



CHAPTER 1

Layer 2 VPN Concepts

This chapter provides an overview of Prime Fulfillment Layer 2 VPN concepts. It contains the following sections.

- [Layer 2 Terminology Conventions, page 1-1](#)
- [L2VPN Service Provisioning, page 1-5](#)
- [FlexUNI/EVC Ethernet Service Provisioning, page 1-10](#)
- [VPLS Service Provisioning, page 1-15](#)

Layer 2 Terminology Conventions

Layer 2 service provisioning for the Prime Fulfillment consists of the Layer 2 Virtual Private Network (L2VPN) Service, EVC Service and the Virtual Private LAN Service (VPLS). The purpose of this section is to clarify the terminologies used in Prime Fulfillment, as well as in the industry at large, for these services.

There are three sets of terminologies in use:

- The current Metro Ethernet Forum (MEF) terminology
- The former MEF terminology
- Prime Fulfillment terminology (which is close to the former MEF terminology)

MEF Terminology Conventions

In general, for L2VPN services, Prime Fulfillment supports four general Ethernet service type constructs:

- Ethernet Line (E-Line). Provides a point-to-point Ethernet Virtual Circuit (EVC).
- Ethernet LAN (ELAN). Provides a multipoint-to-multipoint EVC.
- EVC with PW-core. This offers EPL and EVPL.
- EVC with VPLS-core. This offers ELAN and EPLAN

Two Ethernet services are available for each type. These are distinguished by the means of service identification used at the user-to-network interface (UNIs), as follows:

- Port based. All-to-one bundling. These are referred to as “private.”

- VLAN-based. These services are multiplexed. The EVC is identified by a VLAN ID. These are referred to as “virtual private.”

Table 1-1 summarizes these relationships.

Table 1-1 Ethernet Service Definitions

Service Type	Port-Based	VLAN-Based
E-Line	Ethernet Private Line (EPL)	Ethernet Virtual Private Line (EVPL)
ELAN	Ethernet Private LAN (EPLAN)	Ethernet Virtual Private LAN (EVPLAN)

In addition to E-Line and ELAN services, two additional service types are available for Layer 2:

- Frame Relay over MPLS (FRoMPLS)
- ATM over MPLS (ATMoMPLS)

These service types are not covered in the current MEF documentation, in spite of the fact that the MEF has merged with the Frame Relay forum.

Formerly, another terminology was used by the MEF for Layer 2 services. Table 1-3 maps the older terminology to the current one.

Table 1-2 MEF Ethernet Service Term Mappings

Current MEF Term	Former MEF Term
L2VPN over MPLS Core	
Ethernet Private Line (EPL)	Ethernet Wire Service (EWS)
Ethernet Virtual Private Line (EVPL)	Ethernet Relay Service (ERS)
ATM over MPLS (ATMoMPLS)	ATM over MPLS (ATMoMPLS)
Frame Relay over MPLS (FRoMPLS)	Frame Relay over MPLS (FRoMPLS)
VPLS over MPLS Core	
Ethernet Private LAN (EPLAN)	Ethernet Wire Service (EWS) or Ethernet Multipoint Service (EMS)
Ethernet Virtual Private LAN (EVPLAN)	Ethernet Relay Service (ERS) or Ethernet Relay Multipoint Service (ERMS)
VPLS over Ethernet Core	
Ethernet Private LAN (EPLAN)	Ethernet Wire Service (EWS)
Ethernet Virtual Private LAN (EVPLAN)	Ethernet Relay Service (ERS)

For additional information about MEF conventions and additional useful background information on Metro Ethernet standards, see the MEF website at the following URL:

<http://metroethernetforum.org>

In particular, see the document *Metro Ethernet Services Definitions Phase 2* available under the section Information Center > MEF Technical Specifications on the MEF website for a useful presentation of Metro Ethernet terms and definitions.

Mapping MEF Terminologies to Network Technologies

The MEF terminology only describes the outside characteristics of a service, that is, what the service looks like from the perspective of a customer looking in towards the user-to-network-interface (UNI) device. It does not describe how the service is implemented.

For details about how these service are implemented, see the following URL:

<http://www.cisco.com/go/ce>

In particular, see the documentation on that site on the subject of Cisco IP Next-Generation Network (NGN) Carrier Ethernet Design. The IP NGN Carrier Ethernet Design represents key elements of the Cisco IP NGN architecture that enable a best-in-class implementation for consistent service delivery optimized to meet the specific demands of each service. It is the end-to-end service transport foundation from the network access to the IP/MPLS core. This design provides integrated linkages with the service and application layer components to offer a converged, intelligent, reliable, and scalable network model to meet current and future network service requirements.

