

“HANDS-ON SESSION BORDER GATEWAY PROTOCOL”

Ali ALSabeh, J. Crichigno
Department of Integrated Information Technology
University of South Carolina



NSF Award 1829698

“CyberTraining CIP: Cyberinfrastructure Expertise on High-throughput Networks for Big Science Data Transfers”

Lab Series: Border Gateway Protocol

- **Ali AlSabeH** is a first-year PhD student in the College of Engineering and Computing at the University of South Carolina (USC). His area of research focuses on malware analysis, Software Defined Networking (SDN), and data plane programming. Prior to joining USC he received his bachelor's degree in computer science from the Lebanese University (LU), and his master's degree from the American University of Beirut (AUB).
- **Shahrin Sharif:** is a M.Sc. student at the University of South Carolina in the United States of America. Previously she worked as a teaching assistant in the School of Engineering at American International University-Bangladesh.



LAB SERIES: BORDER GATEWAY PROTOCOL

BGP

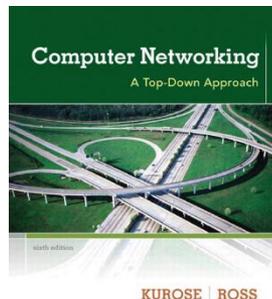
- During the “2019 Technology Exchange Conference” (December’19), it was mentioned the need for BGP training
 - A main technical event in the global R&E community for scientists, researchers, engineers, operators, and students

BGP

- During the “2019 Technology Exchange Conference” (December’19), it was mentioned the need for BGP training
- BGP is a topic taught at different levels
 - Professional certifications, two-year technical colleges, four-year universities, graduate school

BGP

- During the “2019 Technology Exchange Conference” (December’19), it was mentioned the need for BGP training
- BGP is a topic taught at different levels
 - Professional certifications, two-year technical colleges, four-year universities, graduate school
- BGP is very complex
- Even after having read books and RFCs, students (instructors) may find it difficult to fully master BGP without having practiced it
- As critical protocol for the Internet (the protocol that glues the Internet together), it is important to understand it



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
- Lab 11: Configuring Multiprotocol BGP
- Lab 12: IP Spoofing and Mitigation Techniques
- Lab 13: BGP Hijacking

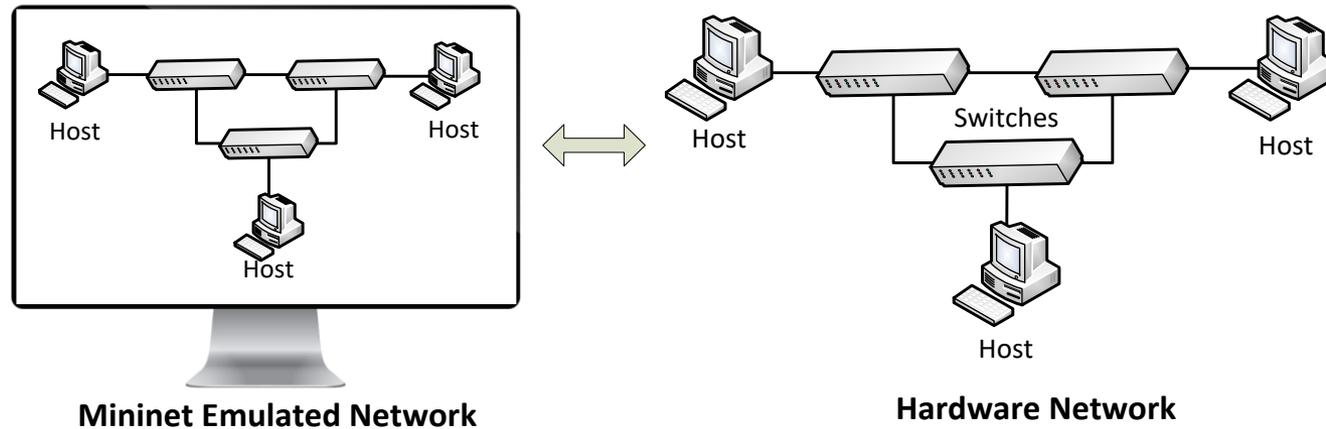
Organization of Lab Manuals

- Each lab starts with a section *Overview*
 - Objectives
 - 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

LAB 1: INTRODUCTION TO MININET

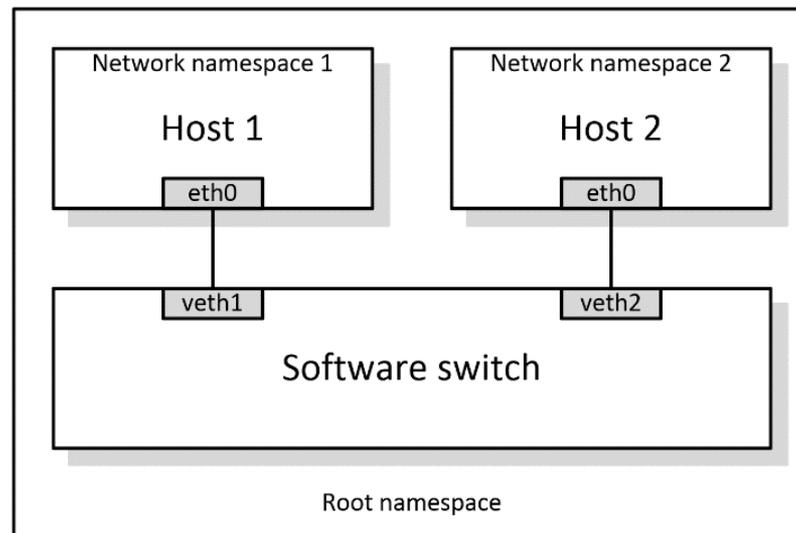
What is Mininet?

- A virtual testbed capable of recreating realistic scenarios
- It enables the development, testing of network protocols
- Inexpensive solution, real protocol stack, reasonably accurate



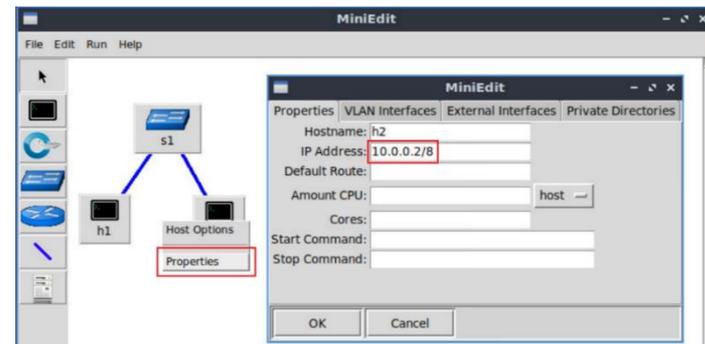
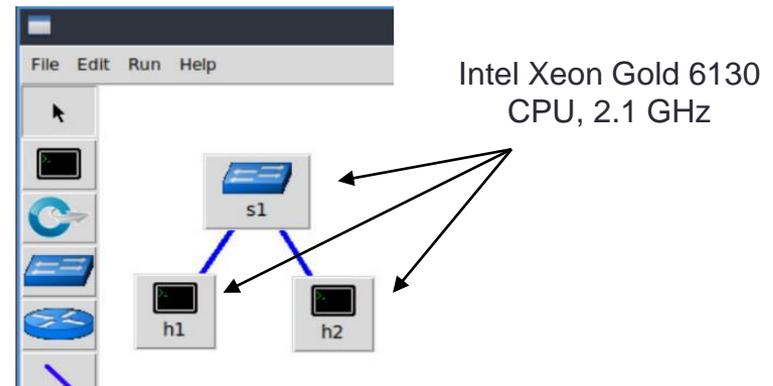
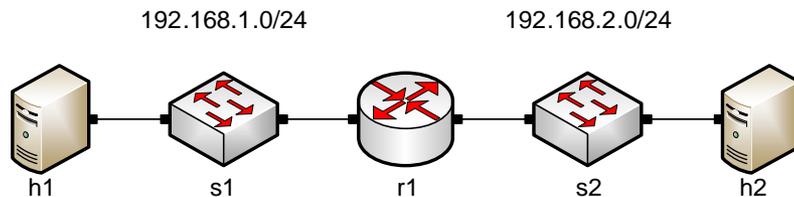
What is Mininet?

