

# Label Assignment and Distribution

---

## Overview

This chapter describes the label assignment and distribution for unicast IP routing over MPLS.

It contains the following topics:

- Label Distribution Protocol (LDP) Role in Unicast IP Routing
- Typical Label Distribution in Packet-mode MPLS
- Convergence in Packet-mode MPLS
- Typical Label Distribution over LC-ATM Interfaces and VC-Merge
- MPLS Label Allocation, Distribution and Retention Modes
- LDP Neighbor Discovery
- Penultimate Hop Popping
- Summary

## Objectives

Upon completion of this chapter, you will be able to perform the following tasks:

- Describe the need for LDP/TDP in MPLS network.
- Describe the LDP/TDP neighbor discovery and session establishment procedures.
- Explain the needs for different LDP/TDP label distribution modes.
- Explain the difference between independent and ordered control.
- Describe various LDP/TDP retention modes (conservative and liberal).

- Explain the functions and benefits of penultimate-hop-popping.

# LDP Role in Unicast IP Routing

---

## Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Explain the need for the Label Distribution Protocol (LDP) in unicast IP routing MPLS application.
- Describe the LDP's interaction with other Label Switch Router (LSR) components.

# MPLS Unicast IP Routing

- **MPLS introduces a new 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 into 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.**

© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1-5

One of the applications of Multi-protocol Label Switching (MPLS) is Unicast IP routing. A label is assigned to destination IP networks and later used to label packets sent toward those destinations.

---

**Note** When using MPLS terminology we can say that a Forwarding Equivalence Class equals an IP destination networks.

---

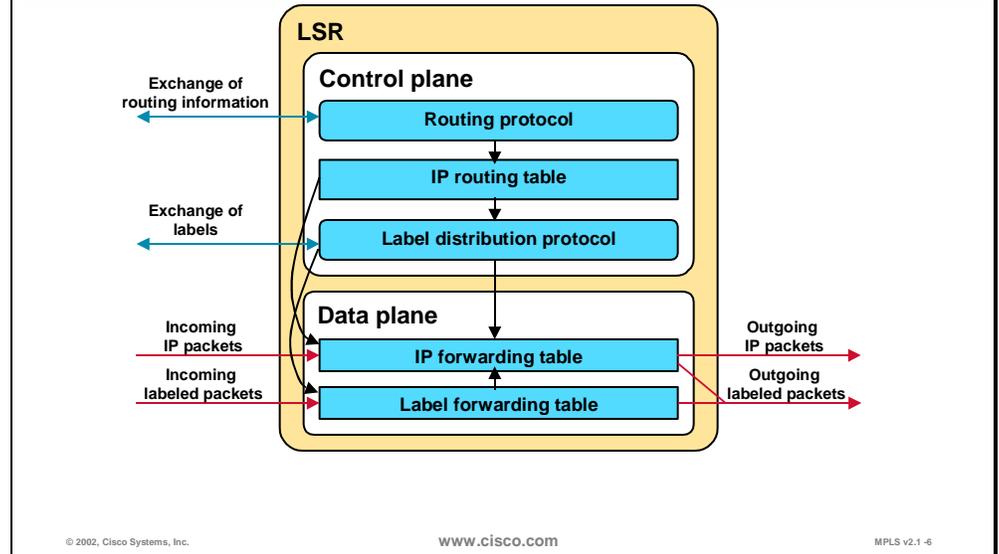
Standard or vendor-specific routing protocols are used to advertise IP routing information. MPLS adds a new piece of information that must be exchanged between adjacent routers.

There are two possible approaches to propagating this additional information (labels) between adjacent routers:

- Extend the functionality of existing routing protocols
- Create a new protocol dedicated to exchanging labels

The first approach would require much more time and effort because of a large number of different routing protocols (OSPF, IS-IS, EIGRP, IGRP, RIP, etc.) as well as causing interoperability problems when interconnecting routers that support this new functionality and those that do not. Therefore, the second option was selected.

# MPLS Unicast IP Routing Architecture



The figure shows the building blocks of routers that perform traditional IP forwarding.

The Control Plane consists of a routing protocol that exchanges routing information and maintains the contents of the main routing table. When combined with Cisco Express Forwarding (CEF) the IP forwarding table in the Data Plane takes care of forwarding the packets through the router.

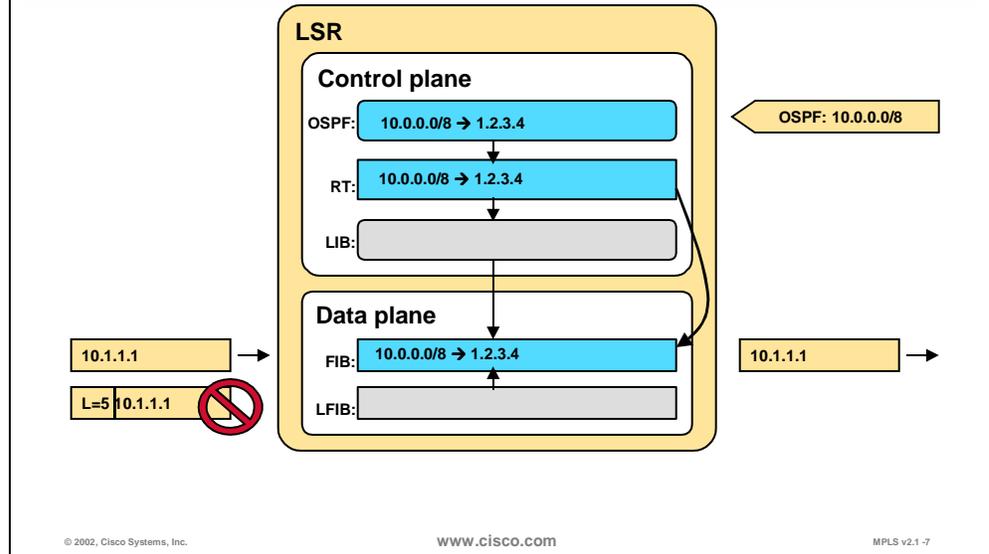
The new Label Distribution Protocol (LDP) in the Control Plane exchanges labels and stores them in the Label Information Base (LIB). This information is then used in the Data Plane to provide the MPLS functionality:

- A label is added to the IP forwarding table (Forwarding Information Base or FIB) to map an IP prefix to a next-hop label.
- A locally generated label is added to the Label Forwarding Information Base (LFIB) and mapped to a next-hop label.

The following forwarding scenarios are possible once MPLS is enabled on a router:

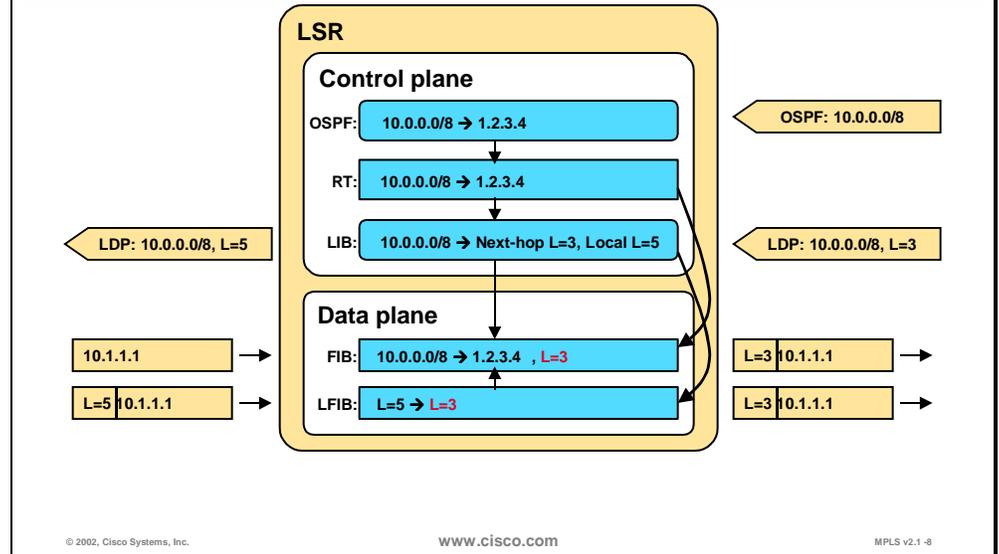
- An incoming IP packet is forwarded by using the FIB table and sent out as an IP packet (the usual CEF switching).
- An incoming IP packet is forwarded by using the FIB table and sent out as a labeled IP packet (the default action if there is a label assigned to the destination IP network).
- An incoming labeled packet is forwarded by using the LFIB table and sent out as a labeled IP packet.

# MPLS Unicast IP Routing: Example



The figure shows a scenario where IP packets are successfully forwarded using the FIB table. Labeled packets on the other hand are not forwarded due to a lack of information in the LFIB table. Normal MPLS functionality prevents this from happening since no adjacent router is going to use a label unless this router has previously advertised it. The example simply illustrates that label switching only tries to use the LFIB table if the incoming packet is labeled, even if the destination address is reachable by using the FIB table.

# MPLS Unicast IP Routing: Example (Cont.)



The figure shows a router where OSPF is used to exchange IP routing information and LDP is used to exchange labels.

An incoming IP packet is forwarded using the FIB table where a next-hop label dictates that the outgoing packet should be labeled with label 3.

An incoming labeled packet is forwarded using the LFIB table where the incoming (locally significant) label 5 is swapped with the next-hop label 3.

# Summary

A new protocol is introduced into MPLS-enabled networks to exchange labels assigned to IP destination networks. The Label Distribution Protocol (LDP) exchanges locally significant labels between adjacent routers. Labels, received from LDP peers, are bound to IP destination networks in the FIB table and local labels in the LFIB table.

## Lesson Review

1. Why is the LDP/TDP needed?
2. What is the forwarding equivalence class in MPLS unicast IP forwarding?
3. Where is this forwarding equivalence class taken from?
4. What is the output of the LDP/TDP?

# Typical Label Distribution in Packet- mode MPLS

---

## Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Describe label allocation in packet-mode MPLS environments.
- Describe TDP/LDP label distribution in packet-mode MPLS environments.
- Explain how the MPLS data structures are built based on label allocation and distribution.

# Label Allocation in Packet-Mode MPLS Environment

**Label allocation and distribution in packet-mode MPLS environment 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.

© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1 -13

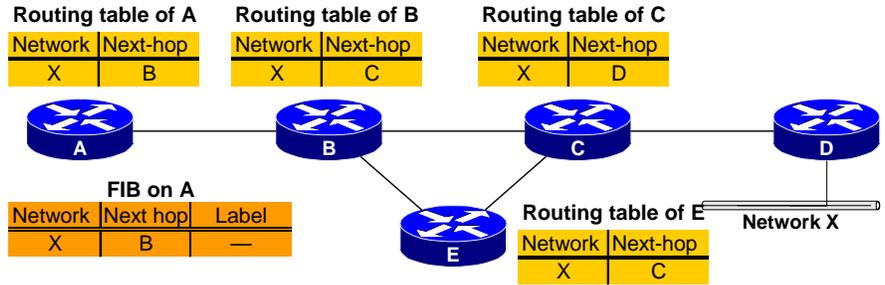
Unicast IP routing and MPLS functionality can be divided into the following modules:

- Routing information exchange using standard or vendor-specific IP routing protocols (OSPF, IS-IS, EIGRP, etc.).
- Generation of local labels. One locally unique label is assigned to each IP destination found in the main routing table and stored in the LIB table.
- Propagation of local labels to adjacent routers where these labels might be used as next-hop labels (stored in the FIB and LFIB tables to enable label switching).

