



# MPLS in communication networks.

«Computer networks and telecommunications» (Additional chapters).



# Icons and Symbols




Router



Workgroup Switch



Edge Label  
Switch  
Router



Network  
Cloud,



Line: Ethernet



Line: Serial




# Foundations of Traditional IP Routing

- Routing protocols are used to distribute Layer 3 routing information.
- Forwarding decision is made based on:
  - Packet header
  - Local routing table
- Routing lookups are independently performed at every hop.



# Traditional IP Routing



- Every router may need full Internet routing information.
- Destination-based routing lookup is needed on every hop.

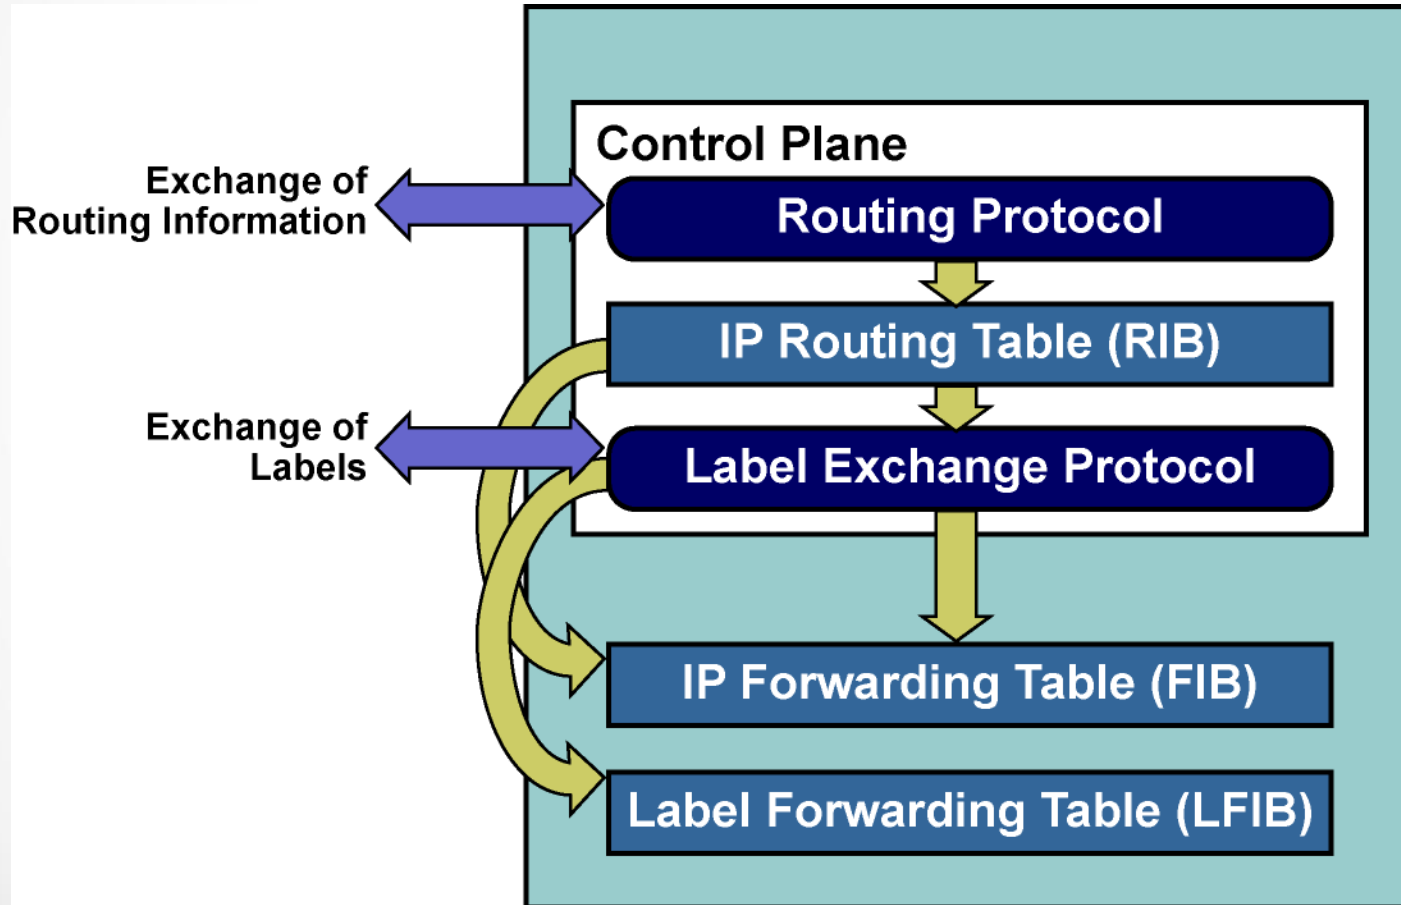


# Benefits of MPLS

- MPLS supports multiple applications including:
  - Unicast and multicast IP routing
  - VPN
  - TE
  - QoS
- MPLS decreases forwarding overhead on core routers.
- MPLS can support forwarding of non-IP protocols.




# MPLS Architecture: Control Plane






# MPLS Architecture: Data Plane





# MPLS Devices: LSRs

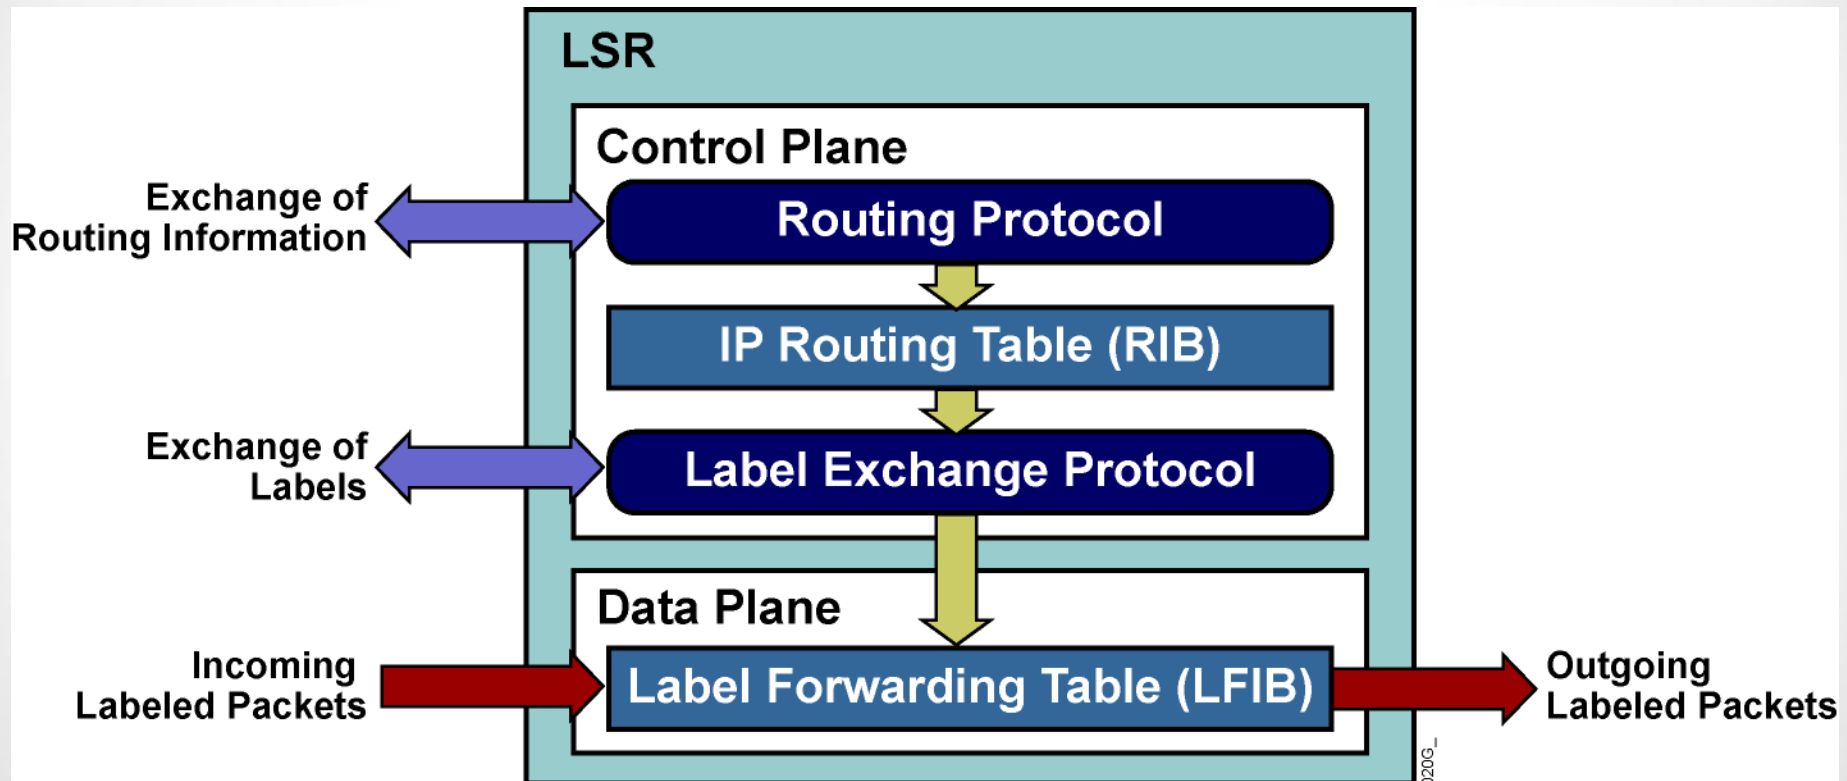


- The **LSR** forwards labeled packets in the MPLS domain.
- The **edge LSR** forwards labeled packets in the MPLS domain, and it forwards IP packets into and out of the MPLS domain.



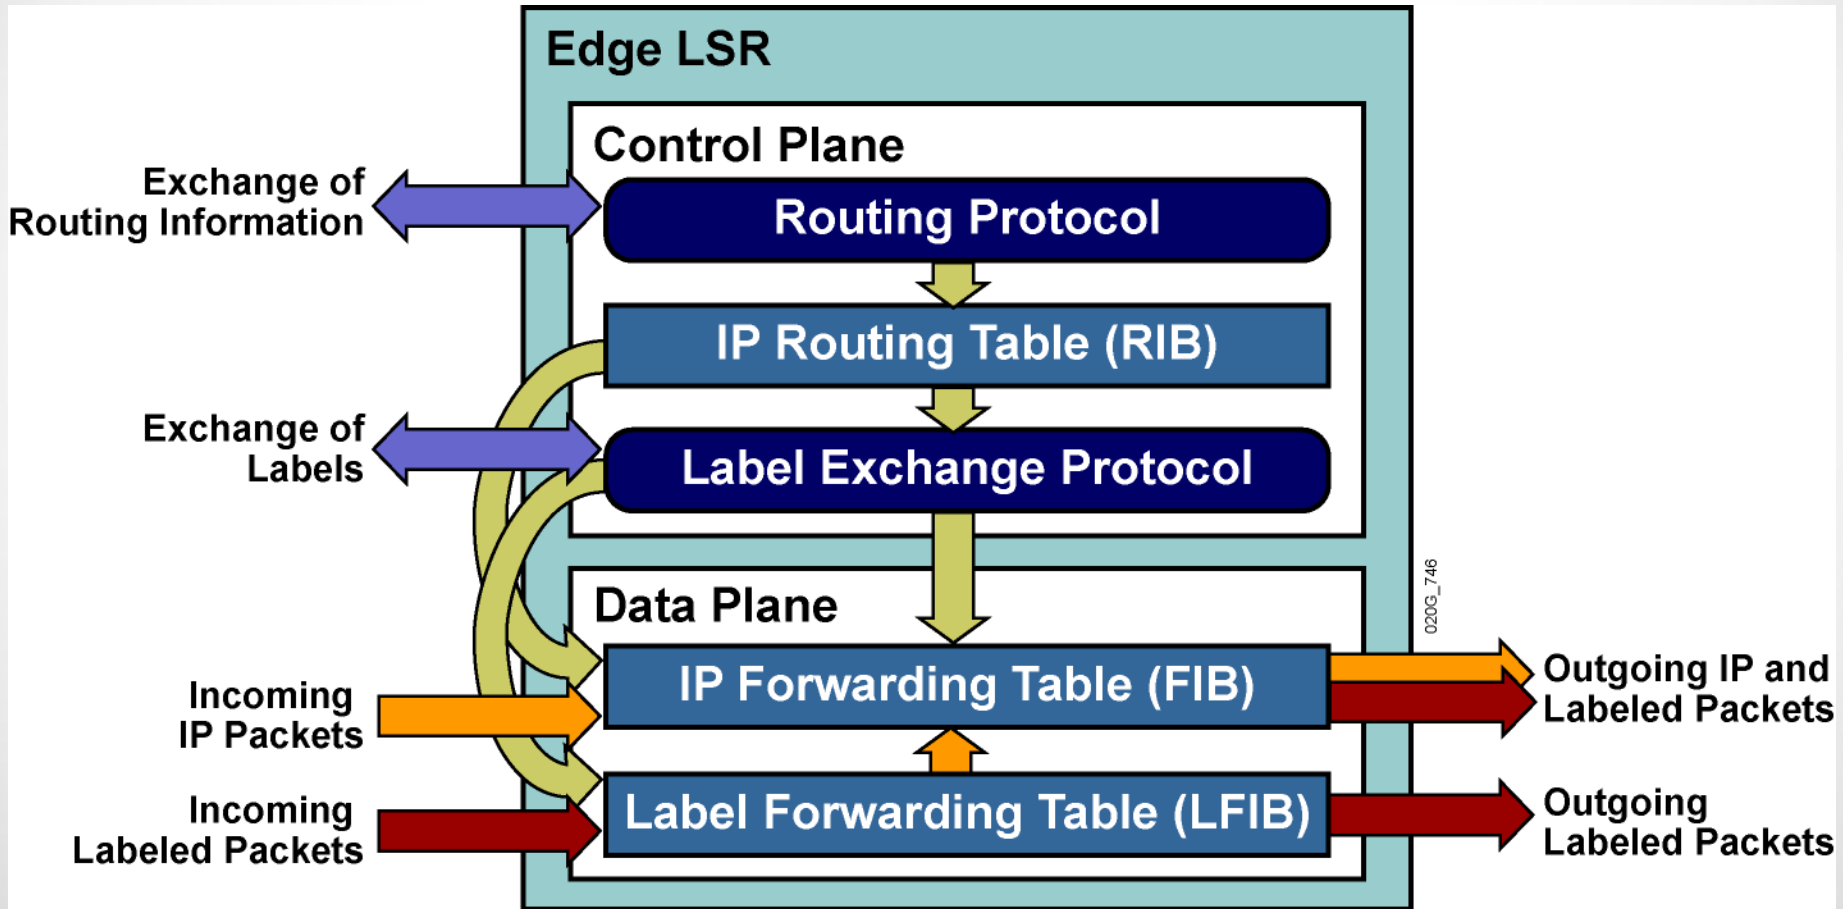


# Label Switch Routers: Architecture of LSRs



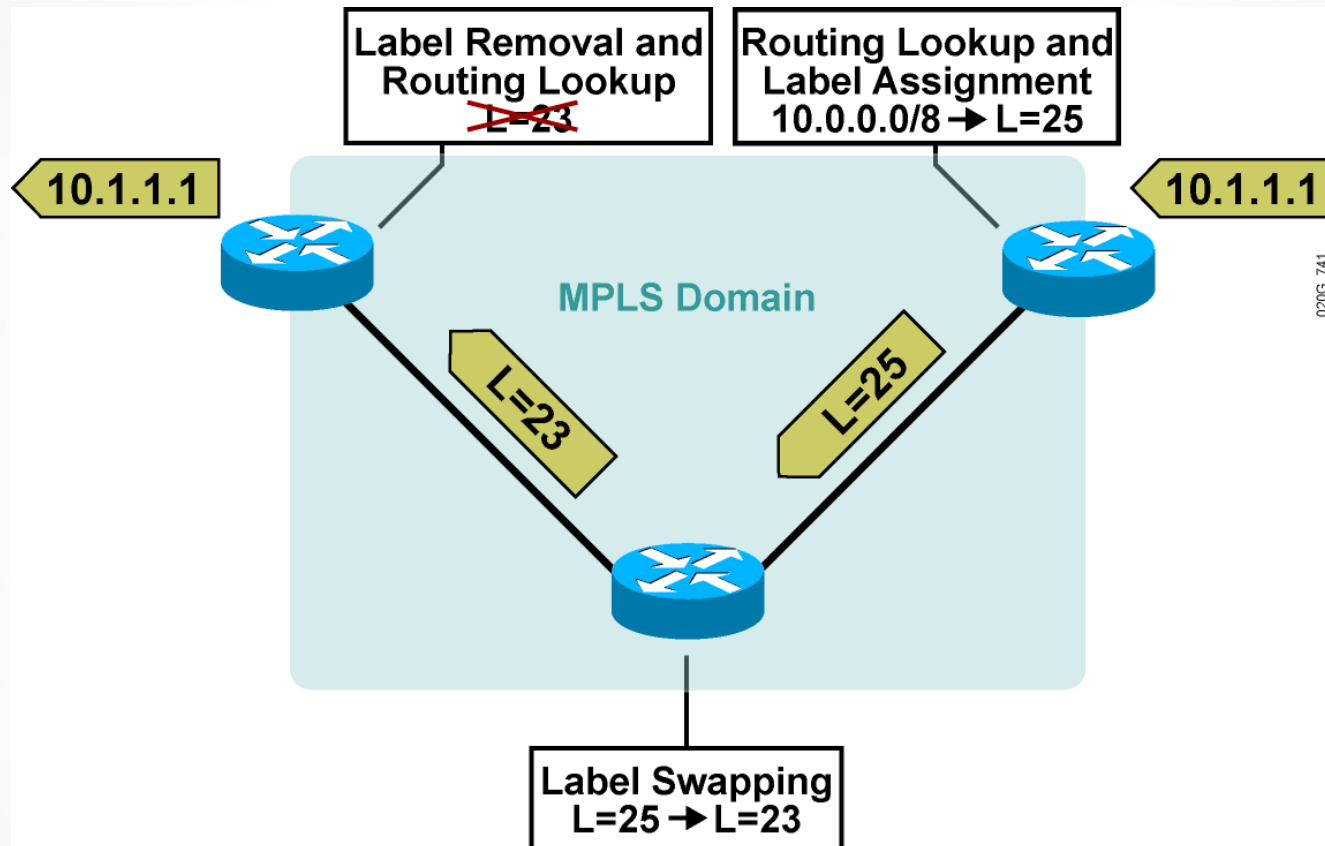


# LSRs: Architecture of Edge LSRs





# Basic MPLS Example



- MPLS core routers swap labels and forward packets based on simple label lookups.
- MPLS edge routers also perform a routing table lookup, and add or remove labels.



# MPLS Concepts

Introducing MPLS Labels and  
Label Stacks



# MPLS Labels

- Labels are 4 byte identifiers used for forwarding decisions
- Define the destination and services for a packet
- Identify a **forwarding equivalence class (FEC)**
- Have local significance
  - Each LSR independently maps a label to an FEC in a **label binding**.
  - Label bindings are exchanged between LSRs.



