MPLS VPN

16-dec-2014

What this lecture is about:

- IP CEF
- MPLS architecture
  - What is MPLS?
  - MPLS labels
  - Packet forwarding in MPLS
- MPLS VPNs

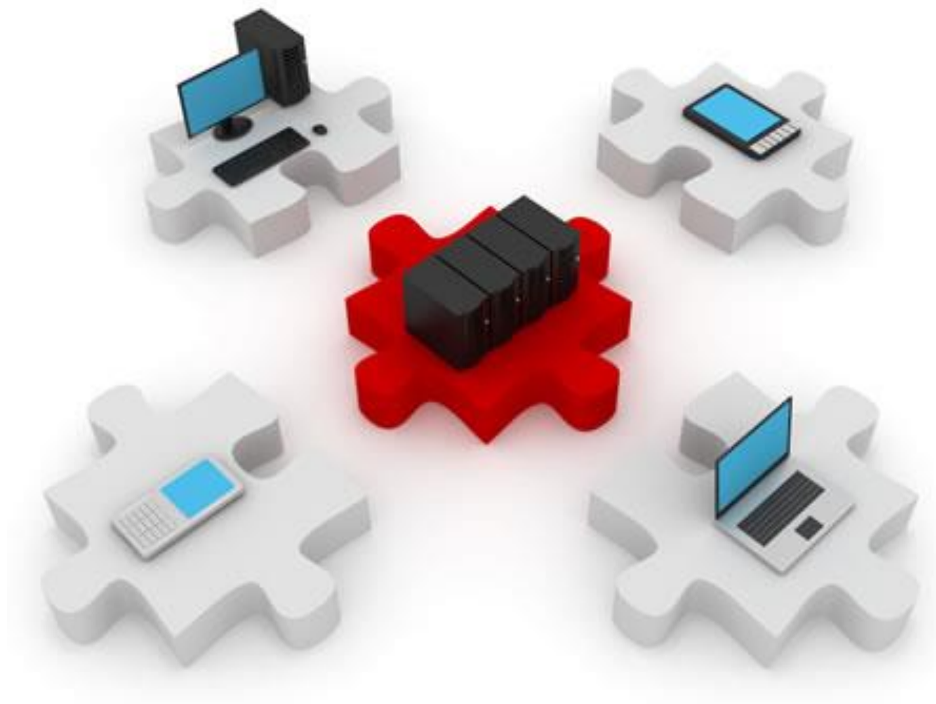

IP CEF & MPLS Overview

How does a router forward packets?

- Process switching, or routing table-driven switching:
  - Full lookup is performed at every packet
  - Slowest method
- Fast switching, or cache-driven switching:
  - Most recent destinations are entered in the cache
  - First packet is always process-switched
- Topology-driven switching:
  - CEF (prebuilt FIB table)

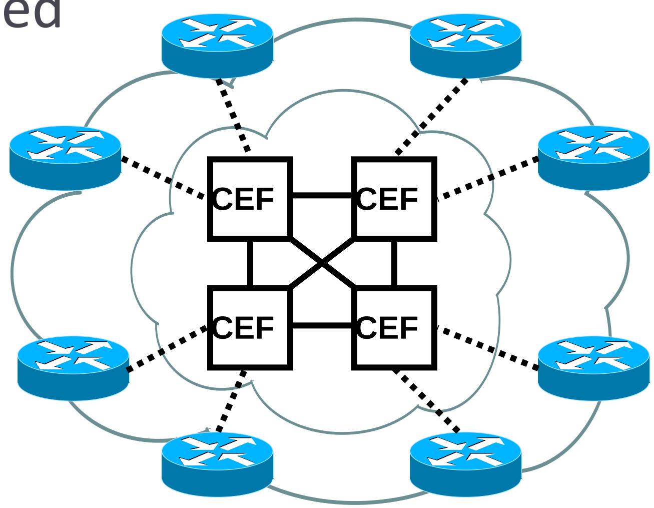

Cisco Express Forwarding

RIB & FIB

- RIB = "Routing Information Base"
  - Is basically the routing table
  - A routing table often requires multiple lookups in order to determine all required parameters for an outgoing packet
- FIB = "Forwarding Information Base"
  - It is a pre-built table similar to the routing table.
  - Provides fast lookup
  - Does not require recursive lookups
- The FIB is constantly synchronized with the routing table
  - Works in conjunction with an adjacency table
    - The AT stores necessary L2 data

Topologies: router types

- With MPLS, we're talking about ISPs and their customers
  - Customers are considered sites (companies) with their own private network.
  - The SP's role is to ensure connectivity between these sites

- C (Customer Router)
  - Belongs to a customer's internal network
- CE (Customer Edge Router)
  - Connects to the SP's network
- PE (Provider Edge Router)
  - Connects to customers' network
- P (Provider Router)
  - Internal SP router

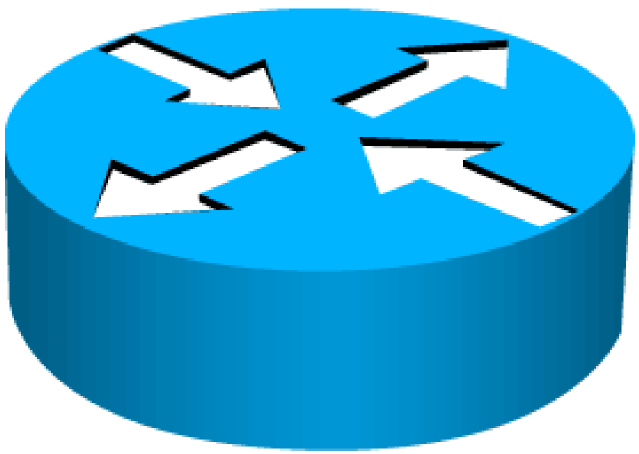

Connecting multiple sites – Full mesh

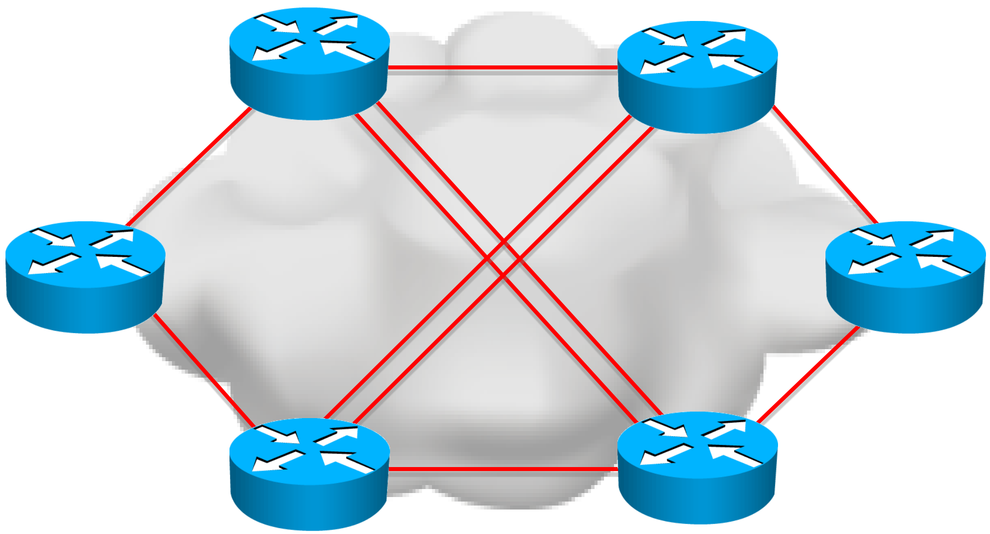

- Provides optimal routing between the sites.
- Dedicated virtual circuit between any two CE routers.
- Very expensive and hard to manage
  - All VCs must be manually configured and maintained

Connecting multiple sites – Partial mesh

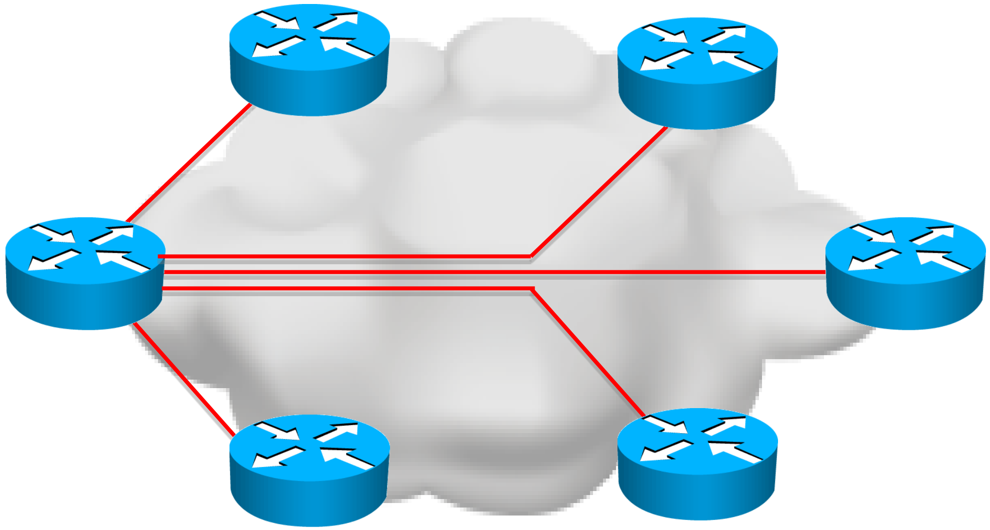

- Also known as hub-and-spoke topology.
- Central point (of failure).
- Suboptimal routing solution.
- Fewer circuits required.
- Less expensive.

Connecting multiple sites – MPLS topology

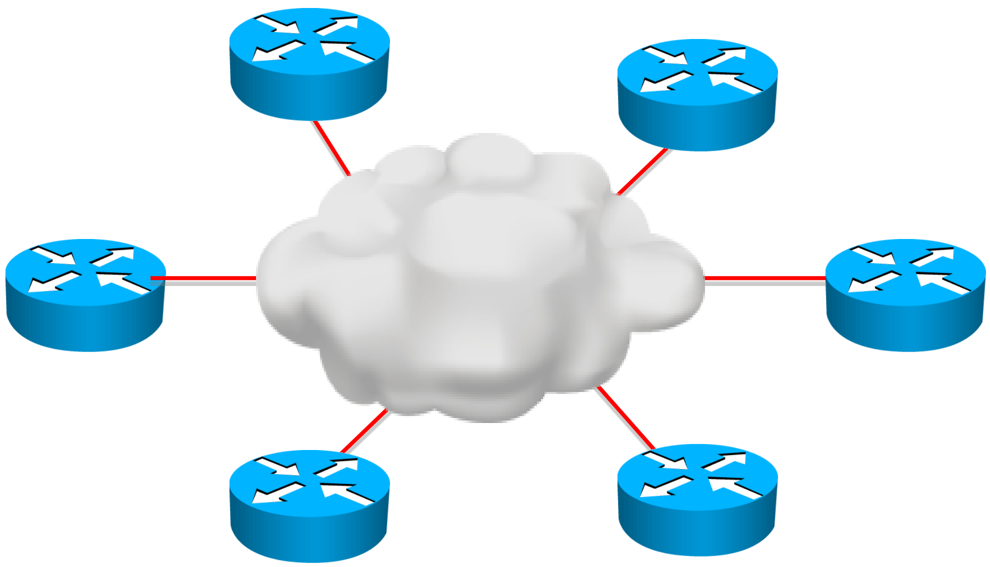

- Optimal routing solution.
- Easy deployment.
  - Each sites only needs to connect to the SP network.
  - Connections between CE and PE routers.
- Packets are switched (not routed) inside the SP's network.

Basic Multiprotocol Label Switching (MPLS) Features

- MPLS reduces routing lookups.
  - MPLS relies on CEF
- MPLS forwards packets based on labels.
  - Label-switching does not involve the routing table
- Labels usually correspond to IP destination networks (equal to traditional IP forwarding).
- Labels can also correspond to other parameters:
  - Layer 3 VPN destination
  - Layer 2 circuit
  - Outgoing interface on the egress router; QoS; Source address; etc
- Currently, MPLS requires IPv4
  - Label switching can work regardless of the L3 protocol
  - However, LDPv6 is only in draft stage at the moment (draft 14 expires April 2015)

MPLS Operation

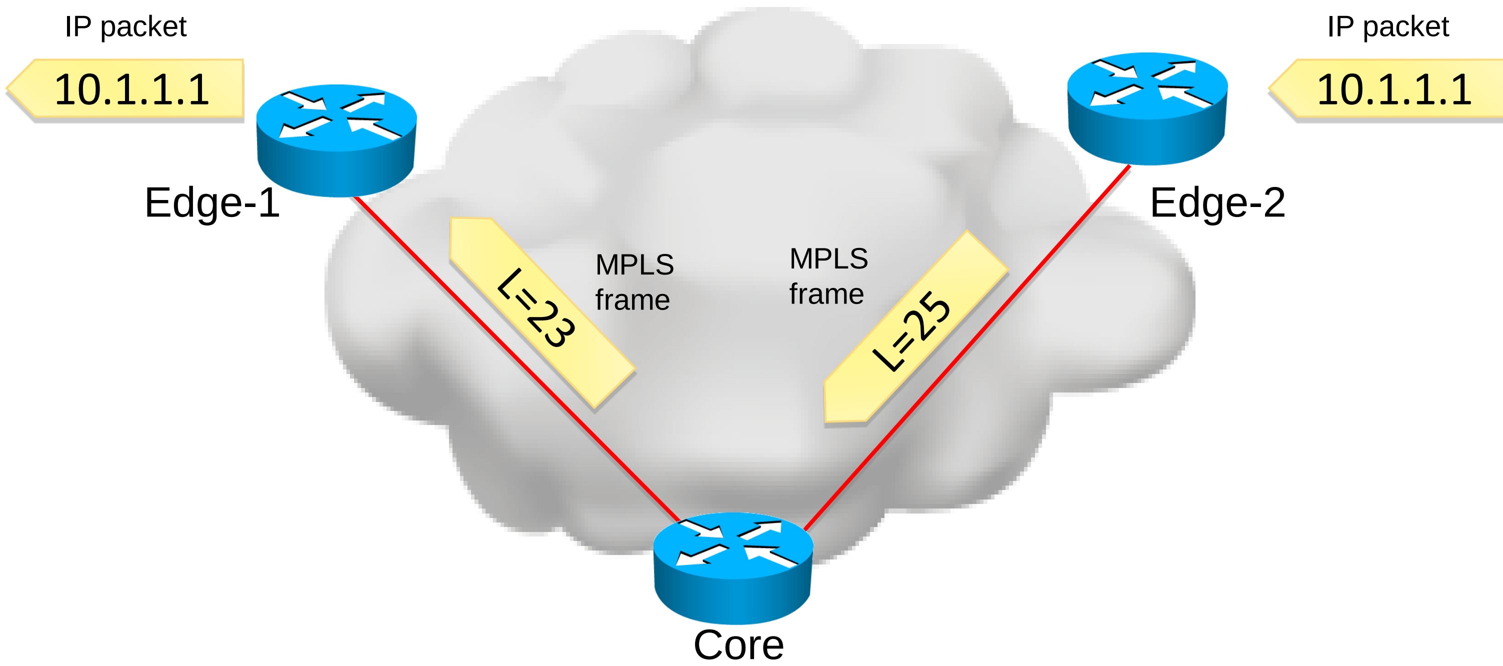

- Only edge routers must perform a routing lookup.
- Core routers switch packets based on simple label lookups and swap labels.
  - No recursive lookups required.
- How do routers know which label to use?
  - Find out later in this lecture.

