

Chapter 4

Layer 2 Services over MPLS Overview

This chapter contains the following sections:

- Layer 2 Services over MPLS Overview on page 485
- Layer 2 Services over MPLS Platform Considerations on page 486
- Layer 2 Services over MPLS References on page 488
- Layer 2 Services over MPLS Implementation on page 489
- Local Cross-Connects Between Layer 2 Interfaces Using MPLS on page 489
- MPLS Shim Interfaces for Layer 2 Services over MPLS on page 490
- Multiple Layer 2 Services over MPLS on page 491
- ATM Layer 2 Services over MPLS on page 492
- HDLC Layer 2 Services over MPLS on page 496

Layer 2 Services over MPLS Overview

Many Internet service providers offer multiple services such as Frame Relay, Asynchronous Transfer Mode (ATM), Ethernet, High-Speed Data Link Control (HDLC), and IP to their customers, but are consolidating to a single, packet-based, optical network from several service-specific legacy layer 2 networks. Although legacy layer 2 network links are disappearing from the Internet service provider's network, the legacy layer 2 network links and services to customers must remain.

Layer 2 services over IP/MPLS enable service providers to emulate the legacy layer 2 network links and services over their IP/MPLS-based network. Layer 2 services over MPLS are especially desirable because MPLS provides features such as traffic engineering and fast reroute. These MPLS features can be used to emulate certain layer 2 service characteristics. From the perspective of the customer edge (CE) devices, all that exists is the layer 2 circuit, even though the circuit actually exists over the service provider's MPLS network.

The JUNOS software currently support the following layer 2 services over MPLS:

- ATM with ATM Adaptation Layer 5 (AAL5) encapsulation
- ATM with virtual channel connection (VCC) cell relay encapsulation
- Ethernet (Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, bridged Ethernet, bridged Ethernet/VLAN, Ethernet/VLAN)




NOTE: For the purposes of configuring layer 2 services over MPLS, Ethernet interfaces and bridged Ethernet interfaces function identically, as do Ethernet/VLAN interfaces and bridged Ethernet/VLAN interfaces. For simplicity, the term *Ethernet* in this chapter refers to both Ethernet interfaces and bridged Ethernet interfaces, and the term *Ethernet/VLAN interfaces* refers to both Ethernet/VLAN interfaces and bridged Ethernet/VLAN interfaces.

- Frame Relay
- HDLC

Figure 113 illustrates layer 2 services being supported over an Internet service provider's MPLS network. Customers using Frame Relay, ATM, or other legacy layer 2 connections to E-series routers are unaware that MPLS tunneling is used.

Figure 113: Layer 2 Services over a Provider's MPLS Network



Layer 2 Services over MPLS Platform Considerations

To configure layer 2 services over MPLS, you must first configure the underlying layer 2 service (ATM, bridged Ethernet, Fast Ethernet, Gigabit Ethernet, 10-Gigabit Ethernet, Frame Relay, or HDLC) and MPLS.

Module Requirements

For information about the modules that support the underlying layer 2 service and MPLS on ERX-14xx models, ERX-7xx models, and the ERX-310 router:

- See *ERX Module Guide, Table 1, Module Combinations* for detailed module specifications.
- See *ERX Module Guide, Appendix A, Module Protocol Support* for information about the modules that support the underlying layer 2 service and MPLS.

For information about the modules that support the underlying layer 2 service and MPLS on the E120 router or the E320 router:

- See *E120 and E320 Module Guide, Table 1, Modules and IOAs* for detailed module specifications.
- See *E120 and E320 Module Guide, Appendix A, IOA Protocol Support* for information about the modules that support the underlying layer 2 service and MPLS.



NOTE: In the current release, the E120 and E320 routers supports all layer 2 services over MPLS shown in the bulleted list in *Layer 2 Services over MPLS Overview* on page 485 *except* ATM with AAL5 encapsulation, ATM with VCC cell relay encapsulation, Frame Relay, and HDLC.

Interface Specifiers

Many of the configuration task examples in this chapter use the *slot/port[.subinterface]* format to specify the physical interface for the underlying layer 2 service. However, the interface specifier format that you use depends on the router that you are using.