# FEC and MPLS Forwarding

- An FEC is a group of packets forwarded:
  - In the same manner
  - Over the same path
  - With the same forwarding treatment
- MPLS packet forwarding consists of:
  - Assigning a packet to a specific FEC
  - Determining the next hop of each FEC
- MPLS forwarding is connection-oriented.





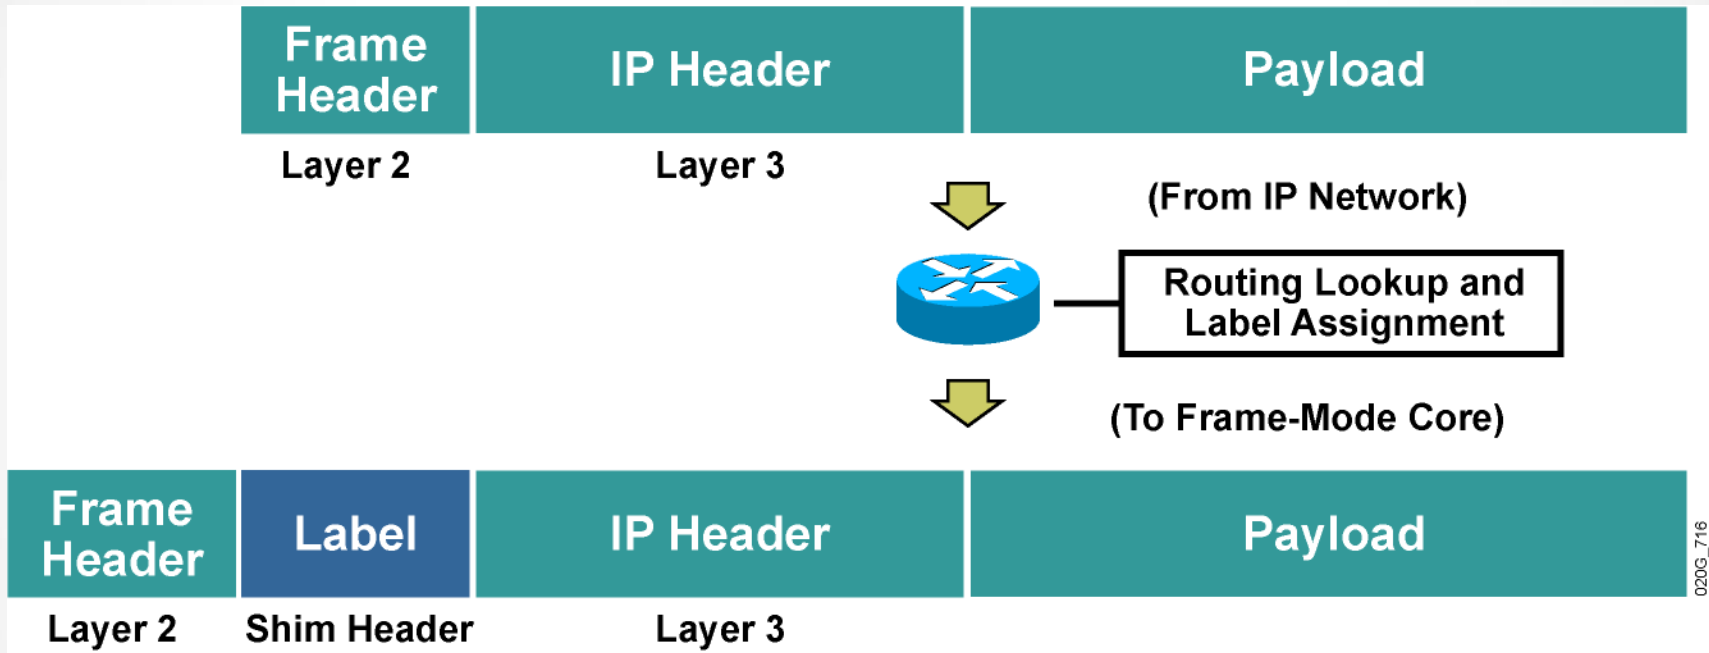
# MPLS Labels

- MPLS technology is intended to be used anywhere regardless of Layer 1 media and Layer 2 encapsulation.
- **Frame-mode MPLS** is MPLS over a frame-based Layer 2 encapsulation
  - The label is inserted between the Layer 2 and Layer 3 headers.
- **Cell-mode MPLS** is MPLS over ATM.
  - The fields in the ATM header are used as the label.





# MPLS Labels: Frame-Mode MPLS



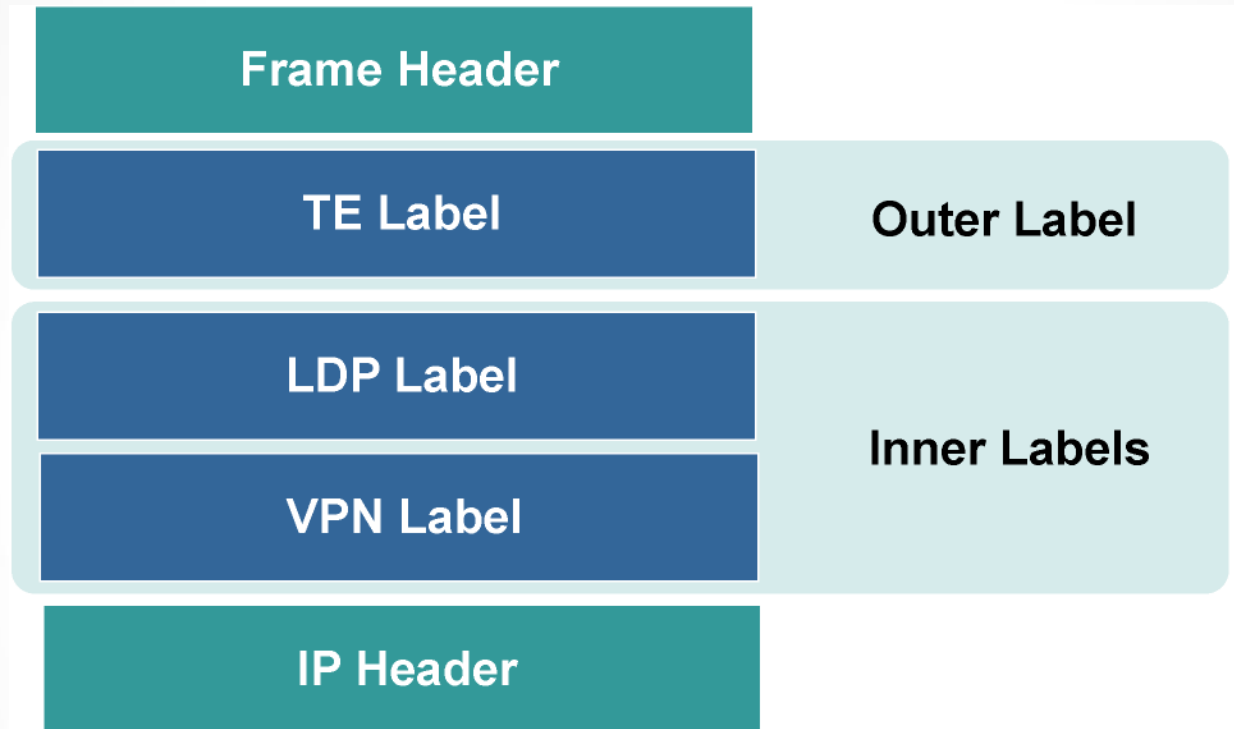


# MPLS Label Stack

- Usually only one label is assigned to a packet, but multiple labels in a **label stack** are supported.
- These scenarios may produce more than one label:
  - MPLS VPNs (two labels): The top label points to the egress router, 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).



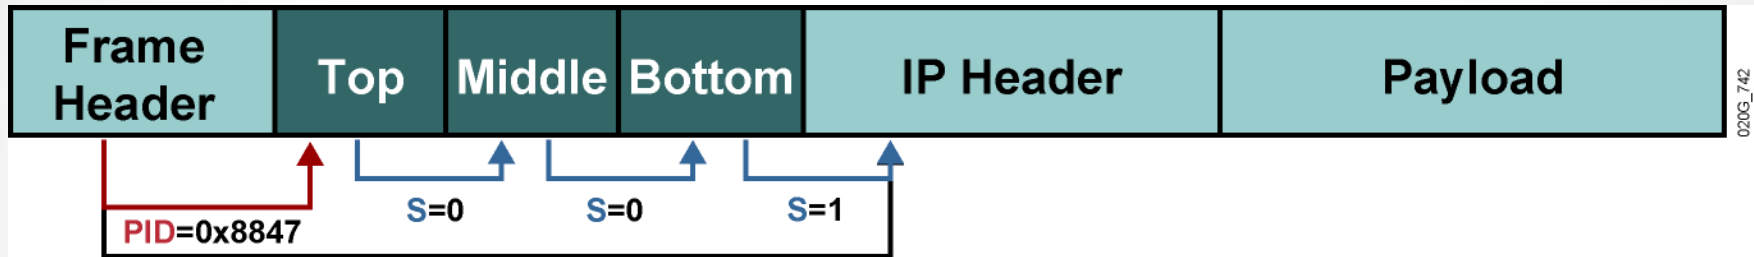
# Example: MPLS Label Stack



- The outer label is used for switching the packet in the MPLS network (points to the TE destination).
- Inner labels are used to separate packets at egress points (points to egress router and identifies VPN).



# Example: MPLS Label Stack Format



- The **PID** in a Layer 2 header specifies that the payload starts with a label (labels) followed by an IP header.
- The **bottom-of-stack bit** indicates whether the label is the last label in the stack.
- The receiving router uses the top label only.

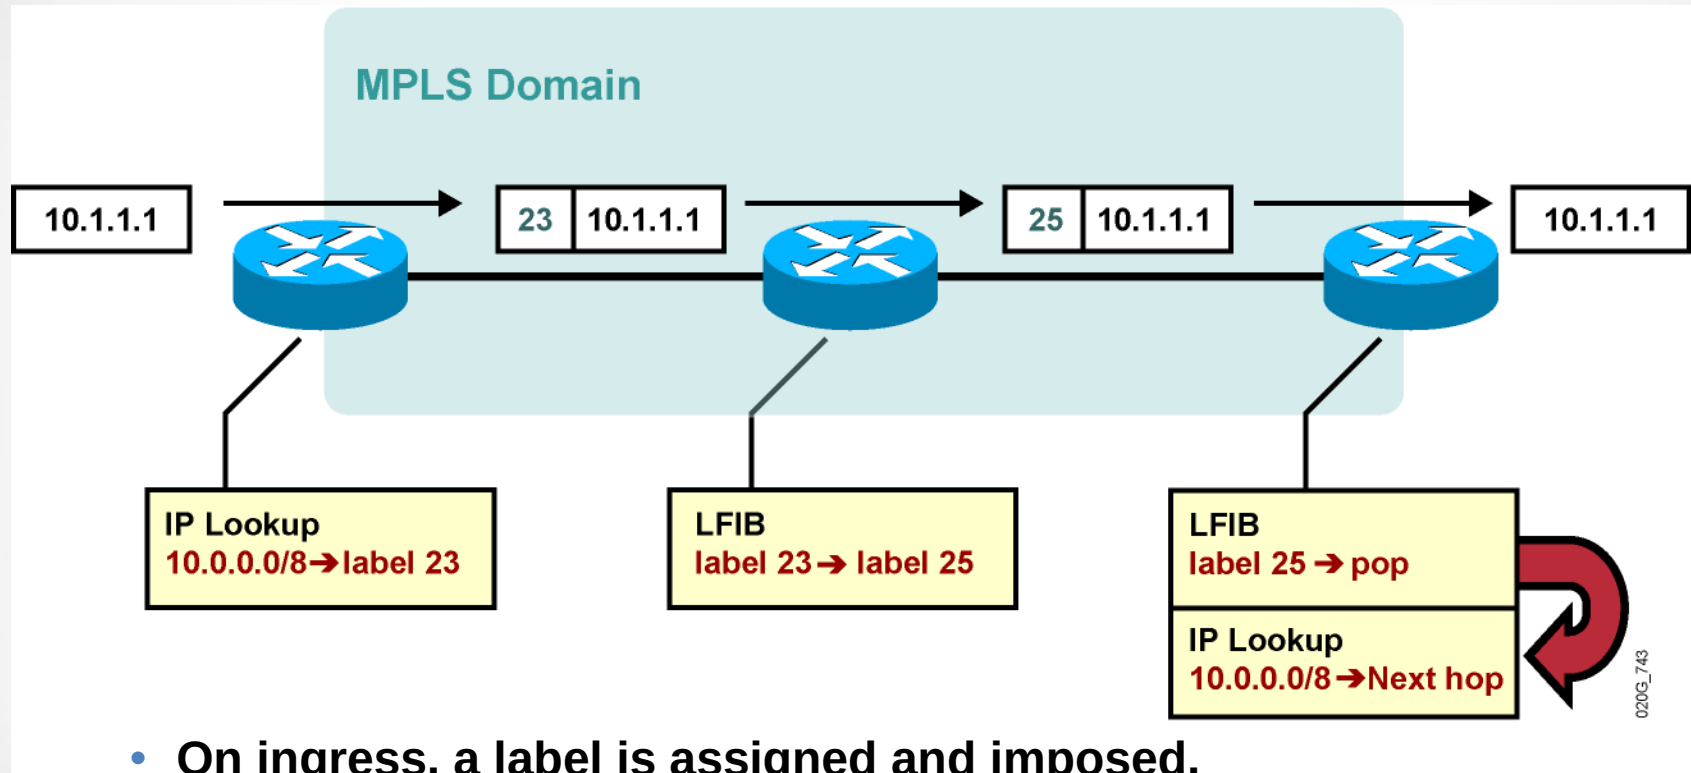


# MPLS Label Operations

- An LSR can perform these functions:
  - Insert (impose or push) a label or a stack of labels on ingress edge LSR
  - Swap a label with a next-hop label or a stack of labels in the core
  - Remove (pop) a label on egress edge LSR



# MPLS Label Operations: Frame Mode



- On ingress, a label is assigned and imposed.
- 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.



# Identifying MPLS Applications



# MPLS Applications

- MPLS is already used in many different applications:
  - Unicast IP routing
  - Multicast IP routing
  - MPLS TE
  - QoS
  - MPLS VPNs





# MPLS Unicast IP Routing

- Basic MPLS service supports unicast IP routing.
- MPLS unicast IP routing provides enhancement over traditional IP routing.
  - The ability to use labels for packet forwarding:
    - Label-based forwarding provides greater efficiency.
    - ***The FEC corresponds to a destination address stored in the IP routing table.***
    - Labels support connection-oriented services.
  - The capability to carry a stack of labels assigned to a packet:
    - Label stacks allow implementation of enhanced applications.



# MPLS Multicast IP Routing

- MPLS can also support multicast IP routing:
  - A dedicated protocol is not needed to support multicast traffic across an MPLS domain.
  - The FEC is equal to a destination multicast address stored in the multicast routing table.



# MPLS VPNs

- MPLS VPNs are highly scalable and support IP services such as:
  - Multicast
  - Quality of QoS
  - Telephony support within a VPN
  - Centralized services including content and web hosting to a VPN
- **Networks are learned via an IGP from a customer or via BGP from other MPLS backbone routers.**
- Labels are propagated via MP-BGP. Two labels are used:
  - The top label points to the egress router.
  - The second label identifies the outgoing interface on the egress router or a routing table where a routing lookup is performed.
- FEC is equivalent to a VPN site descriptor or VPN routing table.



# MPLS TE

- MPLS TE supports constraints-based routing
- MPLS TE enables the network administrator to
  - Control traffic flow in the network
  - Reduce congestion in the network
  - Make best use of network resources
- MPLS TE requires OSPF or ISIS with extensions to hold the entire network topology in their databases.
- OSPF and IS-IS should also have some additional information about network resources and constraints.
- **RSVP** is used to establish TE tunnels and to propagate labels.

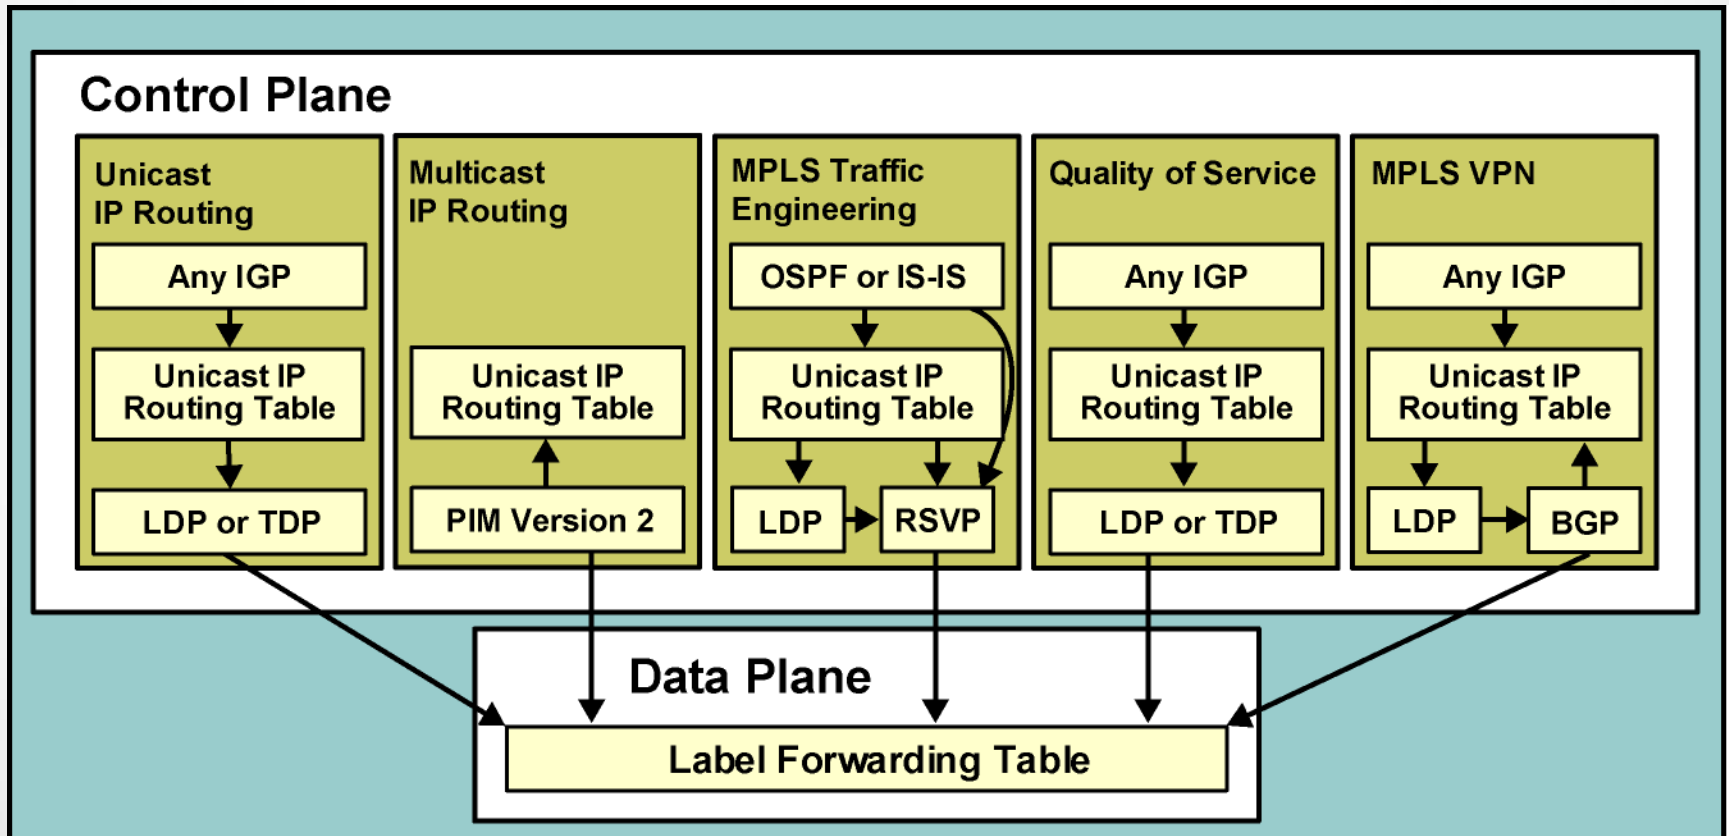


# MPLS QoS

- MPLS QoS provides differentiated types of service across an MPLS network.
- MPLS QoS offers:
  - Packet classification
  - Congestion avoidance
  - Congestion management.
- MPLS QoS is an extension to unicast IP routing that provides differentiated services.
- Extensions to **LDP** are used to propagate different labels for different classes.
- **In this case the FEC** is a combination of a destination network and a class of service.



