

Hands-on Border Gateway Protocol (BGP)

Ali ALSabeh, Jorge Crichigno

University of South Carolina

<http://ce.sc.edu/cyberinfra>

aalsabeh@email.sc.edu, jcrichigno@cec.sc.edu

NSF CC* Workshop - Online

April 15, 2021

Lab Series: Border Gateway Protocol

- Lab 1: Introduction to Mininet
- Lab 2: Introduction to Free Range Routing (FRR)
- Lab 3: Introduction to BGP
- Lab 4: Configure and verify EBGP
- Lab 5: BGP Authentication
- Lab 6: Configure BGP with Default Route
- Lab 7: Using AS_PATH BGP Attribute
- Lab 8: Configuring IBGP and EBGP Sessions, Local Preference, and MED
- Lab 9: IBGP, Next Hop and Full Mesh Topology
- Lab 10: BGP Route Reflection

Organization of Lab Manuals

Each lab starts with a section *Overview*

- Objectives
- Lab topology
- Lab settings: passwords, device names
- Roadmap: organization of the lab

Section 1

- Background information of the topic being covered (e.g., fundamentals of BGP)
- Section 1 is optional (i.e., the reader can skip this section and move to lab directions)

Section 2... n

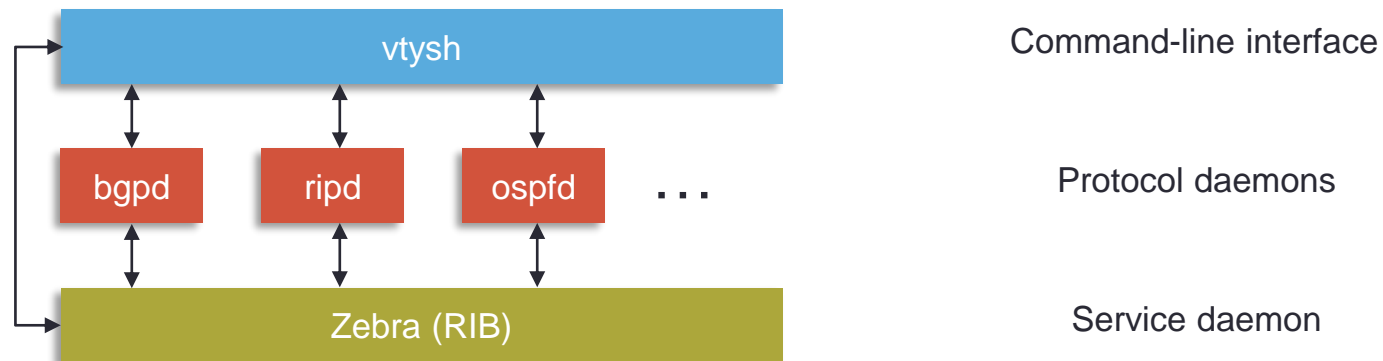
- Step-by-step directions

Introduction to FRR

Lab activities are described in Lab 2, BGP Lab Series

What is FRR?

- FRR is an open source routing protocol stack¹
- The configuration is similar to other vendors
 - Command-line shell and configuration file
- Protocols are implemented as independent processes
- Zebra is the process that controls the routing information base (RIB)