For ERX-7xx models, ERX-14xx models, and ERX-310 routers, use the *slot/port[.subinterface]* format. For example, the following command specifies ATM 1483 subinterface 10 on slot 0, port 1 of an ERX-7xx model, ERX-14xx model, or ERX-310 router.

```
host1(config)#interface atm 0/1.10
```

For the E120 router and the E320 router, use the *slot/adapter/port[.subinterface]* format, which includes an identifier for the bay in which the I/O adapter (IOA) resides. In the software, adapter 0 identifies the right IOA bay (E120 router) and the upper IOA bay (E320 router); adapter 1 identifies the left IOA bay (E120 router) and the lower IOA bay (E320 router). For example, the following command specifies ATM 1483 subinterface 20 on slot 5, adapter 0, port 0 of an E320 router.

```
host1(config)#interface atm 5/0/0.20
```

Related Topics

- For more information about supported interface types and specifiers on E-series routers, see *Interface Types and Specifiers* in *JUNOS Command Reference Guide, About This Guide*.

Layer 2 Services over MPLS References

For information about layer 2 services, consult the following resources:

- Encapsulation Methods for Transport of ATM Over MPLS Networks—draft-ietf-pwe3-atm-encap-07.txt (April 2005 expiration)
- Encapsulation Methods for Transport of Ethernet Frames Over IP/MPLS Networks—draft-ietf-pwe3-ethernet-encap-05.txt (June 2004 expiration)
- Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks—draft-martini-l2circuit-encap-mpls-08.txt (March 2005 expiration)
- Encapsulation Methods for Transport of PPP/HDLC Over IP and MPLS Networks—draft-ietf-pwe3-hdlc-ppp-encap-mpls-03.txt (October 2004 expiration)
- Framework for Pseudo Wire Emulation Edge-to-Edge (PWE3)—draft-ietf-pwe3-arch-06.txt (April 2004 expiration)
- IEEE 802.3ad (Link Aggregation)
- Pseudowire Setup and Maintenance Using LDP—draft-ietf-pwe3-control-protocol-08.txt (January 2005 expiration)
- Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)—draft-ietf-pwe3-requirements-08.txt (June 2004 expiration)
- Transport of Layer 2 Frames Over MPLS—draft-martini-l2circuit-trans-mpls-11.txt (October 2003 expiration)



NOTE: IETF drafts are valid for only 6 months from the date of issuance. They must be considered as works in progress. Please refer to the IETF Web site at <http://www.ietf.org> for the latest drafts.

For information about configuring supported layer 2 interfaces, consult the following resources:

- *JUNOS Physical Layer Configuration Guide, Chapter 5, Configuring Ethernet Interfaces*
- *JUNOS Link Layer Configuration Guide, Chapter 1, Configuring ATM*
- *JUNOS Link Layer Configuration Guide, Chapter 2, Configuring Frame Relay*
- *JUNOS Link Layer Configuration Guide, Chapter 6, Configuring Packet over SONET*
- *JUNOS Link Layer Configuration Guide, Chapter 9, Configuring Bridged Ethernet*

For information about configuring supported serial interfaces, which are also referred to as HDLC channels, see the following resources:

- *JUNOS Physical Layer Configuration Guide, Chapter 1, Configuring Channelized T3 Interfaces*
- *JUNOS Physical Layer Configuration Guide, Chapter 2, Configuring T3 and E3 Interfaces*
- *JUNOS Physical Layer Configuration Guide, Chapter 4, Configuring Channelized OCx/STMx Interfaces*

For information about configuring MPLS, see *Chapter 2, Configuring MPLS*.

Layer 2 Services over MPLS Implementation

When layer 2 services are configured over MPLS, layer 2 traffic is encapsulated in MPLS frames and sent over MPLS tunnels. A virtual circuit (VC) label that indicates a specific layer 2 connection, such as a Frame Relay data-link connection identifier (DLCI), is pushed into the label stack between the tunnel label and the layer 2 data.

A service-specific control word may be placed between the layer 2 data and the VC label. The control word is used for frame sequencing and carrying service-specific information, such as Frame Relay forward explicit congestion notification (FECN) and backward explicit congestion notification (BECN) information. At the tunnel end, the VC label is used to find the layer 2 interface over which the traffic is sent. The control word, if present, is used to convert the encapsulated layer 2 traffic into its native format.

