Configuration and Management of Networks

Pedro Amaral

BGP – Connecting to an ISP

Enterprise session initiation requirement:

- One-way: Connectivity from an enterprise network toward the Internet is the only connectivity required.
- Two-way: Connectivity from the Internet to an enterprise network is also required.

Solutions:

- One-way: private IP address space with address translation
- Two-way: public IP address space (in combination with private) and proper routing

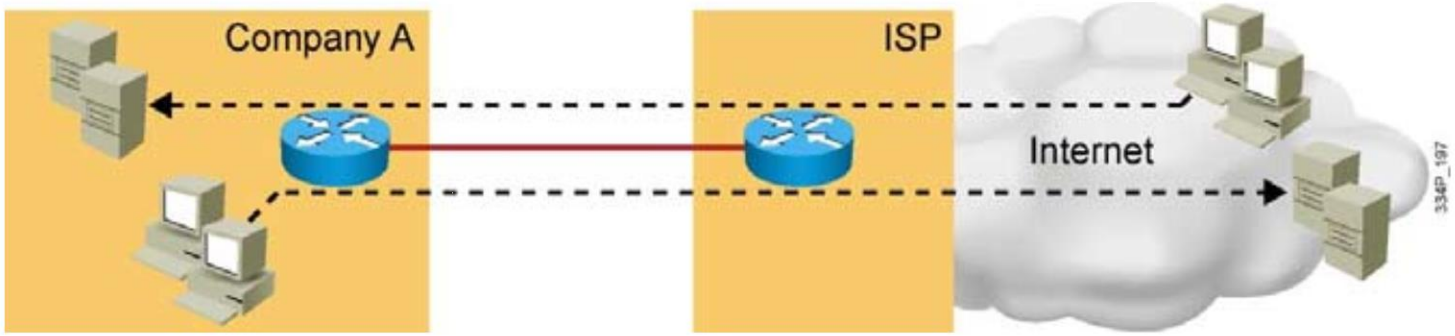

BGP – Connecting to an ISP

- Public IP address space (subpool or whole /24 subnet)
- Link type and bandwidth availability
- Routing options
- Connection redundancy
- Independency in regard to an ISP:
  - Public IP address space
  - AS number

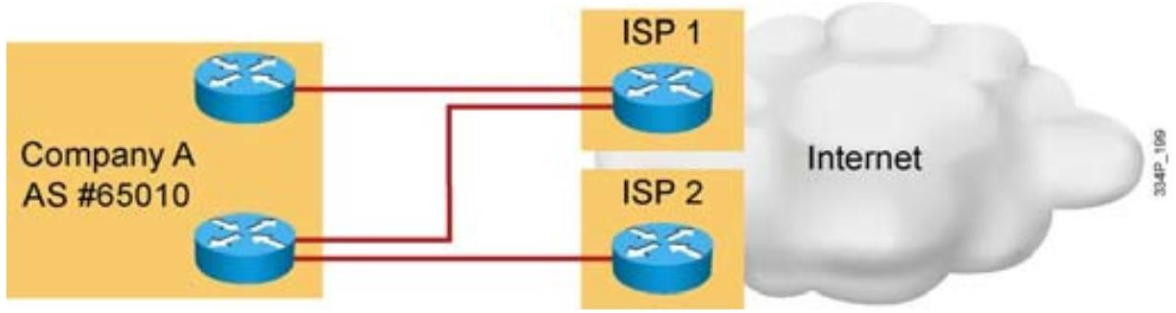

BGP – Connecting to an ISP

- The customer uses the default route toward the ISP.
- The service provider uses static route(s) for customer public network
- No automatic adjustment to any changes in the network.

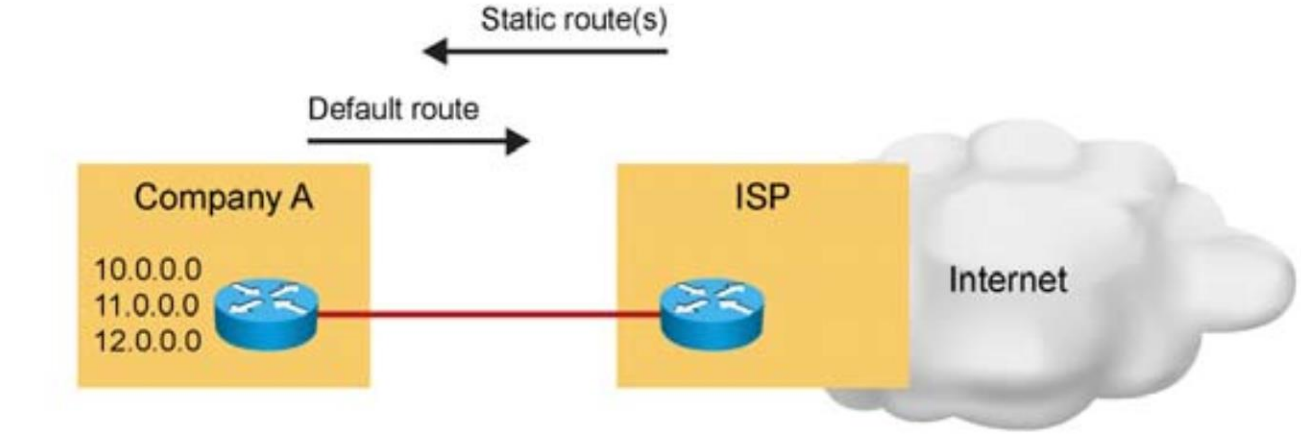

R1#

```
router eigrp 110
 network 10.0.0.0
ip default-network 0.0.0.0
ip route 0.0.0.0 0.0.0.0 serial 0/0
```

PE#

```
ip route 10.0.0.0 255.0.0.0 serial 1/0
ip route 11.0.0.0 255.0.0.0 serial 1/0
ip route 12.0.0.0 255.0.0.0 serial 1/0
```

BGP – Connecting to an ISP

- The customer deploys BGP to announce its public networks
- The ISP announces a default route, a subset of Internet routes, or a complete Internet routing table
- Typically used for inter-AS routing

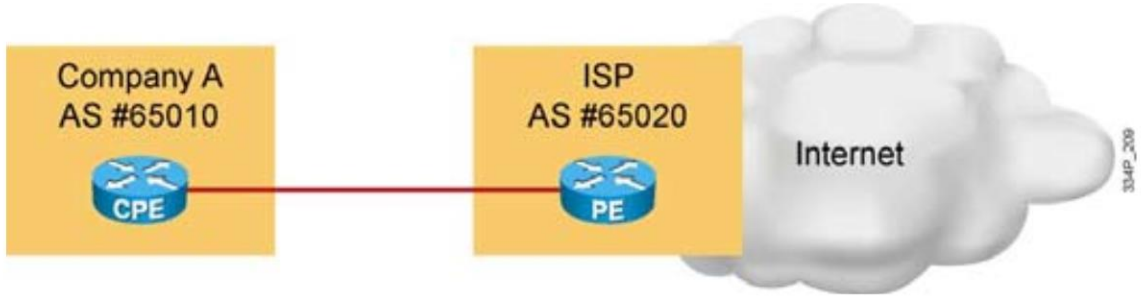

CPE—Customer to Provider Edge router

BGP – Connecting to an ISP

- Single-homed
- Dual-homed
- Multihomed
- Dual-multihomed

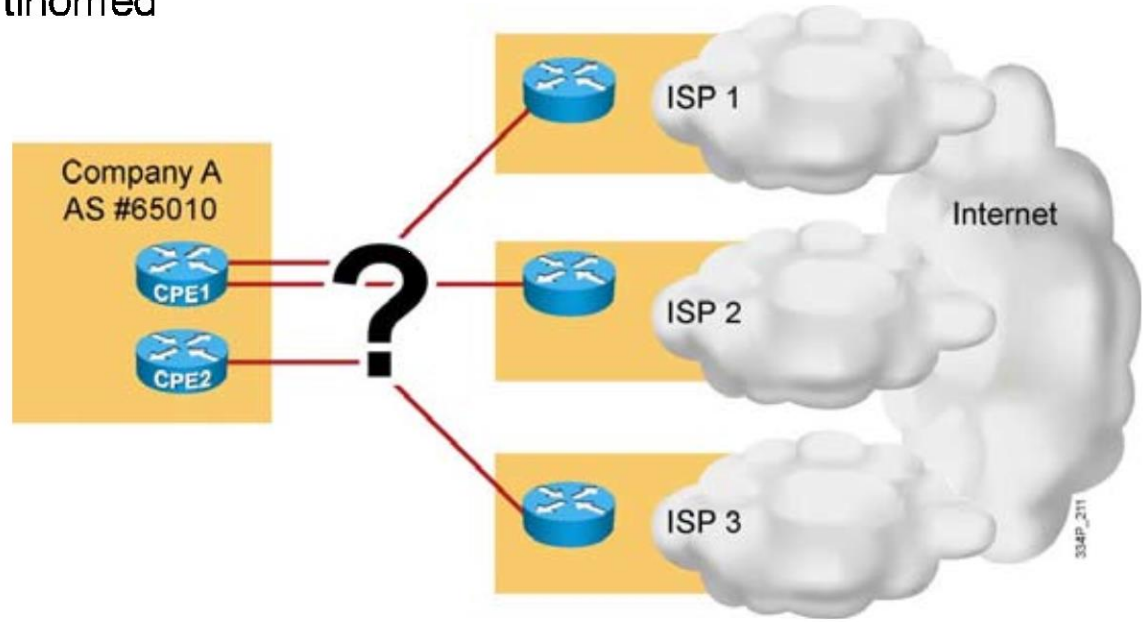

BGP – Multihoming from 2 providers

Characteristics:

- Connected to two or more different ISPs
- Can use a single router or multiple edge routers
- Dynamic routing with BGP

