
TOMORROW
starts here.
Cisco
Cisco live!

MPLS for R&S CCIE Candidates

Johnny Bass

CCIE #6458

Cisco live!

About the Presenter

- Johnny Bass
- Networking industry since the late 1980s
- CCIE R&S #6458
- CCSI 97168
- Cisco 360 R&S Master Instructor
- Course director for several programs, including Cisco 360 Route Switch, for Global Knowledge

Cisco live!

Why Are We Here?

- Show of hands, how many of you are currently supporting MPLS?
- Show of hands, how many of you actually have MPLS configured on a router?
- Of those, how many of you work for a service provider?
- How many of you attended this session last year?

Cisco live!

MPLS and the CCIE R&S Exam (V5.0)

4.1.1	Implement and Troubleshoot MPLS Operations
4.1.1 a	Label stack, LSR, LSP
4.1.1 b	LDP
4.1.1 c	MPLS Ping, MPLS Traceroute
4.1.2	Implement and Troubleshoot Basic MPLS L3VPN
4.1.2 a	L3VPN, CE, PE, P
4.1.2 b	Extranet (Route leaking)

Agenda

- **Multiprotocol Label Switching Review**
- Virtual Private Networks (without VRFs)
- Virtual Routing and Forwarding tables (VRF Lite)
- Virtual Private Networks (with VRFs)
- MPLS advanced topics (CCIE twists)
- Troubleshooting
- Q&A

Cisco live!

MPLS History

- **Original drivers towards label switching:**
  - Designed to make routers faster
    - ATM switches were faster than routers
    - Fixed length label lookup faster than longest match used by IP routing
    - Allow a device to do the same job as a router with performance of ATM switch
  - Enabled IP + ATM integration
    - Mapping of IP to ATM had become very complex, hence simplify by replacing ATM signaling protocols with IP control protocols

Cisco live!

MPLS History

- 1996 Ipsilon Networks proposed “IP Switching” technology based on ATM switches
- Cisco Systems introduced a related proposal, not restricted to ATM transmission, called "Tag Switching“
- Cisco Tag Switching was a proprietary protocol, and was renamed "Label Switching“ and handed over to the Internet Engineering Task Force (IETF) for open standardization

Cisco live!

