



Technical Specification

MEF 10

Ethernet Services Attributes Phase 1

(Obsoletes MEF 1 and MEF 5)

November 2004

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 2005. All Rights Reserved.

Table of Contents

1. Abstract..... 1

2. Terminology..... 1

3. Scope..... 4

 3.1 Scope of Phase 1 5

 3.2 Scope of Future Phases 5

4. Compliance Levels 5

5. Introduction..... 5

6. Ethernet Virtual Connection Service Attributes 6

 6.1 Ethernet Virtual Connection Type Service Attribute 7

 6.1.1 Point-to-Point EVC 7

 6.1.2 Multipoint-to-Multipoint EVC..... 7

 6.2 EVC ID Service Attribute..... 8

 6.3 UNI List Service Attribute 8

 6.4 Service Frame Delivery Service Attributes 8

 6.4.1 Types of Service Frame..... 8

 6.4.1.1 Unicast Service Frame 8

 6.4.1.2 Multicast Service Frame..... 8

 6.4.1.3 Broadcast Service Frame 8

 6.4.1.4 Layer 2 Control Protocol Service Frame 9

 6.4.2 Service Frame Disposition 9

 6.4.3 Service Frame Transparency 10

 6.5 CE-VLAN Tag Preservation Service Attributes 10

 6.5.1 CE-VLAN ID Preservation Service Attribute..... 10

 6.5.2 CE-VLAN CoS Preservation Service Attribute 11

 6.6 EVC Layer 2 Control Protocol Processing Service Attribute 11

 6.7 EVC Related Performance Service Attributes..... 12

 6.7.1 Frame Delay Performance for a Point-to-Point EVC..... 12

 6.7.2 Frame Delay Performance for a Multipoint-to-Multipoint EVC 14

 6.7.3 Frame Delay Variation Performance a Point-to-Point EVC 14

 6.7.4 Frame Delay Variation Performance for a Multipoint-to-Multipoint EVC 16

 6.7.5 Frame Loss Ratio Performance for a Point-to-Point EVC 16

 6.7.6 Frame Loss Ratio Performance for a Multipoint-to-Multipoint EVC..... 17

7. UNI Service Attributes 17

 7.1 UNI Identifier Service Attribute..... 17

 7.2 Physical Layer Service Attribute 17

 7.3 MAC Layer Service Attribute 18

 7.4 Service Multiplexing Service Attribute..... 18

 7.5 Identifying an EVC at the UNI..... 19

 7.5.1 Customer Edge VLAN ID..... 19

 7.5.2 UNI EVC ID Service Attribute 19

7.6	CE-VLAN ID/EVC Map Service Attribute.....	19
7.6.1	Basic Concept.....	19
7.6.2	CE-VLAN ID Significance	20
7.7	Maximum Number of EVCs Service Attribute	21
7.8	Bundling Service Attribute.....	21
7.9	All to One Bundling Service Attribute	22
7.10	Bandwidth Profiles Service Attributes	23
7.10.1	Standard Bandwidth Profile Parameters	23
7.10.2	Enforcement of Bandwidth Profile Parameters.....	24
7.10.2.1	<i>Bandwidth Profile Algorithm</i>	24
7.10.2.2	<i>Service Frame Disposition</i>	26
7.10.3	Bandwidth Profile Per Ingress UNI Service Attribute	26
7.10.4	Bandwidth Profile Per EVC Service Attribute.....	26
7.10.5	Bandwidth Profile Per Class of Service Service Attribute.....	27
7.10.6	Simultaneous Application of the Bandwidth Profile Application Models.....	27
7.11	Security	28
7.12	UNI Layer 2 Control Protocol Processing Service Attribute	28
7.12.1	Discard	28
7.12.2	Peer.....	28
7.12.3	Pass to EVC.....	28
8.	Ethernet Service Framework.....	29
8.1	Ethernet Service Types	29
8.2	Service Attributes	29
8.3	Service Attribute Parameters	30
8.4	Ethernet Service Framework Summary.....	30
9.	References.....	32

List of Figures

Figure 1	– Ethernet Services Model.....	6
Figure 2	– Point-to-Point EVCs.....	7
Figure 3	– Multipoint-to-Multipoint EVC	8
Figure 4	– Frame Delay for Service Frame	13
Figure 5	– Frame Delay Variation Parameters	15
Figure 6	– Example of Service Multiplexing on UNI A.....	18
Figure 7	– Example of a CE-VLAN ID/EVC Map.....	20
Figure 8	– Example of CE-VLAN ID/EVC Maps at Two UNIs.....	21
Figure 9	– Example of Bundling.....	22
Figure 10	– Example of a Simple Description of Bundling.....	22
Figure 11	– The Bandwidth Profile Algorithm.....	25
Figure 12	– Ingress Bandwidth Profile per Ingress UNI	26
Figure 13	– Ingress Bandwidth Profile per EVC	27
Figure 14	– Ingress Bandwidth Profile per EVC and CE-VLAN CoS.....	27
Figure 15	– Ethernet Service Framework	29

List of Tables

Table 1 – List of Standardized Layer 2 Control Protocols	9
Table 2 – CE-VLAN ID Preservation.....	11
Table 3 – Frame Delay Performance Parameters	13
Table 4 – Frame Delay Variation Parameters.....	16
Table 5 – Frame Loss Ratio Performance Parameters.....	17
Table 6 – Possible Physical Layer Characteristics	18
Table 7 – Valid Combinations of Service Multiplexing, Bundling, and All to One Bundling	23
Table 8 – Service Frame Disposition.....	26
Table 9 – UNI Service Attributes	31
Table 10 – EVC Service Attributes	31

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. In addition, a framework for defining specific instances of Ethernet Services is described. This document supersedes and replaces MEF 1[5] and MEF 5[6].