MPLS Switching Overview

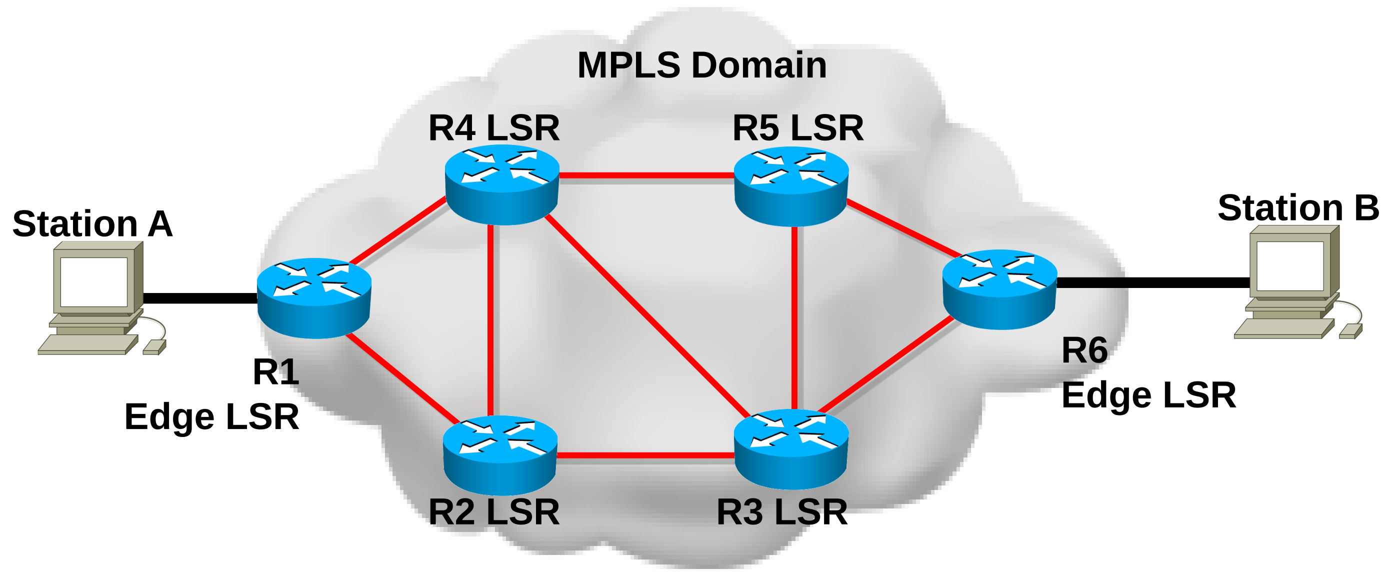

- LSR: Label Switching Router
  - **Core LSRs** forward only based on label information
  - **Edge LSRs** perform routing table lookup and label
- When IP packets enter the MPLS domain, the Edge LSR converts IP packets into MPLS packets by adding a label.
- When packets leave the MPLS domain, the Edge LSR converts the packets back from MPLS packets to IP packets by removing the label

MPLS Switching Overview

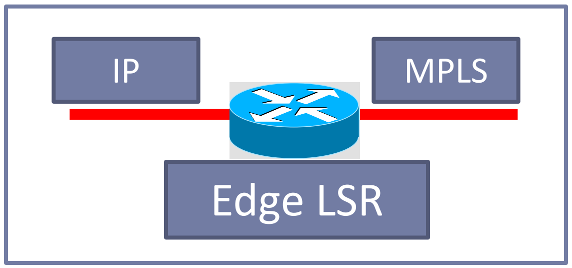

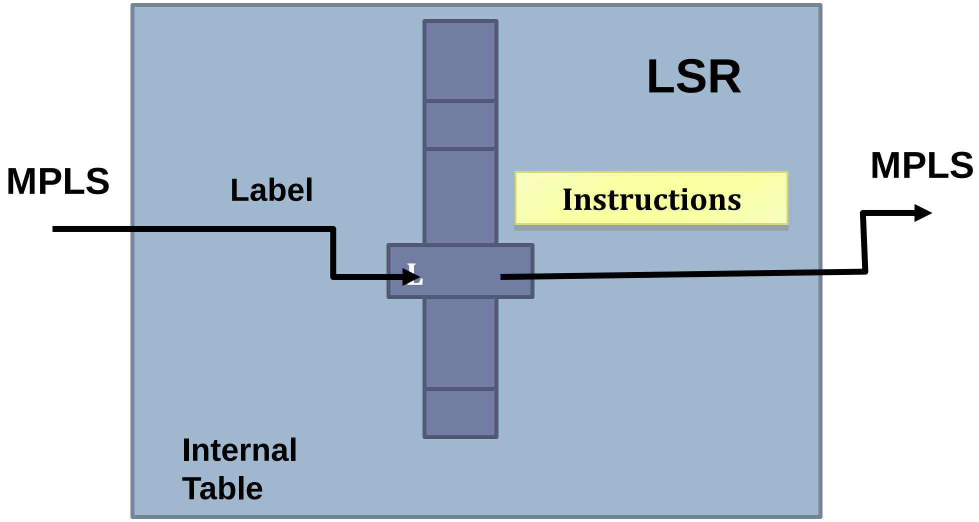

- Edge LSRs can receive an IP packet and insert an MPLS label before forwarding it.

  Or:
- Edge LSRs can receive a labeled packet and remove the label to forward it as an IP packet.

- Core LSRs only analyse the MPLS label to decide:
  - which next hop should be used;
  - which MPLS label should be used with this next hop.

MPLS Characteristics

- MPLS technology is intended to be used anywhere, regardless of Layer 1 media and Layer 2 protocol.
- MPLS uses a 32-bit MPLS field that is inserted between Layer 2 and Layer 3 headers (frame mode MPLS).
  - The Ethertype field announces the MPLS label: 0x8847
  - Which Ethertype field value announces an IP header?
- MPLS has two modes:
  - Frame-mode MPLS (inserts a label between L2 and L3)
  - Cell-mode MPLS (uses the ATM header as label)
    - ATM cannot simply insert a field because it is limited to its 53-byte cell size.
- MPLS uses an LFIB (Label FIB) to map labels to next-hop addresses.

Major Components of MPLS Architecture

- Control plane:
  - Exchanges routing information and labels
  - Contains complex mechanisms, such as OSPF, EIGRP, IS-IS, and BGP, to exchange routing information
  - Exchanges labels using protocols like TDP and LDP
- Data plane:
  - Forwards packets based on labels
  - Has a simple forwarding engine
  - Uses an LFIB (Label Forwarding Information Base)

Control Plane Components Example

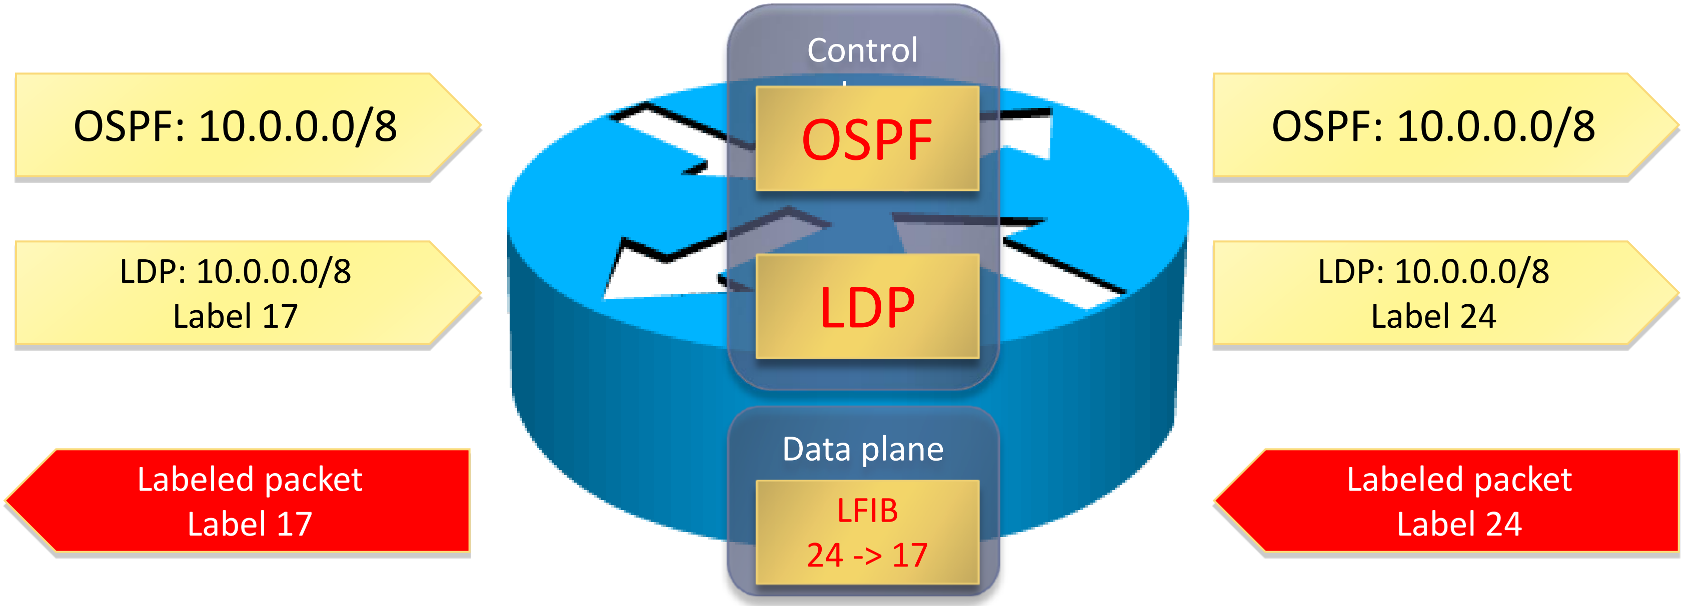

- A Layer3 routing protocol is required to propagate routing information.
- A label exchange mechanism is required to propagate labels for all Layer 3 destinations.
- Information from control plane is sent to the data plane.

Control Plane Components Example

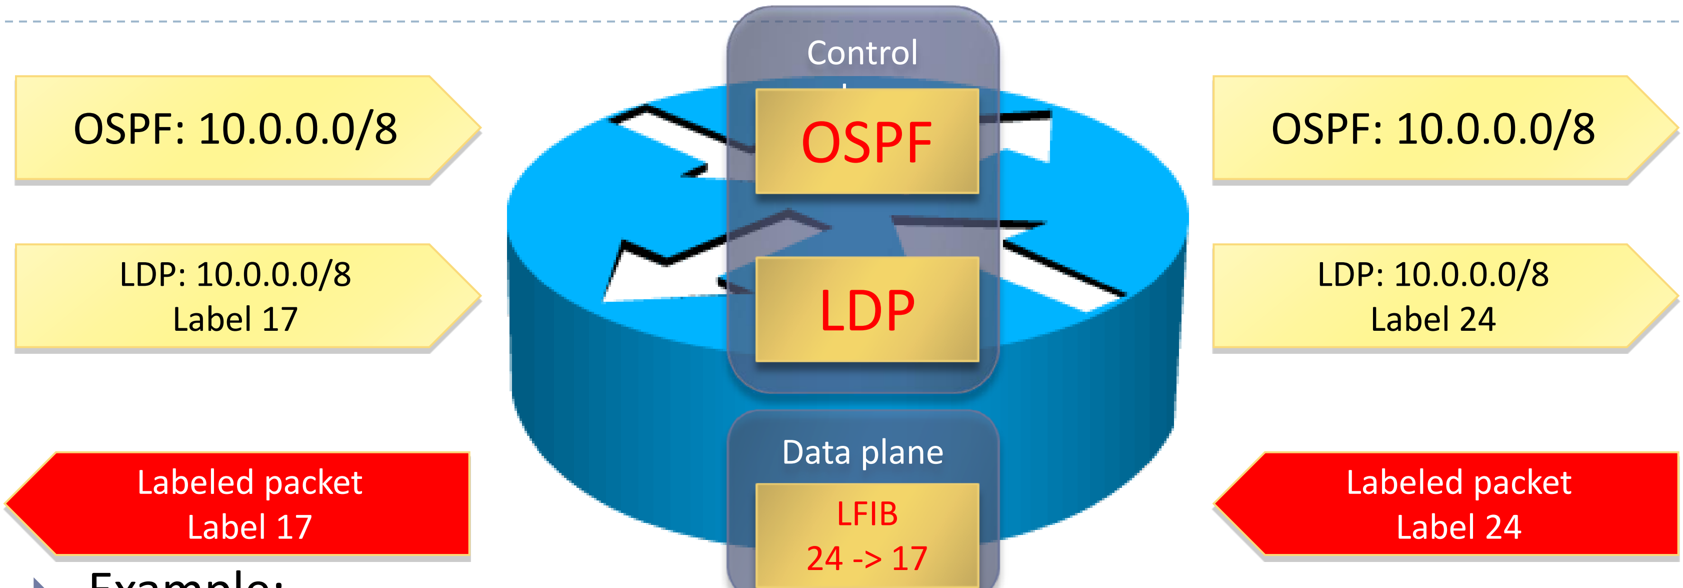

- Example:
  - A routing protocol (OSPF) receives and forwards a routing update for the 10.0.0.0/8 network.
  - A label distribution protocol (LDP) receives label 17 to use for the 10.0.0.0/8 network.
    - It generates a local label of 24
    - It sends the 24 label to all upstream neighbors
    - The upstream neighbors will now use label 24 for all packets destined to 10.0.0.0/8

LFIB example

- The LFIB stores one local and one remote label per prefix.
  - The local label is announced to neighbors that might send packets to the prefix.
  - The remote label is used when the router itself forwards packets towards the prefix.

- Example:

```
Router#show mpls forwarding-table
Local   Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag     tag or VC   or Tunnel Id      switched   interface
16      Pop tag     10.1.1.12/30      636        Se3/0      point2point
17      Pop tag     10.10.10.1/32     0          Se3/0      point2point
21      Pop tag     10.1.1.16/30      0          Se3/0      point2point
22      16          10.10.10.5/32     0          Se3/0      point2point
23      Pop tag     10.10.10.2/32     0          Se4/0      point2point
```

- An outgoing pop tag indicates that any labeled packet with the "Local" tag will be sent untagged.

Functions of LSRs

Component	Function
Control plane	– Exchanges routing information – Exchanges labels
Data plane	– Forwards packets (LSRs and Edge LSRs)

Label Switch Routers (LSRs)

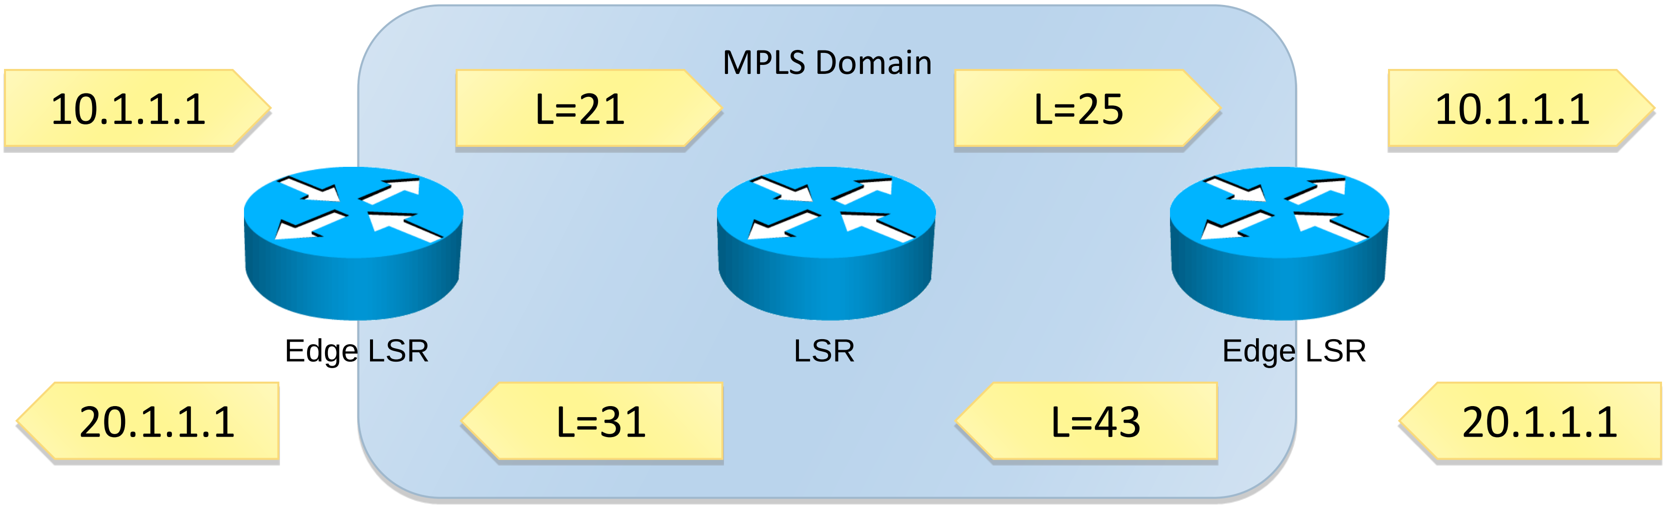

- LSR primarily forwards labeled packets (swap labels).
- Edge LSRs:
  - Label IP packets (impose labels) and forward them into the MPLS domain.
  - Remove labels (pop labels) and forward IP packets out of the MPLS domain.
- Labels are locally significant only.
- Multiple interfaces and multiple routers cand share the same labels.
  - But not multiple destinations on the same router.

Component Architecture of Edge LSR

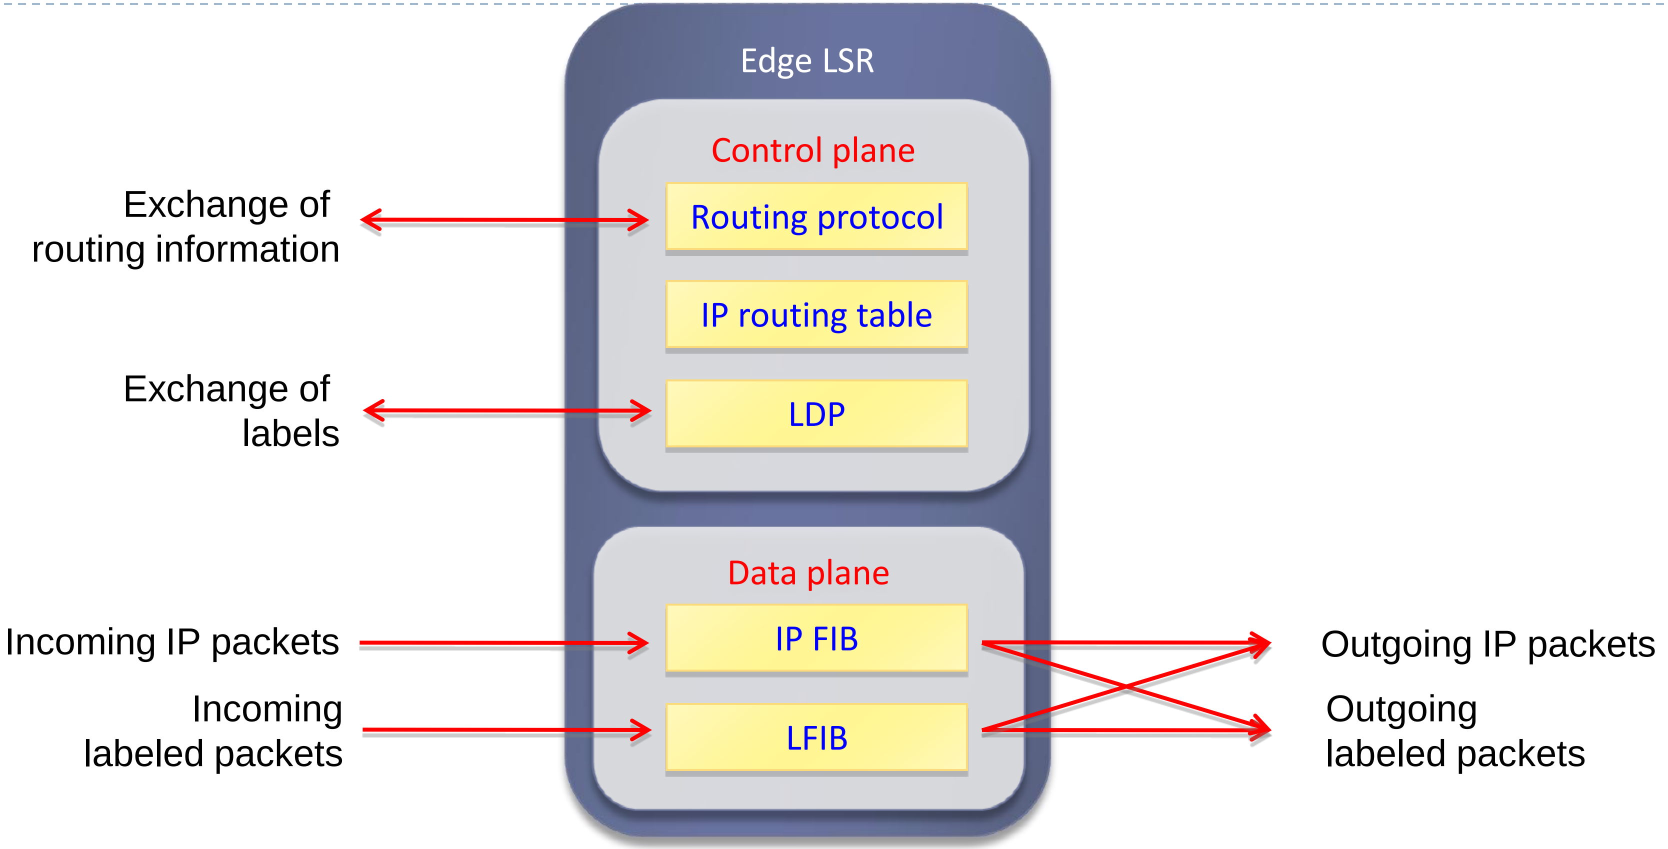

- Scenarios:
  - Receive IP packet > send IP packet (when?)
  - Receive IP packet > send labeled packet

