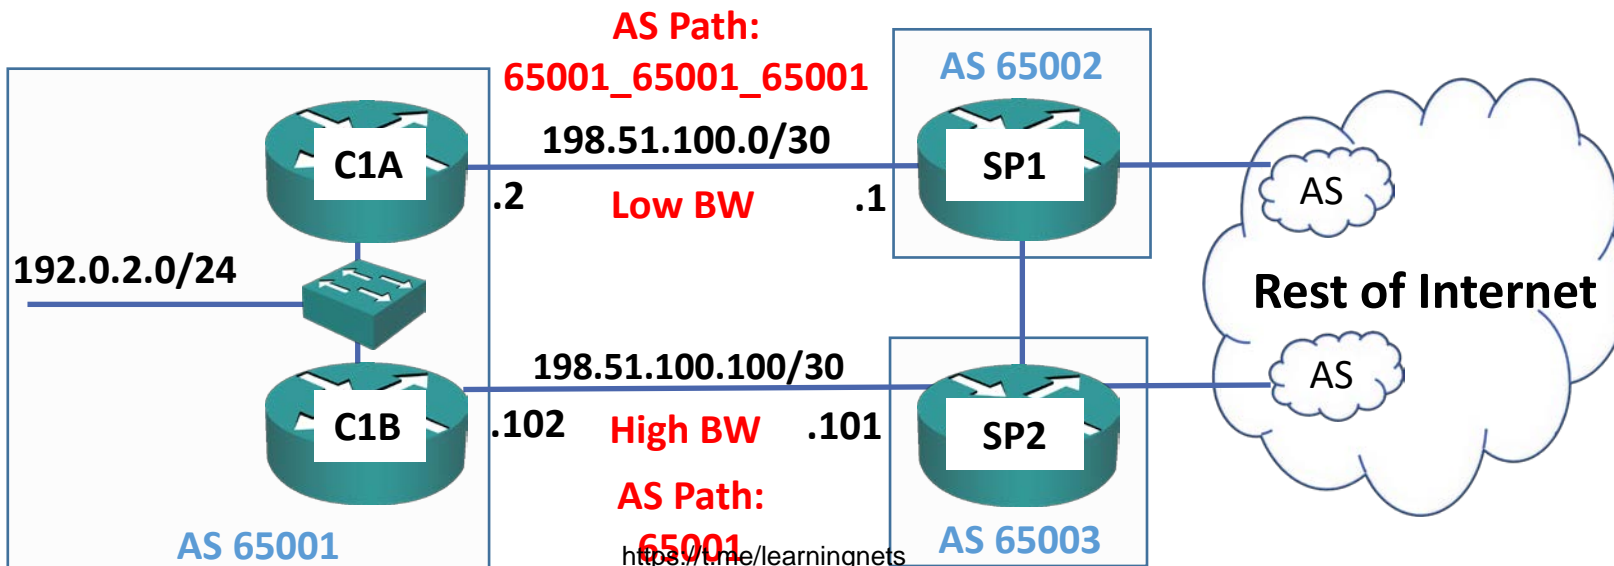
```
SP2#show ip route bgp
```

```
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
```



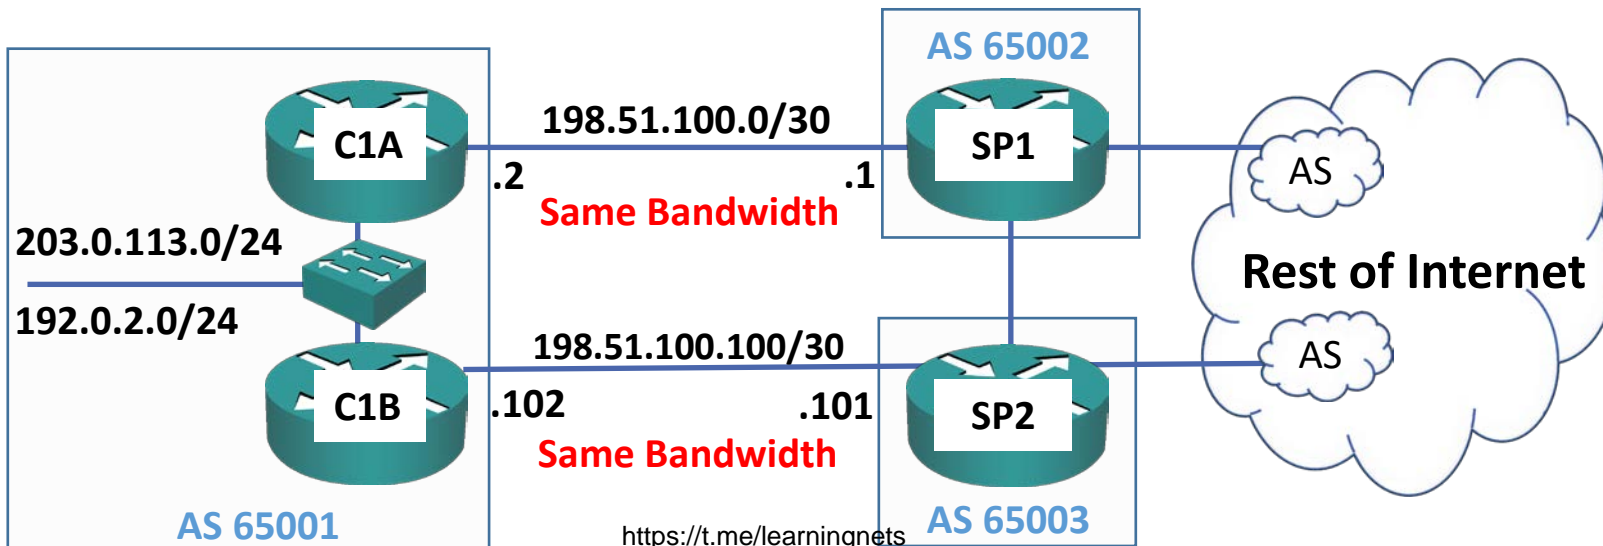
```
Gateway of last resort is not set
```

```
B 192.0.2.0 [20/0] via 198.51.100.102, 00:25:07
```