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 a reference point and a specification of the disposition of each Service Frame based on its level of compliance with the Bandwidth Profile. In this document the reference point is 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.
CBS	Committed Burst Size
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
CF	Coupling Flag
CIR	Committed Information Rate
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) the EVC to which the Service Frame is mapped or b) the combination of the EVC to which the Service Frame is mapped and a set of one or more CE-VLAN CoS values.

CM	Color Mode
Color Mode	CM is a Bandwidth Profile parameter. The Color Mode parameter indicates whether the color-aware or color-blind property is employed by the Bandwidth Profile. It takes a value of “color-blind” or “color-aware” only.
Color-aware	A Bandwidth Profile property where a pre-determined level of Bandwidth Profile compliance for each Service Frame is taken into account when determining the level of compliance for each Service Frame.
Color-blind	A Bandwidth Profile property where a pre-determined level of Bandwidth Profile compliance for each Service Frame, if present, is ignored when determining the level of compliance for each Service Frame.
Committed Burst Size	CBS is a Bandwidth Profile parameter. It limits the maximum number of bytes available for a burst of ingress Service Frames sent at the UNI speed to remain CIR-conformant.
Committed Information Rate	CIR is a Bandwidth Profile parameter. It defines the average rate in bits/s of ingress Service Frames up to which the network delivers Service Frames and meets the performance objectives defined by the CoS Service Attribute.
Coupling Flag	CF is a Bandwidth Profile parameter. The Coupling Flag allows the choice between two modes of operations of the rate enforcement algorithm. It takes a value of 0 or 1 only.
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 Service Frame that is either tagged or priority tagged.
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.
EBS	Excess Burst Size
Egress Service Frame	A Service Frame sent from the Service Provider network to the CE.
EIR	Excess Information Rate
E-LAN Service	Ethernet LAN Service
E-Line Service	Ethernet Line Service
Ethernet LAN Service	An Ethernet Service Type distinguished by its use of a Multi-point-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
Excess Burst Size	EBS is a Bandwidth Profile parameter. It limits the maximum number of bytes available for a burst of ingress Service Frames sent at the UNI speed to remain EIR-conformant.
Excess Information Rate	EIR is a Bandwidth Profile parameter. It defines the average rate in bits/s of ingress Service Frames up to which the network may deliver Service Frames without any performance objectives.
FD	Frame Delay
FDV	Frame Delay Variation
FLR	Frame Loss Ratio
Frame	Short for Ethernet frame.
Frame Delay	The time required to transmit a Service Frame from source to destination across the metro Ethernet network.
Frame Delay Performance	A measure of the delays experienced by different Service Frames belonging to the same CoS instance.
Frame Delay Variation	The difference in delay of two Service Frames.
Frame Delay Variation Performance	A measure of the variation in the delays experienced by different Service Frames belonging to the same CoS instance.
Frame Loss Ratio Performance	Frame Loss Ratio is a measure of the number of lost frames inside the MEN. Frame Loss Ratio is expressed as a percentage.
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 is 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. 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.

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 Level Agreement	The contract between the Subscriber and Service Provider specifying the agreed to service level commitments and related business agreements.
Service Level Specification	The technical specification of the service level being offered by the Service Provider to 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).
SLA	Service Level Agreement
SLS	Service Level Specification
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 Ethernet Service attributes. 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. This document supersedes and replaces MEF 1, *Ethernet Services Model, Phase 1* [5] and MEF 5, *Traffic Management Specification: Phase 1* [6].

The goals of this Technical Specification 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. It follows as a corollary that vendors of Service Provider network equipment will make use of this information for implementing functions that complement the functions in the CE.

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 6.4) 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 [2],[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 additional IEEE 802.3 protocols.

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[1]. All key words must be in upper case, bold text.

5. Introduction

This document provides the model and framework for Ethernet Services. 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 [2] frame transmitted across the UNI toward the Service Provider or an Ethernet [2] frame transmitted across the UNI toward the Subscriber. The Service Frame consists of the first bit of the Destination MAC Address through the last bit of the Frame Check Sequence. Since the protocol as seen by the CE operating at the UNI **MUST** be standard Ethernet [2], 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 [2]. 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 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. They are 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 mapped (see Section 7.6) 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 mapped 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 multicast 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 the Multipoint-to-Multipoint EVC.




Figure 3 – Multipoint-to-Multipoint EVC

6.2 EVC ID Service Attribute

The EVC ID is an arbitrary string administered by the Service Provider that is used to identify an EVC within the MEN. The EVC ID **MUST** be unique across all EVCs in the MEN. It is intended for management and control purposes. The EVC ID is not carried in any field in the Service Frame. As an example, the Acme Service Provider might use “EVC-0001898-ACME-MEGAMART” to represent the 1898th EVC in the MEN and the customer for the EVC is MegaMart.

6.3 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.

6.4 Service Frame Delivery Service Attributes

6.4.1 Types of Service Frame

There are several types of Service Frame.

6.4.1.1 *Unicast Service Frame*

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

6.4.1.2 *Multicast Service Frame*

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

6.4.1.3 *Broadcast Service Frame*

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

6.4.1.4 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 6.6 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. [7] contains some recommendations for the delivery of specific Layer 2 Control protocols.

