MPLS Workshop

Jan 15th to Jan 19th 2009

Abdul Rahim , arahim@cisco.com

Outline:

- MPLS Concepts and Overview
- Label Distribution Protocol
- Configuring MPLS
- MPLS in BGP Transit AS

MPLS Concepts & Overview

Traditional IP Forwarding

- Traditional IP Forwarding is based on following:

  - Routing Protocols are used to distribute L3 routing information

  - Forwarding is based on destination addresses only

  - Routing lookups are performed on every hop

Traditional IP Forwarding (contd.)

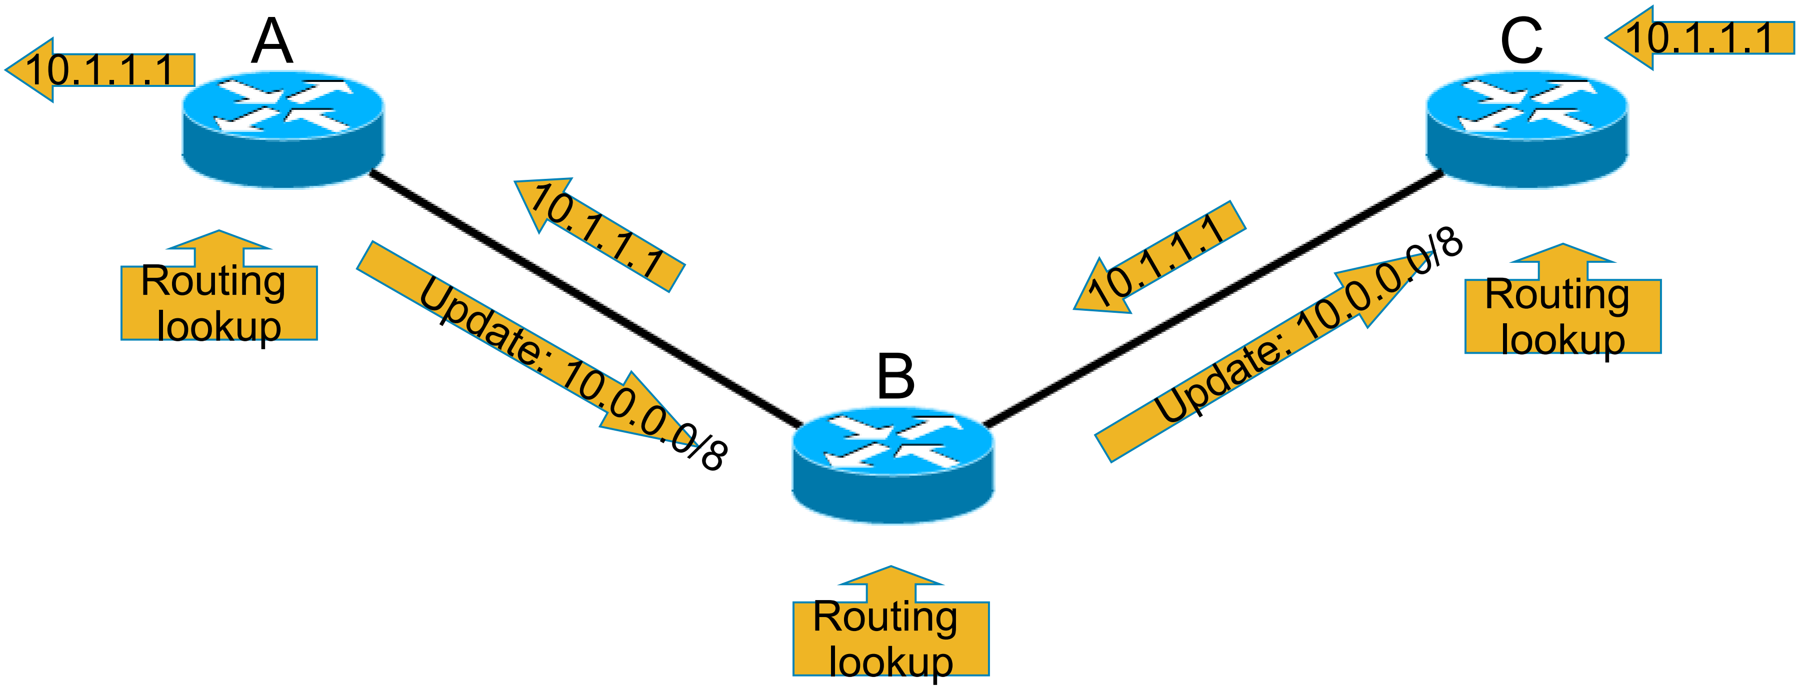

- Destination-based IP lookup is performed on every hop
- Every router needs to carry full route

IP Limitations:
Traffic Engineering

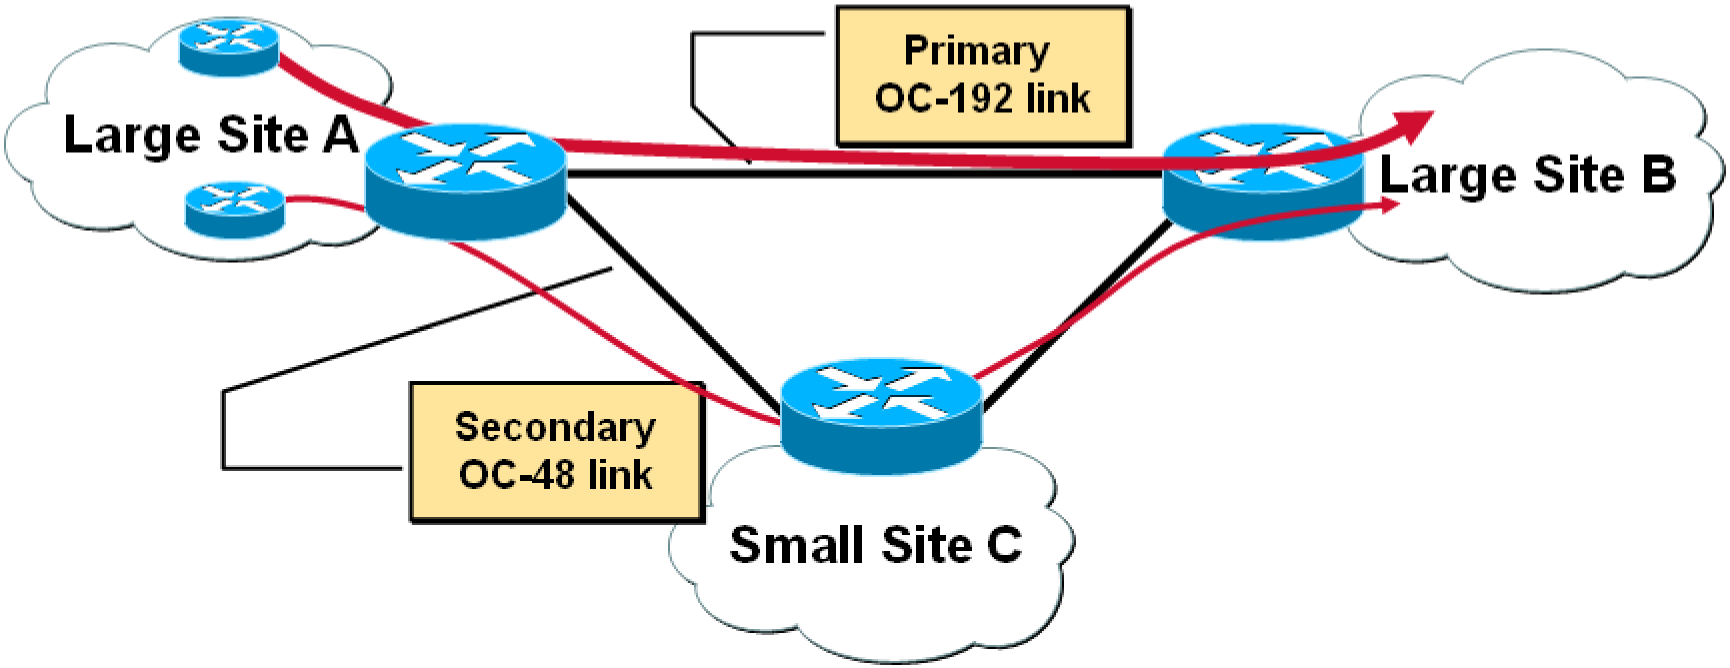

- Traffic can be forwarded based on other parameters like source address
- Load-sharing across un-equal cost paths can be achieved

IP over ATM

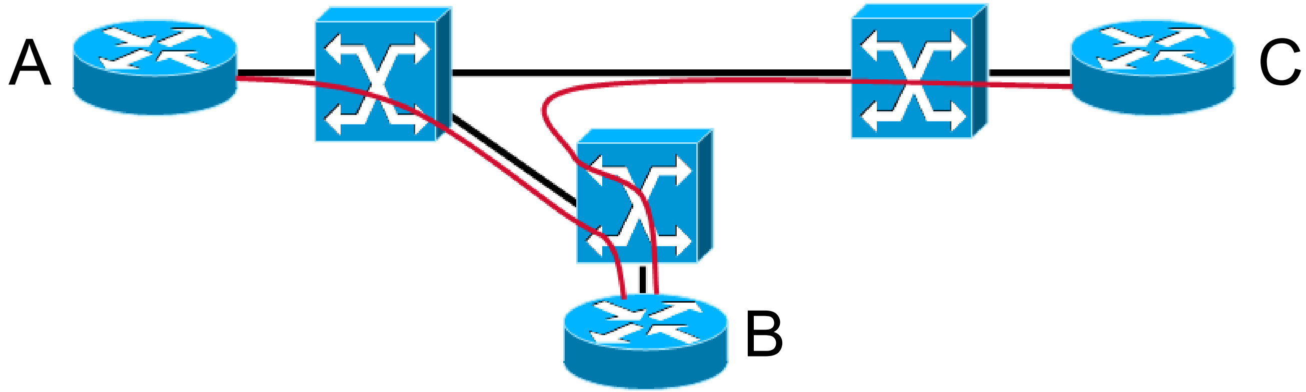

- Layer 2 topology may be different from Layer 3 topology, resulting in suboptimal paths and link utilization.
- Layer 2 devices have no knowledge of Layer 3 routing information—virtual circuits must be manually established.
- Even if the two topologies overlap, the hub and spoke topology is usually used because of easier management.

Basic MPLS Concepts

- MPLS is a new forwarding mechanism in which packets are forwarded based on labels
- Labels may corresponds to IP destination networks (equal to traditional IP forwarding)
- Labels may also corresponds to other parameters such as source address
- MPLS was designed to support forwarding of other protocols as well

Basic MPLS Concepts

MPLS (Multi Protocol Label Forwarding) uses Labels to forward packets.

Labels usually correspond to IP destination networks (equal to traditional IP forwarding).

Labels can also correspond to other parameters, such as ATM VC (in case of ATM over MPLS), VLAN id (in case of Ethernet over MPLS) or even QoS parameters

Now packet forwarding is no longer strictly tied to IP destination address

MPLS is an IETF standard based on RFC 3031, 3032

Basic MPLS Concepts

Labels assigned to packets can be based on:

- Destination prefix
- Traffic Engineering tunnel
- VPN-ID, ATM VC, VLAN ID
- Class of Service

MPLS Forwarding

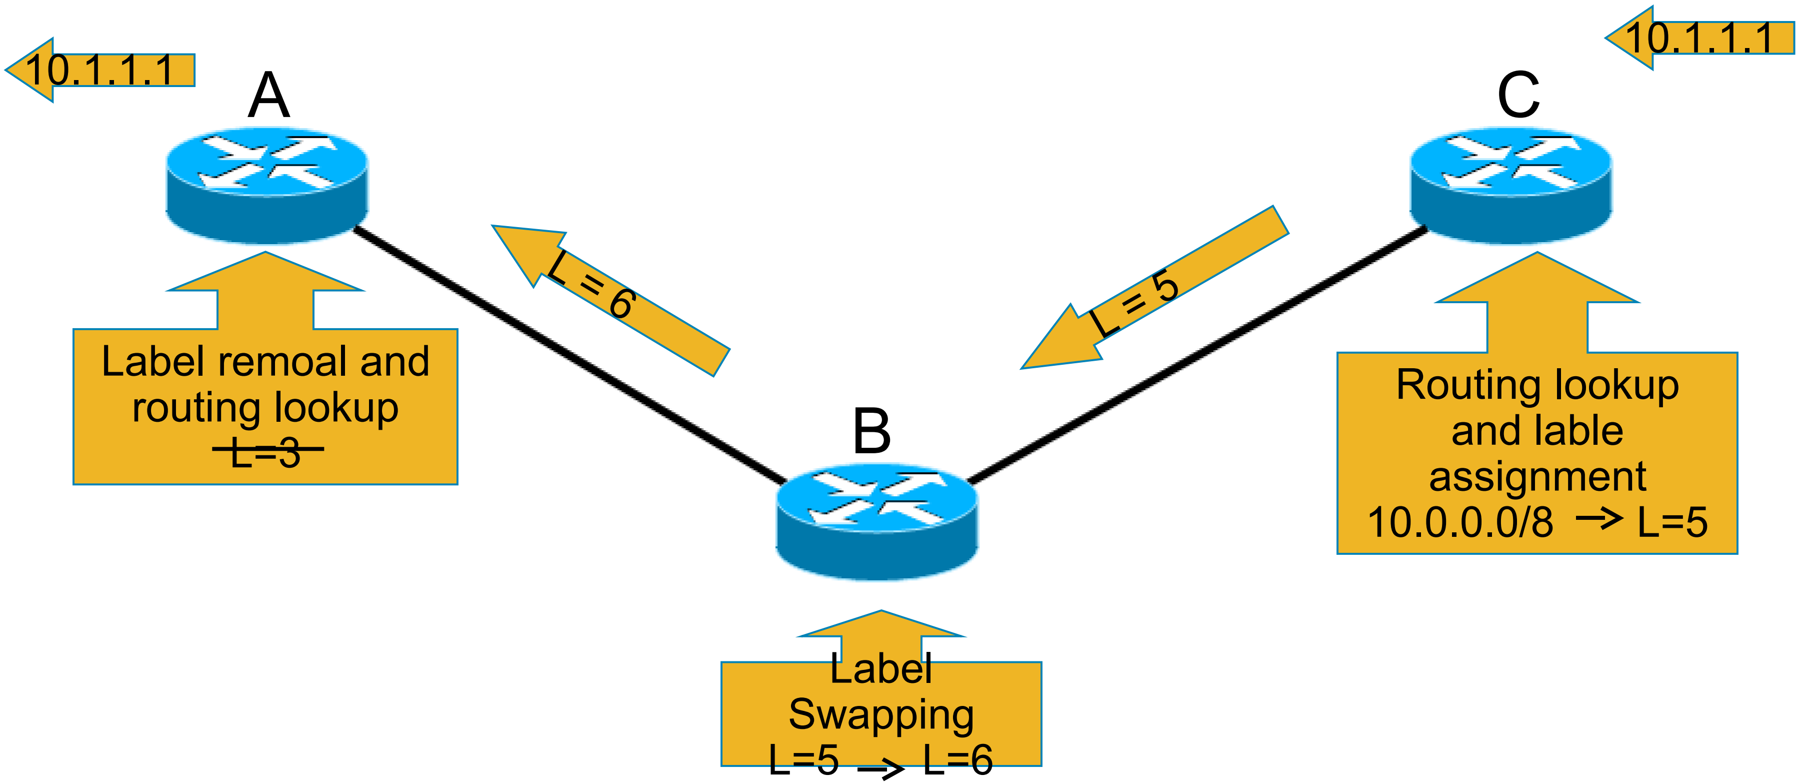

- Only edge routers must perform the routing lookup
- Core routers switch packets based on simple lable lookup and swap

MPLS versus IP over ATM

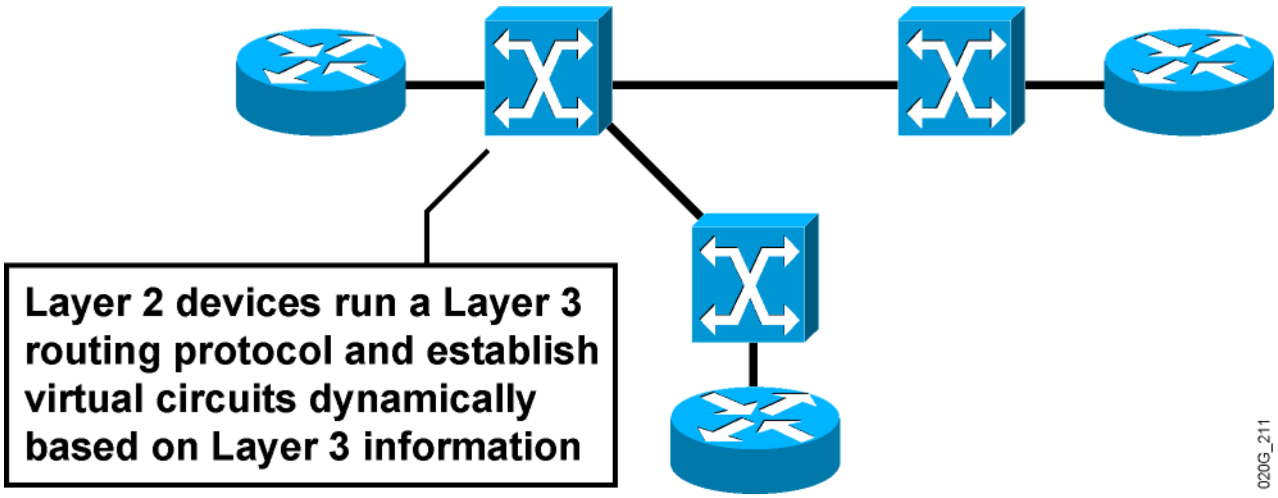

Layer 2 devices are IP-aware and run a routing protocol.

MPLS Versus IP over ATM (contd.)

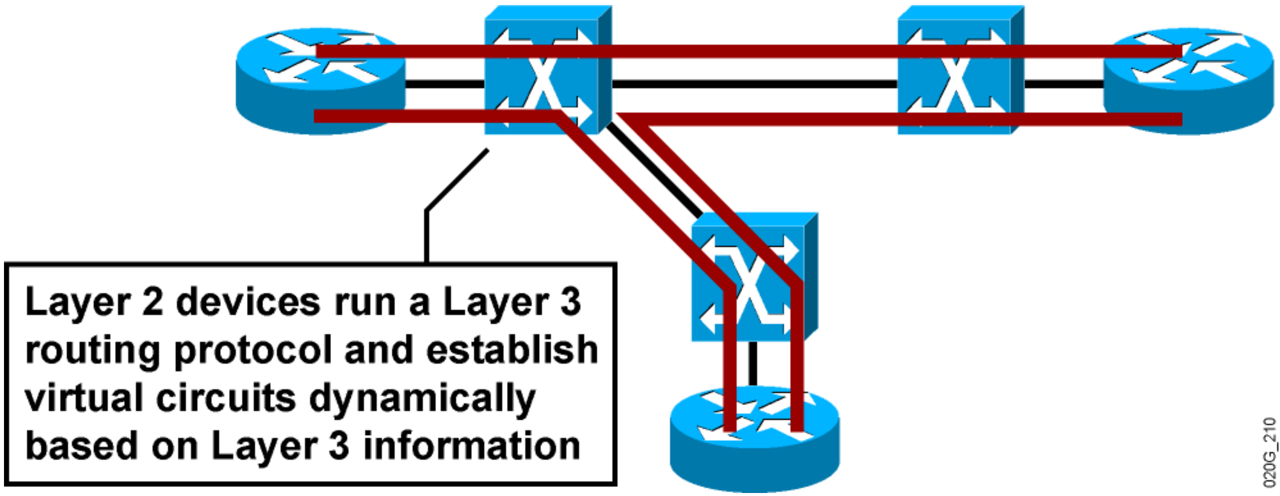

- There is no need to manually establish virtual circuits.

MPLS Versus IP over ATM

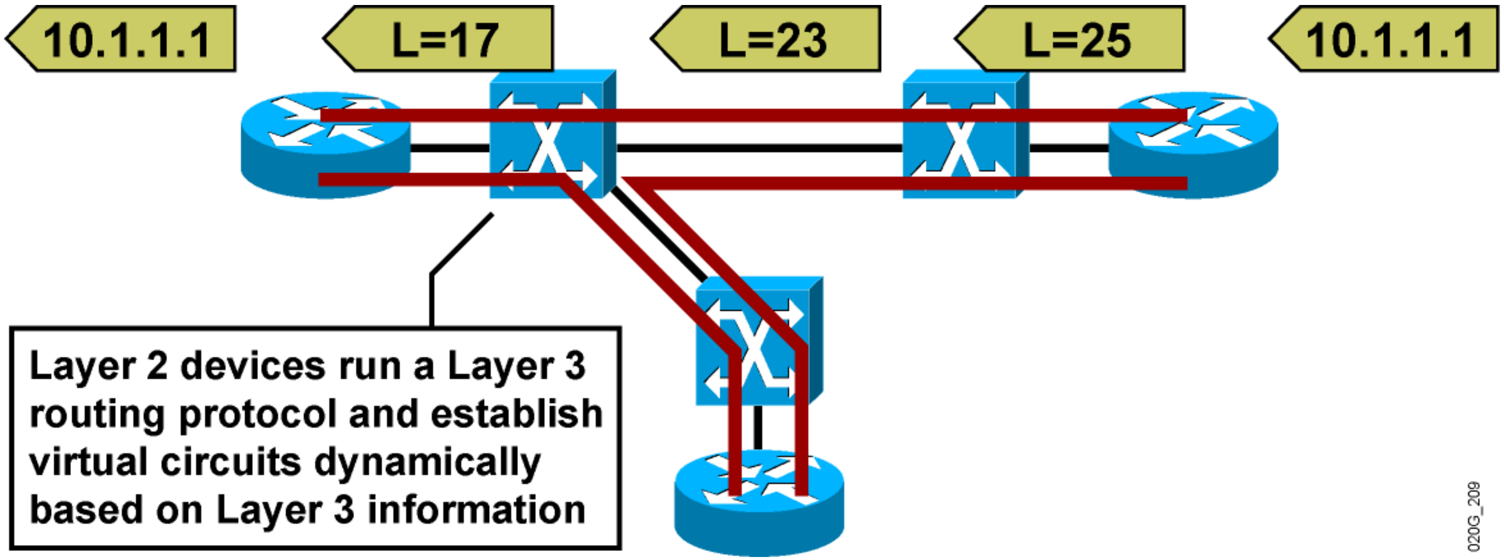

- Layer 2 devices are IP-aware and run a routing protocol.
- There is no need to manually establish virtual circuits.
- MPLS provides a virtual full-mesh topology.

Traffic Engineering with MPLS

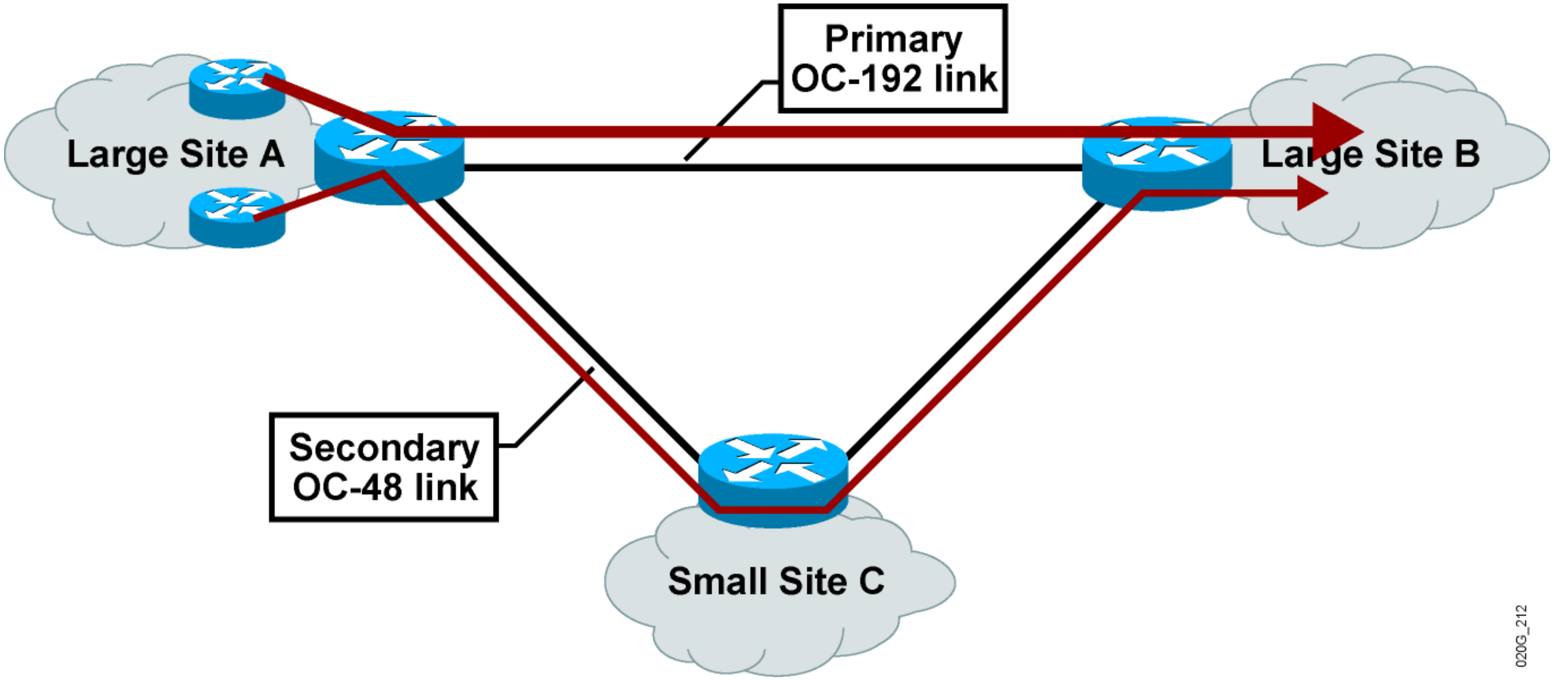

Traffic can be forwarded based on labels via the primary and secondary links

Load sharing across unequal paths can be achieved.

We shall cover MPLS Traffic Engineering in detail in later modules

MPLS Benefits

- Separates Control plane & the forwarding plane
- Only ingress router needs to look up the network layer & make routing decision. Other LSRs only swap labels
- <u>Source Based routing</u>: e.g. explicit routes in MPLS-TE
- <u>Scalability</u>: Hierarchy of Routing (via label stacking)
- <u>AnyThing over MPLS (AToM):</u> Labels are common binding between different Layer 2 technologies like ATM, Ethernet.

MPLS Architecture

- MPLS has two major components:
  - Control plane—exchanges Layer 3 routing information and labels
  - Data plane—forwards packets based on labels
- Control plane contains complex mechanisms to exchange routing information, such as Open Shortest Path First (OSPF), Enhanced Interior Gateway Routing Protocol (EIGRP), Intermediate System-to-Intermediate System (IS-IS), and BGP, and to exchange labels, such as Tag Distribution Protocol (TDP), label distribution protocol (LDP), BGP, and Resource Reservation Protocol (RSVP).
- Data plane has a simple forwarding engine.
- Control plane maintains contents of the label-switching table (label forwarding information base, or LFIB).

MPLS Architecture

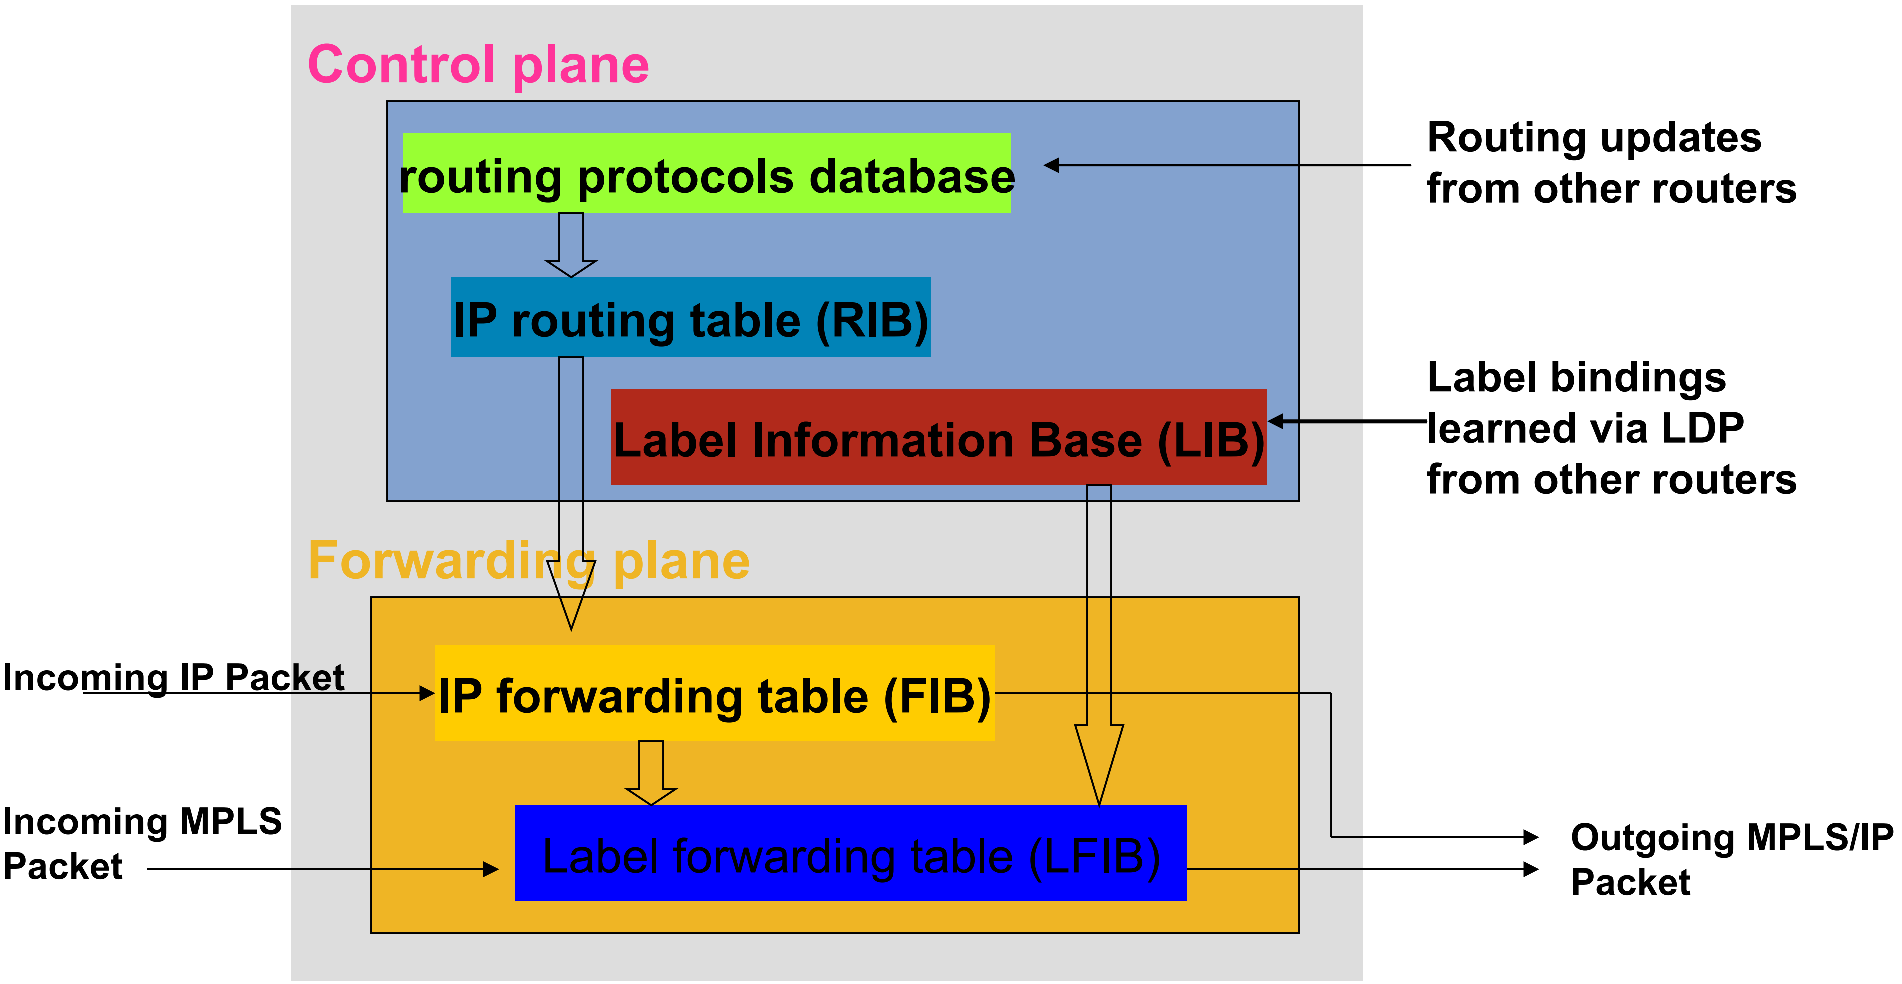

Population of RIB/FIB/LIB/LFIB in an MPLS router

MPLS Architecture

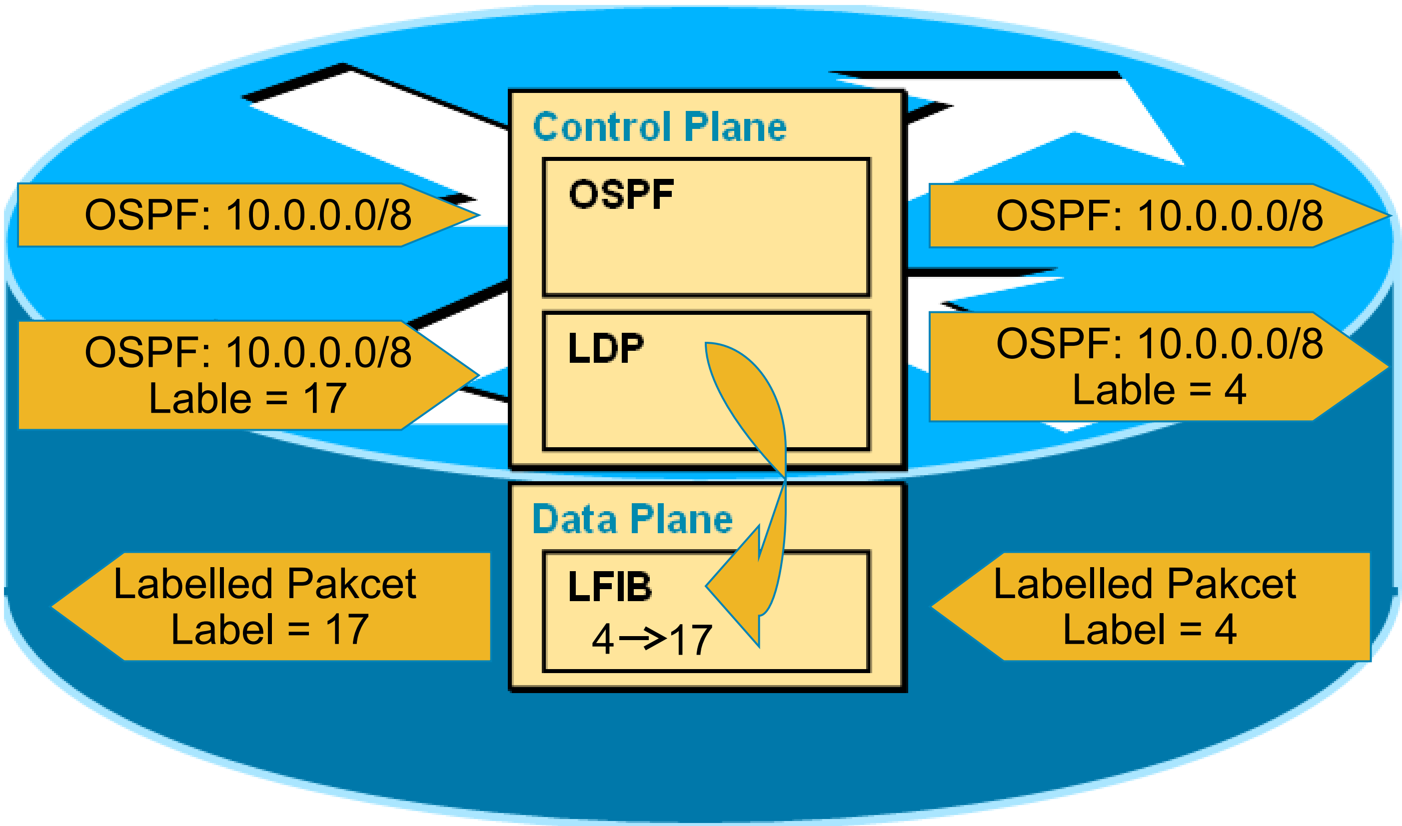

- Router functionality is divided into two major parts: Control plane and Data plane

MPLS Terminology: FEC

- FEC (Forwarding Equivalence Class)

  Group of IP packets forwarded in the same manner (e.g. over same forwarding path)

- A FEC can represent a: Destination IP prefix, VPN ID, ATM VC, VLAN ID, Traffic Engineering tunnel, Class of Service.

