eBGP Behaviour Summary

1) BGP has a security mechanism where it will only peer with an eBGP neighbor if it is directly connected – the TTL is set to 1 on eBGP packets by default
2) The advertising router changes the BGP next-hop address to its BGP source address.
3) The advertising router prepends its AS number to the AS Path.

   The receiving router rejects the update if its own AS number is in the AS Path. This is the AS Path loop prevention check.

1. eBGP TTL Security

- eBGP neighbors are typically directly connected to each other
- BGP has a security mechanism where it will only accept packets from an eBGP neighbor if it is directly connected
- The TTL is set to 1 on eBGP packets by default, and a router will not accept BGP packets from an eBGP neighbor if they have a TTL above 1
- Packets received with a TTL above 1 would indicate a bogus BGP neighbour multiple hops away that is attempting to spoof a legitimate neighbour

eBGP TTL Security

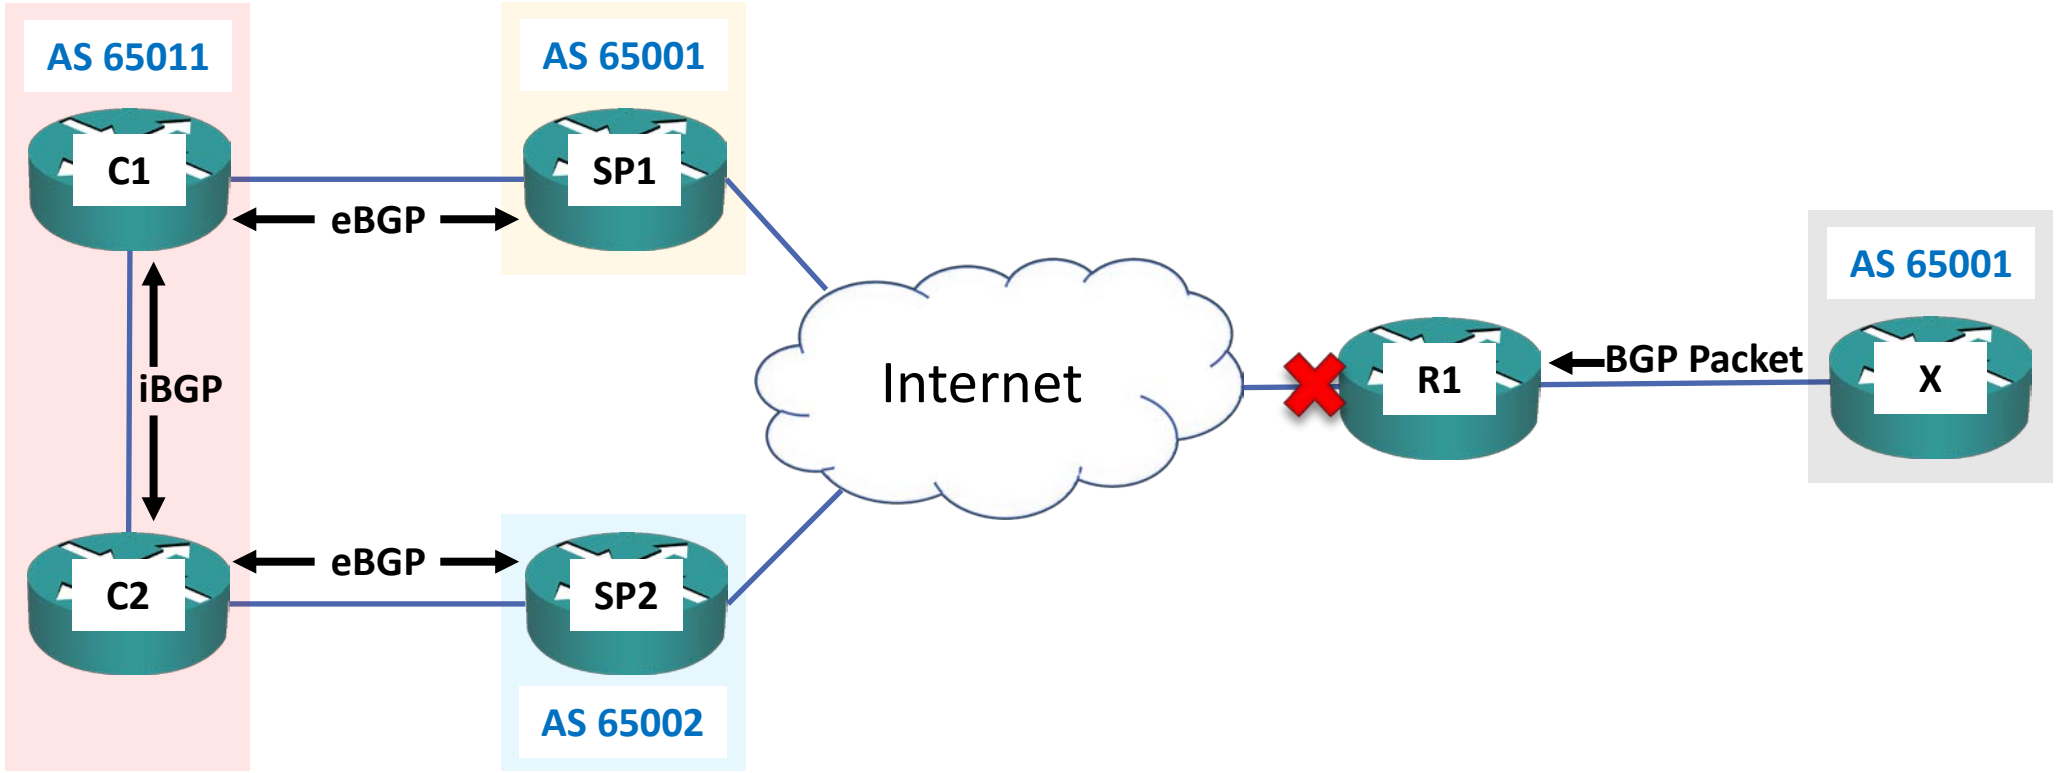

eBGP Neighbors

- If there are redundant links to the eBGP neighbor then a loopback address should be used
- The loopback address is an additional logical hop away and packets with the default TTL of 1 will be dropped

eBGP between Loopback Addresses