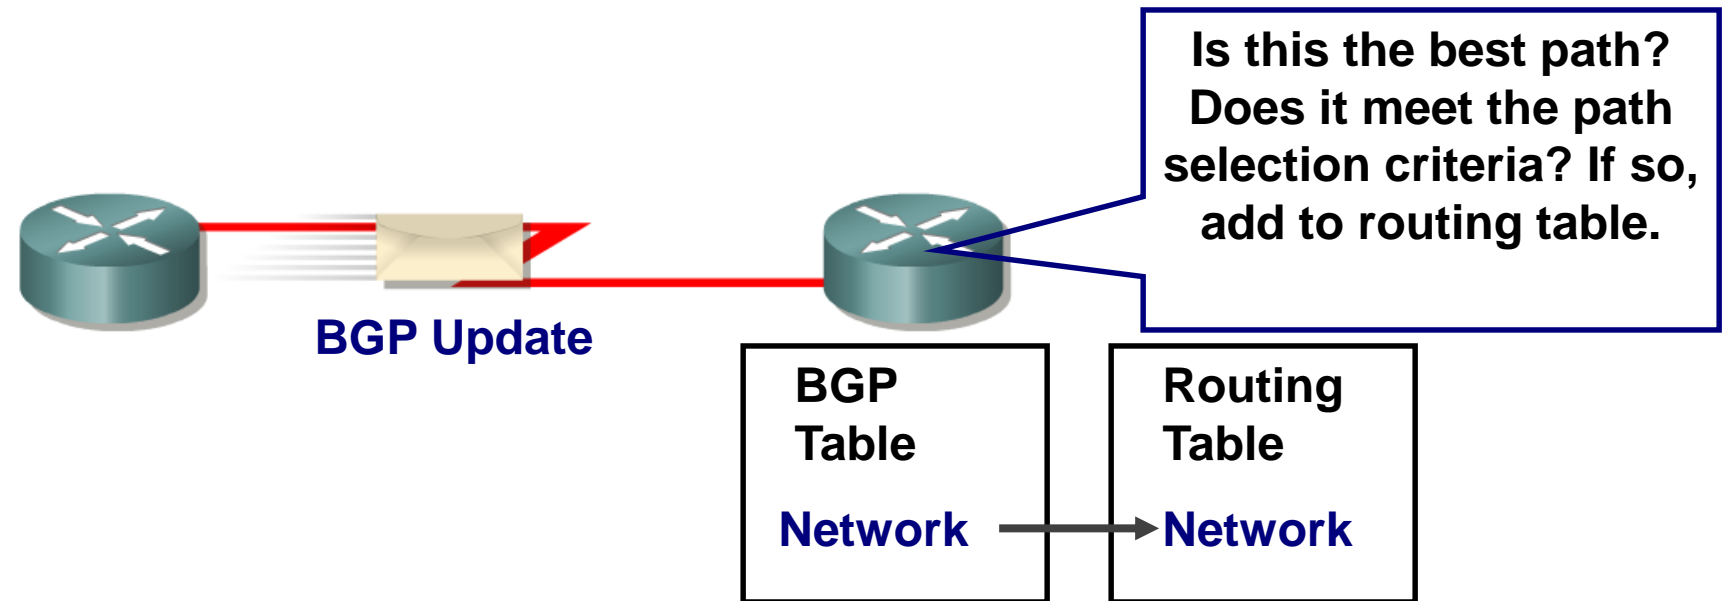


¹FRRouting website, <https://frrouting.org>

A Short Overview of BGP

BGP – Best Path

- The main goal is to provide interdomain routing
- BGP selects one path as the best path
- It places the selected path in its routing table and propagates the path to its neighbors



Slide credits: Rick Graziani, "Implementing Cisco IP Routing," Cisco Press, 2015

BGP – Best Path

1. Prefer the path with **largest weight** (configured locally; set to 0 (default) for routes not originated by this router)
2. If weights are the same, prefer the path with **largest local preference** (set to 100 by default)
3. If the local preferences are the same, prefer the **path that was originated by BGP running on this router**
4. If no route was originated, prefer the route that has the **shortest AS_path**
5. If paths have the same AS_path length, prefer the path with the **lowest origin type** (IGP is lower than EGP, and EGP is lower than Incomplete. Prefer routes introduced with the network command over redistributed)
6. If the origin codes are the same, prefer the path with the **lowest MED attribute** (set to 0 by default)
7. If the paths have the same MED, prefer the **external path over the internal path**
8. If the paths are still the same, prefer the path through the **closest IGP neighbor**
9. Prefer the path with the **lowest IP address**, as specified by the BGP router ID

- Without route manipulation, the most common reason for path selection is Step 4
- If a network administrator does not like the path with fewest ASes, he/she can manipulate weight or local preference to change which outbound path BGP selects

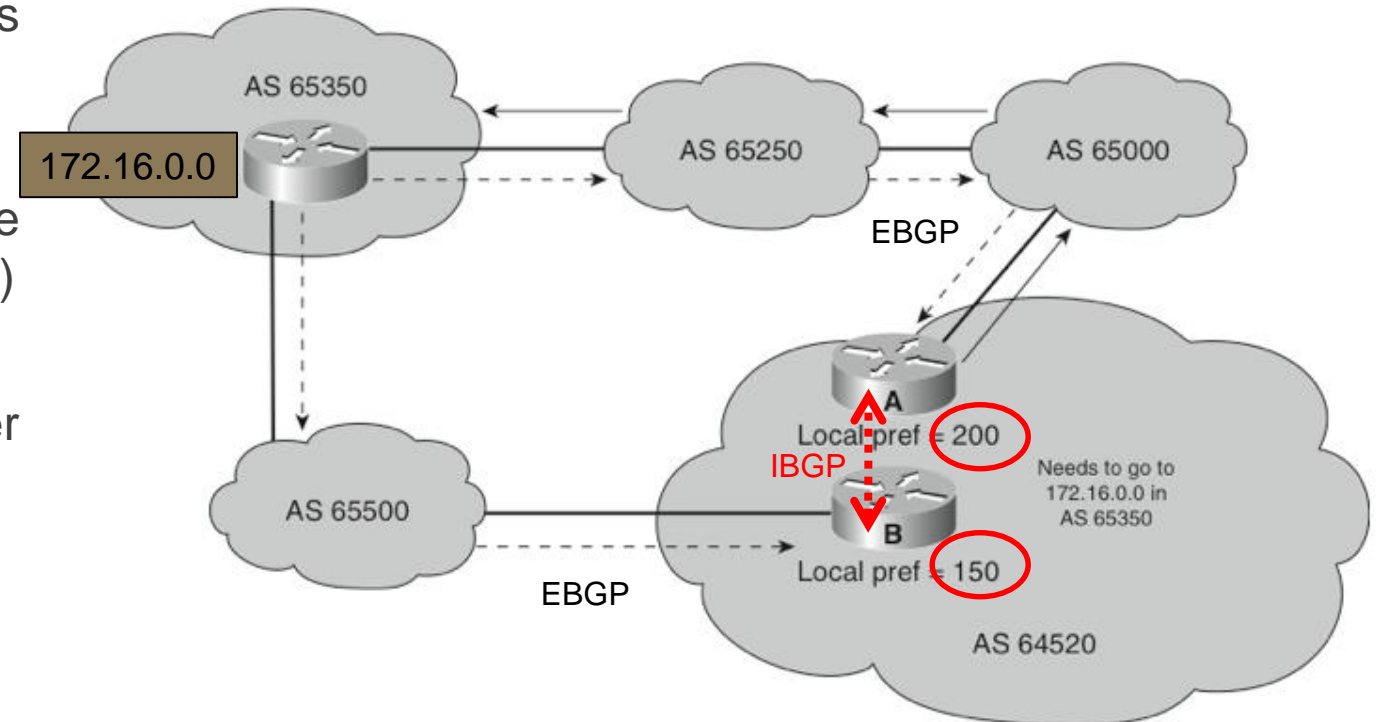
Rick Graziani, "Implementing Cisco IP Routing," Cisco Press, 2015