MPLS the Beginning

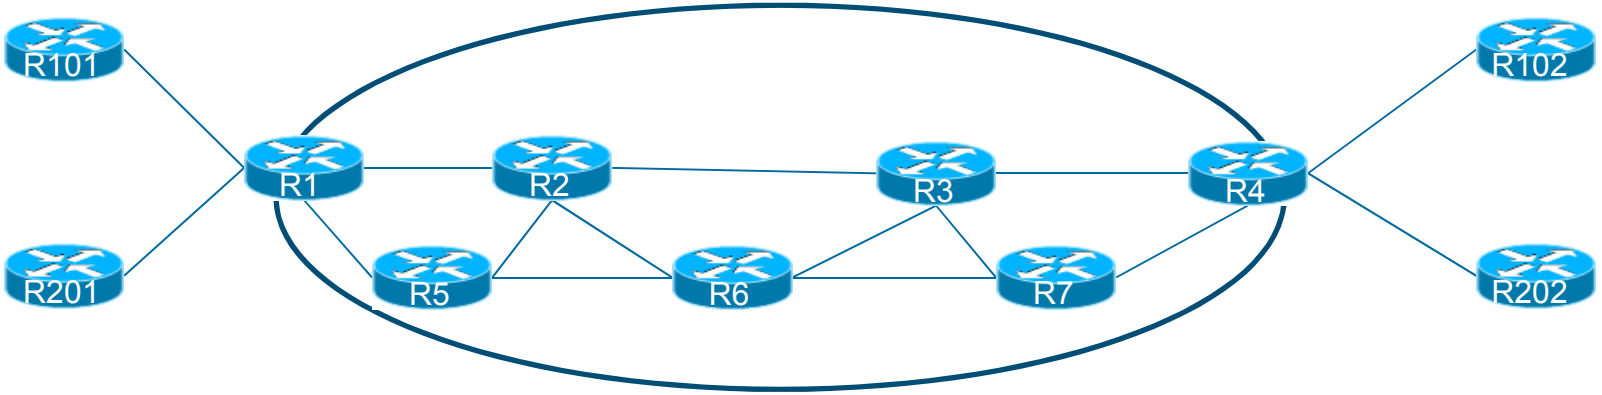

MPLS the Beginning – continued

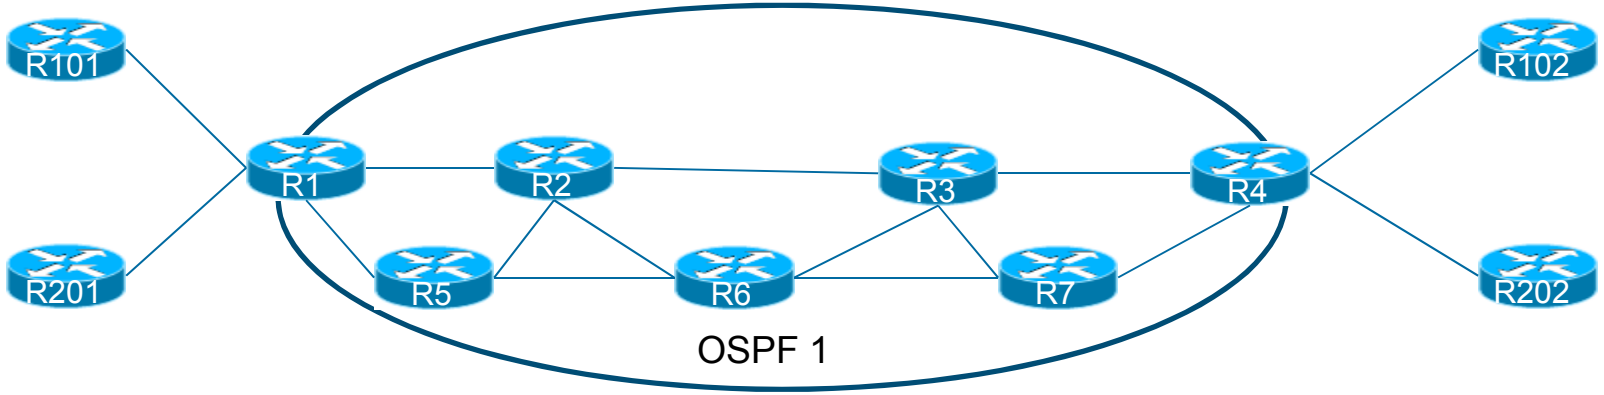

MPLS the Beginning – continued

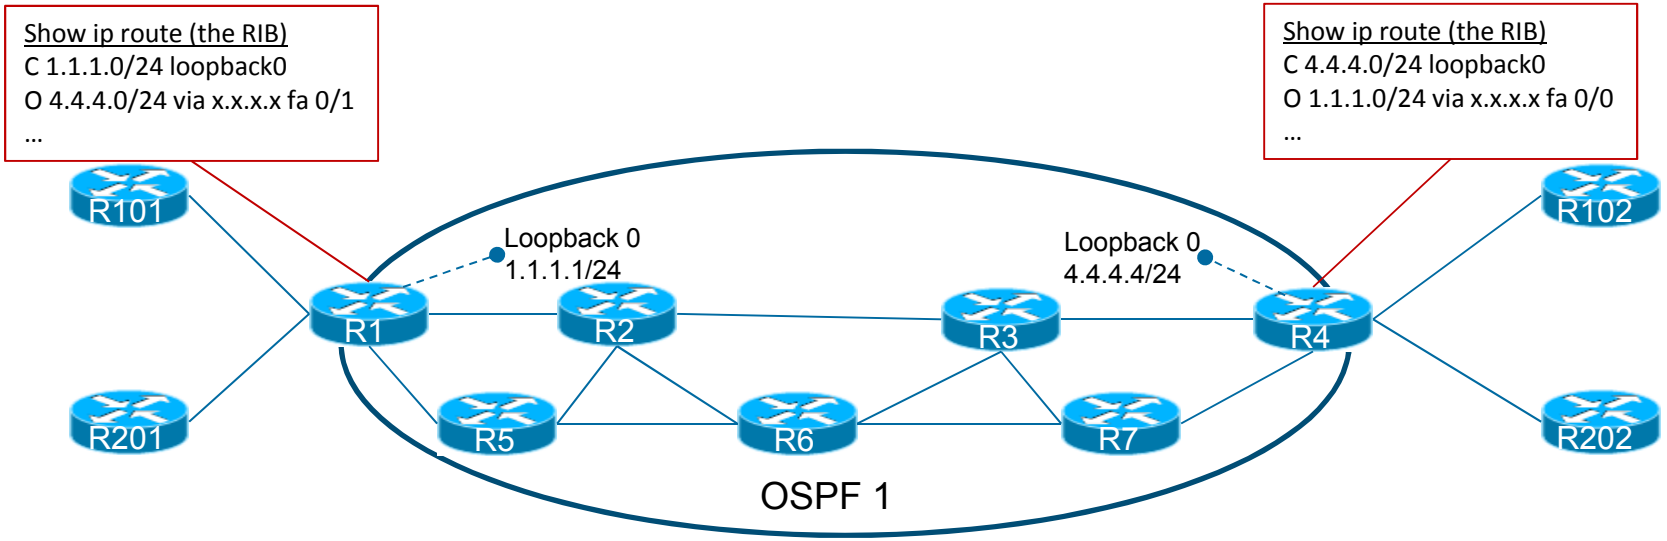

MPLS the Beginning – continued

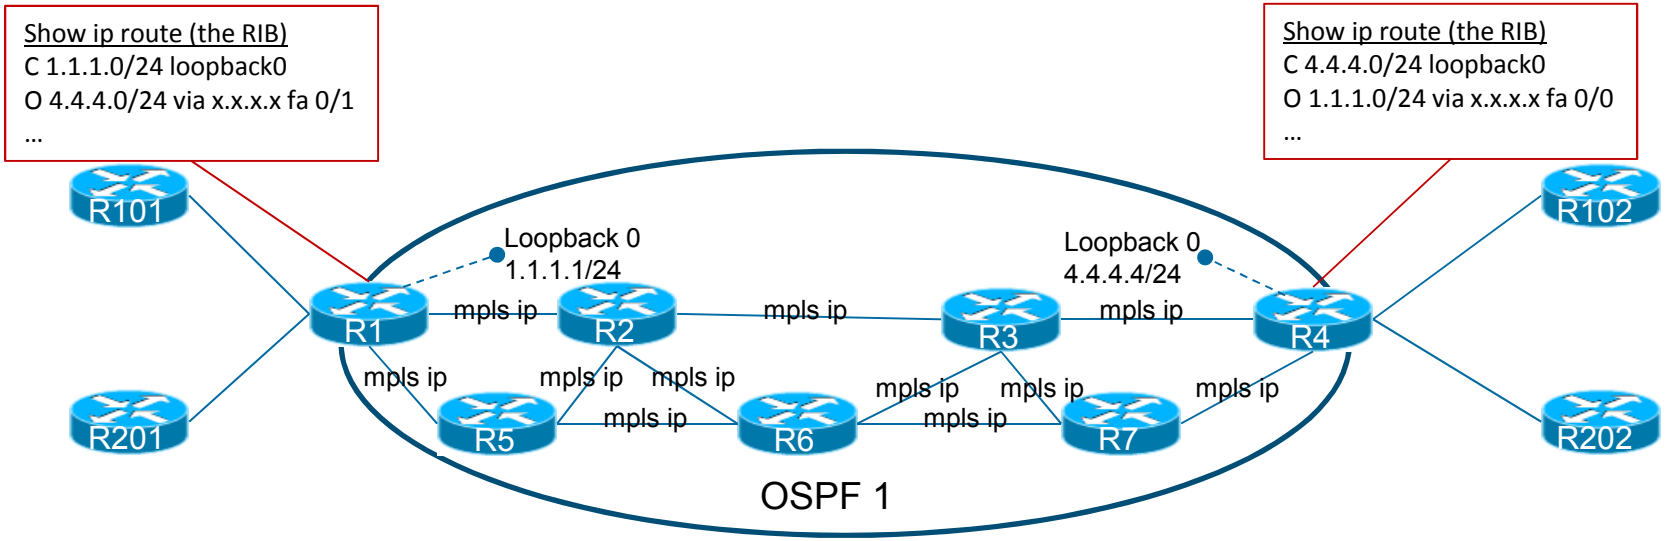

MPLS the Beginning – continued

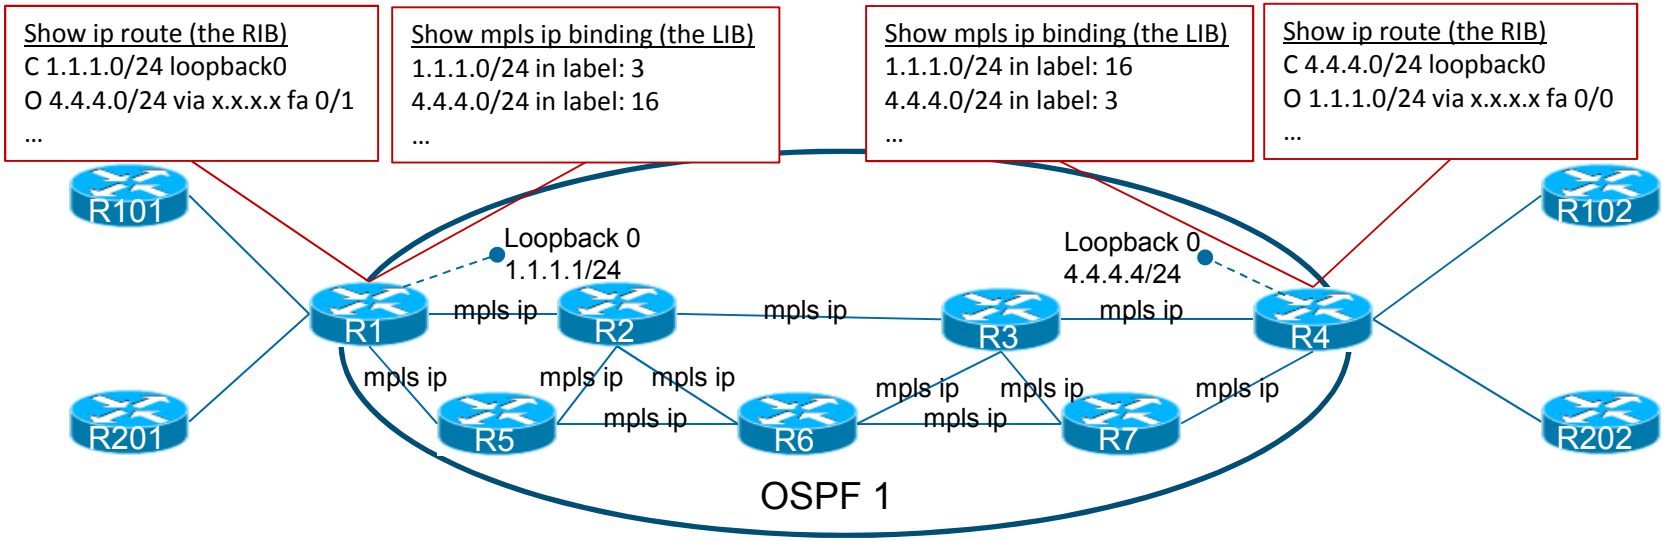

MPLS the Beginning – continued

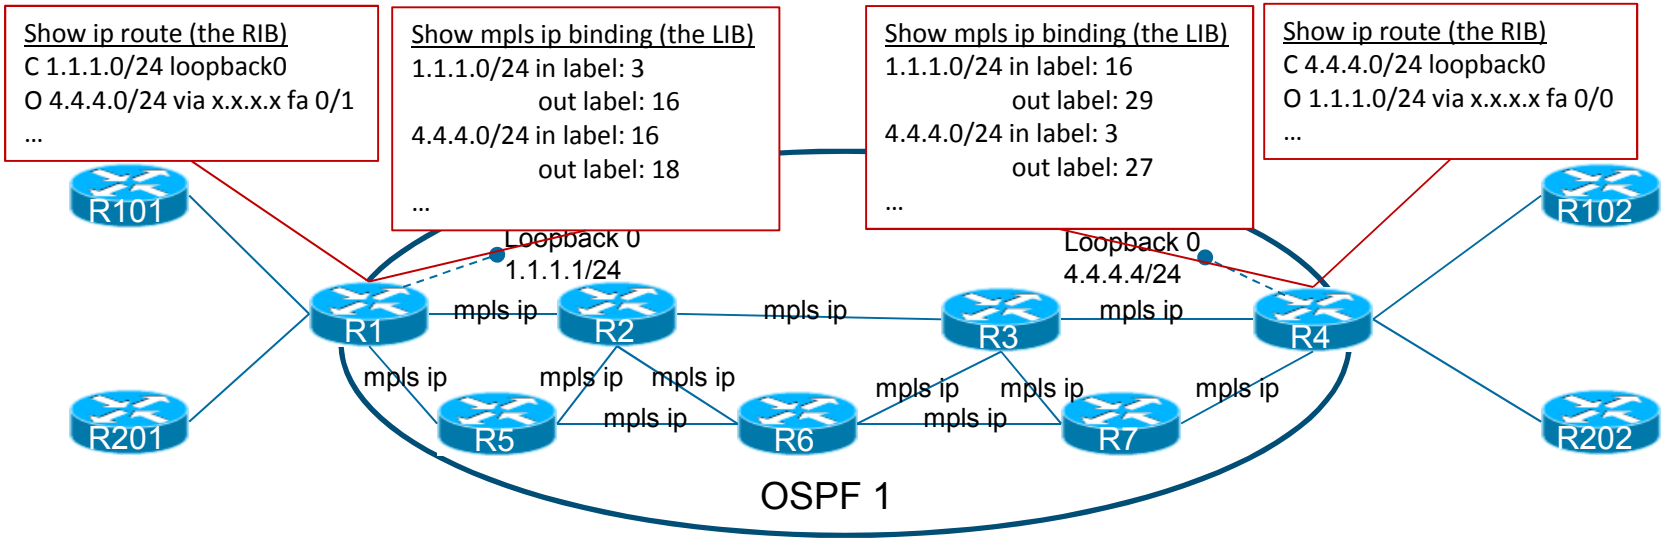

MPLS the Beginning – continued

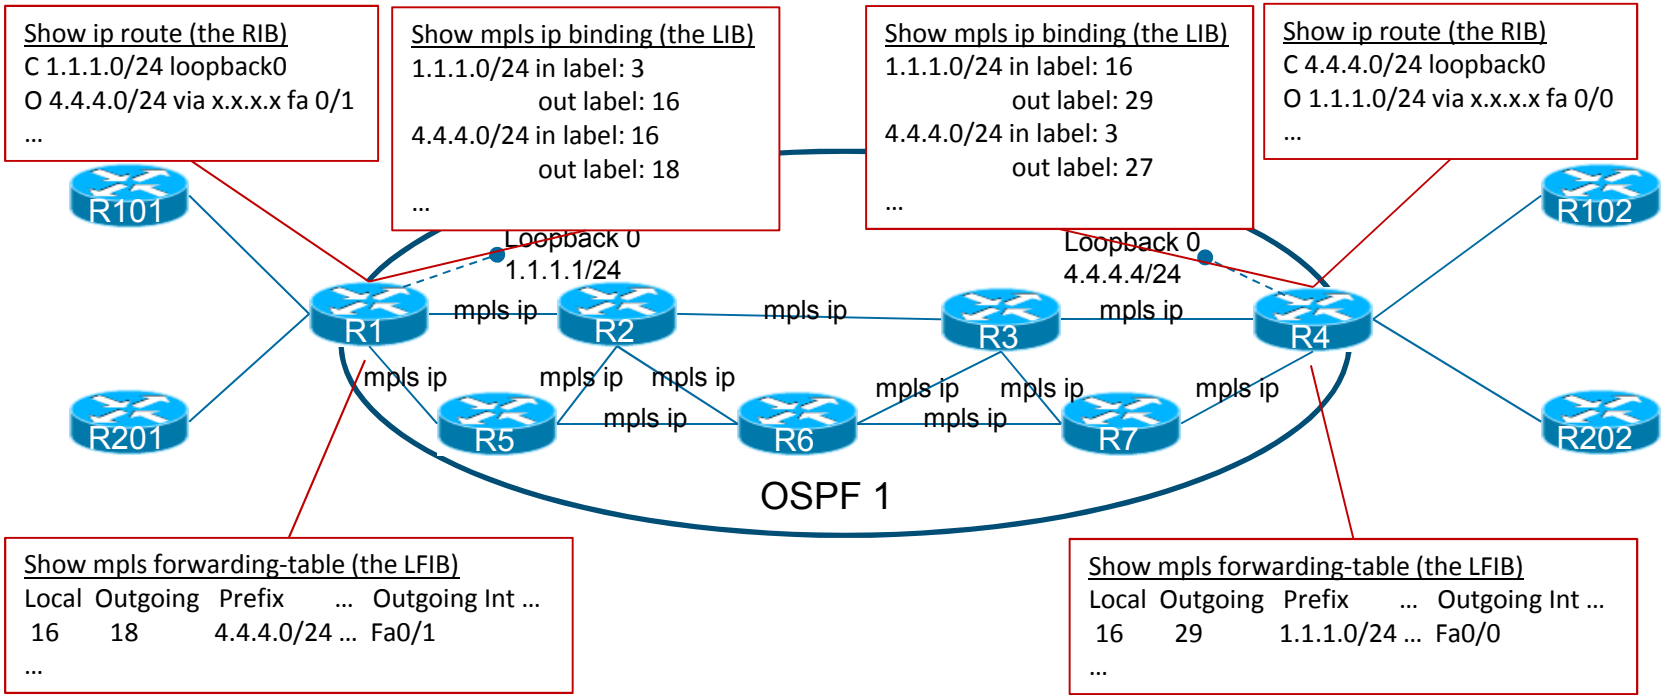

MPLS the Beginning – continued

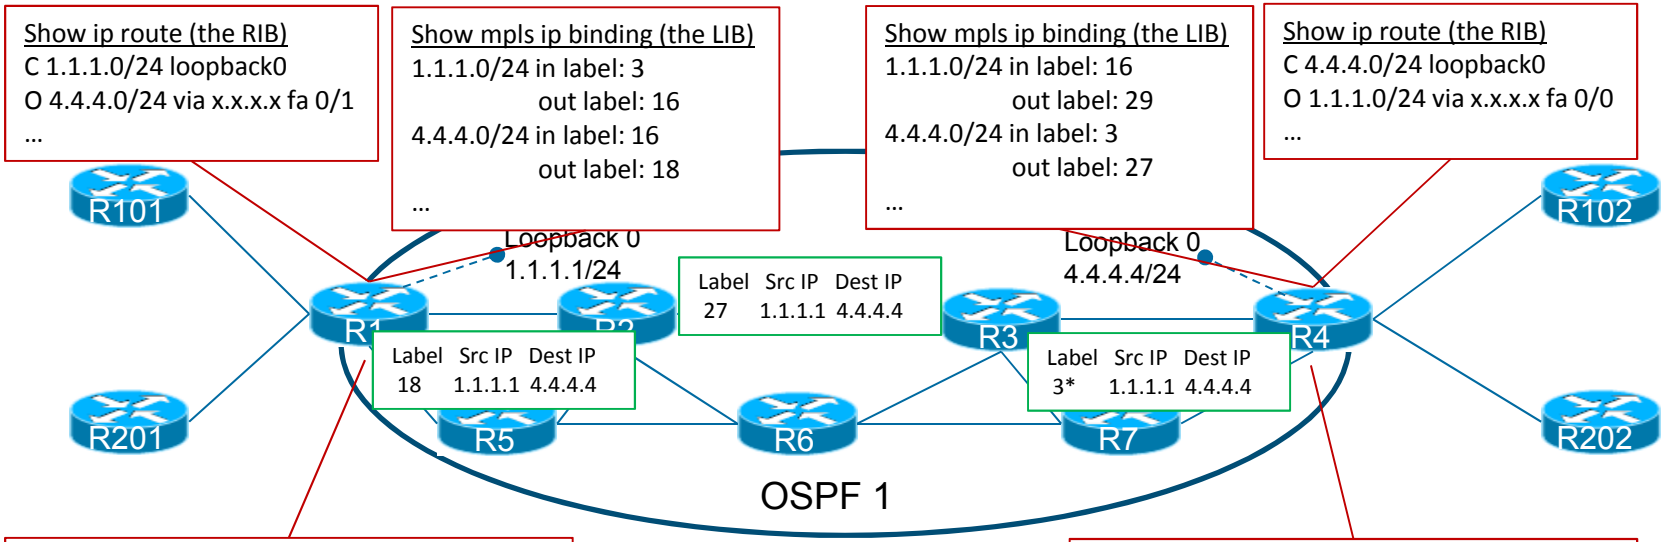

* 3 is the implicit null label, with PHP it gets popped before the end router

MPLS the Beginning - continued

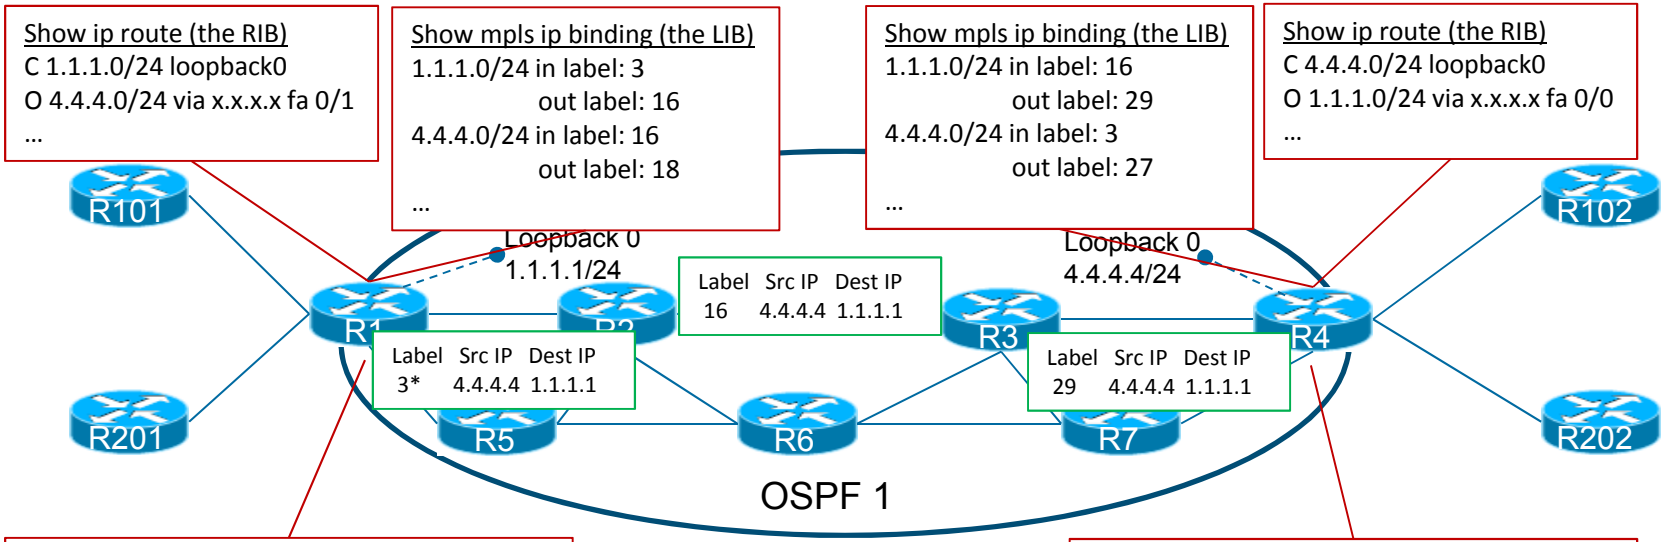

* 3 is the implicit null label, with PHP it gets popped before the end router

Agenda

- Multiprotocol Label Switching Review
- **Virtual Private Networks (without VRFs)**
- Virtual Routing and Forwarding tables (VRF Lite)
- Virtual Private Networks (with VRFs)
- MPLS advanced topics (CCIE twists)
- Future potential adds
- Troubleshooting
- Q&A

Cisco live!

