

Metro Ethernet Services – A Technical Overview

Ralph Santitoro

Introduction

This whitepaper provides a comprehensive technical overview of Ethernet services, based on the work (as of April 2003) of the Metro Ethernet Forum (MEF) Technical Committee. The paper is intended to help buyers and users of Ethernet services understand the various types and characteristics of Ethernet services, and to help service providers clearly communicate their service capabilities. Throughout this paper, buyers and users will be collectively referred to as subscribers. This paper will be updated as new work emerges from the MEF Technical Committee.

Background

Metro Ethernet services are now offered by a wide range of service providers. Some providers have extended Ethernet services beyond the metropolitan area and across the wide area. Thousands of subscribers already use Ethernet services and their numbers are growing rapidly. These subscribers have been attracted by the benefits of Ethernet services, including:

- Ease of use
- Cost Effectiveness
- Flexibility

Ease of Use

Ethernet services are provided over a standard, widely available and well-understood Ethernet interface. Virtually all networking equipment and hosts connect to the network using Ethernet so using an Ethernet service to interconnect such devices simplifies network operations, administration, management and provisioning (OAM&P).

Cost Effectiveness

Ethernet services can reduce subscribers' capital expense (CapEx) and operation expense (OpEx) in three ways.

- First, due to its broad usage in almost all networking products, the Ethernet interface itself is inexpensive.
- Second, Ethernet services can often cost less than competing services due to lower equipment, service and operational costs.
- Third, many Ethernet services allow subscribers to add bandwidth more incrementally, e.g., in 1 Mbps increments. This allows subscribers to add

bandwidth as needed so they pay for only what they need.

Flexibility

Many Ethernet services allow subscribers to network their business in ways that are either more complex or impossible with alternative services. For example, a single Ethernet service interface can connect multiple enterprise locations for their Intranet VPNs, connect business partners or suppliers via Extranet VPNs and provide a high speed Internet connection to an Internet Service Provider. With managed Ethernet services, subscribers are also able to add or change bandwidth in minutes instead of days or weeks when using other access network services. Additionally, these changes do not require the subscriber to purchase new equipment and coordinate a visit with a service provider technician.

Many Ethernet services allow subscribers to add bandwidth in minutes or hours instead of weeks or months

What is an Ethernet Service?

All Ethernet services share some common attributes, but there are differences. The basic model for Ethernet services is shown in Figure 1. Ethernet Service is provided by the Metro Ethernet Network (MEN) provider. Customer Equipment (CE) attaches to the network at the User-Network Interface (UNI) using a standard 10Mbps, 100Mbps, 1Gbps or 10Gbps Ethernet interface.




Figure 1 – Basic Model

Note that when discussing subscriber applications, this paper will often refer to the subscriber's network connection as a "site" or "subscriber" connection. However, it is possible to have multiple subscribers (UNIs) connect to the MEN from a single location (site).

Finally, the services are defined from a subscriber-perspective (referred to as "retail" services). Such services can be supported over a variety of transport technologies and protocols in the MEN such as SONET, DWDM, MPLS, GFP, etc. However, from a subscriber-perspective, the network connection at the subscriber side of the UNI is Ethernet.

Ethernet Virtual Connection

One key Ethernet service attribute is the Ethernet Virtual Connection (EVC). An EVC is defined by the MEF as "an association of two or more UNIs", where the UNI is a standard Ethernet interface that is the point of demarcation between the Customer Equipment and service provider's MEN.

In simple terms, an EVC performs two functions:

- Connects two or more subscriber sites (UNIs) enabling the transfer of Ethernet service frames between them.
- Prevents data transfer between subscriber sites that are not part of the same EVC. This capability enables an EVC to provide data privacy and security similar to a Frame Relay or ATM Permanent Virtual Circuit (PVC).

Two basic rules govern delivery of Ethernet frames over an EVC. First, a service frame must never be delivered back to the UNI from which it originated. Second, service frames must be delivered with the Ethernet MAC addresses and frame contents unchanged, i.e., the Ethernet frame remains intact from source to destination(s). Contrast this to a typical routed network where the Ethernet frame headers are removed and discarded.

Based on these characteristics, an EVC can be used to construct a Layer 2 Private Line or Virtual Private Network (VPN).¹

The MEF has defined two types of EVCs.

- Point-to-Point

¹ The term "Layer 2 VPNs" helps distinguish EVCs from "IP VPNs".

- Multipoint-to-Multipoint

... an EVC can be used to construct a Layer 2 Private Line or Virtual Private Network (VPN)

Beyond these common characteristics, Ethernet services may vary in many ways. The rest of this paper discusses different types of Ethernet services and some of the important characteristics that distinguish them from other service offerings.

Ethernet Service Definition Framework

To help subscribers better understand the variations among Ethernet services, the MEF has developed the Ethernet Service Definition Framework. The goals of this framework are to:

1. Define and name common Ethernet Service Types.
2. Define the attributes and associated parameters used to define specific Ethernet Services.




Figure 2: Ethernet Service Definition Framework

The MEF has currently defined two Ethernet Service Types:

- Ethernet Line (E-Line) Service type
 - point-to-point service
- Ethernet LAN (E-LAN) Service type
 - multipoint-to-multipoint service

The service types are really "umbrella" categories, since specific services created from one service type may differ substantially from each other. To fully specify an Ethernet Service, a provider must define the service type and UNI and EVC service attributes associated with the service type. These service attributes can be grouped under the following categories:

- Ethernet Physical Interface
- Traffic Parameters
- Performance Parameters
- Class of Service

- Service Frame Delivery
- VLAN Tag Support
- Service Multiplexing
- Bundling
- Security Filters

Ethernet Service Types

The MEF has defined two basic service types discussed below. Other service types may be defined in the future.

Ethernet Line Service type

The Ethernet Line Service (E-Line Service) provides a point-to-point Ethernet Virtual Connection (EVC) between two UNIs as illustrated in Figure 3. The E-Line Service is used for Ethernet point-to-point connectivity.

In its simplest form, an E-Line Service can provide symmetrical bandwidth for data sent in either direction with no performance assurances, e.g., best effort service between two 10Mbps UNIs. In more sophisticated forms, an E-Line Service may provide a CIR (Committed Information Rate) and associated CBS (Committed Burst Size), EIR (Excess Information Rate) and associated EBS (Excess Burst Size) and delay, jitter, and loss performance assurances between two different speed UNIs.




Figure 3: E-Line Service using Point-to-Point EVC

Service multiplexing of more than one EVC may occur at none, one or both of the UNIs (Refer to the Service Multiplexing section). For example, more than one point-to-point EVC (E-Line Service) may be offered on the same physical port at one of the UNIs.

An E-Line Service can provide point-to-point EVCs between UNIs analogous to using Frame Relay PVCs to interconnect sites as illustrated in Figure 4.




Figure 4: Frame Relay analogy of E-Line Service

An E-Line Service can also provide a point-to-point connection between UNIs analogous to a TDM private line service. Such a service interconnects two UNIs and provides full transparency for service frames between the UNIs such that the service frame’s header and payload are identical at both the source and destination UNI. Such a service would also have some fundamental characteristics such as minimal Frame Delay, Frame Jitter and Frame Loss and no Service Multiplexing, i.e., a separate UNI (physical interface) is required for each EVC as illustrated in Figure 5.

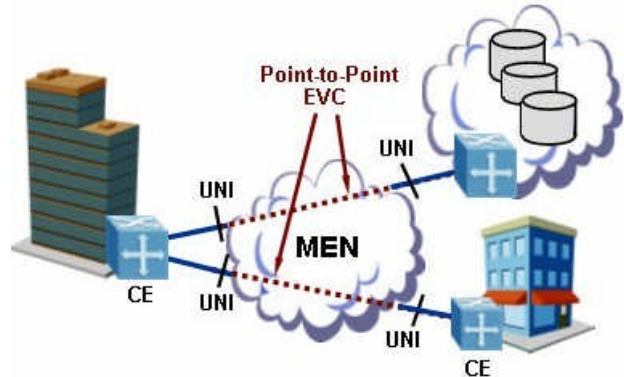


Figure 5: Private line analogy using E-Line Service

In summary, an E-Line Service can be used to construct services analogous to Frame Relay or private leased lines. However, the range of Ethernet bandwidth and connectivity options is much greater.