Component Architecture of Edge LSR

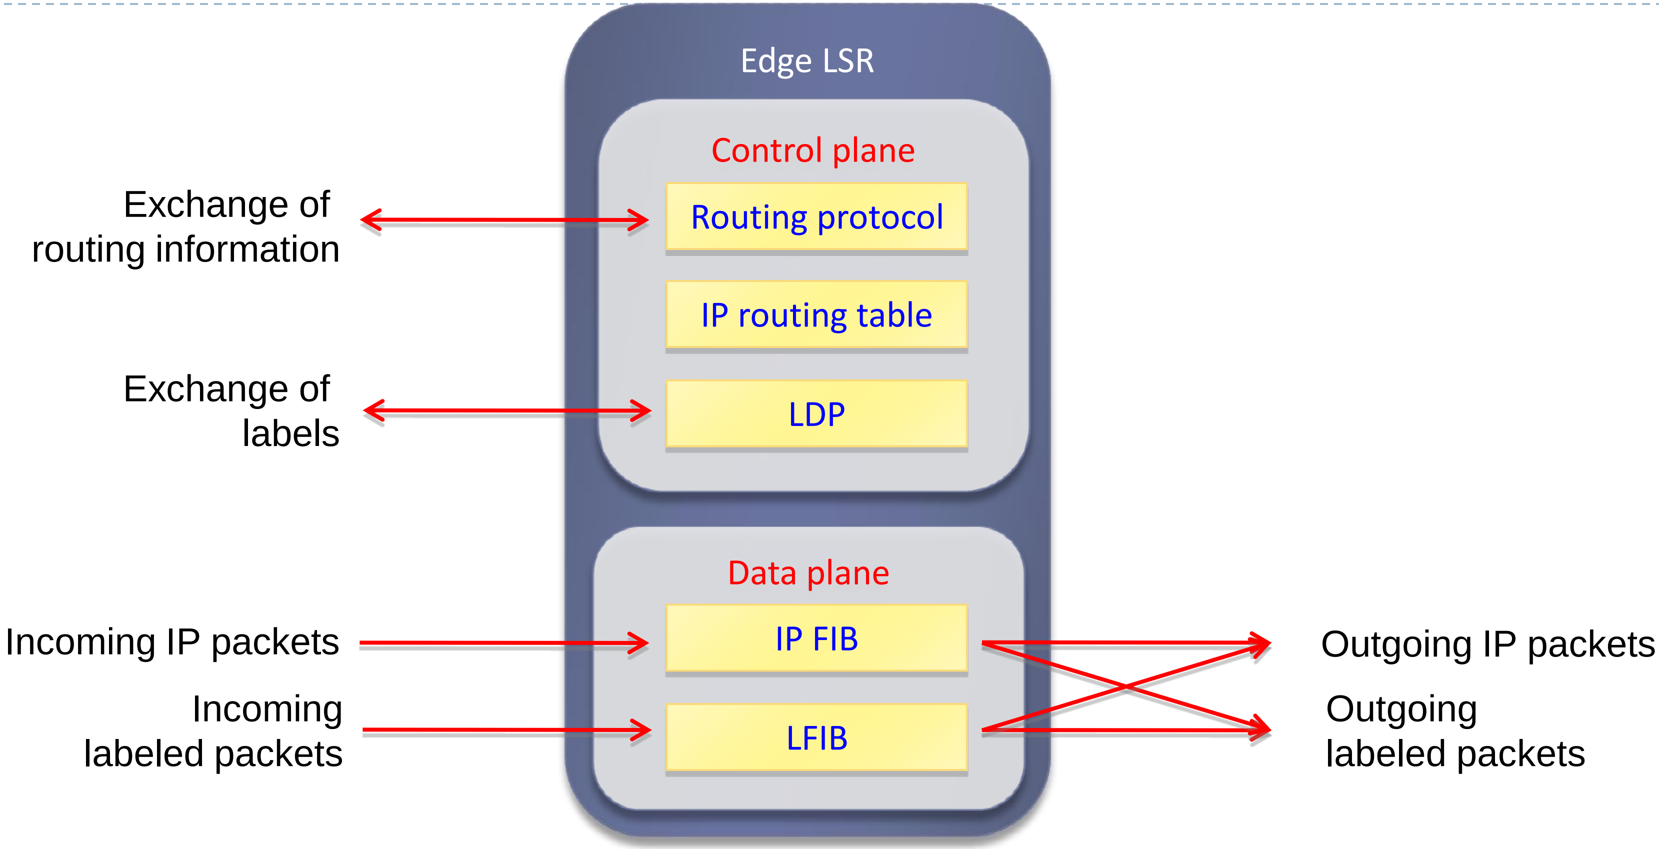

- Scenarios:
  - Receive labeled packet > swap label and send labeled packet
  - Receive labeled packet > remove label and send IP packet

Component Architecture of Edge LSR

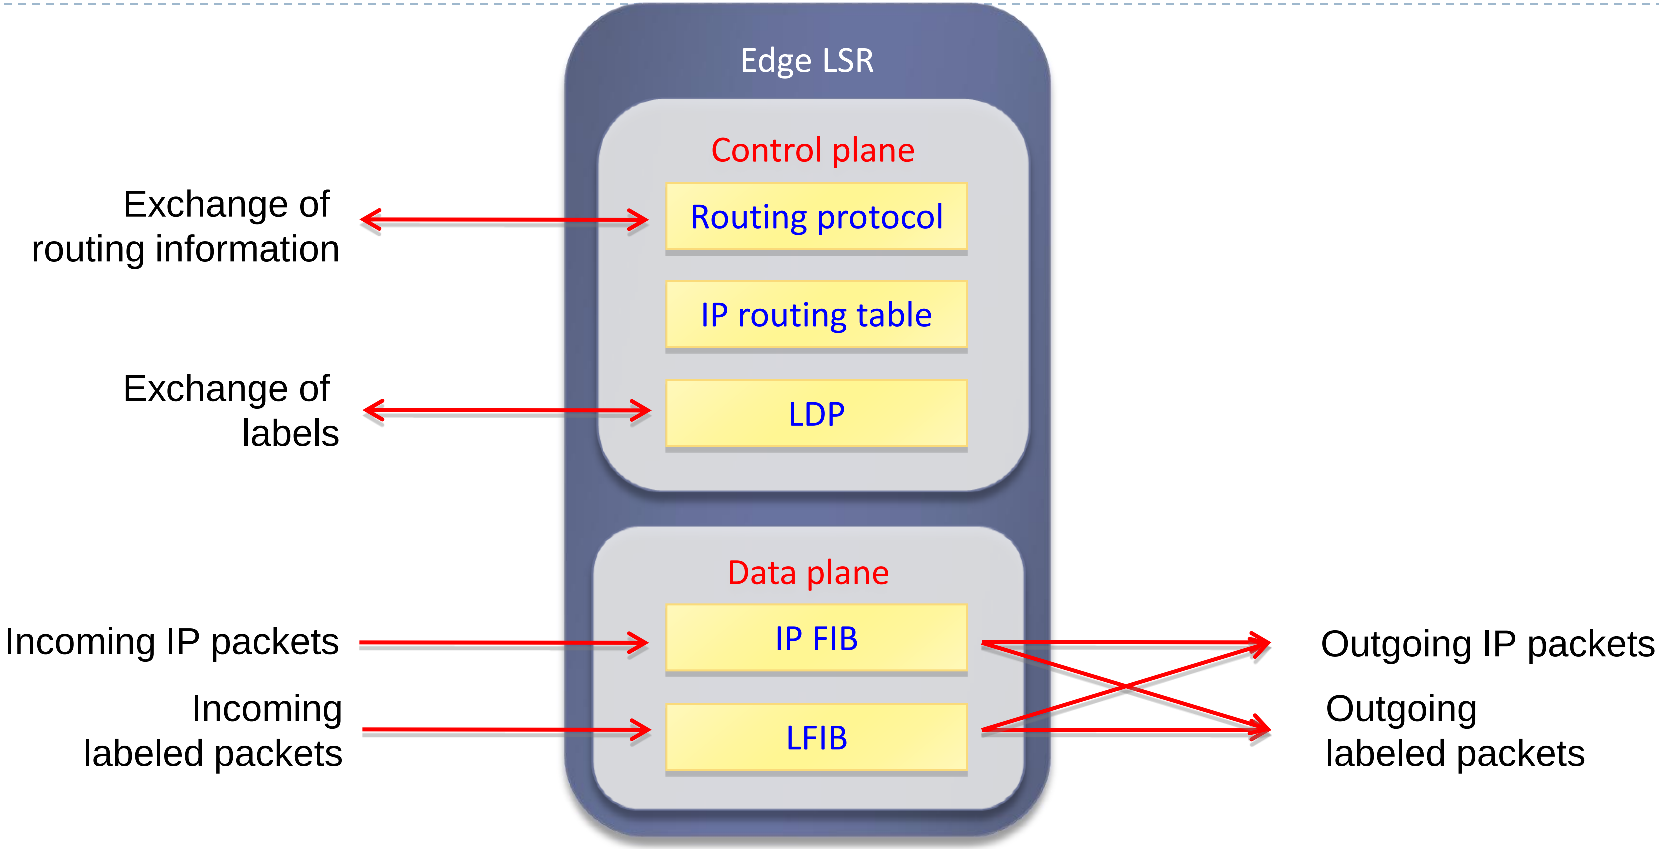

- Possible error scenarios:
  - Receive labeled packet > drop packet because destination is not in LFIB
  - Receive IP packet > drop packet because destination is not in FIB

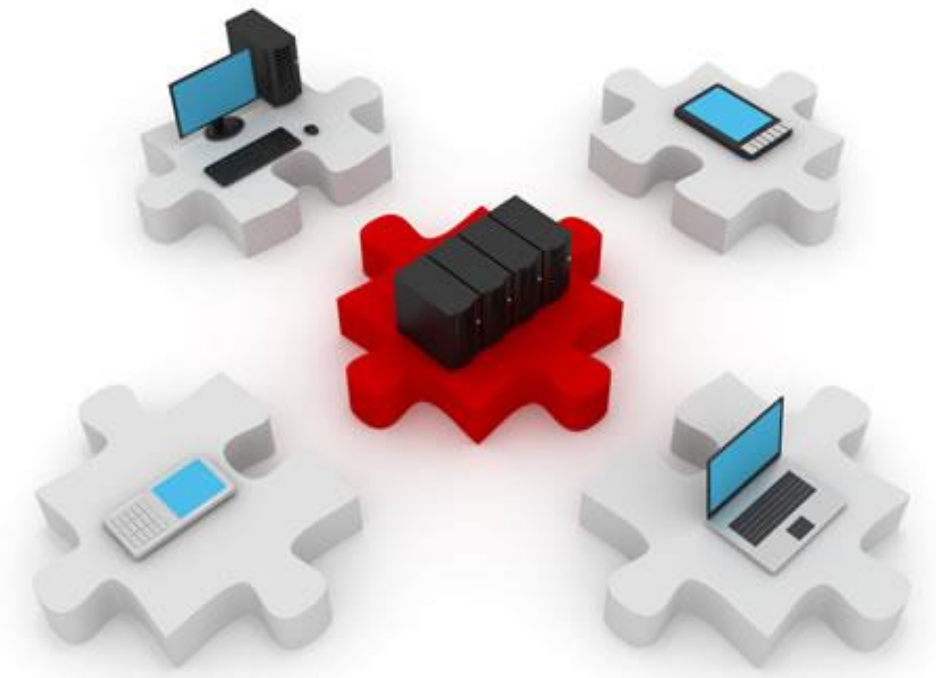

MPLS Labels

Label Format

Field	Description
20-bit label	The actual label. Values 0 to 15 are reserved.
3-bit experimental (EXP) field	Used by Cisco to define a class of service (CoS) in order to assign a value for QoS.
1-bit bottom-of-stack indicator	MPLS allows multiple labels to be inserted. The bottom-of-stack bit determines if this label is the last label in the packet. If this bit is set (1), the setting indicates that this label is the last label.
8-bit Time to Live (TTL) field	Has the same purpose as the TTL field in the IP header.

Frame Mode MPLS Operation

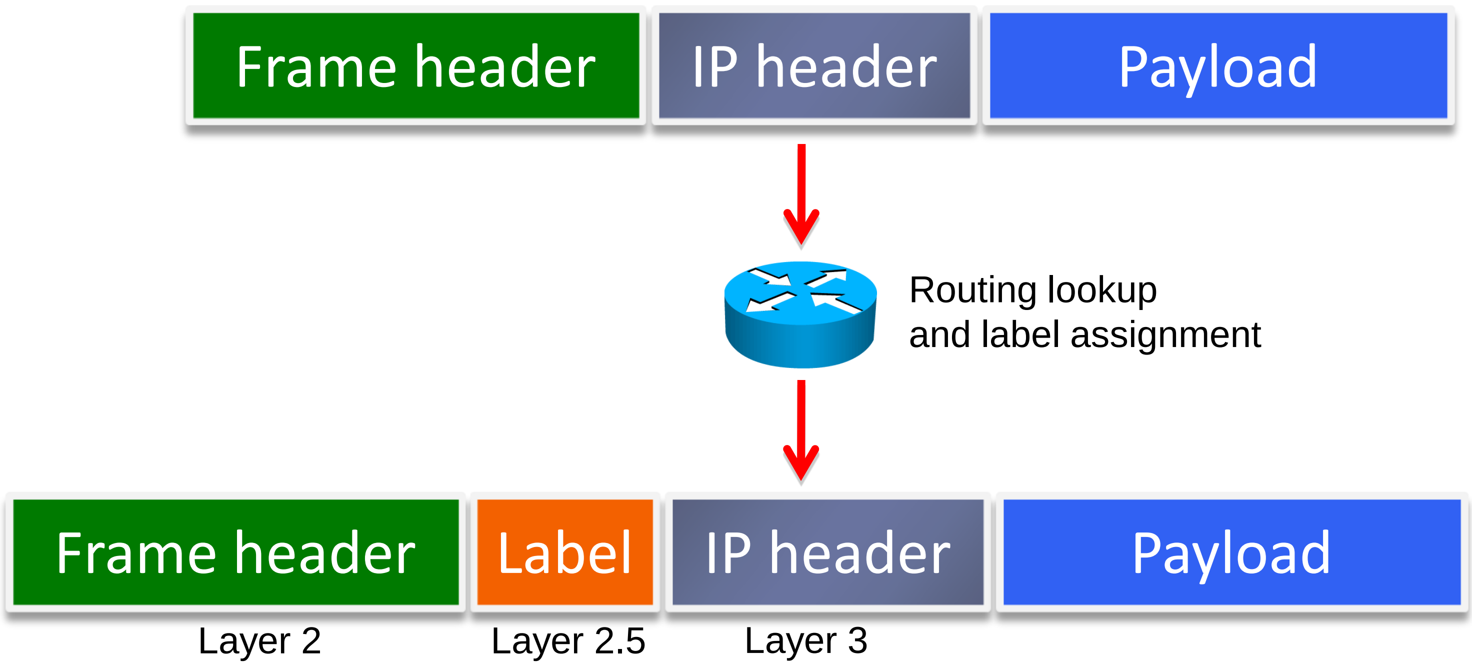

- An MPLS label is announced by the frame's Ethertype field.
- An MPLS label does not store the encapsulated protocol
  - How does an edge LSR that removes the last label know what protocol lies inside?

Label Stack

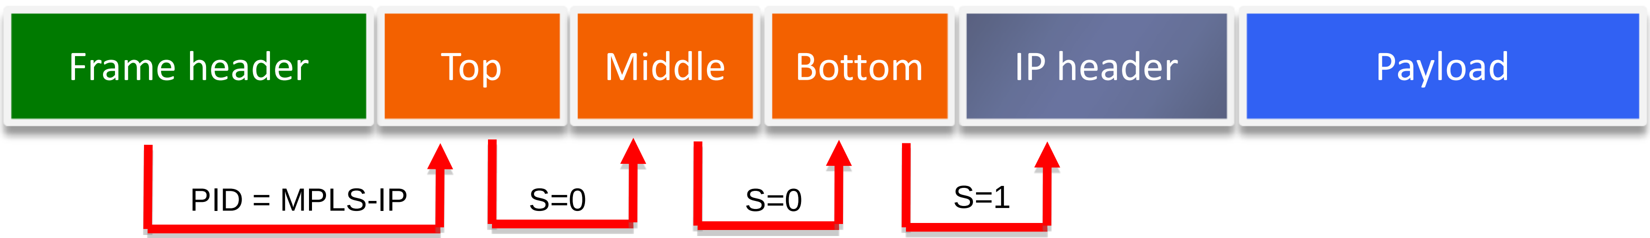

- There may be more than one label in an MPLS packet.
- Only the outermost label is used to route/switch packets in the MPLS domain.
- The bottom-of-stack bit indicates whether the next header is another label or a Layer 3 header.
- Other labels allow services like:
  - MPLS VPNs
  - Traffic engineering (TE)

Label Allocation in a Frame Mode MPLS Environment

- Label allocation and distribution in a frame mode MPLS network follows these steps:
  1. IP routing protocols build the IP routing table.
  2. Each LSR independently assigns a label to every destination in the IP routing table.
  3. LSRs announce their assigned labels to all other LSRs.
  4. Every LSR builds LIB, LFIB, and FIB data structures based on the received labels.