MAC Addresses ⁴	Description
01-80-C2-00-00-00 through 01-80-C2-00-00-0F	Bridge Block of protocols
01-80-C2-00-00-20 through 01-80-C2-00-00-2F	GARP Block of protocols
01-80-C2-00-00-10	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.

6.4.2 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 identified by 01-80-C2-00-00-01), that is always discarded at the UNI. (See Section 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.

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

⁴Hexadecimal canonical format

More details about the disposition of Layer 2 Control Protocol Service Frames are presented in Sections 6.6 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. See the EVC Related Performance Service Attributes in Section 6.7.

6.4.3 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 identical at ingress and egress. See Section 6.5.

6.5 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.

6.5.1 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) ⁵
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 ⁶

Table 2 – CE-VLAN ID Preservation

When an EVC includes a UNI at which more than one CE-VLAN ID is mapped to the EVC by the CE-VLAN ID/EVC Map (see Sections 7.8 and 7.9), the EVC **MUST** have the CE-VLAN ID Preservation Service Attribute.

An obvious benefit of the CE-VLAN ID Preservation feature is enhanced operational simplicity. For example, for a Subscriber connecting multiple campuses using IEEE 802.1Q bridges, the feature obviates the task of renumbering VLANs in different corporate campuses.

6.5.2 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 the properties of the CE-VLAN ID/EVC map (see Section 7.6).

6.6 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 processed⁷ 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.

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. If a Service Frame carrying a Layer 2 Control Protocol is tunneled, it **MUST** be tunneled on the EVC that is identified by the CE-VLAN/EVC Map for the CE-VLAN ID indicated by the Service Frame carrying the Layer 2 Control Protocol. See Section 7.6.

⁵ 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.

⁶ Note that, depending on the internal network implementation, a Service Provider may not be able to preserve all values of VLAN ID for a given EVC. 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.

⁷ For example, the Subscriber’s Ethernet information can be encapsulated in another frame separate from the control protocol frame.

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.)

6.7 EVC Related Performance Service Attributes

Service Frame delivery performance is specified for all Service Frames transported within 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.

The EVC Related Performance Service Attributes specify the Service Frame delivery performance. Three performance attributes are considered in this specification. Those are Frame Delay Performance, Frame Delay Variation Performance, and Frame Loss Ratio Performance. If defined, the Performance Attributes **MUST** apply to all Service Frames with the level of Bandwidth Profile conformance determined to be Green, and associated with a particular Class of Service Identifier on a Point-to-Point EVC that arrive at the UNI during a time interval T . Performance Attributes **MUST NOT** apply to Service Frames with the level of conformance determined to be Yellow or Red. Typically, the Frame Loss Ratio Performance will be degraded for Service Frames determined to be Yellow. Service Frames determined to be Red will be discarded. (See Section 7.10.2.2.) Performance Attributes for a Multipoint-to-Multipoint EVC are beyond the scope of Phase 1 of this Technical Specification.

6.7.1 Frame Delay Performance for a Point-to-Point EVC

The Frame Delay for a Service Frame is defined as the time elapsed from reception at the ingress UNI of the first bit of the ingress Service Frame until the transmission of the last bit of the Service Frame at the egress UNI. This delay is illustrated in Figure 4. Note that this definition of Frame Delay for a Service Frame is the one-way⁸ delay that includes the delays encountered as a result of transmission across the ingress and egress UNIs as well as that introduced by the MEN.

⁸ One-way delay is difficult to measure and therefore one way delay may be approximated from two way measurements. However these techniques are beyond the scope of this document.




Figure 4 – Frame Delay for Service Frame

Frame Delay Performance for a particular Class of Service instance on a Point-to-Point EVC for a time interval T **SHALL** be defined as the P-Percentile of the delay for all Service Frames with Bandwidth Profile compliance determined to be Green, successfully delivered between the UNI pairs during the interval T . The Unicast, Broadcast, Multicast and Layer 2 Control Protocol Service Frame Delivery service attributes define which Service Frames should be successfully delivered.

To restate the definition mathematically, let S_T be the set of Frame Delay values for all successfully delivered Service Frames declared Green whose first bit arrived at the ingress UNI during the interval T . S_T can be expressed as, $S_T = \{d_1, d_2, \dots, d_N\}$ where d_i is the Frame Delay of the i^{th} Service Frame. Then the Frame Delay Performance, \bar{d}_T can be expressed as

$$\bar{d}_T = \min \left\{ d \mid P \leq \frac{100}{N} \sum_{j=1}^N I(d, d_j) \right\}$$

where,

$$I(d, d_j) = \begin{cases} 1 & \text{if } d > d_j \\ 0 & \text{otherwise} \end{cases}$$

The parameters of the Frame Delay Performance and are given in Table 3.

Parameter	Description
T	The time interval
P	The percentile of the Frame Delay Performance
\hat{d}	Frame Delay Performance Objective

Table 3 – Frame Delay Performance Parameters

Given T , P , and a Frame Delay Performance objective \hat{d} , expressed in time units, the Frame Delay Performance **SHALL** be defined as met over the time interval T if and only if $\bar{d}_T \leq \hat{d}$.

6.7.2 Frame Delay Performance for a Multipoint-to-Multipoint EVC

For further study.

6.7.3 Frame Delay Variation Performance a Point-to-Point EVC

Frame Delay Variation (FDV) is a measure of the variation in the Frame Delay between a pair of Service Frames. FDV Performance is applicable to all successfully delivered Service Frames with Bandwidth Profile compliance determined to be Green for a particular Class of Service Identifier on a Point-to-Point EVC for a time interval T . The Unicast, Broadcast, Multicast, and Layer 2 Control Protocol Service Frame Delivery service attributes define which Service Frames should be successfully delivered.