MPLS Terminology:
Label Switch Path (LSP)

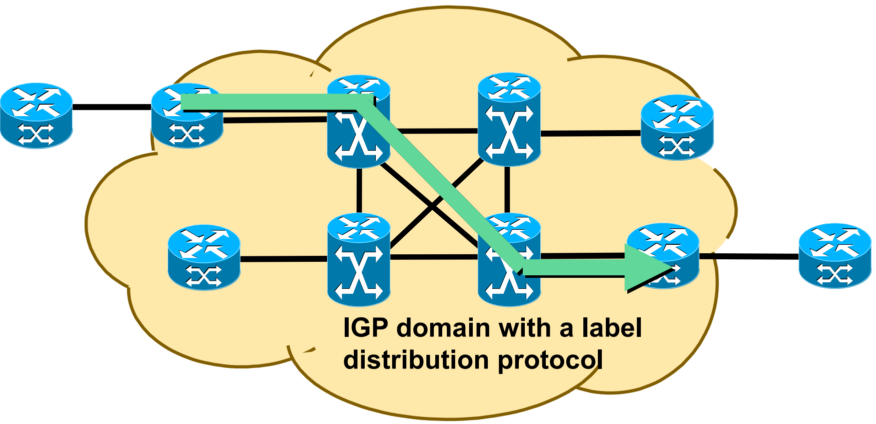

LSP follows IGP shortest path

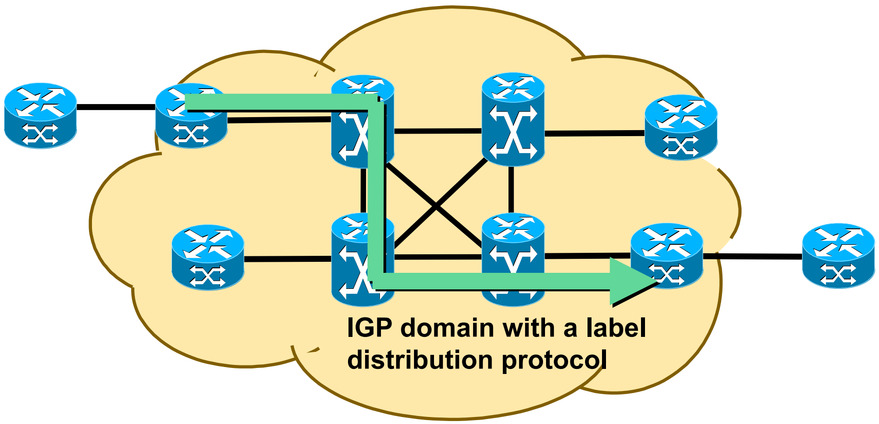

LSP diverges from IGP shortest path

- LSPs are derived from IGP routing information
- LSPs may diverge from IGP shortest path
  - LSP tunnels (explicit routing) with TE
- LSPs are unidirectional
  - Return traffic takes another LSP

MPLS terminology: LSR, LER

- **Label Edge Router (LER)**

Label Switch Router (LSR)

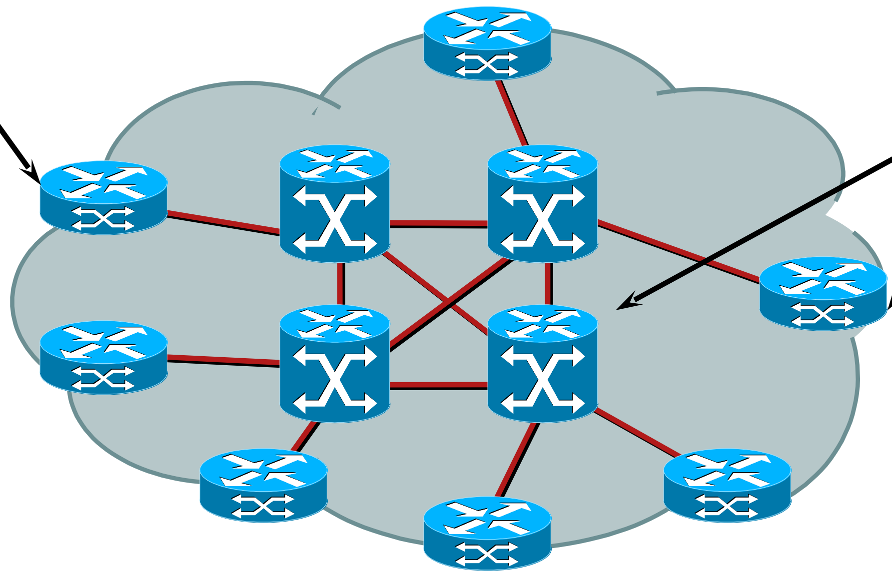

- LSR (Label Switch Router) is any netwrok router/switch running MPLS label switching
- LER (Label Edge Router) is an edge LSR. Also referred to as PE (Provider Edge) router

Label Switch Router

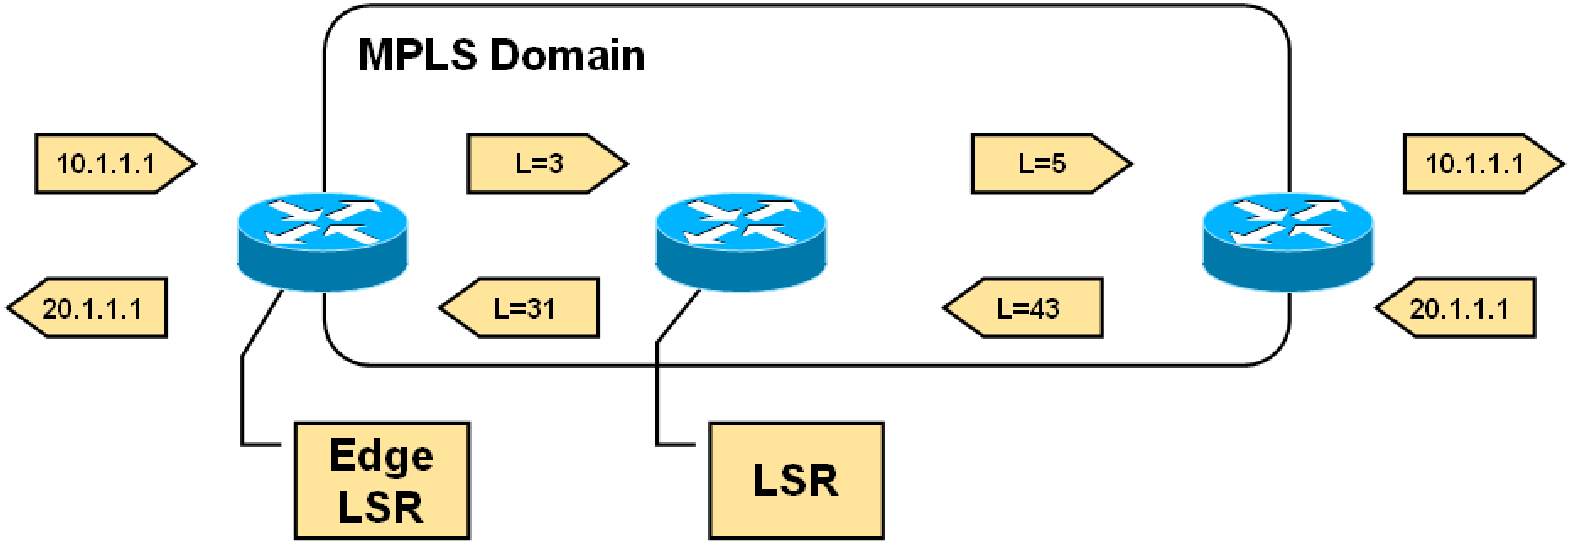

- **Label switch router (LSR)** primarily forwards labeled packets (label swapping)
- **Edge LSR** primarily labels IP packets and forwards them into MPLS domain, or removes labels and forwards IP packets out of the MPLS domain

Architecture of LSR

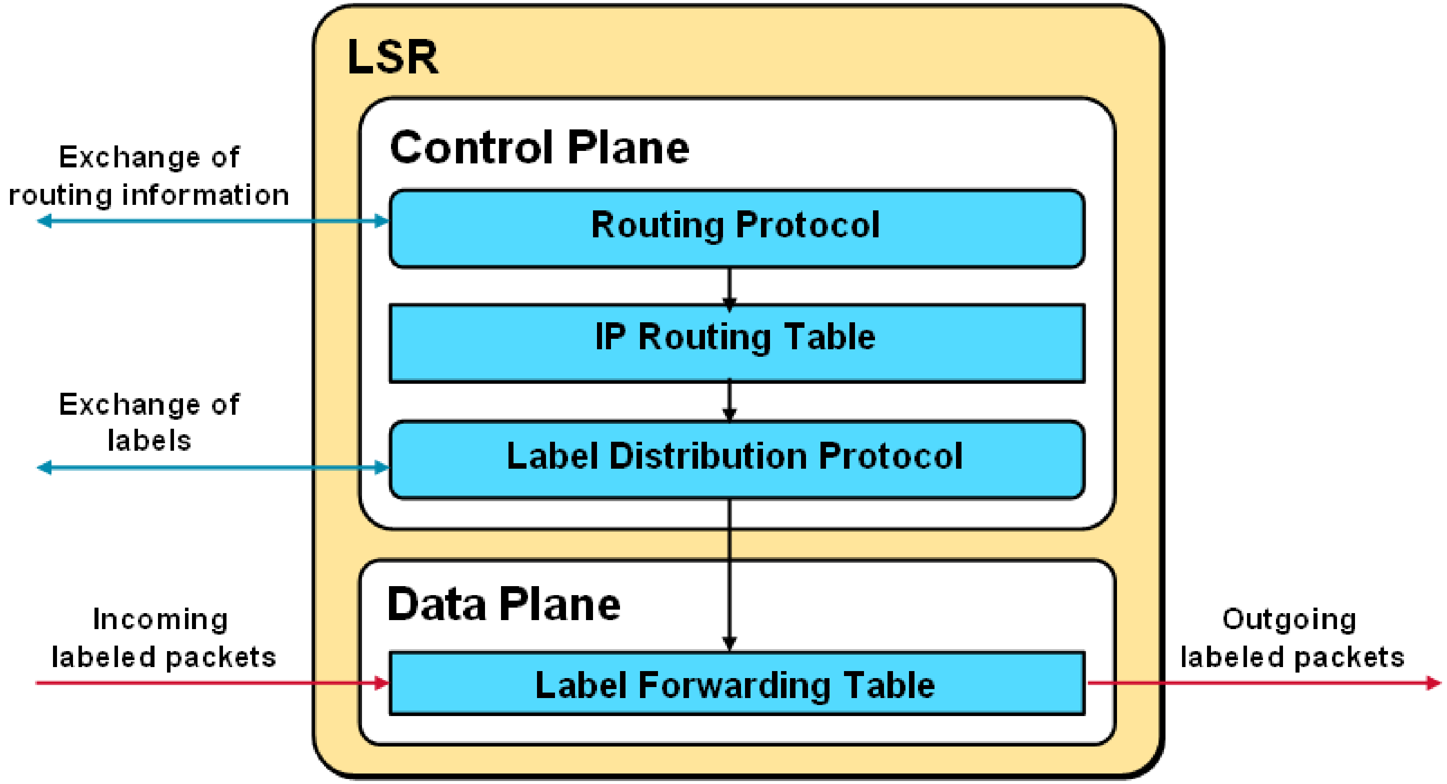

LSRs primarily forward labeled packets

Architecture of Edge LSR

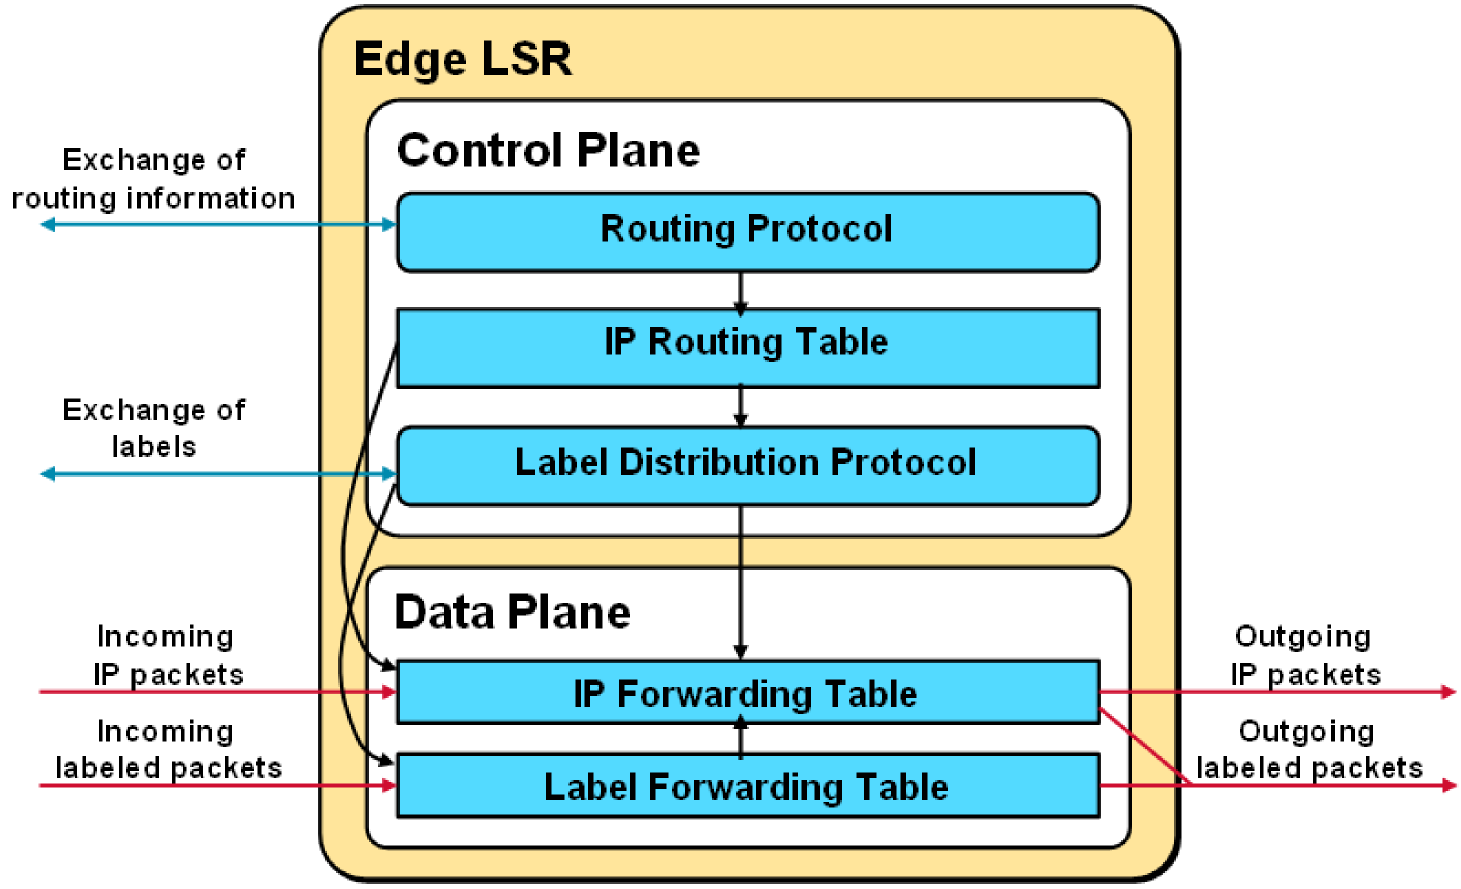

MPLS terminology: Upstream and Downstream LSRs

- Rtr-C is the downstream neighbor of Rtr-B for destination 171.68.10/24
- Rtr-B is the downstream neighbor of Rtr-A for destination 171.68.10/24
- LSRs know their downstream neighbors through the IP routing protocol

  Next-hop address is the downstream neighbor

MPLS Labels

MPLS uses a 32-bit label field that is inserted between Layer 2 and Layer 3 headers
(frame-mode).

MPLS with ATM uses the VPI, VCI fields of the ATM header as the label (cell-mode).

MPLS Labels:
Label Format (Shim header)

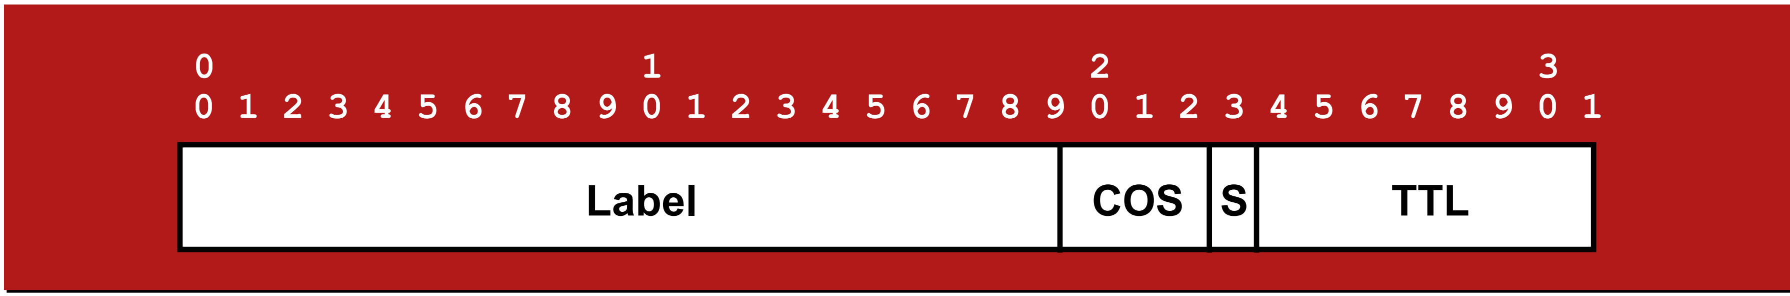

Label = 20 bits
COS/EXP = Class of Service, 3 bits
S = Bottom of Stack, 1 bit
TTL = Time to Live (Loop detection)

MPLS Labels:
Special Label values

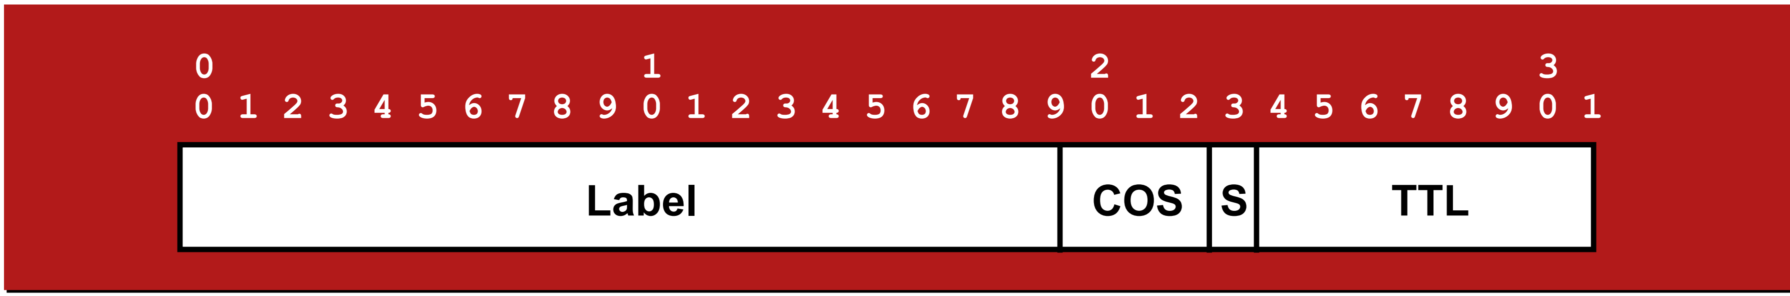

- <u>SPECIAL LABEL VALUES</u>
- 0 – IPv4 Explicit Null
- 1 – Router Alert
- 2 – IPv6 Explicit Null
- 3 – Implicit Null

Frame-Mode MPLS

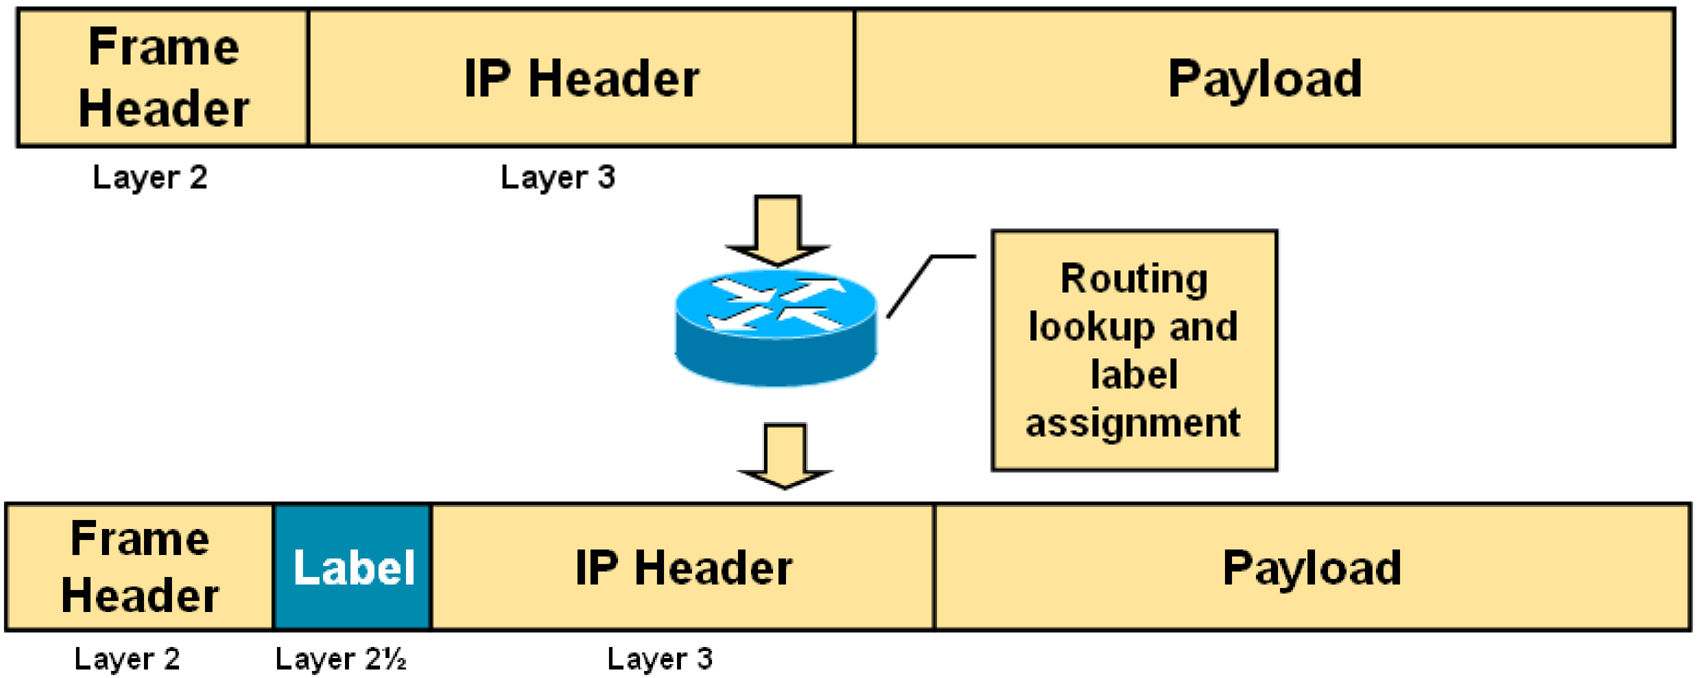

MPLS Labels:
Frame Mode Label Encapsulation

PPP Header (Packet over SONET/SDH)

PPP Header	Shim Header	Layer 3 Header

LAN MAC Label Header

MAC Header	Shim Header	Layer 3 Header

- Shim header is used with Ethernet, 802.3, or PPP frames
- Sits between the Layer 2 and Layer 3 header
- L2 frame has

  ethertype=0x8847 to indicate frame carrying MPLS unicast packet

  ethertype=0x8848 to indicate frame carrying MPLS multicast packet

MPLS Labels: Cell mode Label Encapsulation

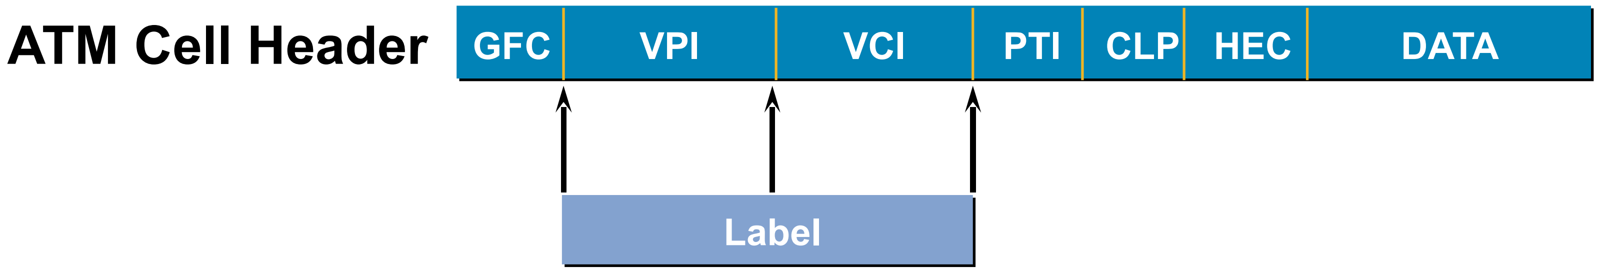

- ATM switches forward cells, not packets

- In case of label stack:
- First level label could be in VPI
- Second level label could be in VCI

MPLS Labels: Label Assignment

- Labels have local significance

  Each LSR binds his own label mappings

- LSR assigns labels to prefixes learnt in the routing table
- < Label, prefix, prefix mask > are exchanged between adjacent LSRs

MPLS Label Stack

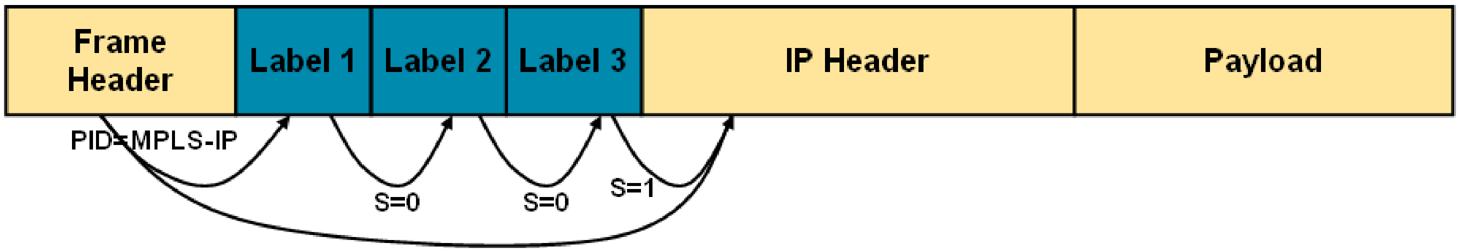

- Protocol identifier in a Layer 2 header specifies that the payload starts with a label (labels) and is followed by an IP header
- Bottom-of-stack bit indicates whether the next header is another label or a Layer 3 header
- Receiving router uses the top label only