“.. an E-Line Service can be used to construct services analogous to Frame Relay or private leased line”

Ethernet LAN Service type

The Ethernet LAN Service (E-LAN Service) provides multipoint connectivity, i.e., it may connect two² or more UNIs as illustrated in Figure 6. Subscriber data sent from one UNI can be received at one or more of the other UNIs. Each site (UNI) is connected to a multipoint EVC. As new sites (UNIs) are added, they are connected to the same multipoint EVC thus simplifying provisioning and service activation. From a Subscriber standpoint, an E-LAN Service makes the MEN look like a LAN.

An E-LAN Service can be used to create a broad range of services. In its simplest form, an E-LAN Service can provide a best effort service with no performance assurances. In more sophisticated forms, an E-LAN Service may define a CIR (Committed Information Rate) and associated CBS (Committed Burst Size), EIR (Excess Information Rate) and associated EBS (Excess Burst Size) (refer to Bandwidth Profile section) and delay, jitter, and loss performance assurances for the service.




Figure 6: E-LAN Service using Multipoint EVC

An E-LAN Service may support service multiplexing of EVCs at none, one or more of the UNIs (Refer to Service Multiplexing section). For example, an E-LAN Service (Multipoint-to-Multipoint EVC) and an E-Line Service (Point-to-Point EVC) may be offered at one of the UNIs. In this example, the E-LAN Service may be used to interconnect other subscriber sites while the E-Line Service is used to connect to the Internet with both services offered via EVC service multiplexing at the same UNI.

² Note that an E-LAN Service with only two UNIs (sites) still uses a multipoint EVC but with only 2 UNIs in the multipoint connection. Unlike a Point-to-Point EVC which is limited to 2 UNIs, a multipoint EVC can have additional UNIs added to the EVC.

An E-LAN Service may include a configured CIR, EIR and associated burst sizes as part of the UNI Bandwidth Profile (refer to Bandwidth Profile section). The port speed at each UNI may be different. For example, in Figure 6, UNIs 1, 2 and 3 may each have a 100Mbps Ethernet interface with a 10Mbps CIR. UNI 4 may have a 1Gbps Ethernet interface with a 40Mbps CIR.

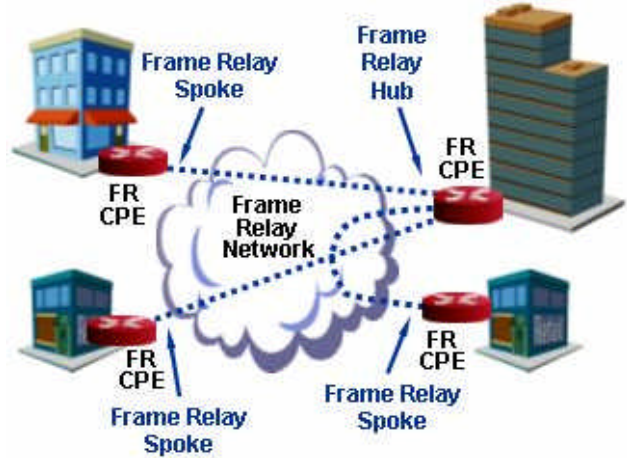


Figure 7: Frame Relay analogy to E-LAN Service

Now contrast the E-LAN Service with a typical hub and spoke Frame Relay network topology (refer to Figure 7). Frame Relay PVCs are point-to-point connections and Frame Relay creates a multipoint service via multiple point-to-point PVC connections. As new sites are added, a new PVC must be added between the new “spoke” site and the “hub” site requiring provisioning at both sites instead of just at the new “spoke” site.

“From a Subscriber standpoint, an E-LAN Service makes the MEN look like a LAN.”

E-LAN Service in point-to-point configuration

An E-LAN Service can be used to connect only two UNIs (sites). While this may appear similar to an E-Line Service, there are significant differences.




Figure 8: Adding a site using E-Line Service

With an E-Line Service, when a new UNI (site) is added, a new EVC must be added to connect the new UNI to one of the existing UNIs. In Figure 8, a new site (UNI) is added and a new EVC must be added to all sites to achieve full connectivity when using the E-Line Service. The Frame Relay analogy would be to add a Frame Relay PVCs between each site.




Figure 9: Adding a site using an E-LAN Service

With an E-LAN Service (refer to Figure 9), only the new UNI needs to be added to the multipoint EVC. No additional EVCs are required since the E-LAN Service uses a multipoint-to-multipoint EVC. An E-LAN Service also allows the new site (UNI) to communicate with all other UNIs. With an E-Line Service, this would require separate EVCs to all UNIs. Hence, an E-LAN Service requires only one EVC to achieve multi-site connectivity.

In summary, an E-LAN Service can interconnect large numbers of sites with less complexity than meshed or hub and spoke connections implemented using point-to-point networking technologies such as Frame Relay or ATM. Furthermore, an E-LAN Service can be used to

create a broad range of services such as Private LAN and Virtual Private LAN services.

“.. an E-LAN Service requires only one EVC to achieve multi-site connectivity.”

Ethernet Service Attributes

The Ethernet Service Attributes define the capabilities of the Ethernet Service Type. As previously mentioned, some Service Attributes apply to the UNI while others apply to the EVC. This distinction will be pointed out for the different service attributes.

Ethernet Physical Interface

At the UNI, the Ethernet physical interface has several service attributes. They are described in the following subsections.

Physical Medium

The Physical Medium UNI service attribute specifies the physical interface as defined by the IEEE 802.3-2000 standard. Example Physical Media includes 10BaseT, 100BaseT and 1000BaseSX.

Speed

The Speed UNI service attribute specifies the standard Ethernet speeds of 10Mbps, 100Mbps, 1Gbps and 10Gbps.

Mode

The Mode UNI service attribute specifies whether the UNI supports full or half duplex or can perform auto speed negotiation.

MAC Layer

The MAC Layer UNI service attribute specifies which MAC layer is supported. The currently supported MAC layers are specified in IEEE 802.3-2002.

Bandwidth Profile

The MEF has defined the Bandwidth Profile service attribute that can be applied at the UNI or for an EVC. A Bandwidth Profile is a limit on the rate at which Ethernet frames can traverse the UNI. There can be separate Bandwidth Profiles for frames ingressing into the network and for frames egressing from the network. The Committed Information Rate for a Frame Relay PVC is an example of a Bandwidth Profile.

The MEF has defined the following three Bandwidth Profile service attributes:

- Ingress Bandwidth Profile Per Ingress UNI
- Ingress Bandwidth Profile Per EVC
- Ingress Bandwidth Profile Per CoS Identifier

The Bandwidth Profile service attribute consist of four traffic parameters described in the following sections. These parameters affect the bandwidth or throughput delivered by the service. It is important to understand what these parameters mean and more importantly, how they affect the service offering.

A bandwidth profile for an Ethernet service consists of the following traffic parameters:

- CIR (Committed Information Rate)
- CBS (Committed Burst Size)
- EIR (Excess Information Rate)
- EBS (Excess Burst Size)

A service may support up to three different types of Bandwidth Profiles <CIR, CBS, EIR, EBS> at the UNI. One could apply a bandwidth profile per UNI, per EVC at the UNI or per CoS Identifier (Refer to Class of Service Identifiers section) for a given EVC at the UNI.

Service Frame Color

Before discussing the traffic parameters, the concept of service frame color should be introduced since it the result of different levels of traffic conformance to the bandwidth profile.

The “color” of the service frame is used to determine the bandwidth profile conformance of a particular service frame. A service may have two or three colors depending upon the configuration of the traffic parameters.

A service frame is marked “green” if it is conformant with CIR and CBS in the bandwidth profile, i.e., the average service frame rate and maximum service frame size is less than or equal to the CIR and CBS, respectively. This is referred to as being “CIR-conformant”.