- Mininet nodes are network namespaces
 - Each node has different / separate virtual interfaces, routing tables
- Nodes use the underlying protocol stack of the host device
- Nodes are connected via virtual Ethernet (veth) links, which behave as Ethernet links



Pod Design with Mininet

- A pod is a set of virtual equipment (routers, switches, etc.)
- A pod is created every time a lab reservation is made
- For the Border Gateway Protocol Lab series, pods are embedded into Mininet



LAB 2: INTRODUCTION TO FRR

What is FRR?

- FRR is an open source routing protocol stack¹
- It provides similar routing capabilities to other vendors
- It runs natively on Linux and other platforms

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

What is FRR?

- The configuration is similar to other vendors
 - Command-line shell and configuration file

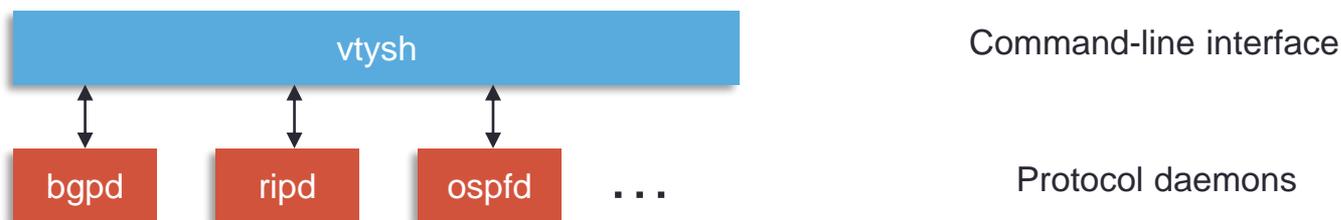


vtysh

Command-line interface

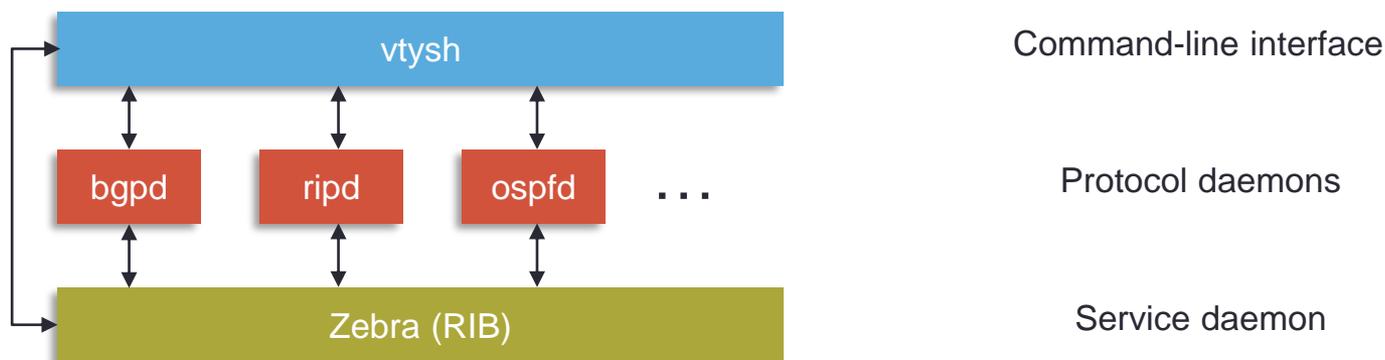
What is FRR?

- The configuration is similar to other vendors
 - Command-line shell and configuration file
- Protocols are implemented as independent processes



What is FRR?

- 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)



FRR and Mininet Integration

- Mininet provides network emulation, allowing all network software at any layer to be simply run as is
- The set of commands provided by FRR are inherited and can be run using Mininet's command-line interface

```
Host: r2
root@frr-pc:/etc/routers/r2# zebra
root@frr-pc:/etc/routers/r2# staticd
root@frr-pc:/etc/routers/r2# ripd
root@frr-pc:/etc/routers/r2# ospfd
root@frr-pc:/etc/routers/r2# bgpd
root@frr-pc:/etc/routers/r2# vtysh

Hello, this is FRRouting (version 7.2-dev).
Copyright 1996-2005 Kunihiro Ishiguro, et al.

frr-pc#
```



Lab 2 Topology

- Two hosts: h1 and h2; two switches: s1 and s2; two routers: r1 and r2
- Static routing is configured on both routers so that the hosts can ping

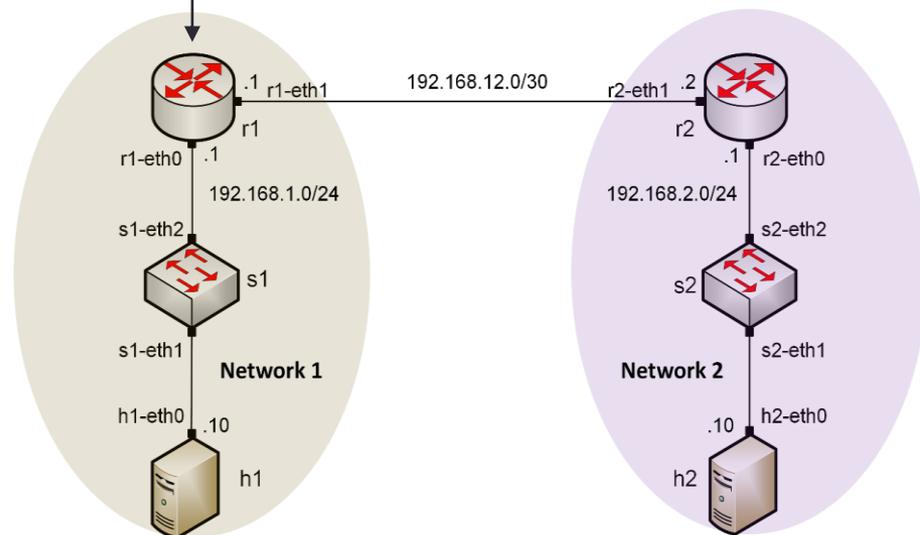
Router r1

```

"Host: r1"
root@frr-pc:/etc/routers/r1# zebra
root@frr-pc:/etc/routers/r1# staticd
root@frr-pc:/etc/routers/r1# vtysh

Hello, this is FRRouting (version 7.4-dev).
Copyright 1996-2005 Kunihiro Ishiguro, et al.

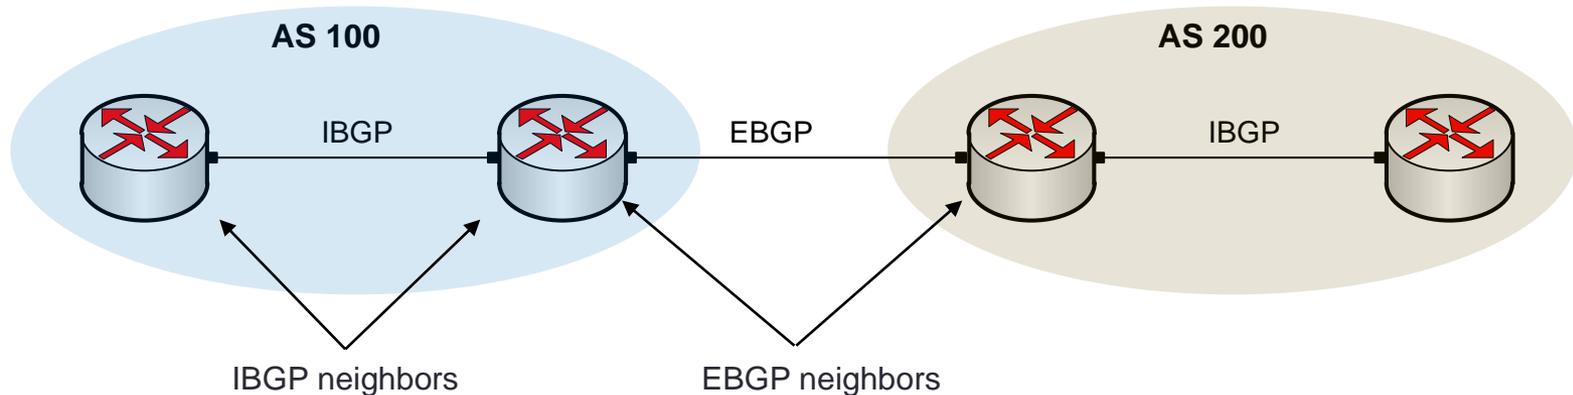
frr-pc# configure terminal
frr-pc(config)# ip route 192.168.2.0/24 192.168.12.2
frr-pc(config)#
  