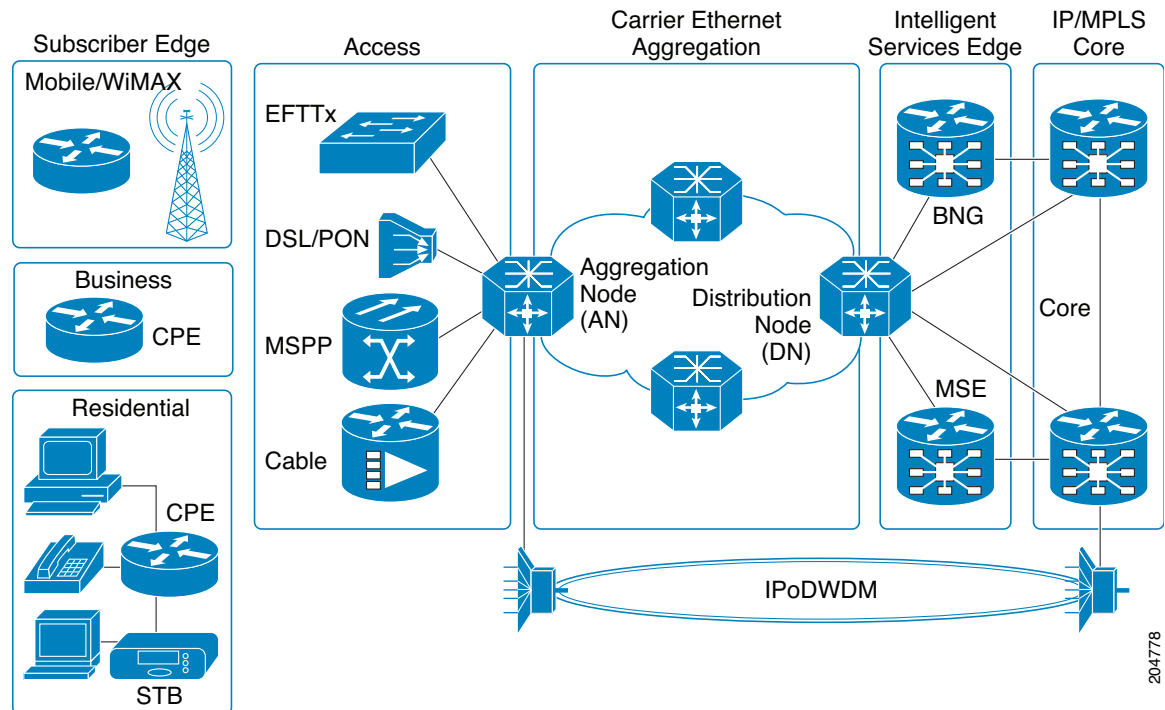
The IP NGN Carrier Ethernet Design (see [Figure 1-1](#)) provides a platform-independent architecture and Ethernet-based services model across all Carrier Ethernet platforms. This allows service providers to optimize service transport with the intelligence of appropriate networking technologies (such as Ethernet, IP, MPLS, multicast, pseudowire, or hierarchical private virtual LAN services) to meet their business and quality-of-experience goals.



Note

Real-world network implementations may implement only a subset of this very scalable architecture.

Figure 1-1 IP NGN Carrier Ethernet Design



204778

Prime Fulfillment Terminology and Supported Network Types

This section discusses the Prime Fulfillment terminology for Layer 2 services and supported network types. Prime Fulfillment can provision the following service types:

- E-Line (EPL/EWS and EVPL/ERS)
- ELAN (EPLAN and EVP-LAN/ERMS)
- FRoMPLS
- ATMoMPLS

Prime Fulfillment also supports provisioning Ethernet services on a network that consists only of Ethernet switches (no MPLS), and this is referred to in Prime Fulfillment terminology as VPLS with L2 core.



Note

For E-Line and ELAN services, we recommend using the EVC service policy type (see the appropriate chapters in this guide for how to create EVC policies and service requests). You might have existing services that have been provisioned using the L2VPN and VPLS service policy types. These are still supported and can be maintained with those service types, but new services should use the EVC service policy type. For ATM and FRoMPLS services, use the L2VPN service policy, as before.

In the Prime Fulfillment GUI and throughout this user guide, the naming conventions for these Ethernet services appear. These align closely with the earlier MEF conventions. This is expected to be updated in future releases of Prime Fulfillment. The equivalent terms used by the MEF forum are summarized in [Table 1-3](#), for reference.