A service frame is marked “yellow” if it is not CIR-conformant but conformant with the EIR and EBS in the bandwidth profile, i.e., the average service frame rate is greater than the CIR but less than the EIR and the maximum service frame size is less than the EBS. This is referred to as being “EIR-conformant”.

A service frame is marked “red” and discarded if it is neither CIR-conformant nor EIR-conformant.

The MEF Technical Committee is currently working on how colors are marked in service frames.

CIR and CBS

The Committed Information Rate (CIR) is the average rate up to which service frames are delivered per the service performance objectives, e.g., delay, loss, etc. The CIR is an average rate because all service frames are sent at the UNI speed, e.g., 10Mbps, and not at the CIR, e.g., 2Mbps. CBS is the size up to which service frames may be sent and be CIR-conformant.

Service frames whose average rate is greater than the CIR or those which send more than CBS bytes are not CIR-conformant and may be discarded or colored to indicate non-conformance depending upon whether the service frames are EIR-conformant or not.

A CIR may be specified to be less than or equal to the UNI speed. If multiple bandwidth profiles are applied at the UNI, the sum of all CIRs must be less than or equal to the UNI speed.

A CIR of zero indicates that the service provides no bandwidth or performance assurances for delivery of subscriber service frames. This is often referred to as a “best effort” service.

EIR and EBS

The Excess Information Rate (EIR) specifies the average rate, greater than or equal to the CIR, up to which service frames are delivered without any performance objectives. The EIR is an average rate because all service frames are sent at the UNI speed, e.g., 10Mbps, and not at the EIR, e.g., 8Mbps. EBS is the size up to which service frames may be sent and be EIR-conformant.

Service frames whose average rate is greater than the EIR or those which send more than EBS bytes are not EIR-conformant and may be discarded or colored to indicate non-conformance depending upon the service being offered.

The EIR may be specified to be less than or equal to the UNI speed. When, non-zero, the EIR is greater than or equal to the CIR.

Performance Parameters

The performance parameters affect the service quality experienced by the subscriber. Performance parameters consist of the following:

- Availability
- Frame Delay
- Frame Jitter
- Frame Loss

Availability

The MEF Technical Committee is currently defining parameters and metrics for availability. This section will be updated as the work progresses further.

Frame Delay

Frame Delay is a critical parameter and can have a significant impact on the QoS for real-time applications services such as IP telephony.

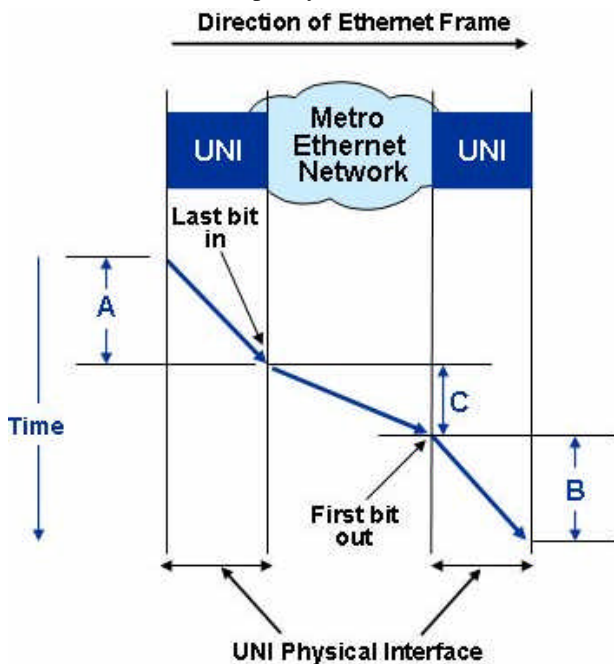


Figure 10: Network Delay Partitioning

Frame Delay can be broken down into three parts as illustrated in Figure 10 as represented by A, B and C. The delay introduced by A and B are dependent upon the line rate at the UNI, e.g., 10Mbps, and the Ethernet service frame size, e.g., 1518 bytes. For example, both A and B introduce 1.214ms of transmission delay for a standard service frame size of 1518 bytes and a 10Mbps UNI at both CEs. C is the amount of delay introduced by the Metro Ethernet Network and is statistically characterized by the Metro Ethernet Network provider

measured over a time interval. Frame Delay is represented by $A + B + C$ where A and B can be calculated while C is specified over a measurement interval. Note that the service frame size must also be specified in order to calculate A and B.

Frame Delay is defined as the maximum delay measured for a percentile of successfully delivered CIR-conformant (green) service frames over a time interval.

For example, the delay is measured between two 10Mbps UNIs using a 5 minute measurement interval and percentile of 95%. During the measurement interval, 1000 service frames were successfully delivered. The maximum delay for 95% of the 1000 successfully delivered service frames was measured to be 15ms. Therefore, $C = 15\text{ms}$. This results in a Frame Delay of:

$$\text{Frame Delay} = A + B + C = 1.214\text{ms} + 1.214\text{ms} + 15\text{ms} = 17.43\text{ms}$$

Services requiring stringent delay performance may provide a higher percentile, e.g., 99th percentile, used in the delay calculation. In general, the percentile is 95% or greater based on current industry practices.

The Frame Delay parameter is used in the CoS service attribute.

“Frame Delay is a critical parameter ... for real-time applications such as IP telephony”

Frame Jitter

Jitter, also known as delay variation, is a critical parameter for real-time applications such as IP telephony or IP video. These real-time applications require a low and bounded delay variation to function properly. While jitter is a critical parameter for real-time applications, jitter has essentially no negative QoS effect on non-real-time data applications.

Frame Jitter can be derived from the Frame Delay measurement. Over the population of frame delay samples used in the Frame Delay calculation, the service frame with the lowest service frame delay is subtracted from Frame Delay value (maximum frame delay in the sample population). This is the Frame Jitter. Note that Frame Jitter only applies to all CIR-conformant (green) service frames. Frame Jitter can be calculated as follows:

Frame Jitter = Frame Delay value – Service Frame with lowest delay in Frame Delay population

Using the example in Figure 10, the Frame Delay over a 5 minute measurement interval and 95th percentile was calculated to be 17.43 ms. Over the population used in the Frame Delay calculation, the service frame with the lowest delay was measured to be 15 ms. Therefore, the Frame Jitter is 2.43 ms.

$$\text{Frame Jitter} = 17.43\text{ms} - 15\text{ms} = 2.43\text{ms}$$

The Frame Jitter parameter is used in the CoS Service Attribute.

Frame Loss

Frame loss is defined the percentage of CIR-conformant (green) service frames not delivered between UNIs over a measurement interval. Note that the MEF Technical Committee has currently defined Frame Loss for point-to-point EVCs and is working on the definition for multipoint-to-multipoint EVCs.

$$\text{Frame Loss} = \left[1 - \frac{\text{Number of Service Frames delivered to destination UNI in the EVC}}{\text{Total Service Frames sent to destination UNI in the EVC}} \right] \times 100$$

For example, in Figure 11, over a point-to-point EVC, 1000 service frames were transmitted from the source UNI to the destination UNI and during a 5 minute measurement interval. Over the measurement interval, 990 service frames were delivered successfully to the destination UNI. In this example, the Frame Loss would be as follows:

$$\text{Frame Loss} = \left[1 - \frac{990 \text{ service frames delivered}}{1000 \text{ total service frames to be delivered}} \right] \times 100 = 1\%$$

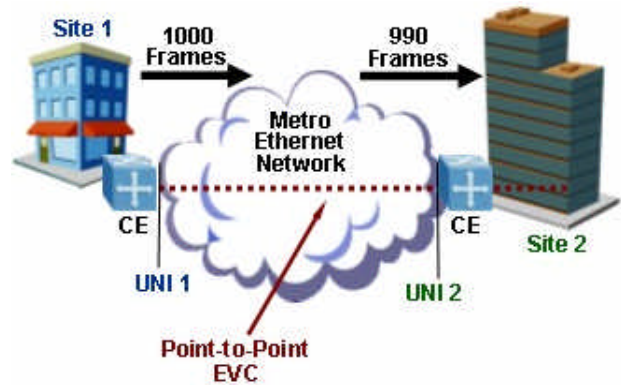


Figure 11: Frame Loss Example for Point-to-Point EVC

Frame Loss has a different impact on the QoS, depending upon the application, service or higher layer protocols used by the service. For example, a 1% packet loss for a Voice over IP (VoIP) application may be acceptable. A 3% packet loss, however, will result in unacceptable voice quality. Streaming media applications can tolerate varying degrees of packet loss and compensate by adjusting the transmit rate as packet loss is detected. TCP-based applications, such as Internet web browser HTTP requests can tolerate varying degrees of packet loss because the TCP protocol will retransmit lost packets. However, increasingly excessive packet loss will negatively affect the subscriber’s QoS.