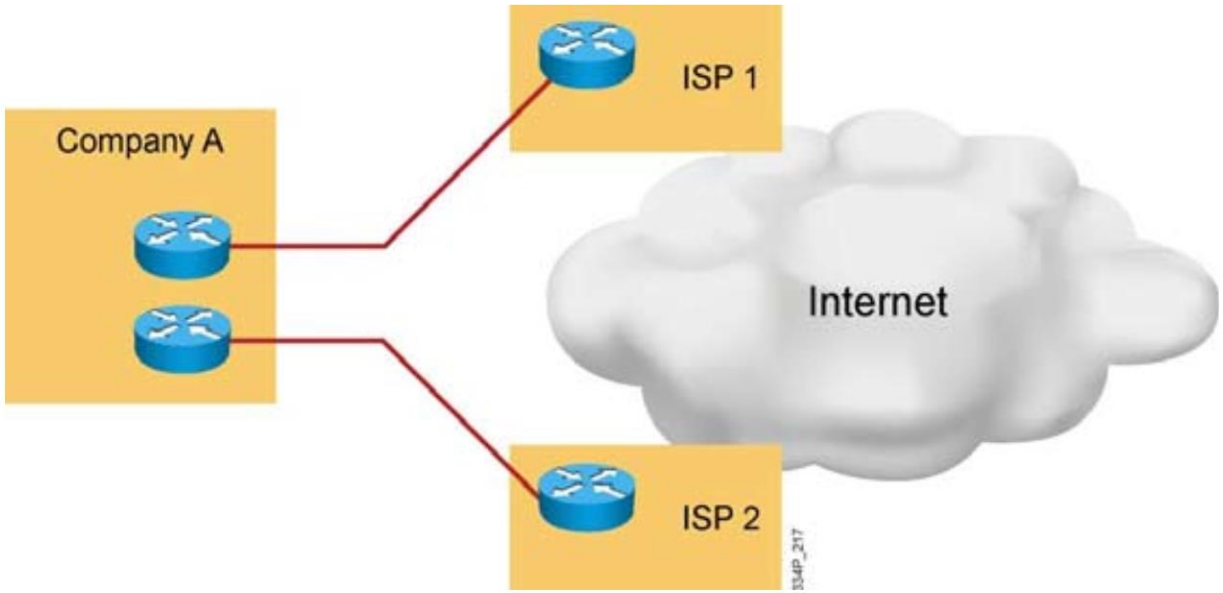

BGP – Connecting to an ISP summary

- Connecting an enterprise network to an ISP requires, at a minimum, a public IP address pool, a proper link to the ISP, consideration of redundancy requirements, and the proper routing protocol.
- To exchange routing updates with an ISP, the customer can use different options. Static routes and BGP are the options that are most commonly used.
- The way in which the customer connects to an ISP depends on the redundancy requirements, where a single-homed connectivity has no redundancy, and the dual-multihomed connectivity has the most redundancy built in.

BGP – Terminology

- Autonomous system (AS): a collection of networks under a single administrative domain
- Interdomain routing: routing between the customer and the ISP
- Internal routing: uses IGP protocol (RIP, OSPF, EIGRP, and so on) to exchange routing information inside the AS
- External routing: uses EGP protocol (BGP) to exchange routes between autonomous systems
- Two BGP implementations:
  - Internal BGP (IBGP): when BGP is used inside an AS
  - External BGP (EBGP): when BGP is used between autonomous systems

BGP – Connecting to an ISP

Multihoming options with respect to connections

- Using a single connection to an individual ISP
- Using multiple connections to an individual ISP

Multihoming options with respect to routing

- Default routes from all providers
- Default routes and partial Internet routing from the providers
- Full Internet routing from the providers

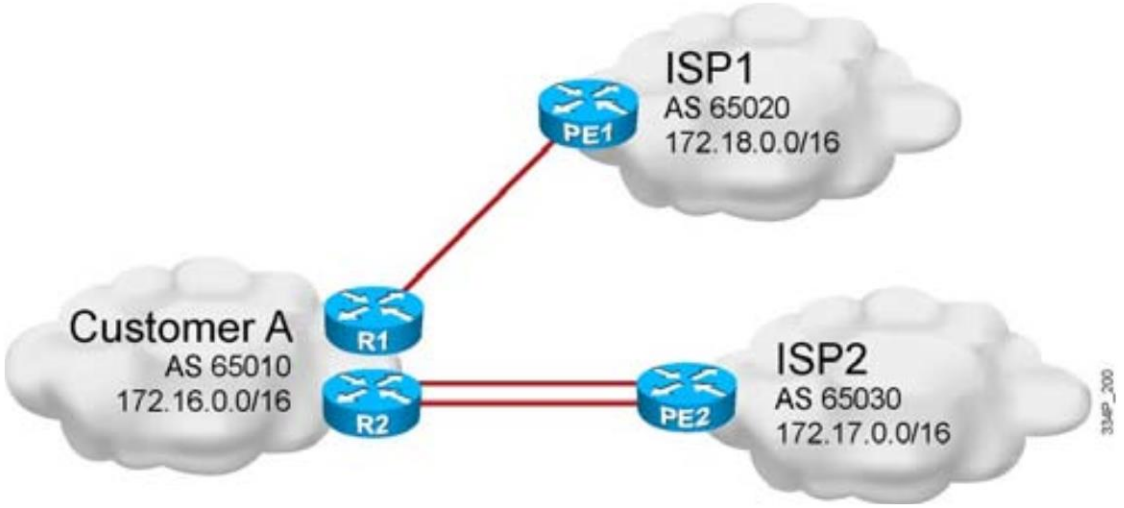

BGP – Default routes

- Customer A receives the default route from each ISP.

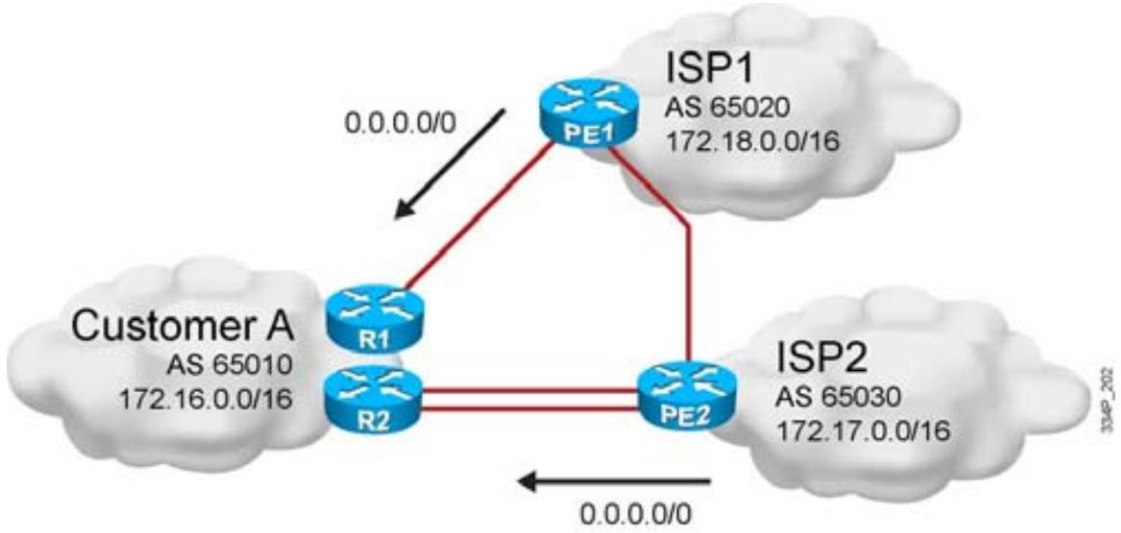

Only one default route will be used (the one with smallest IGP cost)

BGP – Default routes and partial tables

- Customer A receives the default route from each ISP.
- Customer A receives a partial routing table from each ISP.

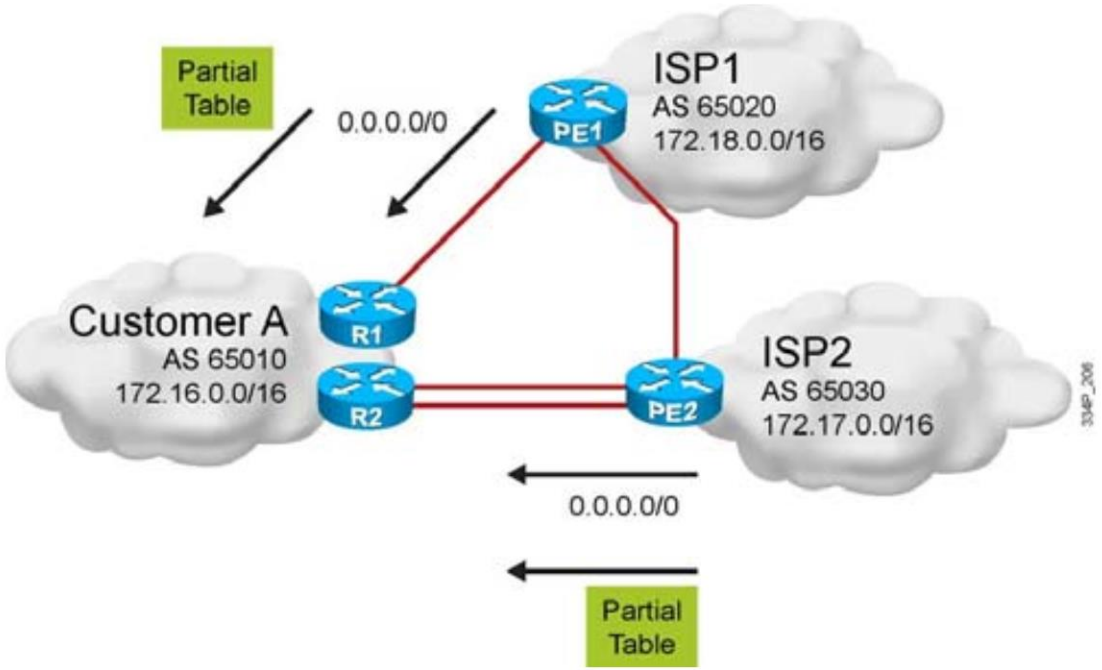

BGP – Default routes and partial tables

- The partial table is used to forward traffic to the correct ISP.
- If the destination is unknown, then a default route to one of the ISPs is used.

