BROADBAND AND HIGH SPEED NETWORKS

7

Multiprotocol Label Switching (MPLS)

MPLS DATA STRUCTURES

- An MPLS LSR maintains *three* important data structures:
  1. IP Routing Table.
  2. Label Information Base (LIB).
  3. Label Forwarding Information Base (LFIB).

1. IP ROUTING TABLE

- To determine the next-hop forwarding information for the LFIB.
- The content of the IP routing table caused by exchanging information among all routing protocol peers within the network via a *routing algorithm* such as OSPF.
- The result computed by the *routing algorithm* is then stored in the IP routing table.
- The routing table will reflect the layer 3 topology of the network, as well as, the shortest path to reach any given network node.

2. LABEL INFORMATION BASE (LIB)

- LIB stores all labels that have been advertised by other LSRs in the MPLS network.
- LSRs exchange the mapping of labels to their associated FECs.
- This is done using a variety of *label distribution protocols* such as:
  - Label Distribution Protocol (LDP).
  - Resource Reservation Protocol (RSVP).
  - Multiprotocol BGP (MP-BGP).

3. LABEL FORWARDING INFORMATION BASE (LFIB)

- Used by the actual packet forwarding process.
- MPLS LSR that wants to forward a packet will consult LIFB to determine *which label to swap on the packet* as well as *which next-hop interface it must forward the packet on to*.
- The information in the LFIB is the result of examining the IP routing table and the contents of the LIB and determining which IP destinations can be mapped to which MPLS labels.

