



Technical Specification

MEF 1

Ethernet Services Model, Phase 1

10 November 2003

Disclaimer

The information in this publication is freely available for reproduction and use by any recipient and is believed to be accurate as of its publication date. Such information is subject to change without notice and the Metro Ethernet Forum (MEF) is not responsible for any errors. The MEF does not assume responsibility to update or correct any information in this publication. No representation or warranty, expressed or implied, is made by the MEF concerning the completeness, accuracy, or applicability of any information contained herein and no liability of any kind shall be assumed by the MEF as a result of reliance upon such information.

The information contained herein is intended to be used without modification by the recipient or user of this document. The MEF is not responsible or liable for any modifications to this document made by any other party.

The receipt or any use of this document or its contents does not in any way create, by implication or otherwise:

- (a) any express or implied license or right to or under any patent, copyright, trademark or trade secret rights held or claimed by any MEF member company which are or may be associated with the ideas, techniques, concepts or expressions contained herein; nor
- (b) any warranty or representation that any MEF member companies will announce any product(s) and/or service(s) related thereto, or if such announcements are made, that such announced product(s) and/or service(s) embody any or all of the ideas, technologies, or concepts contained herein; nor
- (c) any form of relationship between any MEF member companies and the recipient or user of this document.

Implementation or use of specific Metro Ethernet standards or recommendations and MEF specifications will be voluntary, and no company shall be obliged to implement them by virtue of participation in the Metro Ethernet Forum. The MEF is a non-profit international organization accelerating industry cooperation on Metro Ethernet technology. The MEF does not, expressly or otherwise, endorse or promote any specific products or services.

© The Metro Ethernet Forum 2003. All Rights Reserved.

Table of Contents

1. Abstract..... 1

2. Terminology..... 1

3. Scope..... 2

 3.1 Scope of Phase 1 3

 3.2 Scope of Future Phases..... 3

4. Compliance Levels 3

5. Introduction..... 3

6. Ethernet Virtual Connection Service Attributes 4

 6.1 Ethernet Virtual Connection Type Service Attribute 5

 6.1.1 Point-to-Point EVC 5

 6.1.2 Multipoint-to-Multipoint EVC..... 5

 6.2 UNI List Service Attribute 6

 6.3 Service Frame Delivery Service Attributes 6

 6.3.1 Types of Service Frame..... 6

 6.3.1.1 Unicast Service Frame 6

 6.3.1.2 Multicast Service Frame..... 6

 6.3.1.3 Broadcast Service Frame 6

 6.3.1.4 Layer 2 Control Protocol Service Frame 6

 6.3.2 Service Frame Disposition 7

 6.3.3 Service Frame Transparency 7

 6.4 CE-VLAN Tag Preservation Service Attributes 8

 6.4.1 CE-VLAN ID Preservation Service Attribute..... 8

 6.4.2 CE-VLAN CoS Preservation Service Attribute 9

 6.5 EVC Layer 2 Control Protocol Processing Service Attribute 9

 6.6 EVC Related Performance Service Attributes..... 9

7. UNI Service Attributes 10

 7.1 UNI Identifier Service Attribute..... 10

 7.2 Physical Layer Service Attribute..... 10

 7.3 MAC Layer Service Attribute 10

 7.4 Service Multiplexing Service Attribute..... 10

 7.5 Identifying an EVC at the UNI..... 11

 7.5.1 Customer Edge VLAN ID..... 11

 7.5.2 UNI EVC ID Service Attribute 11

 7.6 CE-VLAN ID/EVC Map Service Attribute..... 12

 7.6.1 Basic Concept..... 12

 7.6.2 CE-VLAN ID Significance 13

 7.7 Maximum Number of EVCs Service Attribute 13

 7.8 Bundling Service Attribute..... 13

 7.9 All to One Bundling Service Attribute 14

 7.10 Bandwidth Profiles Service Attributes 15

7.10.1	Standard Bandwidth Profile Parameters	15
7.10.2	Bandwidth Profile Per Ingress UNI Service Attribute	15
7.10.3	Bandwidth Profile Per EVC Service Attribute.....	16
7.10.4	Bandwidth Profile Per Class of Service Service Attribute.....	16
7.10.5	Simultaneous Application of the Bandwidth Profile Application Models.....	16
7.10.6	Actions Taken Upon Bandwidth Profile Violation	16
7.11	Security	17
7.12	UNI Layer 2 Control Protocol Processing Service Attribute	17
7.12.1	Discard	17
7.12.2	Peer.....	17
7.12.3	Pass to EVC.....	17
8.	Ethernet Service Framework.....	18
8.1	Ethernet Service Types	18
8.2	Service Attributes	19
8.3	Service Attribute Parameters	19
8.4	Ethernet Service Framework Summary.....	19
9.	References	20

List of Figures

Figure 1	– Ethernet Services Model.....	4
Figure 2	– Point-to-Point EVCs.....	5
Figure 3	– Multipoint-to-Multipoint EVC	6
Figure 4	– Example of Service Multiplexing on UNI A.....	11
Figure 5	– Example of a CE-VLAN ID/EVC Map.....	12
Figure 6	– Example of CE-VLAN ID/EVC Maps at Two UNIs.....	13
Figure 7	– Example of Bundling.....	14
Figure 8	– Example of a Simple Description of Bundling.....	14
Figure 9	– Example of Ingress Token Bucket Policing per UNI.....	15
Figure 10	– Example of Ingress Token Bucket Policing per EVC	16
Figure 11	– Ethernet Service Framework	18

List of Tables

Table 1	– List of Standardized Layer 2 Control Protocols	7
Table 2	– CE-VLAN ID Preservation.....	8
Table 3	– Possible Physical Layer Characteristics	10
Table 4	– Valid Combinations of Service Multiplexing, Bundling, and All to One Bundling	14
Table 5	– UNI Service Attributes	20
Table 6	– EVC Service Attributes	20

1. Abstract

The attributes of Ethernet Services observable at a User Network Interface (UNI) and from User Network Interface to User Network Interface (UNI to UNI) are defined. All service attributes are defined in this document except for Bandwidth Profile traffic parameters and Performance service attributes and associated parameters. These are to be defined in [1] but are referenced in this document. Finally, a framework for defining specific instances of Ethernet Services is described.

