MPLS Concepts

Overview

This module explains the features of Multi-protocol Label Switching (MPLS) compared to traditional ATM and hop-by-hop IP routing. MPLS concepts and terminology as well as MPLS label format and Label Switch Router (LSR) architecture and operations are explained.

The module contains the following topics:

- Drawbacks of Traditional IP Routing
- Basic MPLS Concepts
- MPLS Labels and Label Stack
- MPLS Applications
- Differences Between Tag Switching and MPLS

Objectives

Upon completion of this module, the learner will be able to perform the following tasks:

- Identify the drawbacks of traditional IP routing
- Describe basic MPLS concepts and LSR types
- Describe how different MPLS applications coexist on the same platform using the same underlying technology
- List the standard bodies that are working on MPLS technology and the relationship between Tag Switching and MPLS

Drawbacks of Traditional IP Routing

Objectives

Upon completion of this lesson, the learner will be able to identify the drawbacks of traditional IP routing.

Traditional IP Forwarding

- **Traditional IP forwarding is based on the following:**
  - **Routing protocols are used to distribute Layer 3 (L3) routing information**
  - **Forwarding is based on the destination address only**
  - **Routing lookups are performed on every hop**

 www.cisco.com MPLS v2.1 -5

Before explaining basic MPLS functionality, three drawbacks of traditional IP forwarding should be highlighted:

- Routing protocols are used on all devices to distribute the routing information.
- Regardless of the routing protocol, routers always forward packets based on the destination address only. The only exception is policy-based routing (PBR) that bypasses the destination-based routing lookup.
- Routing lookups are performed on every router. Each router in the network makes an independent decision when forwarding packets.

MPLS helps reduce the number of routing lookups, possibly changes the forwarding criteria, and eliminates the need to run a particular routing protocol on all the devices.

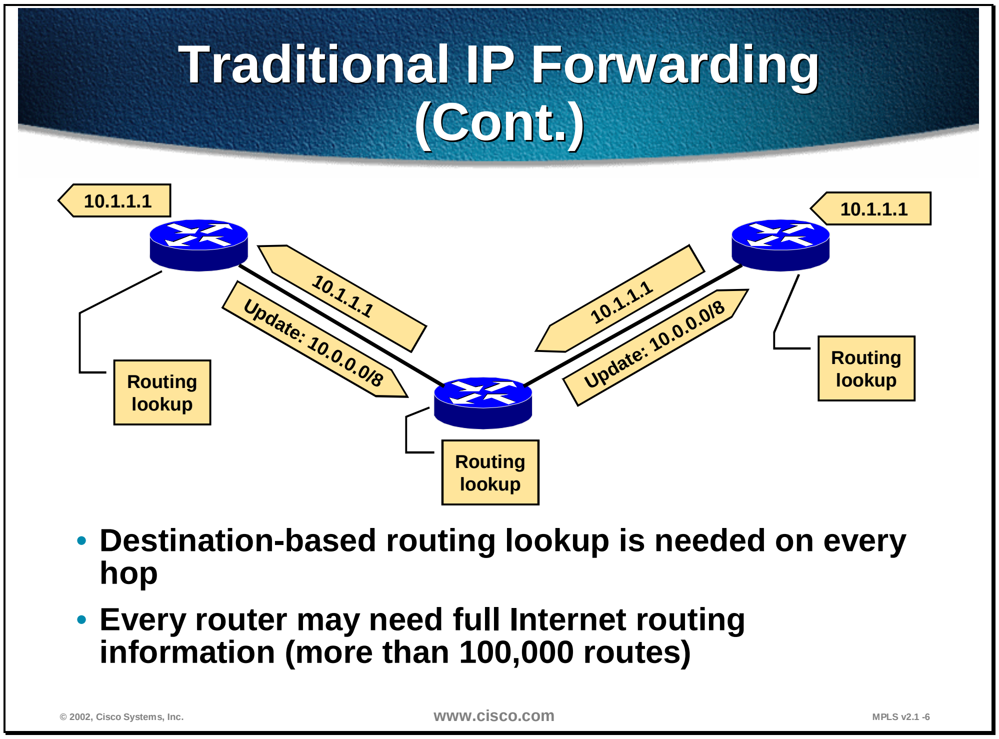

This figure illustrates how routers in a service provider network forward packets based on their destination addresses. The figure also shows that all the routers need to run a routing protocol (BGP) to get all the Internet routing information.

Every router in the path performs a destination-based routing lookup in a large forwarding table. Forwarding complexity is usually related to the size of the forwarding table and the switching mechanism.

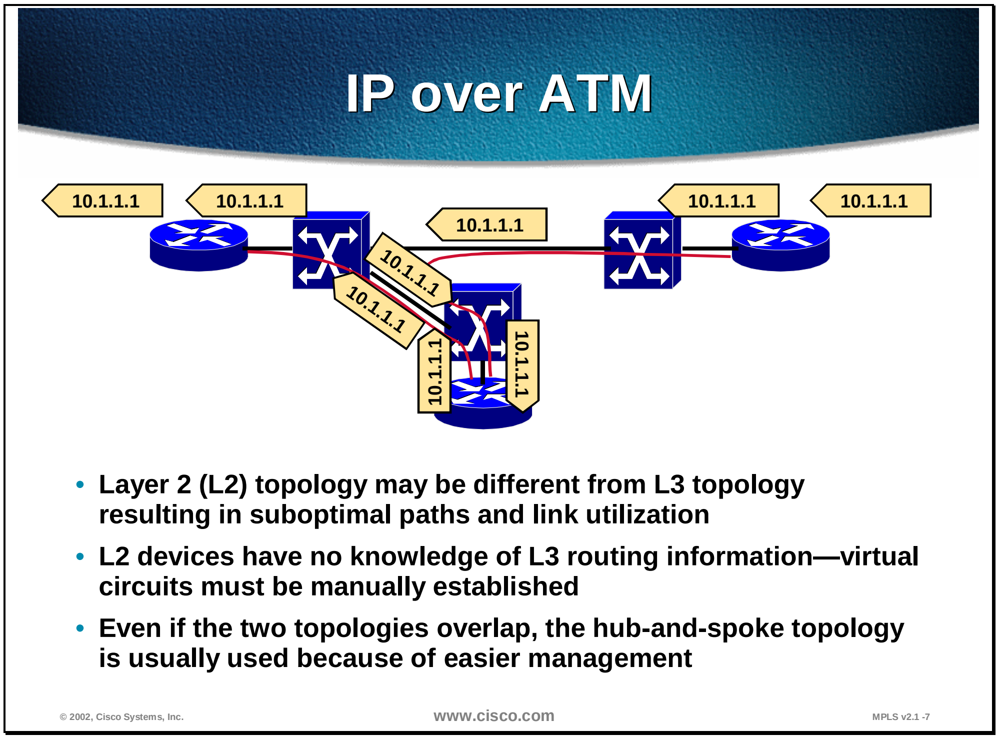

This figure illustrates a worst-case scenario where Layer 2 (L2) and Layer 3 (L3) topologies do not overlap. The result is that a single packet could be propagated with three L2 hops but instead requires 7 hops. The reason is that L2 devices have static information about how to interconnect L3 devices. Routers use a routing protocol to propagate L3 routing information through the intermediary router.

Even in another L2 topology, where the forwarding to the hub router was more optimal, the packet forwarding from the rightmost router to the leftmost router would still require unnecessary hops. The only possible solution to get optimal forwarding from any router to any other router would be to have a full mesh of virtual circuits. But this is rarely used due to its complexity.

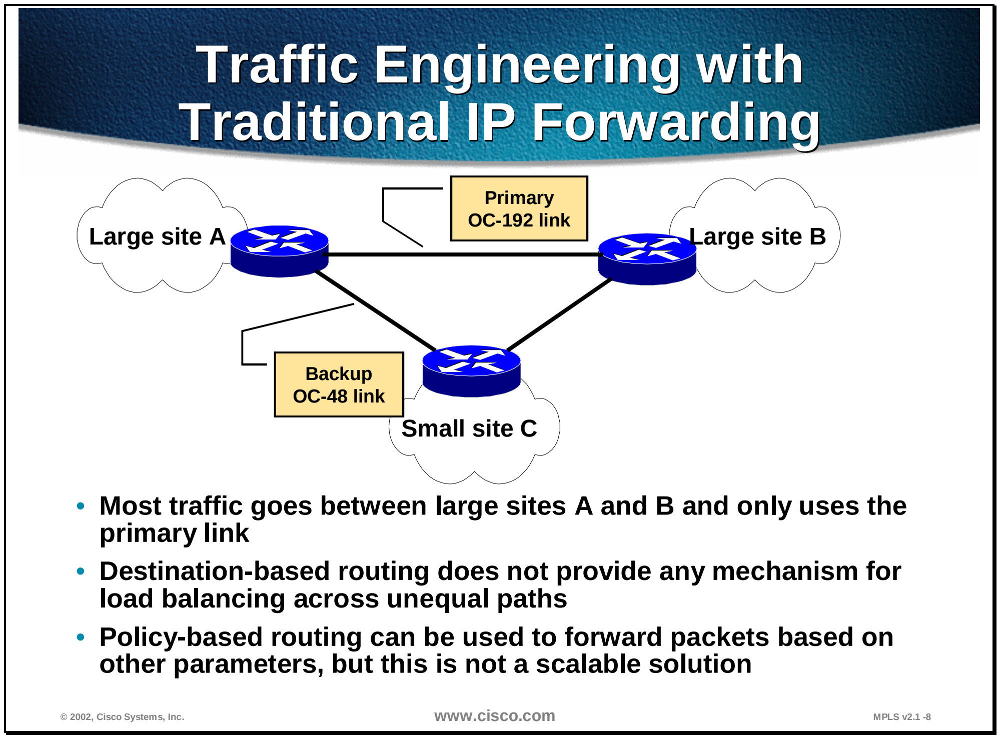

This figure illustrates a topology with unequal links. Traffic patterns illustrate that most of the traffic goes between sites A and B.

Traditional IP forwarding does not have a scalable mechanism to allow the utilization of the backup link (unequal load balancing).

Policy-based routing could be used to select some packets and route those along the backup link. But this is not possible in high volume traffic due to performance limitations.

Summary

After completing this lesson, you should be able to identify the following drawbacks of traditional IP routing:

- Routing protocols with full routing information are required on all routers.
- Routers make a destination-based forwarding decision only.
- Routers must make a routing lookup on every hop.

Lesson Review

1. List major drawbacks of traditional IP routing.
2. Based on what information do routers forward IP packets?
3. What mechanism can be used to forward packets based on other parameters?
4. Why is this mechanism not suitable for large networks?

Basic MPLS Concepts

Objectives

Upon completion of this lesson, the learner will be able to perform the following tasks:

- Describe MPLS architecture
- Describe the MPLS approach to IP routing
- Describe the difference between data plane and control plane in MPLS
- Describe the difference between packet-mode and cell-mode MPLS
- List LSR types
- Describe LSR architecture

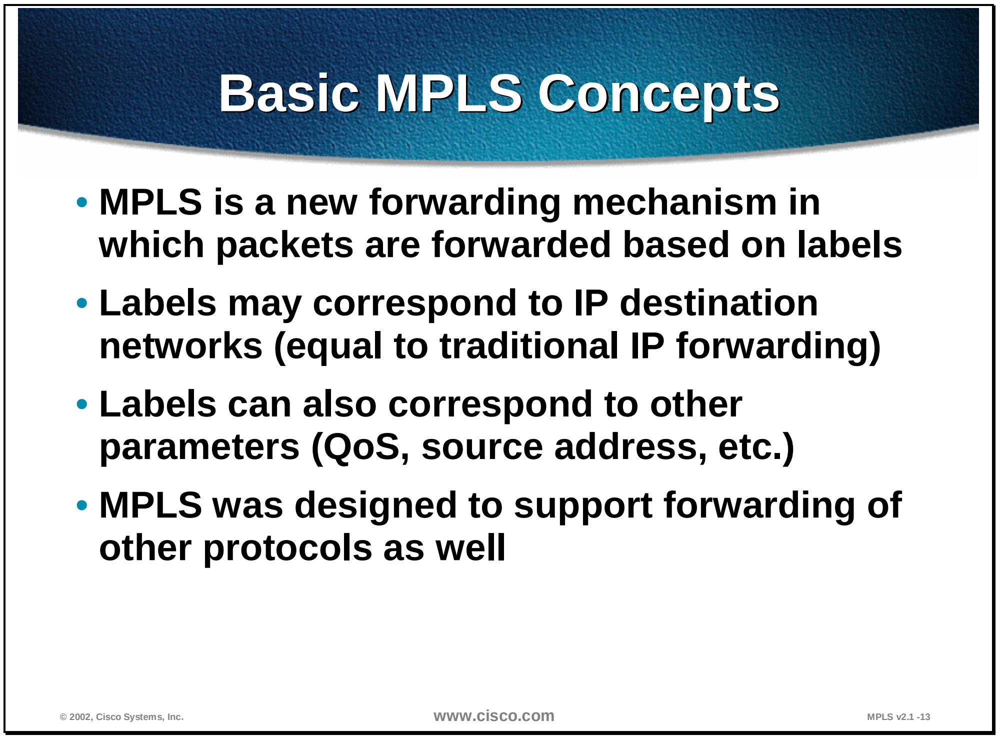

MPLS is a new switching mechanism that uses labels (numbers) to forward packets.

Labels usually correspond to L3 destination addresses (equal to destination-based routing). Labels can also correspond to other parameters (Quality of Service [QoS], source address, etc.).

MPLS was designed to support other protocol stacks than IP as well. Label switching is performed regardless of the L3 protocol.

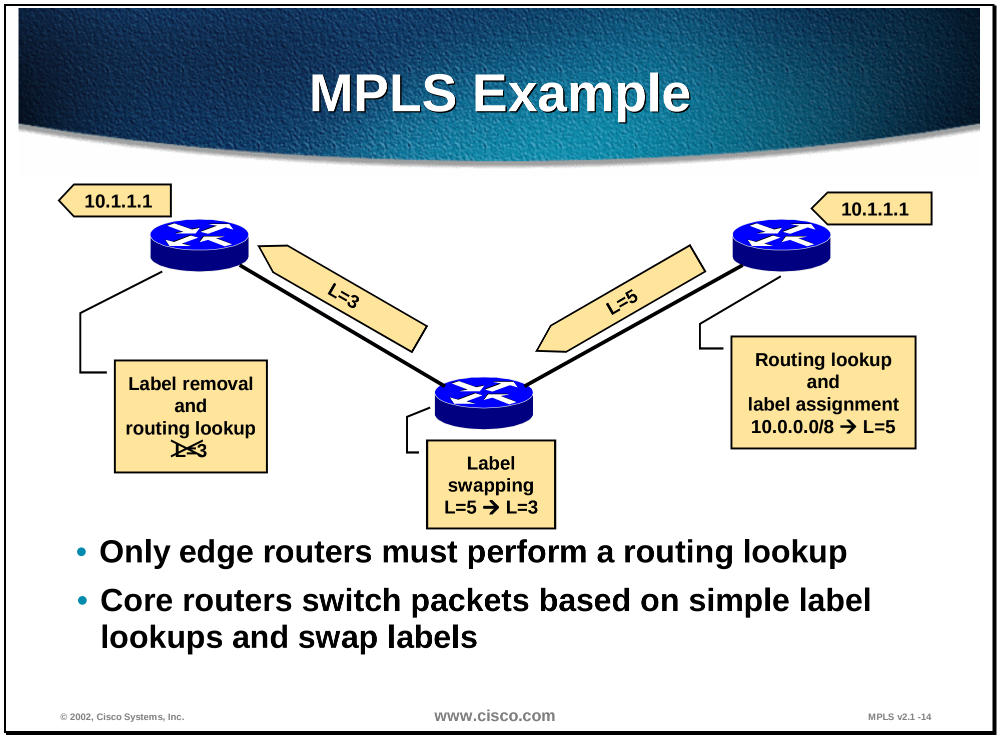

This figure illustrates a situation where the intermediary router does not have to perform a time-consuming routing lookup. Instead this router simply swaps a label with another label (5 is replaced by 3) and forwards the packet based on the received label (3).

