

Metro Ethernet WAN Services and Architectures

Abstract

Data networking service providers worldwide are deploying *Ethernet* technology to offer wide-area network (WAN) services to business and residential customers within metropolitan regions as well as to offer long distance WAN connectivity. Though Ethernet is a widely used LAN technology, its adoption for WAN services has generated much interest as well as some confusion in the data networking industry. To provide a common framework for metropolitan-area (or metro) and wide area Ethernet, this paper presents a service-centric approach to Ethernet WAN services and defines four canonical services. Service-related attributes that enable and enhance Ethernet WAN services are also described in detail. Reference network architectures, including those being discussed in standards groups are also included for completeness.

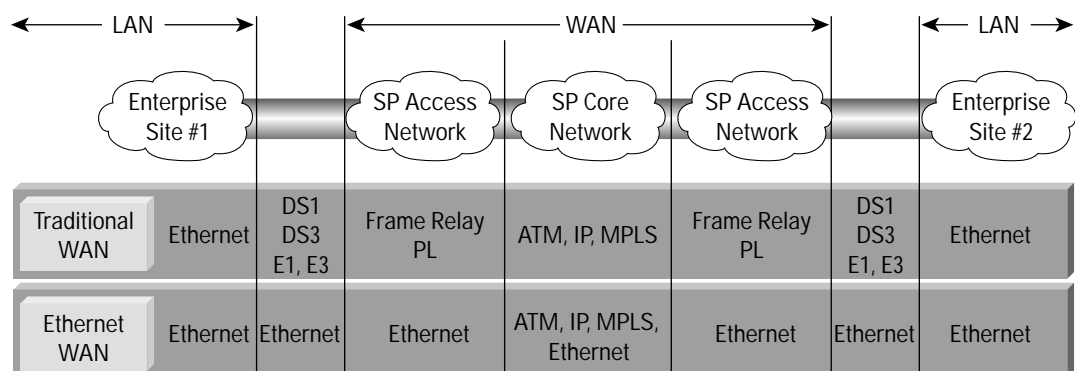
Introduction

Metro Ethernet services are gaining significant popularity among data networking service providers and their enterprise customers. Traditionally, Ethernet-based networks have been deployed in enterprise LANs because of their simplicity, low equipment cost, high speed and multivendor interoperability. Service providers have also recognized these benefits of Ethernet and have begun to offer Ethernet-based WAN services as alternatives to traditional WAN services such as Frame Relay and private line (for

example, DS1, DS3). Though metro Ethernet services are being initially offered in metropolitan areas, modern Ethernet technology allows distances greater than a metro, and thus inter-metro Ethernet WAN service offerings are also being planned by major telecommunication service providers.

Partitioning of LAN and WAN topologies and the technology options available for deploying WAN services are shown in Figure 1. As seen from the figure, Ethernet technology is being deployed in the WAN access network for offering Ethernet WAN services. In the WAN core

Figure 1
 LAN and WAN
 Technologies





network, Ethernet as well as non-Ethernet technologies may be deployed, including ATM, SONET/SDH, IP, and IP/MPLS. In the case of the latter, *service interworking* between Ethernet WAN access and non-Ethernet WAN core may be required. The service user always sees an Ethernet interface, independent of the service provider's WAN core.

Traditional as well as non-traditional service providers are building Ethernet networks to provide Ethernet WAN services. These include the incumbent local exchange carriers (ILECs), inter-exchange carriers (IXCs), Postal Telephone and Telegraph (PTT) carriers, municipalities and utilities, and Ethernet service providers. The high-capacity fiber network build-outs of the last decade have also motivated service providers to offer high-bandwidth services to better use their network infrastructure and to generate new revenue streams.

Though Ethernet service providers in the United States have struggled for various reasons, they appear to have a reasonable business model in Europe and the Asia-Pacific region. FastWeb is a promising example of a European Ethernet service provider that offers triple-play services (Internet access, IP telephony, and IP video services) in cities such as Milan, Italy. For more information, see www.fastweb.com. Customer demographics in these regions are quite favorable for offering profitable Ethernet WAN and other vertical services. The reason is that both businesses and residences densely populate European and Asian cities. Moreover, penetration of broadband data services and cable video services is relatively low in Europe and Asia when compared to the United States. Many high-rise buildings in these cities contain several floors of retail shops, several floors of small and mid-sized businesses, and remaining floors of residential dwellings. By bringing Ethernet services to these buildings, the Ethernet service provider has a sizable customer footprint to recover its infrastructure investment in a reasonable timeframe. Also, Ethernet service providers in these regions tend to offer data, voice, and video services for a faster return on investment (ROI).

The remainder of this paper is organized as follows:

- Service provider and enterprise benefits of Ethernet-based WAN services
- Recent advances in Ethernet
- Definition of Ethernet WAN services and their service-enabling and service-enhancing attributes
- Network architectures that enable Ethernet WAN services
- References and Acronyms

Why Use Ethernet?

Only a few years ago, Ethernet technology was one of many LAN technologies. In LANs, Ethernet was competing with technologies such as FDDI, Token Ring, and Token Bus. In WANs, Frame Relay and ATM became dominant Layer 2 technologies, and Ethernet was not a consideration until only recently. The success of Ethernet as a Layer 2 data networking technology of choice in the LAN environment was primarily due to continued improvement of the Ethernet control-plane by the IEEE 802.1 Working Group [1], tremendous increase in Ethernet connection speed, and steady decrease in the deployment and operational cost of Ethernet equipment. Some key milestones in the evolution of the Ethernet technology are: Ethernet switching, Fast Ethernet (100 Mb/s), resilient platforms for mission-critical services, virtual LAN (VLAN) support, longer reach physical layer devices, and Gigabit and 10-Gigabit Ethernet speeds. Thus the Ethernet technology has marched on, primarily driven by the enterprise LAN engine, and is now poised to succeed in the WAN access services.

As the enterprise LAN bandwidth improved significantly due to Ethernet, the disparity between LAN and WAN bandwidths steadily worsened. The demand for Ethernet WAN is primarily driven by the enterprise need to bridge this bandwidth gap and by the multiservice capability it offers. For example, improved network response time is a



key requirement for emerging enterprise information technology (IT) applications such as distributed network storage, data redundancy and recovery, voice over IP, e-learning, multimedia-based corporate communication, and video conferencing.

Traditional WAN data services are primarily based on the time-division multiplexing (TDM) network infrastructure. WAN speeds are essentially bounded by the speeds offered by DS1 (up to 1.5 Mbps), E1 (up to 2 Mbps), E3 (up to 32 Mbps), and DS3 (up to 45 Mbps) circuits. Layer 1 private line service and Layer 2 Frame Relay service are the two dominant WAN services offered today. Inter-site connectivity and Internet access have been the primary drivers for these traditional WAN services.

Ethernet technology enables WAN connectivity at much higher bandwidths (e.g., at 1 Gb/s) than that offered by the traditional WAN technologies. Enterprises seeking to improve WAN bandwidth in a cost-effective manner are attracted to Ethernet WAN services for several reasons:

- No special learning curve is needed to set up Ethernet WAN connections with the service provider.
- No special customer equipment or port adapters are required. Customer equipment can be either a router (with an Ethernet interface), an Ethernet switch or bridge, or a host.
- WAN bandwidths of up to 1 Gbps are available, with speeds of 10 Gbps to follow soon. Flexible bandwidth granularities are possible (for example, 32 Kbps at the lower speed range and 1 Mbps at the upper speed range).
- Service prices, in terms of dollars per megabits per second, are quite favorable when compared to the prices for TDM connections.
- Multipoint connectivity is provided by Ethernet, thus alleviating the need to manage multiple point-to-point connections on the customer equipment.

Service providers are also keen to offer Ethernet WAN services. The main reasons for this are:

- Ethernet technology is mature, standards-based and highly multivendor interoperable.
- Cost of deployment of Ethernet network infrastructure tends to be low, due to the availability of low-cost Ethernet switches. The overall switch cost should continue to remain lower due to further integration of network functions in switch ASICs and due to economies of scale achieved by mass deployment of Ethernet switches in enterprises.
- WAN Ethernet provides new revenue streams by using the existing fiber-rich infrastructure and also provides opportunity for easier uptake.
- With the improved WAN response time, many enterprises are expected to subscribe to value-added (often higher margin) managed services; for example, in areas of voice, storage, security, content delivery, and application hosting. The overall customer base is also expected to increase, thus further enhancing the revenue and profit potential.
- Provisioning costs are significantly reduced as the wide range of bandwidth and Ethernet familiarity within the customer base alleviates the need for a truck-roll for service activation and service upgrades.
- Ethernet's inherent multipoint connectivity alleviates the need to provision and manage multiple point-to-point connections.