Configuring IBGP and EBGP Sessions, Local Preference, and MED

Lab activities are described in Lab 8, BGP Lab Series

The Local Preference Attribute

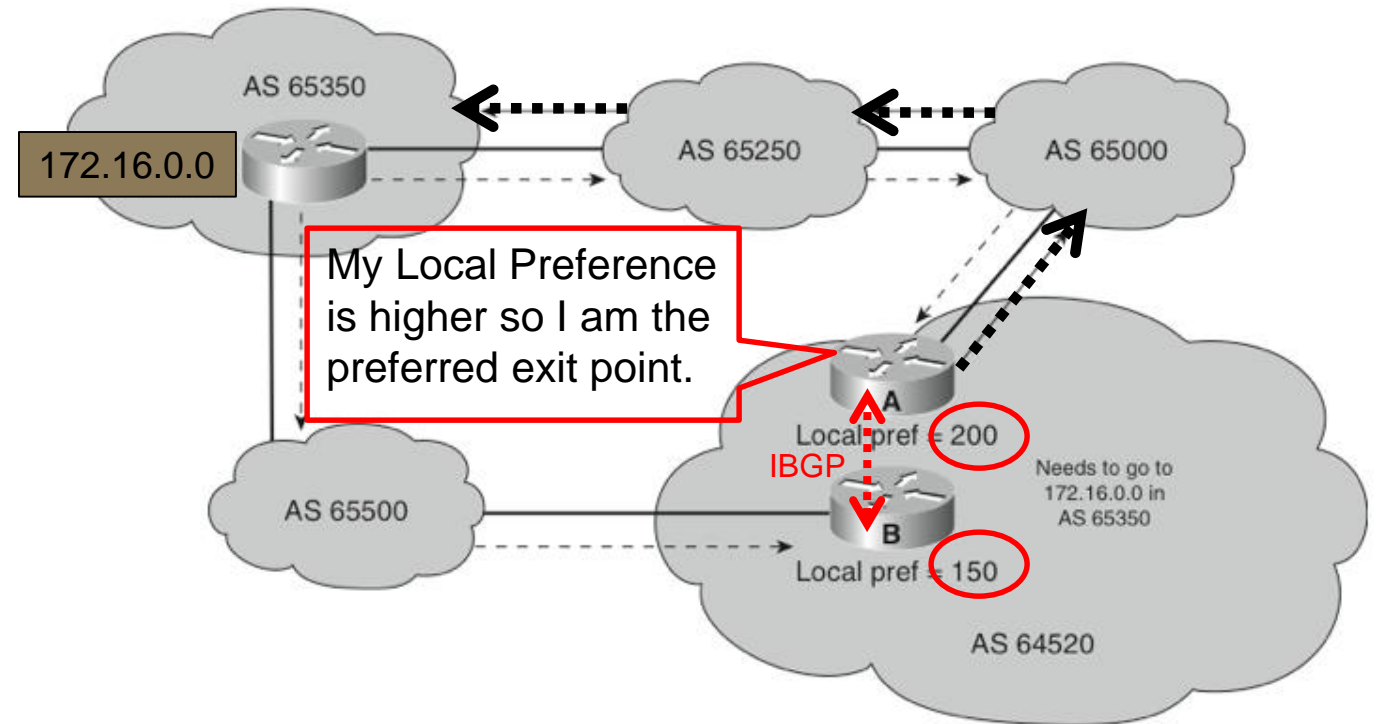
- A well-known discretionary attribute
- Indicates to routers in the AS which path is preferred **to exit the AS** (higher is better)
- Configured on a router
- Exchanged only among routers within the same AS (passed only via IBGP not via EBGP)
- Default value on a Cisco/FRR router is 100
- Local Preference takes precedence over AS_PATH