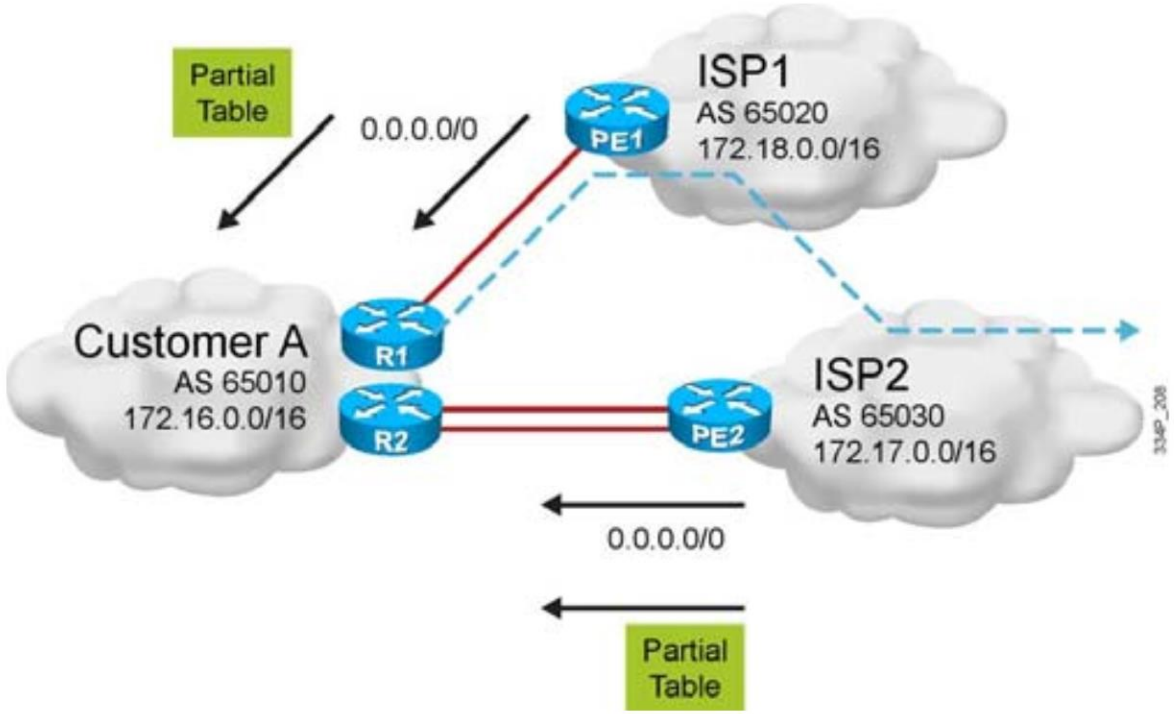

BGP – Full Table

- Customer A receives a full routing table from each ISP.
- Requires that enough memory and CPU resources are available.

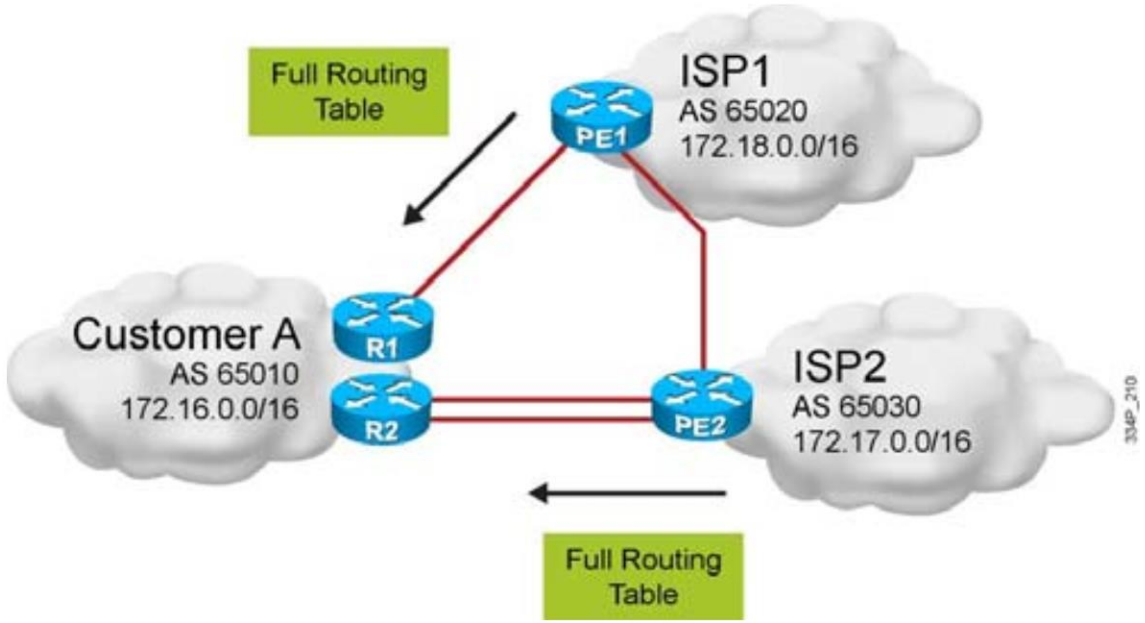

BGP – Routing between ASes

- Autonomous system (AS) – a collection of networks under a single technical administration.
  - 16-bit numbers (as of January 2009 32-bit numbers are available)
  - Ranging from 1 to 65535
  - Private AS: 64512–65535
- Internet Assigned Numbers Authority (IANA) allocates AS numbers.
- IGPs operate within an AS.
- BGP is used between autonomous systems.

BGP

- BGP is used to provide an interdomain routing system.
- BGP guarantees the exchange of loop-free routing information.
- BGP works differently than IGPs.
  - BGP is a policy-based routing protocol.
  - Control traffic flow using multiple BGP path attributes.

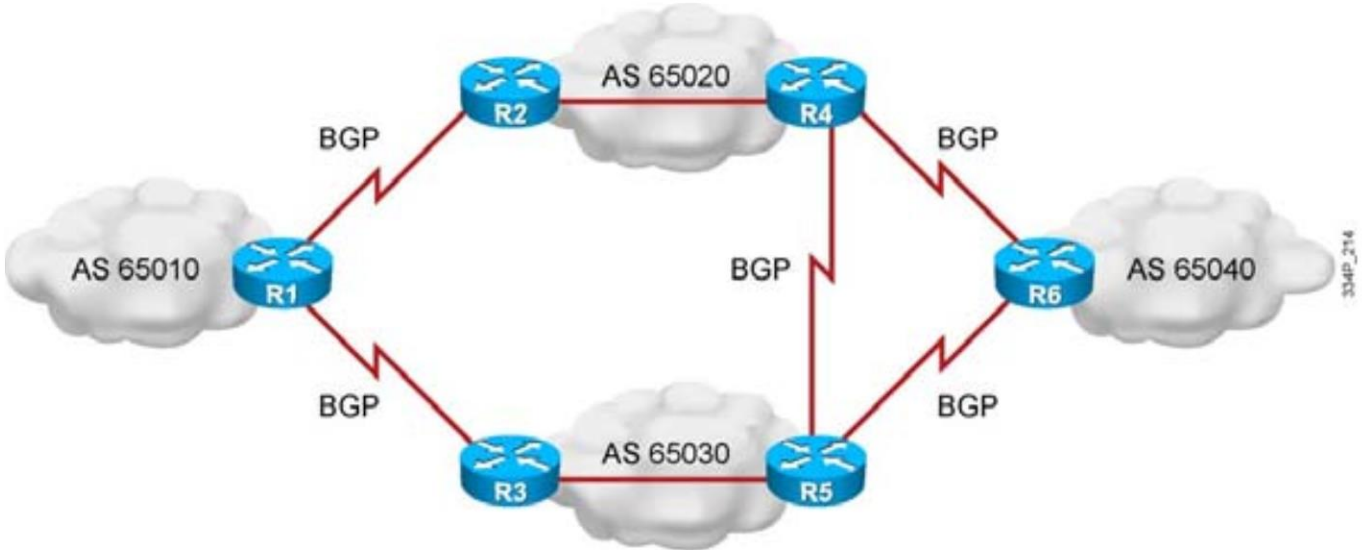

BGP

- BGP announces:
  - Paths
  - Networks that are reachable at the end of the path
- The path is described by using attributes.
- The administrator can define data flow through autonomous systems.

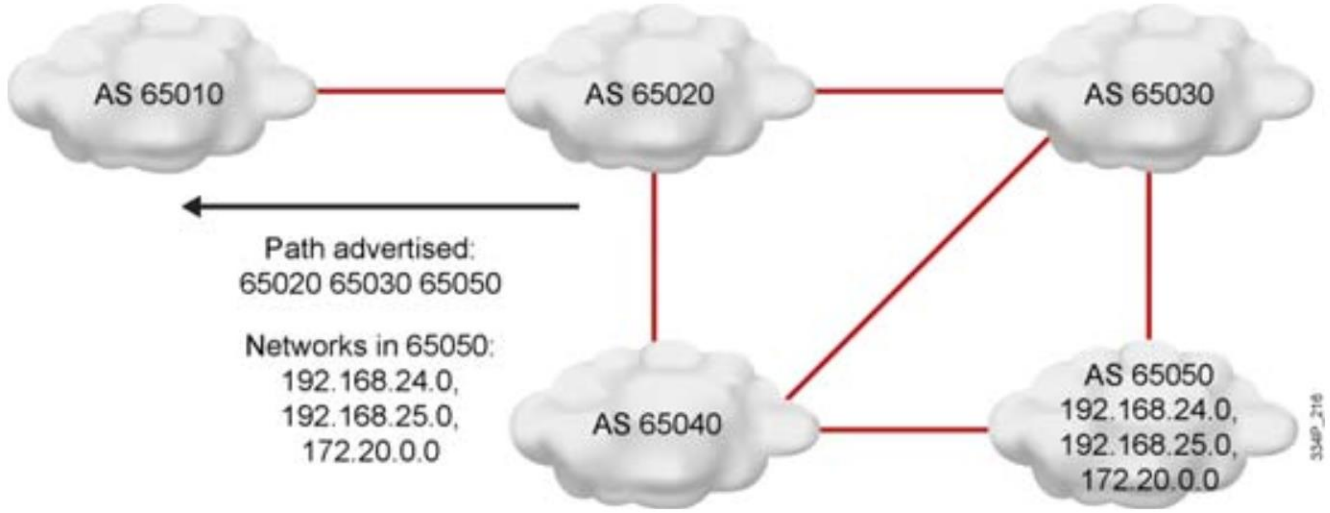

BGP

BGP is a path vector protocol with the following properties:

- Reliable updates: BGP runs on top of TCP (port 179)
- Incremental, triggered updates only
- Periodic keepalive messages to verify TCP connectivity
- Rich metrics (called path vectors or attributes)
- Designed to scale to huge internetworks (for example, the Internet)

It has enhancements over distance vector protocols.