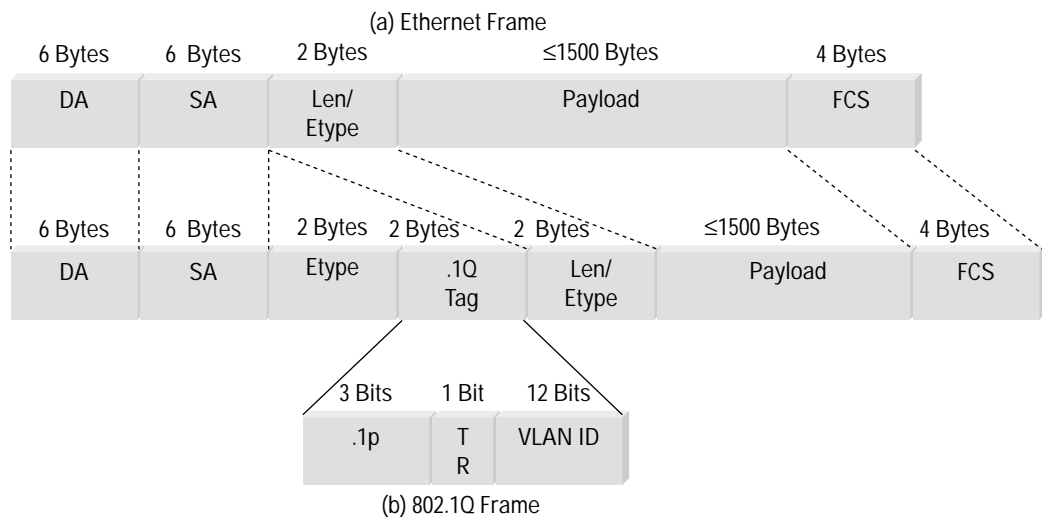


Recent Advances in Ethernet

Ethernet started as a shared, 10 Mbps, half-duplex media access technology. As Ethernet began to proliferate, there has been a sustained effort over the last two decades to significantly enhance it. This section discusses these advancements in Ethernet technology.

Figure 2(a) depicts the Ethernet frame format. It has a media access control (MAC) header that consists of, among other things, the source MAC address, the destination MAC address, and Length/EtherType. The maximum size of an Ethernet frame, including header and payload, is 1518 bytes. Each Ethernet frame is broadcasted to all devices on the shared Ethernet network segment. The device, whose MAC address matches the frame's destination MAC address, accepts the frame while remaining devices ignore that frame. Frames transmitted by distinct Ethernet devices often experience collisions, and these collisions negatively impact the overall network throughput. Several mechanisms have been adopted to improve Ethernet: reducing the collision domain by bridging and switching, connecting bridges in a mesh topology for network resiliency in the event of a bridge or link failure, and limiting the broadcast domain via VLAN frame tagging.

Figure 2
Frame Format (a) Ethernet Frame (b) 802.1Q Frame



Limiting the Ethernet Collision Domain

The size of the Ethernet collision domain is reduced by segmenting it into two sub-domains. Ethernet bridges are used for this purpose. A bridge learns about the devices (for example, their MAC addresses) on each segment by snooping each frame's source MAC address. It builds a MAC-to-segment map table. After the topology is fully learned, the bridge forwards only inter-segment frames and thus helps reduce the number of collisions on a given segment. When the bridge receives a frame with an unknown destination MAC address, it floods the frame on all its segments. To completely remove collision domains, each bridge segment can be limited to only one device. That is, each device is connected to one full-duplex port of an Ethernet bridge (referred to as an Ethernet switch in such cases).



Resiliency in Ethernet-Bridged Networks

It is generally desirable to connect bridges in a mesh topology to ensure network resiliency in the event of a bridge or a link failure. To maintain loop-free operation, the IEEE 802.1D *Spanning-Tree Protocol* has been defined [1]. Spanning-Tree Protocol uses special control frames, called *bridge protocol data units* (BPDUs), to identify whether a bridge port is in a *forwarding* state or in a *blocking* state. When the bridge topology is changed (for example, when a bridge is added or deleted or when a failure occurs), the Spanning-Tree Protocol provides sub-minute convergence to a loop-free topology. The Spanning-Tree Protocol protocol has been recently refined by the IEEE 802.1w Working Group to provide faster convergence [1]. This refined protocol, achieving sub-second convergence, is referred to as the *Rapid Spanning-Tree Protocol*.

Managing Ethernet Broadcast Domain

Though switching plus full-duplex operation has all but eliminated the collision domain, an Ethernet switch can still experience congestion due to broadcast/multicast traffic and due to frames with unknown destination MAC addresses. In either case, the switch floods the frame on all its ports. To limit Ethernet broadcast domains, the concept of VLAN has been defined. A VLAN is a logical representation of an Ethernet broadcast domain. Associating one or more ports of each switch with the corresponding VLAN creates a VLAN-based broadcast domain. In each VLAN, broadcast traffic is then limited to switch ports that are associated with that particular VLAN.

In a VLAN-aware switch, each Ethernet frame's MAC header is extended with a 16-bit IEEE 802.1Q tag that includes a 12-bit VLAN ID [1]. The tag also includes three 802.1p *priority* bits. One bit of the 802.1Q tag is used to indicate Token Ring encapsulation. The structure of an Ethernet frame with an 802.1Q tag is shown in Figure 2(b). These extended-header frames are commonly referred to as .1Q frames or *tagged* frames. Note that the .1p bits are used to provide priority treatment to a tagged frame, and up to eight levels of priority may be indicated. For efficient inter-switch forwarding of tagged frames, a .1Q trunk port is configured and associated with one or more VLANs. Thus, trunking allows a single Ethernet port to carry frames with multiple VLAN IDs.

With the popularity of VLANs, it soon became clear that the Spanning-Tree Protocol/Rapid Spanning-Tree Protocol needed to be made VLAN-aware. This led to the definition of *multiple spanning-tree protocol* (MSTP) by the IEEE 802.1s Working Group [1]. MSTP allows Ethernet switches to participate in multiple spanning trees. This allows better utilization of the network links. With a single spanning tree, links connected to blocked ports are idle. With multiple spanning trees, blocking is done on a per VLAN basis, which allows traffic to be distributed over all the links in the network.

Additional Enhancements

Two additional Ethernet enhancements worth mentioning are link aggregation and authentication. Link aggregation allows multiple Ethernet links to be logically grouped to form a higher-speed aggregated link. The specific mechanism for aggregation is defined by the *Link Aggregation Control Protocol* (LACP) and is standardized by the IEEE 802.3ad Working Group [1]. In addition, the link aggregation enables very fast fail over in the event of a failure of one link in the group. Thus, link aggregation improves network scalability and network resiliency.

To enhance security, the IEEE 802.1x protocol is defined to provide authentication capability [1]. An 802.1x-aware Ethernet switch activates its port only after the host attached to that port is authenticated using information such as the host, or user name and password. Device authentication reduces the likelihood of an intruder breaking into a network and hence reduces the possibility of theft-of-service and denial-of-service (DoS) attacks.



The above-mentioned control-plane advances in Ethernet, plus improved throughput and speed and reduced cost, have made Ethernet the dominant LAN technology. These enhancements have also provided a solid foundation for extending Ethernet to the WAN. There still remain some challenges for Ethernet WAN, with network scalability being a major concern:

- Since the VLAN ID is only 12 bits long and is globally significant, the .1Q-based WAN architecture provides only 4096 distinct logical partitions for keeping service user traffic segregated within a Layer 2 Ethernet domain.
- Spanning tree convergence of a large Ethernet network (consisting of many Ethernet switches) has yet to be proven to the satisfaction of many telecommunications service providers.
- The size of MAC learning tables is also of concern, as many thousands of MAC addresses may need to be learned by an Ethernet switch.

Ethernet WAN Services

Detailed technical descriptions of Ethernet WAN services are presented in this section. In particular, four key services with their essential service attributes are proposed. Two of the four services are the Ethernet analog of existing WAN services.

The four Ethernet WAN services described below may be thought of as basic or *canonical* services. The marketing names used by service providers to sell canonical services to service users may be different, but these marketed services may be mapped to the four canonical services proposed here. Additionally, a service provider may market specific service bundles to service users (for example, a managed firewall service) that include one or more canonical services plus additional vertical services offered by the service provider or its affiliates.

A key property of the canonical services is that an existing customer equipment with a standard Ethernet interface can connect to the service provider network and use the service with *no* hardware or software change. These Ethernet WAN services are the first WAN services with over 100 million existing devices that are capable of accessing the services.

The four canonical Ethernet WAN services are now defined. See also Table 1 for their high-level descriptions.

1. *Ethernet Relay Service* (ERS) is the Ethernet analog of the Frame Relay service. Here routers or hosts are used as customer equipment devices to establish point-to-point connections between two sites. Each connection is a logical one and is identified using a VLAN ID. Multiple logical connections can be multiplexed on a single physical connection to the service provider network, allowing multiple remote sites to be reached from the single physical connection to the service provider network.
2. *Ethernet Relay Multipoint Service* (ERMS) is an important extension of ERS that allows multipoint-to-multipoint connections between multiple sites. This is a direct result of the multipoint capability of Ethernet switches. Each connection in ERMS is a logical one and is identified using a VLAN ID. Multiplexed multipoint-to-multipoint and point-to-point connections are allowed. ERMS supports routers or hosts as CE devices.
3. *Ethernet Wire Service* (EWS) is the Ethernet analog of the private line (or private wire) service. Here, routers, bridges, or hosts may be used as customer equipment devices to establish point-to-point connections between two sites. A single remote site can be reached from a single physical connection to the service provider network. EWS ensures that all traffic delivered to the destination customer equipment is unaltered.



