



Technical Specification
MEF 10.3

Ethernet Services Attributes Phase 3

October 2013

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 MEF 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 MEF 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 MEF Forum 2013. All Rights Reserved.

Table of Contents

1.	List of Contributing Members	1
2.	Abstract.....	1
3.	Terminology and Acronyms.....	1
4.	Scope	10
4.1	New Material beyond MEF 10.2 and MEF 10.2.1	10
4.2	Revisions to Material in MEF 10.2 and MEF 10.2.1	11
5.	Compliance Levels.....	12
6.	Numerical Prefix Conventions	12
7.	Introduction	13
8.	Ethernet Virtual Connection Service Attributes	15
8.1	Ethernet Virtual Connection Type Service Attribute	15
8.1.1	Point-to-Point EVC	15
8.1.2	Multipoint EVCs	16
8.1.2.1	<i>Rooted-Multipoint EVC.....</i>	<i>16</i>
8.1.2.2	<i>Multipoint-to-Multipoint EVC</i>	<i>17</i>
8.2	EVC ID Service Attribute.....	17
8.3	UNI List Service Attribute	18
8.4	Maximum Number of UNIs Service Attribute	18
8.5	Service Frame Delivery Service Attributes	18
8.5.1	Types of Service Frame	18
8.5.1.1	<i>Layer 2 Control Protocol Service Frame.....</i>	<i>18</i>
8.5.1.2	<i>SOAM Service Frame.....</i>	<i>19</i>
8.5.1.3	<i>Data Service Frame</i>	<i>19</i>
8.5.2	Data Service Frame Disposition.....	19
8.5.3	Data Service Frame Transparency.....	20
8.6	CE-VLAN Tag Preservation Service Attributes.....	22
8.6.1	CE-VLAN ID Preservation Service Attribute	22
8.6.2	CE-VLAN CoS Preservation Service Attribute	23
8.7	EVC Layer 2 Control Protocol Processing Service Attribute	23
8.8	EVC Related Performance Service Attributes	24
8.8.1	One-way Frame Delay Performance for an EVC.....	26
8.8.2	One-way Inter-Frame Delay Variation Performance for an EVC	30
8.8.3	One-way Frame Loss Ratio Performance for an EVC.....	33
8.8.4	One-way Availability Performance for an EVC	34
8.8.5	One-way Resiliency Performance for an EVC.....	40
8.8.6	One-way Group Availability Performance for an EVC.....	44
8.9	EVC Maximum Service Frame Size Service Attribute.....	45
9.	UNI Service Attributes	46
9.1	UNI ID Service Attribute	46
9.2	Physical Layer Service Attribute	46
9.3	Synchronous Mode Service Attribute	47
9.4	Number of Links Service Attribute.....	47

9.5	UNI Resiliency Service Attribute	47
9.6	Service Frame Format Service Attribute	48
9.7	UNI Maximum Service Frame Size Service Attribute.....	48
9.8	Service Multiplexing Service Attribute.....	48
9.9	Customer Edge VLAN ID	49
9.10	CE-VLAN ID/EVC Map Service Attribute.....	50
9.10.1	Basic Concept.....	50
9.10.2	CE-VLAN ID Significance	51
9.10.3	Describing the Contents of the CE-VLAN ID/EVC Map.....	52
9.11	Maximum Number of EVCs Service Attribute	52
9.12	Bundling Service Attribute	52
9.13	All to One Bundling Service Attribute.....	53
9.14	Ingress Bandwidth Profile per UNI Service Attribute	53
9.15	Egress Bandwidth Profile per UNI Service Attribute	54
9.16	Link OAM Service Attribute	55
9.17	UNI MEG Service Attribute.....	55
9.18	E-LMI Service Attribute.....	55
9.19	UNI Layer 2 Control Protocol Processing Service Attribute	56
10.	EVC per UNI Service Attributes	56
10.1	UNI EVC ID Service Attribute.....	56
10.2	Class of Service Identifier Service Attribute	56
10.2.1	Class of Service Identifier for Data Service Frames.....	57
10.2.1.1	<i>Class of Service Identifier Based on EVC</i>	<i>57</i>
10.2.1.2	<i>Class of Service Identifier Based on Priority Code Point Field.....</i>	<i>57</i>
10.2.1.3	<i>Class of Service Identifier Based on Internet Protocol.....</i>	<i>58</i>
10.2.2	Class of Service Identifier for Layer 2 Control Protocol Service Frames	60
10.2.3	Class of Service Identifier for SOAM Service Frames.....	60
10.3	Color Identifier Service Attribute	61
10.3.1	Color Identifier Based on EVC	62
10.3.2	Color Identifier Based on CE-VLAN Tag DEI.....	62
10.3.3	Color Identifier Based on CE-VLAN CoS.....	62
10.3.4	Color Identifier Based on Internet Protocol	63
10.3.5	Color Identifier Constraints	64
10.4	Egress Equivalence Class Identifier Service Attribute.....	64
10.4.1	Egress Equivalence Class Identifier for Data Service Frames	64
10.4.1.1	<i>Egress Equivalence Class Identifier Based on CE-VLAN CoS.....</i>	<i>65</i>
10.4.1.2	<i>Egress Equivalence Class Identifier Based on Internet Protocol.....</i>	<i>65</i>
10.4.2	Egress Equivalence Class Identifier for Layer 2 Control Protocol Service Frames	66
10.4.3	Egress Equivalence Class Identifier for SOAM Service Frames	67
10.5	Ingress Bandwidth Profile per EVC Service Attribute	67
10.6	Ingress Bandwidth Profile per Class of Service Identifier Service Attribute	68
10.7	Egress Bandwidth Profile per EVC Service Attribute	69
10.8	Egress Bandwidth Profile per Egress Equivalence Class Identifier Service Attribute	69
10.9	Source MAC Address Limit Service Attribute.....	70
10.10	Test MEG Service Attribute	71
10.11	Subscriber MEG MIP Service Attribute.....	71
11.	Multiple EVC Related Performance Service Attributes	72

11.1 One-way Multiple EVC Group Availability.....	72
12. Bandwidth Profile at the UNI	73
12.1 Bandwidth Profile Parameters and Algorithm.....	75
12.2 Bandwidth Profile without Sharing.....	81
12.3 Ingress Bandwidth Profiles Service Attributes.....	83
12.4 Egress Bandwidth Profiles Service Attributes.....	84
12.5 Simultaneous Application of the Ingress Bandwidth Profile Application Models.....	84
12.6 Simultaneous Application of the Egress Bandwidth Profile Application Models.....	85
13. Ethernet Service Framework	85
14. References	88
Appendix A Examples (Informative).....	90
A.1 CE-VLAN ID Preservation Service Attribute	90
A.1.1 CE-VLAN ID Preservation Enabled	90
A.1.2 CE-VLAN ID Preservation Disabled	91
A.2 Examples of the Use of the CE-VLAN ID/EVC Map and EVCs.....	94
A.2.1 Untagged UNIs.....	94
A.2.2 Use of Rooted-Multipoint EVC	94
A.2.3 Redundant Higher Layer Service Access.....	95
A.3 Traffic Shaping Example.....	95
A.4 Examples of Availability Metrics for Multipoint EVCs	98
A.4.1 UNI-oriented Availability Example	99
A.4.2 EVC-oriented Availability Example.....	100
Appendix B Effect of Service Frame Overhead (Informative)	100
Appendix C Bandwidth Profiles (Informative)	102
C.1 Visual Representation of the Bandwidth Profile	102
C.2 Bandwidth Profile Use Cases	104
C.2.1 Sharing Excess Bandwidth.....	104
C.2.2 Uncoupled Bandwidth Sharing	105
C.2.3 Sharing Bandwidth among Class of Service Names	106
C.2.4 Active/Standby EVCs.....	108
C.3 Bandwidth Profile Implementation	110
C.3.1 Eliminating Bias Toward Short Service Frames	111

List of Figures

Figure 1 – Ethernet Services Model.....	13
Figure 2 – Point-to-Point EVCs.....	15
Figure 3 – Rooted-Multipoint EVC	17
Figure 4 – Multipoint-to-Multipoint EVC.....	17
Figure 5 – Tagged In to Tagged Out Data Service Frame Transparency	20
Figure 6 – Tagged In to Untagged Out Data Service Frame Transparency	21
Figure 7 – Untagged In to Tagged Out Data Service Frame Transparency	21
Figure 8 – Untagged In to Untagged Out Data Service Frame Transparency.....	22
Figure 9 - Frame Delay for Service Frame.....	26

Figure 10 – One-way Inter-Frame Delay Variation Definition	30
Figure 11 – Flowchart Definition of $A_{\langle i,j \rangle}(\Delta t_k)$	37
Figure 12 – Example of the Determination of $A_{\langle i,j \rangle}(\Delta t_k)$	38
Figure 13 – Example of the Impact of a Maintenance Interval	38
Figure 14 – Hierarchy of Time Showing the Resiliency Attributes	41
Figure 15 – Determining and Counting Consecutive High Loss Intervals over T	42
Figure 16 – Example of Counting High Loss Intervals and Consecutive High Loss Intervals	43
Figure 17 – Example of Service Multiplexing on UNI A	49
Figure 18 – Example of a CE-VLAN ID/EVC Map	50
Figure 19 – Example of CE-VLAN ID/EVC Maps at Two UNIs	51
Figure 20 – Example of Bundling	52
Figure 21 – Example of a Simple Description of Bundling	53
Figure 22 – Ingress Bandwidth Profile per UNI	54
Figure 23 – Egress Bandwidth Profile per UNI	55
Figure 24 – Ingress Bandwidth Profile per EVC	68
Figure 25 – Ingress Bandwidth Profile per Class of Service Identifier	69
Figure 26 – Frame Discard for Source MAC Address Limit	71
Figure 27 – The Bandwidth Profile Algorithm for an Envelope	81
Figure 28 – Bandwidth Profile Algorithm for one Bandwidth Profile Flow in an Envelope	83
Figure 29 – CE-VLAN ID/EVC Map Notation	90
Figure 30 – Example 1: CE-VLAN ID Preservation Enabled with All to One Bundling Enabled	91
Figure 31 – Example 2: CE-VLAN ID Preservation Enabled with Bundling Enabled and Multiple CE-VLAN IDs mapped to EVC2	91
Figure 32 – CE-VLAN ID/EVC Map Notation	92
Figure 33 – Example 3: CE-VLAN ID Preservation Disabled	93
Figure 34 – Untagged UNIs	94
Figure 35 – Use of a Rooted-Multipoint EVC	95
Figure 36 – Redundant Higher Layer Service Access	95
Figure 37 – Periodic Algorithm	97
Figure 38 – New Frame Algorithm	98
Figure 39 – Multipoint EVC Example	99
Figure 40 – IEEE Std 802.3 Packet Format (simplified)	100
Figure 41 – Ratio of Information Rate to UNI Line Rate for Back to Back Service Frames	101
Figure 42 – Behavior of the Bandwidth Profile Algorithm	104
Figure 43 – Bandwidth Profile Algorithm for Sharing Excess Bandwidth	105
Figure 44 – Bandwidth Profile Algorithm for Uncoupled Bandwidth Sharing	106
Figure 45 – Bandwidth Profile Algorithm for Bandwidth Sharing among Class of Service Identifiers	108
Figure 46 – Active/Standby EVCs Use Case	108
Figure 47 – Mapping of 6 Bandwidth Profile Flows to 3 Envelopes for the Active/Standby Use Case	109
Figure 48 – Bandwidth Profile Algorithm for the Envelope of H Bandwidth Profile Flows	110
Figure 49 – Service Frame Profiling Results for the Algorithm of Section 12.1	112
Figure 50 – Service Frame Profiling Results for Modified Algorithm (Time Measured Last Bit to Last Bit)	112

Figure 51 – Service Frame Profiling Results for Modified Algorithm (Time Measured First Bit to First Bit)	112
---	-----

List of Tables

Table 1 – Terminology and Acronyms	9
Table 2 – Numerical Prefix Conventions	13
Table 3 – List of Standardized Layer 2 Control Protocol Destination MAC Addresses	19
Table 4 – Behavior for CE-VLAN ID Preservation for a Service Frame	22
Table 5 – CE-VLAN ID Preservation Service Attribute for an EVC	23
Table 6 – One-way Frame Delay Performance Parameters	29
Table 7 – One-way Inter-Frame Delay Variation Parameters	32
Table 8 – One-way Frame Loss Ratio Performance Parameters	34
Table 9 – One-way Availability Performance Parameters for an EVC	40
Table 10 – One-way Resiliency Performance Parameters for an EVC	43
Table 11 – One-way Group Availability Performance Parameters for an EVC	45
Table 12 – Valid Combinations of Service Multiplexing, Bundling, and All to One Bundling ...	53
Table 13 – PCP Class of Service Identifier Example for EVC1 at the UNI	58
Table 14 – PCP Class of Service Identifier Example for EVC2 at the UNI	58
Table 15 – Internet Protocol Class of Service Identifier Example for an EVC at the UNI	59
Table 16 – IPv4 Class of Service Identifier Example for an EVC at the UNI	59
Table 17 – IPv6 Class of Service Identifier Example for an EVC at the UNI	59
Table 18 – CE-VLAN CoS Color Identifier Example for EVC 1 at the UNI	62
Table 19 – CE-VLAN CoS Color Identifier Example for EVC 2 at the UNI	62
Table 20 – Combined Mappings Consistent with Table 4 in MEF 23.1 [22]	63
Table 21 – Internet Protocol Color Identifier Example	63
Table 22 – IPv4 Only Color Identifier Example	64
Table 23 – Internet Protocol Egress Equivalence Class Identifier Example for an EVC at the UNI	66
Table 24 – IPv4 Only Egress Equivalence Class Identifier Example for an EVC at the UNI	66
Table 25 – One-way Multiple EVC Group Availability Performance Parameters	73
Table 26 – UNI Service Attributes	86
Table 27 – EVC Service Attributes	87
Table 28 – EVC per UNI Service Attributes	88
Table 29 – Parameter Values for Sharing Bandwidth among Class of Service Identifiers	107
Table 30 – Ingress Bandwidth Profile Parameter Values for the H Bandwidth Profile Flows ...	110
Table 31 – Service Frame Profiling Results for the Last 250 Service Frames	113

1. List of Contributing Members

The following members of the MEF participated in the development of this document and have requested to be included in this list.

ADVA Optical Networking SE	EXFO Inc.
at&t	Nokia Solutions and Networks
BTI Systems Inc.	NTT Advanced Technology Corporation
CENX	Omnitron Systems Technology, Inc.
Ciena Corporation	Pulse Communications (Pulsecom)
Colt	Verizon

2. 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 10.2 [15] and MEF 10.2.1 [16].

3. Terminology and Acronyms

This section defines the terms used in this document. In many cases, the normative definitions to terms are found in other documents. In these cases, the third column is used to provide the reference that is controlling, in other MEF or external documents.

Term	Definition	Source
All to One Bundling	A UNI Service Attribute in which all CE-VLAN IDs are mapped to a single EVC.	This document
Available Time	The set of small time intervals Δt contained in T that do not intersect a Maintenance Interval and have $A_{\langle i,j \rangle}(\Delta t) = 1$.	This document
Bandwidth Profile	A characterization of the lengths and arrival times for Service Frames at a reference point.	This document
Bandwidth Profile Flow	A set of Service Frames at a UNI that meet a specific set of criteria.	This document
BPDU	Bridge Protocol Data Unit	IEEE 802.1Q – 2011 [2]
Broadcast Data Service Frame	A Data Service Frame that has the broadcast Destination MAC Address.	This document
Bundling	A UNI Service Attribute in which more than one CE-VLAN ID can be mapped to an EVC.	This document

Term	Definition	Source
Carrier Ethernet Network	A network from a Service Provider or network Operator supporting the MEF service and architecture models. ¹	MEF 12.1 [17]
CBS	Committed Burst Size	This document
CE	Customer Edge	This document
CEN	Carrier Ethernet Network	MEF 12.1 [17]
CE-VLAN CoS	Customer Edge VLAN CoS	This document
CE-VLAN CoS Preservation	An EVC Service Attribute that, when Enabled, requires an egress Service Frame resulting from an ingress Service Frame that contains a CE-VLAN CoS to have the identical CE-VLAN CoS.	This document
CE-VLAN ID	Customer Edge VLAN ID	This document
CE-VLAN ID Preservation	An EVC Service Attribute that, when Enabled, requires an egress Service Frame resulting from an ingress Service Frame to have an identical CE-VLAN ID.	This document
CE-VLAN ID/EVC Map	An association of CE-VLAN IDs with EVCs at a UNI.	This document
CE-VLAN Tag	Customer Edge VLAN Tag	This document
CF	Coupling Flag	This document
CHLI	Consecutive High Loss Interval	This document
CIR	Committed Information Rate	This document
CIR_{max}	The Bandwidth Profile parameter that limits the rate of tokens added to the committed token bucket.	This document
Class of Service Identifier	The mechanism and/or values of the parameters in the mechanism to be used to identify the Class of Service Name that applies to a Service Frame.	MEF 23.1 [22]

¹ This document addresses MEF services only from a Service Provider.