BGP Databases

- BGP neighbor table
  - List of BGP neighbors
- BGP table
  - List of all networks learned from each BGP neighbor
  - Multiple paths to same destination network can be present
  - Each path is associated with BGP attributes
- IP routing table (forwarding database)
  - List of best paths to destination networks used to forward traffic

BGP Summary

- BGP is typically used for interdomain routing.
- Three common ways to perform multihoming with BGP are as follows:
  - Each ISP passes only a default route.
  - Each ISP passes only a default route and specific provider-owned routes.
  - Each ISP passes all routes.
- BGP is the external routing protocol used between autonomous systems. Forwarding is based on policies and not on best path.
- BGP routers exchange network reachability information called path vectors, made up of path attributes.
- A router running BGP keeps its own tables to store BGP information that it receives from and sends to other routers, including a neighbor table, a BGP table, and an IP routing table.

BGP

- A “BGP neighbor” is also known as a “BGP peer”
- TCP connection

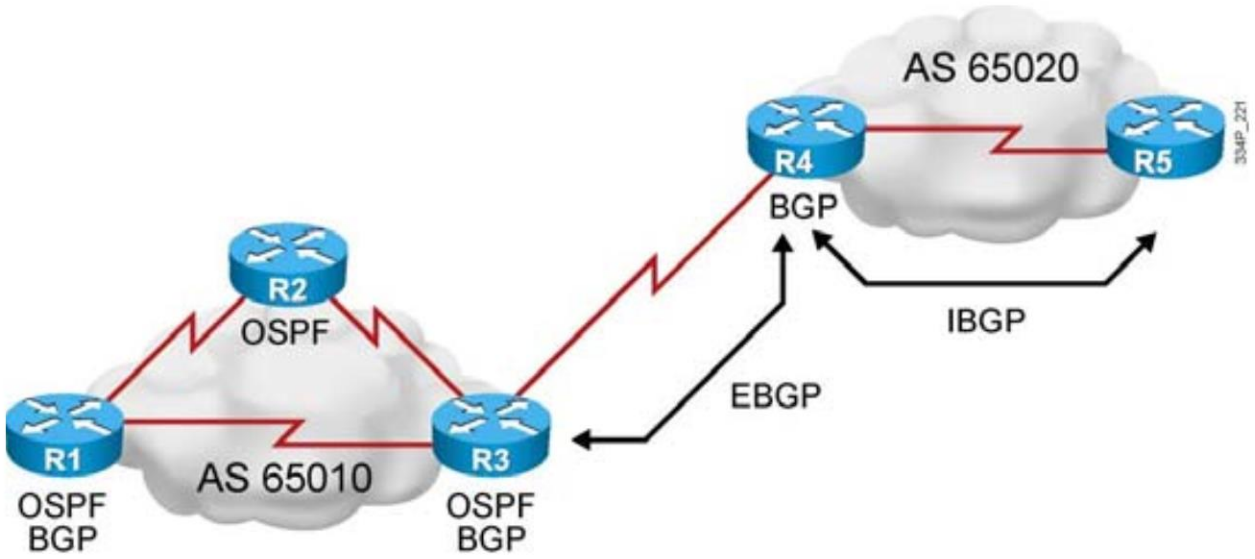

BGP

- A BGP neighbor in same AS
- Requirements:
  - Same AS number
  - Defined neighbors
  - Reachability

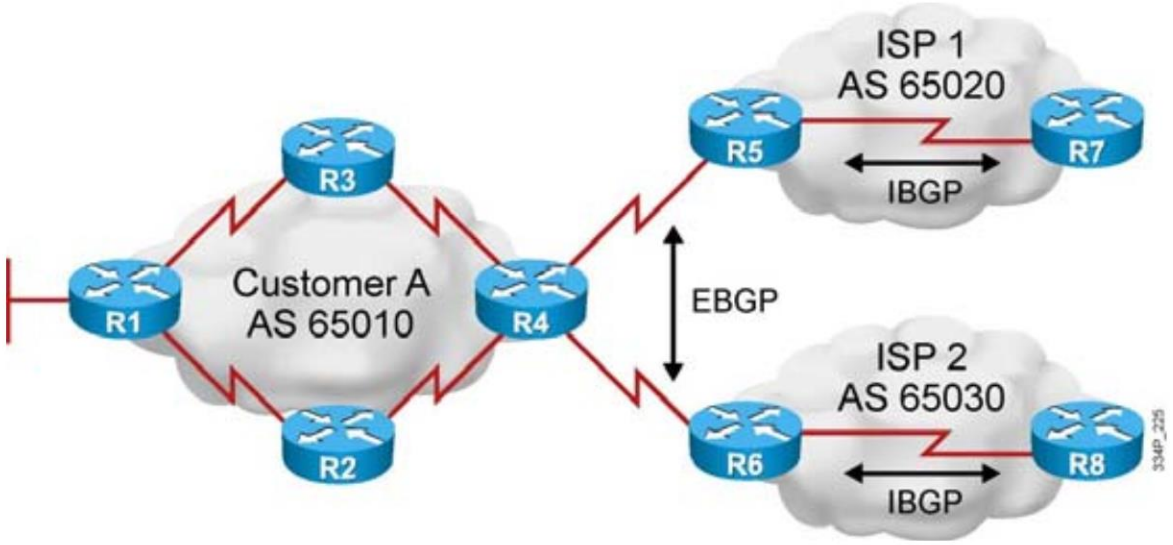

BGP - Basic Configuration

- Define the BGP process
- Establish a EBGP neighbor relationship
- Advertise the networks

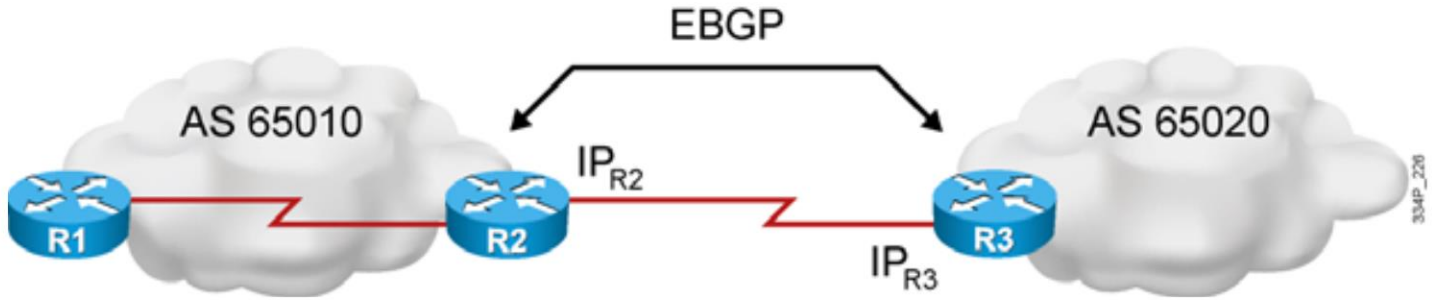

BGP

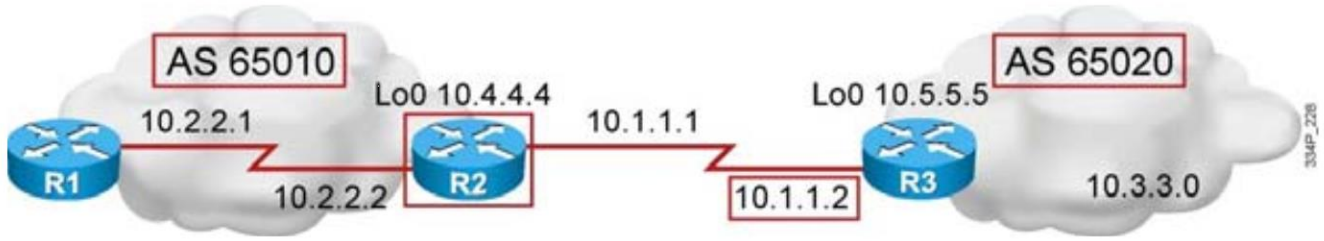


```
R2(config)#
router bgp 65010
```

- Define the BGP process locally with a local AS number.

```
R2(config-router)#
neighbor 10.1.1.2 remote-as 65020
```

Activate EBGP session to the neighbor

- Remote router IP Address and AS number

BGP

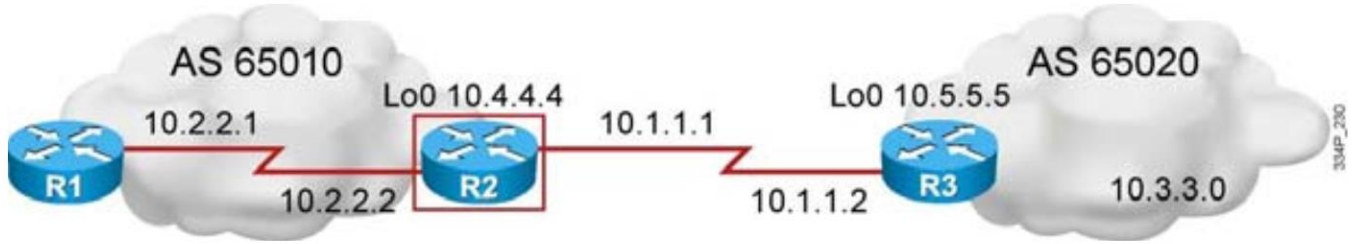

Option 1:

```
R2(config-router)#
network 10.2.2.0 mask 255.255.255.0
network 10.4.4.0 mask 255.255.255.0
```

- Configure the local networks to be advertised and include them in BGP

Option 2:

- Redistribution from IGP to BGP

BGP

```
R2(config-router)#
 network 192.168.1.1 mask 255.255.255.0
```

- The router looks for 192.168.1.1/24 in the routing table, but cannot find it, so it will not announce anything.

```
R2(config-router)#
 network 192.168.0.0 mask 255.255.0.0
```

- The router looks for 192.168.0.0/16 in the routing table.
- If the exact route is not in the table, you can add a static route to null0 so that the route can be announced.

```
R2(config-router)#
 network 192.168.1.0
```

- The router looks for a C class 192.168.1.0 network in the routing table.