The following data structures contain label information:

- **Label Information Base (LIB)** in the Control Plane is the database used by LDP where an IP prefix is assigned a locally significant label mapped to a next-hop label learned from the downstream neighbor.
- **Label Forwarding Information Base (LFIB)** in the Data Plane is the database used to forward labeled packets. Local labels, previously advertised to upstream neighbors, are mapped to next-hop labels, previously received from downstream neighbors.
- **Forwarding Information Base (FIB)** in the Data Plane is the database used to forward unlabeled IP packets. A forwarded packet is labeled if a next-hop label is available for a specific destination IP network. Otherwise, a forwarded packet is not labeled.

# Building the IP Routing Table



- IP routing protocols are used to build IP routing tables on all LSRs.
- Forwarding tables (FIB) are built based on IP routing tables with no labeling information.

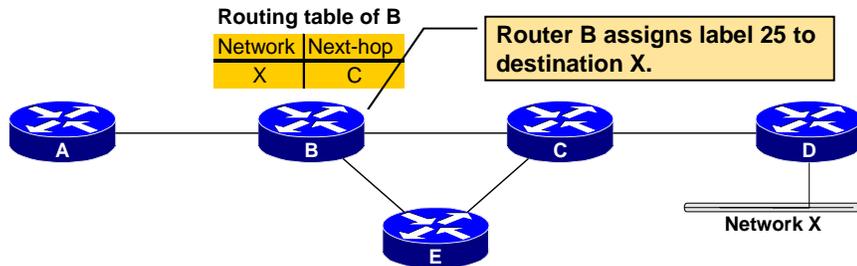
© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1-14

The figure illustrates how all routers learn about network X via an IGP such as OSPF, IS-IS, EIGRP, etc. The FIB table on router A contains the network X that is mapped to the IP next-hop address B. At this time a next-hop label is not available which means that all packets are forwarded in a traditional fashion (as unlabelled IP packets)

# Allocating Labels



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

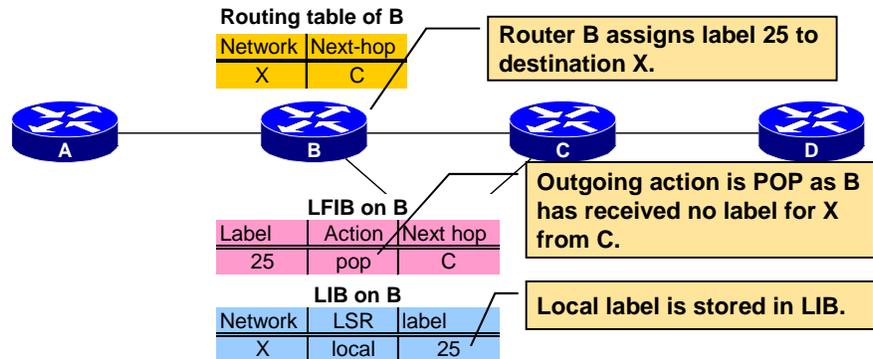
© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1 -15

The figure shows how router B generates a locally significant and locally unique label 25 assigned to IP network X. Router B generates this label regardless of other routers (asynchronous allocation of labels).

# LIB and LFIB Set-up



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

© 2002, Cisco Systems, Inc.

www.cisco.com

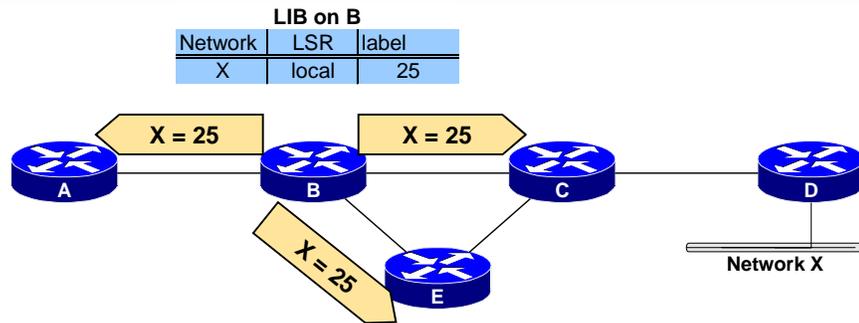
MPLS v2.1-16

Once a label is assigned to an IP prefix, it is stored in two tables:

- The LIB table is used to maintain the mapping between the IP prefix (network X), local label (25) and the next-hop label (not available yet).
- LFIB table is modified to contain the local label mapped to the *pop* action (label removal). The *pop* action is used until the next-hop label is received from the downstream neighbor.

**Note** The *pop* action results in labels being removed rather than swapped. Therefore, if a labeled packet is received, the label is removed and the remainder of the packet (which might be labeled if the incoming packet contained a label stack) is forwarded to the appropriate IP next-hop. A similar action is the *untagged* action, which is equivalent to the *pop* action with one exception - the resulting packet **must be unlabeled**, otherwise it will be discarded.

# Label Distribution



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

© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1 -17

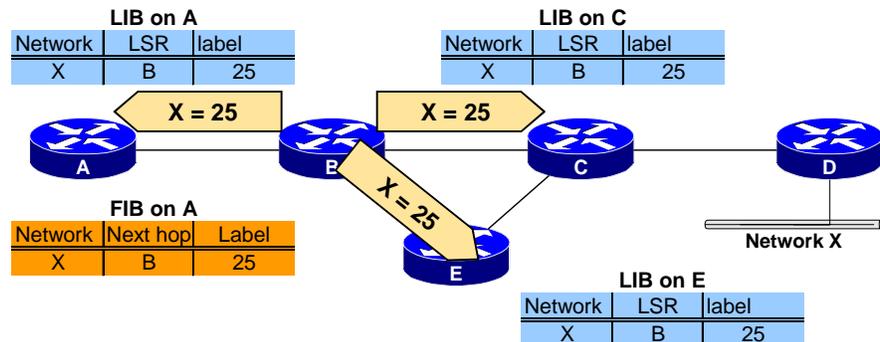
The figure illustrates the next step after a local label has been generated. Router B propagates this label to all adjacent neighbors where this label can be used as a next-hop label.

---

**Note** As router B cannot predict which routers might use it as the downstream neighbor, it sends its label mappings to all LDP neighbors.

---

# 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.

© 2002, Cisco Systems, Inc.

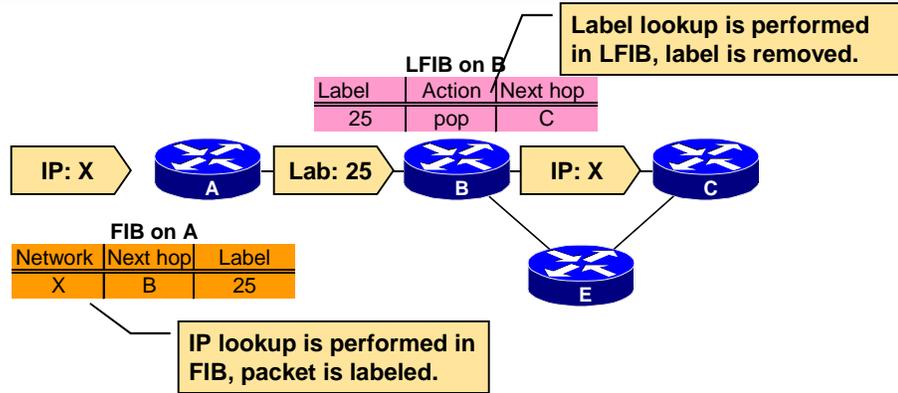
www.cisco.com

MPLS v2.1-18

Upon receiving an LDP update, router A can fill in the missing piece in its LIB, LFIB and FIB tables:

- Label 25 is stored in the LIB table as the label for network X received from LSR B.
- Label 25 is attached to the IP forwarding entry in the FIB table to enable the MPLS edge functionality (incoming IP packets are forwarded as labeled packets).
- Local label in the LFIB table is mapped to outgoing label 25 instead of the *pop* action (incoming labeled packets can be forwarded as labeled packets).

# Interim Packet Propagation



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

© 2002, Cisco Systems, Inc.

www.cisco.com

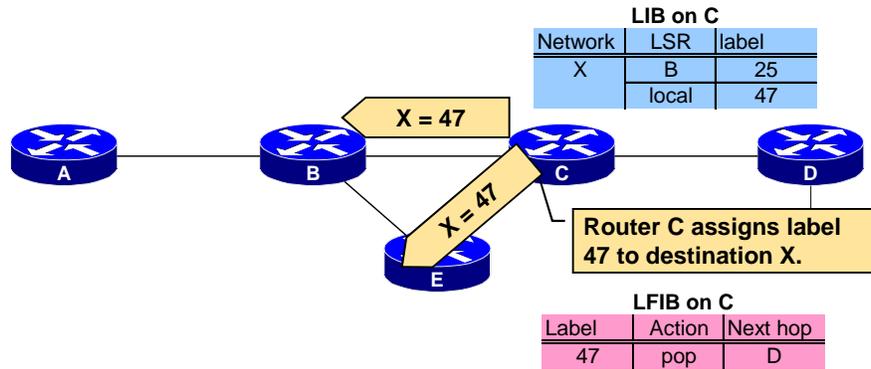
MPLS v2.1 -19

The figure shows how an unlabeled IP packet is forwarded based on the information found in the FIB table on router A. Label 25 found in the FIB table is used to label the packet.

Router B must remove the label, because the LSR B has not yet received any next-hop label (the action in its LFIB is *pop*).

Router A performs an IP lookup (CEF switching), whereas router B performs a label lookup (label switching) although the label was removed and a normal IP packet was sent out of router B.

# Further Label Allocation



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

© 2002, Cisco Systems, Inc.

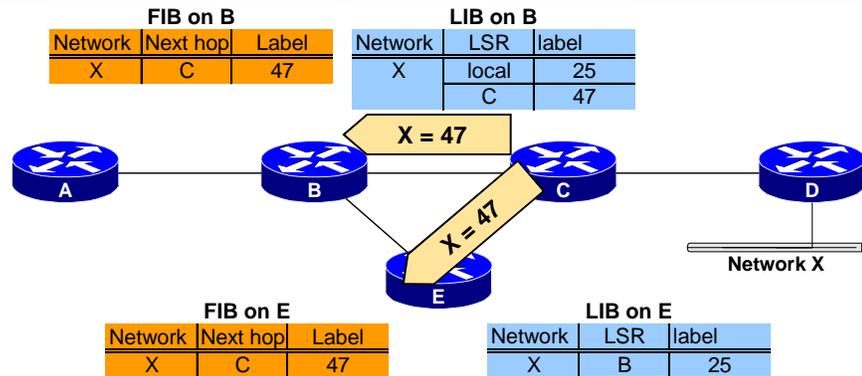
www.cisco.com

MPLS v2.1 -20

Since all routers in an MPLS domain asynchronously do the same as routers A and B, a label-switching path (LSP) is generated spanning from router A to router D.