2. Terminology

All to One Bundling	A UNI attribute in which all CE-VLAN IDs are associated with a single EVC.
Bandwidth Profile	A characterization of ingress Service Frame arrival times and lengths at the UNI.
Broadcast Service Frame	A Service Frame that has the broadcast destination MAC address.
Bundling	A UNI attribute in which more than one CE-VLAN ID can be associated with an EVC.
CE	Customer Edge
CE-VLAN CoS	Customer Edge VLAN CoS
CE-VLAN ID	Customer Edge VLAN ID
CE-VLAN ID Preservation	An EVC attribute in which the CE-VLAN ID of an egress Service Frame is identical in value to the CE-VLAN ID of the corresponding ingress Service Frame.
CE-VLAN ID/EVC Map	An association of CE-VLAN IDs with EVCs at a UNI.
CE-VLAN Tag	Customer Edge VLAN Tag
Class of Service	A set of Service Frames that have a commitment from the Service Provider to receive a particular level of performance.
Class of Service Identifier	Information derivable from a Service Frame that allows the identification of the Class of Service instance that applies to the Service Frame.
Customer Edge	Equipment on the Subscriber side of the UNI.
Customer Edge VLAN CoS	The user_priority bits in the IEEE 802.1Q Tag in a tagged Service Frame.
Customer Edge VLAN ID	The identifier derivable from the content of a Service Frame that allows the Service Frame to be associated with an EVC at the UNI.
Customer Edge VLAN Tag	The IEEE 802.1Q Tag in a tagged Service Frame.
E-LAN Service	Ethernet LAN Service
E-Line Service	Ethernet Line Service
Egress Service Frame	A Service Frame sent from the Service Provider network to the CE.
Ethernet LAN Service	An Ethernet Service Type distinguished by its use of a Multipoint-to-Multipoint EVC.
Ethernet Line Service	An Ethernet Service Type distinguished by its use of a Point-to-Point EVC.

Ethernet Virtual Connection	An association of two or more UNIs that limits the exchange of Service Frames to UNIs in the Ethernet Virtual Connection.
EVC	Ethernet Virtual Connection
Frame	Short for Ethernet frame.
Ingress Service Frame	A Service Frame sent from the CE into the Service Provider network.
Layer 2 Control Protocol Service Frame	A Service Frame that is used for Layer 2 control, e.g., Spanning Tree Protocol.
Layer 2 Control Protocol Tunneling	The process by which a Layer 2 Control Protocol Service Frame is passed through the Service Provider network without being processed and delivered unchanged to the proper UNI(s).
Multicast Service Frame	A Service Frame that has a multicast destination MAC address.
Multipoint-to-Multipoint EVC	An EVC with two or more UNIs.
Point-to-Point EVC	An EVC with exactly 2 UNIs.
Service Frame	An Ethernet frame transmitted across the UNI toward the Service Provider or an Ethernet frame transmitted across the UNI toward the Subscriber.
Service Multiplexing	A UNI service attribute in which the UNI can be in more than one EVC instance.
Service Provider	The organization providing Ethernet Service(s).
Subscriber	The organization purchasing and/or using Ethernet Services.
UNI	User Network Interface
Unicast Service Frame	A Service Frame that has a unicast destination MAC address.
User Network Interface	The physical demarcation point between the responsibility of the Service Provider and the responsibility of the Subscriber.

3. Scope

This document describes a model for specifying Ethernet Services. The Ethernet Services are modeled from the point of view of the Subscriber's equipment referred to as the Customer Edge (CE) that is used to access the service. The basic elements of Ethernet Services are defined. In addition, a number of Service Attributes are defined that may be offered as part of an Ethernet Service including the definition of Service Level Specification (details to be defined in [1]).

The goals of the model are two-fold. The first goal is to provide sufficient technical specificity to allow a Subscriber to successfully plan and integrate Ethernet Services into his or her overall networking infrastructure. The second goal is to provide enough detail so that Customer Edge equipment vendors can implement capabilities into their products so that they can be used to successfully access Ethernet Services.

Traffic and performance parameters for Ethernet service level specifications are to be defined in [1]. The definitions of specific instances of Ethernet Service can be based on the content of this service model document.

3.1 Scope of Phase 1

As the title implies, this is Phase 1 of the Technical Specification. The scope of this phase is limited as follows:

- The services considered here are only those based on Ethernet. The various attributes are such that the service given to a particular Ethernet Service Frame (see Section 6.3) is determined by only the contents of the Ethernet protocol and/or lower layers.
- From the Subscriber equipment point of view, the protocol operating at the UNI between the Subscriber's equipment and the Metro Ethernet Network is a Standard Ethernet [3] protocol (PHY and MAC layers).
- The services considered here are limited to services between two or more UNIs. Future phases of this document may define service attributes for other interfaces to the MEN.
- It is assumed that the configuration of both the Service Provider and Subscriber equipment to create and access a service is done administratively. Similarly, resolution of configuration conflicts between the CE and the MEN are done administratively. Control and Management aspects are beyond the scope of Phase 1 of this Technical Specification.

3.2 Scope of Future Phases

Subsequent phases of this Technical specification may address additional service attributes. Possible topics may include but are not limited to services dependent on higher layer protocols such as IP and additional protocols at the UNI such as Ethernet over SONET/SDH and the protocols being developed in IEEE 802.3ah.

