

BGP Configuration



- Only one instance of BGP can be configured on a router.
- The autonomous system number identifies the autonomous system to which the router belongs.

```
R2(config)#router bgp 65002
```

BGP Identifier



- The BGP Identifier is equivalent to the Router ID in IGP
- By default, the highest IP address on a loopback interface is used, otherwise the highest IP address on a physical interface will be used

```
R2(config)#router bgp 65002
```

```
R2(config-router)#bgp router-id 192.168.0.2
```

The Neighbor Command




- Peering between routers works differently for BGP than IGPs
- BGP routers within an AS might not be directly physically connected
- Link-local multicast Hello messages cannot be used to form adjacencies
- You have to manually specify BGP neighbours
- Targeted TCP sessions using port 179 are used to establish peering and exchange routes

How BGP Works



- Administrator configures BGP and specifies the neighbour IP address
- There must be a route to the neighbour IP address in the routing table
- The neighbour relationship is set up using a unicast TCP session on port 179




R1: "I'm 172.16.0.1, I want to form a BGP peer relationship with you 10.0.0.1"

How BGP Works



- R3 replies within the TCP session and agrees to set up the BGP peer relationship if it has a matching neighbour statement for R1
- The routers exchange routes advertised in BGP, as configured by the administrator



Matching Neighbor Statements




- R3 will not accept BGP packets from 172.16.0.1 in AS 65002 unless it has a neighbor statement for 172.16.0.1 in AS 65002

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

```
R1(config-router)#neighbor 10.0.0.1 remote-as 65002
```

```
R3(config)#router bgp 65002
```

```
R3(config-router)#neighbor 172.16.0.1 remote-as 65002
```



The Neighbor Command




- A description is optional

```
R3(config)#router bgp 65002
```

```
R3(config-router)#neighbor 172.16.0.1 remote-as 65002
```

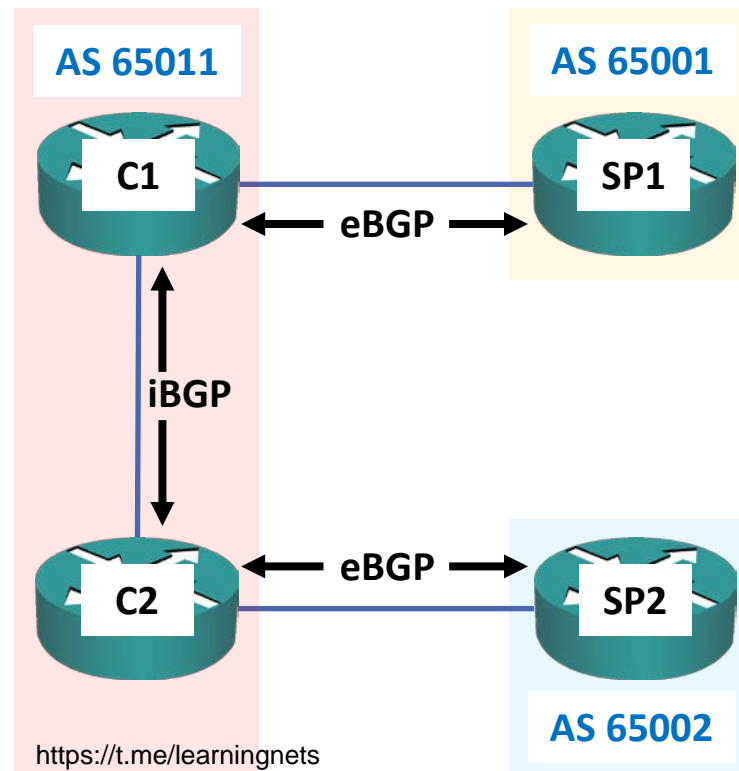
```
R3(config-router)#neighbor 172.16.0.1 description R1
```



iBGP and eBGP



- iBGP (internal BGP) neighbors are in the same AS as each other
- eBGP (external BGP) neighbors are in different ASs
- Some different rules are used for iBGP and eBGP neighbours



iBGP and eBGP



- The local AS number is compared to the AS number listed in neighbor statements to determine if the neighbor is an internal (iBGP) or external (eBGP) neighbor.