The figure illustrates how an LDP update from router C is sent to adjacent routers including router B.

# Receiving Label Advertisement



- 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 (FIB).

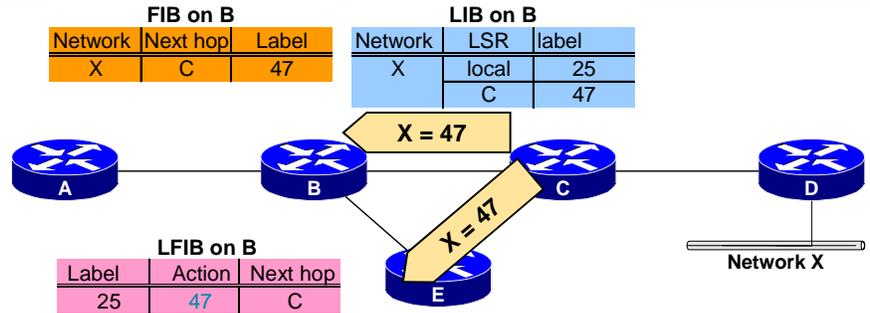
© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1 -21

Router B can now map the entry for network X in its FIB and the local label 25 in its LFIB to the next-hop label 47 received from the downstream neighbor router C.

# Populating LFIB



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

© 2002, Cisco Systems, Inc.

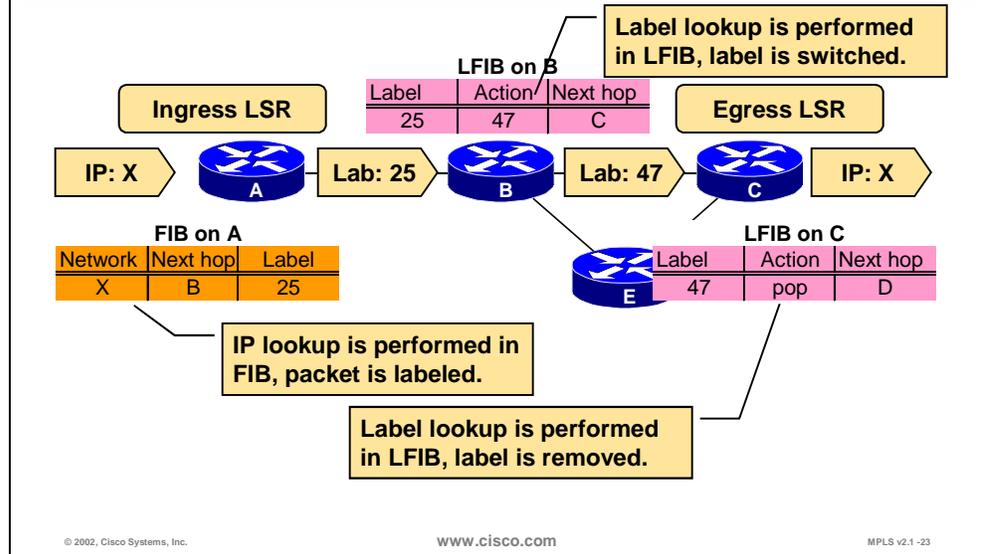
www.cisco.com

MPLS v2.1-22

Once router C advertises label 47 to adjacent routers the label switching path (LSP) for network X has two hops:

- On router A network X is mapped to the next-hop label 25.
- On router B label 25 is mapped to label 47.
- Router C still has no next-hop label. Label 47 is therefore still mapped to the *pop* action.

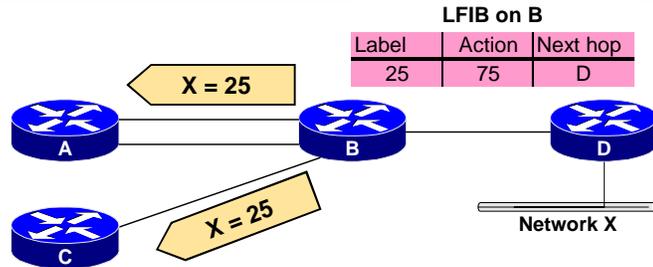
# Packet Propagation Across MPLS Network



The figure illustrates how IP packets are propagated across an MPLS domain:

- Router A labels a packet destined for network X by using the next-hop label 25 (CEF switching using the FIB table).
- Router B swaps label 25 with label 47 and forwards the packet to router C (label switching using the LFIB table).
- Router C removes the label and forwards the packet to router D (label switching using the LFIB table).

# Per-platform Label Allocation



- LFIB on a router usually does not contain an incoming interface.
- The same label can be used on any interface – **per-platform label allocation**.
- LSR announces a label to adjacent LSR only once – even if there are parallel links between them.

© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1-24

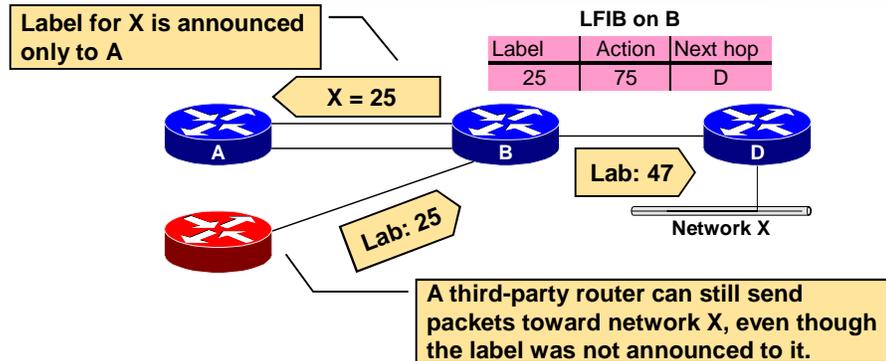
There are two possible approaches when assigning labels to networks:

- **Per-platform label allocation** where one single label is assigned to a destination network and announced to all neighbors. The label must be locally unique and valid on all incoming interfaces. This is the default operation in frame-mode MPLS.
- **Per-interface label allocation** where local labels are assigned to IP destination prefixes on a per-interface basis. These labels must be unique on a per-interface basis.

The figure illustrates how one single label (25) is assigned to a network and used on all interfaces. The same label is propagated to both routers A and C.

The figure also shows how one label is sent across one LDP session between routers A and B even though there are two parallel links between the two routers.

# Benefits and Drawbacks of Per-platform Label Allocation



## Benefits:

- Smaller LFIB
- Quicker label exchange

## Drawbacks:

- Insecure – any neighbor LSR can send packets with any label in LFIB

© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1 -25

A potential drawback of per-platform label allocation is illustrated in the figure, which shows how an adjacent router can send a labeled packet with a label that has not been previously advertised to this router (label spoofing).

On the other hand, per-platform label allocation results in a smaller LIB and LFIB table and faster label exchange.

# Summary

## Summary of Packet-mode MPLS Label Allocation and Distribution

**Packet-mode MPLS label allocation and distribution follows these rules:**

- Every LSR assigns a label for every destination in the IP routing table.
- Labels are assigned once per LSR (per-platform).
- Every LSR advertises its label assignments to all neighbors.
- Every LSR stores all advertised labels in the LIB.
- Labels received from next-hop LSRs are used to populate label information in FIB and the outgoing label in LFIB.

© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1-26

To summarize the operation of MPLS, the following steps are required:

- Allocation of local labels on a per-platform basis. Every destination in the main routing table is assigned a locally significant and locally unique label.
- Local labels are advertised to adjacent routers.
- Labels received from downstream neighbors are stored in the LIB table. They are used in the FIB table to forward and label IP packets and in the LFIB table to forward labeled packets.

## Lesson Review

1. Where are received labels stored?
2. Which forwarding table is used to forward a labeled packet?
3. Which forwarding table is used to forward an unlabeled packet?
4. What happens if the next-hop label is not in the LFIB table?
5. How many labels are usually assigned to one destination network?

# Convergence in Packet-mode MPLS

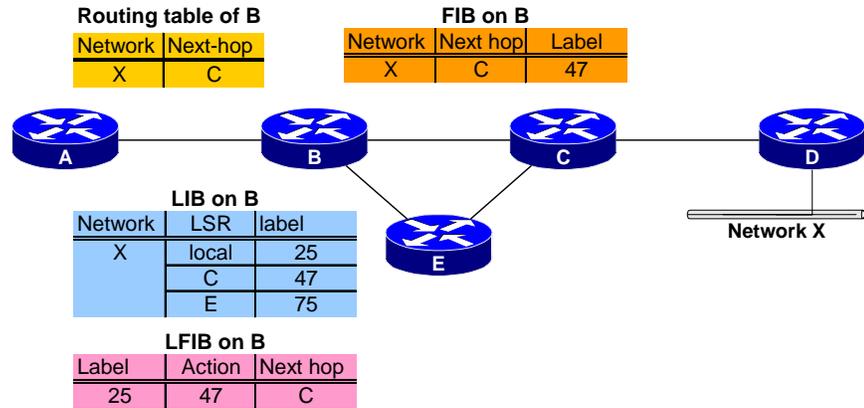
---

## Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Identify TDP/LDP convergence issues.
- Describe the interaction between routing protocol convergence and LDP convergence.
- Describe the packet-mode MPLS convergence after link failure and link recovery.

# Steady State Description



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

© 2002, Cisco Systems, Inc.

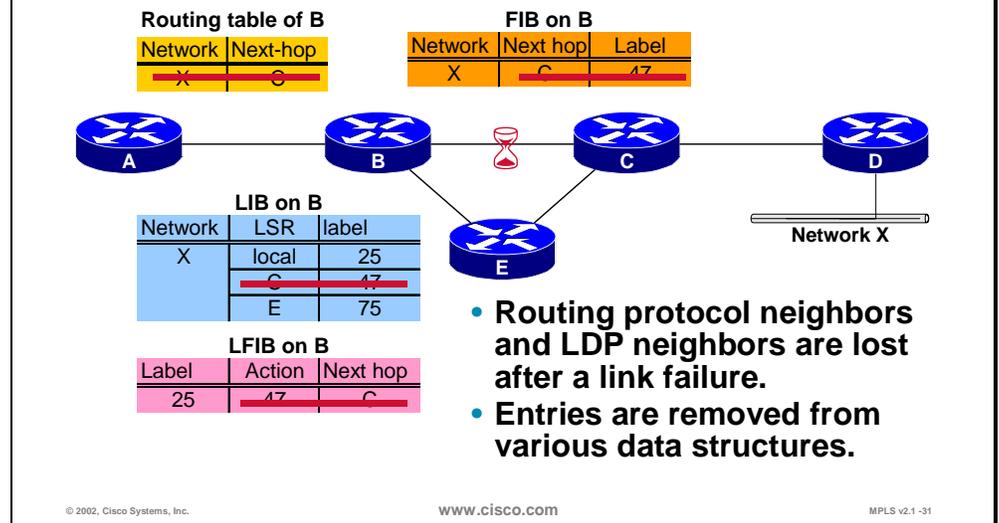
www.cisco.com

MPLS v2.1-30

MPLS is fully functional once the IGP and LDP have populated all the tables: the main IP routing table, the LIB table, the FIB table and the LFIB table.

Although it takes longer for LDP to exchange labels (as compared to IGPs), a network can use the FIB table in the mean time, so there is no additional “down-time” while LDP is exchanging labels between adjacent LSRs.