Rick Graziani, "Implementing Cisco IP Routing," Cisco Press, 2015

The Local Preference Attribute

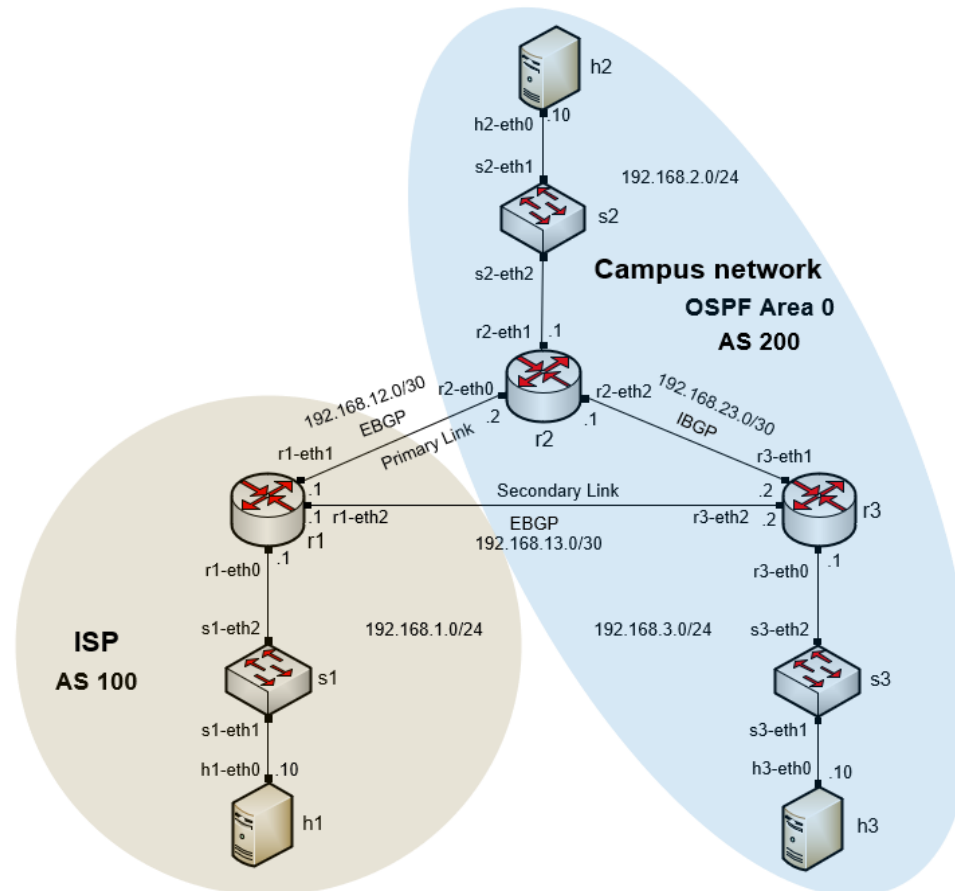
- AS 64520 receives updates about network 172.16.0.0 from two directions:
 - via AS 65500 (65500, 65350)
 - via AS 65000 (65000, 65250, 65350)
- Local preference:
 - On Router A for network 172.16.0.0 is 200
 - On Router B for network 172.16.0.0 is 150
- Local preference information is exchanged within AS 64520 via IBGP
- All traffic in AS 64520 addressed to network 172.16.0.0 is sent to Router A as an exit point from AS 64520