Term	Definition	Source
Class of Service Name	A designation given to one or more sets of performance objectives and associated parameters by the Service Provider.	MEF 23.1 [22]
CM	Color Mode	This document
Color Identifier	The mechanism and/or values of the parameters in the mechanism used to identify the Color that applies to the Service Frame at a given UNI.	MEF 23.1 [22]
Color Mode	The Bandwidth Profile parameter that 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.	This document
Color-aware	A Bandwidth Profile property where the level of compliance for each Service Frame is dependent on the value of the Frame's Color Identifier.	This document
Color-blind	A Bandwidth Profile property where the level of compliance for each Service Frame is not dependent on the value of the Frame's Color Identifier.	This document
Committed Burst Size	The Bandwidth Profile parameter that limits the maximum number of bytes available for a burst of Service Frames sent at the UNI line rate that will be declared Green by the Bandwidth Profile.	This document
Committed Information Rate	The Bandwidth Profile parameter that limits the average rate in bits per second of Service Frames that will be declared Green by the Bandwidth Profile.	This document
Consecutive High Loss Interval	A sequence of small time intervals contained in T , each with a high frame loss ratio.	This document
CoS	Class of Service	This document
CoS_Name	A parameter used in Performance Metrics that specifies the Class of Service Name for the metric	This document
Coupling Flag	The Bandwidth Profile parameter that determines whether or not overflow tokens not used for Service Frames declared Green can be used as Yellow tokens.	This document
Customer Edge	Equipment on the Subscriber side of the UNI.	This document

Term	Definition	Source
Customer Edge VLAN Class of Service	The Priority Code Point bits in the IEEE Std 802.1Q – 2011 Customer VLAN Tag in a Tagged Service Frame.	This document
Customer Edge VLAN ID	The identifier derivable from the content of a Service Frame that allows the Service Frame to be mapped to an EVC at the UNI.	This document
Customer Edge VLAN Tag	The IEEE Std 802.1Q – 2011 Customer VLAN Tag in a Tagged Service Frame.	This document
Data Service Frame	A Service Frame that is neither a Layer 2 Control Protocol Service Frame nor a SOAM Service Frame.	This document
DEI	Drop Eligible Indicator	IEEE Std 802.1Q – 2011 [2]
DSCP	Differentiated Services Code Point	RFC 2474 [9]
DTE	Data Termination Equipment	IEEE Std 802.3 – 2012 [3]
EBS	Excess Burst Size	This document
Egress Bandwidth Profile	A Service Attribute that specifies the length and arrival time characteristics of egress Service Frames at the egress UNI.	This document
Egress Equivalence Class Identifier	The mechanism and/or values of the parameters in the mechanism that can be used to specify an Egress Bandwidth Profile Flow for egress Service Frames.	This document
Egress Service Frame	A Service Frame sent from the Service Provider network to the CE.	This document
EIR	Excess Information Rate	This document
EIR_{max}	The Bandwidth Profile parameter that limits the rate of tokens added to the excess token bucket.	This document
E-LMI	Ethernet Local Management Interface	MEF 16 [18]
Envelope	A set of n Bandwidth Profile Flows in which each Bandwidth Profile Flow is assigned a unique rank between 1 (lowest) and n (highest).	This document

Term	Definition	Source
ESMC	Ethernet Synchronization Message Channel	ITU G.8264 [5]
Ethernet Virtual Connection	An association of two or more UNIs that limits the exchange of Service Frames to UNIs in the Ethernet Virtual Connection.	This document
EVC	Ethernet Virtual Connection	This document
EVC Maximum Service Frame Size	An EVC Service Attribute that specifies the maximum size of a Service Frame allowed for that EVC.	This document
Excess Burst Size	The Bandwidth Profile parameter that limits the maximum number of bytes available for a burst of Service Frames sent at the UNI line rate that will be declared Yellow by the Bandwidth Profile.	This document
Excess Information Rate	The Bandwidth Profile parameter that limits the average rate in bits per second of Service Frames that will be declared Yellow by the Bandwidth Profile.	This document
FD	Frame Delay	This document
FDR	Frame Delay Range	This document
Frame	Short for Ethernet frame.	This document
Frame Delay	The time elapsed from the transmission at the ingress UNI of the first bit of the corresponding ingress Service Frame until the reception of the last bit of the Service Frame at the egress UNI.	This document
Frame Delay Range	The Frame Delay Performance minus the minimum Service Frame delay.	This document
High Loss Interval	A small time interval contained in T with a high frame loss ratio.	This document
HLI	High Loss Interval	This document
IFDV	Inter-Frame Delay Variation	This document

Term	Definition	Source
Information Rate	The average bit rate of Ethernet service frames at the measurement point starting with the first MAC address bit and ending with the last FCS bit.	ITU Y.1564 [7]
Ingress Bandwidth Profile	A characterization of ingress Service Frame arrival times and lengths at the ingress UNI and a specification of disposition of each Service Frame based on its level of compliance with the characterization.	This document
Ingress Service Frame	A Service Frame sent from the CE into the Service Provider network.	This document
Inter-Frame Delay Variation	The difference between the one-way delays of a pair of selected Service Frames.	This document
L2CP Service Frame	Layer 2 Control Protocol Service Frame	This document
LAG	Link Aggregation Group	IEEE Std 802.1AX – 2008 [1]
Layer 2 Control Protocol Service Frame	A Service Frame that could be used in a recognized Layer 2 Control Protocol.	This document
Maintenance Interval	A time interval agreed upon by both the Subscriber and Service Provider during which the service may not perform well or at all.	This document
Maximum Number of UNIs	An EVC Service Attribute that specifies the maximum number of UNIs that may be in an EVC.	This document
Multicast Data Service Frame	A Data Service Frame that has a multicast Destination MAC Address.	This document
Multipoint EVC	An EVC that associates two or more UNIs.	This document
Multipoint-to-Multipoint EVC	An EVC with two or more UNIs with all UNIs having the Root Role. 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.	This document
One-way Availability Performance	A characterization of the percentage of time, during the time interval T , that Service Frame loss is small.	This document

Term	Definition	Source
One-way Frame Delay Performance	A characterization of the delays experienced by different Service Frames with the same Class of Service Name.	This document
One-way Frame Delay Range Performance	A characterization of the extent of delay variability experienced by different Service Frames with the same Class of Service Name.	This document
One-way Inter-Frame Delay Variation Performance	A characterization of the variation in the delays experienced by different Service Frames mapped to the same EVC that have the same Class of Service Name.	This document
One-way Frame Loss Ratio Performance	A characterization of the number of lost frames between the ingress UNI and the egress UNI for Service Frames that have the same Class of Service Name. Frame Loss Ratio Performance is expressed as a percentage.	This document
One-way Mean Frame Delay Performance	The arithmetic mean of delays experienced by a set of Service Frames.	This document
One-way Resiliency Performance	The number of High Loss Intervals and/or Consecutive High Loss Intervals in T	This document
Ordered Pair of UNIs	A directional UNI pair of the form <Ingress UNI, Egress UNI>, selected from the UNI list for the EVC of interest.	This document
PCP	Priority Code Point	IEEE Std 802.1Q – 2011 [2]
Performance Metric	A quantitative characterization of Service Frame delivery quality.	This document
Point-to-Point EVC	An EVC with exactly 2 UNIs.	This document
Priority Tagged Service Frame	A Service Frame with a TPID = 0x8100 following the Source Address and the corresponding VLAN ID value is 0x000 in the tag following the TPID.	This document
QL	Quality Level of clock used in Synchronous Ethernet.	ITU G.8264 [5]
Qualified Service Frames	The set of frames that comply with specific criteria, such as the arrival time at the Ingress UNI and Bandwidth Profile compliance, on which a performance attribute is based.	This document

Term	Definition	Source
Rooted-Multipoint EVC	A multipoint EVC in which each UNI has either the Root Role or the Leaf Role. Ingress Service Frames at a Root Role UNI can be delivered to one or more of any of the other UNIs in the EVC. Ingress Service Frames at a Leaf Role UNI can be delivered to only one or more Root Role UNIs in the EVC.	This document
Service Frame	The first bit of the Destination MAC Address through the last bit of the Frame Check Sequence of an IEEE 802.3 Packet transmitted across the UNI.	This document
Service Level Agreement	The contract between the Subscriber and Service Provider specifying the service level commitments and related business agreements.	This document
Service Level Specification	The technical specification of the service level being offered by the Service Provider to the Subscriber.	This document
Service Multiplexing	A UNI Service Attribute in which the UNI can be in more than one EVC instance.	This document
Service Provider	The organization providing Ethernet Service(s).	This document
SES	Severely Errored Seconds	ITU Y.1563 [6]
SLA	Service Level Agreement	This document
SLS	Service Level Specification	This document
SOAM Service Frame	A Service Frame whose MAC Destination Address does not indicate it to be an L2CP Service Frame and whose Ethertype = 0x8902.	This document
S-Tagged Service Frame	A Service Frame with a TPID = 0x88a8 following the Source Address.	IEEE Std 802.1Q – 2011 [2]
Subscriber	The organization purchasing and/or using Ethernet Services.	This document
<i>T</i>	The time interval over which the SLS applies	This document
Tagged Service Frame	A Service Frame that is either a VLAN Tagged Service Frame or a Priority Tagged Service Frame.	This document

Term	Definition	Source
TCI	Tag Control Information	IEEE Std 802.1Q – 2011 [2]
TPID	Tag Protocol Identifier	IEEE Std 802.1Q – 2011 [2]
Unavailable Time	The set of small time intervals Δt contained in T that do not intersect a Maintenance Interval and have $A_{(i,j)}(\Delta t) = 0$.	This document
UNI	User Network Interface	This document
UNI Line Rate	The MAC data rate ² at the UNI.	This document
UNI Maximum Service Frame Size	A UNI Service Attribute that specifies the maximum size of a Service Frame allowed at the UNI.	This document
Unicast Data Service Frame	A Data Service Frame that has a unicast Destination MAC Address.	This document
Untagged Service Frame	A Service Frame with the two bytes following the Source Address field containing neither the value 0x8100 nor the value 0x88a8.	This document
User Network Interface	The physical demarcation point between the responsibility of the Service Provider and the responsibility of the Subscriber.	This document
Utilized Line Rate	The average bit rate of the Ethernet line at the measurement point, including the bits a) allocable to the minimum-duration period of each Inter-Packet gap (but not the number of bits allocable to the part of each Inter-Packet gap longer than the minimum), b) in the preamble, c) in the start of frame delimiter and d) in the Ethernet Service Frame starting with the first MAC address bit and ending with the last FCS bit.	ITU Y.1564 [7]
VLAN Tagged Service Frame	A Service Frame with a TPID = 0x8100 following the Source Address and the corresponding VLAN ID value is not 0x000 in the tag following the TPID.	This document

Table 1 – Terminology and Acronyms

² See IEEE 802.3 - 2012 [3] for the definition of MAC data rate.

4. Scope

This document describes Ethernet Service attributes for services provided to the Subscriber by the Service Provider. 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 10.2, *Ethernet Services Attributes Phase 2* [15] and MEF 10.2.1, *Performance Attributes Amendment to MEF 10.2* [16].

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.

Management of the services is not addressed in this document. See MEF 7.2 [14], MEF 17 [19], MEF 30.1 [25], MEF 31 [26], MEF 35.1 [28], MEF 36 [29], MEF 38, [30] and MEF 40 [31] for the management perspective of the Ethernet service attributes.

4.1 New Material beyond MEF 10.2 and MEF 10.2.1

New topics in this document include:

- Behavior for Service Frames with TPID = 0x88e7 (I-tagged frames in IEEE Std 802.1Q – 2011 [2]) is specified (behavior for Service Frames with TPID = 0x88a8 is beyond the scope of this document),
- SOAM Service Frame (Section 8.5.1.2),
- One-way Group Availability Performance for an EVC: A characterization of availability of sets of ordered UNI pairs associated by an EVC (Section 8.8.6),
- Synchronous Mode Service Attribute: Support of Synchronous Ethernet at the UNI (Section 9.3),
- Number of Links Service Attribute: The number of links supporting the UNI (Section 9.4),
- UNI Resiliency Service Attribute: Support of Link Aggregation when there are two links supporting the UNI (Section 9.5),
- Link OAM Service Attribute: Support of Link OAM at the UNI (Section 9.16),
- UNI MEG Service Attribute: Support of the UNI Maintenance Entity Group at the UNI (Section 9.17),

- E-LMI Service Attribute: Support of Ethernet Local Management Interface at the UNI (Section 9.18),
- Color Identifier Service Attribute: The mechanism for indicating the color of Service Frames for an EVC at a UNI (Section 10.3),
- Source MAC Address Limit Service Attribute: A limit on the number of source MAC addresses that can be used for an EVC at a UNI (Section 10.9),
- Test MEG Service Attribute: Support of the Test Maintenance Entity Group for an EVC at a UNI (Section 10.10),
- Subscriber MEG MIP Attribute: Support of the Subscriber Maintenance Entity Group, Maintenance Intermediate Point for an EVC at a UNI (Section 10.11),
- One-way Multiple EVC Group Availability Performance: A characterization of availability of sets of ordered UNI pairs associated by more than one EVC (Section 11.1), and
- New Bandwidth Profile Algorithm: A backward compatible generalization of the Bandwidth Profile algorithm in MEF 10.2 that allows token sharing among Bandwidth Profile Flows (Section 12).

4.2 Revisions to Material in MEF 10.2 and MEF 10.2.1

Revisions to previous material include:

- All requirements are numbered,
- Values "Yes/No" and "True/False" are replaced with "Enabled/Disabled,"
- The terms "Class of Service Instance" and "Class of Service Frame Set" have been removed and replaced by "Class of Service Name,"
- The service attributes are grouped into EVC, UNI, and EVC per UNI categories (Sections 8, 9, and 10),
- The terminology of UNI Root type and UNI Leaf type are changed to UNI Root Role and UNI Leaf Role,
- When CE-VLAN ID Preservation is Enabled for an EVC, Ingress VLAN Tagged Data Service Frames whose CE-VLAN ID equals the CE-VLAN ID used for Untagged and Priority Tagged Service Frames are now exempted from having their CE-VLAN ID Preserved (Section 8.6.1),
- The Service Frame Delivery Service Attributes are revised to align them with the requirements of the Rooted-Multipoint EVC (Section 8.5),

- Details for the EVC Layer 2 Control Processing Service Attribute are replaced with a reference to MEF 6.1.1 [13] (Section 8.7),
- The definition of Qualified Service Frames is revised to improve clarity and logical consistency (Section 8.8),
- EVC Maximum Service Frame Size Service Attribute replaces the EVC Maximum Transmission Unit Size and makes the maximum size of the Service Frame format dependent (Section 8.9),
- Requirements limiting the length of identifiers are changed to 45 characters and requirements are added that such identifiers be RFC 2579 [10] DisplayStrings,
- UNI Maximum Service Frame Size Service Attribute replaces the UNI Maximum Transmission Unit Size and makes the maximum size of the Service Frame format dependent (Section 9.7),
- Details for the UNI Layer 2 Control Processing Service Attribute are replaced with a reference to MEF 6.1.1 [13] (Section 9.19),
- The description of the Class of Service Identifier is revised to improve clarity and align with MEF 23.1 [22] (Section 10.2),
- The concept of Egress Equivalence Class is introduced for specifying Egress Bandwidth Profiles (Section 10.4), and
- "Metro Ethernet Network" (MEN) is replaced by "Carrier Ethernet Network" (CEN). These two terms are equivalent and thus this change is editorial not technical.

5. 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 [8]. All key words must be in upper case, bold text.

Items that are **REQUIRED** (contain the words **MUST** or **MUST NOT**) are labeled as [Rx] for required. Items that are **RECOMMENDED** (contain the words **SHOULD** or **SHOULD NOT**) are labeled as [Dx] for desirable. Items that are **OPTIONAL** (contain the words **MAY** or **OPTIONAL**) are labeled as [Ox] for optional.

6. Numerical Prefix Conventions

This document uses the prefix notation to indicate multiplier values as shown in Table 2.

Decimal		Binary	
Symbol	Value	Symbol	Value
k	10^3	Ki	2^{10}
M	10^6	Mi	2^{20}
G	10^9	Gi	2^{30}
T	10^{12}	Ti	2^{40}
P	10^{15}	Pi	2^{50}
E	10^{18}	Ei	2^{60}
Z	10^{21}	Zi	2^{70}
Y	10^{24}	Yi	2^{80}

Table 2 – Numerical Prefix Conventions

7. Introduction

This document provides the model and framework for Ethernet Services. The model is built on the reference model as shown in Figure 1. In Figure 1, "Carrier Ethernet Network" is defined as a network from a Service Provider or network Operator supporting the MEF service and architecture models. This document addresses MEF services only from a Service Provider and thus the Subscriber sees a single CEN that is provided by a single Service Provider.