In larger networks the result of MPLS labeling is that only the edge routers perform a routing lookup. All the core routers forward packets based on the labels.

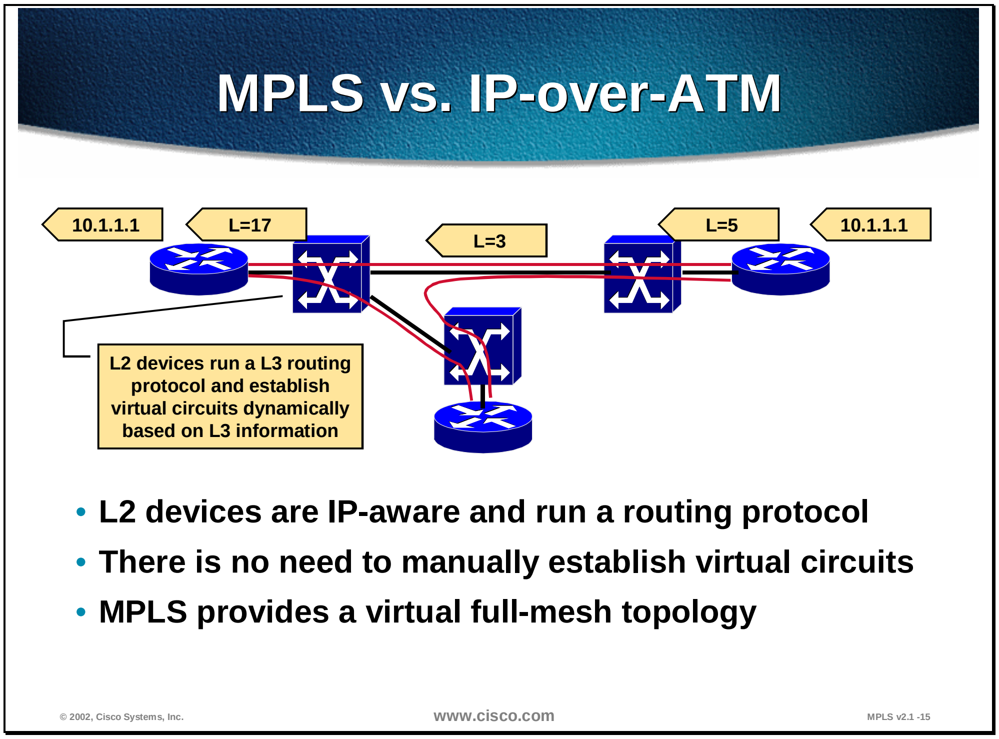

This figure illustrates how MPLS is used in ATM networks to provide optimal routing across L2 ATM switches. In order for MPLS to work with ATM switches, the switches must be somewhat L3 aware. The ATM switches must run a L3 routing protocol and thus have knowledge about how to reach L3 subnets.

The ATM switches automatically create a full mesh of virtual circuits based on L3 routing information. The VCs are used to forward the user data segmented into ATM cells. Another benefit of this setup is that there is no longer a need to manually establish virtual circuits.

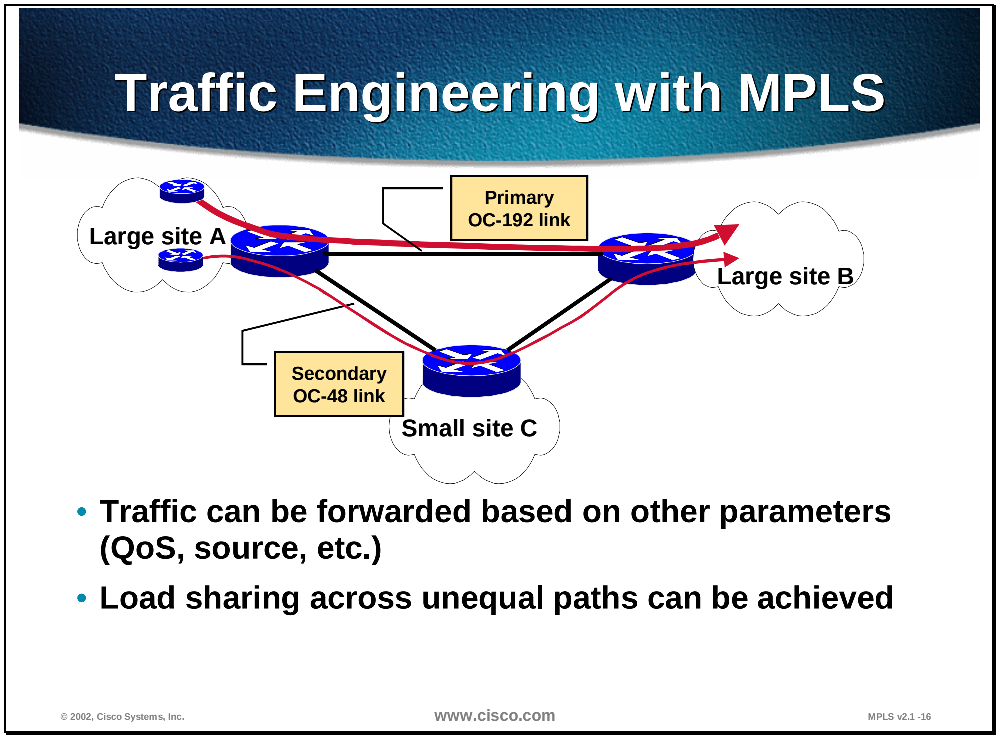

MPLS also supports traffic engineering. Traffic engineered tunnels can be created based on traffic analysis to provide load balancing across unequal paths.

Multiple traffic engineering tunnels can lead to the same destination but can use different paths. Traditional IP forwarding would force all traffic to use the same path based on the destination-based forwarding decision. Traffic engineering determines the path at the source based on additional parameters (available resources and constraints in the network).

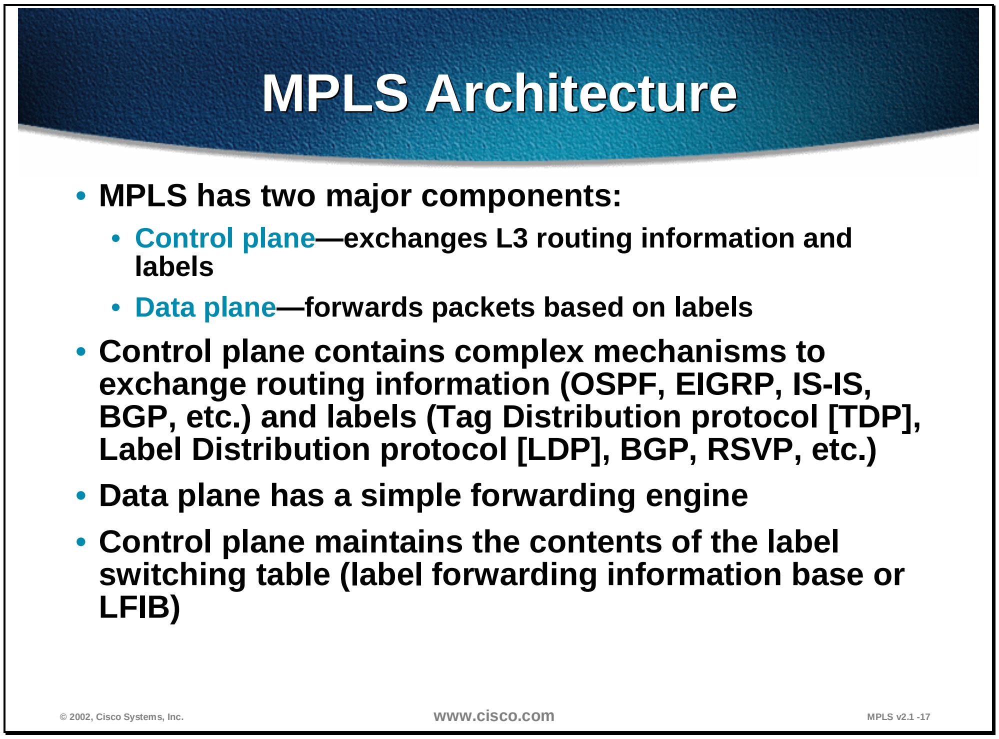

To better understand the inner workings of MPLS its two major components have to be introduced:

- **Control plane:** Takes care of the routing information exchange and the label exchange between adjacent devices
- **Data plane:** Takes care of forwarding either based on destination addresses or labels

There is a large number of different routing protocols such as OSPF, IGRP, EIGRP, IS-IS, RIP, BGP, etc. that can be used in the control plane.

The control plane also requires protocols to exchange labels, such as:

- Tag Distribution Protocol [TDP] (MPLS)
- Label Distribution Protocol [LDP] (MPLS)
- BGP (MPLS virtual private networks [VPNs])
- Resource-Reservation Protocol [RSVP] (MPLS Traffic Engineering [MPLS-TE])
- CR-LDP (MPLS-TE)

The data plane however, is a simple label-based forwarding engine that is independent of the type of routing protocol or label exchange protocol. A Label Forwarding Information Base (LFIB) is used to forward packets based on labels. The LFIB table is populated by the label exchange protocols used in the control plane.

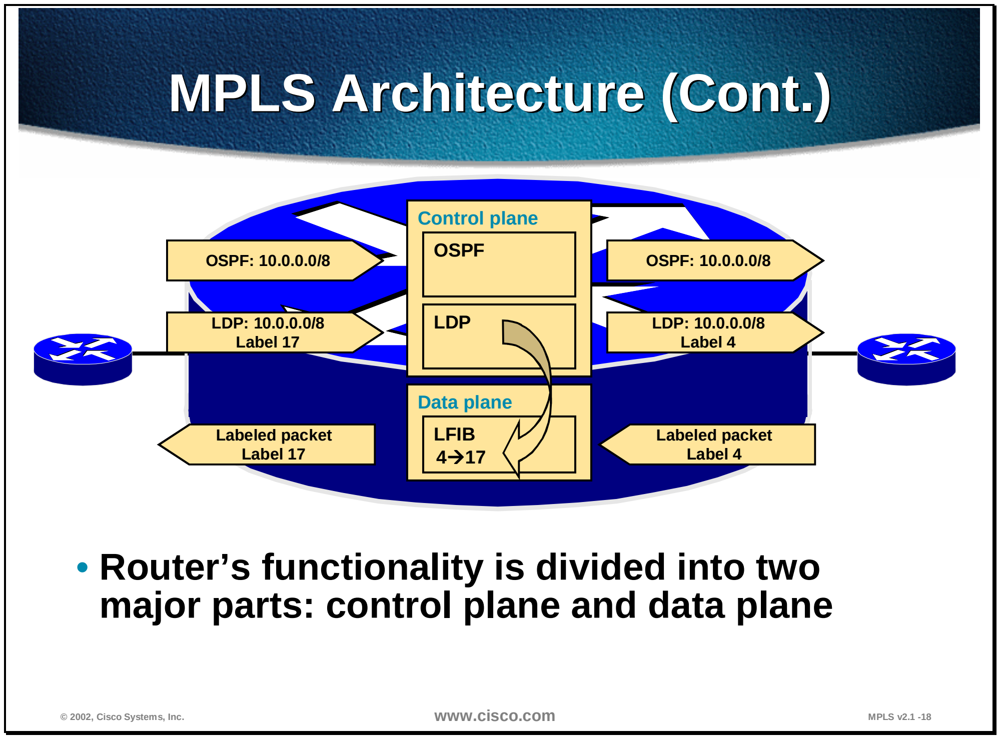

A simple MPLS implements destination-based forwarding that uses labels to make forwarding decisions.

A L3 routing protocol is still needed to propagate L3 routing information. A label exchange mechanism is simply an add-on to propagate labels that are used for L3 destinations.

This figure illustrates the two components of the control plane:

- OSPF that receives IP network 10.0.0.0/8 from the left neighbor and forwards it to the right neighbor.
- LDP that receives label 17 from the left neighbor to be used for packets with a destination address 10.x.x.x when forwarded to that neighbor. A local label 4 is generated and sent to upstream neighbors so these neighbors can label packets with the appropriate label. LDP inserts an entry into Data Plane's LFIB table where label 4 is mapped to label 17.

The data plane then forwards all packets with label 4 through the appropriate interfaces and replaces the label with label 17.

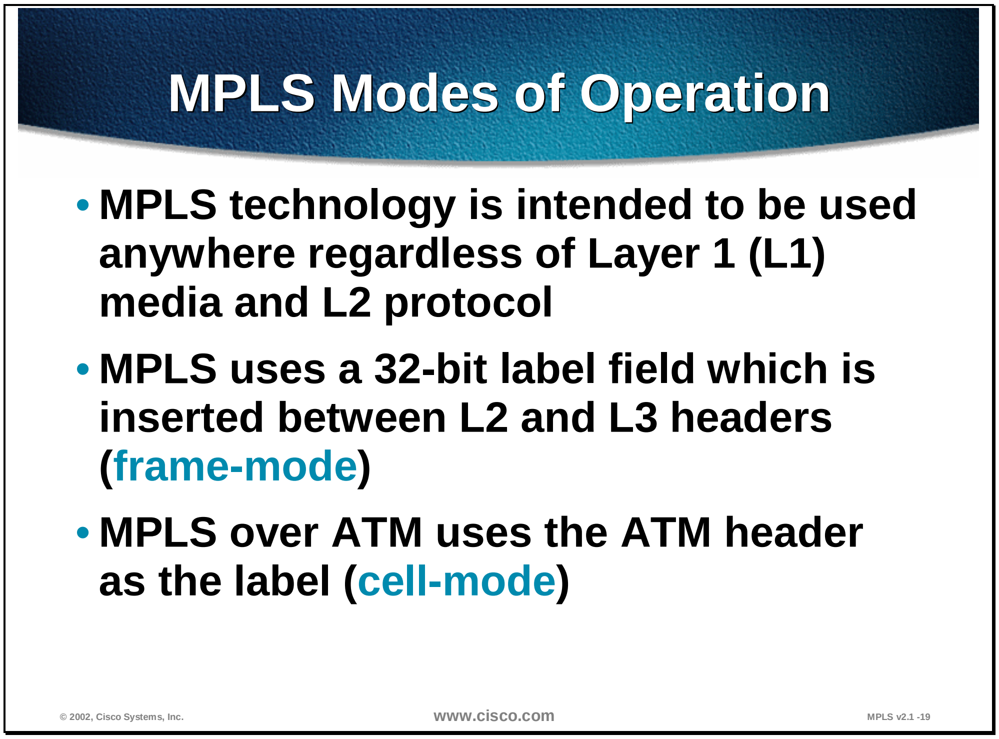

MPLS is designed for use on virtually any media and L2 encapsulation. Most L2 encapsulations are frame-based and MPLS simply inserts a 32-bit label between the L2 and L3 headers ("frame-mode" MPLS).

