Hands-on Workshop on
Open vSwitch and Software-defined Networking

Ali AlSabeh, Jorge Crichigno
University of South Carolina
http://ce.sc.edu/cyberinfra
aalsabeh@email.sc.edu, jcrichigno@cec.sc.edu

WASTC 2021 virtual Faculty Development Weeks (vFDW)
June 22, 2021

Interconnection between legacy networks and SDN networks

Potential Drawbacks of SDN

- SDN networks have many advantages over traditional networks
  - Ease of network management
  - Enforcement of security policies
  - Customized network behavior
- However, SDN is typically not fully deployed in networks due to several reasons
  - Limited budget for new network infrastructure
  - Fear of downtime during the transition to SDN
  - Limited training opportunities in SDN technology

R. Amin, M. Reisslein, N. Shah, Hybrid SDN networks: A survey of existing approaches, IEEE CST, Vol. 20, Issue: 4, 2018.

Potential Drawbacks of SDN

- One possible solution to address these concerns is to deploy a limited number of SDN-enabled devices alongside the traditional (legacy) network devices
- Incrementally replacing traditional network devices by SDN devices
- The network can be converted in stages, targeting specific network areas for conversion and rolling out the changes incrementally
  - A network containing a mix of SDN and legacy network devices is referred to as a hybrid SDN network

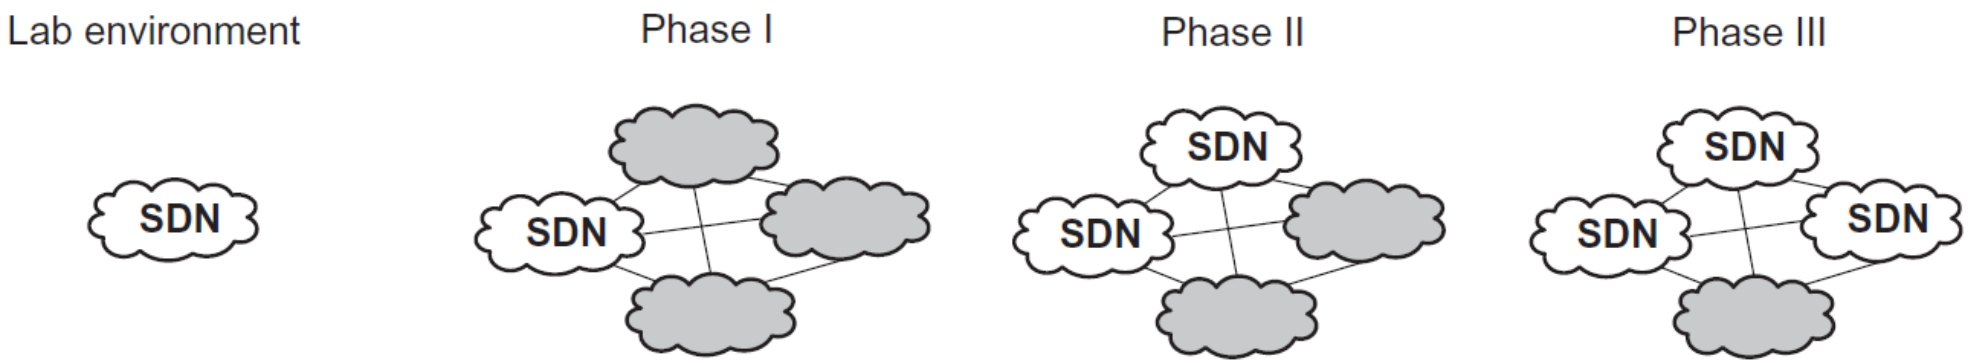

Advantages of Hybrid SDN Networks

- Hybrid SDN networks ease these budget concerns
  - Cost of replacing legacy devices by SDN devices
  - Train engineers to design, configure, and operate the SDN network
- SDN provides fine-grained control for data traffic flows
  - If this is required for a small network portion, then SDN can be implemented in that portion only
- Scenarios where two SDN networks are interconnected by legacy network devices require hybrid SDN network mechanisms

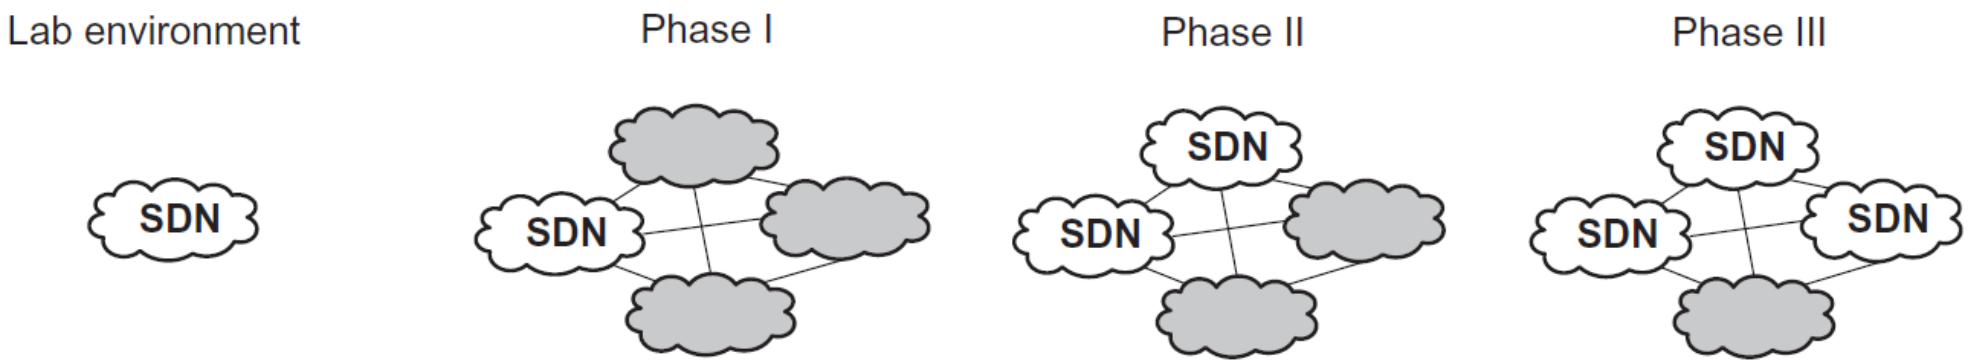

Border Gateway Protocol

- The Border Gateway Protocol (BGP) version 4 is the standard inter-autonomous system (AS) protocol in today's Internet
  - An AS is a group of routers typically under the same administrative control (e.g., ISP, company)
  - BGP is the *"glue that holds the Internet together"*
- In BGP, pairs of routers exchange routing information over TCP (default port 179)