4. *Ethernet Multipoint Service* (EMS) is the WAN analog of the multipoint Ethernet LAN capability. A single connection is established among multiple sites to deliver frames without alteration. The network delivers each unicast frame to its designated destination site based on the learned MAC address. Broadcast and multicast frames or frames with unknown MAC addresses are replicated to all sites except the originating site.

Note: Ethernet WAN service commonly referred to as *transparent LAN service* (TLS) can be considered as a special case of EWS or EMS. In particular, point-to-point TLS is a special case of EWS and multipoint TLS is a special case of EMS.

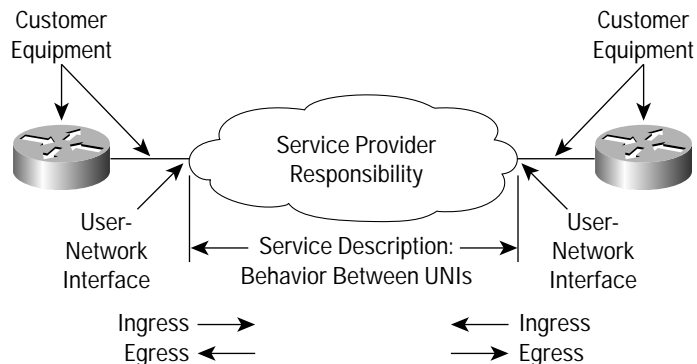
Table 1 High-Level View of the Canonical Ethernet WAN Services

Ethernet Service	Customer Equipment	Basis for Frame Delivery	Broadcast
Ethernet Relay Service	Router, host	VLAN tag	No
Ethernet Relay Multipoint Service	Router, host	VLAN tag + destination MAC address	Yes
Ethernet Wire Service	Router, bridge, host	Physical network connection point	No
Ethernet Multipoint Service	Router, bridge, host	Physical network connection point + destination MAC address	Yes

Service Model

A service is what the service user’s equipment, called *customer equipment*, sees. Implementation details of a service are invisible to the service user’s equipment. Service descriptions treat the service provider network as an “opaque cloud,” as illustrated in Figure 3. They may thus be viewed as requirements on the implementation within the service provider network. An Ethernet service model is defined by two fundamental characteristics: User-Network Interface (UNI) and Ethernet Virtual Connection (EVC).

Figure 3
Service Model





User-Network Interface

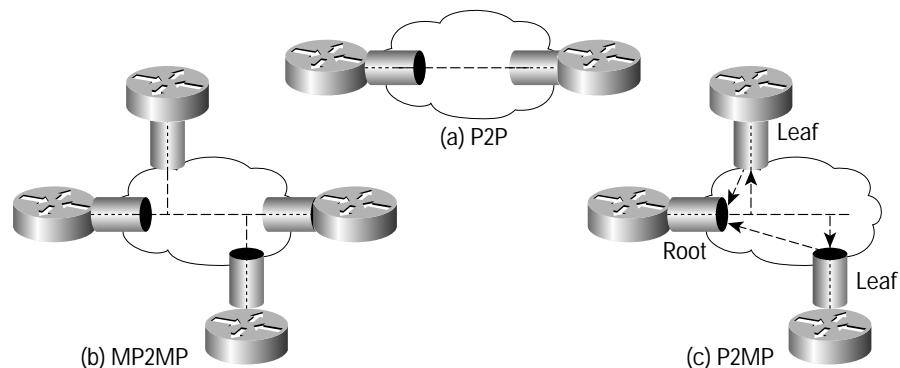
A UNI represents the demarcation point between a service provider and a service user (see Figure 3). For Ethernet WAN service, UNIs are based on standard Ethernet. A service provider could define the UNI in slightly different ways without impacting the service description. For example, in Figure 3, the implication is that the cable connecting the customer equipment to the service provider network is the responsibility of the service user. In other words, the Ethernet UNI could be an RJ-45 socket on a service provider owned device or a patch panel. If the cable was deemed to be the responsibility of the service provider, then the Ethernet UNI could be an RJ-45 connector on the customer equipment. An Ethernet UNI is specified by its physical medium, speed (10 Mbps, 100 Mbps, 1 Gbps, etc.) and mode (half- or full-duplex).

Ethernet Virtual Connection

An Ethernet Virtual Connection is an association of two or more UNIs. A frame sent into an EVC can be delivered to one or more of the UNIs in the EVC other than the ingress UNI. It must never be delivered back to the ingress UNI. It must never be delivered to a UNI not in the EVC. There are three types of EVC as described below:

- In a *point-to-point* EVC, exactly two UNIs are associated. An ingress frame at one UNI can only be an egress frame at the other UNI. A point-to-point EVC is shown in Figure 4(a).
- In a *multipoint-to-multipoint* EVC, two or more UNIs are associated. A multipoint-to-multipoint EVC, even with the two UNIs, is allowed in that, unlike a point-to-multipoint EVC, another UNI can be added to the multipoint-to-multipoint EVC. An ingress frame at one of the UNIs can be an egress frame at one or more of the other UNIs. A multipoint-to-multipoint EVC is shown in Figure 4(b). Broadcast and multicast frames are delivered to all UNIs except the ingress UNI.
- In a *point-to-multipoint* EVC, two or more UNIs are associated. One UNI is designated as the *root* UNI, and the remaining UNIs are designated as *leaf* UNIs. An ingress frame at the root UNI can be an egress frame at one or more of the leaf UNIs. An ingress frame at a leaf UNI can be an egress frame at only the root UNI. In other words, frames cannot be exchanged among leaf UNIs. Inclusion of the P2MP EVC is for the sake of completeness. Ethernet WAN services described here do not depend on P2MP EVCs. A point-to-multipoint Ethernet Virtual Connection is shown in Figure 4(c). Broadcast and multicast frames that ingress at the root UNI are delivered to all leaf UNIs.

Figure 4
EVC Types





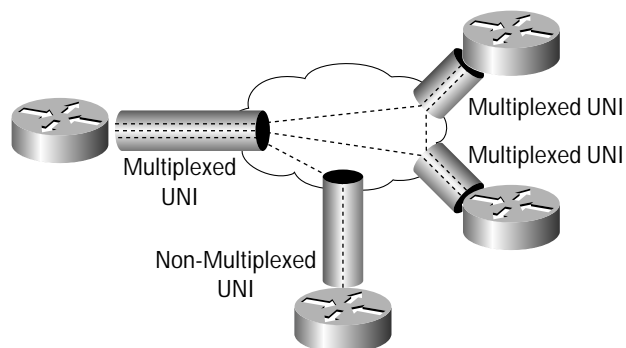
Service Enabling Attributes

There are a number of service attributes that may apply to a particular EVC or to a UNI in an EVC. This sub-section focuses on key service attributes that help establish basic characteristics of Ethernet WAN services. Service attributes that further enhance Ethernet WAN services are discussed in a later section.

Service Multiplexing

Service multiplexing allows a UNI to be in multiple EVCs. Such a UNI is referred to as a *service-multiplexed UNI*. When a UNI is in a single EVC, it is referred to as a *non-multiplexed UNI*. The mixing of multiplexed and non-multiplexed UNIs in a single EVC may be constrained by the additional features of the EVC. Figure 5 shows an example of multiplexed and non-multiplexed UNIs. A service multiplexed UNI is of significant value to a service user that wants to use a single UNI to reach multiple sites because it allows the efficient use of customer equipment ports and reduces the number of UNIs that need to be purchased and managed. Multiplexing point-to-point virtual connections on a single physical interface is one of the values of Frame Relay. That same value applies here.

Figure 5
Service Multiplexed and Non-Multiplexed UNIs



A VLAN configured on the customer equipment port that is connected to the UNI is referred to as the customer equipment-VLAN (CE-VLAN). In the case that the customer equipment is a router, the CE-VLAN is a sub-interface on the port connected to the UNI. At each UNI there is a mapping between CE-VLANs and EVCs called the CE-VLAN/EVC Map. In the simple case, when the Bundling feature (see below) is not invoked, exactly one CE-VLAN can be mapped to an EVC. Its 802.1Q tag identifies the CE-VLAN for a tagged frame whereas an untagged frame identifies a null (or “4097th”) CE-VLAN.

Table 2 is an example of a CE-VLAN/EVC Map. In this example, an ingress frame with CE-VLAN Tag 47 is transported according to the properties and features of EVC₁. An untagged ingress frame is transported according to the properties and features of EVC₃. An egress frame coming from EVC₂ is given CE-VLAN Tag 1343.

In general, it is necessary for the service user and the service provider to agree upon the CE-VLAN/EVC map at the UNI. One way to implement this is to have the service provider dictate the mapping. This is what is done with the mapping between *data link connection identifiers* (DLCIs) and permanent virtual connections (PVCs) for Frame Relay.