BGP

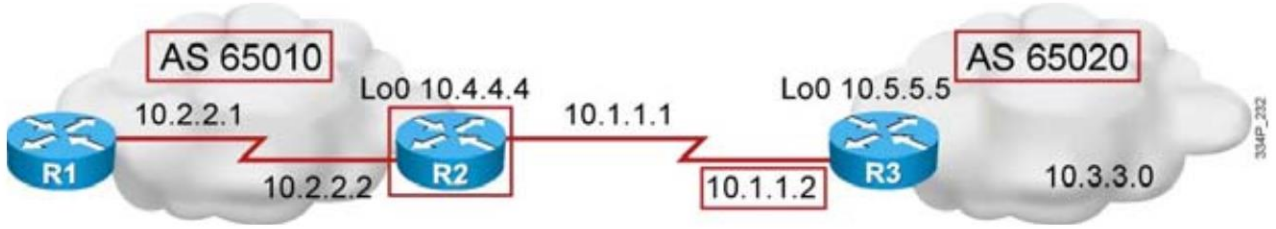

Option 1

```
R2#
!
<output omitted>
!
router bgp 65010
 neighbor 10.1.1.2 remote-as 65020
 network 10.2.2.0 mask 255.255.255.0
 network 10.4.4.0 mask 255.255.255.0
!
```

Option 2

```
R2#
!
<output omitted>
router ospf 10
 network 10.2.2.0 mask 255.255.255.0
 network 10.4.4.0 mask 255.255.255.0
!
router bgp 65010
 neighbor 10.1.1.2 remote-as 65020
 redistribute ospf
!
```

BGP EBGP and IBGP

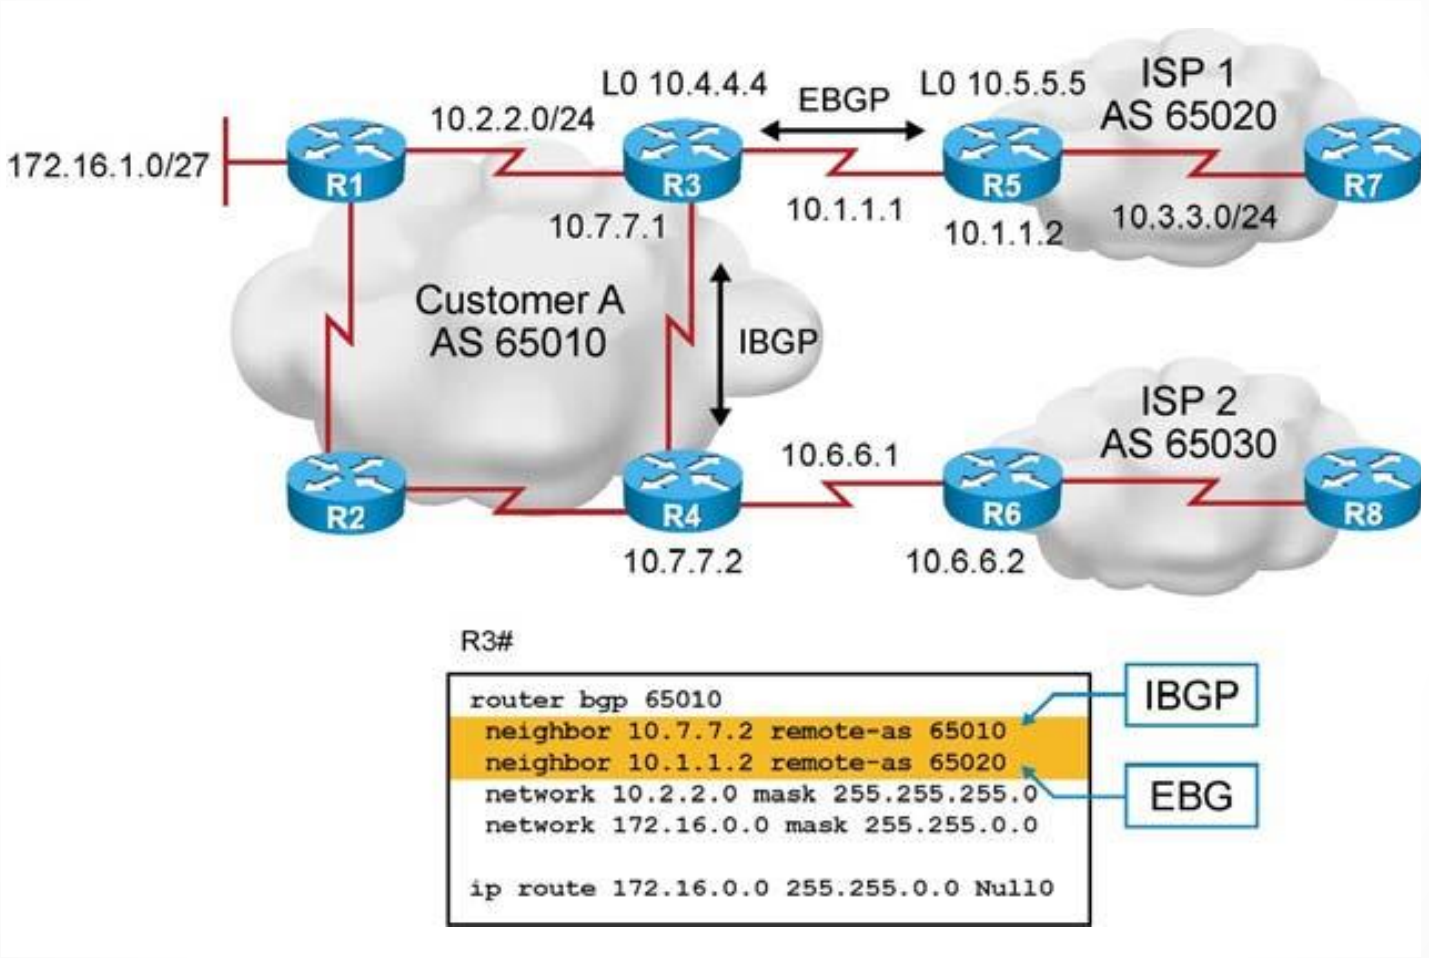

BGP neighbour shutdown

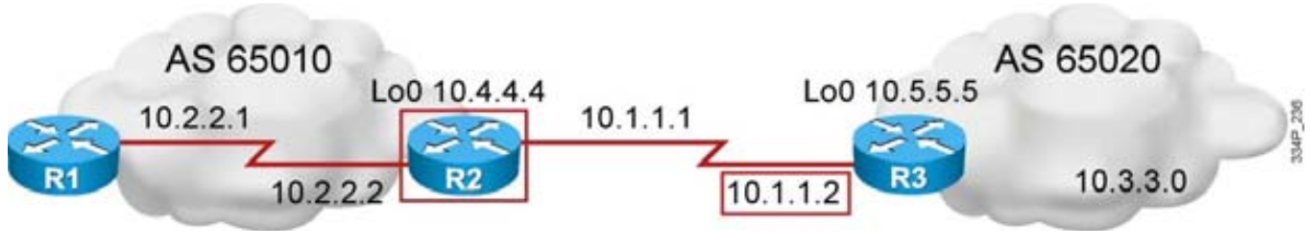


```
R2(config-router)#
neighbor 10.1.1.2 shutdown
```

- Administratively brings down a BGP neighbor
- Used for maintenance/policy changes to prevent route flapping

```
R2(config-router)#
no neighbor 10.1.1.2 shutdown
```

- Reenables a BGP neighbor that has been administratively shut down

BGP

- An IBGP neighbor relationship is established.
- What happens if the link between R3 and R4 goes down?
- Which IP address should be used to establish an IBGP session?

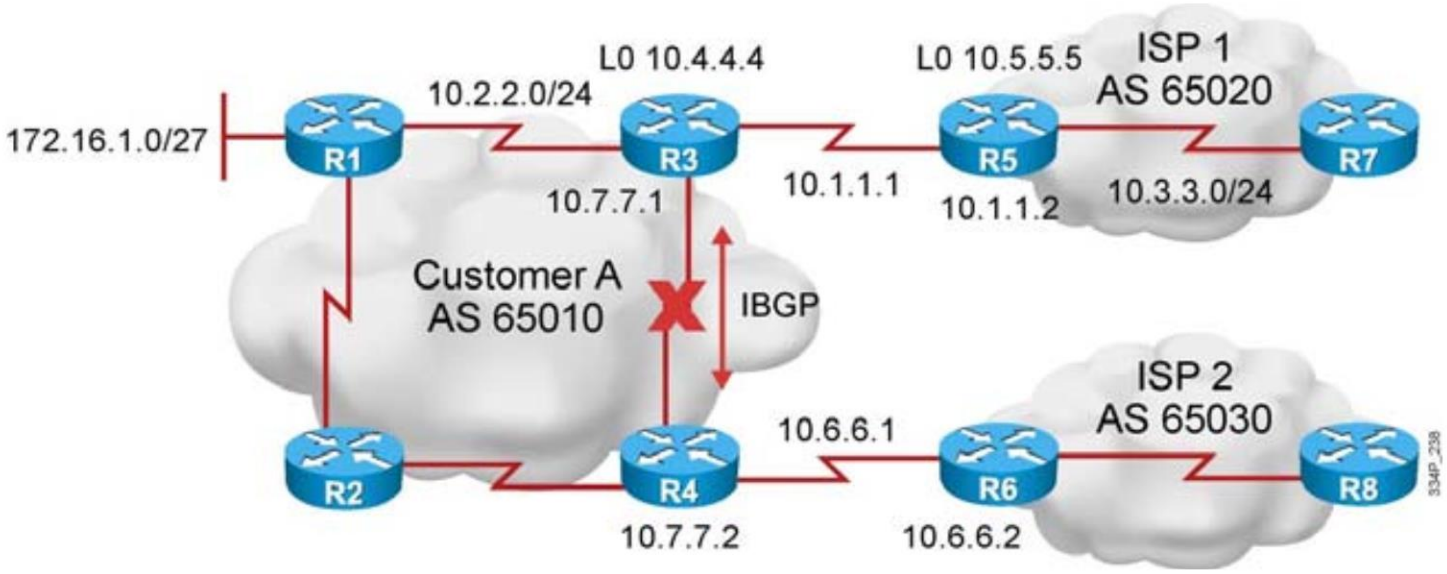

BGP

Create a BGP packet:

- The destination IP address defined by the neighbor statement
- The source IP address defined by the outbound interface

The source address of the received BGP packet is compared to list of neighbor statements:

- If a match is found in the list of neighbors, a relationship is established.
- If no match is found in the list of neighbors, the packet is ignored.