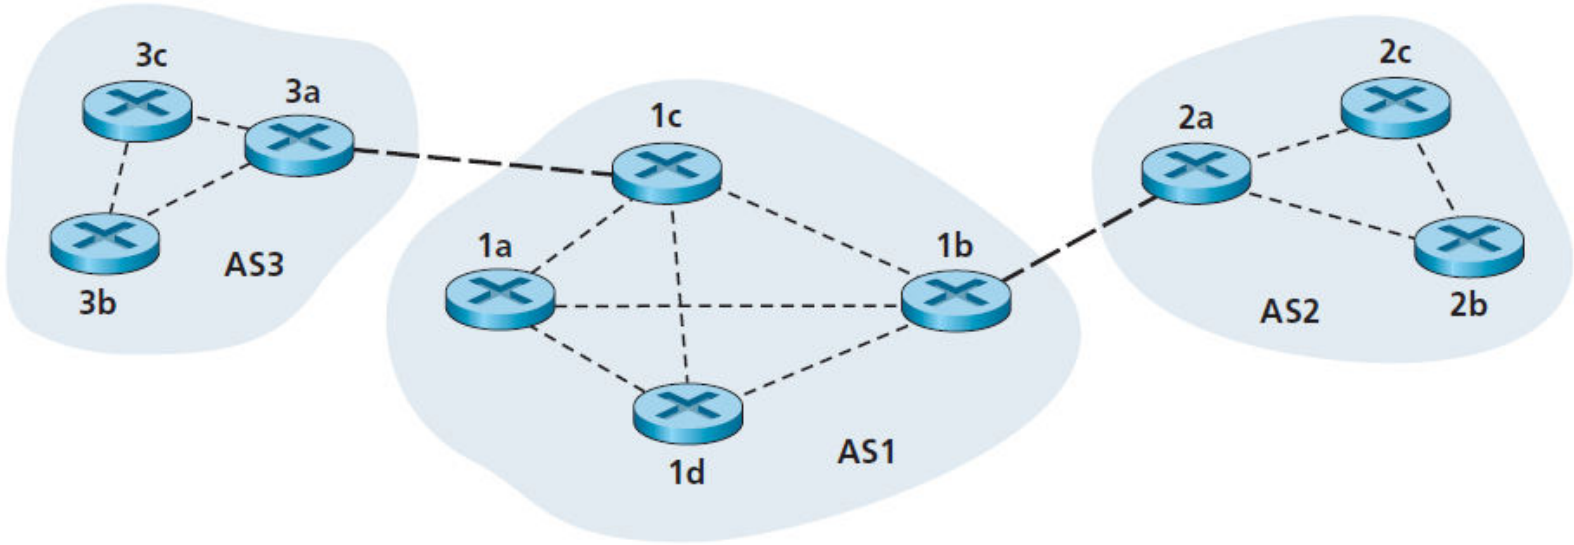

J. Kurose, K. Ross, Computer Networking: A Top-Down Approach, 7th Edition, Pearson/Addison Wesley, 2017.

Border Gateway Protocol

- BGP provides each AS a means to:
  - eBGP: obtain subnet reachability information from neighboring ASes
  - iBGP: propagate reachability information to AS-internal routers
  - Determine “good” routes to other networks based on reachability information and policy
- Allows subnet to advertise its existence to rest of Internet: “I am here”