4. Compliance Levels

The key words "**MUST**", "**MUST NOT**", "**REQUIRED**", "**SHALL**", "**SHALL NOT**", "**SHOULD**", "**SHOULD NOT**", "**RECOMMENDED**", "**MAY**", and "**OPTIONAL**" in this document are to be interpreted as described in RFC 2119.[2] All key words must be in upper case, bold text.

5. Introduction

This document does not define Ethernet services, but provides the model and framework from which services can be defined and built upon (see Section 8). The model is built on the reference model as shown in Figure 1.




Figure 1 – Ethernet Services Model

The technical definition of a service is in terms of what is seen by each Customer Edge (CE). This includes the User Network Interface (UNI), which is the physical demarcation point between the responsibility of the Service Provider and the responsibility of the Subscriber. A UNI **MUST** be dedicated to a single Subscriber.¹

The CE and MEN exchange Service Frames across the UNI. A Service Frame is an Ethernet [3] frame transmitted across the UNI toward the Service Provider or an Ethernet [3] frame transmitted across the UNI toward the Subscriber. The Service Frame does not include the preamble. Since the protocol as seen by the CE operating at the UNI **MUST** be Standard Ethernet [3], the maximum length Service Frame **MUST** be 1518 bytes when there is no IEEE 802.1Q Tag and it **MUST** be 1522 bytes when there is an IEEE 802.1Q Tag [3]. Subsequent phases of this Technical Specification may specify a larger maximum length.

There are no assumptions about the details of the Metro Ethernet Network. It could consist of a single switch or an agglomeration of networks based on many different technologies.

Connectivity between UNIs is specified by the Ethernet Virtual Connection (EVC). There are a number of types of EVC and a number of service attributes that an EVC can have. These are described in Section 6.

There are a number of different service attributes for a UNI. These are described in Section 7.

Section 8 contains a framework for defining a service. Attributes used in this framework include Ethernet Virtual Connection type, traffic parameters, Service Frame delivery, and performance.

6. Ethernet Virtual Connection Service Attributes

A fundamental aspect of Ethernet Services is the Ethernet Virtual Connection (EVC). An EVC is an instance of an association of two or more UNIs. These UNIs are said to be “in the EVC.” A given UNI can support more than one EVC via the Service Multiplexing attribute as described in Section 7.4.

¹ Multiplexing traffic from multiple Subscribers onto a single link can be a valuable function but is an internal MEN function and is not visible at the UNI.

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 NOT** be delivered back to the ingress UNI. It **MUST NOT** be delivered to a UNI not in the EVC. An EVC is always bi-directional in the sense that Service Frames can originate at any UNI in an EVC.

6.1 Ethernet Virtual Connection Type Service Attribute

There are two types of EVC as described in the following subsections.

6.1.1 Point-to-Point EVC

In a Point-to-Point EVC, exactly two UNIs **MUST** be associated with one another. An ingress Service Frame to the EVC at one UNI **MUST NOT** result in an egress Service Frame at a UNI other than the other UNI in the EVC. (An ingress Service Frame is sent from the CE into the Service Provider network. An egress Service Frame is sent from the Service Provider network to the CE.) The rules under which a Service Frame is delivered to the destination UNI are specific to the particular service definition. Figure 2 illustrates two Point-to-Point EVCs.




Figure 2 – Point-to-Point EVCs

6.1.2 Multipoint-to-Multipoint EVC

In a Multipoint-to-Multipoint EVC, two² or more UNIs **MUST** be associated with one another. An ingress Service Frame to the EVC at one of the UNIs **MUST NOT** result in an egress Service Frame at a UNI that is not in the EVC. The rules under which a frame is delivered to a UNI in the EVC are specific to the particular service definition. Typically, a single broadcast or multi-cast ingress Service Frame (as determined from the destination MAC address) at a given UNI would be replicated in the Metro Ethernet Network and a single copy would be delivered to each of the other UNIs in the EVC. This kind of delivery would also typically apply to a Service Frame for which the MEN has not yet learned an association of the destination MAC address with an EVC, UNI pair. Figure 3 illustrates a Multipoint-to-Multipoint EVC.

² A Multipoint-to-Multipoint EVC with two UNIs is different from a Point-to-Point EVC because one or more additional UNIs can be added to it.




Figure 3 – Multipoint-to-Multipoint EVC

6.2 UNI List Service Attribute

The UNI List for an EVC is a list of UNI Identifiers (see Section 7.1). The list **MUST** have exactly one UNI Identifier for each UNI in the EVC.

6.3 Service Frame Delivery Service Attributes

6.3.1 Types of Service Frame

There are several types of Service Frame.

6.3.1.1 Unicast Service Frame

This is a Service Frame that has a unicast destination MAC address.

6.3.1.2 Multicast Service Frame

This is a Service Frame that has a multicast destination MAC address.

6.3.1.3 Broadcast Service Frame

This is a Service Frame with the broadcast destination MAC address.

6.3.1.4 Layer 2 Control Protocol Service Frame

Given that there are several Layer 2 protocols used for various control purposes, it is important that Metro Ethernet Networks be able to process such information effectively.³ A Layer 2 Control Protocol frame **MUST** be identified by its destination MAC address. For the list of standardized addresses addressed by this Technical Specification, see Table 1. For each MAC address in Table 1, the disposition of a Service Frame with the address as the destination MAC address **MUST** be specified as per sections 6.5 and 7.12. Disposition of Service Frames carrying Layer 2 Control Protocols **MAY** be different for different protocols that use the same destination MAC address, e.g., by destination MAC address and Ethertype.

³ This capability will be especially important for Subscribers who choose to deploy 802.1Q bridges (as opposed to routers) as CEs.

MAC Addresses	Description
0x0180c2000000 through 0x0180c200000f	Bridge Block of protocols
0x0180c2000020 through 0x0180c200002f	GARP Block of protocols
0x0180c2000010	All Bridges Protocol

Table 1 – List of Standardized Layer 2 Control Protocols

A Service Provider **MAY** define additional addresses as Layer 2 Control in addition to those in Table 1.

