MPLS AToM — Overview

This document provides an introduction to MPLS AToM and includes the following sections:

Documentation Specifics

This documentation set includes the following sections:

- *Start Here: MPLS AToM: Transport, Platform, and Release Specifics*
- *MPLS AToM: Overview* (this document)
- *MPLS AToM: Configuring*
- *MPLS AToM: Commands*

Note *Start Here: MPLS AToM: Transport, Platform, and Release Specifics* details the features that are supported in each release and on each platform. Not all MPLS AToM features are supported in each Cisco IOS software release for each platform. Read the entire chapter before reading the other MPLS AToM chapters.

The other chapters provide overview, configuration, and command reference information for MPLS AToM features.

Feature Overview

Any Transport over MPLS (AToM) is a solution for transporting Layer 2 packets over an MPLS backbone. AToM enables service providers to supply connectivity between customer sites with existing data link layer (Layer 2) networks by using a single, integrated, packet-based network infrastructure — a Cisco MPLS network. Instead of separate networks with network management environments, service providers can deliver Layer 2 connections over an MPLS backbone.

With Cisco AToM technology, provisioning and connecting is straightforward. A customer using Ethernet in a building or campus in one location can connect through a service provider offering Ethernet over MPLS to the customer's Ethernet networks in remote locations.

AToM provides a common framework to encapsulate and transport supported Layer 2 traffic types over an MPLS network core. Service providers can use a single MPLS network infrastructure to offer customers connectivity for supported Layer 2 traffic, as well as customers' IP traffic in Layer 3 VPNs.

AToM supports the following transport types:

- ATM AAL5 over MPLS
- ATM Cell Relay over MPLS
- Ethernet VLAN over MPLS
- Frame Relay over MPLS
- PPP over MPLS
- HDLC over MPLS

ATM AAL5 over MPLS

How ATM AAL5 SDUs Move Between PE Routers

ATM AAL5 over MPLS encapsulates ATM AAL5 service data units (SDUs) in MPLS packets and forwards them across the MPLS network. Each AAL5 SDU is transported as a single packet. The following steps outline the process of encapsulating the SDU.

Ingress PE router

1. An ingress provider edge (PE) router receives an ATM AAL5 SDU and removes the header.
2. The PE router copies the control word elements from the header to the corresponding fields in the control word of the SDU. The control word contains:
   - Explicit forward congestion indication (EFCI) bit—Used by ATM switches to indicate congestion experienced by forwarded data cells.
   - Cell loss priority (CLP) bit—Indicates whether a cell should be dropped if it encounters extreme congestion as it moves through the ATM network.
3. The PE router adds a virtual circuit (VC) label and a label switched path (LSP) tunnel label to the packet for normal MPLS routing through the MPLS backbone. The core routers use the LSP tunnel label to move the packet through the MPLS backbone. A core router does not distinguish ATM AAL5 traffic from other types of traffic. The packet is handled just like other packets in the MPLS backbone.

Egress PE router

1. At the other edge of the MPLS backbone, the egress PE router receives the packet and copies the control word elements from the control word to the header.
2. The PE router removes the VC label and LSP tunnel label if one is present. If no LSP tunnel label is present, it is because the penultimate router removed that label.
3. The PE router adds an AAL5 header and sends the packet out the appropriate customer-facing interface.

Figure 1 illustrates this process.