Note: Label allocation, label imposing, label swapping, and label popping usually happen in the service provider network, not the customer (enterprise) network. Customer routers never see a label.

1. Building the IP routing table

Routing table on A

Network	Next Hop
X	B

Routing table on B

Network	Next Hop
X	C

Routing table on C

Network	Next Hop
X	D

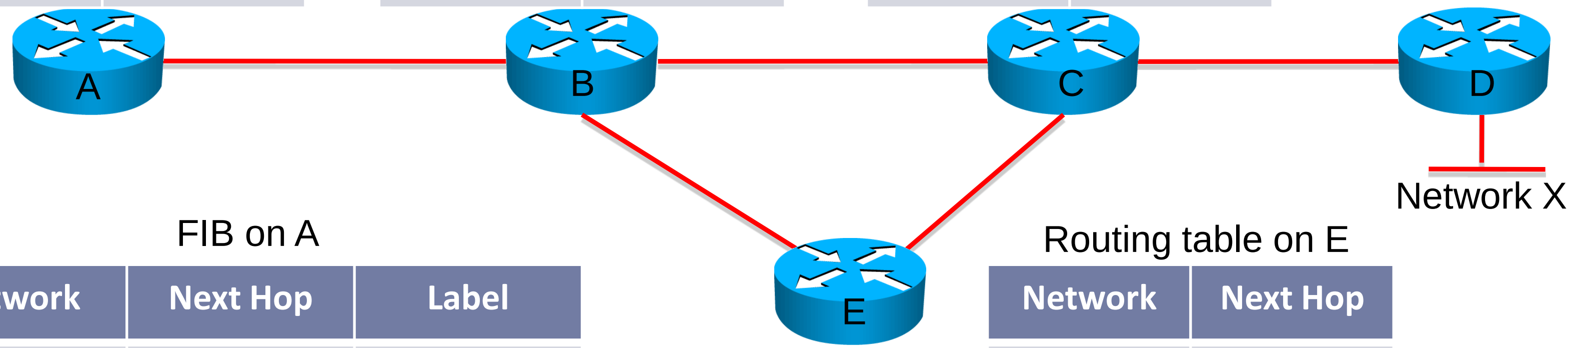

FIB on A

Network	Next Hop	Label
X	B	-

Routing table on E

Network	Next Hop
X	C

- IP routing protocols are used to build IP routing tables on all LSRs.
- FIBs are built based on IP routing tables, with no labeling information.

Allocating labels

Routing table on B

Network	Next Hop
X	C

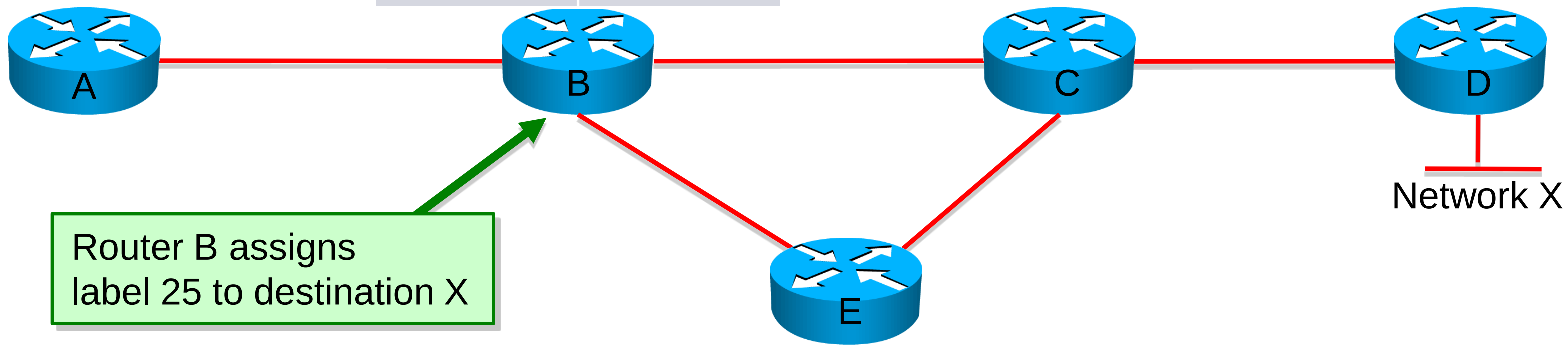

- Every LSR allocates a label for every destination in the IP routing table.
- Labels have local significance.
- Label allocations are asynchronous.
  - Regardless of other routers

LIB and LFIB Setup

Routing table on B

Network	Next Hop
X	C

LIB on B

Network	LSR	Label
X	Local	25

LFIB on B

Label	Action	Next Hop
25	untagged	C

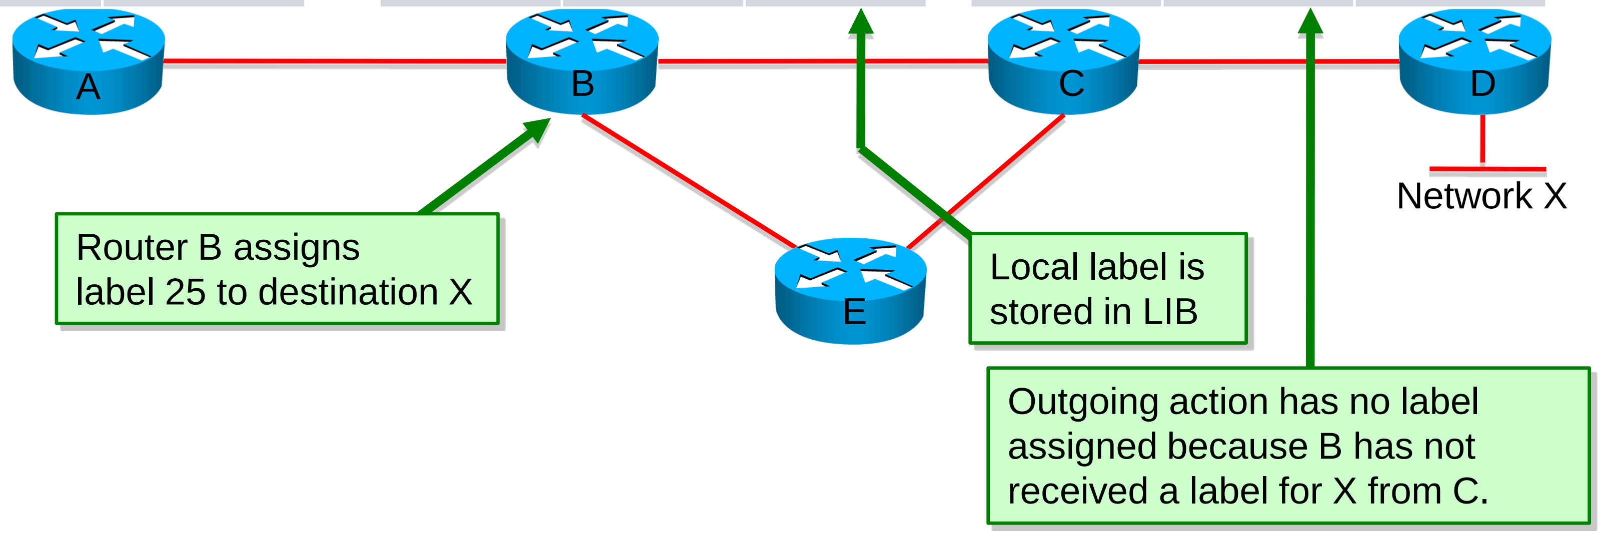

- LIB and LFIB structures have to be initialized on the LSR that is allocating the label.
- Untagged action removes the label from the frame and causes the router to send a pure IP packet.

Label distribution and advertisement

LIB on B

Network	LSR	Label
X	local	25

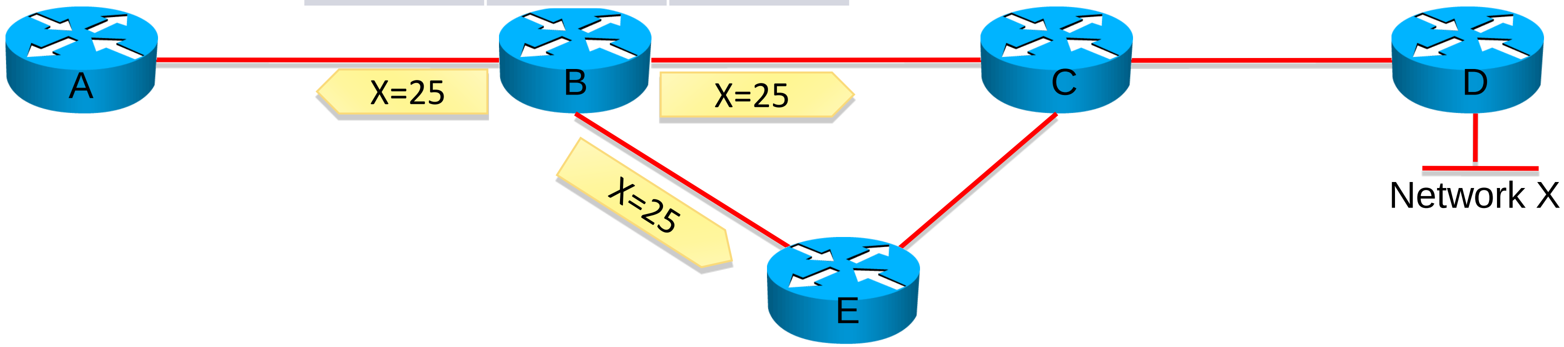

- The allocated label is advertised to all neighbor LSRs, regardless of whether the neighbors are upstream or downstream LSRs for the destination.
- These neighbors will use the received label value when sending packets towards the destination.

Receiving Label Advertisement

LIB on A

Network	LSR	Label
X	local	25

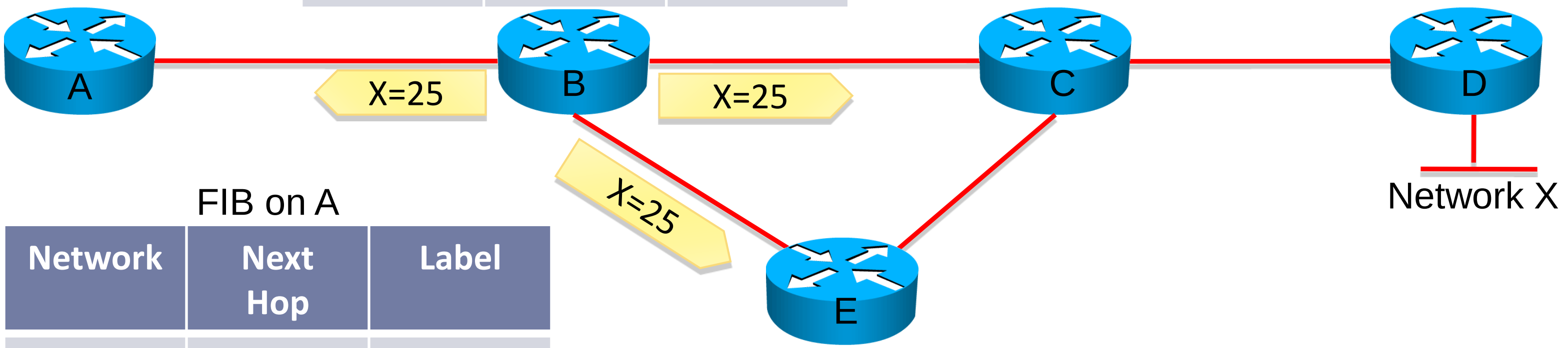

FIB on A

Network	Next Hop	Label
X	B	25

- Every LSR stores the received label in the LSR's LIB (Label Information Base).
- Edge LSRs that receive the label from their next hop also store the label information in the FIB.

Interim packet propagation

LFIB on B

Label	Action	Next Hop
25	untagged	C

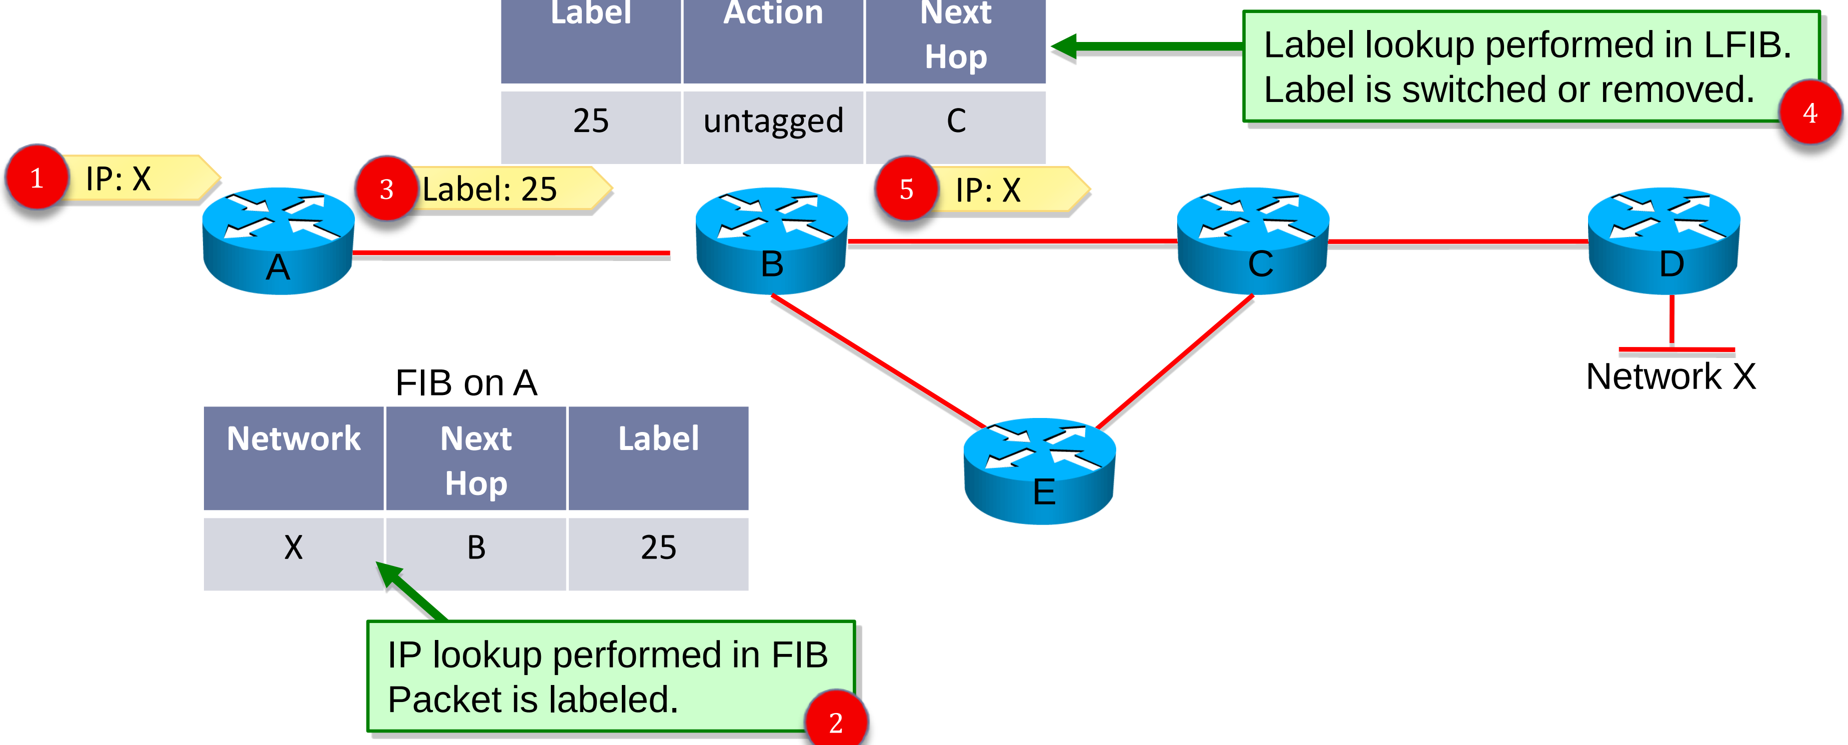

FIB on A

Network	Next Hop	Label
X	B	25

- Forwarded IP packets are labeled only on the path segments where the labels have already been assigned.
- If a label has not been received yet from the downstream neighbor, the packet is sent untagged.