6.3.2 Service Frame Disposition

The disposition of an Ingress Service Frame is described by one of the following:

- **Discard:** The Service Frame is discarded. An example is a Service Frame containing a particular Layer 2 Control protocol, (e.g., IEEE 802.3X), that is always discarded at the UNI. (See 7.12.) All ingress Service Frames with an invalid FCS **MUST** be discarded by the MEN.
- **Deliver Unconditionally:** No matter what the content (assuming correct FCS) of the Service Frame, it is delivered across the other (egress) UNI(s). This might be the behavior of a Point-to-Point EVC.
- **Deliver Conditionally:** The Service Frame is delivered across an egress UNI if certain conditions are met. An example of such a condition is that the destination MAC address is known by the Metro Ethernet Network to be “at” the destination UNI. Another example is broadcast throttling where some Service Frames with the broadcast destination MAC address are dropped to limit the amount of such traffic. When this option is in force the conditions **MUST** be specified.

More details about the disposition of Layer 2 Control Protocol Service Frames are presented in 6.5 and 7.12.

Note that this is a description of the ideal service. Service Frames that should be delivered might be discarded due to network failure or congestion conditions. The actual behavior of the service compared to the ideal is the measure of the performance. See Section 6.6.

6.3.3 Service Frame Transparency

All fields of each egress Service Frame **MUST** be identical to the same fields of the corresponding ingress Service Frame except as follows:

- The egress Service Frame **MAY** have an IEEE 802.1Q Tag while the corresponding ingress Service Frame does not. In this case the egress Service Frame **MUST** have a recalculated FCS.

- The egress Service Frame **MAY** not have an IEEE 802.1Q Tag while the corresponding ingress Service Frame does have a Tag. In this case the egress Service Frame **MUST** have a recalculated FCS.
- If both the egress Service frame and corresponding ingress Service Frame have an IEEE 802.1Q Tag, the content of the Tag in the egress Service Frame **MAY** be different from the content of the Tag in the corresponding ingress Service Frame. If the contents of the ingress and egress tags are different, the egress Service Frame **MUST** have a recalculated FCS.

However, specific attributes of an EVC **MAY** enforce the condition that additional fields must be unchanged from ingress to egress. See Section 6.4.

6.4 CE-VLAN Tag Preservation Service Attributes

Service Frames at the UNI may contain an IEEE 802.1Q Tag. Such a Tag is referred to as a Customer Edge VLAN Tag (CE-VLAN Tag). The portion of the CE-VLAN Tag that identifies a VLAN indicates the Customer Edge VLAN ID (CE-VLAN ID). (See Section 7.5.) The portion of the CE-VLAN Tag that contains the user_priority bits is called the Customer Edge VLAN CoS (CE-VLAN CoS). An EVC **MAY** have two attributes related to CE-VLAN Tag Preservation as described in the following two subsections.

6.4.1 CE-VLAN ID Preservation Service Attribute

In an EVC with CE-VLAN ID Preservation, the relationship between the ingress Service Frame and its corresponding egress Service Frame(s) described in Table 2 **MUST** be maintained.

Ingress Service Frame	Egress Service Frame(s) ⁴
No IEEE 802.1Q Tag	No IEEE 802.1Q Tag
Contains IEEE 802.1Q Tag	Contains IEEE 802.1Q Tag with VLAN ID equal to the VLAN ID of the Tag on the ingress Service Frame ⁵

Table 2 – CE-VLAN ID Preservation

When an EVC includes a UNI at which Bundling occurs for the EVC (see Section 7.8), the EVC **MUST** have the CE-VLAN ID Preservation Service Attribute. When an EVC includes a UNI at which All to One Bundling occurs (see 7.9), the EVC **MUST** have the CE-VLAN ID Preservation Service Attribute.

⁴ Note that in the case of a Multipoint-to-Multipoint EVC, a single ingress Service Frame can result in more than one egress Service Frame.

⁵ Note that, depending on the internal network implementation, a Service Provider may not be able to preserve all values of VLAN ID. If this is the case, the CE-VLAN ID/EVC map (see Section 7.6) would not map Service Frames with the VLAN ID to an EVC that has the CE-VLAN ID Preservation attribute.

An obvious benefit of the CE-VLAN ID Preservation feature is enhanced operational simplicity; for Subscribers using IEEE 802.1Q bridges, the feature obviates the task of renumbering VLANs in different corporate campuses.

6.4.2 CE-VLAN CoS Preservation Service Attribute

In an EVC with CE-VLAN CoS Preservation, an egress Service Frame resulting from an ingress Service Frame that contains a CE_VLAN CoS **MUST** have the identical CE-VLAN CoS. CE-VLAN CoS Preservation is independent of Bundling and All to One Bundling.

6.5 EVC Layer 2 Control Protocol Processing Service Attribute

In some cases, 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 Service Provider network without being processed⁶ and delivered to the proper UNI or UNIs. The tunneling capability can be extremely useful, for example, when the Subscriber chooses to attach bridges to all UNIs and thus BPDUs need to be carried across the Network. When a Layer 2 control protocol is tunneled, the Service Frame at each egress UNI **MUST** be identical to the corresponding ingress Service Frame. Since the Layer 2 control protocols in Table 1 are all untagged, this means that at the egress UNI they **MUST** also be untagged.

For a given EVC at a given UNI, the Service Provider defines which Layer 2 control protocols will be tunneled via the EVC and which will be discarded. Since the Layer 2 control protocols in Table 1 are all untagged, if they are tunneled, they **MUST** be tunneled on the EVC that is mapped to the untagged CE-VLAN ID in the CE-VLAN ID/EVC Map. See Section 7.6.

