



UCF / FLR Workshop on Networking Topics

Session 3: Routing using BGP Attributes



Jorge Crichigno, Ali AlSabeih
University of South Carolina
<http://ce.sc.edu/cyberinfra>

University of Central Florida (UCF)
Florida LambdaRail (FLR)
The Engagement and Performance Operations Center (EPOC)
Energy Sciences Network (ESnet)
University of South Carolina (USC)

Orlando, Florida
February 16th, 2023

BGP – Best Path

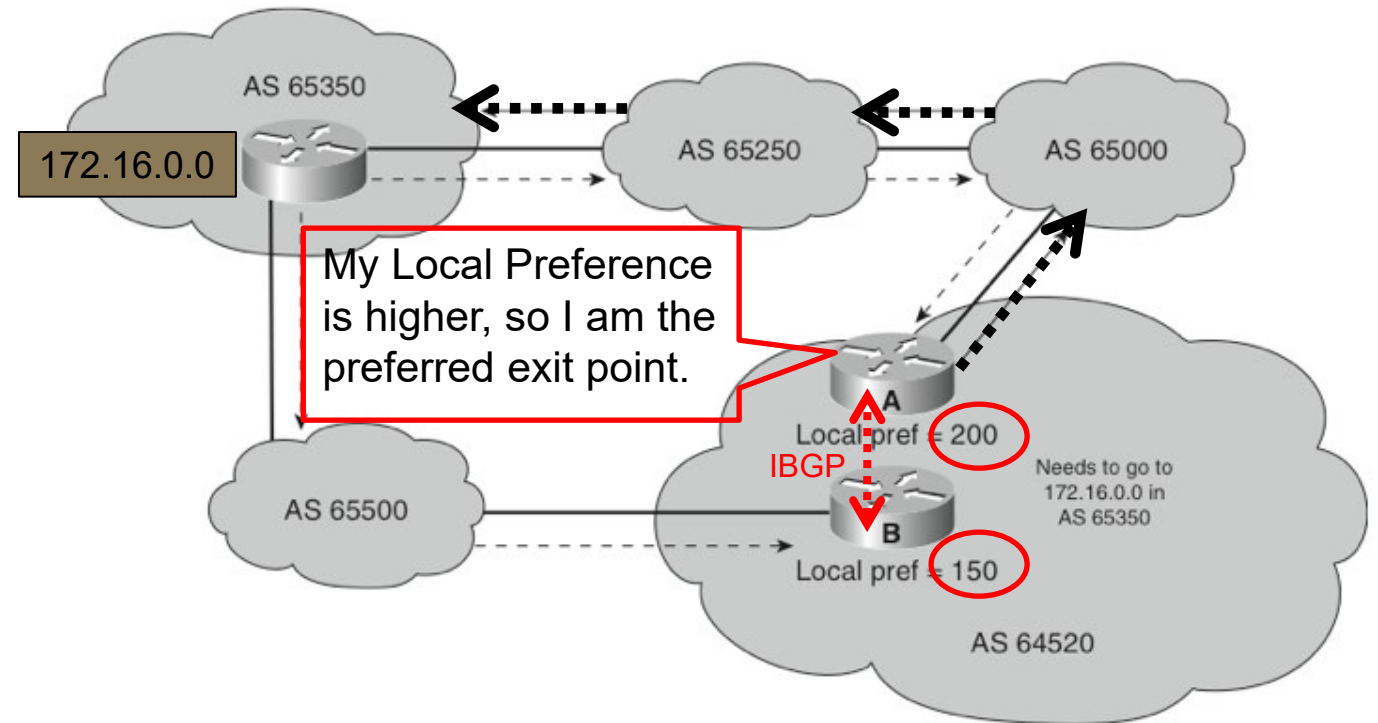
1. Prefer the path with **highest weight** (configured locally; set to 0 (default) for routes not originated by the router)
2. If weights are the same, prefer the path with **highest local preference** (set to 100 by default)
3. If the local preferences are the same, prefer the **path that was originated by BGP running on the router or redistributed from an Interior Gateway Protocol (IGP)**
4. If no route was originated, prefer the path with the **shortest AS_PATH**
5. If the paths have the same AS_PATH length, prefer the path with the **lowest origin type** (IGP is lower than Exterior Gateway Protocol (EGP), and EGP is lower than Incomplete)
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 (EBGP) over the Internal path (IBGP)**
8. If the paths are still the same, prefer the path through the **closest IGP neighbor (lowest IGP metric)**
9. If both paths are external, prefer the path that was **received first (oldest one)**
10. If the paths are still the same, prefer the path from the **BGP router with the lowest router ID**
11. If the router ID is the same for multiple paths, prefer the path with the **lowest IP address**

BGP – Best Path

- Without route manipulation, the most common reason for path selection is Step 4
 - If no route was originated, prefer the path with the **shortest AS_PATH**
- If multiple paths have the same number of autonomous systems to traverse, the second most common decision point is Step 7
 - If the paths have the same MED, prefer **EBGP over IBGP**

