



# ISC L2VPN and VPLS Concepts

---

## Overview of ISC L2VPN and VPLS

Layer 2 service provisioning for the IP Solution Center (ISC) 3.2 consists of the Layer 2 Virtual Private Network (L2VPN) service and the Virtual Private LAN Service (VPLS). L2VPN services are point-to-point and include Ethernet, Frame Relay, and ATM services. VPLS services are multipoint and include Ethernet services over an MPLS core or an Ethernet core.

This chapter contains the following sections.

- [AToM L2VPN, page 1-1](#)
- [ATM over MPLS \(ATMoMPLS\), page 1-5](#)
- [Frame Relay over MPLS \(FRoMPLS\), page 1-6](#)
- [VPLS for an MPLS-Based Provider Core, page 1-7](#)
- [VPLS for an Ethernet-Based Provider Core, page 1-9](#)

## AToM L2VPN

Cisco's Any Transport over MPLS (AToM) enables service providers to deliver profitable, comprehensive services to their customers. The L2VPN service provisioning available in ISC 3.2 is in the following four areas:

- [Point-to-Point Ethernet Wire Service \(EWS\), page 1-2](#)
- [Point-to-Point Ethernet Relay Service \(ERS\), page 1-2](#)
- [ATM over MPLS \(ATMoMPLS\), page 1-5](#)
- [Frame Relay over MPLS \(FRoMPLS\), page 1-6](#)

The EWS and ERS services are delivered with the Cisco Metro Ethernet offering.

The same network architecture can simultaneously provide both ERS and EWS connections to diverse customers. Additionally, this Metro Ethernet infrastructure may be used for access to higher-level services, such as IP-based virtual private networking, public internet communications, Voice over IP, or a combination of all applications.

## Point-to-Point Ethernet Wire Service (EWS)

An Ethernet Virtual Circuit (EVC) connects two physical User-to-Network Interfaces (UNI) such that the connection appears like a virtual private line to the customer. VLAN transparency and control protocol tunneling are supplied by the implementation of 802.1Q-in-Q tag-stacking technology. Packets received on one UNI are transported directly to the other corresponding UNI.

## Point-to-Point Ethernet Relay Service (ERS)

An Ethernet Virtual Circuit (EVC) is used to logically connect end-points, but multiple EVCs may exist per single UNI. Each EVC is distinguished by 802.1q VLAN tag identification. The ERS network acts as if the Ethernet frames have crossed a switched network, and certain control traffic is not carried between ends of the EVC. ERS is analogous to Frame Relay where the CE-VLAN tag plays the role of a Data-Link Connection Identifier (DLCI).

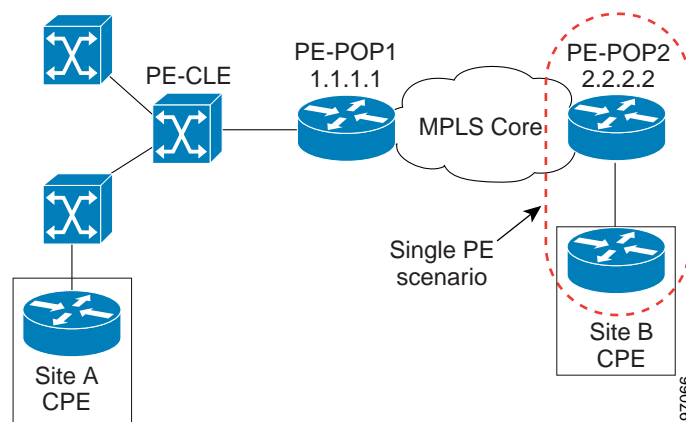
## Topology for L2VPN Ethernet Over MPLS (ERS and EWS)

Ethernet Over MPLS (EoMPLS) is a tunneling mechanism that allows the service provider to tunnel customer Layer 2 traffic through a Layer 3 MPLS network. It is important to remember that EoMPLS is a point-to-point solution only.

The following figures provide a reference for how EoMPLS is utilized. Ethernet Services can be distributed to the end-customer in two ways.