Note that if a Layer 2 Control Protocol is to be tunneled, then all UNIs in the EVC **MUST** be configured to pass the Layer 2 Control Protocol to the EVC. (See Section 7.12.3.)

6.6 EVC Related Performance Service Attributes

Service Frame delivery performance is specified for all Service Frames on an EVC with a particular Class of Service instance. The Class of Service instance is identified by a Class of Service Identifier associated with each Service Frame. The Class of Service Identifier **MUST** be derived from either:

- The EVC to which the Service Frame is mapped, or
- The combination of the EVC to which the Service Frame is mapped and a set of one or more CE-VLAN CoS values.

⁶ For example the Subscriber's Ethernet information can be encapsulated in another frame separate from the control protocol frame.

7. UNI Service Attributes

A UNI can have a number of characteristics that will be important to the way that the CE sees a service. One of the key aspects of a service description will be the allowable mix of UNIs with different characteristics in an EVC. For example, a specific (simple) service might require all UNIs to have the same speed at the physical layer. A more sophisticated service may allow a wide variety of speeds.

7.1 UNI Identifier Service Attribute

The UNI Identifier attribute is a value that is assigned to the UNI by the Service Provider. It MAY have any string as a value. The UNI Identifier MUST be unique among all UNIs for the MEN. As an example, the Service Provider might use "SCPOP1-Node3-Slot2-Port1" as a UNI Identifier and this could signify Port 1 in Slot 2 of Node 3 in Santa Clara POP1.

7.2 Physical Layer Service Attribute

For a UNI, the Speed (in bits per second), Mode, and Physical Medium **MUST** be one of the combinations shown in Table 3. Typically there are no constraints in mixing UNIs with different physical media in the same EVC.⁷ Constraints on the mix of speeds and modes **MAY** be imposed for some services.

Speed	Mode	Physical Medium
10 Mbps	Full duplex	All Ethernet physical media compatible with Speed and Mode listed in [3].
100 Mbps	Full duplex	
10/100 Mbps	Auto negotiation	
1 Gbps	Full duplex	
10 Gbps	Full duplex	

Table 3 – Possible Physical Layer Characteristics

7.3 MAC Layer Service Attribute

A UNI **MUST** support the IEEE 802.3-2002 [3] frame formats.

7.4 Service Multiplexing Service Attribute

A UNI with the Service Multiplexing attribute **MUST** be configurable to support multiple EVCs.⁸ Point-to-Point EVCs and Multipoint-to-Multipoint EVCs **MAY** be multiplexed in any combination at a UNI. Figure 4 shows an example of Service Multiplexing. In this example, CE A is attached to the network via a Gigabit Ethernet UNI. CEs B, C, and D are attached via 100

⁷ An exception might be wireless when the service requires stringent requirements on packet loss.

⁸ Since the UNI is dedicated to a single Subscriber, only one Subscriber can access the EVCs at the UNI.

Mbps Ethernet UNIs. Using Service Multiplexing, instances of Point-to-Point EVCs to each of B, C, and D can be implemented at A without requiring 3 physical ports on the CE at A.




Figure 4 – Example of Service Multiplexing on UNI A

7.5 Identifying an EVC at the UNI

7.5.1 Customer Edge VLAN ID

The EVC for a Service Frame **MUST** be identified by the Customer Edge VLAN ID (CE-VLAN ID). There are 4095 CE-VLAN IDs numbered 1 through 4095. The CE-VLAN ID for a Service Frame with an IEEE 802.1Q tag **MUST** be the value of the VLAN ID, if not 0, in the tag. Untagged and priority tagged⁹ Service Frames **MUST** have the same CE-VLAN ID whose value **MUST** be configurable to any value in the range 1, 2, ..., 4094. When the CE-VLAN ID Preservation Service Attribute (see Section 6.4.1) is not in force for an EVC, egress Service Frames with the CE-VLAN ID value for untagged and priority tagged Service Frames **MUST** be untagged.

More than one CE-VLAN ID may point to the same EVC as described in Section 7.8.

Note that certain of the VLAN ID values in IEEE 802.1Q Tags are reserved for special purposes in IEEE 802.1Q bridges and thus the number of VLANs in a subscriber network is less than 4095. Nevertheless, Service Frames with any VLAN ID value as well as untagged Service Frames can exist at the UNI. Consequently the CE-VLAN ID can have 4095 values. However, less than 4095 EVCs **MAY** be supported at a UNI. See Section 7.6.

7.5.2 UNI EVC ID Service Attribute

The UNI EVC ID is an arbitrary string administered by the Service Provider that is used to identify an EVC at the UNI. It is intended for management and control purposes. As an example, the

⁹ A priority tagged Service Frame is a Service Frame with an IEEE 802.1Q tag whose VLAN ID in the tag equals 0.

Service Provider might use “SCPOP1-Node3-Slot2-Port1 – 0017” as a UNI EVC ID and this could signify the 17th EVC provisioned at the UNI identified by Port 1 in Slot 2 of Node 3 in Santa Clara POP1.

7.6 CE-VLAN ID/EVC Map Service Attribute

7.6.1 Basic Concept

At each UNI there **MUST** be a mapping of each CE-VLAN ID to at most one EVC called the CE-VLAN ID/EVC Map. Note that a given CE-VLAN ID **MAY** not be mapped to any EVC. In the simple case, when the Bundling and All to One Bundling attributes (as defined in Sections 7.8 and 7.9) are not invoked, exactly one CE-VLAN ID **MUST** be mapped to at most one EVC. Figure 5 is an example of a CE-VLAN ID/EVC Map.