Table 2 Example of a CE-VLAN/EVC Map

CE-VLAN	EVC
Tag = 47	EVC ₁
Tag = 1343	EVC ₂
Untagged	EVC ₃

VLAN Transparency

In an EVC with *VLAN Transparency*, the CE-VLAN of an egress frame is always identical to the CE-VLAN on the corresponding ingress frame. An obvious benefit of the VLAN Transparency feature is enhanced operational simplicity; for service users using IEEE 802.1Q bridges, the feature obviates the task of renumbering VLANs in different corporate campuses. Formally, an EVC with the VLAN Transparency feature has the following two properties:

1. The CE-VLAN/EVC Map for the EVC is identical at all UNIs in the EVC.
2. The CE-VLAN tag (including untagged) of an egress frame is always identical to the CE-VLAN tag on the corresponding ingress frame. In other words, the service provider network does not change the tag value. When an EVC has the Bundling feature (see below), Property 1 does not imply Property 2.

Table 3 shows an example of three point-to-point EVCs forming a full mesh among three UNIs. In this example, EVC₁ and EVC₂ have VLAN Transparency while EVC₃ does not.

Table 3 Example of CE-VLAN/EVC Maps for a Full Mesh of EVCs Among Three UNIs

UNI A		UNI B		UNI C	
CE-VLAN	EVC	CE-VLAN	EVC	CE-VLAN	EVC
Tag = 47	EVC ₁	Tag = 47	EVC ₁	Untagged	EVC ₂
Tag = 113	EVC ₃	Untagged	EVC ₂	Tag = 47	EVC ₃

Bundling

When an EVC has the *Bundling* feature, more than one CE-VLAN can map to an EVC at a UNI. Table 4 shows an example of Bundling. In this example, EVC₁ has the bundling feature (and the VLAN Transparency feature). It is difficult to imagine how an EVC with Bundling, but not VLAN Transparency, could be used. *All-to-one Bundling* is a special case of Bundling where all CE-VLANs map to a single EVC at the UNI. A common approach to implementing All-to-One Bundling is through 802.1Q tag stacking where an additional 802.1Q tag is appended to the MAC header of an already tagged frame. Tag stacking is also referred to as 802.1Q tunneling, 802.1Q encapsulation, or simply Q-in-Q. Hierarchical tagging (for example, more than two tags) is also possible. Note that tag stacking is not a standard, and hence vendor interoperability is not guaranteed.



Table 4 Example of Bundling

UNI A		UNI B		UNI C	
CE-VLAN	EVC	CE-VLAN	EVC	CE-VLAN	EVC
Tag = 47,48,49	EVC ₁	Tag = 47,48,49	EVC ₁	Untagged	EVC ₂
Tag = 113	EVC ₃	Untagged	EVC ₂	Tag = 47	EVC ₃

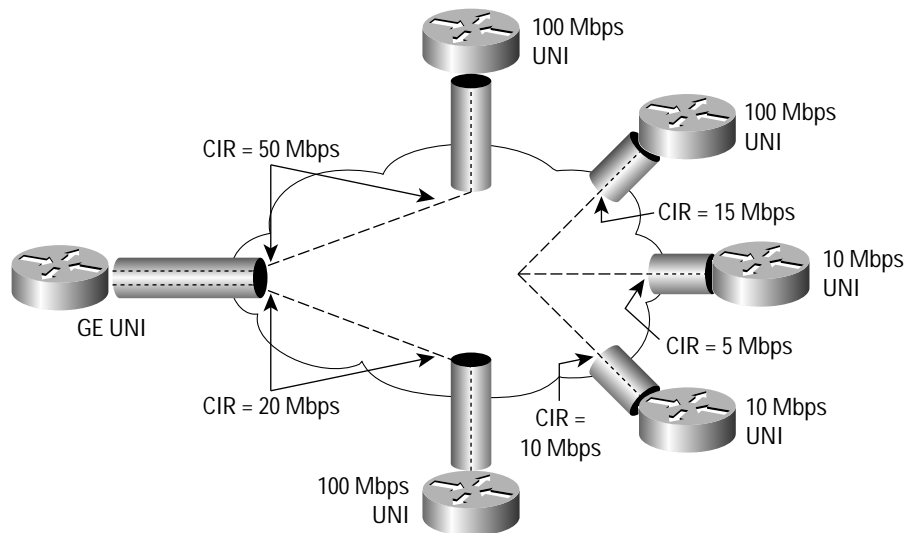
Bandwidth Profile Delivery

A *bandwidth profile* is a limit on the rate at which frames or bytes can traverse the UNI. There can be separate bandwidth profiles for ingress into the network and for egress from the network. The committed information rate/excess information rate (CIR/EIR) of Frame Relay is an example of a bandwidth profile [2].

There are two models for the application of a bandwidth profile:

- *Per UNI*: In this model, a bandwidth profile is applied to all traffic across the UNI.
- *Per Connection*: In this model, a bandwidth profile is applied to all traffic for an instance of an EVC at a UNI. Thus, if a UNI had five Ethernet virtual connections, there could be five ingress bandwidth profiles and five egress bandwidth profiles, one for each connection. Figure 6 illustrates per-connection bandwidth profiles.

Figure 6
Bandwidth Profiles for Ethernet WAN Services





The definition of a bandwidth profile includes the action taken when a frame violates the profile. There are several possible actions:

- Discard the frame.
- Mark the frame for priority discard. This is analogous to the setting of the *cell loss priority (CLP)* bit in ATM or to the setting of the *discard eligibility* bit in Frame Relay.
- Smooth the traffic. This is done at some point in the network, probably near the ingress UNI, by buffering the offending frames and holding them briefly before forwarding. This is done in such a way that the resulting flow of frames leaving the smoothing point complies with the bandwidth profile.

Layer 2 Control Protocol Handling

There are a significant number of Layer 2 protocols used for various control purposes in Ethernet. Several well-known, standardized protocols are listed in Table 5. It is important that Ethernet WANs be able to process this information effectively. This capability will be especially important for service users who choose to deploy 802.1Q bridges (rather than routers) as customer equipment. There are two ways of handling a Layer 2 Control Protocol (L2CP):

- *Control protocols handled on a UNI basis:* Each such protocol is handled on a UNI basis and is independent of the EVCs on the UNI. The network side of the UNI either participates in the protocol as a peer or discards frames received and never sources those frames. Examples of control protocols handled on a UNI basis are IEEE 802.3x and IEEE 802.1x protocols (see Table 5).
- *Control protocols handled on an EVC basis:* Under some conditions, it is desirable to carry Layer 2 control protocols across the service provider network. This is called Layer 2 control protocol *tunneling* because the frame must be passed through the switches in the service provider network without being processed by those switches. In some cases, it will be useful to examine the contents of control frames such as BPDUs but this is not the same as processing them to determine a spanning tree. For example, tunneling of spanning tree BPDUs is desirable when the service user is attaching bridges to all UNIs.

Table 5 Several Standardized Layer 2 Control Protocols

Layer 2 Control Protocol	Destination MAC Address
IEEE 802.1D, 802.1s, and 802.1w Spanning Tree Protocols (STP, RSTP, and MSTP)	01-80-C2-00-00-00
IEEE 802.3x Flow Control	01-80-C2-00-00-01
IEEE 802.3ad Link Aggregation Control Protocol (LACP)	01-80-C2-00-00-02
IEEE 802.1x Port Authentication	01-80-C2-00-00-03
Generic Attribute Registration Protocol (GARP)	01-80-C2-00-00-2X



Technical Descriptions of Canonical Ethernet WAN Services

This section provides a technical description of canonical Ethernet WAN services. The following service-enabling attributes are used to describe ERS, ERMS, EWS, and EMS: service multiplexing, VLAN transparency, L2CP handling, and all-to-one bundling. See the previous section for their detailed descriptions.

ERS is the Ethernet analog of point-to-point Frame Relay service. Here, routers or hosts are used as customer equipment devices to establish point-to-point connections between sites. Instead of data link connection identifier (DLCI) used in Frame Relay, ERS uses VLAN IDs to identify different point-to-point Ethernet connections. Like Frame Relay, service multiplexing is key in ERS as multiple point-to-point connections are provisioned on a single physical interface. Similarly, like the DLCI coordination between service provider and service users, coordination of VLAN IDs is required for ERS. Layer 2 control protocols are either peered or discarded by the network-side UNI. In particular, all BPDUs are dropped, as router customer equipment does not send any BPDUs. Two typical uses of ERS are:

- *Hub-and-spoke enterprise connectivity:* At the hub location, a single high-speed UNI is provisioned and multiple instances of ERS are service multiplexed, one ERS to each branch location. This saves the use of multiple physical ports on the customer equipment at the hub and avoids managing and paying for multiple UNIs at the hub location.
- *Internet service provider (ISP) to customer connectivity:* An ISP can sell Internet access in significant bandwidth increments, for example, 5–25 Mbps, to many customers. Rather than consume a Fast Ethernet port on the ISP router for each customer and underuse each port, multiple instances of ERS are service multiplexed via a Gigabit Ethernet UNI with each ERS going to a different ISP customer. The benefit to the ISP from the improvement in router efficiency is substantial.