Further label allocation

LIB on C

Network	LSR	Label
X	B	25
	local	47

LFIB on C

Label	Action	Next Hop
47	pop	D

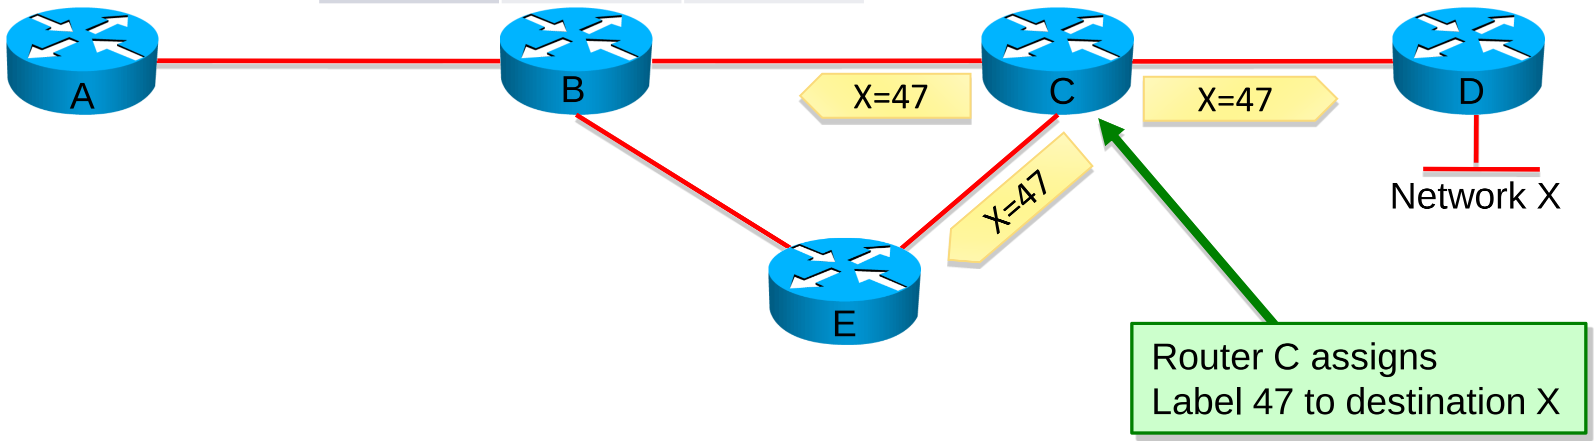

- Every LSR will eventually assign a label for every destination.
- A path of MPLS-enabled routers between two PE routers that has propagated all necessary routes and labels is called an LSP (Label Switched Path).

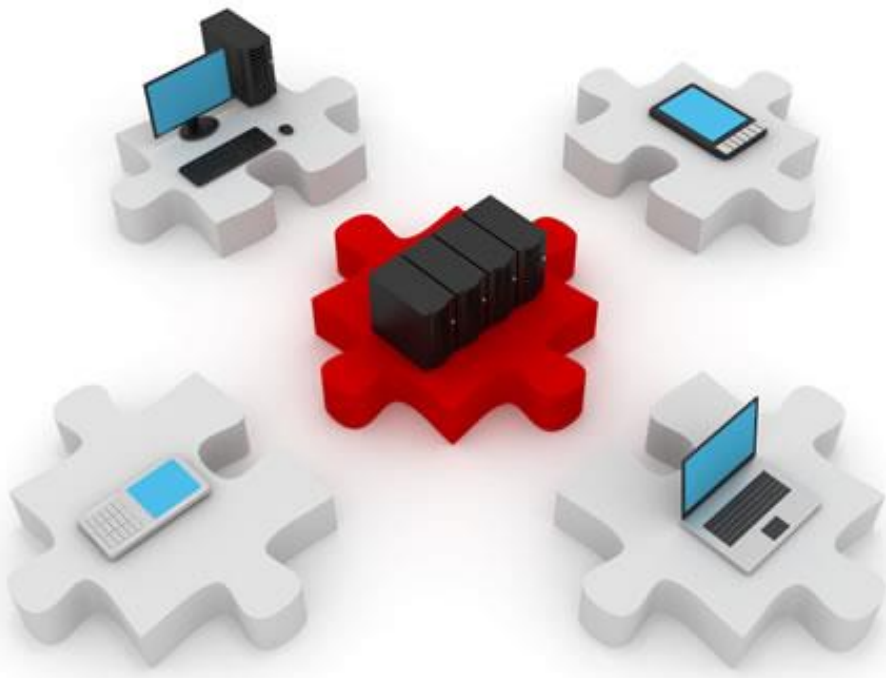

MPLS VPN

VPN Taxonomy

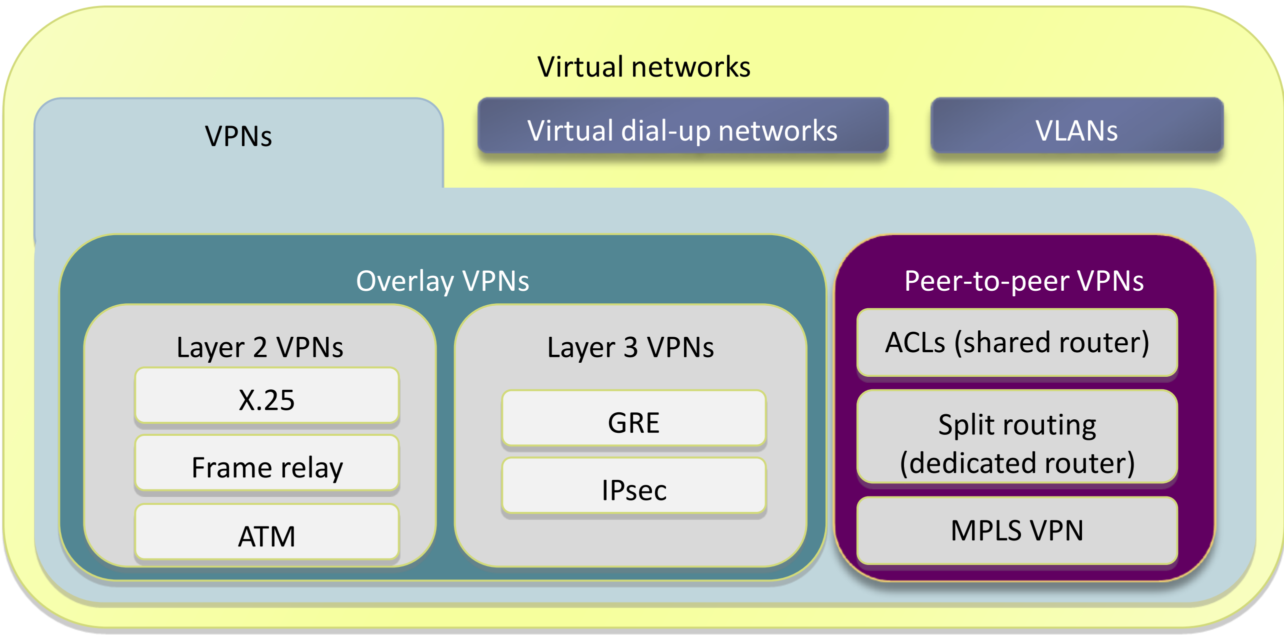

- There are two types of VPN topologies
  - **Overlay VPNs**: the SP provides virtual point-to-point links
    - Customers send their routes through their own tunnels.
  - **Peer-to-peer VPNs**: the SP participates in customer routing
    - The SP is aware and transports the customers' routes.

Overlay VPNs

- Layer 1 Overlay VPN
  - Mentioned for historical reasons only (ISDN, E1/T1)
- Layer 2 Overlay VPN
  - Traditional switched WAN
  - Implmented with X.25, Frame Relay, ATM.
  - SP is responsible for transport of Layer 2 frames
  - Customer is responsible for all higher layers
- Layer 3 Overlay VPN
  - SP network is invisible to customer routers
  - Uses IP tunneling: GRE or IPsec (or both)
  - Routing protocols run directly between customer routers
  - SP provides point-to-point data transport between customer sites
    - The SP is not aware of customer routes.

Layer 3 Overlay VPNs

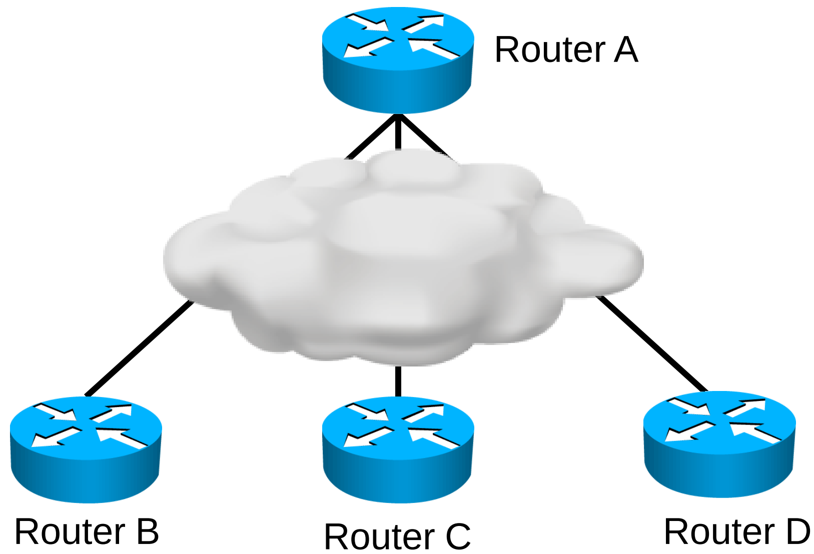

- The service provider infrastructure appears as point-to-point links to customer routes.
- Routing protocols run directly between customer routers.
  - Adjacencies are established over the SP network
- The use of tunnels allows the use of private addresses (RFC 1918).
  - Interconnecting sites with private addressing can be done without NAT

Peer-to-peer VPNs

- The SP and the customers use the same network protocol (IPv4, for example)
- The SP's core carries all customer routes
  - PE routers exchange routing information with CE routers
  - CE routers establish L3 adjacencies only with the PE routers
    - This greately reduces the overhead of full or partial mesh topologies
  - PE routers exchange routing information required for sites to communicate.
    - The SP has to run a routing protocol capable of carrying customer routes.
- The SP's network is a public address space
  - But it carries customer routes that are very likely to use private addressing: first problem.

Peer-to-peer VPNs

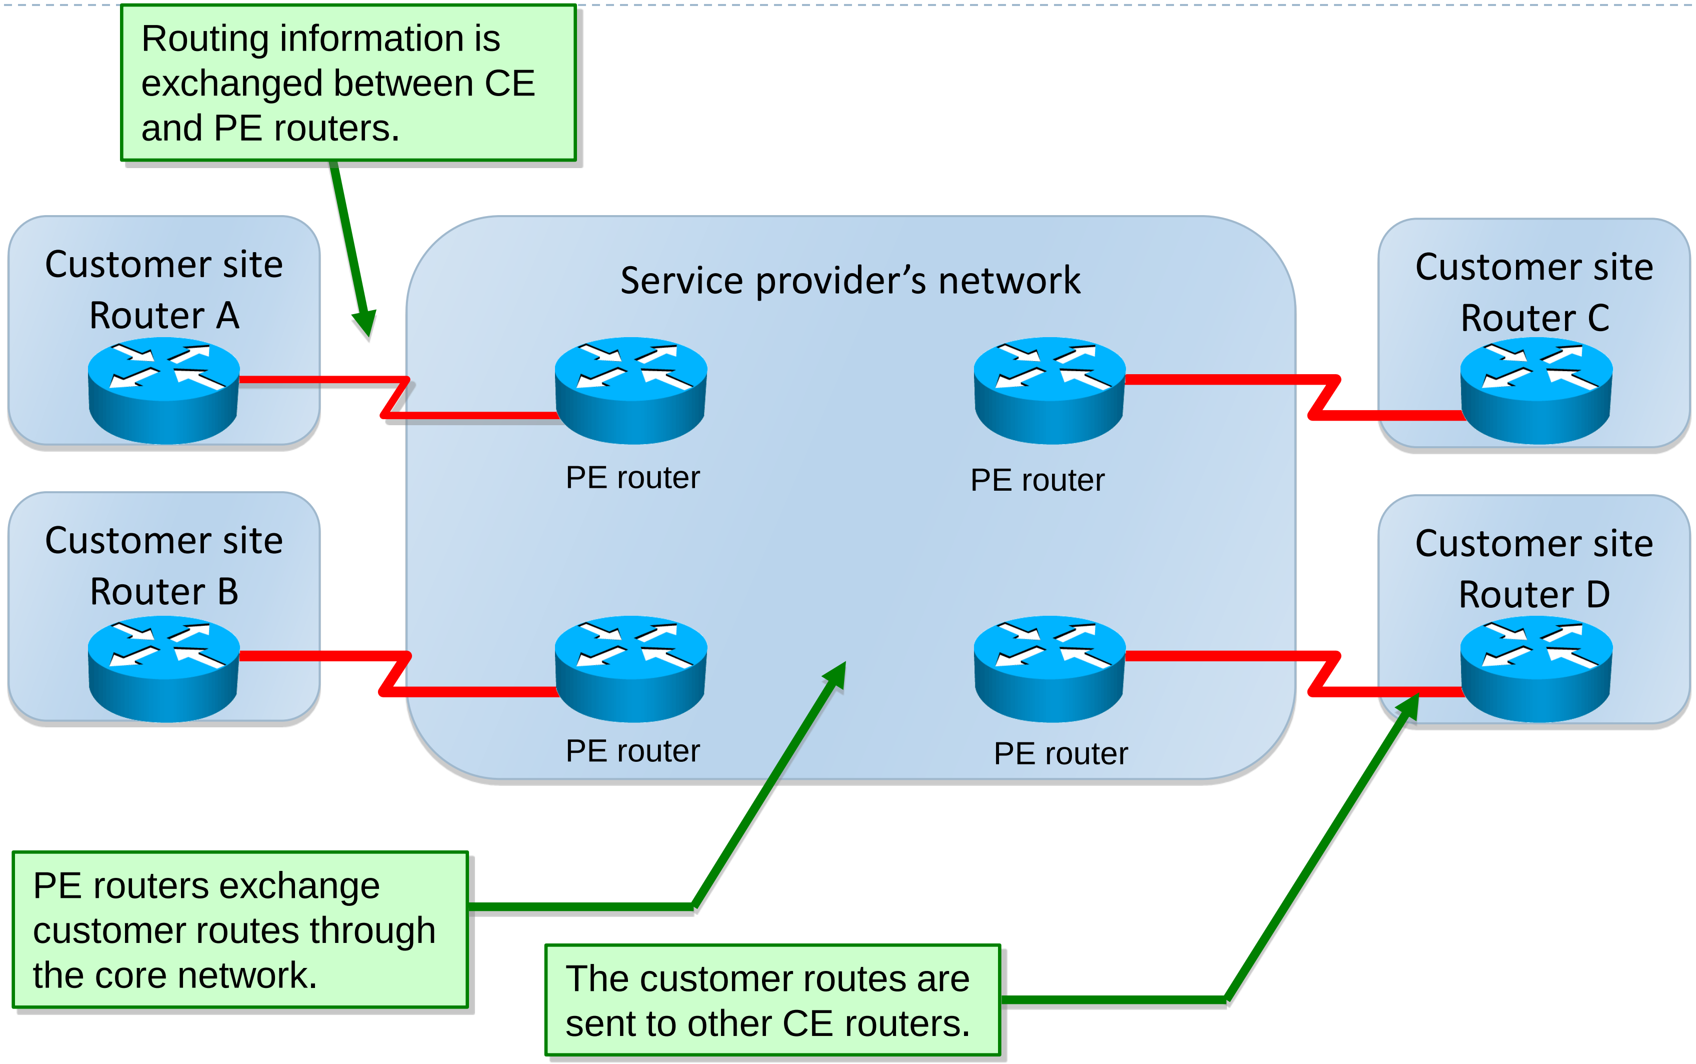

Benefits and disadvantages of the Overlay VPN Implementation Model

- Benefits:
  - Well-known and easy to implement.
  - Service provider does not participate in customer routing.
  - Customer network and service provider network are well-isolated.
- Disadvantages:
  - Implementing optimum routing requires a full mesh of VCs.
  - VCs have to be configured manually.
  - Bandwidth must be reserved on a site-to-site basis.
  - Overlay VPNs always increase encapsulation overhead (IPsec or GRE).
    - 24 to 80+ bytes per packet

Benefits and Disadvantages of the Peer-to-Peer VPN Implementation Model

- Benefits:
  - Guarantees optimum routing between customer sites.
  - Extremely scalable
  - Only sites are configured, not links between them.

- Disadvantages:
  - The service provider participates in customer routing.
  - The service provider becomes responsible for customer convergence.
  - PE routers carry all routes from all customers.
  - The service provider needs detailed IP routing knowledge
    - And skilled employees.

Non-SP Related Drawbacks of Peer-to-Peer VPNs

- If multiple customers share the same PE router:
  - All customers share the same (provider-assigned or public) address space.
  - High maintenance costs are associated with packet filters.
  - Performance is lower—each packet has to pass a packet filter.

- Dedicated PE router:
  - All customers share the same address space.
  - Each customer requires a dedicated PE router.
  - Much more expensive.