Because MPLS labels are unidirectional, two VC labels are required for each layer 2 connection. The VC labels are distributed by the Label Distribution Protocol (LDP) in downstream-unsolicited (DU) mode between the two routers. The layer 2 connection status signaling may be emulated by advertising and withdrawing the VC labels. For example, if the Frame Relay subinterface between customer premises equipment (CPE) and a provider edge (PE) router goes down, the corresponding VC label is withdrawn by the PE router. When the remote PE router at the other end receives the label withdrawal, it translates the label withdrawal into LMI notifications to its CPE. When the Frame Relay subinterface comes back, a VC label is advertised, and the remote PE router again translates it into LMI notifications.

Local Cross-Connects Between Layer 2 Interfaces Using MPLS

You can configure layer 2 services over MPLS to transmit data between two layer 2 interfaces that reside on the same E-series router. In this configuration, which is referred to as a local cross-connect, traffic that arrives at the router's ingress interface is switched out the egress interface, instead of going through an MPLS core network.

A local cross-connect enables the router to function as a layer 2 switch. It operates with any supported layer 2 service. To configure local cross-connects, you must use the **mpls-relay** command.

Related Topics

- For a list of supported layer 2 services, see *Layer 2 Services over MPLS Overview* on page 485.
- For a configuration example that shows how to create local cross-connects between Ethernet/VLAN interfaces, see *Configuring Local Cross-Connects Between Ethernet/VLAN Interfaces* on page 503.

MPLS Shim Interfaces for Layer 2 Services over MPLS

An MPLS shim interface is stacked on a layer 2 interface to do either of the following:

- Create a layer 2 circuit by cross-connecting the layer 2 interface to an LSP corresponding to the VC label.
- Create a local cross-connect by cross-connecting the layer 2 interface to another layer 2 interface.

You can create or remove MPLS shim interfaces with the **mpls-relay** or **route interface** commands. Shim interfaces are also created in or removed implicitly from a load-balancing group by the **member interface** command when you configure the group. Each MPLS shim interface exists in a particular virtual router.

Each MPLS shim interface points to a single MPLS next hop. When layer 2 frames arrive on the layer 2 interface below the MPLS shim interface, they are encapsulated in an MPLS packet and forwarded to that MPLS next hop. The details of the encapsulation are determined by the attributes of the shim interface.

The MPLS next hop to which the shim interface points is set by an MPLS signaling protocol, which adds a special entry with implicit null in label to the interface label-space MPLS forwarding table of the shim interface. For traffic arriving from the core, the MPLS signaling protocol adds a normal entry with a real in label to the platform label-space MPLS forwarding table whose next hop points to the MPLS shim interface.

You can configure the following attributes for each MPLS shim interface:

- The administrative state, enabled or disabled, configured with the **mpls-relay disable** command.
- The IP address of the remote PE router for the layer 2 circuit, configured with the **mpls-relay** command. The shim interface is cross-connected to an LSP corresponding to the VC label received from the specified remote PE router, or to another shim interface in the local cross-connects case. For local cross-connects the IP address is local to the PE router.
- The name of an RSVP-TE base tunnel to be used for the layer 2 circuit, configured with the **route interface** command.
- The group ID of the shim interface, configured using the **group-id** option of the **mpls-relay** and **route interface** commands. Even though you can configure the group ID, the JUNOS software does not currently use it.

- Whether the control word is used, configured with the **control-word** and **no-control-word** options of the **mpls-relay** and **route interface** commands. The layer 2 interface determines the default preference if this option is not configured. Some layer 2 interfaces require a control word; others do not support it.
- Whether sending nonzero sequence numbers in the control word is preferred, configured with the **sequencing** and **no-sequencing** options of the **mpls-relay** and **route interface** commands. The layer 2 interface determines the default preference if this option is not configured. Even when preferred, the sequence numbers might not be sent if the control word is not used due to configuration. You can only configure whether the numbers are sent. MPLS always accepts zero sequence numbers and checks the order of nonzero sequence numbers in received MPLS packets that are forwarded to an MPLS shim interface. Out-of-order packets are discarded.
- The number of the load-balancing group of which the shim interface is a member. This attribute is set to the current load-balancing group when the shim interface is implicitly created with the **member interface** command.