```



LAB 3: INTRODUCTION TO BGP

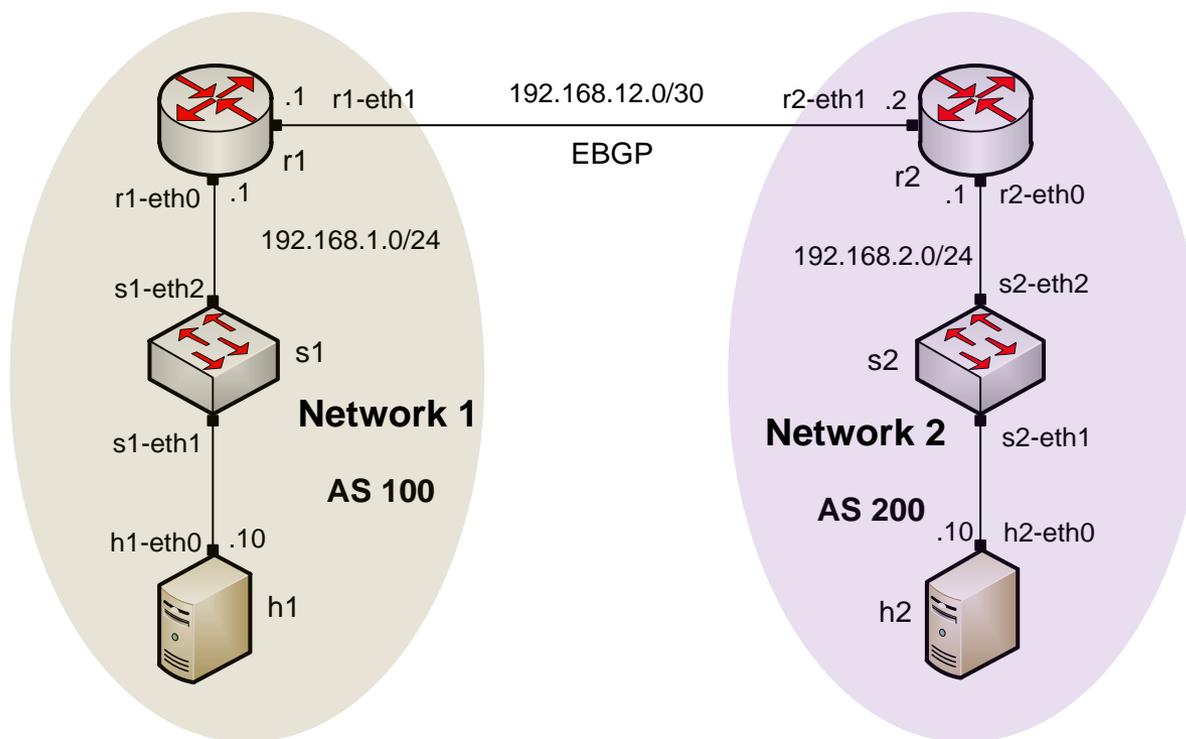
What is BGP?

- The Internet can be viewed as a collection of networks or Autonomous Systems (ASes) that are interconnected
- BGP is an exterior gateway protocol designed to exchange routing and reachability information among ASes on the Internet



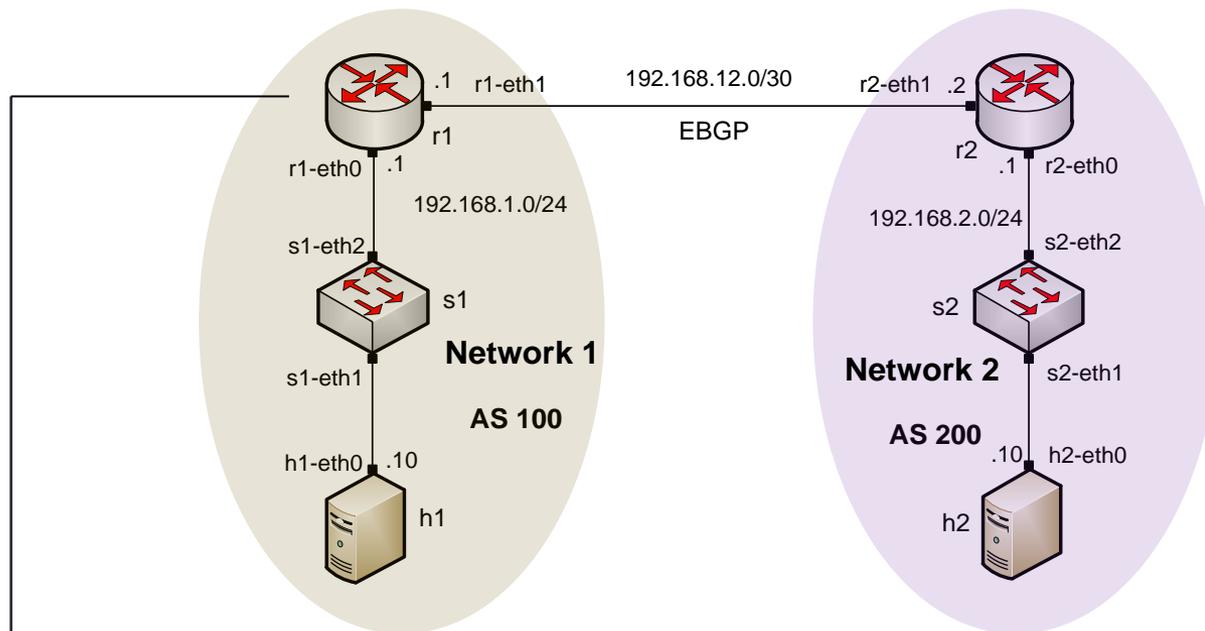
Lab 3 Topology

- Network 1 in AS 100 consists of a host, a switch, and a router
- Network 2 in AS 200 consists of a host, a switch, and a router