The Frame Delay Variation Performance **SHALL** be defined as the P-percentile of the difference between the Frame delays of a Service Frame pair that satisfies the following characteristics:

- The two Service Frames that comprise the pair arrive at the ingress UNI within the time interval T , and
- The two Service Frames that comprise the pair arrive at the ingress UNI exactly Dt time units apart.

This definition is in agreement with the IP packet delay variation definition given in [8] where the delay variation is defined as the difference between the one-way delay of two packets selected according to some selection function and are within a given interval $[T_1, T_2]$.

The choice of the value for Dt can be related to the application timing information. As an example for voice applications where voice frames are generated at regular intervals, Dt may be chosen to be few multiples of the inter-frame time.

Let a_i be the time of the arrival of the first bit of the i^{th} Service Frame at the ingress UNI, then the two frames i and j are selected according to the selection criterion:

$$\{a_j - a_i = Dt \quad \text{and} \quad j > i\}$$

Let r_i be the time frame i is successfully received (last bit of the frame) at the egress UNI, then the difference in the delays encountered by frame i and frame j is given by:

$$Dd_{ij} = d_i - d_j = (r_i - a_i) - (r_j - a_j) = (a_j - a_i) - (r_j - r_i)$$

With d_j being the delay of the j^{th} frame, a positive value for Dd_{ij} implies that the two frames are closer together at the egress UNI while a negative value for Dd_{ij} implies that the two frames are further apart at the egress UNI. If either or both frames are lost or not delivered due to, for example, FCS violation, then the value Dd_{ij} is not defined and does not contribute to the evaluation of the Frame Delay Variation.

Figure 5 shows a depiction of the different times that are related to Frame Delay Variation Performance.




Figure 5 – Frame Delay Variation Parameters

Let $S_{DR} = \{Dd_{ij} : i, j \text{ such that } a_j - a_i = Dt, a_i \hat{=} T, \text{ and } a_j \hat{=} T\}$ be the set of all delay variations for all eligible pairs of Service Frames. Let K be the number of elements in S_{DR} . Define \tilde{d}_{DR} to be the P-percentile of the set

$$\tilde{d}_{DR} = \min\{d : P \leq \frac{100}{K} \sum I(d, Dd_{ij})\}$$

where;

$$I(d, Dd_{ij}) = \begin{cases} 1 & \text{if } d > Dd_{ij} \\ 0 & \text{otherwise} \end{cases}$$

and the sum is carried out over all the values in the set S_{DR} .

Frame Delay Variation Performance depends on the choice of the value for Dt . Values for both Dt and T typically should be chosen to achieve a reasonable level of statistical accuracy.

For the SLS, the Frame Delay Variation entry **MUST** specify a set of parameters and an objective. The parameters of the Frame Delay Variation Performance are given in Table 4.

Parameter	Description
T	The interval
P	Frame Delay Variation Performance percentile
D_t	The separation between frame pairs for which Frame Delay Variation Performance is defined
d	Frame Delay Variation Performance Objective

Table 4 – Frame Delay Variation Parameters

Given T , P , D_t , and d , the Frame Delay Variation Performance **SHALL** be defined as met over the time interval T if and only if $\tilde{d}_{DT} \leq d$.

6.7.4 Frame Delay Variation Performance for a Multipoint-to-Multipoint EVC

For further study.

6.7.5 Frame Loss Ratio Performance for a Point-to-Point EVC

The definition of Frame Loss Ratio Performance for a particular Class of Service instance on a Point-to-Point EVC is based on the number of Service Frames that arrive at an ingress UNI during the interval T and that should be delivered to the egress UNI according to the Service Frame Delivery service attributes (see Sections 6.4.2, 6.6, and 7.12) and whose level of Bandwidth Profile compliance is determined to be Green. The Frame Loss Ratio Performance **SHALL** be defined as the ratio, expressed as a percentage, of the number of such Service Frames not delivered divided by the number of such Service Frames. Note that Layer 2 Control Protocol Service Frames that are peered or discarded at the ingress UNI are not counted as lost frames.

Frame Loss Ratio can be expressed mathematically as follows. Let I_T be the number of Service Frames that arrive at an ingress UNI during the interval T and that should be delivered to the egress UNI according to the Service Frame Delivery service attributes (see Sections 6.4.2, 6.6, and 7.12) and whose level of Bandwidth Profile compliance is determined to be Green. Let E_T be the number of such Service Frames that are delivered during the interval T . Then

$$FLR_T = \frac{I_T - E_T}{I_T} \cdot 100\%$$

For the SLS, the Frame Loss Ratio Performance entry **MUST** specify a set of parameters and an objective. The parameters and objective of the Frame Loss Ratio Performance are given in Table 5.

Parameter	Description
T	The time interval
\hat{L}	Frame Loss Ratio Performance objective

Table 5 – Frame Loss Ratio Performance Parameters

Given T , the Frame Loss Ratio Performance **SHALL** be defined as met over the time interval T if and only if $FLR_T \leq \hat{L}$.

6.7.6 Frame Loss Ratio Performance for a Multipoint-to-Multipoint EVC

For further study.

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 assigned to the UNI by the Service Provider. It **MUST** be a string and the string **MAY** have any 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 6. Typically there are no constraints in mixing UNIs with different physical media in the same EVC.⁹ Constraints on the mix of speeds and modes **MAY** be imposed for some services.

⁹ An exception might be wireless when the service requires stringent requirements on packet loss.