ATM is a special case where fixed-length cells are used and a label cannot be inserted on every cell. MPLS uses the virtual path identifier/ virtual channel identifier (VPI/VCI) fields in the ATM header as a label ("cell-mode" MPLS).

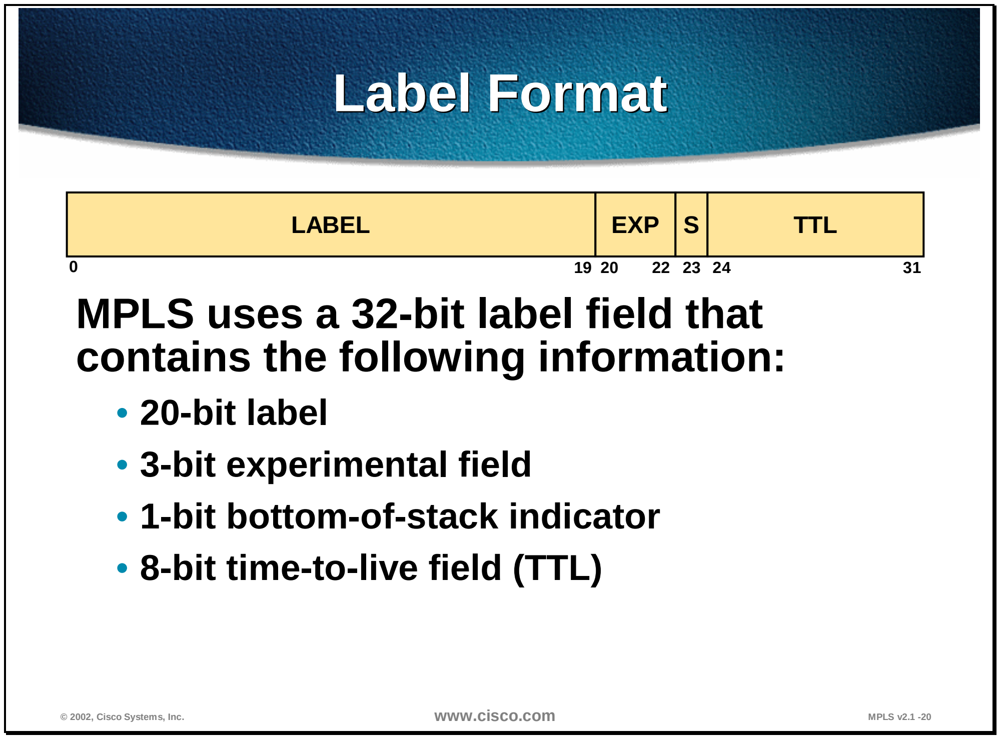

A 32-bit label contains the following fields:

- **20-bit label:** The actual label
- **3-bit experimental field:** It is used to define a class of service (i.e. IP precedence)
- **Bottom-of-stack bit:** MPLS allows multiple labels to be inserted; this bit is used to determine if this is the last label in the packet
- **8-bit time-to-live (TTL) field:** It has the same purpose as the TTL field in the IP header

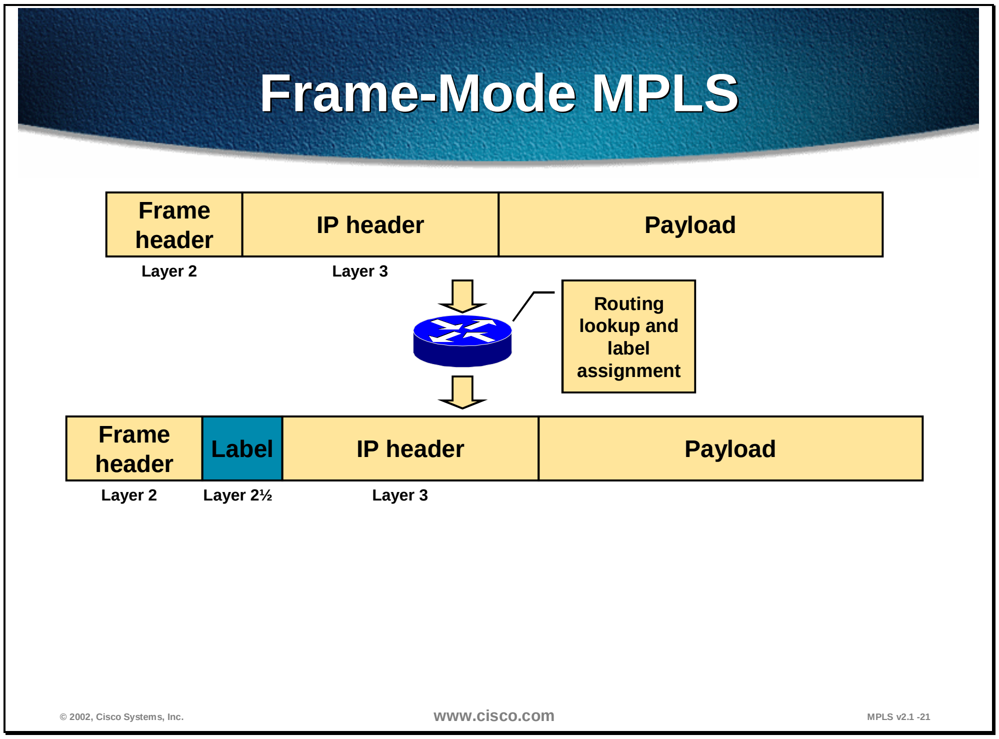

This figure illustrates an edge router that receives a normal IP packet. The router then performs the following actions:

Step 1 Routing lookup to determine the outgoing interface

Step 2 Label is assigned and inserted between L2 frame header and L3 packet header if the outgoing interface is enabled for MPLS and a next-hop label for the destination exists

Step 3 The labeled packet is sent

Other routers in the core simply forward packets based on the label.

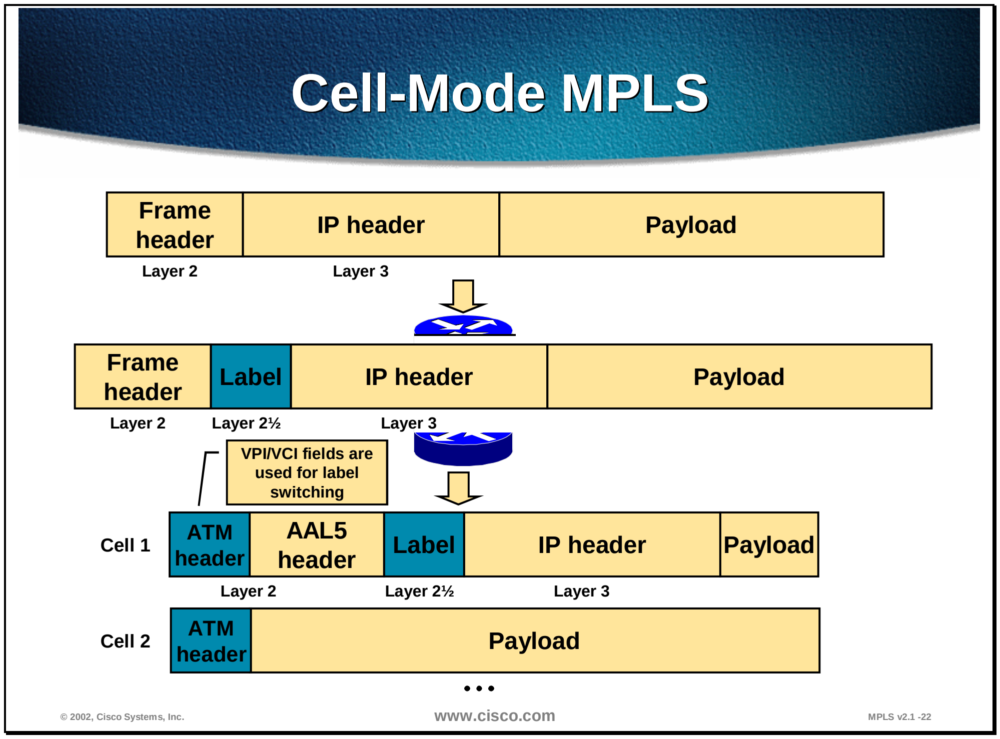

Cell-mode MPLS uses the ATM header's VPI/VCI field for forwarding decisions while the 32-bit label is still preserved in the frame but not used in the ATM network. The original label is only present in the first cell of a packet.

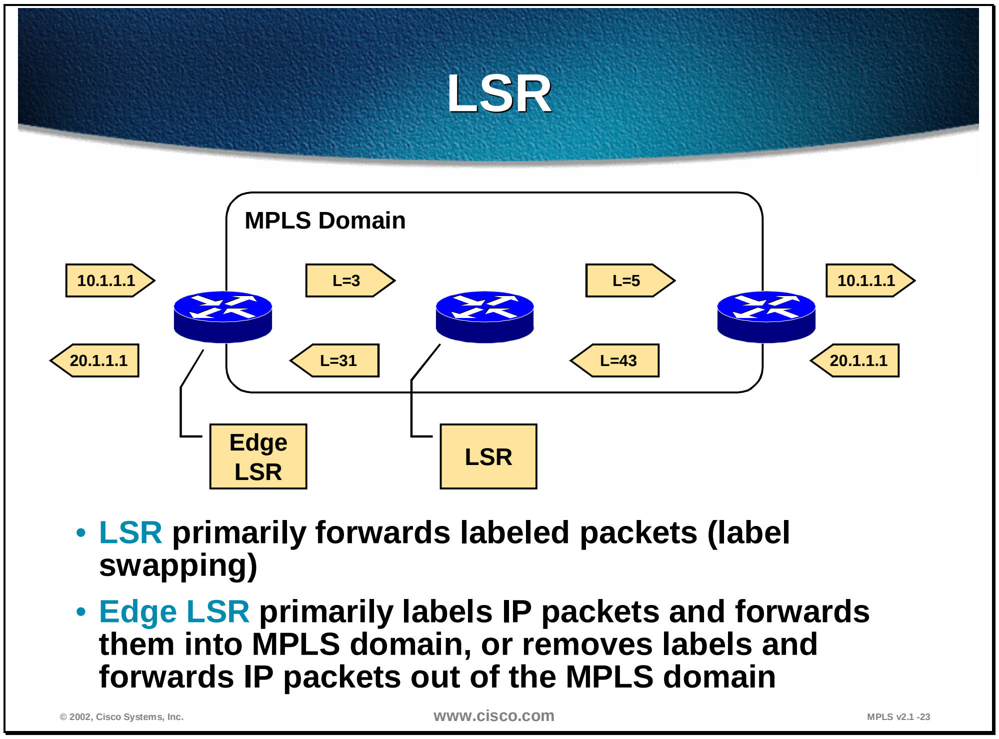

Before proceeding with a detailed description of MPLS, some of the terminology that is used in this course is presented:

- **LSR:** A device that primarily forwards packets based on labels
- **Edge LSR:** A device that primarily labels packets or removes labels

LSRs and Edge LSRs are usually devices that are capable of doing both label switching and IP routing. Their names are based on their position in an MPLS domain. Routers that have all interfaces enabled for MPLS are called LSRs because they mostly forward labeled packets. Routers that have some interfaces that are not enabled for MPLS are usually at the edge of an MPLS domain (autonomous system). These routers also forward packets based on IP destination addresses and label them if the outgoing interface is enabled for MPLS.

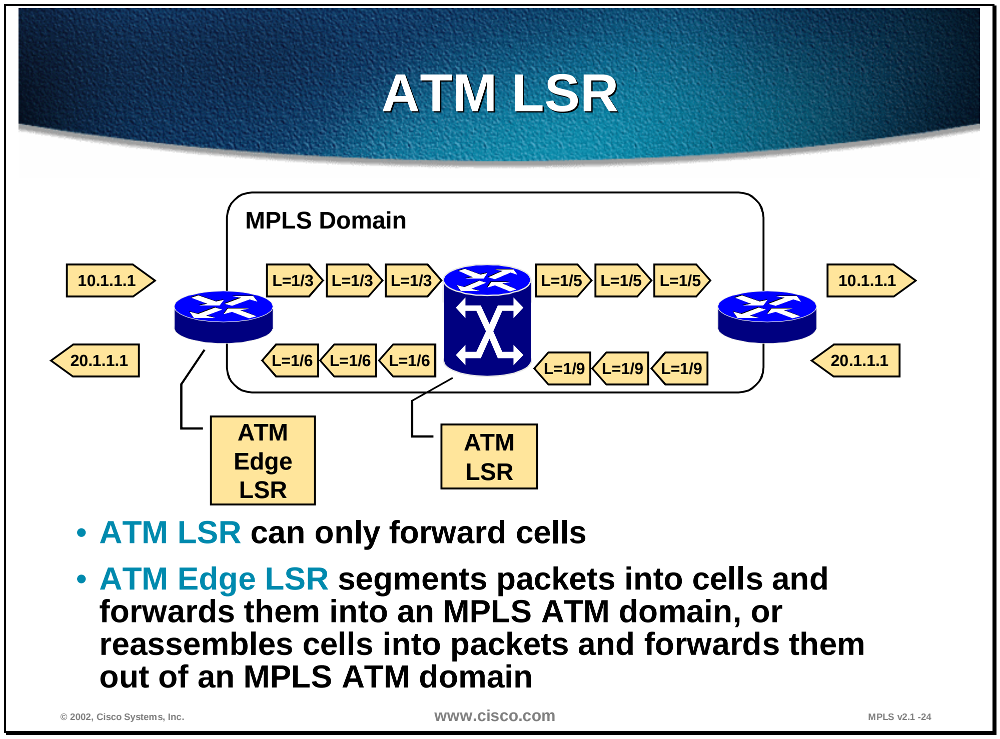

LSRs that perform cell-mode MPLS are called:

- **ATM LSR** if they are ATM switches. All interfaces are enabled for MPLS and forwarding is done of cells <u>only,</u> based on labels.
- **ATM Edge LSR** if they are routers connected to an MPLS-enabled ATM network. Some interfaces are traditional IP interfaces where IP packets are exchanged. The IP packets are segmented into cells and forwarded into the MPLS ATM domain on cell-mode MPLS enabled interfaces. Cell traffic in the other direction is reassembled into IP packets.