Rick Graziani, "Implementing Cisco IP Routing," Cisco Press, 2015

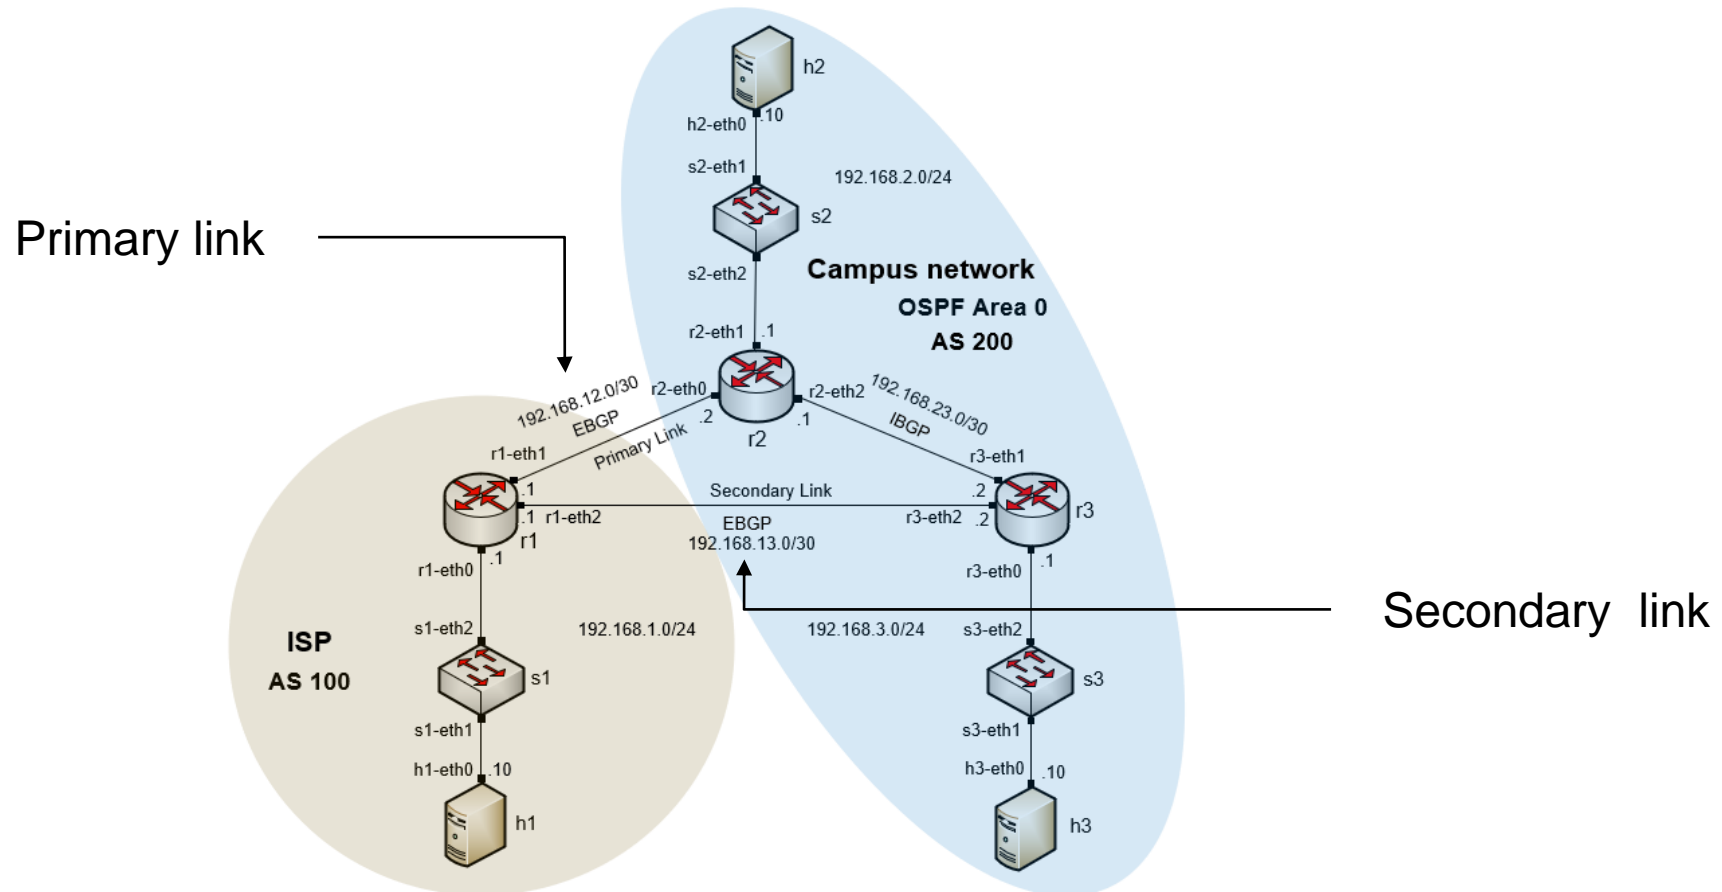
Lab 8 Topology

- Configure OSPF as Internal Gateway Protocol (IGP) in AS 200
- Configure BGP as External Gateway Protocol (EGP)



Lab 8 Topology

- Configure LOCAL_PREF and MED attributes to favor the primary link over the secondary one



BGP Table Version

- Internal version number of the table
- This number is incremented whenever the table changes

Router r3

```
frr-pc# show ip bgp
BGP table version is 3, local router ID is 192.168.23.2, vrf id 0
Default local pref 100, local AS 200
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, r RIB-failure, S Stale, R Removed
Nexthop codes: @NNN nexthop's vrf id, < announce-nh-self
Origin codes:  i - IGP, e - EGP, ? - incomplete

  Network          Next Hop             Metric LocPrf Weight Path
  *> i192.168.1.0/24 192.168.12.1         0     100     0 100 i
  *> i192.168.2.0/24 192.168.23.1         0     100     0  i
  *> 192.168.3.0/24  0.0.0.0              0                   32768  i
```

Status Code

- The status is displayed at the beginning of each line in the table

Code	Meaning
s	Table entry is suppressed
d	Table entry is dampened
h	Table entry history
*	Table entry is valid
>	Table entry is the best entry to use for this network
i	Table entry was learned via an internal BGP session
r	Table entry is a RIB-failure
S	Table entry is stale
=	Table entry has multipath to use for this network
b	Table entry has a backup path to use for this network
x	The table entry has a best external route to use for this network


Router r3

```
frr-pc# show ip bgp
BGP table version is 3, local router ID is 192.168.23.2, vrf id 0
Default local pref 100, local AS 200
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, r RIB-failure, S Stale, R Removed
Next hop codes: @NNN nexthop's vrf id, < announce-nh-self
Origin codes:  i - IGP, e - EGP, ? - incomplete

   Network        Next Hop        Metric LocPrf Weight Path
  *> i192.168.1.0/24  192.168.12.1      0     100      0 100 i
  *> i192.168.2.0/24  192.168.23.1      0     100      0  i
  *> i192.168.3.0/24  0.0.0.0           0           32768  i
```

Next-hop Attribute

- A well-known mandatory attribute
- Indicates the next-hop IP address that is to be used to reach a destination
- Unlike IGPs, BGP routes AS by AS, not router by router
- The default next-hop is the next AS
- The next-hop address for a network from another AS is an IP address of the entry point of the next AS along the path to that destination network