The Frame Loss parameter is used in the CoS Service Attribute.

“Frame loss has a different impact on the QoS, depending upon the application, service or higher layer protocols used”

Class of Service Identifiers

Metro Ethernet networks may offer different classes of service (CoS) to subscribers identified via various CoS Identifiers (CoS IDs) such as:

- Physical Port
- CE-VLAN CoS (802.1p)
- DiffServ / IP TOS

The service provider will enforce different traffic parameters, e.g., CIR, for each class of service. Each class of service will offer different levels of performance as specified in the performance parameters per class of service, e.g., delay, jitter and loss. If a service provider

supports multiple classes of service between UNIs, the traffic and performance parameters must be specified for each class.

The following subsections will explore each of the aforementioned CoS identifiers.

Physical Port

In this case, a single class of service is provided per physical port. All traffic ingressing or egressing the port receives the same CoS. This is the simplest form to implement but has the least amount of flexibility. The method is also costly for subscribers who need multiple classes of service for their traffic. If the subscriber requires multiple classes of service for their traffic, separate physical ports would be required, each providing the different CoS.

A single set of traffic and performance parameters apply to a port-based implementation, i.e., a single CIR, CBS, EIR and EBS for the interface, and delay, jitter and loss for the service.

CE-VLAN CoS (802.1p)

The MEF has defined the CE-VLAN CoS as the CoS (802.1p) bits in the IEEE 802.1Q tag in a tagged Service Frame. When using the CE-VLAN CoS, up to 8 classes of service can be indicated. If the service provider supports CE-VLAN CoS to determine the class of service, the service provider should specify the bandwidth profile and performance parameters for each CoS.

The class of service may be based on forwarding (emission) priority, i.e., service frames with CE-VLAN CoS 7 get forwarded ahead of service frames with CE-VLAN CoS 6. The CoS may also use more sophisticated DiffServ-based behaviors applied to the service frames for a given CE-VLAN CoS value. For example, CE-VLAN CoS 6 may get DiffServ Expedited Forwarding behavior and CE-VLAN CoS 5/4/3 get DiffServ Assured Forwarding behavior where CE-VLAN CoS 5 has lowest drop precedence and CE-VLAN CoS 3 has highest drop precedence. (Refer to [DiffServ], [EF PHB] and [AF PHB]).

Note that an Ethernet Service that uses the subscriber's CE-VLAN CoS values to determine the class of service may or may not preserve the subscriber's CE-VLAN CoS bits in the VLAN tag at the UNI (See VLAN Tag Support section). Services that provide VLAN tag translation may also provide a class of service such that

multiple CE-VLAN CoS values are mapped to the same class of service.

DiffServ / IP TOS values

DiffServ or IP TOS values can be used to determine the class of service. IP TOS, in general, is used to provide 8 classes of service known as IP precedence. IP precedence is very similar to the 802.1p definition in IEEE 802.1Q when CoS is provided based on forwarding (emission) priority.

DiffServ, by contrast, has defined several per-hop behaviors (PHBs) that provide more robust QoS capabilities when compared to the simple forwarding-based priority of IP TOS and 802.1p. DiffServ uses the same field in the IP header (2nd byte) as IP TOS but redefines the meaning of the bits. DiffServ provides 64 different values (called DiffServ codepoints or DSCPs) that can be used to determine the class of service. Standardized DiffServ PHBs include Expedited Forwarding (EF) for a low delay, low loss service, four classes of Assured Forwarding (AF) for bursty real-time and non-real-time services, Class Selector (CS) for some backward compatibility with IP TOS, and Default Forwarding (DF) for best effort services.

Unlike CE-VLAN CoS (802.1p), DiffServ and IP TOS require the subscriber and provider's networking equipment to inspect the IP packet header in the Ethernet frame's payload to determine the DSCP or TOS value. Essentially all routers and Ethernet switches support this capability, except for the low end consumer or small office versions. If the device cannot inspect the DSCP in the IP packet header, then a mapping function between DiffServ, IP TOS and 802.1p must be performed by the last / first IP-capable device so the CoS can be determined.

Note that routing functions are not required on the Ethernet switch to support a DSCP/IP TOS-based classification method. The switch simply needs to be able to classify the DiffServ/TOS Field in the IP header in the Ethernet frame's payload in addition to inspecting the Ethernet frames 802.1Q tag.

Up to 64 different traffic and performance parameters can be applied to a DiffServ-based implementation, i.e., a separate CIR, CBS, EIR, EBS, delay, jitter and loss for each of the 64 CoS levels defined by the DiffServ values. In general, the 4 standard DiffServ PHBs would be implemented, namely, Expedited Forwarding, Assured Forwarding, Class Selector and Default Forwarding. This would result in up to 13 possible classes of service

(1 EF, 4 AF, 7 CS and 1 DF) to be implemented. Like 802.1p, an IP TOS-based implementation can create up to 8 classes of service.

Finally, the Class of Service EVC service attribute defines the class of service offered over an EVC based on the following parameters:

- Class of Service Identifier
- Frame Delay
- Frame Jitter
- Frame Loss

For example, a service offers a “Premium” class of service in the metro network. For this service, the Class of Service EVC service attribute could be specified as in Table 1.

Class of Service parameters	Example Value
Class of Service Identifier	CE-VLAN CoS (802.1p) = 6
Frame Delay	< 10ms
Frame Jitter	< 1 ms
Frame Loss	< 0.01% (99 th percentile)

Table 1: Example CoS EVC service attribute

Service Frame Delivery

An Ethernet Virtual Connection (EVC) allows Ethernet service frames to be exchanged between UNIs that are connected via the same EVC. Some frames are subscriber data service frames while others are Ethernet control service frames. There are many possible ways to determine which frames are delivered and, in the case of a multipoint EVC, to which UNIs they should be delivered. Several parameters can be used to specify Ethernet service frame delivery.

Some Ethernet Services deliver all types of service frames while others have some restrictions. Service providers specify the types of service frames supported (and the actions that are taken) and those that are not supported (discarded). The following subsections provide some different types of service frames and how they may be supported.

Unicast Service Frame Delivery

The unicast service frame is defined by the destination MAC address. The unicast service frame address may be “known” (already learned by the network) or “unknown”. This EVC service attribute specifies whether unicast service frames are Discarded, Delivered

Unconditionally or Delivered Conditionally for each ordered UNI pair. If the service frames are delivered conditionally, the conditions would be specified.

Multicast Service Frame Delivery

IETF RFC 1112 defines the Internet multicast range to be destination MAC addresses 01-00-5E-00-00-00 through 01-00-5E-7F-FF-FF. This EVC service attribute specifies whether multicast service frames are Discarded, Delivered Unconditionally or Delivered Conditionally for each ordered UNI pair. If the service frames are delivered conditionally, the conditions would be specified.

Broadcast Frame Delivery

IEEE 802.3 defines the Broadcast address as a destination MAC address of FF-FF-FF-FF-FF-FF. This EVC service attribute specifies whether broadcast service frames are Discarded, Delivered Unconditionally or Delivered Conditionally for each ordered UNI pair. If the service frames are delivered conditionally, the conditions would be specified.

Layer 2 Control Protocol Processing

This service attribute can be applied at the UNI or per EVC. There are many layer 2 control protocols that may be used in the network. Table 2 provides a partial list of standardized protocols currently in use. Depending upon the service offering, the provider may process or discard these protocols at the UNI or pass them to the EVC. The provider may also discard or tunnel these protocols across an EVC.