Architecture of LSRs

- LSRs, regardless of the type, perform the following three functions:
  - Exchange routing information
  - Exchange labels
  - Forward packets (LSRs and edge LSRs) or cells (ATM LSRs and ATM edge LSRs)
- The first two functions are part of the control plane
- The last function is part of the data plane

 www.cisco.com MPLS v2.1 -25

LSRs of all types must perform the following functions:

- Exchange L3 routing information (ATM LSRs must also exchange L3 routing information)
- Exchange labels
- Forward packets or cells

Frame-mode and cell-mode MPLS use a different data plane:

- Frame-mode MPLS forwards packets based on the 32-bit label
- Cell-mode MPLS forwards packets based on labels encoded into the VPI/VCI fields in the ATM header

The control plane performs the following functions:

- Exchange routing information regardless of the type of LSR
- Exchange labels according to the type of MPLS (frame-mode or cell-mode)

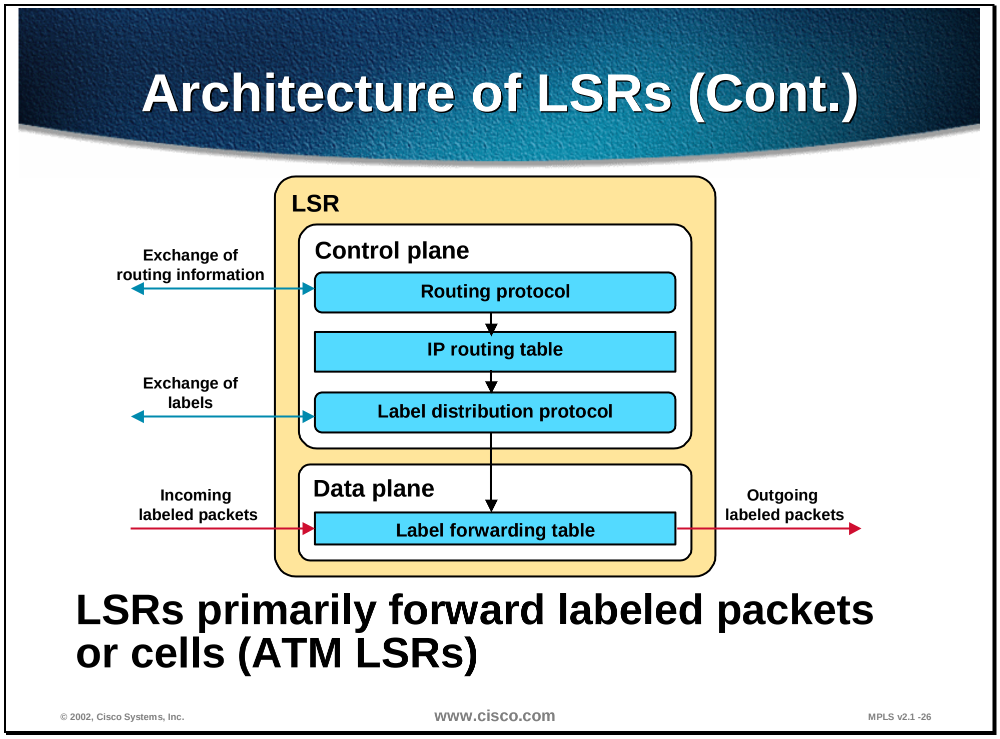

The primary function of an LSR is to forward labeled packets. Therefore, every LSR needs a L3 routing protocol (OSPF, EIGRP, IS-IS, etc.) and a label exchange protocol (LDP, TDP, etc.).

The label exchange protocol populates the LFIB table in the data plane that is used to forward labeled packets.

Note LSRs may not be able to forward unlabeled packets either because they are ATM LSRs, or they do not have all the routing information.

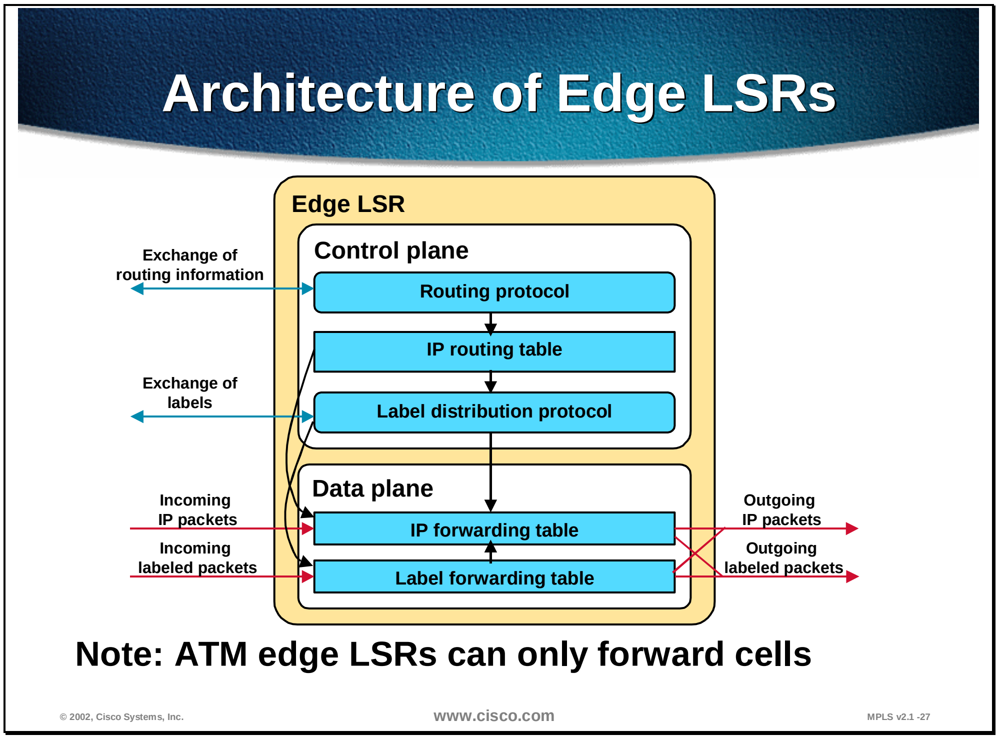

Edge LSRs also forward IP packets based on their IP destination addresses and optionally label them if a label exists.

The following combinations are possible:

- A received IP packet is forwarded based on the IP destination address and sent as an IP packet
- A received IP packet is forwarded based on the IP destination address and sent as a labeled packet
- A received labeled packet is forwarded based on the label; the label is changed and the packet is sent
- A received labeled packet is forwarded based on the label; the label is removed and the packet is sent out is an IP packet

The following scenarios are possible if the network is misconfigured:

- A received labeled packet is dropped if the label is not found in the LFIB table even if the IP destination exists in the FIB table
- A received IP packet is dropped if the destination is not found in the FIB table even if there is a label-switched path (LSP) available for the destination

Summary

MPLS architecture is divided into two parts:

- Control plane that takes care of routing information and label propagation
- Data plane that takes care of the forwarding of packets

MPLS has two modes:

- Frame-mode MPLS that is used on all frame-based media
- Cell-mode MPLS that is used in MPLS-enabled ATM networks

MPLS networks use the following devices:

- LSR to forward packets based on a 32-bit label
- Edge LSR to forward labeled packets or label IP packets or remove labels
- ATM LSRs to forward cells based on labels encoded into the VPI/VCI fields in the ATM header
- ATM Edge LSRs that segment labeled or unlabeled packets into ATM cells where a label is encoded into VPI/VCI fields in the ATM header

Lesson Review

1. What are the major drawbacks of traditional IP forwarding and how does MPLS solve them?
2. What functions does an LSR perform?
3. List the types of LSRs.
4. Name the two modes of MPLS.
5. Explain the difference between an LSR and an Edge LSR.
6. Explain the difference between an LSR and an ATM LSR.

MPLS Labels and Label Stack

Objectives

Upon completion of this lesson, the learner will be able to perform the following tasks:

- Describe the format of MPLS label
- Explain the concept of the MPLS Label Stack
- Describe the way MPLS labels are used in Packet-mode and ATM environment

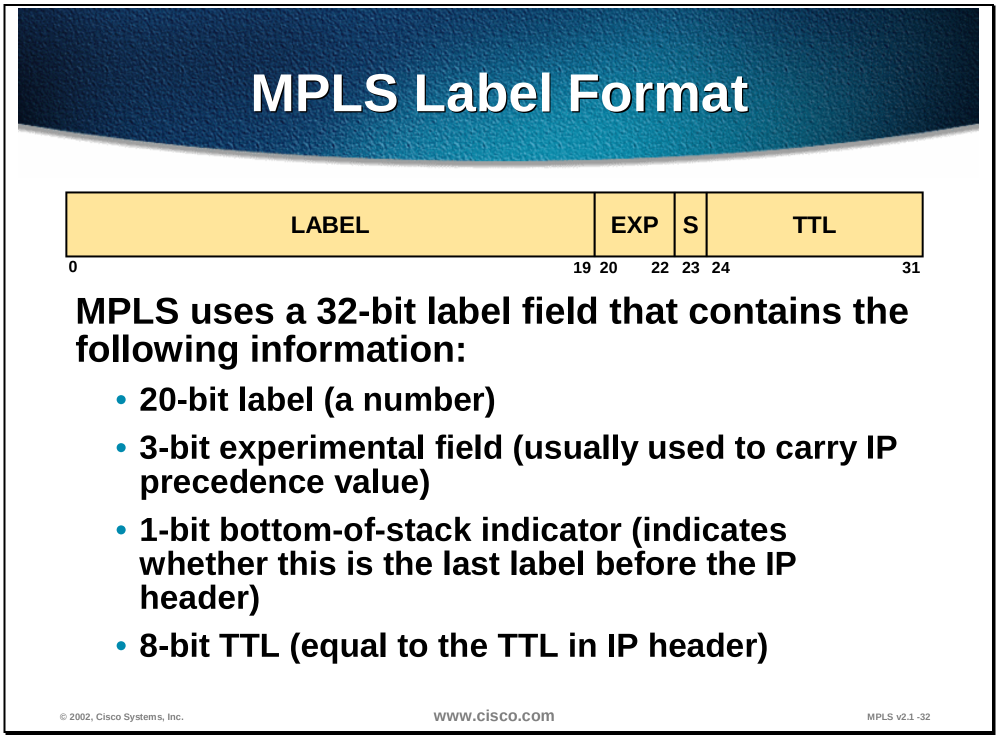

MPLS uses a 32-bit label that is inserted between the L2 and L3 header. An MPLS label contains four fields:

- The actual label
- Experimental field
- Bottom-of-stack bit
- TTL field

MPLS Labels

- Labels are inserted between the L2 (frame) header and the L3 (packet) header
- There can be more than one label (label stack)
- Bottom-of-stack bit indicates if the label is the last label in the label stack
- TTL field is used to prevent indefinite looping of packets
- Experimental bits are usually used to carry the IP precedence value

 www.cisco.com MPLS v2.1 -33

Labels are inserted between the L2 (frame) header and the L3 (packet) header. In some MPLS applications (e.g. MPLS VPN) more than one label is required. In those cases the multiple labels form label stack.

Each label contains the following fields:

- **20-bit label:** The actual label, which is a simple 20-bit number that has local significance and changes on every hop.
- **3-bit experimental field:** Currently used to define a class of service such by reflecting the IP precedence of the encapsulated IP packet. Cisco routers automatically assign the IP precedence value to this field.
- **Bottom-of-stack bit:** MPLS allows multiple labels to be inserted. The bottom-of-stack bit is used to determine if this is the last label in the packet. This bit is set to "1" in the last label in the packet.
- **8-bit TTL field:** It has the same purpose as the TTL field in the IP header. This field is decreased on every hop.

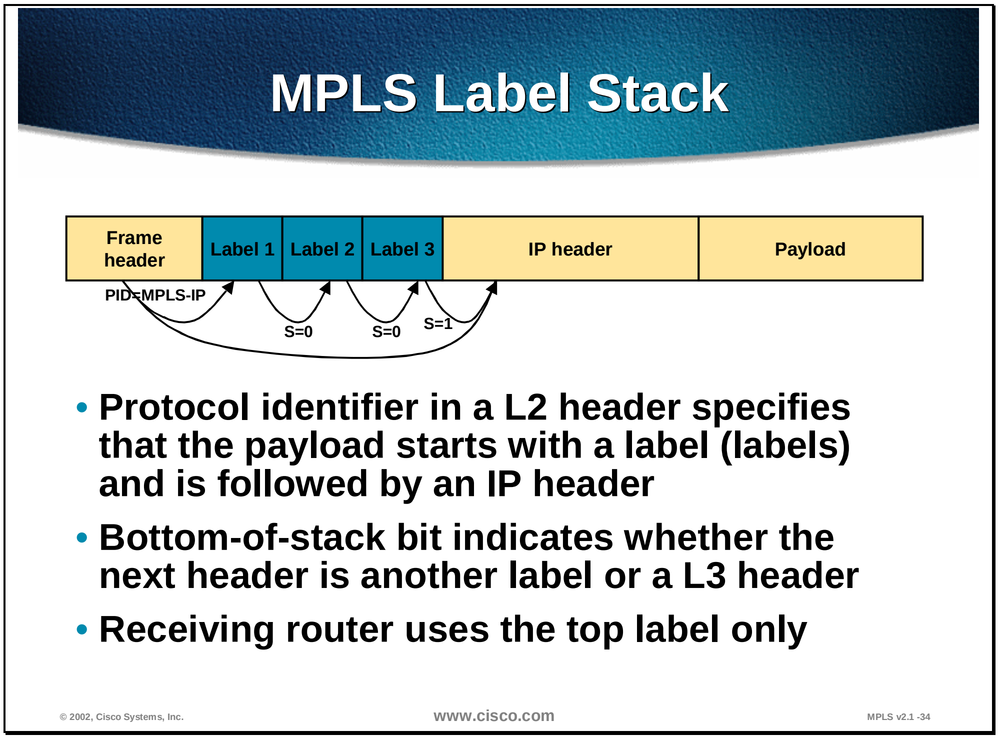

A label does not contain any information about the L3 protocol being carried in a packet. A new protocol identifier is used for every MPLS enabled L3 protocol.

This list shows the ethertype values used to identify L3 protocols with most L2 encapsulations:

- **Unlabeled IP unicast:** PID=0x0800 identifies that the frame payload is an IP packet.
- **Labeled IP unicast:** PID=0x8847 identifies that the frame payload is a unicast IP packet with at least one label preceding the IP header. The bottom-of-bit indicates when the IP header actually starts.
- **Labeled IP multicast:** PID=0x8848 identifies that the frame payload is a multicast IP packet with at least one label preceding the IP header. The bottom-of-bit indicates when the IP header actually starts.

A router that receives a frame where the PID indicates that it is a labeled packet uses only the top label in stack for forwarding decisions.