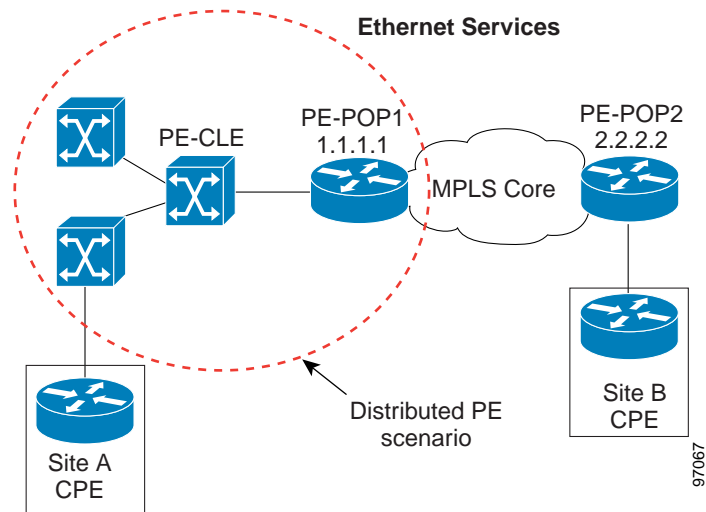
- Single PE scenario: The customer is directly connected to an Ethernet port on the PE-POP in [Figure 1-1](#).

**Figure 1-1** Single PE scenario



- Distributed PE scenario: The end-customer is connected through an Access Domain to the PE in [Figure 1-2](#). That is, there is a Layer 2 switching environment in the middle of CE and PE.

Figure 1-2 Distributed PE Scenario



In both cases, a VLAN is assigned in one of the following ways:

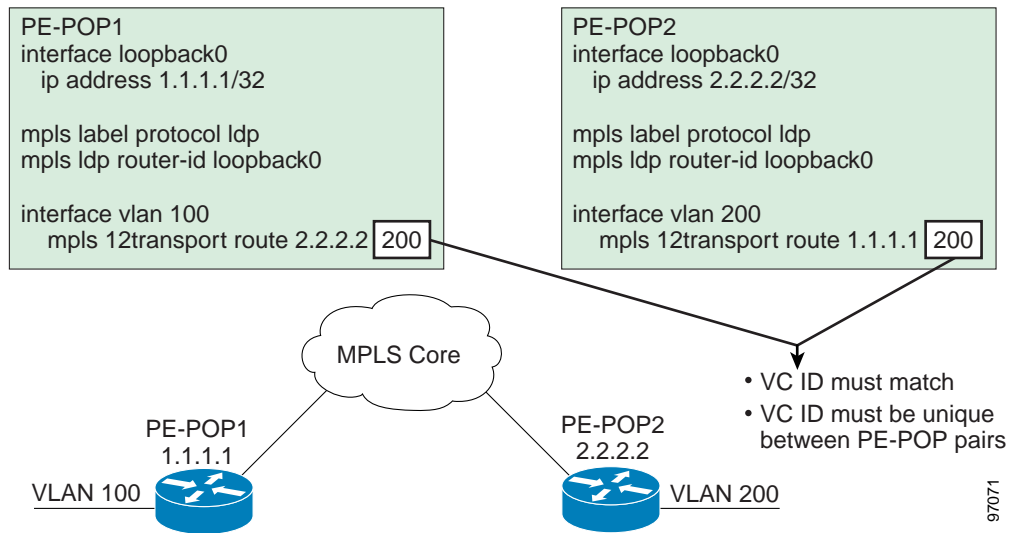
- Automatically assigned by ISC from the VLAN pool that is predefined by the user.
- Manually assigned by the user through the GUI or the North Bound Interface (NBI).