Protocol	Destination MAC Address
IEEE 802.3x MAC Control Frames	01-80-C2-00-00-01
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
Spanning Tree Protocol (STP)	01-80-C2-00-00-00
A protocol to be multicast to all bridges in a bridged LAN	01-80-C2-00-00-10

Table 2: Standardized Layer 2 Control Protocols

In general, all Ethernet Services support Unicast, Multicast and Broadcast service frames.

An E-LAN Service will support address learning and unicast. Ethernet frames with an unknown unicast, multicast or broadcast address will be delivered to all UNIs associated with the Ethernet Virtual Connection (EVC), while frames with a known unicast address will


be delivered only to the UNI where that MAC address has been learned.

“In general, all Ethernet Services will support Unicast, Multicast and Broadcast service frames.”

VLAN Tag Support

VLAN tag support provides another important set of capabilities that affect service frame delivery and performance. Since Ethernet service frames may be 802.1Q tagged or untagged, it is important to understand what happens to both tagged and untagged frames, and whether the VLAN ID in a tagged service frame is used to determine service frame delivery. Since VLAN support varies significantly between Ethernet Services, it is important to understand the variations.

Note that the UNI pairs for an EVC could support different VLAN tag types. For example, one UNI may support only untagged service frames, while the other UNI may only support tagged service frames. Furthermore, another UNI may support both tagged and untagged service frames. Refer to Figure 12. The utility of this is explained in the subsequent section on Service Multiplexing.



For UNIs that support VLAN tagging, the subscriber must know how VLAN tags are supported by the service and whether they are preserved or mapped.

Provider versus Customer VLAN tag

A provider may add an additional VLAN tag to the service frame to isolate the subscribers' VLAN tags. One proprietary approach, called VLAN tag stacking

(also referred to as Q-in-Q) inserts a second provider VLAN tag into the subscriber's service frame Ethernet header. Another proprietary approach called, MAC-in-MAC, adds an additional provider Ethernet MAC header (including an additional VLAN tag) to the subscriber's service frame.

To distinguish the subscriber's VLAN tag from the provider inserted VLAN tag (when using Q-in-Q or MAC-in-MAC), the MEF has defined the term CE-VLAN ID (Customer Edge VLAN ID) to represent the subscriber's VLAN ID. The CE-VLAN tag also contains the 802.1p field which the MEF has termed CE-VLAN CoS which refers to the subscriber's 802.1p field.

CE-VLAN Service Attributes

The MEF has defined the two service attributes regarding CE-VLAN tag support.

- CE-VLAN ID Preservation
- CE-VLAN CoS Preservation

The CE-VLAN tag consists of both the CE-VLAN ID and the CE-VLAN CoS so a service may preserve one, both or neither of these.

CE-VLAN ID Preservation

The CE-VLAN ID Preservation is an EVC service attribute that defines whether the CE-VLAN ID is preserved (unmodified) across the EVC or not (in which case it would be mapped to another value). CE-VLAN ID preservation also implies that there is no constraint on the subscriber's choice of VLAN ID or the number of VLAN IDs that can be used on one interface.

CE-VLAN ID preservation is useful for services such as LAN extension (Refer to the Example Service Offerings section) because the CE-VLAN IDs may be used in the subscriber's network and would need to be preserved. Note that a service supporting CE-VLAN ID Preservation may also support untagged Ethernet service frames and send them along the same EVC unaltered.

CE-VLAN CoS Preservation

CE-VLAN CoS preservation is an EVC service attribute that defines whether the CE-VLAN CoS bits, i.e., 802.1p bits, are preserved (unmodified) across the EVC or not (in the latter case they would be mapped to another value).

CE-VLAN CoS Preservation is also useful for services such as LAN extension (Refer to the Example Service Offerings section) because the CE-VLAN CoS bits may

be used in the subscriber’s network and would need to be preserved.

Mapping VLAN IDs

CE-VLAN IDs must be mapped when one UNI supports tagging and the other UNI does not support tagging. In these cases, the CE-VLAN ID used to identify an EVC is locally significant to each UNI. To address this, the MEF has defined two service attributes. One is the CE-VLAN ID / EVC Map, which provides a mapping table between the CE-VLAN IDs at the UNI to the EVC to which they belong. The other is called the UNI List which provides a list of UNIs associated with an EVC. The UNI List service attribute for an E-Line Service would consist of two UNIs while for an E-LAN Service, the UNI List would consist of two or more UNIs.

When a UNI does not support VLAN tags, any Ethernet service frames delivered at the UNI will be delivered without VLAN tags. If the originating UNI supports VLAN tags and the service frame was sent to the UNI with a CE-VLAN tag, the provider will remove the CE-VLAN tag before delivering the service frames to the UNI that does not support VLAN tagging. For service frames sent from a UNI supporting untagged service frames to a UNI supporting tagged service frames, the service provider will insert the proper CE-VLAN tag before delivery to the UNI supporting tagged service frames as defined in the CE-VLAN ID / EVC Map service attribute.

UNI Capability	VLAN Tag Support		
	Untagged	Tagged	Tagged / Untagged
VLAN tags prohibited	v	N/A	N/A
VLAN tags mapped	v	v	v
VLAN tags preserved	N/A	v	v ³

Table 3: VLAN Tag support possibilities on a UNI

Table 3 provides a summary of the possible combinations of VLAN tag support and VLAN tag interpretation at the UNI. Some services may support only one of these possible combinations, while others may support more than one and allow the subscribers to choose.

³ Untagged service frames could be supported at a UNI that supports VLAN tag preservation.

“Since VLAN support varies significantly between Ethernet Services, it is important to understand the variations.”

Service Multiplexing

The Service Multiplexing service attribute is used to support multiple EVCs at the UNI. Figure 13 shows an example of service multiplexing. In this example, UNI A is a Gigabit Ethernet UNI that supports service multiplexing. UNIs B, C and D are 100 Mbps UNIs. Using service multiplexing, three point-to-point EVCs are set up at UNI A, namely, EVC 1, EVC 2 and EVC 3. Service multiplexing at UNI A eliminates the need for three different physical interfaces (UNIs).

Since only one EVC is used at the UNIs B, C and D, these UNIs need not support service multiplexing and may or may not support VLAN tags, depending on what the service supports and the subscriber requires. For example, on EVC 1 from service multiplexed UNI A that supports tagged services frames to UNI B that does not support tagged service frames, the Metro Ethernet network would remove the CE-VLAN tags from Ethernet frames sent from UNI A to UNI B and add the CE-VLAN tags for frames sent from UNI B to UNI A.

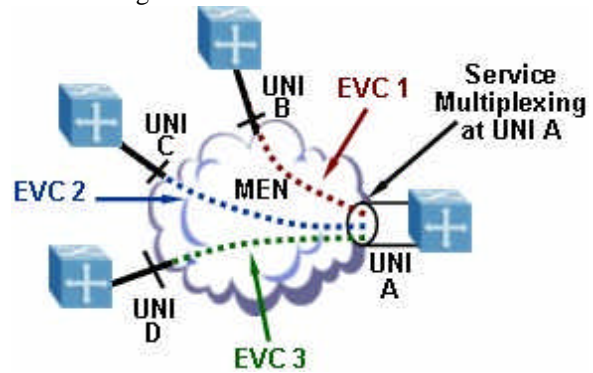


Figure 13: Service Multiplexing with Point-to-Point EVCs

Service Multiplexing Benefits

In summary, service multiplexing allows one UNI (physical interface) to support multiple EVCs. Compared to the alternative of a separate physical interface for each EVC, there are several benefits of a service-multiplexed interface.

Lowers Equipment Cost

Service multiplexing minimizes the number of subscriber router or switch ports and maximizes the density of

port/slot utilization. This generally reduces the subscriber’s equipment cost and may also help delay or eliminate the need for an equipment upgrade.

Minimizes space, power and cabling

Service multiplexing minimizes space, power and cabling. Compared to multiple non-multiplexed UNIs, service-multiplexed UNIs reduce the amount of rack space and power required for the subscriber and service provider equipment and reduces the number of cross connects between them.