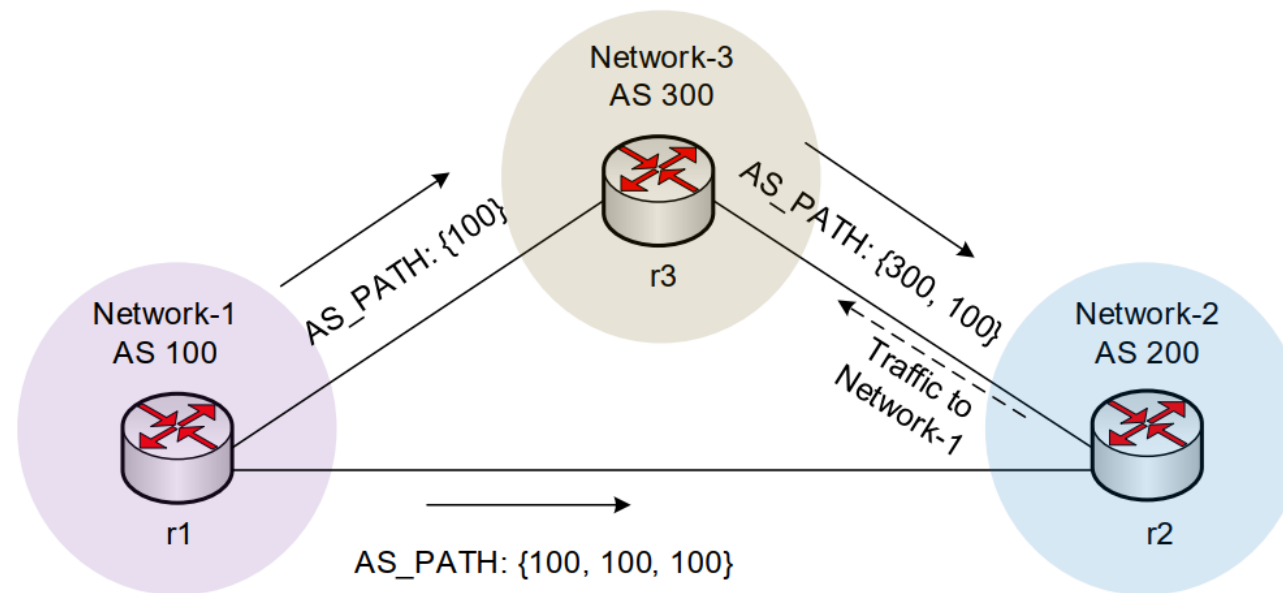
The Local Preference Attribute

- Indicates to routers in the AS which path is preferred **to exit the AS** (higher is better)
- 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



The AS_PATH Attribute

- The AS_PATH attribute includes all the ASes that need to be traversed to reach a destination
- When there are multiple paths to the same destination, BGP prefers the path with the lowest number of ASes
- Prepending is sometimes used to deprioritize a route by artificially increasing the AS-PATH length attribute
- NOTE: AS_PATH prepending in an excessive manner may lead a route to become vulnerable to disruption or misdirection; it may incentivize ASes to choose another origin if one were to suddenly appear by mistake or not¹.