In EoMPLS, ISC creates a point-to-point tunnel and then targets the EoMPLS tunnel to the peer PE-POP router through which the remote site can be reached. The remote PE-POP is identified by its loopback address (in Figure 1-3, PE-POP1 and PE-POP2 have 1.1.1.1 and 2.2.2.2 as loopback addresses). In Figure 1-3, Site A has been allocated a VLAN-100 and Site-B a VLAN-200. You can have different VLAN IDs at either end of the circuit because the VLANs have local significance only (that is, within the Ethernet access domain which is delimited by the PE-PoP).

For the PE-POP that is serving Site A, a VLAN interface (Layer 3 interface) is created to terminate all L2 traffic for the customer, and an EoMPLS tunnel is configured on this interface.<sup>1</sup>


1. This configuration is based on the Cisco 7600 Optical Services Router. Other routers, such as the Cisco 7200, have different configurations.

Figure 1-3 Ethernet over MPLS Configuration



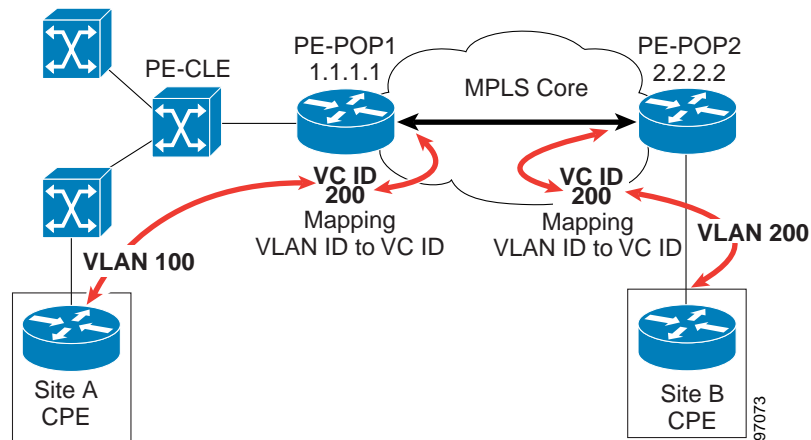
The VC ID that defines the EoMPLS tunnel is 200 as shown in Figure 1-3. Note that the VC ID has to be the same on both ends of the EoMPLS tunnel. On each PE-POP, there is mapping done between the VLANs to the EoMPLS tunnel (Figure 1-4). For the overall connection, this mapping is: VLAN ID <-> VC ID <-> VLAN ID.

Figure 1-4 EoMPLS Tunnel



This VLAN-VC ID mapping lets the service provider reuse VLAN IDs in Access Domains (see Figure 1-5). The VLAN IDs allocated and used at each access domain do not have to be the same.

Figure 1-5 VLAN-VC ID Mapping



## ATM over MPLS (ATMoMPLS)

With Cisco ATM over MPLS (ATMoMPLS), Cisco supports ATM Adaptation Layer 5 (AAL5) transport and Cell Relay over MPLS.

### AAL5

AAL5 allows you to transport AAL5 PDUs from various customers over an MPLS backbone. ATM AAL5 extends the usability of the MPLS backbone by enabling it to offer Layer 2 services in addition to already existing Layer 3 services. You can enable the MPLS backbone network to accept AAL5 PDUs by configuring the provider edge (PE) routers at both ends of the MPLS backbone.

To transport AAL5 PDUs over MPLS, a virtual circuit is set up from the ingress PE router to the egress PE router. This virtual circuit transports the AAL5 PDUs from one PE router to the other. Each AAL5 PDU is transported as a single packet.

### Cell Relay over MPLS


Cell Relay over MPLS allows you to transport ATM cells from various customers over an MPLS backbone. ATM Cell Relay extends the usability of the MPLS backbone by enabling it to offer Layer 2 services in addition to already existing Layer 3 services. You can enable the MPLS backbone network to accept ATM cells by configuring the provider edge (PE) routers at both ends of the MPLS backbone.

To transport ATM cells over MPLS, a virtual circuit is set up from the ingress PE router to the egress PE router. This virtual circuit transports the ATM cells from one PE router to the other. Each MPLS packet can contain one or more ATM cells.

## Topology for ATMoMPLS

Only the single PE scenario is supported as shown in [Figure 1-6](#).