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 as-path prepend 65001 65001
C1A(config)#route-map DEMO permit 20
C1A(config)#router bgp 65001
C1A(config-router)#neighbor 198.51.100.1 route-map DEMO out
```

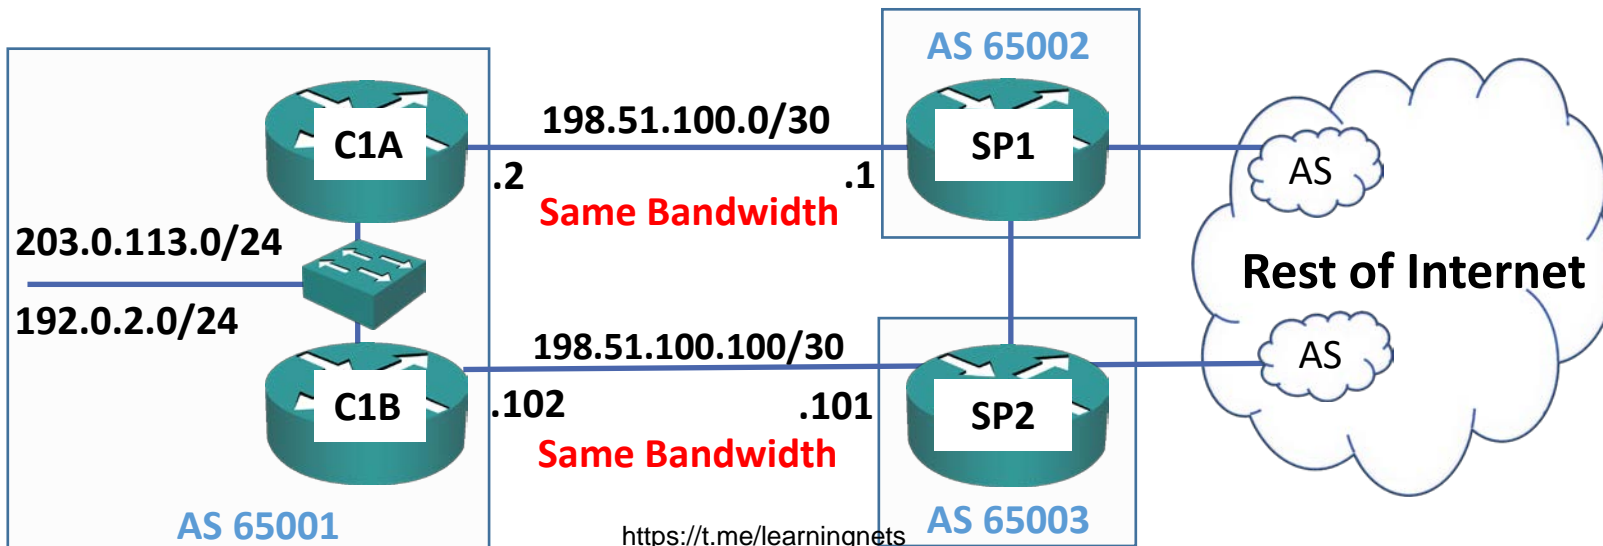


Traffic from neighbors to the 203.0.113.0/24 network goes via the top link to C1A

Traffic to the 192.0.2.0/24 network goes via the bottom link to C1B

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 as-path prepend 65001 65001
C1B(config)#route-map DEMO permit 20
C1B(config)#router bgp 65001
C1B(config-router)#neighbor 198.51.100.101 route-map DEMO out
```



Traffic from neighbors to the 203.0.113.0/24 network goes via the top link to C1A

Traffic to the 192.0.2.0/24 network goes via the bottom link to C1B