BGP IGP with loopback interfaces

- A loopback interface can be used as the source and destination IP address of all BGP updates between neighbors.
- The **neighbor update-source** command is normally used only with IBGP neighbors.

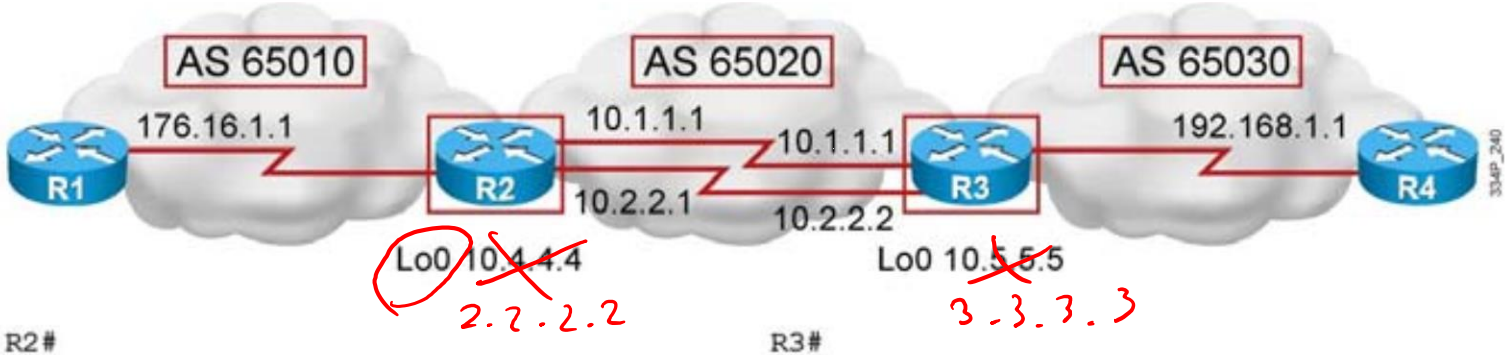

R2#

```
router bgp 65020
neighbor 172.16.1.1 remote-as 65010
neighbor 3.3.3.3 remote-as 65020
neighbor 3.3.3.3 update-source Loopback0
!
router eigrp 1
network 10.0.0.0
network 2.0.0.0
```

R3#

```
router bgp 65020
neighbor 192.168.1.1 remote-as 65030
neighbor 2.2.2.2 remote-as 65020
neighbor 2.2.2.2 update-source Loopback0
!
router eigrp 1
network 10.0.0.0
network 3.0.0.0
```

BGP

- IBGP does not modify next hop.

R3

Route	Next hop
11.0.0.0/8	10.1.1.2

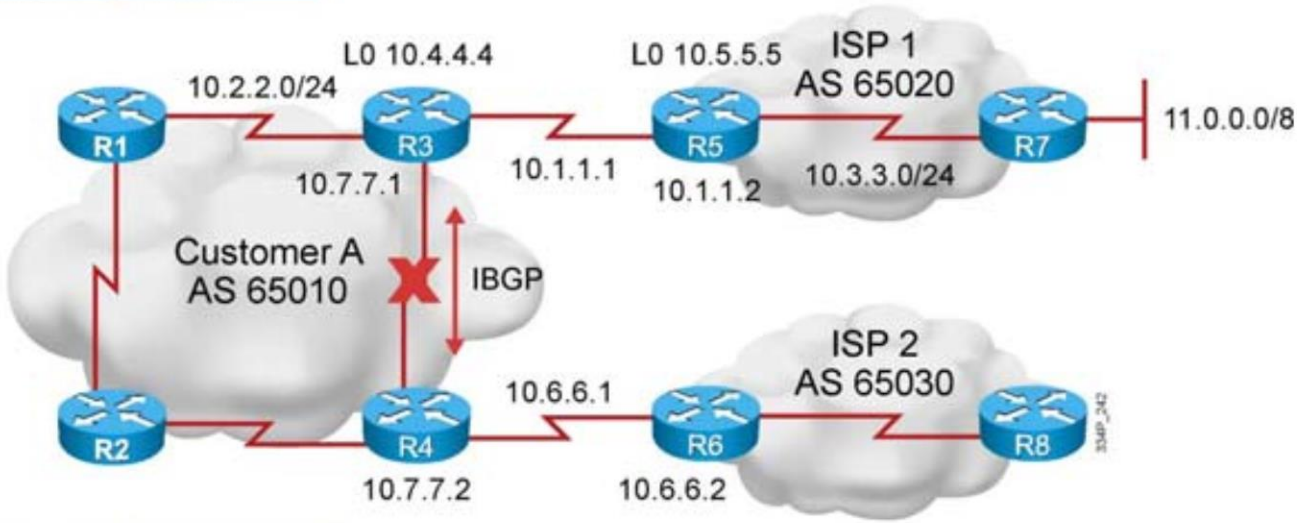

R4

Route	Next hop
11.0.0.0/8	10.1.1.2

BGP

R3

Route	Next hop
11.0.0.0/8	10.1.1.2

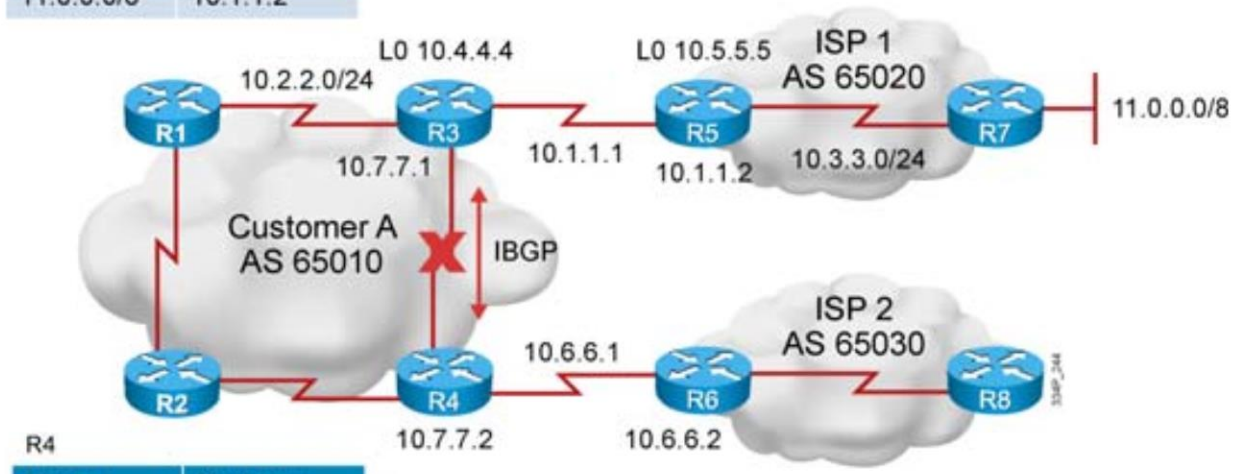

R4

Route	Next hop
11.0.0.0/8	10.4.4.4

R3 (config-router) #

```
neighbor 10.8.8.8 next-hop-self
```

- Forces all updates for neighbor R4 to be advertised with this router as the next hop—the same IP address as for the source of the BGP packet.

BGP Authentication

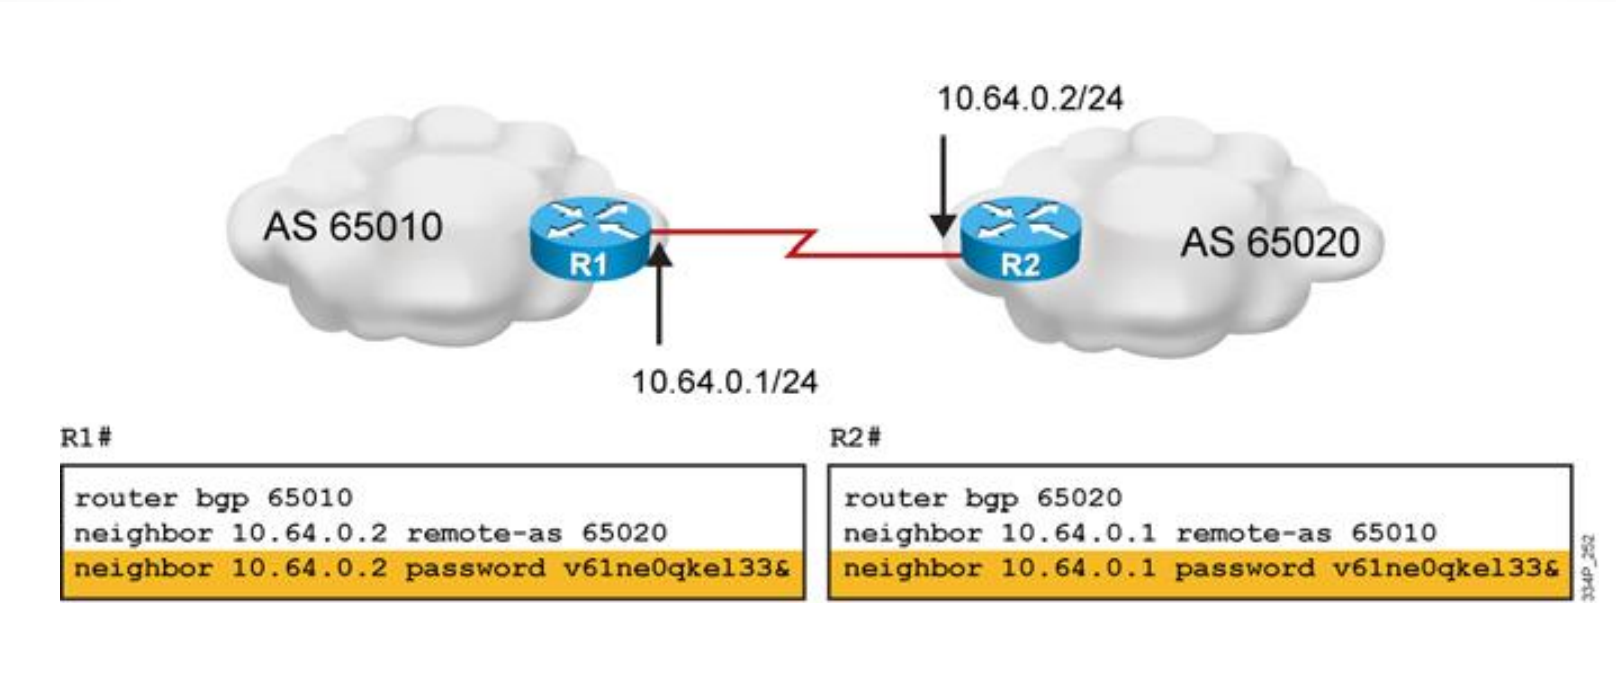

BGP – Path selection process

- The BGP table can have several paths for each network to choose from
- BGP is not designed to perform load balancing:
  - Paths are chosen because of policy.
  - Paths are not chosen based upon bandwidth.
- The BGP selection process eliminates any multiple paths until a single best path remains.