FLACKBOX
www.flackbox.com

```
SP1A(config)#ip route 172.16.0.2 255.255.255.255 192.0.2.2
SP1A(config)#ip route 172.16.0.2 255.255.255.255 198.51.100.2
SP1A(config)#router bgp 65001
SP1A(config-router)#neighbor 172.16.0.2 remote-as 65002
SP1A(config-router)#neighbor 172.16.0.2 update-source loopback 0
SP1A(config-router)#neighbor 172.16.0.2 ebgp-multihop 2
```

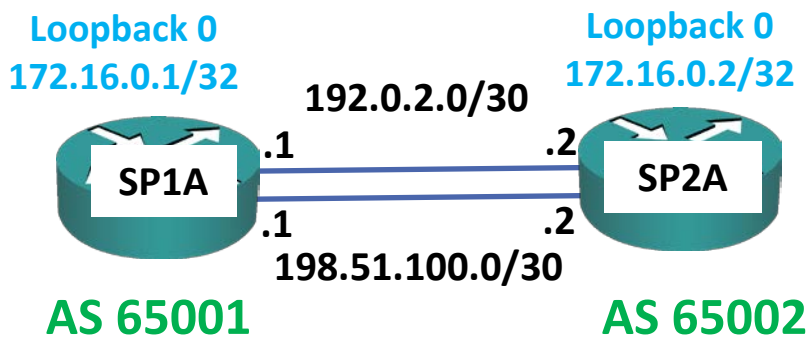

eBGP TTL Security Circumvention

- An attacker could adjust the TTL of sent BGP packets so that they appear to be originating from a directly-connected peer. (If they know the BGP router they want to attack is 3 hops away, they could send BGP packets with a starting TTL of 3)

BGP TTL Security

- The TTL Security feature inverts the direction in which the TTL is counted. Packets are sent with the maximum TTL of 255 (instead of the default EBGP multihop TTL of 1)
- Only BGP messages with an IP TTL greater than or equal to 255 (which is the maximum allowed TTL) minus the specified hop count are accepted.
- This ensures the neighbour is exactly where expected.
- Must be configured on both neighbours
- TTL security and EBGP multihop (the default) are mutually exclusive.

```
R1(config-router)#neighbor 172.16.0.2 ttl-security
hops 2
```