Figure 1 ATM AAL5 Packets as They Traverse the MPLS Backbone

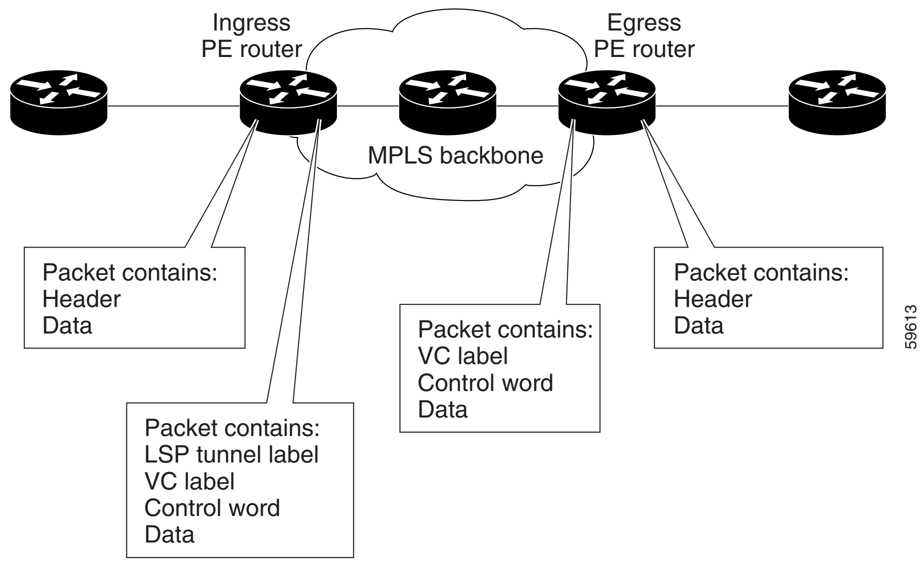

AAL5 Packets Containing OAM Cells

The Cisco 7200 and 7500 series routers support the transport of F5 end-to-end operational, administrative, and maintenance (OAM) cells. Only Mode 0 is supported. F5 OAM cells are transported over the MPLS backbone with the payload. The OAM cell fits into the payload of a single AAL5 packet.

Notes:

- PVC switching is not supported in OAM encapsulation.
- Both PE routers must be configured with the same VPI/VCI value.
- OAM transparency is not supported on the Cisco 12000 series routers.

OAM Cell Emulation

Supported Platforms:

This functionality is supported on the following platforms:

- Cisco 7200 series routers
- Cisco 7500 series routers

If a PE router does not support the transport of OAM cells across an LSP, you can use OAM cell emulation to locally terminate or loopback the OAM cells. You configure OAM cell emulation on both PE routers, which emulates a VC by forming two unidirectional LSPs. You use the **oam-ac emulation-enable** command on both PE routers to enable OAM cell emulation.

After OAM cell emulation is enabled on a router, you can configure and manage the ATM VC in the same manner as you would a terminated VC. A VC that has been configured with OAM cell emulation can send loopback cells at configured intervals toward the local CE router. The endpoint can be either of the following:

- End-to-end loopback, which sends OAM cells to the local CE router.

- Segment loopback, which responds to OAM cells to a device along the path between the PE and CE routers.

The OAM cells include the following:

- Alarm indication signal (AIS)
- Remote defect indication (RDI)

These cells identify and report defects along a VC. When a physical link or interface failure occurs, intermediate nodes insert OAM AIS cells into all the downstream devices affected by the failure. When a router receives an AIS cell, it marks the ATM VC down and sends an RDI cell to let the remote end know about the failure.

See the Configure OAM Cell Emulation for ATM AAL5 over MPLS section for information on configuring OAM cell emulation.

ATM Cell Relay over MPLS

ATM Cell Relay over MPLS transports single ATM cells over the MPLS backbone. The AToM circuit is configuring on permanent virtual circuits. In this release, only PVC mode, single cell relay is supported.

How ATM Cells Move Between PE Routers

ATM Cell Relay over MPLS encapsulates ATM cells in MPLS packets and forwards them across the MPLS network. Each MPLS packet contains one ATM cell. In other words, each ATM cell is transported as a single packet. The following steps outline the process of encapsulating the ATM cell.

Ingress PE Router

1. The ingress PE router receives an ATM cell and removes the header. The following items are not removed from the ATM cell:
   - The control word. The control word contains:

     Explicit forward congestion indication (EFCI) bit **—** Used by ATM switches to indicate congestion experienced by forwarded data cells.