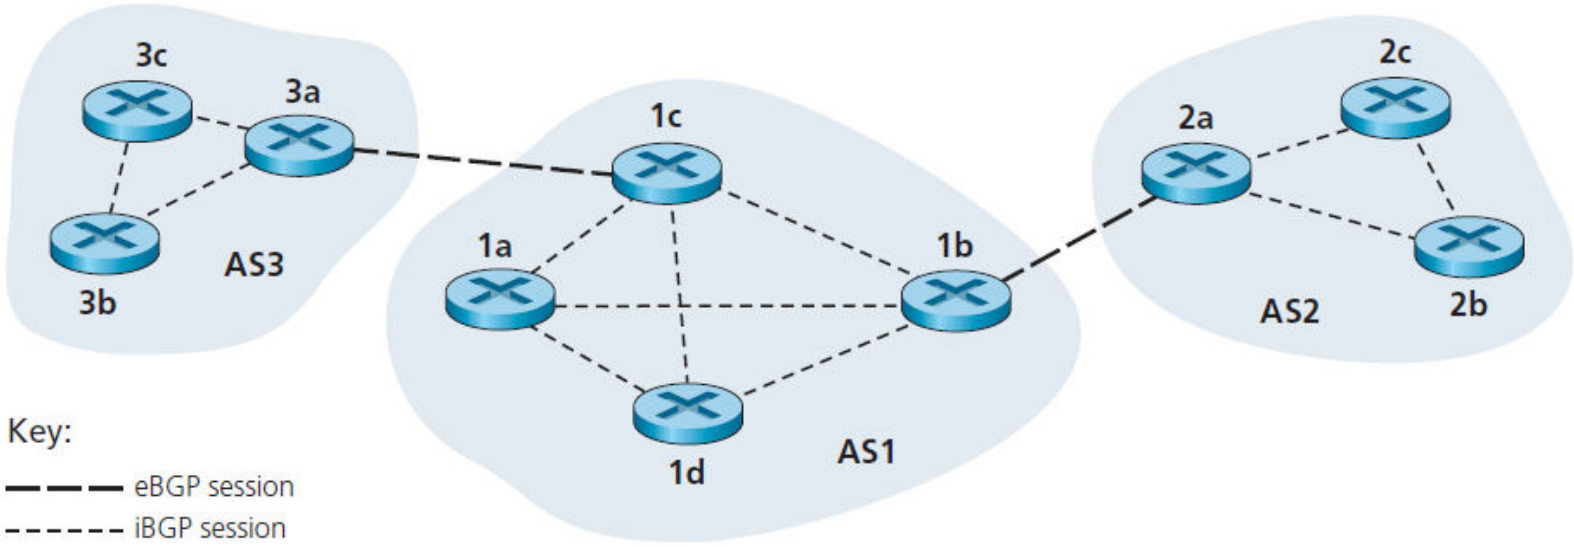

J. Kurose, K. Ross, Computer Networking: A Top-Down Approach, 7th Edition, Pearson/Addison Wesley, 2017.

Lab 8: Interconnection between legacy networks and SDN networks

Lab 8: Interconnecting Legacy and SDN networks

- Two legacy networks connected to an SDN network
- SDN network consists of switches controlled by an ONOS controller
- The ONOS controller interacts with an application referred to as SDN-IP
- SDN-IP allows the SDN network to i) exchange BGP information with an iBGP router; and ii) translates routing information to SDN flow rules

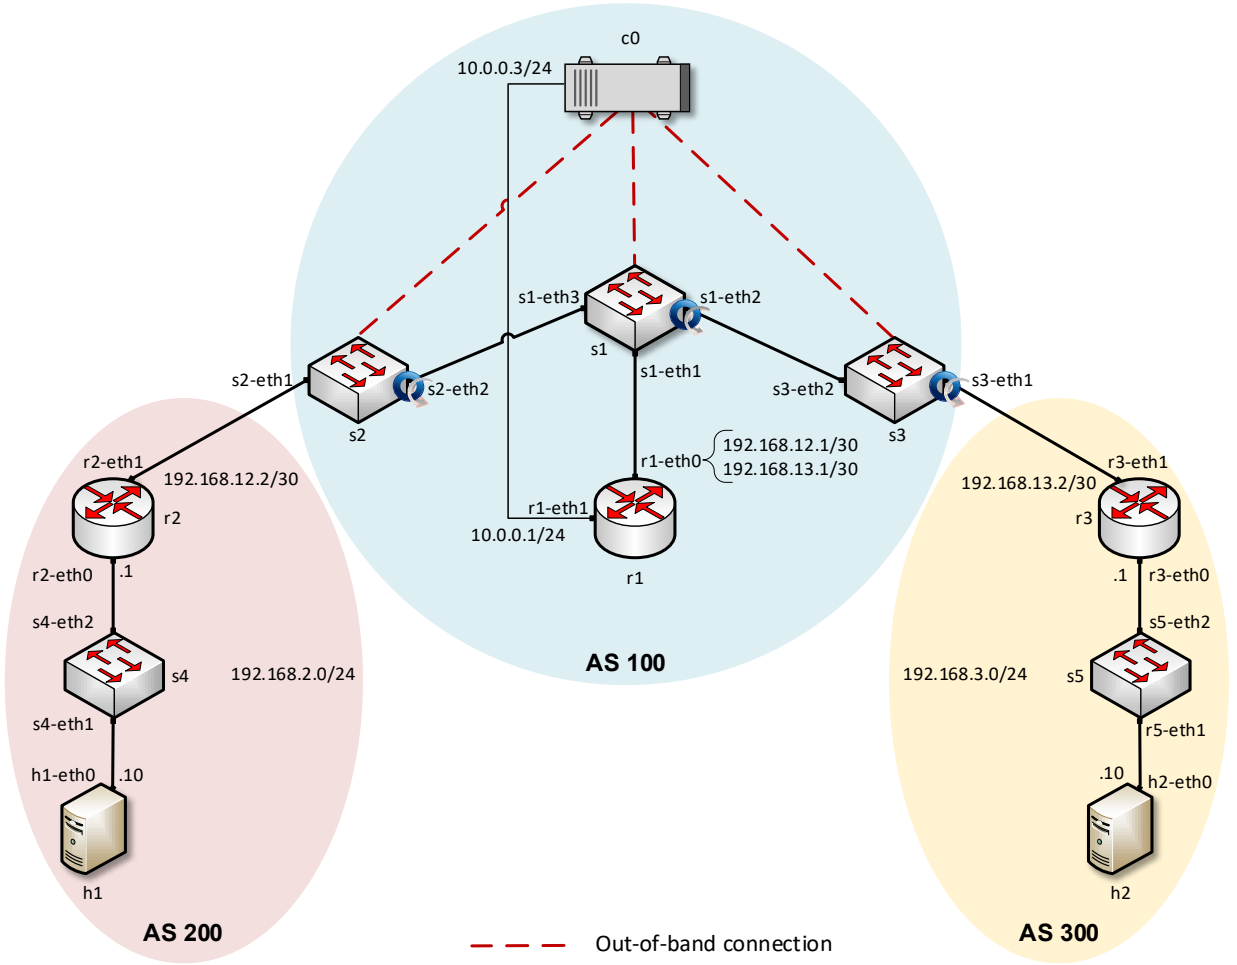

Lab 8: Interconnecting Legacy and SDN networks

1. Routers r2 and r3 attempt to connect to r1 (eBGP)
2. Switches s2 and s3 forward incoming packets to controller
3. Controller installs rules to forward those BGP packets to r1
4. Router r1 exchanges BGP information to controller (iBGP)
5. Controller translates BGP information into rules
6. Rules are installed in s1, s2, and s3 (remote networks)
7. Full connectivity is established