MPLS Labels: Label Stack

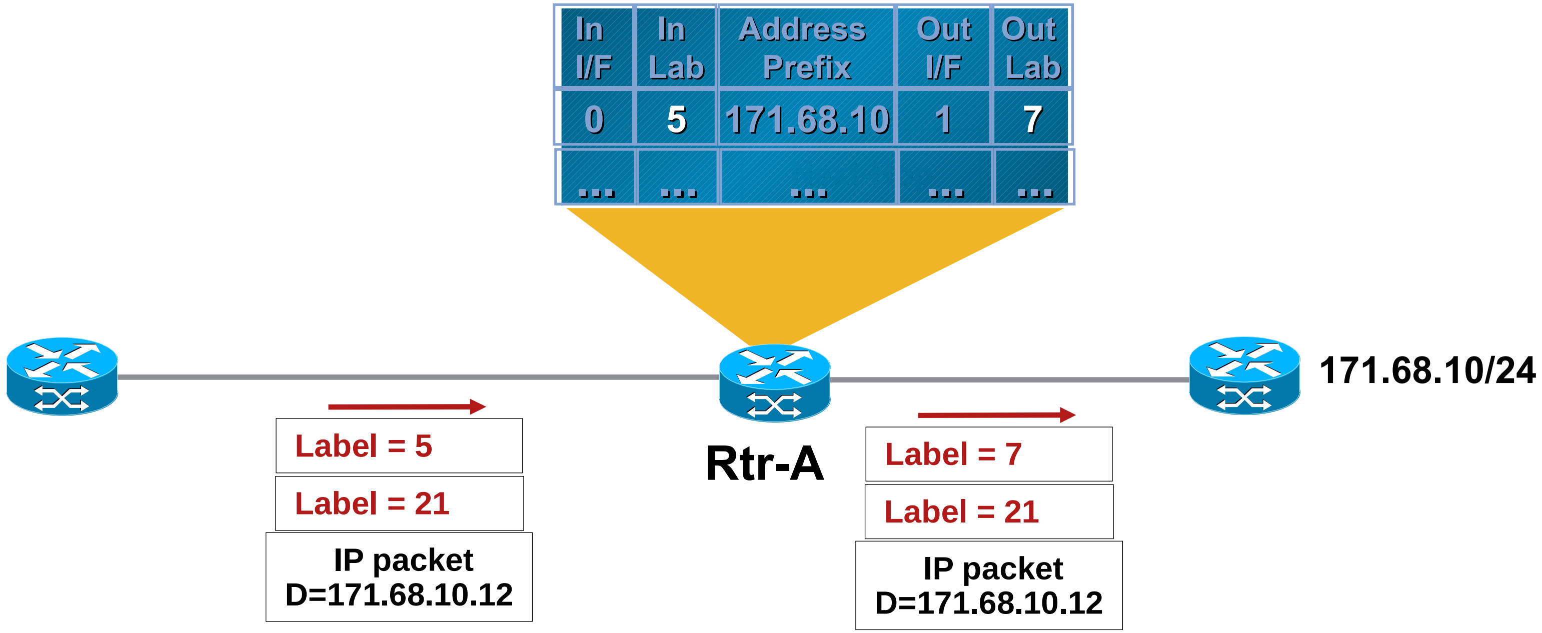

- Rtr-A forwards the labelled packet based on the label at the top of the label stack

MPLS Label Stack

- Usually only one label assigned to a packet.
- The following scenarios may produce more than one label:
  - MPLS VPNs (two labels—the top label points to the egress routers and the second label identifies the VPN)
  - MPLS TE (two or more labels—the top label points to the endpoint of the traffic engineering tunnel and the second label points to the destination)
  - MPLS VPNs combined with MPLS TE (three or more labels)

MPLS Labels: Label Distribution modes

- **Unsolicited**

  Downstream LSR advertises Label Binding to all adjacent LSRs, irrespective of whether they demand the Label binding or not

  Example: LDP, MP-iBGP

- **On-demand**

  Downstream LSR advertises Label Binding to those adjacent LSRs, who demand the Label binding

  Example: RSVP-TE, ATM

MPLS Labels: Unsolicited example

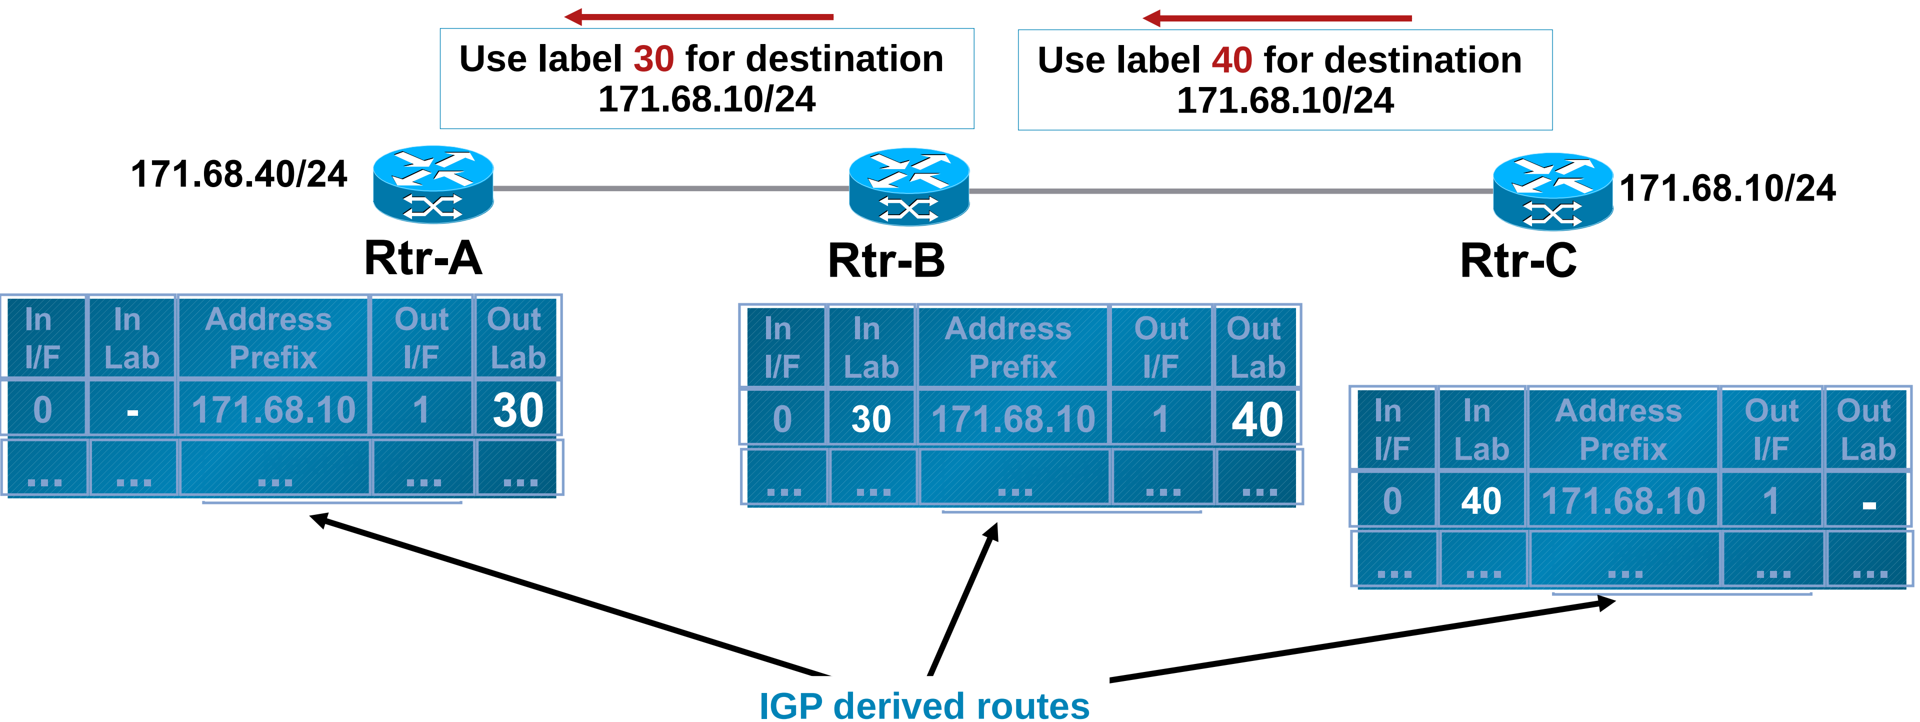

Rtr-A:

In I/F	In Lab	Address Prefix	Out I/F	Out Lab
0	-	171.68.10	1	30
...	...	...	...	...

Rtr-B:

In I/F	In Lab	Address Prefix	Out I/F	Out Lab
0	30	171.68.10	1	40
...	...	...	...	...

Rtr-C:

In I/F	In Lab	Address Prefix	Out I/F	Out Lab
0	40	171.68.10	1	-
...	...	...	...	...

- LSRs distribute labels to the upstream neighbors

MPLS Labels:
Downstream on-demand example

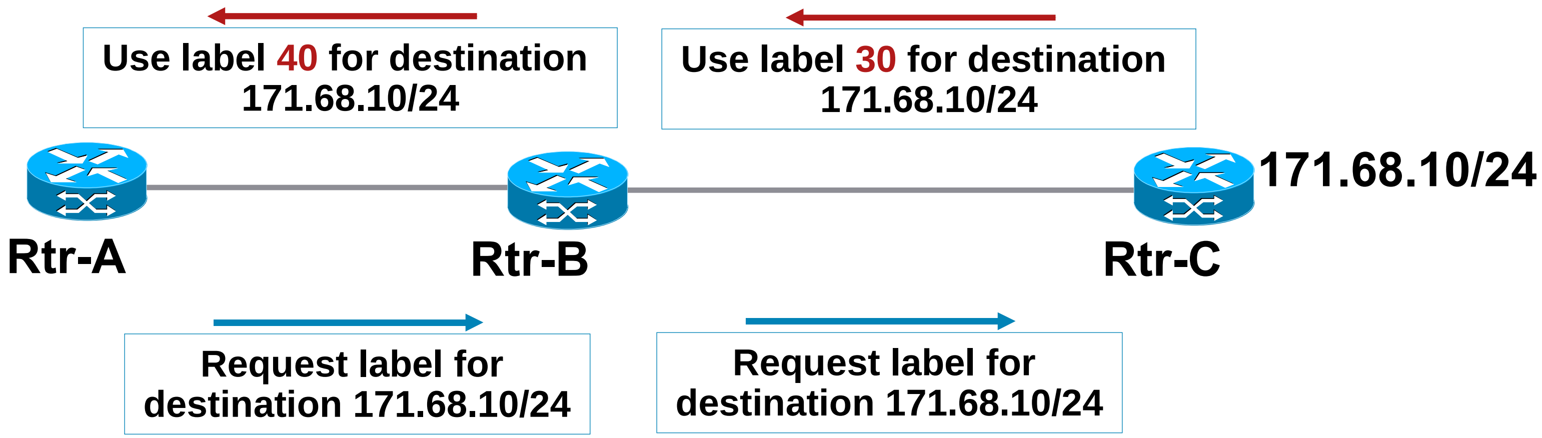

- Upstream LSRs request labels to downstream neighbors
- Downstream LSRs distribute labels upon request

MPLS Labels: Control modes

- **Independent LSP control**

  LSR binds a Label to a FEC independently, whether or not the LSR has received a Label the next-hop for the FEC

  The LSR then advertises the Label to its neighbor

  Example: LDP

- **Ordered LSP control**

  LSR only binds and advertise a label for a particular FEC if:

    it is the egress LSR for that FEC or

    it has already received a label binding from its next-hop

  Example: RSVP-TE

MPLS Labels: Retention modes

- Liberal retention
  - LSR retains labels from all neighbors

    In case, the next-hop LSR disappears, LSR already has the Out Label for the next best next-hops

    Quick convergence

    Requires more memory and label space

    Example: LDP
- Conservative retention
  - LSR retains labels only from next-hops neighbors

    LSR discards all labels for FECs which are not routing next-hops

    Free memory and label space

    Example: ATM cell mode

MPLS Labels: Penultimate Hop Popping

- The label at the top of the stack is removed (popped) by the upstream neighbor of the egress LSR
- The egress LSR requests the "popping" through the label distribution protocol

  Egress LSR advertises *implicit-null* label
- One lookup is saved in the egress LSR

MPLS Labels: Penultimate Hop Popping Example

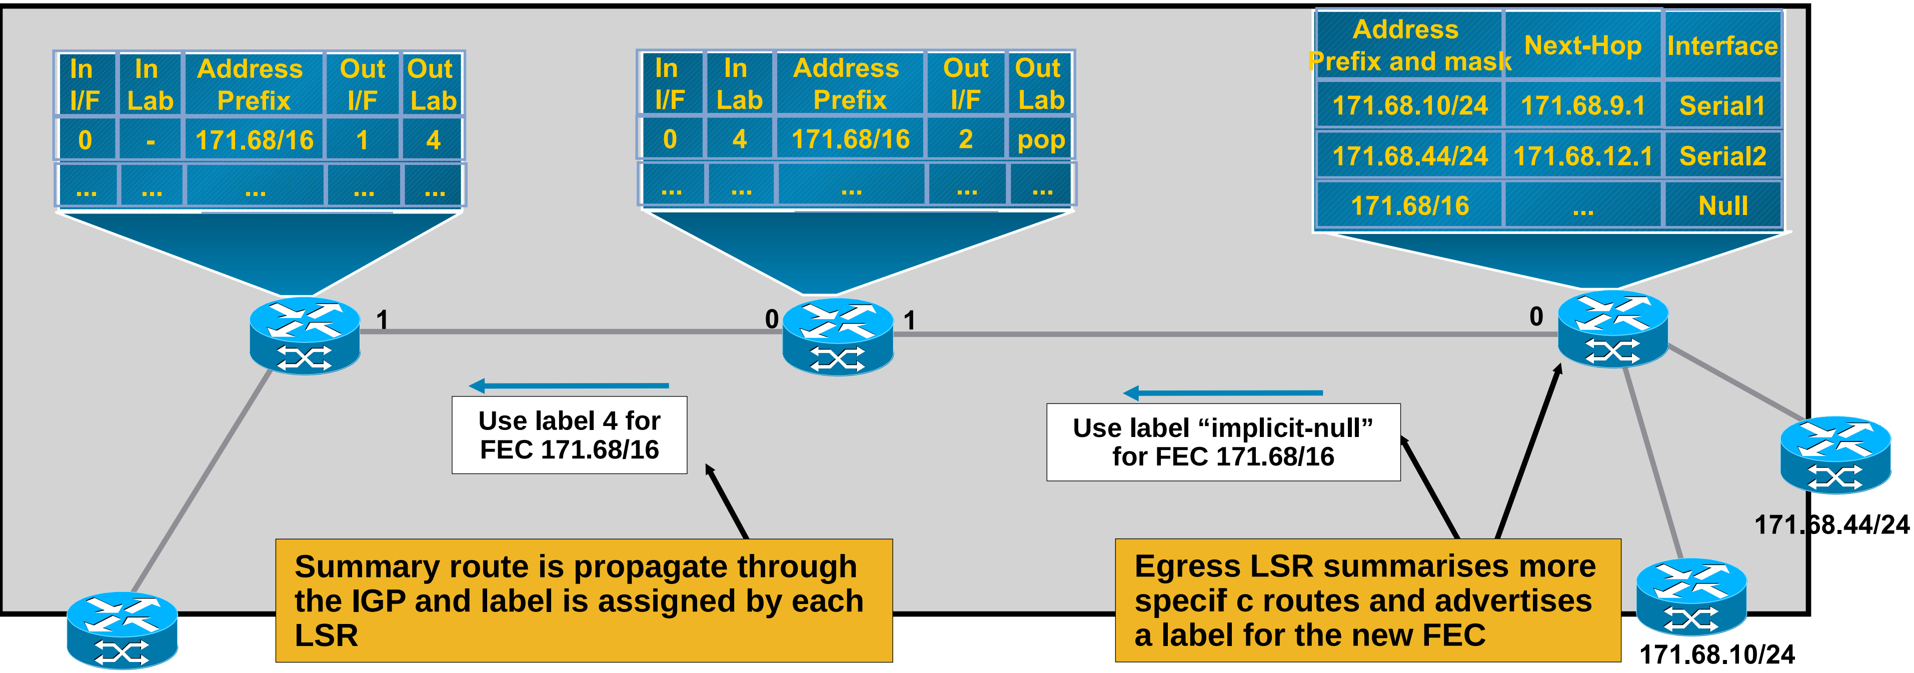

Egress LSR needs to do an IP lookup for finding more specific route

Egress LSR need NOT to receive a labelled packet

labelled will have to be popped anyway

MPLS Labels:
Penultimate Hop Popping Example (contd.)

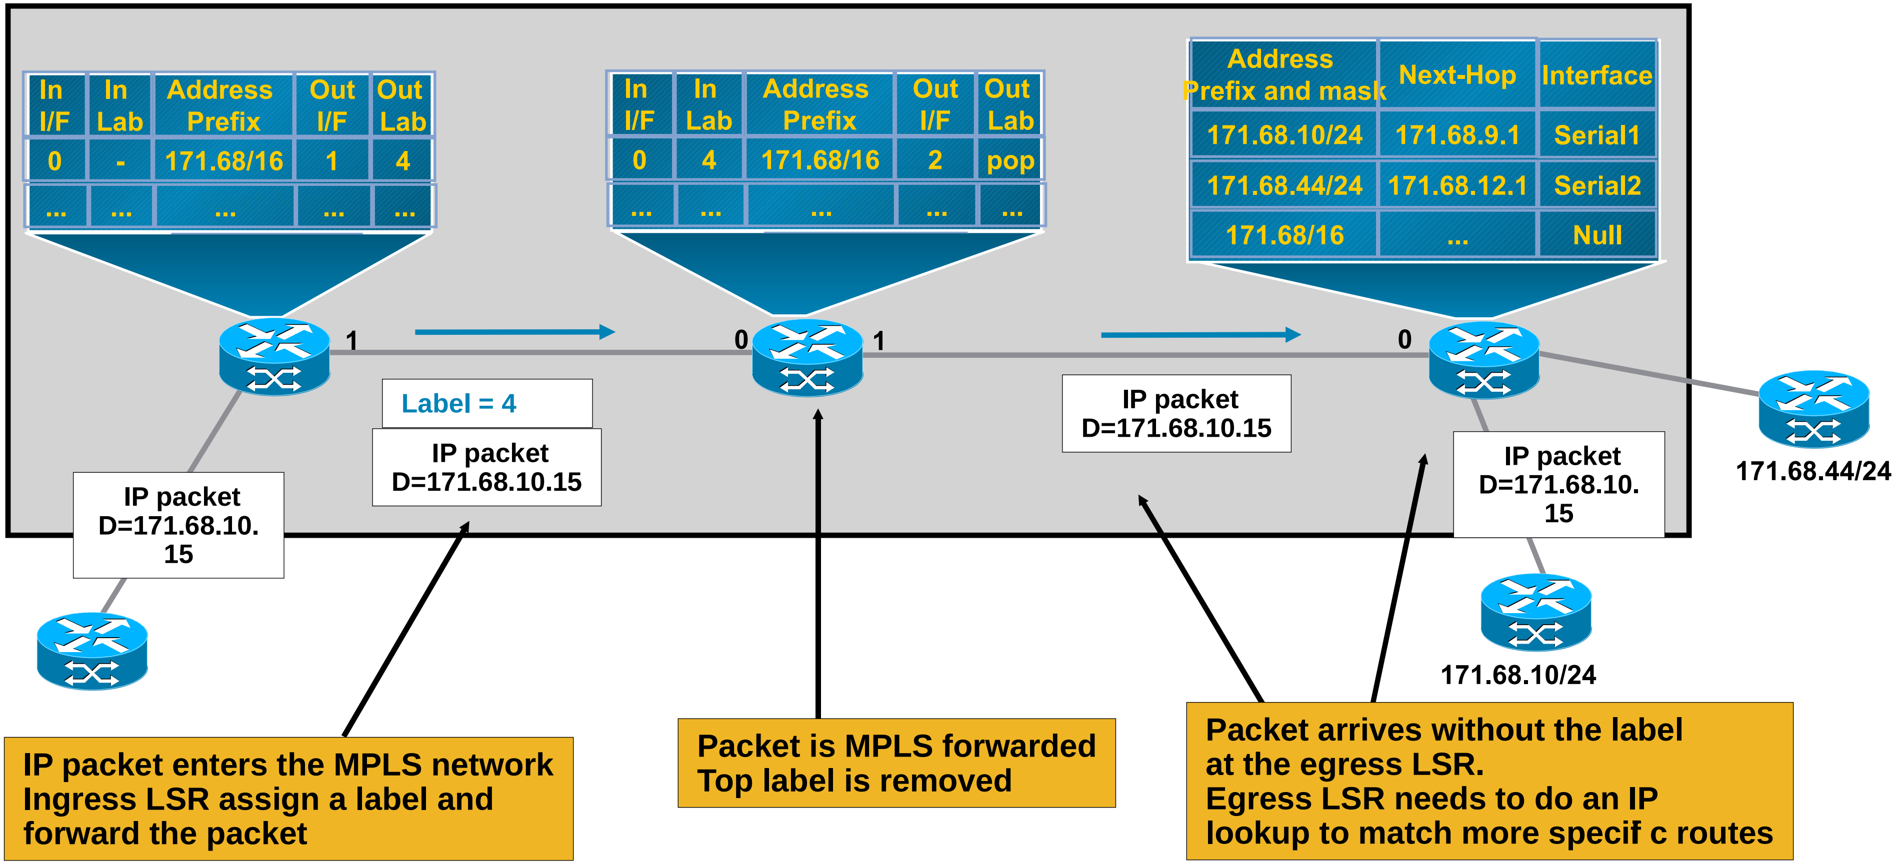

Label Switch Routers: Architecture of LSRs

- LSRs, regardless of the type, perform these functions:
  - Exchange routing information
  - Exchange labels
  - Forward packets or cells

  The first two functions are part of the control plane.

  The last function is part of the data plane.

Label Switch Routers: Exchanging Routing updates

In Label	Address Prefix	Out I'face	Out Label
	128.89	1	
	171.69	1	
	...	...	

In Label	Address Prefix	Out I'face	Out Label
	128.89	0	
	171.69	1	
	...	...	

In Label	Address Prefix	Out I'face	Out Label
	128.89	0	
	...	...	

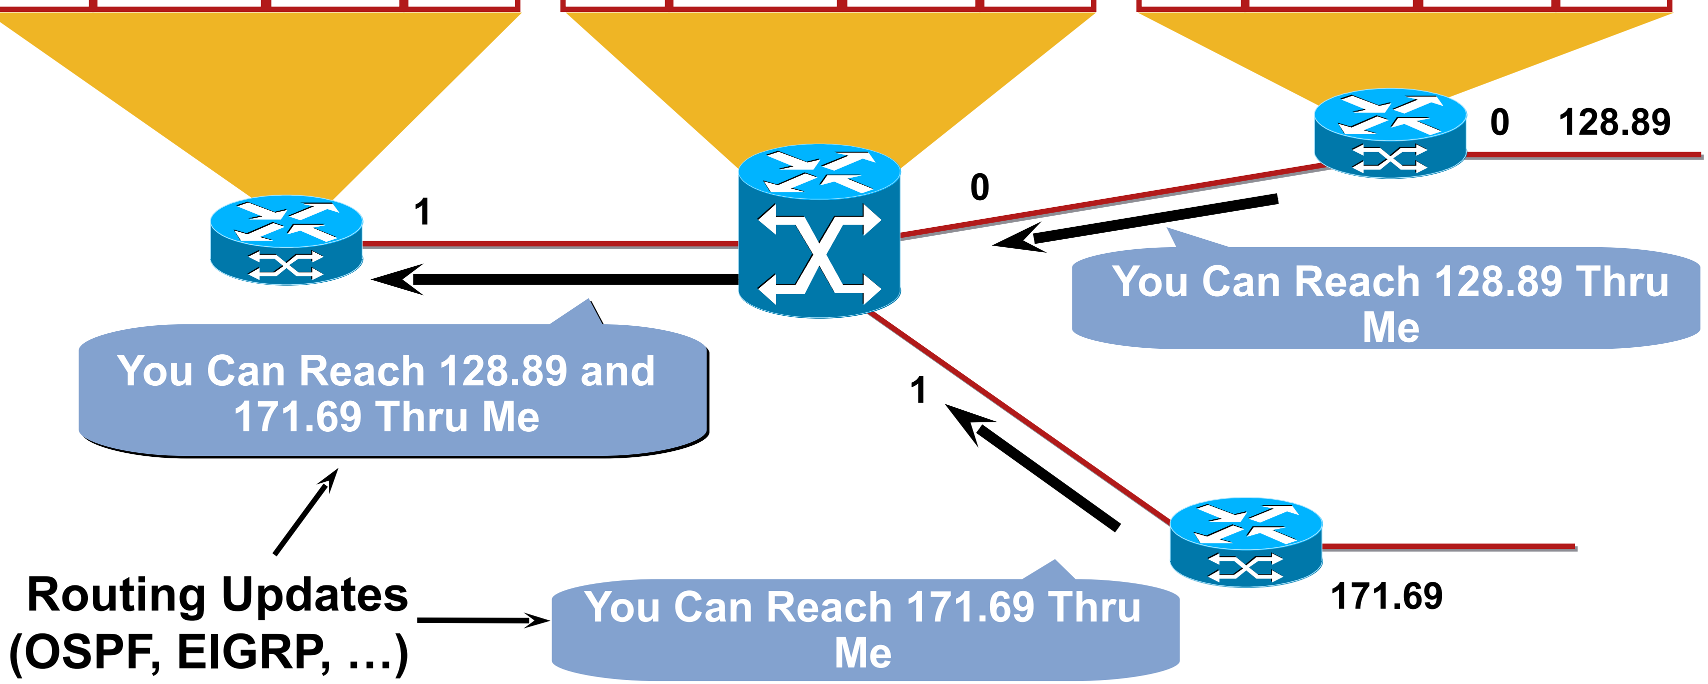

Label Switch Routers: Exchanging and Assigning Labels

In Label	Addres s Prefix	Out I'face	Out Label
-	128.89	1	4
-	171.69	1	5
...	...	...	...

In Label	Addres s Prefix	Out I'face	Out Label
4	128.89	0	9
5	171.69	1	7
...	...	...	...

In Label	Addres s Prefix	Out I'face	Out Label
9	128.89	0	-
...	...	...	...

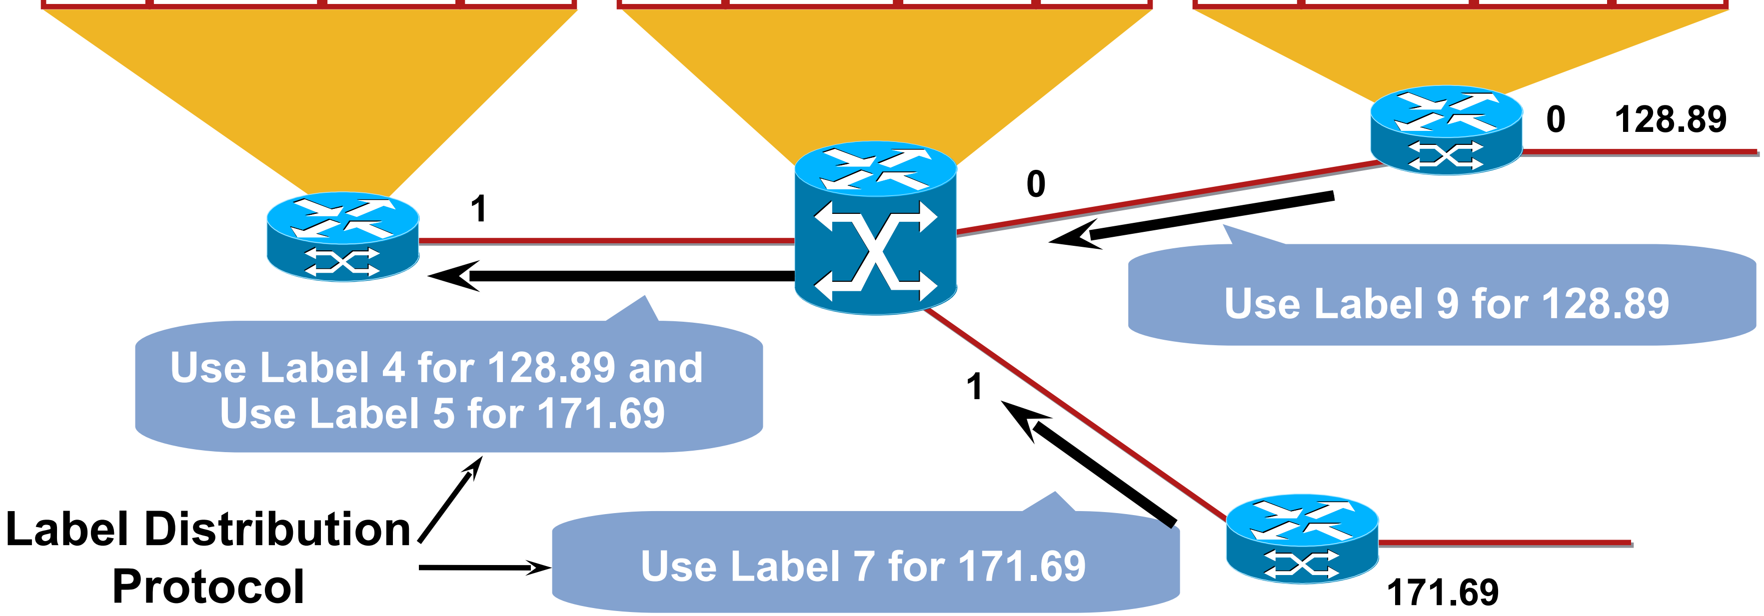

- In Label is the local label generated by the LSR
- Out Label is the remote label advertised by the adjacent LSR, which is the IGP next hop

Label Switch Routers: Forwarding Packets

In Label	Address Prefix	Out I'face	Out Label
-	128.89	1	4
-	171.69	1	5
...	...	...	...

In Label	Address Prefix	Out I'face	Out Label
4	128.89	0	9
5	171.69	1	7
...	...	...	...

In Label	Address Prefix	Out I'face	Out Label
9	128.89	0	-
...	...	...	...

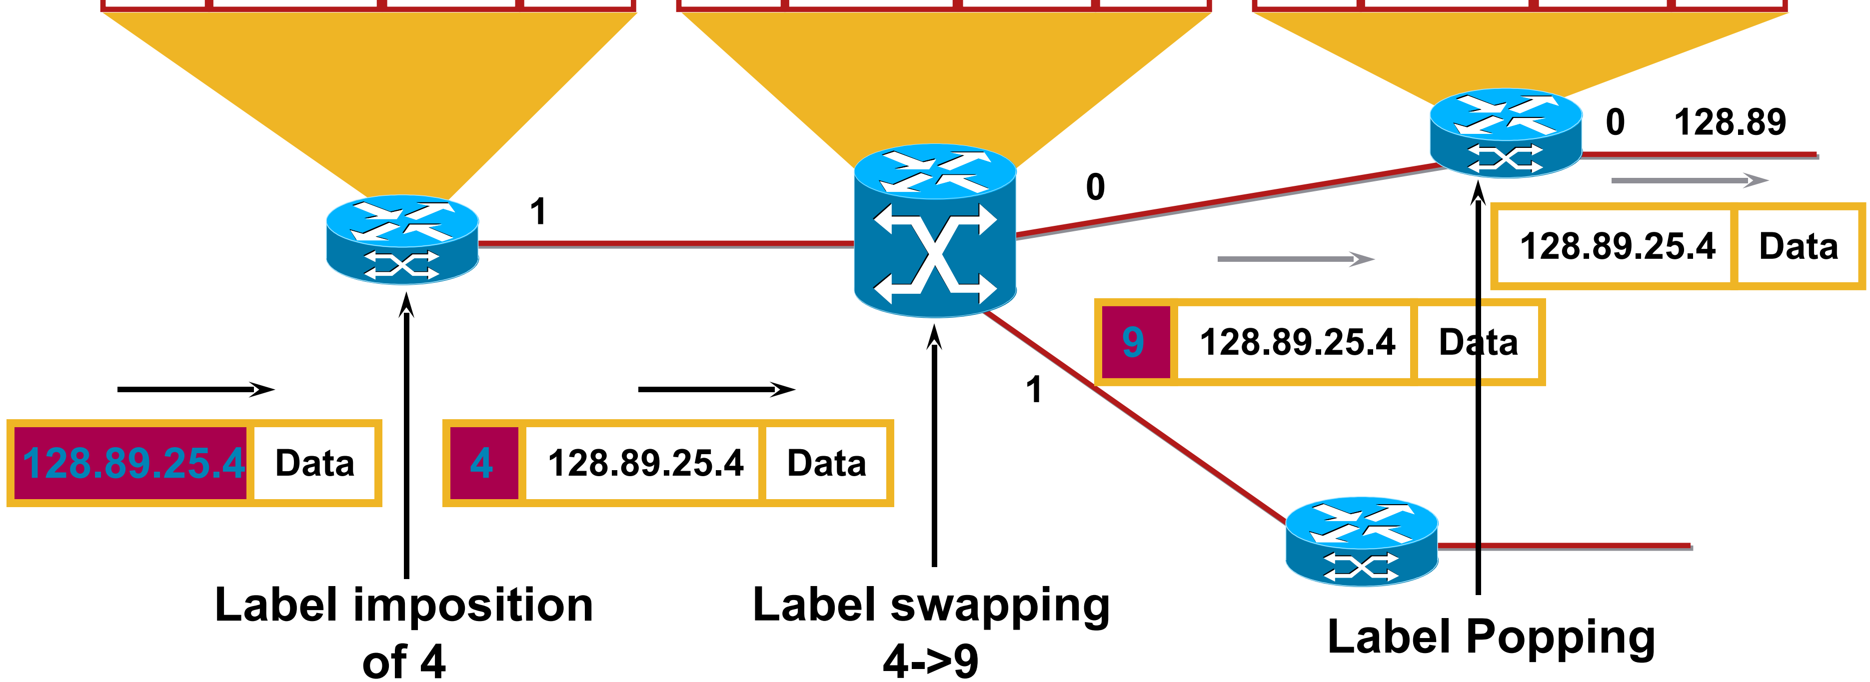

Label Switch Routers: Label functions

An LSR can perform the following functions:

- Insert (impose) a label or a stack of labels on ingress
- Swap a label with a next-hop label or a stack of labels in the core
- Remove (pop) a label on egress

ATM LSRs can swap a label with only one label (VPI/VCI fields change).