MPLS VPN Architecture

- An MPLS VPN combines the best features of overlay VPN and a peer-to-peer VPN models:
  - PE routers participate in customer routing, guaranteeing optimum routing between sites.
  - Very scalable with regards to the number of customers.
  - PE routers carry a separate set of routes for each customer (similar to the dedicated PE router approach).
  - Customers can use overlapping addresses.
    - Unique prefixes are added to each IPv4 route to distinguish between different customers that use the same private addressing scheme

Sample MPLS VPN architecture

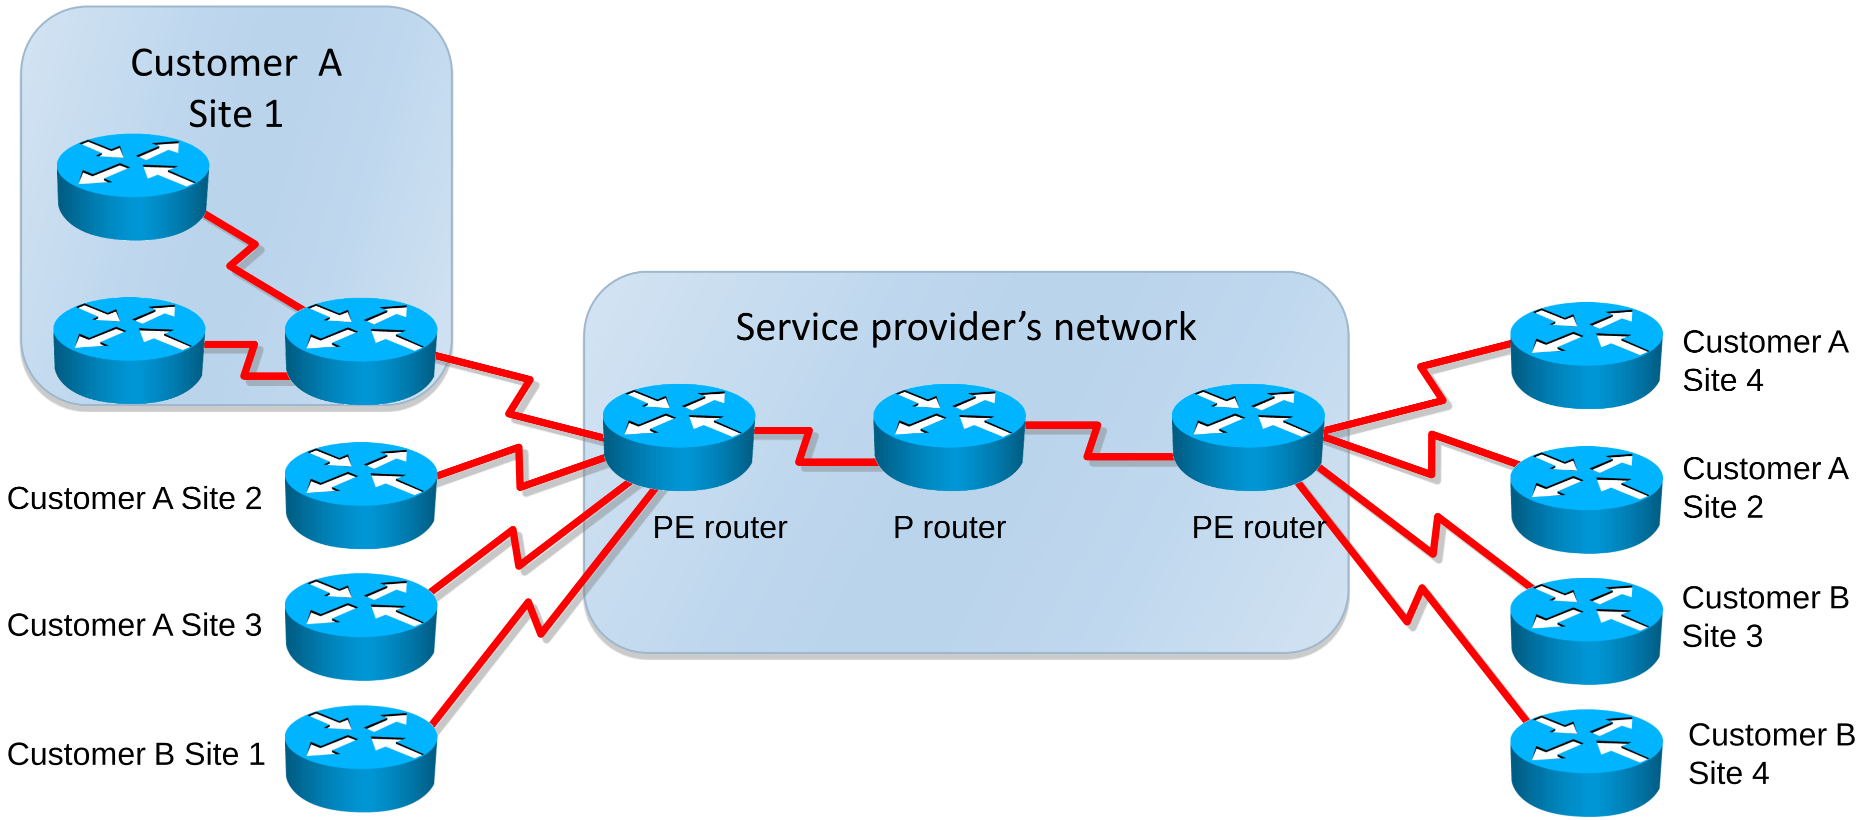

- PE routers transport customer routes
  - P routers simply provide fast transport, without routing knowledge

PE router architecture in MPLS VPN

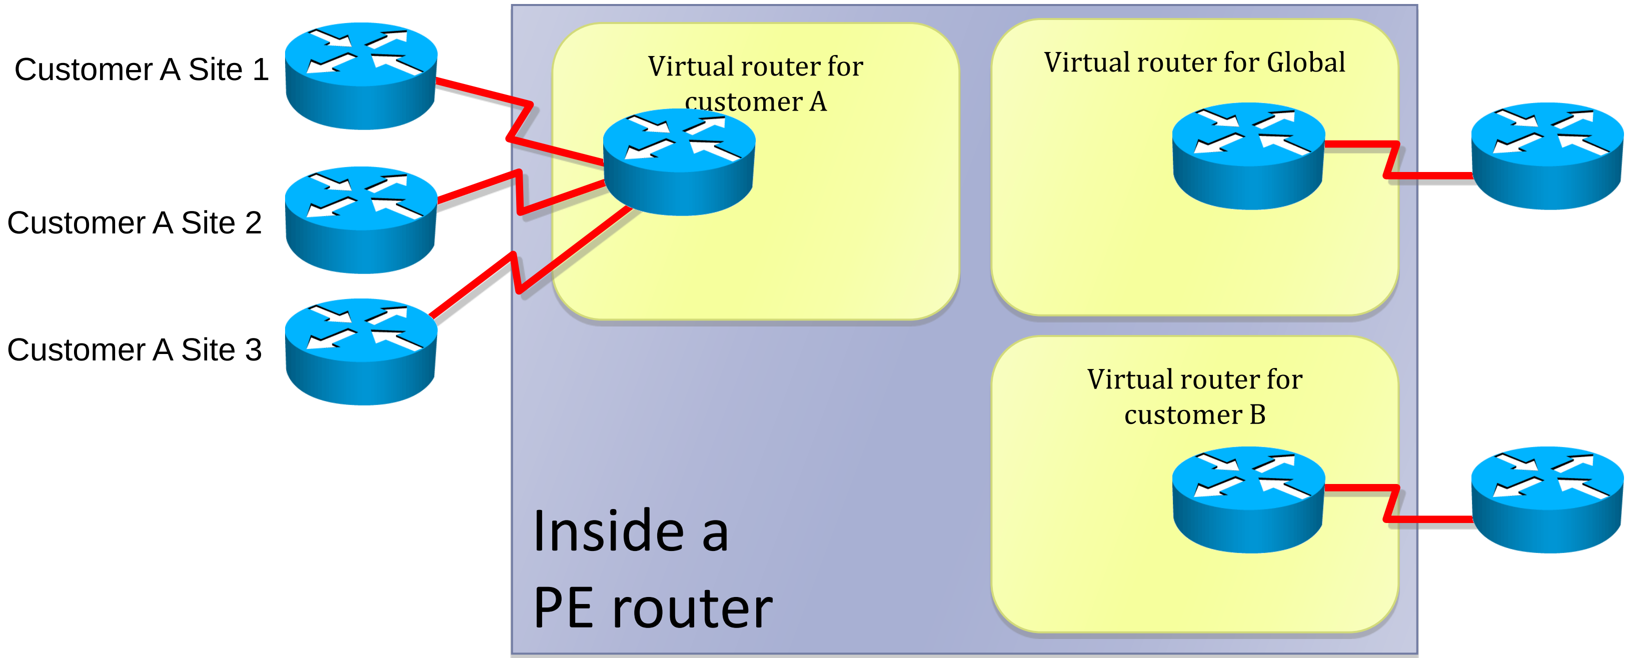

- A PE router is internally divided into multiple virtual routers
  - Each virtual router connects one customer
  - Each customer is assigned an independent Virtual Routing and Forwarding (VRF) table
  - Each VRF corresponds to a dedicated PE router in the traditional peer-to-peer model

Propagation of Routing Information Across the P-Network

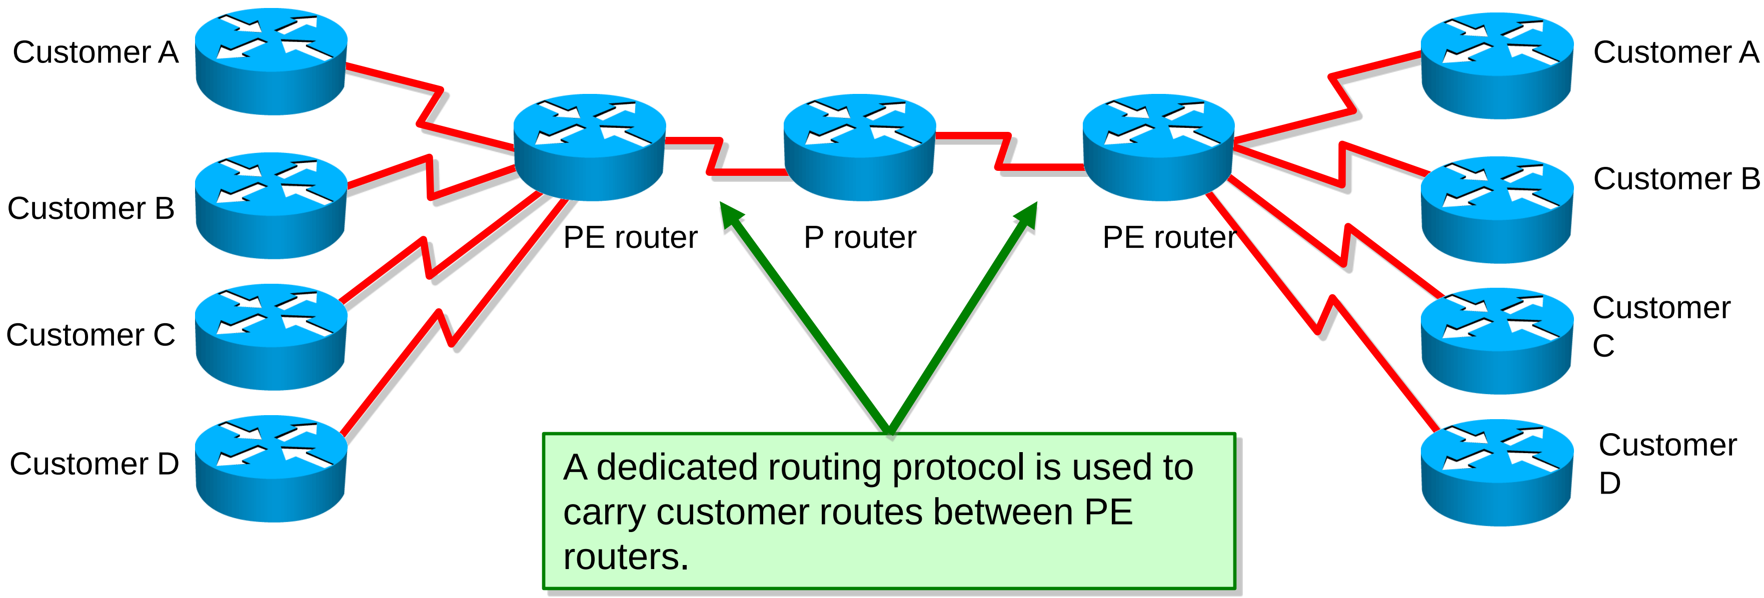

- The number of customer routes can be very large; BGP is the only routing protocol that can scale to such a number.
- BGP is used to exchange customer routes directly between PE routers.
  - Only one routing protocol is required for any number of customers

Route Distinguishers

Question?	How is information about overlapping subnets of two customers propagated via a single routing protocol?
Answer:	Extend the customer addresses to make them unique.

- The 64-bit RD is prepended to an IPv4 address to make the address globally unique.
- The resulting address is a VPNv4 address.
- VPNv4 addresses are exchanged between PE routers via BGP.
- BGP that supports address families other than IPv4 addresses is called multiprotocol BGP (MPBGP).

Using RDs in an MPLS VPN

- The RD has no special meaning.
- The RD is used only to make potentially overlapping IPv4 addresses globally unique.
- This design cannot support all topologies that are required by the customer.

VoIP Service on an MPLS VPN

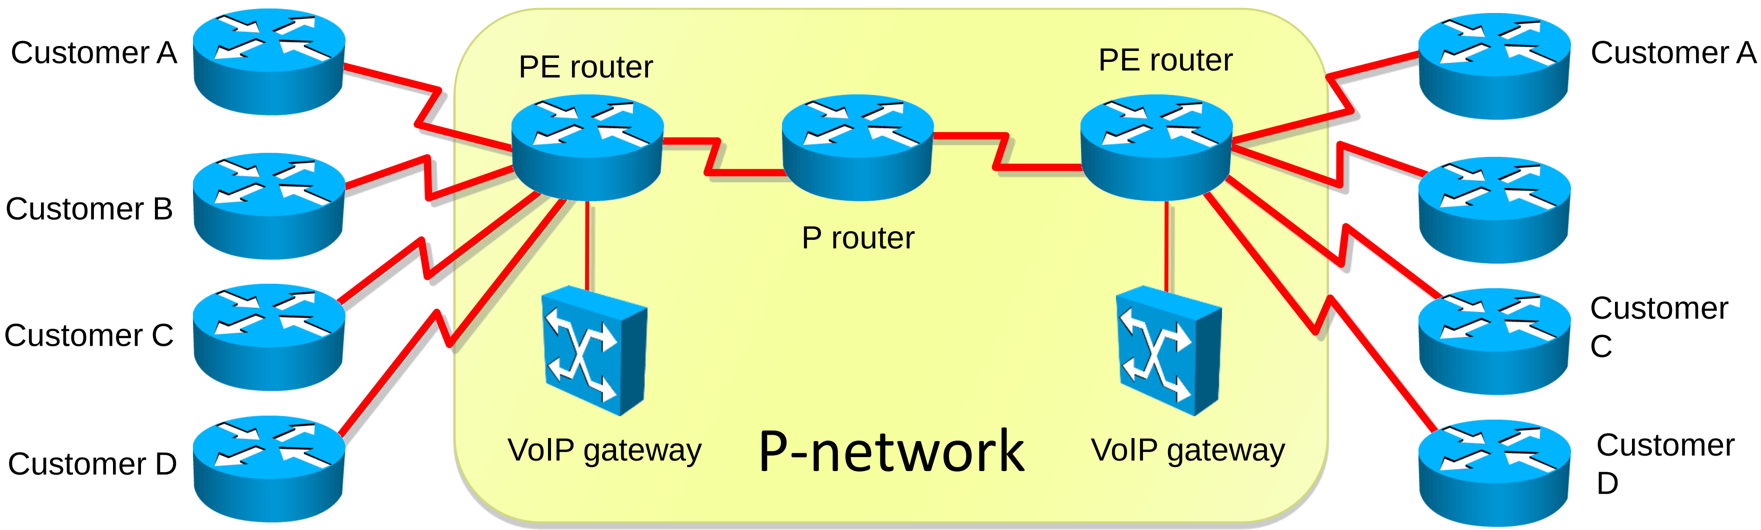

- Requirements:
  - All sites of one customer need to communicate.
  - Central sites of both customers need to communicate with VoIP gateways and other central sites.
  - Other sites from different customers do not communicate with each other.