Lab 3 Configuration

- Establishing BGP neighborhood



Router r1

```

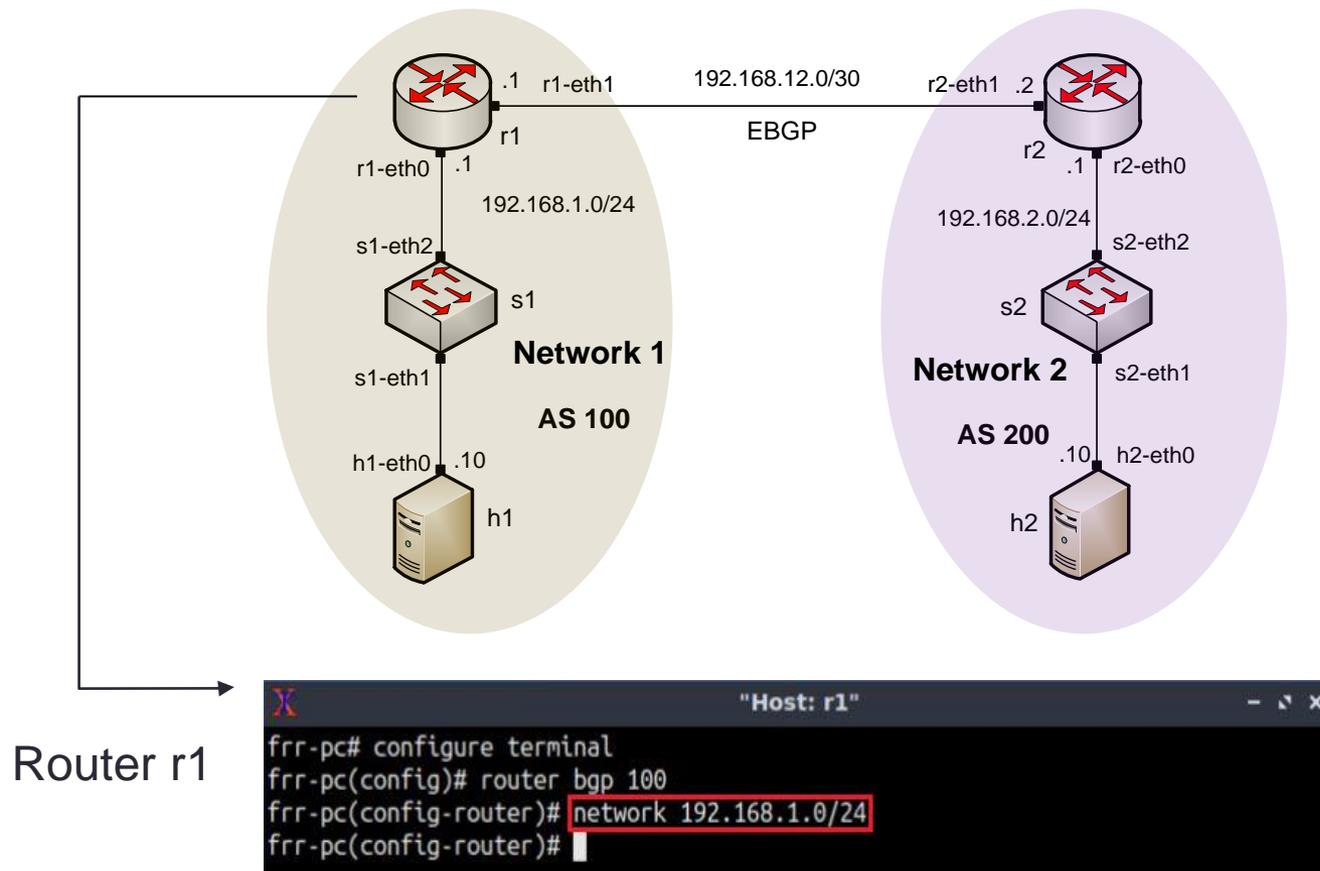
Host: r1
root@frr-pc:/etc/routers/r1# bgpd
root@frr-pc:/etc/routers/r1# vtysh

Hello, this is FRRouting (version 7.2-dev).
Copyright 1996-2005 Kunihiro Ishiguro, et al.

frr-pc# configure terminal
frr-pc(config)# router bgp 100
frr-pc(config-router)# neighbor 192.168.12.2 remote-as 200
frr-pc(config-router)#
  
```

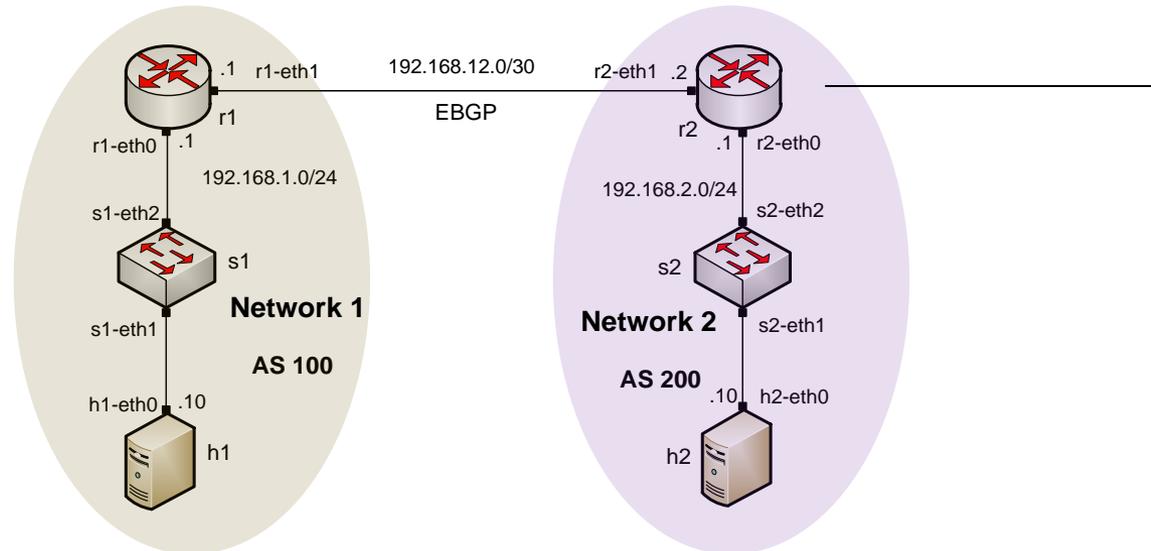
Lab 3 Configuration

- Advertising a network in BGP



Lab 3 Configuration

- Routing table: lists the routes learned from different routing protocols



```

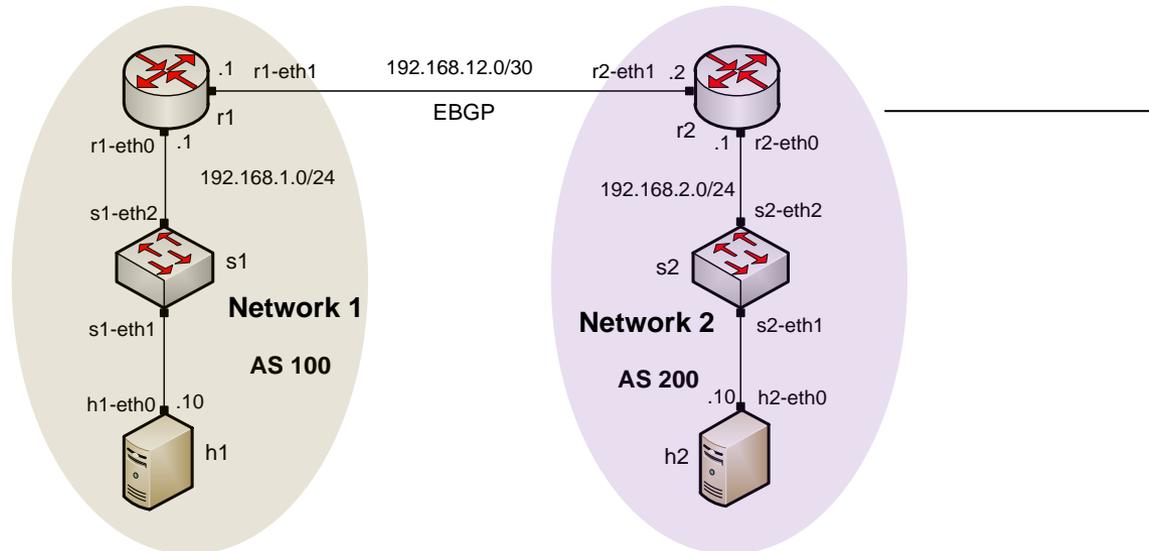
Host: r2
frr-pc# show ip route
Codes: K - kernel route, C - connected, S - static, R - RIP,
       O - OSPF, I - IS-IS, B - BGP, E - EIGRP, N - NHRP,
       T - Table, v - VNC, V - VNC-Direct, A - Babel, D - SHARP,
       F - PBR, f - OpenFabric,
       > - selected route, * - FIB route, q - queued route, r - rejected route

B>* 192.168.1.0/24 [20/0] via 192.168.12.1, r2-eth1, 00:00:52
C>* 192.168.2.0/24 is directly connected, r2-eth0, 00:18:36
C>* 192.168.12.0/30 is directly connected, r2-eth1, 00:18:02
frr-pc#
  