     Cell loss priority (CLP) bit — indicates whether a cell should be dropped if it encounters extreme congestion as it moves through the ATM network.
   - The virtual path identifier (VPI) and virtual channel identifier (VCI). The VPI and VCI identify the next destination of a cell as it passes through a series of ATM switches on its way to its destination. ATM switches use the VPI/VCI fields to identify the next virtual channel link (VCL) that a cell needs to transit on its way to its final destination.
2. The PE router adds a VC label and an LSP tunnel label to the packet for normal MPLS routing through the MPLS backbone. The P routers use the LSP tunnel label to move the packet through the MPLS backbone. A core router does not distinguish ATM Cell Relay traffic from other types of traffic. The packet is handled just like other packets in the MPLS backbone.

Egress PE Router

1. At the other edge of the MPLS backbone, the egress PE router receives the packet and removes the LSP tunnel label if one is present. If no LSP tunnel label is present, it is because the penultimate router removed that label. The PE router also removes the control word and VC label from the packet.

2. The PE router adds an ATM header and sends the packet out the appropriate customer facing interface.

Figure 2 illustrates this process.

Figure 2 ATM Cell Packets as They Traverse the MPLS Backbone

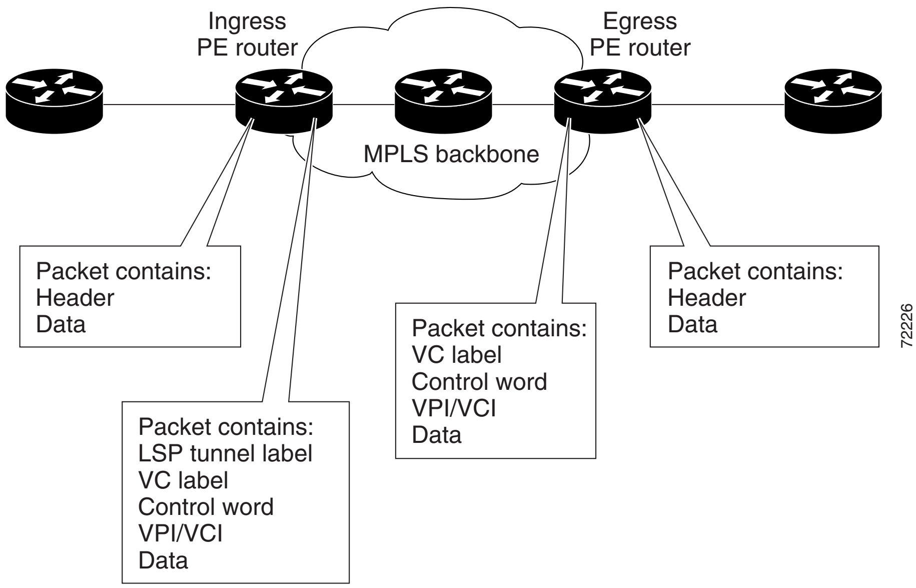

ATM Packets Containing OAM Cells

If F5 end-to-end operational, administrative, and maintenance (OAM) cells are included in a packet, they are transported over the MPLS backbone with the payload. The OAM cell fits into the payload of a single packet. The Cisco 7200 and 7500 series routers support the transport of F5 end-to-end OAM cells. Only

Ethernet over MPLS

How Ethernet PDUs Move Between PE Routers

Ethernet over MPLS works by encapsulating Ethernet PDUs in MPLS packets and forwarding them across the MPLS network. Each PDU is transported as a single packet. The following steps outline the process of encapsulating the PDU.

Ingress PE Router:

1. The ingress PE router receives a PDU and removes the preamble, start of frame delimiter (SFD), and the frame check sequence (FCS). The rest of the header remains the same.
2. The PE router copies the control word from the header, even though it is not used. The PE router adds a VC label and LSP tunnel label for normal MPLS routing through the MPLS backbone. The core routers use the LSP tunnel label to move the packet through the MPLS backbone. A core router does not distinguish Ethernet traffic from other types of traffic. The packet is handled just like other packets in the MPLS backbone.

Egress PE Router

1. At the other edge of the MPLS backbone, the egress PE router receives the packet and removes the LSP tunnel label if one is present. If no LSP tunnel label is present, it is because the penultimate router removed that label. The PE router also removes the control word and VC label from the packet.
2. The PE router updates the header if necessary and sends the packet out the appropriate customer facing interface.