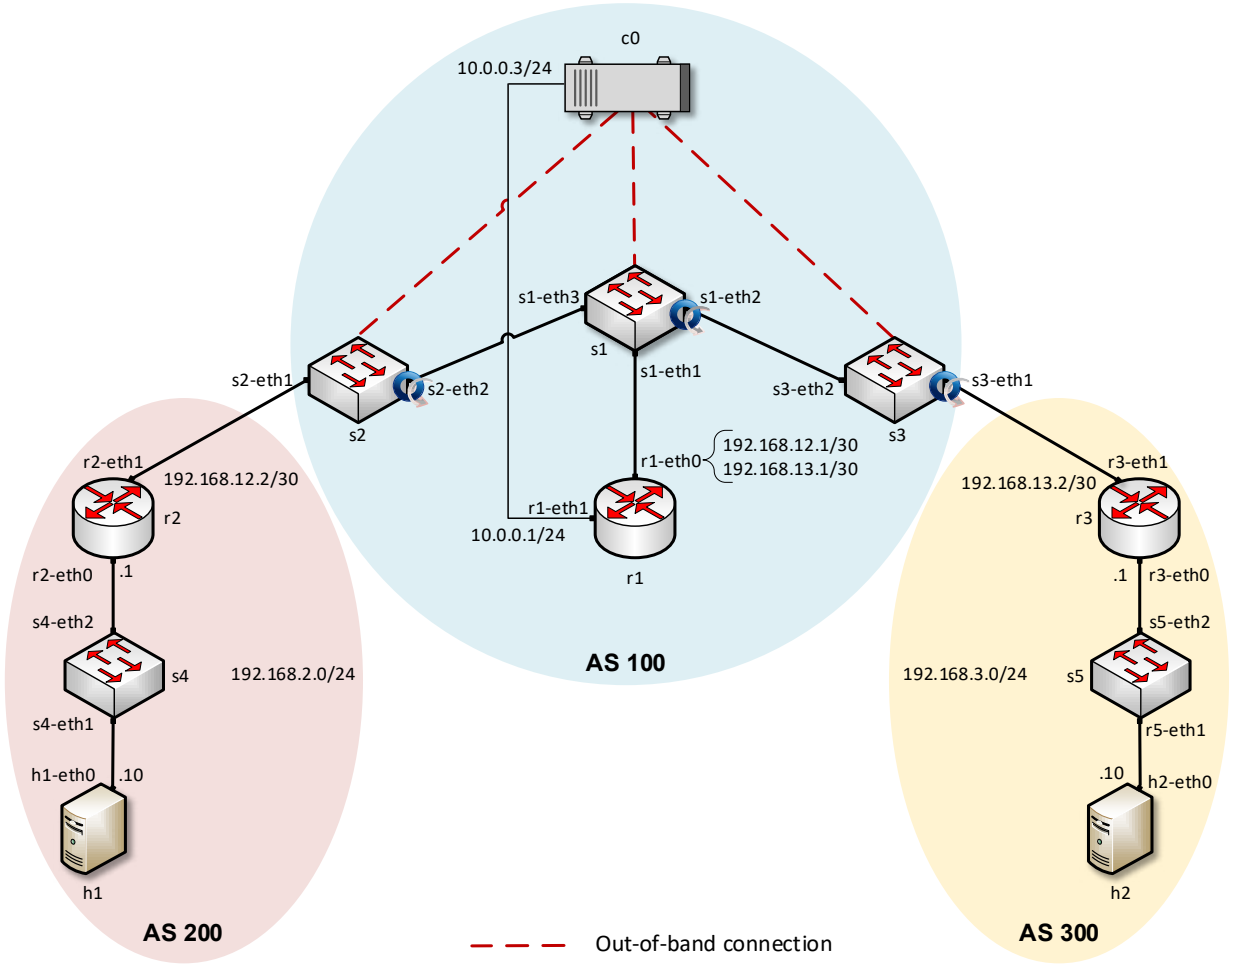

Configuring BGP in Legacy Networks

Router r2

```
"Host: r2"
admin# exit
root@admin:/etc/routers/r2# bgpd
root@admin:/etc/routers/r2# vtysh

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

admin# configure terminal
admin(config)# router bgp 200
admin(config-router)# neighbor 192.168.12.1 remote-as 100
admin(config-router)# network 192.168.2.0/24
admin(config-router)#
```

Router r3

```
"Host: r3"
admin# exit
root@admin:/etc/routers/r3# bgpd
root@admin:/etc/routers/r3# vtysh

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

admin# configure terminal
admin(config)# router bgp 300
admin(config-router)# neighbor 192.168.13.1 remote-as 100
admin(config-router)# network 192.168.3.0/24
admin(config-router)# end
admin#
```