```

Router r2

Lab 3 Configuration

- BGP table: lists the routes learned from BGP routing protocol



```

Host: r2
frr-pc# show ip bgp
BGP table version is 2, local router ID is 192.168.12.2, 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         0 100 i
*> 192.168.2.0/24  0.0.0.0         0         32768 i

Displayed 2 routes and 2 total paths
frr-pc#
  
```

Router r2

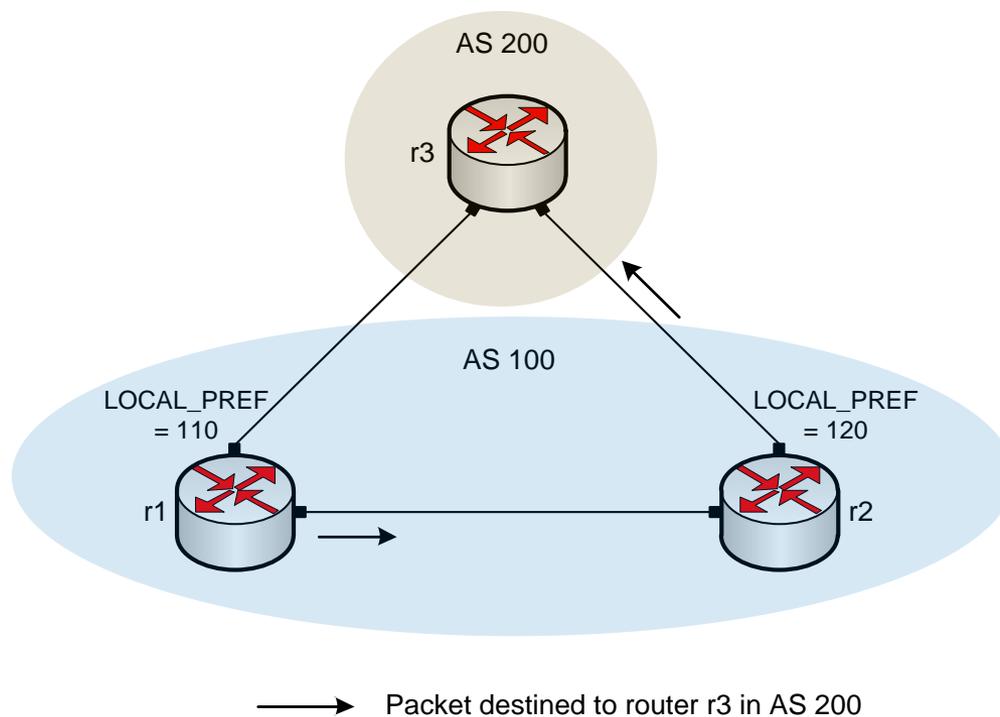
LAB 8: CONFIGURING IBGP AND EBGP SESSIONS, LOCAL PREFERENCE, AND MED

BGP attributes

- A BGP router includes a number of BGP attributes when advertising a prefix across a BGP session
- These attributes help BGP select the best path when there are multiple paths to the same destination