Table 1-3 Ethernet Service Term Mappings

Term Used in Prime Fulfillment 6.2 GUI and This User Guide	Current MEF Equivalent Term
L2VPN over MPLS Core	
Ethernet Wire Service (EWS)	Ethernet Private Line (EPL)
Ethernet Relay Service (ERS)	Ethernet Virtual Private Line (EVPL)
ATM over MPLS (ATMoMPLS)	—
Frame Relay over MPLS (FRoMPLS)	—
VPLS Over MPLS Core	
Ethernet Wire Service (EWS) or Ethernet Multipoint Service (EMS)	Ethernet Private LAN (EPLAN)
Ethernet Relay Service (ERS) or Ethernet Relay Multipoint Service (ERMS)	Ethernet Virtual Private LAN (EVPLAN)
VPLS over Ethernet Core	
Ethernet Wire Service (EWS)	Ethernet Private LAN (EPLAN)
Ethernet Relay Service (ERS)	Ethernet Virtual Private LAN (EVPLAN)

L2VPN Service Provisioning

This section provides an overview of Prime Fulfillment provisioning for L2VPN services that provide Layer 2 point-to-point connectivity over an MPLS core. Cisco's Any Transport over MPLS (AToM) supports these services. These implementations, in turn, support service types, as follows:

- Ethernet Wire Service (EWS). The MEF term for this service is EPL.
- Ethernet Relay Service (ERS). The MEF term for this service is EVPL.
- ATM over MPLS (ATMoMPLS)
- Frame Relay over MPLS (FRoMPLS)

Instructions on creating policies and service requests for these services are provided in other chapters of the guide. For more information, see the following sections:

- [Point-to-Point Ethernet \(EWS and ERS\) \(EPL and EVPL\)](#), page 1-5
- [ATM over MPLS \(ATMoMPLS\)](#), page 1-8
- [Frame Relay over MPLS \(FRoMPLS\)](#), page 1-9

Point-to-Point Ethernet (EWS and ERS) (EPL and EVPL)

The EWS and ERS services (also known as EPL and EVPL, respectively, in MEF terminology) are delivered with the Cisco Metro Ethernet offering. The same network architecture can simultaneously provide both ERS (EPL) and EWS (EVPL) connections to diverse customers. Additionally, this Metro Ethernet infrastructure can 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.

Ethernet Wire Service (EWS or EPL)

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 tunnelling 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.

The MEF term for this service is EPL.

Ethernet Relay Service (ERS or EVPL)

An Ethernet Virtual Circuit (EVC) is used to logically connect endpoints, but multiple EVCs could 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).

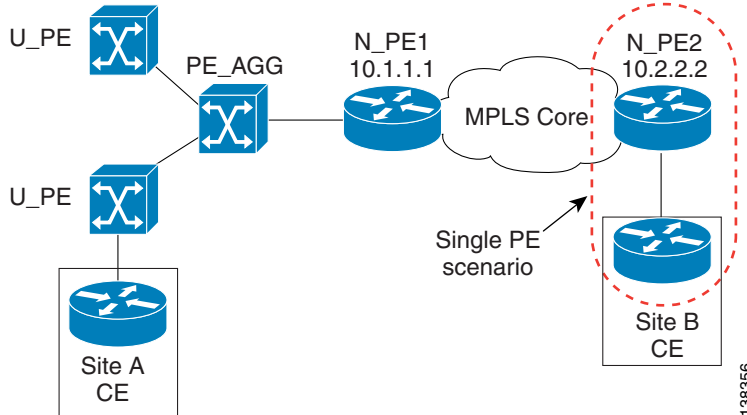
The MEF term for this service is EVPL.

Topology for L2VPN Ethernet Over MPLS (ERS and EWS) (EPL and (EVPL)

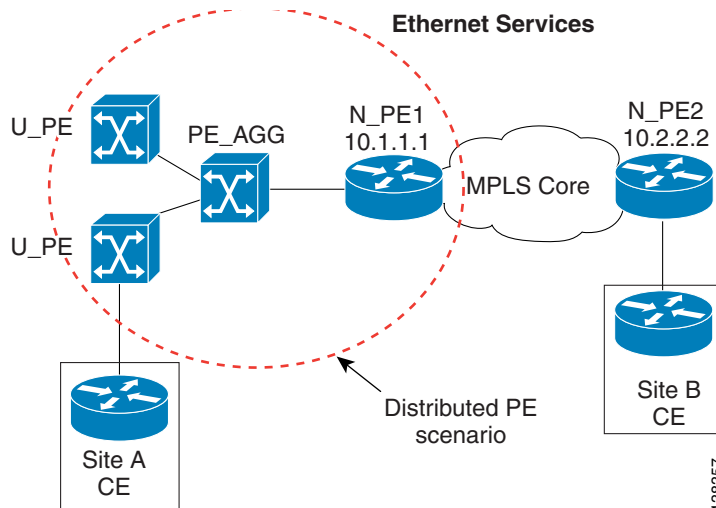
Ethernet Over MPLS (EoMPLS) is a tunnelling mechanism that allows the service provider to tunnel customer Layer 2 traffic through a Layer 3 MPLS network. It is important to remember that EoMPLS is a point-to-point solution only.

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 N-PE in [Figure 1-2](#).

Figure 1-2 Single PE scenario

- Distributed PE scenario—The end customer is connected through an Access Domain to the N-PE in [Figure 1-3](#). That is, there is a Layer 2 switching environment in the middle of CE and N-PE.