eBGP TTL

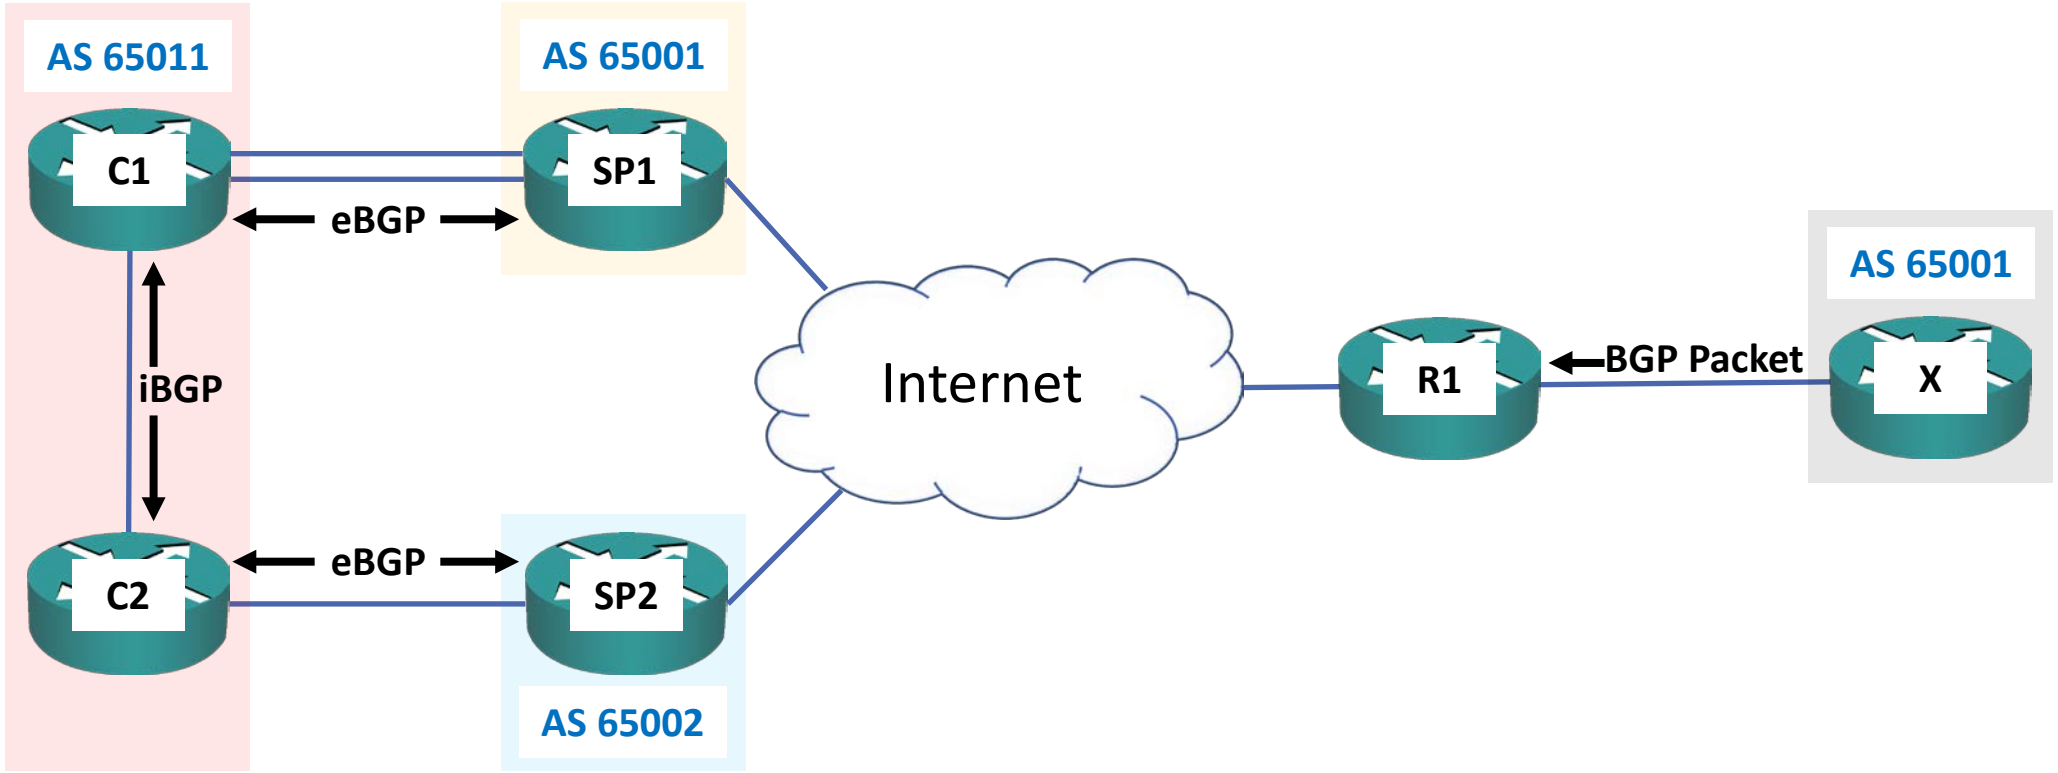

2. BGP Next Hop Address

- BGP is an AS-by-AS routing protocol, not a physical router-by-router routing protocol.
- The next hop does not mean the next physical router, it means the IP address to reach the next AS.
- When a route is received from an eBGP neighbour, the default next hop is the IP address of the neighbor that sent the update.
- The next hop advertised by eBGP is carried over into iBGP and advertised to iBGP neighbours by default.