Label Switch Routers: Cell Mode

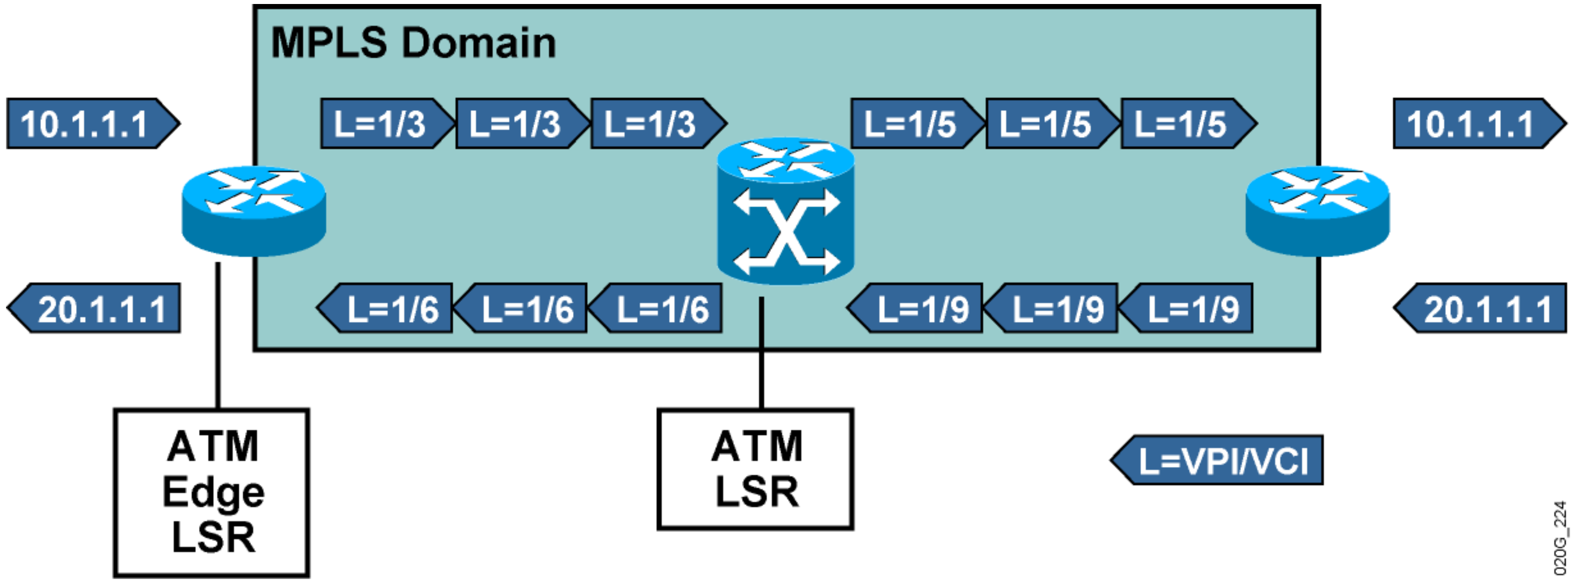

ATM LSR can forward only cells.

ATM edge LSR segments packets into cells and forwards them into an MPLS ATM domain, or reassembles cells into packets and forwards them out of an MPLS ATM domain.

MPLS Forwarding Plane: LFIB and Outgoing Labels

```
RSP-PE-SOUTH-5#show mpls forwarding-table 10.13.1.11
Local  Outgoing    Prefix            Bytes tag  Outgoing   Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
59     46          10.13.1.11/32     0          Se10/0/0   point2point
RSP-PE-SOUTH-5#
```

- Outgoing label tells what treatment the packet is going to get. It could also be -

  Pop - Pops the topmost label

  Untagged - Untag the incoming MPLS packet

  Aggregate - Untag and then do a FIB lookup

  0 - Nullify the top label (first 20bits)

- Label values 0-15 are reserved.

MPLS Forwarding Plane: Outgoing label types

```
PE1#sh mpls forwarding-table
Local   Outgoing     Prefix              Bytes tag  Outgoing    Next Hop
tag     tag or VC    or Tunnel Id        switched   interface
16      2002         10.13.1.22/32       0          Et0/0       10.13.1.5

        2002         10.13.1.22/32       0          Et1/0       10.13.1.9

17      2001         10.13.1.62/32       0          Et0/0       10.13.1.5

        2001         10.13.1.62/32       0          Et1/0       10.13.1.9

18      Pop tag      10.13.1.101/32      0          Et1/0       10.13.1.9

        Pop tag      10.13.1.101/32      0          Et0/0       10.13.1.5

19      Pop tag      10.13.2.4/30        0          Et1/0       10.13.1.9

        Pop tag      10.13.2.4/30        0          Et0/0       10.13.1.5

20      Untagged     5.5.5.5/32[V]       0          Se2/0       point2point

21      Pop tag      10.13.21.4/30       0          Et1/0       10.13.1.9

        Pop tag      10.13.21.4/30       0          Et0/0       10.13.1.5
```

V means it is a VPN prefix

MPLS Forwarding Plane: Outgoing label types (cont.)

- Untagged

  Convert the incoming MPLS packet to an IP packet and forward it.

- Pop

  Pop the top label from the label stack present in an incoming MPLS packet and forward it as an MPLS packet. If there was only one label in the stack, then forward it as an IP packet. SAME as imp-null label.

- Aggregate

  Convert the incoming MPLS packet to an IP packet and then do a FIB lookup for it to find out the outgoing interface.

- 0 (zero)

  Same as exp-null label. Simplify fills 0 in the first 20 bits of label; helps to preserve the EXP value of the top label.

MPLS and BGP

- Labels are assigned to FECs which are derived from IP routing protocols (IGP)
- Labels are NOT assigned to BGP routes
- BGP routes use recursive routing to find next-hop reachability
- Labels are assigned to BGP next-hops
- This saves CPU/Memory, label space and stability on core LSRs

  Core LSRs are preserved from BGP instability
- We can assign labels to BGP learnt routes based on RFC 3107

MPLS and BGP (cont.)

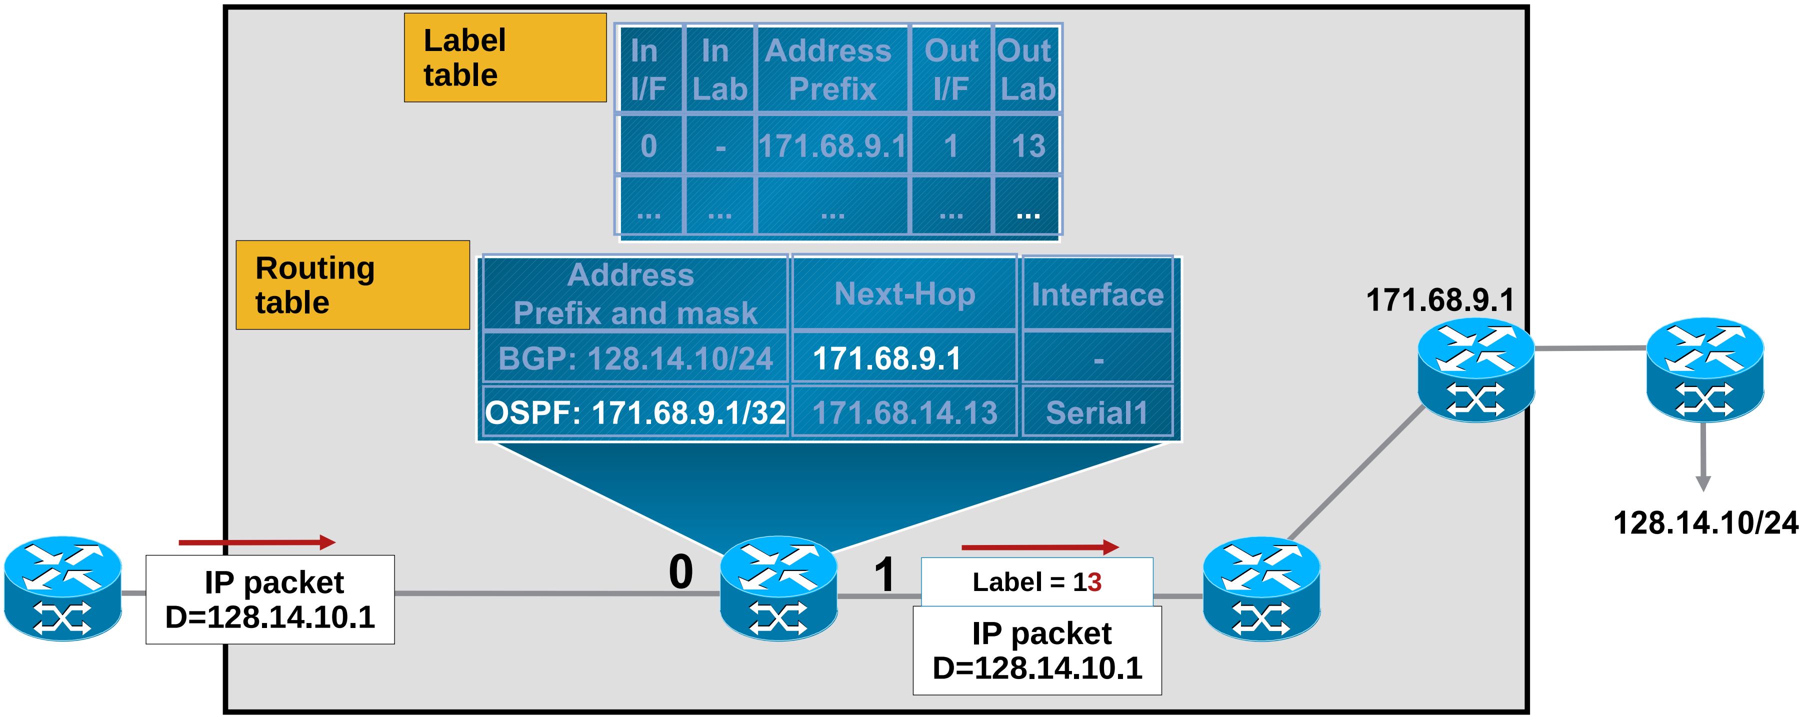

Label table

In I/F	In Lab	Address Prefix	Out I/F	Out Lab
0	-	171.68.9.1	1	13
...	...	...	...	...

Routing table

Address Prefix and mask	Next-Hop	Interface
BGP: 128.14.10/24	171.68.9.1	-
OSPF: 171.68.9.1/32	171.68.14.13	Serial1

Ingress LSR receives IP packet
Destination is given by BGP
BGP has next-hop known in the IGP
Label is available for BGP next-hop, through IGP route
Packet will traverse the core using IGP (BGP next-hop) label

Label Distribution Protocols

Several protocols for label exchange

- **TDP/LDP**
- **RSVP**
- **BGP**

MPLS Forwarding

An LSR can perform the following functions:

- Insert (impose) a label or a stack of labels on ingress.
- Swap a label with a next-hop label or a stack of labels in the core.
- Remove (pop) a label on egress.

MPLS Forwarding

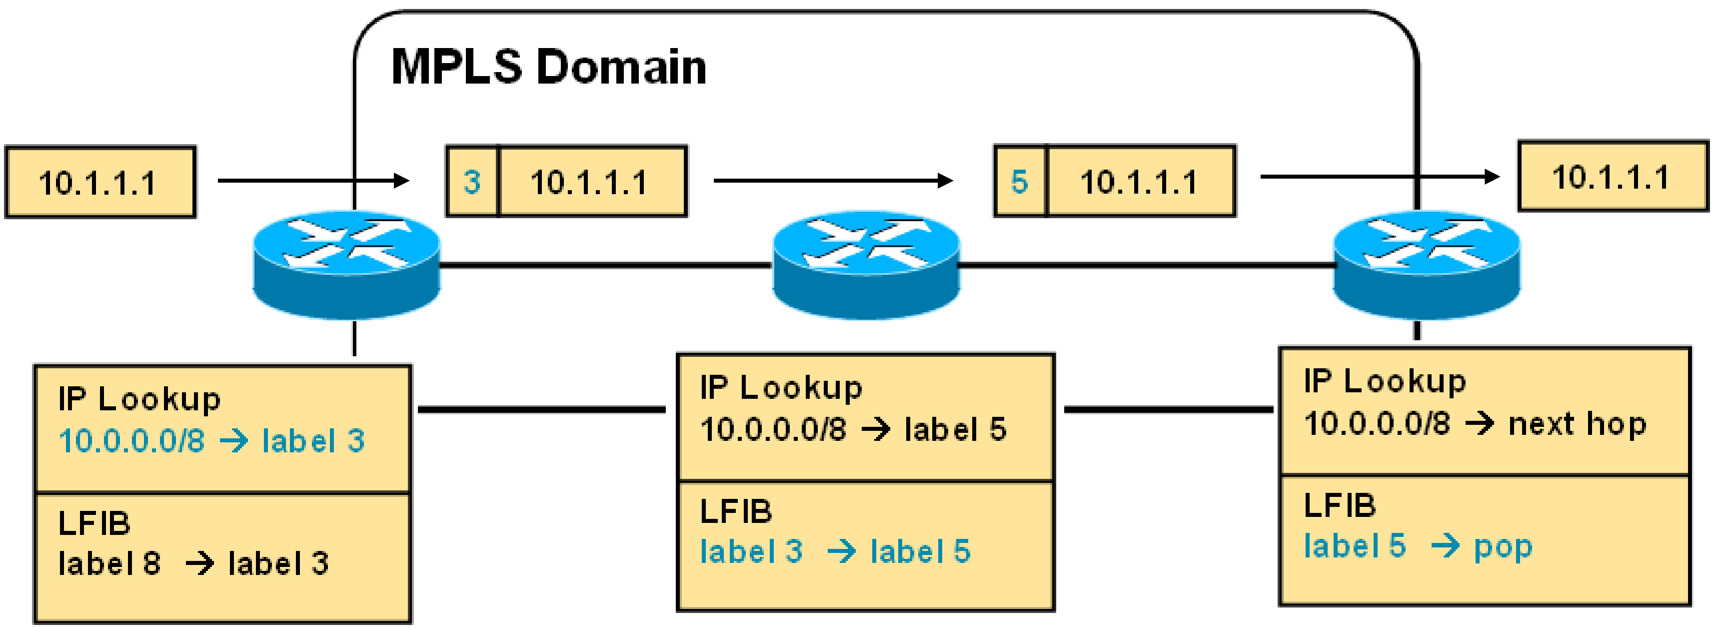

- On ingress a label is assigned and imposed by the IP routing process.
- LSRs in the core swap labels based on the contents of the label forwarding table.
- On egress the label is removed and a routing lookup is used to forward the packet.

MPLS Applications

MPLS is already used in many different applications:

- Unicast IP routing
- Multicast IP routing
- Traffic Engineering (MPLS TE)
- QoS
- Virtual private networks (MPLS VPN)

Regardless of the application, the functionality is always split into the control plane and the data plane:

- The applications differ only in the control plane.
- They all use a common label-switching data plane.
- Edge LSR Layer 3 data planes may differ.
- In general, a label is assigned to a forwarding equivalence class (FEC).

Unicast IP Routing

- Two mechanisms are needed on the control plane:
  - IP routing protocol (OSPF, IS-IS, EIGRP, ...)
  - Label distribution protocol (LDP or TDP)
- A routing protocol carries the information about the reachability of networks.
- The label distribution protocol binds labels to networks learned via a routing protocol.
- The forwarding equivalence class (FEC) is equal to a destination network, stored in the IP routing table.

MPLS TE

- MPLS traffic engineering requires OSPF or IS-IS with extensions for MPLS TE as the IGP.
- OSPF and IS-IS with extensions hold the entire topology in their databases.
- OSPF and IS-IS should also have some additional information about network resources and constraints.
- RSVP or CR-LDP is used to establish traffic engineering tunnels (TE tunnels) and propagate labels.

Virtual Private Networks

- Networks are learned via an IGP (OSPF, EBGP, RIP version 2 [RIPv2] or static) from a customer or via BGP from other internal routers.
- Labels are propagated via MP-BGP.
- Two labels are used:
  - Top label points to the egress router (assigned through LDP or TDP).
  - Second label identifies the outgoing interface on the egress router or a routing table where a routing lookup is performed.
- FEC is equal to a VPN site descriptor or VPN routing table.

Label Assignment and Distribution

MPLS Unicast IP Routing Architecture

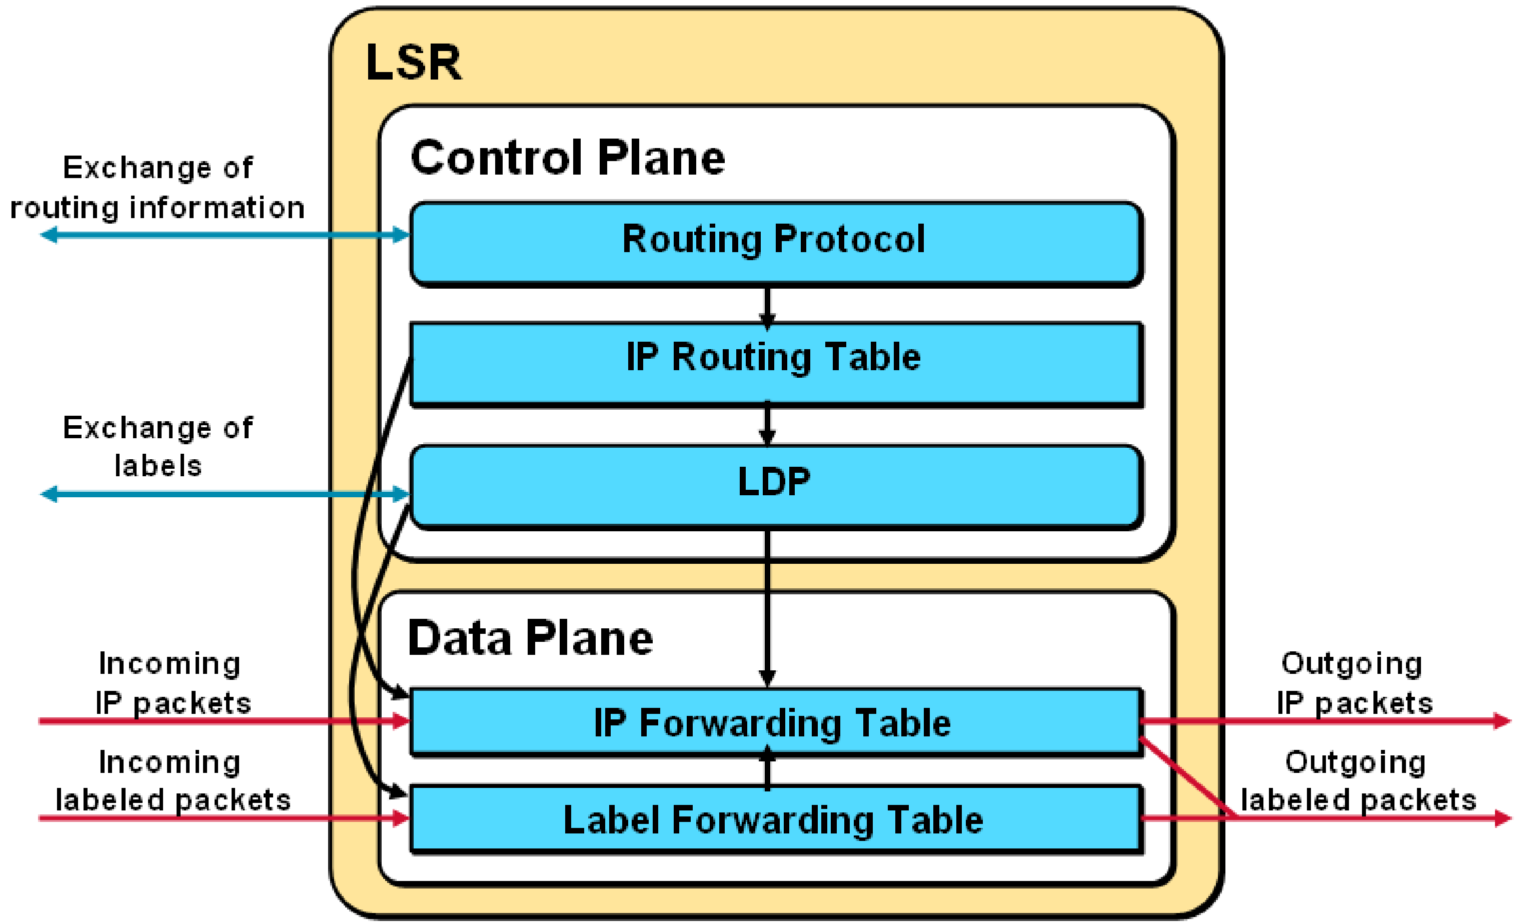

MPLS Unicast IP Routing Example

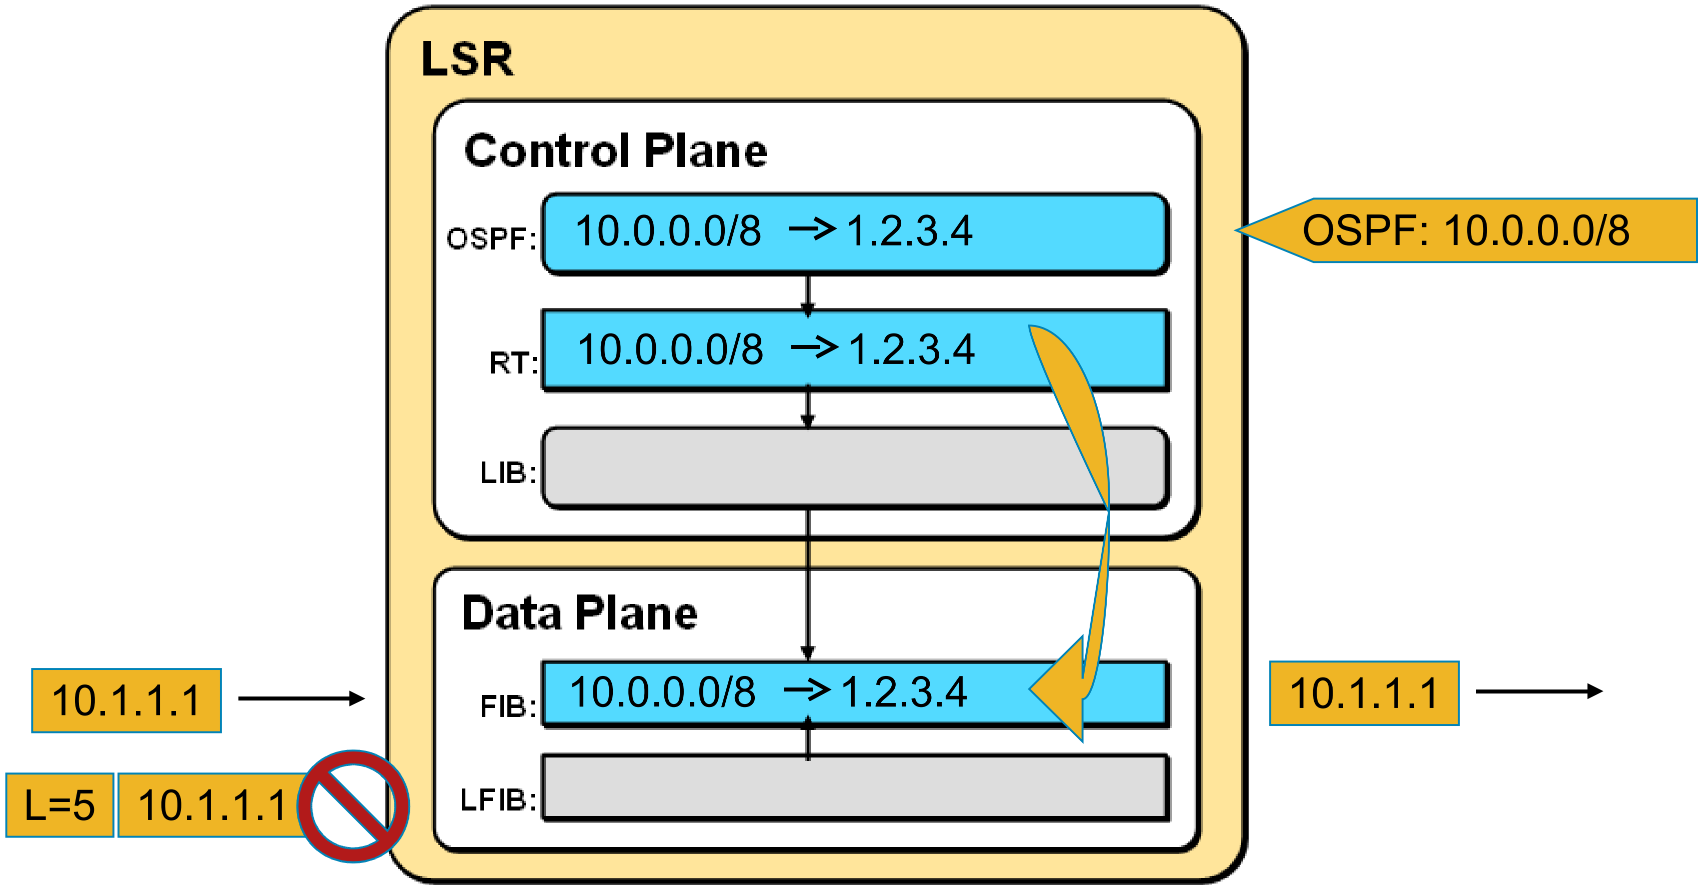

MPLS Unicast IP Routing Example

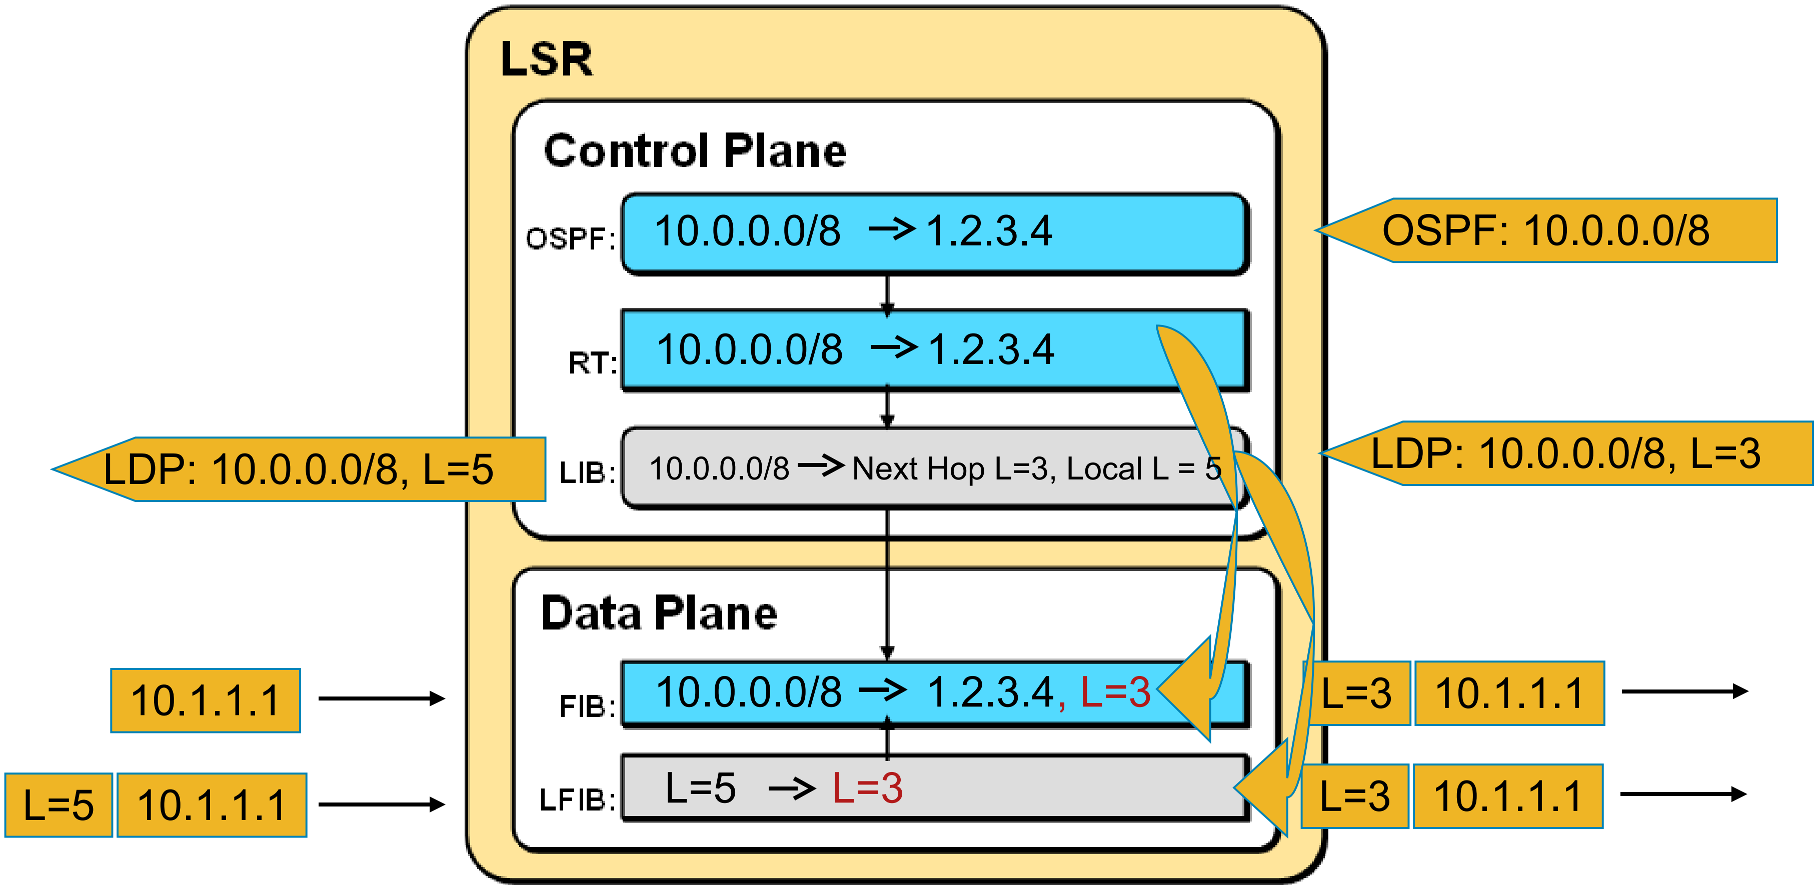

Building the IP Routing Table

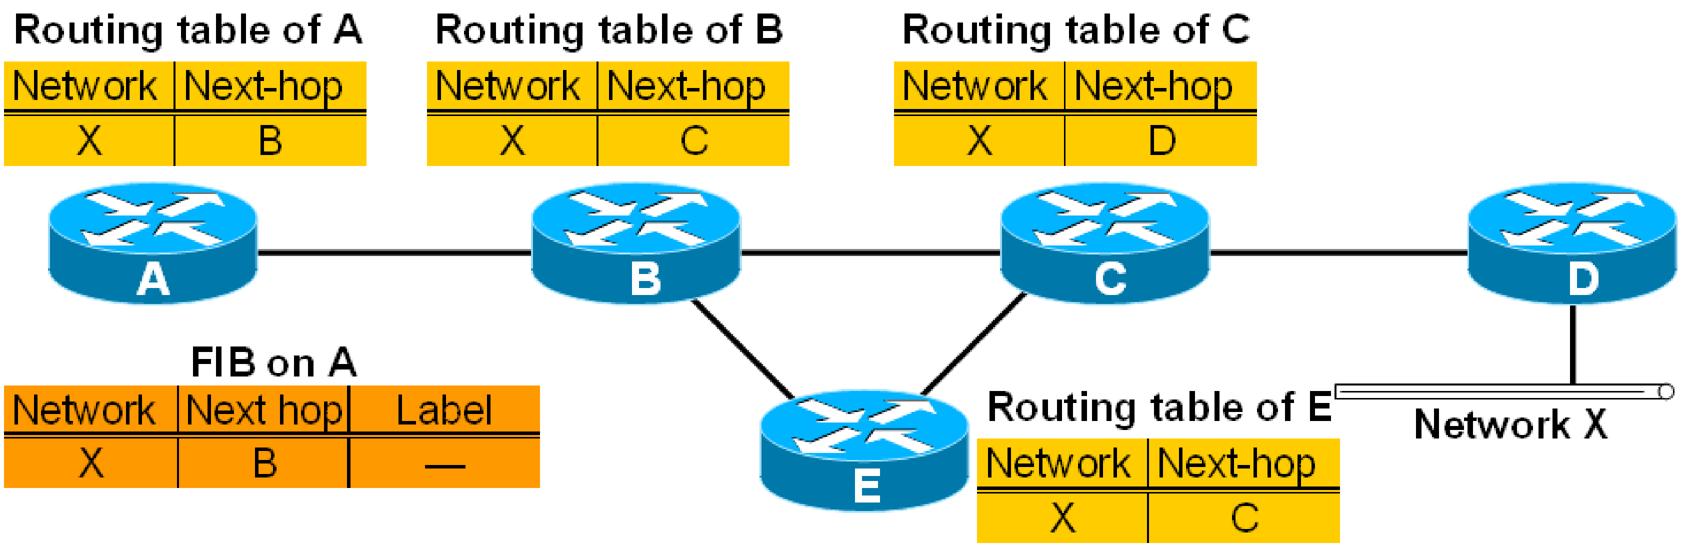

- IP routing protocols are used to build IP routing tables on all LSRs.
- FIBs are built based on IP routing tables with no labeling information.