Figure 1-6 Configuring AAL5 and Cell Relay over MPLS



## Frame Relay over MPLS (FRoMPLS)

With Cisco AToM for Frame Relay, customer Frame Relay traffic can be encapsulated in MPLS packets and forwarded to destinations required by the customer. Cisco AToM allows service providers to quickly add new sites with less effort than typical Frame Relay provisioning.

Frame Relay over MPLS enables a service provider to transport Frame Relay frames across an MPLS backbone. This extends the reachability of Frame Relay and allows service providers to aggregate frame transport across a common packet backbone. The service provider can integrate an existing Frame Relay environment with the packet backbone to improve operational efficiency and to implement the high-speed packet interfaces to scale the Frame Relay implementations.

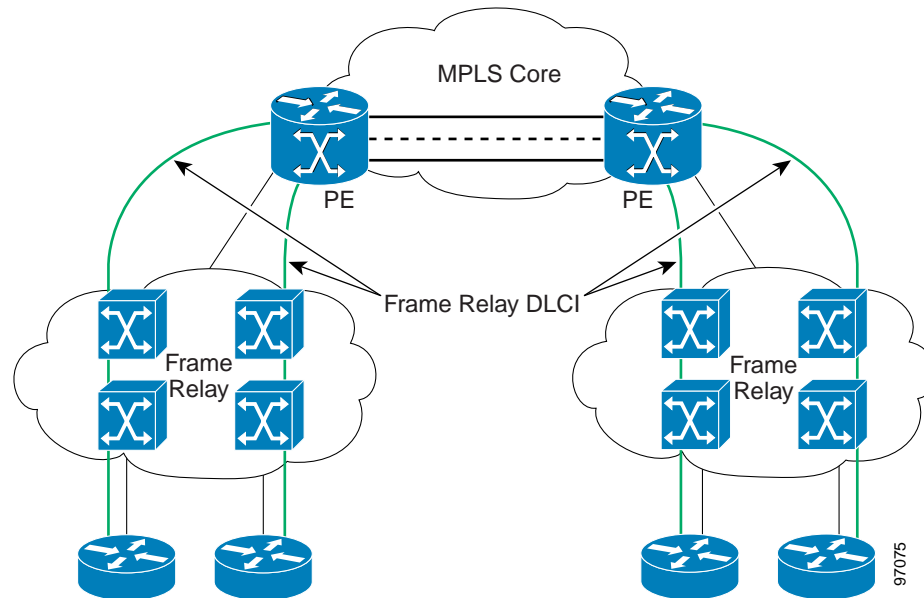
Transporting Frame Relay frames across MPLS networks provides a number of benefits, including:

- Frame Relay extended service.
- Aggregation to a higher-speed backbone, such as OC-192, to scale Frame Relay implementations.
- Improved operational efficiency—the MPLS backbone becomes the single network that integrates the various existing networks and services.

## Topology for FROverMPLS

Only the single PE scenario is supported as shown in [Figure 1-7](#).

Figure 1-7 FRoverMPLS



## VPLS for an MPLS-Based Provider Core

VPLS is a multipoint Layer 2 VPN that connects two or more customer devices using EoMPLS bridging techniques. VPLS EoMPLS is an MPLS-based provider core, that is, the PE routers have to cooperate to forward customer Ethernet traffic for a given VPLS instance in the core.

A VPLS essentially emulates an Ethernet switch from a user's perspective. All connections are peers within the VPLS and have direct communications. The architecture is actually that of a distributed switch.

Multiple attachment circuits have to be joined together by the provider core. The provider core has to simulate a virtual bridge that connects these multiple attachment circuits together. To achieve this, all PE routers participating in a VPLS instance form emulated VCs among them.