Next Hop

- eBGP2 advertises 203.0.113.0/24 to iBGP2 with a next hop address of 198.51.100.2
- iBGP2 advertises 203.0.113.0/24 to iBGP1 with a next hop address of 198.51.100.2
- iBGP1 advertises 203.0.113.0/24 to eBGP1 with a next hop address of 192.0.2.1

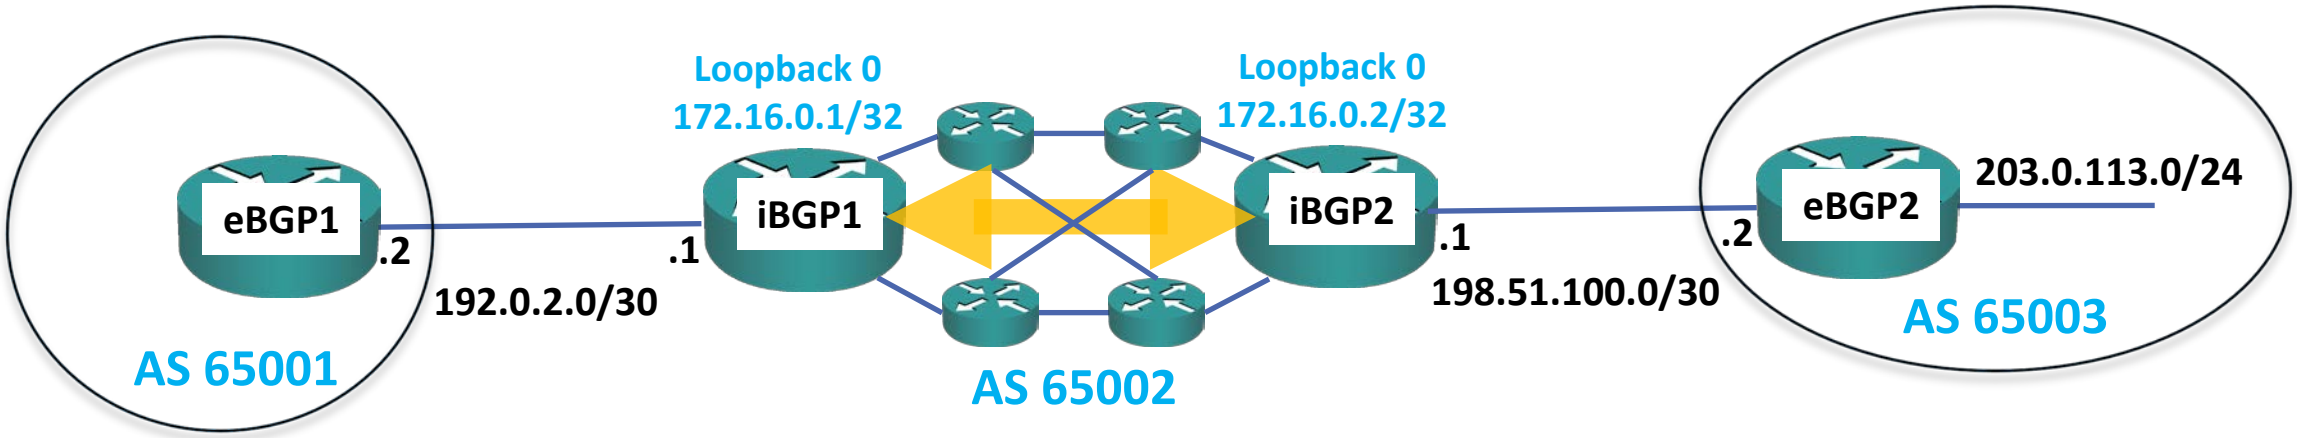

Next-Hop-Self

- iBGP2 advertises 203.0.113.0/24 to iBGP1 with a next hop address of 198.51.100.2
- iBGP1 will not have reachability if it does not have a route to 198.51.100.2
- Ensure 198.51.100.0/30 is included in the IGP (passive interface), or use ‘next-hop-self’