Figure 1-3 Distributed PE Scenario

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

- Automatically assigned by Prime Fulfillment 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, Prime Fulfillment creates a point-to-point tunnel and then targets the EoMPLS tunnel to the peer N-PE router through which the remote site can be reached. The remote N-PE is identified by its loopback address. In [Figure 1-4](#), N-PE1 and N-PE2 have 10.1.1.1 and 10.2.2.2 as loopback addresses. In [Figure 1-4](#), 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 N-PE).

For the N-PE 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.

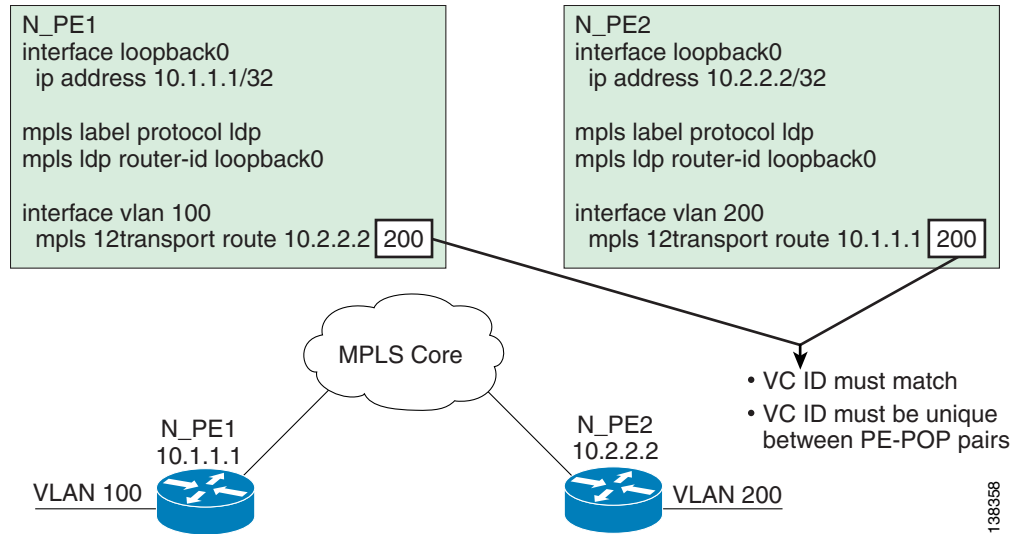


Note

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


The VC ID that defines the EoMPLS tunnel is 200. (See Figure 1-4.)

Figure 1-4 Ethernet over MPLS Configuration



Note that the VC ID has to be the same on both ends of the EoMPLS tunnel. On each N-PE, there is mapping done between the VLANs to the EoMPLS tunnel. (See Figure 1-5.)

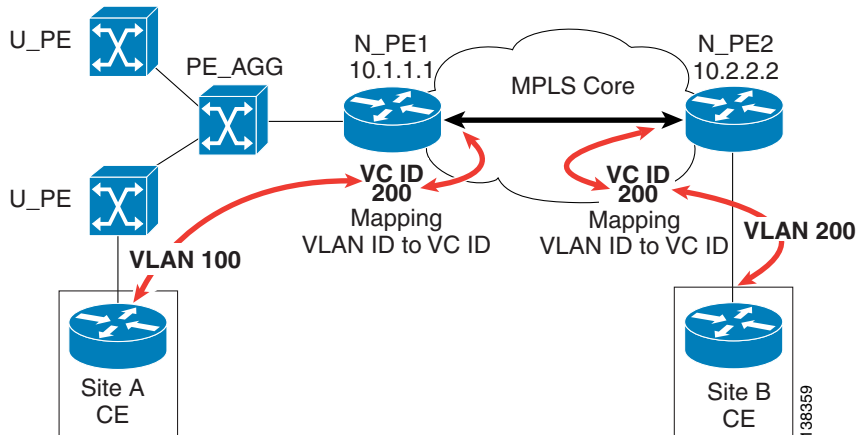
Figure 1-5 EoMPLS Tunnel



For the overall connection, this mapping is: VLAN ID <-> VC ID <-> VLAN ID.

This VLAN-VC ID mapping lets the service provider reuse VLAN IDs in Access Domains. (See Figure 1-6.)

Figure 1-6 VLAN-VC ID Mapping



The VLAN IDs allocated and used at each access domain do not have to be the same.

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. The encapsulation type is AAL0.

Topology for ATMoMPLS

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

Figure 1-7 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 FRoMPLS

Only the single PE scenario is supported. (See [Figure 1-8](#).)