You can enable statistics collection for the MPLS shim interfaces.

Related Topics

- For information about the creation or implicit removal of shim interfaces from a load-balancing group by the **member interface** command when you configure the group, see *Configuring CE-Side Load Balancing for Martini Layer 2 Transport* on page 510.
- For information about collecting statistics for the MPLS shim interfaces, see *Chapter 6, Monitoring Layer 2 Services over MPLS*.

Multiple Layer 2 Services over MPLS

When you configure an MPLS shim interface over an ATM, Frame Relay, or HDLC layer 2 interface, no other interface (for example, PPP or IP) can be stacked above the layer 2 interface.

By contrast, when you configure an MPLS shim interface over any Ethernet or Ethernet/VLAN interface, both the MPLS shim interface and other interfaces (such as IP, PPP, or PPPoE) can be stacked above the layer 2 interface.

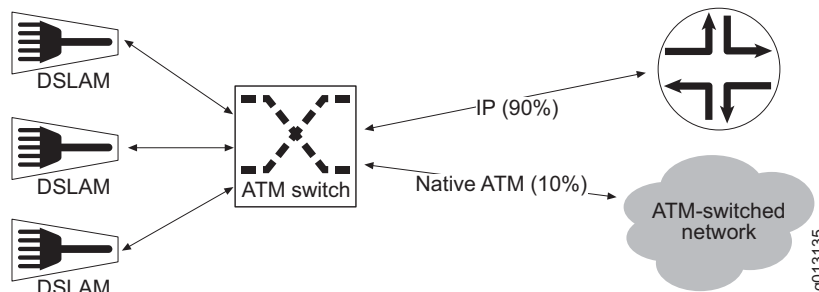
When you configure both an MPLS shim interface and another interface over a layer 2 interface, traffic for a protocol explicitly configured in the interface stack is forwarded to that protocol layer for further processing. Traffic for any nonconfigured protocols is forwarded to the MPLS shim interface on the other side of the connection.

When the MPLS shim interface is the only layer stacked above the layer 2 interface, as is the case with ATM, Frame Relay, and HDLC, then all traffic is forwarded to the MPLS shim interface and across the MPLS tunnel.

ATM Layer 2 Services over MPLS


ATM layer 2 services over MPLS provide ATM switch-like functionality for E-series routers. This feature is useful for customers who run IP in the majority of their network but still have to carry a small amount of non-IP traffic, as in the example shown in Figure 114.

Figure 114: Common ISP Network



In this scenario, it is not economical to have special ATM switches in front of E-series routers just to accommodate the small percentage of non-IP traffic. The ATM layer 2 services over MPLS feature lets you replace the ATM traffic selector switch with an E-series router, as shown in Figure 115. The two routers pass traffic between two interfaces through an MPLS tunnel using Martini encapsulation, regardless of the contents of the packets.

Figure 115: E-series Router Replacing Remote ATM Switch



ATM layer 2 services over MPLS supports two encapsulation methods on E-series routers:

- AAL5 relay encapsulation
- VCC cell relay encapsulation

The following sections describe each of these encapsulation methods.


AAL5 Encapsulation

JUNOS software supports the AAL5 relay method of encapsulation that is specified in the Martini draft. This method is also referred to as AAL5 service data unit (SDU) encapsulation.

ATM Martini encapsulation emulates ATM switch behavior by creating a pseudowire between pairs of ATM virtual circuits. When the router receives AAL5 packets on one of those circuits, it reassembles them, encapsulates them using Martini encapsulation, and forwards them to an MPLS tunnel. At the end of the tunnel, the packet is de-encapsulated, segmented back, and sent to a selected ATM VC.

In Figure 116, an MPLS tunnel connects two E-series routers, and ATM cross-connects provide a pseudowire between the ATM VCs on the two routers. All AAL5 packets on the pseudowire are encapsulated. The egress VC does not need the same ATM address as the ingress circuit.

Figure 116: AAL5 Pseudowire and MPLS Tunnel



To use AAL5 SDU encapsulation, you must use the **aal5all** encapsulation keyword when you configure ATM subinterfaces.

OAM Cells

The E-series router performs a similar operation for Operation, Administration, and Maintenance (OAM) cells, except that they do not need reassembly.