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.

[R1] A UNI **MUST** be dedicated to a single Subscriber.³

The CE and CEN exchange Service Frames across the UNI. A Service Frame is an Ethernet [3] frame transmitted across the UNI toward the Service Provider (called an ingress Service Frame) or an Ethernet [3] frame transmitted across the UNI toward the Subscriber (called an egress Service Frame). The Service Frame consists of the first bit of the Destination MAC Address through

³ This document describes attributes as seen by the Subscriber. MEF 26.1 [23], MEF 28 [24], and MEF 33 [26] describe attributes offered by an Operator to a Service Provider. See MEF 26.1 [23] for the details of the Operator/Service Provider model.

the last bit of the Frame Check Sequence. The allowed formats for the Service Frame are detailed in Section 9.6.

Throughout this document:

- When the field following the Source Address field is a TPID⁴ with the value 0x8100 and the corresponding VLAN ID value is not 0x000, the Service Frame is said to be a VLAN Tagged Service Frame.
- When the field following the Source Address field is a TPID⁴ with the value 0x8100 and the corresponding VLAN ID value is 0x000, the Service Frame is said to be a Priority Tagged Service Frame.
- A Service Frame that is either a VLAN Tagged Service Frame or a Priority Tagged Service Frame is said to be a Tagged Service Frame.
- When the field following the Source Address field is a TPID⁴ with the value 0x88a8, the Service Frame is said to be an S-Tagged Service Frame.
- When the two bytes following the Source Address field of the Service Frame do not contain the value 0x8100 or the value 0x88a8, the Service Frame is said to be an Untagged Service Frame.

Note that the definition of Untagged Service Frame means that a Service Frame with the two bytes following the Source Address field containing 0x88e7 (a Backbone Service Instance Tag per IEEE Std 802.1Q – 2011 [2]) is an Untagged Service Frame.

Also note that the behavior for S-Tagged Service Frames is beyond the scope of this document and the term "S-Tagged Service Frame" does not appear in the remainder of this document. Consequently, the behavior experienced by S-Tagged Service Frames can vary from Service Provider to Service Provider. A Subscriber who wants to use S-Tagged Service Frames is urged to check with his or her Service Provider to determine the behavior for such Service Frames.

There are no assumptions about the details of the Carrier Ethernet Network. It could consist of a single switch or an agglomeration of networks based on many different technologies. For an architectural perspective of the Carrier Ethernet Network see MEF 12.1. [17]

The UNIs between which Service Frames can be exchanged are specified by the Ethernet Virtual Connection (EVC). There are a number of types of EVCs and a number of service attributes that an EVC can have. These are described in Section 8.

There are a number of service attributes for a UNI. These are described in Section 9.

Some service attributes apply to an EVC per UNI. These are described in Section 10.

Performance Service Attributes that apply to multiple EVCs are described in Section 11.

⁴ TPID, Tag Protocol Identifier, is defined in IEEE Std 802.1Q – 2011 [2].

The details of the Bandwidth Profile parameters and algorithm are described in Section 12.

Section 13 contains a framework for defining a service that is based on the service attributes.

Extensive examples and supplementary explanations are contained in Appendix A, Appendix B, and Appendix C.

8. 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. The UNIs associated by an EVC 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 9.8.

An ingress Service Frame that is mapped to the EVC (see Section 9.10) can be delivered to zero or more UNIs in the EVC; and only to UNIs in the EVC other than the ingress UNI.

- [R2] If an egress Service Frame mapped to an EVC at a UNI results from an ingress Service Frame mapped to an EVC at a UNI, it **MUST** be the same EVC that associates the two UNIs.
- [R3] If an egress Service Frame mapped to an EVC at a UNI results from an ingress Service Frame mapped to an EVC at a UNI, the two UNIs **MUST** be different from each other.

An EVC always supports bi-directional transmission of Service Frames. That is, each UNI associated by the EVC always supports ingress and egress Service Frames for that EVC. Note that ingress Service Frames can originate at any UNI in the EVC.

8.1 Ethernet Virtual Connection Type Service Attribute

There are three types of EVC. They are as described in Sections 8.1.1, 8.1.2.1, and 8.1.2.2.

8.1.1 Point-to-Point EVC

In a Point-to-Point EVC, exactly two UNIs are associated with one another. 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

8.1.2 Multipoint EVCs

In a Multipoint EVC, two or more UNIs are associated with one another. A Multipoint-to-Multipoint EVC associating two UNIs is different from a Point-to-Point EVC because additional UNIs can be added to the Multipoint-to-Multipoint EVC.

8.1.2.1 Rooted-Multipoint EVC

Each UNI associated by a Rooted-Multipoint EVC has either a Root Role or a Leaf Role. At a UNI having the Root Role, ingress Service Frames do not have any restrictions as to the Roles of egress UNIs for their resulting egress Service Frames. In contrast, at a UNI having the Leaf Role, ingress Service Frames are restricted to egress UNIs having the Root Role for their resulting egress Service Frames.

- [R4] In a Rooted-Multipoint EVC, one or more of the UNIs **MUST** have the Root Role.
- [R5] In a Rooted-Multipoint EVC each of the UNIs not having the Root Role **MUST** have the Leaf Role.
- [R6] If an egress Service Frame mapped to a Rooted-Multipoint EVC at a UNI with the Leaf Role results from an ingress Service Frame mapped to the EVC at a UNI, the ingress UNI **MUST** have the Root Role.

If an ingress Service Frame at a UNI with the Leaf Role is delivered to an egress UNI with the Leaf Role, [R6] is violated. Thus [R6] means that ingress Service Frames at a UNI with the Leaf Role can be delivered to only UNIs that have the Root Role. Ingress Service Frames at a UNI with the Root Role can be delivered to any of the other UNIs associated by the Rooted-Multipoint 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 ingress Service Frame with a broadcast or multicast Destination Address at a Root Role UNI would be replicated in the Carrier 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 CEN has not yet learned an association of the destination MAC address with an <EVC, UNI> pair. Figure 3 illustrates a Rooted-Multipoint EVC with one Root UNI.

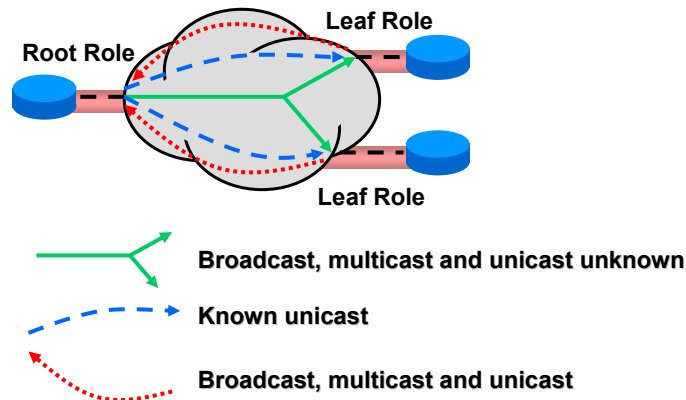


Figure 3 – Rooted-Multipoint EVC

8.1.2.2 Multipoint-to-Multipoint EVC

Each UNI associated by a Multipoint-to-Multipoint EVC has the Root Role. In a Multipoint-to-Multipoint 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 ingress Service Frame with a broadcast or multicast Destination Address at a given UNI would be replicated in the Carrier 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 CEN has not yet learned an association of the destination MAC address with an <EVC, UNI> pair. Figure 4 illustrates a Multipoint-to-Multipoint EVC.




Figure 4 – Multipoint-to-Multipoint EVC

8.2 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 CEN.

- [R7] The EVC ID **MUST** be unique across all EVCs in the CEN.
- [R8] The EVC ID **MUST** contain no more than 45 characters.⁵
- [R9] The EVC ID **MUST** be a non-null RFC 2579 [10] DisplayString but not contain the characters 0x00 through 0x1f.

⁵ The limit of 45 characters is intended to establish limits on field lengths in existing or future protocols that will carry the identifier.

The EVC ID is intended for management and control purposes.⁶ As an example, the Acme Service Provider might use “EVC-0001898-ACME-MEGAMART” to represent the 1898th EVC in the CEN and the customer for the EVC is MegaMart.

8.3 UNI List Service Attribute

The UNI List for an EVC is a list of pairs of the form <UNI ID (see Section 9.1), UNI Role>.

- [R10] The list **MUST** have exactly one pair of the form <UNI ID, UNI Role> for each UNI in the EVC.
- [R11] The UNI Role **MUST** have the value of either “Root” or “Leaf.”
- [R12] If the type of EVC is Point-to-Point or Multipoint-to-Multipoint, then the UNI Role **MUST** equal “Root.”

8.4 Maximum Number of UNIs Service Attribute

The Maximum Number of UNIs service attribute specifies the maximum number of UNIs allowed in the UNI List service attribute.

- [R13] For a Point-to-Point EVC, the Maximum Number of UNIs **MUST** be two.
- [R14] For a Multipoint EVC, the Maximum Number of UNIs **MUST** be three or greater.

8.5 Service Frame Delivery Service Attributes

8.5.1 Types of Service Frame

8.5.1.1 Layer 2 Control Protocol Service Frame

A Layer 2 Control Protocol Service Frame (L2CP Service Frame) is a Service Frame that could be used in a recognized Layer 2 Control Protocol. Given that there are several Layer 2 protocols used for various control purposes, it is important that Carrier Ethernet Networks be able to process such information effectively.⁷

- [R15] A Service Frame whose destination MAC address is one of the addresses listed in Table 3 **MUST** be treated as a Layer 2 Control Protocol Service Frame.

Some Layer 2 Control protocols share the same destination MAC address and are identified by additional fields such as the Ethertype. Therefore, disposition of Service Frames carrying Layer 2 Control Protocols can be different for different protocols that use the same destination MAC

⁶ For example, see Section 8.1 of MEF 30.1. [25]

⁷ This capability will be especially important for Subscribers who choose to deploy IEEE Std 802.1Q – 2011 [2] bridges (as opposed to routers) as CEs.

address. MEF 6.1.1 [13] contains some requirements and recommendations for the handling of specific Layer 2 Control protocols.

MAC Destination Addresses ⁸	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	MRP Block of protocols

Table 3 – List of Standardized Layer 2 Control Protocol Destination MAC Addresses

- [O1]** A Service Provider **MAY** define additional addresses for identifying Layer 2 Control Protocol Service Frames in addition to those in Table 3.

8.5.1.2 SOAM Service Frame

A Service Frame whose MAC Destination Address does not indicate the frame to be an L2CP Service Frame and whose Ethertype = 0x8902 is defined to be a SOAM Service Frame. (See Table 21-1 of IEEE Std 802.1Q – 2011 [2].) Sections 9.17, 10.10, and 10.11 contain requirements for the handling of certain SOAM Service Frames. Other requirements for handling SOAM Service Frames can be found in MEF 17 [19], MEF 30.1 [25], and MEF 31 [26].

8.5.1.3 Data Service Frame

A Service Frame that is not a Layer 2 Control Protocol Service Frame and not a SOAM Service Frame is defined to be a Data Service Frame. A Data Service Frame with a unicast Destination MAC Address is defined to be a Unicast Data Service Frame.⁹ A Data Service Frame with a multicast Destination MAC Address is defined to be a Multicast Data Service Frame.⁹ A Data Service Frame with a broadcast Destination MAC Address is defined to be a Broadcast Data Service Frame.⁹

8.5.2 Data Service Frame Disposition

This EVC Service Attribute applies to any ingress Data Service Frame that is not discarded per [R69], [R70], [R89], [R128], [R156], [R157], [D1], or [D2]. There are two possible values and corresponding dispositions:

- Discard: The Data Service Frame is discarded.
- Deliver Unconditionally: The Data Service Frame is delivered across the other (egress) UNI(s) provided [R6] is satisfied. This might be the behavior of a Point-to-Point EVC.
- Deliver Conditionally: The Data Service Frame is delivered across an egress UNI if certain conditions are met, provided [R6] is satisfied. An example of such a condition is that the destination MAC address is known by the Carrier Ethernet Network to be “at” the destination UNI. Another example is broadcast throttling

⁸Hexadecimal canonical format

⁹ See Section 9.6 for details about the format of the Service Frame.

where some Broadcast Data Service Frames are dropped to limit the amount of such traffic. A final example is multicast throttling where some Multicast Data Service Frames are dropped to limit the amount of such traffic.

- [R16] When Deliver Conditionally is Enabled the conditions that determine whether a Data Service Frame is delivered or discarded **MUST** be specified.
- [R17] The Data Service Frame disposition (Discard, Deliver Unconditionally or Deliver Conditionally) **MUST** be specified for Unicast Data Service Frames.
- [R18] The Data Service Frame disposition (Discard, Deliver Unconditionally or Deliver Conditionally) **MUST** be specified for Multicast Data Service Frames.
- [R19] The Service Frame disposition (Discard, Deliver Unconditionally or Deliver Conditionally) **MUST** be specified for Broadcast Data Service Frames.

[R17], [R18], and [R19] mean that the disposition of Unicast Data Service Frames can be different than the disposition of Multicast Data Service Frames which in turn can be different than the disposition of Broadcast Data Service Frames.

Note that this is a description of the ideal service. Data 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 8.8 and Section 11.

8.5.3 Data Service Frame Transparency

The requirements on the format of Service Frames are in Section 9.6.

- [R20] If the egress Data Service Frame resulting from an ingress Tagged Data Service Frame is also a Tagged Service Frame then fields in the egress Data Service Frame **MUST** be unchanged as indicated in Figure 5.

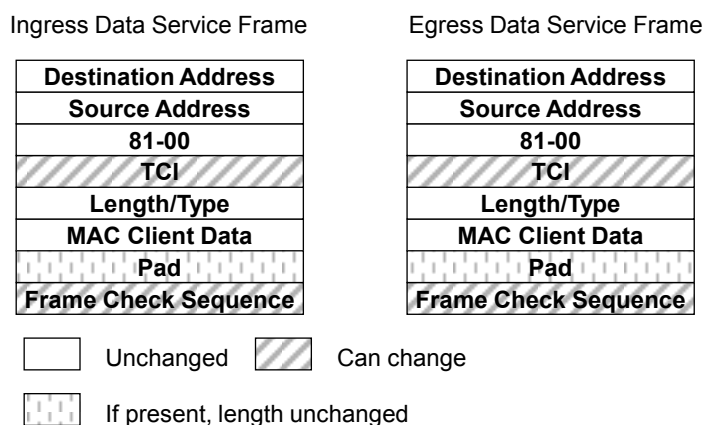


Figure 5 – Tagged In to Tagged Out Data Service Frame Transparency

- [R21] If the egress Data Service Frame resulting from an ingress Tagged Data Service Frame is an Untagged Service Frame then fields in the egress Data Service Frame **MUST** be unchanged as indicated in Figure 6.

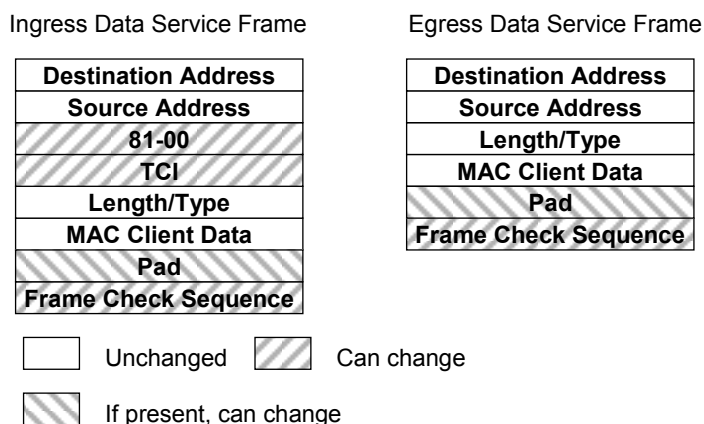


Figure 6 – Tagged In to Untagged Out Data Service Frame Transparency

- [R22] If the egress Data Service Frame resulting from an ingress Untagged Data Service Frame is a Tagged Data Service Frame then fields in the egress Data Service Frame **MUST** be unchanged as indicated in Figure 7.

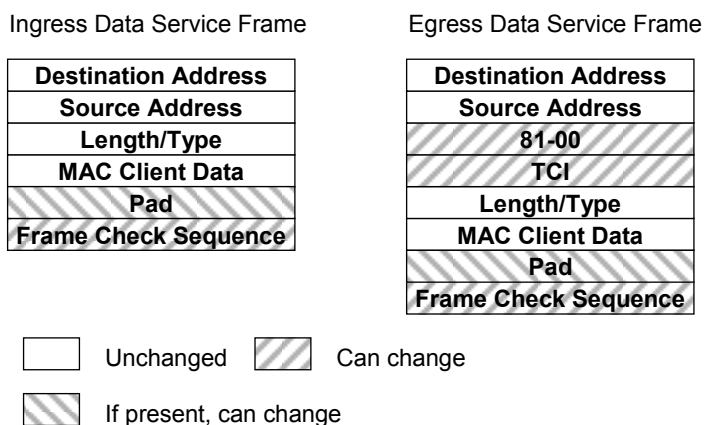


Figure 7 – Untagged In to Tagged Out Data Service Frame Transparency

- [R23] If the egress Data Service Frame resulting from an ingress Untagged Data Service Frame is also an Untagged Service Frame then fields in the egress Data Service Frame **MUST** be unchanged as indicated in Figure 8.