MPLS DATA STRUCTURES - EXAMPLE

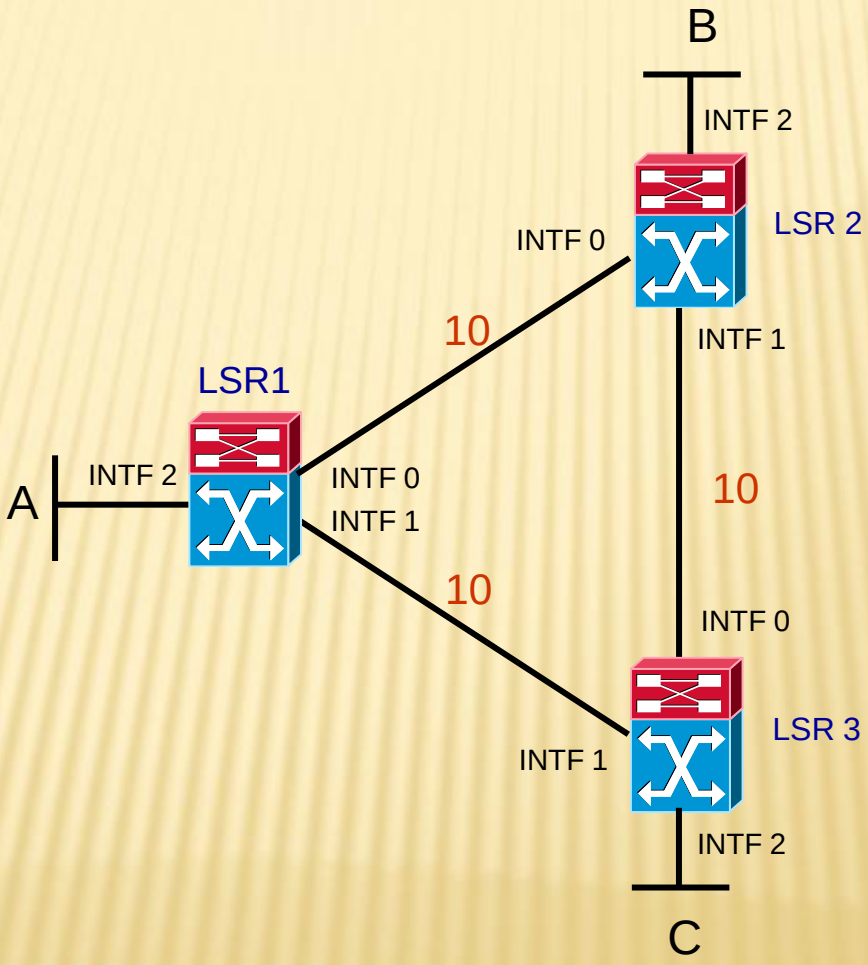

MPLS DATA STRUCTURES - EXAMPLE

Routing Table LSR 1

DEST	METRIC	NEXT HOP	INTERFACE
A	0	Connected	2
B*	10	LSR 2	0
B	20	LSR 3	1
C*	10	LSR 3	1
C	20	LSR 2	0

* Optimum route to destination

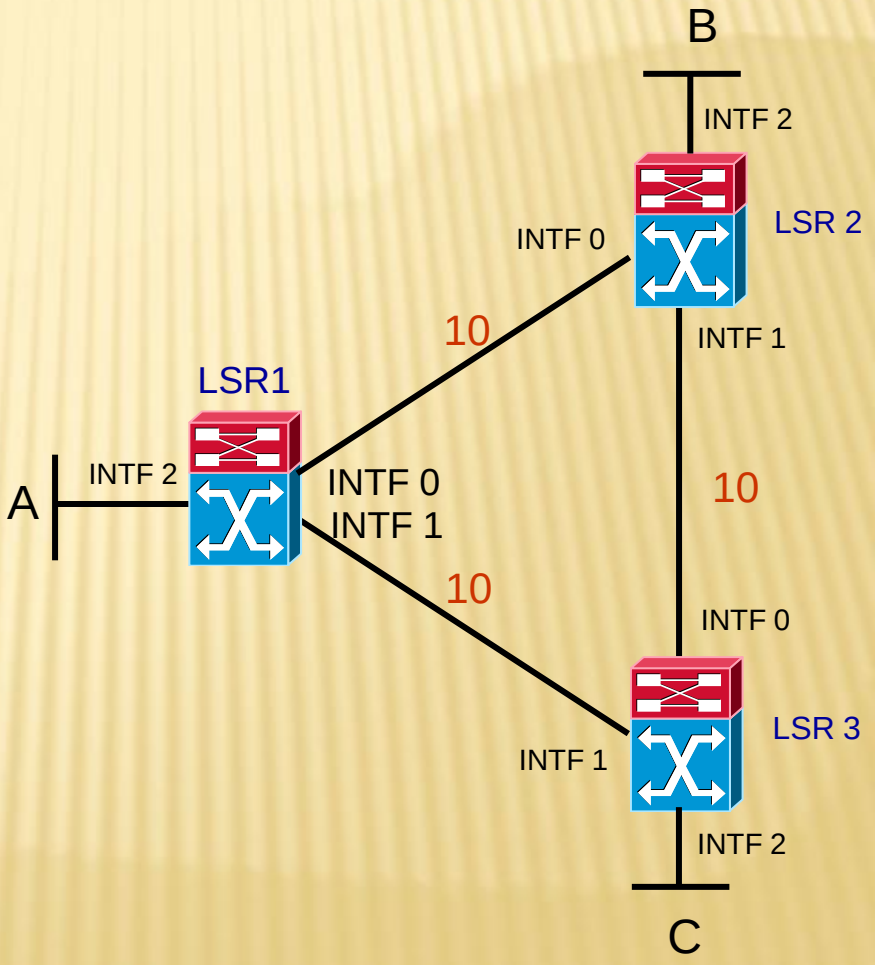

MPLS DATA STRUCTURES - EXAMPLE

LIB LSR 1

DEST	LABEL	ADV. LSR	BINDING
A	13	LSR 1	Local
B	7	LSR 2	Remote
B	8	LSR 3	Remote
B	6	LSR 1	Local
C	9	LSR 2	Remote
C	2	LSR 3	Remote
C	11	LSR 1	Local