Figure 1-8 Frame Relay over MPLS



EVC Ethernet Service Provisioning

This section describes EVC Ethernet and EVC ATM-Interworking support in Cisco Prime Fulfillment 6.2. It contains the following topics:

- [Overview, page 1-10](#)
- [EVC Features, page 1-11](#)
- [Platform Support for EVC in Cisco Prime Fulfillment 6.2, page 1-12](#)
- [Device Roles with EVC, page 1-14](#)
- [Topology Overview for EVC, page 1-14](#)
- [A Note on Checking of Configurations, page 1-15](#)
- [EVC ATM-Ethernet Interworking Service Provisioning, page 1-15](#)



Note

For Ethernet (E-Line and ELAN) services, use of the EVC policy and service request is recommended. If you are provisioning services using the EVC syntax, or plan to do so in the future, use the EVC service. Existing services that have been provisioned using the L2VPN and VPLS service policy types are still supported and can be maintained with those service types. For ATM and FRoMPLS services, use the L2VPN service policy, as before.

Overview

Flexible user network interface (FlexUNI) is a generic approach for creating Ethernet services in Prime Fulfillment. It can, if supported by the hardware, be used for all Ethernet provisioning. (For information on what platforms support EVC see [Platform Support for EVC in Cisco](#))

[Prime Fulfillment 6.2, page 1-12.](#)) The EVC policy is flexible and generic and allows for service designers to provide greater service offerings than available through traditional Prime Fulfillment L2VPN and VPLS services.

Certain line cards have interfaces that support the Cisco IOS Ethernet Virtual Circuit (EVC) syntax. These interfaces can be configured with either EVC infrastructure features or with switch-port command-line interface commands. FlexUNI optionally supports the EVC CLI syntax/infrastructure. For this reason, the FlexUNI policies and service requests are referred to by the umbrella term “EVC.” However, it is important to note that EVC policies and service request are not tied to the new EVC syntax. Service endpoints can use non-EVC syntax also.

Services leveraging the EVC infrastructure are varied in nature, and there is not always a clear delineation between different services. This is because EVC provides great flexibility in the way these services can be delivered. This can make it challenging to define the services. For example, a traditional ERS could be delivered in several ways by variations of the Class on the platform.

The EVC policy and associated service request offer a generic and flexible service construct to support device capabilities. This policy is flexible enough to cater to different service offerings using the EVC architecture. It allows service designers to utilize most of the EVC features in a flexible manner, to match the hardware and platform capabilities.

The EVC policy can be used to create only an EVC service request and not any other existing Prime Fulfillment service request types, such as L2VPN, VPLS, and so on. Likewise, an EVC service request can be created using only an EVC policy and not any other existing Prime Fulfillment policies.

The EVC infrastructure provides several benefits to Carrier Ethernet (CE) deployments, including:

- Flexible frame matching.
- Flexible VLAN tag manipulation and/or translation.
- Multiple services on the same port.
- Flexible service mapping.
- VLAN scaling and locally significant VLANs.

EVC supports a variety of network configurations, such as the following:

- Provisioning of Ethernet access as an EVC-capable EWS interface on the N-PE.
- Interconnecting Ethernet accesses terminating on a single Cisco 7600 N-PE on one or multiple ports in a bridge domain.
- Interconnecting Ethernet accesses terminating on multiple Cisco 7600 N-PEs in a VPLS service.
- EVC service support on Cisco ASR 9000 Series Routers running IOS XR.
- Services that combine the existing services with the Ethernet access, including the ERS/EWS interworking service.
- Provisioning of E-Line services, in which one or both N-PE interfaces are EVC.

EVC Features

This section summarizes the features supported by the EVC policy and service request in Prime Fulfillment:

- Choice of topology:
 - Customer edge device (CE) directly connected.
 - CE connected through Ethernet access devices.

- Choice of platforms:
 - EVC on all N-PEs.
 - EVC on none of the N-PEs.
 - Mix of EVC and the old infrastructure. This allows both the old and new platforms to co-exist, in order to ensure continued support for deployed platforms.
- Choice of connectivity across the MPLS core (with or without bridge-domain):
 - Pseudo wires
 - VPLS
 - Local (local connects)
- Flexible VLAN handling mechanism that deals with up to two levels of VLAN tags:
 - VLAN matching for service classification. This provides the ability to match both outer and inner VLAN tags, or the ability to match a range of inner VLAN tags.
 - VLAN manipulations, such as pop outer tag, pop inner tag, push outer tag, push inner tag, and VLAN translations (1:1, 2:1, 1:2, 2:2).
- Flexible forwarding options:
 - Configure a pseudowire on the MPLS core directly under a service instance (for E-Line only).
 - Configure a pseudowire on the MPLS core under a switch virtual interface (SVI) by associating it to a bridge domain.



Note The appropriate VLAN manipulations are applicable to pseudowire in both cases.

- Associate traffic from different interfaces and/or VLANs onto a single bridge domain, with appropriate VLAN manipulations for VPLS.
- Associate traffic from different interfaces and/or VLANs onto a single bridge domain with appropriate VLAN manipulations for local connects.