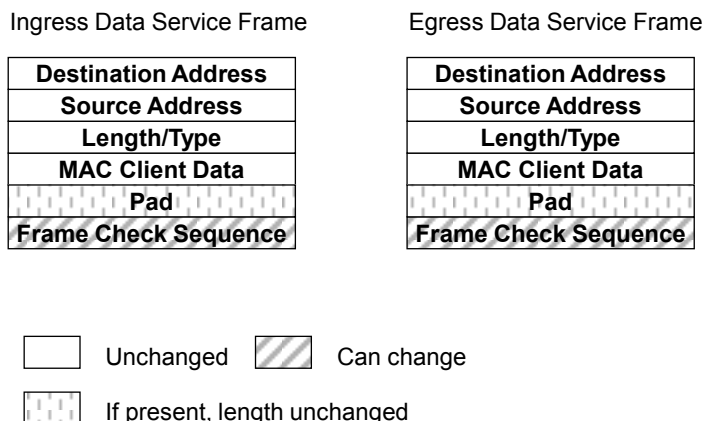


Figure 8 – Untagged In to Untagged Out Data Service Frame Transparency

However, specific attributes of an EVC can enforce conditions that additional fields must be identical at ingress and egress. See Section 8.6.

8.6 CE-VLAN Tag Preservation Service Attributes

Service Frames at the UNI may contain an IEEE Std 802.1Q – 2011 Customer VLAN Tag [2]. 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 is defined to be the VLAN ID. The VLAN ID indicates the Customer Edge VLAN ID (CE-VLAN ID) as specified in Section 9.9. The portion of the CE-VLAN Tag that contains the Priority Code Point bits is defined to be the Customer Edge VLAN Class of Service (CE-VLAN CoS). An EVC has two attributes related to CE-VLAN Tag Preservation as described in the following two subsections.

8.6.1 CE-VLAN ID Preservation Service Attribute

CE-VLAN ID Preservation can be either Enabled or Disabled.

A Service Frame is defined to have its CE-VLAN ID preserved when the relationship between the ingress Service Frame and its corresponding egress Service Frame(s) is as described in Table 4.

Ingress Service Frame	Egress Service Frame(s) ¹⁰
Untagged Service Frame	Untagged Service Frame
Tagged Service Frame	Tagged Service Frame with VLAN ID equal to the VLAN ID of the Tag on the ingress Service Frame

Table 4 – Behavior for CE-VLAN ID Preservation for a Service Frame

[R24] An EVC with CE-VLAN ID Preservation Enabled **MUST** preserve the CE-VLAN ID for Service Frames as described in Table 5. (See Section 9.13 for the description of All to One Bundling.)

¹⁰ Note that in the case of a Multipoint EVC, a single ingress Service Frame can result in more than one egress Service Frame.

CE-VLAN ID/EVC Map Characteristic	Service Frames with CE-VLAN ID Preserved
All to One Bundling Enabled at all UNIs ¹¹	All Data Service Frames
All other cases	All VLAN Tagged Data Service Frames with VLAN ID in the range 1 – 4094 excluding the CE-VLAN ID used for Untagged and Priority Tagged Service Frames

Table 5 – CE-VLAN ID Preservation Service Attribute for an EVC

- [R25]** 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 9.12 and 9.13), the EVC **MUST** have CE-VLAN ID Preservation Enabled.

Note that when the CE-VLAN ID configured for Untagged and Priority Tagged Service Frames (see Section 9.9) is mapped to an EVC with the CE-VLAN ID Preservation Enabled, there is no mandate for CE-VLAN ID Preservation for ingress Untagged and Priority Tagged Service Frames at this UNI except in the case of All to One Bundling Enabled.

An obvious benefit of CE-VLAN ID Preservation Enabled is enhanced operational simplicity. For example, for a Subscriber connecting multiple campuses using IEEE Std 802.1Q – 2011 bridges, the feature obviates the task of renumbering VLANs in different corporate campuses.

8.6.2 CE-VLAN CoS Preservation Service Attribute

CE-VLAN CoS Preservation can be either Enabled or Disabled.

- [R26]** In an EVC with CE-VLAN CoS Preservation Enabled, an egress Service Frame resulting from an ingress Service Frame that contains a CE-VLAN CoS **MUST** have the identical CE-VLAN CoS for all possible values of CE-VLAN CoS.

Since the CE-VLAN CoS refers to the PCP field in the CE-VLAN Tag of a Service Frame, [R26] means that the egress Service Frame is a Tagged Service Frame and the PCP field value is identical to that of the corresponding ingress Tagged Service Frame when CE-VLAN CoS Preservation is Enabled. Note that [R26] does not apply to ingress Untagged Service Frames.

When CE-VLAN CoS Preservation is Enabled and the Color Identifier (Section 10.3) is based on the CE-VLAN Tag PCP field, the Service Provider cannot change the Color of a Service Frame. However, basing the Color Identifier on the CE-VLAN Tag DEI would allow the Service Provider to change the Color of a Service Frame.

8.7 EVC Layer 2 Control Protocol Processing Service Attribute

Requirements for the processing of L2CP Service Frames are contained in MEF 6.1.1.[13]

¹¹ Per [R83], if one of the UNIs associated by an EVC has All to One Bundling Enabled, then all of the UNIs associated by the EVC have All to One Bundling Enabled.

8.8 EVC Related Performance Service Attributes

The performance for an EVC is specified using one or more Performance Metrics. A Performance Metric is a quantitative characterization of Service Frame delivery quality. This section specifies the following Service Frame Performance Metrics:

- One-way Frame Delay Performance (Section 8.8.1),
- One-way Frame Delay Range Performance (Section 8.8.1),
- One-way Mean Frame Delay Performance (Section 8.8.1),
- One-way Inter-Frame Delay Variation Performance (Section 8.8.2),
- One-way Frame Loss Ratio Performance (Section 8.8.3)
- One-way Availability Performance (Section 8.8.4),
- One-way Resiliency Performance expressed as High Loss Intervals (Section 8.8.5),
- One-way Resiliency Performance expressed as Consecutive High Loss Intervals (Section 8.8.5), and
- One-way Group Availability Performance (Section 8.8.6).

The technical specification of the service level being offered by the Service Provider to the Subscriber is defined to be the Service Level Specification (SLS). For any given SLS, performance objectives related to the Performance Metrics may or may not be specified.

[R27] If an SLS contains an objective for a given performance attribute, then the SLS **MUST** specify the related parameters for that objective.

An SLS can contain multiple objectives for a given Performance Metric, but one or more of the parameters associated with each objective for a given Performance Metric need to be different from each other. For example, an SLS could contain two objectives for One-way Frame Delay Performance each corresponding to a different value of P_d (see Section 8.8.1).

However, certain parameters defined in Section 8.8.4 can assume only a single value in an SLS per the following three requirements:

[R28] An SLS **MUST** use a single value for Δt for all Performance Metrics where Δt is used.

[R29] An SLS **MUST** use a single value for C for all Performance Metrics where C is used.

[R30] An SLS **MUST** use a single value for n for all Performance Metrics where n is used.

Performance Attributes, with the exception of Availability,¹² apply to Qualified Service Frames, which are frames that meet the following criteria for a given Class of Service Name (See Section 10.2):

[R31] An SLS **MUST** define Qualified Service Frames as every Service Frame that satisfies the following criteria for a given EVC, a given ordered UNI pair $\langle i, j \rangle$ associated by the EVC, a given time interval T , and a given Class of Service Name:

- Each Service Frame **MUST** ingress at UNI i ,
- Each Service Frame **MUST** map to the given EVC at UNI i via the CE-VLAN ID/EVC Map. (See Section 9.10.),
- Each Service Frame **MUST NOT** be discarded per requirements [R69], [R70], [R89], [R128], [D1], [D2], or any of the conditions specified per [R16].
- The first bit of each Service Frame **MUST** arrive at UNI i within the time interval T , and within a time interval Δt smaller than T that has been designated as part of Available Time for $\langle i, j \rangle$ (see Section 8.8.4),
- Each Service Frame **MUST** have the given Class of Service Name (See Section 10.2 for the description of how the Class of Service Name is determined from the Class of Service Identifier for a Service Frame.),
- Each Service Frame that is subject to an Ingress Bandwidth Profile **MUST** have an Ingress Bandwidth Profile compliance of Green, and
- Each Service Frame that is not subject to an Ingress Bandwidth Profile **MUST** meet one of the following two conditions:
 - The Service Frame has no Color Identifier, or
 - The Service Frame has a Color (see Section 10.3) that is Green per the color indication requirements of MEF 23.1[22].

Note that the determination of Qualified Service Frames depends on the values of Δt , C and n via Section 8.8.4. Thus any SLS that contains a performance objective for a Performance Metric based on Qualified Service Frames will need to specify values for these parameters. The only Performance Metric that does not depend on Qualified Service Frames is Availability but this Performance Metric depends directly on Δt , C and n . Thus it follows that an SLS that contains one or more performance objectives will need to specify values for Δt , C and n .

[R32] An SLS **MUST** use a single value of T for all Performance Metrics.

¹² Availability is used to define Qualified Service Frames.

Performance objectives are defined for non-empty subsets of the ordered pairs of UNIs that are associated by the given EVC. Appendix A.4 provides examples on how to structure these subsets to be UNI-oriented and EVC-oriented.

8.8.1 One-way Frame Delay Performance for an EVC

This section defines three Performance Metrics: the One-way Frame Delay Performance corresponding to a percentile of the distribution, the One-way Mean Frame Delay, and the One-way Frame Delay Range.

The One-way Frame Delay for an egress Service Frame at a given UNI in the EVC is defined as the time elapsed from the transmission at the ingress UNI of the first bit of the corresponding ingress Service Frame until the reception of the last bit of the Service Frame at the given UNI. Recall that the Service Frame consists of the first bit of the Destination Address through the last bit of the Frame Check Sequence. This delay definition is illustrated in Figure 9.

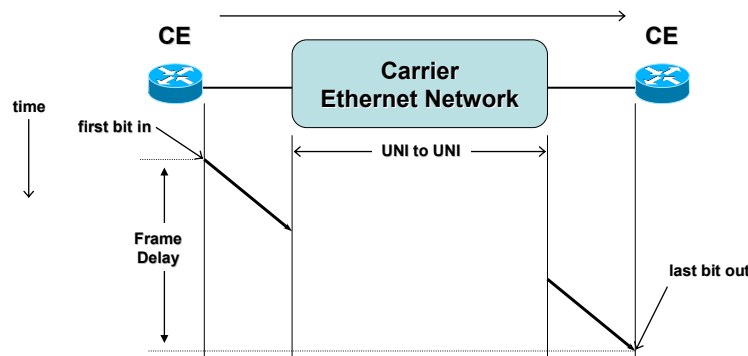


Figure 9 - Frame Delay for Service Frame

Note that this definition of Frame Delay for a Service Frame is the one-way¹³ delay that includes the delays encountered as a result of transmission of the Service Frame across the ingress and egress UNIs as well as that introduced by the CEN.

Each such Performance Metric is based on a subset of the ordered pairs of UNIs in the EVC for a time interval T .

[R33] For a given non-empty set of ordered pairs of UNIs S , a time interval T , and a given Class of Service Name, the SLS **MUST** define each Frame Delay Performance Metric as follows:

- Let the UNIs in the EVC be numbered from 1 to m . And let S be a non-empty subset of the ordered UNI pairs in the EVC. That is $S \subseteq \{ \langle i, j \rangle \mid i = 1, \dots, m, j = 1, \dots, m, i \neq j \}, S \neq \emptyset$.

¹³ One-way delay is difficult to measure and therefore one way delay may be approximated from two way measurements. However measurement techniques are beyond the scope of this document.

- Let $\bar{d}_{Td}^{(i,j)}$ represent the P_d -Percentile of one-way delay for all Service Frames delivered to UNI j resulting from an ingress Qualified Service Frame at UNI i . If there are no such egress Service Frames at UNI j , then $\bar{d}_{Td}^{(i,j)} = \text{Undefined}$.
- Then the One-way Frame Delay Performance Metric **MUST** be defined as the maximum value of all of the defined values $\bar{d}_{Td}^{(i,j)}$ for $\langle i, j \rangle \in S$, unless all $\bar{d}_{Td}^{(i,j)}$ are Undefined in which case the performance is Undefined.
- Let $\bar{d}_{TR}^{(i,j)} = \bar{d}_{Tr}^{(i,j)} - \bar{d}_{\min}^{(i,j)}$. The term $\bar{d}_{Tr}^{(i,j)}$ represents the P_r -Percentile of the one-way delay for all Service Frames delivered to UNI j resulting from an ingress Qualified Service Frame at UNI i . The term $\bar{d}_{\min}^{(i,j)}$ is the minimum of the one-way delays for all Service Frames delivered to UNI j resulting from an ingress Qualified Service Frame at the UNI i . If there are no such egress Service Frames at UNI j resulting from ingress Qualified Service Frames at UNI i , then $\bar{d}_{TR}^{(i,j)} = \text{Undefined}$.
- Then the One-way Frame Delay Range Performance Metric **MUST** be defined as the maximum value of all of the defined values of $\bar{d}_{TR}^{(i,j)}$ for $\langle i, j \rangle \in S$, unless all $\bar{d}_{TR}^{(i,j)}$ are Undefined in which case the performance is Undefined.
- Let $\bar{\mu}_T^{(i,j)}$ represent the arithmetic mean of one-way delay for all Service Frames delivered to UNI j resulting from an ingress Qualified Service Frame at UNI i . If there are no such egress Service Frames at UNI j resulting from ingress Qualified Service Frames at UNI i , then $\bar{\mu}_T^{(i,j)} = \text{Undefined}$.
- Then the One-way Mean Frame Delay Performance Metric **MUST** be defined as the maximum value of all of the values $\bar{\mu}_T^{(i,j)}$ for $\langle i, j \rangle \in S$, unless all $\bar{\mu}_T^{(i,j)}$ are Undefined in which case the performance is Undefined.

To restate the Frame Delay definition mathematically, let the UNIs in the EVC be numbered from 1 to m and let $D_T^{(i,j)}$ be the set of one-way Frame Delay values for all egress Service Frames at UNI j resulting from ingress Qualified Service Frames at UNI i . $D_T^{(i,j)}$ can be expressed as $D_T^{(i,j)} = \{d_1^{(i,j)}, d_2^{(i,j)}, \dots, d_{N_T^{(i,j)}}^{(i,j)}\}$, where $d_k^{(i,j)}$ is the one-way Frame Delay of the k^{th} Service Frame. Define $\bar{d}_{Td}^{(i,j)}$ for $P_d > 0$ as

$$\bar{d}_{Td}^{(i,j)} = \begin{cases} \min \left\{ d \in D_T^{(i,j)} \mid P_d \leq \frac{100}{N_T^{(i,j)}} \sum_{k=1}^{N_T^{(i,j)}} I(d, d_k^{(i,j)}) \right\} & \text{if } N_T^{(i,j)} \geq 1 \\ \text{Undefined} & \text{otherwise} \end{cases}$$

where $I(\cdot, d_k)$ is an indicator function defined by

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

$\bar{d}_{Td}^{(i,j)}$ is the minimal delay during the time interval T that P_d percent of the frames do not exceed.