Figure 3 illustrates this process.

Figure 3 Ethernet Packets as They Traverse the MPLS Backbone

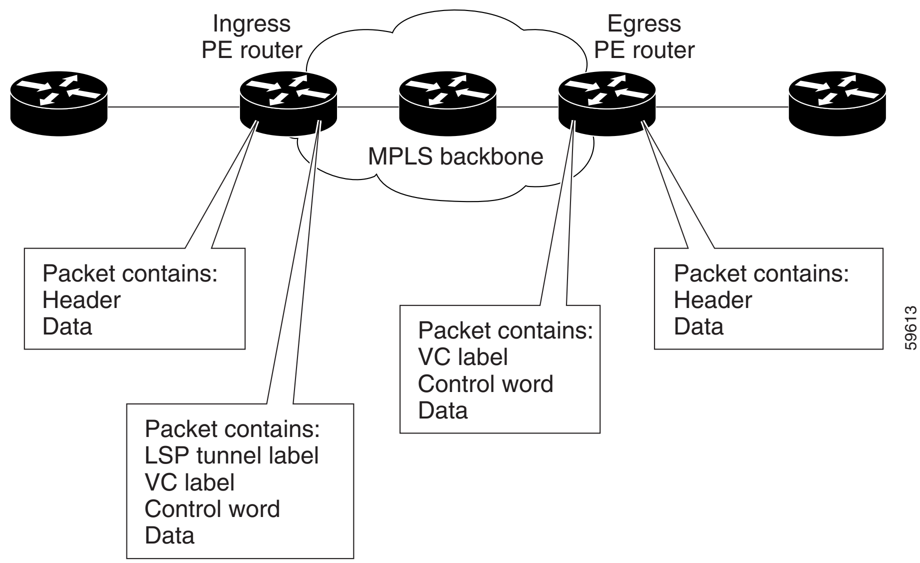

Frame Relay over MPLS

How Frame Relay PDUs Move Between PE Routers

Frame Relay over MPLS encapsulates Frame Relay protocol data units (PDUs) in MPLS packets and forwards them across the MPLS network. The process of transporting the PDU differs, depending on whether you set up DLCI-to-DLCI connections or port-to-port connections. The following sections explain both processes.

How Frame Relay Packets Move Between PE Routers with DLCI-to-DLCI Connections

The following steps outline the process of encapsulating the PDU in a Frame Relay configuration with DLCI-to-DLCI connections.

Ingress PE router

1. An ingress PE router receives a Frame Relay PDU and removes the Frame Relay header and the frame check sequence (FCS).
2. The PE router copies the control word elements from the Frame Relay header to the corresponding fields in the control word of the Frame Relay PDU. The control word elements include:

 - Backward explicit congestion notification (BECN)
 - Forward explicit congestion notification (FECN)
 - Discard eligibility (DE)
 - Command/response

3. The PE router adds a VC label and an LSP tunnel label to the packet for normal MPLS routing through the MPLS backbone. The core routers use the LSP tunnel label to move the packet through the MPLS backbone. A core router does not distinguish Frame Relay traffic from other types of traffic. The packet is handled just like other packets in the MPLS backbone.

Egress PE router

1. At the other edge of the MPLS backbone, the egress PE router receives the packet and copies the control word elements from the control word to the Frame Relay header.
2. The egress PE router removes the VC label and LSP tunnel label if one is present. If no LSP tunnel label is present, it is because the penultimate router removed that label.
3. The PE router adds the Frame Relay header and sends the packet out the appropriate customer-facing interface.

Figure 4 illustrates this process.