# Interactions Between MPLS Applications





# Label Assignment and Distribution

Discovering LDP Neighbors



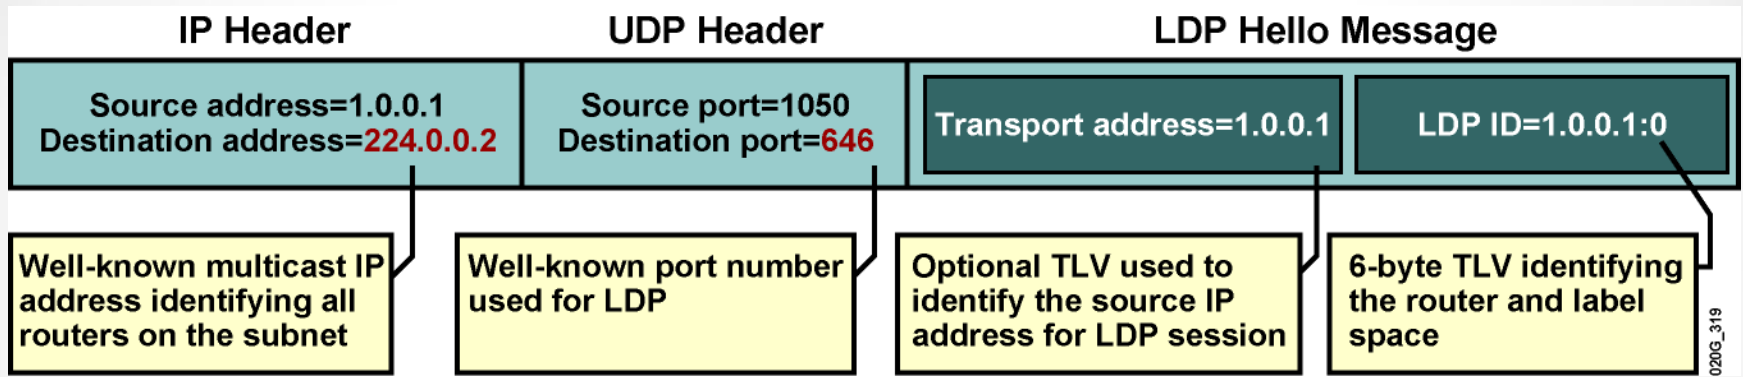
# LDP Neighbor Session Establishment

- LDP establishes a session in two steps:
  - Hello messages are periodically sent on all MPLS-enabled interfaces.
  - MPLS-enabled routers respond to received hello messages by attempting to establish a session with the source of the hello messages.
- **LDP link hello** message is a UDP packet sent to the “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.






# LDP Link Hello Message



- Hello messages are sent to all routers reachable through an interface.
- LDP uses well-known port number 646 with UDP for hello messages.
- A 6-byte LDP identifier (TLV) identifies the router (first 4 bytes) and label space (last 2 bytes).
- The source address used for an LDP session can be set by adding the transport address TLV to the hello message.




# Label Space: Per-Platform



- The label forwarding information base (LFIB) on an LSR does not contain an incoming interface.
- The same label can be used on any interface and is announced to all adjacent LSRs.
- The label is announced to adjacent LSRs only once and can be used on any link.
- Per-platform label space is less secure than per-interface label space.



# Negotiating Label Space

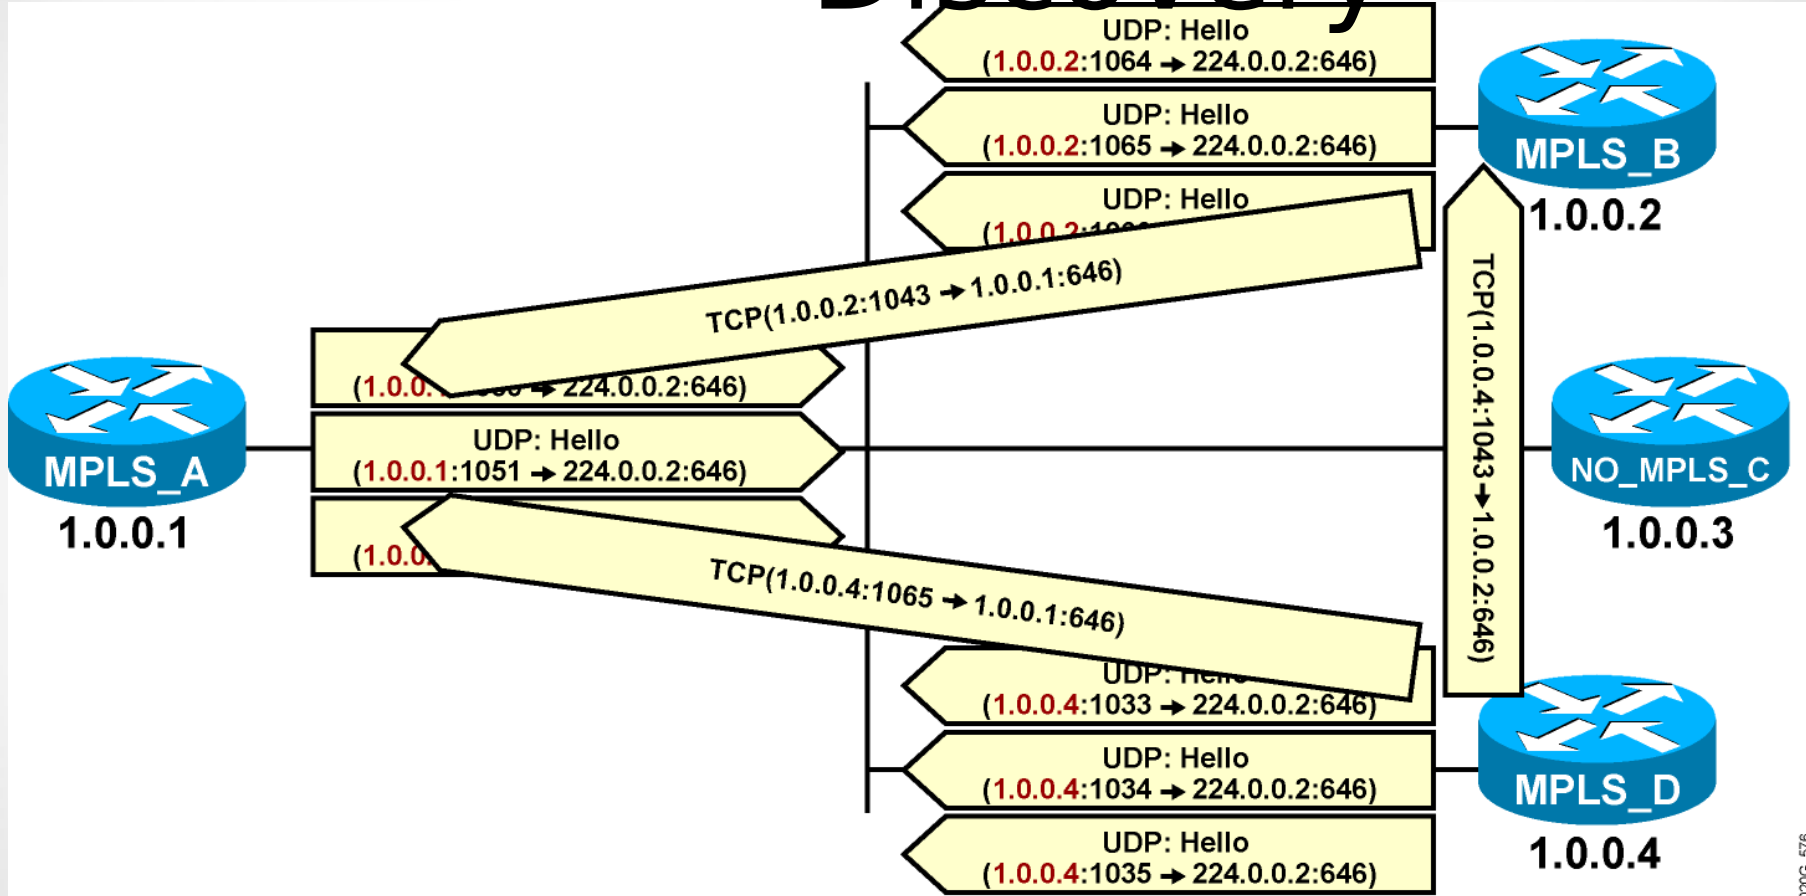


020G\_781

- 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



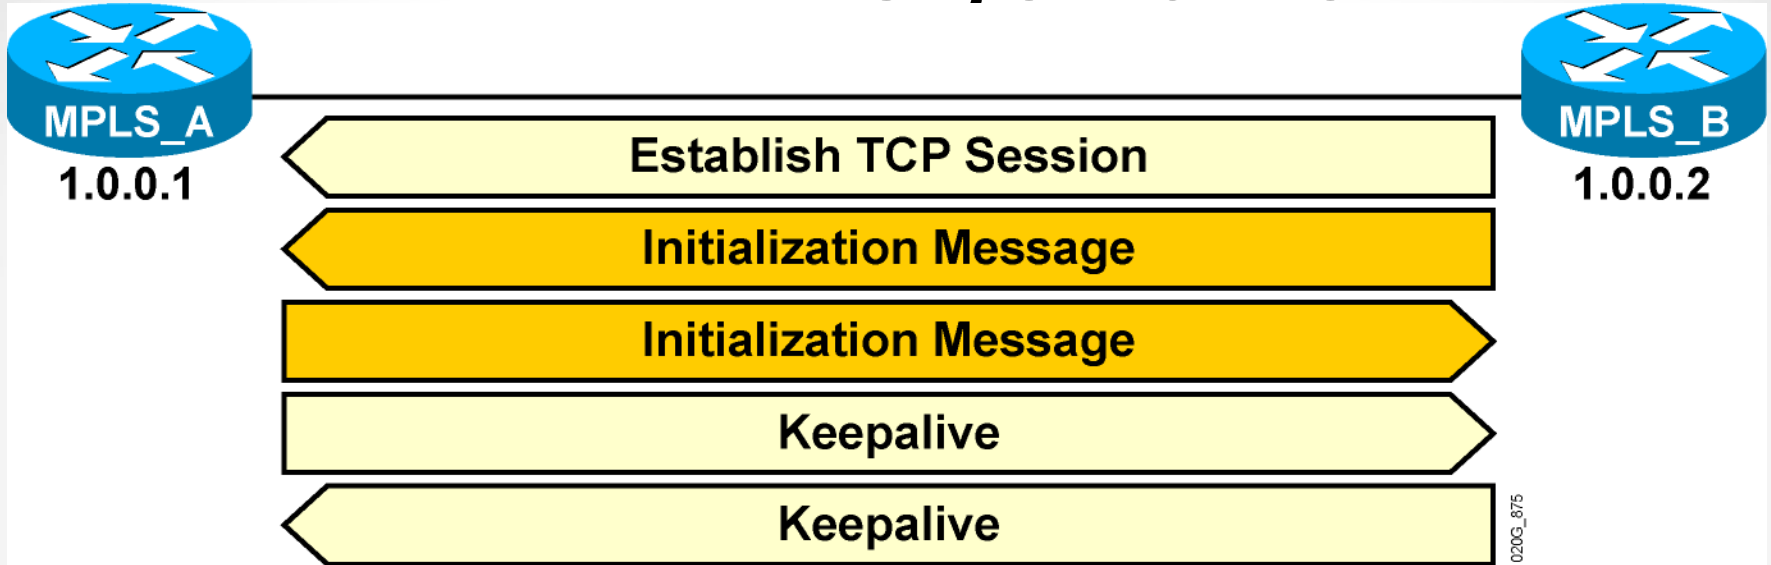
# LDP Neighbor Discovery



- An LDP session is established from the router with the higher IP address.



# LDP Session Negotiation



- Peers first exchange initialization messages.
- The session is ready to exchange label mappings after receiving the first keepalive.



# LDP packet examples

Filter: ldp

No.	Time	Source	Destination	Protocol	Length	Info
468	572.653000	192.168.1.11	224.0.0.2	LDP	76	Hello Message
469	572.673000	192.168.1.1	224.0.0.2	LDP	76	Hello Message
470	572.693000	192.168.1.11	224.0.0.2	LDP	76	Hello Message
474	572.743000	2.2.2.2	1.1.1.1	LDP	90	Initialization Message
475	572.753000	1.1.1.1	2.2.2.2	LDP	98	Initialization Message Keep Alive Message
476	572.783000	2.2.2.2	1.1.1.1	LDP	98	Keep Alive Message Address Message
477	572.793000	1.1.1.1	2.2.2.2	LDP	86	Address Message
478	573.104000	2.2.2.2	1.1.1.1	LDP	98	Label Mapping Message

Frame 470: 76 bytes on wire (608 bits), 76 bytes captured (608 bits) on interface 0

- Ethernet II, Src: HuaweiTe\_a3:23:f0 (54:89:98:a3:23:f0), Dst: IPv4mcast\_00:00:02 (01:00:5e:00:00:02)
- Internet Protocol Version 4, Src: 192.168.1.11 (192.168.1.11), Dst: 224.0.0.2 (224.0.0.2)
- User Datagram Protocol, Src Port: ldp (646), Dst Port: ldp (646)
  - Source port: ldp (646)
  - Destination port: ldp (646)
  - Length: 42
  - Checksum: 0x46e4 [validation disabled]
- Label Distribution Protocol
  - Version: 1
  - PDU Length: 30
  - LSR ID: 2.2.2.2 (2.2.2.2)
  - Label space ID: 0
- Hello Message
  - 0... .... = U bit: Unknown bit not set
  - Message Type: Hello Message (0x100)
  - Message Length: 20
  - Message ID: 0x000000a1
- Common Hello Parameters TLV
  - 00.. .... = TLV Unknown bits: Known TLV, do not Forward (0x00)
  - TLV Type: Common Hello Parameters TLV (0x400)
  - TLV Length: 4
  - Hold Time: 15
  - 0... .... = Targeted Hello: Link Hello
  - .0.. .... = Hello Requested: source does not request periodic hellos
  - ..0. .... = GSM Flag: Not set
  - ...0 0000 0000 0000 = Reserved: 0x0000
- IPv4 Transport Address TLV
  - 00.. .... = TLV Unknown bits: Known TLV, do not Forward (0x00)
  - TLV Type: IPv4 Transport Address TLV (0x401)
  - TLV Length: 4
  - IPv4 Transport Address: 2.2.2.2 (2.2.2.2)



# LDP packet examples

No.	Time	Source	Destination	Protocol	Length	Info
468	572.653000	192.168.1.11	224.0.0.2	LDP	76	Hello Message
469	572.673000	192.168.1.1	224.0.0.2	LDP	76	Hello Message
470	572.693000	192.168.1.11	224.0.0.2	LDP	76	Hello Message
474	572.743000	2.2.2.2	1.1.1.1	LDP	90	Initialization Message
475	572.753000	1.1.1.1	2.2.2.2	LDP	98	Initialization Message Keep Alive Message
476	572.783000	2.2.2.2	1.1.1.1	LDP	98	Keep Alive Message Address Message
477	572.793000	1.1.1.1	2.2.2.2	LDP	86	Address Message
478	573.104000	2.2.2.2	1.1.1.1	LDP	98	Label Mapping Message

Frame 474: 90 bytes on wire (720 bits), 90 bytes captured (720 bits) on interface 0

- Ethernet II, Src: HuaweiTe\_a3:23:f0 (54:89:98:a3:23:f0), Dst: HuaweiTe\_13:27:82 (54:89:98:13:27:82)
- Internet Protocol Version 4, Src: 2.2.2.2 (2.2.2.2), Dst: 1.1.1.1 (1.1.1.1)
- Transmission Control Protocol, Src Port: 61573 (61573), Dst Port: ldp (646), Seq: 1, Ack: 1, Len: 36
- Label Distribution Protocol
  - Version: 1
  - PDU Length: 32
  - LSR ID: 2.2.2.2 (2.2.2.2)
  - Label Space ID: 0
  - Initialization Message
    - 0... .... = U bit: Unknown bit not set
    - Message Type: Initialization Message (0x200)
    - Message Length: 22
    - Message ID: 0x000000a2
  - Common Session Parameters TLV
    - 00.. .... = TLV Unknown bits: Known TLV, do not Forward (0x00)
    - TLV Type: Common Session Parameters TLV (0x500)
    - TLV Length: 14
  - Parameters
    - Session Protocol Version: 1
    - Session KeepAlive Time: 45
    - 0... .... = Session Label Advertisement Discipline: Downstream Unsolicited proposed
    - .0.. .... = Session Loop Detection: Loop Detection Disabled
    - Session Path Vector Limit: 0
    - Session Max PDU Length: 4096
    - Session Receiver LSR Identifier: 1.1.1.1 (1.1.1.1)
    - Session Receiver Label Space Identifier: 0



# LDP packet examples

No.	Time	Source	Destination	Protocol	Length	Info
468	572.653000	192.168.1.11	224.0.0.2	LDP	76	Hello Message
469	572.673000	192.168.1.1	224.0.0.2	LDP	76	Hello Message
470	572.693000	192.168.1.11	224.0.0.2	LDP	76	Hello Message
474	572.743000	2.2.2.2	1.1.1.1	LDP	90	Initialization Message
475	572.753000	1.1.1.1	2.2.2.2	LDP	98	Initialization Message Keep Alive Message
476	572.783000	2.2.2.2	1.1.1.1	LDP	98	Keep Alive Message Address Message
477	572.793000	1.1.1.1	2.2.2.2	LDP	86	Address Message
478	573.104000	2.2.2.2	1.1.1.1	LDP	98	Label Mapping Message

Frame 476: 98 bytes on wire (784 bits), 98 bytes captured (784 bits) on interface 0

