

BGP Route Propagation



- BGP routers will form an adjacency and be BGP peers after successfully configuring the neighbor relationship
- When an adjacency is formed between IGP routers, they automatically exchange routes for the networks on their IGP enabled interfaces and routes learned from other IGP neighbors
- BGP is not enabled on interfaces and does not automatically advertise routes, you need to additionally manually configure this after the neighbor configuration

The Network Command



- The `network` command works differently for BGP than it does for IGPs
- In IGPs, it means 'enable the IGP on all interfaces with an IP address which falls within this range, and advertise the network prefixes configured on those interfaces'

The Network Command



- In BGP, it has the more intuitive meaning of basically ‘advertise this network’. (Actually ‘insert it into the local BGP table to be advertised to BGP neighbours’ to be more precise.)
- There is no ‘enabling BGP on an interface’. BGP uses targeted TCP sessions to form peers, not multicast Hellos

BGP Route Propagation



- IGP only advertise the router's best paths to each network to neighbors - the paths in its routing table, not all paths known in the IGP database
- Similarly, BGP only advertises the router's best path to each network to its peers (not all paths known in the BGP table)

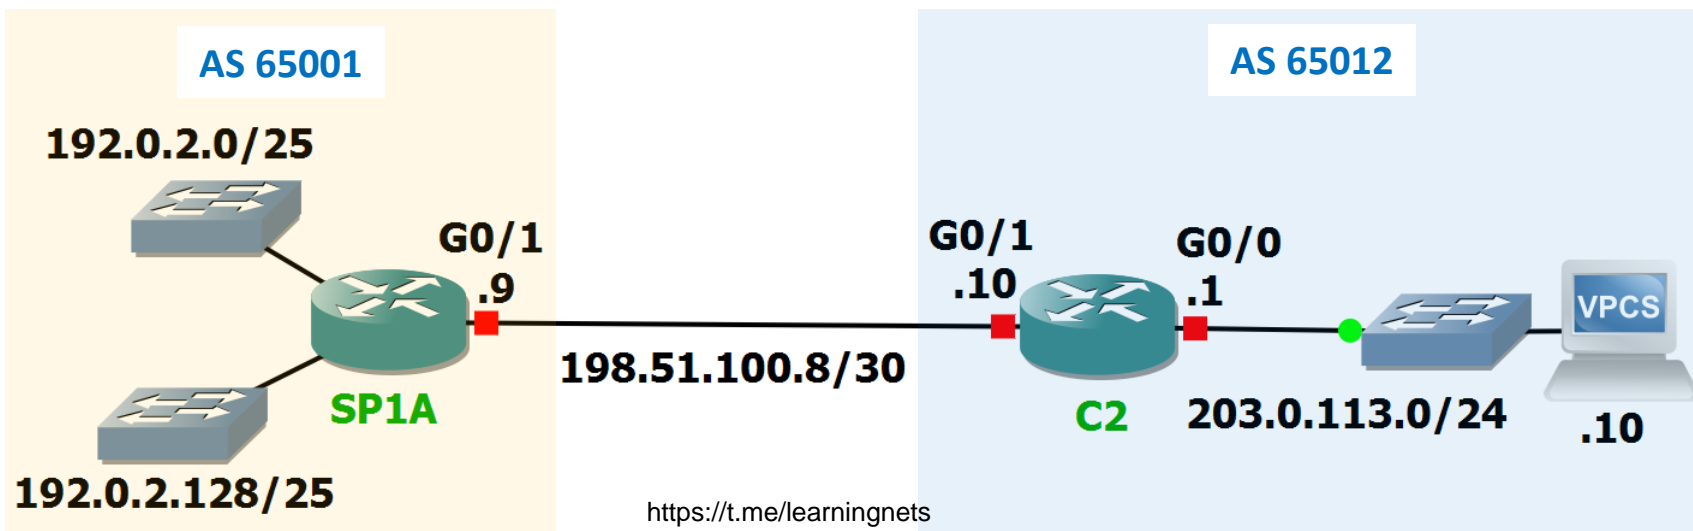
Neighbor Configuration

```
C2(config)#router bgp 65012
```

```
C2(config-router)#neighbor 198.51.100.9 remote-as 65001
```

```
SP1A#show ip bgp
```

```
SP1A#
```