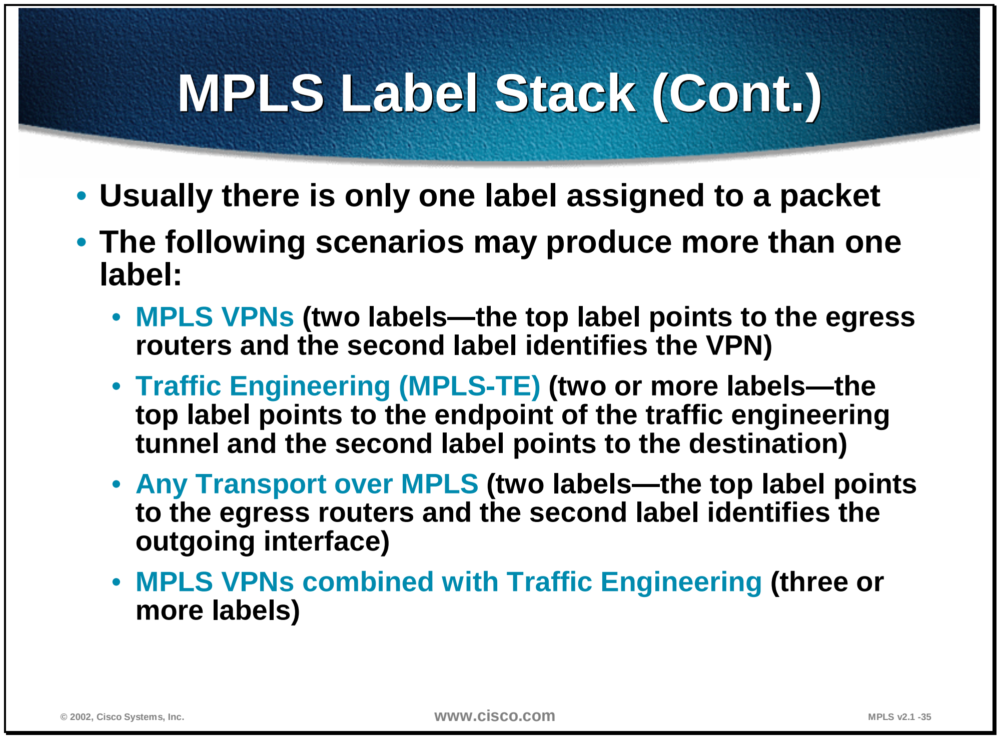

As previously noted, MPLS supports multiple labels in one packet. Simple MPLS uses just one label in each packet. The following applications may add additional labels to packets:

- MPLS VPNs use multiprotocol BGP to propagate a second label that is used in addition to the one propagated by TDP or LDP.
- MPLS-TE uses RSVP to establish label-switched tunnels. RSVP also propagates labels that are used in addition to the one propagated by LDP or TDP.
- Any Transport over MPLS (AToM) uses a directed multihop LDP session between the edge routers to propagate a second label that is used in addition to the one propagated by the per link LDP- or TDP-sessions.
- A combination of the above mentioned mechanisms with some other features might result in three or more labels being inserted into one packet.

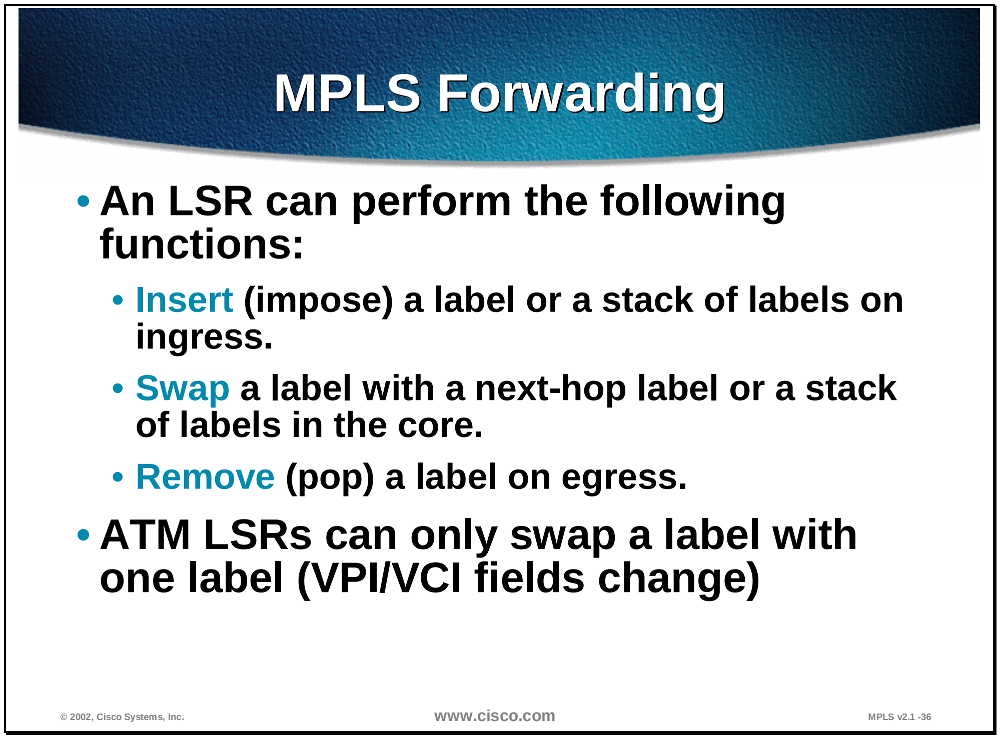

An IP packet going through an MPLS domain experiences the following:

- A label or a stack of labels is inserted (imposed) on an Edge LSR
- The top label is swapped with a next-hop label or a stack of labels on an LSR
- The top label is removed on the LSP endpoint (usually one hop before the egress Edge LSR or on the egress edge LSR itself)

ATM LSRs only support the swapping of one label (normal ATM operation).

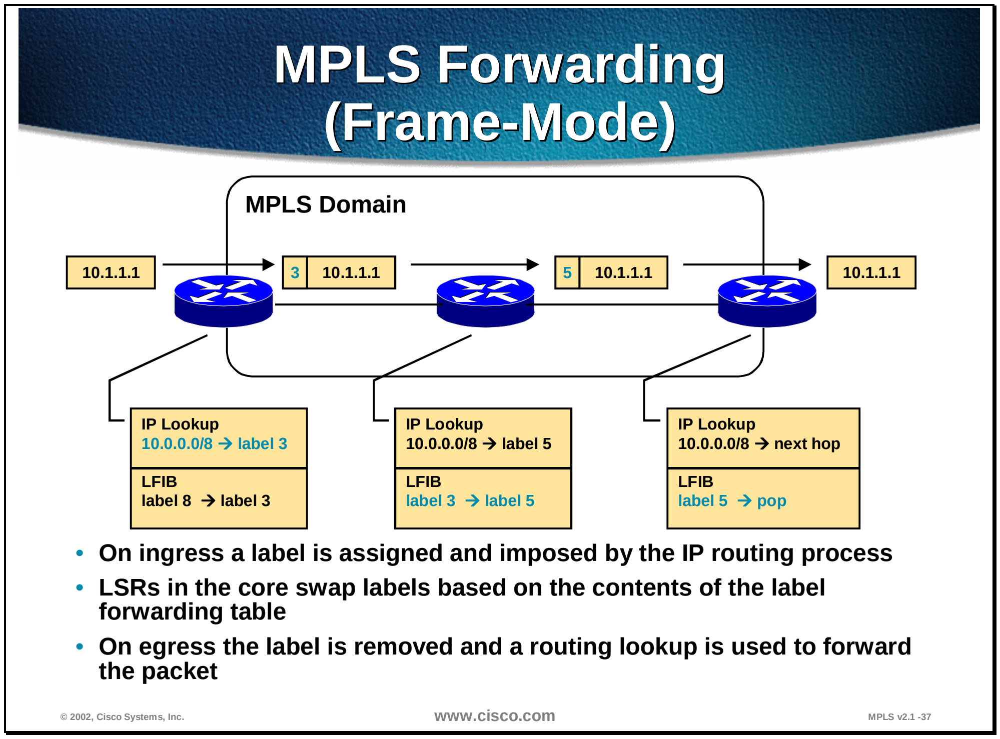

This figure illustrates an MPLS network using frame-mode MPLS. All LSRs are capable of forwarding IP packets or labeled packets. The ingress edge LSR performs a routing lookup and assigns a label. The middle router simply swaps the label. The egress LSR removes the label (penultimate hop popping is covered later) and optionally performs a routing lookup.

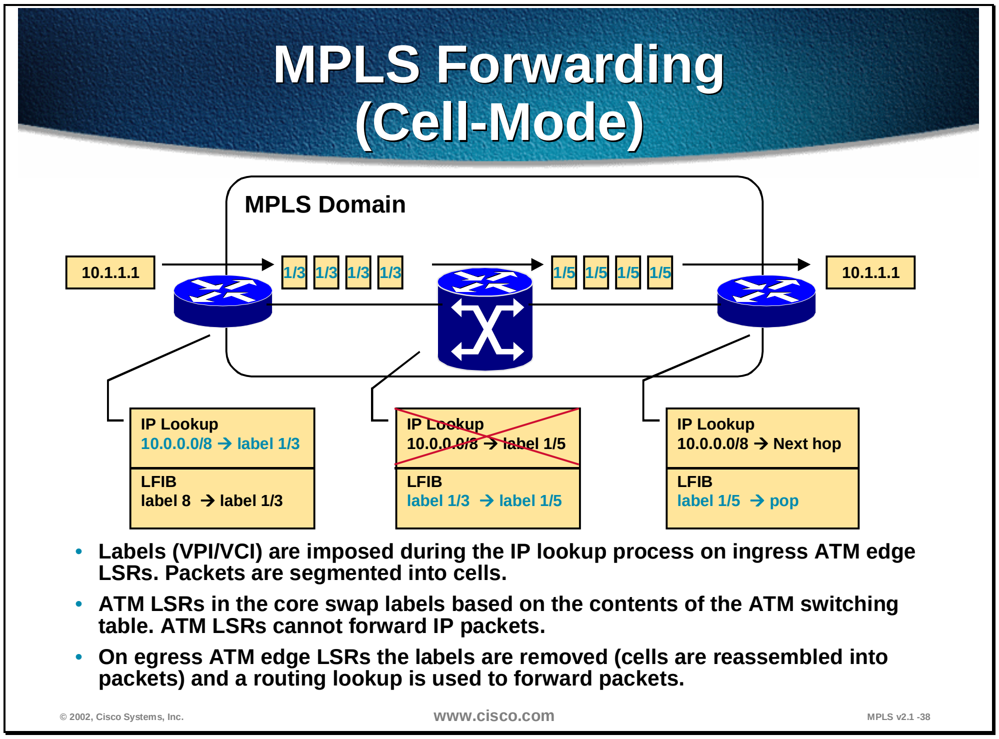

Cell-mode MPLS is similar to frame-mode MPLS. The difference is that ATM LSRs (ATM switches) are not capable of forwarding IP packets:

The ingress ATM edge LSR (router) performs an IP routing table lookup and finds the outgoing interface and the label to use. It segments the IP packet into cells and assigns each cell's VPI/VCI field the label value.

The ATM LSR in the core (ATM switch) receives ATM cells and uses the VPI/VCI field to forward those cells. This is equivalent to ATM switching. The ATM LSR does not have any concept of packets or frames. It can only handle individual cells. Therefore it can never do any forwarding of IP packets.

On the egress ATM edge LSR (router), the cells are reassembled to form an AAL5 frame. The label values in each cell are lost during the reassembly. Also the label between the AAL5 header and the IP header is removed and the IP packet is forwarded.

Summary

MPLS uses a 32-bit label that contains the following fields:

- 20-bit label
- Three experimental bits
- Bottom-of-stack bit
- 8-bit time-to-live (TTL) field

MPLS supports multiple labels in a single packet (label stack). Bottom-of-stack bit is used to determine the last label in the stack.

Lesson Review

1. What fields does a label have?
2. How does a receiving router know if the packet is labeled or not?
3. How does a receiving router know if there is another label?
4. Why is more than one label needed?
5. What are the major differences between frame-mode and cell-mode MPLS?

MPLS Applications

Objectives

Upon completion of this lesson, the learner will be able to perform the following tasks:

- Identify various MPLS applications
- Describe the overall structure of each MPLS application
- Explain the interactions between several MPLS applications running on the same platform

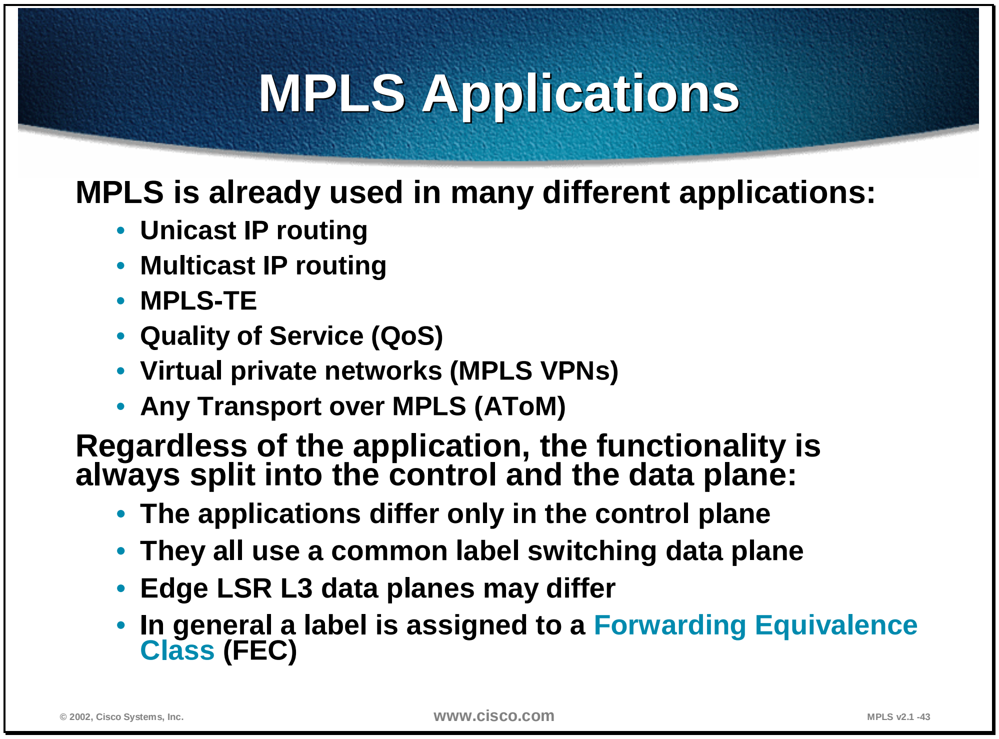

MPLS can be used in different applications:

- Unicast IP routing is the most common application for MPLS
- Multicast IP routing is treated separately because of different forwarding requirements
- MPLS-TE is an add-on to MPLS that provides better and more intelligent link utilization
- Differentiated QoS can also be provided with MPLS
- MPLS VPNs are implemented using labels to allow overlapping address space between VPNs
- AToM is allowing transport of L2 frames (or cells) across an MPLS cloud

The data plane is the same regardless of the application. The control plane however needs appropriate mechanisms to exchange routing information and labels.