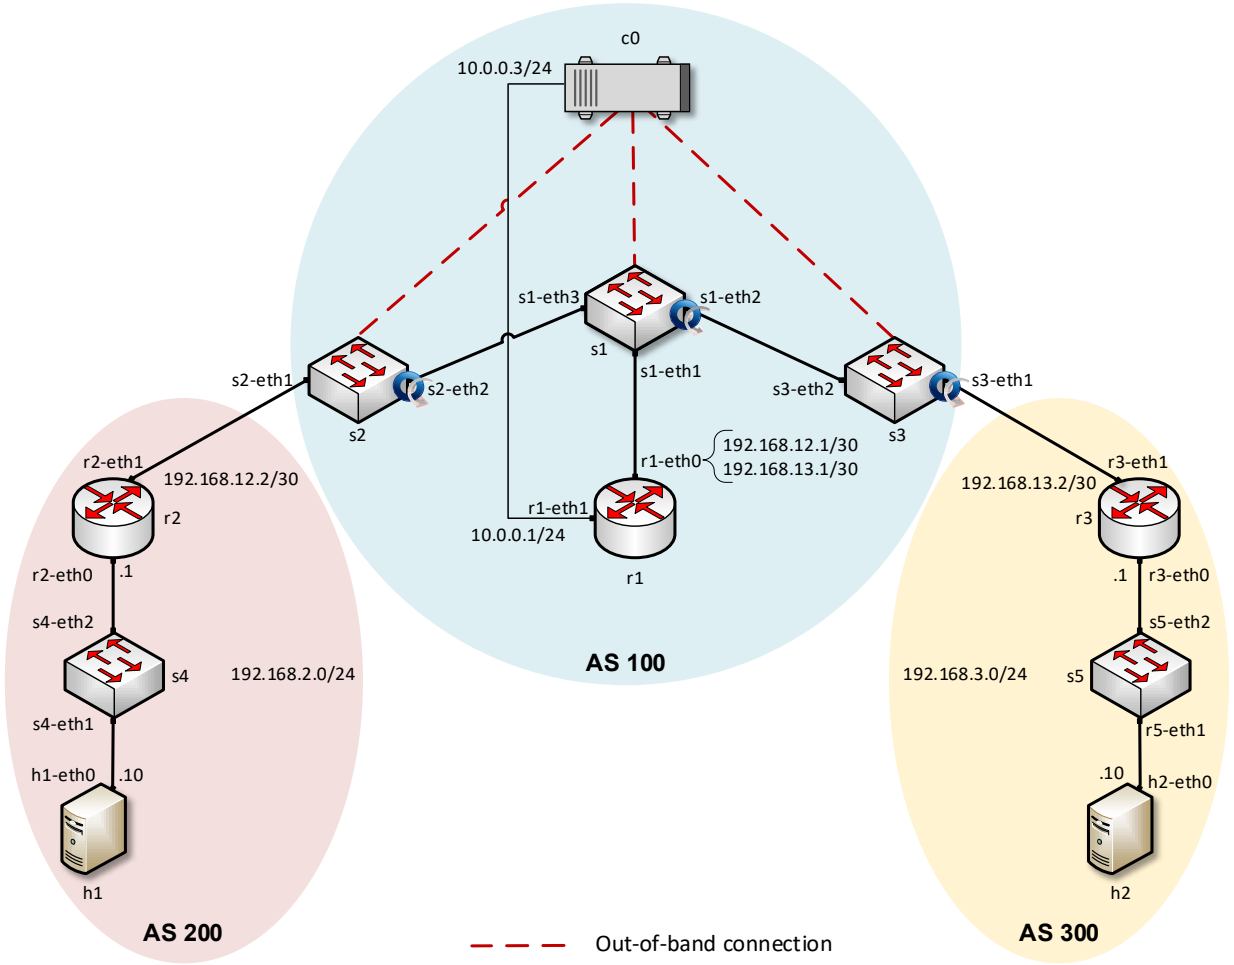

Configure BGP in SDN Network

Router r1

```
"Host: r1"
root@admin:/etc/routers/r1# bgpd
root@admin:/etc/routers/r1# vtysh

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

admin# configure terminal
admin(config)# router bgp 100
admin(config-router)# neighbor 192.168.12.2 remote-as 200
admin(config-router)# neighbor 192.168.13.2 remote-as 300
admin(config-router)# neighbor 10.0.0.3 remote-as 100
admin(config-router)# neighbor 10.0.0.3 port 2000
admin(config-router)#
```