The Network Command

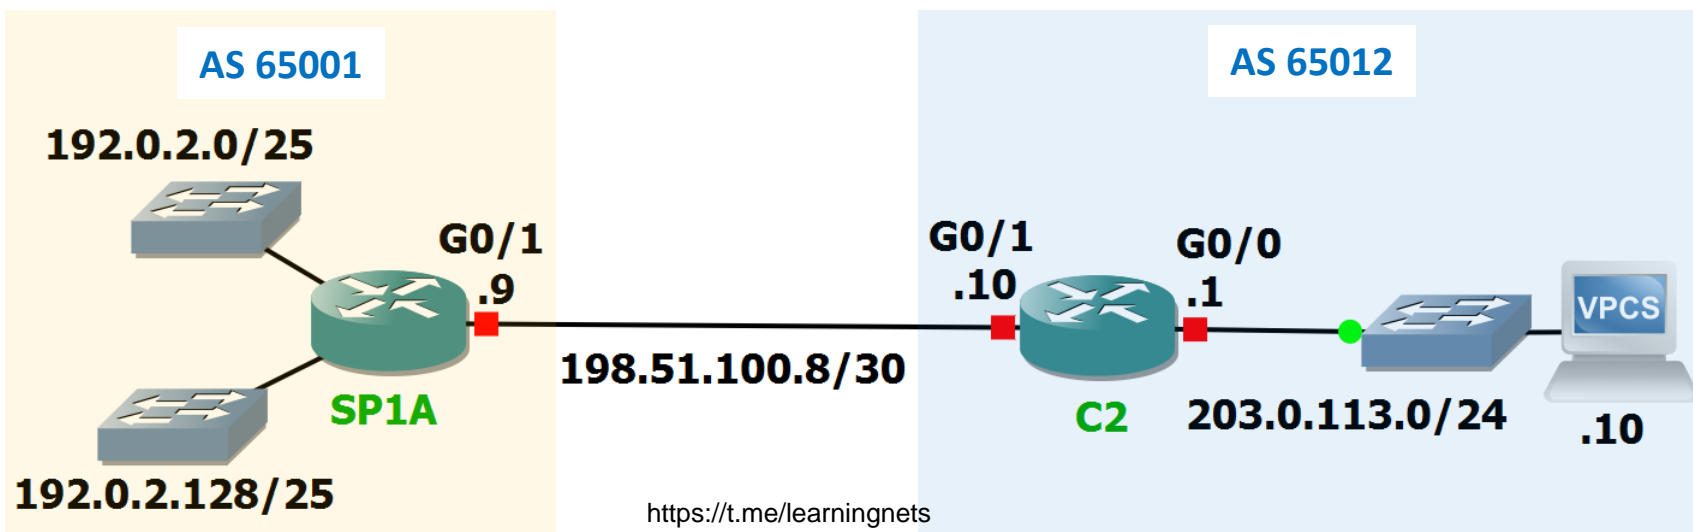
```
C2(config)#router bgp 65012
```

```
C2(config-router)#network 203.0.113.0 mask 255.255.255.0
```

```
SP1A(config)#router bgp 65001
```

```
SP1A(config-router)#network 192.0.2.0 mask 255.255.255.128
```

```
SP1A(config-router)#network 192.0.2.128 mask 255.255.255.128
```