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

```
R1(config-router)#neighbor 10.0.0.1 remote-as 65002
```

```
R1(config-router)#neighbor 192.168.0.1 remote-as 65003
```

iBGP and eBGP



- Routes learned via iBGP have an Administrative Distance of 200
- Routes learned via eBGP have an Administrative Distance of 20

iBGP Split Horizon



- A BGP router will propagate:
- Routes received from iBGP neighbors to eBGP neighbors
- Routes received from eBGP neighbors to iBGP neighbors
- Routes received from eBGP neighbors to eBGP neighbors


iBGP Split Horizon



- A BGP router will not propagate a route it received from an iBGP neighbor to another iBGP router
- This is the BGP split horizon rule to prevent routing loops
- It means iBGP routers must be in a full mesh topology
- (Route Reflectors and BGP Confederations can be used to avoid the full mesh requirement)


iBGP Split Horizon – Potential Loops

- A loop could be formed if an iBGP router propagated a route to another iBGP router



iBGP Split Horizon – Need For Full Mesh


- AS 65001 will never learn the route to 203.0.113.0/24 if iBGP1 and iBGP2 are not iBGP neighbors



iBGP Neighbors



- BGP neighbours can be multiple physical hops apart and/or have multiple redundant paths to reach each other
- This can be the case with iBGP neighbours. eBGP neighbours are typically directly connected