Figure 4 Frame Relay Packets as They Traverse the MPLS Backbone

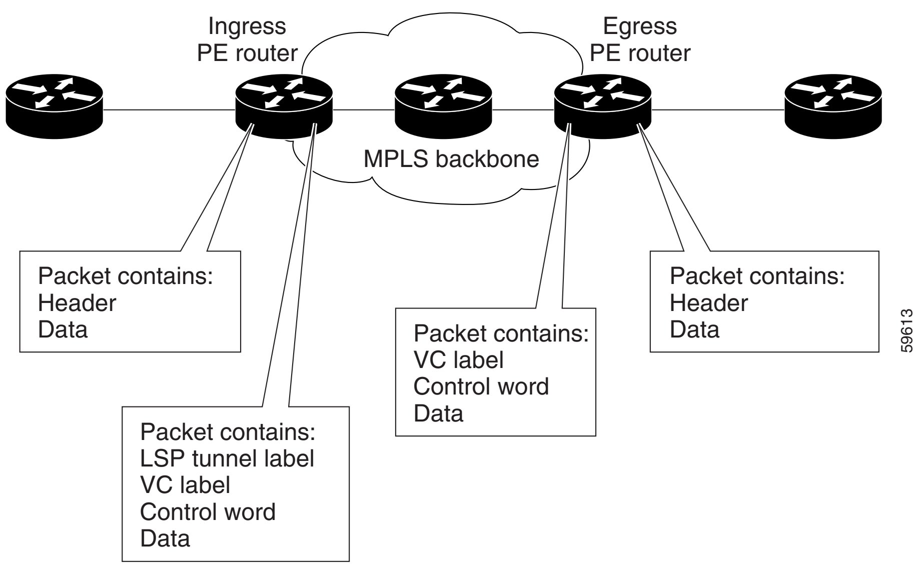

How Frame Relay Packets Move Between PE Routers with Port-to-Port Connections

When you set up a port-to-port connection between PE routers, you use HDLC mode to transport the Frame Relay encapsulated packets. In HDLC mode, the whole HDLC packet is transported. Only the HDLC flags and FCS bits are removed. The contents of the packet are not used or changed, including the FECN, BECN, and DE bits. For more information about the HDLC packets, see the "How HDLC Packets Move Between PE Routers" section on page 22.

Local Management Interface and Frame Relay over MPLS

Local Management Interface (LMI) is a protocol that communicates status information about permanent virtual circuits (PVCs). When a PVC is added, deleted, or changed, the LMI notifies the endpoint of the status change. LMI also provides a polling mechanism that verifies that a link is up.

How LMI Works

To determine the PVC status, LMI checks that a PVC is available from the reporting device to the Frame Relay end-user device. If PVC is available, LMI reports that the status is "Active." A status of Active means that all interfaces, line protocols and core segments are operational between the reporting device and the Frame Relay end-user device. If any of those components is not available, the LMI reports a status of "Inactive."

Note Only the data circuit-terminating equipment (DCE) and network-to-network interface (NNI) interface types can report LMI status.

Figure 5 is a sample topology that helps illustrate how LMI works.

Figure 5 Sample Topology

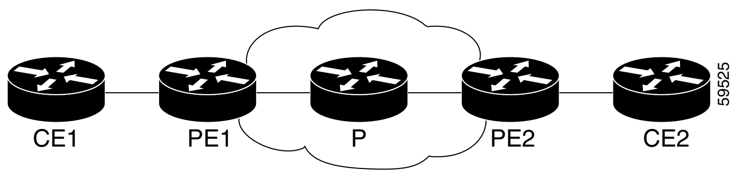

In Figure 5, note the following:

- CE1 and PE1 and PE2 and CE2 are Frame Relay LMI peers.
- CE1 and CE2 can be Frame Relay switches or end-user devices.
- Each Frame Relay PVC is composed of multiple segments.
- The DLCI value is local to each segment and is changed as traffic is switched from segment to segment. Two Frame Relay PVC segments exist in Figure 5; one is between PE1 and CE1 and the other is between PE2 and CE2.

How the LMI protocol behaves depends on whether you have DLCI-to-DLCI or port-to-port connections.

DLCI-to-DLCI Connections

If you have DLCI-to-DLCI connections, LMI runs locally on the Frame Relay ports between the PE and CE devices.

- CE1 sends an active status to PE1 if the PVC for CE1 is available. If CE1 is a switch, LMI checks that the PVC is available from CE1 to the user device attached to CE1.
- PE1 sends an active status to CE1 if the following conditions are met:
  - A PVC for PE1 is available.
  - PE1 has received an MPLS label from the remote PE router.
  - An MPLS tunnel label exists between PE1 and the remote PE.