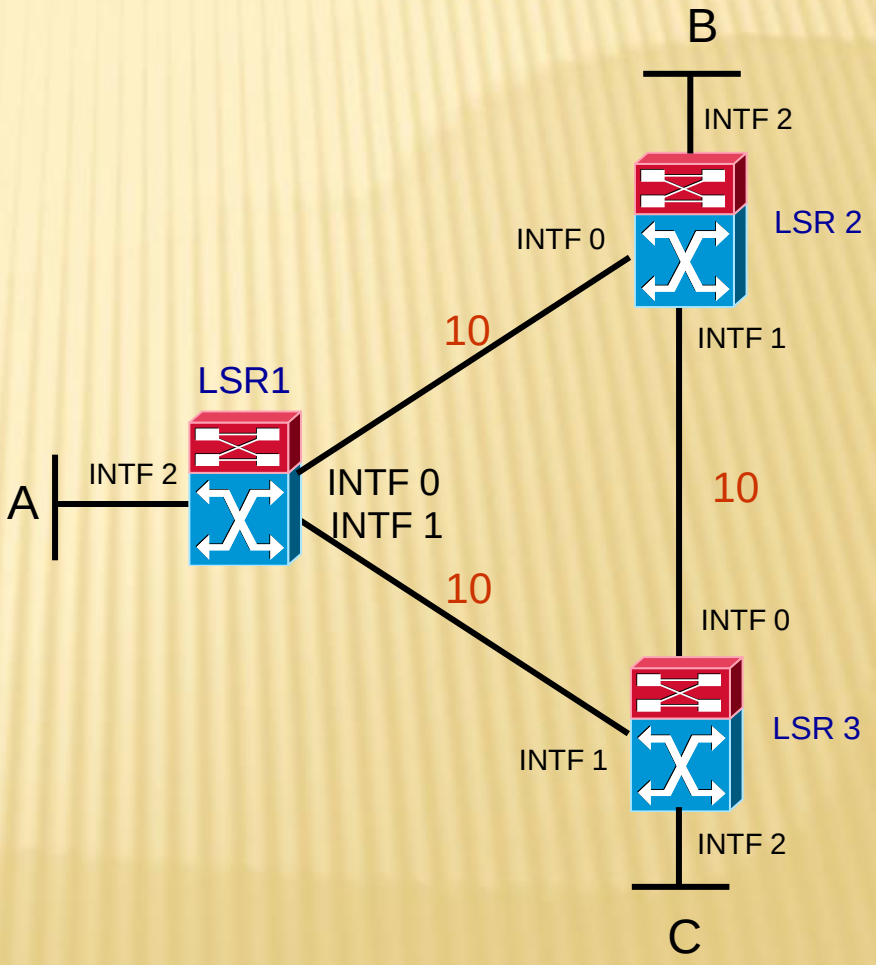

MPLS DATA STRUCTURES - EXAMPLE

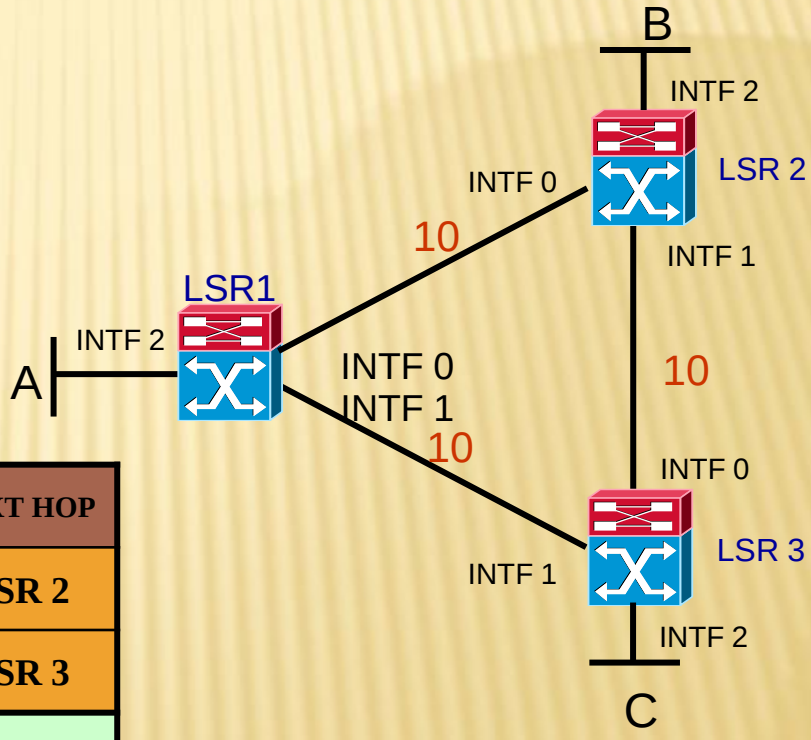

LFIB LSR 1

IN	OUT	FEC	INTERFACE	NEXT HOP
-	7	B	0	LSR 2
-	2	C	1	LSR 3
6	7	B	0	LSR 2
11	2	C	1	LSR 3
13	-	A	-	Routing Process

Mapping of incoming unlabeled packets to labels

Labeled packets exchanged between LSR3 and LSR2

Labeled packets destined for network A will be passed on to the IP routing process