A Virtual Forwarding Instance (VFI) is created on the PE router for each VPLS instance. PE routers make packet-forwarding decisions by looking up the VFI of a particular VPLS instance. The VFI acts like a virtual bridge for a given VPLS instance. More than one attachment circuit belonging to a given VPLS can be connected to this VFI. The PE router establishes emulated VCs to all the other PE routers in that VPLS instance and attaches these emulated VCs to the VFI. Packet forwarding decisions are based on the data structures maintained in the VFI. All the PE routers in the VPLS domain use the same VC-ID for establishing the emulated VCs. This VC-ID is also called the VPN-ID in the context of the VPLS VPN.

## Multipoint EWS for an MPLS-Based Provider Core

With multipoint EWS, the PE router forwards all Ethernet packets received from an attachment circuit, including tagged, untagged, and Bridge Protocol Data Unit (BPDU) to either:

- Another attachment circuit or an emulated VC if the destination MAC address is found in the L2 forwarding table (VFI).

- All other attachment circuits and emulated VCs belonging to the same VPLS instance if the destination MAC address is a multicast/broadcast address or not found in the L2 forwarding table.

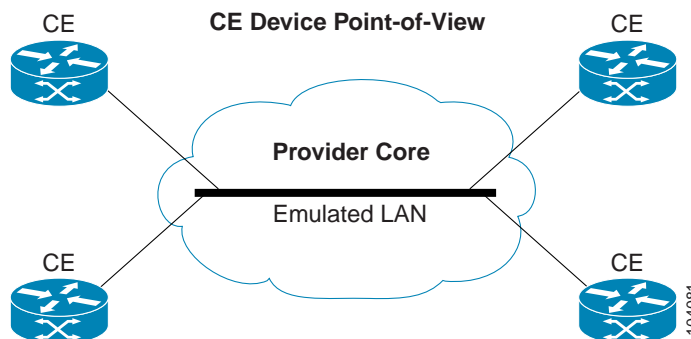
## Multipoint ERS for an MPLS-Based Provider Core

With multipoint ERS, the PE router forwards all Ethernet packets with a particular VLAN tag received from an attachment circuit, excluding BPDU, to another attachment circuit or an emulated VC if the destination MAC address is found in the L2 forwarding table (VFI). If the destination MAC address is not found or if it is a broadcast/multicast packet, then it is sent on all other attachment circuits and emulated VCs belonging to the VPLS instance. The demultiplexing VLAN tag used to identify a VPLS domain is removed prior to forwarding the packet to the outgoing Ethernet interfaces or emulated VCs because it only has local significance.

## Topology for MPLS-Based VPLS

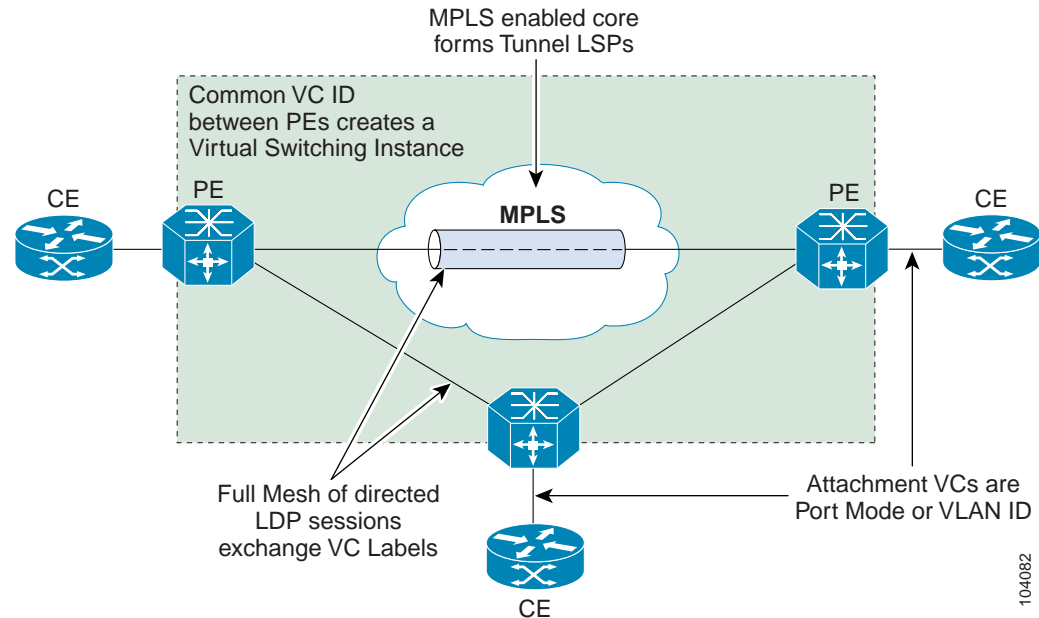
From a customer point of view there is no topology for VPLS. All the CE devices are connected to a logical bridge emulated by the provider core. Therefore, the CE devices see a single emulated LAN as shown in [Figure 1-8](#).