- CE2 reports an Active status to PE2. If CE2 is a switch, LMI checks that the PVC is available from PE1 to the end user device attached to CE2.

For data terminal equipment (DTE)/DCE configurations, the following LMI behavior exists:

The Frame Relay device accessing the network (DTE) does the polling. The network device (DCE) responds to the LMI polls. Therefore, if a problem exists on the DTE side, the DCE is not aware of the problem, because it does not poll.

Port-to-Port Connections

If you have port-to-port connections, the PE routers do not participate in the LMI status-checking procedures. LMI operates between the customer edge (CE) routers only. The CE routers must be configured as DCE-DTE or NNI-NNI.

HDLC over MPLS

How HDLC Packets Move Between PE Routers

HDLC over MPLS encapsulates HDLC protocol data units (PDUs) in MPLS packets and forwards them across the MPLS network. The PE routers do not participate in any protocol negotiation or authentication. The following steps outline the process of encapsulating the PDU.

Ingress PE Router

1. An ingress PE router receives an HDLC PDU and removes the flags and the frame check sequence (FCS).
2. The PE router copies the control field to the PDU, even though the control field is not used. The PE router adds a VC label and LSP tunnel label for normal MPLS routing through the MPLS backbone. The core routers use the LSP tunnel label to move the packet through the MPLS backbone. A core router does not distinguish HDLC traffic from other types of traffic. The packet is handled just like other packets in the MPLS backbone.

Egress PE Router

1. At the other edge of the MPLS backbone, the PE router receives the packet and removes the VC label and the LSP tunnel label if one is present. If no LSP tunnel label is present, it is because the penultimate router removed that label.
2. The PE router adds the flags and FCS and sends the packet out the appropriate customer facing interface.

Figure 6 illustrates this process.

Figure 6 HDLC Packets as They Traverse the MPLS Backbone

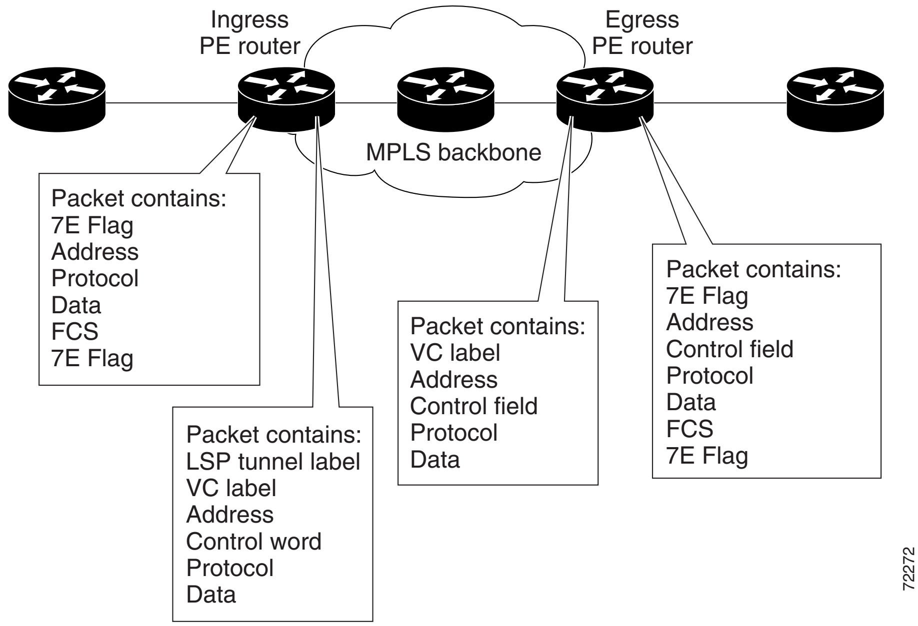

PPP over MPLS

How PPP Packets Move Between PE Routers

PPP over MPLS encapsulates PPP PDUs in MPLS packets and forwards them across the MPLS network. The PE routers do not participate in any protocol negotiation or authentication. The following steps outline the process of encapsulating the PDU.