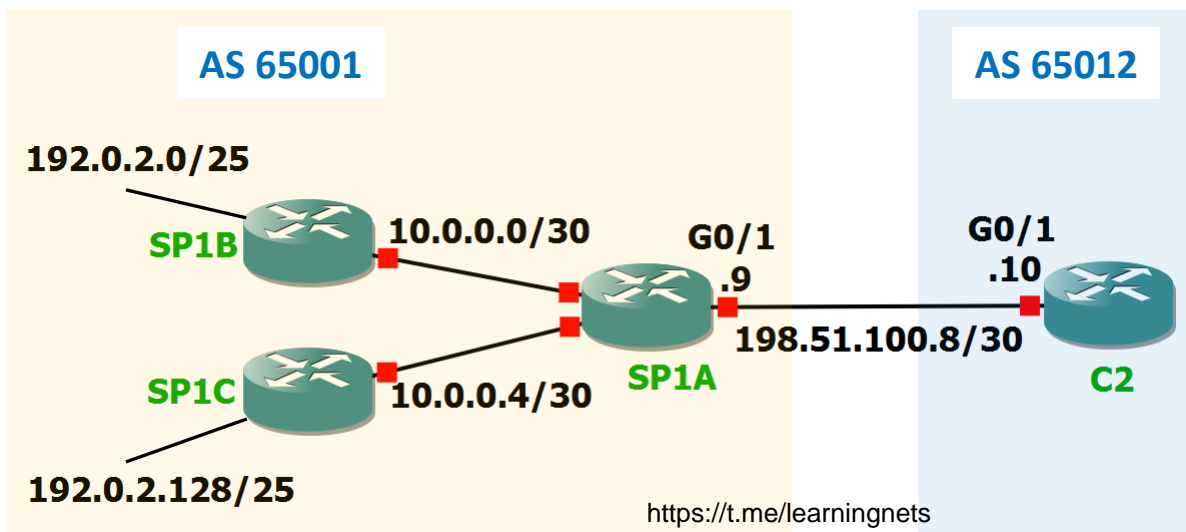
The Network Command

- The BGP session and route advertisements are separate components.
- The BGP router does not need to be directly connected to the advertised network, it just needs to have it in its routing table.

```
SP1A(config)#router bgp 65001
```

```
SP1A(config-router)#network 192.0.2.0 mask 255.255.255.128
```

```
SP1A(config-router)#network 192.0.2.128 mask 255.255.255.128
```



The Mask Keyword



- The 'mask' keyword is optional on the network command
- If you do not enter a mask, it uses the default mask for that address class
- Default masks are omitted from most 'show' command output

```
C2(config)#router bgp 65012
```

```
C2(config-router)#network 203.0.113.0 mask 255.255.255.0
```

OR

```
C2(config-router)#network 203.0.113.0
```

The Network Command



- Routes in a network statement will only be advertised in BGP if there is an exact match in the routing table
- For example, if you have 203.0.113.0/30, 203.0.113.4/30, 203.0.113.8/30 etc. you **cannot** advertise them all with 203.0.113.0/24

```
C2(config)#router bgp 65012
```

```
C2(config-router)#network 203.0.113.0 mask 255.255.255.0
```

The Network Command



`network 203.0.113.0 0.0.0.255` in an IGP would advertise all 203.0.113.x/30 networks.

`network 203.0.113.0 mask 255.255.255.0` **will not advertise anything in BGP.**

Add a network statement for each network you want to advertise.

```
R1#sh ip route
```

```
! Truncated
```

```
    203.0.113.0/24 is variably subnetted, 7 subnets, 3 masks
```

```
C      203.0.113.0/30 is directly connected, FastEthernet0/0
```

```
L      203.0.113.1/32 is directly connected, FastEthernet0/0
```

```
C      203.0.113.4/30 is directly connected, FastEthernet1/0
```

```
L      203.0.113.5/32 is directly connected, FastEthernet1/0
```

```
C      203.0.113.8/30 is directly connected, FastEthernet2/0
```

```
L      203.0.113.9/32 is directly connected, FastEthernet2/0
```

The Network Command



```
C2(config)#router bgp 65012
```

```
C2(config-router)#network 203.0.113.0 mask 255.255.255.252
```

```
C2(config-router)#network 203.0.113.4 mask 255.255.255.252
```

```
C2(config-router)#network 203.0.113.8 mask 255.255.255.252
```

Null Route for BGP Summarisation



- You can summarise by adding a matching null route in the routing table

```
ip route 203.0.113.0 255.255.255.0 null0
router bgp 65001
  network 203.0.113.0 mask 255.255.255.0
```