ERMS extends the ERS by allowing multipoint-to-multipoint EVCs between multiple sites. Here routers or hosts are used as CE devices. Like ERS, Service Multiplexing is a key service feature supported by ERMS, and this necessitates VLAN co-ordination between service provider and service user. Layer-2 control protocols are peered or discarded by the network-side ERMS UNI. The benefits of ERMS are:

- A single high-speed physical interface is used for multiple logical interfaces, thus avoiding the management and cost of multiple UNIs.
- Each multipoint EVC alleviates the need to manage multiple point-to-point EVCs (as would be the case in ERS).
- For IP connectivity, IP addresses are more efficiently utilized with a multipoint EVC as it requires only a single IP subnet (whereas, in ERS, a subnet may be needed for each point-to-point EVC).

ERMS is useful in situations such as the following:

- An enterprise hub location can connect to its Ethernet sites with a single multipoint EVC and also connect its Frame Relay sites with point-to-point EVCs over a single Ethernet UNI. Note that, in this case, the service provider must support service interworking between Ethernet and Frame Relay (see the discussion below on Frame Relay service interworking).
- An enterprise extranet can use multipoint-to-multipoint and point-to-point EVCs to connect its partners and suppliers over a single Ethernet UNI. For instance, a car manufacturer can use a multipoint-to-multipoint EVC to connect its dealers and point-to-point EVCs to connect each of its suppliers. If multiple sites of the same supplier need to be connected, a multipoint-to-multipoint EVC can be employed as well.



- An enterprise can use a multiplexed Ethernet UNI to connect its multiple branch offices using a multipoint-to-multipoint EVC and to connect to its ISP using a point-to-point EVC. Note that special consideration to network security may be warranted for traffic to and from the Internet.

EWS is the Ethernet analog of private line (or private wire) service. Customer equipment can be a router, bridge, or a host. Frames arrive at the destination UNI unaltered. All-to-One Bundling and VLAN Transparency are employed at the expense of service multiplexing. In the case of a bridge CE, the customer's BPDUs are tunneled through the service provider's Layer 2 network. EWS is ideal for connecting a few sites together and the coordination of CE-VLAN values is not a concern. In the case where the service user wants to build a flat network, EWS can connect two customer equipment bridges and allow the VLAN architecture of the service user to be extended between sites.

EMS is the WAN analog of the multipoint Ethernet LAN capability. A single multipoint EVC provides connectivity to multiple sites. Customer equipment can be a router, bridge, or a host. Traffic from the originating CE reaches one or more destination CE unaltered. Like EWS, All-to-One Bundling and VLAN Transparency are employed at the expense of service multiplexing. In the case of a bridge CE, a customer's BPDUs are replicated and tunneled through the service provider's Layer 2 network to all UNIs other than the ingress UNI. EMS is a good way to connect several campuses of the same enterprise with CE routers. CE bridges can also be used but at the possible risk of suffering the adverse effects of a large Layer 2 network.

It should be noted that there are differences in point-to-point and multipoint-to-multipoint Ethernet services. For instance, in the case of a multipoint-to-multipoint EVC, all UNIs can communicate with each other as peers. Defining the CIR for each UNI in this case may require an *a priori* understanding of the unicast and broadcast/multicast traffic patterns in that multipoint-to-multipoint EVC.

Table 6 provides a summary of technical descriptions of the four Ethernet WAN services.

Table 6 Technical Description of Canonical Ethernet WAN Services

Ethernet Service	Customer Equipment	EVC	L2CP Handling	Service Multiplexing	VLAN Transparency	Bundling (All-to-One)
ERS	Router, host	Point-to-point	Discard or peer at UNI	Yes	No	No
ERMS	Router, host	Multipoint-to-multipoint	Discard or peer at UNI	Yes	No	No
EWS	Router, bridge, host	Point-to-point	Discard, peer or tunnel	No	Yes	Yes
EMS	Router, bridge, host	Multipoint	Discard, peer or tunnel	No	Yes	Yes



Service Enhancing Attributes

Service attributes that further enhance Ethernet WAN services are described in this section.

MAC Address Limiting

When frames are delivered based on *learned* MAC addresses, each learned MAC address consumes service provider network resources. Therefore, it may be desirable to ensure that the number of MAC addresses learned for each customer is limited. At each instance of an EVC on a UNI, the source MAC addresses in ingress frames is the measured quantity. These are called *Observed MAC Addresses*. The MAC Address Limit is defined for each instance of an EVC at a UNI. At any given time, there is a list of *Allowed MAC Addresses*. An ingress frame with an observed MAC address that is not in the list may be discarded depending on the method of constructing the list and the number of entries in the list.

Service Differentiation

Service differentiation provides frames with differentiated levels of service. The difference in service level will be reflected in the UNI-to-UNI latency, latency jitter, and loss. Several classes of service may be offered by a service provider; for instance, one for best-effort data, another for mission-critical data, and a third for voice or video traffic.

There are several ways a frame receives differentiated service:

- Classification by UNI: All frames carried by a UNI are treated with an assigned class of service.
- Classification by EVC: All frames carried by an EVC are treated with an assigned class of service.
- Classification by 802.1p bits: Each frame is classified according to the priority bits in the CE-VLAN tag on the UNI, including a specific class for untagged frames. If the service provider wants to charge differently for different classes of service, then this method of classification will probably require a bandwidth profile for each instance of the triplet <EVC/UNI/.1p>.
- Classification by differentiated services code point (DSCP) bits: When frames carry IP packets, the classification is done by the value of the DSCP.

Security

No amount of security makes the network or the service unbreakable. Depending on the *trust* level, different levels of security are applied to different portions of the network. Of course, the perception of trust varies among different parties. For instance, from a service provider's perspective, a customer equipment device is considered *untrusted* while the service provider's WAN core network is considered *trusted*. A service provider-managed network device placed in a building may be considered *semi-trusted*. On the other hand, a service user with high security requirements may consider the entire WAN network untrusted.

Generally speaking, there are three main aspects of security for WAN services: Denial-of-Service (DoS) attacks, Theft of Service, and Loss of Privacy. Also, an act to violate security may be intentional or unintentional. There are many ways to address the different aspects of security, each providing a varying degree of service protection. For instance, MAC address limits may be enabled to prevent DoS attacks that flood frames with unknown MAC addresses into the WAN network. Untrusted customer equipment devices should be authenticated, implicitly or explicitly, to reduce the risk of theft of service and loss of privacy. Physical security (for example, locked cages) may also be necessary in certain situations.



Service Resiliency

Service resiliency ensures the availability of a service in the event of an equipment or link failure. A service provider may deploy service-resilient WAN architectures to differentiate its service offerings. There are several ways to provide resiliency to service users and these include:

- Customer equipment dual-homed to a single service provider switch (link aggregation may be enabled in this case to provide service resiliency.)
- Customer equipment dual-homed to two different service provider switches at a single site
- Customer equipment dual-homed to two different service provider switches at two different sites
- Customer equipment dual-homed to two different service providers
- Redundant customer equipment, each single- or dual-homed to the service provider network

When a bridge customer equipment is dual homed (or redundant bridge customer equipments are connected) to a service provider network, BPDU tunneling may be required to ensure loop-free service user topology.

Service Provisioning and Monitoring

Service provisioning includes service activation, service modification, and service monitoring. Service users' familiarity with Ethernet provides a unique opportunity to service providers to deploy a subscriber self-provisioning system. This eliminates truck rolls and calls to the customer care center and thus provides significant cost savings to service providers and accuracy and responsiveness to the service user.

Service monitoring requires additional control-plane intelligence as discussed below:

- *Ethernet Link Management Interface (LMI)*: This would be a new protocol between the service provider network and the customer equipment that would let the customer equipment learn about EVCs available at the UNI. Ethernet LMI will require new software but not new hardware for the customer equipment. A customer equipment with Ethernet LMI enabled could learn, for example, the CE-VLAN/EVC map as well as other EVC properties such as the ingress bandwidth profile and the class of service.
- *EVC integrity assurance*: The goal of this attribute is to allow a customer equipment to quickly learn when an EVC is no longer carrying frames. For example, in the case of a point-to-point EVC with non-multiplexed UNIs, this indication could be accomplished by bringing the physical layer down (hard) at both UNIs when the service provider network detects the broken EVC.

Ethernet Interworking Network-to-Network Interface

Ethernet Interworking NNI is important in situations where a carrier serves another carrier to provide end-to-end Ethernet WAN service. For instance, an Ethernet service provider that has a presence in only metropolitan areas must subscribe to the services of a national carrier to provide metro-to-metro Ethernet WAN services.