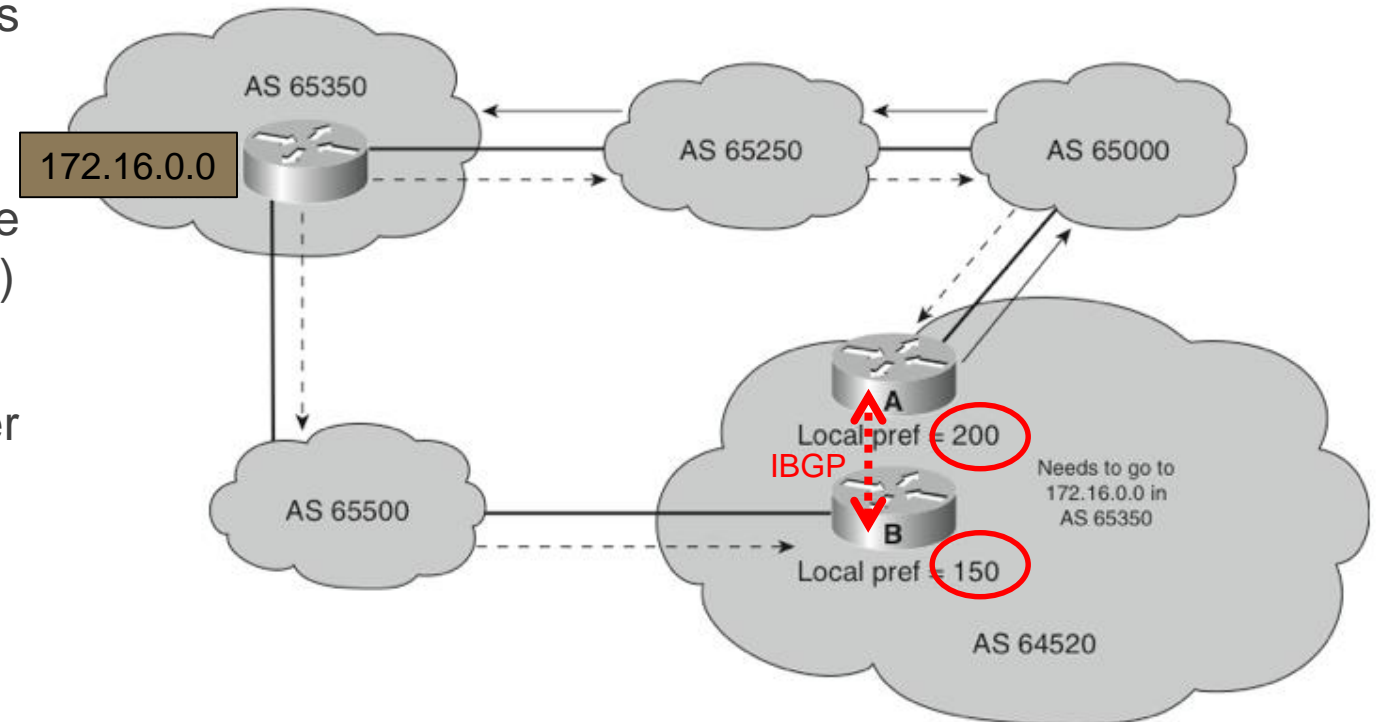
BGP table router r1

```
frr-pc# show ip bgp
BGP table version is 2, local router ID is 192.168.12.1, vrf id 0
Default local pref 100, local AS 100
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, r RIB-failure, S Stale, R Removed
Nexthop codes: @NNN nexthop's vrf id, < announce-nh-self
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop           Metric LocPrf Weight Path
*> 192.168.1.0/24  0.0.0.0            0         32768 i
*> 192.168.2.0/24 192.168.12.2       0         0 200 i
```


The Local Preference Attribute


- A well-known discretionary attribute
- Indicates to routers in the AS which path is preferred **to exit the AS** (higher is better)
- Configured on a router
- Exchanged only among routers within the same AS (passed only via IBGP not via EBGp)
- Default value on a Cisco/FRR router is 100
- Local Preference takes precedence over AS_PATH



Rick Graziani, "Implementing Cisco IP Routing," Cisco Press, 2015

The Local Preference Attribute

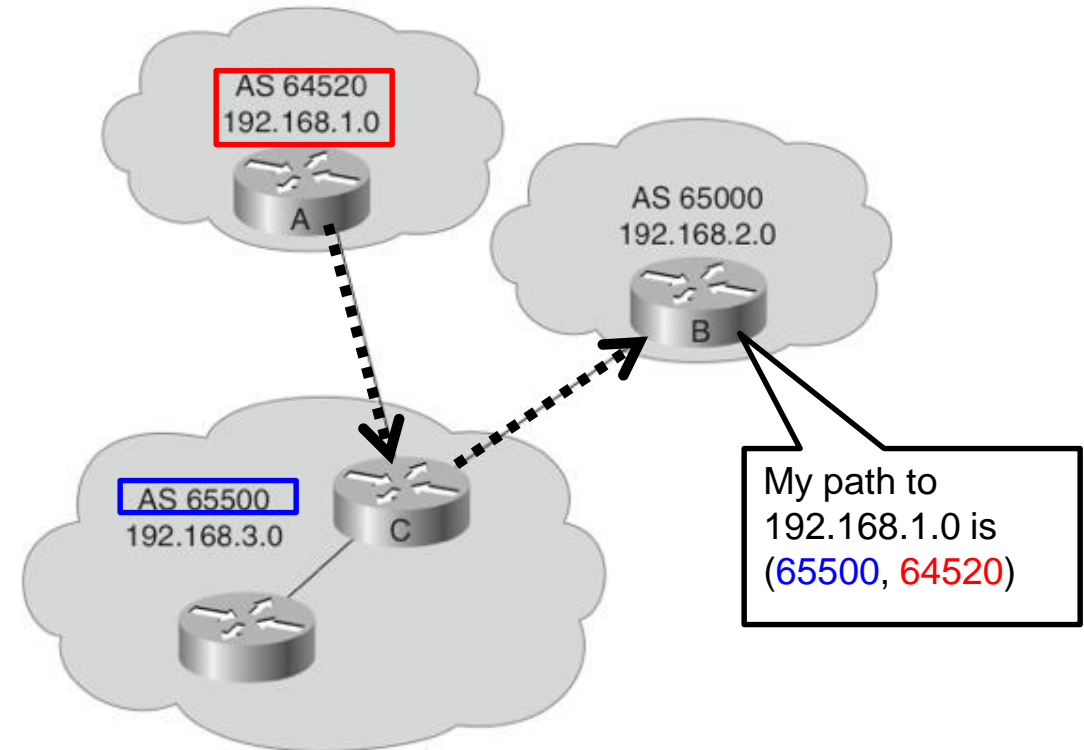
- AS 64520 receives updates about network 172.16.0.0 from two directions:
 - via AS 65500 (65500, 65350)
 - via AS 65000 (65000, 65250, 65350)
- Local preference:
 - On Router A for network 172.16.0.0 is 200
 - On Router B for network 172.16.0.0 is 150
- Local preference information is exchanged within AS 64520 via IBGP
- All traffic in AS 64520 addressed to network 172.16.0.0 is sent to Router A as an exit point from AS 64520