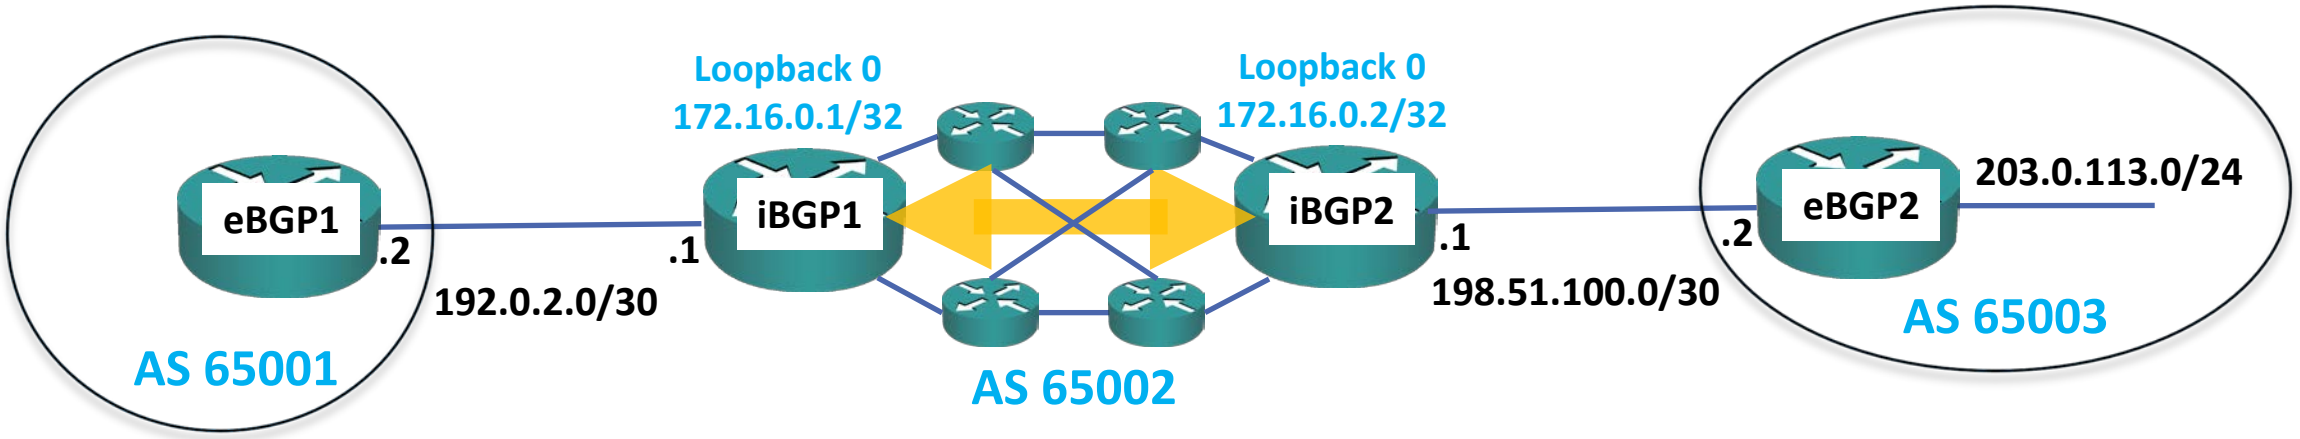

Next-Hop-Self

- Next-hop-self forces all updates for that neighbor to be advertised with ‘this router’ as the next hop.
- The next hop is set to the source IP address of the BGP packet.