Then a one-way Frame Delay Performance Metric for an EVC can be expressed as

$$\bar{d}_{T,S} = \begin{cases} \max \{ \bar{d}_{Td}^{(i,j)} \mid \langle i, j \rangle \in S \text{ and where } N_T^{(i,j)} > 0 \} \\ Undefined & \text{when all } N_T^{(i,j)} = 0 \text{ for all } \langle i, j \rangle \in S \end{cases}$$

Another parameter is the objective for the difference between the P_r percentile delay and

$\bar{d}_{\min}^{(i,j)} = \min \{ d \in D_T^{(i,j)} \}$, expressed as

$$\bar{d}_{TR}^{(i,j)} = \begin{cases} (\bar{d}_{Tr}^{(i,j)} - \bar{d}_{\min}^{(i,j)}) & \text{if } N_T^{(i,j)} > 0 \\ Undefined & \text{if } N_T^{(i,j)} = 0 \end{cases}$$

where

$$\bar{d}_{Tr}^{(i,j)} = \begin{cases} \min \left\{ d \in D_T^{(i,j)} \mid P_r \leq \frac{100}{N_T^{(i,j)}} \sum_{k=1}^{N_T^{(i,j)}} I(d, d_k^{(i,j)}) \right\} & \text{if } N_T^{(i,j)} \geq 1 \\ Undefined & \text{otherwise} \end{cases}$$

Then a one-way Frame Delay Range Performance Metric for an EVC can be expressed as

$$\bar{d}_{TR,S} = \begin{cases} \max \{ \bar{d}_{TR}^{(i,j)} \mid \langle i, j \rangle \in S \text{ and where } N_T^{(i,j)} > 0 \} \\ Undefined & \text{when all } N_T^{(i,j)} = 0 \text{ for all } \langle i, j \rangle \in S \end{cases}$$

Another One-way Frame Delay attribute is the arithmetic mean of $D_T^{(i,j)}$, which can be expressed as

$$\bar{\mu}_T^{(i,j)} = \begin{cases} \frac{1}{N_T^{(i,j)}} \sum_{k=1}^{N_T^{(i,j)}} (d_k^{(i,j)}) & \text{if } N_T^{(i,j)} > 0 \\ Undefined & \text{if } N_T^{(i,j)} = 0 \end{cases}$$

Then a One-way Mean Frame Delay Performance Metric for an EVC can be expressed as

$$\bar{\mu}_{T,S} = \begin{cases} \max \left\{ \bar{\mu}_T^{\langle i,j \rangle} \mid \langle i,j \rangle \in S \text{ and where } N_T^{\langle i,j \rangle} > 0 \right\} \\ \text{Undefined} & \text{when all } N_T^{\langle i,j \rangle} = 0 \text{ for all } \langle i,j \rangle \in S \end{cases}$$

The parameters of these three One-way Frame Delay Performance Metrics are given in Table 6.

Parameter	Description
T	The time interval
S	Non-empty subset of the ordered UNI pairs in the EVC
CoS_Name	The Class of Service Name
P_d	A specific percentile > 0 for One-way Frame Delay Performance
P_r	A specific percentile > 0 for One-way Frame Delay Range Performance
\hat{d}	One-way Frame Delay Performance Objective corresponding to P_d
\hat{d}_R	One-way Frame Delay Range Performance Objective corresponding to P_r
$\hat{\mu}$	One-way Mean Frame Delay Performance Objective

Table 6 – One-way Frame Delay Performance Parameters

- [R34]** Given T , S , CoS_Name , P_d , and a one-way Frame Delay Performance objective \hat{d} , expressed in time units, an SLS **MUST** define the one-way Frame Delay Performance objective as met over the time interval T for the subset S if and only if $\bar{d}_{T,S} \leq \hat{d}$ or $\bar{d}_{T,S}$ is Undefined.
- [R35]** Given T , S , CoS_Name , P_r , and a One-way Frame Delay Range Performance objective \hat{d}_R , expressed in time units, an SLS **MUST** define the one-way Frame Delay Range Performance objective as met over the time interval T for the subset S if and only if $\bar{d}_{TR,S} \leq \hat{d}_R$ or $\bar{d}_{TR,S}$ is Undefined.
- [R36]** Given T , S , CoS_Name , and a One-way Mean Frame Delay Performance objective $\hat{\mu}$, expressed in time units, an SLS **MUST** define the Mean Frame Delay Performance as met over the time interval T for the subset S if and only if $\bar{\mu}_{T,S} \leq \hat{\mu}$ or $\bar{\mu}_{T,S}$ is Undefined.

Recall that if any of $\bar{d}_{Td}^{\langle i,j \rangle}$, $\bar{d}_{Tr}^{\langle i,j \rangle}$, or $\bar{\mu}_T^{\langle i,j \rangle}$ are Undefined for time interval T and ordered pair $\langle i,j \rangle$, then the performance for that ordered pair is to be excluded from calculations on the performance of pairs in S .

As noted in [R33] and Table 6, the subset S has to be a non-empty subset of the ordered pairs of UNIs in the EVC. This is the only requirement on S for Point-to-Point and Multipoint-to-Multipoint EVCs. For example, S can contain just one of the ordered UNI pairs in a Point-to-Point EVC. Rooted-Multipoint EVCs have an additional requirement on S per [R37].

- [R37] For a Rooted-Multipoint EVC, S **MUST** be a non-empty subset of the ordered pairs of UNIs in the EVC such that all ordered pairs in S contain at least one UNI that has the Root Role.

8.8.2 One-way Inter-Frame Delay Variation Performance for an EVC

One-way Inter-Frame Delay Variation (IFDV) is the difference between the one-way delays of a pair of selected Service Frames. This definition is borrowed from RFC 3393 [11] where IP packet delay variation is defined.

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

$$|a_k - a_l| = \Delta\tau$$

Let r_q be the time frame q is successfully received (last bit of the frame) at the egress UNI and let d_q be the delay for frame q , then the difference in the delays encountered by frame k and frame l is given by $d_k - d_l$. Define

$$\Delta d_{kl} = |d_k - d_l| = |(r_k - a_k) - (r_l - a_l)| = |(a_l - a_k) - (r_l - r_k)|$$

A positive value for $d_k - d_l$ implies that the two frames are closer together at the egress UNI while a negative value 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 Δd_{kl} is not defined and does not contribute to the evaluation of the One-way Inter-Frame Delay Variation.

Figure 10 shows a depiction of the different times that are related to One-way Inter-Frame Delay Variation Performance.




Figure 10 – One-way Inter-Frame Delay Variation Definition

- [R38] For a given non-empty set of ordered pairs of UNIs S , a time interval T , and a given Class of Service Name, the SLS **MUST** define the One-way Inter-Frame Delay Variation Performance Metric as follows:

- Let the UNIs in the EVC be numbered from 1 to m . And let S be a non-empty subset of the ordered UNI pairs in the EVC. That is $S \subseteq \{\langle i, j \rangle \mid i = 1, \dots, m, j = 1, \dots, m, i \neq j\}, S \neq \emptyset$.
- Let $\Delta \tilde{d}_T^{\langle i, j \rangle}$ be the P_v -percentile of the Δd_{kl} 's of all Qualified Service Frame pairs where each Qualified Service Frame ingresses at UNI i and results in an egress Service Frame at UNI j and whose difference in the arrival times of the first bit of each Service Frame in the pair at UNI i was $\Delta \tau$.
- If there are no such pairs of Service Frames for UNI i and UNI j , then $\Delta \tilde{d}_T^{\langle i, j \rangle} = \text{Undefined}$.
- Then the One-way Inter-Frame Delay Variation Performance Metric **MUST** be the maximum of the values $\Delta \tilde{d}_T^{\langle i, j \rangle}$ for $\langle i, j \rangle \in S$, unless all $\Delta \tilde{d}_T^{\langle i, j \rangle}$ are Undefined in which case the performance is Undefined.

This definition is in agreement with the IP packet delay variation definition given in RFC3393 [11] 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 $\Delta \tau$ can be related to the application timing information. As an example for voice applications where voice frames are generated at regular intervals, $\Delta \tau$ may be chosen to be few multiples of the inter-frame time.

To restate the definition mathematically, let the UNIs in the EVC be numbered from 1 to m . And let S be a non-empty subset of the ordered UNI pairs in the EVC. That is

$$S \subseteq \{\langle i, j \rangle \mid i = 1, \dots, m, j = 1, \dots, m, i \neq j\}, S \neq \emptyset$$

Let

$$V_T^{\langle i, j \rangle} = \{\Delta d_1^{\langle i, j \rangle}, \Delta d_2^{\langle i, j \rangle}, \dots, \Delta d_{N_T^{\langle i, j \rangle}}^{\langle i, j \rangle}\}$$

be the set of all absolute value of delay variations for all eligible pairs of Qualified Service Frames from UNI i to UNI j where the difference in the arrival times of the first bit of each Service Frame at the ingress UNI was exactly $\Delta \tau$. Define

$$\Delta \tilde{d}_T^{\langle i, j \rangle} = \begin{cases} \min \left\{ d \in V_T^{\langle i, j \rangle} \mid P_v \leq \frac{100}{N_T^{\langle i, j \rangle}} \sum_{k=1}^{N_T^{\langle i, j \rangle}} I(d, \Delta d_k^{\langle i, j \rangle}) \right\} & \text{if } N_T^{\langle i, j \rangle} \geq 1 \\ \text{Undefined} & \text{otherwise} \end{cases}$$

Where $I(d, \Delta d)$ is an indicator function defined by

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

Then a One-way Inter-Frame Delay Variation Performance Metric for an EVC can be expressed as

$$\Delta \tilde{d}_{T,S} = \begin{cases} \max \{ \Delta \tilde{d}_T^{\langle i,j \rangle} \mid \langle i,j \rangle \in S \text{ and where } N_T^{\langle i,j \rangle} \geq 1 \} \\ \text{Undefined} & \text{when all } N_T^{\langle i,j \rangle} = 0 \text{ for all } \langle i,j \rangle \in S \end{cases}$$

The parameters and objective for a One-way Inter-Frame Delay Variation Performance Metric are given in Table 7.

Parameter	Description
T	The time interval
S	Non-empty subset of the ordered UNI pairs in the EVC
CoS_Name	The Class of Service Name
P_v	One-way Inter-Frame Delay Variation Performance percentile
$\Delta \tau$	The separation between frame pairs for which One-way Inter-Frame Delay Variation Performance is defined
\tilde{d}	One-way Inter-Frame Delay Variation Performance Objective

Table 7 – One-way Inter-Frame Delay Variation Parameters

[R39] Given T , S , CoS_Name , P_v , $\Delta \tau$, and \tilde{d} , the One-way Inter-Frame Delay Variation Performance **MUST** be defined by an SLS as met over the time interval T for the subset S if and only if $\Delta \tilde{d}_{T,S} \leq \tilde{d}$ or $\Delta \tilde{d}_{T,S}$ is Undefined.

Recall that if the One-way Inter-Frame Delay Variation is Undefined for time interval T and ordered pair $\langle i,j \rangle$, then the performance for that ordered pair is excluded from calculations on the performance of pairs in S .

As noted in [R38] and Table 7, the subset S has to be a non-empty subset of the ordered pairs of UNIs in the EVC. This is the only requirement on S for Point-to-Point and Multipoint-to-Multipoint EVCs. For example, S can contain just one of the ordered UNI pairs in a Point-to-Point EVC. Rooted-Multipoint EVCs have an additional requirement on S per [R40].

[R40] For a Rooted-Multipoint EVC, S **MUST** be a non-empty subset of the ordered pairs of UNIs in the EVC such that all ordered pairs in S contain at least one UNI that has the Root Role.

8.8.3 One-way Frame Loss Ratio Performance for an EVC

[R41] For a given non-empty set of ordered pairs of UNIs S , a time interval T , and a given Class of Service Name, the SLS **MUST** define the One-way Frame Loss Ratio Performance Metric as follows

- Let the UNIs in the EVC be numbered from 1 to m . And let S be a non-empty subset of the ordered UNI pairs in the EVC. That is $S \subseteq \{\langle i, j \rangle \mid i = 1, \dots, m, j = 1, \dots, m, i \neq j\}, S \neq \emptyset$.
- Let $I_T^{\langle i, j \rangle}$ denote the number of ingress Qualified Service Frames at ingress UNI i that should have been delivered to UNI j according to the Service Frame Delivery service attributes (see Section 8.5).
- Let $E_T^{\langle i, j \rangle}$ be the number of unique (not duplicate) egress Service Frames where each frame is the first egress Service Frame delivered to UNI j that results from a Service Frame counted in $I_T^{\langle i, j \rangle}$.

- Define $FLR_T^{\langle i, j \rangle} = \begin{cases} \left(\frac{I_T^{\langle i, j \rangle} - E_T^{\langle i, j \rangle}}{I_T^{\langle i, j \rangle}} \right) \times 100 & \text{if } I_T^{\langle i, j \rangle} \geq 1 \\ Undefined & \text{otherwise} \end{cases}$.

- Then the One-way Frame Loss Ratio Performance Metric **MUST** be defined as $FLR_{T,S} = \begin{cases} \max \{ FLR_T^{\langle i, j \rangle} \mid \langle i, j \rangle \in S \text{ and where } I_T^{\langle i, j \rangle} \geq 1 \} \\ Undefined & \text{when all } I_T^{\langle i, j \rangle} = 0 \text{ for all } \langle i, j \rangle \in S \end{cases}$.

In the case of a Multipoint-to-Multipoint or a Rooted-Multipoint EVC, the Service Provider and the Subscriber can agree to define $FLR_T^{\langle i, j \rangle}$ as

$$FLR_T^{\langle i, j \rangle} = \begin{cases} \left(\frac{\tilde{I}_T^{\langle i, j \rangle} - E_T^{\langle i, j \rangle}}{\tilde{I}_T^{\langle i, j \rangle}} \right) \times 100 & \text{if } \tilde{I}_T^{\langle i, j \rangle} \geq 1 \\ Undefined & \text{otherwise} \end{cases}$$

where $\tilde{I}_T^{\langle i, j \rangle} = I_T^{\langle i, j \rangle} -$ the number of frames counted in $I_T^{\langle i, j \rangle}$ discarded by the Service Provider, in order to conform to either the UNI Line Rate of UNI j or an Egress Bandwidth Profile (if one is used) at UNI j . Such Service Frame drops may occur anywhere in the network, not just close to UNI j . One example of this could be where an Egress Bandwidth Profile is applied on a link within the network. Another example of this could be where excessive frames exceed the line rate on a link within the network. Good traffic engineering principles would suggest dropping such excessive frames as close to the ingress as possible. This adjustment is meant to ac-

count for a focused overload of traffic sent to UNI j from multiple ingress UNIs. The details of such an adjustment are beyond the scope of this document.

The parameters and objective of a One-way Frame Loss Ratio Performance Metric are given in Table 8.

Parameter	Description
T	The time interval
S	Non-empty subset of the ordered UNI pairs in the EVC
CoS_Name	The Class of Service Name
\hat{L}	One-way Frame Loss Ratio Performance objective

Table 8 – One-way Frame Loss Ratio Performance Parameters

- [R42] Given T , S , CoS_Name , and a One-way Frame Loss Ratio Performance objective \hat{L} , the One-way Frame Loss Performance **MUST** be defined as met over the time interval T for the subset S if and only if $FLR_{T,S} \leq \hat{L}$ or $FLR_{T,S}$ is Undefined.

Recall that if the One-way Frame Loss Ratio Performance is Undefined for time interval T and ordered pair $\langle i, j \rangle$, then the performance for that ordered pair is excluded from calculations on the performance of pairs in S .

As noted in [R41] and Table 8, the subset S has to be a non-empty subset of the ordered pairs of UNIs in the EVC. This is the only requirement on S for Point-to-Point and Multipoint-to-Multipoint EVCs. For example, S can contain just one of the ordered UNI pairs in a Point-to-Point EVC. Rooted-Multipoint EVCs have an additional requirement on S per [R43].

- [R43] For a Rooted-Multipoint EVC, S **MUST** be a non-empty subset of the ordered pairs of UNIs in the EVC such that all ordered pairs in S contain at least one UNI that has the Root Role.

8.8.4 One-way Availability Performance for an EVC

Availability Performance is the percentage of time within a specified time interval during which the Service Frame loss is small. (The precise definition is presented in the following paragraphs.) As an example, a service provider can define the Availability Performance to be measured over a month and the value for the Availability Performance objective to be 99.9%. In a month with 30 days and no Maintenance Interval (defined below) this objective will allow the service to be Unavailable for approximately 43 minutes out of the whole month.

Informally, Availability Performance is based on Service Frame loss during a sequence of consecutive small time intervals. When the previous sequence was defined as Available, if the frame loss is high for each small time interval in the current sequence, then the small time interval at the beginning of the current sequence is defined as Unavailable; otherwise it is defined as Available. On the other hand, when the previous sequence was defined as Unavailable, if frame loss is low for each small time interval in the current sequence, then the small time interval at the be-

ginning of the current sequence is defined as Available; otherwise, it is defined as Unavailable. The formal definition follows.

For a time interval T , and a given Class of Service Name, Availability from ingress UNI i to egress UNI j is based on the following three parameters:

- Δt , a time interval much smaller than T ,
- C , a frame loss ratio threshold which if exceeded suggests unavailability,
- n , the number of consecutive small time intervals, Δt , over which to assess availability.

Each Δt_k in T is defined to be either “Available” or “Unavailable” and this is represented by $A_{\langle i,j \rangle}(\Delta t_k)$ where $A_{\langle i,j \rangle}(\Delta t_k) = 1$ means that Δt_k is Available and $A_{\langle i,j \rangle}(\Delta t_k) = 0$ means that Δt_k is Unavailable.

The definition of $A_{\langle i,j \rangle}(\Delta t_k)$ is based on the frame loss ratio function, $flr_{\langle i,j \rangle}(\Delta t_k)$, which is defined as follows.