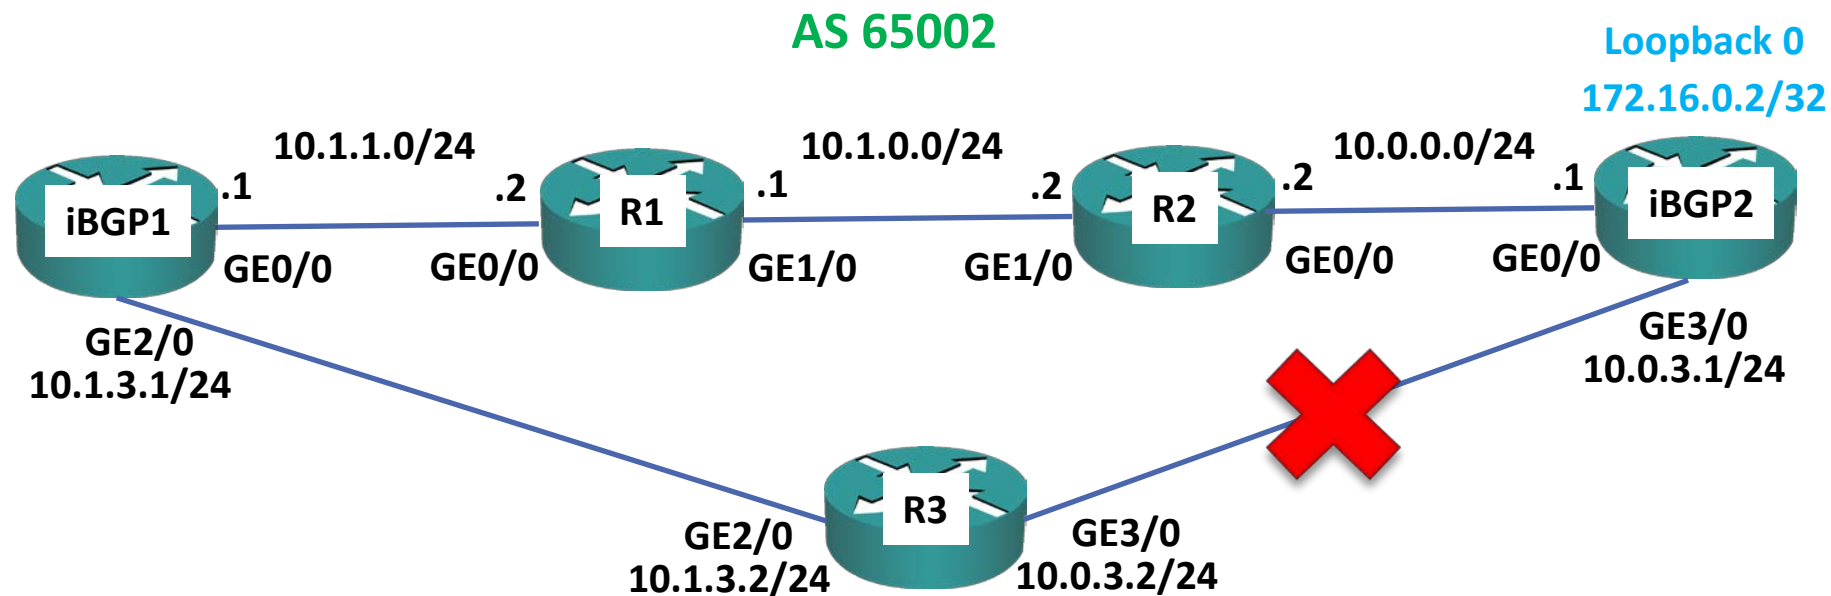
iBGP Neighbors



- Loopback addresses (which are advertised in the IGP) are typically used as the address in iBGP neighbour statements
- If the IP address of a physical interface is used for the BGP session and that physical interface goes down, then so will the BGP session even though there may be an alternate path to the router
- This way BGP peers can continue to reach each other if a physical interface goes down