Ingress PE Router

1. An ingress PE router receives a PPP PDU and removes the flags, address, control field, and the frame check sequence (FCS).
2. The PE router adds a VC label and LSP tunnel label to the packet for normal MPLS routing through the MPLS backbone. The core routers use the LSP tunnel label to move the packet through the MPLS backbone. A core router does not distinguish PPP traffic from other types of traffic. The packet is handled just like other packets in the MPLS backbone.

Egress PE Router

1. At the other edge of the MPLS backbone, the egress PE router receives the packet and removes the VC label and LSP tunnel label if one is present. If no LSP tunnel label is present, it is because the penultimate router removed that label.
2. The PE router adds the flags, address, control field, and FCS and sends the packet out the appropriate customer facing interface.

Figure 7 illustrates this process.

Figure 7 PPP Packets as They Traverse the MPLS Backbone

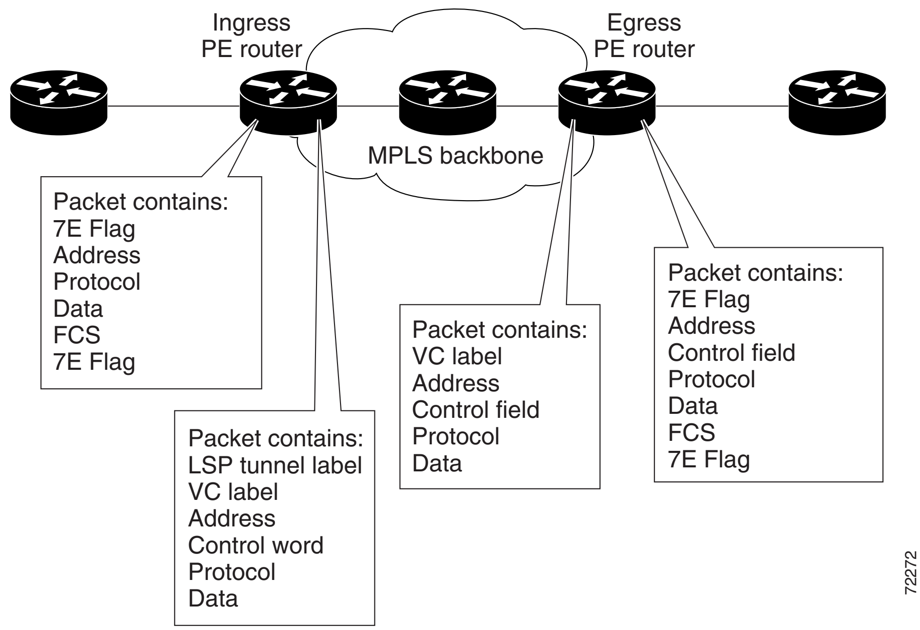

Benefits

The following list explains some of the benefits of enabling Layer 2 packets to be sent in the MPLS network:

- The AToM product set accommodates many types of Layer 2 packets, including Ethernet and Frame Relay, across multiple Cisco router platforms, such as the Cisco 7200 and 7500 series routers. This enables the service provider to transport all types of traffic over the backbone and accommodate all types of customers.
- AToM adheres to the standards developed for transporting Layer 2 packets over MPLS. (See the "Supported Standards, MIBs, and RFCs" section for the specific standards that AToM follows.) This benefits the service provider who wants to incorporate industry-standard methodologies in the network. Other Layer 2 solutions are proprietary, which can limit the service provider's ability to expand the network and can force the service provider to use only one vendor's equipment.
- Upgrading to AToM is transparent to the customer. Because the service provider network is separate from the customer network, the service provider can upgrade to AToM without disruption of service to the customer. The customers assume that they are using a traditional Layer 2 backbone.

What To Do Next

See the following MPLS AToM documentation for more information:

- *Start Here: MPLS AToM: Transport, Platform, and Release Specifics*
- *MPLS AToM: Configuring*
- *MPLS AToM: Commands*