MPLS for BGP

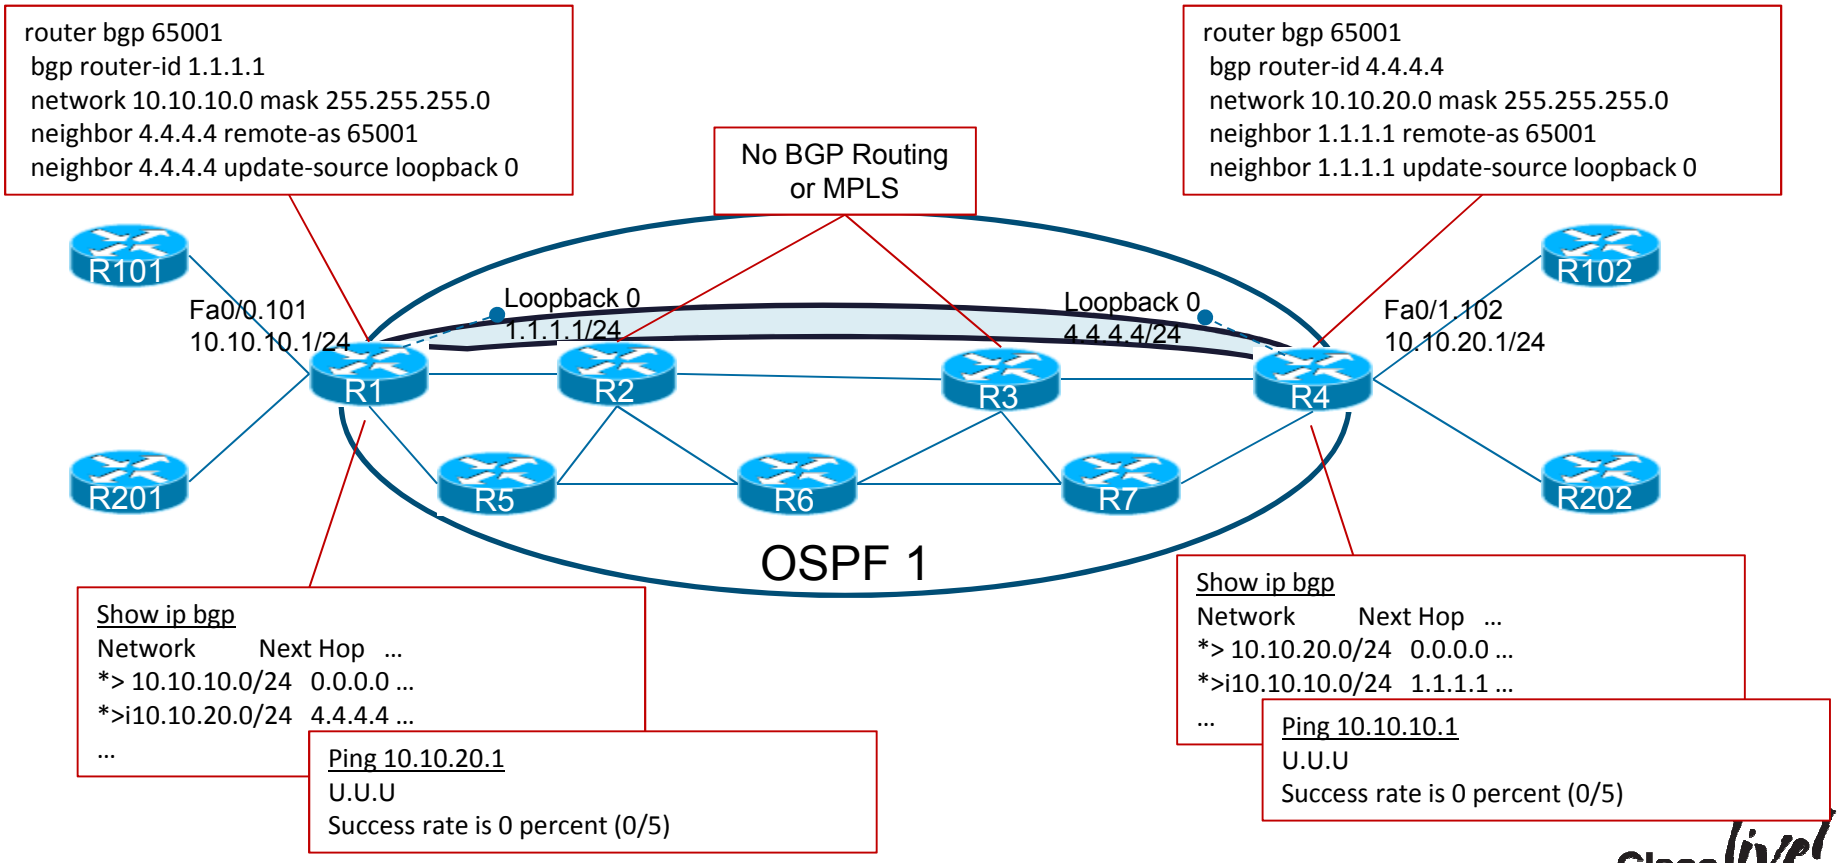

Cisco live!

MPLS for BGP – continued

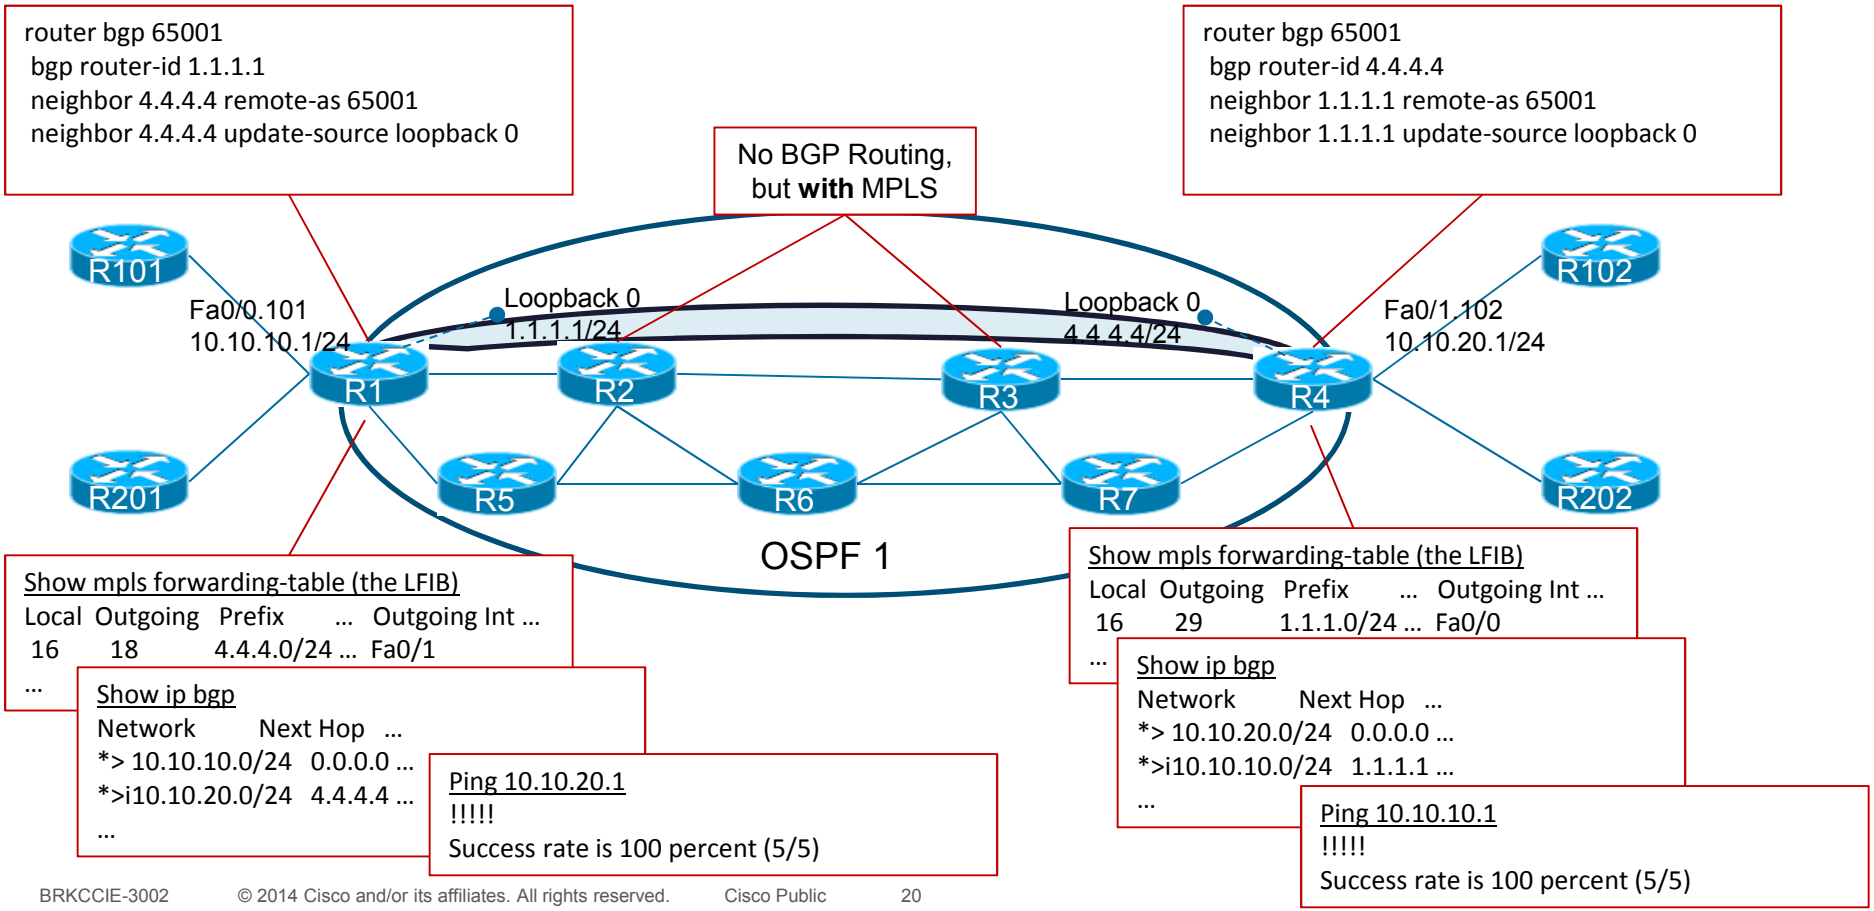

MPLS for BGP – continued

- BGP Generates a new label for the peering
  - Additional labels for each route BGP passes (by default)
    - Use **mpls label mode** command in 15.x to change to per VRF label

Cisco live!

MPLS for BGP – continued

```
router bgp 65001
 bgp router-id 1.1.1.1
 neighbor 4.4.4.4 remote-as 65001
 neighbor 4.4.4.4 update-source loopback 0
```

```
router bgp 65001
 bgp router-id 4.4.4.4
 neighbor 1.1.1.1 remote-as 65001
 neighbor 1.1.1.1 update-source loopback 0
```

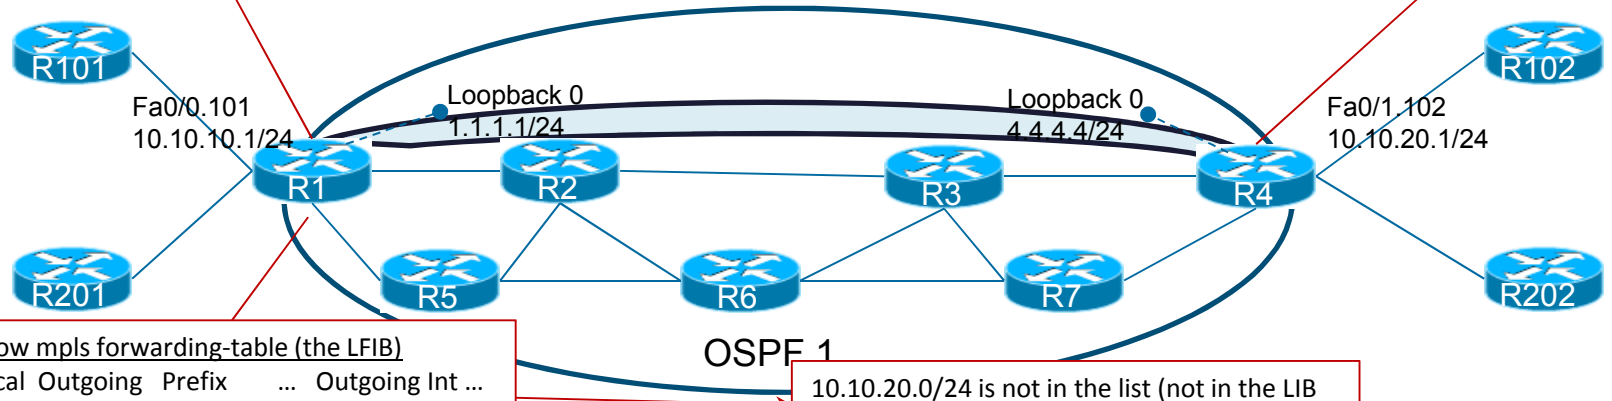

Show mpls forwarding-table (the LFIB)

Local	Outgoing	Prefix	…	Outgoing Int …
16	18	4.4.4.0/24	…	Fa0/1
…				

10.10.20.0/24 is not in the list (not in the LIB or LFIB).

Show ip cef 10.10.20.0

```
10.10.20.0/24,  ….
 tag information from 4.4.4.0/24, shared
   local tag: 21
   fast tag rewrite with Fa0/1, 10.10.12.2, tags imposed: {18}
 …
```

Show mpls forwarding-table 10.10.20.0 255.255.255.0

Local	Outgoing	Prefix	…	Outgoing Int …
16	18	10.10.20.0/24	…	Fa0/1
…				

Agenda

- Multiprotocol Label Switching Review
- Virtual Private Networks (without VRFs)
- **Virtual Routing and Forwarding tables (VRF Lite)**
- Virtual Private Networks (with VRFs)
- MPLS advanced topics (CCIE twists)
- Troubleshooting
- Q&A

Cisco live!

Virtual Routing and Forwarding tables (VRF Lite)