iBGP Neighbors

- If an interface goes down, then the IP address on that interface is down
- A routing protocol will not find an alternative path to that IP address, the IP address is dead



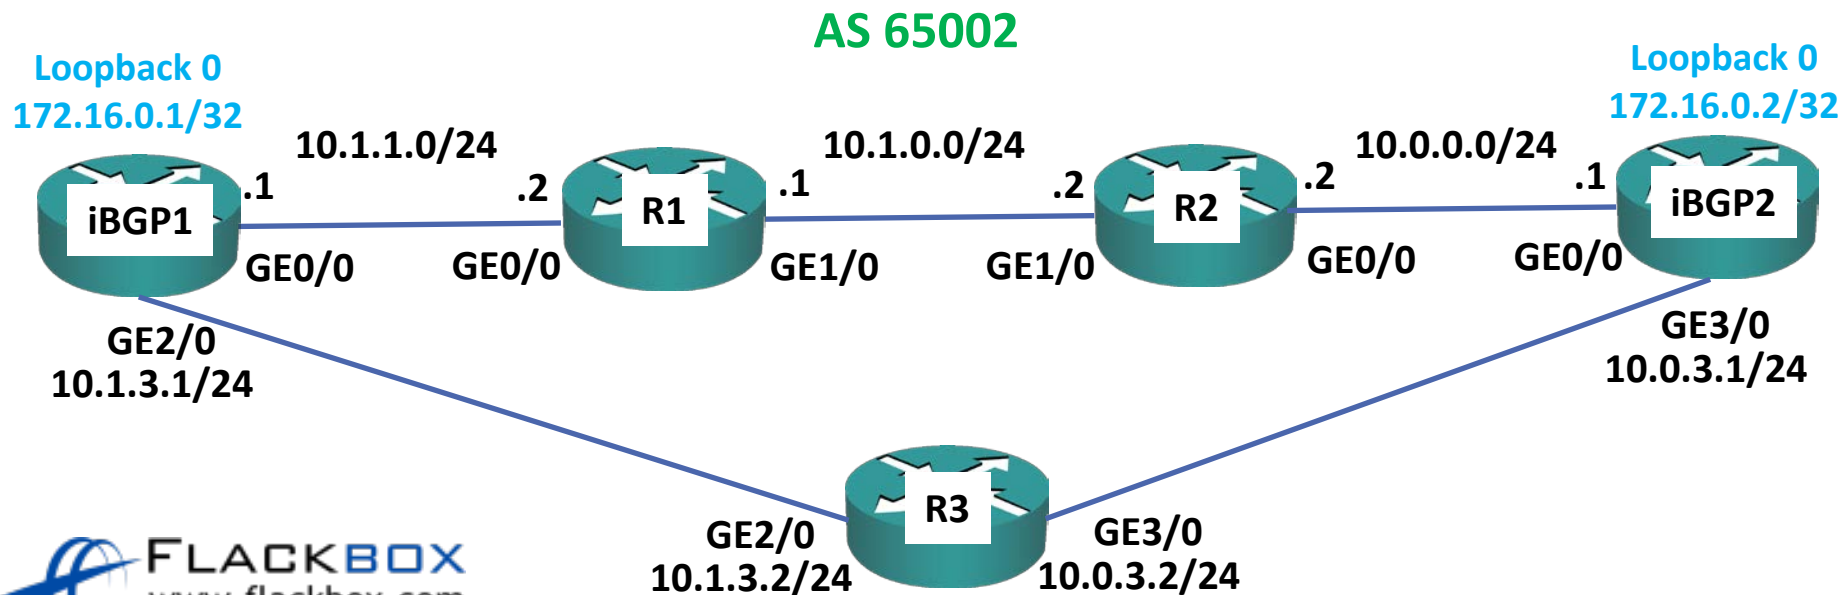
iBGP Neighbors

```
iBGP1(config)#router bgp 65002
```

```
iBGP1(config-router)#neighbor 172.16.0.2 remote-as 65002
```

```
iBGP2(config)#router bgp 65002
```

```
iBGP2(config-router)#neighbor 172.16.0.1 remote-as 65002
```