Allocating Labels

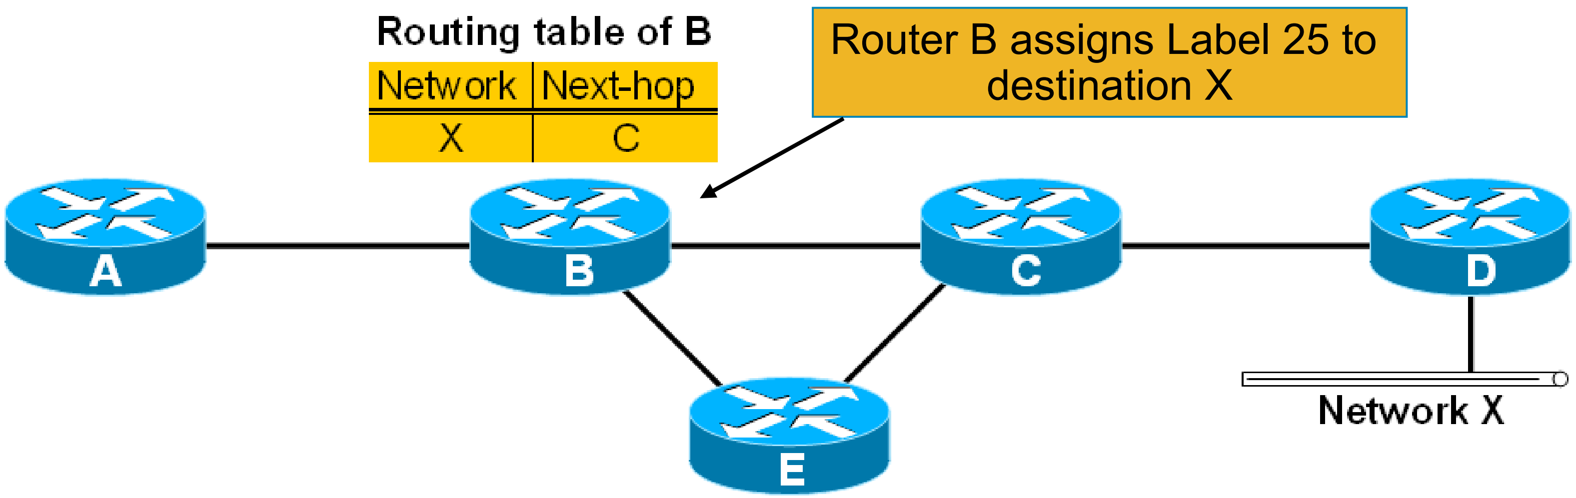

- Every LSR allocates a label for every destination in the IP routing table.
- Labels have local significance.
- Label allocations are asynchronous.

LIB and LFIB Setup

Routing table of B

Network	Next-hop
X	C

Router B assigns Label 25 to destination X

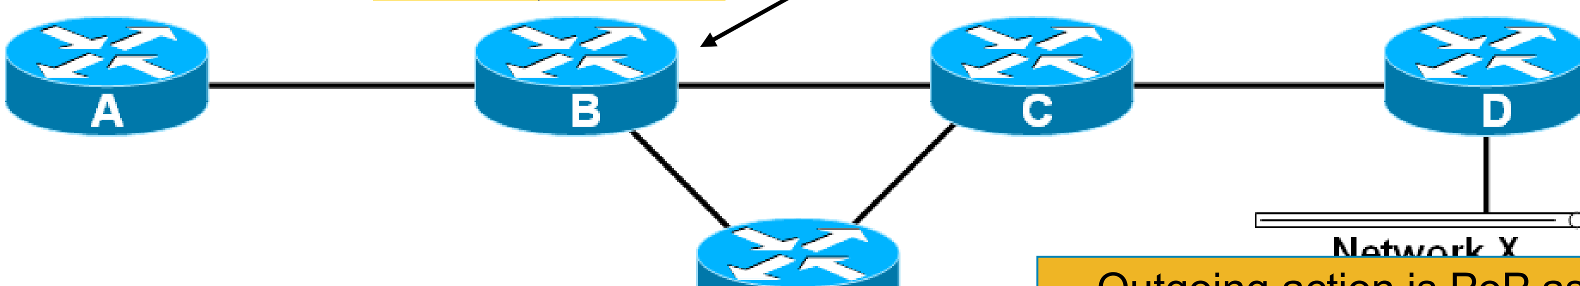

LFIB on B

Label	Action	Next hop
25	pop	C

Outgoing action is PoP as Router B has not received any label for X from C

LIB on B

Network	LSR	label
X	local	25

Local Label is Stored in the LIB

Label Distribution

LIB on B

Network	LSR	label
X	local	25

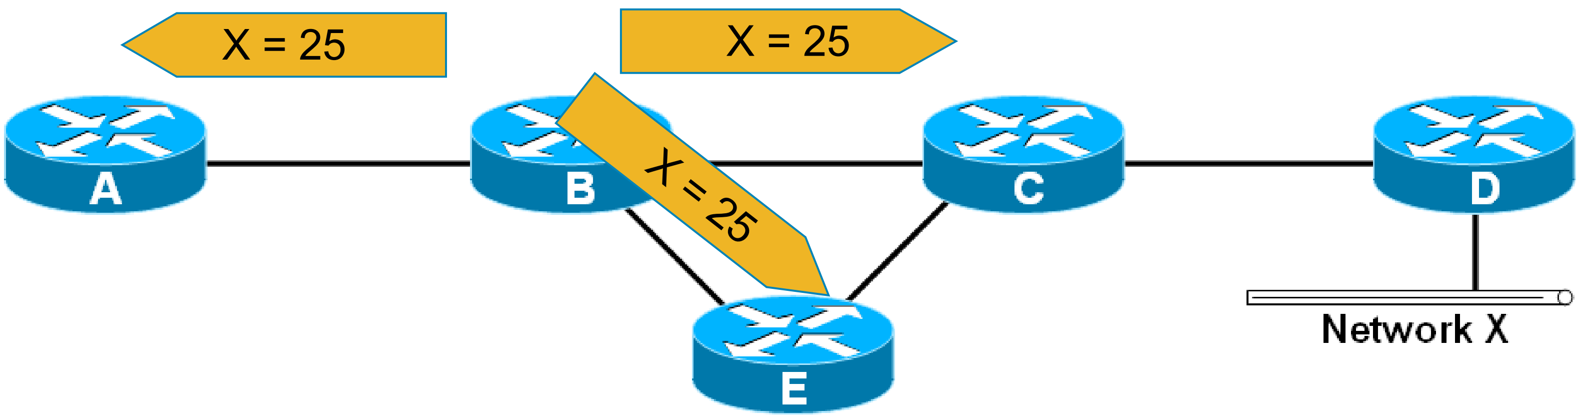

The allocated label is advertised to all neighbor LSRs, regardless of whether the neighbors are upstream or downstream LSRs for the destination.

Receiving Label Advertisement

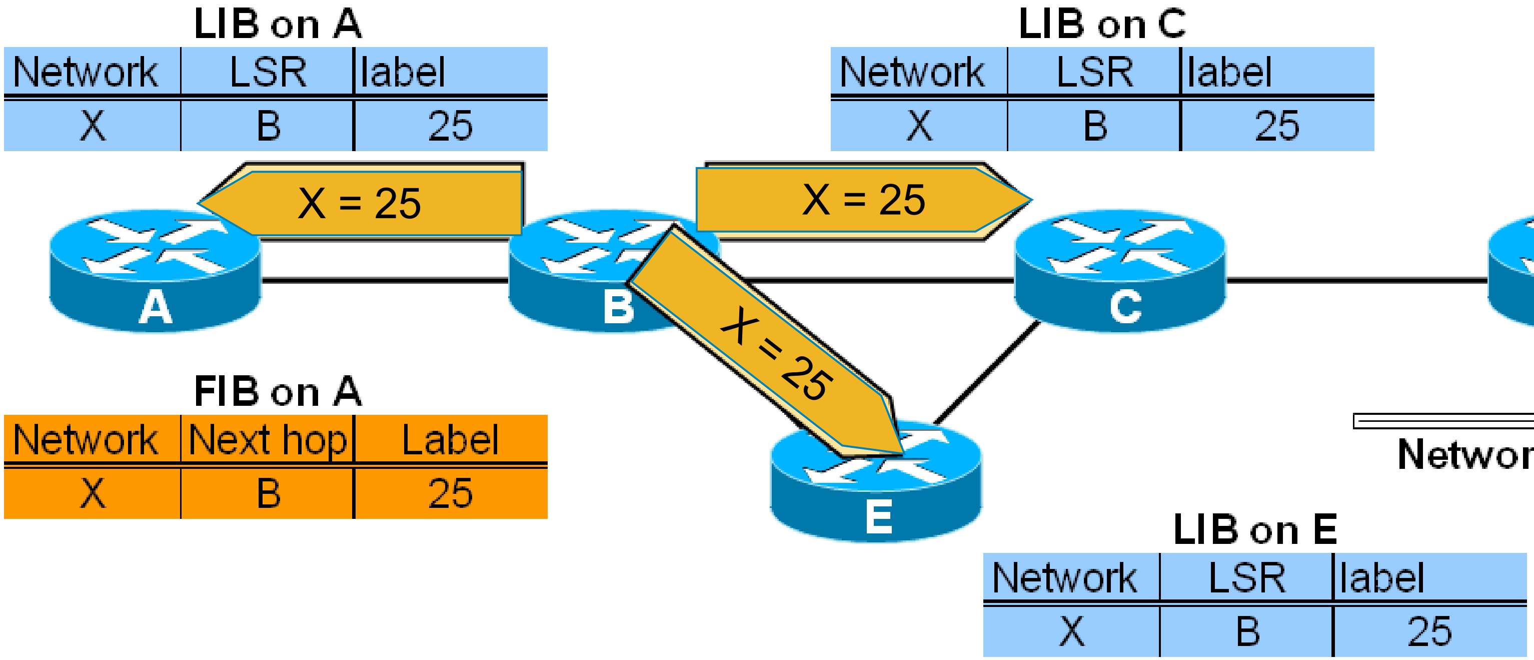

- Every LSR stores the received label in its LIB.
- Edge LSRs that receive the label from their next-hop also store the label information in the FIB.

Populating LFIB

FIB on B

Network	Next hop	Label
X	C	47

LIB on B

Network	LSR	label
X	local	25
	C	47

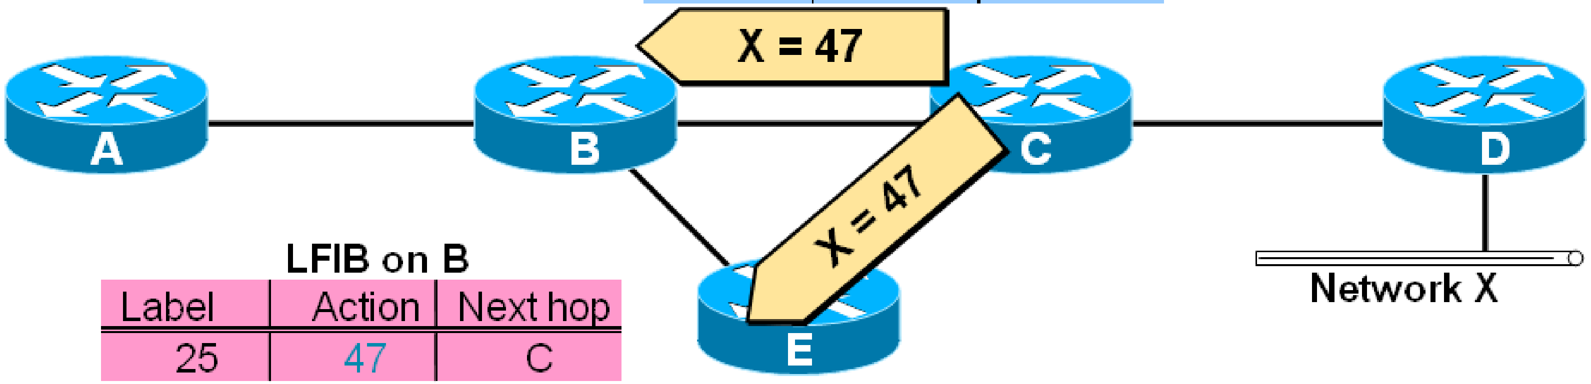

LFIB on B

Label	Action	Next hop
25	47	C

- **Router B has already assigned a label to X and created an entry in the LFIB.**
- **The outgoing label is inserted in the LFIB after the label is received from the next-hop LSR.**

Packet Propogation through MPLS Network

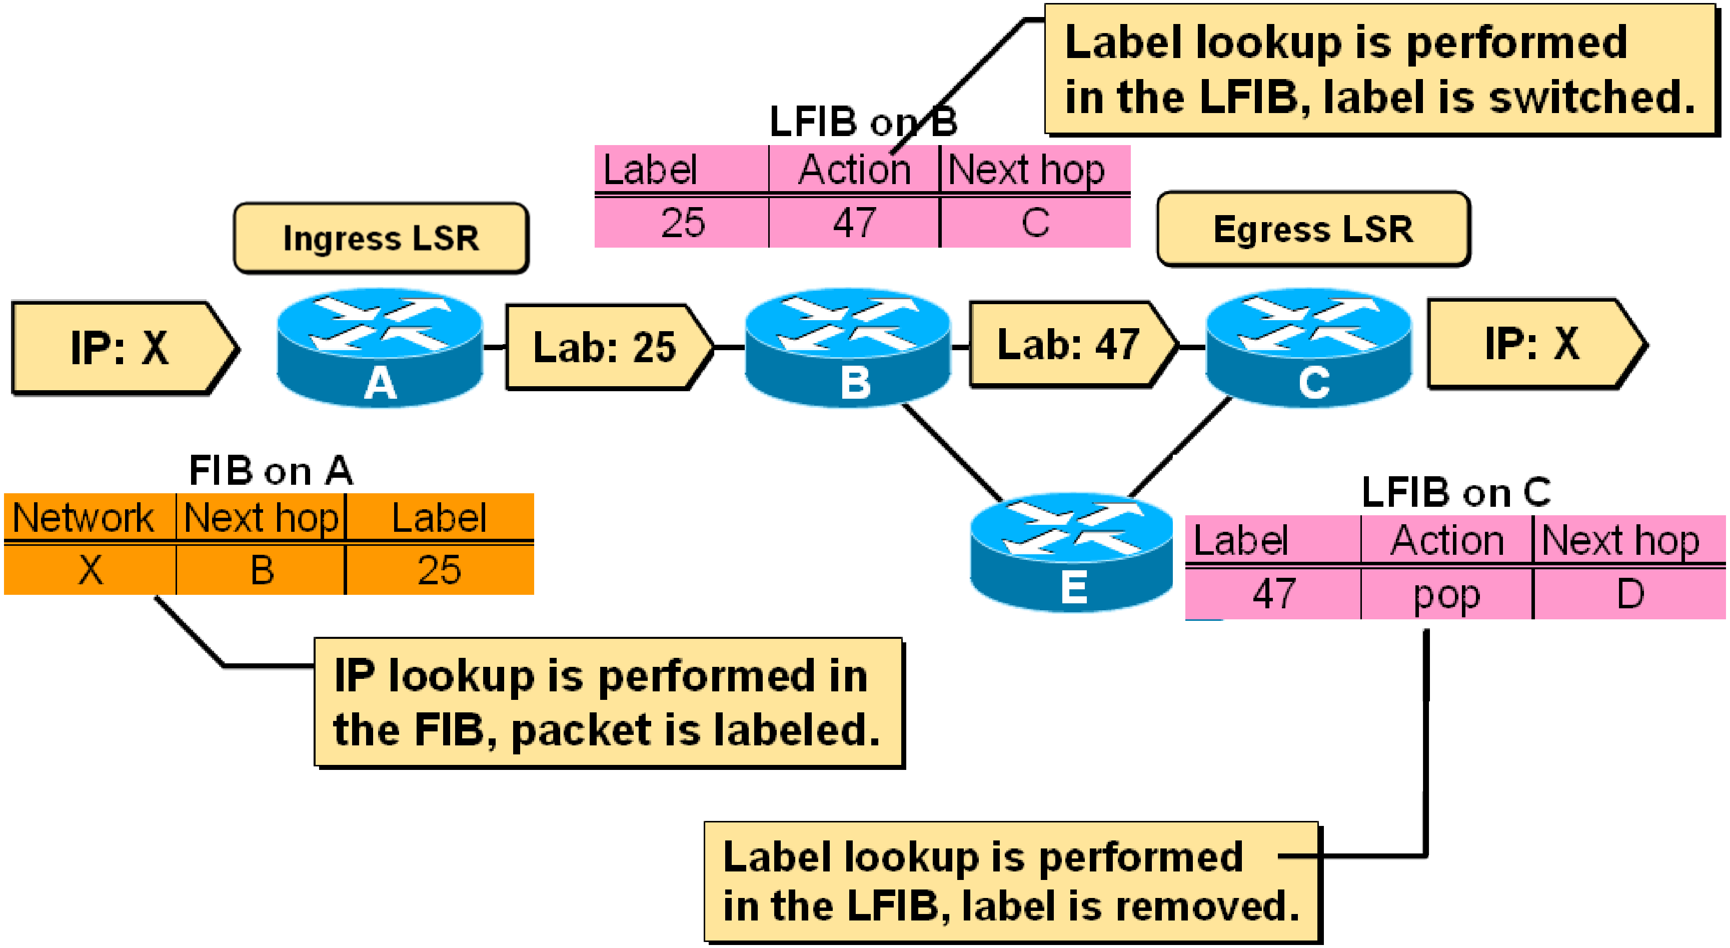

Convergence: Steady State

Routing table of B

Network	Next-hop
X	C

FIB on B

Network	Next hop	Label
X	C	47

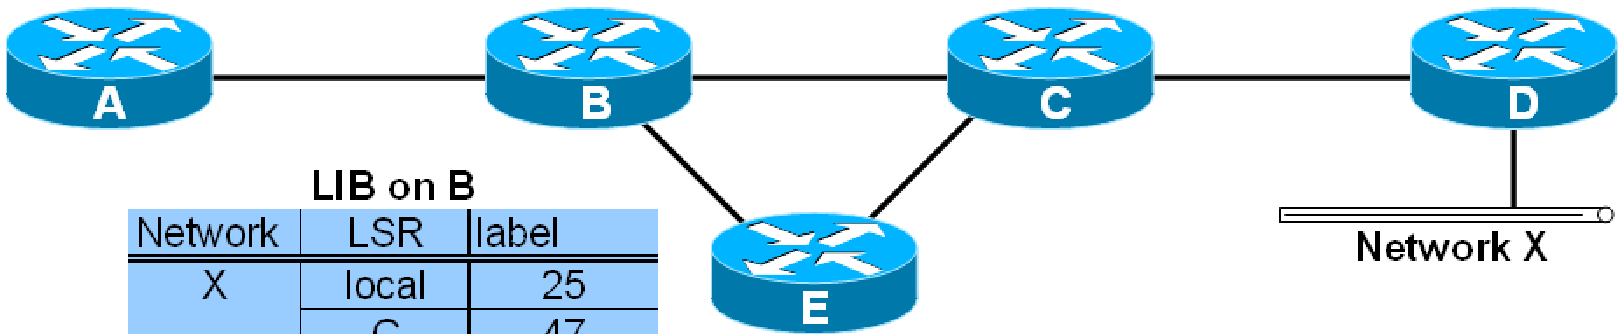

LIB on B

Network	LSR	label
X	local	25
	C	47
	E	75

LFIB on B

Label	Action	Next hop
25	47	C

- **After the LSRs have exchanged the labels, LIB, LFIB and FIB data structures are completely populated.**

Link Failure Action

Routing table of B

Network	Next-hop
~~X~~	~~C~~

FIB on B

Network	Next hop	Label
X	~~C~~	~~47~~

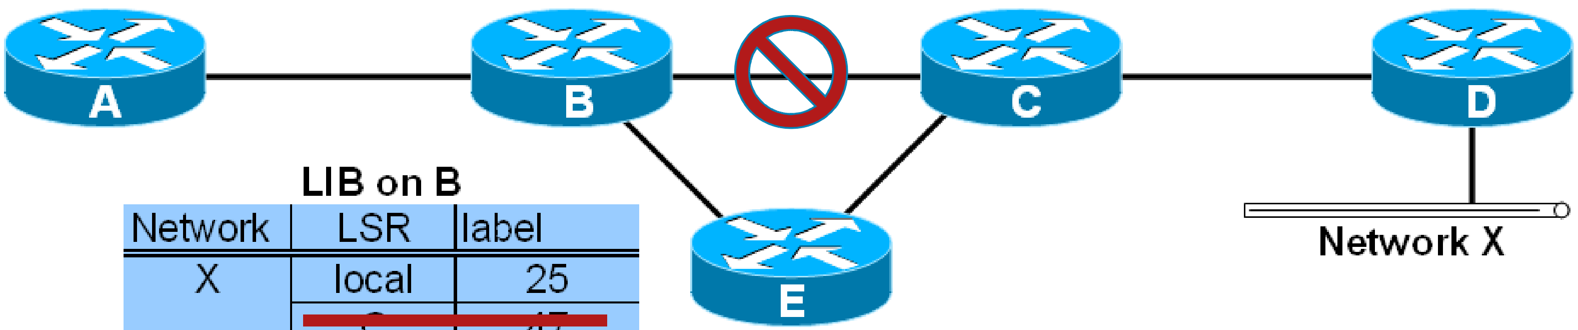

LIB on B

Network	LSR	label
X	local	25
	~~C~~	~~47~~
	E	75

LFIB on B

Label	Action	Next hop
25	~~47~~	~~C~~

- **Routing protocol neighbors and LDP neighbors are lost after a link failure.**
- **Entries are removed from various data structures.**

Routing Protocol Convergence

Routing table of B

Network	Next-hop
X	E

FIB on B

Network	Next hop	Label
X	E	—

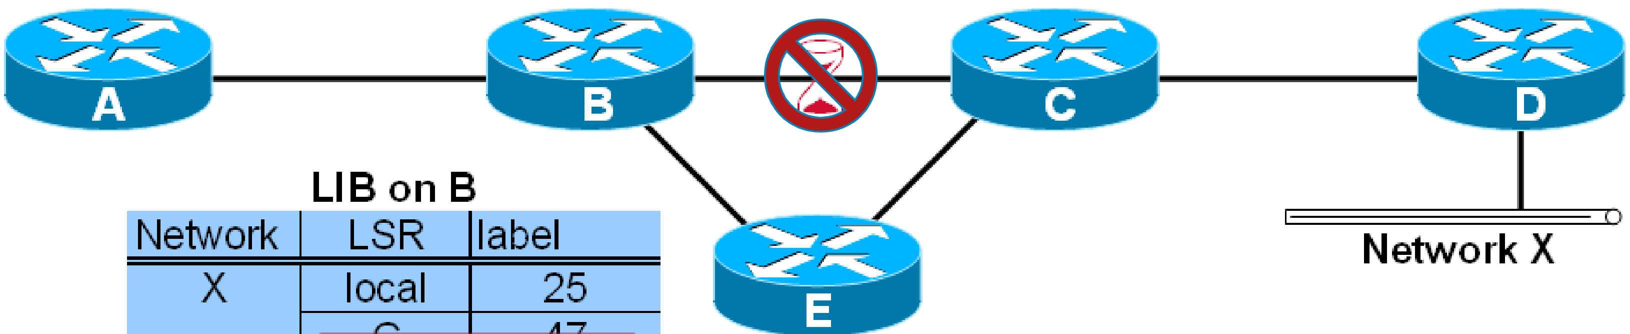

LIB on B

Network	LSR	label
X	local	25
	~~C~~	~~47~~
	E	75

LFIB on B

Label	Action	Next hop
25	~~47~~	~~C~~

Routing protocols rebuild the IP routing table and the IP forwarding table.

MPLS Convergence

Routing table of B

Network	Next-hop
X	E

FIB on B

Network	Next hop	Label
X	E	75

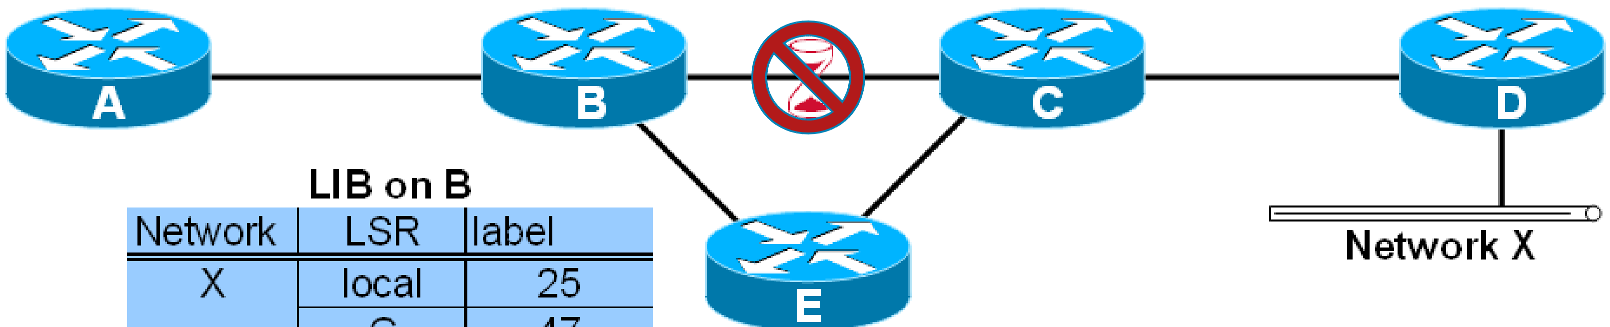

LIB on B

Network	LSR	label
X	local	25
	~~C~~	~~47~~
	E	75

LFIB on B

Label	Action	Next hop
25	75	E

The LFIB and labeling information in the FIB are rebuilt immediately after the routing protocol convergence, based on labels stored in the LIB.

MPLS Convergence After Link Failure

- MPLS convergence in packet-mode MPLS does not affect the overall convergence time.
- MPLS convergence occurs immediately after the routing protocol convergence, based on labels already stored in the LIB.

Link Recovery Actions

Routing table of B

Network	Next-hop
X	E

FIB on B

Network	Next hop	Label
X	E	75

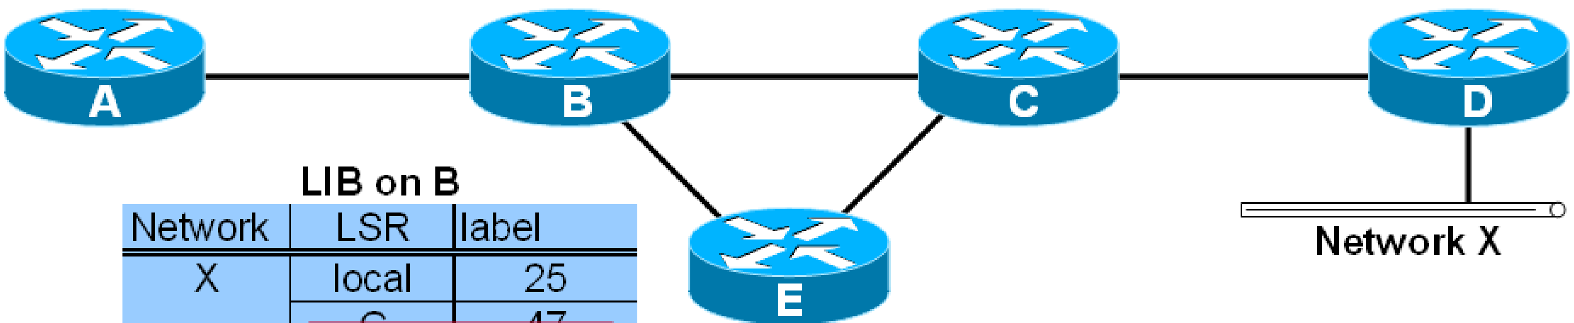

LIB on B

Network	LSR	label
X	local	25
	~~C~~	~~47~~
	E	75

LFIB on B

Label	Action	Next hop
25	75	E

- **Routing protocol neighbors are discovered after link recovery.**

Routing Convergence after Link Recovery

Routing table of B

Network	Next-hop
X	C

FIB on B

Network	Next hop	Label
X	C	-

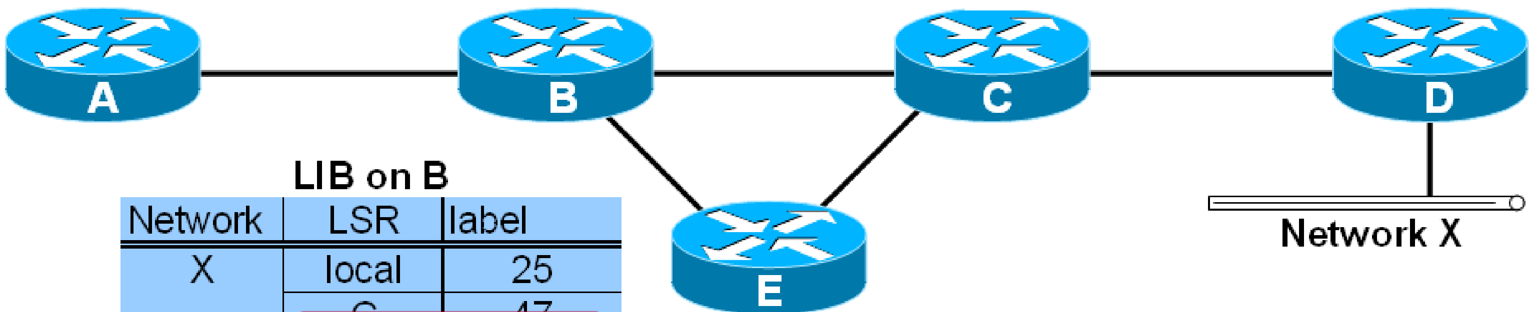

LIB on B

Network	LSR	label
X	local	25
	~~C~~	~~47~~
	E	75

LFIB on B

Label	Action	Next hop
25	pop	C

- **IP routing protocols rebuild the IP routing table.**
- **The FIB and the LFIB are also rebuilt, but the label information might be lacking.**

MPLS Convergence after Link Recovery

- Routing protocol convergence optimizes the forwarding path after a link recovery.
- The LIB might not contain the label from the new next hop by the time the IP convergence is complete.
- End-to-end MPLS connectivity might be intermittently broken after link recovery.
- Use MPLS traffic engineering for make-before-break recovery.

Label Distribution Protocol

LDP Concepts

- Label Distribution Protocol
- LDP works between adjacent/non-adjacent peers
- LDP sessions are established between peers
- LDP messages sent in the form of TLVs

  <Type, Length, Value>
- Standardized via RFC 3036

TDP/LDP Transport

- Uses TCP for reliable transport
- Well-known TCP port
  - LDP (port 646)
  - TDP (port 711)
- LSR with higher LDP router-id opens a connection to port 646 of other LSR
- Design Choice:
  - One TDP/LDP session per TCP connection

LDP Identifier

- Identifies tag space
- 6 bytes (4 bytes =>IP address, 2 bytes =>Label space ID)

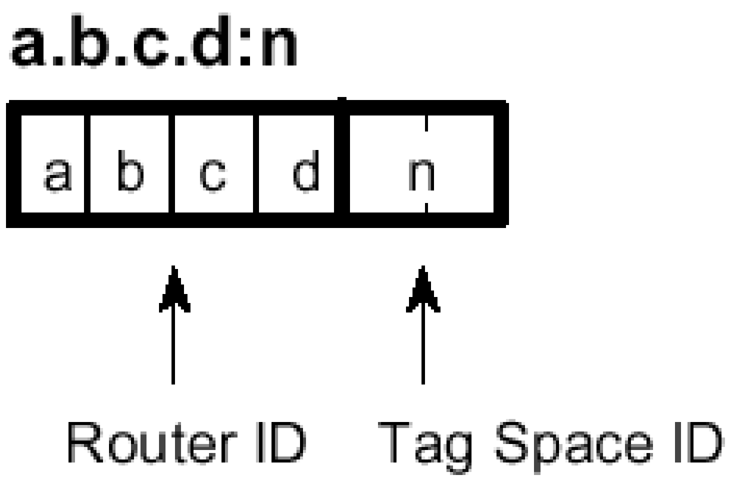

LDP Identifier: Label Space

- LSRs establish one LDP session per label space.
  - Per-platform label space requires only one LDP session, even if there are multiple parallel links between a pair of LSRs.
- Per-platform label space is announced by setting the label space ID to 0, for example:
  - LDP ID = 1.0.0.1:0
- A combination of frame-mode and cell-mode MPLS, or multiple cell-mode links, results in multiple LDP sessions.