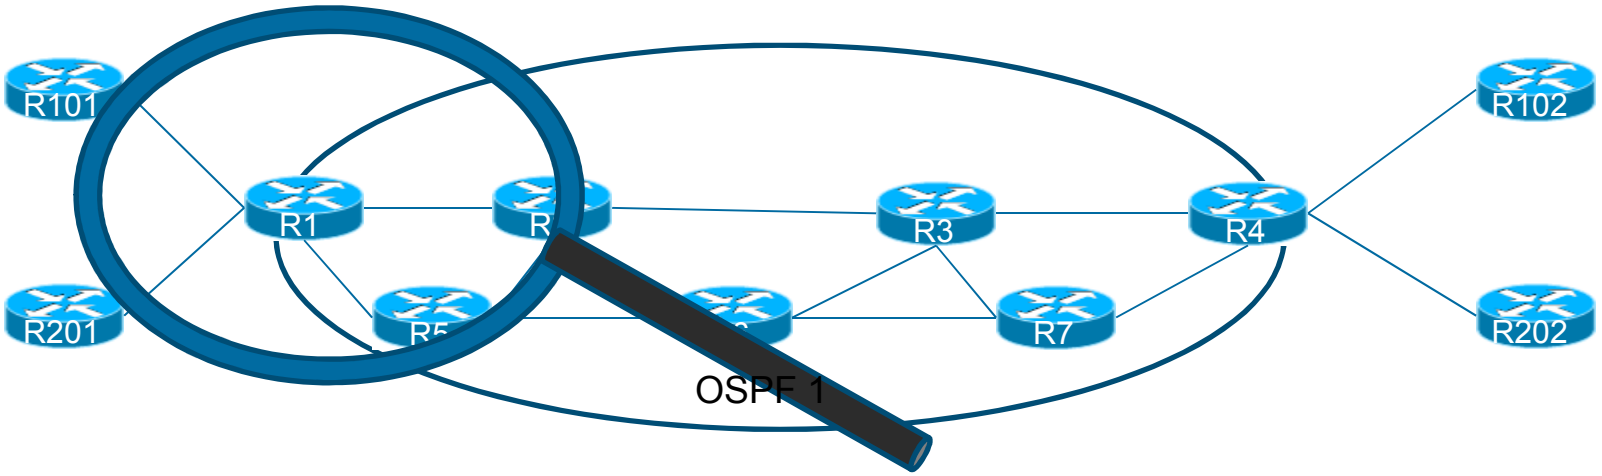

VRF Lite – continued

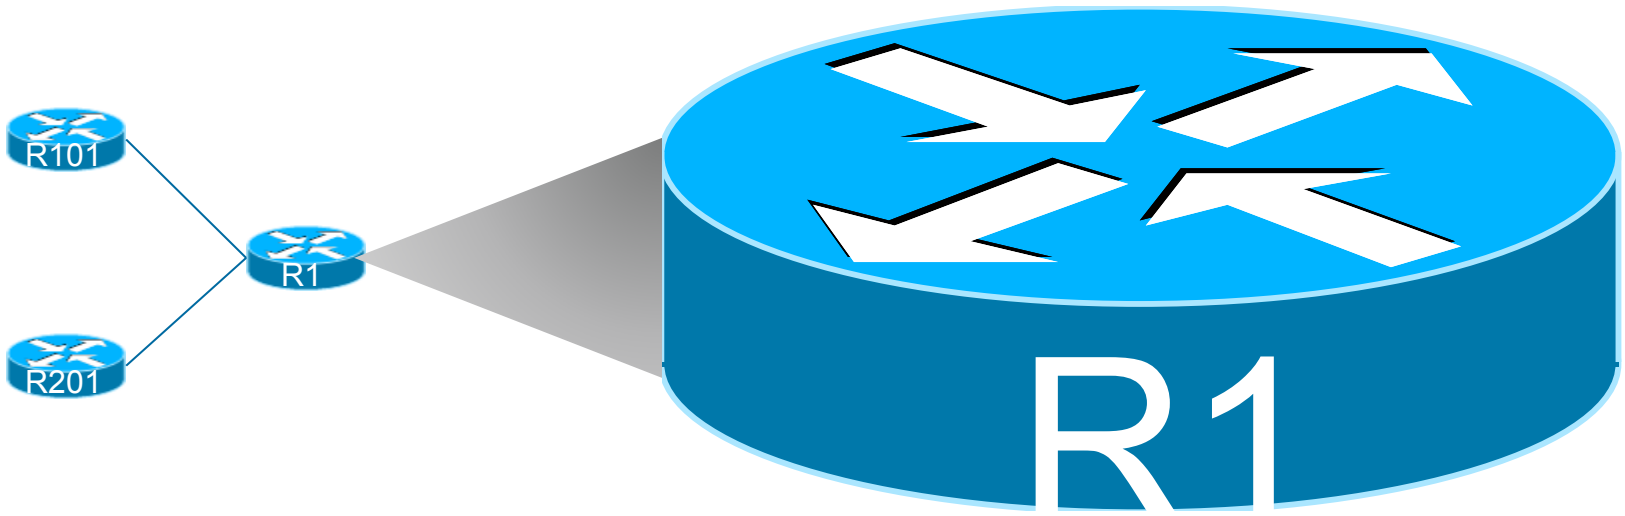

VRF Lite – continued

```
ip vrf CustomerA

Interface FastEthermet0/0.101
 encapsulation dot1q 101
 ip vrf forward CustomerA
 ip address 10.10.10.1 255.255.255.0

Router eigrp 1
 address-family ipv4 vrf CustomerA
 autonomous-system 100
 network 10.0.0.0
```

```
ip vrf CustomerB

Interface FastEthernet0/0.201
 encapsulation dot1q 201
 ip vrf forward CustomerB
 ip address 10.20.10.1 255.255.255.0

Router ospf 2 vrf CustomerB
 router-id 11.11.11.11
 network 10.0.0.0 0.255.255.255 area 10
```

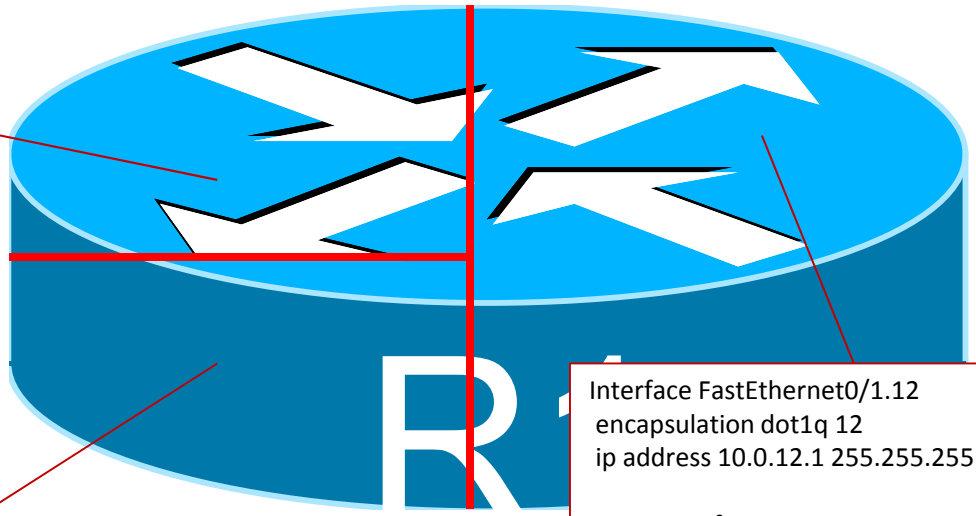

```
Interface FastEthernet0/1.12
 encapsulation dot1q 12
 ip address 10.0.12.1 255.255.255.0

Router ospf 1
 router-id 1.1.1.1
 network 10.0.0.0 0.255.255.255 area 0
 network 1.1.1.1 0.0.0.0 area 0
```

VRF Lite – continued

Show ip route vrf CustomerA
Routing Table: CustomerA
…
Gateway of last resort is not set

10.0.0.0/24 is subnetted, 3 subnets
D 10.10.30.0 [90/4125962] via 10.10.10.2, 01:50:08, FastEthernet0/0.101
D 10.10.60.0 [90/4125962] via 10.10.10.2, 01:50:08, FastEthernet0/0.101
C 10.10.10.0 is directly connected, FastEthernet0/0.101

Show ip route vrf CustomerB
Routing Table: CustomerB
…
Gateway of last resort is not set

10.0.0.0/24 is subnetted, 3 subnets
O 10.20.30.0 [110/762] via 10.20.10.2, 00:28:08, FastEthernet0/0.201
O IA 10.20.60.0 [110/762] via 10.20.10.2, 00:28:08, FastEthernet0/0.201
C 10.20.10.0 is directly connected, FastEthernet0/0.201

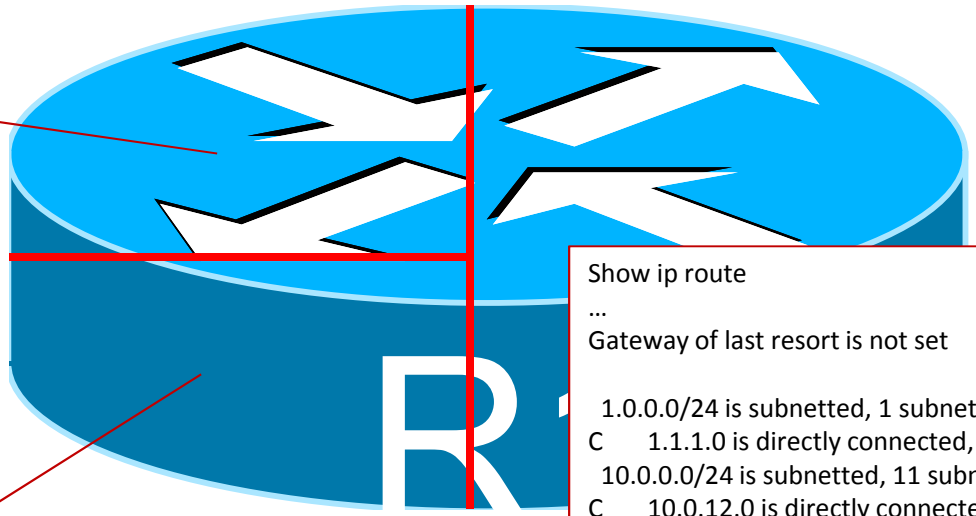

Show ip route
…
Gateway of last resort is not set

1.0.0.0/24 is subnetted, 1 subnets
C 1.1.1.0 is directly connected, Loopback0
10.0.0.0/24 is subnetted, 11 subnets
C 10.0.12.0 is directly connected, FastEthernet0/1.12
C 10.0.15.0 is directly connected, FastEthernet0/1.15
O 10.….

VRF Lite – continued

```
router rip
  address-family ipv4 vrf name
            version 2
            network x.x.x.x
            redistribute protocol
            distribute-list  list …
```

```
router eigrp as-number
 address-family ipv4 vrf name
            autonomous-system #
            network x.x.x.x
```

```
router bgp as-#
 address-family ipv4 vrf name
            neighbor n.n.n.n remote-as as-#
            network x.x.x.x
```

```
router ospf # vrf name
            router-id i.i.i.i
            network x.x.x.x w.w.w.w area #
```

```
Ip route vrf name t.t.t.t m.m.m.m exgress-int | next-hop [global]
```


Cisco live!

Agenda

- Multiprotocol Label Switching Review
- Virtual Private Networks (without VRFs)
- Virtual Routing and Forwarding tables (VRF Lite)
- **Virtual Private Networks (with VRFs)**
- MPLS advanced topics (CCIE twists)
- Future potential adds
- Troubleshooting
- Q&A

Cisco live!

Layer 3 VPN

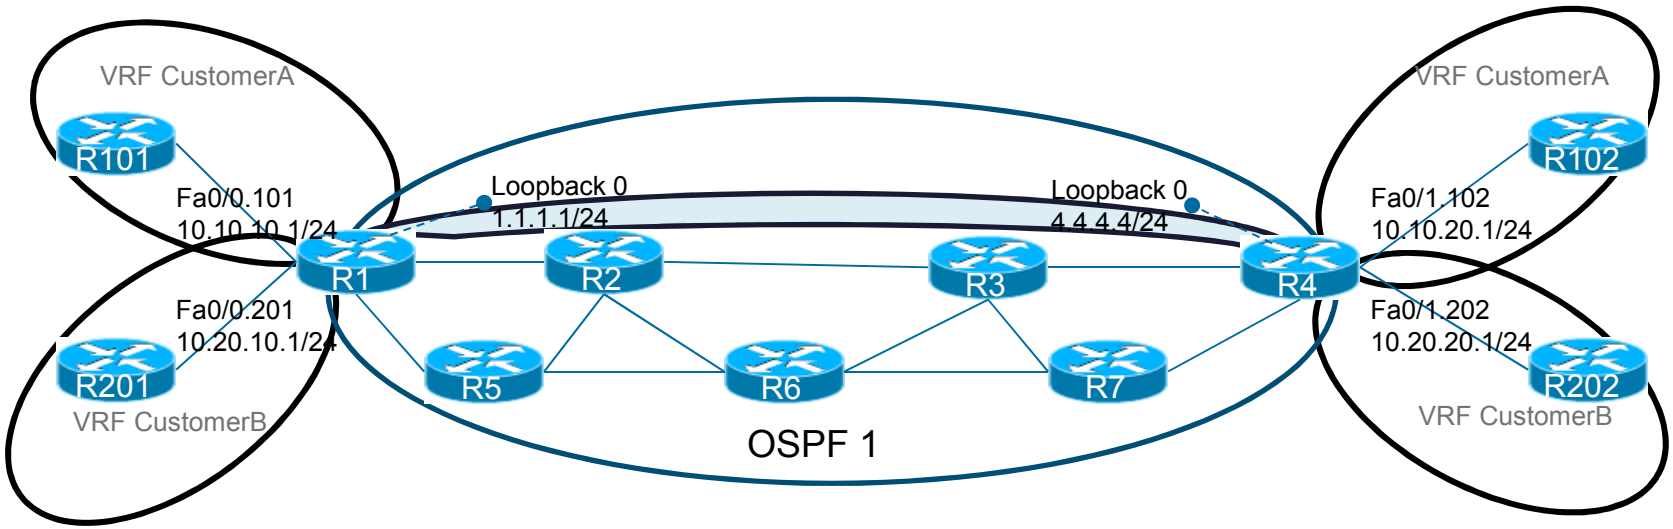

Layer 3 VPN – continued

```
ip vrf CustomerA
 rd 1:100
 route-target export 1:100
 route-target import 1:100
```

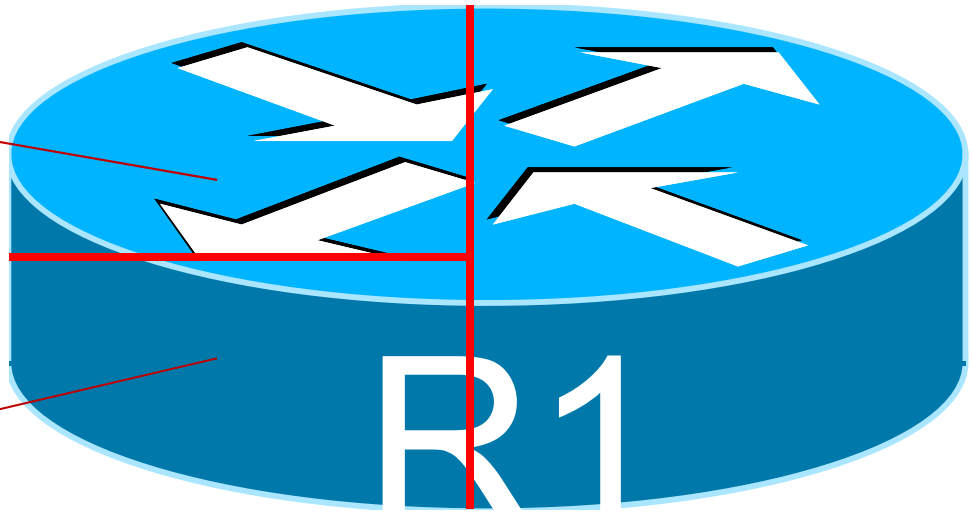


```
ip vrf CustomerB
 rd 1:200
 route-target export 1:200
 route-target import 1:200
```