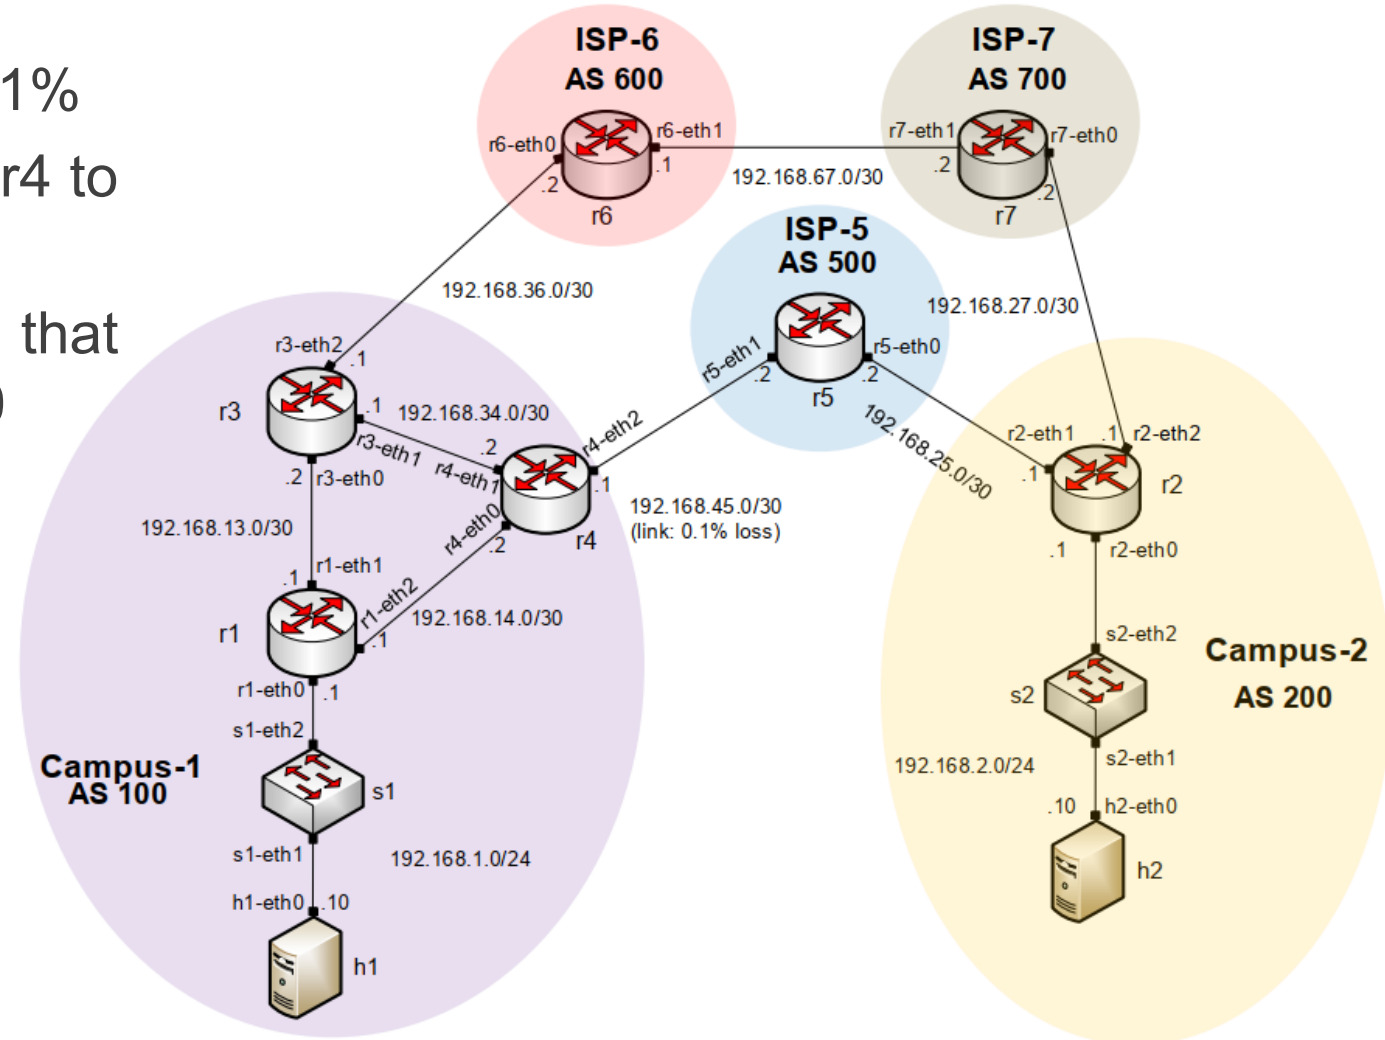


1. D. Madoury, "Excessive BGP AS-PATH prepending is a self-inflicted vulnerability," Jul. 2019. URL: <https://tinyurl.com/bdh2c86a>.

Lab 11: Configuring Local Preference and AS_PATH Prepending

Lab Topology

- All links have a capacity of 10Gbps
- Link r4-r5 has a packet loss rate = 0.1%
- Configure LOCAL_PREF on r3 and r4 to route traffic **out** of AS 100 via r3
- Set AS_PATH prepending on r4 so that traffic **into** AS 100 occurs via AS 600