Label Space and number of LDP sessions

Frame Mode

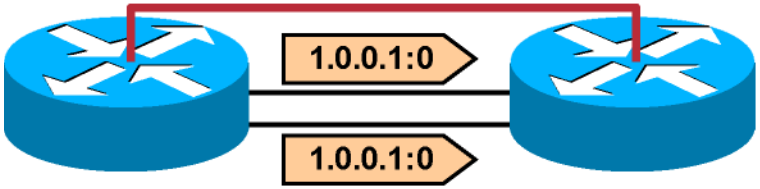

Frame Mode

Cell Mode

Mixed Mode

020G_318

Label Space and number of LDP sessions (Cont.)

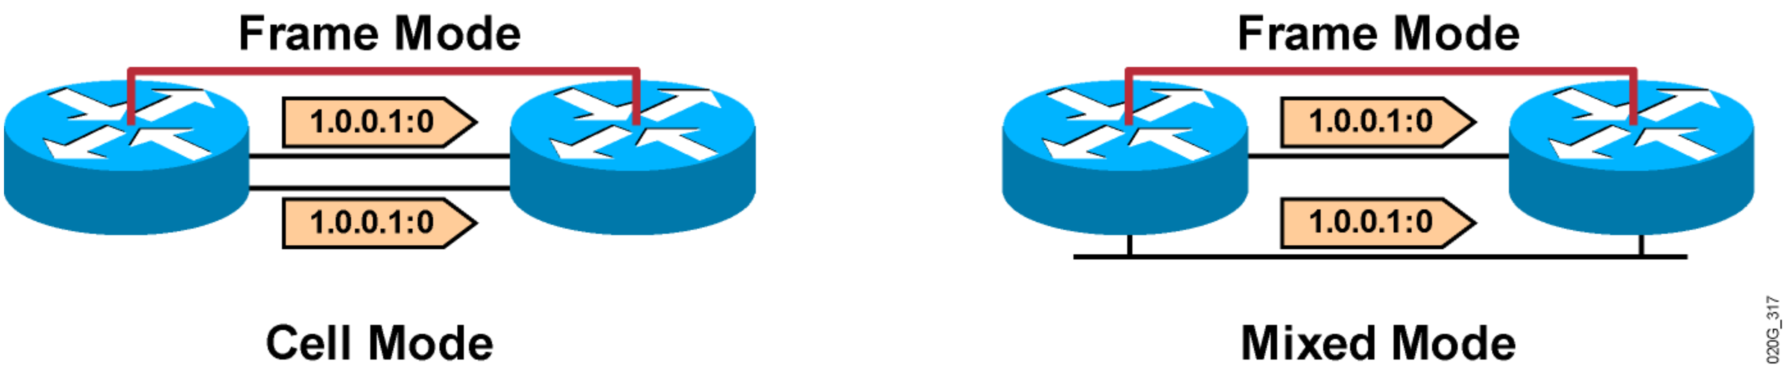

Label Space and number of LDP sessions (Cont.)

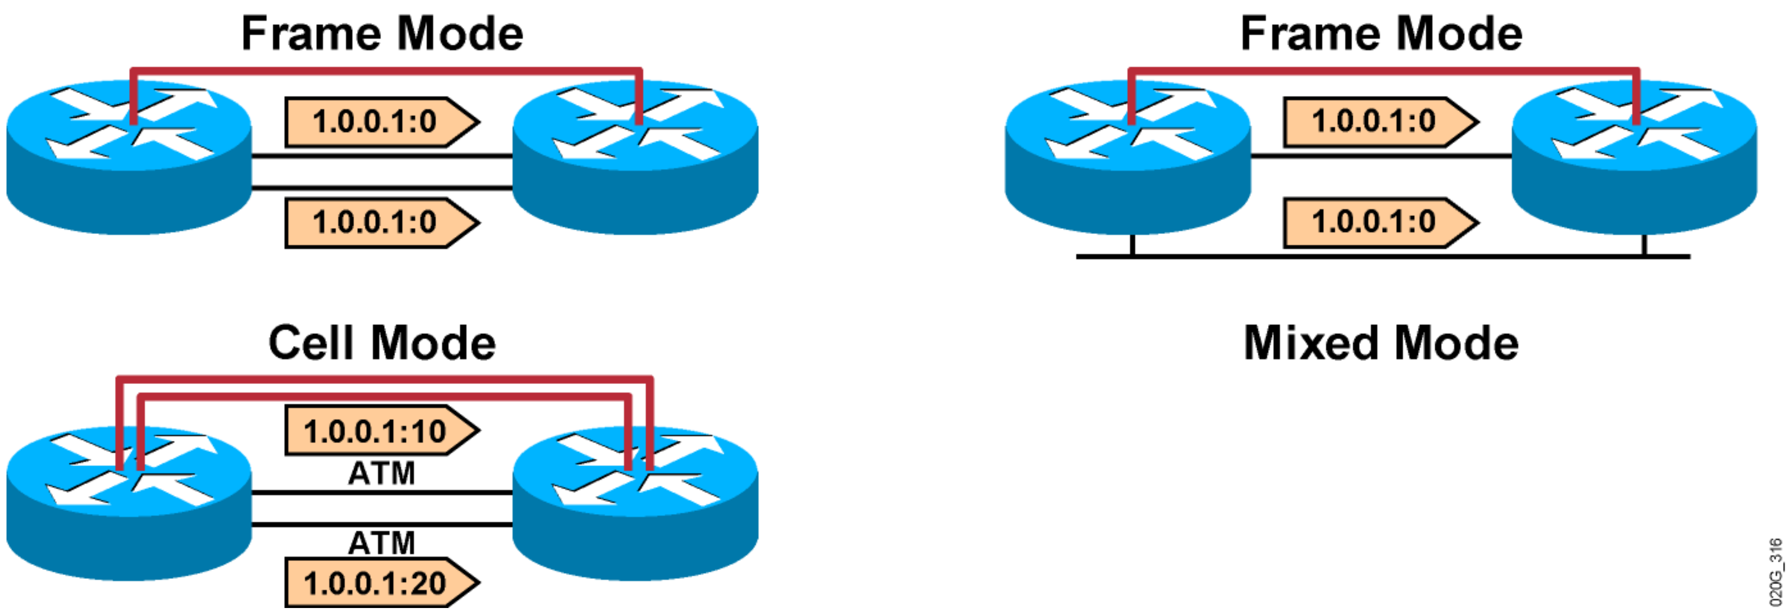

Label Space and number of LDP sessions (Cont.)

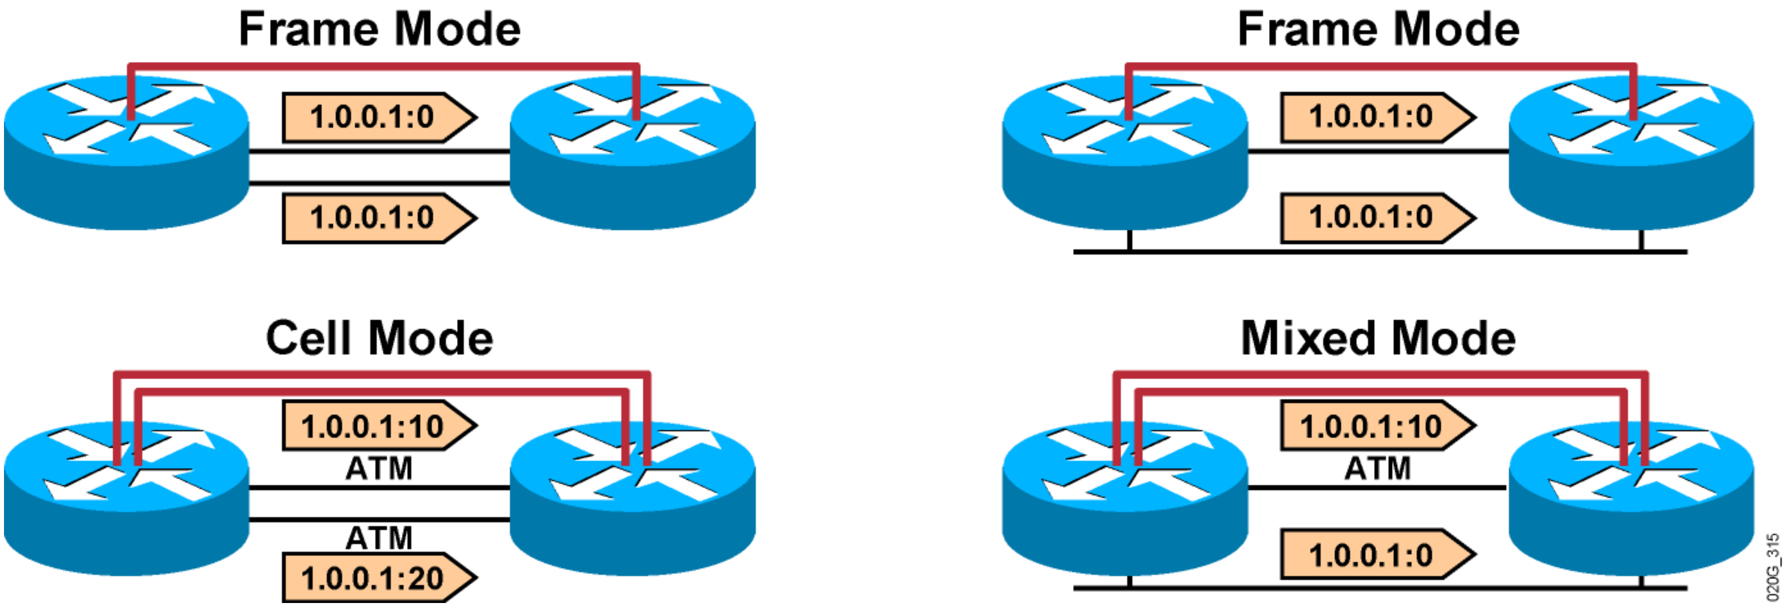

- One LDP session is established for each announced LDP identifier (router ID + label space).
- The number of LDP sessions is determined by the number of different label spaces.
- The bottom right example is not common, because ATM LSRs do not use Ethernet for packet forwarding, and frame-mode MPLS across ATM uses per-platform label space.

LDP Protocol Data Units

- **All LDP information is sent in the form of PDUs over the TCP connection**

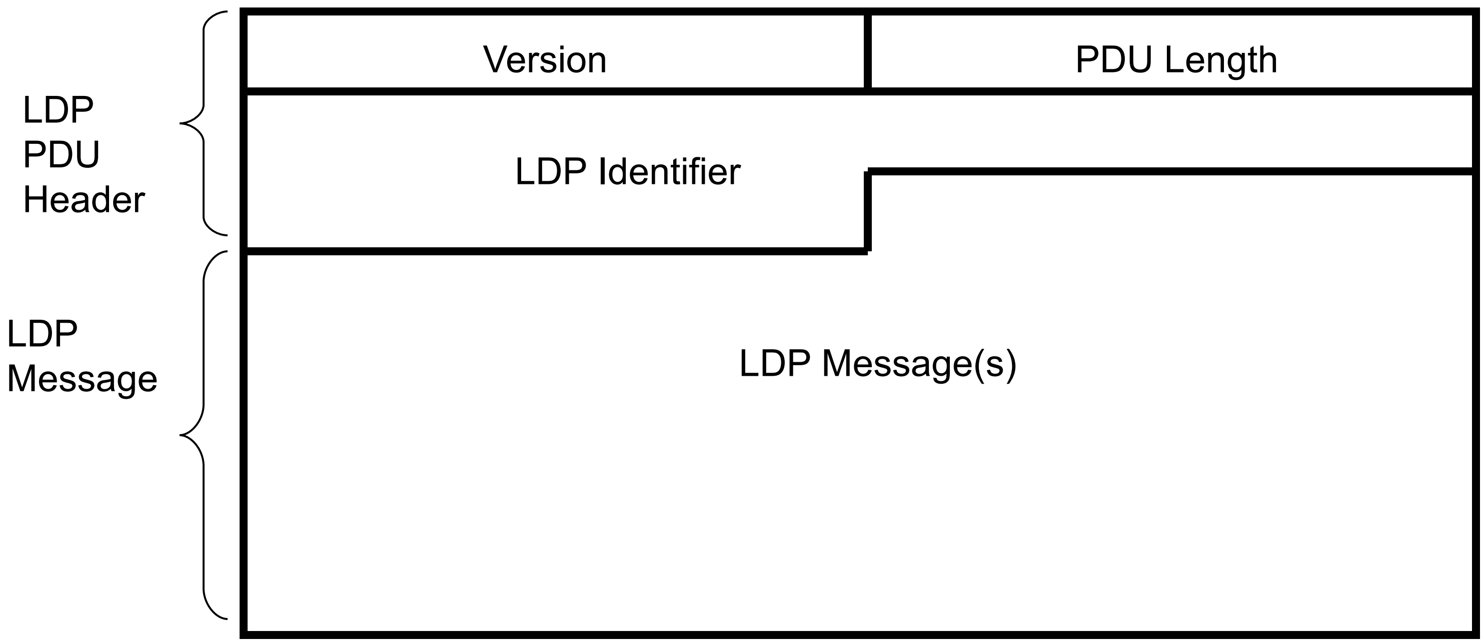

•**Version** => LDP version. Current LDP version is 1
•**PDU Length** (excludes Version and PDU Length fields) => total length of PDU in bytes.
•**LDP Identifier** => discussed earlier
•**LDP Messages** => one or more LDP messages

LDP Message

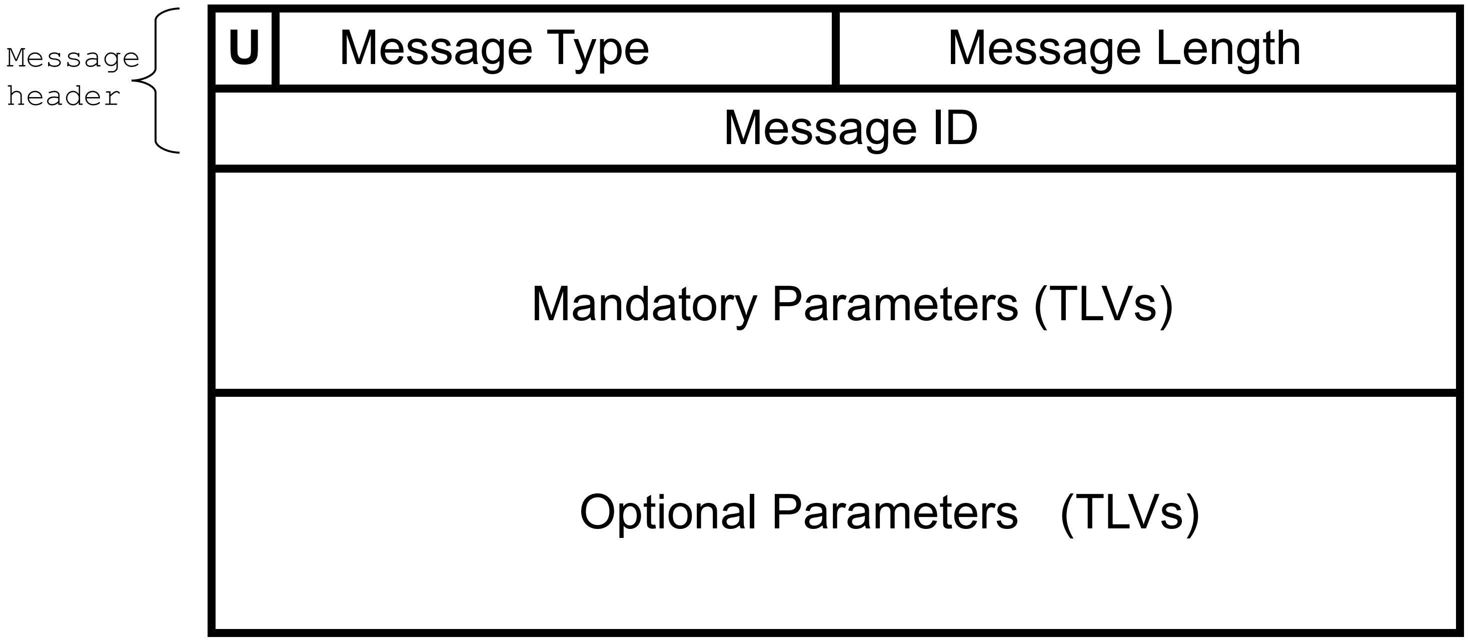

U bit is the Unknown Message bit. If the received message is of unknown type, then if:

U=0, send Notification Message to the originator of this message

U=1, silently ignore the unknown message

LDP Message TLVs

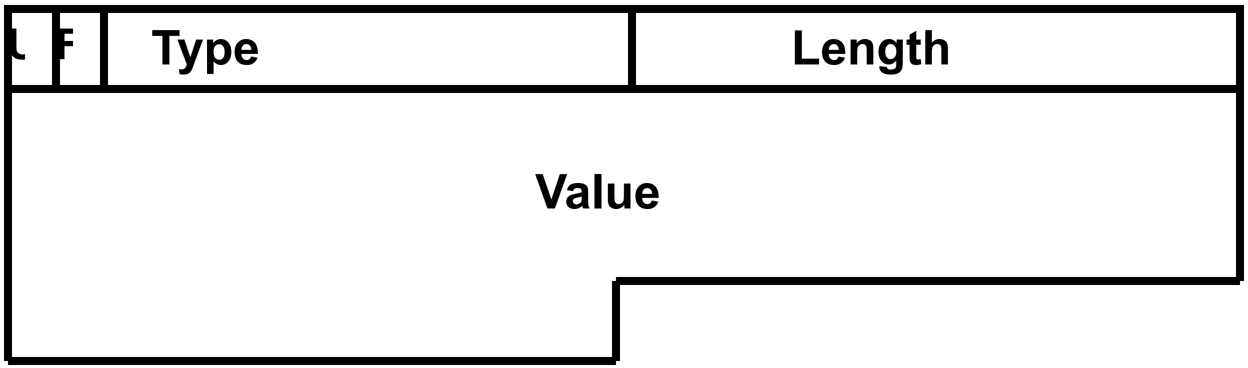

U bit is the Unknown TLV bit. If the received TLV is of unknown type, then if:

U=0, send Notification Message to the originator of this message and ignore the entire message

U=1, silently ignore the unknown TLV and process the rest of the message

F bit is the Forward unknown TLV bit. F bit is only applicable when the U=1

F=0,the unknown TLV is not forwarded with its LDP message

F=1,the unknown TLV is forwarded with its LDP message

LDP Messages Types

- DISCOVERY messages
- ADJACENCY messages deal with initialization, keepalive & shutdown of sessions
- LABEL ADVERTISEMENT messages deal with label binding, requests, withdrawal & release
- NOTIFICATION messages provide advisory information & signal errors

Discovery Message

- Used to discover and maintain the presence of new peers using HELLO messages
- Hello packets (UDP) sent to all-routers multicast address (224.0.0.2)
- Direct unicast hello is sent to non-adjacent neighbors
- Once session is established, HELLO messages serve as link integrity messages
- Session is bi-directional

Adjacency Messages

- **INITIALIZATION**
  - Two LSRs negotiate on various parameters & options
  - These include keepalive timer values, Label ranges, Unsolicited vs. On-demand label advertisement, Ordered vs. Independent mode, Liberal vs. Conservative Label retention
- **KEEPALIVE**
  - LDP message that indicates that neighbor is alive

Label Advertisement related messages

- **LABEL RELEASE**

  An LSR releases a Label Binding that it previously got from it's LDP peer. Used in Conservative Label Retention mode

- **LABEL REQUEST**

  Used by an upstream LSR to request a Label binding for a prefix from the downstream LDP peer. Used in downstream on-demand mode

- **LABEL ABORT REQUEST**

  Send to abort the LABEL REQUEST message

- **LABEL MAPPING**

  Are the TLV object containing <Label, prefix> information

- **LABEL WITHDRAWAL**

  Used to revoke a previously advertised label binding

Notification message

- NOTIFICATION

  Used for Error Notification and Advisory

LDP Session Establishment

LDP establishes a session by performing the following:

- Hello messages are periodically sent on all interfaces that are enabled for MPLS.
- If there is another router connected to that interface, that it also has MPLS enabled, it will respond by trying to establish a session with the source of the hello messages.

UDP is used for hello messages. It is targeted at “all routers on this subnet” multicast address (224.0.0.2).

TCP is used to establish the session.

Both TCP and UDP use well-known LDP port number 646 (711 for TDP).

LDP Hello Message

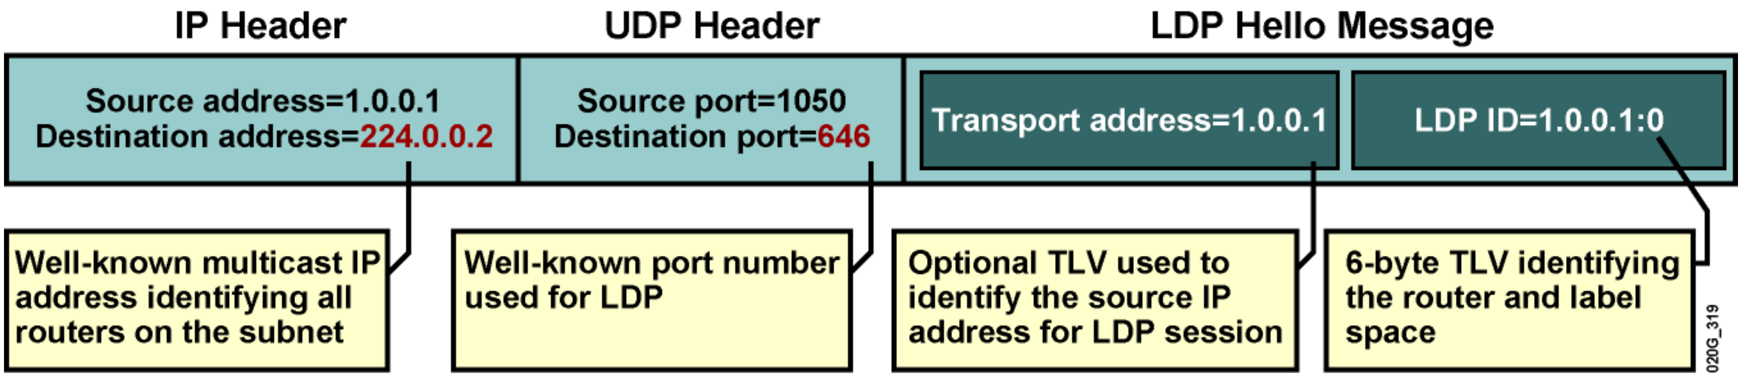

- Hello messages are targeted at all routers reachable through an interface.
- LDP uses well-known (UDP and TCP) port number 646.
- The source address used for an LDP session can be set by adding the transport address TLV to the hello message.
- A 6-byte LDP identifier (TLV) identifies the router (first four bytes) and label space (last two bytes).

LDP Neighbor Discovery

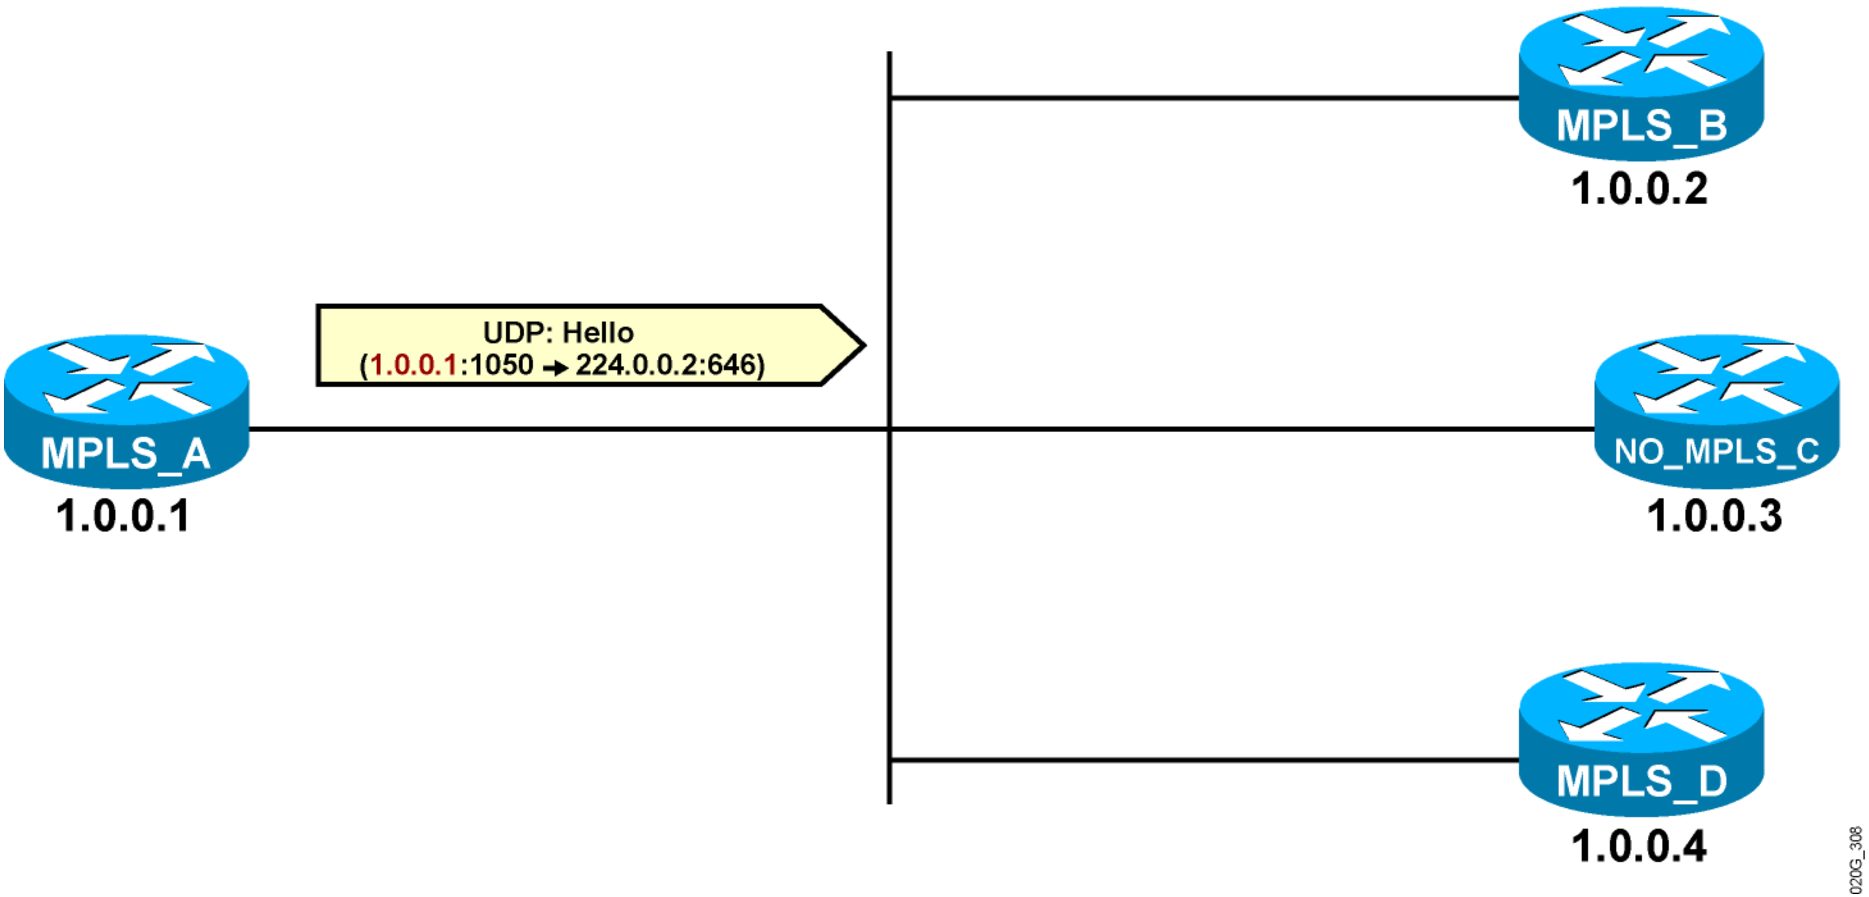

LDP Neighbor Discovery

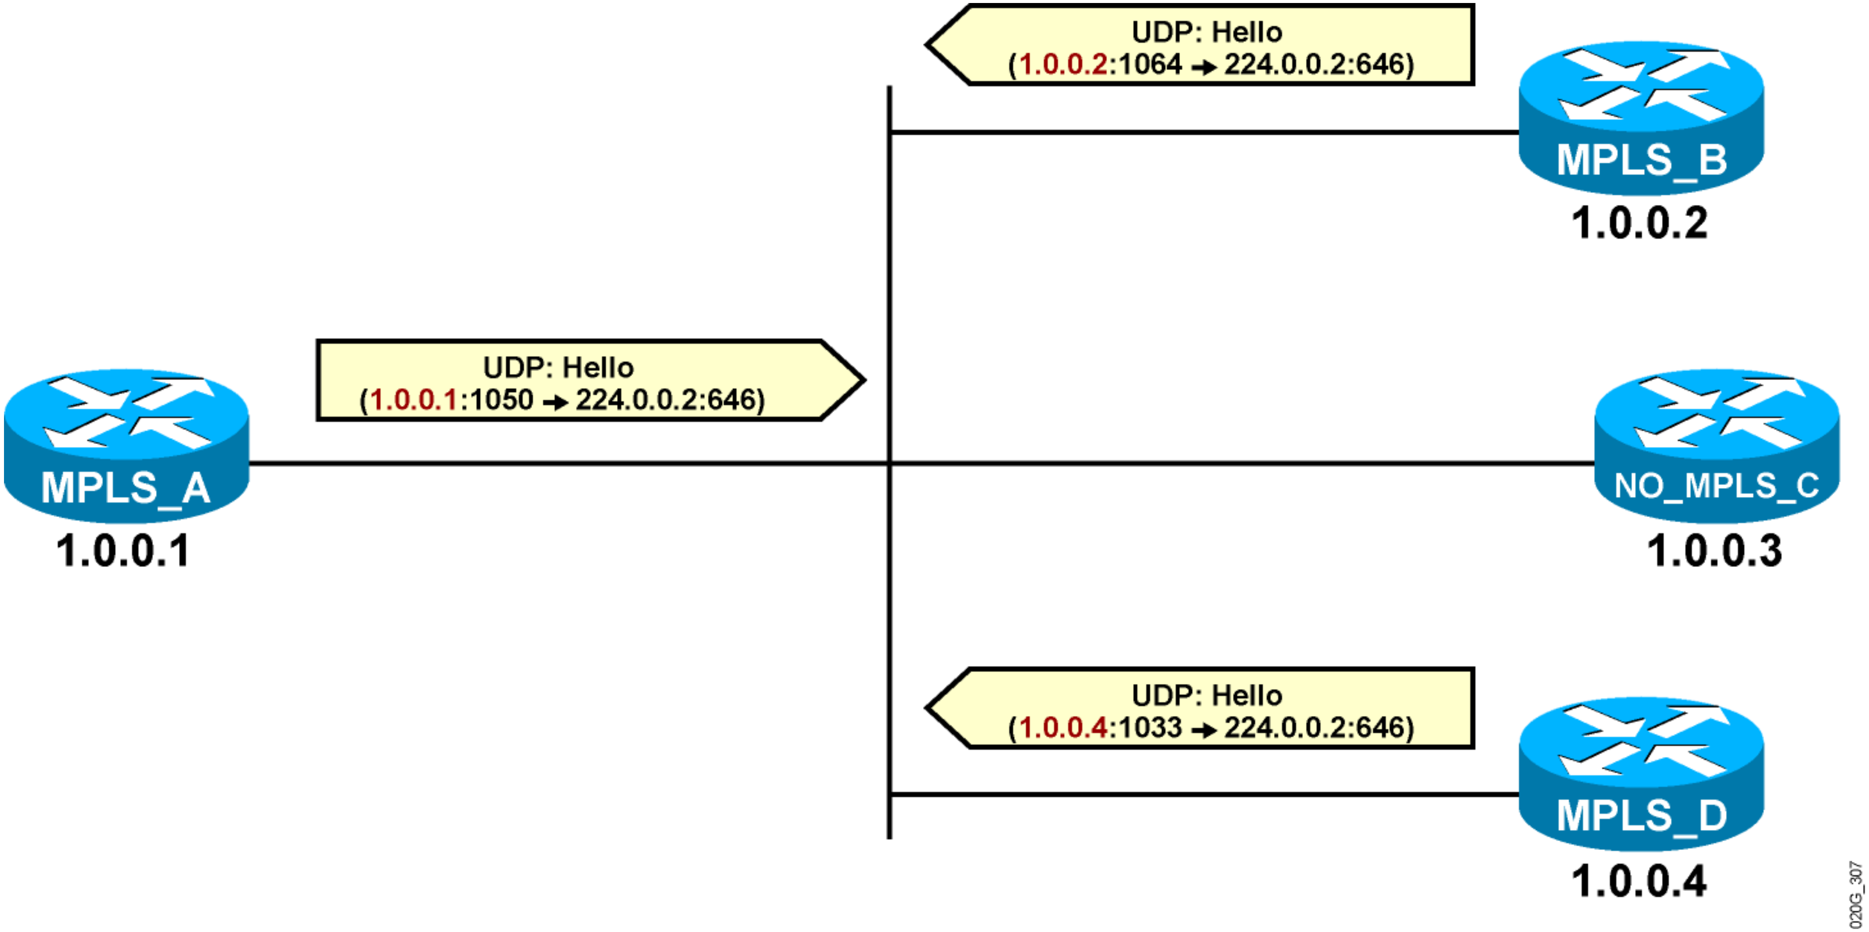

LDP Session: Transport Connection

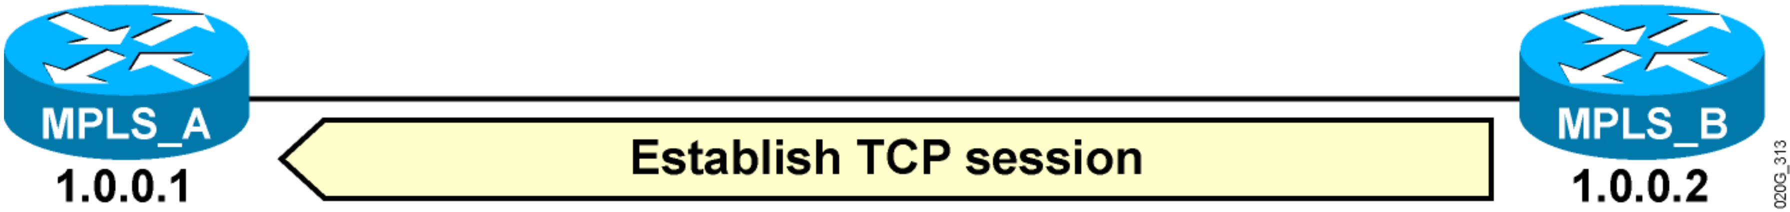

- Once LDP peers receive hellos, they establish a TCP connection
- Peer with higher LDP router-id is active LSR and the peer with lower LDP router-id is the passive LSR
- Active LSR tries to open a TCP connection to the well-known LDP port number 646 of the passive LSR, while the passive LSR waits for the active LSR to initiate the connection

LDP Session: Session Initialization

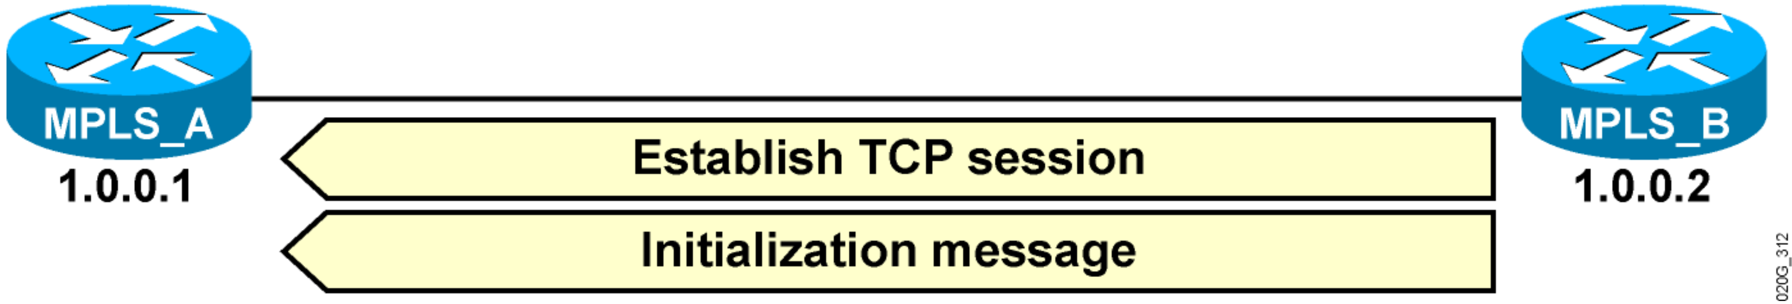

Active LDP peer (1.0.0.2) sends Initialization message to passive LDP peer

Initialization message contains important parameters:

- Session keepalive time (default=180 sec)
- Label distribution method: Downstream unsolicited
- Max PDU length
- Receiver's LDP Identifier
- Whether Loop Detection is enabled
- Some optional parameters

LDP Session: Session Initialization (cont.)