Platform Support for EVC in Cisco Prime Fulfillment 6.2

EVC services are supported on both IOS and IOS XR platforms. See the [Supported Devices Table for Prime Fulfillment 6.2](#) for the most current list of hardware and software platforms supported in this release.

http://www.cisco.com/en/US/docs/net_mgmt/prime/fulfillment/6.2/supported/devices/supported_devices_table.xls

Two example platform scenarios are covered in the next sections. Note that the UNI characteristics and the EVC capabilities of the N-PE are not inter-dependent.

Example 1

EVC service requests allow operators to add links with either an EVC-capable interface and/or the nonEVC-capable interface on the N-PE. For example, an operator can add three links to an EVC service request (with VPLS connectivity) with the following configurations:

- Link one has a Cisco 67xx interface and an IOS 12.2 (33)SRB image on a Cisco 7600 N-PE.
- Link two has a Cisco 67xx interface and an IOS 12.2 (33)SRB image on a Cisco 7600 N-PE.

- Link three has an ES20-based interface and an IOS 12.2 (33)SRB image on a Cisco 7600 N-PE.

Example 2

As far as Layer 2 access nodes are concerned, configurations on the UNI/NNI of a U-PE and/or PE-AGG are not influenced by the EVC capability on the N-PE. However, if a selected named physical circuit (NPC) with N-PE interface is configured with EVC, it cannot be provisioned for traditional configuration. An error will be generated while saving the service request.

On the other hand, if a selected NPC with N-PE interface is configured without EVC, it cannot be provisioned for EVC configuration. An error will be generated while saving the service request.

For example, if for link one of an EVC service request, if the encapsulation is selected as dot1Q, the interface can share other L2 ERS/VPLS ERMS UNIs on the same U-PE/PE-AGG.

If the N-PE interface that is part of the NPC being picked is already configured with non-EVC features (using an existing L2VPN or VPLS service request), you cannot configure EVC on it.



Note

If “Dot1Q Tunnel” is selected as the encapsulation type, the port cannot be shared with other services.

IOS XR Platform Support

EVC services are supported on Cisco ASR 9000 Series Routers running IOS XR.



Note

See the [Cisco Prime Fulfillment Installation Guide 6.2](#) for the most current list of hardware and software platforms supported in this release.

The following EVC features are supported on IOS XR platforms:

- E-line connections. If an ASR 9000 is added on a direct link, only DOT1Q encapsulation is supported for E-Line services. When using L2 access nodes with NPCs, all supported encapsulations are available.
- ELAN connections.
- Flexible frame matching.
- Flexible VLAN tag manipulation/translation.
- VLAN scaling and locally significant VLANs.
- The ability to create L2 and L3 services under the same physical interface, restricted only to subinterface.
- All the Layer 2 ports on Cisco ASR 9000 devices are trunk ports, hence only trunk port-based configuration is supported.

The following EVC services are not supported on IOS XR platforms:

- The N-PE Pseudo-wire on SVI attribute is not supported. SVI interfaces are not available on the devices, which restricts support for Standard UNI and Port Security configuration when the UNI is configured on an N-PE.
- xconnect commands are not directly supported under interface configuration. Support for these commands has been moved to different a hierarchy in IOS XR.
- When the UNI is configured on an N-PE device, EWS service is not supported. Since the Cisco ASR 9000 device is a router, all the Layer 2 ports are default trunk. There is no option for configuring access ports, which restricts support for access port-based services.

**Note**

Unless otherwise noted, all of the EVC policy and service request features documented in this guide are applicable for both IOS and IOS XR platforms.

Device Roles with EVC

Presently, Prime Fulfillment has U-PE, PE-AGG and N-PE devices. The basic PE device role association of Prime Fulfillment continues for EVC policy and service requests. In this release of Prime Fulfillment, there are no changes made to the PE role assignment. A device having EVC capabilities will not call for a change in the existing role assignment in Prime Fulfillment. However, EVC capabilities in Prime Fulfillment are supported only for interfaces on N-PE and not on PE-AGG or U-PE devices.

**Note**

Prime Fulfillment does not support customer edge devices (CEs) for EVC. If the access port contains any DSLAMS, non-Cisco Ethernet devices and/or other Cisco devices that are not supported by Prime Fulfillment, such nodes and beyond are not in the scope of Prime Fulfillment. In such cases, from the Prime Fulfillment perspective, the interface on the first Prime Fulfillment-managed device is the UNI.