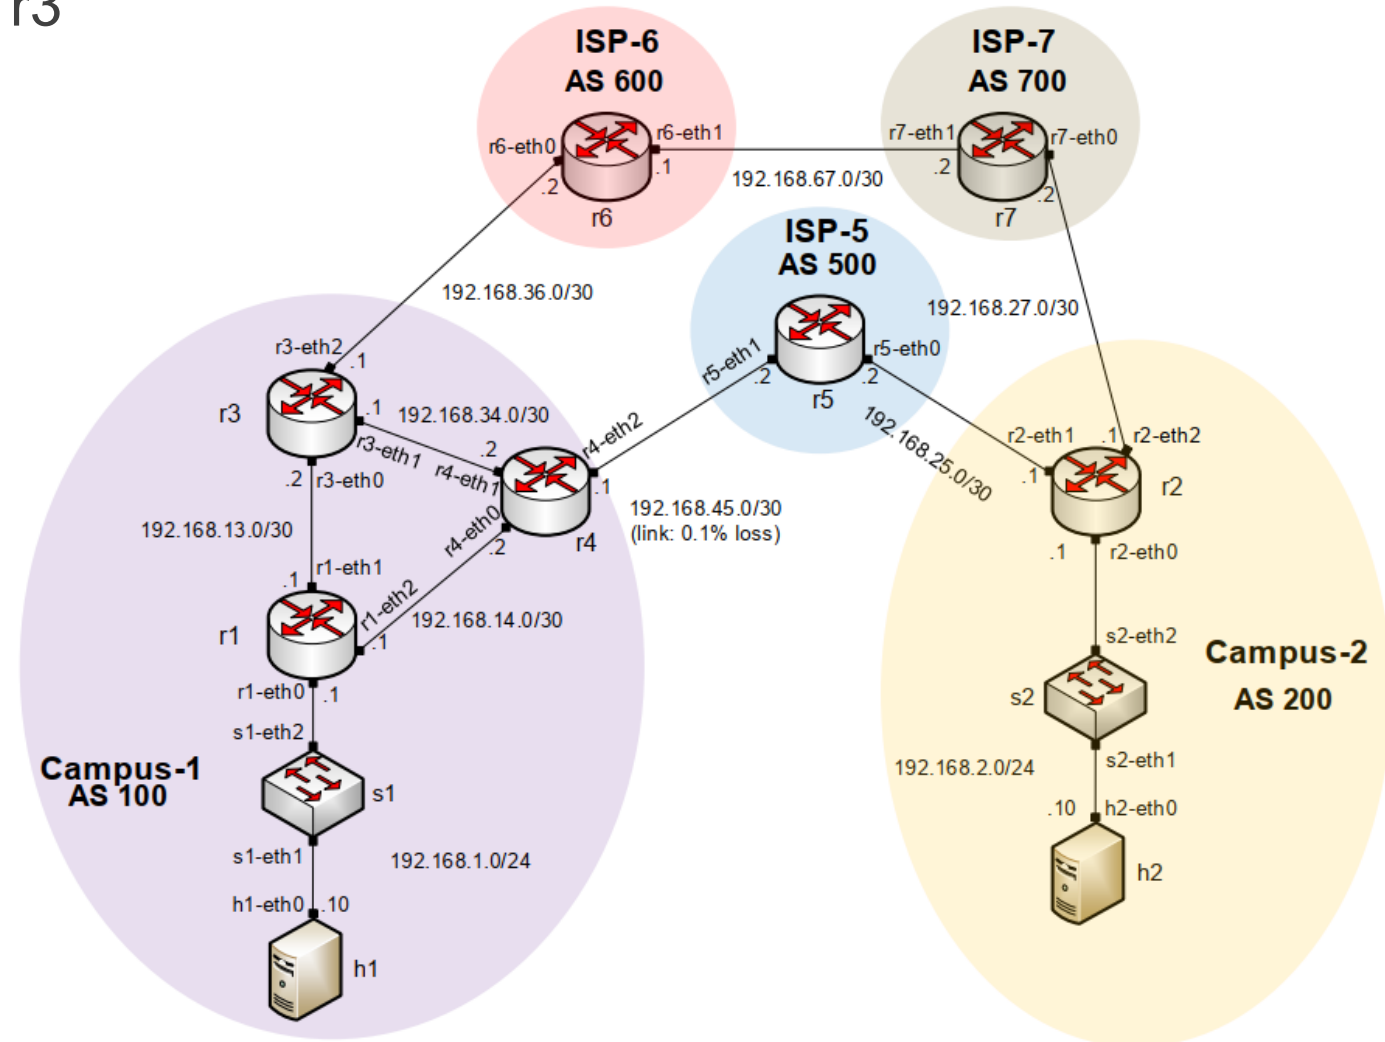


Lab Topology

- Configure LOCAL_PREF in routers r3 (150) and r4 (125) (higher is better)
- Outbound traffic via r3

```
"Host: r3"
frr-pc# configure terminal
frr-pc(config)# route-map primary_in permit 10
frr-pc(config-route-map)# set local-preference 150
frr-pc(config-route-map)# exit
frr-pc(config)# router bgp 100
frr-pc(config-router)# neighbor 192.168.36.2 route-map primary_in in