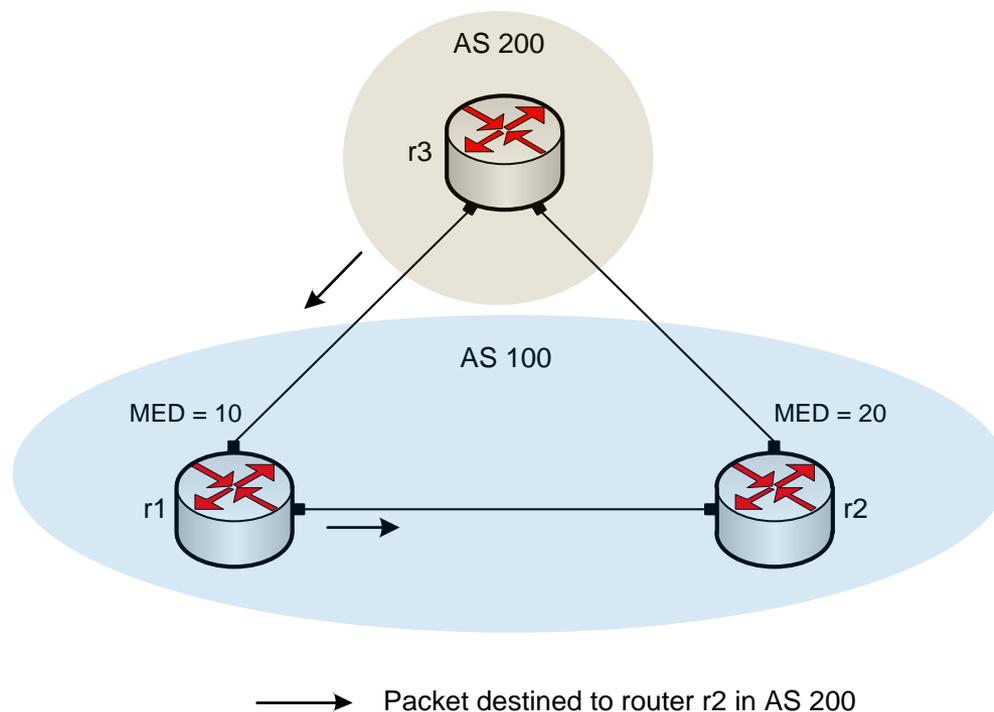
LOCAL_PREF attribute

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



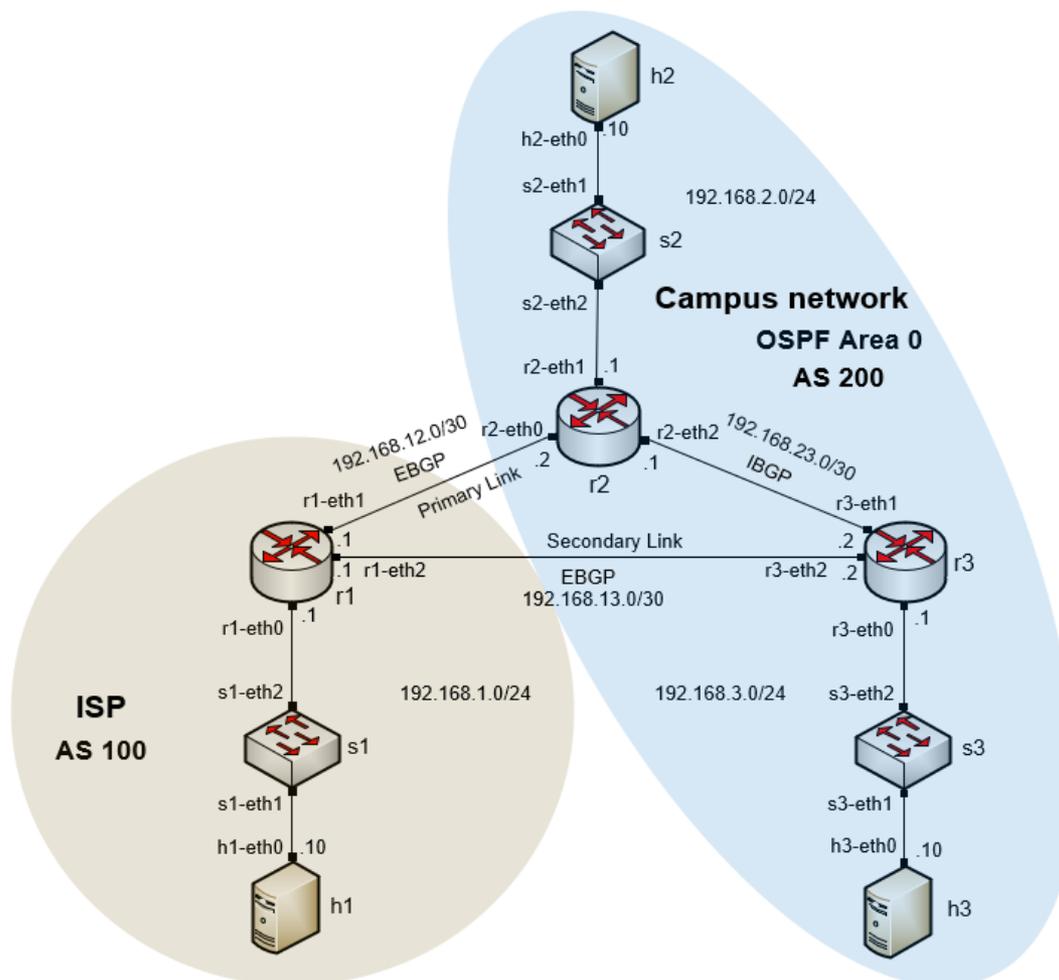
MED attribute

- 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



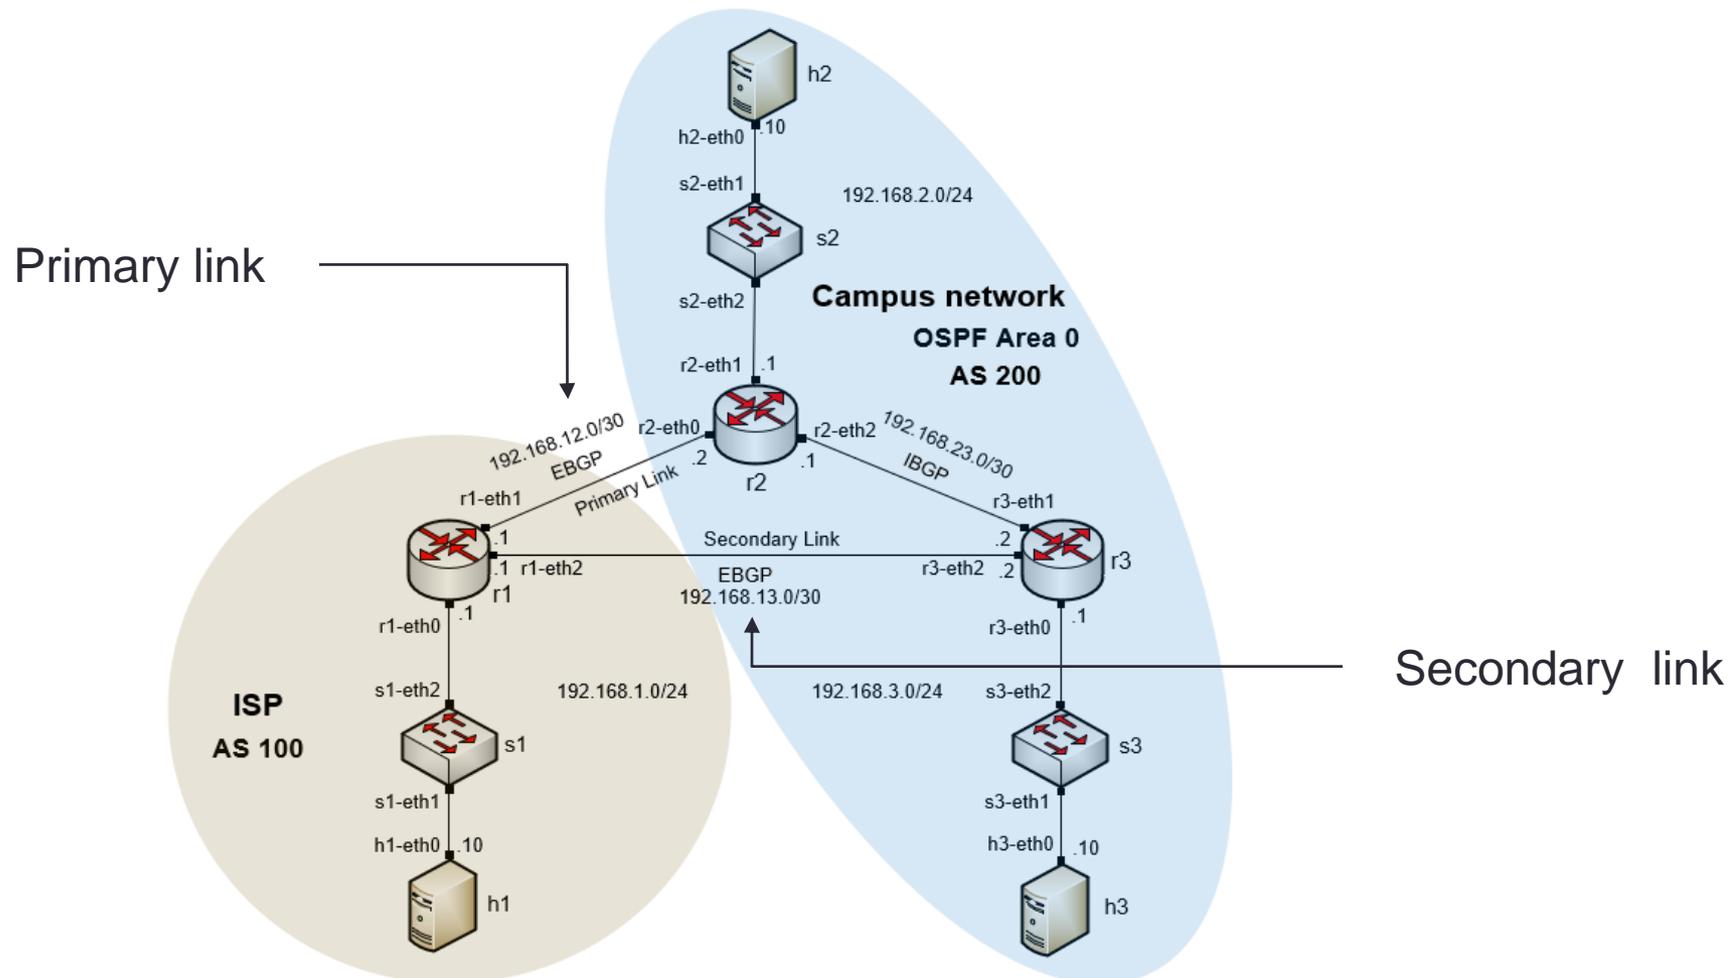
Lab 8 Topology

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



Lab 8 Topology

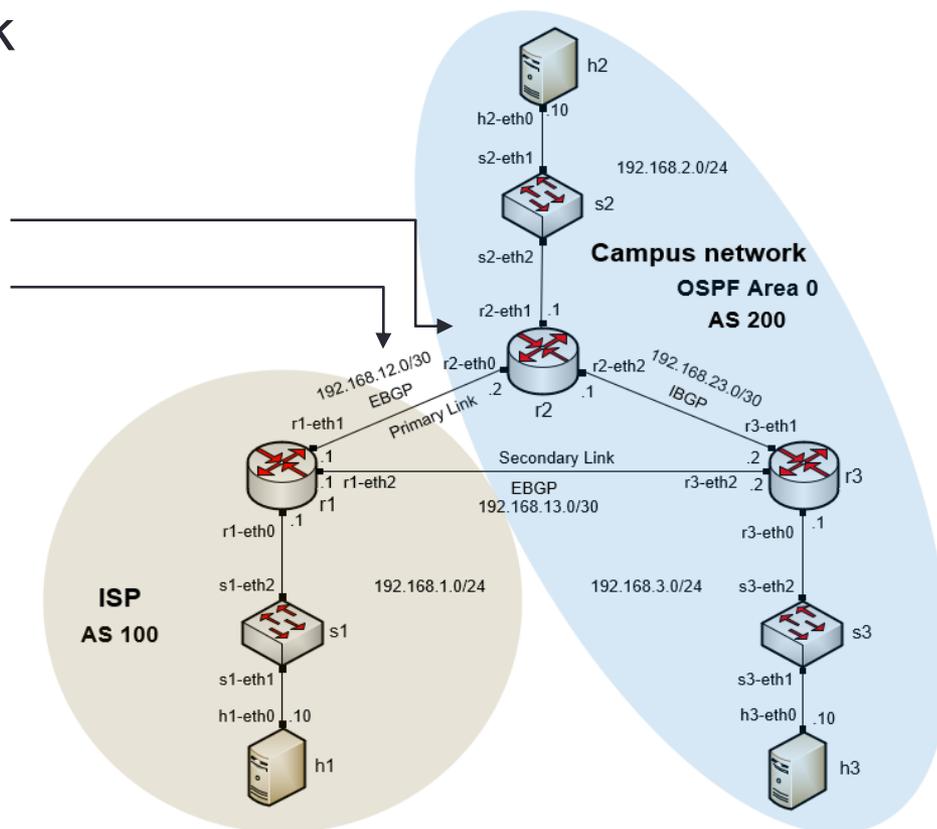
- Configure LOCAL_PREF and MED attributes to favor the primary link over the secondary one