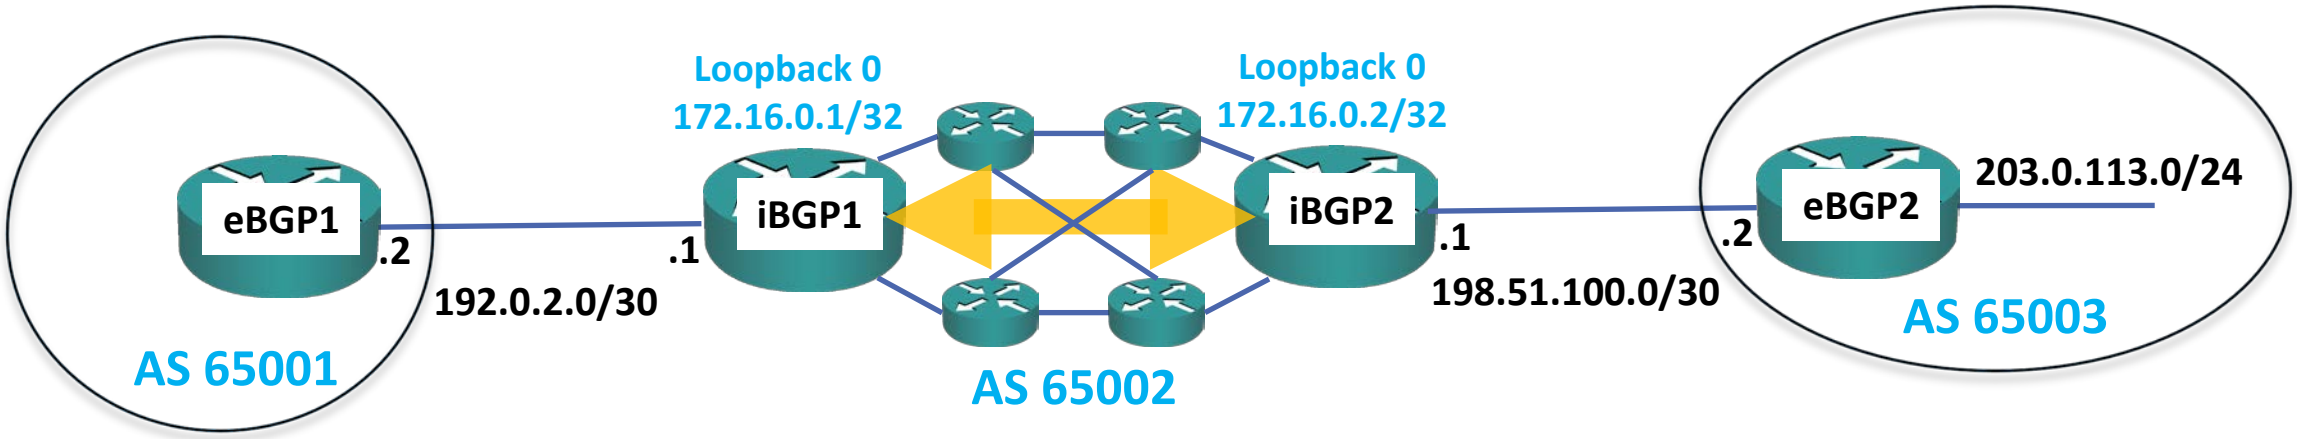

Next Hop

```
iBGP2(config-router)#neighbor 172.16.0.1 next-hop-self
```

- eBGP2 advertises 203.0.113.0/24 to iBGP2 with a next hop address of 198.51.100.2
- iBGP2 advertises 203.0.113.0/24 to iBGP1 with a next hop address of 172.16.0.2
- iBGP1 advertises 203.0.113.0/24 to eBGP1 with a next hop address of 192.0.2.1