The router passes the following OAM cells transparently:

- F5 alarm indication signal (AIS) segment and end-to-end
- F5 remote defect indication (RDI) segment and end-to-end
- F5 loopback segment and end-to-end
- Resource management
- F5 continuity check segment and end-to-end

In addition, F4 OAM cell forwarding is supported.

JUNOS software does not allow for setting a segment endpoint on an ATM cross-connect interface. Segment OAM cells are forwarded to the egress interface in the same manner as end-to-end cells.

QoS Classification

Packets are subject to normal quality of service (QoS) classification according to service categories assigned to ATM virtual circuits that make up the connection.

Limitations

The JUNOS implementation of the Martini draft has the following limitations:

- Only AAL5 packets and OAM cells are forwarded.
- There is no equivalent of VP switching.
- Point-to-multipoint connections are not supported.
- Automatic connection setup using user-to-network interface (UNI) signaling and private network-to-network interface (PNNI) is not supported.
- The ATM MIB cross-connected table is not supported.
- Connections between ATM circuits and non-ATM interfaces are not supported.

Control Word Support

Martini draft encapsulation includes a control word with the following fields:

- T bit for transport type
- E bit for explicit forward congestion indication (EFCI)
- C bit for cell loss priority (CLP) = 1 indication
- U bit for command/response indication based on AAL5 common part convergence sublayer user-to-user indication (CPCS-UU)
- Optional sequence number

The current JUNOS implementation supports the T bit and optional sequence number fields.

VCC Cell Relay Encapsulation

E-series routers support virtual channel connection (VCC) cell relay encapsulation for ATM layer 2 services over MPLS. VCC cell relay encapsulation enables a router to emulate ATM switch behavior by forwarding individual ATM cells over an MPLS pseudowire (also referred to as an MPLS tunnel) created between two ATM VCCs, or as part of a local ATM cross connect between two ATM 1483 subinterfaces on the same router. The JUNOS implementation conforms to the required N-to-1 cell mode encapsulation method described in the Martini draft, with the provision that only a single ATM virtual circuit (VC) can be mapped to an MPLS pseudowire.

VCC cell relay encapsulation is useful for voice-over-ATM applications that use ATM Adaptation Layer 2 (AAL2)–encapsulated voice transmission.

AALO Raw Cell Mode

VCC cell relay encapsulation supports ATM Adaptation Layer 0 (AALO) encapsulation, also referred to as raw cell mode or null encapsulation. AALO is often used to carry signaling ATM cells, which the router treats as raw cells.

When you create an ATM PVC as part of a VCC cell relay configuration, you must use the **aal0** encapsulation.

Cell Concatenation Parameters

VCC cell relay encapsulation supports the concatenation (aggregation) of multiple ATM cells in a single encapsulated packet that is transmitted on the MPLS pseudowire.

You can use the **atm cell-packing** and **atm mcpt-timers** commands to configure the following parameters that control how the router performs cell concatenation:

- Maximum number of ATM cells that the router can concatenate in a single packet.
- Values (in microseconds) for each of the three ATM Martini cell packing timers maintained on the router. These timers define the time threshold that the router uses to concatenate ATM cells and transmit the cells in an MPLS packet on the pseudowire.
- Numeric identifier (1, 2, or 3) that indicates which of the three ATM Martini cell packing timers you want to use to detect timeout of the cell collection threshold.

Based on this configuration, the router attempts to concatenate the specified maximum number of ATM cells into an MPLS packet within the time interval allowed by the ATM Martini cell packing timer you selected. When the timer detects that the allotted time interval has expired, the router forwards the MPLS packet even if it contains fewer than the specified maximum number of aggregated cells per packet.

Cell Concatenation and Latency

Cell concatenation increases network latency, which is undesirable for voice-over-ATM applications. To minimize these effects, use care in choosing values for the ATM Martini cell packing timers.

We recommend that for voice-over-ATM configurations, you select timeout values between 6 microseconds and 3 x 6 microseconds. Values within this range are generally low enough to maintain a reasonable cell delay and high enough to take advantage of the cell concatenation mechanism.

If traffic shaping is enabled on the egress router, the JUNOS implementation of VCC cell relay encapsulation preserves the spacing between cells.