Connectivity Requirements for VoIP Service

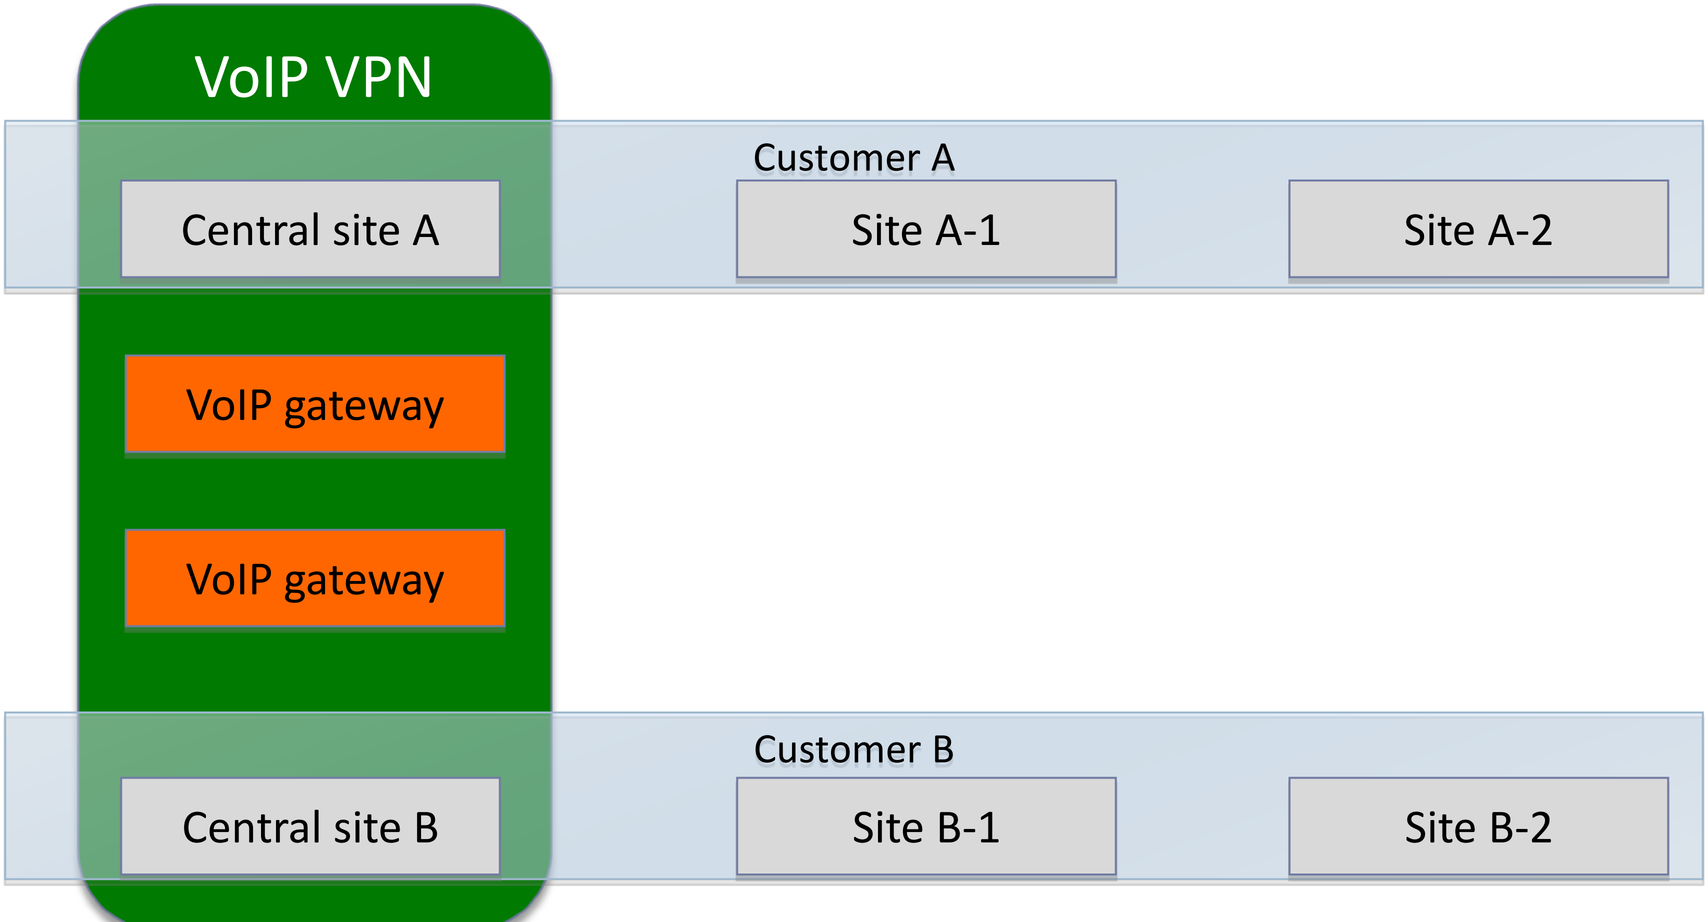

Route Targets

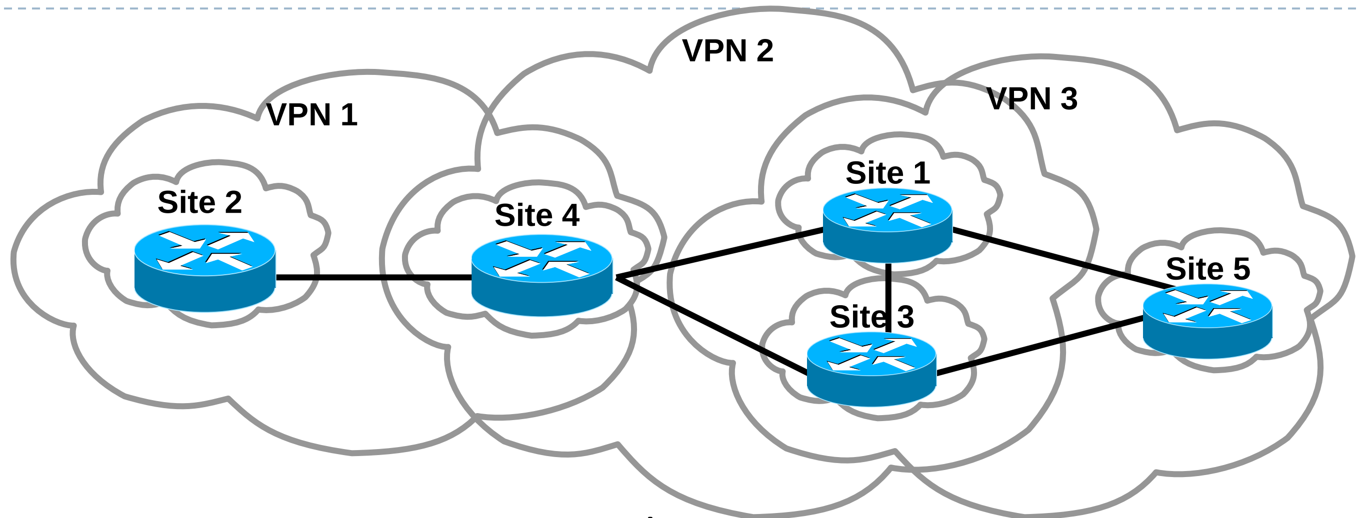

- Some sites participate in more than one VPN.
  - For example the VoIP VPN and the inter-site VPN
- The RD only identifies routes from the same customer
  - But if the customer participates in more than one VPN, all its routes will have the same RD.
- RTs were introduced in the MPLS VPN architecture to support complex VPN topologies.
- RTs are additional attributes that attach to VPNv4 BGP routes to indicate VPN membership.

How Do RTs Work?

- Export RTs:
  - Identify VPN membership
  - Appended to the customer route when the route is converted into a VPNv4 route
  - Configured separately for each VRF instance.

- Import RTs:
  - Associate with each virtual routing table
  - Select which routes are inserted into the virtual routing table

MPLS VPN Routing Criteria

- Designers imposed these criteria on MPLS VPNs:
  - CE routers can only run standard IP routing software.
  - CE routers have no MPLS knowledge
  - CE routers have no VPN knowledge
  - Only PE routers need to support MPLS VPN services and Internet routing.
  - P routers have no VPN routes.

MPLS VPN Routing: CE Router Perspective

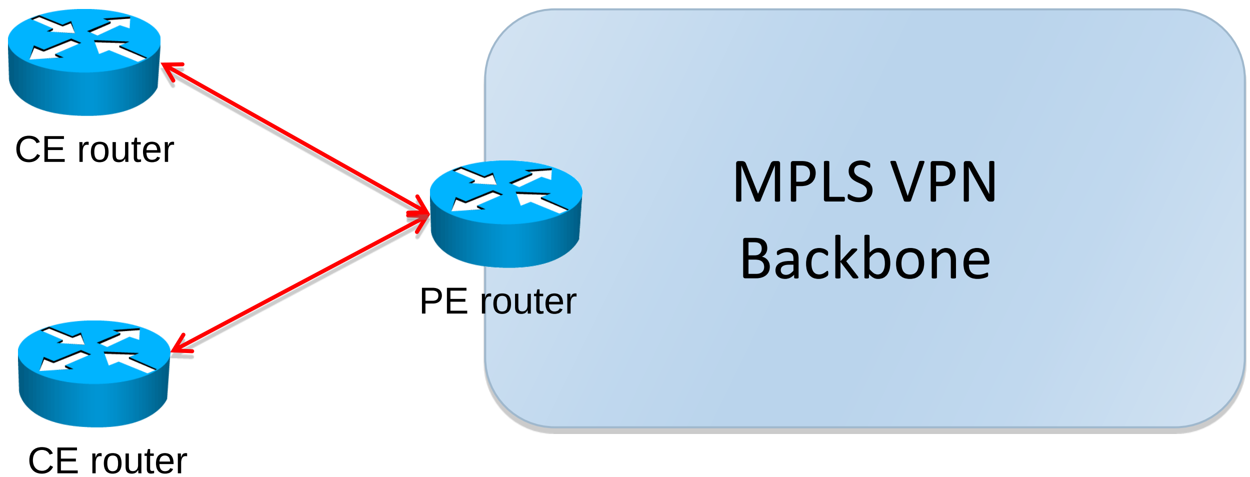

- The CE routers run standard IP routing software and exchange routing updates with the PE router.
- The PE router appears as another router in the C-network.

PE-CE Routing Protocols

- PE-CE routing protocols are configured for individual VRFs.
- Supported protocols include BGP, OSPF, static, RIP, and EIGRP.
- Routing configuration on the CE router has no VRF information.
- Customer configuration is the same as if the customer is routing between devices in the C-network.

MPLS VPN Routing: P Router Perspective

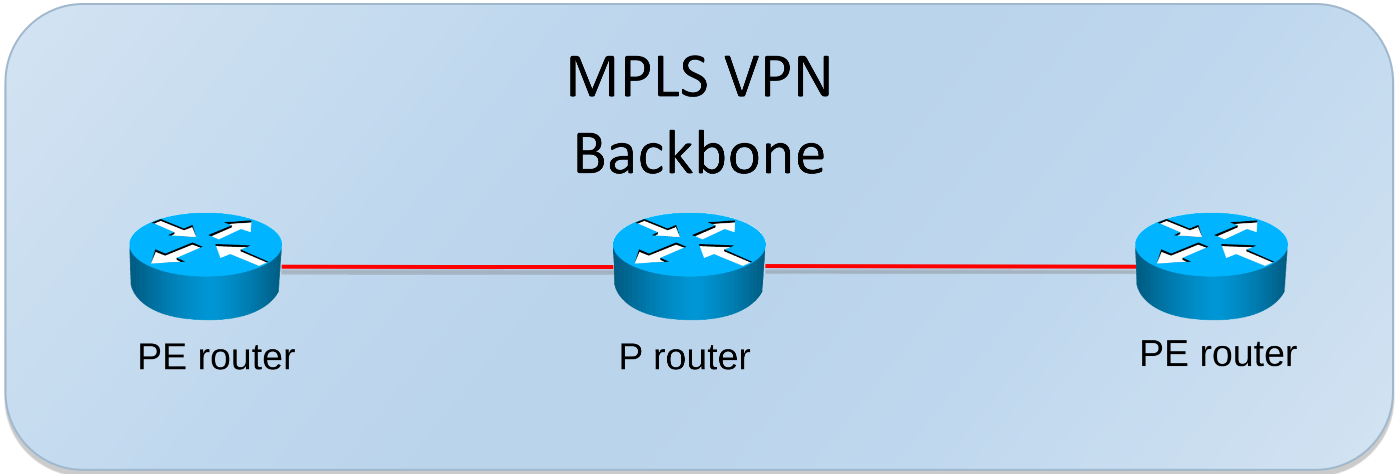

- P routers perform as follows:
  - Do not participate in MPLS VPN routing and do not carry VPN routes
  - Run backbone IGP with the PE routers and exchange information about global subnetworks
  - Use MPLS labels to quickly forwards labeled packets

MPLS VPN Routing: PE Router Perspective

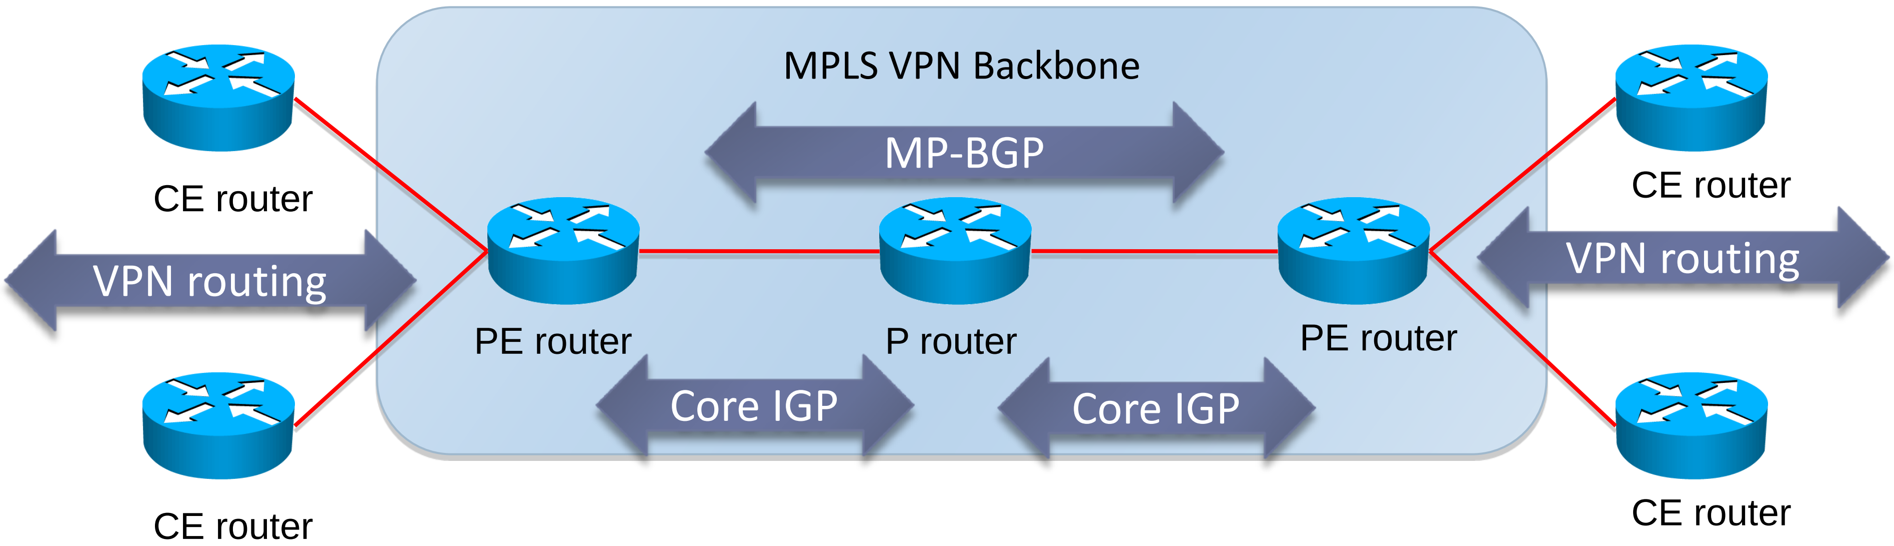

- PE routers exchange the following:
  - VPN routes with CE routers via per-VPN routing protocols
  - Core routes with P routers and PE routers via core IGP
  - VPNv4 routes with other PE routers via MPBGP sessions