Simplifies new service activation

Service multiplexing allows new EVCs to be established without the need for a site visit for equipment installation, cross connects or patch cables.

Service multiplexing ..., minimizes space, power and cabling, and simplifies new service activation.

Bundling

The Bundling service attribute enables two or more CE-VLAN IDs to be mapped to a single EVC at a UNI. With bundling, the provider and subscriber must agree on the CE-VLAN IDs used at the UNI and the mapping between each CE-VLAN ID and a specific EVC. A service provider might allow the subscriber to select the CE-VLAN IDs and mapping (perhaps within some restricted range), or the service provider may provide the CE-VLAN ID values.

A special case of Bundling occurs when every CE-VLAN ID at the UNI (interface) maps to a single EVC. This service attribute is called All to One Bundling.

Security Filters

Some service providers may allow a subscriber to specify additional filtering of Ethernet frames for added security or traffic management. For example, the service provider might allow a subscriber to specify a list of Ethernet MAC addresses that should be granted access on a given UNI, sometimes referred to as an “Access Control List.” The service provider would then discard frames with source MAC addresses not on that list. Such security filtering capabilities may vary between providers. This section will be updated as MEF work in this area progresses.

Services Framework Summary

A broad range of Ethernet Services can be constructed using the Ethernet Services Framework by selecting Ethernet Service Attributes and applying different parameter values.

Table 4 provides a summary of the Ethernet Service Attributes and their associated parameters for UNIs while Table 5 provides these for EVCs.

UNI Service Attribute	Parameter Value or Range of Values
Physical Medium	IEEE 802.3-2002 Physical Interface
Speed	10 Mbps, 100 Mbps, 1 Gbps, or 10 Gbps
Mode	Full Duplex, or Auto negotiation
MAC Layer	IEEE 802.3-2002
Service Multiplexing	Yes ⁴ or No
CE-VLAN ID / EVC Map	Mapping table of CE-VLAN IDs to EVC
Bundling	Yes or No ⁵
All to One Bundling	Yes ⁶ or No ⁷
Ingress Bandwidth Profile Per Ingress UNI	No or <CIR, CBS, EIR, EBS>
Ingress Bandwidth Profile Per EVC	No or <CIR, CBS, EIR, EBS>
Ingress Bandwidth Profile Per CoS Identifier	No or <CIR, CBS, EIR, EBS>
Layer 2 Control Protocol Processing	Peer, Discard or Pass to EVC IEEE 802.3x MAC Control Frames
	Peer, Discard or Pass to EVC Link Aggregation Control Protocol (LACP)
	Peer, Discard or Pass to EVC IEEE 802.1x Port Authentication
	Peer, Discard or Pass to EVC Generic Attribute Registration Protocol (GARP)
	Peer, Discard or Pass to EVC Spanning Tree Protocol (STP)
	Peer, Discard or Pass to EVC a protocol multicasted to all bridges in a bridged LAN

Table 4: UNI Service Attribute Summary

⁴ If Yes, then All to One Bundling must be No.

⁵ Must be No if All to One Bundling is Yes and Yes if All to One Bundling is No.

⁶ If Yes, then Service Multiplexing and Bundling must be No.

⁷ Must be No if Bundling is Yes.

EVC Service Attribute	Type of Parameter Value
EVC Type	Point-to-Point or Multipoint-to-Multipoint
UNI List	Provides the list of UNIs associated with an EVC.
CE-VLAN ID Preservation	Yes or No
CE-VLAN CoS Preservation	Yes or No
Unicast Frame Delivery	Deliver Unconditionally or Deliver Conditionally
Multicast Frame Delivery	Deliver Unconditionally or Deliver Conditionally
Broadcast Frame Delivery	Deliver Unconditionally or Deliver Conditionally
Layer 2 Control Protocol Processing	Discard ⁸ or Tunnel IEEE 802.3x MAC Control Frames
	Discard or Tunnel Link Aggregation Control Protocol (LACP)
	Discard or Tunnel IEEE 802.1x Port Authentication
	Discard or Tunnel Generic Attribute Registration Protocol (GARP)
	Discard or Tunnel Spanning Tree Protocol (STP)
Service Performance	Discard or Tunnel a protocol multicasted to all bridges in a bridged LAN
	<CoS Identifier, Frame Delay, Frame Jitter, Frame Loss>.

Table 5: EVC Service Attribute Summary

Example Service Offerings

The following sections describe useful, popular and potential mass-market Ethernet service offerings.

Dedicated Internet Access

Subscribers are continually seeking higher speed Internet connections to support their business objectives. An Ethernet Virtual Connection can provide an ideal way to connect the subscriber’s site to the local point-of-presence (POP) of an Internet Service Provider (ISP). The most common service for Internet Access is a point-to-point E-Line Service as shown in Figure 14.

In the simplest scenario, untagged service frames may be used at the subscriber’s site. A subscriber may want to use the Border Gateway Protocol (BGP) for multi-

homing to two or more ISPs⁹. In this case, the subscriber would use a separate E-Line Service to each ISP. If the subscriber wants to use the same UNI to support both Internet access and an Intranet or Extranet connection in the Metro, then separate EVCs would also be used.




Figure 14: Dedicated Internet Access

The ISP typically service multiplexes subscribers over a high-speed Ethernet UNI. For example, in Figure 14, the ISP may have a 1Gbps UNI (UNI 3) while subscriber UNIs 1 and 2 may be 100Mbps. In this example, there is no service multiplexing at subscriber UNIs 1 and 2. Service multiplexing is only performed at the ISP’s UNI (UNI 3) so in effect, subscriber UNIs 1 and 2 have dedicated Ethernet connections to the ISP POP. Table 6 provides an example service level specification (SLS) for the Dedicated Internet Access service.

UNI Service Attribute	Service Attribute Values and Parameters
Physical Medium	IEEE 802.3-2002 Physical Interface
Speed	UNIs 1 and 2: 100Mbps UNI 3: 1Gbps
Mode	UNIs 1 and 2: 100Mbps FDX fixed UNI 3: 1Gbps FDX
MAC Layer	IEEE 802.3-2002
Service Multiplexing	No at UNIs 1 and 2 Yes at UNI 3
CE-VLAN ID / EVC Map	N/A since only untagged frames used over the EVC
Bundling	No
All to One Bundling	No

⁸ When a BPDU is discarded at the UNI, the Layer 2 Control Protocol Processing at the EVC is Not Applicable (N/A)

⁹ ISP will likely use a Service-Multiplexed UNI in order to support multiple subscribers on a single port.

UNI Service Attribute	Service Attribute Values and Parameters
Ingress Bandwidth Profile Per EVC	UNIs 1 and 2: CIR=50Mbps, CBS=2MB, EIR=100Mbps, EBS=4MB UNI 3: CIR=500Mbps, CBS=20MB, EIR=1Gbps, EBS=40MB
Layer 2 Control Protocol Processing	Discard 802.3x MAC Control Frames
	Discard Link Aggregation Control Protocol (LACP)
	Discard 802.1x Port Authentication
	Discard Generic Attribute Registration Protocol (GARP)
	Discard Spanning Tree Protocol
	Discard a protocol multicasted to all bridges in a bridged LAN
EVC Service Attribute	Service Attribute Values and Parameters
EVC Type	Point-to-Point
UNI List	EVC 1: UNI 1, UNI 3 EVC 2: UNI 2, UNI 3
CE-VLAN ID Preservation	No. Mapped VLAN ID for use with multi-homed ISPs (if required)
CE-VLAN CoS Preservation	No
Unicast Frame Delivery	Deliver Unconditionally for each UNI pair
Multicast Frame Delivery	Deliver Unconditionally for each UNI pair
Broadcast Frame Delivery	Deliver Unconditionally for each UNI pair
Layer 2 Control Protocol Processing	N/A ⁸ - IEEE 802.3x MAC Control Frames
	N/A - Link Aggregation Control Protocol (LACP)
	N/A - IEEE 802.1x Port Authentication
	N/A - Generic Attribute Registration Protocol (GARP)
Service Performance	Only 1 CoS supported. Frame Delay < 30ms (95 th percentile), Frame Jitter: N/S ¹⁰ , Frame Loss < 0.1%

Table 6: Example Dedicated Internet Access SLS

LAN Extension

Subscribers with multiple sites in a metro area often want to interconnect them at high speeds so all sites appear to be on the same Local Area Network (LAN) and have equivalent performance and access to resources such as servers and storage. This is commonly referred to as a LAN Extension. A LAN Extension implies connecting

the subscriber’s LANs together without any intermediate routing between UNIs (sites). In some cases, this is simpler and cheaper than routing, although it would not typically scale well for very large networks.