Null Route for BGP Summarisation



`network 203.0.113.0 mask 255.255.255.0` will now work.

```
R1#sh ip route
```

```
! Truncated
```

```
    203.0.113.0/24 is variably subnetted, 7 subnets, 3 masks
```

```
S      203.0.113.0/24 is directly connected, Null0
```

```
C      203.0.113.0/30 is directly connected, FastEthernet0/0
```

```
L      203.0.113.1/32 is directly connected, FastEthernet0/0
```

```
C      203.0.113.4/30 is directly connected, FastEthernet1/0
```

```
L      203.0.113.5/32 is directly connected, FastEthernet1/0
```

```
C      203.0.113.8/30 is directly connected, FastEthernet2/0
```

```
L      203.0.113.9/32 is directly connected, FastEthernet2/0
```

The aggregate-address Command



- An alternative summarization configuration method is using 'aggregate-address'
- At least one component route must be in the BGP table
- A Null Route is automatically created

```
router bgp 65001
  network 203.0.113.0 mask 255.255.255.252
  aggregate-address 203.0.113.0 255.255.255.0 summary-only
```

Redistribution



- Routes can also be redistributed into BGP from static routes or an IGP

```
R1(config)#router bgp 65002
```

```
R1(config-router)#redistribute static
```

```
R1(config-router)#redistribute ospf 1
```

Verification – show run | section bgp


```
R1#show run | section bgp
router bgp 65002
  bgp log-neighbor-changes
  network 192.0.2.0
  network 203.0.113.64 mask 255.255.255.252
  network 203.0.113.72 mask 255.255.255.248
  neighbor 172.16.0.2 remote-as 65002
  neighbor 172.16.0.2 update-source Loopback0
  neighbor 203.0.113.65 remote-as 65001
```

Verification – show ip bgp

```
C2#show ip bgp
Status codes: * valid, > best
! truncated
```

- The Next Hop is shown as 0.0.0.0 when the BGP router is directly connected to the advertised network

	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	192.0.2.0/25	198.51.100.9	0		0	65001 i
*>	192.0.2.128/25	198.51.100.9	0		0	65001 i
*>	203.0.113.0	0.0.0.0	0		32768	i



Verification – show ip bgp



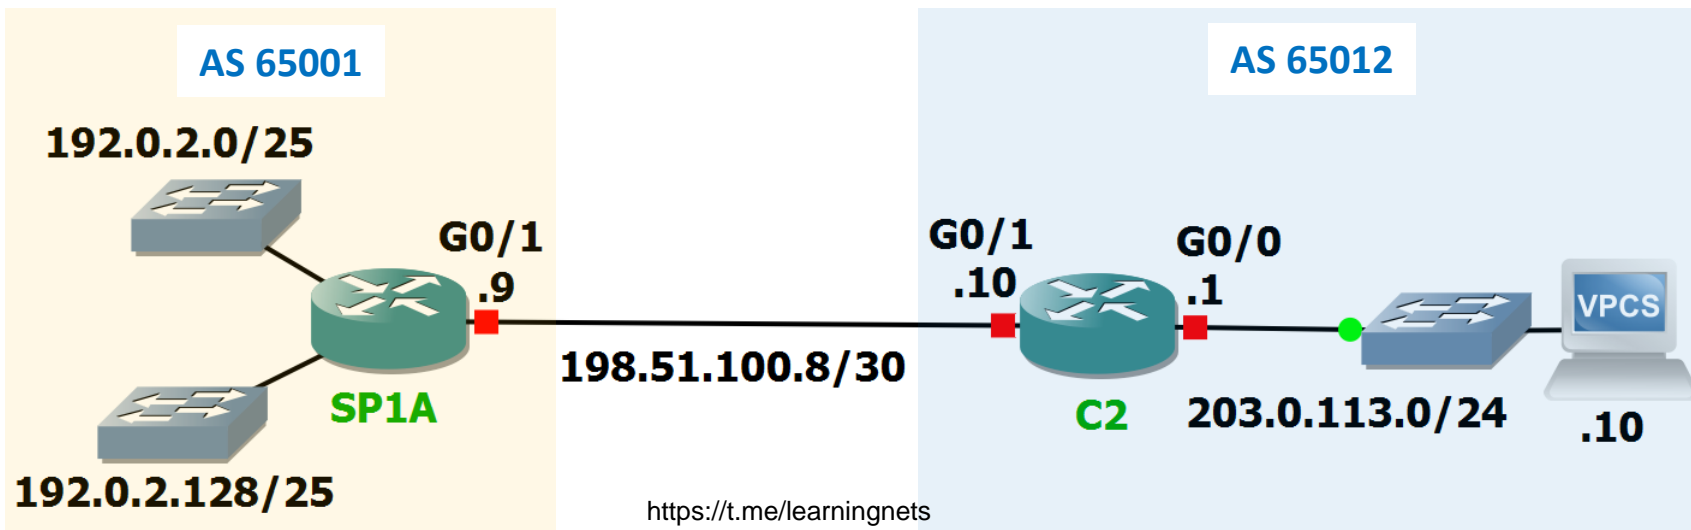
```
SP1A#show ip bgp
```

```
Status codes: * valid, > best
```

```
! Truncated
```