Lab 8 Configuration

- Create a route map to set the LOCAL_PREF attribute of the primary link

Router r2
Primary link



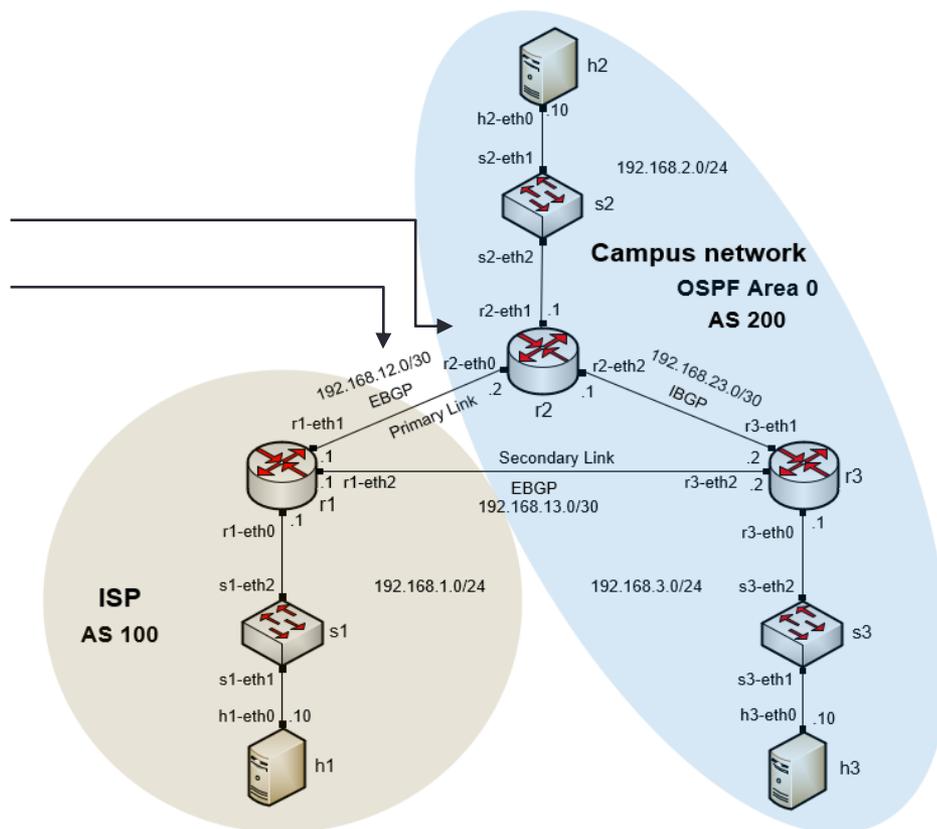
```

Host: r2"
frr-pc# configure terminal
frr-pc(config)# route-map primary_in permit 10
frr-pc(config-route-map)#
  
```

Lab 8 Configuration

- Set the LOCAL_PREF value of the primary link to 150

Router r2
Primary link



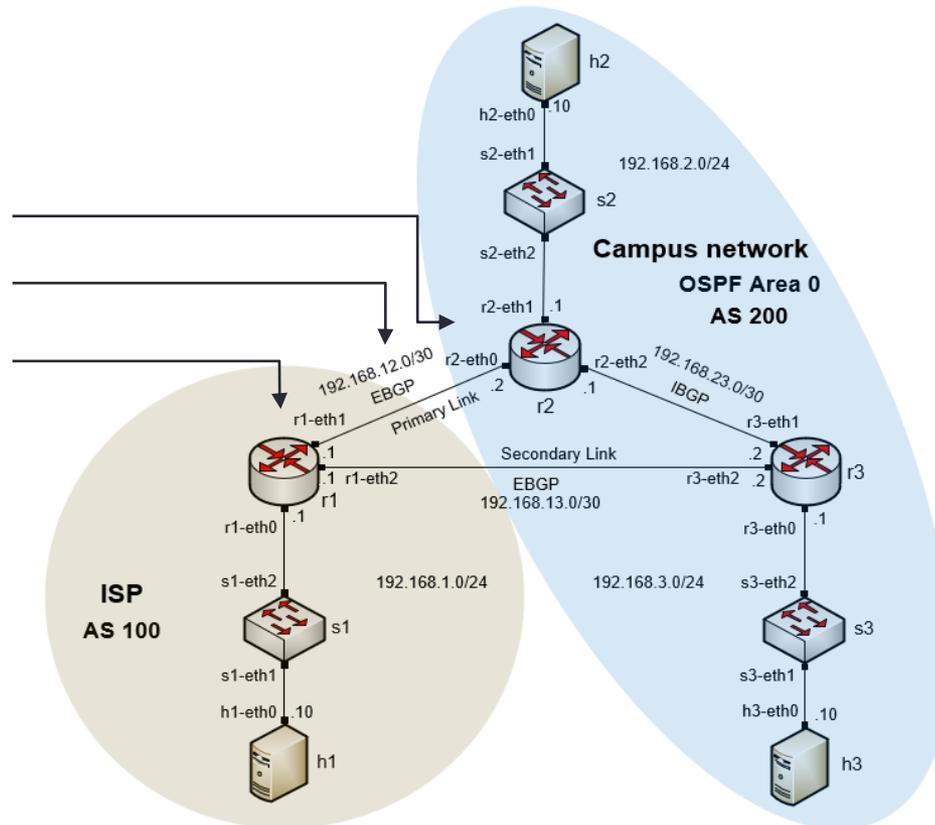
```

Host: r2
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)#
  
```

Lab 8 Configuration

- Assign the route-map to the BGP neighbor (router r1)