Control Word Support

VCC cell relay encapsulation requires use of a control word. The control word contains the T, E, C, and U bits, as well as an optional sequence number field.

The JUNOS implementation of VCC cell relay encapsulation supports the T bit, which is always set to indicate raw ATM cells, and the optional sequence number. The E, C, and U bits have no meaning for VCC cell relay configurations because the router forwards a complete ATM cell with as much header information as possible.

Unsupported Features

VCC cell relay encapsulation on JUNOS routers does not support the following features:

- Mapping multiple ATM VCs to a single MPLS pseudowire
- Optional Martini one-to-one cell encapsulation method with cell headers removed

Related Topics

- For information about AAL5 SDU encapsulation, see Encapsulation Methods for Transport of ATM Over MPLS Networks—draft-ietf-pwe3-atm-encap-07.txt (April 2005 expiration).
- For configuration examples of AAL5 SDU encapsulation using the **aal5all** encapsulation keyword when you configure ATM subinterfaces, see *Configuring Local ATM Cross-Connects with AAL5 Encapsulation* on page 505.
- For information about, VCC cell relay encapsulation, see Encapsulation Methods for Transport of ATM Over MPLS Networks—draft-ietf-pwe3-atm-encap-07.txt (April 2005 expiration).
- For VCC cell relay encapsulation configuration instructions, see *Configuring an MPLS Pseudowire with VCC Cell Relay Encapsulation* on page 506.
- *Control Word Support* on page 494

HDLC Layer 2 Services over MPLS

E-series routers support the creation of HDLC layer 2 circuits across an MPLS network. An HDLC layer 2 circuit can carry any standard HDLC traffic (including PPP) or Cisco HDLC traffic between two CE devices across an MPLS network. In an HDLC layer 2 circuit configuration, an E-series router functions as one of the PE routers.

You can configure an HDLC layer 2 circuit between two serial interfaces, between two packet over SONET (POS) interfaces, or between a serial interface and a POS interface. The interfaces at either end of the circuit can operate at the same speed or at different speeds. For example, you can configure an HDLC layer 2 circuit between a serial interface on a T1 circuit and a POS interface on an OC3 circuit.

Interface Stacking

When you configure an MPLS shim interface above an HDLC layer 2 interface, which is in turn stacked above a serial or POS interface, no other interfaces (for example, PPP) can be stacked above the HDLC interface. In other words, the HDLC interface can have only one upper interface.

In practice, this means that you cannot issue the **mpls-relay** command (to create the HDLC layer 2 circuit) and then issue an **encapsulation** command (such as **encapsulation ppp**) for the same HDLC interface. If you attempt to do so, the router prevents the configuration and displays an error message.

This behavior contrasts with that of bridged Ethernet and Ethernet interfaces (with and without VLANs), which allow configuration of both an MPLS shim interface and another interface (such as IP, PPP, or PPPoE) above the layer 2 interface.

Encapsulation

The JUNOS implementation of HDLC layer 2 circuits supports encapsulation of either HDLC frames or PPP frames within MPLS frames. By default, the router uses VC-type HDLC signaling and HDLC encapsulation to encapsulate HDLC frames in MPLS.

However, if you want the router to encapsulate PPP frames directly in MPLS without the HDLC header, you can include the optional **relay-format ppp** keywords in the **mpls-relay** or **route interface** command to cause the router to use VC-type PPP signaling and PPP encapsulation instead of the default VC-type HDLC signaling and HDLC encapsulation. This option, which is available only for serial and POS interfaces, is useful if the traffic carried on the serial or POS interface contains actual PPP packets and not, for example, Cisco HDLC packets.

Control Word Support

The router always sends a control word for HDLC layer 2 circuits, regardless of whether or not sequencing is enabled.

Local Cross-Connects

You can configure an HDLC layer 2 circuit in a local cross-connect configuration between serial or POS interfaces within the same router. In this configuration, the pairs of HDLC interfaces are directly cross-connected to each other. The cross-connected interfaces can be different types and operate at different speeds; for example, you can cross-connect a serial interface on a T1 circuit and a POS interface on an OC3 circuit.

Related Topics

- *Configuring HDLC Layer 2 Services on page 509.*