# Link Failure Actions



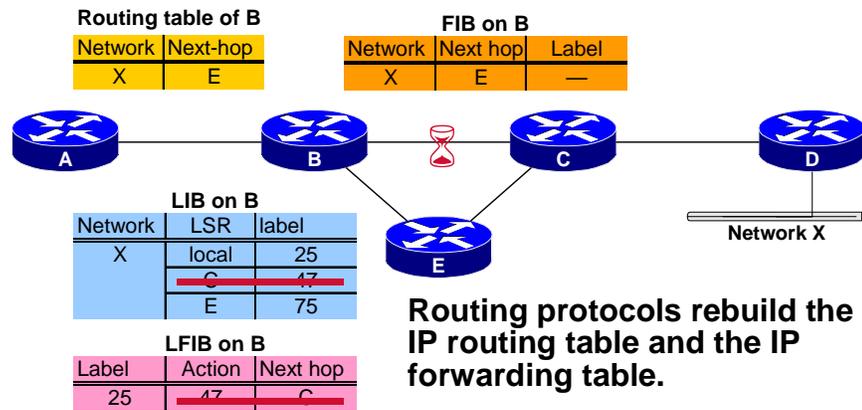
The figure illustrates how a link failure is handled in an MPLS domain:

- The overall convergence fully depends on the convergence of the IGP used in the MPLS domain.
- Once router B determines that router E should be used to reach network X, the label learned from router E can be used to label-switch packets.

LDP stores all labels in the LIB table even if they are not used because the IGP decided to use another path. This is seen in the figure where two next-hop labels were available in the LIB table:

- Label 17 was learned form router C and is currently unavailable so it has to be removed form the LIB table.
- Label 75 was learned form router E and can now be used the moment the IGP decides that router E is the next-hop for network X.

# Routing Protocol Convergence



© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1 -32

The figure illustrates how two entries are removed when the link between routers B and C fails:

- Router B removes the entry from the FIB table once the IGP determines that the next-hop is no longer reachable.
- Router B removes the entry from the LIB table and the LFIB table once the LDP determines that router C is no longer reachable.

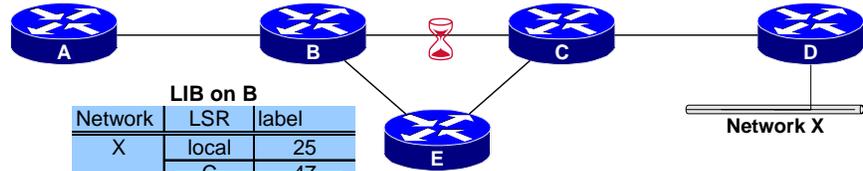
# MPLS Convergence

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
	<del>C</del>	<del>47</del>
	E	75

LFIB on B

Label	Action	Next hop
25	75	E

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

© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1-33

After the IGP has determined that there is another path available, a new entry is created in the FIB table. This new entry points towards router E and there is already a label available for network X via router E. This information is then used in the FIB and the LFIB table to reroute the LSP to router E.

## MPLS Convergence After a Link Failure

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

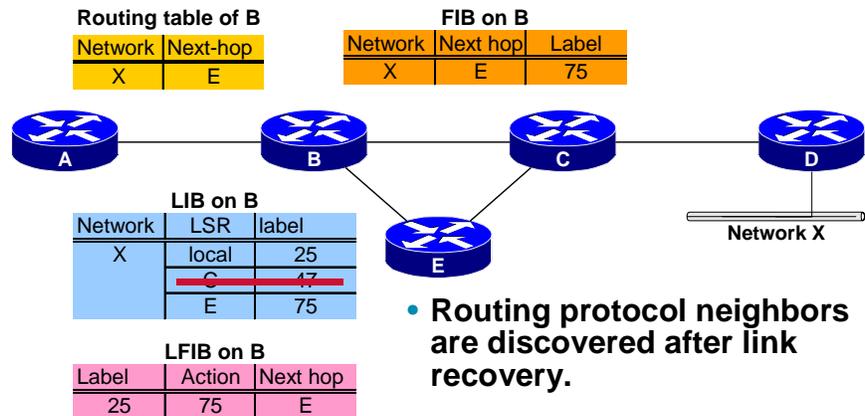
© 2002, Cisco Systems, Inc.

[www.cisco.com](http://www.cisco.com)

MPLS v2.1 -34

The overall convergence in an MPLS network is not impacted by the LDP convergence when a link fails. Frame-mode MPLS uses Liberal Retention Mode that enables routers to store all received labels even if they are not being used. These labels can be used after the network convergence to enable immediate establishment of an alternate label switched path (LSP).

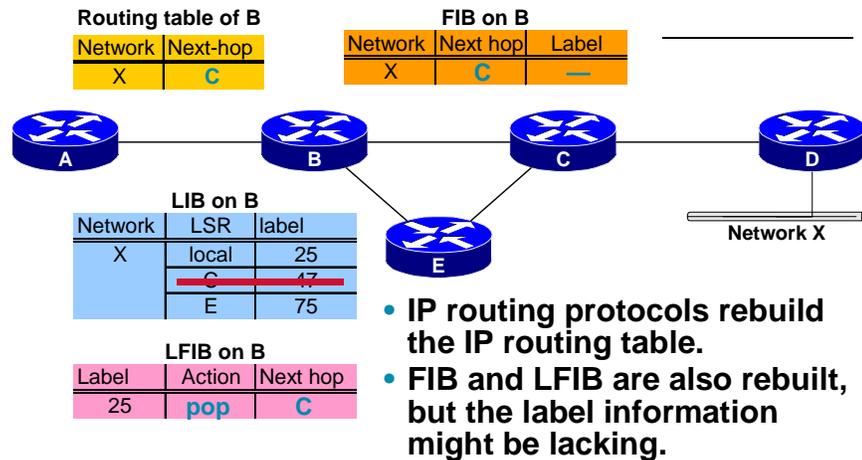
# Link Recovery Actions



The example in the figure illustrates the actions of routers when the link between routers B and C becomes available again.

The IGP determines that the link is available again and changes the next-hop address for network X to point to router C. However, when router B also tries to set the next-hop label for network X, it has to wait for the LDP session between routers B and C to re-establish.

# IP Routing Convergence After Link Recovery



© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1-36

*Pop* action is used in the LFIB on router B while the LDP establishes the session between routers B and C. This process adds to the overall convergence time in an MPLS domain. The downtime for network X is not influenced by LDP convergence because normal IP forwarding is used until the new next-hop label is available.

**Note** While this behavior has no significant impact on traditional IP routing, it can significantly influence the MPLS VPN networks, because the VPN traffic cannot be forwarded before the LDP session is fully operational.

## MPLS Convergence After a Link Recovery

- Routing protocol convergence optimizes the forwarding path after a link recovery.
- 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.**

© 2002, Cisco Systems, Inc.

[www.cisco.com](http://www.cisco.com)

MPLS v2.1-37

Link recovery requires that an LDP session be re-established, which adds to the convergence time of LDP. Networks may be temporarily unreachable due to the convergence limitations of routing protocols.

MPLS Traffic Engineering can be used to prevent longer downtime when a link fails or recovers.

# Summary

The convergence time in MPLS networks is influenced mainly by the convergence time of the IGP that is used in the networks. LDP convergence, however, must be considered from two perspectives:

- Upon link failure, LDP usually already has labels for secondary paths. LSP convergence, therefore, depends solely on the IGP convergence.
- Upon link recovery, the LDP session must be re-established adding to the overall convergence of LSPs.

## Lesson Review

1. What is the impact of LDP/TDP convergence time on overall convergence when a link is lost?
2. What is the impact of LDP/TDP convergence time on overall convergence when a link is restored?

# Typical Label Distribution over LC- ATM Interfaces and VC-Merge

---

## Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Describe ATM-related issues that dictate additional needs for label-distribution procedures over LC-ATM interfaces.
- Describe TDP/LDP label distribution in cell-mode MPLS environments.
- Identify the need for per-edge-LSR VC allocation over ATM network.
- Describe the functions, benefits and drawbacks of VC-Merge.

## Issues of Cell-Mode MPLS Environments

- **MPLS label is encoded as the VPI/VCI value in cell-mode MPLS environments.**
- **Each VPI/VCI combination represents a Virtual Circuit in ATM.**
- **The number of Virtual Circuits supported by router and switch hardware is severely limited.**
- **Conclusion:** labels in cell-mode MPLS are a scarce resource.

© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1 -42

Cell-mode MPLS is significantly different from frame-mode MPLS because of some ATM specific requirements:

- ATM uses cells and not frames. A single packet may be encapsulated into multiple cells. Cells are of a fixed length, which means that normal labels cannot be used because they would increase the size of a cell. The VPI/VCI field in the ATM header is used as an MPLS label. An LSP is, therefore, called a virtual circuit (VC) in ATM terminology.
- ATM switches and routers usually have a limited number of virtual circuits that they can use. MPLS establishes a full mesh of LSPs (VCs), which can result in an extremely large number of LSPs.

Additional mechanisms must be used due to the limitations of ATM hardware.

# Building the IP Routing Table

Routing table of A

Network	Next-hop
X	C

Routing table of C

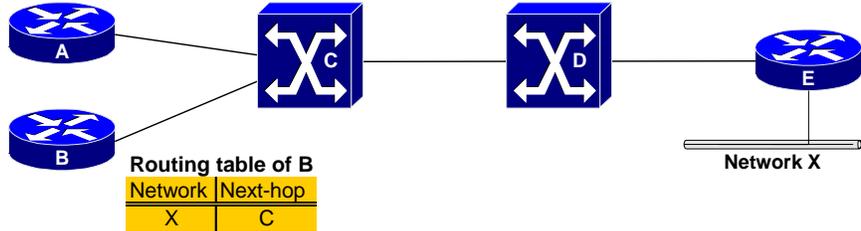
Network	Next-hop
X	D

Routing table of D

Network	Next-hop
X	E

Routing table of E

Network	Next-hop
X	conn



- IP routing protocols are used to build IP routing tables on all LSRs.
- The routing tables are built as if the ATM switches were regular routers.

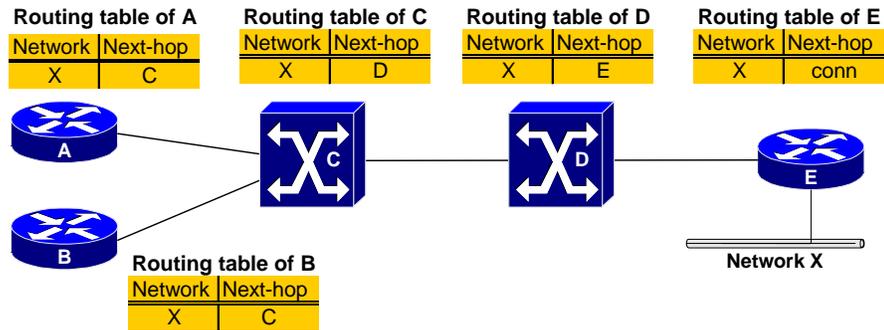
© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1-43

The figure shows how IP and MPLS-aware ATM switches exchange IP routing information with routers. On the control plane, each ATM switch acts as an IP router and is thus seen as an extra IP hop in the network.