- Ethernet II, Src: HuaweiTe\_a3:23:f0 (54:89:98:a3:23:f0), Dst: HuaweiTe\_13:27:82 (54:89:98:13:27:82)
- Internet Protocol Version 4, Src: 2.2.2.2 (2.2.2.2), Dst: 1.1.1.1 (1.1.1.1)
- Transmission Control Protocol, Src Port: 61573 (61573), Dst Port: ldp (646), Seq: 37, Ack: 45, Len: 44
- Label Distribution Protocol
  - Version: 1
  - PDU Length: 40
  - LSR ID: 2.2.2.2 (2.2.2.2)
  - Label space ID: 0
  - Keep Alive Message
    - 0... .. = U bit: Unknown bit not set
    - Message Type: Keep Alive Message (0x201)
    - Message Length: 4
    - Message ID: 0x000000a3
  - Address Message
    - 0... .. = U bit: Unknown bit not set
    - Message Type: Address Message (0x300)
    - Message Length: 22
    - Message ID: 0x000000a4
  - Address List TLV
    - 00.. .... = TLV Unknown bits: Known TLV, do not Forward (0x00)
    - TLV Type: Address List TLV (0x101)
    - TLV Length: 14
    - Address Family: IPv4 (1)
    - Addresses
      - Address 1: 192.168.1.11
      - Address 2: 2.2.2.2
      - Address 3: 192.168.0.2

```
0000 54 89 98 13 27 82 54 89 98 a3 23 f0 08 00 45 c0  T...'.T. ..#...E.
0010 00 54 00 b0 00 00 ff 06 b4 2e 02 02 02 02 01 01  .T.....
0020 01 01 f0 85 02 86 7f c6 43 92 95 4c 1f 0a 50 18  .....C..L..P.
0030 20 00 8c de 00 00 00 01 00 28 02 02 02 02 00 00  .....(.
0040 02 01 00 04 00 00 00 a3 03 00 00 16 00 00 00 a4  .....
0050 01 01 00 00 00 01 c0 38 01 0b 02 02 02 02 c0 38
```

File: "C:\DOCUME~1\ensp\LOCALS~1\Temp\wir... P... Profile: Default





# LDP packet examples

No.	Time	Source	Destination	Protocol	Length	Info
600	655.712000	192.168.1.1	224.0.0.2	LDP	76	Hello Message
602	658.817000	192.168.1.11	224.0.0.2	LDP	76	Hello Message
605	660.800000	192.168.1.1	224.0.0.2	LDP	76	Hello Message
610	662.532000	2.2.2.2	1.1.1.1	LDP	72	Keep Alive Message
612	662.943000	1.1.1.1	2.2.2.2	LDP	72	Keep Alive Message
613	663.003000	2.2.2.2	1.1.1.1	LDP	98	Label Mapping Message
614	663.063000	1.1.1.1	2.2.2.2	LDP	98	Label Mapping Message
616	663.754000	192.168.1.11	224.0.0.2	LDP	76	Hello Message

Ethernet II, Src: Huawei10\_a3:23:TU (54:89:98:a3:23:TU), Dst: Huawei10\_13:27:82 (54:89:98:13:27:82)

Internet Protocol Version 4, Src: 2.2.2.2 (2.2.2.2), Dst: 1.1.1.1 (1.1.1.1)

Transmission Control Protocol, Src Port: 61573 (61573), Dst Port: ldp (646), Seq: 277, Ack: 273, Len: 44

Label Distribution Protocol

- Version: 1
- PDU Length: 40
- LSR ID: 2.2.2.2 (2.2.2.2)
- Label space ID: 0
- Label Mapping Message
  - 0... .. = U bit: unknown bit not set
  - Message Type: Label Mapping Message (0x400)
  - Message Length: 30
  - Message ID: 0x000000d8
- Forwarding Equivalence Classes TLV
  - 00.. .... = TLV unknown bits: Known TLV, do not Forward (0x00)
  - TLV Type: Forwarding Equivalence Classes TLV (0x100)
  - TLV Length: 8
- FEC Elements
  - FEC Element 1
    - FEC Element Type: Prefix FEC (2)
    - FEC Element Address Type: IPv4 (1)
    - FEC Element Length: 32
    - Prefix: 3.3.3.3
- Generic Label TLV
  - 00.. .... = TLV unknown bits: Known TLV, do not Forward (0x00)
  - TLV Type: Generic Label TLV (0x200)
  - TLV Length: 4
  - Generic Label: 1025
- Unknown TLV type (0x0900)
  - 11.. .... = TLV unknown bits: Unknown TLV, do Forward (0x03)
  - TLV Type: Unknown TLV type (0x900)
  - TLV Length: 2
  - TLV Value: 05dc



# Label Assignment and Distribution

Introducing Typical Label  
Distribution in Frame-Mode  
MPLS

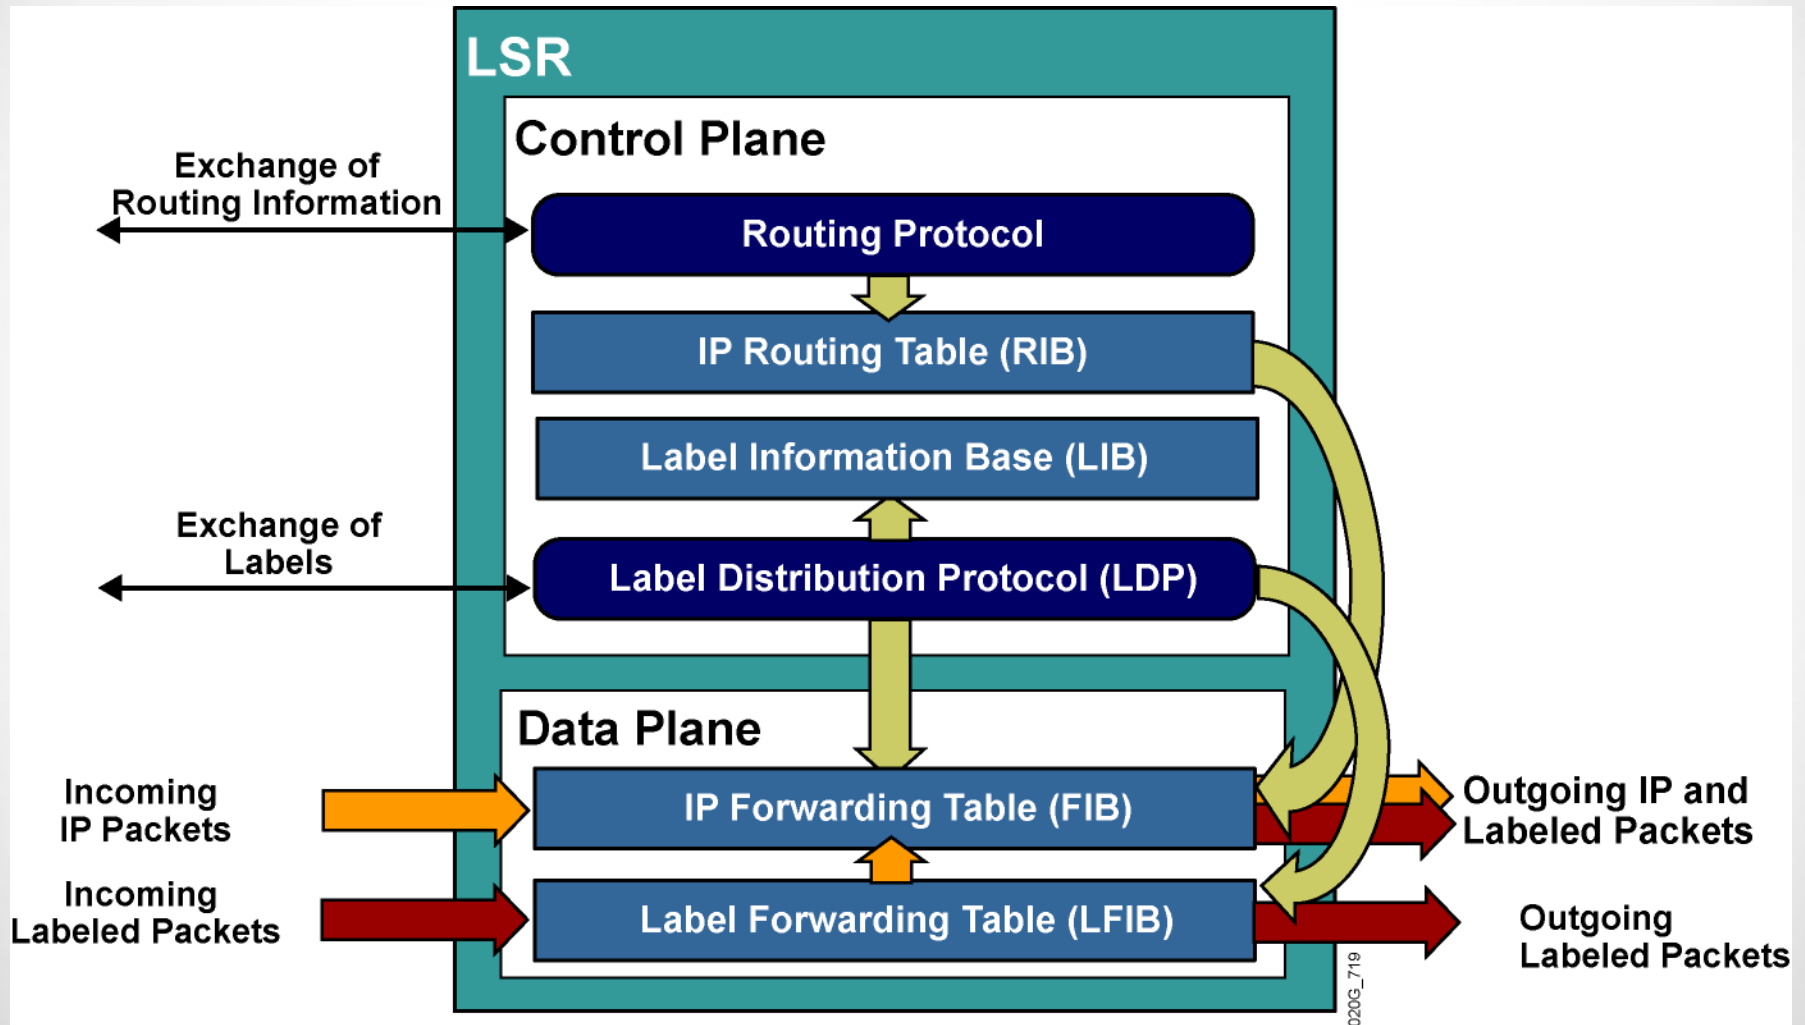


# MPLS Unicast IP Routing Architecture

- MPLS introduces a label field that is used for forwarding decisions.
- Although labels are locally significant, they have to be advertised to directly reachable peers.
  - One option would be to include this parameter in existing IP routing protocols.
  - The other option is to create a new protocol to exchange labels.
- The second option has been used because there are too many existing IP routing protocols that would have to be modified to carry labels.




# MPLS Unicast IP Routing Architecture (Cont.)



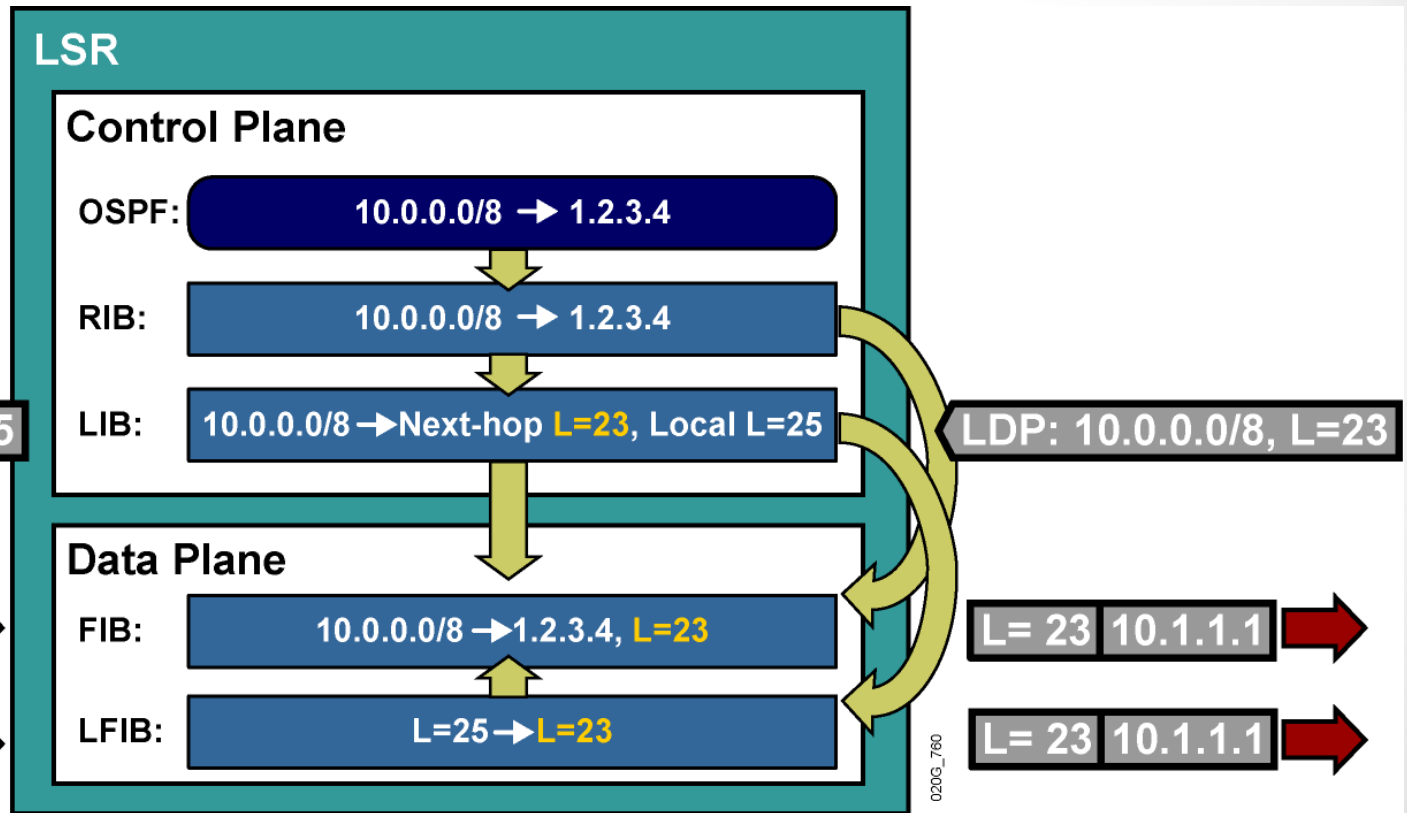


# MPLS Unicast IP Routing Architecture (Cont.)





# MPLS Unicast IP Routing Architecture (Cont.)



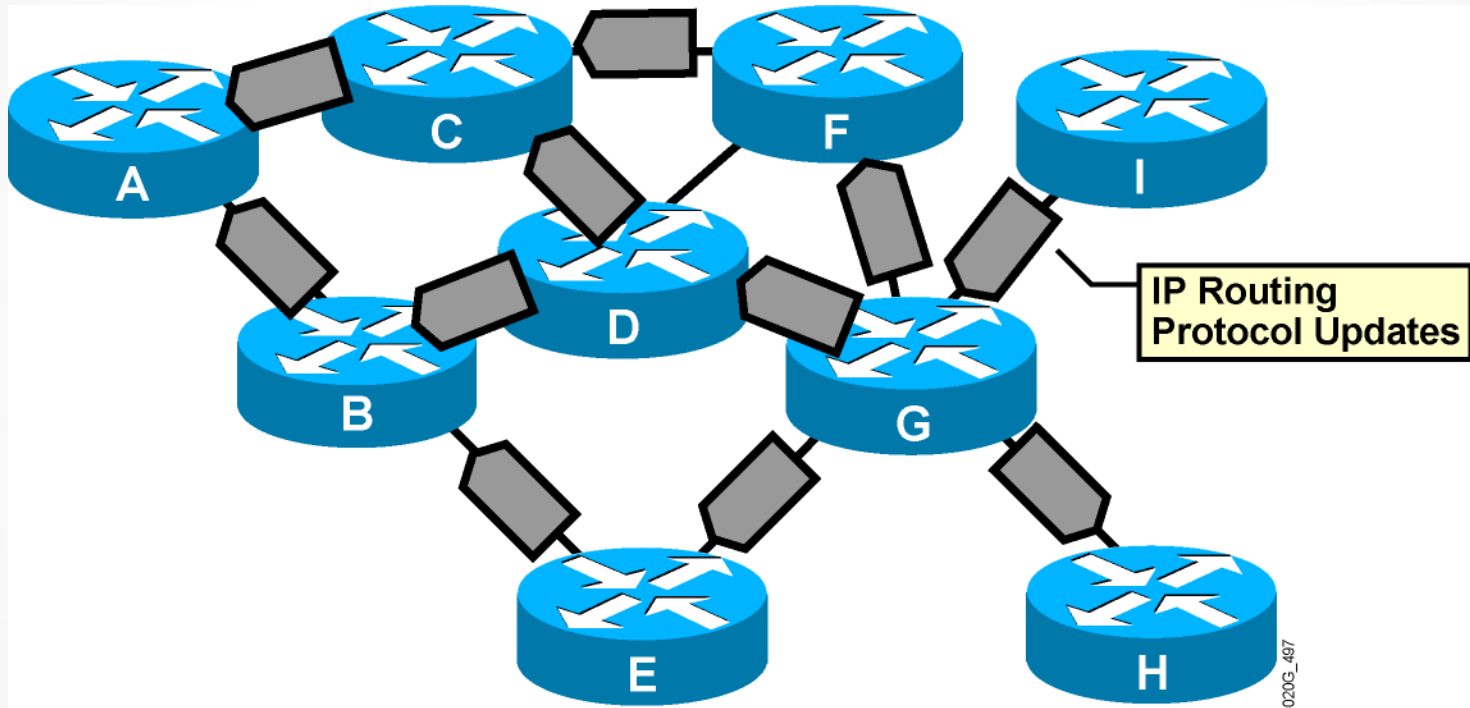


# Label-Switched Path

- An **LSP** is a sequence of LSRs that forwards labeled packets of a certain forwarding equivalence class.
  - MPLS unicast IP forwarding builds LSPs based on the output of IP routing protocols.
  - LDP advertises labels only for individual segments in the LSP.
- LSPs are **unidirectional**.
  - Return traffic uses a different LSP (usually the reverse path because most routing protocols provide symmetrical routing).
- An LSP can take a different path from the one chosen by an IP routing protocol (MPLS TE).