Rick Graziani, "Implementing Cisco IP Routing," Cisco Press, 2015

The AS-Path Attribute


- Well-known mandatory attribute
- Whenever a route update passes through an AS, the AS number is *prepend*ed to that update
- Router A: advertises network 192.168.1.0 in AS 64520.
- Router C: prepends its own AS number to it and advertises the route to Router B
- Router B: From Router B's perspective, the path to reach 192.168.1.0 is:
 - 65500, 64520



Rick Graziani, "Implementing Cisco IP Routing," Cisco Press, 2015

Origin Attribute

- A well-known mandatory attribute
- Defines the origin of the path information
- The origin attribute can be one of three values:
- **IGP (“i”)**
 - The route is interior to the originating AS
 - Normally when the **network command** is used
- **EGP (“e”)**
 - The route is learned via **EGP**
 - EGP is legacy and no longer supported
- **Incomplete (“?”)**
 - The route’s origin is unknown / some other means
 - It usually occurs when a route is **redistributed into BGP**



BGP table router r1

```
frr-pc# show ip bgp
BGP table version is 2, local router ID is 192.168.12.1, vrf id 0
Default local pref 100, local AS 100
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, r RIB-failure, S Stale, R Removed
Next hop codes: @NNN next hop's vrf id, < announce-nh-self
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network        Next Hop           Metric LocPrf Weight Path
*> 192.168.1.0/24  0.0.0.0             0         32768 i
*> 192.168.2.0/24 192.168.12.2       0         0 200 i
```

Administrative Distance

- A router may run multiple routing protocols / static routes
- if BGP and OSPF are configured on a router, both protocols may provide different best paths (analogous to google-maps and mapquest)
- How does the router know which protocol to choose?
 - The route with lower Administrative Distance is installed in the routing table

Route Source	Administrative Distance
Connected	0
Static	1
EIGRP summary route	5
External BGP	20
Internal EIGRP	90
IGRP	100
OSPF	110
IS-IS	115
RIP	120
External EIGRP	170
Internal BGP	200


Router r3

```
frr-pc# show ip route
Codes: K - kernel route, C - connected, S - static, R - RIP,
       O - OSPF, I - IS-IS, B - BGP, E - EIGRP, N - NHRP,
       T - Table, v - VNC, V - VNC-Direct, A - Babel, D - SHARP,
       F - PBR, f - OpenFabric,
       > - selected route, * - FIB route, q - queued route, r - rejected route

B>* 192.168.1.0/24 [20/0] via 192.168.13.1, r3-eth2, 00:34:48
B 192.168.2.0/24 [200/0] via 192.168.23.1, r3-eth1, 00:34:38
O>* 192.168.2.0/24 [110/20] via 192.168.23.1, r3-eth1, 00:49:27
O 192.168.3.0/24 [110/10] is directly connected, r3-eth0, 00:49:04
C>* 192.168.3.0/24 is directly connected, r3-eth0, 00:52:03
C* 192.168.13.0/30 is directly connected, r3-eth2, 00:52:03
O 192.168.23.0/30 [110/10] is directly connected, r3-eth1, 00:49:32
C>* 192.168.23.0/30 is directly connected, r3-eth1, 00:52:03
```

The Med Attribute


- An optional nontransitive attribute, also called *metric*
- Indicates to external neighbors the preferred path *into* an AS
- By default, a router compares the MED only for paths from neighbors in the **same AS**. Lowest Wins!
- MED is sent to EBGP peers:
 - Those routers propagate the MED within their AS
 - But do not pass it on to the next AS



Rick Graziani, "Implementing Cisco IP Routing," Cisco Press, 2015

The Med Attribute

- Router B has set the MED to 150 for 172.20.0.0
- Router C has set the MED to 200 for 172.20.0.0
- Router A receives EBGP updates from routers B and C
 - Chooses Router B as the best next hop to get to AS 65500 because of the lower MED