# Building the IP Forwarding Table



- Unlabeled IP packets cannot be propagated across LC-ATM interfaces.
- Forwarding tables are not built until the labels are assigned to destinations in IP routing tables.

© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1 -44

Since ATM switches are not able to forward IP packets, labels cannot be asynchronously assigned and distributed. Instead, an ordered sequence of requests is initiated by the router on the upstream side of the ATM network.

# Requesting a Label

Routing table of A

Network	Next-hop
X	C

Routing table of C

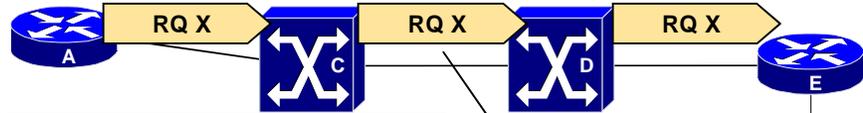
Network	Next-hop
X	D

Routing table of D

Network	Next-hop
X	E

Routing table of E

Network	Next-hop
X	conn



An ATM switch can only allocate an incoming label if it already has a corresponding outgoing label.

A switch requests a label from its next-hop.

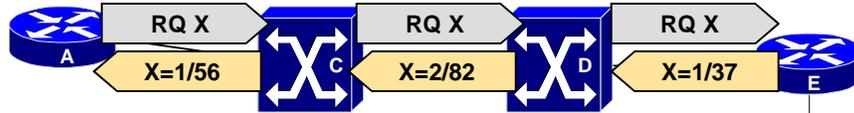
- Labels have to be explicitly requested over LC-ATM interfaces.
- A router requests a label for every destination in the routing table with the next-hop reachable over an LC-ATM interface.

The request from router A is sent to the ATM switch C. Because the ATM switch cannot perform IP lookup, the switch is not allowed to reply with the local label unless it already has the next-hop label. If the switch does not have the next-hop label it must forward the request to the next downstream neighbor.

Once the request reaches router E a reply can be sent because the cell-mode part of the network ends on router E (which therefore acts as an ATM edge LSR).

# Allocating a Label

Routing table of A		Routing table of C		Routing table of D		Routing table of E	
Network	Next-hop	Network	Next-hop	Network	Next-hop	Network	Next-hop
X	C	X	D	X	E	X	conn



LFIB on C			LIB on D			LIB on E		
Label	Action	Next hop	Network	LSR	label	Network	LSR	label
1/56	2/82	D	X	E	1/37	X	local	1/37
				local	2/82			

LFIB on D			LFIB on E		
Label	Action	Next hop	Label	Action	Next hop
2/82	1/37	—	1/37	pop	—

LFIB is actually the ATM switching matrix.

ATM LSR can allocate an incoming label after receiving outgoing label. It replies with the allocated label to the incoming request.

Egress ATM edge LSR allocates a label and replies to the request.

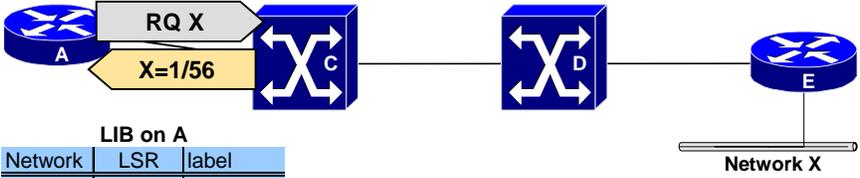
Router E replies with its local label 1/37. The ATM switch D can now generate and use its local label 2/82. Switch C receives the next-hop label from switch D and forwards its own local label 1/56 to router A.

As seen in the figure, an ordered sequence of downstream requests is followed by an ordered sequence of upstream replies. This type of operation is called a “Downstream on demand” allocation of labels.

# Processing Label Allocation Reply

Routing table of A

Network	Next-hop
X	C



LIB on A

Network	LSR	label
X	C	1/56

FIB on A

Network	Next hop	Label
X	C	1/56

**Ingress ATM edge LSR requesting a label inserts the received label in its LIB, FIB and (optionally) LFIB.**

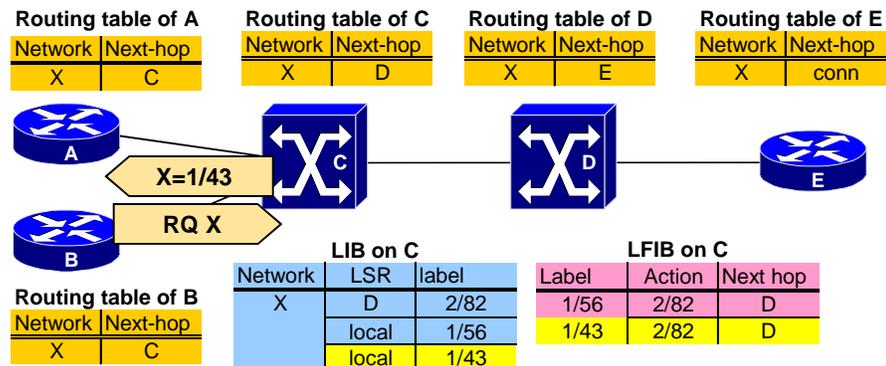
© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1-47

Processing of LDP replies on router A (also ATM edge LSR) is similar to the processing in frame-mode MPLS; the received label is stored in the LIB, the FIB and the LFIB tables.

# Allocation Requests from Additional Upstream LSRs



Every upstream LSR will request a label for downstream destinations from an ATM LSR.

ATM LSR could reuse an already allocated downstream label for the second upstream label.

© 2002, Cisco Systems, Inc.

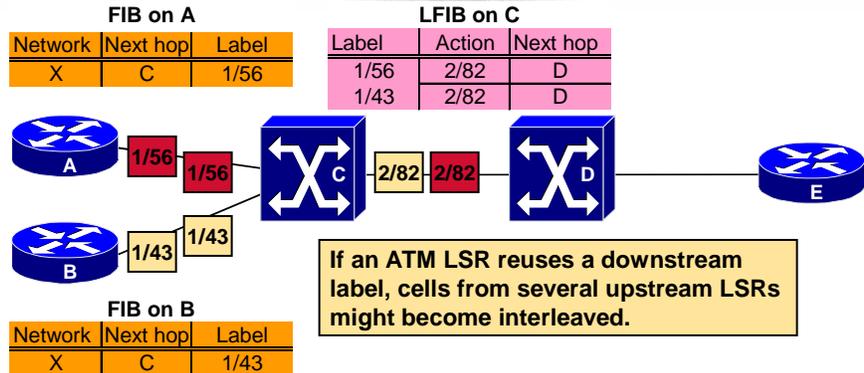
www.cisco.com

MPLS v2.1 -48

The figure shows how another router requests a label for the same destination that router A has previously requested. The ATM switch C already has a next-hop label for network X and can therefore, immediately reply to router B.

The figure also shows that the switch used a different local label than the one sent to router A because ATM switches use per-interface VPI/VCI values and can now also use per-interface label space.

# Cell Interleave Issue



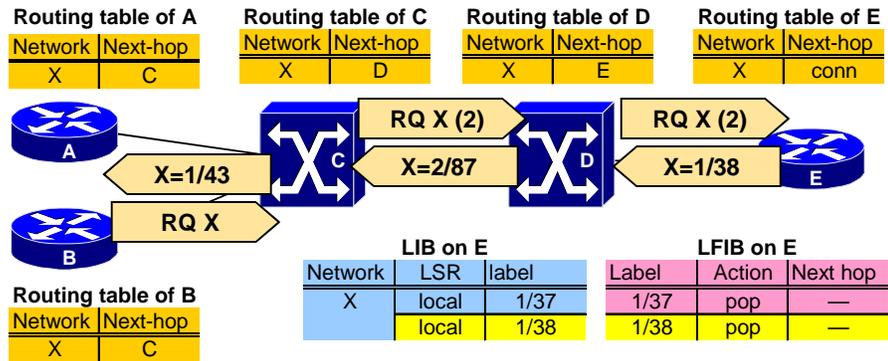
- **Solution #1 - allocate a separate downstream label for every upstream request.**
- **Solution #2 - prevent cell interleave by blocking incoming cells until a whole frame is collected.**

By analyzing the previous two figures it can be noted that an unusual situation has developed. Two virtual circuits from routers A and B (1/56 and 1/43) merge into one (2/82). Standard ATM virtual switching hardware does not support this and as a result segmented packets from two sources may become interleaved between the ATM switches C and D. The receiving router E is then unable to correctly reassemble those cells into two packets.

There are two possible solutions to this problem:

- Allocate a new downstream label for each request. This solution would result in an even greater number of labels.
- Buffer cells of the second packet until all cells of the first packet are forwarded. This solution results in an increased delay of packets because of buffering.

# Additional Label Allocation



ATM LSR requests a new label from downstream LSRs for every upstream request.

ATM egress router has to allocate a unique label for every ATM ingress router for every destination.

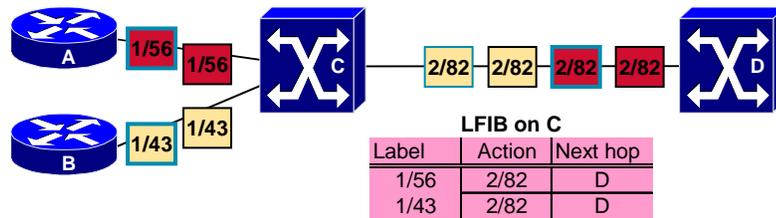
© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1-50

The figure illustrates the first option where an additional LSP is created for the same destination network X for every upstream ATM edge LSR. ATM switch C now has two next-hop labels for network X, one for source router A, the other for source router B.

# Virtual Circuit Merge



- VC-Merge is a solution in which incoming cells are blocked until the last cell in a frame arrives.
- All buffered cells are then forwarded to the next-hop ATM LSR.

© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1-51

The figure illustrates the second option where the ATM switch C buffers cells coming from router B until the last cell of the packet coming from router A is forwarded.

This option reduces the number of labels (VCs) needed in the ATM network but increases the average delay across the network.

# Benefits and Drawbacks of VC-Merge

## Benefits of VC-Merge

- Merging ATM LSR can reuse the same downstream label for multiple upstream LSRs.

## Drawbacks of VC-Merge

- Buffering requirements increase on the ATM LSR.
- Jitter and delay across ATM network increase.
- ATM network is effectively transformed into a frame-based network.

© 2002, Cisco Systems, Inc.

[www.cisco.com](http://www.cisco.com)

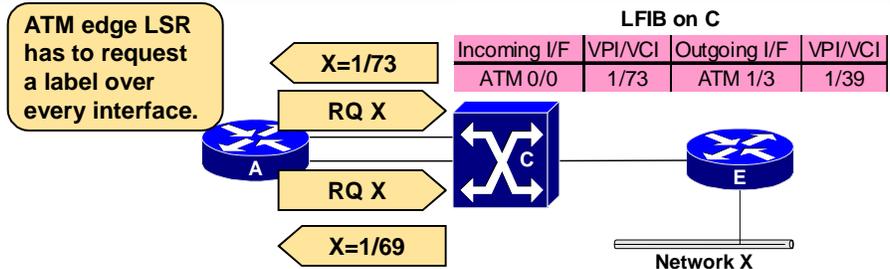
MPLS v2.1 -52

The option that buffers cells is usually referred to as “VC-Merge” because it implements the merging of two virtual circuits into one while preventing cell interleaving.