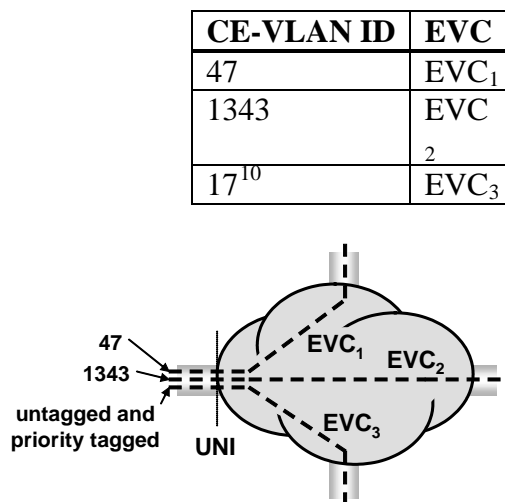


Figure 5 – Example of a CE-VLAN ID/EVC Map

In this example, an ingress Service Frame with CE-VLAN ID 47 is transported according to the properties and attributes of EVC₁. An untagged or priority tagged ingress Service Frame is transported according to the properties and attributes of EVC₃. An egress frame coming from EVC₂ is given CE-VLAN ID 1343.

When an instance of the CE-VLAN ID/EVC Map does not contain an entry for a given CE-VLAN ID, any ingress Service Frame at the UNI with this instance of the map with this CE-VLAN ID **MUST** be discarded by the MEN. And, an egress Service Frame **MUST NOT** have a CE-VLAN ID with this value at the UNI with this instance of the CE-VLAN ID/EVC Map.

In some scenarios, it may be necessary for the Subscriber and the Service Provider to agree upon the CE-VLAN ID/EVC Map at the UNI. One way to implement this is to have the Service Provider dictate the mapping. This is what is frequently done with the mapping between DLCIs and

¹⁰ In this example, the CE-VLAN ID for untagged and priority tagged Service Frames is configured to 17.

permanent virtual connections for Frame Relay. Also note that for a given UNI, the CE-VLAN ID/EVC Map may be constrained by the range of CE-VLAN ID values that can be supported by the CE and the range of CE-VLAN ID values that can be supported by the Service Provider.¹¹

7.6.2 CE-VLAN ID Significance

CE-VLAN ID values **MAY** only be significant at a given UNI. Restated, the CE-VLAN ID/EVC mapping for a given UNI in an EVC **MAY** be different from the mapping at another UNI in the EVC. Figure 6 shows valid CE-VLAN ID/EVC Maps for three EVCs between two UNIs. Note that when the CE-VLAN ID Preservation attribute (Section 6.4.1) applies to an EVC, the mappings for the EVC are identical as is the case for EVC₁ in Figure 6. (Otherwise the CE-VLAN ID cannot be preserved).

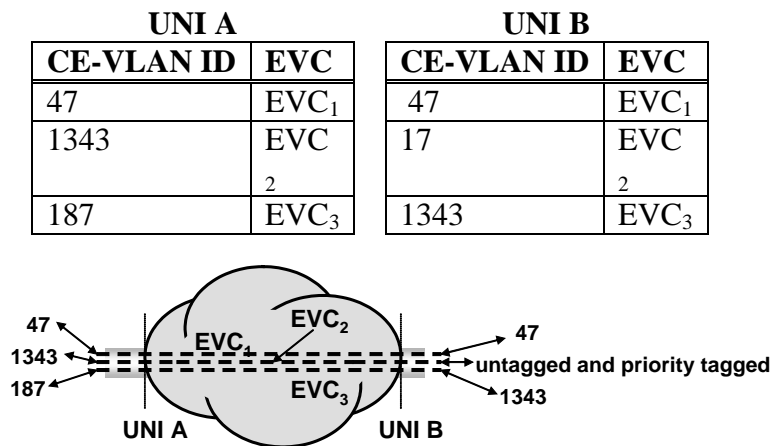


Figure 6 – Example of CE-VLAN ID/EVC Maps at Two UNIs

7.7 Maximum Number of EVCs Service Attribute

This attribute defines the maximum number of EVCs that the UNI can support. It **MUST** have a value of at least one.

7.8 Bundling Service Attribute

When a UNI has the Bundling attribute, it **MUST** be configurable so that more than one CE-VLAN ID can map to an EVC at the UNI. An EVC with more than one CE-VLAN ID mapping to it **MUST** have the CE-VLAN ID Preservation Service Attribute (see Section 6.4.1) and the list of CE-VLAN IDs mapped to the EVC **MUST** be the same at each UNI in the EVC. Figure 7 shows an example of Bundling. In this example, UNI A and UNI B have the bundling feature as

¹¹ In later Phases of this Technical Specification, dynamic EVC setup via a signaling protocol across the UNI may be specified. In that case, it may be desirable to specify the range of CE-VLAN ID values supported by the Service Provider as a UNI attribute. In this phase of this Technical Specification, resolving the CE-VLAN ID/EVC Map is assumed to be done administratively and thus this specifying of a CE-VLAN ID range is not needed.

seen from the mapping for EVC₁. (EVC₁ has the CE-VLAN ID Preservation feature.). Note that Bundling is compatible with Service Multiplexing. In Figure 7, UNI A and UNI B are examples of Service Multiplexing and Bundling on the same UNI.

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

Figure 7 – Example of Bundling

This model does not constrain the way that the Service Provider and Subscriber communicate the contents of the CE-VLAN ID/EVC map. For example, a Service Provider could simply describe bundling as shown in Figure 8.

Description		Actual Map	
CE-VLAN ID	EVC	CE-VLAN ID	EVC
2000	EVC ₁	2000	EVC ₁
2001	EVC ₃	2001	EVC ₃
All others	EVC ₄	1, ..., 1999, 2002, ..., 4095	EVC ₄

Figure 8 – Example of a Simple Description of Bundling

7.9 All to One Bundling Service Attribute

When a UNI has the All to One Bundling attribute, all CE-VLAN IDs **MUST** map to a single EVC at the UNI. The EVC at the UNI **MUST** have the CE-VLAN ID Preservation Service Attribute (see Section 6.4.1) and the list of CE-VLAN IDs mapped to the EVC **MUST** be the same at each UNI in the EVC. It follows that such a UNI cannot have Service Multiplexing.