Speed	Mode	Physical Medium
10 Mbps	Full duplex	All Ethernet physical media compatible with Speed and Mode specified in [2] or [3].
100 Mbps	Full duplex	
10/100 Mbps Auto-Negotiation	Full duplex	
1 Gbps	Full duplex	
10 Gbps	Full duplex	

Table 6 – Possible Physical Layer Characteristics

7.3 MAC Layer Service Attribute

The protocols running at the UNI **MUST** support the IEEE 802.3-2002 [2] frame formats.

7.4 Service Multiplexing Service Attribute

A UNI with the Service Multiplexing attribute **MUST** be configurable to support multiple EVCs.¹⁰ Point-to-Point EVCs and Multipoint-to-Multipoint EVCs **MAY** be multiplexed in any combination at a UNI. Figure 6 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 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 6 – Example of Service Multiplexing on UNI A

¹⁰ Since the UNI is dedicated to a single Subscriber, only one Subscriber can access the EVCs at the UNI.

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 tagged¹¹ Service Frames **MUST** have the same CE-VLAN ID and that value **MUST** be configurable to any value in the range 1, 2, ..., 4094. When the CE-VLAN ID Preservation Service Attribute is not in force for an EVC to which the CE-VLAN ID for untagged and priority tagged Service Frames is mapped, egress Service Frames for this EVC **MUST** be untagged. When CE-VLAN ID Preservation Service Attribute is in force for an EVC to which the CE-VLAN ID for untagged and priority tagged Service Frames is mapped, the format of an egress Service Frame for this EVC depends on the format of the corresponding Service Frame at the ingress UNI as described in Section 6.5.1.

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 a string formed by the concatenation of the UNI ID and the EVC ID that is used to identify an EVC at the UNI. It is intended for management and control purposes.

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. The collection of all of these mappings is 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 7 is an example of a CE-VLAN ID/EVC Map. In this example and all of the following examples, the entry in the EVC column can be any suitable identifier for the EVC, e.g., the EVC ID (Section 6.2) or the UNI EVC ID (Section 7.5.2).

¹¹ A priority tagged Service Frame is a Service Frame with an IEEE 802.1Q tag in which the VLAN ID in the tag equals 0.

CE-VLAN ID	EVC
47	EVC ₁
1343	EVC ₂
17 ¹²	EVC ₃




Figure 7 – 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 CE-VLAN ID **MUST** be discarded by the MEN. Also, an egress Service Frame **MUST NOT** have a CE-VLAN ID with this value at the UNI while using 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 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.¹³

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 8 shows valid CE-VLAN ID/EVC Maps for three EVCs between two UNIs. Note that when the CE-VLAN ID Preservation attribute (Section 6.5.1) applies to an EVC,

¹² In this example, the CE-VLAN ID for untagged and priority tagged Service Frames is configured to 17.

¹³ 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.

the mappings for the EVC are identical as is the case for EVC₁ in Figure 8. (Otherwise the CE-VLAN ID cannot be preserved).

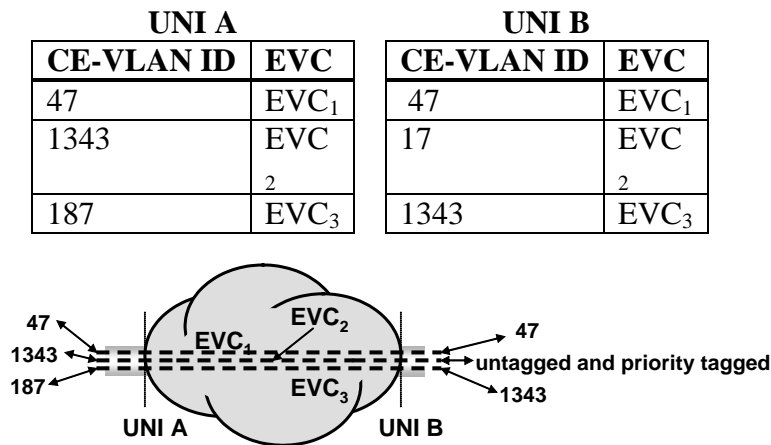


Figure 8 – Example of CE-VLAN ID/EVC Maps at Two UNIs

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 a particular 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 6.5.1) and the list of CE-VLAN IDs mapped to the EVC **MUST** be the same at each UNI in the EVC. Figure 9 shows an example of Bundling. In this example, UNI A and UNI B have the bundling feature as seen from the mapping for EVC₁. (EVC₁ has the CE-VLAN ID Preservation feature.). Note that Bundling is compatible with Service Multiplexing. In Figure 9, UNI A and UNI B are examples of Service Multiplexing and Bundling on the same UNI.