Topology Overview for EVC

This section provides examples of various topologies supported with EVC. As mentioned in the note at the end of section [Device Roles with EVC, page 1-14](#), Prime Fulfillment does not support customer edge devices (CEs) with EVC. References to the term “CE” in the following topology variations (such as “CE directly connected” and so on) is only to indicate how the customer or third-party devices connect to the N-PE. For all the cases involving EVC, the CE is not supported in Prime Fulfillment. Also, any provider device that is not supported by Prime Fulfillment, and which is used in the access circuit, marks the boundary for the scope of Prime Fulfillment, beyond which no devices (that is, towards the CE, and including the unsupported node) is managed by Prime Fulfillment.

CE Directly Connected and EVC

With this combination, the UNI is the interface on a supported line card, with EVC capability configured. Prime Fulfillment does not configure Prime Fulfillment’s standard UNI functionality (for example, port-security, storm control, and Layer 2 Protocol Tunneling). This is because of lack of command support on the EVC-capable hardware. Operators can use templates to configure relevant platform supported parameters to realize any of these features not provided by Prime Fulfillment. Prime Fulfillment configures only the service instance with VLAN manipulations and pseudowire, VPLS, or local-connect on the UNI. NPCs are not needed while creating such links because NPCs are only required when there are access nodes between the N-PE and CE. Other intermediate Ethernet access nodes are not involved in this topology.

CE Directly Connected and No EVC

This is similar to the UNI on N-PE case in Prime Fulfillment. The EVC service request can be used to create such links with older Cisco 7600 platforms (that is, N-PE interfaces without EVC capability), but with plans of adding one or more future links with EVC support. If not, one could use the existing

ERS/EWS/ERMS/EMS functionality in Prime Fulfillment. NPCs are not needed while creating such links because NPCs are only required when there are access nodes between the N-PE and CE. Other intermediate Ethernet access nodes are not involved in this topology.

CE Not Directly Connected and EVC

This topology involves the following configurations:

- UNI on a U-PE or PE-AGG to which the CE is connected.
- Ethernet U-PE and/or PE-AGGs.
- N-PE with EVC-capable interface on the CE-facing side.

All service-specific parameters, such as port-security, L2 Protocol Tunneling, storm control, and so on, are applicable to the UNI (Standard UNI) in such links. The U-PE and/or PE-AGG configurations will also have no change in CLIs. However, the EVC commands are applicable only on the N-PE (on the CE-facing interface). NPCs are used while creating such links.

CE Not Directly Connected and No EVC

This link is identical to an attachment circuit in existing Prime Fulfillment implementations. This has a standard UNI as in existing Prime Fulfillment services. NPCs are used while creating such links.

A Note on Checking of Configurations

Prime Fulfillment attempts to provision all configurations generated by an EVC service request. Prime Fulfillment does not perform any prior checks to verify if the CLIs are compatible with the specific devices being provisioned. This is to ensure flexibility of support for device/platform features, which could change over time. Hence, it is important for the service designer or operator to carefully create the EVC policies and service requests.

EVC ATM-Ethernet Interworking Service Provisioning

Prime Fulfillment supports interworking of a service with ATM and Ethernet protocols across the MPLS core or local switching. ATM-Ethernet interworking is supported through the following features:

- Creation of an EVC policy of type “ATM-Ethernet Interworking.” The ATM-Ethernet Interworking policy type supports a choice of MPLS core options:
 - Pseudowire
 - Local (local connects)
- Provisioning of ATM-Ethernet interworking using a single EVC service request.
- Combination of EVC and non-EVC syntax, that is, the creation of an L2 circuit consisting of an L2 syntax and EVC syntax.

VPLS Service Provisioning

VPLS services are multipoint. They provide multipoint connectivity over an MPLS or an Ethernet core. These implementations, in turn, support service types, as follows:

- VPLS over MPLS core:
 - Ethernet Wire Service (EWS). This is also sometimes referred to as EMS, or Ethernet Multipoint Service. The MEF term for this service is EPLAN.
 - Ethernet Relay Service (ERS). This is also sometimes referred to ERMS, or Ethernet Relay Multipoint Service. The MEF term for this service is EVP-LAN.
- VPLS over Ethernet core:
 - Ethernet Wire Service (EWS). The MEF term for this service is EPLAN.
 - Ethernet Relay Service (ERS). The MEF term for this service is EVPLAN.

Instructions on creating policies and service requests for these services are provided in other chapters of the guide.

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.