All to One Bundling is a special case of Bundling but it is sufficiently important to be called out as a separate attribute.

Table 4 shows the valid combinations of the bundling and Service Multiplexing attributes.

	Valid Combination 1	Valid Combination 2	Valid Combination 3
Service Multiplexing	⊖	⊖	○
Bundling	○	⊖	○
All to One Bundling	○	○	⊖

Table 4 – Valid Combinations of Service Multiplexing, Bundling, and All to One Bundling

7.10 Bandwidth Profiles Service Attributes

A Bandwidth Profile is a characterization of the lengths and arrival times for ingress Service Frames at the UNI. When a Bandwidth Profile is applied to a given sequence of ingress Service Frames, each Service Frame in the sequence is declared to be compliant or not compliant with the Bandwidth Profile. A Bandwidth Profile includes a specification of the disposition of ingress Service Frames that do not comply with the Bandwidth Profile. Phase 1 of this Technical Specification only defines discard actions. The discard actions are to be defined in [1]. Subsequent phases may specify additional actions.

Typically a Bandwidth Profile defines Service Frame traffic that is less than the full bandwidth of the UNI. Thus the Bandwidth Profile can be considered to be analogous to the traffic policing of Frame Relay.[4]

Bandwidth Profiles are associated with the UNI. This allows different Bandwidth Profiles at each UNI in an EVC as in Section 7.10.3. For example, on a Multipoint-to-Multipoint EVC, a different Bandwidth Profile could apply at each UNI in the EVC. To describe this situation on an EVC basis would require the specification of a vector of Bandwidth Profiles, one for each UNI. Thus, to simplify the description, Bandwidth Profiles are specified as a UNI service attribute.

7.10.1 Standard Bandwidth Profile Parameters

The normative definition of a Bandwidth Profile will be provided in [1]. This definition is expected to be based on Token Bucket algorithms. The parameters for these algorithms will be specified in [1].

7.10.2 Bandwidth Profile Per Ingress UNI Service Attribute

In this model, a single Bandwidth Profile **MUST** be applied to all ingress Service Frames at the UNI. Figure 9 illustrates an example of the application of ingress policing with a per ingress UNI Token Bucket.




Figure 9 – Example of Ingress Token Bucket Policing per UNI

7.10.3 Bandwidth Profile Per EVC Service Attribute

In this model, a single Bandwidth Profile **MUST** be applied to all ingress Service Frames for an instance of an EVC at the UNI. Thus, if a UNI has 5 Ethernet Virtual Connections, there could be 5 ingress Bandwidth Profiles, one for each EVC. Figure 10 illustrates an example of the application of ingress policing with per EVC Token Buckets.




Figure 10 – Example of Ingress Token Bucket Policing per EVC

7.10.4 Bandwidth Profile Per Class of Service Service Attribute

In this model, a single Bandwidth Profile **MUST** be applied to all ingress Service Frames with a specific Class of Service Identifier. Class of Service Identifiers are specified in Section 6.6. For example, within a given EVC, two different Classes of Service could be available with the CE-VLAN CoS identifying the class for each Service Frame. In this case, a separate Bandwidth Profile for each class could be applied. That is, there would be two bandwidth profiles per EVC and the bandwidth profile to apply to a Service Frame would be determined by its Class of Service Identifier.

7.10.5 Simultaneous Application of the Bandwidth Profile Application Models

Multiple models of Bandwidth Profile application **MAY** exist simultaneously at a UNI. However, a UNI **MUST** be configured such that only a single Bandwidth Profile applies to any given Service Frame.

7.10.6 Actions Taken Upon Bandwidth Profile Violation

The definition of a Bandwidth Profile includes the specification of 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 bit in ATM.¹²

The details of the action taken are to be defined in [1].

7.11 Security

In Phase 1, the Ethernet Virtual Connection is the fundamental service construct that defines how the separation between Subscribers' traffic is maintained. Additional security constructs and service attributes may be addressed in subsequent phases.

7.12 UNI Layer 2 Control Protocol Processing Service Attribute

There are three alternatives for processing a given Layer 2 Control Protocol (see Table 1) at a UNI as described in the following subsections.¹³

7.12.1 Discard

When this alternative is in force, the MEN **MUST** discard all ingress Service Frames carrying the Layer 2 Control Protocol and the MEN **MUST NOT** generate any egress Service Frames carrying the Layer 2 Control Protocol. Note that when this alternative is in force for the Layer 2 Control Protocol, the Layer 2 Control Protocol cannot be processed by an EVC. See Section 6.5.

7.12.2 Peer

When this alternative is in force, the MEN **MUST** act as a peer of the CE in the operation of the Layer 2 Control Protocol. From the CE point of view, the MEN is a single device that is running the Layer 2 Control Protocol. Where the protocol is terminated in the MEN is an internal network implementation issue and beyond the scope of this Technical Specification. Note that when this alternative is in force for the Layer 2 Control Protocol, the Layer 2 Control Protocol cannot be processed by an EVC. See Section 6.5.

7.12.3 Pass to EVC

When this alternative is in force, the tunneling of the Layer 2 Control Protocol **MUST** be determined by the Layer 2 Control Protocol Processing attribute of the EVC (tunneled or discarded)

¹² The implementation of priority discard is an internal network issue and beyond the scope of this specification. However, an example is the setting of an indication field in a protocol that encapsulates the Service Frame while it is in the Service Provider network. Such an encapsulation could be the outer VLAN tag in a double tagged frame and the indication field could be the user_priority bits of the tag.

¹³ There is a potential for an additional alternative beyond the three described in this document. This alternative would have the MEN both participate and tunnel a given Layer 2 control protocol. This may be included in a future phase of this Technical Specification.

to which the untagged CE-VLAN ID is mapped by the CE-VLAN ID/EVC map.¹⁴ See Section 6.5.