The major benefit of VC-Merge is that it minimizes the number of labels (VPI/VCI values) needed in the ATM part of the network.

The major drawback is the increase in the average delay and jitter in the ATM network. ATM networks under heavy load become more like frame-based networks.

# Per-interface Label Allocation



- LFIB on an ATM switch (ATM switching matrix) always contains the incoming interface.
- Labels have to be assigned for individual interfaces – **per-interface label allocation**.
- The same label can be reused (with a different meaning) on different interfaces.

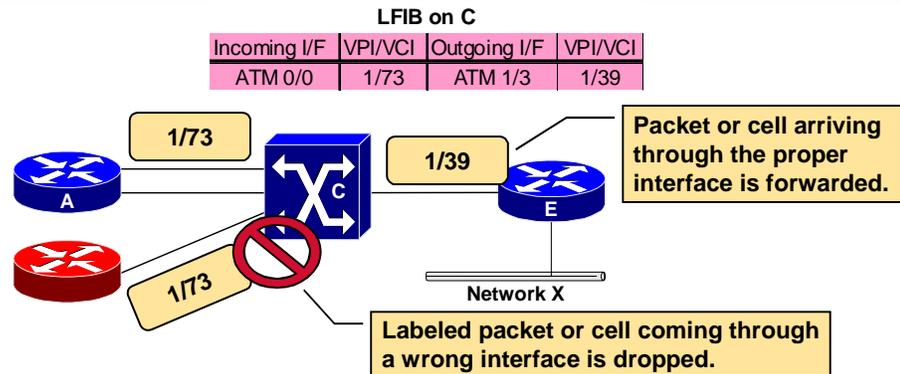
© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1-53

Cell-mode MPLS defaults to using per-interface label space because ATM switches support per-interface VPI/VCI values to encode labels. Therefore, if a single router has two parallel links to the same ATM switch, two LDP sessions are established and two separate labels are requested.

# Security of Per-interface Label Allocation



One of the benefits of per-interface label space is that it prevents label spoofing. The figure shows how the red (bottom) router tried to send a cell with a label that was advertised only to router A. The switch failed to forward the cell because it came in through the wrong interface.

To summarize the forwarding differences between frame-mode and cell-mode MPLS:

- Frame-mode MPLS forwards packets based solely on the label.
- Cell-mode MPLS forwards cells based on the incoming interface and the label (VPI/VCI field).

# Summary

MPLS-aware ATM switches use an IP routing protocol and LDP to exchange routing information and labels. VPI/VCI fields in the ATM header are used to encode MPLS labels.

ATM LSRs use the downstream on demand allocation of labels where an ordered sequence of requests is sent to the other end of the ATM network (per-edge-LSR allocation of VCs) or the first ATM switch that already has the next-hop label (VC-Merge).

VC-Merge is used to minimize the number of required labels because most router and ATM switches only support a limited number of virtual circuits. VC-merge on the other hand introduces greater delay to packets because its cells may be buffered in ATM switches.

## Lesson Review

1. What is ATM switching matrix called in MPLS terminology?
2. Why is it necessary to have the next-hop label before propagating the local label?
3. What are the benefits of VC-Merge?
4. What are the drawbacks of VC-Merge?



# MPLS Label Allocation, Distribution and Retention Modes

---

## Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Describe the difference between per-interface and per-platform label space.
- Describe TDP/LDP unsolicited downstream and downstream-on-demand label distribution.
- Describe the difference between ordered and independent label allocation control.
- Describe the difference between conservative and liberal retention mode.
- Identify parameter sets used in Cisco IOS MPLS implementation.

# Label Distribution Parameters

## MPLS architecture defines several label allocation and distribution parameters:

- **Per-interface or per-platform label space**
- **Unsolicited downstream and downstream on demand label distribution**
- **Ordered or Independent Label Switch Path control**
- **Liberal or conservative label retention**

© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1 -59

This section describes the options in label allocation, propagation and retention.

The two label space options are:

- **Per-interface label space** where labels must be unique for a specific input interface.
- **Per-platform label space** where labels must be unique for the entire platform (router).

The two options of label generation and distribution are:

- **Unsolicited downstream** distribution of labels is used on frame-mode MPLS where all routers can asynchronously generate their local labels and propagate them to adjacent routers.
- **Downstream on demand** distribution of labels is used where ATM LSRs have to request a label for destinations found in the IP routing table.

Another aspect of label distribution focuses on the (a)synchronicity of label distribution:

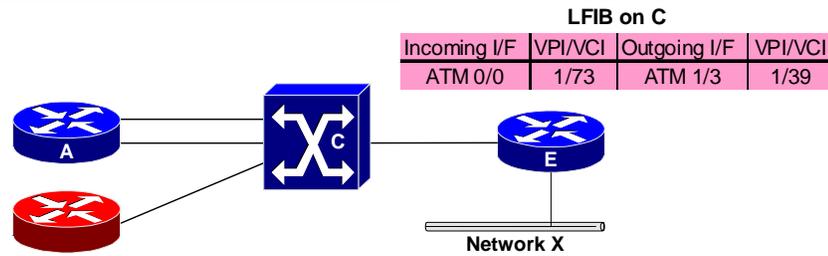
- Frame-mode MPLS uses **Independent control mode** where all routers can start propagating labels independently of each other.
- Cell-mode MPLS requires LSRs to already have the next-hop label if they want to generate and propagate their own local labels. This is called **Ordered control mode**.

The last aspect of label distribution looks at labels that are received but not used:

- Frame mode MPLS may result in multiple labels being received but only one being used. Unused labels are kept and this is usually referred to as **Liberal retention mode**.

- Cell-mode MPLS only keeps labels that it previously requested. This mode is called **Conservative retention mode**.

# Label Space: Per Interface



- **LFIB on an LSR contains incoming interface.**
- **Labels have to be assigned for individual interfaces.**
- **The same label can be reused (with a different meaning) on different interfaces.**
- **Label allocation is secure – LSRs cannot send packets with labels that were not assigned to them.**

© 2002, Cisco Systems, Inc.

www.cisco.com

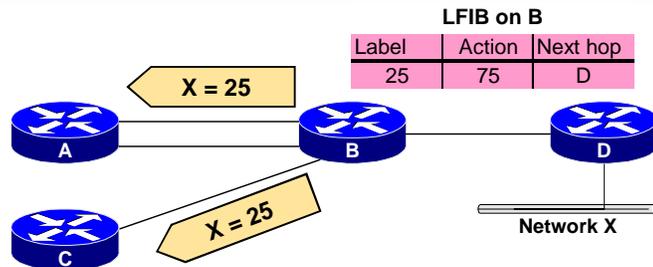
MPLS v2.1-60

The figure shows the main difference between the LFIB table used with frame-mode MPLS and the LFIB table (ATM switching matrix) used with cell-mode MPLS.

The LFIB table used with cell-mode MPLS maps a local label bound to an input interface to a next-hop label pointing to the outgoing interface. The label assigned to an input interface can be reused on another interface and it can have a different meaning (assigned to a different destination).

Per-interface label space prevents label spoofing by not allowing cell forwarding for labels (VPI/VCI values) that are not bound to the interface where the cell was received.

# Label Space: Per Platform



- LFIB on a 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-platforms label-space is less secure than per-interface label space.

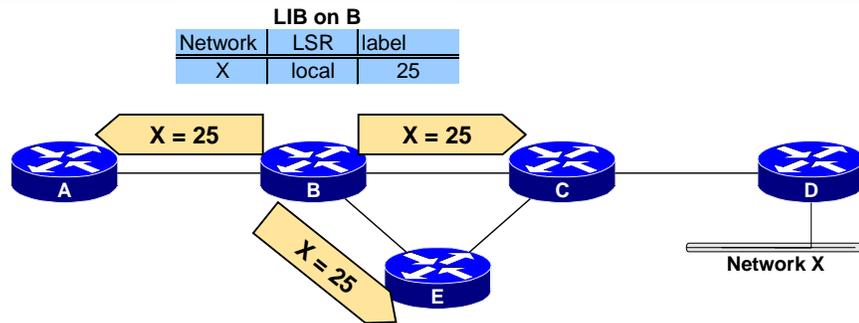
© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1-61

Per-platform label space is used with frame-mode MPLS where one single label is assigned to a destination network and sent to all LDP peers. This label can then be used on any incoming interface. The per-platform label space minimizes the number of LDP sessions and allows upstream label switch paths to span parallel links because the same label is used on all of them. However, per-platform label space is less secure than the per-interface label space because untrusted routers can use labels that were never allocated to them.

# Label Distribution: Unsolicited Downstream



**Label for a prefix is allocated and advertised to all neighbor LSRs, regardless of whether the neighbors are upstream or downstream LSRs for the destination.**

© 2002, Cisco Systems, Inc.

www.cisco.com

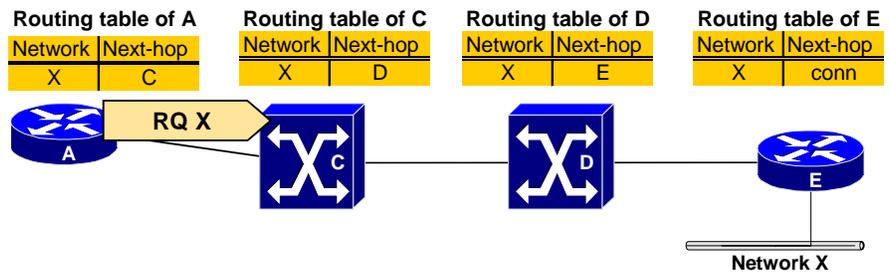
MPLS v2.1 -62

Unsolicited downstream allocation of labels is a method where each router independently assigns a label to a destination IP prefix and sends this mapping to all LDP peers. There is no control mechanism to govern the propagation of labels in an ordered fashion.

The figure illustrates how router B creates a local label and sends it to all neighbors. The same action is taken on other routers after the IGP has put network X into the main routing table.

Each neighbor then decides whether to use the label (if router B is the closest next-hop for network X), keep it in the LIB table or ignore it.

# Label Distribution: Downstream on Demand



- A LSR will only assign a label to a prefix when asked for a label by an upstream LSR.
- Label distribution is a hop-by-hop parameter – different label distribution mechanisms can coexist in an MPLS network.

© 2002, Cisco Systems, Inc.

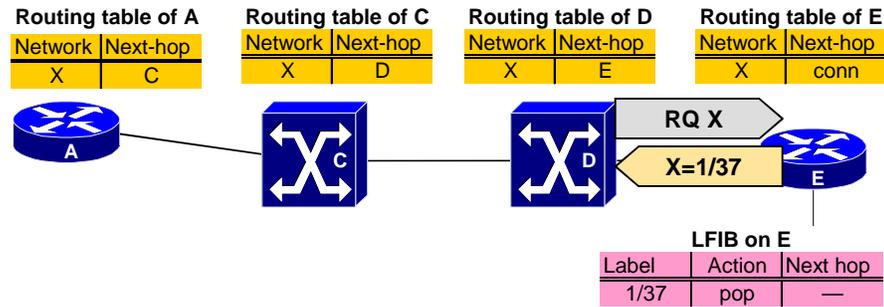
www.cisco.com

MPLS v2.1-63

Downstream on demand allocation of labels requires each LSR to specifically request a label from its downstream neighbor. The figure shows how router A requests a next-hop label from its downstream LDP peer.