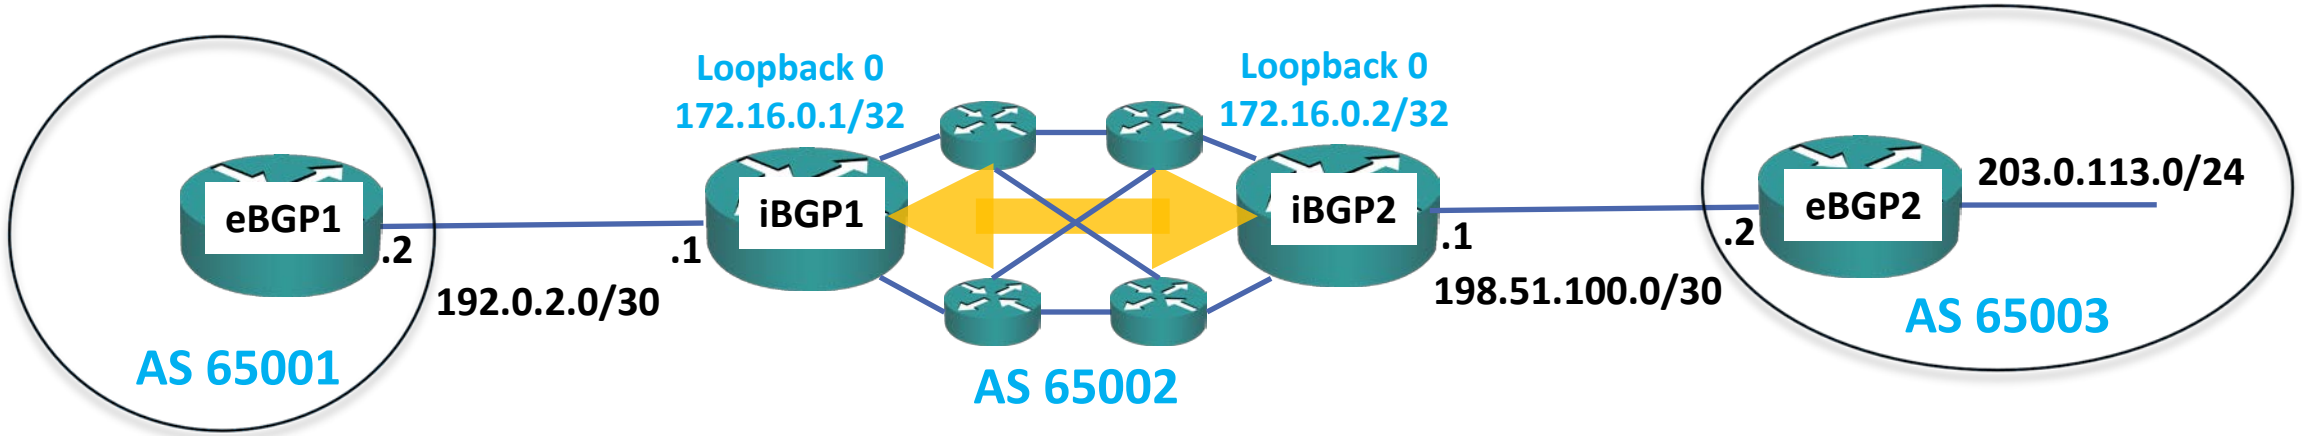

3. eBGP Loop Prevention

- The receiving router rejects the update if its own AS number is in the AS Path.