8. Ethernet Service Framework

The Ethernet service framework provides the definition and relationship between attributes and their associated parameters used to create an Ethernet Service. An Ethernet Service consists of (Refer to Figure 11):

- One Ethernet Service Type,
- One or more Ethernet Service Attributes and
- One or more parameter values associated with each Ethernet Service Attribute.

The Service Framework associates a service with the UNI characteristics across which the Service is offered to the Subscriber and with the EVC supporting the service. The Ethernet Service Attributes are what define the UNI and EVC characteristics.




Figure 11 – Ethernet Service Framework

8.1 Ethernet Service Types

Ethernet Service Types are the fundamental service constructs. Ethernet Service Types can be used to create a broad range of services. Each Ethernet Service Type has a set of Ethernet Service Attributes that define the service characteristics. These Ethernet Service Attributes in turn have a set of parameters associated with them that provide various options for the different Service Attributes. Refer to Figure 11. Two Ethernet Service Types have been defined. The first, Ethernet Line Service (E-Line Service) uses a Point-to-Point EVC. The second, Ethernet LAN Service (E-LAN Service), uses a Multipoint-to-Multipoint EVC.

¹⁴ The use of the EVC corresponding to the untagged CE-VLAN ID is because all of the Layer 2 Control Protocols in Table 1 are carried in untagged frames. Should a future Layer 2 Control Protocol be supported that is carried in tagged frames, then the EVC should be determined by the mapping for the CE-VLAN ID carried in the tag.

8.2 Service Attributes

The Service Attributes define the capabilities of the Ethernet Service Type. Some or all of the Service Attributes may apply to an Ethernet Service Type. Service Attributes are described in Section 6 and Section 7.

8.3 Service Attribute Parameters

For each Service Attribute there can be one or more parameters that specify the attribute. Parameters can have various types of values including:

- Logical (true or false)
- Integer
- Bandwidth
- Protocol
- Vector of values of multiple types
- Character String.

8.4 Ethernet Service Framework Summary

For a particular Ethernet Service Type, there are two types of Service Attributes, those that apply to a UNI as described in Section 7 and those that apply to an EVC as described in Section 6. The UNI Service Attributes are listed in Table 5 along with the type of parameter value for the attribute. For a given instance of a service, a table like that of Table 5 **MUST** be specified for each UNI in the EVC associated with the service.

Attribute	Type of Parameter Value
UNI Identifier (7.1)	Any string
Physical Medium (7.2)	A Standard Ethernet [3] PHY ¹⁵
Speed (7.2)	10 Mbps, 100 Mbps, 1 Gbps, or 10 Gbps ¹⁵
Mode (7.2)	Full Duplex or Auto negotiation ¹⁵
MAC Layer (7.3)	IEEE 802.3 – 2002 [3]
Service Multiplexing (7.4)	Yes or No ¹⁶
UNI EVC ID (7.5.2)	An arbitrary string for the EVC supporting the service instance
CE-VLAN ID/EVC Map (7.6)	Map as per 7.6

¹⁵ There are interdependencies among the values of these parameters as per the IEEE 802.3 Standard.[3]

¹⁶ Must be No if All to One Bundling is Yes.

Attribute	Type of Parameter Value
Maximum Number of EVCs (7.7)	Integer ³ 1
Bundling (7.8)	Yes or No ¹⁶
All to One Bundling (7.9)	Yes or No ¹⁷
Ingress Bandwidth Profile Per Ingress UNI (7.10.2)	No or as to be defined in [1]
Ingress Bandwidth Profile Per EVC (7.10.3)	No or as to be defined in [1]
Ingress Bandwidth Profile Per Class of Service Identifier (7.10.4)	No or as to be defined in [1]
Layer 2 Control Protocols Processing (7.12)	All entries from Table 1 with each being labeled Discard, Peer, or Pass to EVC

Table 5 – UNI Service Attributes

The EVC Service Attributes are listed in Table 6 along with the type of parameter value for the attribute. For a given instance of a service, a table like that of Table 6 **MUST** be specified for the EVC associated with the service.

Attribute	Type of Parameter Value
EVC Type (6.1)	Point-to-Point or Multipoint-to-Multipoint
UNI List (6.2)	A list of UNI Identifiers (7.1)
CE-VLAN ID Preservation (6.4.1)	Yes or No
CE-VLAN CoS Preservation (6.4.2)	Yes or No
Unicast Service Frame Delivery (6.3.1.1)	Discard, Deliver Unconditionally, or Deliver Conditionally. If Deliver Conditionally is used, then the conditions MUST be specified.
Multicast Service Frame Delivery (6.3.1.2)	Discard, Deliver Unconditionally, or Deliver Conditionally. If Deliver Conditionally is used, then the conditions MUST be specified.
Broadcast Service Frame Delivery (6.3.1.3)	Discard, Deliver Unconditionally, or Deliver Conditionally. If Deliver Conditionally is used, then the conditions MUST be specified.
Layer 2 Control Protocols Processing (6.5)	Entries from Table 1 labeled Tunnel or Discard.
EVC Performance (6.6)	To be defined in [1].

Table 6 – EVC Service Attributes

9. References

[1] Metro Ethernet Forum, *Traffic Management Specification – Phase 1*, to be published.

¹⁷ Must be No if Bundling is Yes or Service Multiplexing is Yes.

- [2] Bradner, S., *Key words for use in RFCs to Indicate Requirement Levels*, RFC 2119, March 1997. (Normative)
- [3] IEEE P 802.3 – 2002, *Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications*, 8 March 2002. (Normative)
- [4] International Telecommunication Union, Recommendation I.370, *Integrated Services Digital Network (ISDN), Overall Network Aspects And Functions, ISDN User-Network Interfaces, Congestion Management For The ISDN Frame Relaying Bearer Service*, 1991.