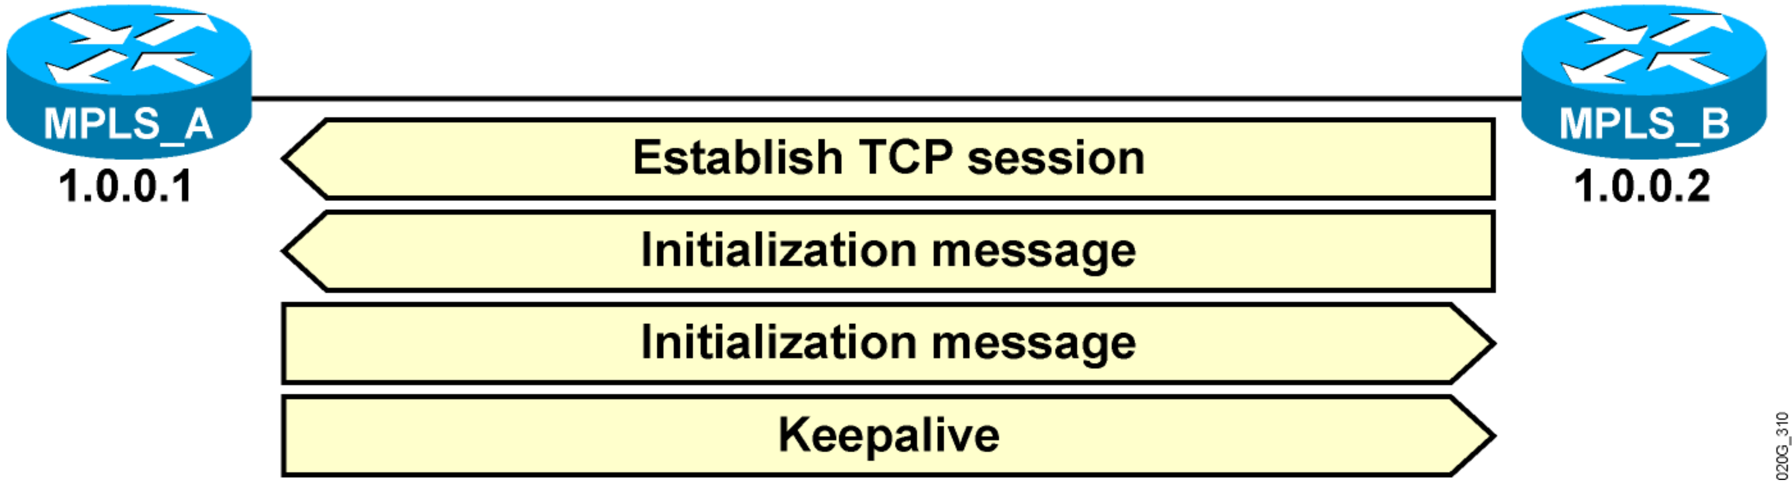

Passive LDP peer sends Initialization message and/or keepalive message to active LDP peer if Initialization message parameters are acceptable

Passive LDP peer could also send Error Notification & close the LDP connection if something was unacceptable

LDP Session:
Session Initialization (cont.)

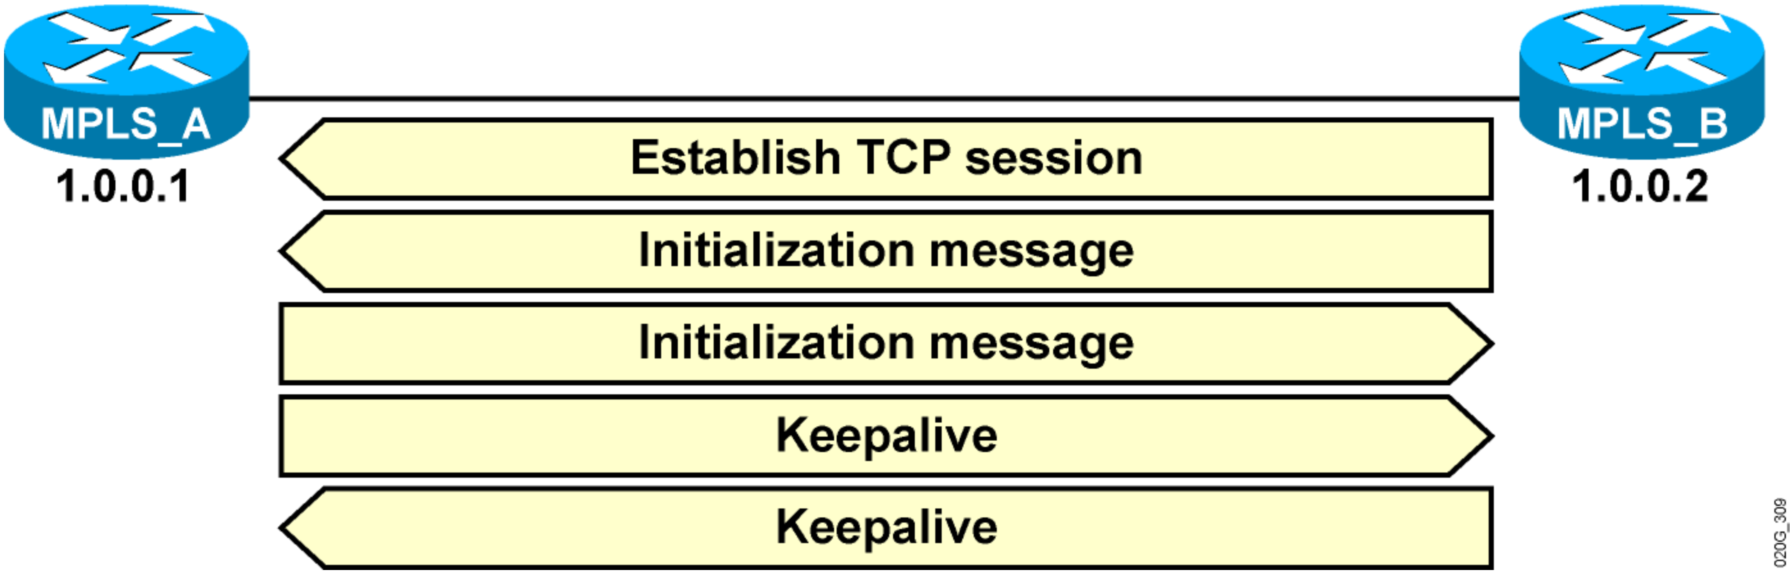

Active LDP peer sends keepalive to passive LDP peer & the LDP session is up

The session is ready to exchange label mappings after receiving the first keepalive.

Targeted LDP Sessions

- LDP neighbor discovery of nonadjacent neighbors differs from normal discovery only in the addressing of hello packets:
  - Hello packets use unicast IP addresses instead of multicast addresses.
- When a neighbor is discovered, the mechanism to establish a session is the same.

Double Lookup Scenario

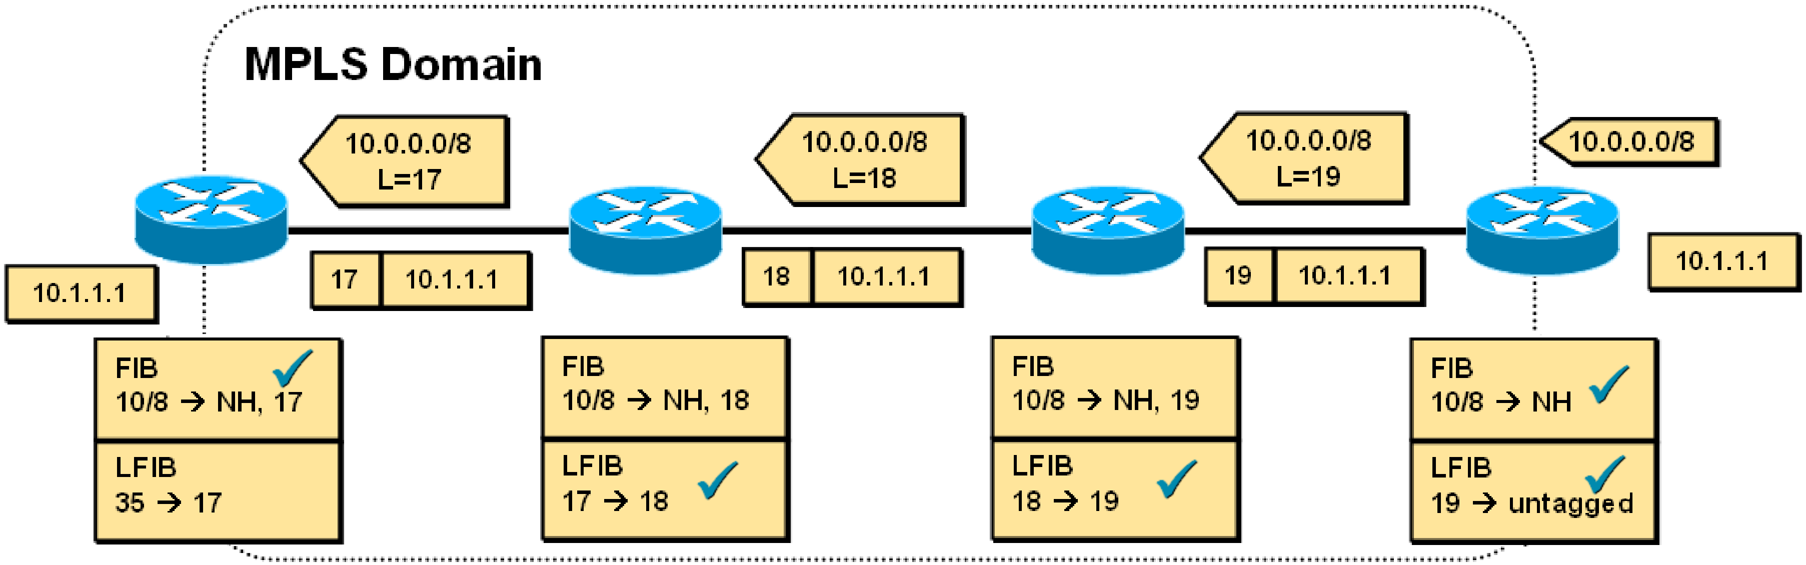

- Double lookup is not an optimal way of forwarding labeled packets.
- A label can be removed one hop earlier.

Double lookup is needed:
1. LFIB: remove the label.
2. FIB: forward the IP packet based on IP next-hop address.

Penultimate Hop Popping

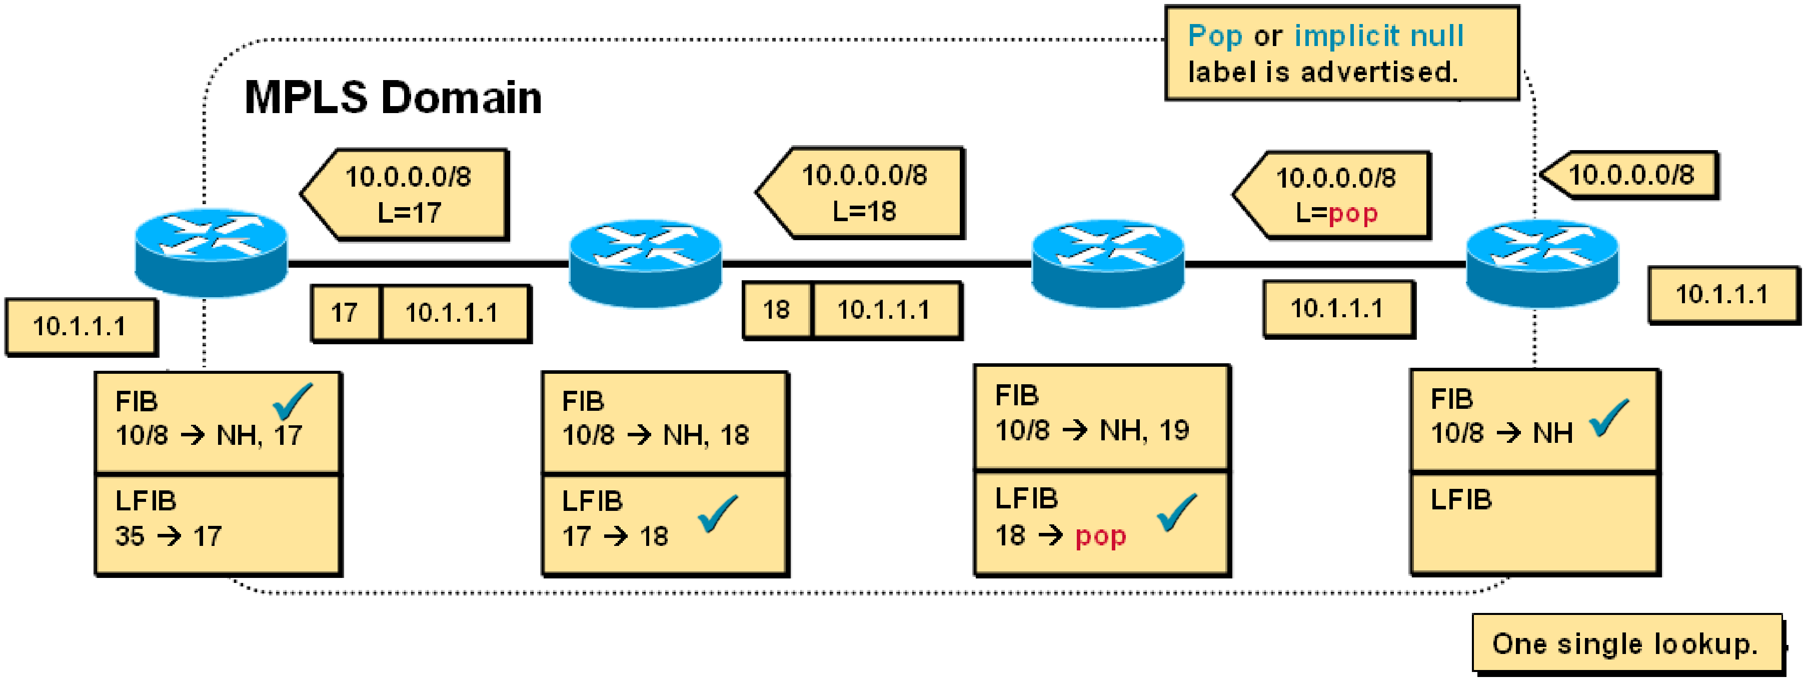

- **A label is removed on the router before the last hop within an MPLS domain.**

LSP Tunnel

- An LSP tunnel is a sequence of LSRs that forward labeled packets of a certain forwarding equivalence class.
- MPLS unicast IP forwarding builds LSP tunnels based the output of IP routing protocols.
- LDP or TDP advertises labels only for individual segments in the LSP tunnel.
- LSP tunnels are unidirectional.
- Return traffic uses a different LSP tunnel (usually the reverse path, as most routing protocols provide symmetrical routing).
- An LSP tunnel can take a different path from the one chosen by an IP routing protocol (MPLS traffic engineering).

LSP Tunnel Building Example

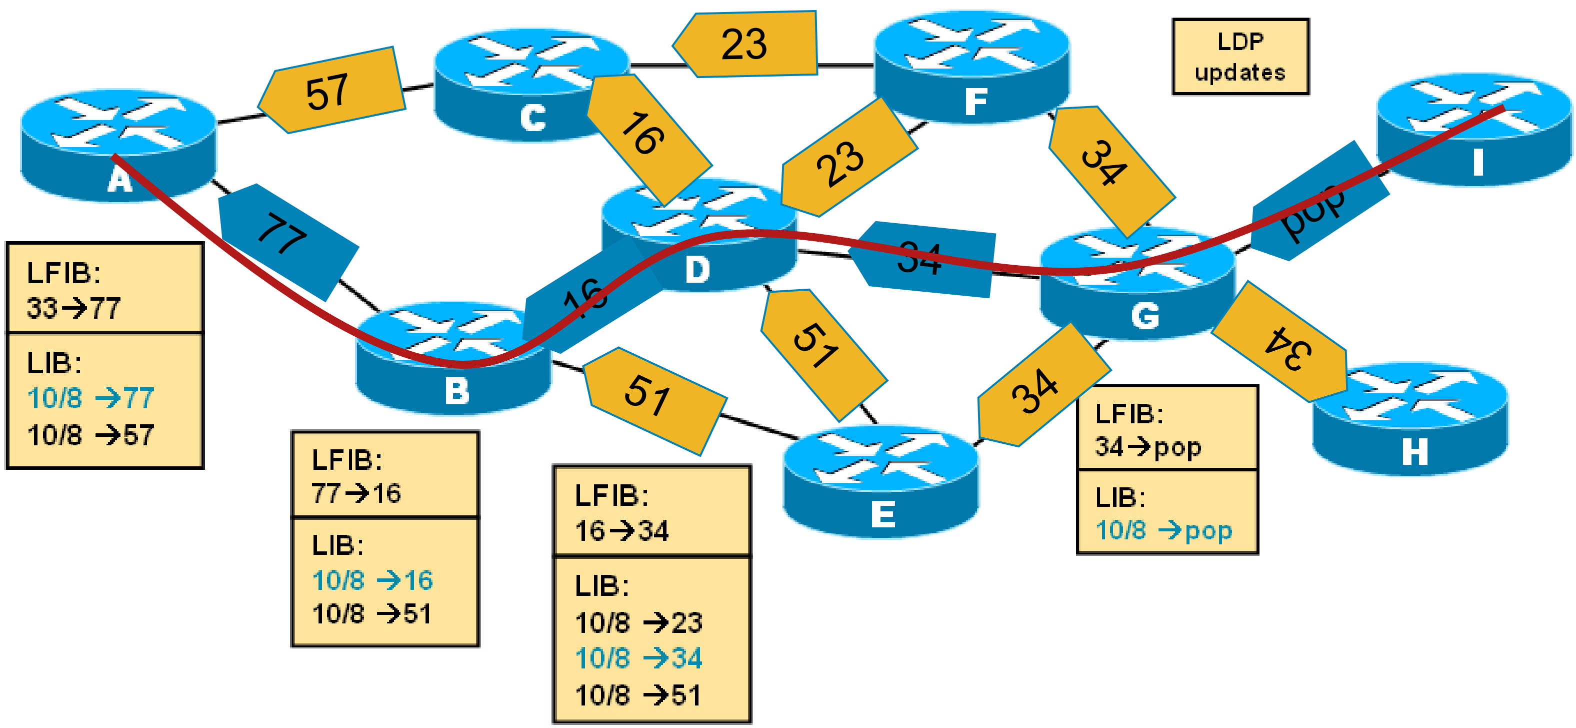

Loop Detection

- LDP or TDP relies on loop detection mechanisms built into IGPs that are used to determine the path.
- If, however, a loop is generated (that is, misconfiguration with static routes), the TTL field in the label header is used to prevent indefinite looping of packets.
- TTL functionality in the label header is equivalent to TTL in the IP headers.
- TTL is usually copied from the IP headers to the label headers (TTL propagation).

Loop Detection

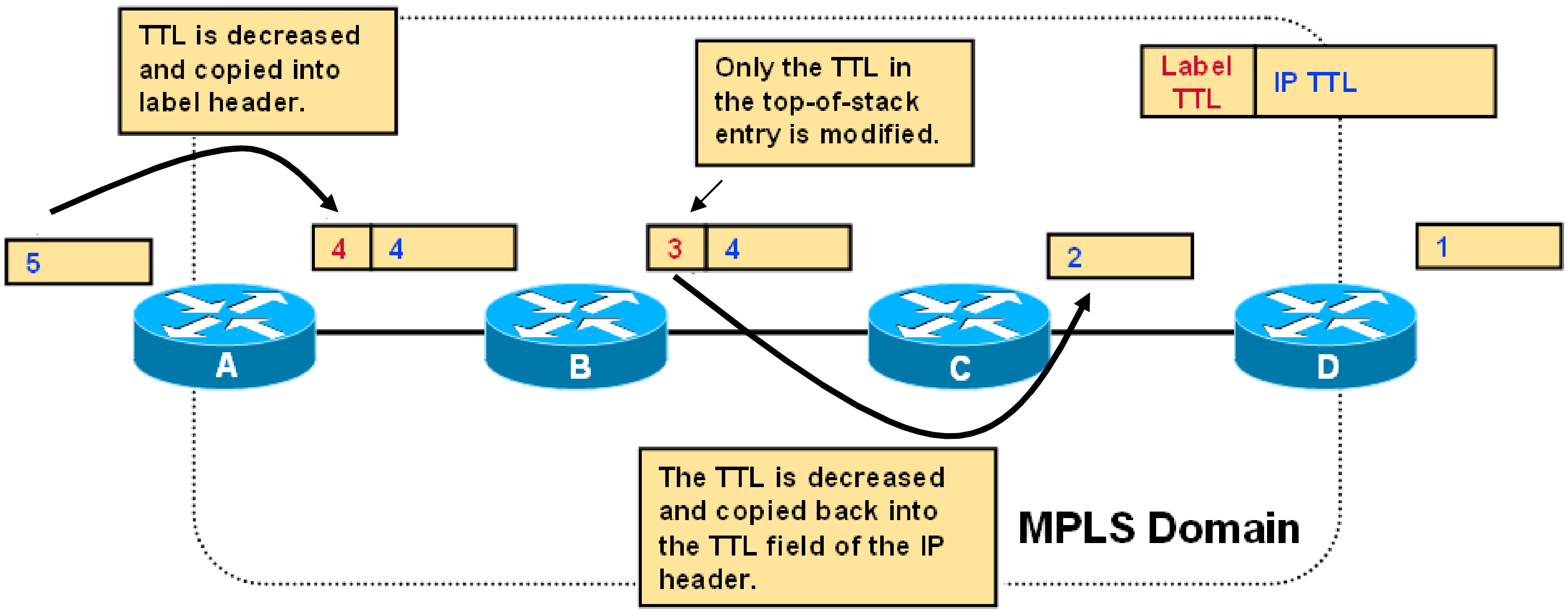

Cisco routers have TTL propagation enabled by default.

- On ingress: TTL is copied from IP header to label header.
- On egress: TTL is copied from label header to IP header.

Disabling TTL Propogation

- TTL propagation can be disabled.
- IP TTL value is not copied into the labels and label TTL is not copied back into IP TTL.
- Instead, the value 255 is assigned to the label header TTL field on the ingress LSR.
- Disabling TTL propagation hides core routers in the MPLS domain.
- Traceroute across an MPLS domain does not show any core routers.

Configuring MPLS

Configuring MPLS

Mandatory:

- Enable CEF switching.
- Configure Tag Distribution Protocol or Label Distribution Protocol on every label-enabled interface.

Optional:

- Configure MTU size for labeled packets.
- Configure IP TTL propagation.
- Configure conditional label advertising.

Configuring LDP

Global

ip cef <distributed>

mpls label protocol <ldp | tdp | both>

tag-switching tdp router-id Loopback0

mpls ldp explicit-null (optional)

no mpls ip propagate-ttl (optional)

Interface

mpls ip **or** **tag-switching ip** (enables this interface for MPLS forwarding)

mpls label protocol ldp

(optional, if you want to run LDP on this interface only, while other interfaces don't run LDP or run another label protocol such as TDP)

Configuring Conditional Label Distribution

Router(config)#

tag-switching advertise-tags for *net-acl* [to *tdp-acl*]

- By default, labels for all destinations are announced to all LDP/TDP neighbors.
- This command enables you to selectively advertise some labels to some LDP/TDP neighbors.
- Conditional label advertisement only works over frame-mode interfaces.
- Parameters:
    - Net-ACL – the IP ACL that selects the destinations for which the labels will be generated.
    - TDP-ACL – the IP ACL that selects the TDP neighbors that will receive the labels.

Conditional Label Distribution Example

- The customer is already running IP infrastructure.
- MPLS is only needed to support MPLS/VPN services.
  - Labels should only be generated for loopback interfaces (BGP next-hops) of all routers.
  - All loopback interfaces are in one contiguous address block (192.168.254.0/24).

Conditional Label Distribution Router Configuration

- Enable conditional label advertisment

```
no tag-switching advertise-tags
!
! Configure conditional advertisments
!
tag-switching advertise-tags for 90 to 91
!
access-list 90 permit ip 192.168.254.0 0.0.0.255
access-list 91 permit ip any
```

Monitoring LDP

- `show mpls interface <x> detail`
- `show mpls ldp discovery`
- `show mpls ldp neighbor`
- `show mpls ip/ldp binding <prefix> <prefix-length>`
- `show mpls forwarding-table <prefix> <prefix-length>`
- `sh ip cef <prefix>`
- `show mpls ldp parameters`

Show mpls interface

mpls-7200a#sh mpls interface

Interface	IP	Tunnel	Operational
Ethernet3/0	Yes (ldp)	No	Yes

mpls-7200a#sh mpls interface ethernet3/0 detail

Interface Ethernet3/0:

IP labeling enabled (ldp)

……..<snip>…….

Fast Switching Vectors:

IP to MPLS Fast Switching Vector

MPLS Turbo Vector

MTU = 1500

Show mpls interface (contd..)

- “sh mpls interface [detail]”

  Lists whether MPLS is enabled and the application that enabled MPLS on the interface

```
PE2#sh mpls interface
Interface              IP            Tunnel    Operational
Serial2/0              Yes (ldp)     No        Yes
PE2#
```

MPLS Enabled

```
PE2#sh mpls interface ser2/0 detail
Interface Serial2/0:
        IP labeling enabled (ldp)
        LSP Tunnel labeling not enabled
        BGP tagging not enabled
        Tagging operational
        Fast Switching Vectors:
          IP to MPLS Fast Switching Vector
          MPLS Turbo Vector
        MTU = 1508
PE2#
```

LDP Enabled

MPLS MTU

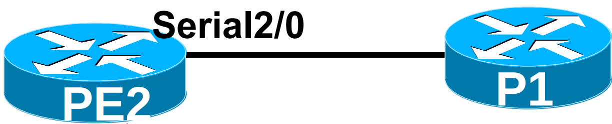


```
!
interface Serial2/0
 description To P1 ser2/0
 ip address 10.13.2.6/30
 mpls label protocol ldp
 tag-switching ip
 tag-switching mtu 1508
!
```

Show mpls interface (contd..)

- This slide is to show that BGPipv4+label (or MP-eBGP) is another application that can enable MPLS; *what's different here -*

```
RSP-PE-SOUTH-6#sh mpls int
Interface              IP            Tunnel    Operational
Fddi1/0/0              Yes (ldp)     No        Yes
ATM1/1/0.108           No            No        Yes
RSP-PE-SOUTH-6#
```

MPLS is Operational.

LDP not enabled

```
RSP-PE-SOUTH-6#sh mpls int ATM1/1/0.108 detail
Interface ATM1/1/0.108:
        IP labeling not enabled
        LSP Tunnel labeling not enabled
        BGP tagging enabled
        Tagging operational
        Optimum Switching Vectors:
          IP to MPLS Feature Vector
          MPLS Feature Vector
        Fast Switching Vectors:
          IP to MPLS Fast Feature Switching Vector
          MPLS Feature Vector
        MTU = 4470
RSP-PE-SOUTH-6#
```

LDP not enabled

BGP+Label Enabled

MPLS MTU

LDP discovery/adjacency: commands and debugs

- `show mpls ldp discovery`
- `debug mpls ldp transport`
- `debug mpls ldp session io`

LDP discovery

Interface eth3/0 configured with LDP

```
mpls-7200a#sh mpls ldp discovery
 Local LDP Identifier:
     4.4.4.4:0
     Discovery Sources:
     Interfaces:
         Ethernet3/0 (ldp): xmit/recv
             LDP Id: 5.5.5.5:0
```

My LDP id

we are transmitting & receiving LDP messages

Neighbor's LDP id

- "debug mpls ldp transport events"

  Should give information regarding whether the HELLOS are advertised/received

LDP adjacency debugs

LDP discovery, connection setup and shutdown events

```
mpls-7200a#debug mpls ldp transport events
debugging for LDP discovery and connection setup / shutdown events
2d11h: ldp: Send ldp hello; Ethernet3/0, src/dst 10.0.3.4/224.0.0.2, inst_id 0
2d11h: ldp: Rcvd ldp hello; Ethernet3/0, from 10.0.3.5 (5.5.5.5:0), intf_id 0, opt 0xC
shutting neighbor
2d11h: %CLNS-5-ADJCHANGE: ISIS: Adjacency to mpls-12008a (Ethernet3/0) Down, hold time expired
2d11h: ldp:Discovery hold timer expired for adj 0x17D45A0, 5.5.5.5:0,will close conn
2d11h: ldp: Discovery hold timer expired for adj 0x17D45A0; 5.5.5.5:0
2d11h: ldp:     adj_addr/adj_xport_addr: 10.0.3.5/5.5.5.5
2d11h: ldp: LDP ptcl SM; close xport request for adj 0x0
2d11h: ldp: Close LDP transport conn for adj 0x17D45A0
2d11h: ldp: Closing ldp conn 4.4.4.4:646 <-> 5.5.5.5:11012, adj 0x17D45A0
2d11h: ldp: Adj 0x17D45A0; state set to closed
2d11h: ldp: Send ldp hello; Ethernet3/0, src/dst 10.0.3.4/224.0.0.2, inst_id 0
```

LDP session i/o debug

LDP session I/O, excluding periodic Keep Alives

```
mpls-7200a#debug mpls ldp session io <all>

bringing neighbor down

2d11h: %CLNS-5-ADJCHANGE: ISIS: Adjacency to mpls-12008a (Ethernet3/0) Down, hold
time expired

2d11h: ldp: Sent notif msg to 5.5.5.5:0 (pp 0x17A0870)

.............

bringing neighbor up

2d11h: %CLNS-5-ADJCHANGE: ISIS: Adjacency to mpls-12008a (Ethernet3/0) Up, new

adjacency

2d11h: ldp: Rcvd init msg from 5.5.5.5 (pp 0x0)

2d11h: ldp: Sent init msg to 5.5.5.5:0 (pp 0x0)

2d11h: ldp: Sent keepalive msg to 5.5.5.5:0 (pp 0x0)

2d11h: ldp: Rcvd keepalive msg from 5.5.5.5:0 (pp 0x0)

2d11h: ldp: Sent address msg to 5.5.5.5:0 (pp 0x186CB38)

2d11h: ldp: Sent label mapping msg to 5.5.5.5:0 (pp 0x186CB38)

...............
```

LDP neighbor

```
mpls-7200a#sh mpls ldp neighbor
    Peer LDP Ident: 5.5.5.5:0; Local LDP Ident 4.4.4.4:0
       TCP connection: 5.5.5.5.11000 - 4.4.4.4.646
       State: Oper; Msgs sent/rcvd: 268/264; Downstream Up time: 03:41:45
       LDP discovery sources:
          Ethernet3/0, Src IP addr: 10.0.3.5
       Addresses bound to peer LDP Ident:
          10.0.3.5        10.0.4.5        10.0.5.5        5.5.5.5
```

LDP neighbor (contd..)