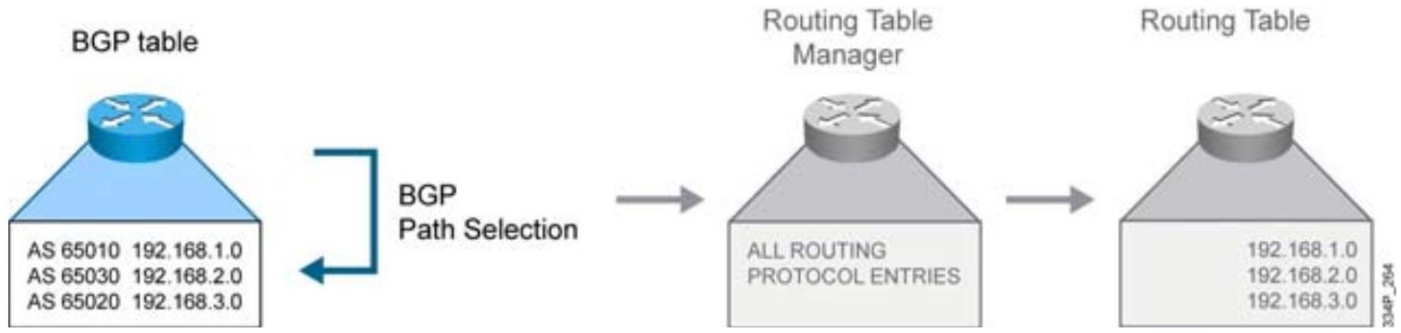

BGP

The best path is submitted to the routing table manager process.

- The best path is evaluated against the routes of other routing protocols for reaching that network.
- The route with the lowest administrative distance from the source will be installed in the routing table.

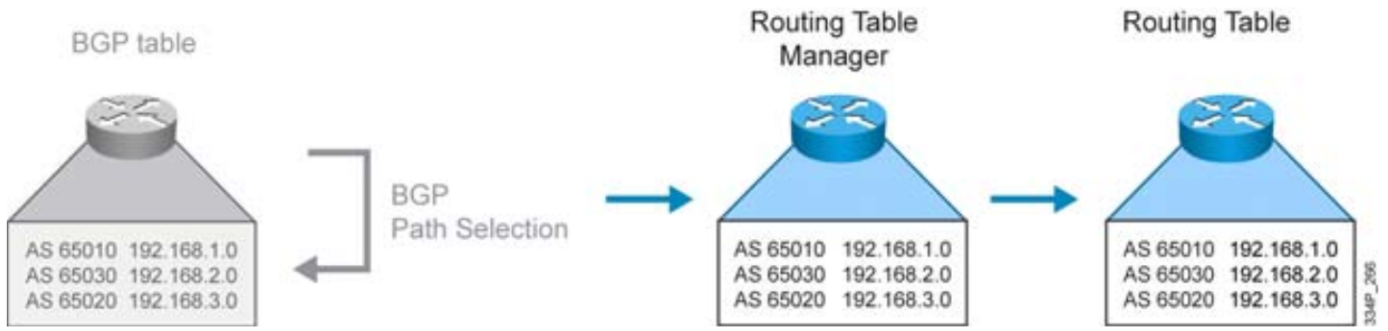

BGP Decision Process

Consider only (synchronized) routes with no AS loops and a valid next hop. The next steps in the evaluation process are:

1.	Prefer highest weight (local to router).
2.	Prefer highest local preference (global within AS).
3.	Prefer route originated by the local router (next hop = 0.0.0.0).
4.	Prefer shortest AS path.
5.	Prefer lowest origin code (IGP < EGP < incomplete).
6.	Prefer lowest MED (exchanged between autonomous systems).
7.	Prefer the EBGP path over the IBGP path.
8.	Prefer the path through the closest IGP neighbor.
9.	Prefer the oldest route for EBGP paths.
10.	Prefer the path with the lowest neighbor BGP router ID.
11.	Prefer the path with the lowest neighbor IP address.

BGP – Local Preference

- Used to select the outbound EBGP path
- Sent to IBGP neighbors only (and only within the AS)
- Stripped in the outgoing EBGP updates except in the EBGP updates with confederation peers
- The local preference attribute is well known and discretionary
- Default value = 100
- Paths with the highest local preference value are preferred

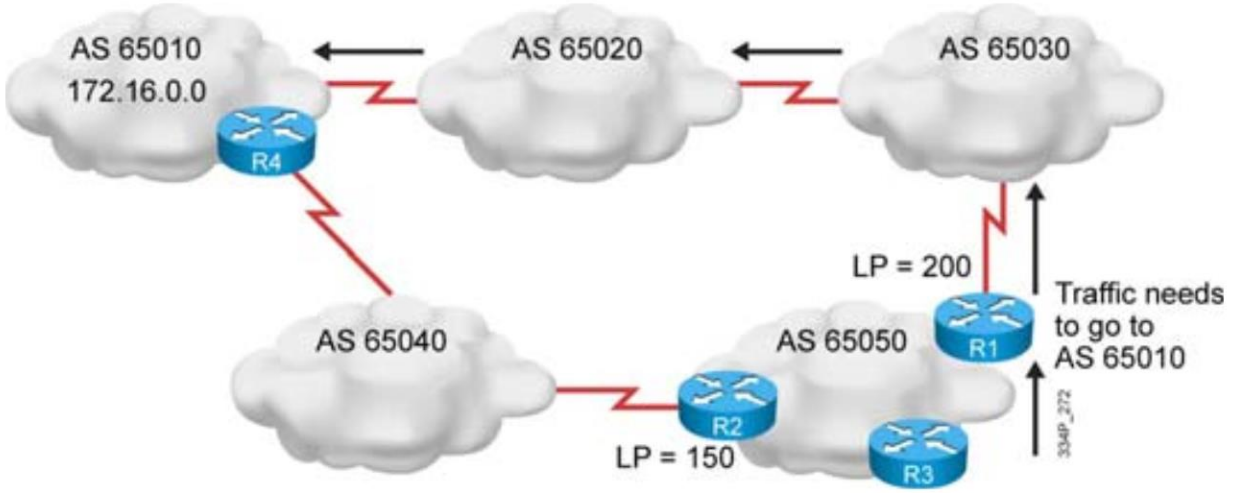

BGP Set with routing map

- Second BGP path selection criteria
- Prefer highest local preference (global within AS)
- Local preference can be set when
  - processing incoming route updates
  - doing redistribution
  - sending outgoing route updates
- BGP local preference can be specified per neighbor by complex criteria with route maps

BGP

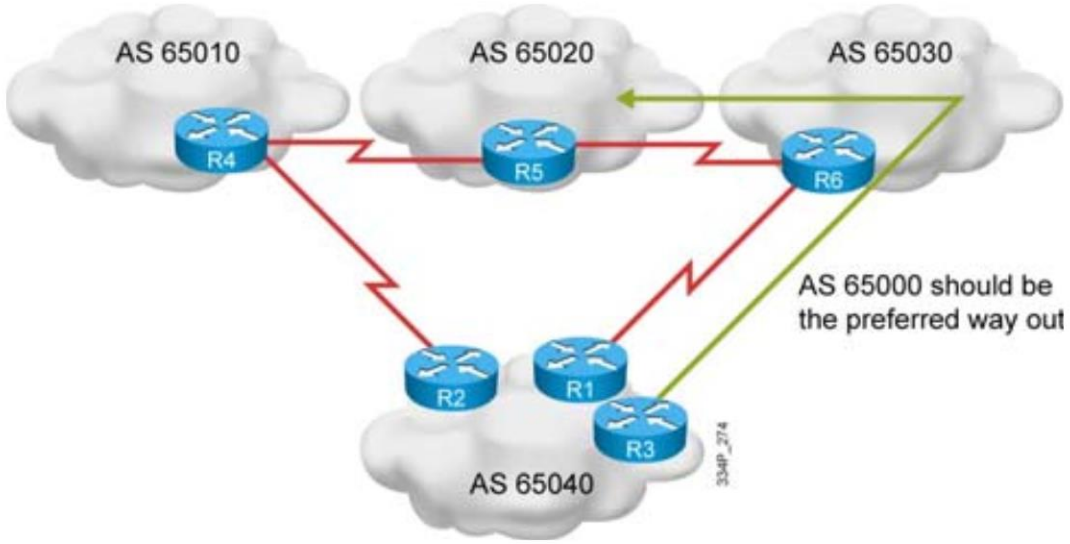

R1#

```
<output omitted>
!
router bgp 64520
 neighbor 10.0.0.1 route-map RM-SET-LP in
!
route-map RM-SET-LP permit 10
 set local-preference 150
```

BGP – Set AS paths with route map

- Fourth BGP path selection criteria
- Prefer shorter AS paths (only length is compared)
- Influences the outbound path selection in a multihomed AS
- Manual manipulation of AS path length—AS path prepending
- AS path prepending can be specified per neighbor by complex criteria with route maps (AS path filters, prefix lists, or other BGP attributes that match the routes in any combination)