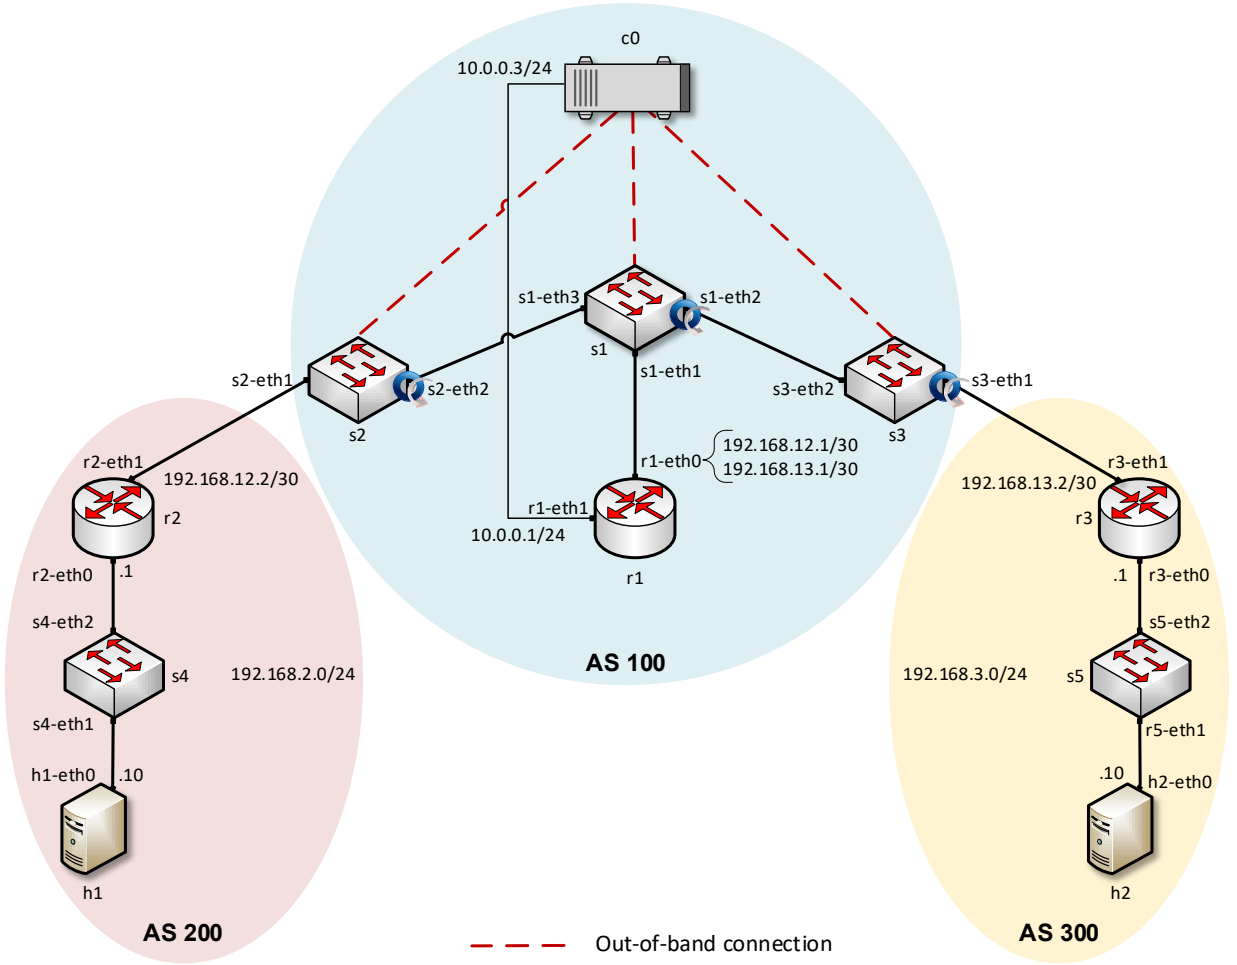

SDN-IP Application

ONOS CLI – BGP neighbors

```
root@admin: /home/sdn/SDN_Labs/lab8
File  Actions  Edit  View  Help
root@admin: /home/sdn/SDN_Labs/lab8
karaf@root > bgp-neighbors                                    14:58:58
BGP neighbor is 192.168.13.1, remote AS 100, local AS 100
  Remote router ID 192.168.13.1, IP /10.0.0.1:48998, BGP version 4, Hol
d time 180
  Remote AFI/SAFI IPv4 Unicast YES Multicast NO, IPv6 Unicast NO Multic
ast NO
  Local router ID 10.0.0.3, IP /10.0.0.3:2000, BGP version 4, Hold tim
e 180
  Local  AFI/SAFI IPv4 Unicast YES Multicast NO, IPv6 Unicast NO Multic
ast NO
  4 Octet AS Capability: Advertised Received
karaf@root >                                                  14:59:00
```