The term “Forwarding Equivalence Class” (FEC) is used to describe the packets that are using the same Labeled Switched Path (LSP) across the network.

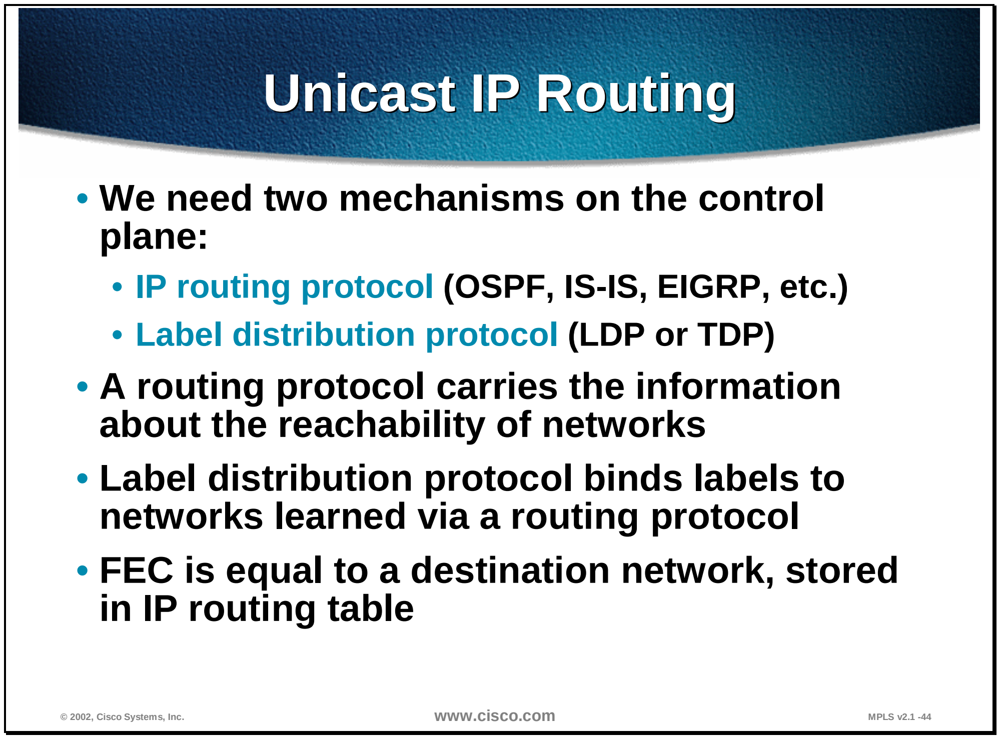

Unicast IP routing setup usually requires two components:

- IP routing protocol (OSPF, EIGRP, IS-IS, etc.)
- Label exchange protocol (TDP or LDP)

These two components are enough to create a full mesh of LSPs.

A label is assigned to every destination network found in the IP forwarding table. Therefore, a FEC corresponds to an IP destination network.

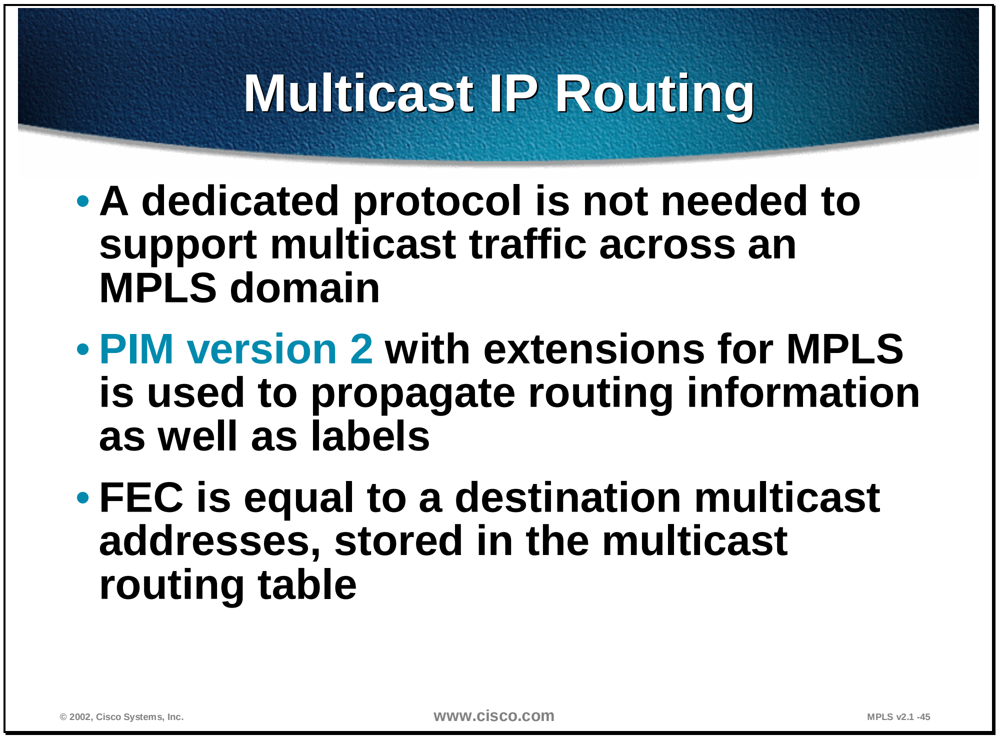

Multicast IP routing can also use MPLS. PIM version 2 with extensions for MPLS is used to propagate routing information and labels.

A FEC is equal to a destination multicast address.

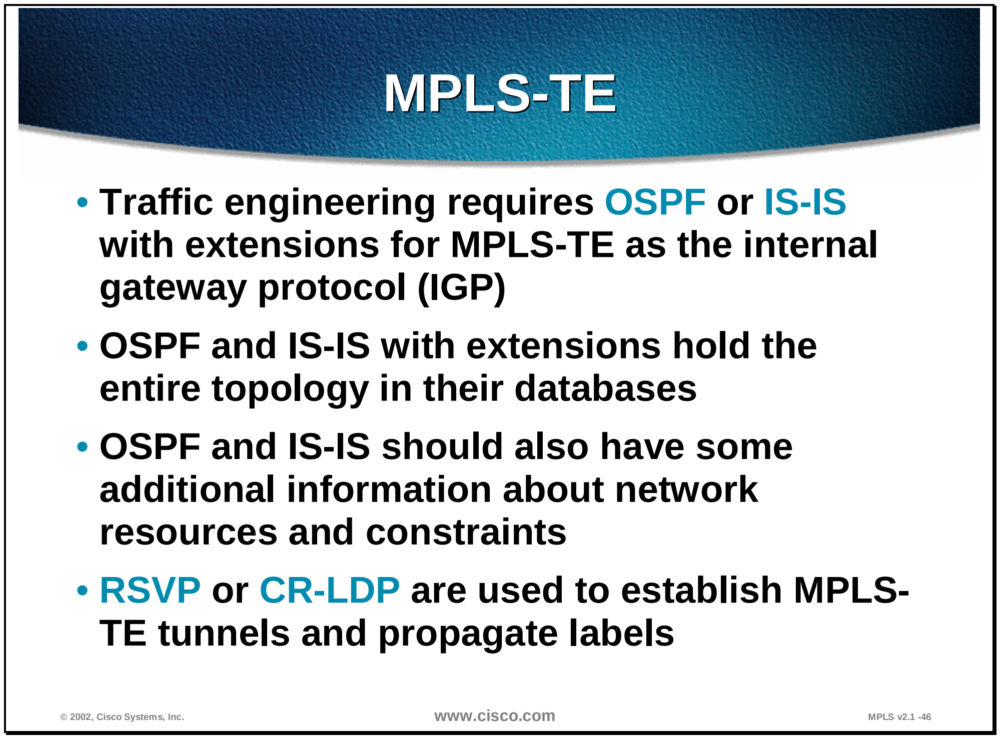

MPLS-TE has special requirements:

- Every LSR must see the entire topology of the network (only OSPF and IS-IS hold the entire topology).
- Every LSR needs additional information about links in the network (available resources and constraints). OSPF and IS-IS have extensions to propagate this additional information.
- Every edge LSR must be able to create an LSP (Label Switched Path) on demand. RSVP or Constraint-based Routing LDP (CR-LDP) is used to create an LSP and to propagate labels for MPLS-TE tunnels.

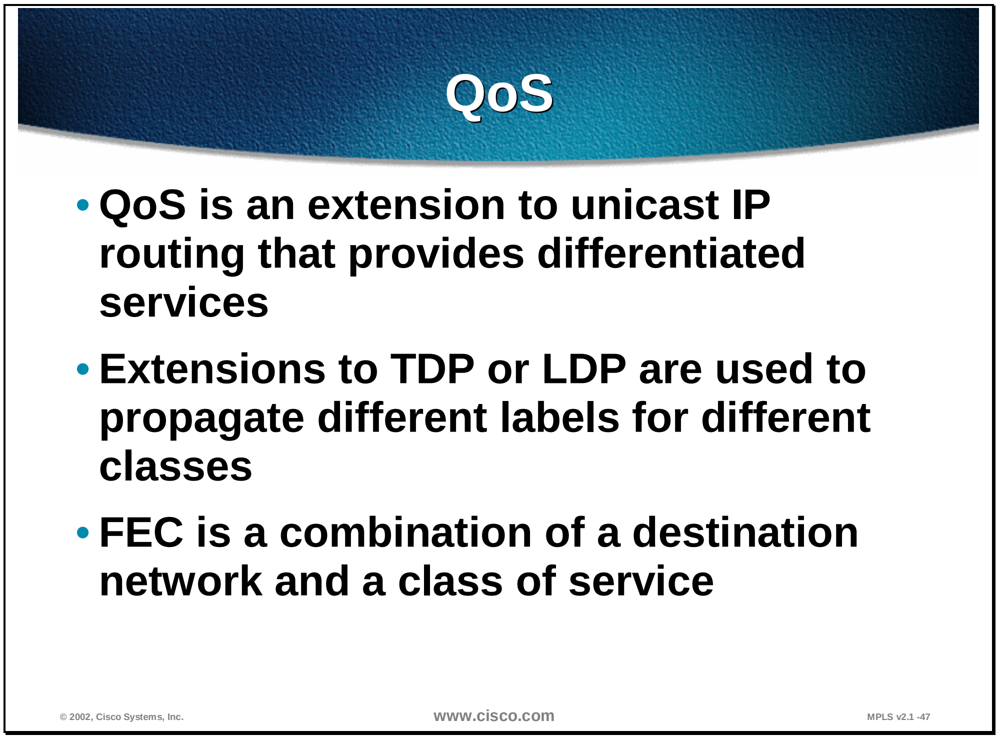

Differentiated QoS is achieved by using MPLS experimental bits or by creating separate LSPs for different classes. Extensions to TDP or LDP are used to create multiple LSPs for the same destination (one for each class).

FEC corresponds to the combination of a destination network and the class of service.

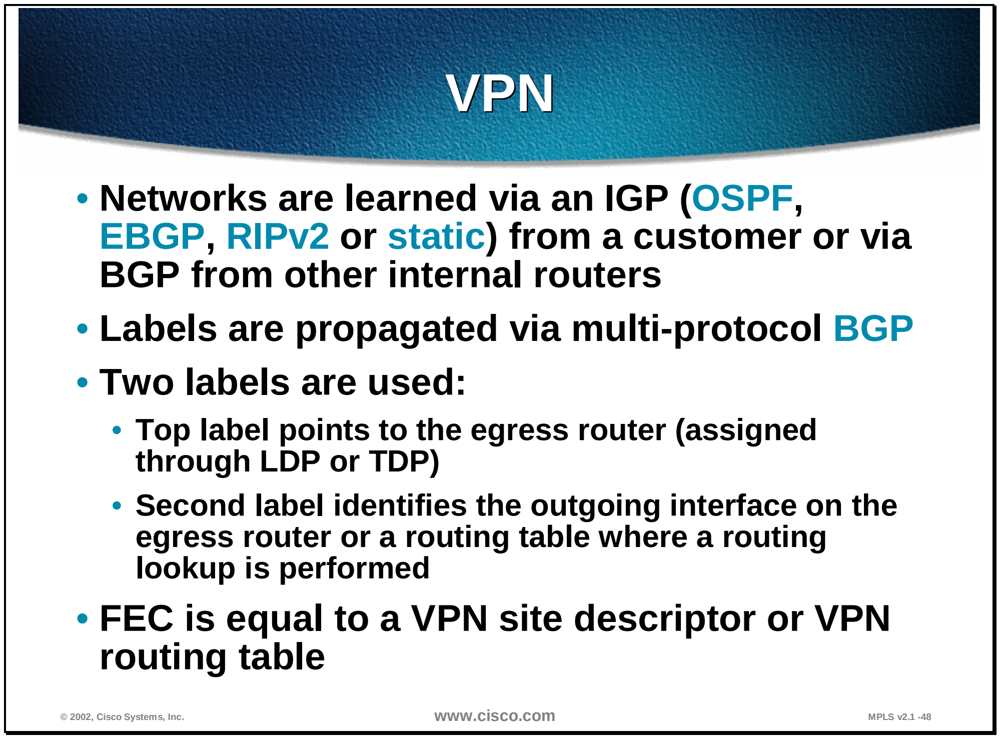

MPLS VPNs use an additional label to determine the VPN and the corresponding VPN destination network. BGP with multiprotocol extensions is used to propagate VPN routing information and labels across the MPLS domain. TDP or LDP is still needed to link edge LSRs with a single LSP.

FEC corresponds to a VPN destination network.