Unsolicited downstream and downstream on demand can be combined because labels are assigned and propagated hop-by-hop. The usual situation is that frame-mode MPLS uses unsolicited downstream propagation and cell-mode MPLS uses downstream on demand.

# LSP Control: Independent Control



- A LSR can always assign a label for a prefix, even if it has no downstream label.
- Independent control can only be used for LSRs with layer-3 capabilities.

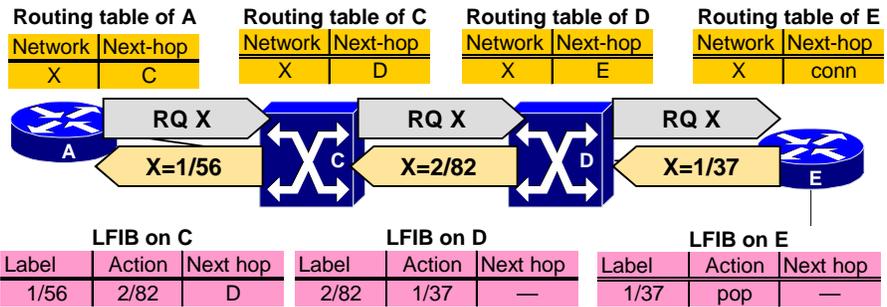
© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1 -64

Independent control mode is usually combined with unsolicited downstream propagation of labels where labels can be created and propagated independently of any other LSR. When using independent control mode, a LSR might be faced with an incoming labeled packet where there is no corresponding outgoing label in the LFIB table. A LSR using independent control mode must therefore be able to perform full layer-3 lookups. Independent control mode can only be used on LSRs with edge LSR functionality.

# LSP Control: Ordered Control



**A LSR can only assign a label if it has already received a label from the next-hop LSR; otherwise it must request a label from the next-hop LSR.**

© 2002, Cisco Systems, Inc.

www.cisco.com

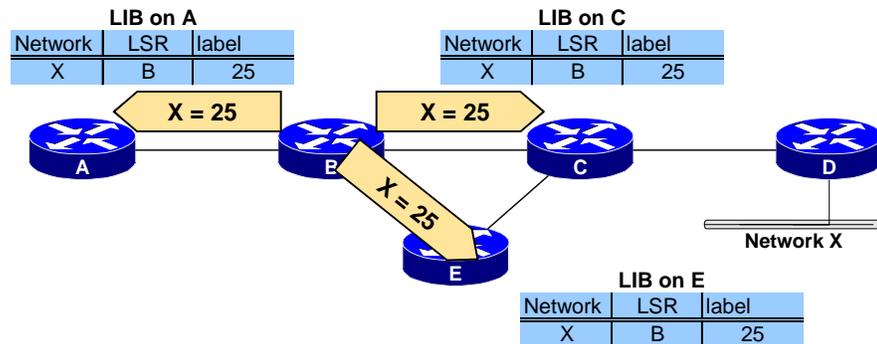
MPLS v2.1-65

Ordered control mode is usually combined with downstream on demand propagation of labels where a local label can be assigned and propagated only if a next-hop label is available. This requirement results in an ordered sequence of downstream requests until an LSR is found that already has a next-hop label or an LSR is reached that uses independent control mode.

While the ordered control mode could be used with frame-mode MPLS, its usage is mandatory on ATM switches which cannot perform layer-3 lookups.

The figure illustrates how both ATM LSRs forward requests until an edge is reached. The edge LSR uses independent control mode and can respond to the request.

# Label Retention: Liberal Retention Mode



- Every LSR stores the received label in its LIB, even when the label is not received from a next-hop LSR.
- Liberal retention mode improves convergence speed.

© 2002, Cisco Systems, Inc.

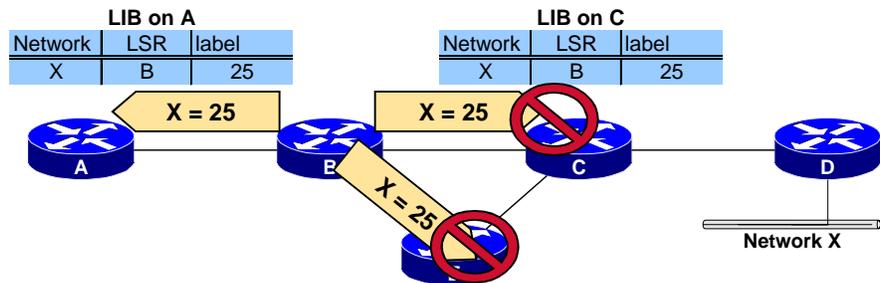
www.cisco.com

MPLS v2.1 -66

Liberal retention mode dictates that every LSR keeps all labels received from LDP peers even if they are not the downstream peers for network X.

The figure shows how router C receives and keeps the label received from router B, for network X, although router D is the downstream peer.

# Label Retention: Conservative Retention Mode



- LSR stores only the labels received from next-hop LSRs; all other labels are ignored.
- Downstream-on-demand distribution is required during the convergence phase.

© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1-67

Conservative retention mode, only keeps labels that can immediately be used.

The figure illustrates how routers B and C do not consider router B to be the next-hop for network X and therefore, drop the labels received from router B.

---

**Note** Conservative retention mode requires downstream on demand label allocation after the network convergence.

---

# Standard Parameter Sets in Cisco IOS MPLS Implementation

## Routers with packet interfaces

- Per-platform label space, unsolicited distribution, liberal retention, independent control

## Routers with ATM interfaces

- Per-interface label space, on-demand distribution, conservative or liberal retention, independent control

## ATM switches

- Per-interface label space, on-demand distribution, conservative retention, ordered control

© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1-68

The following default operation applies to routers using frame-mode MPLS (LSRs):

- Per-platform label space.
- Unsolicited downstream propagation of labels: Every LSR can propagate a label mapping to its neighbors without a request.
- Liberal retention mode: Allows easy failover if a link fails and does not require a combination of unsolicited downstream and downstream on demand.
- Independent control mode: Makes label propagation faster (less time is needed for LDP convergence) because LSRs do not have to wait to get the next-hop label from their downstream neighbors.

The following default operation applies to ATM switches using cell-mode MPLS (ATM LSRs):

- Per-interface label space because it provides better security and is already available with standard ATM switching functionality.
- Downstream on demand propagation of labels because LFIB tables on ATM switches are really ATM switching matrices that require full information before switching can start. Full information includes a next-hop label, which again must be requested.
- Conservative retention mode implicitly achieved by using the downstream on demand propagation of labels; no label is received unless it is requested.
- Ordered control mode is used in combination with downstream on demand propagation of labels to make sure every ATM LSR has all the information needed to create an entry in the LFIB table (ATM switching matrix) including the next-hop label.

The default operation of routers using cell-mode MPLS (ATM edge LSRs) is similar to those of ATM switches except that they use independent control mode as they are the endpoints of the virtual circuits.

# Summary

This section describes the options in label allocation, propagation and retention.

There are two available label space schemes:

- **Per-interface label space** where labels must be unique for a specific input interface.
- **Per-platform label space** where labels must be unique for the entire platform (router).

There are two available label propagation schemes:

- **Unsolicited downstream** distribution of labels is used on frame-mode MPLS where all routers can asynchronously generate their local labels and propagate them to adjacent routers.
- **Downstream on demand** distribution of labels is used where ATM LSRs must request a label for destinations found in the IP routing table.

There are two available types of label propagation control:

- Frame-mode MPLS uses **Independent control mode** where all routers can start propagating labels independently of each other.
- Cell-mode MPLS requires LSRs to already have the next-hop label if they want to generate and propagate their own local labels. This is called **Ordered control mode**.

There are two available label retention schemes:

- Frame mode MPLS may result in multiple labels being received but only one used. Unused labels are kept and this is usually referred to as **Liberal retention mode**.
- Cell-mode MPLS on the other hand only keeps labels which it previously requested. This is called **Conservative retention mode**.

## Lesson Review

1. Why and where is downstream-on-demand label propagation used?
2. What are the benefits and drawbacks of liberal retention mode?
3. What are the benefits and drawbacks of per-platform label space?
4. Why is per-interface label space on ATM LSRs needed?

# LDP Neighbor Discovery

---

## Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Describe LDP/TDP neighbor discovery.
- Describe LDP/TDP session establishment process.

# LDP Session Establishment

- **LDP and TDP use a similar process to establish a session:**
  - **Hello messages** are periodically sent on all interfaces enabled for MPLS.
  - If there is another router on that interface 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).**

© 2002, Cisco Systems, Inc.

www.cisco.com

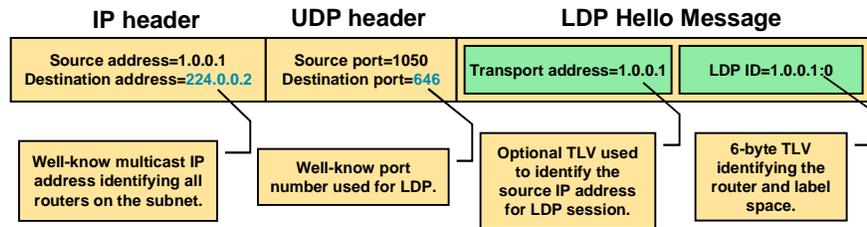
MPLS v2.1-73

Label Distribution Protocol (LDP) is a standard protocol used to exchange labels between adjacent routers. Tag Distribution Protocol is a Cisco-proprietary protocol with the same functionality as LDP.

Although the remainder of the document focuses on LDP it should be noted that TDP, as LDP's predecessor, works in a similar fashion.

LDP periodically send **Hello messages**. The Hello messages use UDP packets with a multicast destination address 224.0.0.2 (“all routers on a subnet”) and destination port number 646 (711 for TDP). If another router is enabled for LDP/TDP it will respond by opening a TCP session with the same destination port number (646 or 711).

# LDP Hello Message



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

© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1-74

The contents of a Hello message are:

- Destination IP address (224.0.0.2) which targets all routers on the subnet.
- Destination port equals LDP's well-known port number 646
- The actual Hello message may optionally contain a Transport Address TLV to instruct the peer to open the TCP session to the Transport Address instead of to the source address that is found in the IP header.
- The LDP identifier is used to uniquely identify the neighbor and the label space. Multiple sessions can be established between a pair of LSRs if they use multiple different label spaces.

## 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 zero (for example, LDP ID=1.0.0.1:0).
- A combination of frame-mode and cell-mode, or multiple cell-mode links, results in multiple LDP sessions.

© 2002, Cisco Systems, Inc.