frr-pc(config-router)# end
frr-pc#
```

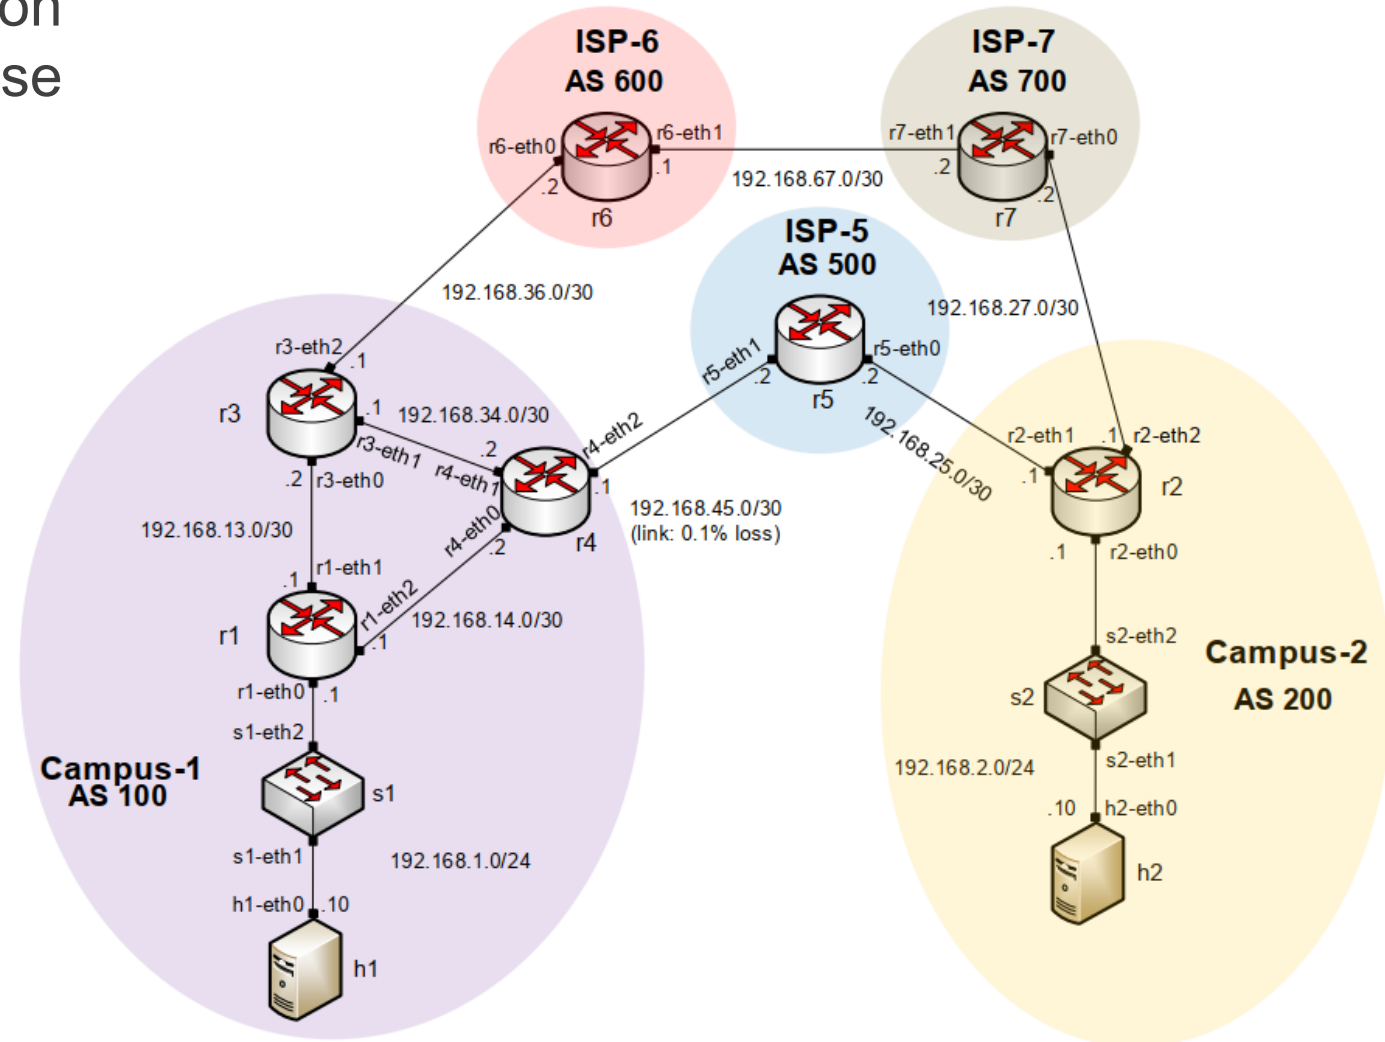
```
"Host: r4"
frr-pc# configure terminal
frr-pc(config)# route-map secondary_in permit 10
frr-pc(config-route-map)# set local-preference 125
frr-pc(config-route-map)# exit
frr-pc(config)# router bgp 100
frr-pc(config-router)# neighbor 192.168.45.2 route-map secondary_in in
frr-pc(config-router)# end
frr-pc#
```