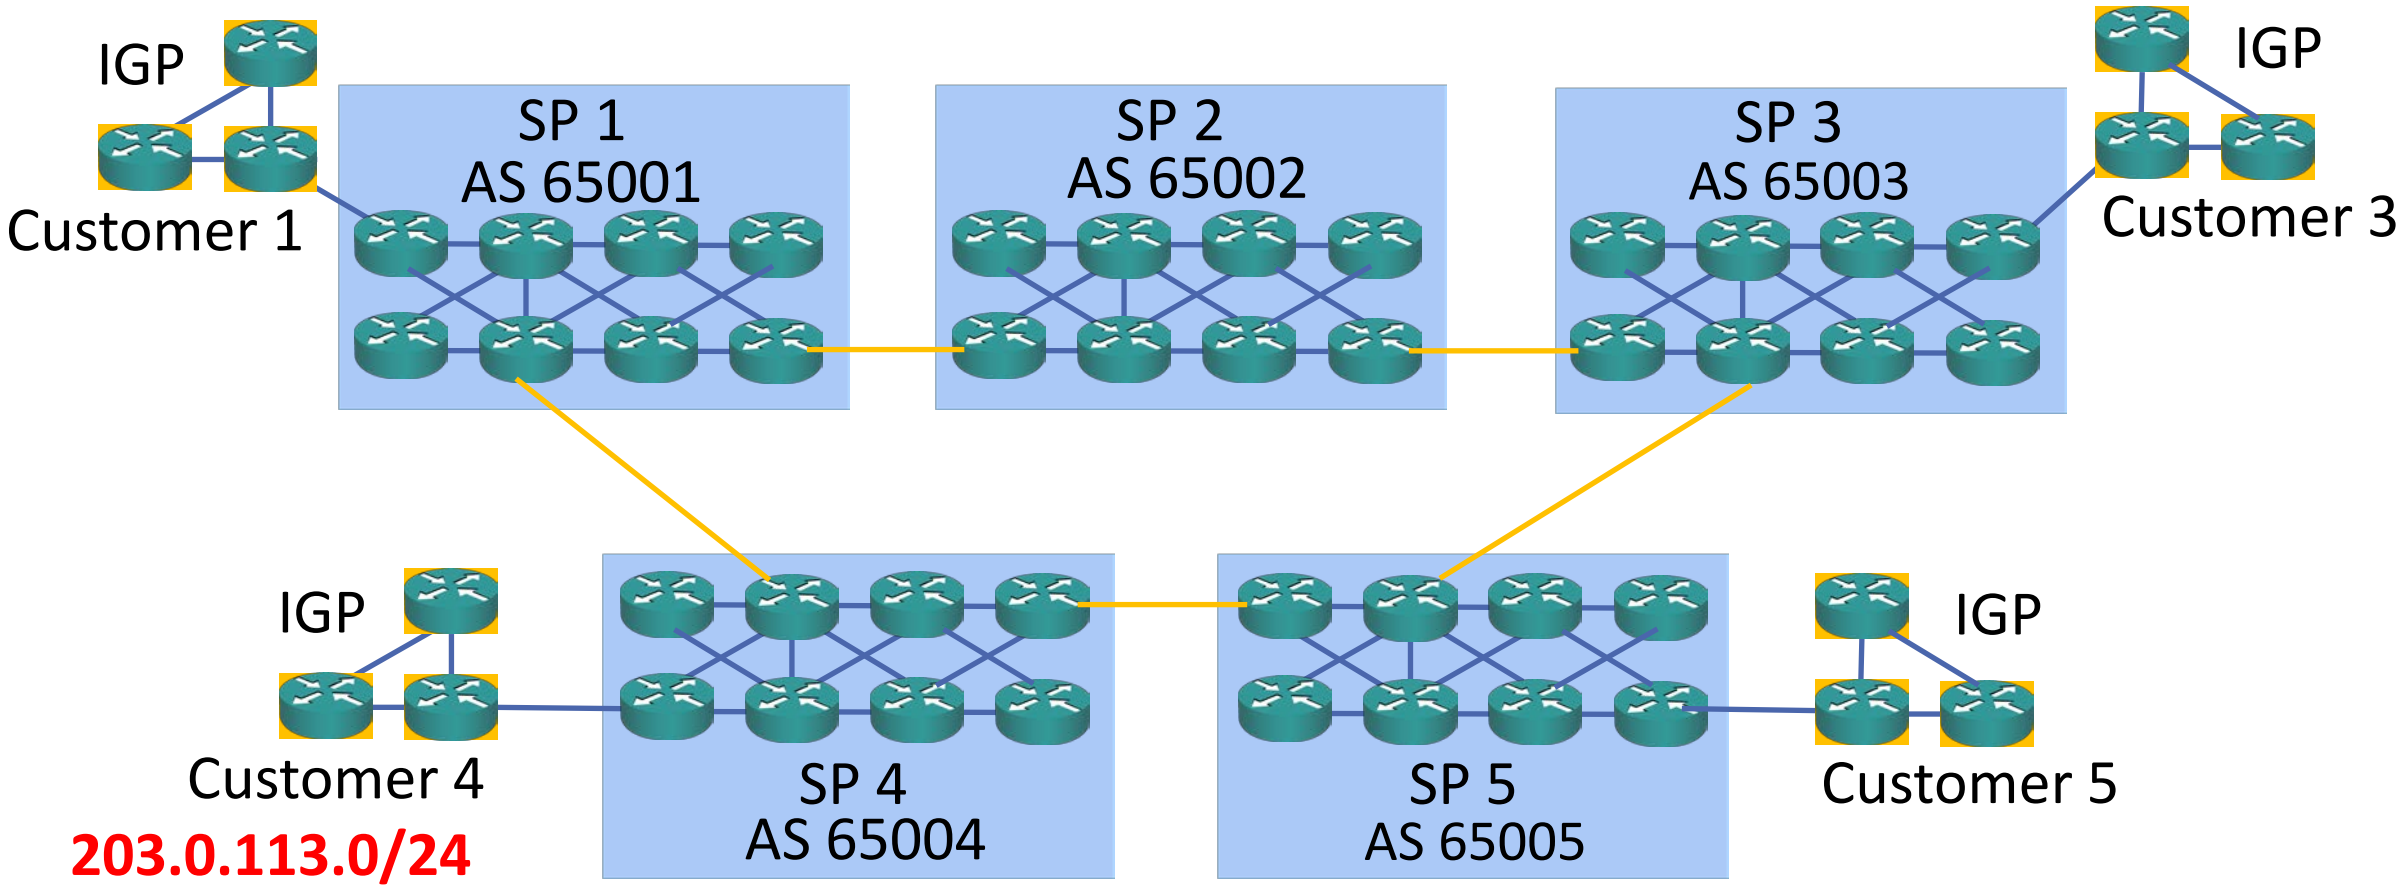

BGP Route Propagation

- The ‘network’ command in an IGP enables the routing protocol on interfaces, and advertises the networks on those interfaces
- 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 it does not automatically advertise routes. Neighbors are manually specified, and you need to manually configure route propagation separately after the neighbor configuration