# LSP Building

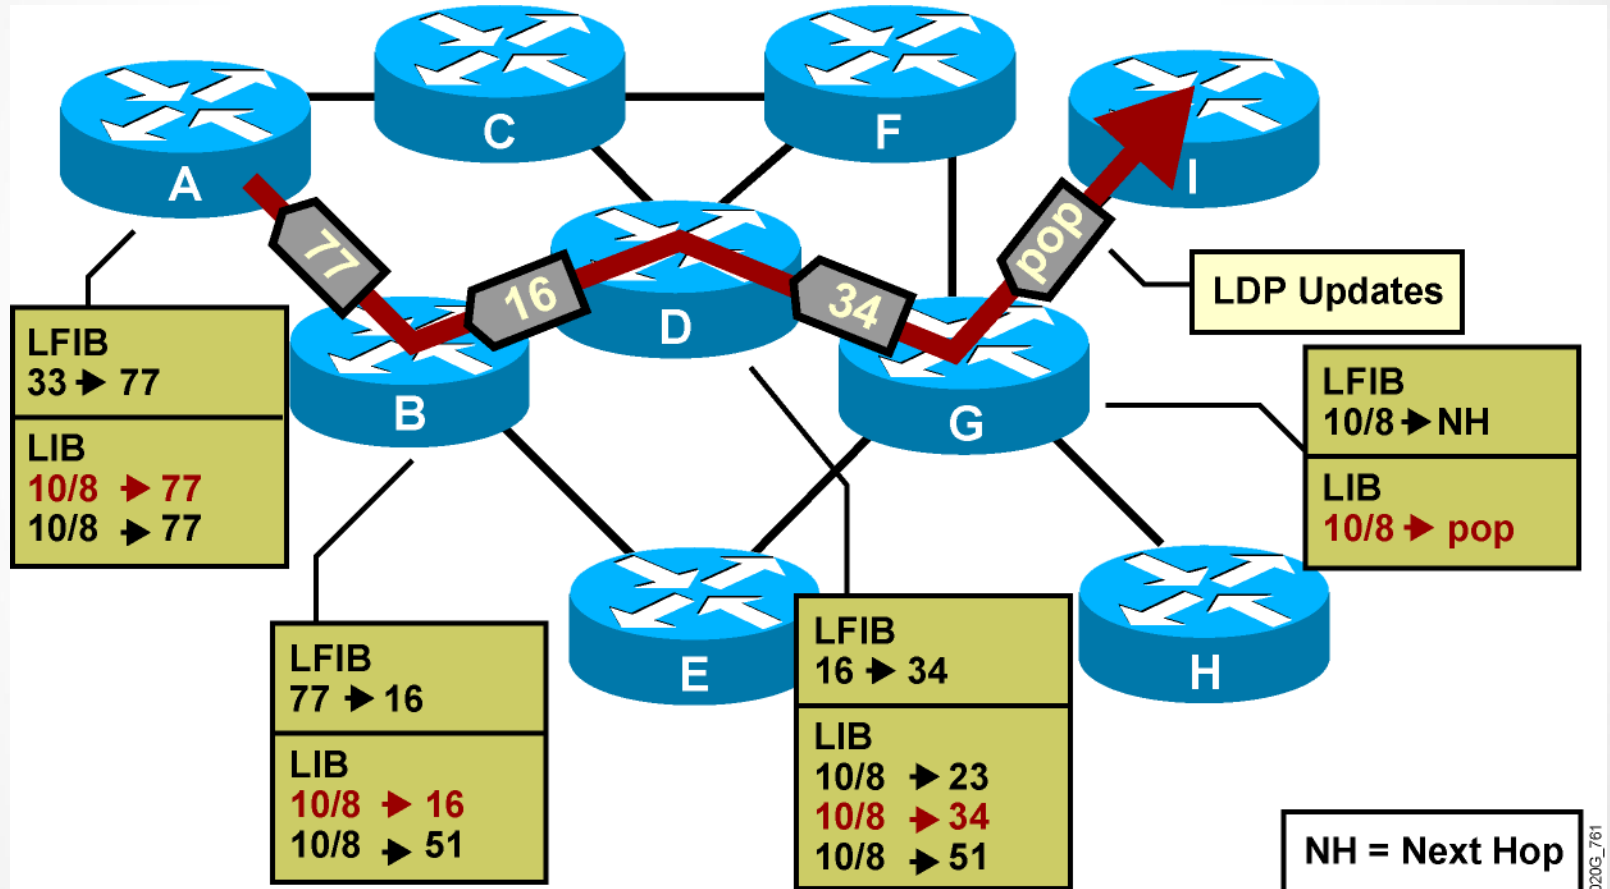


**The IP routing protocol determines the path.**






# LSP Building (Cont.)



LDP propagates labels to convert the path to an LSP.



# PHP: Before



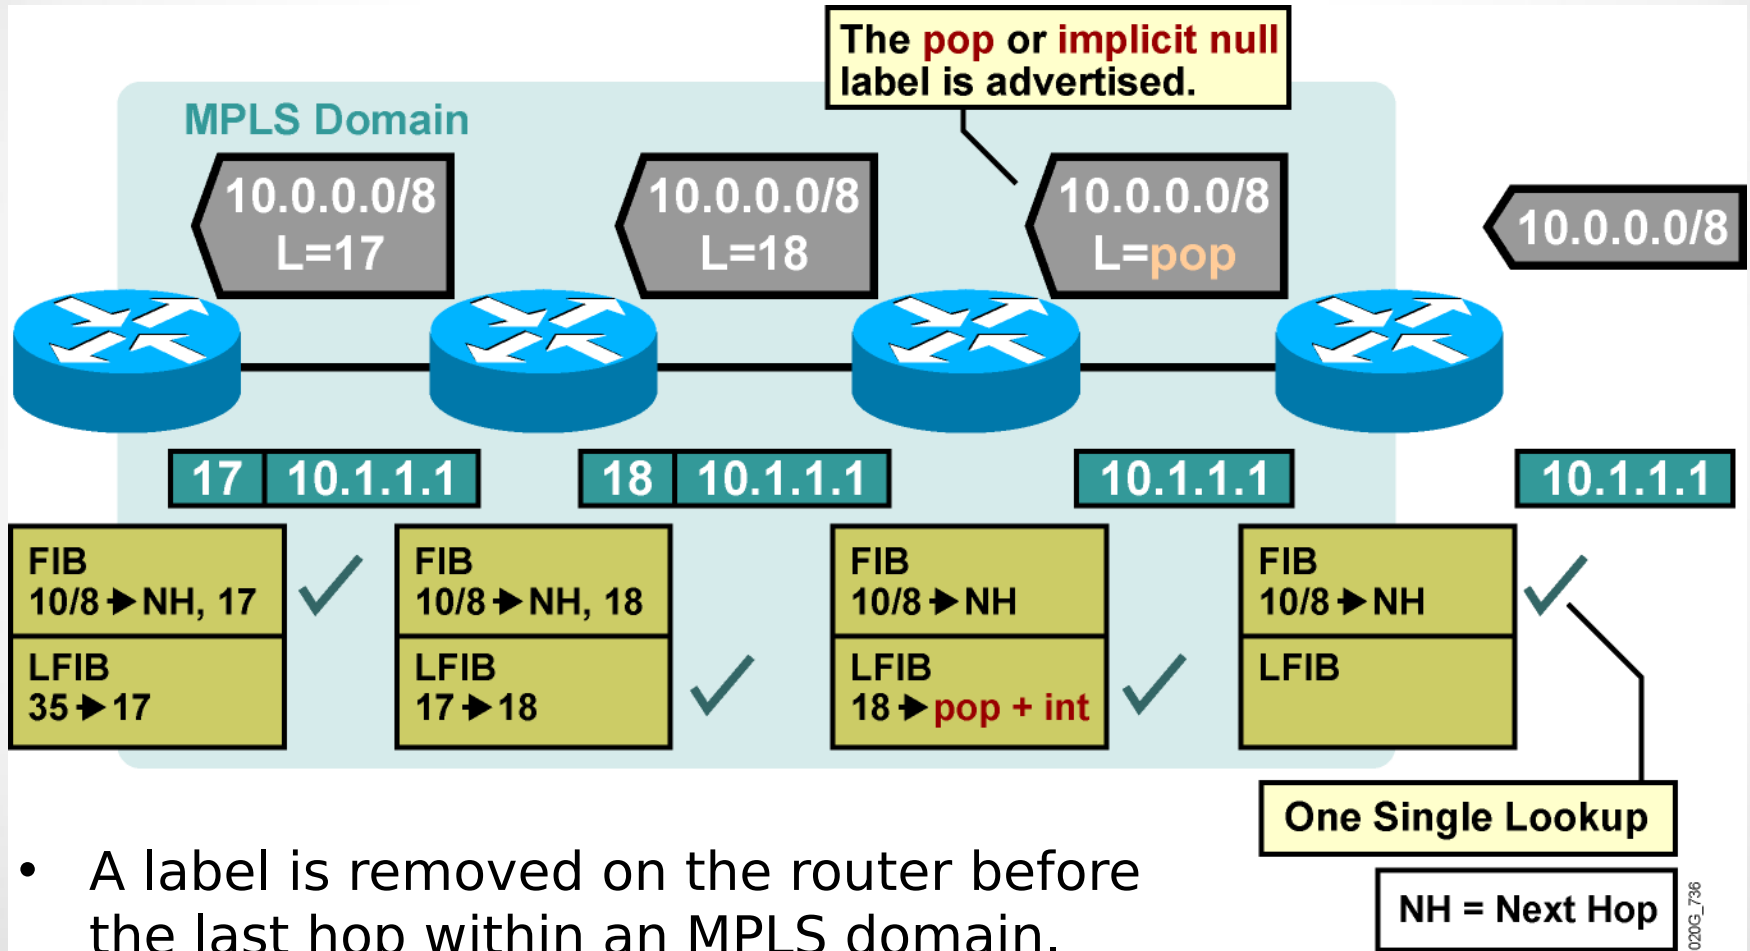
- Double lookup is not an optimal way of forwarding labeled packets.
- A label can be removed one hop earlier.

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

NH = Next Hop



# PHP: After



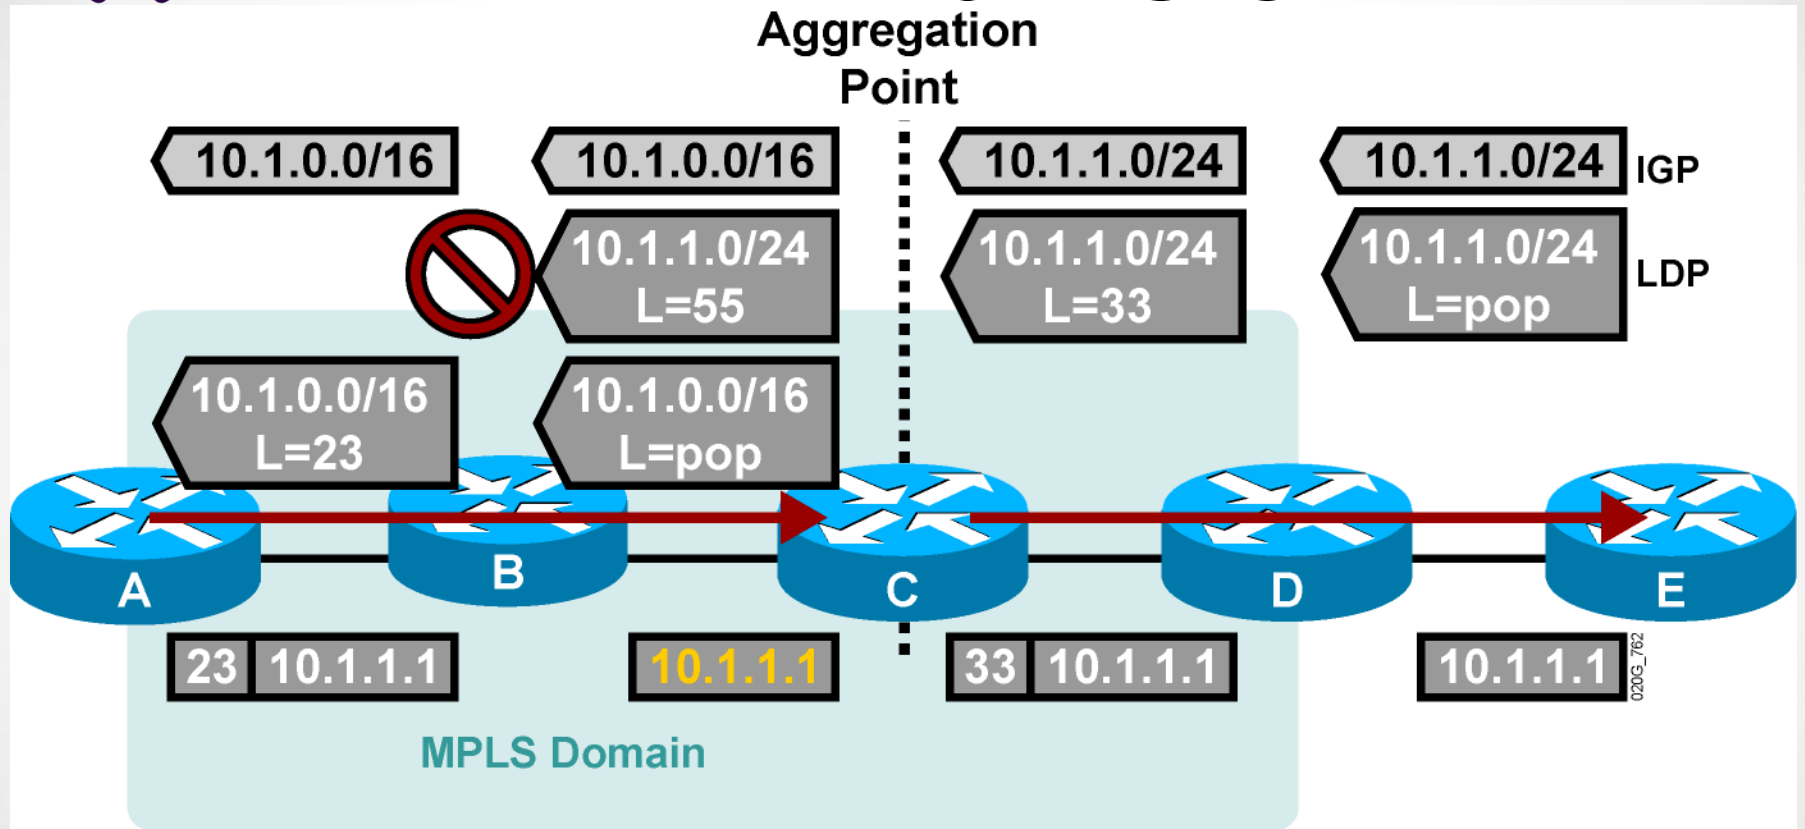


# PHP

- Penultimate hop popping optimizes MPLS performance (one less LFIB lookup).
- PHP does not work on ATM. (virtual path identifier/virtual channel identifier cannot be removed.)
- The **pop** or **implicit null** label uses a reserved value when being advertised to a neighbor.



# Impact of IP Aggregation on LSPs



- IP aggregation breaks an LSP into two segments.
- Router C is forwarding packets based on Layer 3 information.



# Impact of IP Aggregation on LSPs (Cont.)

- IP aggregation breaks an LSP into two segments.
- Aggregation should not be used where endtoend LSPs are required, such as with:
  - MPLS VPNs
  - MPLS TEs
  - MPLS-enabled ATM network
  - Transit BGP where core routers are not running BGP



## Label Allocation in a Frame-Mode MPLS Network

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



# Label Allocation in a Frame-Mode MPLS Network: Building the IP Forwarding Table

Routing Table of A


Network	Next Hop
X	B

Routing Table of B

Network	Next Hop
X	C

Routing Table of C

Network	Next Hop
X	D



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






# Label Allocation in a Frame-Mode MPLS Network: Allocating Labels

Routing Table of B

Network	Next Hop
X	C

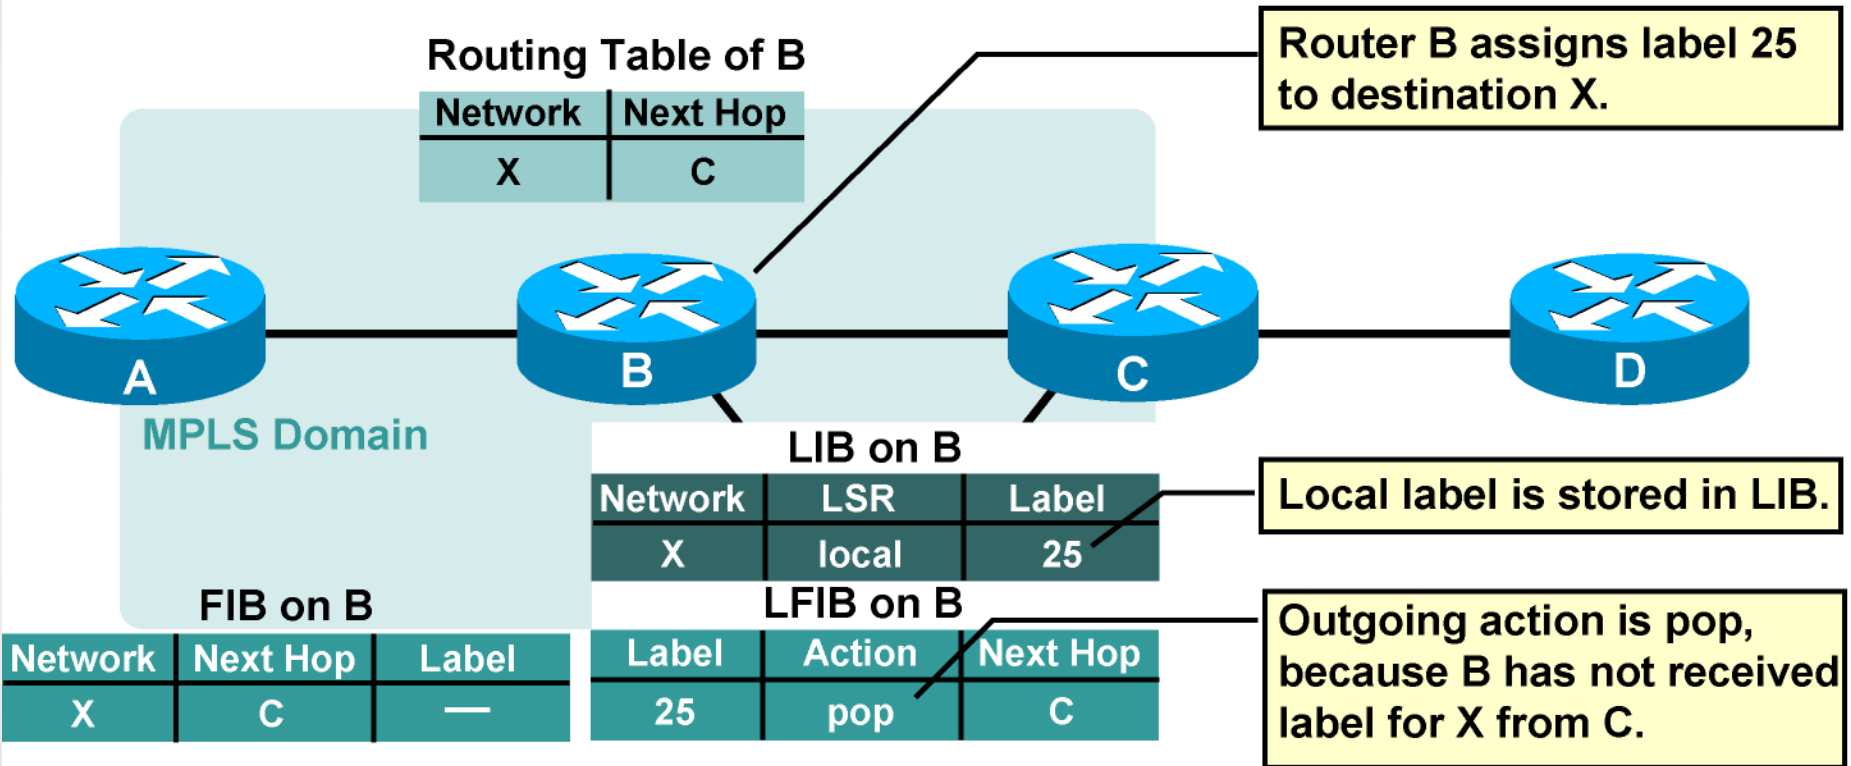
Router B assigns label 25 to destination X.