- LDP session is a TCP session (port = 646)
- Multiple links between two routers still mean single LDP session.

```
PE1#sh mpls ldp neighbor
    Peer LDP Ident: 10.13.1.101:0; Local LDP Ident 10.13.1.61:0
        TCP connection: 10.13.1.101.11031 - 10.13.1.61.646
        State: Oper; Msgs sent/rcvd: 58/60; Downstream
        Up time: 00:39:27
        LDP discovery sources:
          Ethernet0/0, Src IP addr: 10.13.1.5
          Ethernet1/0, Src IP addr: 10.13.1.9
        Addresses bound to peer LDP Ident:
          10.13.1.9       10.13.1.5       10.13.2.5     10.13.1.101
```

```
PE1#sh tcp brief| i 646
43ABB020  10.13.1.101.11031       10.13.1.61.646         ESTAB
PE1#
```

LDP_ID

Unsolicited Label Distribution*

Interfaces on which peer is discovered

Peer's Connected int

LDP binding commands

- "sh mpls ip binding detail"

  Lists all prefixes with labels & LDP neighbors

- "sh mpls ip binding <*prefix*> <mask> det"

  Lists ACLs (if any), *prefix* bindings, and LDP neighbors. Notice "Advertised to:" field.

- "sh mpls ip binding advertisement-acls"

  Lists LDP filter, if there is any, on the first line. Prefixes followed by "Advert acl(s):" are advertised via LDP, others are not.

LIB information

```
mpls-7200a#sh mpls ip binding 12.12.12.12 32
  12.12.12.12/32
        in label:     21
        out label:    19        lsr: 5.5.5.5:0        in use

mpls-7200a#sh mpls ldp binding 12.12.12.12 32
  tib entry: 12.12.12.12/32, rev 48
        local binding:  tag: 21
        remote binding: tsr: 5.5.5.5:0, tag: 19
```

LDP binding related debugs

```
mpls-7200a#debug mpls ldp bindings

shutting neighbor

2d11h: %CLNS-5-ADJCHANGE: ISIS: Adjacency to mpls-12008a (Ethernet3/0) Down, hold
time expired

2d11h: tagcon: tibent(5.5.5.5/32): label imp-null from 5.5.5.5:0 removed

2d11h: tagcon: route_tag_change for: 5.5.5.5/32

        inlabel 16, outlabel withdrwn, nexthop lsr 5.5.5.5:0, reason response to
find_route_tags

2d11h: tagcon: Deassign peer id; 5.5.5.5:0: id 0

2d11h: tagcon: tc_iprouting_table_change: 5.5.5.5/255.255.255.255, event 0x2

2d11h: tagcon: rib change: 5.5.5.5/255.255.255.255; event 0x2; ndb attrflags
0x1000000;

ndb->pdb_index/pdb->index 0x3/0x3

2d11h: tagcon: rib change: 5.5.5.5/255.255.255.255; event 0x2; ndb attrflags
0x1000000;

ndb->pdb_index/pdb->index 0x3/undef
```

LDP Advertisement related debugs

```
mpls-7200a#debug mpls ldp advertisements

shutting neighbor

2d11h: %CLNS-5-ADJCHANGE: ISIS: Adjacency to mpls-12008a (Ethernet3/0) Down,
hold time expired

2d11h: tagcon: Deassign peer id; 5.5.5.5:0: id 0

activating neighbor

2d11h: %CLNS-5-ADJCHANGE: ISIS: Adjacency to mpls-12008a (Ethernet3/0) Up,
new adjacency

2d11h: tagcon: Assign peer id; 5.5.5.5:0: id 0

2d11h: tagcon: peer 5.5.5.5:0 (pp 0x17AFAE0): advertise 4.4.4.4

2d11h: tagcon: Advertise labels: Clear LDP_CTX_TCB_FLAGS_ENULL_RECFG

2d11h: tagcon: peer 5.5.5.5:0 (pp 0x17AFAE0): advertise 4.4.4.4/32, label 3
(imp-null) (#32)
```

LFIB information

- `show mpls forwarding-table <prefix> <prefix-length>`
- `sh ip cef <prefix> internal`

Looking at LFIB

Ethertype=
8847
Label Value in
MPLS shim=
13 Hex=19 dec

Looking at LFIB on 12008a

```
mpls-12008a#sh mpls forwarding 12.12.12.12 32 detail
Local  Outgoing    Prefix            Bytes tag  Outgoing    Next Hop
tag    tag or VC   or Tunnel Id      switched   interface
19     19          12.12.12.12/32    498        Et2/0       10.0.4.11
        MAC/Encaps=14/18, MTU=1500, Tag Stack{19}
        AABBCC000502AABBCC0004028847 00013000
        No output feature configured
 Per-destination load-sharing, slots: 0 2 4 6 8 10 12 14
       19          12.12.12.12/32    498        Et3/0       10.0.5.11
        MAC/Encaps=14/18, MTU=1500, Tag Stack{19}
        AABBCC000503AABBCC0004038847 00013000
        No output feature configured
 Per-destination load-sharing, slots: 1 3 5 7 9 11 13 15
```

Destination MAC=
AABBCC000502
Source MAC=
AABBCC000402

CEF command

```
mpls-12008a#sh ip cef 12.12.12.12 internal

12.12.12.12/32, version 24, epoch 0, per-destination sharing
0 packets, 0 bytes
  tag information set, local tag: 19
  via 10.0.4.11, Ethernet2/0, 0 dependencies
    traffic share 1
    next hop 10.0.4.11, Ethernet2/0
    valid adjacency
    tag rewrite with Et2/0, 10.0.4.11, tags imposed: {19}
  via 10.0.5.11, Ethernet3/0, 0 dependencies
    traffic share 1
    next hop 10.0.5.11, Ethernet3/0
    valid adjacency
    tag rewrite with Et3/0, 10.0.5.11, tags imposed: {19}
```

0 packets, 0 bytes switched through the prefix

……………… (contd..)

tmstats: external 0 packets, 0 bytes

internal 0 packets, 0 bytes

Load distribution: 0 1 0 1 0 1 0 1 0 1 0 1 0 1 0 1 (refcount 1)

Hash	OK	Interface	Address	Packets	Tags imposed
1	Y	Ethernet2/0	10.0.4.11	0	{19}
2	Y	Ethernet3/0	10.0.5.11	0	{19}
3	Y	Ethernet2/0	10.0.4.11	0	{19}
4	Y	Ethernet3/0	10.0.5.11	0	{19}
5	Y	Ethernet2/0	10.0.4.11	0	{19}
6	Y	Ethernet3/0	10.0.5.11	0	{19}
7	Y	Ethernet2/0	10.0.4.11	0	{19}
8	Y	Ethernet3/0	10.0.5.11	0	{19}
9	Y	Ethernet2/0	10.0.4.11	0	{19}
10	Y	Ethernet3/0	10.0.5.11	0	{19}
11	Y	Ethernet2/0	10.0.4.11	0	{19}
12	Y	Ethernet3/0	10.0.5.11	0	{19}
13	Y	Ethernet2/0	10.0.4.11	0	{19}
14	Y	Ethernet3/0	10.0.5.11	0	{19}
15	Y	Ethernet2/0	10.0.4.11	0	{19}
16	Y	Ethernet3/0	10.0.5.11	0	{19}

Monitoring LDP: LDP parameters

```
mpls-7200a#sh mpls ldp parameters
Protocol version: 1
Downstream label generic region: min label: 16; max label: 100000
Session hold time: 180 sec; keep alive interval: 60 sec
Discovery hello: holdtime: 15 sec; interval: 5 sec
Discovery targeted hello: holdtime: 180 sec; interval: 5 sec
Downstream on Demand max hop count: 255
TDP for targeted sessions
LDP initial/maximum backoff: 15/120 sec
LDP loop detection: off
```

Forwarding traffic down the LSP

```
mpls-7200a#sh mpls forwarding-table 12.12.12.12
Local   Outgoing     Prefix             Bytes tag  Outgoing    Next Hop
tag     tag or VC    or Tunnel Id       switched   interface
21      19           12.12.12.12/32     0          Et3/0       10.0.3.5
Note: Bytes tag switched this will increment if packets are being tag
switched using this entry
```

```
mpls-12008a#sh mpls forwarding-table label 19
Local   Outgoing     Prefix             Bytes tag  Outgoing    Next Hop
tag     tag or VC    or Tunnel Id       switched   interface
19      19           12.12.12.12/32     498        Et2/0       10.0.4.11
        19           12.12.12.12/32     1176       Et3/0       10.0.5.11
```

```
mpls-12008b#sh mpls forwarding-table labels 19
Local   Outgoing     Prefix             Bytes tag  Outgoing    Next Hop
tag     tag or VC    or Tunnel Id       switched   interface
19      Pop tag      12.12.12.12/32     4176       Et1/0       10.0.17.12
```


LDP binding and advertisemsnt debugs

- Be Careful on the production routers
- “debug mpls ldp advertisements”

    Useful to see label bindings that are advertised

- “debug mpls ldp binding”

    Useful to see label bindings that are received

- “debug mpls ldp message sent|received”

    Useful for the protocol understanding purposes

MPLS Configuration Example

Enable MPLS on all core interfaces in your network.

```
ip cef

interface hssi 1/0
 mpls ip

interface fastethernet 0/0
 mpls ip
```

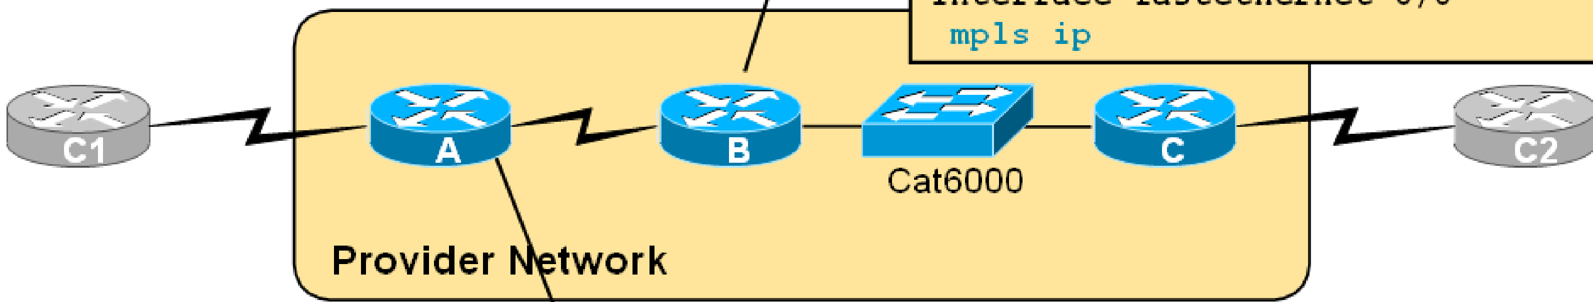


```
ip cef

interface serial 3/1
 ip access-group NoTDP in

interface hssi 1/0
 mpls ip

ip access-list NoTdp deny tcp any any eq 711
ip access-list NoTdp permit ip any any
```

Never run MPLS with your customers.

Use access lists to prevent customers from running TDP with your routers.

MPLS on LAN Configuration Example

Jumbo frames have to be enabled on the switch.

```
set port 1/3 jumbo enable
set port 1/4 jumbo enable
```

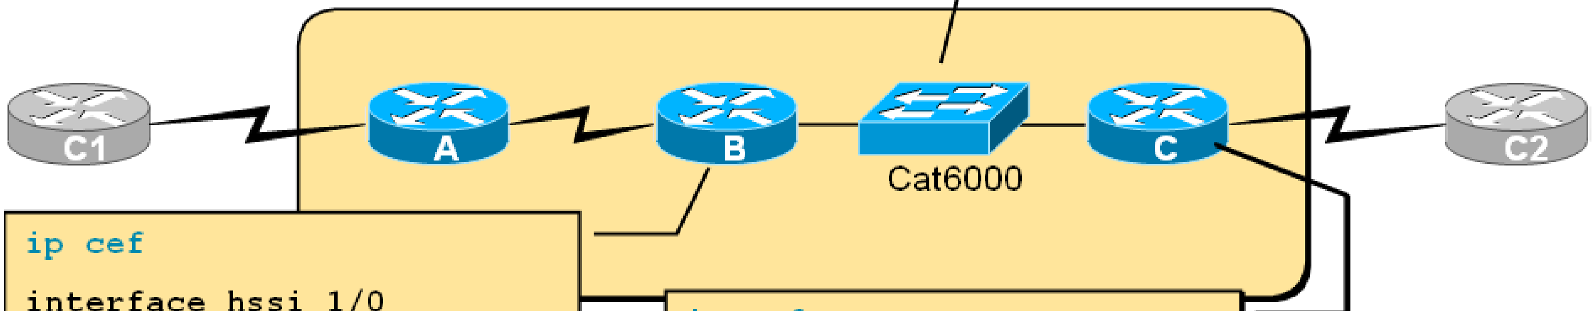


```
ip cef

interface hssi 1/0
 mpls ip

interface fastethernet 0/0
 mpls ip
 mpls mtu 1512
```

```
ip cef

interface fastethernet 0/0
 mpls ip
 mpls mtu 1512
```

MPLS MTU is increased to 1512 to support 1500-byte IP packets and MPLS stack up to three levels deep.

Configuring IP TTL Propogation

router(config)#

no mpls ip propagate-ttl	**12.1(3)T**

- By default, IP TTL is copied into label header at label imposition and label TTL is copied into IP TTL at label removal.
- This command disables IP TTL and label TTL propagation.
  - TTL value of 255 is inserted in the label header.
- The TTL propagation has to be disabled on ingress and egress edge LSR.

sh ip cef detail

```
Router#show ip cef 192.168.20.0 detail
192.168.20.0/24, version 23, cached adjacency to Serial1/0.2
0 packets, 0 bytes
  tag information set
    local tag: 33
    fast tag rewrite with Se1/0.2, point2point, tags imposed: {32}
  via 192.168.3.10, Serial1/0.2, 0 dependencies
    next hop 192.168.3.10, Serial1/0.2
    valid cached adjacency
    tag rewrite with Se1/0.2, point2point, tags imposed: {32}
```


sh mpls ldp neighbor

```
Router#show tag-switching tdp neighbors
Peer TDP Ident: 192.168.3.100:0; Local TDP Ident
192.168.3.102:0
        TCP connection: 192.168.3.100.711 - 192.168.3.102.11000
        State: Oper; PIEs sent/rcvd: 55/53; ; Downstream
        Up time: 00:43:26
        TDP discovery sources:
          Serial1/0.2
        Addresses bound to peer TDP Ident:
          192.168.3.10    192.168.3.14    192.168.3.100
```


sh mpls ldp discovery

```
Router#show tag-switching tdp discovery
Local TDP Identifier:
    192.168.3.102:0
TDP Discovery Sources:
    Interfaces:
        Serial1/0.1: xmit/recv
            TDP Id: 192.168.3.101:0
        Serial1/0.2: xmit/recv
            TDP Id: 192.168.3.100:0
```

sh mpls forwarding table

```
Router#show tag-switching forwarding-table detail
Local   Outgoing     Prefix             Bytes tag  Outgoing    Next Hop
tag     tag or VC    or Tunnel Id       switched   interface
26       Untagged    192.168.3.3/32     0          Se1/0.3     point2point
       MAC/Encaps=0/0, MTU=1504, Tag Stack{}
27       Pop tag     192.168.3.4/32     0          Se0/0.4     point2point
       MAC/Encaps=4/4, MTU=1504, Tag Stack{}
       20618847
28       29          192.168.3.4/32     0          Se1/0.3     point2point
       MAC/Encaps=4/8, MTU=1500, Tag Stack{29}
       18718847 0001D000
```


MPLS LAB Topology

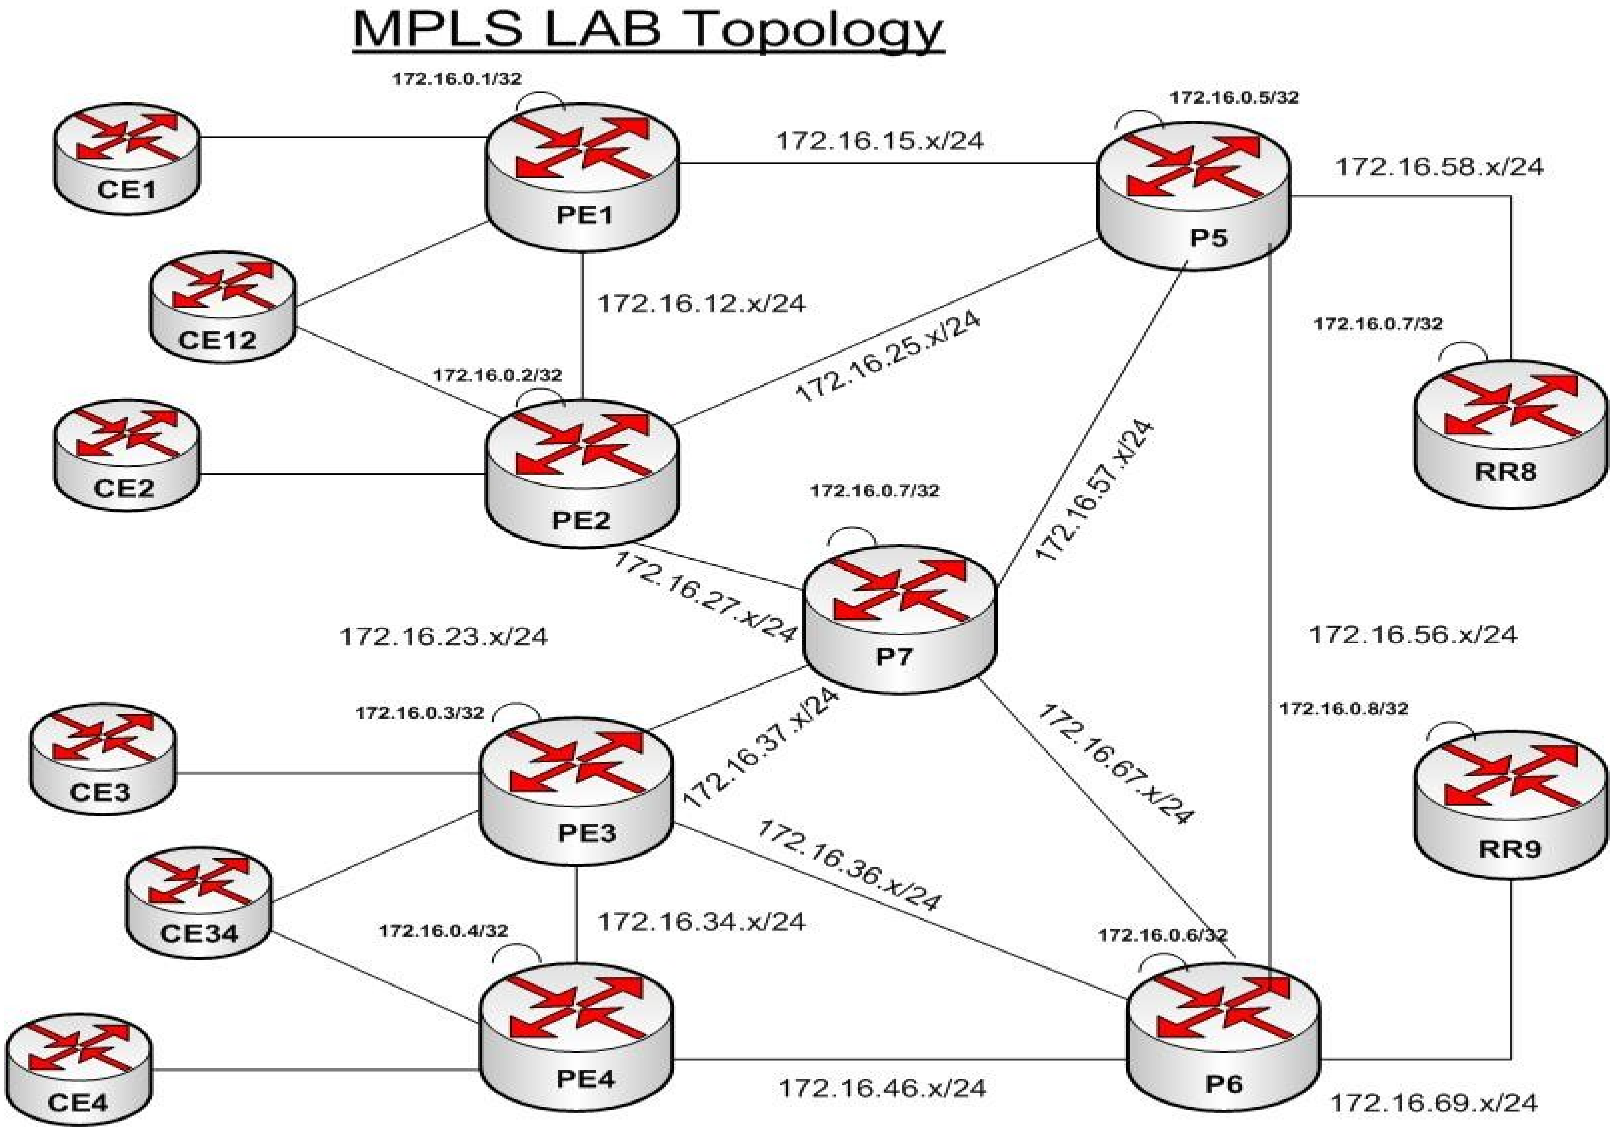

MPLS in BGP Transit AS

MPLS – BGP Interation

- Labels are assigned to FECs.
- FEC in unicast IP routing is equal to a destination prefix found in an IP routing table.
- This is true only for IGP-derived prefixes.
- BGP-derived prefixes are assigned the label that is used for the BGP next-hop address.
- Result: all prefixes learned from an external BGP neighbor use a single label.

Traditional BGP AS Design Requirement

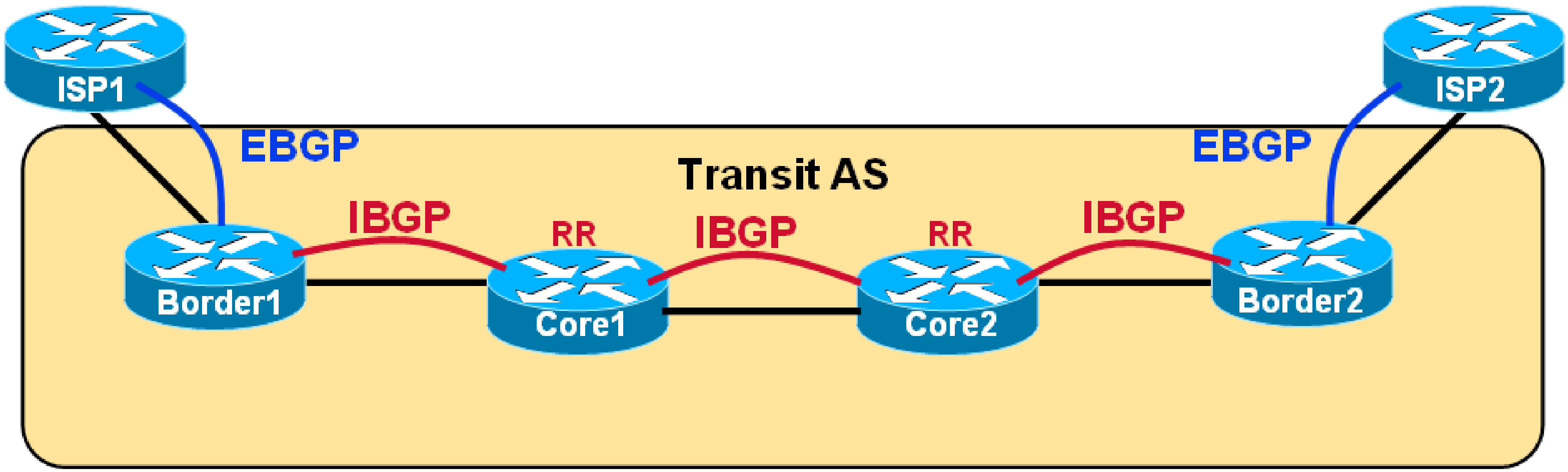

- All core routers are required to run BGP.
- All core routers require full Internet routing information (more than 250 000 networks) to be able to forward IP packets between ISP1 and ISP2.

Simplified BGP Design with MPLS

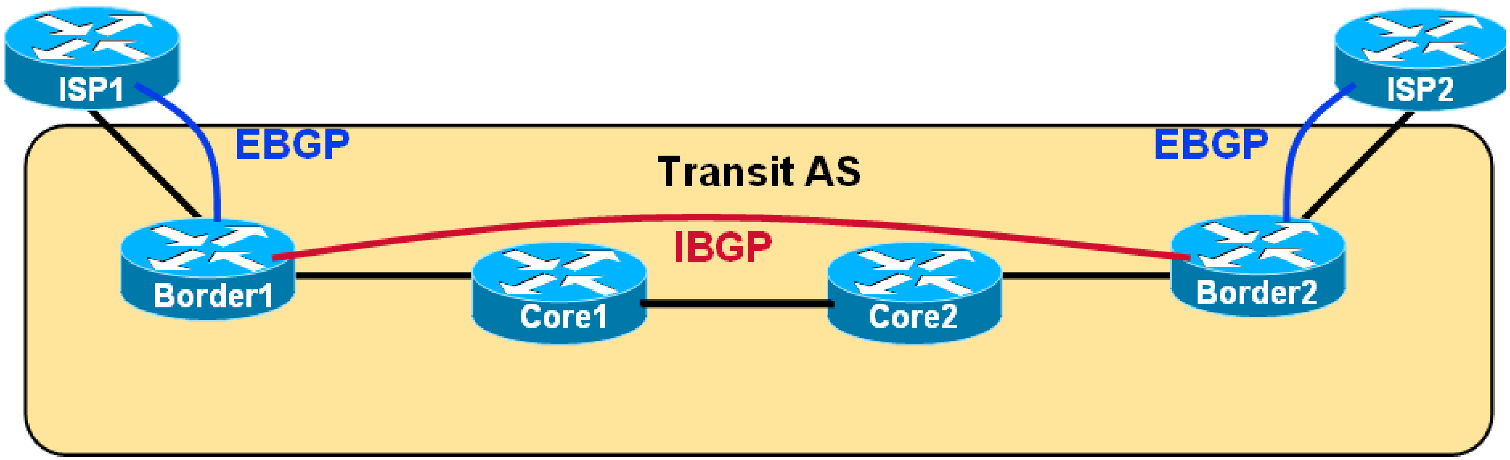

- Only border routers are required to run BGP.
- Core routers run an IGP to learn about BGP next-hop addresses.
- Core routers run LDP or TDP to learn about labels for next-hop addresses.

MPLS based Transit AS Building FIB and LFIB

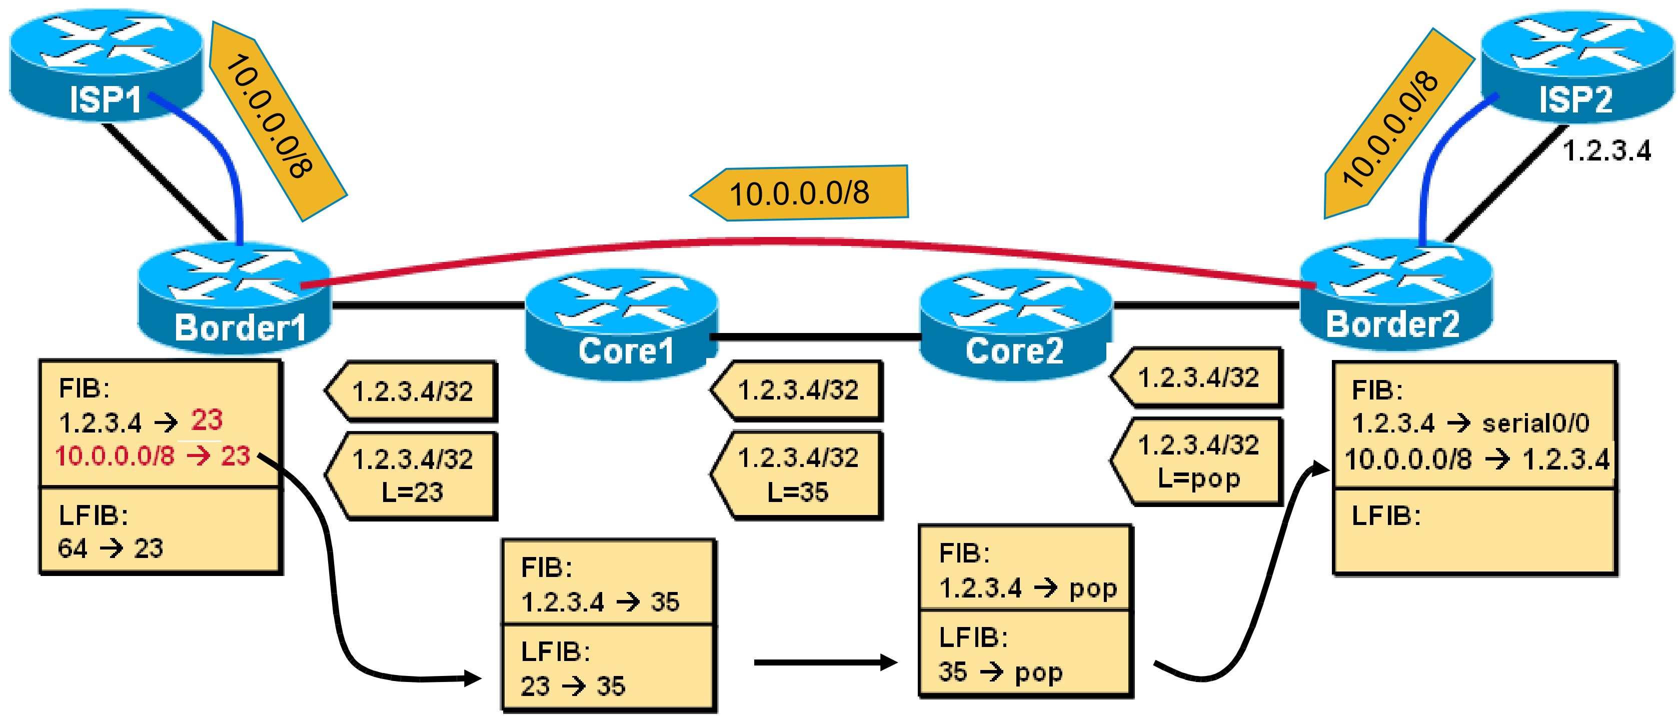

All routers are capable of forwarding packets to external destinations

Border (edge) routers label and forward IP packets.

Core routers forward labeled packets.

Benefit of MPLS-Based Transit AS

- Simplified BGP topology (only AS edge routers are required to run BGP with full Internet routing).
- Core routers do not require a lot of memory 250 000 networks may require more than 50 MB of memory for the BGP table, IP routing table, and CEF's FIB table and distributed FIB tables).
- Changes in the Internet do not impact core routers.
- Allows private addresses (RFC 1918) to be used in the core if TTL propagation is disabled (traceroute across the AS will not show any private addresses).