Lab Topology

- Configure AS_PATH prepending on router r4 to influence router r2 to use the long path
- Inbound traffic from AS 200 via r3


```
"Host: r4"
frr-pc# configure terminal
frr-pc(config)# route-map prepend_out permit 10
frr-pc(config-route-map)# set as-path prepend 100 100
frr-pc(config-route-map)# exit
frr-pc(config)# router bgp 100
frr-pc(config-router)# neighbor 192.168.45.2 route-map prepend_out out
frr-pc(config-router)# end
frr-pc#
```



Additional Slides


The MED Attribute

- 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



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
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
```

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

```
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

O>* 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, 00:49:22
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 Weight Attribute - For “Outbound Route”

- Configured locally and not propagated to any other routers (higher is better)
- This is a vendor-specific attribute – Cisco supports it
- Juniper has a different mechanism to achieve a similar result
- 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 - For “Outbound Route”

- 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



BGP Table

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

```
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

- Displayed at the beginning of each line in the table

```
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

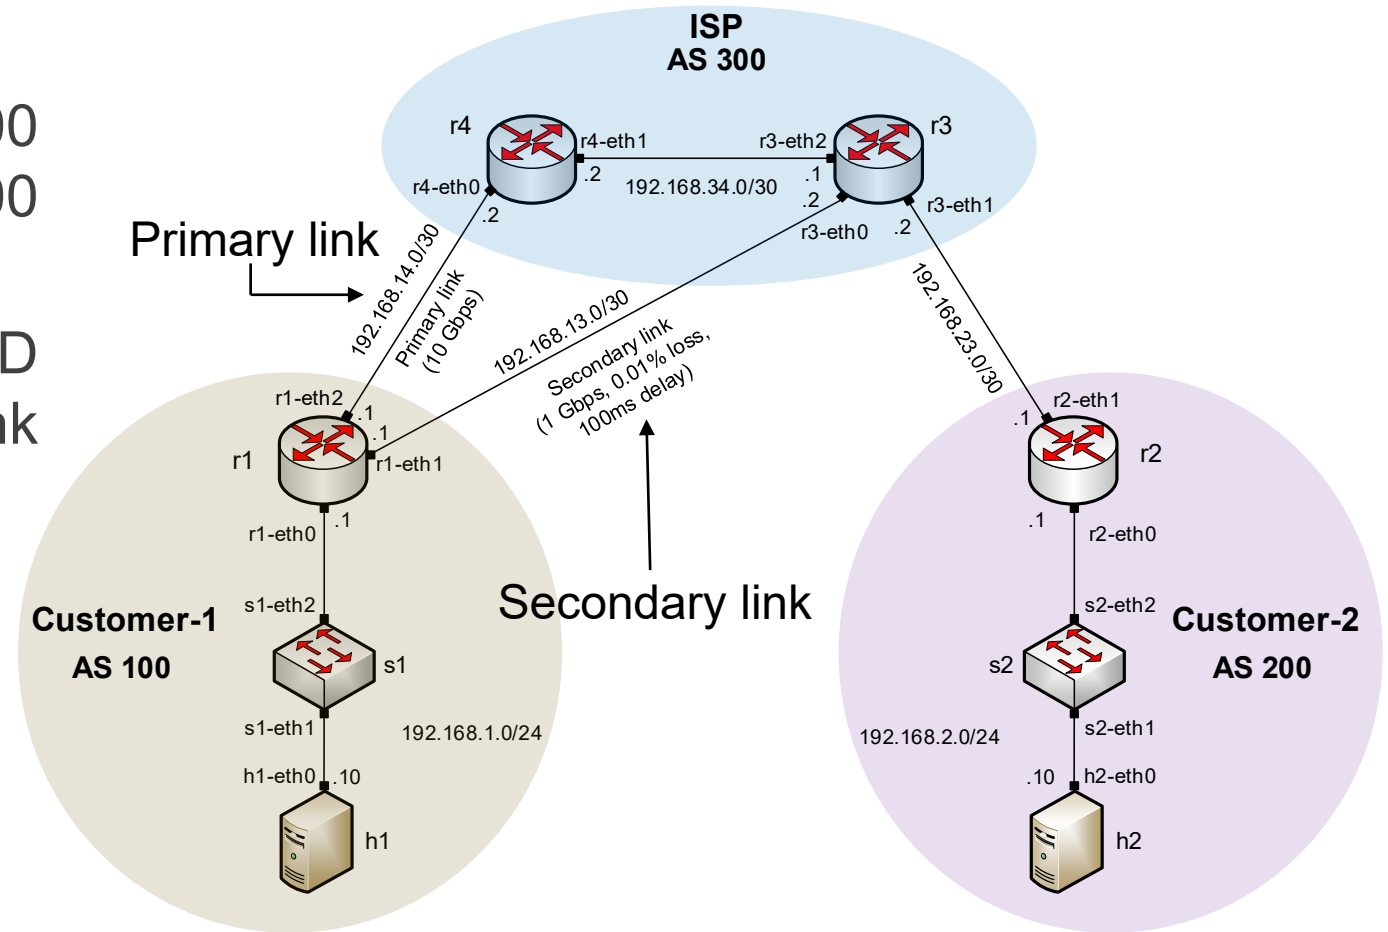
  Network        Next Hop        Metric LocPrf Weight Path
  *> 192.168.1.0/24 192.168.12.1      0     100     0 100 i
  *> 192.168.2.0/24 192.168.13.1      0           0 100 i
  *> 192.168.3.0/24 0.0.0.0           0     100     0  i
```

Code	Meaning
s	Table entry is suppressed
d	Table entry is damped
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

Lab 8.2: Configuring IBGP and EBGP Sessions, Local Preference, and MED

Lab Topology

- Configure IBGP within AS 300
- Configure EBGP between AS 100 and AS 300 and between AS 200 and 300
- Configure LOCAL_PREF and MED attributes to favor the primary link over the secondary one

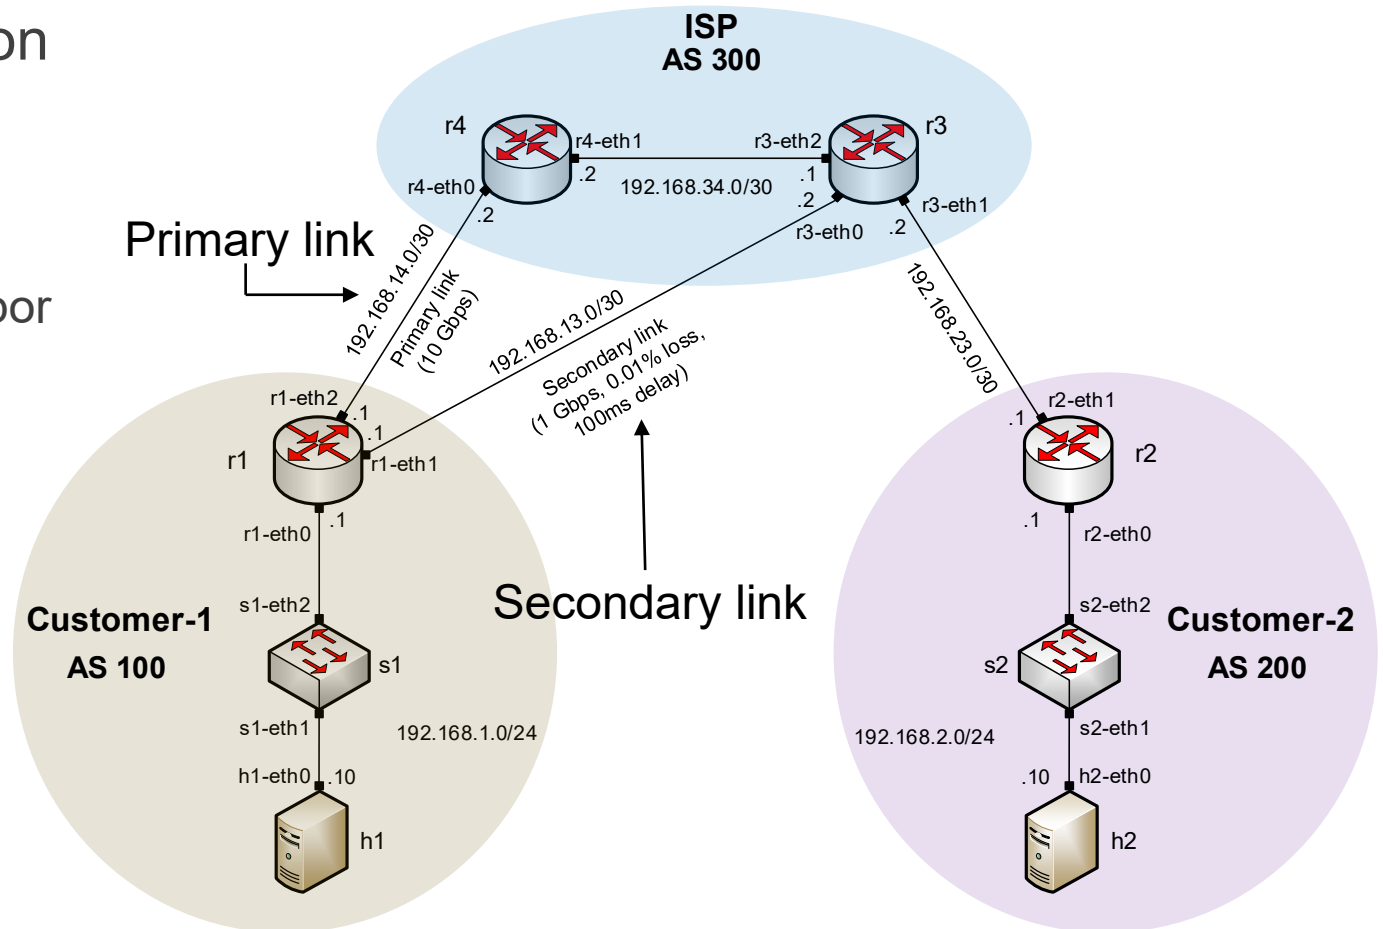


Lab Topology

- LOCAL_PREF attribute configuration on routers r3 and r4
 - Configure a route-map
 - Set the LOCAL_PREF attribute
 - Assign the route-map to the BGP neighbor