Layer 3 VPN – continued

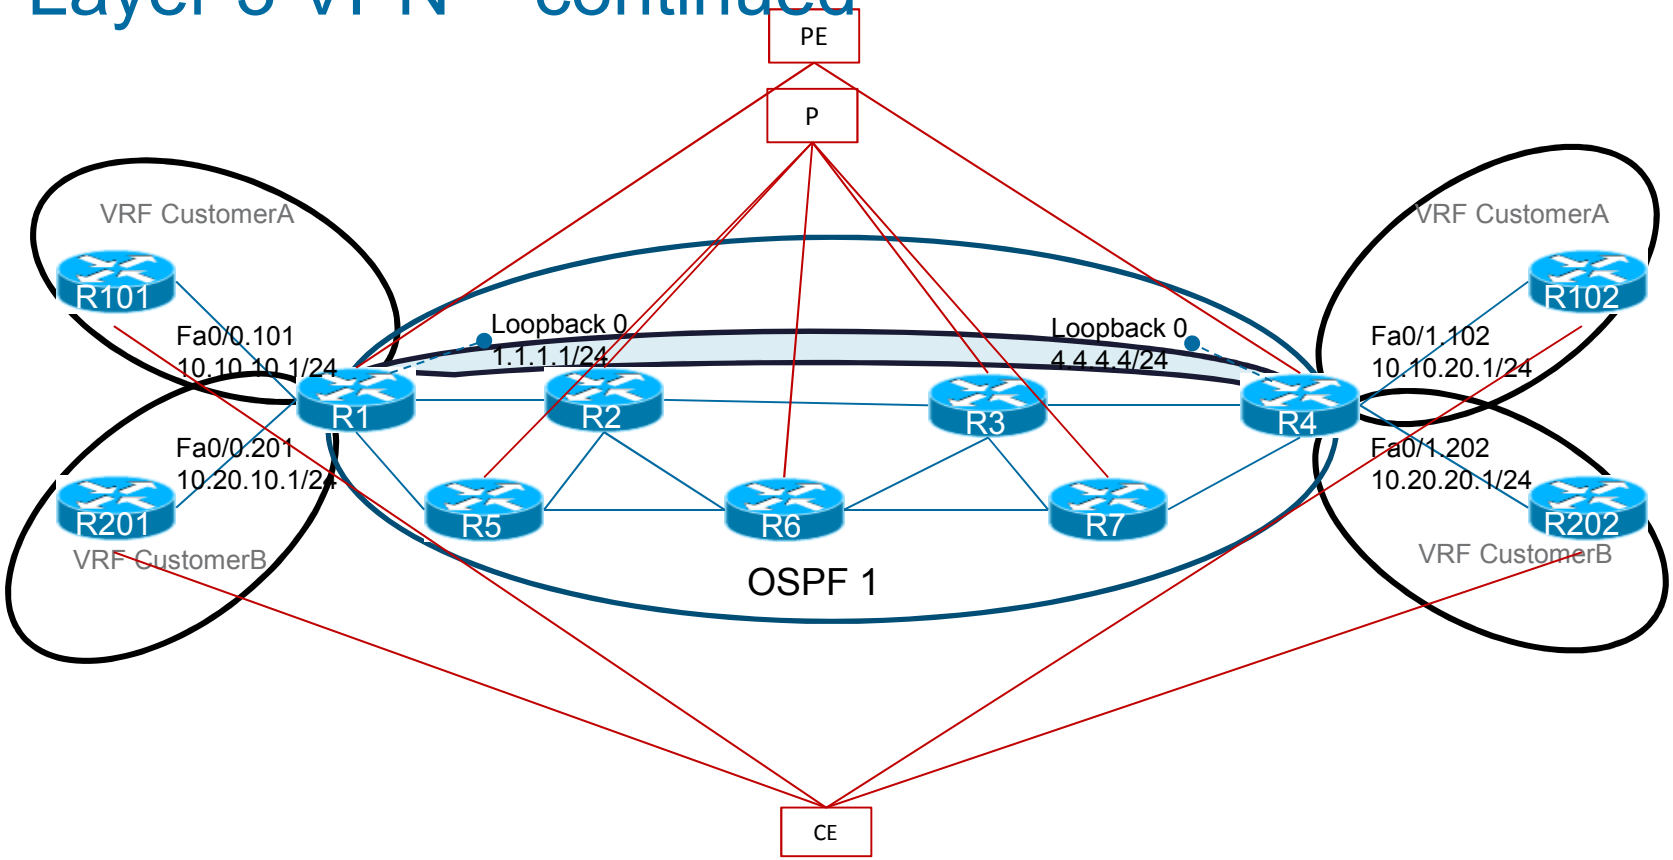

Layer 3 VPN – continued

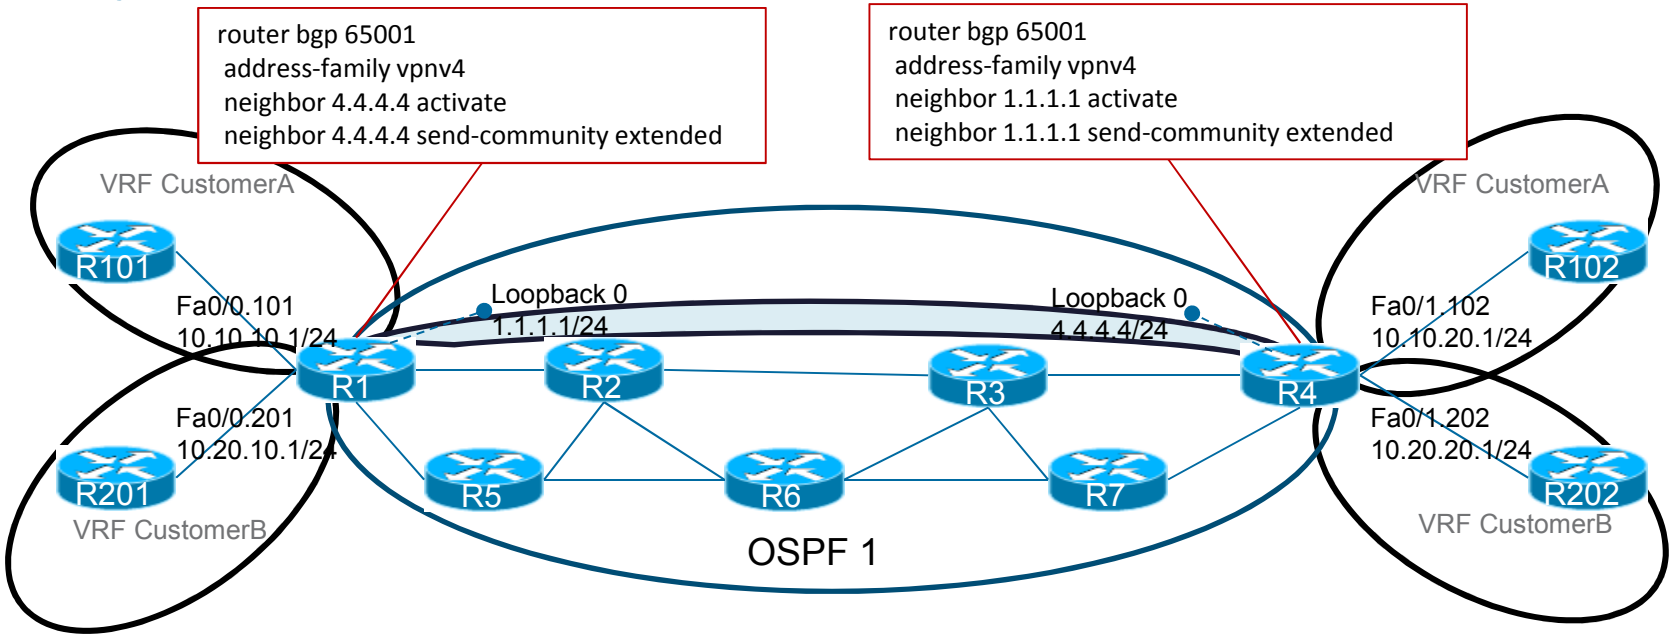

Layer 3 VPN – continued

```
router bgp 65001
address-family vpnv4
 neighbor 4.4.4.4 activate
 neighbor 4.4.4.4 send-community extended
 address-family ipv4 vrf CustomerA
  redistribute eigrp 1
 address-family ipv4 vrf CustomerB
  redistribute ospf 2
```

```
router bgp 65001
address-family vpnv4
 neighbor 1.1.1.1 activate
 neighbor 1.1.1.1 send-community extended
 address-family ipv4 vrf CustomerA
  redistribute eigrp 1
 address-family ipv4 vrf CustomerB
  redistribute ospf 2
```

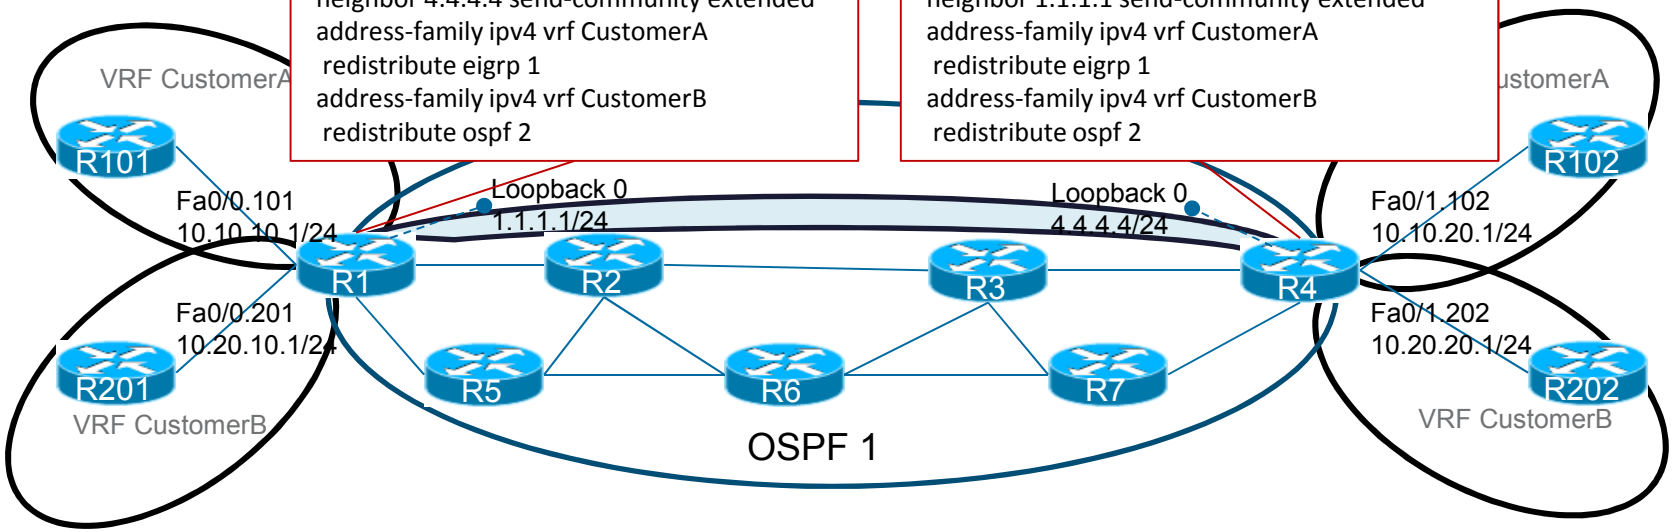

Layer 3 VPN – continued

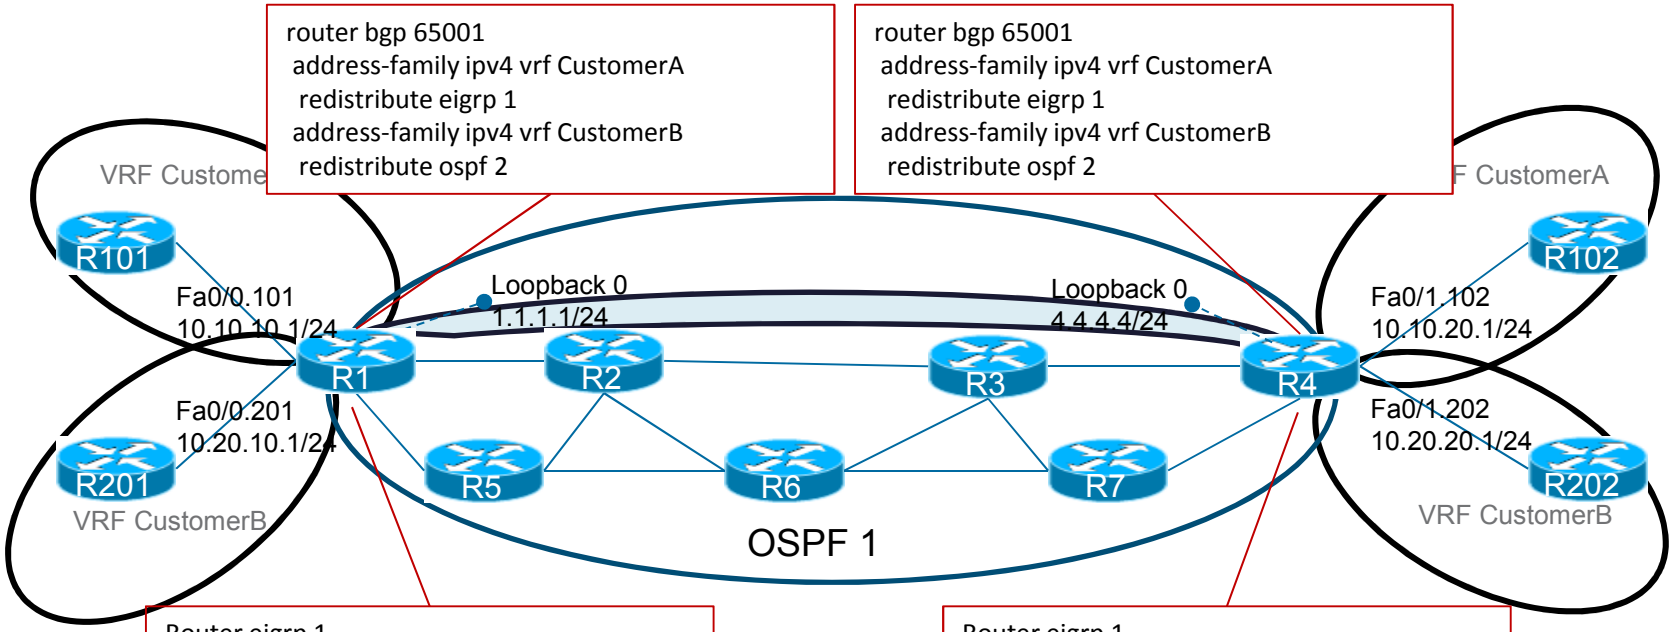

Layer 3 VPN – continued

```
R1#sh ip bgp vpn vrf CustomerA 10.10.20.0
BGP routing table entry for 1:100:10.10.20.0/24, version 8
Paths: (1 available, best #1, table CustomerA)
  …
   4.4.4.4 (metric 3) from 4.4.4.4 (4.4.4.4)
     Origin incomplete, metric 0, localpref 100, valid, external, best
…
     mpls labels in/out nolabel/29
```

```
R1#sh mpls forward 10.10.20.0 255.255.255.0
Local  Outgoing    Prefix          Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id    switched   interface
R1#
```