To connect only two sites, a point-to-point E-Line Service could be used. To connect three or more sites, the subscriber could use multiple E-Line Services or an E-LAN Service.

Since a LAN Extension may use a switch-to-switch connection, it generally requires more transparency than Internet Access. For example, the subscriber may want to run the Spanning Tree Protocol across the interconnected sites thus requiring the Ethernet Service to support BPDU tunneling. If VLANs are used in the subscriber’s network, e.g., to separate departmental traffic, the subscriber may also need to make the VLANs present at multiple sites, requiring support for the subscriber’s CE-VLAN tags to be carried across the Metro Ethernet Network (MEN) connection.

Figure 15 provides an example of LAN Extension involving four sites interconnected across a MEN. Three separate subscriber VLANs are present at the different sites but none are present at all sites. This example achieves the basic goal of LAN Extension across multiple sites without routing the traffic.

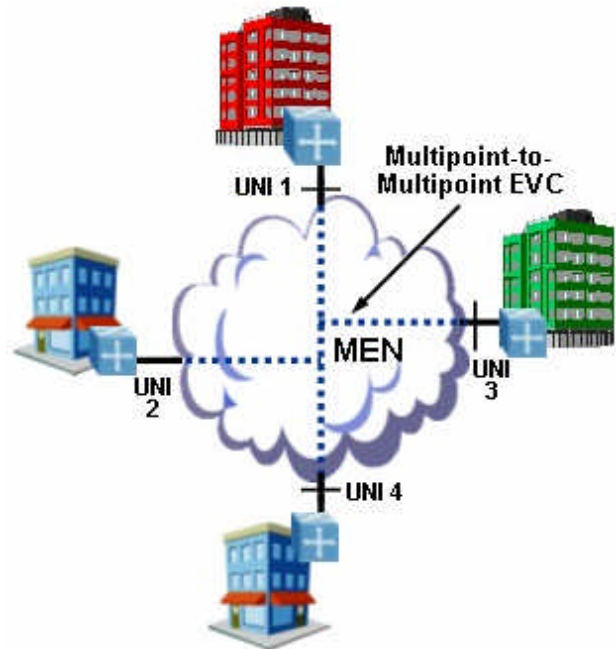


Figure 15: LAN Extension using E-LAN Service

The subscriber can use a single E-LAN Service to connect all four sites and carry all VLANs. Each

¹⁰ Not Specified

interface would support CE-VLAN ID and CE-VLAN CoS preservation, i.e., the subscriber’s VLAN tag and 802.1p bits are not modified by the MEN. In this case, the MEN appears like a single Ethernet segment in which any site can be a member of any VLAN. The advantage with this approach is that the subscriber can configure CE-VLANs across the four sites without any need to coordinate with the service provider.

Table 7 provides an example SLS for LAN Extension using a single E-LAN Service.

UNI Service Attribute	Service Attribute Values and Parameters
Physical Medium	IEEE 802.3-2002 Physical Interface
Speed	10Mbps (all UNIs)
Mode	FDX fixed speed (all UNIs)
MAC Layer	IEEE 802.3-2002
Service Multiplexing	No
CE-VLAN ID / EVC Map	All CE-VLAN IDs map to the single EVC
Bundling	No
All to One Bundling	Yes
Ingress Bandwidth Profile Per Ingress UNI	All UNIs: CIR=5Mbps, CBS=256KB, EIR=10Mbps, EBS=512KB
Layer 2 Control Protocol Processing	Process IEEE 802.3x MAC Control Frames
	Process Link Aggregation Control Protocol (LACP)
	Process IEEE 802.1x Port Authentication
	Pass to EVC Generic Attribute Registration Protocol (GARP)
	Pass to EVC Spanning Tree Protocol
	Pass to EVC a protocol multicasted to all bridges in a bridged LAN
EVC Service Attribute	Service Attribute Values and Parameters
EVC Type	Multipoint-to-Multipoint
UNI List	UNI 1, UNI 2, UNI 3, UNI 4
CE-VLAN ID Preservation	Yes
CE-VLAN CoS Preservation	Yes
Unicast Frame Delivery	Deliver Unconditionally for each UNI pair
Multicast Frame Delivery	Deliver Unconditionally for each UNI pair
Broadcast Frame Delivery	Deliver Unconditionally for each UNI pair

EVC Service Attribute	Service Attribute Values and Parameters
Layer 2 Control Protocol Processing	N/A - IEEE 802.3x MAC Control Frames
	N/A - Link Aggregation Control Protocol
	N/A - IEEE 802.1x Port Authentication
	Tunnel Generic Attribute Registration Protocol (GARP)
	Tunnel Spanning Tree Protocol (STP)
	Tunnel a protocol multicasted to all bridges in a bridged LAN
Service Performance	One CoS for all UNIs Frame Delay < 30ms, Frame Jitter: N/S, Frame Loss < 0.1%

Table 7: Example LAN Extension SLS for E-LAN Service

Intranet / Extranet L2 VPN

Ethernet Services can also be a good choice for routed Intranet connections to remote sites and Extranet connections to suppliers, customers and business partners. Figure 16 shows one Enterprise site, HQ (Headquarters), connecting to three other sites. One of the remote sites is part of the Enterprise’s internal Intranet, while the other two are Extranet locations of a supplier and a business partner. The HQ router interfaces to the Metro Ethernet Network (MEN) using a single Service-Multiplexed UNI and supports three separate point-to-point EVCs.

Such Extranet connections could be used to reach a wide variety of suppliers and partners who are connected to the same Metro Ethernet Network, including various “xSPs” (e.g., Application Service Providers, Managed Service Providers or Storage Service Providers).




Figure 16: Intranet / Extranet L2 VPN

Table 8 provides an example SLS for an Intranet / Extranet L2 VPN service using the E-Line service type.

UNI Service Attribute	Service Attribute Values and Parameters	
Physical Medium	IEEE 802.3-2002 Physical Interface	
Speed	UNI 1: 100 Mbps UNIs 2, 3 and 4: 10 Mbps	
Mode	10Mbps FDX (all sites except HQ), 100Mbps (HQ)	
MAC Layer	IEEE 802.3-2002	
Service Multiplexing	Yes at HQ	
CE-VLAN ID / EVC Map	CE-VLAN ID	EVC
	10, 11, 12	1
	20, 21, 22 30, 31, 32	2 3
Bundling	No	
All to One Bundling	No	
Ingress Bandwidth Profile Per CoS Identifier	<u>CoS 1</u> UNI 1: CIR=10Mbps, CBS=1MB, EIR=80Mbps, EBS=1MB UNIs 2, 3, 4: CIR=2 Mbps, CBS=640KB, EIR=8Mbps, EBS=200KB <u>CoS 2</u> UNI 1: CIR=50Mbps, CBS=2MB, EIR=100Mbps, EBS=1MB UNIs 2, 3, 4: CIR=5 Mbps, CBS=640KB, EIR=10Mbps, EBS=200KB	