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



# Label Allocation in a Frame-Mode MPLS Network: LIB and LFIB Setup



- LIB and LFIB structures have to be initialized on the LSR allocating the label.



# Label Allocation in a Frame-Mode MPLS Network: Labels and Table Setup

Routing Table of A

Network	Next Hop
X	B

Router A assigns label 12 to destination X.



MPLS Domain

LIB on A

Network	LSR	Label
X	local	12

Local label is stored in LIB.

FIB on A

Network	Next Hop	Label
X	B	—

LFIB on A

Label	Action	Next Hop
12	pop	B

Outgoing action is pop, because A has not received label for X from B.


– Router A allocates a label for X independently of router B.



# 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.




# Label Distribution and Advertisement: Receiving Label Advertisement

LIB on A

Network	LSR	Label
X	B	25
	local	12

LIB on C

Network	LSR	Label
X	B	25



FIB on A

Network	Next Hop	Label
X	B	25

LIB on E

Network	LSR	Label
X	B	25

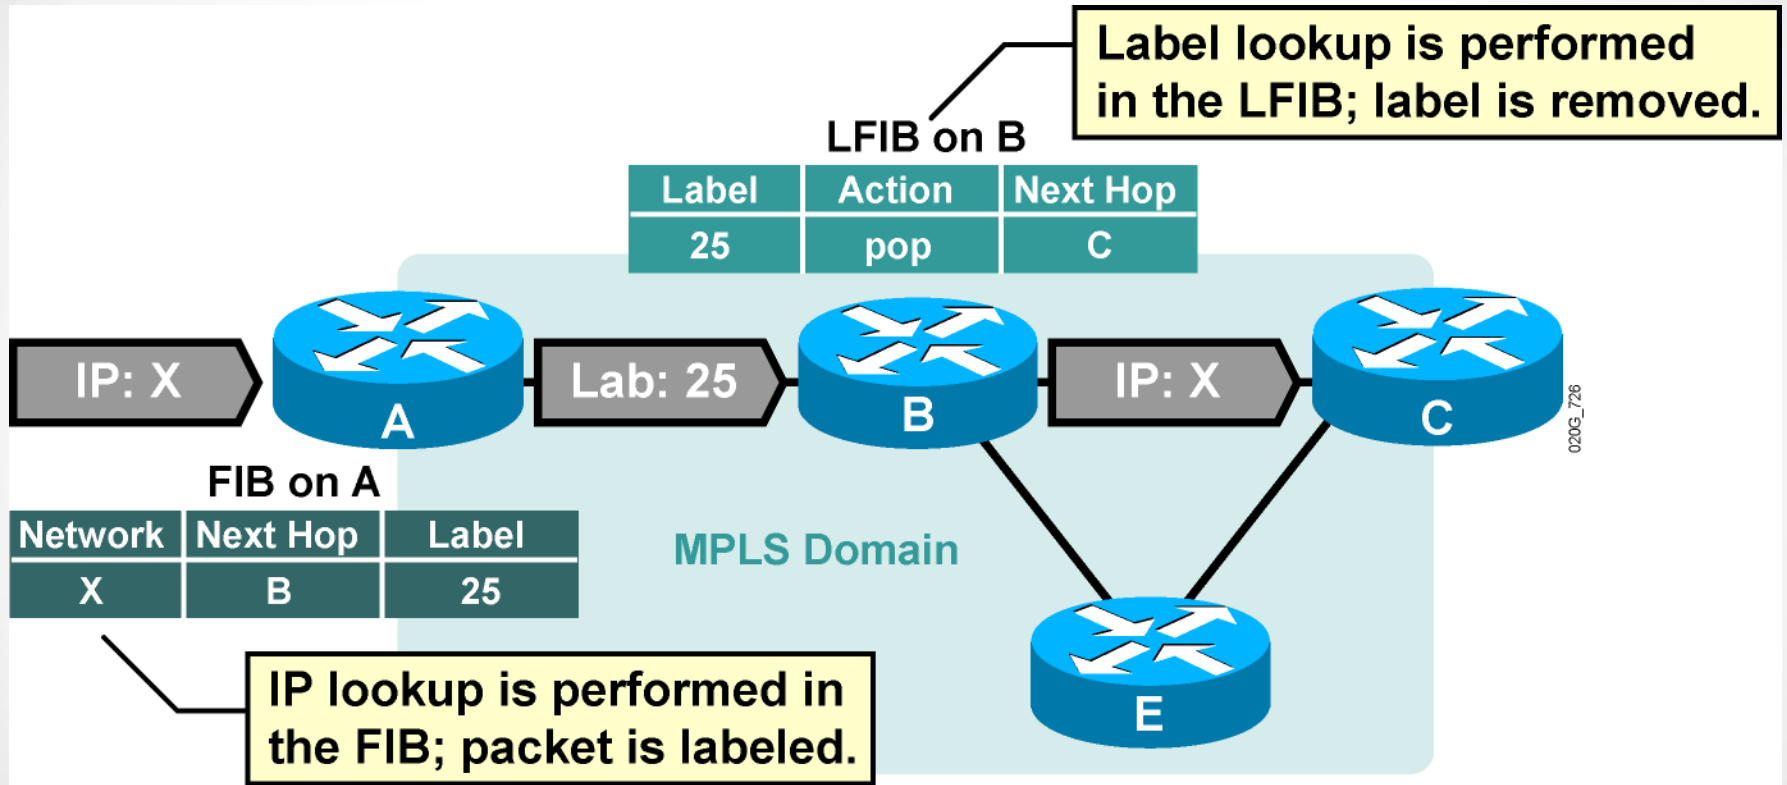
LFIB on A

Label	Action	Next Hop
12	25	B

- 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.



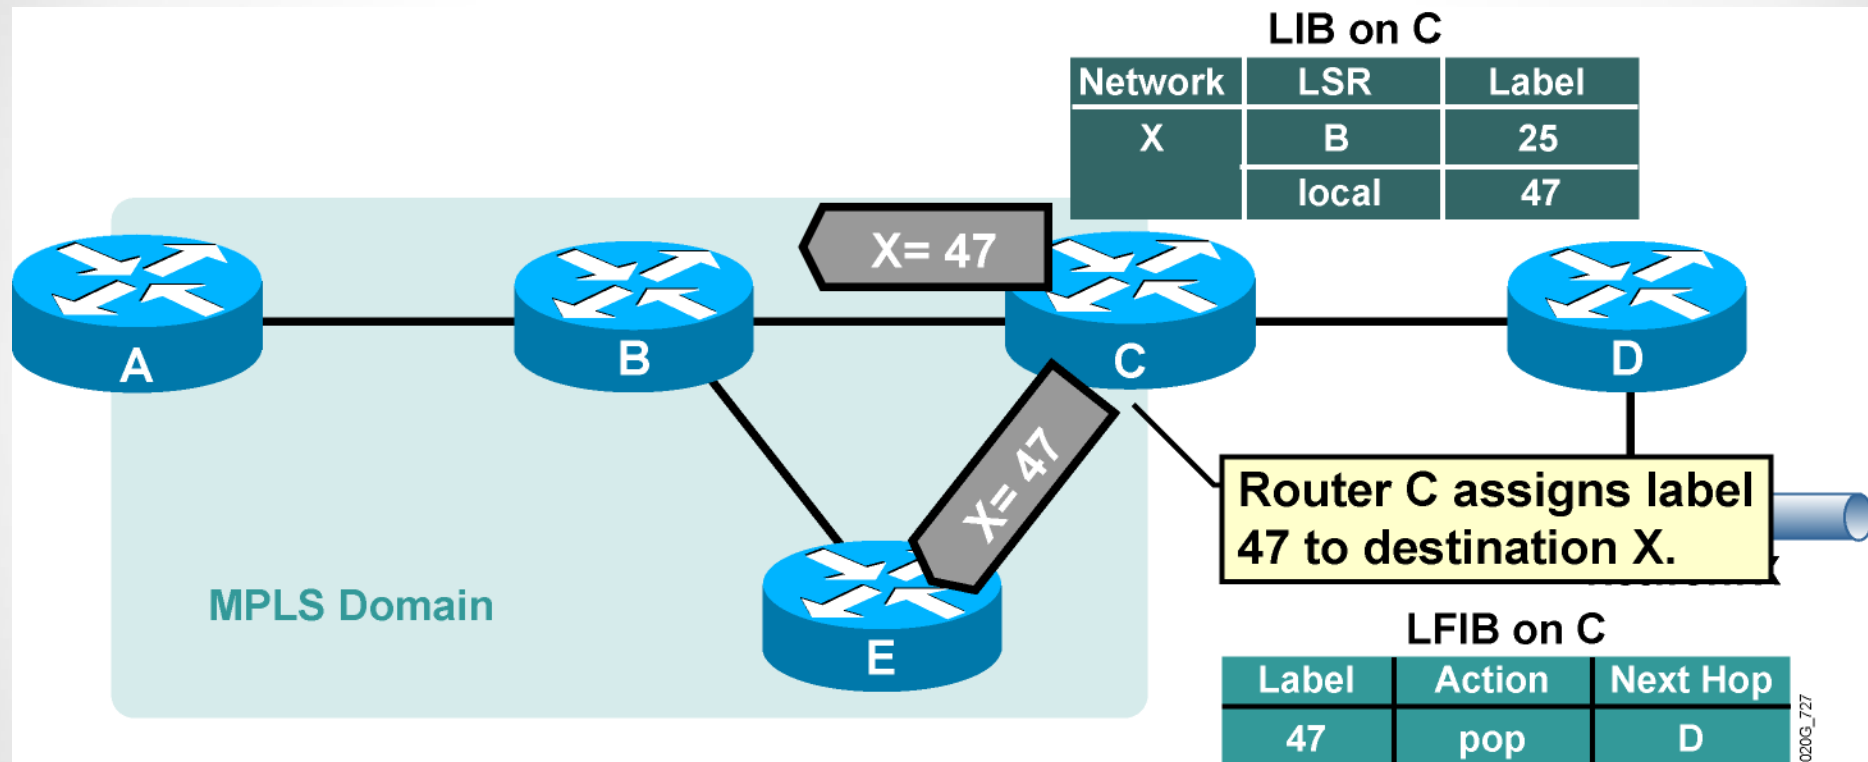
# Label Distribution and Advertisement: Interim Packet Propagation



- Forwarded IP packets are labeled only on the path segments where the labels have already been assigned.



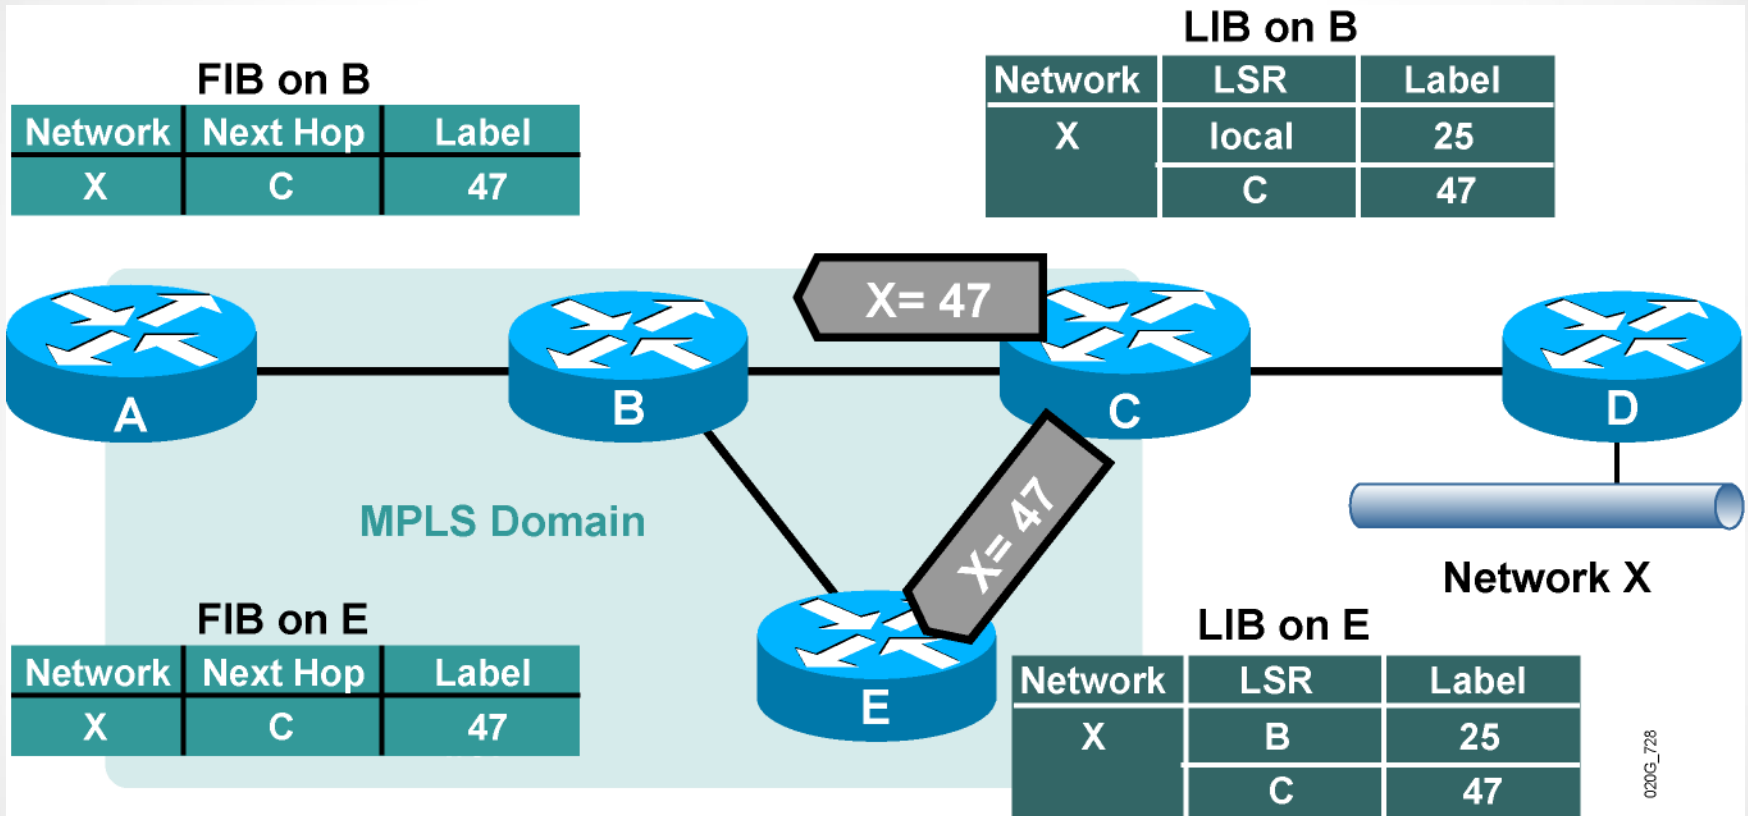
# Label Distribution and Advertisement: Further Label Allocation



- Every LSR will eventually assign a label for every destination.



# Label Distribution and Advertisement: Receiving Label Advertisement




020C\_728

- Every LSR stores received information in its LIB.
- LSRs that receive their label from their next-hop LSR will also populate the IP forwarding table.





# Populating the LFIB



- Router B has already assigned a label to network 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.

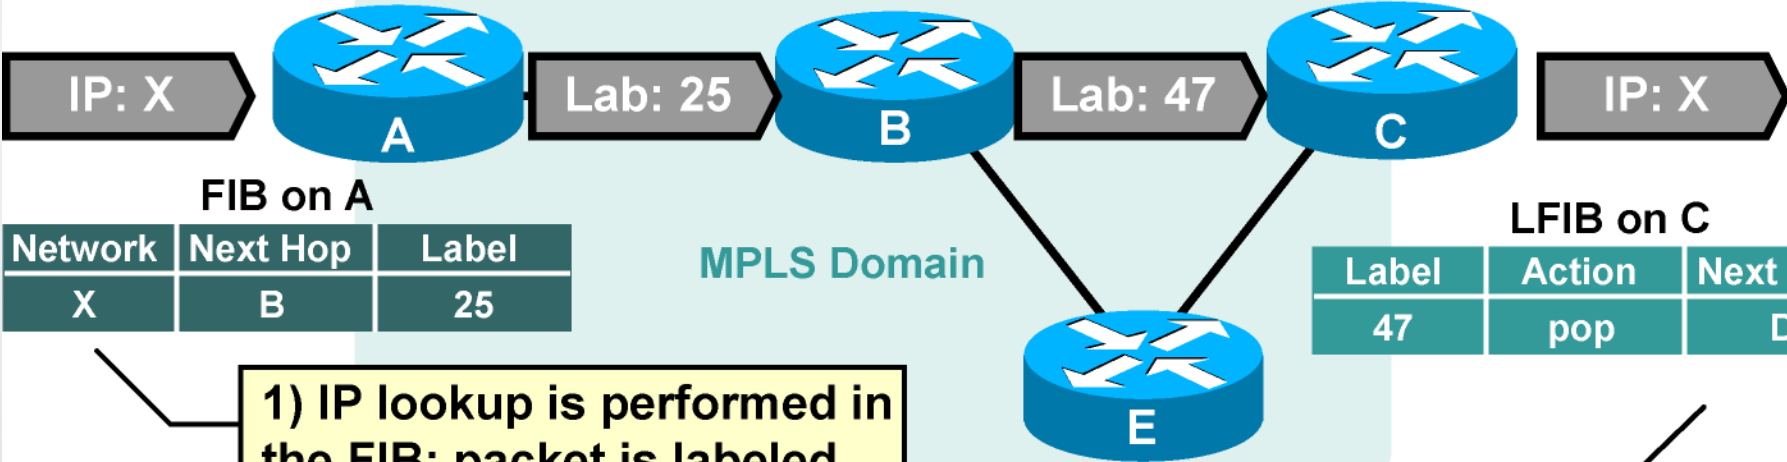


# Packet Propagation Across an MPLS Network

2) Label lookup is performed in the LFIB; label is switched.

LFIB on B

Label	Action	Next Hop
25	47	C



FIB on A

Network	Next Hop	Label
X	B	25

LFIB on C

Label	Action	Next Hop
47	pop	D

1) IP lookup is performed in the FIB; packet is labeled.

3) Label lookup is performed in the LFIB; label is removed.