Ethernet Interworking-NNI would be a new protocol. Several issues to consider when specifying Ethernet Interworking-NNI are:

- Must support multiple Ethernet WAN services
- Must be transparent to the service user
- Must ensure loop-free topology (even with redundant Ethernet Interworking-NNIs)



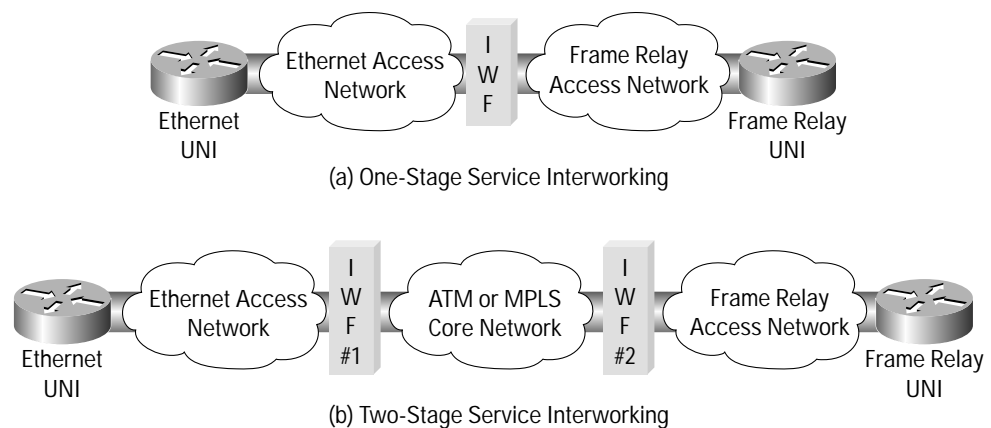
Frame Relay Service Internetworking

Frame Relay is a very popular WAN service. Many enterprises have subscribed to Frame Relay service to connect remote sites with corporate headquarters (HQ). These enterprises are not expected to migrate to Ethernet WAN services quickly. A powerful strategy for an enterprise is to migrate its Frame Relay service at the HQ site to high-speed Ethernet WAN service while maintaining Frame Relay connectivity to the remote sites. Also, service providers have significant investment in their Frame Relay networks and are not expected to carry out a large-scale network build for Ethernet WAN services overnight. Hence, *service internetworking* between Frame Relay and Ethernet is critical for the success of Ethernet WAN services, providing investment protection in Frame Relay and allowing a phased migration to Ethernet.

Service interworking between Ethernet and Frame Relay is depicted in Figure 7(a). The Internet Engineering Task Force (IETF) has standardized, in RFC 2427 [3], encapsulation mechanisms to allow transmission of multiple data networking protocols over Frame Relay. Two encapsulation mechanisms are specified: bridged encapsulation and routed encapsulation. Bridged encapsulation is used to encapsulate Ethernet frames while routed encapsulation is used to encapsulate routed (for example, IP) packets. In a typical application, a CE router with a Frame Relay interface is configured using RFC 2427 routed encapsulation (to encapsulate IP packets). If the service interworking function (IWF) allows interworking of routed encapsulation, no configuration change may be needed on customer equipment routers facing the Frame Relay network. This alleviates any service outage concerns due to equipment misconfiguration and also enables faster migration to Ethernet WAN services. Both of these are significant benefits to the service user and the service provider.

Often, the Ethernet and Frame Relay access networks are connected by an ATM or Multiprotocol Label Switching (MPLS) core network. In this case, two IWFs are required (see Figure 7(b)): Ethernet to ATM (MPLS) and ATM (MPLS) to Frame Relay. Ethernet to ATM interworking is achieved using RFC 2684 [4]. ATM to Frame Relay internetworking is achieved via the Frame Relay Forum specification FRF 8.1 [5]. Since service providers are also rolling out MPLS cores, it follows that there will be multistage interworking over MPLS as well.

Figure 7
Service Interworking Between Ethernet and Frame Relay





Ethernet WAN Network Architectures

This section presents and compares several network architectures for Ethernet WAN services. These architectures use either Ethernet-only or Ethernet plus Layer 3 technology.

Configuration of Customer Equipment and UNI Ports

Regardless of the network architecture deployed, the end-user always accesses Ethernet WAN services via an Ethernet connection to the service provider network, as shown in Figure 8. VLANs of tagged frames on the access link (the customer equipment-to-provider edge link of Figure 8) are referred to as *CE-VLANs* whereas VLANs of frames within the service provider network (on the PE-PE links of Figure 8) are referred to as *PE-VLANs*. Table 7 lists the possible configurations for CE and PE UNI ports. A CE port may be configured to send untagged frames or 802.1Q frames to the UNI. Similarly, a PE's UNI port may be an access port, .1Q trunk port, or Q-in-Q port, as discussed below:

- If the UNI port is an access port, all received frames are untagged and a UNI-configured VLAN is assigned. In this case, the CE-VLAN is a null (or “4097th”) value and the PE-VLAN is the UNI VLAN.
- If the UNI port is a .1Q trunk port, it accepts CE's tagged frames. Each tagged frame's PE-VLAN is the same as its CE-VLAN. This assumes that VLAN *translation* is not enabled at the UNI. Also, untagged Layer-2 control protocol frames received from the customer equipment are either discarded or peered at the UNI.
- If the UNI port is a Q-in-Q port, it accepts tagged or untagged frames from the customer equipment. It assigns a UNI-configured tag, PE-VLAN, to all received frames. If the received frames are already tagged with CE-VLANs, then they become *doubly* tagged at the UNI (with CE-VLAN as the inner tag, and the common PE-VLAN as the outer tag). The Q-in-Q feature is critical for EWS and EMS services as it meets the VLAN transparency and All-to-One Bundling service requirements.

Figure 8
Single-PE Ethernet Network Architecture (T: trunk port, A: access port, Q-in-Q: Q-in-Q port)

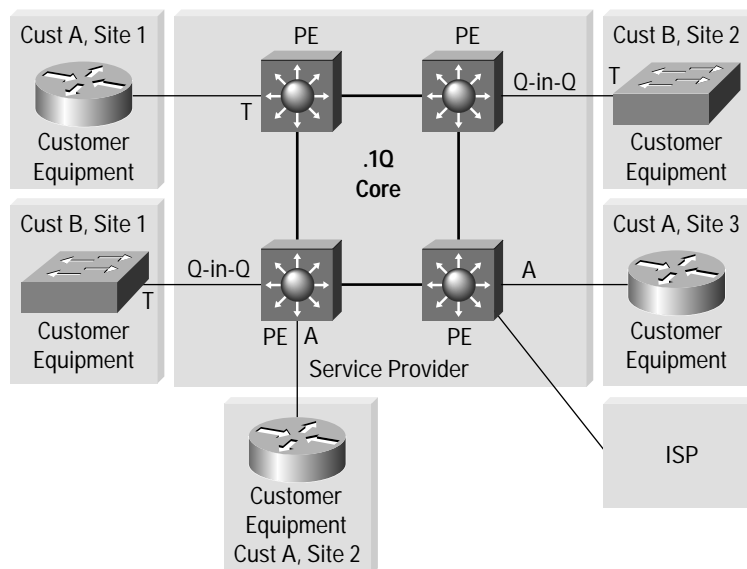




Table 7 Customer Equipment and UNI Port Configuration Options

Customer Equipment Port Configuration	UNI Port Configuration	Supported Service	Comments
Access (untagged)	Access	ERS, ERMS	All received frames are single tagged with the UNI VLAN (for example, PE-VLAN = UNI VLAN and there is no CE-VLAN)
802.1Q	802.1Q (trunk)	ERS, ERMS	All .1Q received frames maintain their original tags (for example, PE-VLAN = CE-VLAN if there is no VLAN translation at the UNI)
Access (untagged)	Q-in-Q	EWS, EMS	All received frames are single tagged with the UNI VLAN (PE-VLAN); there is no CE-VLAN
802.1Q	Q-in-Q	EWS, EMS	Frames are double tagged with CE-VLAN as the inner tag and fixed UNI VLAN (PE-VLAN) as the outer tag

Ethernet-Only Network Architectures

Ethernet-only networks are popular for offering Ethernet WAN services. This is primarily due to the fact that today's Layer 2 switches provide a rich set of service features at an attractive price. Figure 8 depicts a simple network architecture in which multiple Ethernet switches are connected in a sparse mesh topology. This architecture supports the four canonical Ethernet WAN services, namely, ERS, ERMS, EWS, and EMS. Specifically, customer A is subscribed to the ERS service with service multiplexing at the site 1 UNI. Customer B is subscribed to the EWS service. Each service provider switch is designated as a PE switch and is placed at a metro point-of-presence (POP) or at a mini-POP. All network facing links (PE-PE links of Figure 8) are configured as .1Q trunks.

In many instances, it is necessary to place provider switches at or close to customer locations, for instance, in the basement of high-rise buildings. In such cases, a distributed-PE *hub-and-spoke* topology (Figure 9) or distributed-PE *ring* topology (Figure 10) is deployed. The *customer located equipment* (CLE) is referred to as the PE-CLE switch and each POP switch is referred to as the PE-POP switch. Typically, low-cost and lower port-count switches are deployed as PE-CLEs at customer locations.