```
"Host: r3"  
frr-pc# configure terminal  
frr-pc(config)# route-map secondary_in permit 10  
frr-pc(config-route-map)#
```

```
"Host: r3"  
frr-pc# configure terminal  
frr-pc(config)# route-map secondary_in permit 10  
frr-pc(config-route-map)# set local-preference 125  
frr-pc(config-route-map)# exit  
frr-pc(config)# router bgp 300  
frr-pc(config-router)# neighbor 192.168.13.1 route-map secondary_in in  
frr-pc(config-router)#
```



Lab Topology

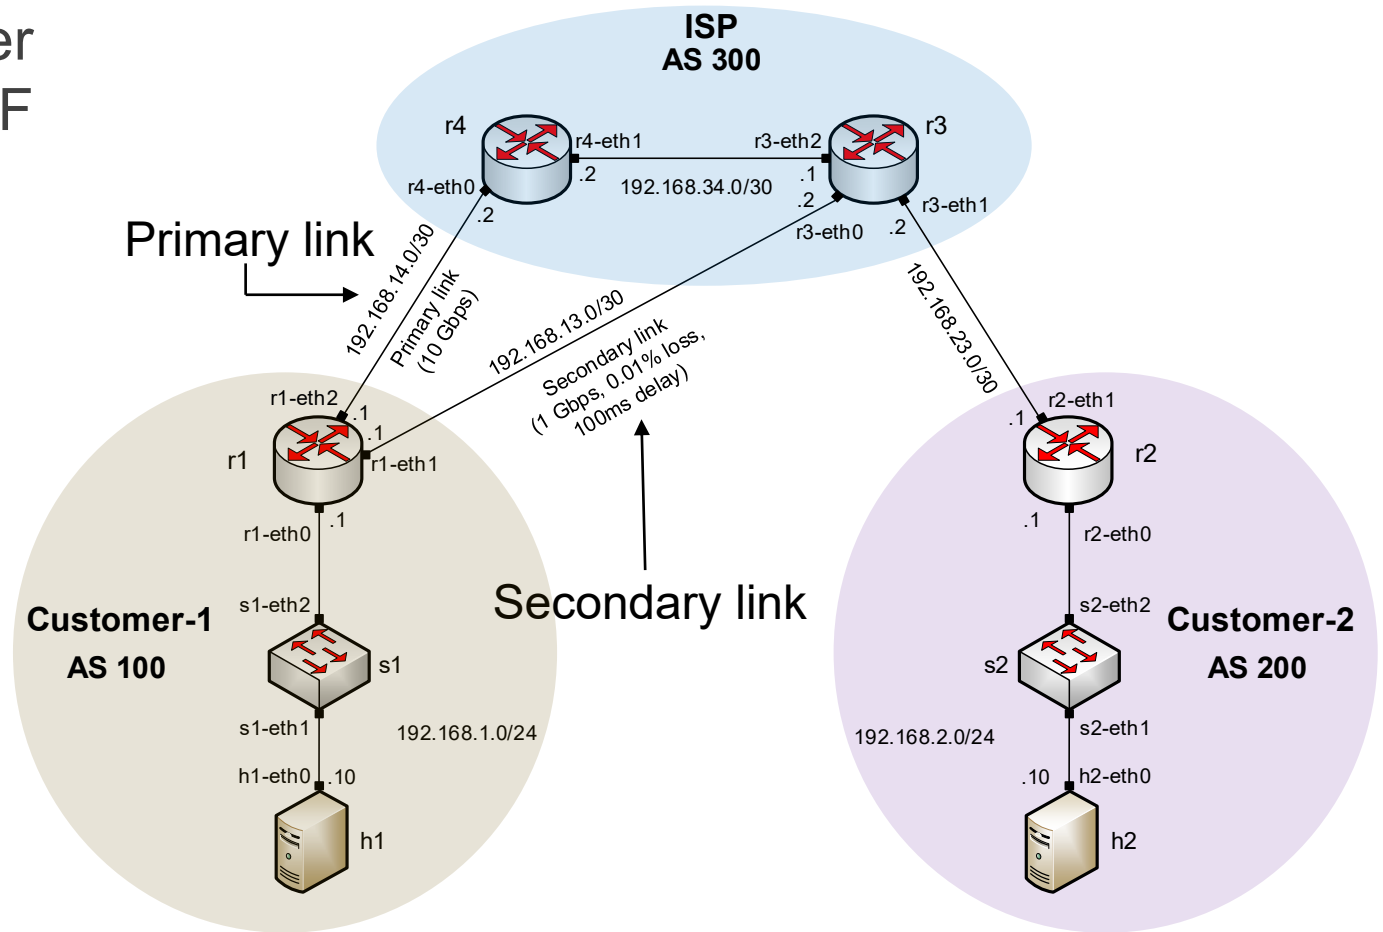
- BGP table of router r3 after configuring the LOCAL_PREF attribute

```

Host: r3
frr-pc# show ip bgp
BGP table version is 8, local router ID is 3.3.3.3, vrf id 0
Default local pref 100, local AS 300
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    192.168.14.1
*                   192.168.13.1
*> 192.168.2.0/24    192.168.23.1
*> 192.168.13.0/30   0.0.0.0
*>i192.168.14.0/30   192.168.34.2
*> 192.168.23.0/30   0.0.0.0
* i192.168.34.0/30   192.168.34.2
*>                   0.0.0.0