For more information, see the following sections:

- [Multipoint EWS \(EPLAN\) for an MPLS-Based Provider Core, page 1-16](#)
- [Multipoint ERS \(EVP-LAN\) for an MPLS-Based Provider Core, page 1-16](#)
- [Topology for MPLS-Based VPLS, page 1-17](#)

Multipoint EWS (EPLAN) for an MPLS-Based Provider Core

With multipoint EWS (also known as EPLAN in MEF terminology), 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 (EVP-LAN) for an MPLS-Based Provider Core

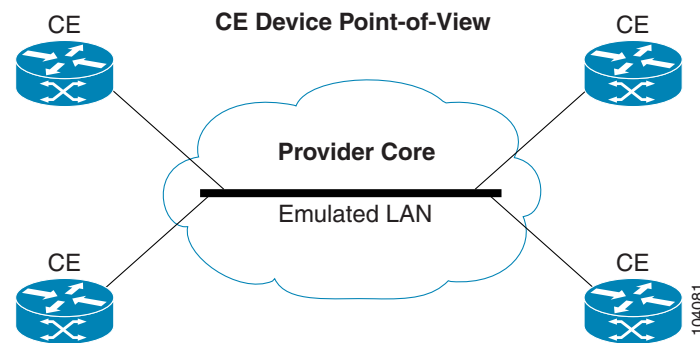
With multipoint ERS (also known as EVP-LAN in MEF terminology), the PE router forwards all Ethernet packets with a particular VLAN tag received from an attachment circuit, excluding BPDU, to another attachment circuit or an emulated VC if the destination MAC address is found in the L2

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

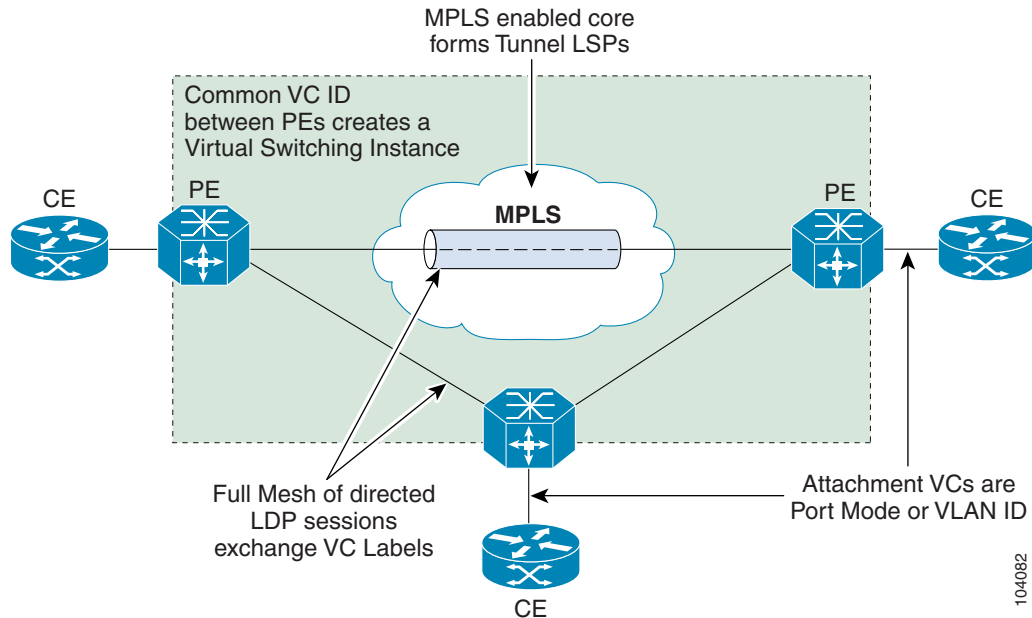
From a customer point of view there is no topology for VPLS. All the CE devices are connected to a logical bridge emulated by the provider core. Therefore, the CE devices see a single emulated LAN. (See Figure 1-9.)

Figure 1-9 *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 1-10.)

Figure 1-10 Full Mesh of Emulated VCs



104082

VPLS for an Ethernet-Based (L2) 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.

For more information on VPLS for an Ethernet-based provided core, see the following sections:

- [Multipoint EWS \(EPLAN\) for an Ethernet-Based Provider Core, page 1-18](#)
- [Multipoint ERS \(EVP-LAN\) for an Ethernet-Based Provider Core, page 1-19](#)
- [Topology for Ethernet-Based VPLS, page 1-19](#)

Multipoint EWS (EPLAN) for an Ethernet-Based Provider Core

Multipoint EWS (also known as EPLAN in MEF terminology) 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 might be present within the customer frames, the service employs a concept of “All to One” bundling.

Multipoint ERS (EVP-LAN) for an Ethernet-Based Provider Core


Multipoint ERS (also known as EVP-LAN in MEF terminology) 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 (EPLAN) and ERS (EVP-LAN) by providing a multipoint connectivity model. The VPLS service does not map an interface or VLAN to a specific point-to-point pseudowire, 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 (EPLAN).

The EWS (EPLAN) service emulates the service attributes of an Ethernet switch and learns source MAC to interface associations, flooding unknown broadcast and multicast frames. [Figure 1-11](#) illustrates an EWS (EPLAN) VPLS topology.

Figure 1-11 VPLS EWS Topology



The Ethernet Relay Service (ERS or EVP-LAN) 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 1-12](#) illustrates an ERS (EVP-LAN) VPLS multipoint topology.

Figure 1-12 VPLS ERS (EVP-LAN) Multipoint Topology