Figure 9
Distributed-PE Ethernet Hub-and-Spoke Network Architecture

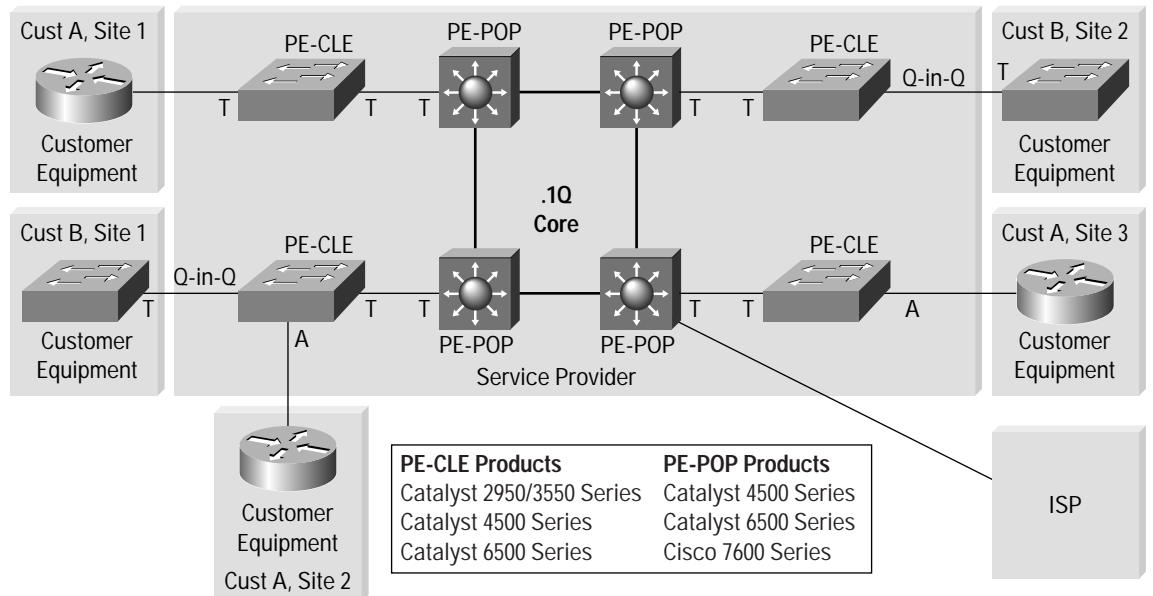
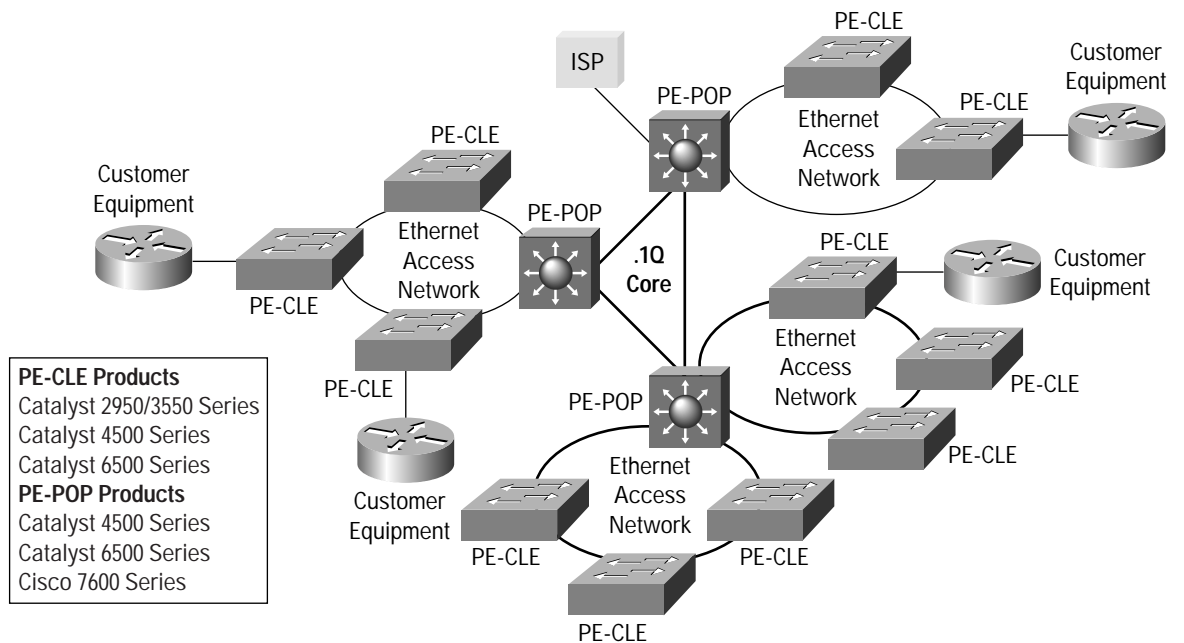


Figure 10
Distributed-PE Ethernet Ring Network Architecture





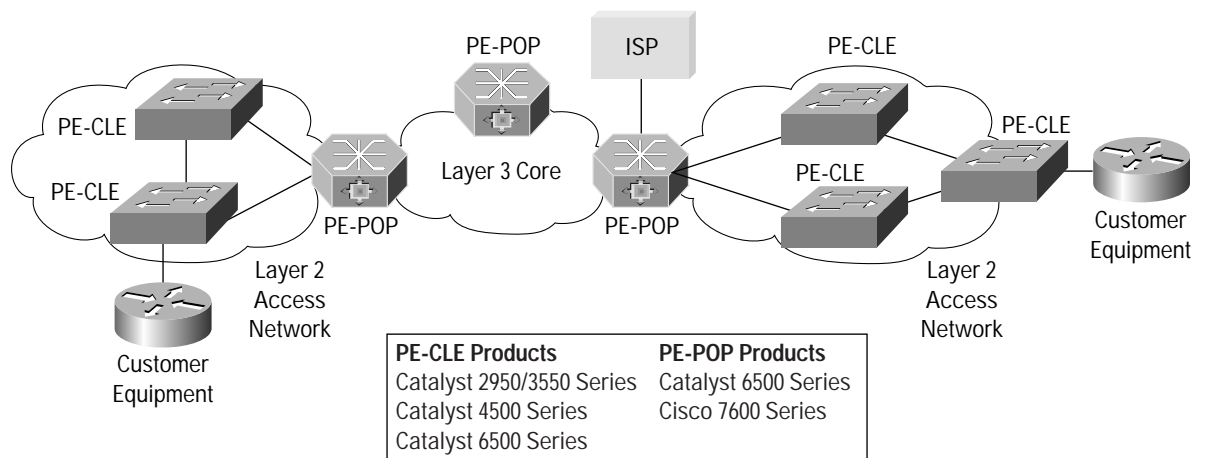
Hybrid Network Architectures

An Ethernet WAN network based on a single Layer 2 domain is ideal for limited networks. It supports a modest number of customers due to the restriction of 4096 PE-VLANs per Layer 2 domain. Also, the stability of spanning tree protocols in a large multi-metro network has not been proven to the satisfaction of some service providers.

To improve network scalability, hybrid architectures have emerged, consisting of multiple Layer 2 domains that are connected via a Layer 3 core (such as an IP or MPLS core). This is illustrated in Figure 11. In a hybrid network, Ethernet-specific concerns (for example, VLAN IDs, spanning tree protocols) are limited to the individual Layer 2 domain. Such a network is inherently scalable because, when a Layer 2 domain becomes large, it can be further partitioned into multiple smaller Layer 2 domains. Ethernet frames between two Layer 2 domains are tunneled through the Layer 3 domain as described below:

- In the case of an IP core, *Layer 2 Tunneling Protocol Version 3 (L2TPv3)* is used to tunnel Ethernet frames [6].
- In the case of an MPLS core, *Ethernet over MPLS (EoMPLS)* signaling and encapsulation protocols are used to tunnel Ethernet frames. Both point-to-point and multipoint-to-multipoint EVCs can be supported. Point-to-point EoMPLS is specified by the IETF Internet Drafts [7],[8], and multipoint-to-multipoint EoMPLS is specified by the Internet Draft RFCs [9],[10].

Figure 11
Multiple Layer 2 (Ethernet) Domains with a Layer 3 Core



In both cases, the VLAN tag may change when a frame is tunneled. This allows independent administration of the VLAN tag values in each Layer 2 domain. It also allows VLAN tag values to be reused across Layer 2 domains. The result is that more EVCs can be supported as more Layer 2 domains are deployed. For example, if all EVCs are point-to-point, then a network with n Layer 2 domains can support at least $4096n/2$ EVCs.