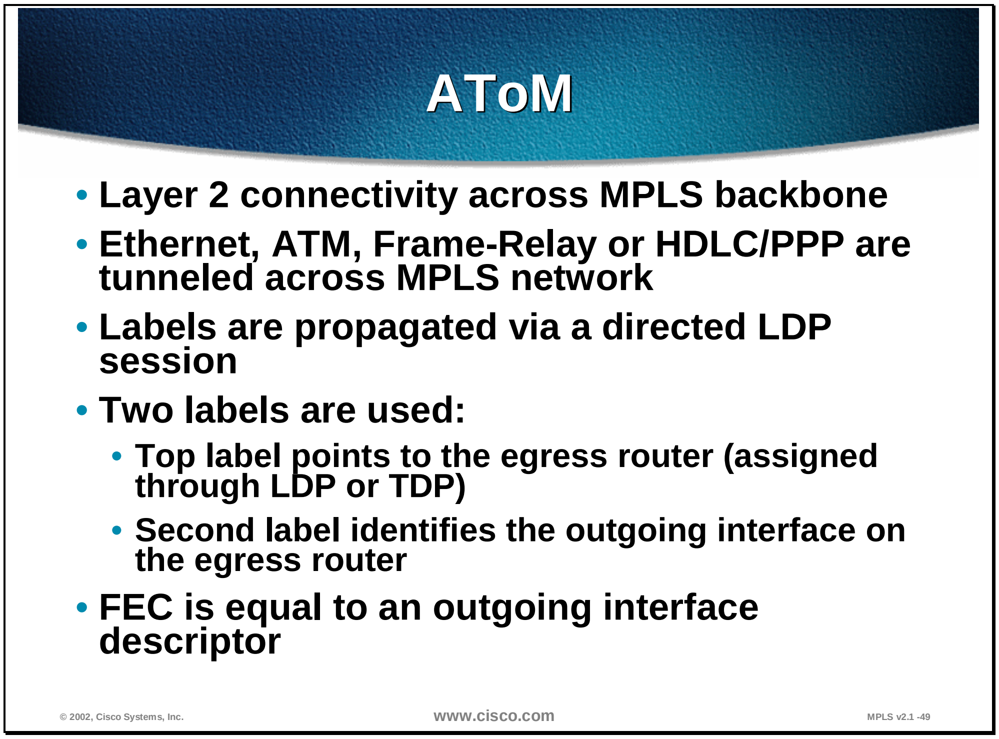

AToM provides forwarding of L2 frames (or cells) across an MPLS backbone. Ethernet, Frame-Relay, High-Level Data Link Control/Point-to-Point Protocol (HDLC/PPP) frames or ATM cells are received by the ingress edge LSR. The L2 frames (or cells) are MPLS encapsulated and assigned a stack of two labels. The top most label will direct the frame to the egress edge LSR. The second label will indicate the outgoing interface on the egress edge LSR.

A directed multihop LDP session between the ingress and egress edge LSRs is used to exchange the second label.

FEC correspond to an outgoing interface.

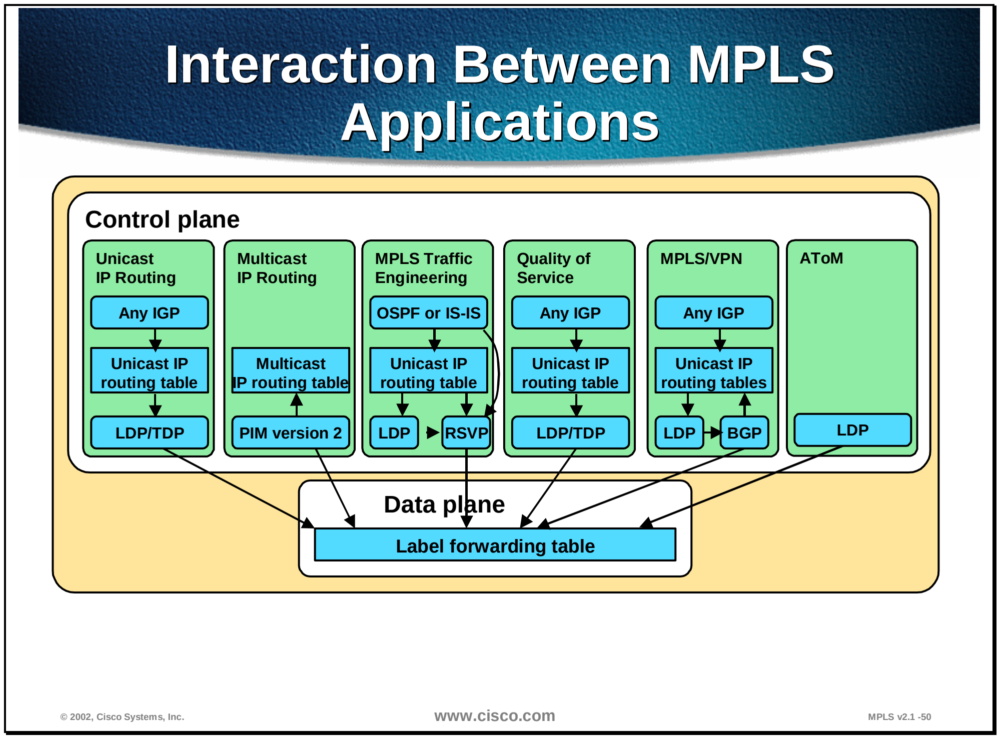

This figure illustrates the complete architecture when all applications are used. Each application may use a different routing protocol and a different label exchange protocol, but all applications use one single label-forwarding engine.

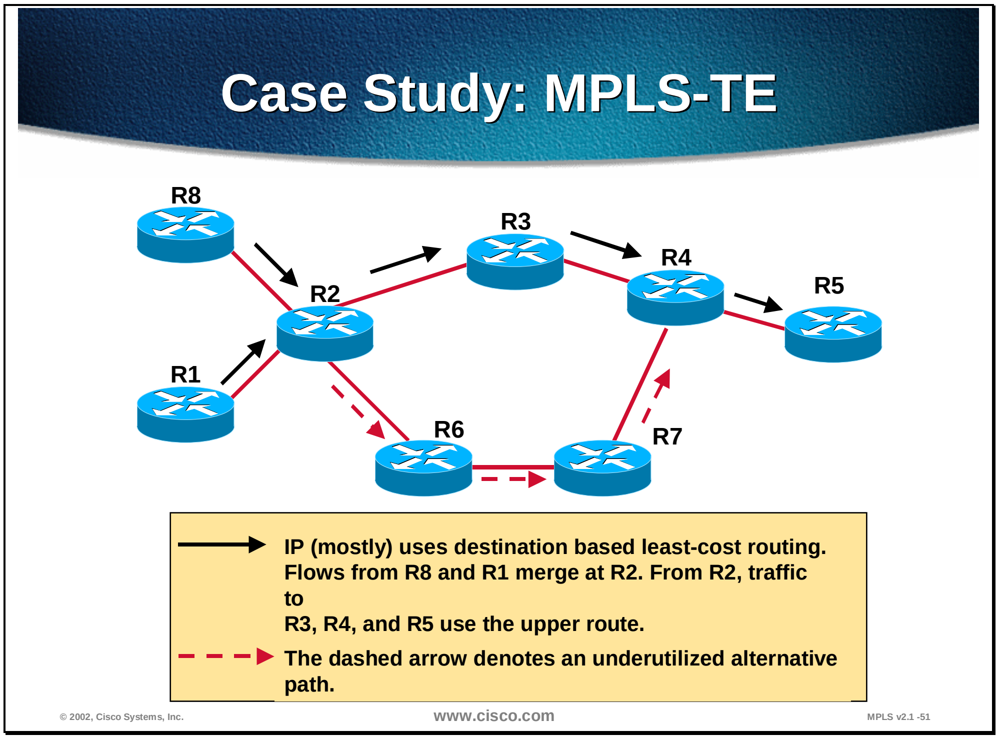

This figure illustrates the result of traditional IP forwarding:

Step 1 Traffic flows from both R1 and R8 towards R5 takes the upper path via R3. This is the result of destination based forwarding in R2.

Step 2 R2 does not make any difference if the packets arrive from R8 or from R1. R2 only cares about the destination.

As a result, the upper path may become over utilized while the lower path (via R6 and R7) may become under utilized.

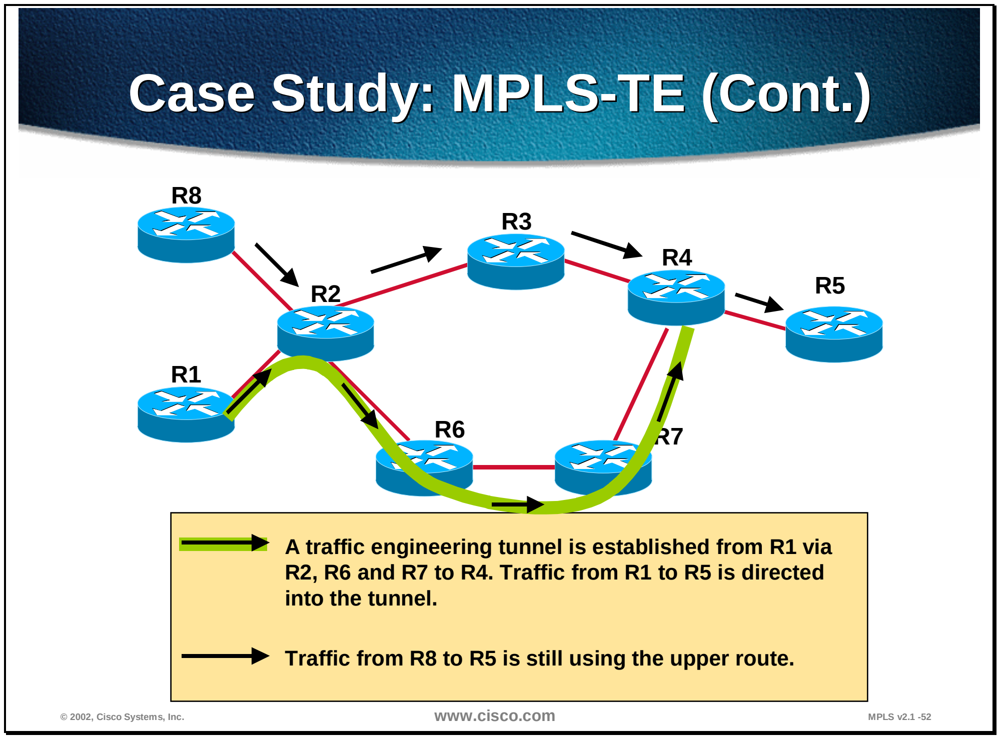

When the network operators are detecting the situation with an over utilized primary path and an under utilized alternative path, they want to move some traffic volume from the over utilized to the under utilized path.

When using traffic engineering to perform this operation, a traffic engineering tunnel is configured from R1 to R4. This tunnel is engineered to take the lower, under utilized, path.

Traffic from R1 to destinations behind R4 can now be directed by R1 into the tunnel. This moves a subset of the volume of traffic that use to take the upper path to now take the lower path. The traffic from R8 is not injected into any tunnel and still takes the upper path.

Case Study: VPN

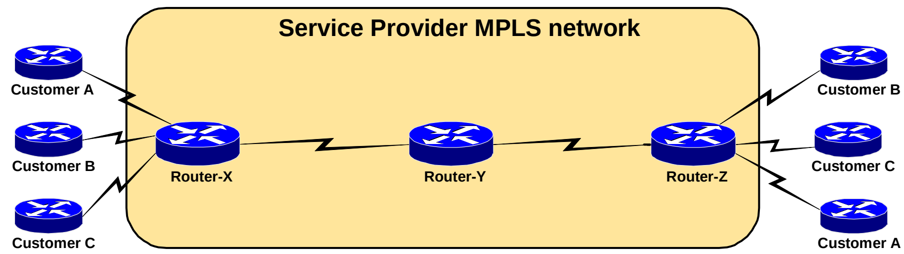

- **The service provider network provides services to customer A, B and C**
- **No leaking of packets between different customers should be possible**

 www.cisco.com MPLS v2.1 -53

In this case study, a service provider is operating an MPLS based network to provide VPN services to customer A, B and C. Customer A should be able to exchange traffic between all customer A sites, customer B between all customer B sites and the same goes for customer C. But there should be no leaking of packets between the customers.

The problem is solved by the MPLS VPN solution:

Step 1 Router-X (ingress edge router) receives an IP packet from customer A

Step 2 This packet is MPLS encapsulated and assigned a label stack of two labels

Step 3 The top most label indicates how the packet should be forwarded to reach router-Z (egress edge router)

Step 4 The second label indicates for router-Z to forward the packet out to the customer-A site

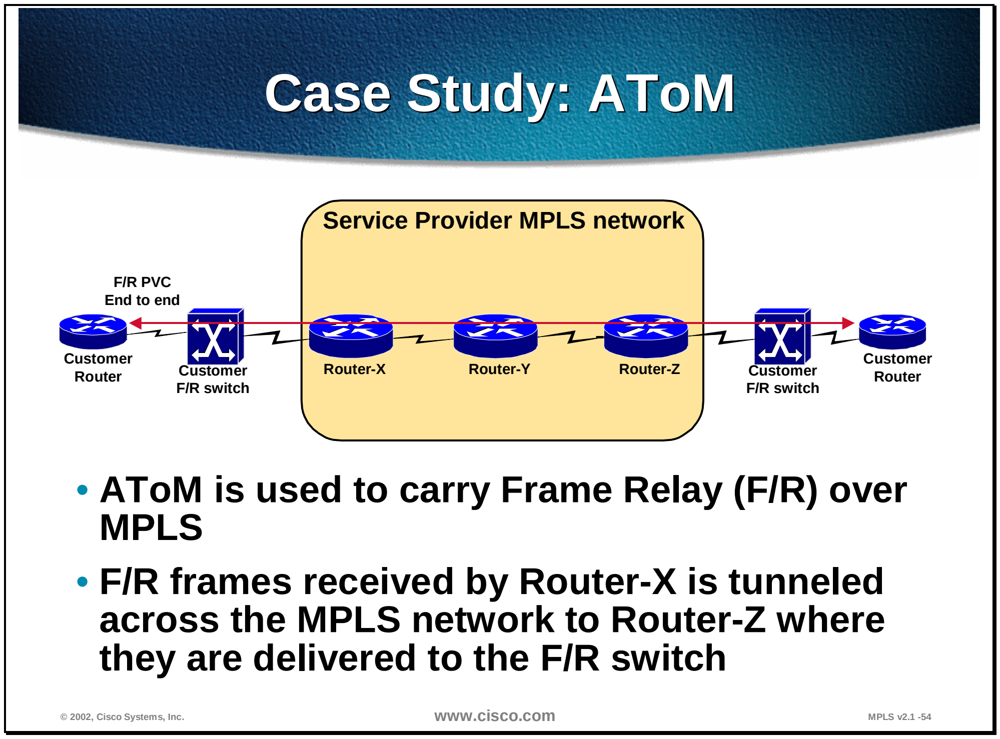

In this case study the service provider, who is operating an MPLS network, is providing forwarding of Frame Relay frames:

Step 1 Router-X (ingress edge router) receives Frame Relay (F/R) frames on a serial interface from the customer's left most F/R switch

Step 2 Router-X MPLS encapsulates them and forward them into the MPLS cloud

Step 3 This forwarding can be done on any media type supported by MPLS

Each MPLS packet is assigned a stack of two labels. The top most label will indicate how the packet should be forwarded to reach router-Z (egress edge router). The second label indicates for router-Z to forward the packet out on to the serial interface to the right most F/R switch.

For the customer's devices, this appears as transparent F/R link. This means that the two F/R switches can be configured to provide a PVC between the two customer routers. The customer routers now appear as L3 neighbors. Any traffic between them (including routing protocols) will be tunneled across the MPLS cloud. Router-X, router-Y and router-Z will not see those packets as IP packets at all.

Summary

MPLS currently supports the following applications:

- Unicast IP routing
- Multicast IP routing
- QoS
- MPLS-TE
- MPLS VPNs
- AToM

Different applications differ in the mechanisms used in the control plane. All applications use the same label forwarding mechanism in the data plane.

Lesson Review

1. Where is MPLS used today?
2. Where can MPLS potentially be used in the future?
3. What do the different applications of MPLS have in common?
4. What are the differences between the various applications of MPLS?