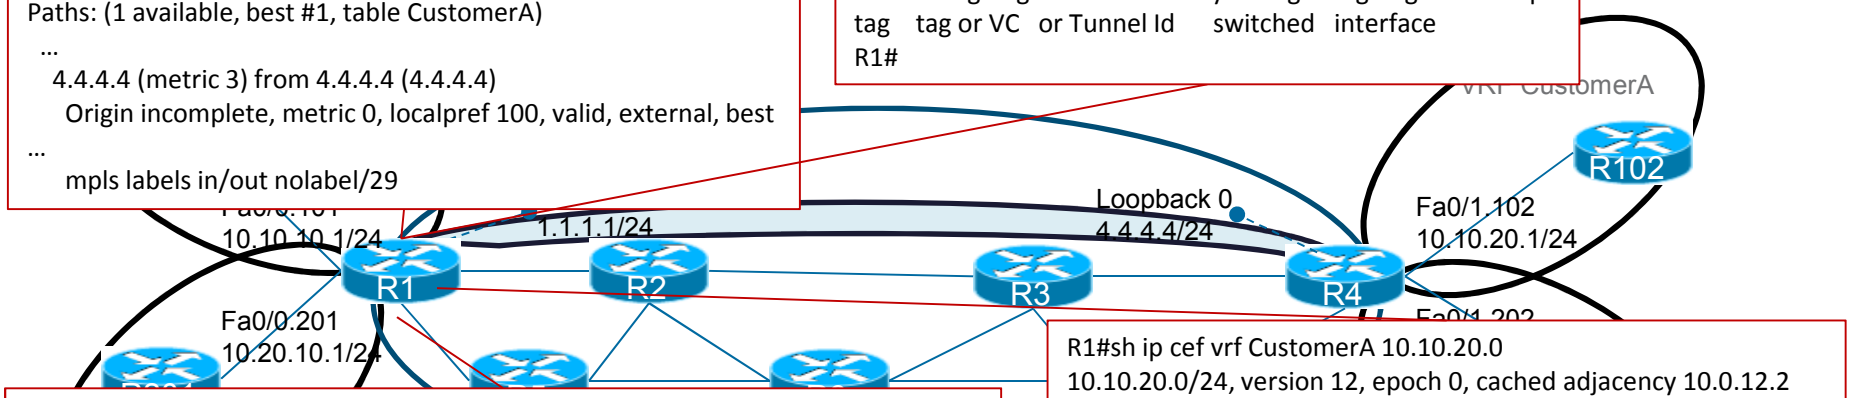


```
R1#show mpls forwarding-table vrf CustomerA 10.10.20.0 255.255.255.0
Local  Outgoing    Prefix          Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id    switched   interface
None   29          10.10.20.0/24   0          Fa0/1.35    144.44.35.5

R1#show mpls forwarding-table vrf CustomerA 10.10.20.0 255.255.255.0 detail
Local  Outgoing    Prefix          Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id    switched   interface
None   29          10.10.20.0/24   0          Fa0/1.12    10.0.12.2
       MAC/Encaps=18/26, MRU=1496, Tag Stack{18 29}
…
```

```
R1#sh ip cef vrf CustomerA 10.10.20.0
10.10.20.0/24, version 12, epoch 0, cached adjacency 10.0.12.2
0 packets, 0 bytes
  tag information set
    local tag: VPN-route-head
    fast tag rewrite with Fa0/1.12, 10.0.12.2, tags imposed: {18 29}
  via 4.4.4.4, 0 dependencies, recursive
    next hop 10.0.12.2, FastEthernet0/1.12 via 10.0.34.0/24
    valid cached adjacency
    tag rewrite with Fa0/1.12, 10.0.12.2, tags imposed: {18 29}
R1#
```

Layer 3 VPN – continued

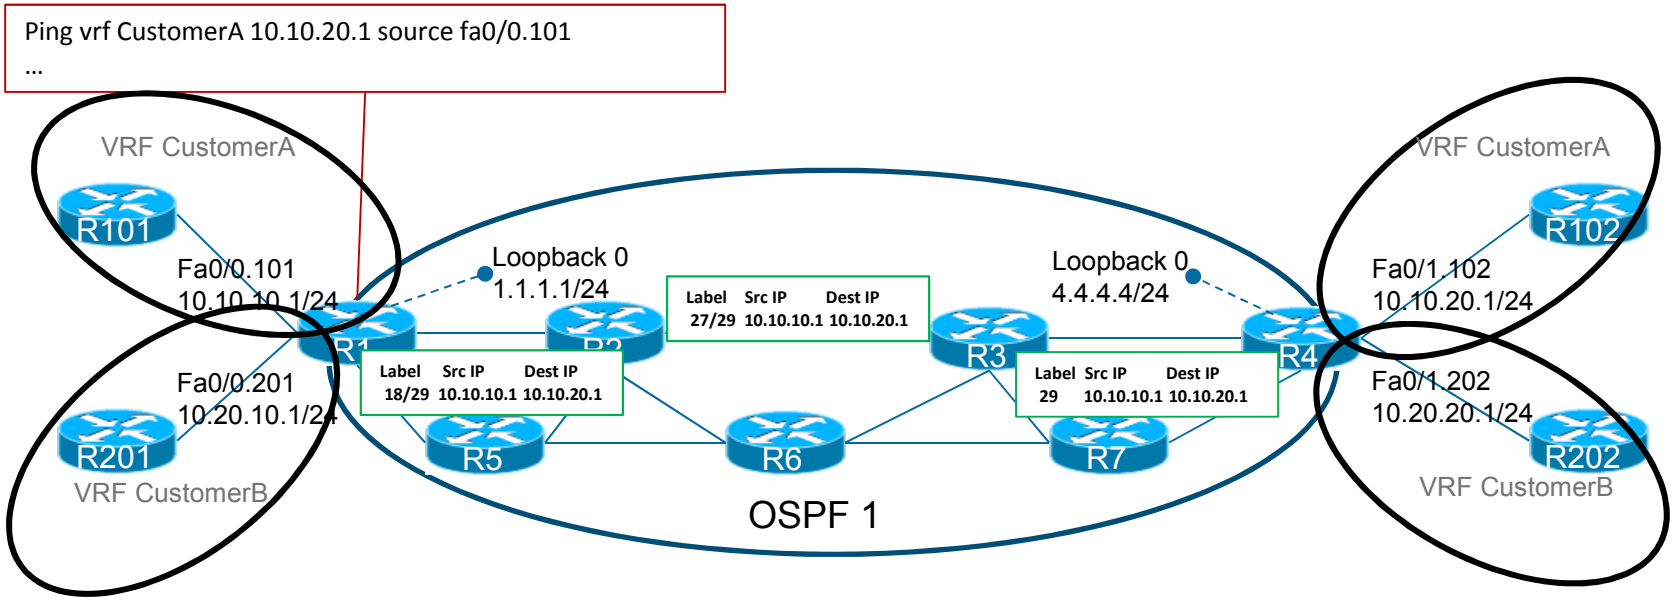

Layer 3 VPN – continued

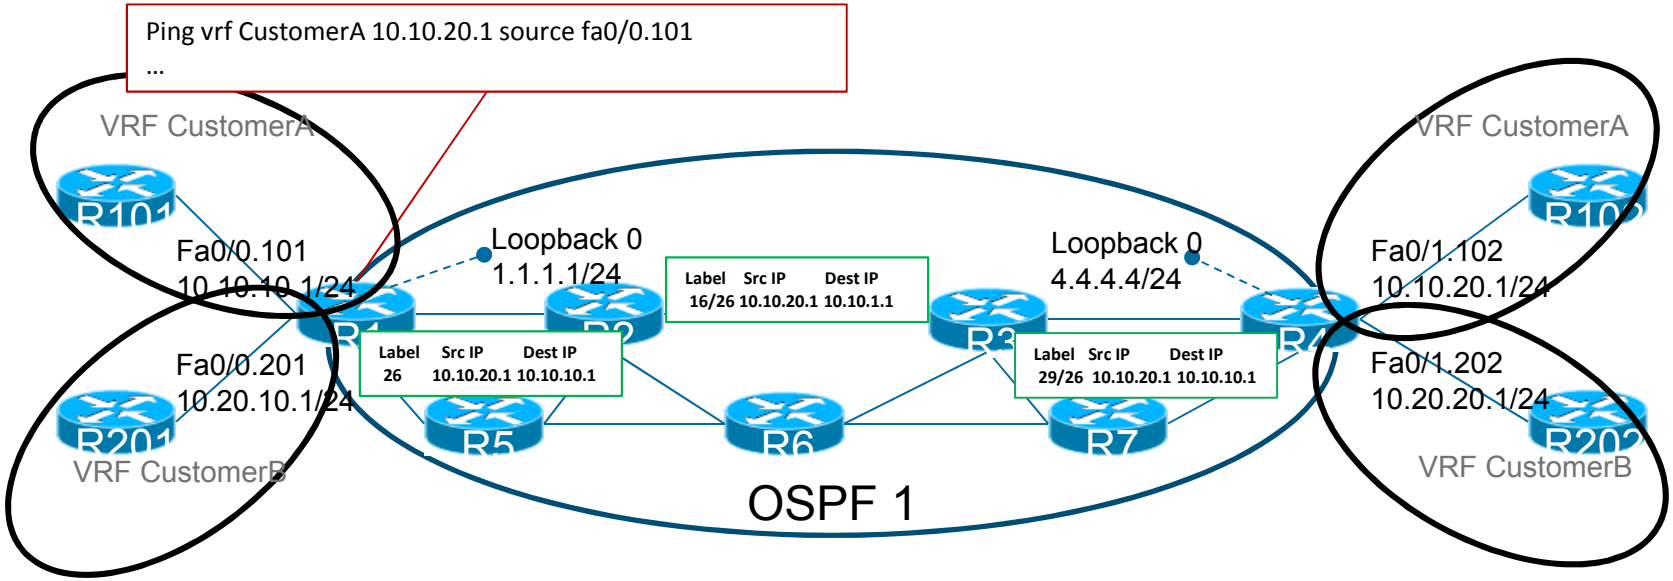

Layer 3 VPN – continued

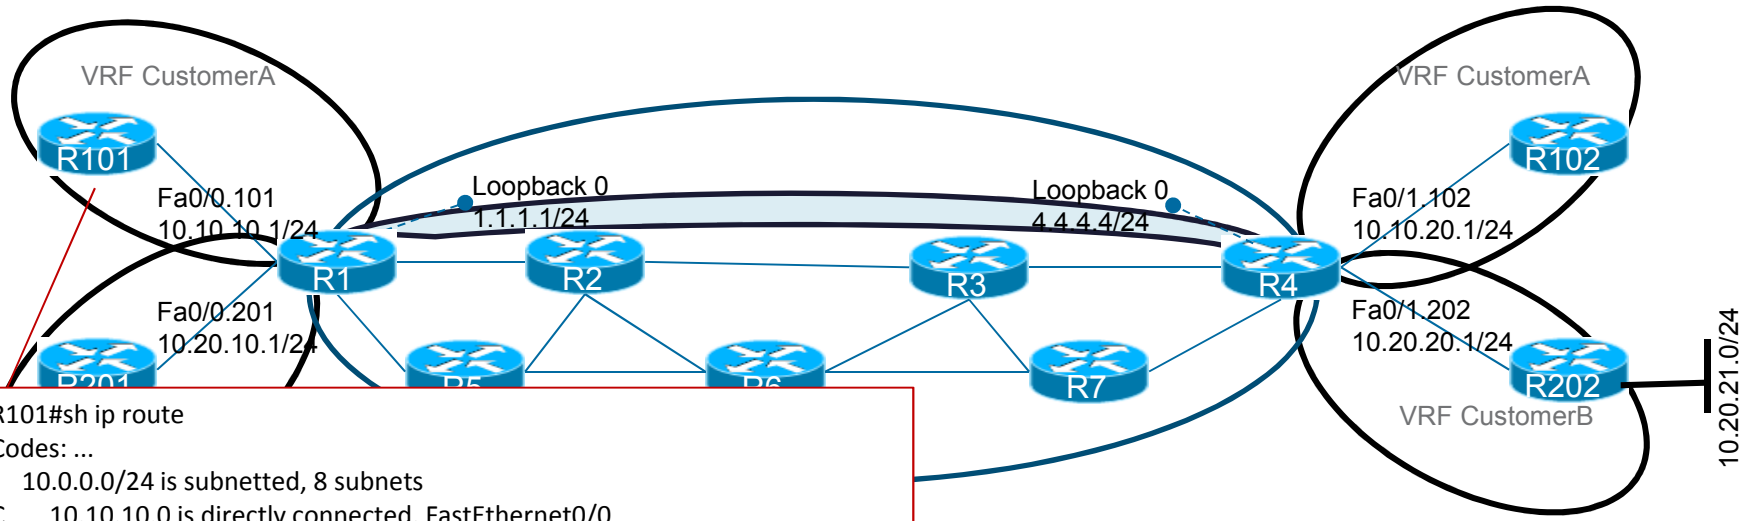


```
R101#sh ip route
Codes: ...
    10.0.0.0/24 is subnetted, 8 subnets
C       10.10.10.0 is directly connected, FastEthernet0/0
C       10.10.11.0 is directly connected, FastEthernet0/1
C       10.10.12.0 is directly connected, FastEthernet0/2
C       10.10.13.0 is directly connected, FastEthernet0/3
D       10.10.20.0 [90/30720] via 10.10.10.1, 00:41:24, FastEthernet0/0
D       10.10.21.0 [90/158720] via 10.10.10.1, 00:41:24, FastEthernet0/0
D       10.10.22.0 [90/158720] via 10.10.10.1, 00:41:24, FastEthernet0/0
D       10.10.23.0 [90/158720] via 10.10.10.1, 00:41:25, FastEthernet0/0
```

Layer 3 VPN – continued

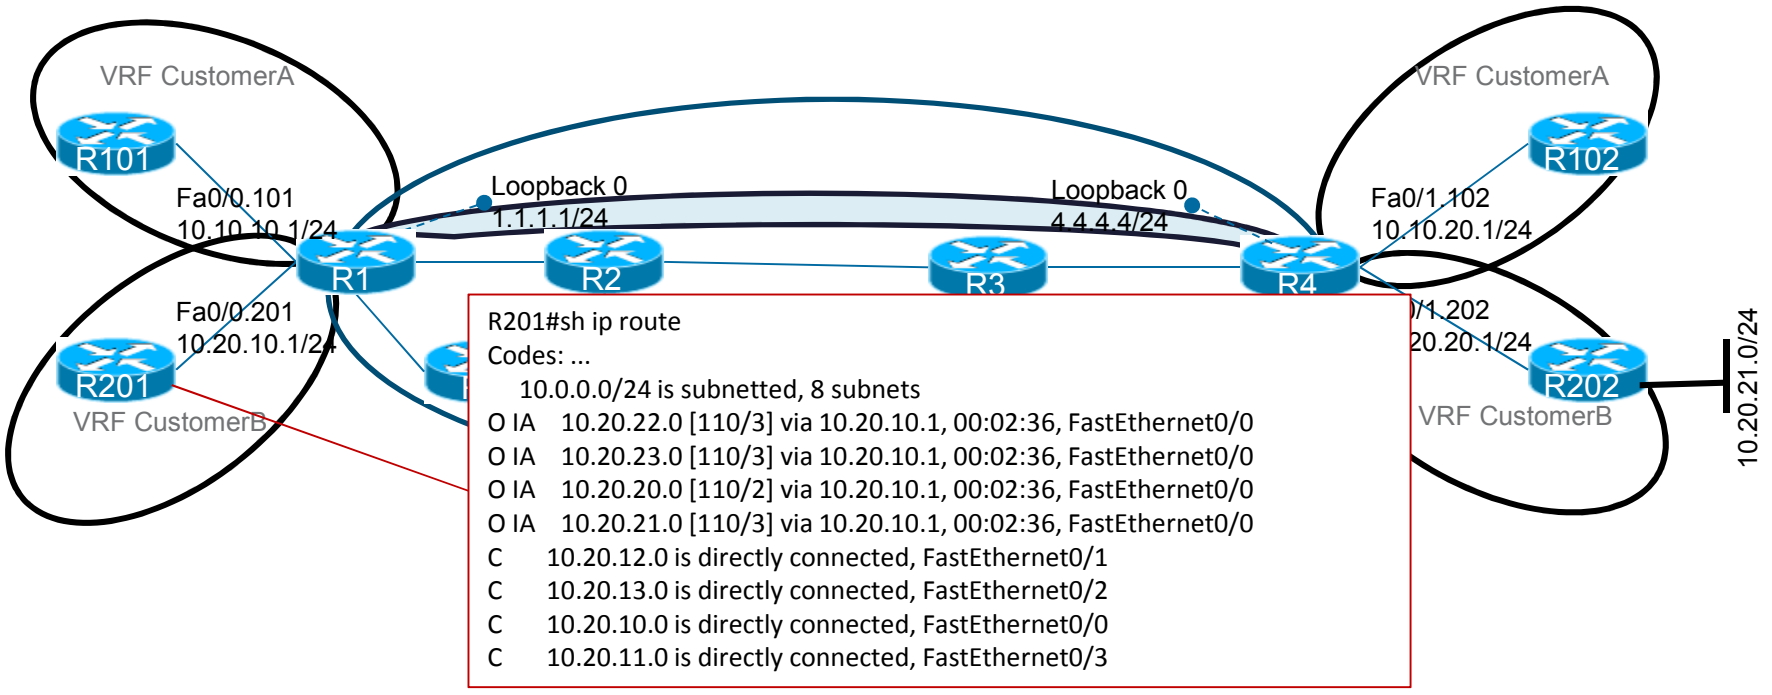

Cisco live!