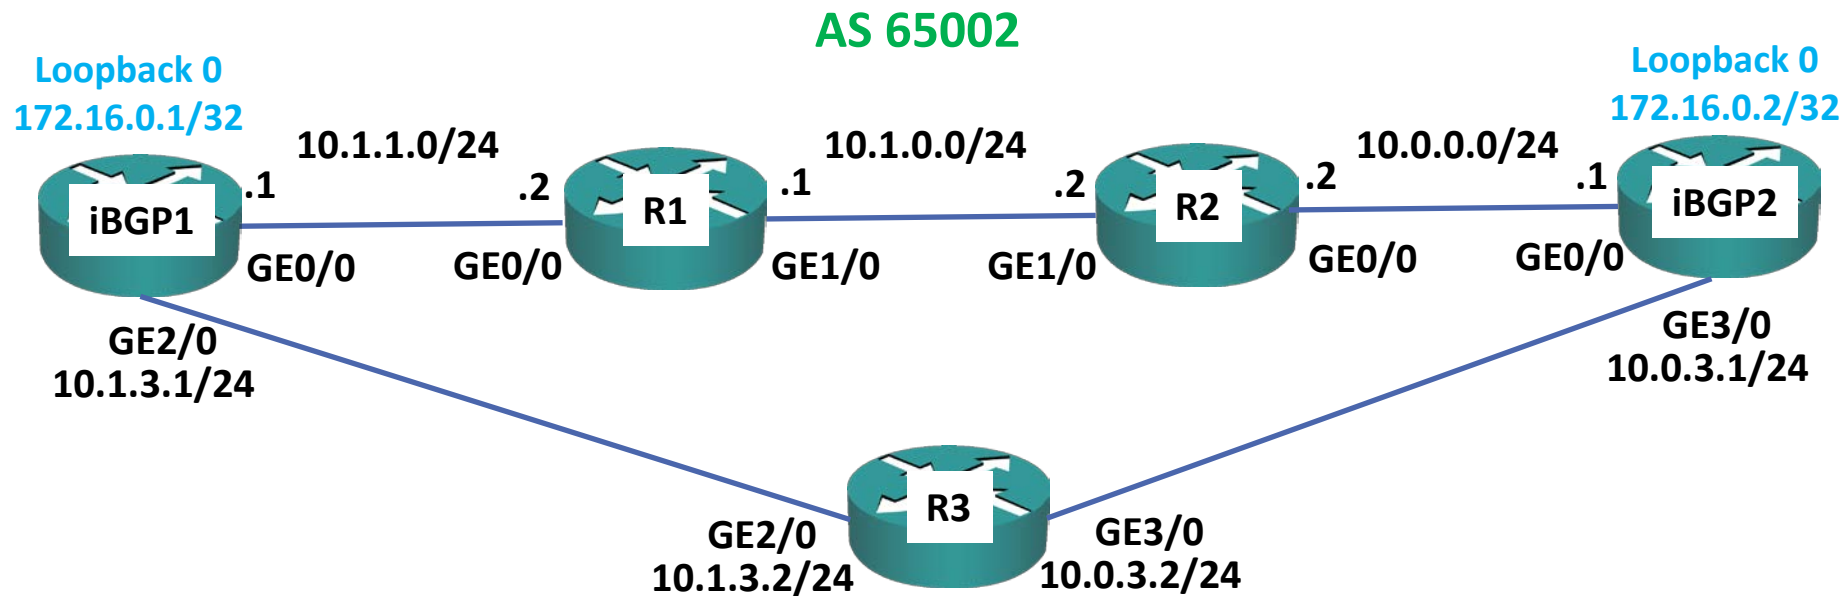
BGP between Loopback Addresses



- BGP has a security mechanism where it will only peer with another router if it has a matching neighbour statement for that peer
- The source address of packets received from the neighbour must match the exact IP address in the neighbour statement
- When a router sends packets from itself it uses the IP address of the exit interface as the source address by default
- This will cause BGP peering to fail between loopback addresses

iBGP Neighbors

- Packets from iBGP1 to iBGP2 have a source address of 10.1.3.1 by default
- iBGP2 does not have a configured neighbour for 10.1.3.1 (it's for 172.16.0.1) so they will not become peers



BGP between Loopback Addresses



```
iBGP2#show ip bgp summary
```

```
BGP router identifier 172.16.0.2, local AS number 65002
```

```
BGP table version is 1, main routing table version 1
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	
State/PfxRcd									
172.16.0.1	4	65002	0	0	1	0	0	never	Idle

BGP between Loopback Addresses

```
iBGP1(config)#router bgp 65002
iBGP1(config-router)#neighbor 172.16.0.2 remote-as 65002
iBGP1(config-router)#neighbor 172.16.0.2 update-source loopback 0
```

```
iBGP2(config)#router bgp 65002
iBGP2(config-router)#neighbor 172.16.0.1 remote-as 65002
iBGP2(config-router)#neighbor 172.16.0.1 update-source loopback 0
```

Verification – show ip bgp summary



```
R1#sh ip bgp summary
```

```
BGP router identifier 172.16.0.1, local AS number 65002
```

```
BGP table version is 1, main routing table version 1
```

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down
State/PfxRcd								
172.16.0.2	4	65002	13	13	1	0	0	00:08:23
203.0.113.65	4	65001	5	4	1	0	0	00:00:51

Verification – show ip bgp neighbors

```
R2#show ip bgp neighbors
BGP neighbor is 172.16.0.1, remote AS 65002, internal link
  BGP version 4, remote router ID 172.16.0.1
  BGP state = Established, up for 00:03:05
  Last read 00:00:23, last write 00:00:25, hold time is 180, keepalive interval is 60 seconds
  Neighbor sessions:
    1 active, is not multisession capable (disabled)
  Neighbor capabilities:
    Route refresh: advertised and received(new)
    Four-octets ASN Capability: advertised and received
    Address family IPv4 Unicast: advertised and received
    Enhanced Refresh Capability: advertised and received
    Multisession Capability:
    Stateful switchover support enabled: NO for session 1
  Message statistics:
    InQ depth is 0
    OutQ depth is 0
```

	Sent	Rcvd
Opens:	1	1
Notifications:	0	0
Updates:	1	1
Keepalives:	5	5
Route Refresh:	0	0
Total:	7	7

Verification – show ip bgp neighbors - Piped

```
R1#show ip bgp neighbors | include (BGP neighbor is|Description|BGP state)
```

```
BGP neighbor is 172.16.1.2, remote AS 65002, external link
```

```
Description: SP2
```

```
BGP state = Established, up for 00:32:19
```

```
BGP neighbor is 192.168.0.4, remote AS 65001, internal link
```

```
Description: SP1A
```

```
BGP state = Established, up for 00:15:16
```

```
BGP neighbor is 198.51.100.2, remote AS 65011, external link
```

```
Description: C1A
```

```
BGP state = Established, up for 00:33:43
```

```
BGP neighbor is 198.51.100.6, remote AS 65012, external link
```

```
Description: C2
```

```
BGP state = Established, up for 00:33:49
```