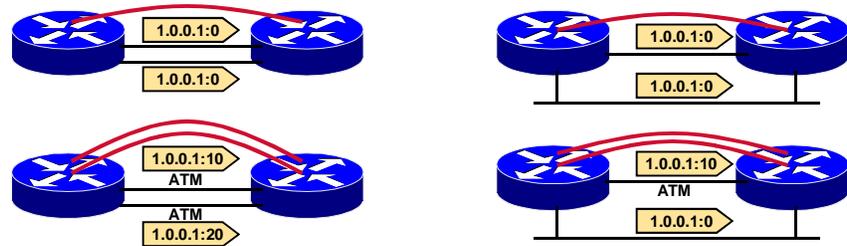
www.cisco.com

MPLS v2.1-75

If a pair of routers is connected over two or more parallel links and use frame-mode MPLS they try to establish multiple sessions using the same LDP identifier. This results in only one session remaining.

If, the two routers use different LDP identifiers, for example one link uses frame-mode MPLS and the other uses cell-mode, they will keep both sessions.

# Label Space Negotiation Example



- 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, since ATM LSRs do not use Ethernet for packet forwarding, and frame-mode MPLS across ATM uses per-platform label space.

© 2002, Cisco Systems, Inc.

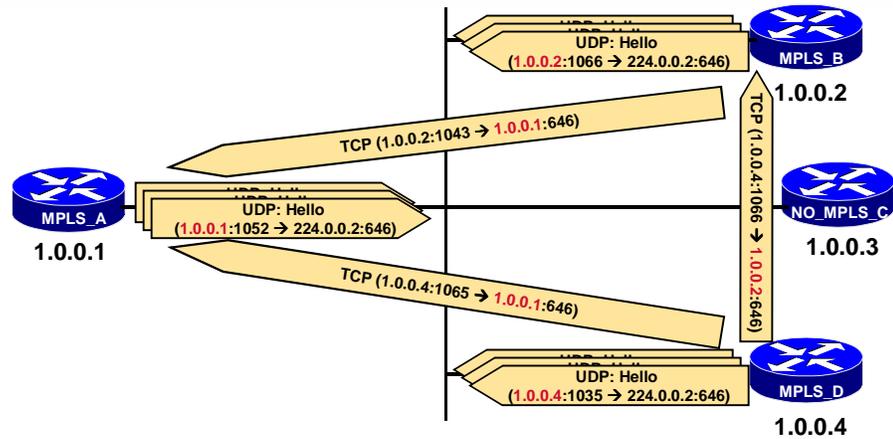
www.cisco.com

MPLS v2.1-76

The figure illustrates four different combinations with two parallel links between a pair of routers.

A general rule can be extracted from the four examples: an LDP session is established per interface except for all frame-mode interfaces where only one LDP session between a pair of LSRs suffices.

# LDP Neighbor Discovery



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

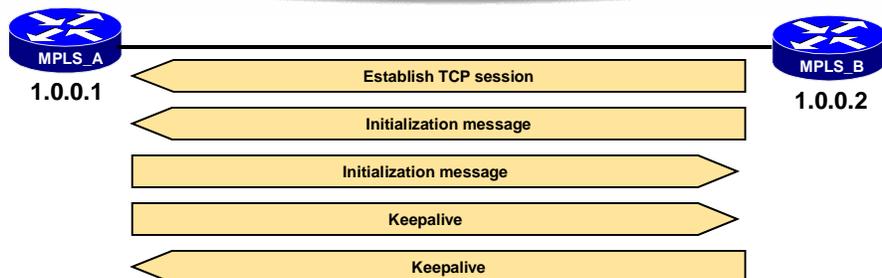
© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1 -77

This figure illustrates how three out of four routers periodically send out LDP Hello messages (the fourth routers does not run MPLS). The routers with higher IP addresses must initiate the TCP session. After the TCP session has established they still keep sending the LDP Hello messages to potentially discover new peers or identify failures.

# LDP Session Negotiation



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

© 2002, Cisco Systems, Inc.

www.cisco.com

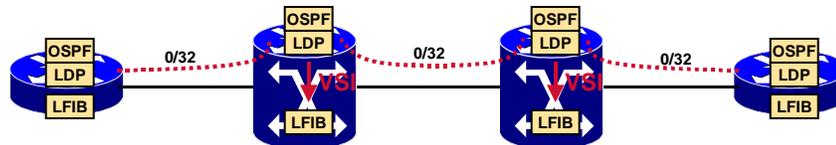
MPLS v2.1-78

The LDP session negotiation can be split into three steps:

- Establishment of the TCP session
- Exchange of initialization messages
- Exchange of initial keepalive messages

After these three steps, the two peers can start exchanging labels for networks that they have in the main routing table.

# LDP Sessions between ATM LSRs



- An IP adjacency between ATM LSRs is established through the control VC (0/32).
- The control VC is used for LDP/TDP as well as for IP routing protocols.
- Virtual Switch Interface (VSI) protocol is used to populate the ATM switching matrix (LFIB) in the data plane of some ATM switches (Cisco-specific implementation).

© 2002, Cisco Systems, Inc.

www.cisco.com

MPLS v2.1-79

The figure illustrates the operation of LDP in ATM networks. ATM LSRs establish the IP adjacency across the MPLS control virtual circuit that by default, uses a VPI/VCI value of 0/32.

An IP routing protocol and LDP or TDP use this control VC to exchange IP routing information and labels.

Some Cisco devices use the Virtual Switch Interface (VSI) protocol to create entries in the LFIB table (ATM switching matrix as part of the data plane) based on the information in the LIB table (control plane). This protocol is used to dynamically create virtual circuits for each IP network.

## LDP/TDP Discovery of Non-adjacent Neighbors

- **LDP/TDP neighbor discovery of non-adjacent neighbors differs from normal discovery only in the addressing of hello packets.**
- **Hello packets use unicast IP addresses instead of multicast addresses.**
- **Once a neighbor is discovered, the mechanism to establish a session is the same.**

© 2002, Cisco Systems, Inc.

[www.cisco.com](http://www.cisco.com)

MPLS v2.1-90

LDP can also be used between non-adjacent routers. LDP Hello messages use unicast IP addresses instead of multicast. The rest of the session negotiation is the same as for adjacent routers.

# Summary

LDP and TDP use a multicast IP address to periodically send Hello messages to all routers reachable through an interface.

The LDP session itself uses TCP to provide reliability.

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

## Lesson Review

1. How do routers find LDP/TDP peers reachable through an interface?
2. Which protocol is used to discover neighboring LSRs?
3. Which protocol is used for the session itself?
4. How do ATM LSRs establish an LDP/TDP session?
5. What is different in an LDP session establishment between non-adjacent peers?

# Penultimate Hop Popping

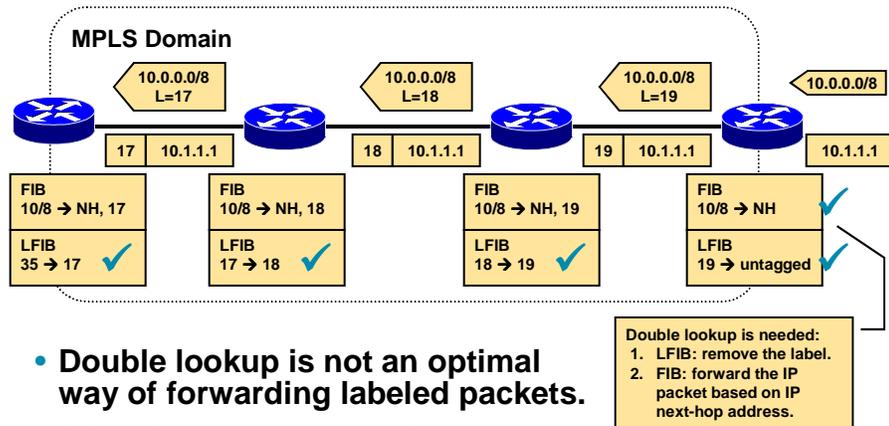
---

## Objectives

Upon completion of this section, you will be able to perform the following tasks:

- Describe penultimate hop popping.
- Describe how LSRs request PHP through TDP/LDP.
- Identify when the PHP could be used in MPLS network.
- Describe the benefits and drawbacks of PHP.

# Double Lookup Scenario



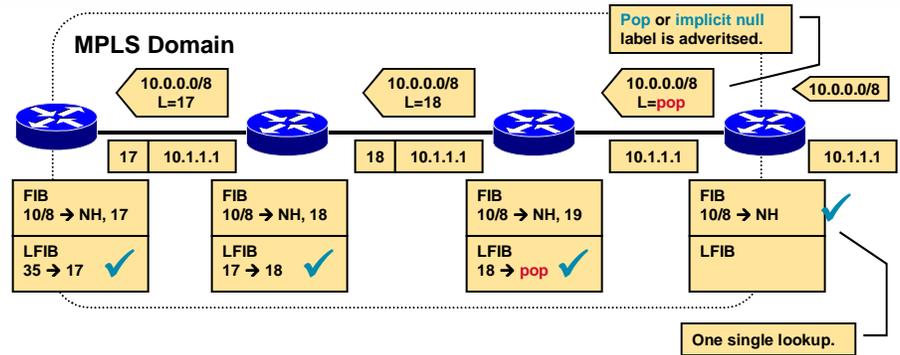
- Double lookup is not an optimal way of forwarding labeled packets.
- A label can be removed one hop earlier.

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

The figure illustrates how labels are propagated and used in a typical frame-mode MPLS network. The checkmarks show which tables are used on individual routers. The egress router in this example must do a lookup in the LFIB table to determine whether the label must be removed and if a further lookup in the FIB table is required.

Penultimate hop popping removes the requirement for double lookup on egress LSRs.

# Penultimate Hop Popping



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

This figure illustrates how a predefined label *pop*, which corresponds to the *pop* action in LFIB, is propagated on the first hop or the last hop, depending on the perspective. The meaning of the label “pop” is to remove the top label in the MPLS label stack instead of swapping it with the next-hop label. The last router before the egress router, therefore, removes the top label.

Penultimate hop popping slightly optimizes MPLS performance by eliminating one LFIB lookup.

## Penultimate Hop Popping (Cont.)

- **Penultimate hop popping optimizes MPLS performance (one less LFIB lookup).**
- **PHP does not work on ATM (VPI/VCI cannot be removed).**
- **Pop or implicit null label uses value 3 when being advertised to a neighbor.**

© 2002, Cisco Systems, Inc.

[www.cisco.com](http://www.cisco.com)

MPLS v2.1 -87

A “pop” label is encoded using a label value of 3. This label instructs upstream routers to remove the label instead of swapping it with label 3.

Penultimate hop popping is not supported on ATM devices because a label is part of the ATM cell payload and cannot be removed by the ATM switching hardware.

# Summary

Penultimate hop popping slightly increases MPLS performance by eliminating one LFIB lookup. PHP uses a well-known label value of 3, which represents a “pop” label. This label instructs neighbors to remove labels when doing label switching on a particular LSP.

Penultimate hop popping can only be used in frame-mode environments. Labels in ATM networks are not removable because they are part of the ATM header (VPI/VCI field).

## Lesson Review

1. What is the main benefit of penultimate hop popping?
2. How does a router know that it has to pop the label?



# Chapter Summary

After completing this chapter, you should be able to perform the following tasks:

- Describe the need for LDP/TDP in an MPLS network.
- Describe the LDP/TDP neighbor discovery and session establishment procedures.
- Explain the needs for different LDP/TDP label distribution modes.
- Explain the difference between independent and ordered control.
- Describe LDP/TDP retention modes (conservative and liberal).
- Explain the functions and benefits of penultimate-hop-popping.