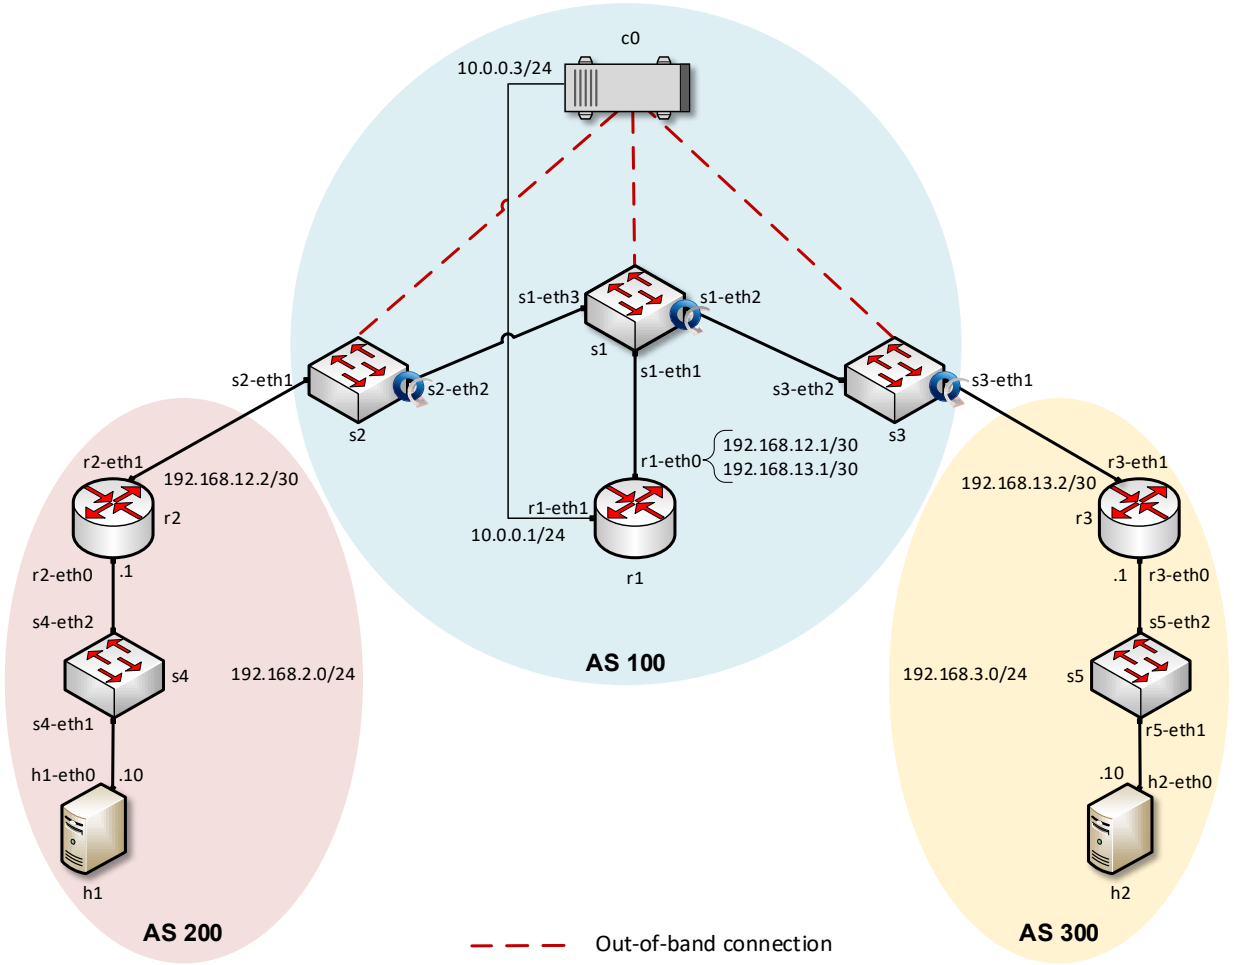

SDN-IP Application

ONOS CLI – advertised routes

```
root@admin: /home/sdn/SDN_Labs/lab8
File  Actions  Edit  View  Help
root@admin: /home/sdn/SDN_Labs/lab8
karaf@root > routes
B: Best route, R: Resolved route

Table: ipv4
B R  Network            Next Hop         Source (Node)
> *  192.168.2.0/24     192.168.12.2     BGP (172.17.0.2)
> *  192.168.3.0/24     192.168.13.2     BGP (172.17.0.2)
   Total: 2

Table: ipv6
B R  Network            Next Hop         Source (Node)
   Total: 0

karaf@root >
```