MPLS DATA STRUCTURES - EXAMPLE

ROUTING TABLE LSR 1

DEST	METRIC	NEXT HOP	INTF
A	0	connected	2
B*	10	LSR2	0
B	20	LSR3	1
C*	10	LSR3	1
C	20	LSR2	0

* Optimum route to destination

LIB LSR 1

DEST	LABEL	ADV. LSR	BINDING
A	13	LSR 1	Local
B	7	LSR 2	Remote
B	8	LSR 3	Remote
B	6	LSR 1	Local
C	9	LSR 2	Remote
C	2	LSR 3	Remote
C	11	LSR 1	Local

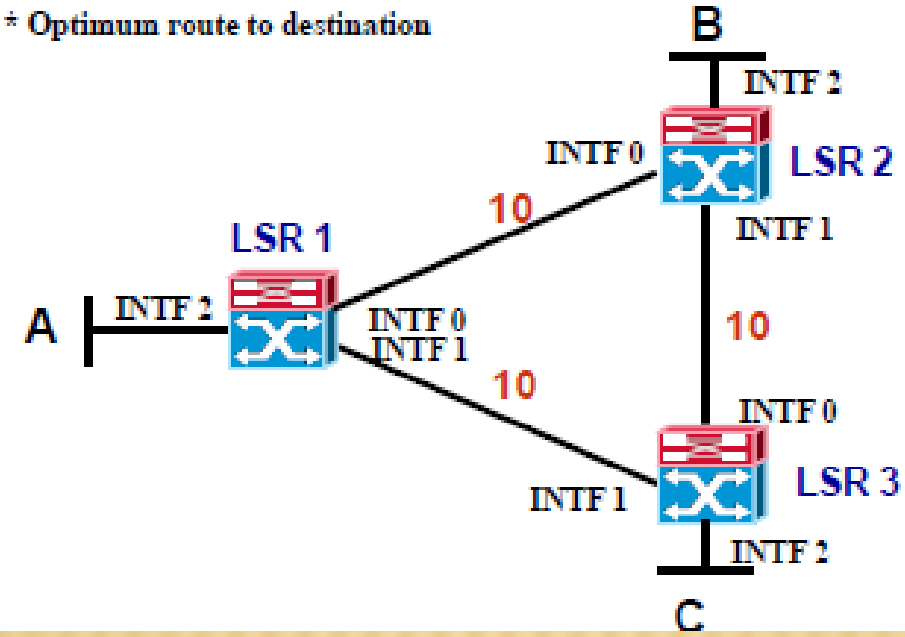

LFIB LSR 1

IN	OUT	FEC	INTF	NEXT HOP
-	7	B	0	LSR 2
-	2	C	1	LSR 3
6	7	B	0	LSR 2
11	2	C	1	LSR 3
13	-	A	2	Routing Process

MPLS DATA STRUCTURES – EXAMPLE NOTES

- LSR 1 has a route to network B via LSR 2. This path has a metric of 10.
- Also, LSR 1 has a route to network B via LSR 3. This path has a metric of 20, and therefore less preferable.
- LSR1 will always choose LSR 2 as the next hop to network B.
- After the routing process has *converged*, the LDP will build a session between all LDP neighbors (LDP Peers) that are reachable by that LSR.
- Each LDP session is established using TCP port 646, and label mapping information is exchanged with the LSR and all of the appropriate LDP peers.
- The downstream LSR must then define which label is mapped to a particular FEC. The label distribution process is called *downstream-on-demand* label allocation.
- "If LSR 1 sends packets destined for network B to LSR 2, use label 7."
- "If LSR 1 sends packets destined for network B to LSR 3, use label 8."

- The LIB will eventually be populated with all possible destinations.
- However, LIB does not yet contain any information about the next-hop address a packet classified with a particular FEC-to-label mapping should take.
- The next-hop can only be found in the IP routing table.
- The next-hop information and related label information are then put in the LFIB.
- The LFIB is used by the forwarding process.
- In LFIB of LSR 1, the *first two entries* are without an incoming label. This indicates that the packets arrived from outside MPLS domain and will have label imposed on them by LSR 1.
- *The last entry* of the LFIB indicates that if a packet arrives with label 13, the label should be removed and the IP packet should be passed on to the routing process in LSR 1.