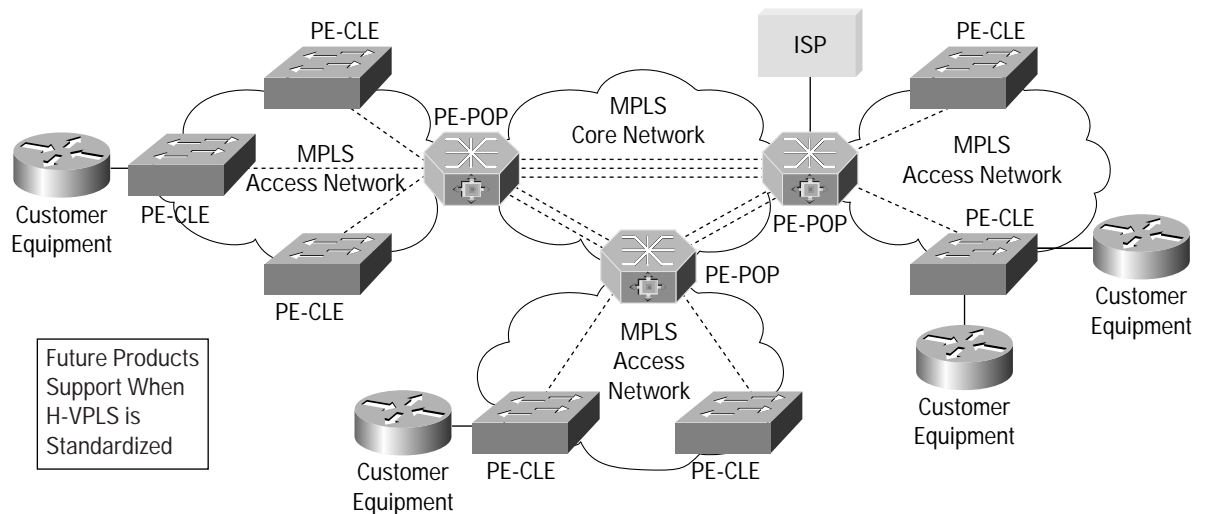
Note that MPLS-based Ethernet WAN network architectures are being standardized in the IETF under the nomenclature *Virtual Private Wire Service (VPWS)* and *Virtual Private LAN Service (VPLS)* [9]. VPWS is defined to support point-to-point Ethernet WAN services whereas VPLS is defined to support multipoint services.



The IETF is also discussing the standardization of MPLS-based distributed-PE architecture. This architecture is referred to as the *Hierarchical VPLS (H-VPLS)* architecture [10]. Two key design options being proposed are:

- *Ethernet access network*: Access networks (PE-CLE-to-PE-POP) are .1Q networks and the core network (PE-POP-to-PE-POP) is MPLS. This is illustrated in Figure 11 with Ethernet Access Network and MPLS core.
- *MPLS access network*: Access networks and the core network are both MPLS. This is shown in Figure 12.

Figure 12
Hierarchical VPLS (H-VPLS) Architecture with MPLS Access Network



The choice between Ethernet access network architecture and MPLS access network architecture depends on many factors, including equipment price and ease of service management. Supporting MPLS on an Ethernet switch increases data-plane and control-plane complexity of the switch, thus potentially increasing its price. Because PE-CLE switches are expected to be deployed in volume, it appears that the H-VPLS architecture with Ethernet access network should have a near-term pricing advantage over the H-VPLS architecture with MPLS access network. On the other hand, the H-VPLS architecture with MPLS access network provides a unified (MPLS-based) control plane and is attractive from the perspective of service management. It may also be perceived as having greater scalability over its Ethernet counterpart.



Table 8 summarizes the various Ethernet WAN architectures.

Table 8 Summary of Several Ethernet WAN Architectures

Topology	Access Link	Access Network	Core Network	EVC Type	Comments
Ethernet end-to-end	Ethernet	Ethernet	Ethernet	Point-to-point, multipoint-to-multipoint	Uses 802.1Q and/or Q-in-Q in the core
IP core with Ethernet access network	Ethernet	Ethernet	IP	Point-to-point, multipoint-to-multipoint	PE-POP uses L2TPv3 to tunnel Ethernet over IP core
H-VPLS with Ethernet access network	Ethernet	Ethernet	MPLS	Point-to-point, multipoint-to-multipoint	PE-POP uses EoMPLS to tunnel Ethernet over MPLS core
H-VPLS with MPLS access network	Ethernet	MPLS	MPLS	Point-to-point, multipoint-to-multipoint	PE-CLE uses EoMPLS to tunnel Ethernet over MPLS; provides unified control plane

References

- [1] IEEE 802 LAN/MAN Standards Committee, www.ieee802.org.
- [2] International Telecommunications Union, Recommendation I.370, *Congestion Management for the ISDN Frame Relaying Bearer Service*, 1999.
- [3] IETF RFC 2427, *Multiprotocol Interconnect over Frame Relay* (obsoletes RFC 1490), www.ietf.org/rfc/rfc2427.txt?number=2427
- [4] IETF RFC 2684, *Multiprotocol Encapsulation over ATM Adaptation Layer 5* (obsoletes RFC 1483), www.ietf.org/rfc/rfc2684.txt?number=2684
- [5] The Frame Relay Forum (www.frforum.com), FRF 8.1, *Frame Relay/ATM PVC Service Interworking Implementation Agreement*, February 2000.
- [6] IETF Internet Draft, *Layer Two Tunneling Protocol (Version 3) "L2TPv3"*, www.ietf.org/internet-drafts/draft-ietf-l2tpext-l2tp-base-03.txt
- [7] IETF Internet Draft, *Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks*, www.ietf.org/internet-drafts/draft-martini-l2circuit-encap-mpls-04.txt
- [8] IETF Internet Draft, *Transport of Layer 2 Frames Over MPLS*, www.ietf.org/internet-drafts/draft-martini-l2circuit-trans-mpls-10.txt
- [9] IETF Internet Draft, *PPVPN L2 Framework*, www.ietf.org/internet-drafts/draft-ietf-ppvpn-l2-framework-01.txt
- [10] IETF Internet Draft, *Virtual Private LAN Services over MPLS*, www.ietf.org/internet-drafts/draft-lasserre-vkompella-ppvpn-vpls-02.txt
- [11] Prashant Gandhi and Bob Klessig, "Metro Ethernet WAN Services and Architectures", International Engineering Consortium's *Annual Review of Communications*, June 2003, www.iec.org



Acronyms

Table 9 Acronyms

Acronym	Definition
802.1D	Spanning Tree Protocol (STP)
802.1Q	16-bit tag used to extend the Ethernet header (includes 12-bit VLAN ID and 3-bit 802.1p class of service bits)
802.1p	Layer-2 Class of Service or priority bits (within 802.1Q tag)
802.1s	Multiple Spanning Tree Protocol (MSTP)
802.1w	Rapid Spanning Tree Protocol (RTSP)
802.1x	Port Authentication
802.3ad	Link Aggregation Control Protocol (LACP)
802.3x	Flow control
ATM	Asynchronous Transfer Mode
BPDU	Bridge Protocol Data Unit
CLP	Cell Loss Priority
DE	Discard Eligible
DLCI	Data Link Connection Identifier
DoS	Denial of Service
DS1	Digital Signal 1 (service based on T1 at 1.5 Mb/s)
DS3	Digital Signal 3 (service based on T3 at 45 Mb/s)
E1	European Signal 1 (service at 2.048 Mb/s)
E3	European Signal 3 (service at 32 Mb/s)
EI-NNI	Ethernet Interworking Network-to-Network Interface
EVC	Ethernet Virtual Connection
FR	Frame Relay
FRF	Frame Relay Forum
H-VPLS	Hierarchical Virtual Private LAN Service
IETF	Internet Engineering Task Force
L2CP	Layer-2 Control Protocol
IP	Internet Protocol
LACP	Link Aggregation Control Protocol (802.3ad)
LAN	Local Area Network
LMI	Link Management Interface
MP2MP	Multipoint-to-multipoint
MSTP	Multiple Spanning Tree Protocol (802.1s)
P2P	Point-to-point
PL	Private Line
RFC	Request for Comments
RTSP	Rapid Spanning Tree Protocol (802.1w)
SDH	Synchronous Digital Hierarchy
SONET	Synchronous Optical Network
SP	Service Provider
STP	Spanning Tree Protocol (802.1D)
UNI	User-to-Network Interface
VLAN	Virtual Local Area Network
VPLS	Virtual Private LAN Service
VPWS	Virtual Private Wire Service
WAN	Wide Area Network



Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters
Cisco Systems, Inc.
Capital Tower
168 Robinson Road
#22-01 to #29-01
Singapore 068912
www.cisco.com
Tel: +65 317 7777
Fax: +65 317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the
Cisco Web site at www.cisco.com/go/offices

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia
Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland
Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland
Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden
Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

All contents are Copyright © 1992–2002, Cisco Systems, Inc. All rights reserved. CCIP, the Cisco Arrow logo, the Cisco *Powered* Network mark, the Cisco Systems Verified logo, Cisco Unity, Follow Me Browsing, FormShare, iQ Breakthrough, iQ Expertise, iQ FastTrack, the iQ logo, iQ Net Readiness Scorecard, Networking Academy, ScriptShare, SMARTnet, TransPath, and Voice LAN are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, Discover All That's Possible, The Fastest Way to Increase Your Internet Quotient, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, LightStream, MGX, MICA, the Networkers logo, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, SlideCast, StrataView Plus, Stratm, SwitchProbe, TeleRouter, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the U.S. and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company.
(0208R) KW/LW3233 12/02