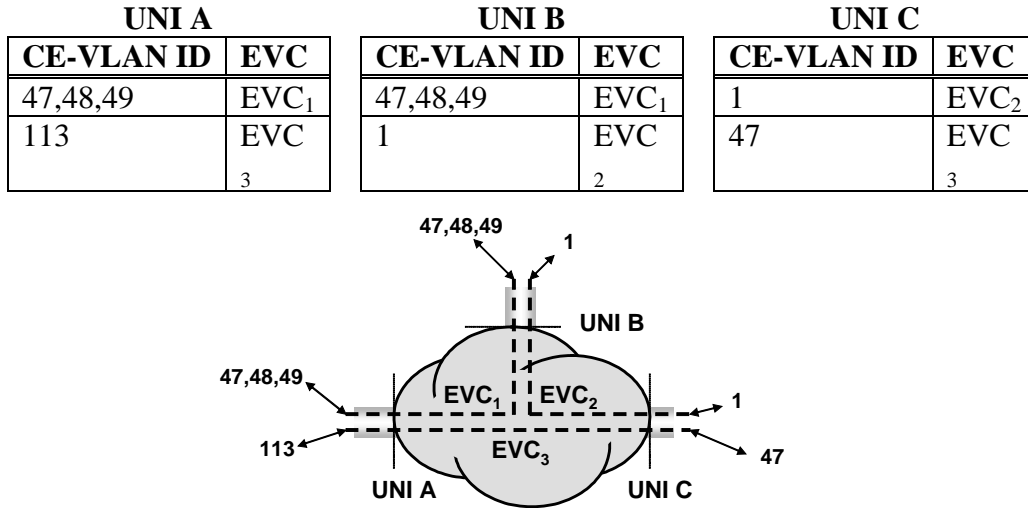


Figure 9 – Example of Bundling

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 10.

Description	
CE-VLAN ID	EVC
2000	EVC ₁
2001	EVC
	3
All others	EVC
	4

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

Figure 10 – 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 set to TRUE, 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 6.5.1) and the list of CE-VLAN IDs mapped to the EVC **MUST** include all CE-VLAN IDs and 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 7 shows the valid combinations of the bundling and Service Multiplexing attributes.

	Valid Combination 1	Valid Combination 2	Valid Combination 3	Valid Combination 4	Valid Combination 5
Service Multiplexing	o	b	b	o	o
Bundling	o	o	b	b	o
All to One Bundling	o	o	o	o	b

Table 7 – 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 a reference point. In this document, the reference point is 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.

The Bandwidth Profile, if present, **SHOULD** be enforced by the provider’s network since it is part of the SLS and is agreed upon between the Subscriber and the Service Provider. QoS policy rules by way of rate enforcement algorithms are typically used to identify the specific actions taken by a Service Provider when Service Frames violate a Bandwidth Profile. The actions may be to drop or recolor the service frame to indicate high discard precedence within the MEN. The outcome of any such rate enforcement, within the Service Provider network is a set of Service Frames that are labeled as Green, Yellow, or Red based on their level of conformance to Bandwidth Profiles.

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 7.10.4. 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.

The Bandwidth Profile defines the set of traffic parameters applicable to a sequence of Service Frames. Associated with the Bandwidth Profile is a rate enforcement algorithm to determine Service Frame compliance with the specified parameters. Rate enforcement also includes the disposition of non-compliant Service Frames by either dropping or marking.

7.10.1 Standard Bandwidth Profile Parameters

The parameters comprising the Bandwidth Profile parameters are:

Committed Information Rate (CIR) expressed as bits per second. **CIR MUST** be ³ 0.

Committed Burst Size (CBS) expressed as bytes. When CIR > 0, **CBS MUST** be greater than or equal to the maximum Service Frame size as specified in [1].

Excess Information Rate (*EIR*) expressed as bits per second. *EIR* **MUST** be ≥ 0

Excess Burst Size (*EBS*) expressed as bytes. When *EIR* > 0, *EBS* **MUST** be greater than or equal to the maximum Service Frame size as specified in [1].

Coupling Flag (*CF*) **MUST** have only one of two possible values, 0 or 1.

Color Mode (*CM*) **MUST** have only one of two possible values, “color-blind” and “color-aware”

Since the coupling Flag has negligible effect in color blind mode, a service definition that uses color blind operation **MAY** be defined without specifying the value of the coupling flag.

7.10.2 Enforcement of Bandwidth Profile Parameters

Each incoming Service Frame is classified to determine which, if any, Bandwidth Profile is applicable to the Service Frame¹⁴. Operation of the Bandwidth Profile algorithm is governed by the six parameters, $\langle CIR, CBS, EIR, EBS, CF, CM \rangle$. The algorithm declares Service Frames to be compliant or non-compliant relative to the Bandwidth Profile. The level of compliance is expressed as one of three colors, Green, Yellow, or Red¹⁵.

The Bandwidth Profile algorithm is said to be in color aware mode when each incoming Service Frame already has a level of compliance (i.e., a color) associated with it and that color is taken into account in determining the level of compliance to the Bandwidth Profile parameters. The Bandwidth Profile algorithm is said to be in color blind mode when the color (if any) already associated with each incoming Service Frame is ignored in determining the level of compliance. Color blind mode support is **REQUIRED** at the UNI. Color aware mode is **OPTIONAL** at the UNI. The color mode of operation **MUST** be determined using the parameter CM.

7.10.2.1 Bandwidth Profile Algorithm

The Bandwidth Profile algorithm is shown in Figure 11. For a sequence of ingress Service Frames, $\{t_j, l_j\}, j \geq 0$, with arrival times t_j and lengths l_j , the level of compliance color assigned to each Service Frame **MUST** be defined according to the algorithm in Figure 11. For this algorithm, $B_c(t_0) = CBS$ and $B_e(t_0) = EBS$. $B_c(t)$ and $B_e(t)$ are the number of bytes in the Committed and Excess token buckets respectively at a given time t .

¹⁴ Recall that a Service Frame is defined as any Ethernet Frame transmitted across the UNI and thus a Layer 2 Control Protocol Ethernet frame is a Service Frame.

¹⁵ The categorization of a Service Frame does not imply any change to the content of the frame. Certain approaches to network implementation may “mark” frames internal to the MEN but such procedures are beyond the scope of this Technical Specification.