UNI Service Attribute	Service Attribute Values and Parameters
Layer 2 Control Protocol Processing	Discard IEEE 802.3x MAC Control Frames
	Discard Link Aggregation Control Protocol
	Discard IEEE 802.1x Port Authentication
	Discard Generic Attribute Registration Protocol (GARP)
	Discard Spanning Tree Protocol
	Discard a protocol multicasted to all bridges in a bridged LAN
EVC Service Attribute	Service Attribute Values and Parameters
EVC Type	Point-to-Point
UNI List	EVC 1: UNI 1, UNI 2 EVC 2: UNI 1, UNI 3 EVC 3: UNI 1, UNI 4
CE-VLAN ID Preservation	No
CE-VLAN CoS Preservation	No
Unicast Frame Delivery	Deliver Unconditionally for each UNI pair
Multicast Frame Delivery	Deliver Unconditionally for each UNI pair
Broadcast Frame Delivery	Deliver Unconditionally for each UNI pair
Layer 2 Control Protocol Processing	N/A - IEEE 802.3x MAC Control Frames
	N/A - Link Aggregation Control Protocol
	N/A - IEEE 802.1x Port Authentication
	N/A - Generic Attribute Registration Protocol (GARP)
	N/A - Spanning Tree Protocol
	N/A - a protocol multicasted to all bridges in a bridged LAN
Service Performance	CoS 1: CoS ID=802.1p 3, Frame Delay <10ms, Frame Jitter: N/S, Frame Loss <0.05% CoS 2: CoS ID=802.1p 0, Frame Delay <30ms, Frame Jitter: N/S, Frame Loss <0.1%

Table 8 : Example Intranet / Extranet L2 VPN SLS

Three possible advantages of using Ethernet Virtual Connections (EVCs) rather than IP VPNs over the Internet are:

1. EVCs are inherently private and secure allowing subscribers to avoid costly and complex IP VPN equipment that is required for connections over the public Internet.
2. EVCs can provide very high bandwidth to support applications that require it. Setting up a dedicated

Ethernet connection can sometimes be more cost-effective than increasing Internet access bandwidth.

- EVCs can provide much higher performance than an IP VPN over the Internet, including lower latency and loss. For some Extranet applications, such as application outsourcing, this performance difference may be critical to user satisfaction and productivity.

Summary

Metro Ethernet Services can support a range of applications more easily, efficiently and cost-effectively than other network services. Using standard Ethernet interfaces, subscribers can set up secure, private Ethernet Virtual Connections across a Metropolitan Area, or even a Wide Area, to connect their sites together and connect to business partners, suppliers and the Internet. Using point-to-point E-Line Services and multipoint E-LAN Services, subscribers can connect to one site or many. With service options such as Service Multiplexing, subscribers can use a single UNI to support multiple connections. Additionally, with many Ethernet services, subscribers can buy just the bandwidth they need today, knowing that they can quickly and easily add bandwidth and set up new connections whenever they need to.

In order to support a wide range of applications and subscriber needs, Ethernet Services come in different types, with different service attributes. As outlined in this paper, the Metro Ethernet Forum is working to define and standardize these service types and attributes, enabling service providers to communicate their offerings clearly and subscriber to better understand and compare different services.

Over time, Ethernet Services will undoubtedly evolve to take advantage of advances in Ethernet technology and provide innovative new service features. However, it is clear that Ethernet Services will continue to provide the benefits of simplicity, cost effectiveness and flexibility that are unmatched for a wide range of applications.

Appendix

Icon Definitions

The following icons are used in this presentation to represent different network elements that can be used in a Metro Ethernet network.



Customer Edge Ethernet (L2 or L3) Switch



Customer Edge Router

Terminology

Term	Definition
ASP	Application Service Provider
CBS	Committed Burst Size
BPDU	Bridge Packet Data Unit
CoS	Class of Service
CE	Customer Edge equipment
CES	Circuit Emulation Services
CIR	Committed Information Rate
CPE	Customer Premise Equipment
DSCP	DiffServ Codepoint
DWDM	Dense Wave Division Multiplexing
EBS	Excess Burst Size
EIR	Excess Information Rate
EVC	Ethernet Virtual Circuit
FDX	Full Duplex
FR	Frame Relay
GFP	Generic Framing Protocol
HDX	Half Duplex
IANA	Internet Assigned Numbers Authority
ISP	Internet Service Provider
LAN	Local Area Network
MEN	Metro Ethernet Network
MPLS	Multi-protocol Label Switching
OAM&P	Operations, Administration, Management and Provisioning.
OWD	One Way Delay
PDU	Protocol Data Unit
POP	Internet Point of Presence
SONET	Synchronous Optical Network
QoS	Quality of Service
SSP	Storage Service Provider
VLAN	Virtual LAN
VPN	Virtual Private Network

References and Resources

Reference	Description
AF PHB	RFC 2597, Assured Forwarding PHB, http://www.ietf.org/rfc/rfc2597.txt
Default Forwarding	RFC 2474, DiffServ Field Definition, http://www.ietf.org/rfc/rfc2474.txt
DiffServ	RFC 2475, Differentiated Services Architecture, http://www.ietf.org/rfc/rfc2475.txt
EF PHB	RFC 3246, An Expedited Forwarding PHB, http://www.ietf.org/rfc/rfc3246.txt
IEEE 802.1x	“Port-Based Network Access Control”, http://standards.ieee.org/reading/ieee/std/lanman/restricted/802.1X-2001.pdf
IEEE 802.3-2002	“CSMA/CD Access Method and Physical Layer Specifications”, http://standards.ieee.org/reading/ieee/std/lanman/restricted/802.3-2002.pdf
IEEE 802.1Q	“Virtual Bridged Local Area Networks”,

Reference	Description
	http://standards.ieee.org/reading/ieee/std/lanman/802.1Q-1998.pdf
MEF 1.0	MEF Technical Specification “Ethernet Services Model, Phase 1”, http://www.metroethernetforum.org/PDFs/Standards/MEF-1.0.pdf
MEF Economic Model	“ Metro Ethernet Services for the Enterprise – The Business Case”, http://www.metroethernetforum.org/MEFBusinessCase101_FINAL.PDF
MEF FAQs	Metro Ethernet Forum Frequently Asked Questions, http://www.metroethernetforum.org/MEF_FAQ061702_Final.pdf
MEF Membership	Metro Ethernet Forum Membership Application, http://www.metroethernetforum.org/MEFMemberApplication10-19-01.pdf
MEF Positioning	Metro Ethernet Forum Positioning Statement, http://www.metroethernetforum.org/MEFPositioningStatementJan-17-02.pdf
MEN Technical Overview	“ Metro Ethernet Networks – A Technical Overview”, http://www.metroethernetforum.org/PDFs/WhitePapers/meftechoverviewwhitepaper102903.pdf
RFC 1112	“Host Extensions for IP Multicasting”, http://www.ietf.org/rfc/rfc1112.txt

Disclaimer

This paper reflects work-in-progress within the MEF, and represents a 75% member majority consensus as voted by 60 members at the October 2003 MEF Technical Committee meeting.

Some technical details may change in due course (by 75% vote) and this paper will be updated as deemed necessary to reflect such changes. The paper does not necessarily represent the views of the authors or their commercial affiliations.

About the Metro Ethernet Forum

The Metro Ethernet Forum (MEF) is a non-profit organization dedicated to accelerating the adoption of optical Ethernet as the technology of choice in metro networks worldwide.

The Forum is comprised of leading service providers, major incumbent local exchange carriers, top network equipment vendors and other prominent networking companies that share an interest in metro Ethernet. As of December 2003, the MEF had over 60 members.

About the Author

Ralph Santituro, Director of Network Architecture at Nortel Networks, has actively participated in the MEF Technical Committee’s work on Ethernet service definitions and traffic management since December 2001. Mr. Santituro is co-author of the three MEF Ethernet service technical specifications and is co-chair of the MEF Technical Marketing Committee. Mr. Santituro can be contacted at +1 805-527-3024 or rsantitro@nortelnetworks.com

Acknowledgements

The author would like to thank Jay Gill, Director of Product Development at Yipes Enterprise Services, for his contributions to the v1.0 and early drafts of this whitepaper.

The author would also like to thank Bob Klessig, Director of Engineering at Cisco Systems and co-chair of the MEF Technical Committee, for providing feedback on the initial draft of this whitepaper and his contributions to the performance parameter definitions.

Updates to this paper

This paper will be updated as new work emerges from the MEF Technical Committee. Updates versions are available at <http://www.metroethernetforum.org/metro-ethernet-services.pdf>