Let $I_{\Delta t}^{(i,j)}$ be the number of ingress Service Frames that meet the following conditions:

- The first bit of each Service Frame arrives at UNI i within the time interval Δt ,
- Each Service Frame is mapped to the EVC via the CE-VLAN ID/EVC Map,
- Each Service Frame should be delivered to UNI j according to the frame delivery service attributes (see Section 8.5),
- Each Service Frame has the given Class of Service Name,
- Each Service Frame is not discarded per [R69], [R70], [R128], [D1], [D2], [O4], or any of the conditions specified per [R16],
- Each Service Frame that is subject to an Ingress Bandwidth Profile has an Ingress Bandwidth Profile compliance of Green, and
- Each Service Frame that is not subject to an Ingress Bandwidth Profile either has no Color Identifier or a Color of Green per the color indication requirements of MEF 23.1 [22].

Let $E_{\Delta t}^{(i,j)}$ be the number of unique (not duplicate) egress Service Frames where each Service Frame is the first unerrored egress Service Frame at UNI j that results from a Service Frame counted in $I_{\Delta t}^{(i,j)}$.

$$\text{Then } flr_{\langle i,j \rangle}(\Delta t) = \begin{cases} \left(\frac{I_{\Delta t}^{\langle i,j \rangle} - E_{\Delta t}^{\langle i,j \rangle}}{I_{\Delta t}^{\langle i,j \rangle}} \right) & \text{if } I_{\Delta t}^{\langle i,j \rangle} \geq 1 \\ 0 & \text{otherwise} \end{cases}.$$

In the case of a Multipoint-to-Multipoint or a Rooted-Multipoint EVC, the Service Provider and the Subscriber can agree to define $flr_{\langle i,j \rangle}(\Delta t)$ as

$$flr_{\langle i,j \rangle}(\Delta t) = \begin{cases} \left(\frac{\tilde{I}_{\Delta t}^{\langle i,j \rangle} - E_{\Delta t}^{\langle i,j \rangle}}{\tilde{I}_{\Delta t}^{\langle i,j \rangle}} \right) & \text{if } \tilde{I}_{\Delta t}^{\langle i,j \rangle} \geq 1 \\ 0 & \text{otherwise} \end{cases}$$

where $\tilde{I}_{\Delta t}^{\langle i,j \rangle} = I_{\Delta t}^{\langle i,j \rangle} -$ the number of Qualified Service Frames discarded by the Service Provider, in order to conform to either the UNI Line Rate of UNI j or an Egress Bandwidth Profile (if one is used) at UNI j . Such Service Frame drops may occur anywhere in the network, not just close to UNI j . One example of this could be where an Egress Bandwidth Profile is applied on a link within the network. Another example of this could be where excessive frames exceed the line rate on a link within the network. Good traffic engineering principles suggest dropping such excessive frames as close to the ingress as possible. This adjustment is meant to account for a focused overload of traffic sent to UNI j from multiple ingress UNIs. The details of such an adjustment are beyond the scope of this document.

Δt_0 is the first short time interval agreed by the Service Provider and Subscriber at or after turn-up of the EVC. $A_{\langle i,j \rangle}(\Delta t_k)$ is defined by the flow diagram in Figure 11 for $k = 0, 1, 2, \dots$.

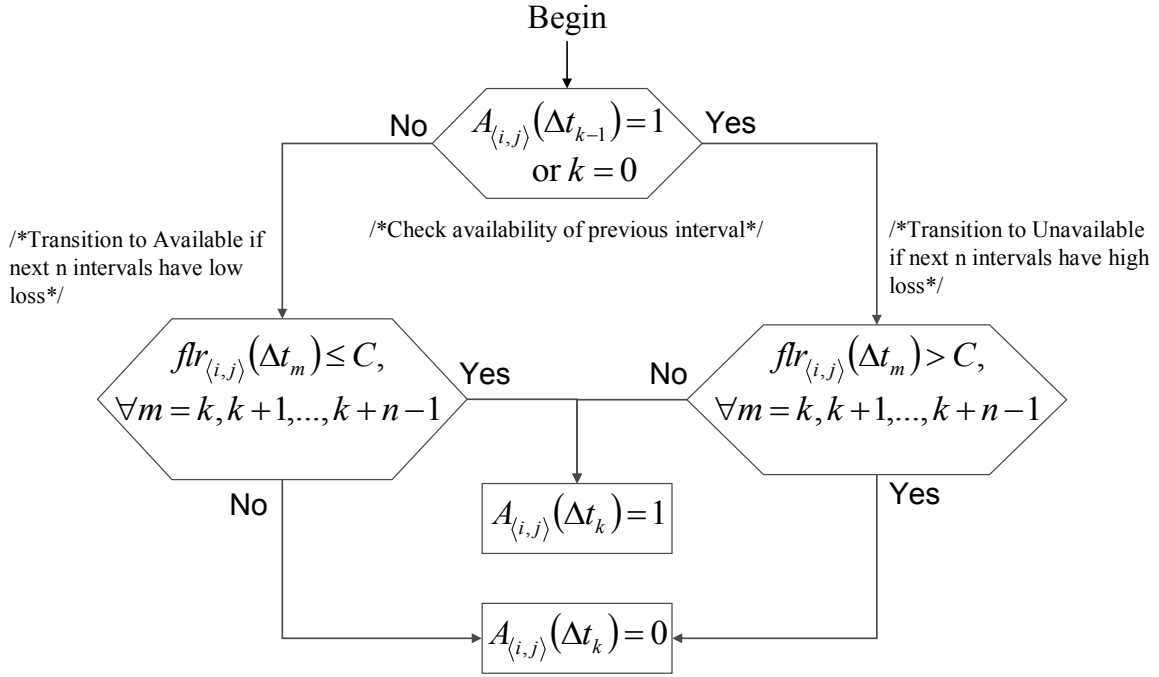


Figure 11 – Flowchart Definition of $A_{\langle i,j \rangle}(\Delta t_k)$

An alternative way of expressing $A_{\langle i,j \rangle}(\Delta t_k)$ for $k = 0$ is

$$A_{\langle i,j \rangle}(\Delta t_0) = \begin{cases} 0 & \text{if } flr_{\langle i,j \rangle}(\Delta t_m) > C, \text{ for all } m = 0, 1, \dots, n-1 \\ 1 & \text{otherwise} \end{cases}$$

and for $k = 1, 2, \dots$ is

$$A_{\langle i,j \rangle}(\Delta t_k) = \begin{cases} 0 & \text{if } A_{\langle i,j \rangle}(\Delta t_{k-1}) = 1 \text{ and } flr_{\langle i,j \rangle}(\Delta t_m) > C, \text{ for all } m = k, k+1, \dots, k+n-1 \\ 1 & \text{if } A_{\langle i,j \rangle}(\Delta t_{k-1}) = 0 \text{ and } flr_{\langle i,j \rangle}(\Delta t_m) \leq C, \text{ for all } m = k, k+1, \dots, k+n-1 \\ A_{\langle i,j \rangle}(\Delta t_{k-1}) & \text{otherwise} \end{cases}$$

In the event of a conflict between the above equations and Figure 11, the content of Figure 11 is controlling.

The availability for Δt_k is based on the frame loss ratio during the short interval and each of the following $n-1$ short intervals and the availability of the previous short time interval. In other words, a sliding window of width $n\Delta t$ is used to determine availability. This use of a sliding window is similar to that of ITU-T Y.1563 [6].

Figure 12 presents an example of the determination of the availability for the small time intervals with a sliding window of 10 small time intervals.




Figure 12 – Example of the Determination of $A_{(i,j)}(\Delta t_k)$

The Availability for a particular Class of Service Name from UNI i to UNI j for a time interval T is based on the percentage of small time intervals that are Available. However the small time intervals that occur during a Maintenance Interval are not included in the calculation of this percentage. A Maintenance Interval is a time interval agreed to by the Service Provider and Subscriber during which the service may not perform well or at all. Examples of a Maintenance Interval include:

- A time interval during which the Service Provider may disable the service for network maintenance such as equipment replacement,
- A time interval during which the Service Provider and Subscriber may perform joint fault isolation testing, and
- A time interval during which the Service Provider may change service features and making the changes may disrupt the service.

Figure 13 shows an example of the elimination of short time intervals for a Maintenance Interval.




Figure 13 – Example of the Impact of a Maintenance Interval

Let W_T be the set of all Δt_k 's contained in T that do not intersect a Maintenance Interval and let $|W_T|$ represent the number of elements in the set W_T .

Available Time for the ordered UNI pair $\langle i, j \rangle$, denoted $AT_{\langle i, j \rangle}$, is defined as

$$AT_{\langle i, j \rangle} = \left\{ \Delta t_k \mid \Delta t_k \in W_T, A_{\langle i, j \rangle}(\Delta t_k) = 1 \right\}$$

Unavailable Time for the ordered UNI pair $\langle i, j \rangle$, denoted $UT_{\langle i, j \rangle}$, is defined as

$$UT_{\langle i, j \rangle} = \left\{ \Delta t_k \mid \Delta t_k \in W_T, A_{\langle i, j \rangle}(\Delta t_k) = 0 \right\}$$

Then the Availability for a particular Class of Service Name from UNI i to UNI j for a time interval T , is defined by

$$A_T^{\langle i, j \rangle} = \begin{cases} \frac{100}{|W_T|} \sum_{\Delta t_k \in W_T} A_{\langle i, j \rangle}(\Delta t_k) & \text{if } |W_T| > 0 \\ 100 & \text{otherwise} \end{cases}$$

In words, Availability Performance is the percentage of Δt_k 's that do not intersect a Maintenance Interval that have $A_{\langle i, j \rangle}(\Delta t_k) = 1$.

Note that the definition of W_T means that the boundaries of T and the boundaries of a Maintenance Interval do not have to align with the boundary of a Δt_k . The definition of Availability Performance excludes any Δt_k that straddles the boundary between two T 's or that straddles the boundary of a Maintenance Interval.

[R44] For a given non-empty set of ordered pairs of UNIs, S , a time interval T , a time interval Δt , and a given Class of Service Name, the SLS **MUST** define the One-way Availability Performance Metric as follows:

- Let the UNIs associated by the EVC be numbered $1, 2, \dots, m$.
- Let S be a non-empty subset of the ordered pairs of UNIs, i.e., $S \subseteq \left\{ \langle i, j \rangle \mid i = 1, 2, \dots, m, j = 1, 2, \dots, m, i \neq j \right\}, S \neq \emptyset$.
- Then the Availability for a particular Class of Service Name for the set S is defined by $A_T^S = \min \left\{ A_T^{\langle i, j \rangle} \mid \langle i, j \rangle \in S \right\}$.

The parameters of a One-way Availability Performance Metric are given in Table 9. Note that there are constraints on Δt , C and n per [R28], [R29], and [R30].

Parameter	Description
T	The time interval
S	Non-empty subset of the ordered UNI pairs in the EVC
CoS_Name	The Class of Service Name
Δt	A time interval much smaller than T
C	Unavailability frame loss ratio threshold
n	Number of consecutive small time intervals for assessing availability
\hat{A}	Availability Performance Objective expressed as a percentage

Table 9 – One-way Availability Performance Parameters for an EVC

[R45] Given T , S , CoS_Name , Δt , C , n , and \hat{A} , the SLS **MUST** define the One-way Availability Performance Objective as being met if and only if $A_T^S \geq \hat{A}$.

As noted in [R44] and Table 9, the subset S has to be a non-empty subset of the ordered pairs of UNIs in the EVC. This is the only requirement on S for Point-to-Point and Multipoint-to-Multipoint EVCs. For example, S can contain just one of the ordered UNI pairs in a Point-to-Point EVC. Rooted-Multipoint EVCs have an additional requirement on S per [R46].

[R46] For a Rooted-Multipoint EVC, S **MUST** be a non-empty subset of the ordered pairs of UNIs in the EVC, such that all ordered pairs in S contain at least one UNI that has the Root Role.

8.8.5 One-way Resiliency Performance for an EVC

This section defines attributes for the Resiliency Performance of an ordered pair of UNIs, $\langle i, j \rangle$. Resiliency Performance is the number of High Loss Intervals and/or Consecutive High Loss Intervals in a time period T . The definitions depend on the Availability status of each Δt to determine whether Resiliency performance counts toward objectives. The Resiliency attributes are similar to the definitions of Severely Errored Seconds (SES) and Consecutive SES in section 9 and Annex B (respectively) of Y.1563 [6], when $\Delta t = 1$ second.

Figure 14 illustrates how the two Resiliency attributes defined here, counts of High Loss Intervals and counts of Consecutive High Loss Intervals, fit into the hierarchy of time and other attributes.

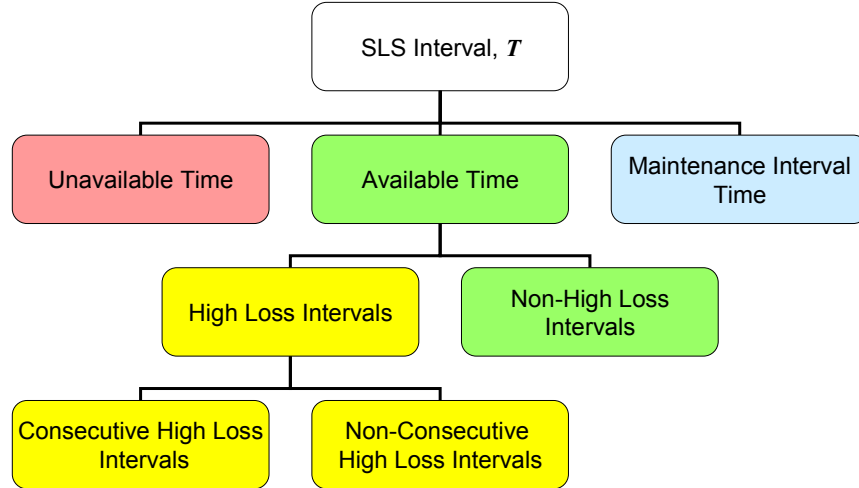


Figure 14 – Hierarchy of Time Showing the Resiliency Attributes

A High Loss Interval (HLI) is a small time interval contained in T (having the same duration as the interval, Δt , defined in Section 8.8.4) with a high frame loss ratio. When sufficient HLIs are adjacent, the interval is designated as a Consecutive High Loss Interval (CHLI). Section 8.8.4 defines terminology for Availability. This section re-uses that terminology and defines the following terms:

- $H_{\langle i,j \rangle}(\Delta t_k)$: The high loss state of Δt_k , equal to 1 if $flr_{\langle i,j \rangle}(\Delta t_k) > C$ and Δt_k is in Available Time, otherwise equal to 0.
- $L_T^{\langle i,j \rangle}$: Count of High Loss Intervals over T .
- \hat{L} : HLI Count Objective for S , T , and a given Class of Service Name.
- p : the minimum integer number of Δt 's in the (sliding) window (with $0 < p < n$) to qualify as a CHLI.
- $B_T^{\langle i,j \rangle}$: Count of CHLIs.
- \hat{B} : CHLI Count Objective for S , T , and a given Class of Service Name.

An equivalent definition for $H_{\langle i,j \rangle}(\Delta t_k)$ is $H_{\langle i,j \rangle}(\Delta t_k) = \begin{cases} 1 & \text{if } \Delta t_k \in AT_{\langle i,j \rangle}, flr_{\langle i,j \rangle}(\Delta t_k) > C \\ 0 & \text{otherwise} \end{cases}$.

For every Δt in T that does not intersect a Maintenance Interval, the flr and Availability state determine the value of $H_{\langle i,j \rangle}(\Delta t_k)$, either 1 or 0 as defined above.

[R47] For the SLS, the count of High Loss Intervals over T **MUST** be determined by $L_T^{\langle i,j \rangle} = \sum_{\Delta t \in T} H_{\langle i,j \rangle}(\Delta t)$.

Note that a counter for H may be implemented in post processing (e.g., in a Management System), outside the Network Element that is monitoring the frame loss rate of each Δt . This could be necessary to correlate with Δt 's in a Maintenance Interval.

When counting CHLI, the threshold p is used similarly to the variable n for the window size in the Availability attribute, and $p < n$.

- [R48]** For the SLS, the Consecutive High Loss Intervals over T **MUST** be determined according to the flow chart in Figure 15.