Rick Graziani, "Implementing Cisco IP Routing," Cisco Press, 2015



UNIVERSITY OF
SOUTH CAROLINA

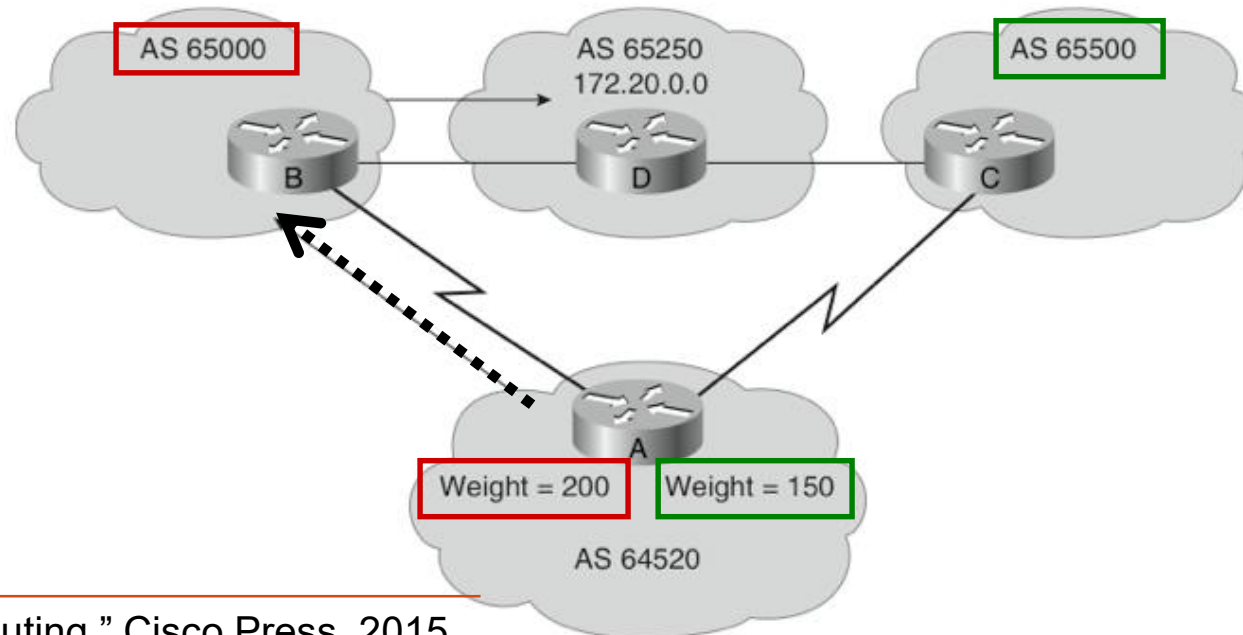
Additional Slides

The Weight Attribute

- Configured locally and ***not propagated to any other routers***
- Higher weight is preferred
- Weight takes precedence over Local Preference
- Value from 0 to 65535
- Default is 32768
- Default is 0 for routes not originated by this router

The Weight Attribute

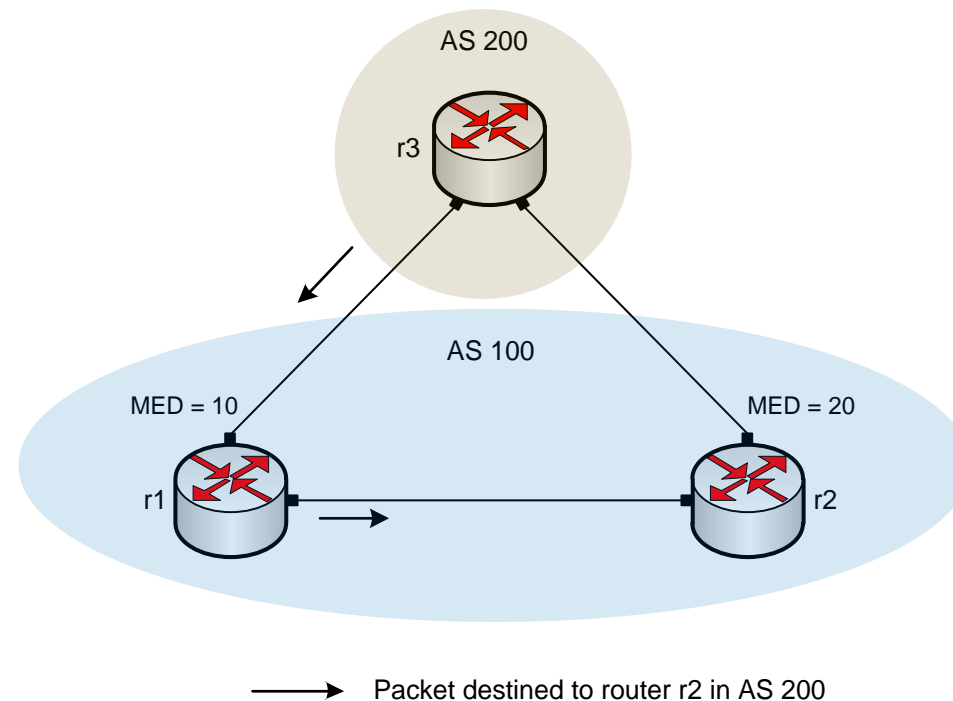
- Router A has two ways to reach 172.20.0.0
 - via Router B (AS 65000)
 - via Router C (AS 65500)
- Router A is configured to set the weight of updates coming from:
 - Router B to 200
 - Router C to 150
- Weight for Router B is higher, so Router A uses Router B as a next hop to reach 172.20.0.0



Rick Graziani, "Implementing Cisco IP Routing," Cisco Press, 2015


MED attribute

- The MED attribute indicates to external neighbors the preferred path into an AS
- It is used when there are multiple entry points to the same AS
- The BGP route with the lowest MED value is preferred



BGP Table

- Example BGP Table



BGP table router r1

```
frr-pc# show ip bgp
BGP table version is 2, local router ID is 192.168.12.1, vrf id 0
Default local pref 100, local AS 100
Status codes: s suppressed, d damped, h history, * valid, > best, = multipath,
               i internal, r RIB-failure, S Stale, R Removed
Nexthop codes: @NNN nexthop's vrf id, < announce-nh-self
Origin codes: i - IGP, e - EGP, ? - incomplete

   Network          Next Hop          Metric LocPrf Weight Path
*> 192.168.1.0/24   0.0.0.0            0         32768 i
*> 192.168.2.0/24  192.168.12.2       0         0 200 i
```

BGP attributes

- When advertising a prefix, A BGP router includes attributes
- Attributes help BGP select the best path

LOCAL_PREF attribute

- The LOCAL_PREF attribute is advertised among Internal BGP (IBGP) routers
- It indicates the preferred path to external Autonomous Systems (ASes) when an AS has multiple routes to another AS
- The BGP route with the highest LOCAL_PREF value is preferred