Router r2
Primary link
Router r1

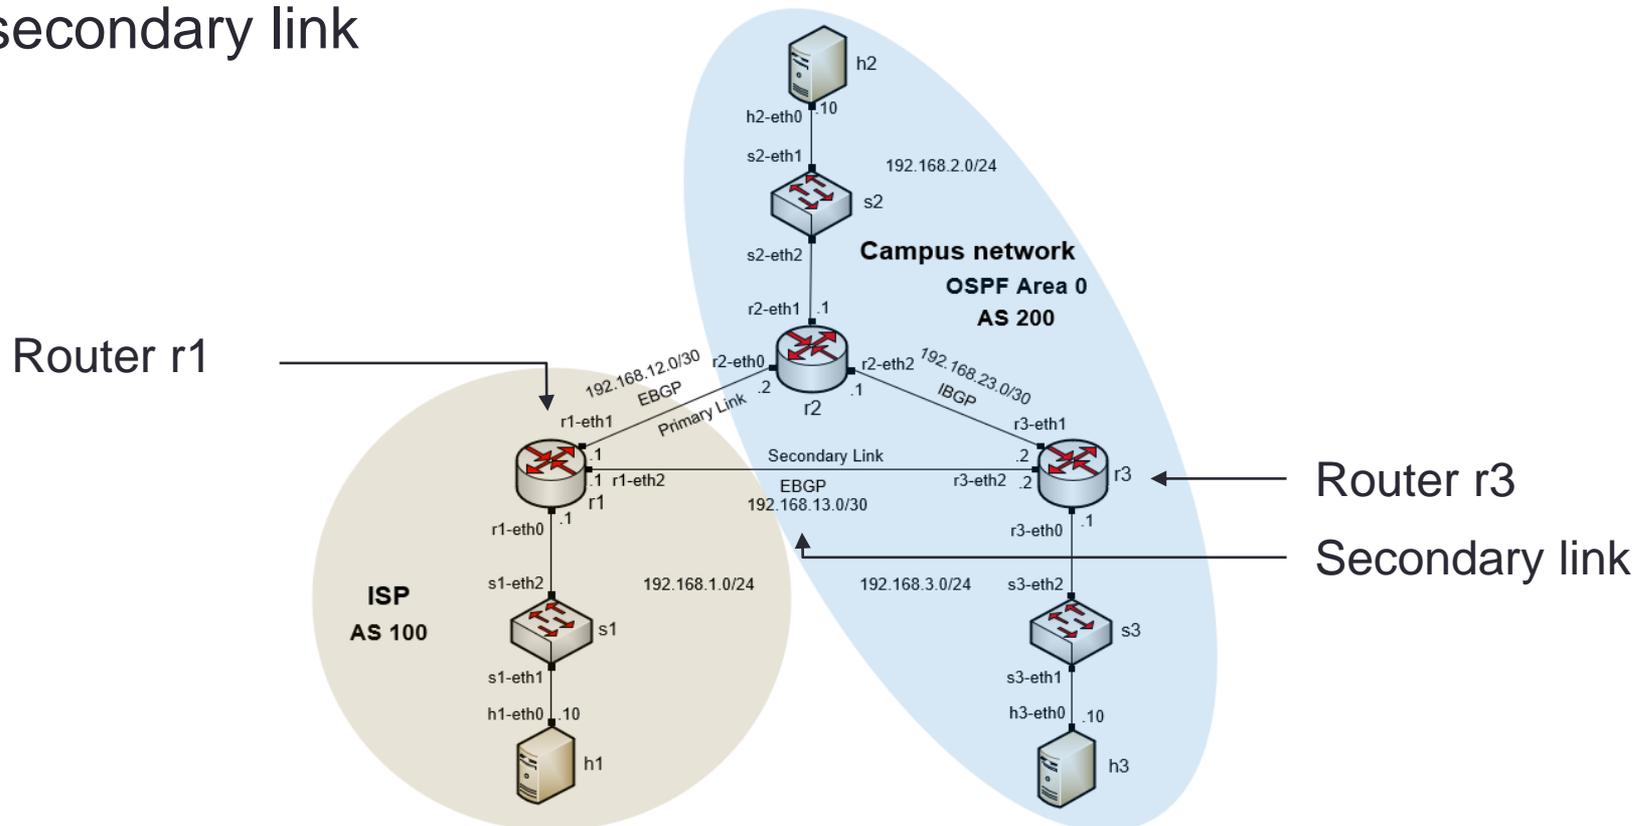


```

X                                     "Host: r2"
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 200
frr-pc(config-router)# neighbor 192.168.12.1 route-map primary_in in
frr-pc(config-router)#
  
```

Lab 8 Configuration

- Similarly, configure LOCAL_PREF on router r3 for the secondary link

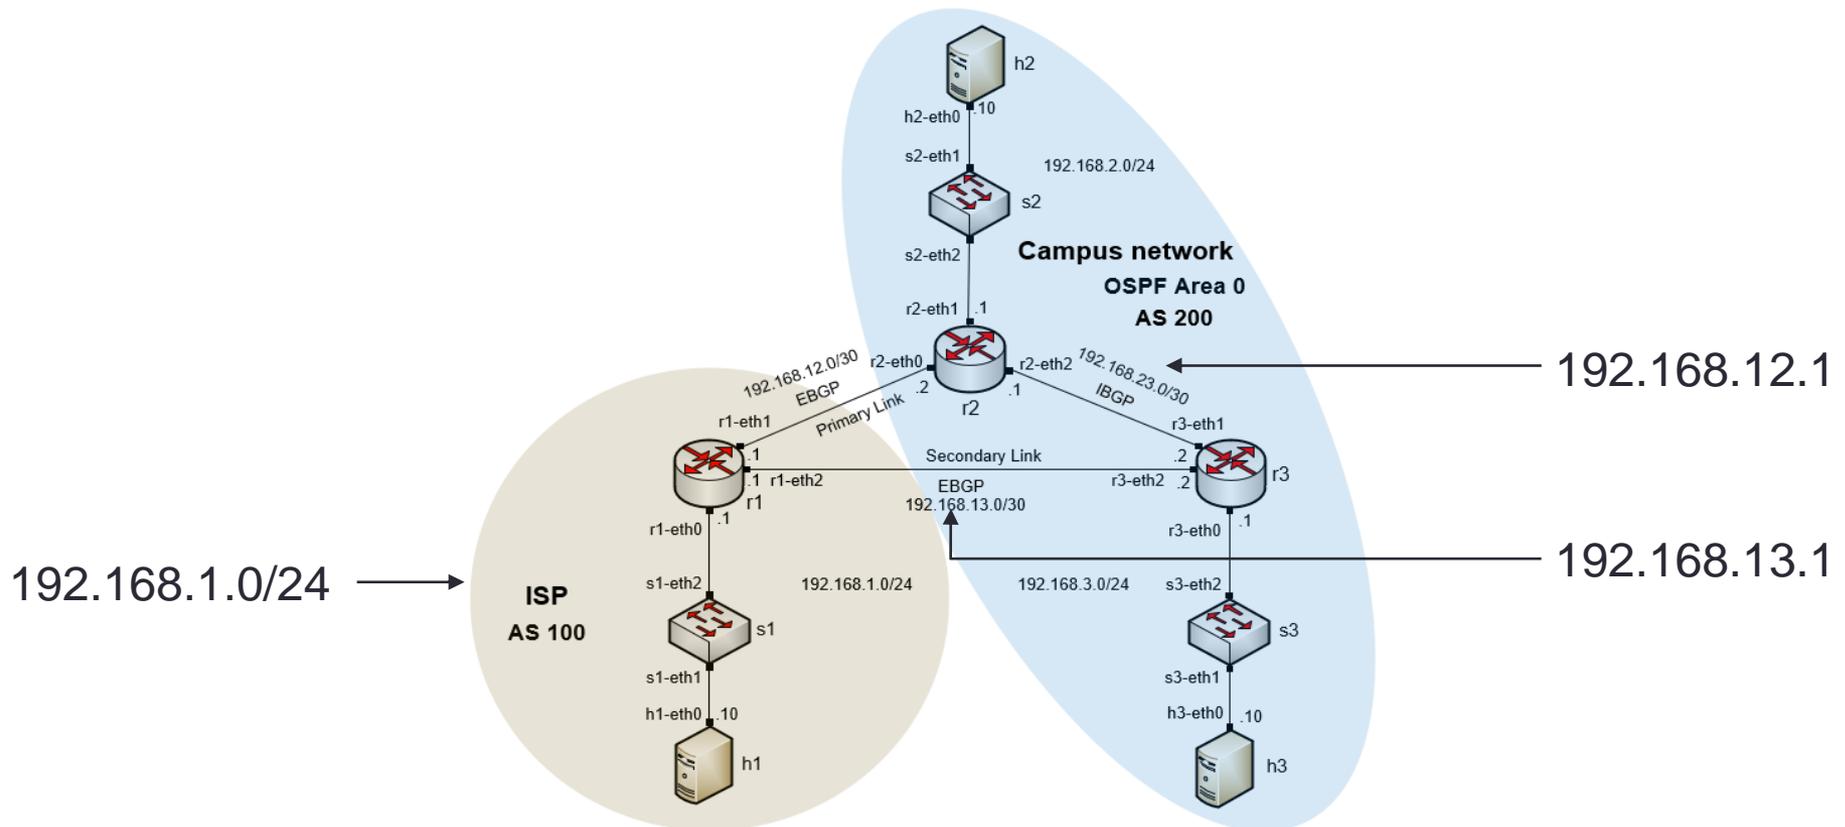


```

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 200
frr-pc(config-router)# neighbor 192.168.13.1 route-map secondary_in in
frr-pc(config-router)# end
  
```

Lab 8 Configuration

- BGP table of router r3



```

Network      Next Hop      Metric LocPrf Weight Path
*->i192.168.1.0/24 192.168.12.1 0 150 0 100 i
*                192.168.13.1 0 125 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

Displayed 3 routes and 4 total paths
frr-pc#

```