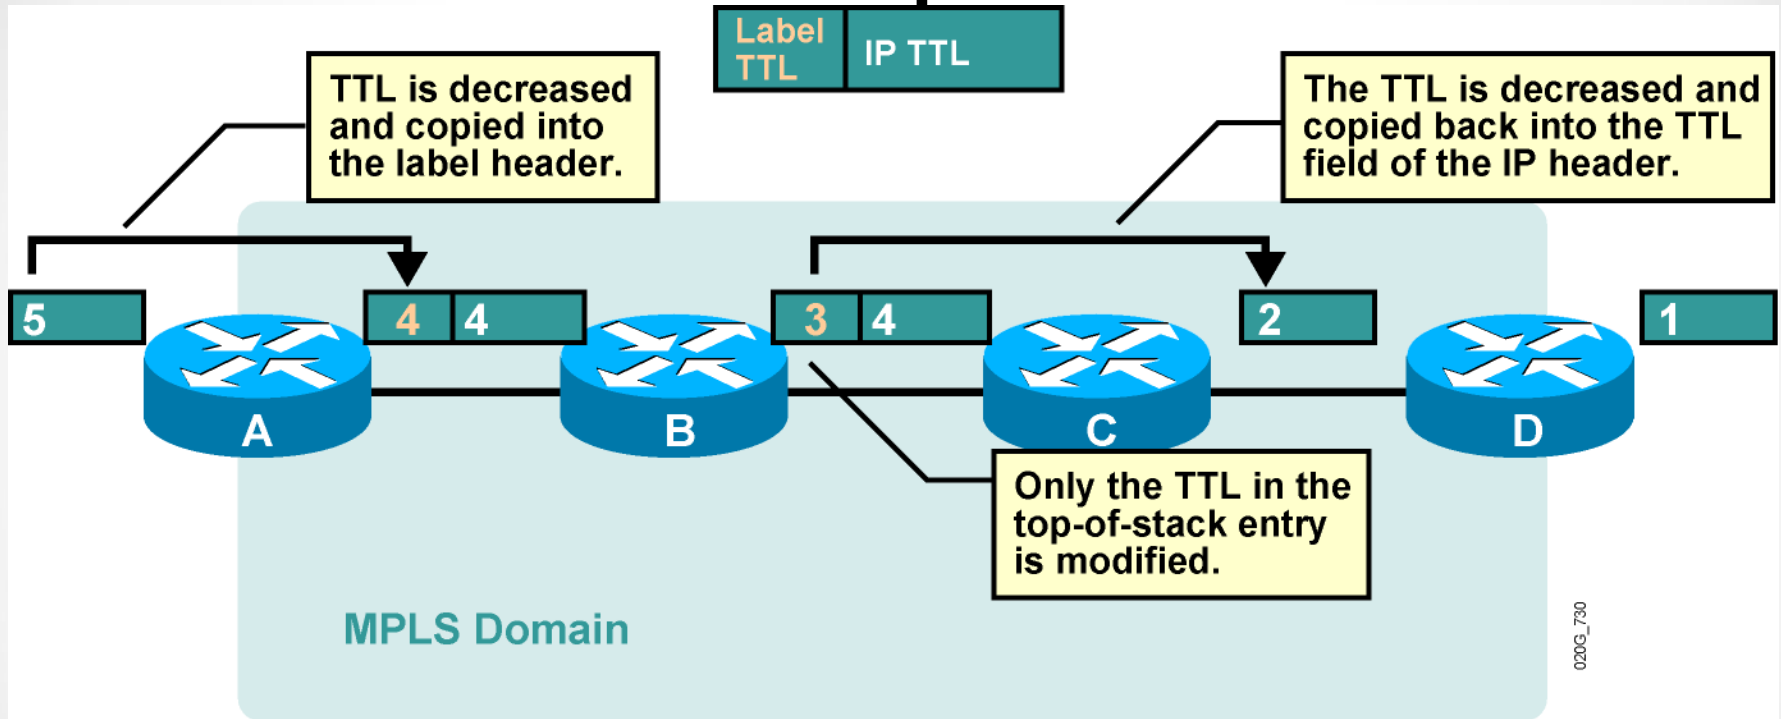


# Loop Detection

- LDP 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**).



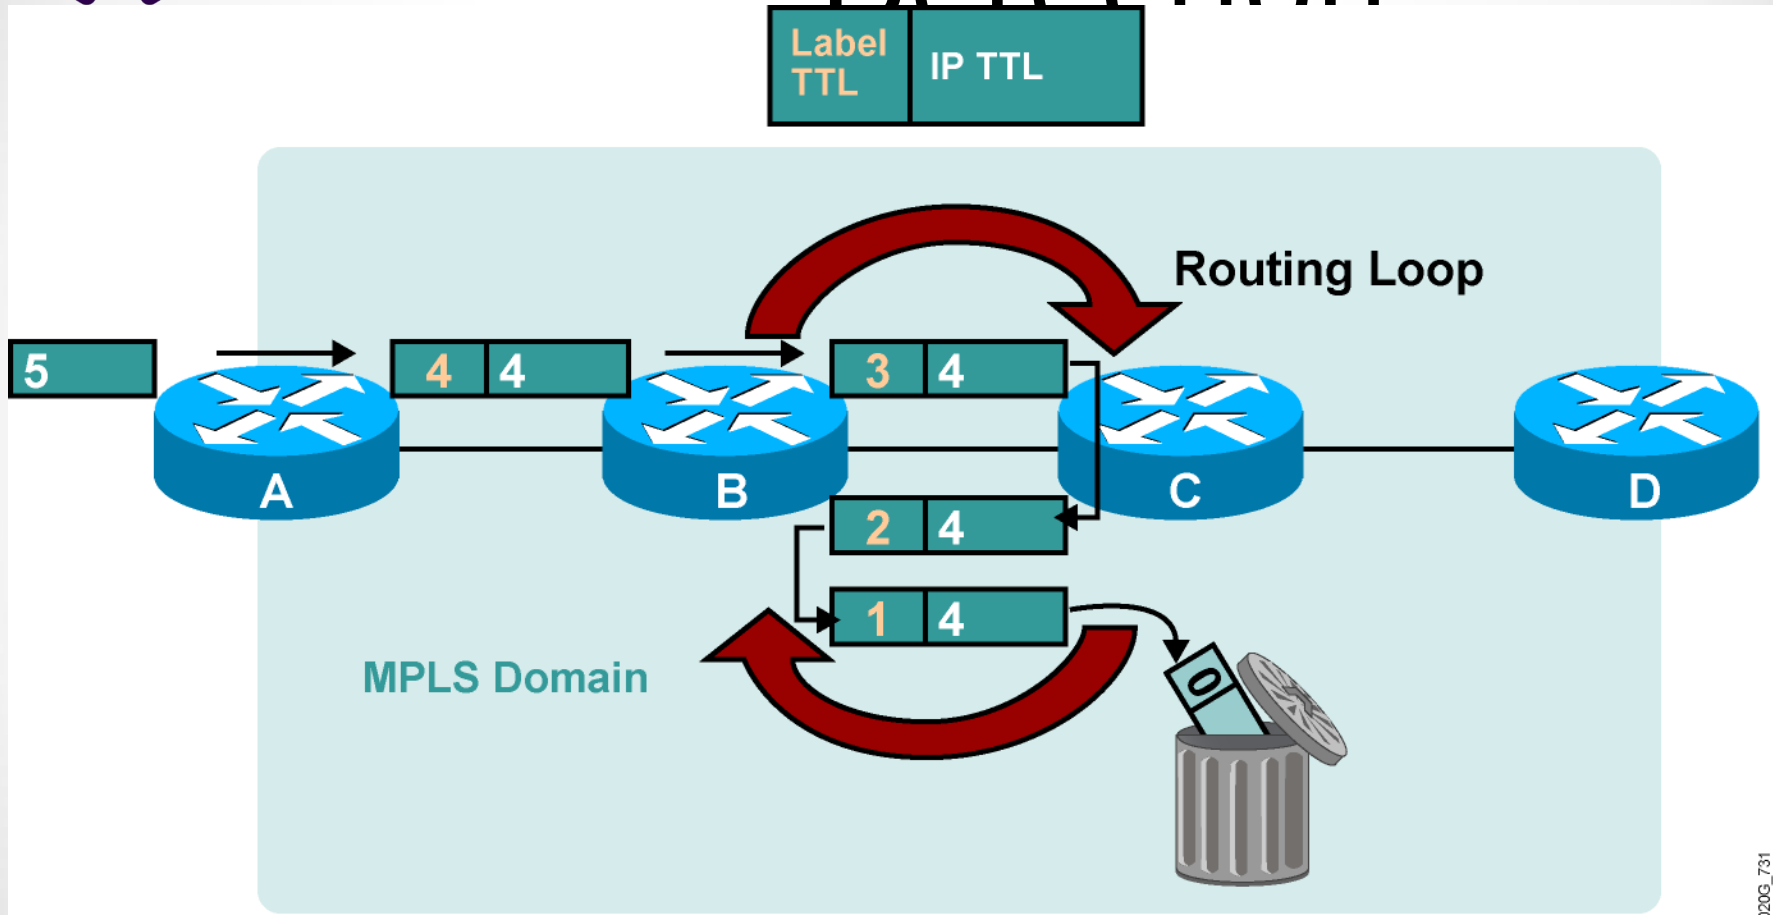
# Normal TTL Operation



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



# TTL and Loop Detection



Labeled packets are dropped when the TTL is decreased to 0.

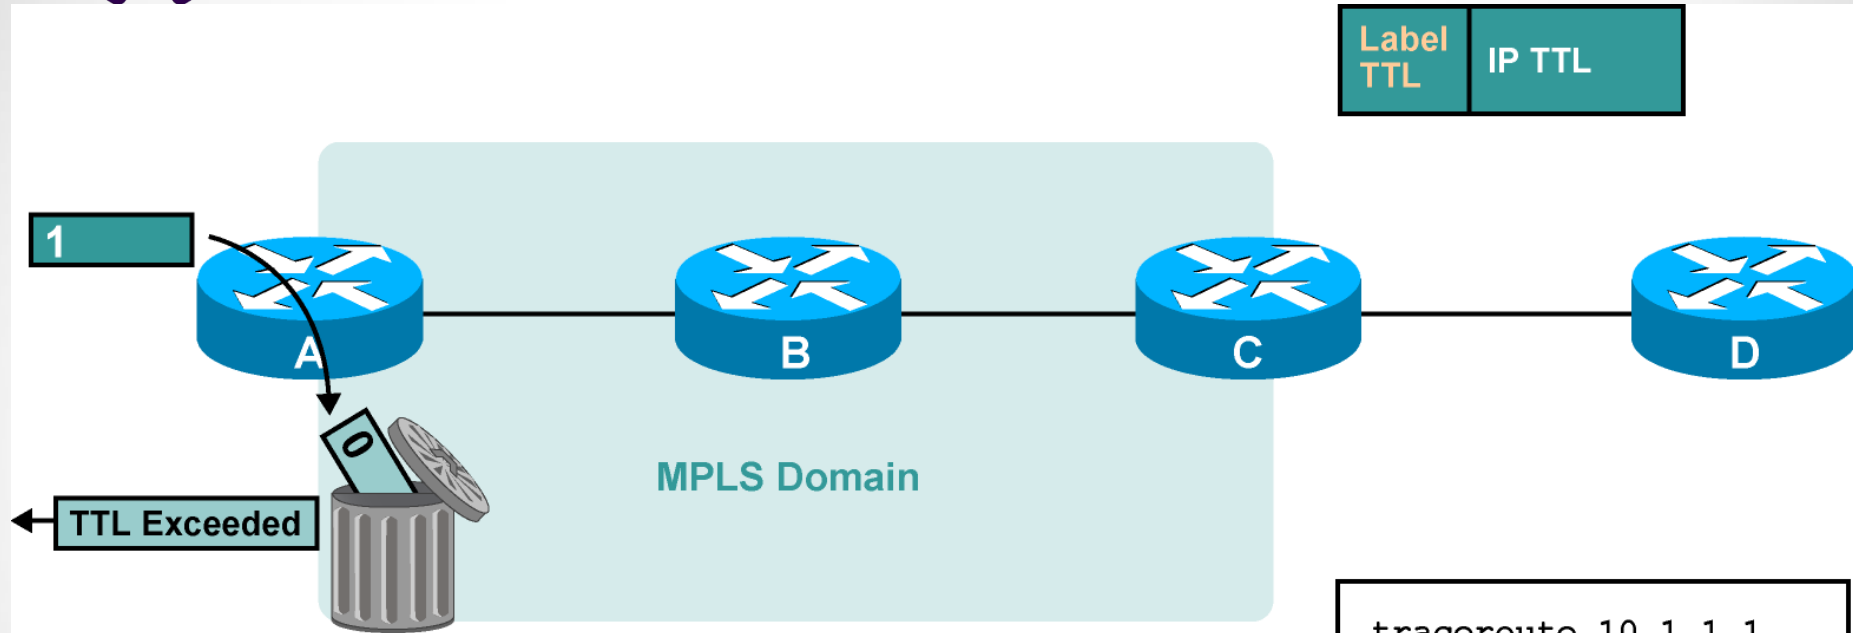


# Disabling TTL Propagation

- TTL propagation can be disabled.
- The **IP TTL** value is not copied into the TTL field of the label, and the **label TTL** is not copied back into the 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.




# Traceroute with Disabled TTL Propagation



- The first traceroute packet (ICMP or UDP) that reaches the network is dropped on router A.
- An ICMP TTL exceeded message is sent to the source from router A.



# Traceroute with Disabled TTL Propagation (Cont.)



- The second traceroute packet that reaches the network is dropped on router D.
- An ICMP TTL exceeded message is sent to the source from router D.

```
traceroute 10.1.1.1
 1  10 ms  A.acme.com
 2  10 ms  D.acme.com
```






# Operating and Configuring Huawei VRP Devices

Operating Huawei VRP Software



# CLI Command Line Views



View the running status and statistics of the device.

Set the system parameters of the device.

Configure interface parameters.

Configure most routing protocol parameters.

```
<Huawei>system-view
```

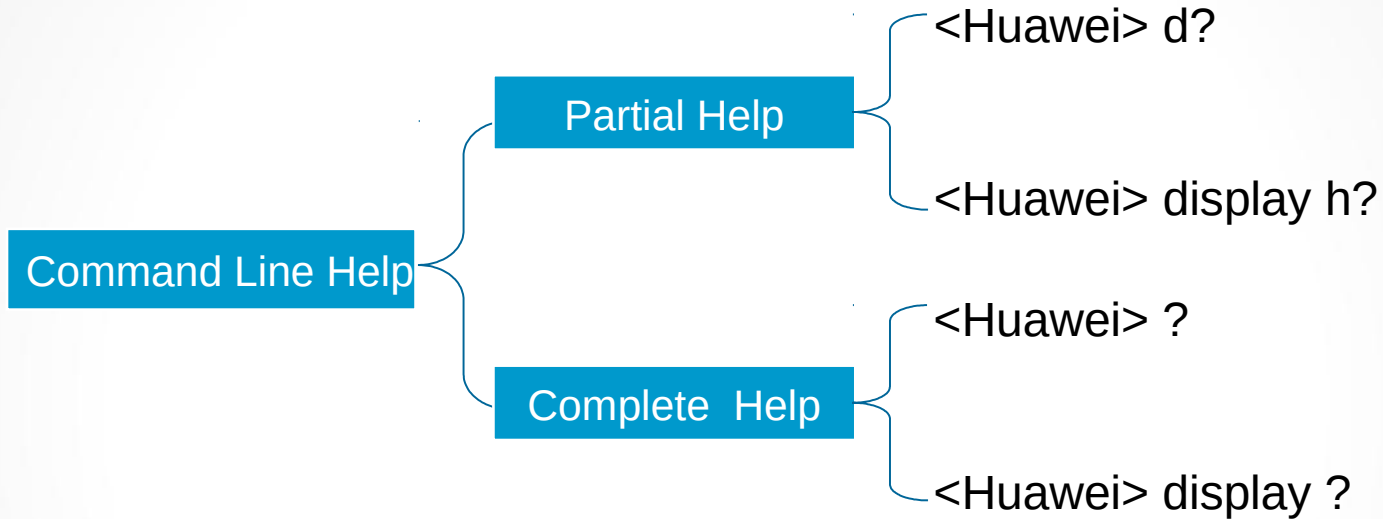
Enter system view, return user view with Ctrl+Z.

```
[Huawei]interface GigabitEthernet 0/0/0
```

```
[Huawei-GigabitEthernet0/0/0]
```



# CLI Help Features



```
[Huawei]d?
```

```
ddns
```

```
dhcpv6
```

```
display
```

```
domain
```

```
dhcp
```

```
diagnose
```

```
dns
```

```
dot1x
```



# CLI Basic Device Setup

Command	Function
sysname	Configures the Device Name..

```
<Huawei>system-view
```

```
Enter system view, return user view with Ctrl+Z.
```

```
[Huawei]sysname RTA
```

```
[RTA]
```

The system name should be assigned to uniquely identify each device within an enterprise network.