End-to-End Routing Information Flow

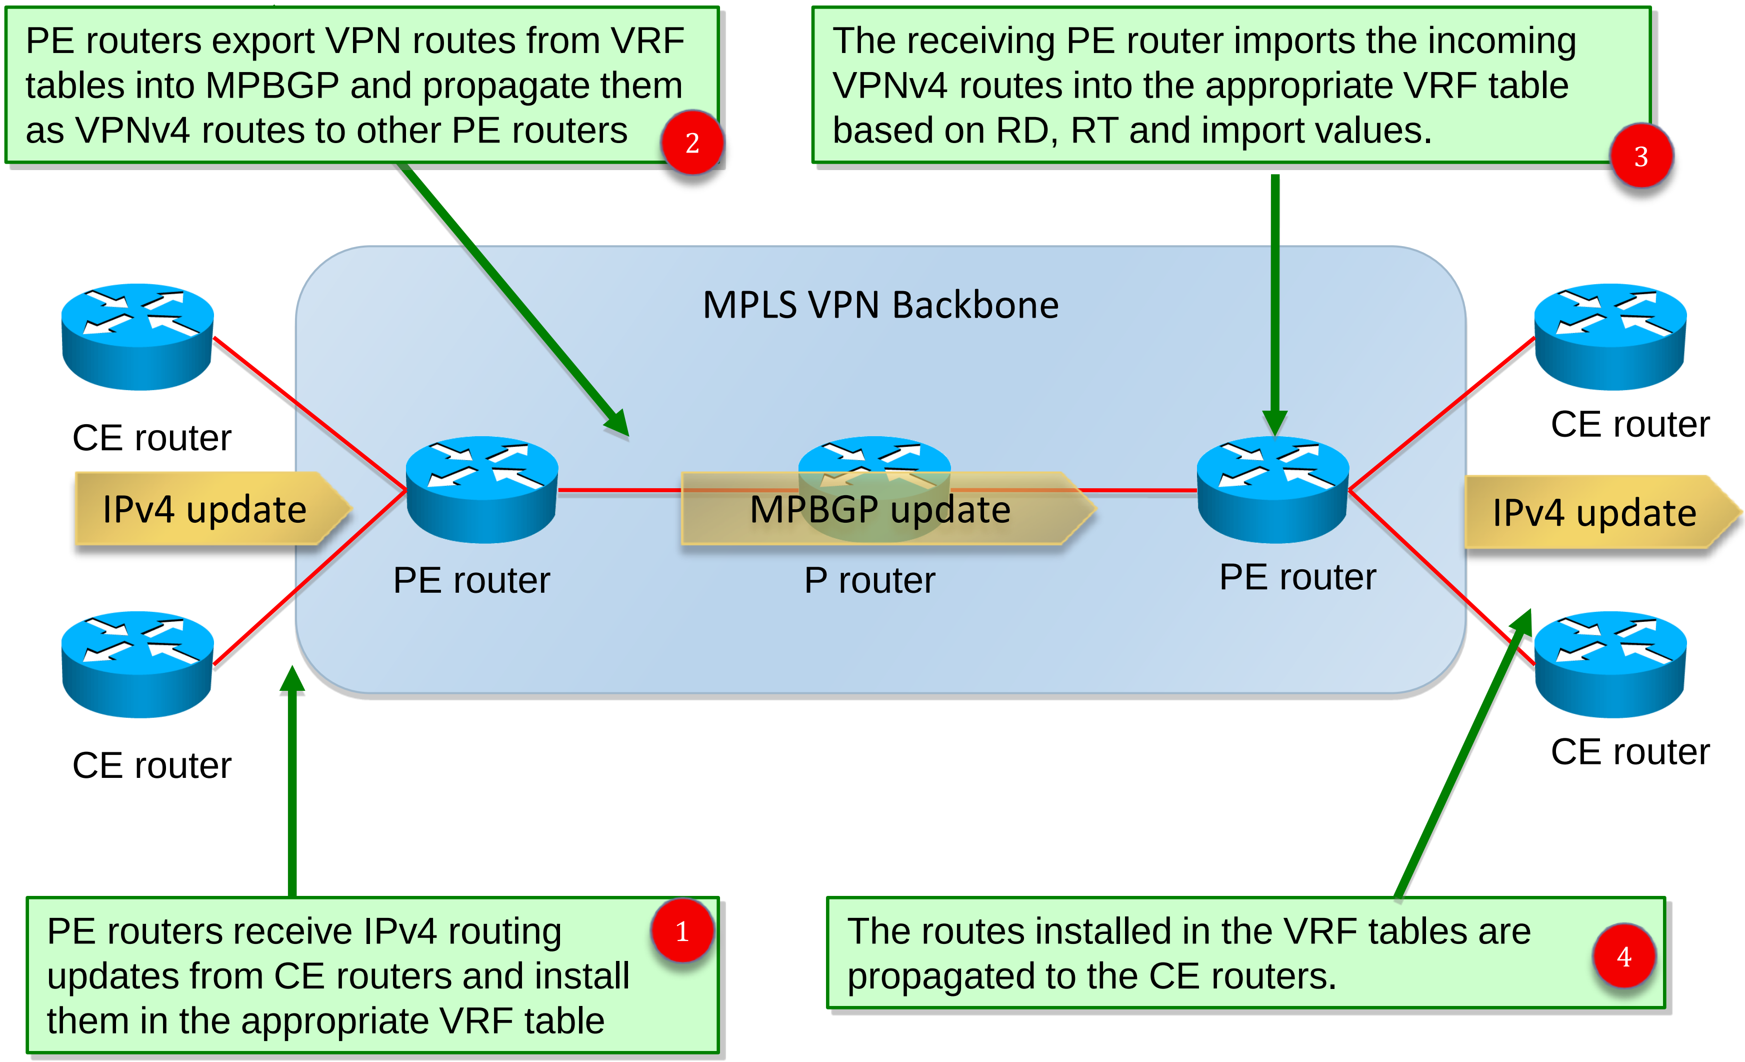

VPN PHP

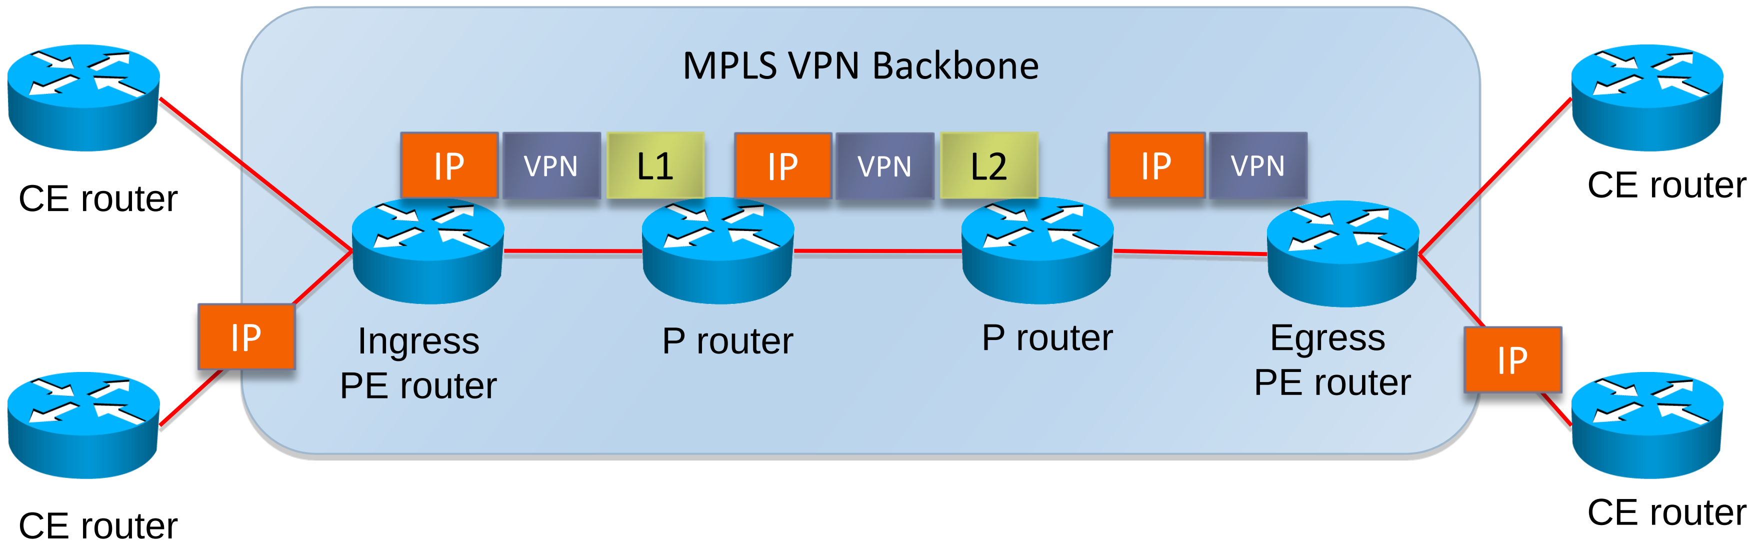

- PHP = Penultimate Hop Popping (try saying that 3 times fast!)
- The last PE router receives a labeled packet and sends an IP packet
  - It performs a label lookup (LFIB) and a routing lookup (FIB)
- The last P router can decide to remove the label before seding the packet to the last PE router.
  - Now, the last PE router only performs a FIB lookup

Penultimate Hop Popping (PHP)

- PHP optimizes MPLS performance by reducing CPU effort on Edge LSRs.
- The Edge LSR advertises a pop or implicit null label (value of 3) to a neighbor.
- The pop tells the neighbor to use PHP.

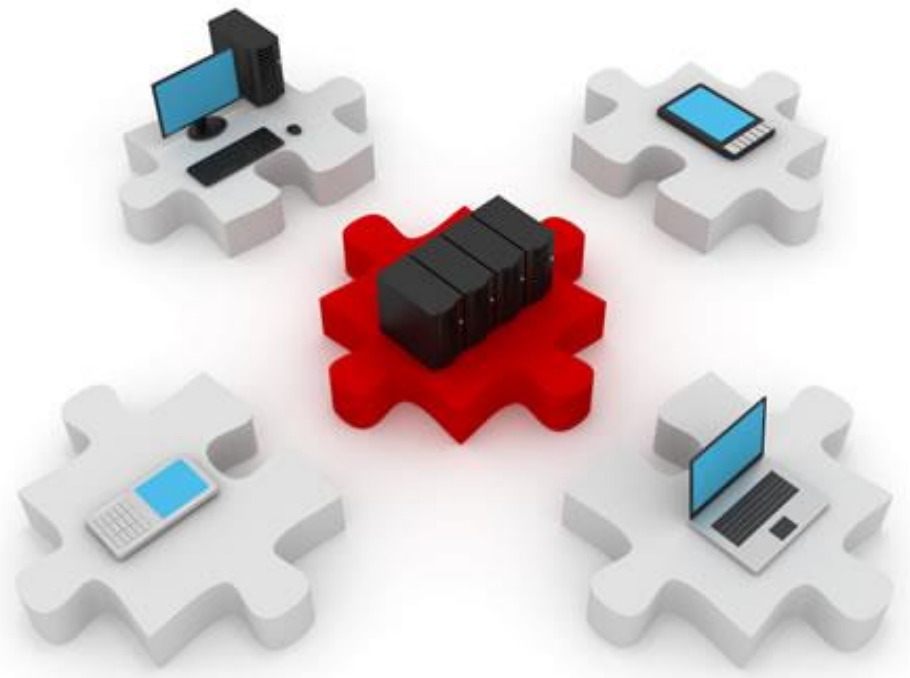

Configuring frame mode MPLS

Lesson 4.3: Implementing Frame Mode MPLS

The Procedure to Configure MPLS

- Configure CEF.
- Configure MPLS on a frame mode interface.
- (Optional) Configure the MTU size in label switching.

Step 1: Configure CEF

- To enable MPLS, you must first configure CEF.
- Configure CEF:
  - Enable CEF switching to create the FIB table.
  - Enable CEF switching on all core interfaces.
- Configure MPLS on a frame mode interface.
- (Optional) Configure the MTU size in label switching.

Commands for Configuring CEF

```
Router(config)#ip cef [distributed]
```

- Starts CEF switching and creates the FIB table
- The **distributed** keyword configures distributed CEF (running on VIP or line cards)
- All CEF-capable interfaces run CEF switching

```
Router(config-if)#ip route-cache cef
```

- Enables CEF switching on an interface
- Usually not needed

Monitoring IP CEF

- Display a summary of the FIB:

```
Router#show ip cef detail
  IP CEF with switching (Table Version 6), flags=0x0
  6 routes, 0 reresolve, 0 unresolved (0 old, 0 new)
  9 leaves, 11 nodes, 12556 bytes, 9 inserts, 0 invalidations
  0 load sharing elements, 0 bytes, 0 references
  2 CEF resets, 0 revisions of existing leaves
  refcounts:  543 leaf, 544 node

Adjacency Table has 4 adjacencies
0.0.0.0/32, version 0, receive
192.168.3.1/32, version 3, cached adjacency to Serial0/0.10
0 packets, 0 bytes
  tag information set
    local tag: 28
    fast tag rewrite with Se0/0.10, point2point, tags imposed: {28}
  via 192.168.3.10, Serial0/0.10, 0 dependencies
    next hop 192.168.3.10, Serial0/0.10
    valid cached adjacency
    tag rewrite with Se0/0.10, point2point, tags imposed: {28}
```

Using `show ip cef` Parameters

Parameter	Description
Unresolved	(Optional) Displays unresolved FIB entries
Summary	(Optional) Displays a summary of the FIB
Network	(Optional) Displays the FIB entry for the specified destination network
Mask	(Optional) Displays the FIB entry for the specified destination network and mask
Longer-prefixes	(Optional) Displays the FIB entries for all the specific destinations
Detail	(Optional) Displays detailed FIB entry information
type number	(Optional) Lists the interface type and number for which to display FIB entries

Step 2: Configure MPLS on Frame Mode Interface

- Configure CEF.
- Configure MPLS on a frame mode interface:
  - Enable label switching on a frame mode interface.
  - Start LDP or TDP label distribution protocol.
- (Optional) Configure the MTU size in label switching.

- TDP is a Cisco-proprietary protocol (TCP port 711)
- LDP is the industry standard (UDP port 646)
- They work the same but are not compatible
  - LDP and TDP manage the allocation and distribution of labels between MPLS neighbors

Commands for Configuring MPLS on a Frame Mode Interface

```
Router(config-if)#mpls ip
```

- Enables label switching on a frame mode interface
- Starts LDP on the interface

```
Router(config-if)#mpls label protocol [tdp | ldp | both]
```

- Starts selected label distribution protocol on the specified interface

Configuring MPLS on a Frame Mode Interface: Example 1

```
ip cef
!
interface Serial 2/1
 mpls ip
!
interface FastEthernet 0/0
 mpls ip
```

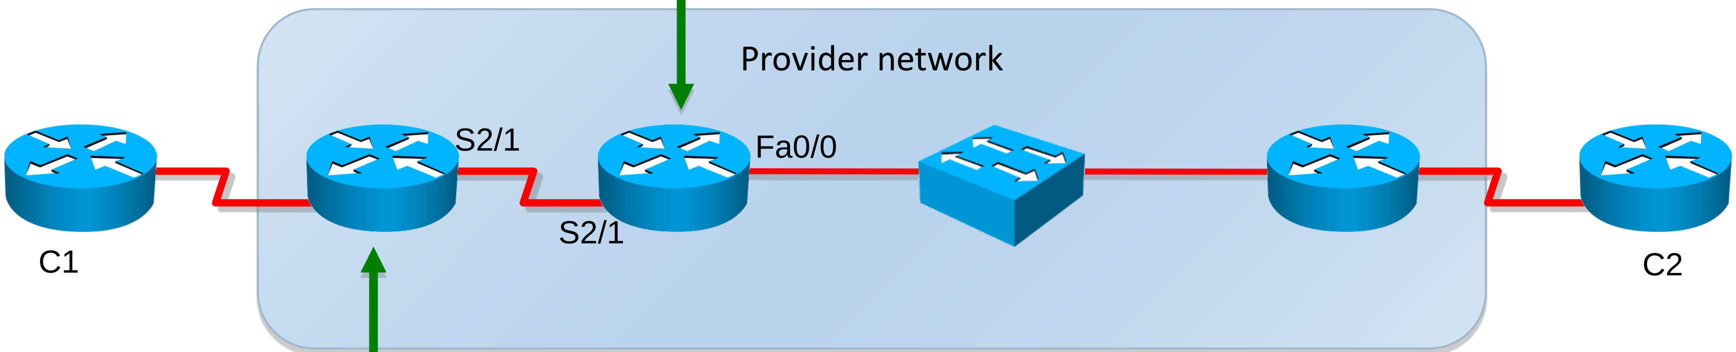


```
ip cef
!
interface Serial 2/1
 mpls ip
!
interface Serial 3/1
 ip access-group NoLDP in
!
ip access-list NoLDP deny udp any any eq 646
ip access-list NoLDP permit ip any any
```

Service providers use ACLs to block customers from using either TDP or LDP with their routers.

Configuring MPLS on a Frame Mode Interface: Example 2

```
ip cef
!
interface Serial 1/0
 mpls ip
 mpls label protocol tdp
!
interface Serial 2/1
 mpls ip
 mpls label protocol ldp
!
interface FastEthernet 0/0
 mpls ip
 mpls label protocol both
```

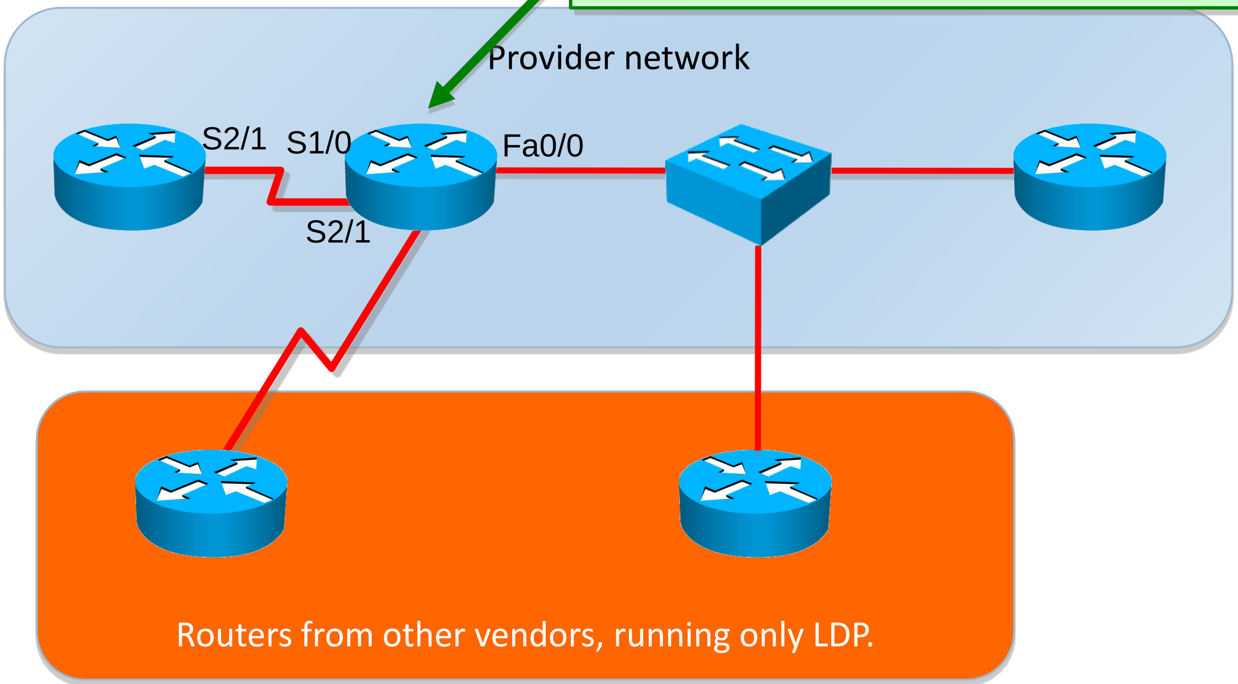

Step 3: Configure the MTU Size in Label Switching

- Configure CEF.
- Configure MPLS on a frame mode interface (**`mpls mtu 1512`**)
- Configure the MTU size in label switching:
  - Increase MTU on LAN interfaces.
  - What is the default MTU size on Ethernet interfaces?
- The MTU size represents the payload of a frame or the size of an IP packet.
- MPLS introduces one or more labels between L2 and L3 headers.
- The MTU must be increased to accommodate the new labels.
  - 1504 for one label, 1508 for two labels, and so on.
  - Switches might also need configuration to support jumbo frames
- The MTU size is automatically increased on WAN interfaces.

Summary

- MPLS provides a blend of Layer 2 switching and Layer 3 routing to forward packets using short, fixed-length labels.

- MPLS provides fast routing for large networks. Only the edge routers perform a routing lookup, and core routers forward packets based on the labels. These two functions mean faster forwarding of packets through the SP network.

- The most recent and preferred Cisco IOS platform switching mechanism is Cisco Express Forwarding (CEF), which incorporates the best of the previous switching mechanisms.

- To support multiple protocols, MPLS divides the classic router architecture into two major components: control plane and data plane.

Summary

- There are four steps for label allocation and distribution in a Unicast IP routing network and MPLS functionality, including label allocation and distribution. The following steps detail what happens:
  - The routers exchange information using routing protocol.
  - Local labels are generated.
  - Local labels are propagated to adjacent routers.
  - Every LSR builds data structures based on received labels.

- When a router receives an IP packet, the lookup done is an IP lookup. When a router receives a labeled packet, the lookup is done in the LFIB table of the router.

- Using Penultimate Hop Popping (PHP), an LSR removes the outermost label of an MPLS-tagged packet before passing the packet to an adjacent Edge LSR. The process reduces the load on the Edge LSR.

Summary

- VPNs allow you to use the shared infrastructure of a SP to implement your private networks. There are two implementation models: overlay and peer-to-peer.
- The MPLS VPN architecture offers SPs a peer-to-peer VPN architecture that combines the best features of overlay VPNs with the best features of peer-to-peer VPNs.
- MPLS VPNs use a 32-bit prefix called the route distinguisher (RD) to convert non-unique 32-bit customer IPv4 addresses into 64-bit unique addresses that can be transported.
- MPLS works by prepending packets with an MPLS header, containing one or more "labels." This is called a label stack.

The end, finally!!!