Agenda

- Multiprotocol Label Switching Review
- Virtual Private Networks (without VRFs)
- Virtual Routing and Forwarding tables (VRF Lite)
- Virtual Private Networks (with VRFs)
- **MPLS advanced topics (CCIE twists)**
- Troubleshooting
- Q&A

Cisco live!

Sham Links

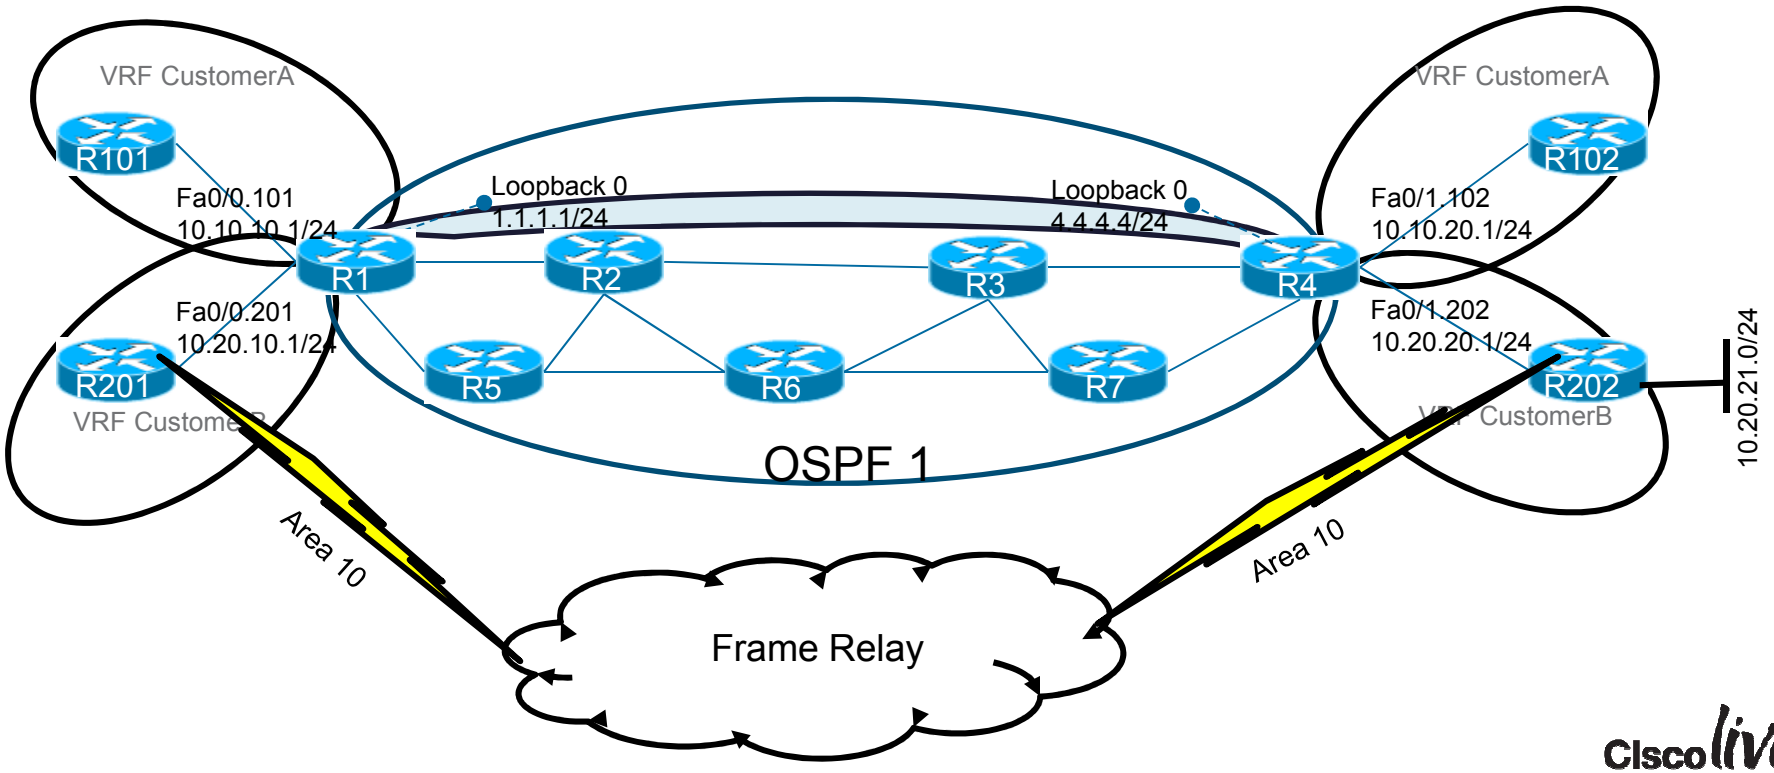

Sham Links – continued

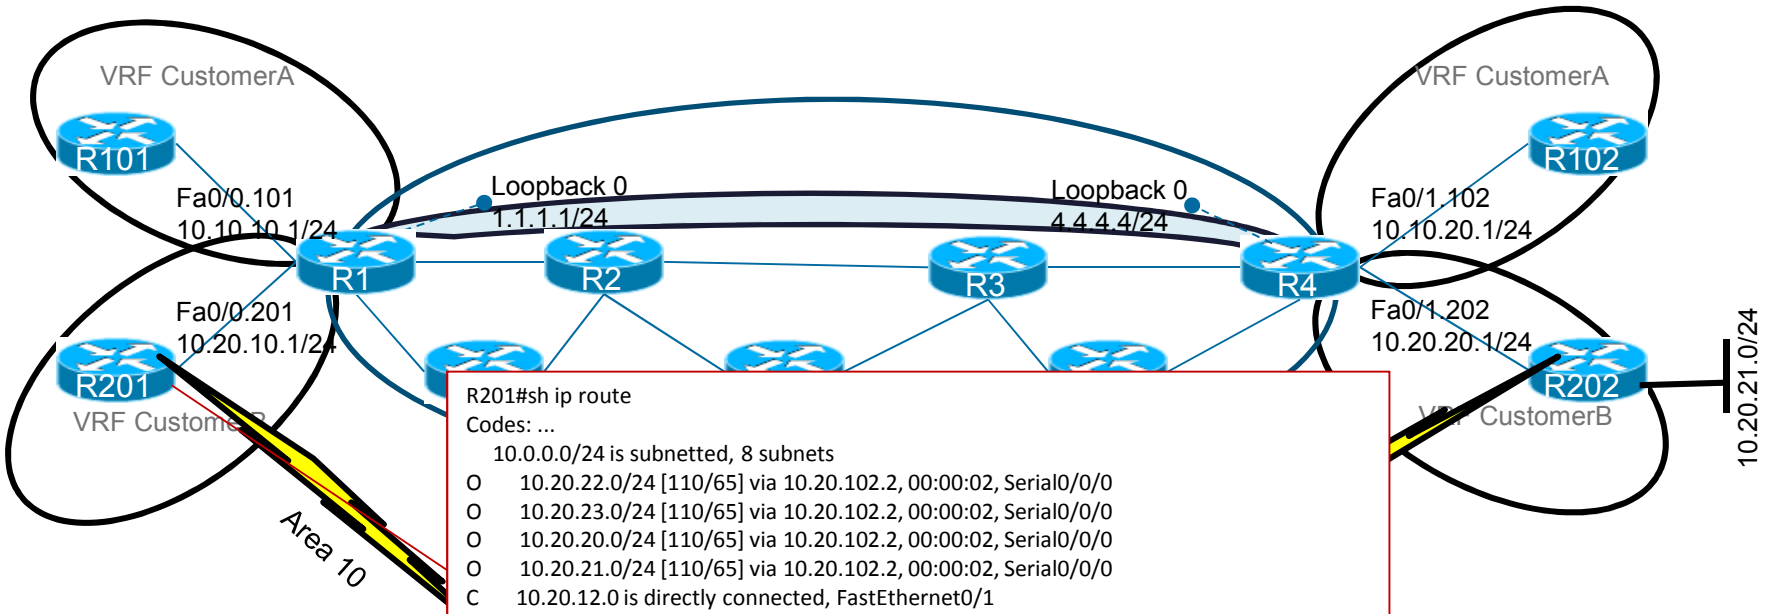


```
R201#sh ip route
Codes: ...
    10.0.0.0/24 is subnetted, 8 subnets
O      10.20.22.0/24 [110/65] via 10.20.102.2, 00:00:02, Serial0/0/0
O      10.20.23.0/24 [110/65] via 10.20.102.2, 00:00:02, Serial0/0/0
O      10.20.20.0/24 [110/65] via 10.20.102.2, 00:00:02, Serial0/0/0
O      10.20.21.0/24 [110/65] via 10.20.102.2, 00:00:02, Serial0/0/0
C      10.20.12.0 is directly connected, FastEthernet0/1
C      10.20.13.0 is directly connected, FastEthernet0/2
C      10.20.10.0 is directly connected, FastEthernet0/0
C      10.20.11.0 is directly connected, FastEthernet0/3
C      10.20.102.0/24 is directly connected, Serial0/0/0
O      10.20.102.2/32 [110/64] via 10.20.102.2, 00:00:03, Serial0/0/0
```

Sham Links – continued

```
R1#
Interface loopback 1
 ip vrf forward CustomerB
 ip address 14.14.14.1 255.255.255.255
router ospf 2 vrf CustomerB
 area 10 sham-link 14.14.14.1 14.14.14.4
router bgp 65001
 address-family ipv4 vrf CustomerB
 network 14.14.14.1 mask 255.255.255.255
```

```
R1#sh ip ospf sham-links
Sham Link OSPF_SL0 to address 14.14.14.4 is up
Area 10 source address 14.14.14.1
  Run as demand circuit
  DoNotAge LSA allowed. Cost of using 1 State POINT_TO_POINT,
  Timer intervals configured, Hello 10, Dead 40, Wait 40,
    Hello due in 00:00:02
    Adjacency State FULL (Hello suppressed)
 Index 2/2, retransmission queue length 0, number of retransmission 0
    First 0x0(0)/0x0(0) Next 0x0(0)/0x0(0)
    Last retransmission scan length is 0, maximum is 0
    Last retransmission scan time is 0 msec, maximum is 0 msec
```

```
R4#
Interface loopback 1
 ip vrf forward CustomerB
 ip address 14.14.14.4 255.255.255.255
router ospf 2 vrf CustomerB
 area 10 sham-link 14.14.14.4 14.14.14.1
router bgp 65001
 address-family ipv4 vrf CustomerB
 network 14.14.14.4 mask 255.255.255.255
```

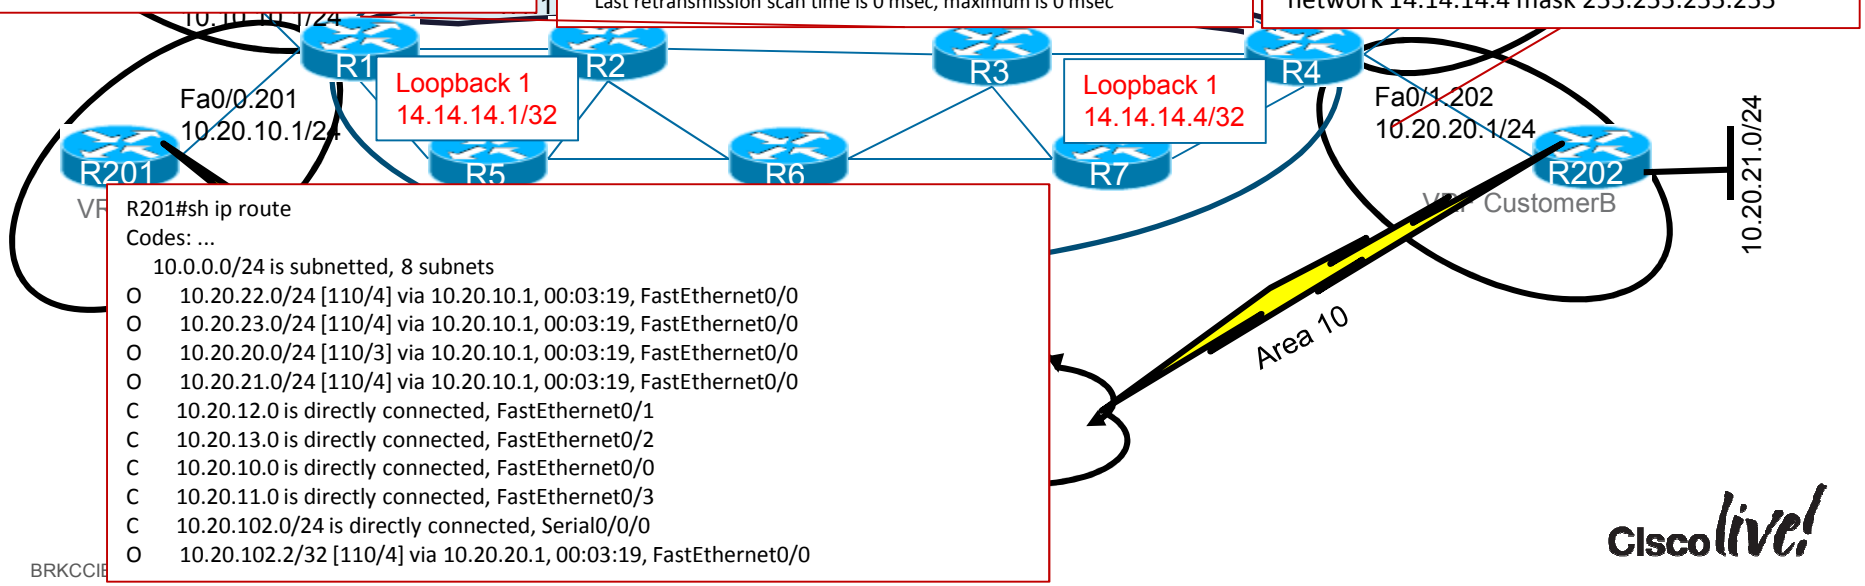


```
R201#sh ip route
Codes: ...
    10.0.0.0/24 is subnetted, 8 subnets
O      10.20.22.0/24 [110/4] via 10.20.10.1, 00:03:19, FastEthernet0/0
O      10.20.23.0/24 [110/4] via 10.20.10.1, 00:03:19, FastEthernet0/0
O      10.20.20.0/24 [110/3] via 10.20.10.1, 00:03:19, FastEthernet0/0
O      10.20.21.0/24 [110/4] via 10.20.10.1, 00:03:19, FastEthernet0/0
C      10.20.12.0 is directly connected, FastEthernet0/1
C      10.20.13.0 is directly connected, FastEthernet0/2
C      10.20.10.0 is directly connected, FastEthernet0/0
C      10.20.11.0 is directly connected, FastEthernet0/3
C      10.20.102.0/24 is directly connected, Serial0/0/0
O      10.20.102.2/32 [110/4] via 10.20.20.1, 00:03:19, FastEthernet0/0
```

The Configs

Cisco live!