## CLI Terminal Attributes

Command	Function
idle-timeout	Sets the timeout duration of the user connection.
screen-length	Sets the number of lines displayed on each terminal screen after a command is executed
history-command max-size	Sets the size of the history command buffer.

```
# Set the size of the history command buffer to 20.  
<Huawei>system-view  
[Huawei]user-interface console 0  
[Huawei-ui-console0]history-command max-size 20  
# Set the timeout duration to 1 minute and 30 seconds.  
[Huawei-ui-console0]idle-timeout 1 30
```



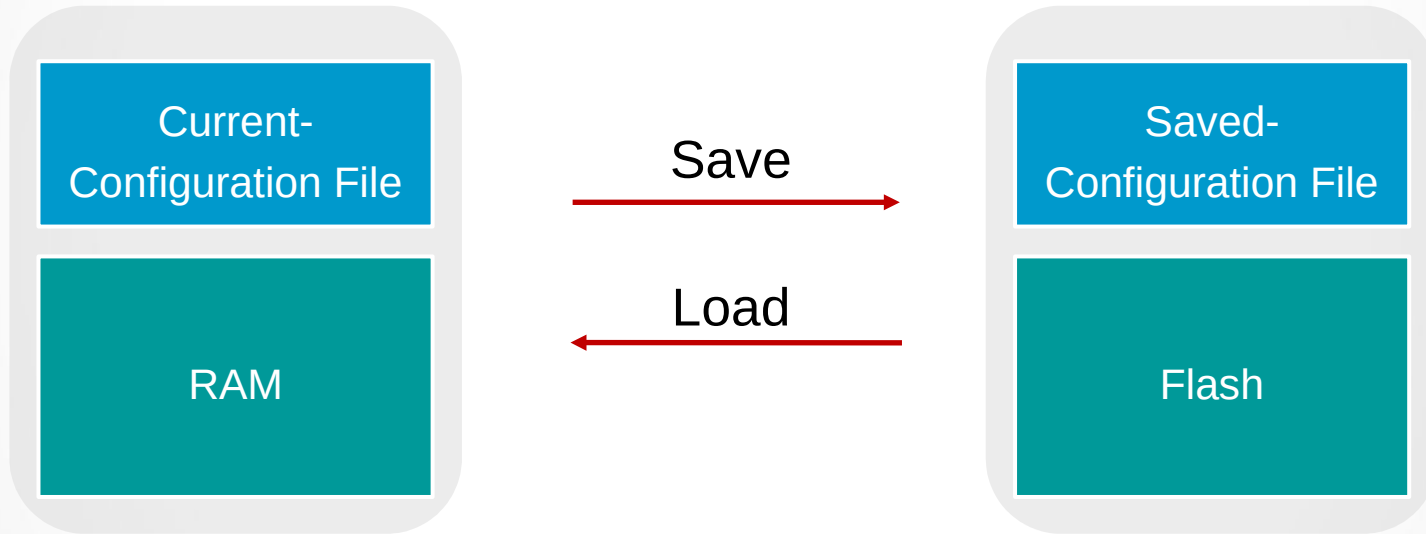
# CLI Interface Configuration



```
# Configure an IP address of 10.0.12.1/24 on interface G0/0/0
and an IP address of 1.1.1.1/32 on loopback interface 0.
<Huawei>system-view
[Huawei]interface GigabitEthernet 0/0/0
[Huawei-GigabitEthernet0/0/0]ip address 10.0.12.1 255.255.255.0
[Huawei-GigabitEthernet0/0/0]interface loopback 0
[Huawei-LoopBack0]ip address 1.1.1.1 32
```



# Configuration File Management System





# Saving the Configuration File

Command	Function
Save	Save the current configuration

```
<Huawei>save
The current configuration will be written to the device.
Are you sure to continue?[Y/N]y
It will take several minutes to save configuration file, please
wait.....
Configuration file had been saved successfully
Note: The configuration file will take effect after being activated
```





# Configuring Loopback interface

```
[R1]interface loopback 0  
[R1-LoopBack0]ip address 2.2.2.2 24
```





# Checking the routing table and the possibility of a connection

```
[R2]ping 10.0.1.2
```

```
  PING 10.0.1.2: 56 data bytes, press CTRL_C to break
```

```
    Reply from 10.0.1.2: bytes=56 Sequence=1 ttl=253 time=36 ms
```

```
    Reply from 10.0.1.2: bytes=56 Sequence=2 ttl=253 time=31 ms
```

```
    Reply from 10.0.1.2: bytes=56 Sequence=3 ttl=253 time=31 ms
```

```
    Reply from 10.0.1.2: bytes=56 Sequence=4 ttl=253 time=31 ms
```

```
● Reply from 10.0.1.2: bytes=56 Sequence=5 ttl=253 time=31 ms
```

```
--- 10.0.1.2 ping statistics ---
```

```
  5 packet(s) transmitted
```

```
  5 packet(s) received
```

```
  0.00% packet loss
```

```
  round-trip min/avg/max = 31/32/36 ms
```



# Checking the status of routing tables

```
[R2]display ip routing-table
```

```
Route Flags: R - relay, D - download to fib
```

```
-----  
Routing Tables: Public
```

```
Destinations : 19
```

```
Routes : 19
```

Destination/Mask	Proto	Pre	Cost	Flags	NextHop	Interface
● 2.2.2.2/32	OSPF	10	1562	D	10.0.12.1	Serial1/0/0
3.3.3.0/24	Direct	0	0	D	3.3.3.3	LoopBack0
3.3.3.3/32	Direct	0	0	D	127.0.0.1	InLoopBack0
3.3.3.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0
4.4.4.4/32	OSPF	10	1562	D	10.0.23.3	Serial2/0/0
10.0.1.0/24	OSPF	10	1563	D	10.0.12.1	Serial1/0/0
10.0.2.0/24	OSPF	10	1563	D	10.0.23.3	Serial2/0/0
10.0.12.0/24	Direct	0	0	D	10.0.12.2	Serial1/0/0
10.0.12.1/32	Direct	0	0	D	10.0.12.1	Serial1/0/0
10.0.12.2/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.12.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.23.0/24	Direct	0	0	D	10.0.23.2	Serial2/0/0
10.0.23.2/32	Direct	0	0	D	127.0.0.1	InLoopBack0
10.0.23.3/32	Direct	0	0	D	10.0.23.3	Serial2/0/0
10.0.23.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0
127.0.0.0/8	Direct	0	0	D	127.0.0.1	InLoopBack0
127.0.0.1/32	Direct	0	0	D	127.0.0.1	InLoopBack0
127.255.255.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0
255.255.255.255/32	Direct	0	0	D	127.0.0.1	InLoopBack0



# Checking the status of LDP sessions between devices

```
[R2]display mpls ldp session
```

```
LDP Session(s) in Public Network
```

```
Codes: LAM(Label Advertisement Mode), SsnAge Unit(DDDD:HH:MM)
```

```
A '*' before a session means the session is being deleted.
```

```
-----  
PeerID                Status      LAM  SsnRole  SsnAge      KASent/Rcv  
-----  
2.2.2.2:0             Operational DU    Active    0000:00:11  46/46  
4.4.4.4:0             Operational DU    Passive   0000:00:10  43/43  
-----
```

```
TOTAL: 2 session(s) Found.
```



# Checking LDP LSP routes

```
[R2]display mpls ldp lsp  
LDP LSP Information
```

DestAddress/Mask	In/OutLabel	UpstreamPeer	NextHop	OutInterface
2.2.2.2/32	NULL/3	-	10.0.12.1	S1/0/0
2.2.2.2/32	1024/3	2.2.2.2	10.0.12.1	S1/0/0
2.2.2.2/32	1024/3	4.4.4.4	10.0.12.1	S1/0/0
*2.2.2.2/32	Liberal/1024		DS/4.4.4.4	
3.3.3.3/32	3/NULL	2.2.2.2	127.0.0.1	InLoop0
3.3.3.3/32	3/NULL	4.4.4.4	127.0.0.1	InLoop0
*3.3.3.3/32	Liberal/1024		DS/2.2.2.2	
*3.3.3.3/32	Liberal/1025		DS/4.4.4.4	
4.4.4.4/32	NULL/3	-	10.0.23.3	S2/0/0
4.4.4.4/32	1025/3	2.2.2.2	10.0.23.3	S2/0/0
4.4.4.4/32	1025/3	4.4.4.4	10.0.23.3	S2/0/0
*4.4.4.4/32	Liberal/1025		DS/2.2.2.2	

```
TOTAL: 8 Normal LSP(s) Found.
```

```
TOTAL: 4 Liberal LSP(s) Found.
```

```
TOTAL: 0 Frr LSP(s) Found.
```

```
A '*' before an LSP means the LSP is not established
```

```
A '*' before a Label means the USCB or DSCB is stale
```

```
A '*' before a UpstreamPeer means the session is in GR state
```

```
A '*' before a DS means the session is in GR state
```

```
A '*' before a NextHop means the LSP is FRR LSP
```



# Switching policy for LSR

Switch the policy to All for all LSR so that all static routes and routes from dynamic routing protocols (such as IGP) can be used to build LDP LSPs routes.

```
[R2]mpls  
[R2-mpls]lsp-trigger all
```

The `lsp-trigger` command sets a policy for establishing LDP LSPs. By default, LDP uses IP host routes with 32-bit addresses (excluding host routes with 32-bit interface addresses) to establish LSPs.

```
lsp-trigger { all | host | ip-prefix ip-prefix-name | none }
```

- all** Indicates that all static and IGP routes trigger the establishment of LSPs. -
- host** Indicates that IP host routes with 32-bit addresses (excluding host routes with 32-bit interface addresses) trigger the establishment of LSPs. -
- ip-prefix** *ip-prefix-name* Specifies the name of the IP prefix list that triggers the establishment of LSPs. The value is an existing IP prefix list name.
- none** Indicates that the establishment of an LSP is not triggered. -



# Configuring the input policy for LDP

If the labels received on the router are not controlled, then it will set a large number of LSP routes, which will require a large amount of memory.

When configuring the input policy for LDP, the router must be configured to receive labels only from the specified router and sets the LSPs route to it, while saving resources.

```
[R1]ip ip-prefix prefix1 permit 10.0.12.0 24
[R1]mpls ldp
[R1-mpls-ldp]inbound peer 3.3.3.3 fec ip-prefix prefix1
[R1-mpls-ldp]quit
[R1]display mpls lsp
```

-----  
LSP Information: LDP LSP  
-----

FEC	In/Out Label	In/Out IF	Vrf Name
2.2.2.2/32	3/NULL	-/-	
10.0.12.0/24	3/NULL	-/-	
10.0.1.0/24	3/NULL	-/-	
2.2.2.0/24	3/NULL	-/-	



# Configuring the input policy for LDP

**inbound peer** { peer-id | peer-group peer-group-name | all } fec { none | host | ip-prefix prefix-name }  
The inbound policy for all peers takes effect: `inbound peer all fec none`

An inbound policy for allowing the local LSR to receive label mapping messages from a specified LDP peer for a specified IGP route is configured.

To apply a policy associated with a single FEC range to an LDP peer group or all LDP peers from which the local LSR receives Label Mapping messages, configure either the peer-group peer-group-name or all parameter in the command.

If two inbound policies are configured in sequence and the peer parameters in the two commands are the same, the second command overwrites the first one. For example, the following two inbound policies are configured:

If an inbound policy for all peers is configured and another inbound policy for a specific peer or peer group is configured, the previous policy has a higher priority, and the later policy does not take effect. For example:

MPLS and MPLS LDP must be enabled globally before an inbound policy is configured.

To delete all inbound policies simultaneously, run the `undo inbound peer all` command

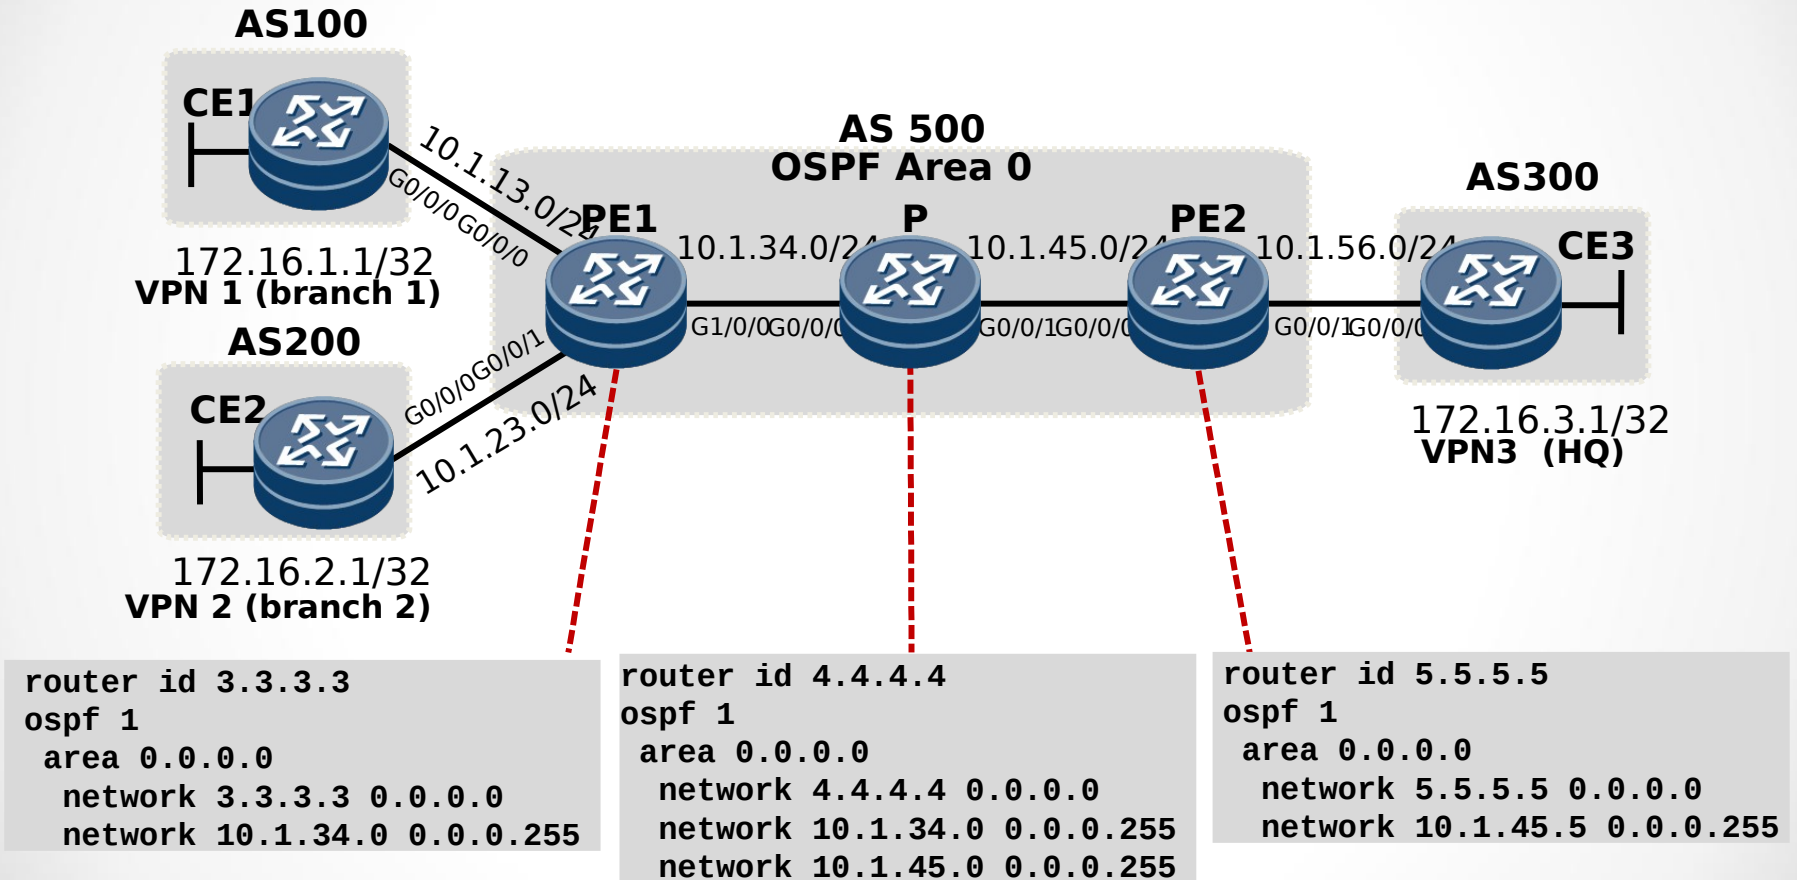




# MPLS Implementation on Huawei VRP Platforms

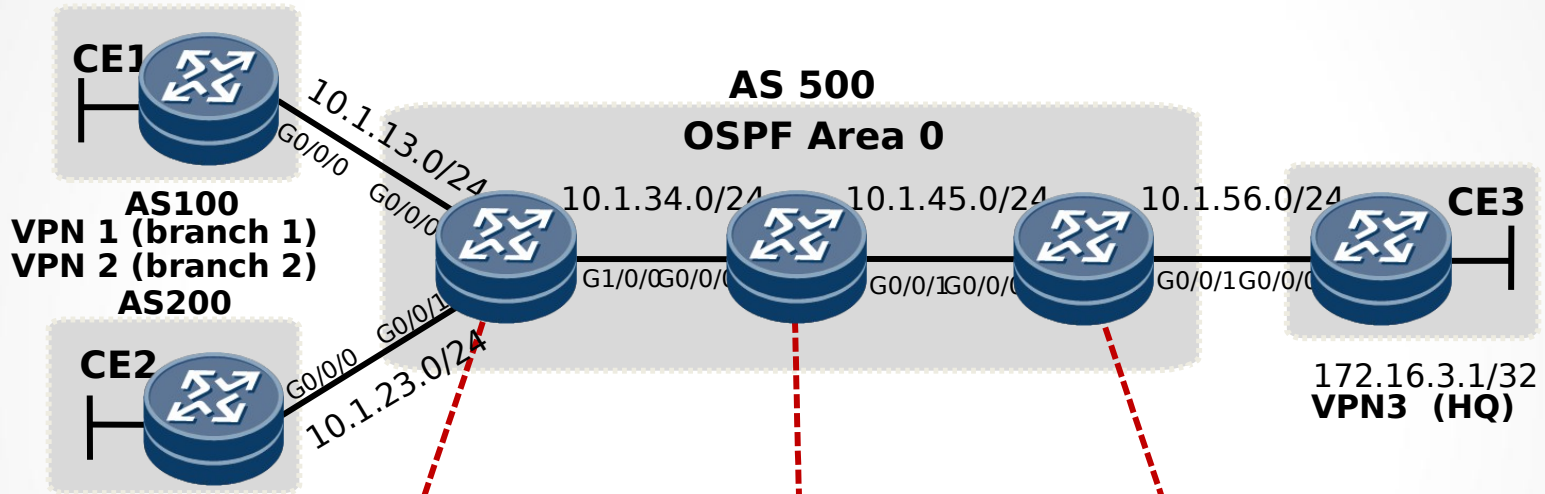


# Configuring IGP on the Backbone Network





# Configuring MPLS



```

mpls lsr-id 3.3.3.3
mpls
mpls ldp
#
interface GigabitEthernet1/0/0
mpls
mpls ldp
  
```

```

mpls lsr-id 4.4.4.4
mpls
mpls ldp
#
interface GigabitEthernet0/0/0
mpls
mpls ldp
#
interface GigabitEthernet0/0/1
mpls
mpls ldp
  
```

```

mpls lsr-id 5.5.5.5
mpls
mpls ldp
#
interface GigabitEthernet0/0/0
mpls
mpls ldp
  
```



# Symptoms of Common Frame-Mode MPLS Issues

- The LDP session does not start.
- Labels are not allocated.
- Labels are not distributed.
- Packets are not labeled, although the labels have been distributed.
- MPLS intermittently breaks after an interface failure.
- Large packets are not propagated across the network.



# LDP Session Startup Issues

- Symptom
  - LDP neighbors are not discovered.
- Diagnosis
  - MPLS is not enabled on the adjacent router.
- Verification
  - Verify with the *display mpls interface* command on the adjacent router.



# LDP Session Startup Issues (Cont.)

- Symptom
  - LDP neighbors are not discovered.
- Diagnosis
  - Packet filter drops LDP neighbor discovery packets.
- Verification
  - Verify access list presence with the show ip interface command.
  - Verify access list contents with the show access-list command.



# LDP Session Startup Issues (Cont.)

- Symptom
  - LDP neighbors are discovered; the LDP session is not established.
  
- Diagnosis
  - The connectivity between loopback interfaces is broken; the LDP session is usually established between loopback interfaces of adjacent LSRs.
  
- Verification
  - Verify connectivity with the extended ping command.



# Intermittent MPLS Failures After Interface Failure

## – Symptom

- The overall MPLS connectivity in a router intermittently breaks after an interface failure.

## – Diagnosis

- The IP address of a physical interface is used for the LDP identifier. Configure a loopback interface on the router.

## – Verification

- Verify the local LDP identifier.





# Packet Propagation Issues

- Symptom
  - Large packets are not propagated across the network.
    - Use of the **extended** ping command with varying packet sizes fails for packet sizes close to 1500 packets.
  - In some cases, MPLS might work, but MPLS VPN will fail.
- Diagnosis
  - There are label MTU issues or switches that do not support jumbo frames in the forwarding path.
- Verification
  - Issue the traceroute command through the forwarding path; identify all LAN segments in the path.
  - Verify the label MTU setting on routers attached to LAN segments.
  - Check for low-end switches in the transit path.