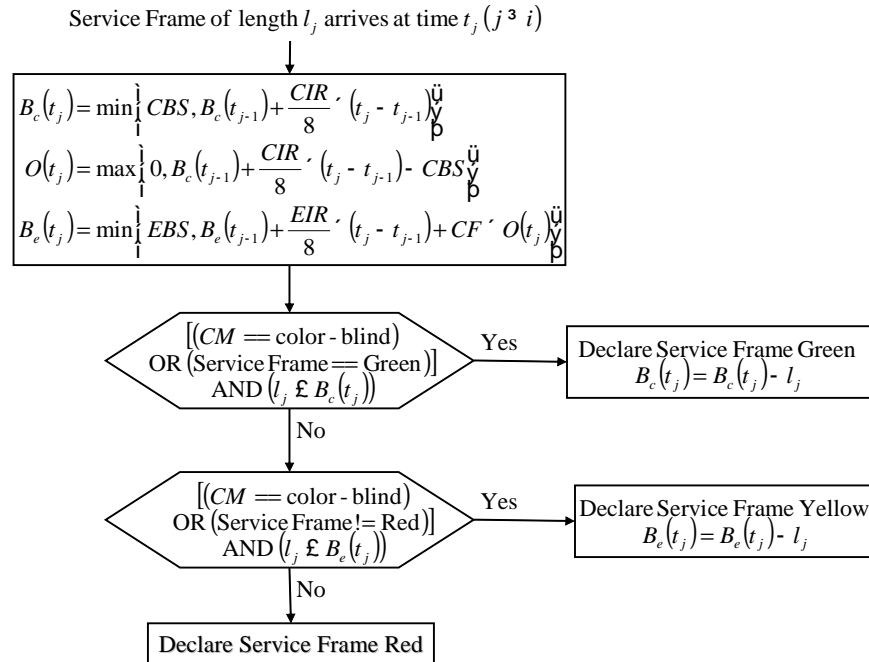


Figure 11 – The Bandwidth Profile Algorithm

Note that the algorithm in Figure 11 does not define an implementation of any network equipment. In fact, since the behavior is described with real numbers for representing time, exactly implementing the behavior is theoretically impossible. However, an implementation should be such that over any time interval $[t_j, t_k]$ the amount of traffic accepted as green, W_G and the amount of traffic accepted as yellow, W_Y are lower bounded below by the values:

$$W_G \geq B_c(t_j) + CIR \cdot (t_k - t_j)$$

$$W_Y \geq B_e(t_j) + EIR \cdot (t_k - t_j)$$

provided that the ingress traffic is greater than these values.

The Coupling Flag CF is set to either 0 or 1. The choice of the value for CF has the effect of controlling the volume of the Service Frames that are declared Yellow that are admitted to the network. When CF is set to 0, the long term average bit rate of Service Frames that are declared Yellow that are admitted to the network is bounded by EIR . When CF is set to 1, the long term average bit rate of Service Frames that are declared Yellow that are admitted to the network is bounded by $CIR + EIR$ depending on volume of the offered Service Frames that are declared Green. In both cases the burst size of the Service Frames that are declared Yellow that are admitted to the network is bounded by EBS .

7.10.2.2 Service Frame Disposition

The second step of Bandwidth Profile rate algorithm is the disposition of the ingress Service Frame based on its level of compliance and **MUST** be as described in Table 8.

Level of Compliance	Service Frame Disposition
Red	Discard
Yellow	Deliver according to the Service Attributes of the service instance but SLS performance objectives do not apply.
Green	Deliver according to the Service Attributes of the service instance and SLS performance objectives apply.

Table 8 – Service Frame Disposition

7.10.3 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 12 illustrates an example of the application of ingress policing with a per ingress UNI. In the example of Figure 12, ingress Service Frames for the three EVCs would all be subject to a single Bandwidth Profile. The Bandwidth Profile per Ingress UNI manages bandwidth non-discriminately for all EVCs at the UNI, i.e. some EVCs may get more bandwidth while others may get less.




Figure 12 – Ingress Bandwidth Profile per Ingress UNI

7.10.4 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 3 Ethernet Virtual Connections, there could

be 3 ingress Bandwidth Profiles, one for each EVC. Figure 13 illustrates an example of the application of Bandwidth Profiles per EVC. In this example, EVC₁ could have CIR=15 Mbps, EVC₂ could have CIR = 10 Mbps, and EVC₃ could have CIR = 20 Mbps.




Figure 13 – Ingress Bandwidth Profile per EVC

7.10.5 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 6.7. Refer to the example in Figure 14. In this example, there are three Class of Service Identifiers within EVC₁, each with a separate Bandwidth Profile.

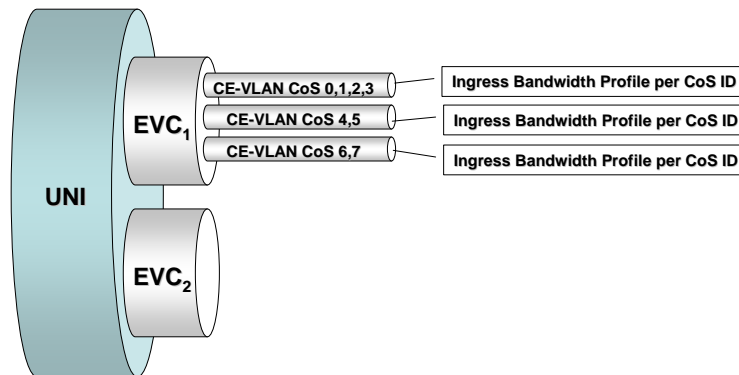


Figure 14 – Ingress Bandwidth Profile per EVC and CE-VLAN CoS

7.10.6 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. For example, in the configuration of Figure 14, there cannot be a Bandwidth Profile for EVC₁. Note also for the configuration in Figure 14, that it is possible to configure a per-

EVC Bandwidth Profile for EVC₂ but there happens to not be a Bandwidth Profile for EVC₂ in this example.