**Figure 1-8** MPLS-Based VPLS Topology



The PE routers must create a full-mesh of emulated virtual circuits (VCs) to simulate the emulated LAN seen by the CE devices. Forming a full-mesh of emulated VCs simplifies the task of emulating a LAN in the provider core. One property of a LAN is to maintain a single broadcast domain. That is, if a broadcast, multicast, or unknown unicast packet is received on one of the attachment circuits, it has to be sent to all other CE devices participating in that VPLS instance. The PE device handles this case by sending such a packet on all other attachment circuits and all the emulated circuits originating from that PE. With a full-mesh of emulated VCs, such a packet will reach all other PE devices in that VPLS instance. See [Figure 1-9](#).

Figure 1-9 Full-mesh of Emulated VCs



## VPLS for an Ethernet-Based Provider Core

With an Ethernet-based provider core, customer traffic forwarding is trivial in the core. VPLS for an Ethernet-based provider core is a multipoint Layer 2 VPN that connects two or more customer devices using 802.1Q-in-Q tag-stacking technology. A VPLS essentially emulates an Ethernet switch from a user's perspective. All connections are peers within the VPLS and have direct communications. The architecture is actually that of a distributed switch.

## Multipoint EWS for an Ethernet-Based Provider Core

Multipoint EWS is a service that emulates a point-to-point Ethernet segment. The EWS service encapsulates all frames that are received on a particular User to Network Interface (UNI) and transports these frames to a single egress UNI without reference to the contents contained within the frame. This service operation means that EWS can be used for untagged or VLAN tagged frames and that the service is transparent to all frames offered. Because the EWS service is unaware that VLAN tags may be present within the customer frames, the service employs a concept of “All to One” bundling.

## Multipoint ERS for an Ethernet-Based Provider Core


Multipoint ERS models the connectivity offered by existing Frame Relay networks by using VLAN indices to identify virtual circuits between sites. ERS does, however, offer a far greater degree of QoS functionality depending upon the service provider's implementation and the customer's acceptance of VLAN indices that are administratively controlled by the service provider. Additionally, ERS service multiplexing capability offers lower cost of ownership for the enterprise as a single interface can support many virtual interfaces.

## Topology for Ethernet-Based VPLS

Ethernet-based VPLS differs from the point-to-point L2VPN definitions of EWS and ERS by providing a multipoint connectivity model. The VPLS service does not map an interface or VLAN to a specific point-to-point pseudo-wire, but instead it models the operation of a virtual Ethernet switch. VPLS uses the customer's MAC address to forward frames to the correct egress UNI within the service provider's network for the EWS.

The EWS service emulates the service attributes of an Ethernet switch and learns source MAC to interface associations, flooding unknown broadcast and multicast frames. [Figure 1-10](#) illustrates an EWS VPLS topology.

*Figure 1-10 VPLS EWS Topology*



The Ethernet Relay Service (ERS) offers the any-to-any connectivity characteristics of EWS and the service multiplexing. This combination enables a single UNI to support a customer's intranet connection and one or more additional EVCs for connection to outside networks, ISPs, or content providers. [Figure 1-11](#) illustrates an ERS VPLS multipoint topology.

Figure 1-11 VPLS ERS Multipoint Topology