	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	192.0.2.0/25	0.0.0.0	0		32768	i
*>	192.0.2.128/25	0.0.0.0	0		32768	i
*>	203.0.113.0	198.51.100.10	0		0 65012	i



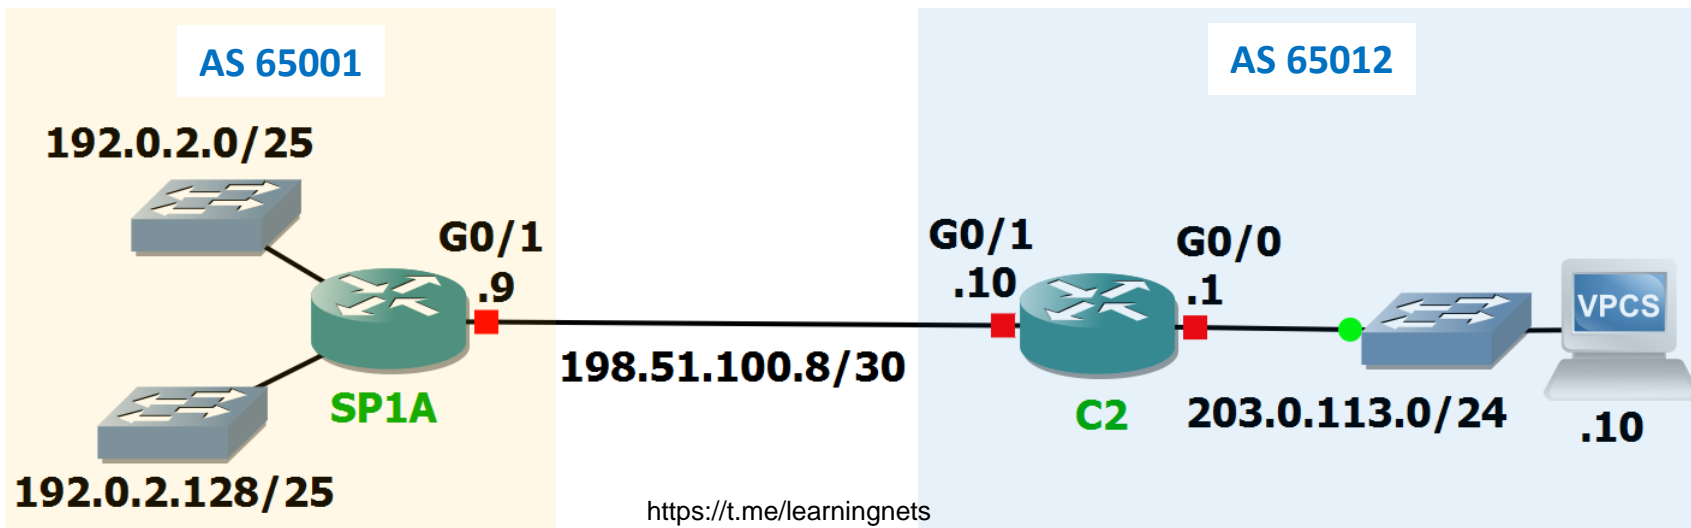
Verification – show ip bgp nei advertised-routes

```
C2#show ip bgp neighbors 198.51.100.9 advertised-routes
BGP table version is 6, local router ID is 192.168.0.3
! Truncated
```



	Network	Next Hop	Metric	LocPrf	Weight	Path
*>	203.0.113.0	0.0.0.0	0		32768	I

Total number of prefixes 1



BGP Operation