BGP – AS path prepending

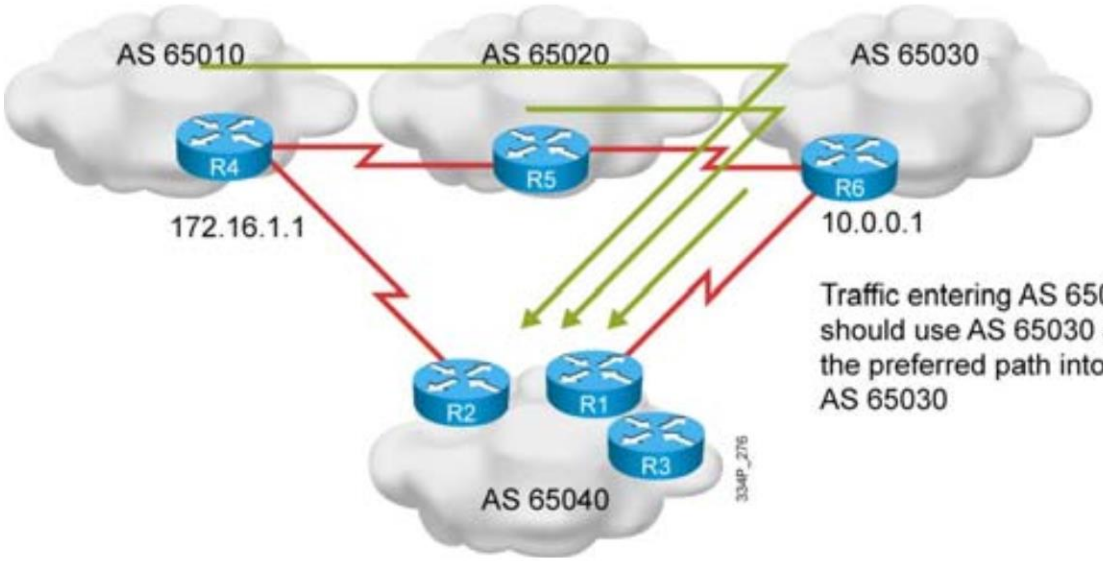

R2

R1#

```
<output omitted>
!
router bgp 65040
 neighbor 172.16.1.1 route-map RM-SET-ASPath out
!
route-map RM-SET-ASPath permit 10
 set as-path prepend 65040 65040 65040
```

BGP – MED Attribute

- The paths with the lowest MED (also called the metric) value are the most desirable.
- MED is used to advertise an exit path to be used by EBGP neighbors to reach networks owned by this AS.
- The MED attribute is optional and nontransitive.

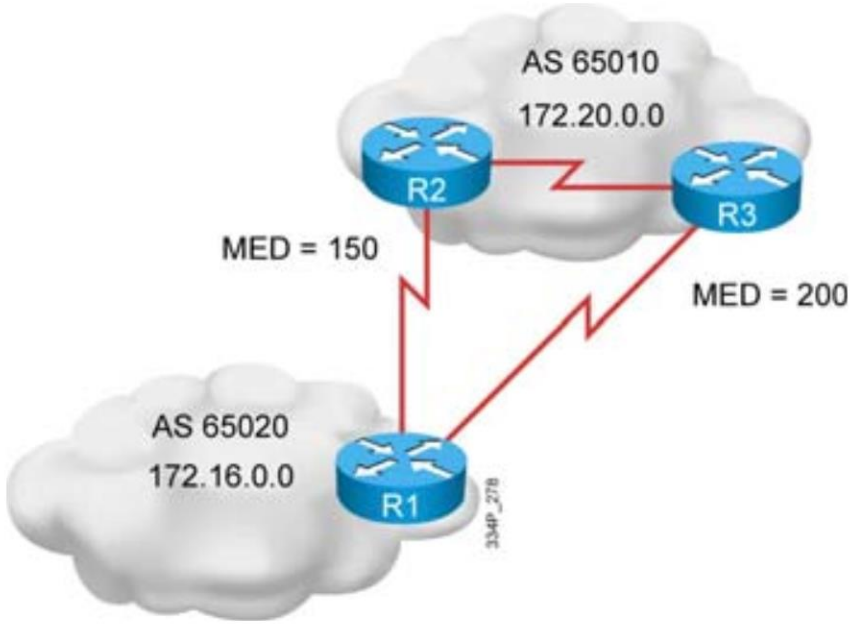

BGP

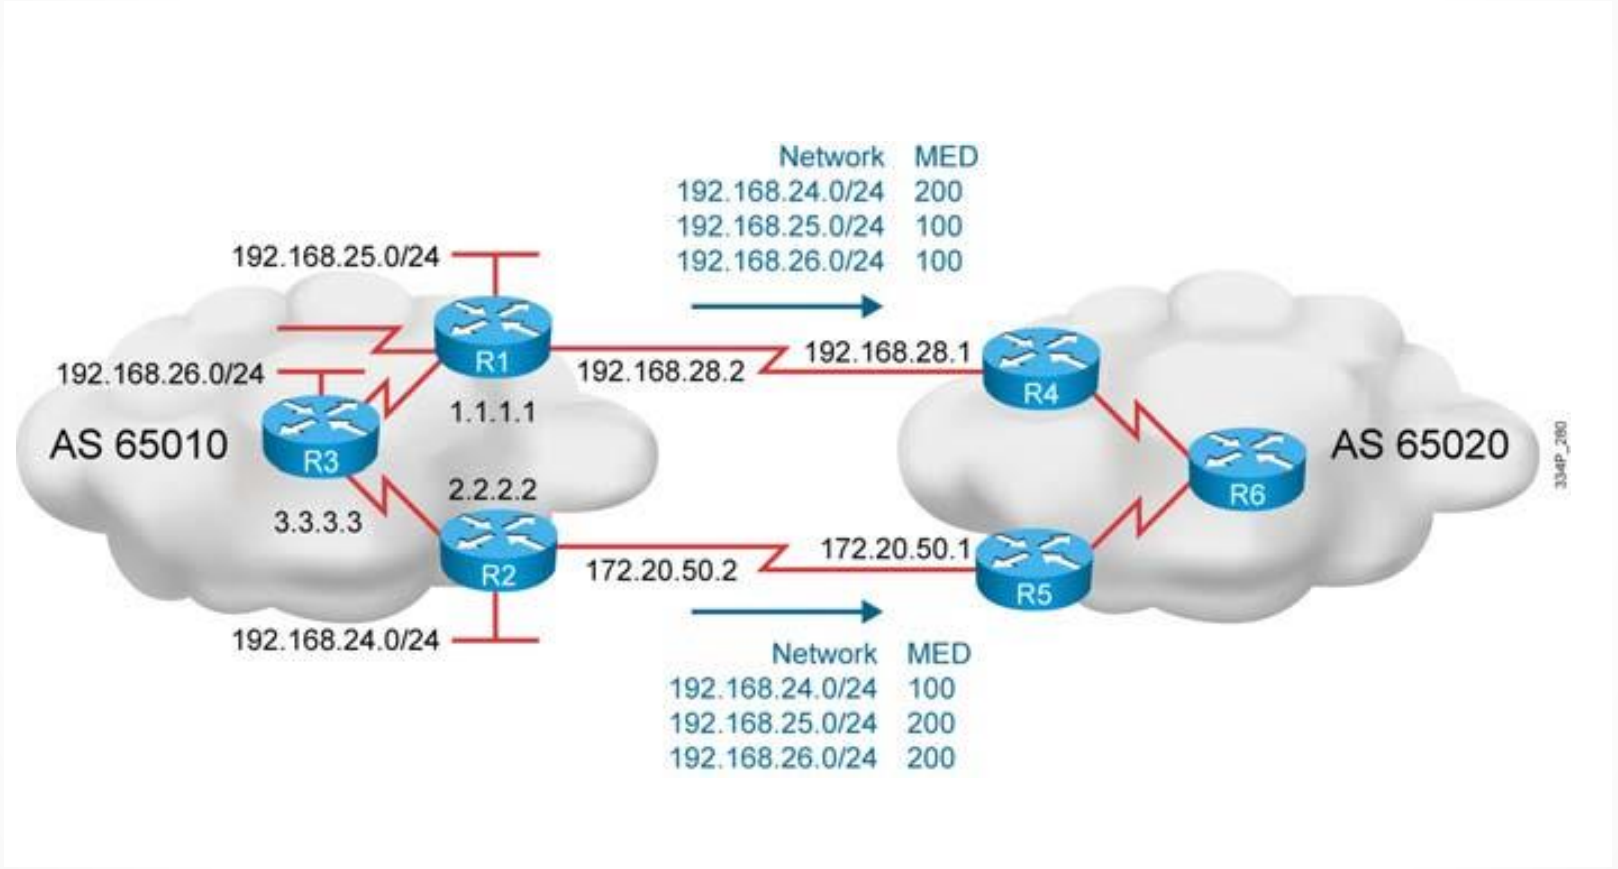
Network MED
192.168.24.0/24 200
192.168.25.0/24 100
192.168.26.0/24 100
192.168.25.0/24
192.168.26.0/24
AS 65010
R1
R3
R2
1.1.1.1
2.2.2.2
3.3.3.3
192.168.28.2
192.168.28.1
R4
R6
R5
AS 65020
172.20.50.2
172.20.50.1
192.168.24.0/24
Network MED
192.168.24.0/24 100
192.168.25.0/24 200
192.168.26.0/24 200
3348P_260

BGP - MED

R1#

```
router bgp 65010
neighbor 2.2.2.2 remote-as 65010
neighbor 3.3.3.3 remote-as 65010
neighbor 2.2.2.2 update-source loopback0
neighbor 3.3.3.3 update-source loopback0
neighbor 192.168.28.1 remote-as 65020
neighbor 192.168.28.1 route-map med_65020 out
!
access-list 66 permit 192.168.25.0.0 0.0.0.255
access-list 66 permit 192.168.26.0.0 0.0.0.255
!
route-map med_65020 permit 10
match ip address 66
set metric 100
!
route-map med_65020 permit 100
set metric 200
```

BGP - MED

R2#

```
router bgp 65010
neighbor 1.1.1.1 remote-as 65010
neighbor 3.3.3.3 remote-as 65010
neighbor 1.1.1.1 update-source loopback0
neighbor 3.3.3.3 update-source loopback0
neighbor 172.20.50.1 remote-as 65020
neighbor 172.20.50.1 route-map med_65020 out
!
access-list 66 permit 192.168.24.0.0 0.0.0.255
!
route-map med_65020 permit 10
match ip address 66
set metric 100
!
route-map med_65020 permit 100
set metric 200
```