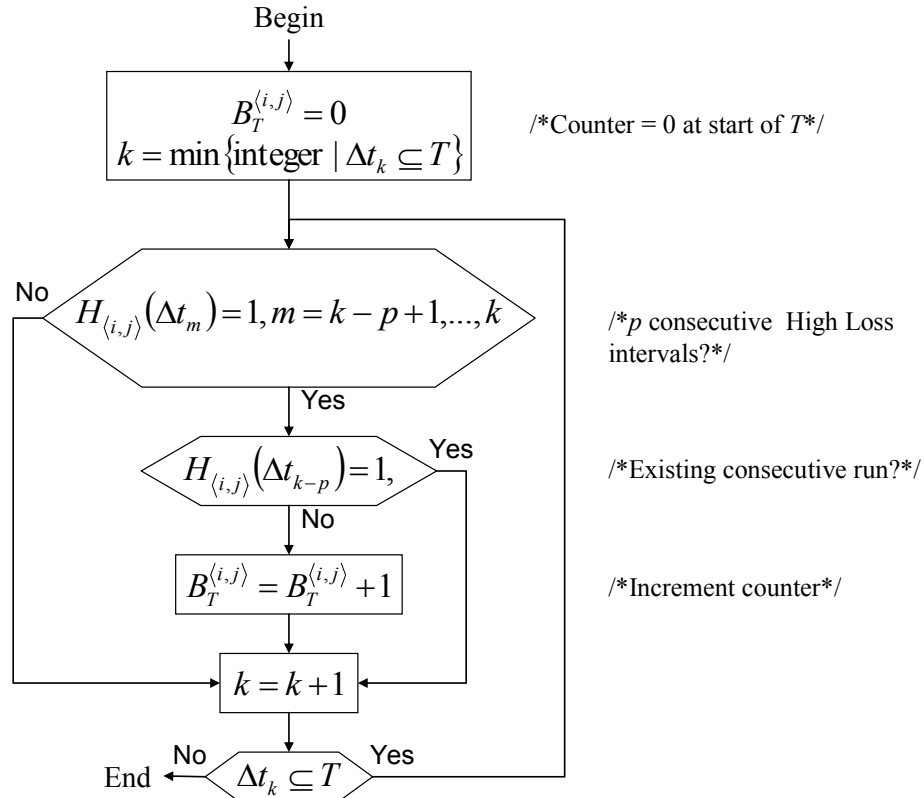


Figure 15 – Determining and Counting Consecutive High Loss Intervals over T

Figure 16 shows an example that depicts the HLI and CHLI counting processes when there is no Maintenance Interval.


$$L_T^S = \max \{L_T^{\langle i,j \rangle} \mid \langle i,j \rangle \in S\} \text{ and } B_T^S = \max \{B_T^{\langle i,j \rangle} \mid \langle i,j \rangle \in S\}.$$

Parameter	Description
T	The time interval
S	Non-empty subset of the ordered UNI pairs in the EVC
CoS_Name	The Class of Service Name
Δt	A time interval much smaller than T
C	Unavailability frame loss ratio threshold
p	Number of consecutive small time intervals for assessing CHLI, where $p < n$
\hat{L}	HLI Performance Objective expressed as an integer
\hat{B}	Consecutive HLI Performance Objective expressed as an integer

Table 10 – One-way Resiliency Performance Parameters for an EVC

- [R49] Given $T, S, CoS_Name, \Delta t, C$, and \hat{L} , the SLS **MUST** define the HLI Performance Objective as being met if and only if $L_T^S \leq \hat{L}$.
- [R50] Given $T, S, CoS_Name, \Delta t, C, p$, and \hat{B} , the SLS **MUST** define the CHLI Performance Objective as being met if and only if $B_T^S \leq \hat{B}$.

MEF 10.3 © The MEF Forum 2013. Any reproduction of this document, or any portion thereof, shall contain the following statement: "Reproduced with permission of the MEF Forum." No user of this document is authorized to modify any of the information contained herein. Page 43

EVCs. For example, S can contain just one of the ordered UNI pairs in a Point-to-Point EVC. Rooted-Multipoint EVCs have an additional requirement on S per [R51].

- [R51]** For a Rooted-Multipoint EVC, S **MUST** be a non-empty subset of the ordered pairs of UNIs in the EVC, such that all ordered pairs in S contain at least one UNI that has the Root Role.

8.8.6 One-way Group Availability Performance for an EVC

For a given Class of Service Name, Group Availability performance is defined for a collection of two or more non-empty sets of ordered pairs of UNIs associated by the EVC denoted by $G = \{S_1, \dots, S_m\}$.¹⁴ It depends on the following parameters:

- T , a time interval,
- Δt , a time interval smaller than T ,
- C , a frame loss ratio threshold,
- K , a positive integer, and
- n , a positive integer.

The Group Availability for the time interval Δt is defined as

$$A_G(\Delta t) = \begin{cases} 1 & \text{if } \sum_{S \in G} A^S(\Delta t, S) \geq K \\ 0 & \text{otherwise} \end{cases}$$

where $A^S(\Delta t, S) = \min \{A_{\langle i, j \rangle}(\Delta t) \mid \langle i, j \rangle \in S\}$ and $A_{\langle i, j \rangle}(\Delta t)$ is as defined in Section 8.8.4. In words, $A^S(\Delta t, S) = 1$ when all ordered UNI pairs in S have been experiencing low loss. Otherwise $A^S(\Delta t, S) = 0$.

- [R52]** The SLS **MUST** define the One-way Group Availability Performance Metric

$$\text{as } A_G = \begin{cases} \frac{100}{|W_T|} \sum_{\Delta t_k \in W_T} A_G(\Delta t_k) & \text{if } |W_T| > 0 \\ 100 & \text{otherwise} \end{cases} \quad \text{where } W_T \text{ is the set of } \Delta t \text{ time intervals contained in } T \text{ that don't intersect a Maintenance Interval as defined in Section 8.8.4.}$$

The One-way Group Availability Performance Metric can be viewed as the percentage of time that at least K S_k 's are Available. For example, consider a Multipoint-to-Multipoint EVC that

¹⁴ Section 11.1 contains a similar Performance Metric for the case where the sets of ordered UNI pairs come from more than one EVC.

associates UNI_A , UNI_B , and UNI_C , $K=1$, and let $G = \{S_1, S_2\}$ with $S_1 = \{\langle UNI_A, UNI_B \rangle, \langle UNI_B, UNI_A \rangle\}$ and $S_2 = \{\langle UNI_A, UNI_C \rangle, \langle UNI_C, UNI_A \rangle\}$. Then A_G represents the percentage of time that UNI_A has good connectivity with at least one of UNI_B or UNI_C .

The parameters of a One-way Group Availability Performance Metric are given in Table 11. Note that there are constraints on Δt , C and n per [R28], [R29], and [R30].

Parameter	Description
T	The time interval
S_1, \dots, S_m	Non-empty subsets of the ordered UNI pairs in the EVC
CoS_Name	The Class of Service Name
Δt	A time interval much smaller than T
C	Unavailability frame loss ratio threshold
K	Minimum number of S 's that are to be Available
n	Number of consecutive small time intervals for assessing availability
\hat{A}_G	Group Availability Performance Objective expressed as a percentage

Table 11 – One-way Group Availability Performance Parameters for an EVC

- [R53] Given T , S_1, \dots, S_m , CoS_Name , Δt , C , K , n , and \hat{A}_G , the SLS **MUST** define the One-way Group Availability Performance Objective as being met if and only if $A_G \geq \hat{A}_G$.
- [R54] For a Rooted-Multipoint EVC, each $S_k, k=1, \dots, m$ **MUST** be a non-empty subset of the ordered pairs of UNIs in the EVC, such that all ordered pairs in $S_k, k=1, \dots, m$ contain at least one UNI that has the Root Role.

8.9 EVC Maximum Service Frame Size Service Attribute

The value for the EVC Maximum Service Frame Size is a positive integer in bytes. Any ingress Tagged Service Frame whose length exceeds the Maximum Service Frame Size is likely to be discarded. Any ingress Untagged Service Frame whose length exceeds the Maximum Service Frame Size minus 4 is likely to be discarded

- [R55] The value of the EVC Maximum Service Frame Size **MUST** be at least 1522 bytes.
- [R56] The value of the EVC Maximum Service Frame Size **MUST** be less than or equal to the smallest of the UNI Maximum Service Frame Sizes (see Section 9.7) for all UNIs associated by the EVC.
- [D1] An ingress Tagged Service Frame that is mapped to the EVC and whose length exceeds the EVC Maximum Service Frame Size **SHOULD** be discarded.

- [D2] An ingress Untagged Service Frame that is mapped to the EVC and whose length exceeds the EVC Maximum Service Frame Size minus 4 **SHOULD** be discarded.

Note that an ingress Service Frame that is discarded per [D1] or [D2] does not consume tokens in a Bandwidth Profile (see Section 12).

9. UNI Service Attributes

This section describes attributes at each UNI.

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.

9.1 UNI ID Service Attribute

The UNI ID Service Attribute is independent of the EVCs at the UNI. It is assigned to the UNI by the Service Provider. It is an arbitrary string subject to the following requirements.

- [R57] The UNI ID **MUST** be unique among all UNIs for the CEN.
- [R58] The UNI ID **MUST** contain no more than 45 characters.⁵
- [R59] The UNI ID **MUST** be a non-null RFC 2579 [10] DisplayString but not contain the characters 0x00 through 0x1f.

As an example, the Service Provider might use "SCPOP1-Node3-Slot2-Port1" as a UNI ID and this could signify Port 1 in Slot 2 of Node 3 in Santa Clara POP1.

Note that [R57] does allow two Service Providers to use the same identifier for different UNIs (one UNI per Service Provider). Of course, using globally unique identifiers for UNIs meets [R57].

9.2 Physical Layer Service Attribute

This attribute is a list of physical layers, one for each physical link implementing the UNI per Section 9.4. Different physical links can use different physical layers.

- [R60] The Physical Layer for each physical link implementing the UNI **MUST** be one of the PHYs listed in IEEE Std 802.3 – 2012 [3] but excluding 1000BASE-PX-D and 1000BASE-PX-U.
- [R61] The Physical Layer **MUST** operate in full duplex mode.

Typically there are no constraints in mixing UNIs with different physical media in the same EVC.¹⁵

9.3 Synchronous Mode Service Attribute

The Synchronous Mode Service Attribute is a list with one item for each of the physical links implementing the UNI per Section 9.4. Each item in the list takes on one of two values: “Enabled” or “Disabled.” When the value of an item is “Enabled,” the bits transmitted from the CEN to the CE on the physical link corresponding to the item can be used by the CE as a bit clock reference.

- [R62] When the value of an item in the Synchronous Mode Service Attribute is “Enabled,” the Service Provider **MUST** specify the quality of the clock reference that is being provided on the physical link corresponding to the list item.

The handling of L2CP frames carrying messages for the Ethernet Synchronization Message Channel (ESMC) as specified in ITU-T G. 8264 [5] is beyond the scope of this document. MEF 22.1 [21] contains an example of the support of ESMC.

9.4 Number of Links Service Attribute

A UNI can be implemented with one or more physical links. This attribute specifies the number of such links.

- [R63] The value for the Number of Links Service Attribute **MUST** be at least 1.

When there are two links, a protection mechanism is required to be identified with UNI Resiliency Service Attribute, see Section 9.5. Protection mechanisms for a UNI with more than two links are beyond the scope of this document.

9.5 UNI Resiliency Service Attribute

The UNI Resiliency Service Attribute value can be “None,” “2-Link Aggregation,” or “Other.”

- [R64] If the value for the Number of Links Service Attribute is 1, then the value of the UNI Resiliency Service Attribute **MUST** be “None.”
- [R65] If the value for the Number of Links Service Attribute is 2, then the value of the UNI Resiliency Service Attribute **MUST** be “2-Link Aggregation.”
- [R66] When the value of the UNI Resiliency Service Attribute is “2-Link Aggregation,” the CEN **MUST** support Link Aggregation as in Clause 5.6.1 of IEEE Std 802.1AX – 2008 [1] with one Link Aggregation Group (LAG) across the links supporting the UNI and with one link in active mode and the other in standby mode.¹⁶

¹⁵ An exception might be wireless when the service requires stringent requirements on packet loss.

¹⁶ It is expected that a future MEF Technical Specification will place requirements on the Destination Address in the L2CP Service Frames used to meet this requirement.

- [R67] If the value for the Number of Links is 3 or greater, then the value of the UNI Resiliency Service Attribute **MUST** be “Other.”

Requirements for when the value of the UNI Resiliency Service Attribute is “Other” are beyond the scope of this document.

9.6 Service Frame Format Service Attribute

- [R68] The format of the Service Frame **MUST** be that of the MAC frame that is specified in Clause 3 of IEEE Std 802.3 – 2012 [3].
- [R69] An ingress Service Frame that is invalid as defined in Clause 3.4 of IEEE Std 802.3 – 2012 [3] **MUST** be discarded.
- [R70] An ingress Service Frame with length that is less than 64 bytes **MUST** be discarded as per Clause 4.2.4.2.2 of IEEE Std 802.3 – 2012 [3].

9.7 UNI Maximum Service Frame Size Service Attribute

The value for the UNI Maximum Service Frame Size is a positive integer in bytes.

- [R71] The value of the UNI Maximum Service Frame Size **MUST** be at least 1522 bytes.

9.8 Service Multiplexing Service Attribute

Service Multiplexing can be either Enabled or Disabled.

- [R72] A UNI with Service Multiplexing Enabled **MUST** be able to support multiple EVCs.¹⁷
- [O2] When a UNI has Service Multiplexing Enabled, Point-to-Point EVCs and Multipoint EVCs **MAY** be multiplexed in any combination at a UNI.

Figure 17 shows an example of Service Multiplexing. In this example, CE A is attached to the CEN 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. UNI A needs to have Service Multiplexing Enabled while the UNIs B, C, and D could have Service Multiplexing Enabled or Disabled.

¹⁷ Since the UNI is dedicated to a single Subscriber, only one Subscriber can access the EVCs at the UNI.




Figure 17 – Example of Service Multiplexing on UNI A

This attribute is independent of the EVCs at the UNI.

9.9 Customer Edge VLAN ID

At the given UNI, the Customer Edge VLAN ID (CE-VLAN ID) can be used to identify an EVC that associates this UNI.

There are 4095 CE-VLAN IDs numbered 1 through 4095. The CE-VLAN ID for a VLAN Tagged Service Frame is the value of the VLAN ID in the tag.

[R73] Untagged and Priority Tagged Service Frames **MUST** have the same CE-VLAN ID.

[R74] The CE-VLAN ID for Untagged and Priority Tagged Service Frames **MUST** be in the range 1, 2 , ..., 4094.¹⁸

When All to One Bundling is Enabled (Section 9.13), the value of the CE-VLAN ID for Untagged and Priority Tagged Service Frames does not affect the behavior of the EVC as seen by the Subscriber and thus can be considered to be not applicable. In such a situation, there is no need for the Subscriber and the Service Provider to agree on the value of this CE-VLAN ID.

[R75] When CE-VLAN ID Preservation is Disabled for an EVC to which the CE-VLAN ID for Untagged and Priority Tagged Service Frames is mapped, egress Service Frames mapped to this EVC at the given UNI **MUST** be Untagged Service Frames.

Note that more than one CE-VLAN ID can map to the same EVC as described in Section 9.10.

Note that the Customer VLAN Tag values 0 and 4095 in IEEE Std 802.1Q – 2011 [2] are reserved for special purposes and thus the number of VLANs in a Subscriber network is typically

¹⁸ 4095 is not included for compatibility with IEEE Std 802.1Q – 2011 [2] where 4095 is reserved for special purposes.

less than 4095. Nevertheless, Tagged 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, fewer than 4095 EVCs can be supported at a UNI. See Section 9.11.

The 4095 CE-VLAN IDs always exist at each UNI and are independent of the EVCs at the UNI. The CE-VLAN ID for Untagged and Priority Tagged Service Frames is also independent of the EVCs at the UNI.

9.10 CE-VLAN ID/EVC Map Service Attribute

9.10.1 Basic Concept

The CE-VLAN ID/EVC Map associates CE-VLAN IDs to EVCs. Note that this mapping applies to both ingress and egress Service Frames. For an ingress Service Frame, the CE-VLAN ID for the Service Frame and the CE-VLAN ID/EVC Map enable the CEN to know how to deliver the Service Frame. For an egress Service Frame, the CE-VLAN ID for the Service Frame and the CE-VLAN ID/EVC Map allow the CE to know which EVC the Service Frame came from.

[R76] At each UNI there **MUST** be a mapping of each CE-VLAN ID to at most one EVC.

Note that [R76] means that a given CE-VLAN ID need not be mapped to an EVC.

[R77] When Bundling is Disabled and All to One Bundling is Disabled (as defined in Sections 9.12 and 9.13), exactly one CE-VLAN ID **MUST** be mapped to each EVC.

Figure 18 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 8.2) or the UNI EVC ID (Section 10.1).

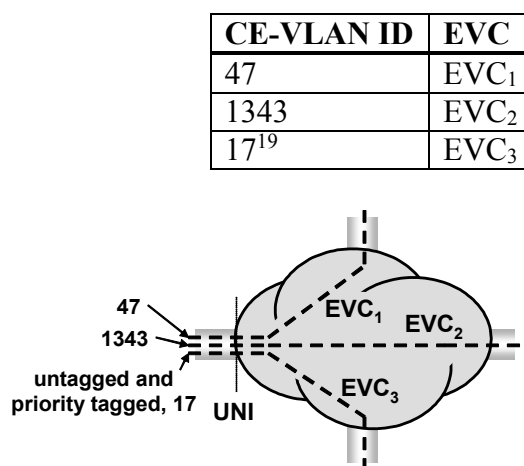


Figure 18 – Example of a CE-VLAN ID/EVC Map

¹⁹ In this example, the CE-VLAN ID for Untagged and Priority Tagged Service Frames is configured to 17.