Displayed 6 routes and 8 total paths
frr-pc#
  
```

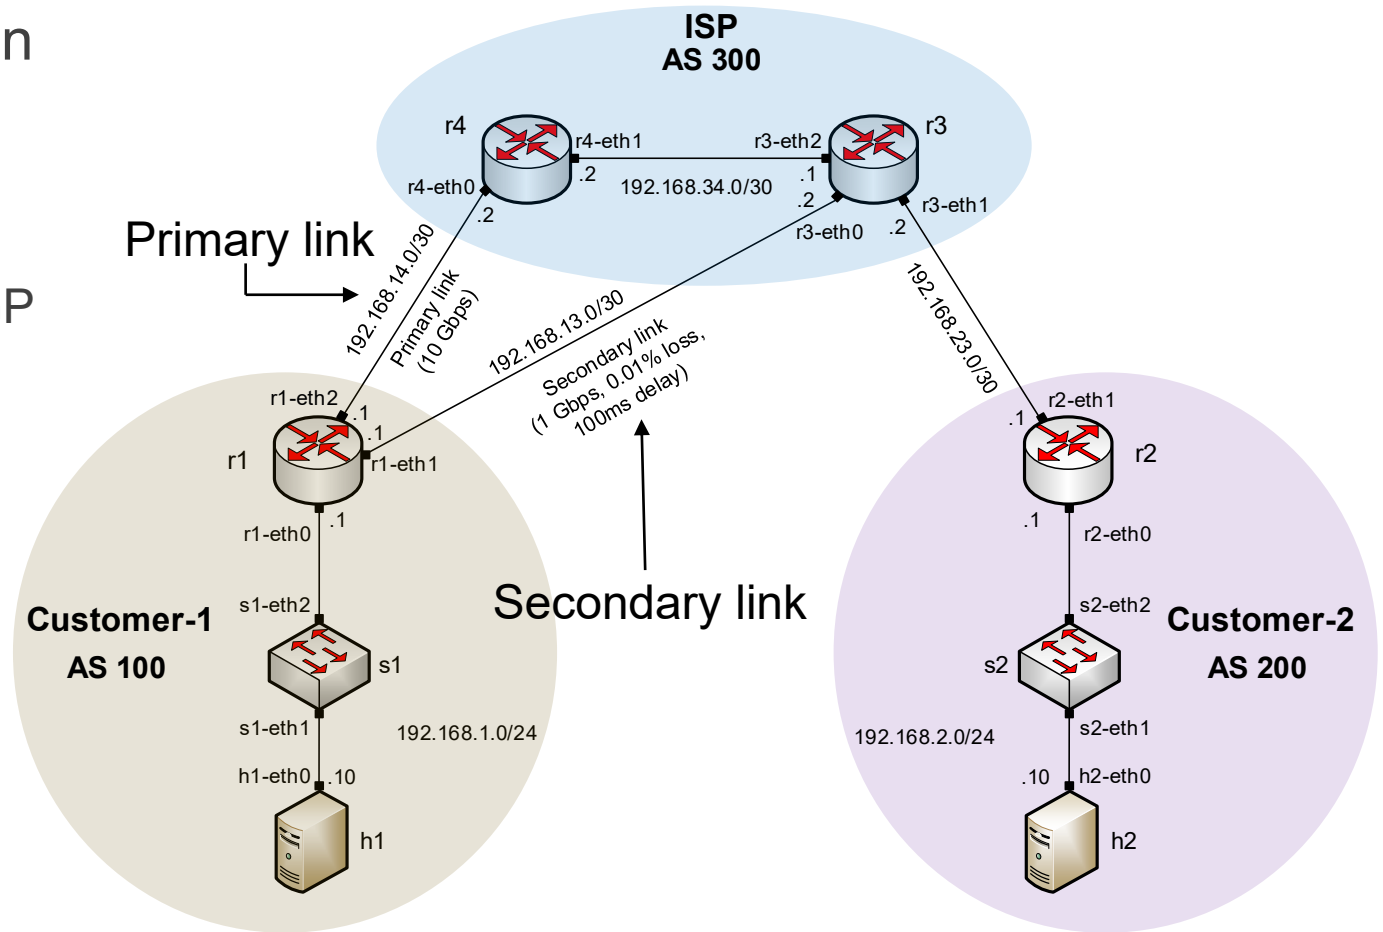


Lab Topology

- MED attribute configuration on routers r3 and r4
 - Configure a route-map
 - Set the MED attribute
 - Assign the route-map to the BGP neighbor

```
Host: r3
frr-pc# configure terminal
frr-pc(config)# route-map secondary_med_out permit 10
frr-pc(config-route-map)# set metric 75
frr-pc(config-route-map)# exit
frr-pc(config)# router bgp 300
frr-pc(config-router)# neighbor 192.168.13.1 route-map seconadry_med_out out
frr-pc(config-router)#
```

```
Host: r4
frr-pc# configure terminal
frr-pc(config)# route-map primary_med_out permit 10
frr-pc(config-route-map)# set metric 50
frr-pc(config-route-map)# exit
frr-pc(config)# router bgp 300
frr-pc(config-router)# neighbor 192.168.14.1 route-map primary_med_out out
frr-pc(config-router)# end
frr-pc#
```



Lab Topology

- BGP table of router r1 after configuring the MED attribute

```

Host: r1
frr-pc# show ip bgp
BGP table version is 15, local router ID is 1.1.1.1, 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
*> 192.168.2.0/24 192.168.14.2    50         0 300 200 i
*                192.168.13.2    75         0 300 200 i
*> 192.168.13.0/30 192.168.14.2    50         0 300 i
*                192.168.13.2    75         0 300 i
*> 192.168.14.0/30 192.168.13.2    75         0 300 i
*                192.168.14.2    50         0 300 i
*> 192.168.23.0/30 192.168.14.2    50         0 300 i
*                192.168.13.2    75         0 300 i
*> 192.168.34.0/30 192.168.14.2    50         0 300 i
*                192.168.13.2    75         0 300 i

Displayed 6 routes and 11 total paths
frr-pc#
  
```