PENULTIMATE HOP POPPING

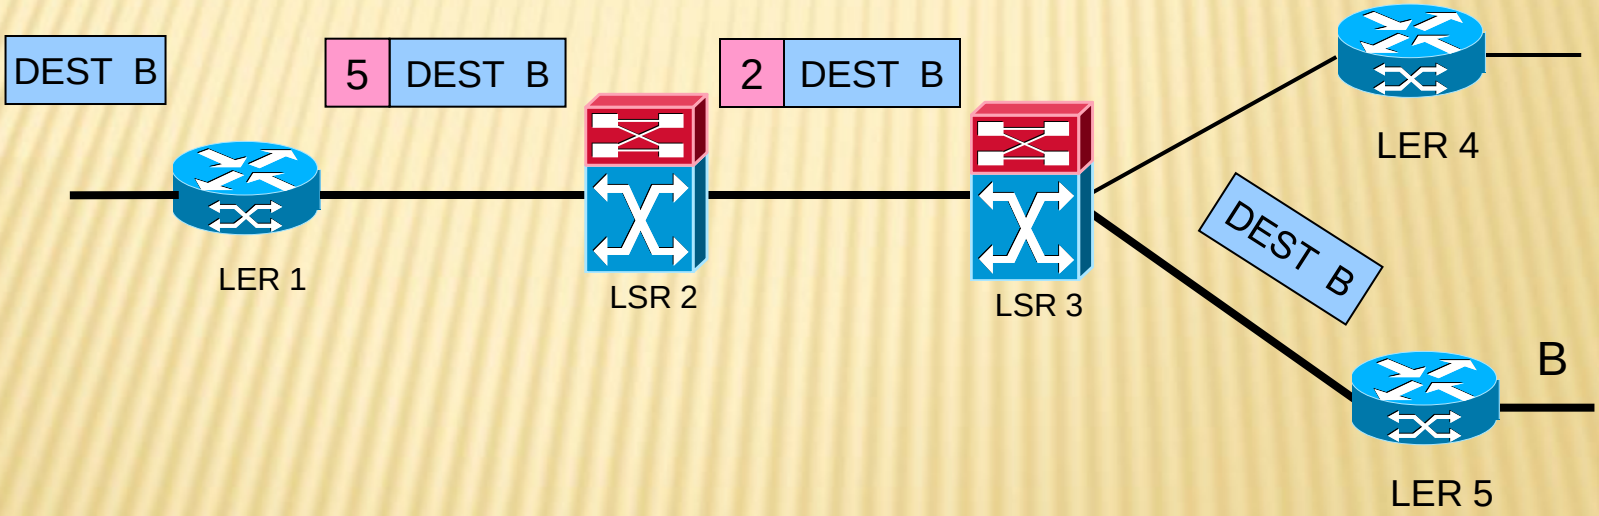

Penultimate Hop Popping (PHP) is a function performed by certain routers in an MPLS enabled network. It refers to the process whereby the outermost label of an MPLS tagged packet is removed by a Label Switch Router (LSR) before the packet is passed to an adjacent Label Edge Router (LER).

MPLS ADVANTAGES OVER ATM

Full Mesh Problem:

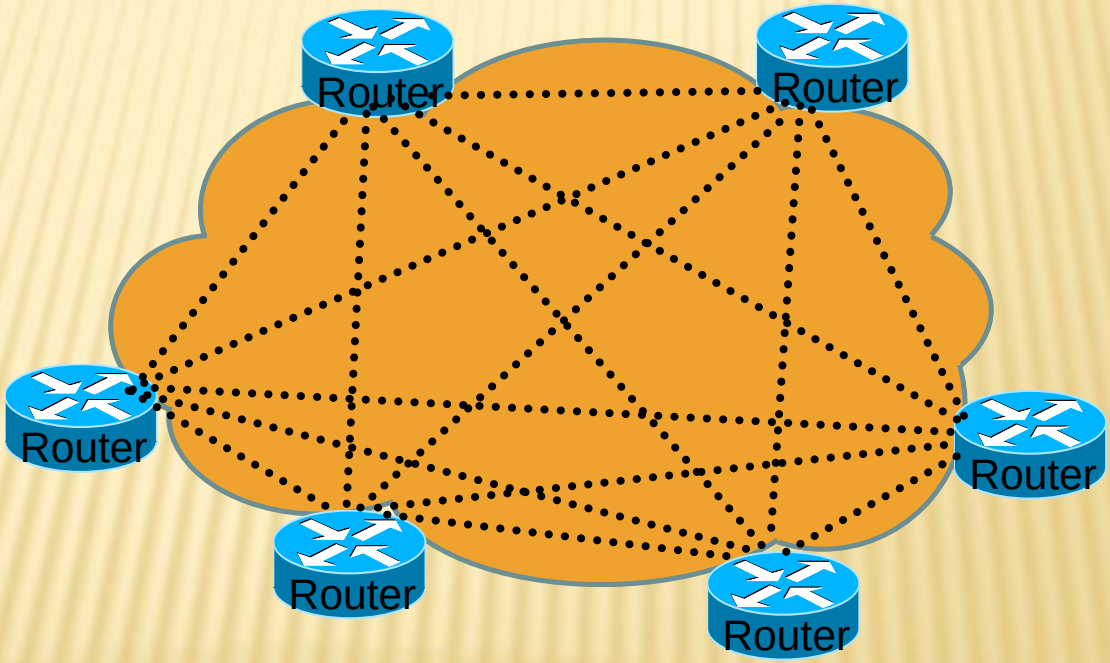

Overlay Model

MPLS ADVANTAGES OVER ATM

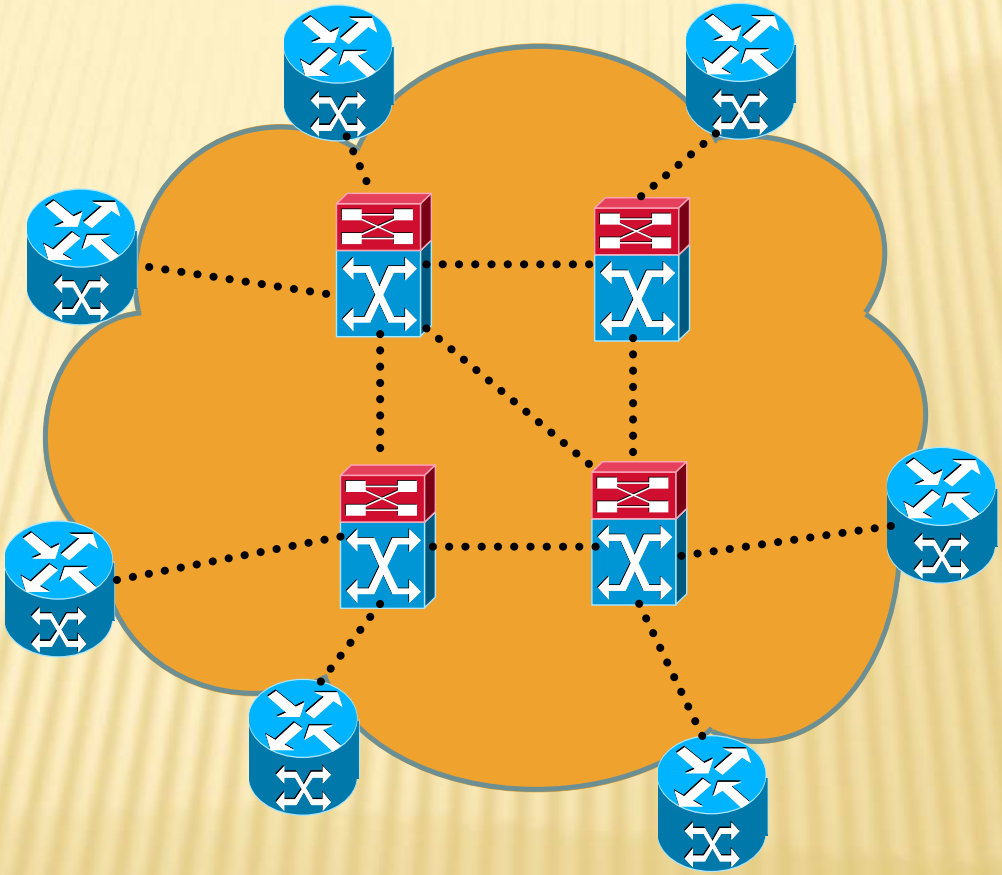

Peer Model

LABEL DISTRIBUTION PROTOCOL

- LDP uses many messages to create LSPs, classified in the following four types:

- Discovery—To identify other LSRs
- Adjacency—To create, maintain, and end sessions between LSRs
- Label advertisement—To request, map, withdraw, and release labels
- Notification—To provide advisory and error information

- Unlike the other LDP messages, the discovery process runs over UDP.
- Each LSR periodically broadcasts a link hello message to the well-known UDP port, 646.
- Each LSR listens on this port for link hello messages from other LSRs. In this manner, each LSR learns about all other LSRs to which it is directly connected, creating link hello adjacencies.
- When an LSR learns about another LSR, it establishes a TCP connection to the peer on well-known TCP port 646 and creates an LDP session on top of the TCP connection.