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.¹⁶

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 6.6.

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 6.6.

7.12.3 Pass to EVC

When this alternative is in force, the disposition of the Layer 2 Control Protocol **MUST** be determined by the Layer 2 Control Protocol Processing attribute of the EVC (tunneled or discarded). The EVC associated with Layer 2 Control Protocol is determined by the CE-VLAN ID of the Service Frame and CE-VLAN ID/EVC Map. See Section 6.6.

¹⁶ 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.

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 15):

- 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 at 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 15 – Ethernet Service Framework

8.1 Ethernet Service Types

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 15. Two Ethernet Service Types have been defined in [7]. 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.

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 9 along with the type of parameter value for the attribute. For a given instance of a service, a table like that of Table 9 **MUST** be specified for each UNI in the EVC associated with the service.

Attribute	Type of Parameter Value
UNI Identifier (Section 7.1)	Any string
Physical Medium (Section 7.2)	A Standard Ethernet PHY ([2] or [3]) ¹⁷
Speed (Section 7.2)	10 Mbps, 100 Mbps, 1 Gbps, or 10 Gbps ¹⁷
Mode (Section 7.2)	Full Duplex or Auto-Negotiation ¹⁷
MAC Layer (Section 7.3)	IEEE 802.3 – 2002 [2]
Service Multiplexing (7.4)	Yes or No ¹⁸
UNI EVC ID (Section 7.5.2)	A string formed by the concatenation of the UNI ID and the EVC ID
CE-VLAN ID for untagged and priority tagged Service Frames (Section 7.5.1)	A number in 1, 2, ..., 4094.
CE-VLAN ID/EVC Map (Section 7.6)	Map as per Section 7.6
Maximum Number of EVCs (Section 7.7)	Integer ³ 1
Bundling (Section 7.8)	Yes or No ¹⁸
All to One Bundling (Section 7.9)	Yes or No ¹⁹

¹⁷ There are interdependencies among the values of these parameters as per the IEEE 802.3 Standard.[2]

¹⁸ Must be No if All to One Bundling is Yes.

¹⁹ Must be No if Bundling is Yes or Service Multiplexing is Yes.

Attribute	Type of Parameter Value
Ingress Bandwidth Profile Per Ingress UNI (Section 7.10.3)	No or parameters as defined in Section 7.10.1
Ingress Bandwidth Profile Per EVC (Section 7.10.4)	No or parameters as defined in Section 7.10.1
Ingress Bandwidth Profile Per Class of Service Identifier (Section 7.10.5)	No or parameters as defined in Section 7.10.1
Layer 2 Control Protocols Processing (Section 7.12)	Entries from Table 1 with each being labeled Discard, Peer, or Pass to EVC

Table 9 – UNI Service Attributes

The EVC Service Attributes are listed in Table 10 along with the type of parameter value for the attribute. For a given instance of a service, a table like that of Table 10 **MUST** be specified for the EVC associated with the service.

Attribute	Type of Parameter Value
EVC Type (Section 6.1)	Point-to-Point or Multipoint-to-Multipoint
EVC ID (Section 6.2)	An arbitrary string, unique across the MEN, for the EVC supporting the service instance
UNI List (6.3)	A list of UNI Identifiers (Section 7.1)
CE-VLAN ID Preservation (6.5.1)	Yes or No
CE-VLAN CoS Preservation (Section 6.5.2)	Yes or No
Unicast Service Frame Delivery (6.4.1.1)	Discard, Deliver Unconditionally, or Deliver Conditionally. If Deliver Conditionally is used, then the conditions MUST be specified.
Multicast Service Frame Delivery (Section 6.4.1.2)	Discard, Deliver Unconditionally, or Deliver Conditionally. If Deliver Conditionally is used, then the conditions MUST be specified.
Broadcast Service Frame Delivery (Section 6.4.1.3)	Discard, Deliver Unconditionally, or Deliver Conditionally. If Deliver Conditionally is used, then the conditions MUST be specified.
Layer 2 Control Protocols Processing (Section 6.6)	Entries from Table 1 labeled Tunnel or Discard.
EVC Performance (Section 6.7)	Performance objectives for Frame Delay Performance, Frame Delay Variation Performance, and Frame Loss Ratio Performance as specified in Section 6.7.

Table 10 – EVC Service Attributes

9. References

- [1] Bradner, S., *Key words for use in RFCs to Indicate Requirement Levels*, RFC 2119, March 1997. (Normative)
- [2] IEEE Std 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)
- [3] IEEE Std 802.3ae – 2002, *IEEE Standard for 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 Amendment: Media Access Control (MAC) Parameters, Physical Layers, and Management Parameters for 10 Gb/s Operation*, 13 June 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.
- [5] Metro Ethernet Forum, MEF 1, *Ethernet Services Model, Phase 1*, 10 November 2003.
- [6] Metro Ethernet Forum, MEF 5, *Traffic Management Specification: Phase 1*, May 2004.
- [7] Metro Ethernet Forum, MEF 6, *Ethernet Services Definitions - Phase 1*, June 2004.
- [8] C. Demichelis and P. Chimento, *IP Packet Delay Variation Metric for IP Performance Metric (IPPM)*, RFC 3393, November 2002.