Router 1

```
R1#sh run
!
hostname R1
!
ip cef
!
ip vrf CustomerA
 rd 1:100
 route-target export 1:100
 route-target import 1:100
!
ip vrf CustomerB
 rd 1:200
 route-target export 1:200
 route-target import 1:200
!
interface Loopback0
 ip address 1.1.1.1 255.255.255.0
 ip ospf network point-to-point
!
interface Loopback1
 ip vrf forwarding CustomerB
 ip address 14.14.14.1 255.255.255.255
!
interface FastEthernet0/0
 no ip address
 duplex auto
 speed auto
!
```

```
interface FastEthernet0/0.101
 encapsulation dot1Q 101
 ip vrf forwarding CustomerA
 ip address 10.10.10.1 255.255.255.0
!
interface FastEthernet0/0.201
 encapsulation dot1Q 201
 ip vrf forwarding CustomerB
 ip address 10.20.10.1 255.255.255.0
!
interface FastEthernet0/1
 no ip address
 duplex auto
 speed auto
!
interface FastEthernet0/1.12
 encapsulation dot1Q 12
 ip address 10.0.12.1 255.255.255.0
 mpls ip
!
interface FastEthernet0/1.15
 encapsulation dot1Q 15
 ip address 10.0.15.1 255.255.255.0
 mpls ip
!
router eigrp 1
 auto-summary
 !
 address-family ipv4 vrf CustomerA
  redistribute bgp 65001 metric 1 1 1 1 1
  network 10.0.0.0
  auto-summary
  autonomous-system 100
 exit-address-family
```

```
router ospf 2 vrf CustomerB
 router-id 11.11.11.11
 log-adjacency-changes
 area 10 sham-link 14.14.14.1 14.14.14.4
 redistribute bgp 65001 subnets
 network 10.0.0.0 0.255.255.255 area 10
!
router ospf 1
 router-id 1.1.1.1
 log-adjacency-changes
 network 1.0.0.0 0.255.255.255 area 0
 network 10.0.0.0 0.255.255.255 area 0
!
router bgp 65001
 no synchronization
 bgp router-id 1.1.1.1
 bgp log-neighbor-changes
 neighbor 4.4.4.4 remote-as 65001
 neighbor 4.4.4.4 update-source Loopback0
 no auto-summary
 !
 address-family vpnv4
  neighbor 4.4.4.4 activate
  neighbor 4.4.4.4 send-community extended
 exit-address-family
 !
 address-family ipv4 vrf CustomerB
  redistribute ospf 2 vrf CustomerB
  redistribute eigrp 100
  no synchronization
  network 14.14.14.1 mask 255.255.255.255
 exit-address-family
 !
 address-family ipv4 vrf CustomerA
  redistribute eigrp 100
  no synchronization
 exit-address-family
```

Router 2, Router 101 and Router 201

```
R2#sh run
!
hostname R2
!
ip cef
!
interface FastEthernet0/0
 ip address 10.0.12.2 255.255.255.0
 mpls ip
!
interface FastEthernet0/1
 ip address 10.0.23.2 255.255.255.0
 mpls ip
!
interface FastEthernet1/0
 ip address 10.0.25.2 255.255.255.0
 mpls ip
!
interface FastEthernet1/1
 ip address 10.0.26.2 255.255.255.0
 mpls ip
!
router ospf 1
 router-id 2.2.2.2
 log-adjacency-changes
 network 10.0.0.0 0.255.255.255 area 0
!
end
```

```
R101#sh run
!
hostname R101
!
ip cef
!
interface FastEthernet0/1
 ip address 10.10.12.1 255.255.255.0
!
interface FastEthernet0/2
 ip address 10.10.13.1 255.255.255.0
!
interface FastEthernet0/3
 ip address 10.10.11.1 255.255.255.0
!
interface FastEthernet0/0
 ip address 10.10.10.2 255.255.255.0
!
!
router eigrp 100
  network 10.0.0.0
  no auto-summary
!
end
```

```
R201#sh run
!
hostname R201
!
ip cef
!
interface FastEthernet0/0
 ip address 10.20.10.2 255.255.255.0
!
interface FastEthernet0/1
 ip address 10.20.11.1 255.255.255.0
!
interface FastEthernet0/2
 ip address 10.20.12.1 255.255.255.0
!
interface FastEthernet0/3
 ip address 10.20.13.1 255.255.255.0
!
interface Serial0/0/0
 ip address 10.20.102.1 255.255.255.0
 encapsulation frame-relay
 ip ospf network point-to-multipoint
!
router ospf 1
 router-id 201.201.201.201
 log-adjacency-changes
 network 10.0.0.0 0.255.255.255 area 10
!
end
```

Static MPLS Label Assignment

- Static bindings between labels and IPv4 prefixes can be configured to support MPLS hop-by-hop forwarding.
- Router(config)# **mpls label range** *min-label max-label* [**static** *min-static-label max-static-label*]
  - Default is no labels reserved for static assignment
- Router(config)# **mpls static binding ipv4** *prefix mask* [**input** | **output** *nexthop*] *label*
  - Bindings specified are installed automatically in the MPLS forwarding table as routing demands.

Cisco live!

MPLS LDP Authentication

- The MPLS LDP—Lossless MD5 Session Authentication feature enables a Label Distribution Protocol (LDP) session to be password-protected without tearing down and reestablishing the LDP session.

```
key chain MyKey
key 1
key-string cisco
send-lifetime 10:00:00 Nov 2 2013 10:00:00 Dec 2 2015
accept-lifetime 00:00:00 Jan 1 1970 duration 1
key 2
key-string cisco2
send-lifetime 00:00:00 Jan 1 1970 duration 1
accept-lifetime 10:00:00 Nov 2 2013 10:00:00 Nov 17 2015
!
mpls ldp password option 1 for nbr-acl key-chain MyKey
```

Disable PHP

- How do you disable PHP?
- Well, you really can't, but…
- Have the egress router send an explicit null, rather than the implicit null label!
  - Implicit nulls are automatically popped, explicit nulls are not!

```
R1(config)# mpls ldp explicit-null
```


Cisco live!

Conditional Label Advertisement

- You can control which routes get labels and which do not and to whom you advertise.

```
access-list 1 per 1.1.1.0 0.0.0.0
access-list 2 per 192.168.12.2 0.0.0.0
!
no mpls ldp advertise-labels
mpls ldp advertise-labels for 1 to 2
```


Cisco live!

Route Leaking Between VRFs

- Without a map:

```
ip vrf vpn1
 rd 100:1
 route-target export 100:1
 route-target import 100:1
 route-target import 200:1
!
ip vrf vpn2
 rd 200:1
 route-target export 200:1
 route-target import 200:1
 route-target import 100:1
!
```

Route Leaking Between VRFs

- With an export map:

```
ip vrf vpn1
 rd 100:1
 route-target both 100:1
 route-target import 17:17
!
ip vrf vpn2
 rd 200:1
 export map 2Routes
 route-target both 200:1
!
access-list 1 permit 1.2.3.0
access-list 1 permit 2.3.4.0
!
route-map 2Routes permit 10
 match ip address 1
 set extcommunity rt 17:17 additive
```

Route Leaking Between VRFs

- With an import map:

```
ip vrf vpn1
 rd 100:1
 route-target both 100:1
 import map AllRoutes
!
ip vrf vpn2
 rd 200:1
 route-target both 200:1
!
access-list 1 permit 1.2.3.0
access-list 1 permit 2.3.4.0
!
ip extcommunity-list standard VPN1 permit rt 100:1
ip extcommunity-list standard VPN2 permit rt 200:1
!
route-map AllRoutes permit 10
 match extcommunity VPN1
!
route-map AllRoutes permit 20
 match ip address 1
 match extcommunity VPN2
```

Generate One label for all routes in a VRF

- By default, Cisco routers generate a label per entry in the BGP table(s)

```
R1(config)# mpls label mode all-vrfs protocol bgp-vpnv4 per-vrf
```


Cisco live!

Agenda

- Multiprotocol Label Switching Review
- Virtual Private Networks (without VRFs)
- Virtual Routing and Forwarding tables (VRF Lite)
- Virtual Private Networks (with VRFs)
- MPLS advanced topics (CCIE twists)
- **Troubleshooting**
- Q&A

Cisco live!

Troubleshooting – Show Commands

- **show mpls ip binding**
- **show mpls ldp bindings**
  - Displays the LIB
- **show mpls forwarding-table**
  - Displays the Label Forwarding Information Base (LFIB),
- **show mpls ldp neighbor**
- **show mpls ldp capabilities**
  - To display the Label Distribution Protocol (LDP) capability information.
- **show mpls static binding**
- **show ip cef**
- **show ip bgp labels**

Cisco live!

Troubleshooting – Debug Commands

- **debug mpls ldp bindings**
  - To display information about addresses and label bindings learned from Label Distribution Protocol (LDP) peers by means of LDP downstream unsolicited label distribution,
- **debug mpls ldp advertisements**
  - To display information about the advertisement of labels and interface addresses to label distribution protocol (LDP) peers.
- **debug mpls events**
  - To display information about significant Multiprotocol Label Switching (MPLS) events.
- **debug mpls packets**
  - To display Multiprotocol Label Switching (MPLS) labeled packets switched by the host router.

Labels, but not passing traffic?

- Check to see how the loopbacks of the PE routers are being advertised.
  - Sneaky CCIE trick
  - /24 on Loopback, but OSPF is advertising a /32
    - Why is this a problem?

  - Routers generate a label for entries in the RIB/FIB
    - PE has the /24, therefore a label for that
    - P routers have a /32 in the RIB/FIB, therefore generates a label for a different route.
    - The first P router doesn't get a label for the /32 route from the PE. If there's no outgoing label, the router will pop the label/labels, exposing the IP header; at that point the P router tries to route based on the IP destination.
    - For the BGP relationship, the P route can get the packet to the PE, but when the traffic is going to a network that was learned through BGP, the P router doesn't have a route and discards.

MPLS Ping

- MPLS LSP ping uses MPLS echo request and reply packets to validate an LSP
- The destination IP address is defined as a 127.*x*.*y*.*z*/8 address
  - Prevents the IP packet from being IP switched to its destination if the LSP is broken

```
R1#ping mpls ipv4 4.4.4.0/24
Sending 5, 100-byte MPLS Echos to 4.4.4.0/24,
     timeout is 2 seconds, send interval is 0 msec:

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
  'L' - labeled output interface, 'B' - unlabeled output interface,
  'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
  'M' - malformed request, 'm' - unsupported tlvs, 'N' - no label entry,
  'P' - no rx intf label prot, 'p' - premature termination of LSP,
  'R' - transit router, 'I' - unknown upstream index,
  'l' - Label switched with FEC change, 'd' - see DDMAP for return code,
  'X' - unknown return code, 'x' - return code 0

Type escape sequence to abort.
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 740/816/868 ms
 Total Time Elapsed 4160 ms
```


Cisco live!

MPLS Traceroute

- MPLS LSP traceroute uses MPLS echo request and reply packets to validate an LSP
- MPLS LSP Traceroute feature uses TTL settings to force expiration of the TTL along an LSP

```
R1#traceroute mpls ipv4 4.4.4.0/24
Tracing MPLS Label Switched Path to 4.4.4.0/24, timeout is 2 seconds

Codes: '!' - success, 'Q' - request not sent, '.' - timeout,
  'L' - labeled output interface, 'B' - unlabeled output interface,
  'D' - DS Map mismatch, 'F' - no FEC mapping, 'f' - FEC mismatch,
  'M' - malformed request, 'm' - unsupported tlvs, 'N' - no label entry,
  'P' - no rx intf label prot, 'p' - premature termination of LSP,
  'R' - transit router, 'I' - unknown upstream index,
  'l' - Label switched with FEC change, 'd' - see DDMAP for return code,
  'X' - unknown return code, 'x' - return code 0

Type escape sequence to abort.
  0 10.0.12.1 MRU 1500 [Labels: 16 Exp: 0]
L 1 10.0.12.2 MRU 1500 [Labels: 16 Exp: 0] 348 ms
L 2 10.0.13.3 MRU 1504 [Labels: implicit-null Exp: 0] 360 ms
! 3 10.0.14.4 384 ms
```

Agenda

- Multiprotocol Label Switching Review
- Virtual Private Networks (without VRFs)
- Virtual Routing and Forwarding tables (VRF Lite)
- Virtual Private Networks (with VRFs)
- MPLS advanced topics (CCIE twists)
- Future potential adds
- Troubleshooting
- **Q&A**

Cisco live!

Q&A

Complete Your Online Session Evaluation

- Give us your feedback and you could win fabulous prizes. Winners announced daily.
- Complete your session evaluation through the Cisco Live mobile app or visit one of the interactive kiosks located throughout the convention center.

Don’t forget: Cisco Live sessions will be available for viewing on-demand after the event at CiscoLive.com/Online

Cisco live!


CISCO