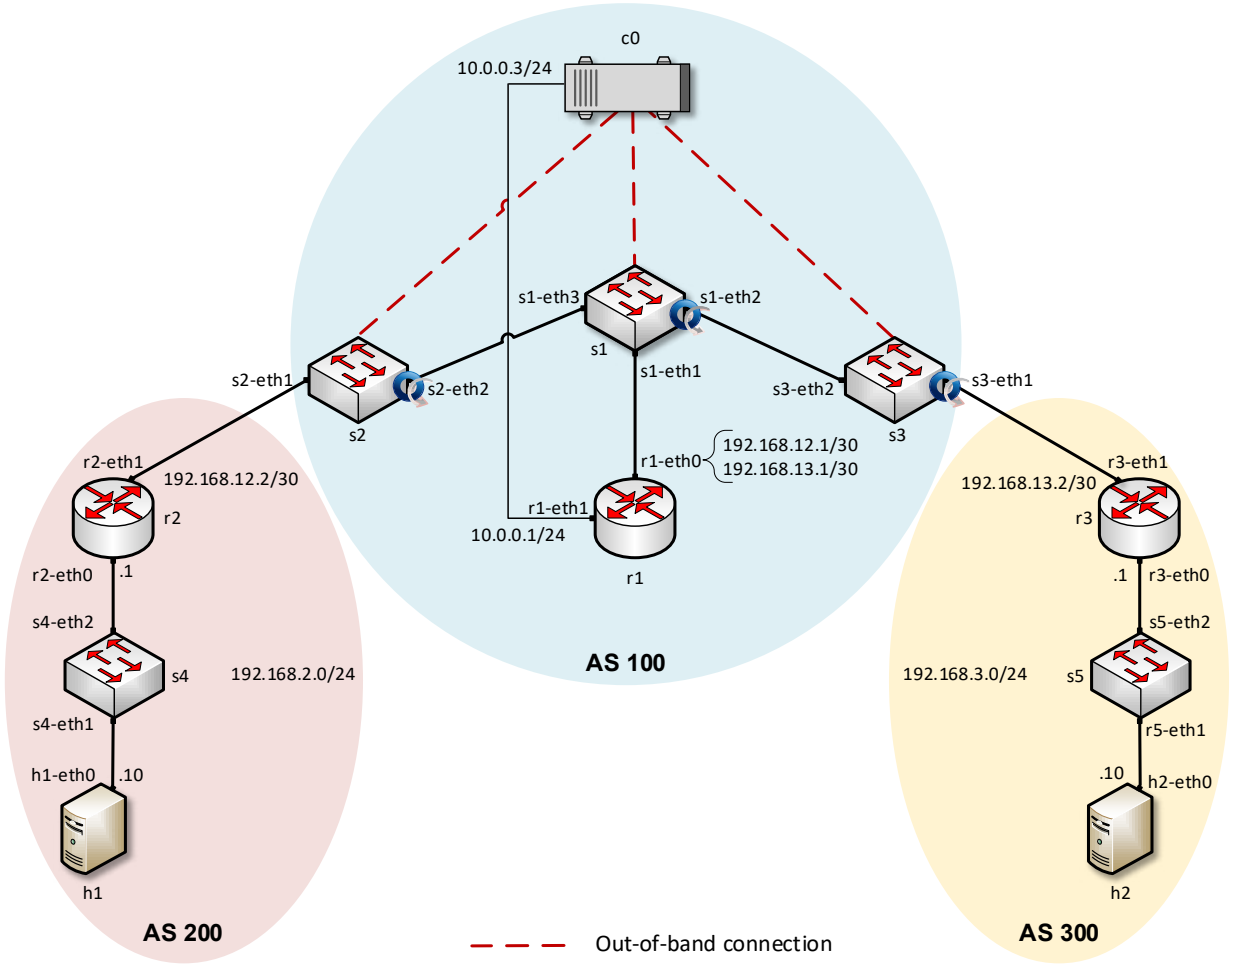

SDN Network

BGP table of router r1

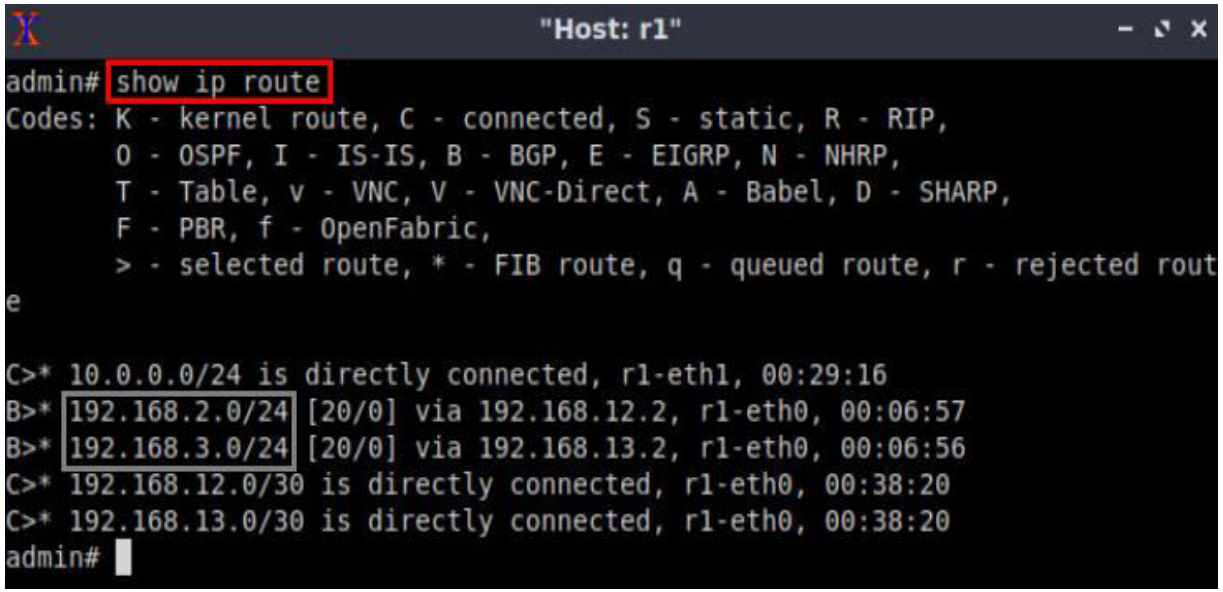

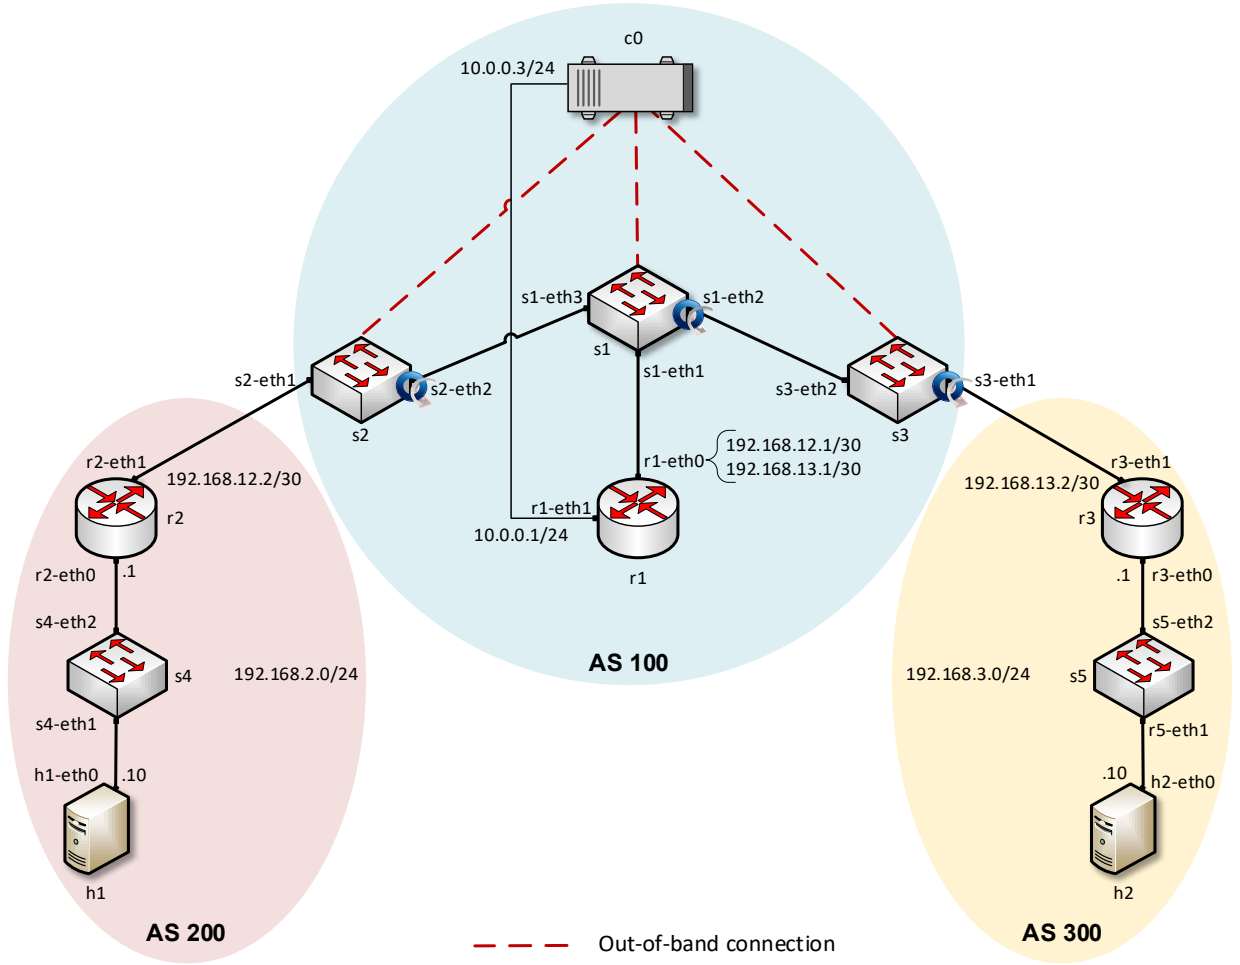

Legacy Network

Routing table of router r2