Differences Between Tag Switching and MPLS

Objectives

Upon completion of this lesson, the learner will be able to perform the following tasks:

- Explain the differences between Tag Switching and MPLS
- List the IETF standards defining MPLS

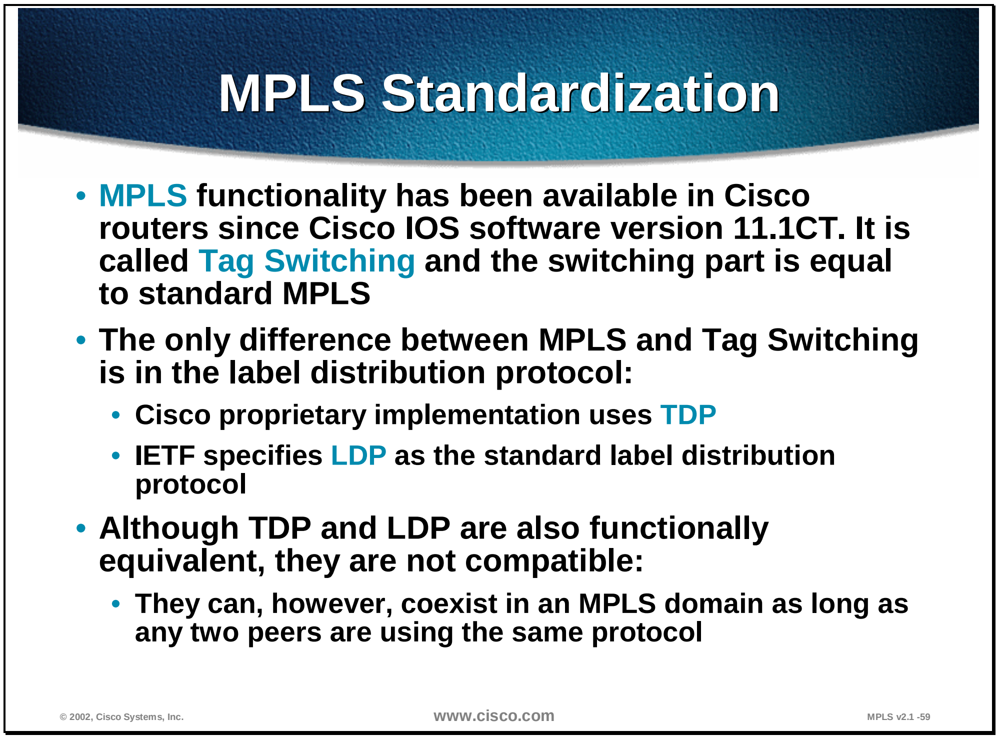

MPLS functionality has been available since Cisco IOS software version 11.1CT and was called Tag Switching. Cisco developed a proprietary label exchange protocol called TDP.

IETF's LDP is a standard label exchange protocol which is very similar to Cisco's TDP.

The only difference between MPLS and Tag Switching is that MPLS uses LDP and Tag Switching uses TDP as their label distribution protocols.

TDP and LDP are not compatible but they can coexist in an MPLS networks as long as any two peers support the same protocol.

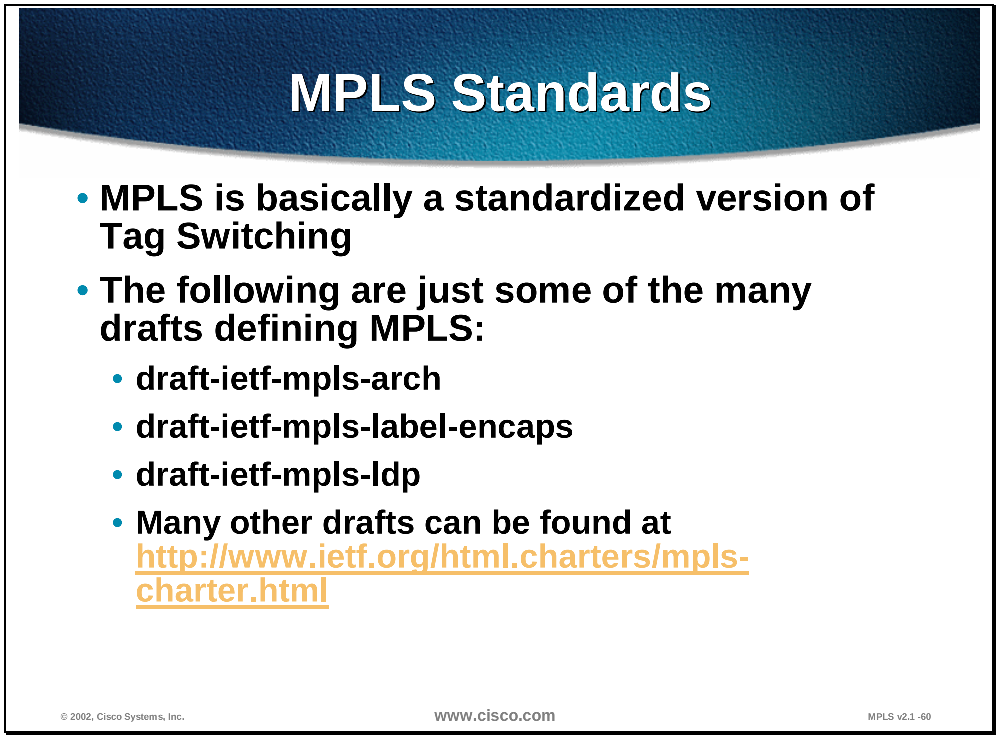

MPLS is a standardized version of Tag Switching. A large number of drafts or RFCs can be found at http://www.ietf.org.

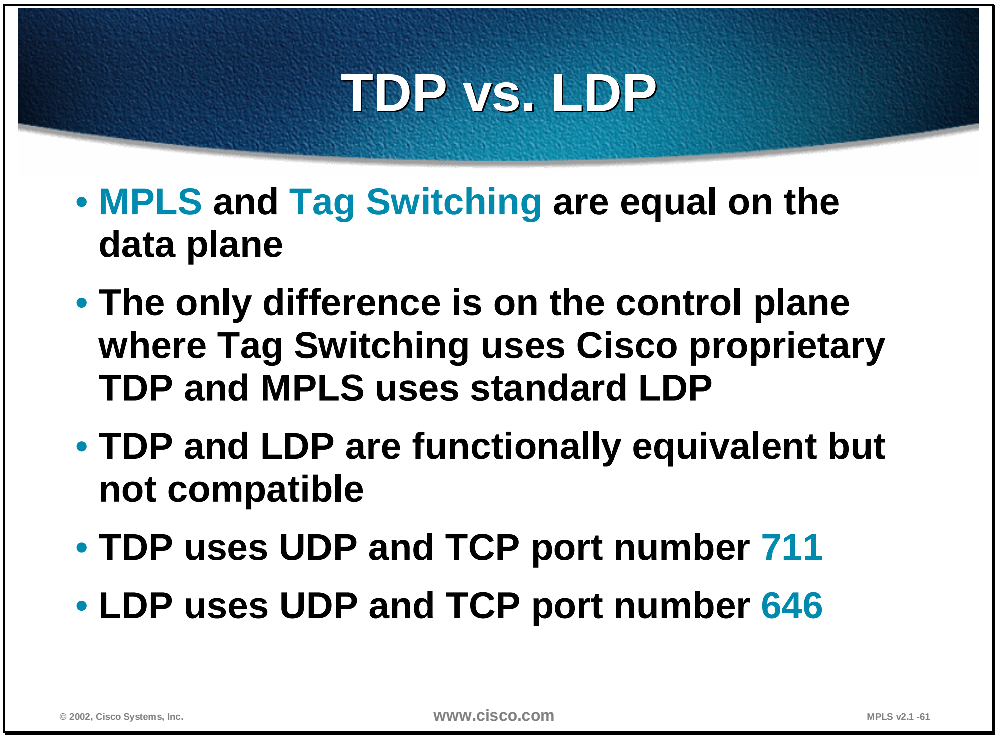

Standard MPLS and Cisco's Tag Switching use the same label format and are therefore, compatible on the data plane. TDP and LDP on the other hand, are different:

- TDP uses TCP and UDP port number 711
- LDP uses TCP and UDP port number 646

TDP and LDP can, generally speaking, perform the same functions. There are only a few exceptions where some new functionality has been added to LDP. But the two protocols are incompatible in the sense that two routers must use a common protocol, otherwise they will not interoperate.

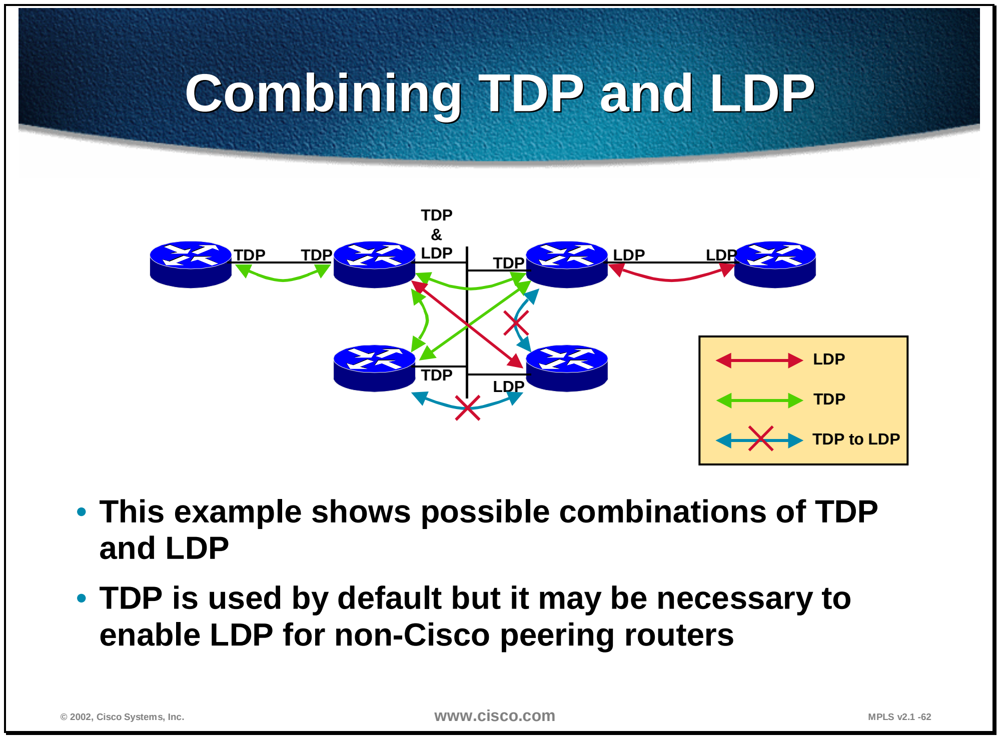

Mixing TDP and LDP in the same MPLS domain is possible during a migration from a pure TDP domain into a pure LDP domain. Many existing networks are using TDP but are gradually changed to use LDP when new functions are required which are only added to LDP and not to the legacy TDP protocol.

This figure illustrates all possible combinations of using TDP and LDP in the same MPLS domain:

- TDP to TDP works
- LDP to LDP works
- TDP to LDP does not work

Due to this incompatibility the recommendation is to avoid using a mix of LDP and TDP on the same link as on the LAN segment in the figure. During a migration phase, a single MPLS domain could very well use LDP on some links and TDP on others, but not LDP and TDP on one and the same link for any sustainable period of time.

An LSR may use both TDP and LDP on the same interface to allow label exchange with other Cisco routers (TDP) and non-Cisco routers (LDP).

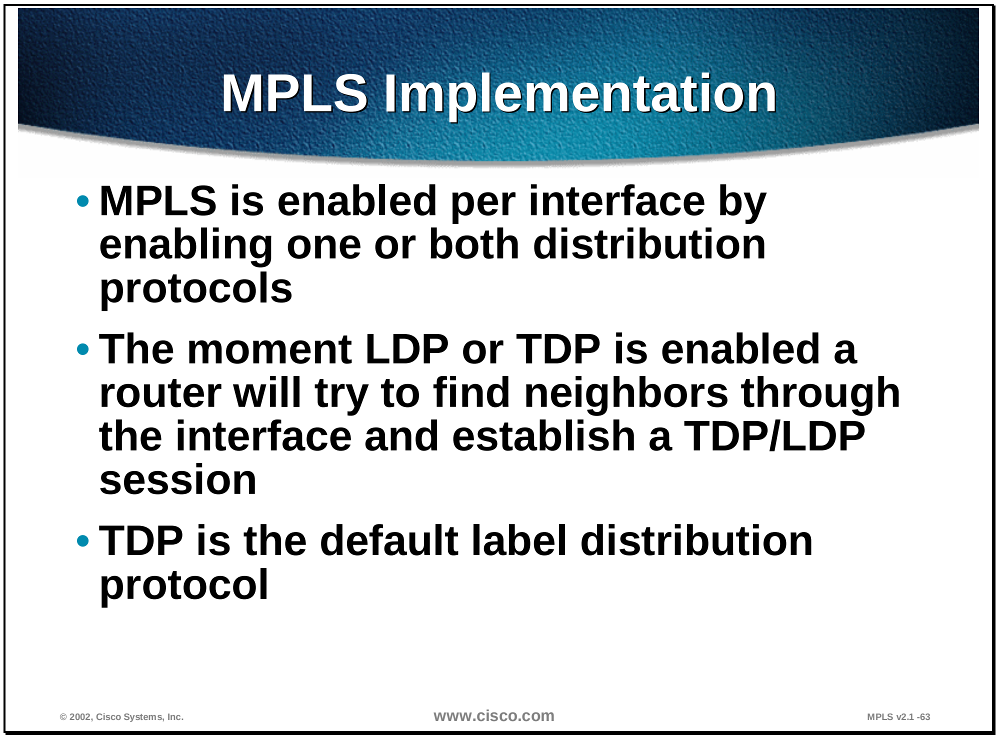

MPLS is enabled per interface by enabling one or both of the label distribution protocols (LDP or TDP).

Cisco routers use TDP by default. TDP is enabled automatically when Tag Switching or MPLS is enabled on an interface. When LDP is needed, it must either be specifically enabled per interface or the default behavior must be changed. Changing the default behavior requires configuration in global configuration mode.

Both TDP and LDP automatically try to find neighbors and establish a session with them.

Summary

- Tag Switching and MPLS are compatible on the data plane but not the control plane.
- TDP is used with Tag Switching and LDP is used with MPLS.

Lesson Review

1. What is the difference between Tag Switching and MPLS on the data plane?
2. What is the difference between Tag Switching and MPLS on the control plane?
3. How can Tag Switching and MPLS be combined?
4. Which label switching mechanism is the default switching mechanism on Cisco routers?

Summary

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

- Identify the drawbacks of traditional IP routing.
- Describe basic MPLS concepts and LSR types.
- Describe how different MPLS applications can coexist on the same platform using the same underlying technology.
- List the standard bodies that are working on MPLS technology and the relationship between Tag Switching and MPLS.