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

[R78] 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 NOT** result in a corresponding egress Service Frame at any UNI in the CEN.

Note that [R78] does not preclude processing ingress SOAM and L2CP Service Frames with a CE-VLAN ID that does not map to an EVC, but such frames do not result in an egress Service Frame at any UNI.

In scenarios where All to One Bundling is Disabled, the Subscriber and the Service Provider need to agree upon the CE-VLAN ID/EVC Map at the UNI. 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.

9.10.2 CE-VLAN ID Significance

When CE-VLAN ID Preservation is Disabled ([R24]), the CE-VLAN ID/EVC mapping for a given EVC at a UNI can be different from the mapping at another UNI in the same EVC. However, in some cases it is mandated to be the same such as when CE-VLAN ID Preservation is Enabled ([R24]) or when Bundling is Enabled ([R81]).

Figure 19 shows valid CE-VLAN ID/EVC Maps for three EVCs between two UNIs. Note that when CE-VLAN ID Preservation (Section 8.6.1) is Enabled for an EVC, the mappings for the EVC are identical as is the case for EVC₁ in Figure 19. Otherwise the CE-VLAN ID cannot be preserved.

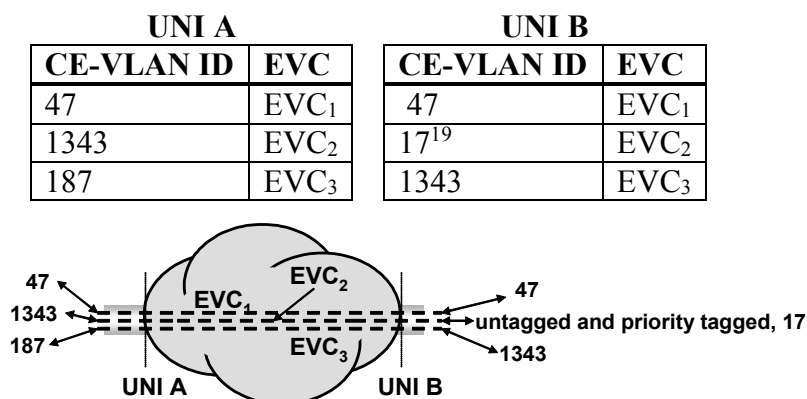


Figure 19 – Example of CE-VLAN ID/EVC Maps at Two UNIs

9.10.3 Describing the Contents of the CE-VLAN ID/EVC Map

The CE-VLAN ID/EVC Map described here is an abstraction. This description does not constrain how the contents can be described in a protocol, database, service order form, etc. For example, shorthand descriptions such as the example of Section 9.12 or the protocol optimizations of the Ethernet Local Management Interface [18] are allowed.

9.11 Maximum Number of EVCs Service Attribute

This attribute defines the maximum number of EVCs that the UNI can support.

[R79] Maximum Number of EVCs **MUST** have a value of at least one.

9.12 Bundling Service Attribute

Bundling can be either Enabled or Disabled.

[R80] When a UNI has Bundling Enabled, it **MUST** be able to support more than one CE-VLAN ID mapping to a particular EVC at the UNI.

The Bundling Service Attribute is independent of the EVCs at the UNI.

Note that when more than one CE-VLAN ID is mapped to an EVC at a UNI, [R25] mandates that the EVC have CE-VLAN ID Preservation Enabled.

[R81] An EVC with more than one CE-VLAN ID mapping to it **MUST** have the same list of CE-VLAN IDs mapping to the EVC at each UNI in the EVC.

Figure 20 shows an example of Bundling. In this example, UNI A and UNI B have Bundling Enabled as seen from the mapping for EVC₁. (EVC₁ has CE-VLAN ID Preservation Enabled.). Note that Bundling Enabled is compatible with Service Multiplexing. In Figure 20, UNI A and UNI B are examples of Service Multiplexing Enabled and Bundling Enabled on the same UNI.

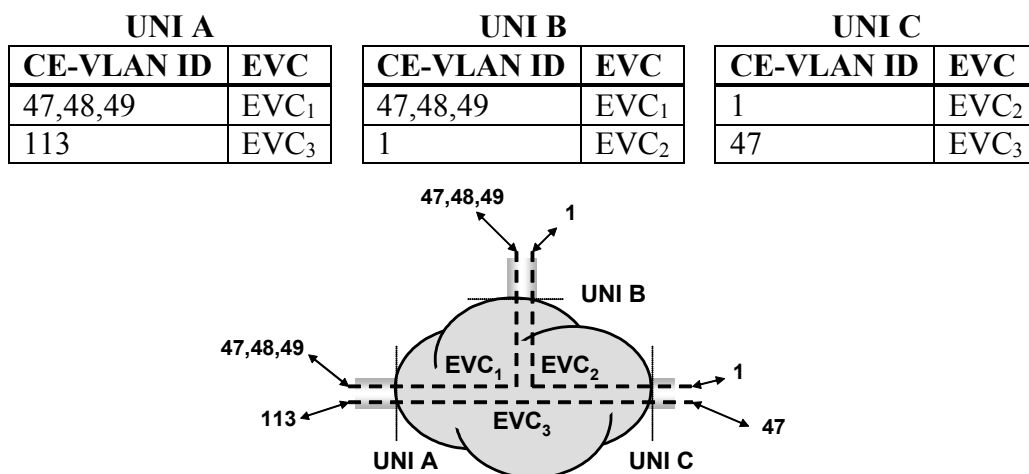


Figure 20 – Example of Bundling

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 21.

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 21 – Example of a Simple Description of Bundling

9.13 All to One Bundling Service Attribute

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

All to One Bundling can be either Enabled or Disabled.

[R82] When a UNI has All to One Bundling Enabled, all CE-VLAN IDs **MUST** map to a single EVC at the UNI.

[R83] When a UNI has All to One Bundling Enabled, all other UNIs in the EVC associating this UNI **MUST** have All to One Bundling Enabled.

Note that [R25] means that the EVC that associates a UNI with All to One Bundling Enabled is required to have CE-VLAN ID Preservation Enabled.

The All to One Bundling service attribute is independent of the EVCs at the UNI.

[R82] means that a UNI with All to One Bundling Enabled cannot have Service Multiplexing Enabled.

Each column of Table 12 shows a valid combination of the Bundling, All to One Bundling and Service Multiplexing attributes.

	Valid Combination 1	Valid Combination 2	Valid Combination 3	Valid Combination 4	Valid Combination 5
Service Multiplexing	Disabled	Enabled	Enabled	Disabled	Disabled
Bundling	Disabled	Disabled	Enabled	Enabled	Disabled
All to One Bundling	Disabled	Disabled	Disabled	Disabled	Enabled

Table 12 – Valid Combinations of Service Multiplexing, Bundling, and All to One Bundling

9.14 Ingress Bandwidth Profile per UNI Service Attribute

[R84] For an Ingress Bandwidth Profile per UNI, a single Bandwidth Profile Flow based on criterion 1 of [R134] **MUST** be contained in a single Envelope. (See Section 12.1 for the definitions of Bandwidth Profile Flow and Envelope.)

As described in Section 12.2, only values for the parameters $\langle CIR^1, CBS^1, EIR^1, EBS^1, CF^1, CM^1 \rangle$ need to be specified for an Ingress Bandwidth Profile per UNI.

Figure 22 illustrates an example of the application of ingress policing with an Ingress Bandwidth Profile per UNI. In the example of Figure 22, ingress Service Frames for the three EVCs would all be subject to a single Ingress Bandwidth Profile. The Ingress Bandwidth Profile per UNI manages bandwidth non-discriminately for all EVCs at the UNI, i.e., some EVCs may get more bandwidth while others may get less.

The Ingress Bandwidth Profile per UNI service attribute is independent of the EVCs at the UNI.




Figure 22 – Ingress Bandwidth Profile per UNI

9.15 Egress Bandwidth Profile per UNI Service Attribute

[R85] For an Egress Bandwidth Profile per UNI, a single Bandwidth Profile Flow based on criterion 4 of [R134] **MUST** be contained in a single Envelope. (See Section 12.1 for the definitions of Bandwidth Profile Flow and Envelope.)

As described in Section 12.2, only values for the parameters $\langle CIR^1, CBS^1, EIR^1, EBS^1, CF^1, CM^1 \rangle$ need to be specified for an Egress Bandwidth Profile per UNI.

The Egress Bandwidth Profile per UNI manages bandwidth non-discriminately for all EVCs at the egress UNI, i.e. some EVCs may get more bandwidth while others may get less. Figure 23 portrays this model of Egress Bandwidth Profile.

The Egress Bandwidth Profile per UNI service attribute is independent of the EVCs at the UNI.




Figure 23 – Egress Bandwidth Profile per UNI

9.16 Link OAM Service Attribute

Link OAM can be either Enabled or Disabled.

- [R86] When the value of the Link OAM Service Attribute is Enabled, the CEN **MUST** support Active DTE mode capabilities as specified in clause 57.2.9 of IEEE Std 802.3 – 2012 [3] on each link in the UNI.

9.17 UNI MEG Service Attribute

UNI MEG can be either Enabled or Disabled.

- [R87] When the value of the UNI MEG Service Attribute is Enabled, the CEN **MUST** meet the mandatory requirements in Section 7.9 of MEF 30.1 [25] that apply to the UNI MEG.

When the value of the UNI MEG Service Attribute is Enabled, several parameter values need to be determined as described in Section 7.9 of MEF 30.1 [25]. An Implementation Agreement can specify such values.

9.18 E-LMI Service Attribute

E-LMI can be either Enabled or Disabled.

- [R88] When the value of the E-LMI Service Attribute is Enabled, the CEN **MUST** meet the mandatory requirements in MEF 16 [18] that apply to the UNI-N.

9.19 UNI Layer 2 Control Protocol Processing Service Attribute

Requirements for the processing of L2CP Service Frames are contained in MEF 6.1.1.[13] Additional requirements for Link Aggregation are in Section 9.5.

10. EVC per UNI Service Attributes

This Section contains attributes that apply to an EVC at a given UNI. For a given EVC, the value of an EVC per UNI Service Attribute can be different at each of the UNIs associated by the EVC.

10.1 UNI EVC ID Service Attribute

The UNI EVC ID is a string formed by the concatenation of the UNI ID (Section 9.1) and the EVC ID (Section 8.2) that is used to identify an EVC at the UNI. It is intended for management and control purposes.

10.2 Class of Service Identifier Service Attribute

Each ingress Service Frame mapped to the given EVC has a single Class of Service Identifier. The Class of Service Identifier can be determined from inspection of the content of the ingress Service Frame.

The Class of Service Identifier in the ingress Service Frame maps to a Class of Service Name. As described in Section 6.2 of MEF 23.1 [22], the Class of Service Name is used to identify the performance objectives and associated parameters that apply to the ingress Service Frame as described in Section 8.8 of this document. Note that the Class of Service Name could be one of the Class of Service Labels standardized in MEF 23.1. [22]

For example, for the given EVC at the UNI, there could be three Class of Service Names offered called Silver, Gold, and Platinum.

- [O3] A Class of Service Identifier **MAY** map to a Class of Service Name that calls for discard.
- [R89] When a Class of Service Identifier maps to a Class of Service Name that calls for discard, any ingress Service Frame with this Class of Service Identifier **MUST** be discarded.

The mapping of Class of Service Identifier to Class of Service Name is not necessarily the same for each EVC at a given UNI, or for each UNI associated by a given EVC.

The requirements regarding the Class of Service Identifier for Data Service Frames and SOAM Service Frames are different than they are for Layer 2 Control Protocol Service Frames as detailed below.

10.2.1 Class of Service Identifier for Data Service Frames

- [R90]** A given EVC at a given UNI **MUST** use the method described in exactly one of the Sections 10.2.1.1, 10.2.1.2, or 10.2.1.3 to determine the Class of Service Identifier from the content of Data Service Frames.

[R90] means that a given EVC at a given UNI either uses the method in 10.2.1.1, or uses the method in 10.2.1.2, or uses the method in 10.2.1.3. Different EVCs at the same UNI can use different methods. Also, different methods can be used by the same EVC at different UNIs.

10.2.1.1 Class of Service Identifier Based on EVC

- [R91]** When Class of Service Identifier is based on EVC, all ingress Data Service Frames mapped to the EVC **MUST** map to the same Class of Service Name at the given UNI.

As an example, consider EVC 1 and EVC 2 at a UNI. Ingress Data Service Frames mapped to EVC 1 have a first Class of Service Identifier that indicates Class of Service Label H. Ingress Data Service Frames mapped to EVC 2 have a second Class of Service Identifier that indicates Class of Service Label M.

10.2.1.2 Class of Service Identifier Based on Priority Code Point Field

Throughout this section, recall that CE-VLAN CoS refers to the Priority Code Point Field in the Customer VLAN Tag in a Tagged Service Frame.

- [R92]** When the Class of Service Identifier is based on the Priority Code Point Field, the CE-VLAN CoS **MUST** determine the Class of Service Name.
- [R93]** When the Class of Service Identifier is based on the Priority Code Point Field, each possible CE-VLAN CoS value **MUST** map to exactly one Class of Service Name.

[R93] means that the sets of CE-VLAN CoS values that each map to a different Class of Service Name are disjoint sets and the union of all such sets is the set of all possible CE-VLAN CoS values. [R93] also means that each ingress Service Frame has a single Class of Service Name that applies to it.

Table 13 shows an example of Class of Service Identifiers based on Priority Code Point Field for EVC 1 at the UNI. In this example, an ingress Tagged Data Service Frame with CE-VLAN CoS = 5 is subject to the performance objectives for Class of Service Name Gold. Table 14 shows an example of Class of Service Identifiers based on Priority Code Point Field for EVC 2 at the UNI. These examples illustrate that the Class of Service Identifiers can be different for different EVCs at the UNI.

CE-VLAN CoS Values	Class of Service Name
4,5,6,7	Gold
0,3	Silver
1,2	Discard

Table 13 – PCP Class of Service Identifier Example for EVC1 at the UNI

CE-VLAN CoS Values	Class of Service Name
7	Platinum
0,1,2,3,4,5,6	Gold

Table 14 – PCP Class of Service Identifier Example for EVC2 at the UNI

- [R94]** When the Class of Service Identifier is based on the Priority Code Point Field, the Class of Service Name for all ingress Untagged Data Service Frames mapped to the given EVC **MUST** be agreed to by the Subscriber and Service Provider.
- [D3]** When the Class of Service Identifier is based on the Priority Code Point Field, the Class of Service Name for ingress Untagged Data Service Frames mapped to the given EVC (see [R94]) **SHOULD** be the same Class of Service Name as that for a Tagged Data Service Frame with CE-VLAN CoS = 0.

[R93] and [R94] along with the fact that the mapping of Class of Service Identifier to Class of Service Names is not necessarily the same for each EVC at given UNI, mean that to determine the Class of Service Name based on the Priority Code Point Field for an ingress Data Service Frame at a given UNI, both the EVC to which the frame is mapped and the CE-VLAN CoS value need to be considered. Stated differently, the Class of Service Name for an ingress Data Service Frame at a given UNI is determined from the EVC plus the value of the CE-VLAN CoS.²⁰

10.2.1.3 Class of Service Identifier Based on Internet Protocol

When the Class of Service Identifier is based on Internet Protocol, the Class of Service Identifier is determined from the DSCP for a Data Service Frame carrying an IPv4 or an IPv6 packet.

- [R95]** When the Class of Service Identifier is based on Internet Protocol, each possible DSCP value **MUST** map to exactly one Class of Service Name for ingress Data Service Frames carrying an IPv4 packet.
- [R96]** When the Class of Service Identifier is based on Internet Protocol, each possible DSCP value **MUST** map to exactly one Class of Service Name for ingress Data Service Frames carrying an IPv6 packet.
- [R97]** When the Class of Service Identifier is based on Internet Protocol, all ingress Data Service Frames mapped to the given EVC that do not contain either an IPv4 or an IPv6 packet **MUST** have the Class of Service Name that is agreed upon by the Subscriber and the Service Provider.

²⁰Per [R94], Untagged Service Frames can be considered to contain a special CE-VLAN CoS value for this purpose.

Table 15 shows an example of using Class of Service Identifier based on Internet Protocol. In this example an ingress Data Service Frame carrying an IPv4 packet with DSCP = 37 would be subject to the Class of Service Name Platinum. Similarly, an ingress Data Service Frame carrying an IPv6 packet with DSCP = 37 would be subject to the performance objectives for Class of Service Name Platinum. In this example Diamond can only apply to ingress Data Service Frames carrying an IPv4 packet and Ruby can only apply to ingress Data Service Frames carrying an IPv6 packet.

IPv4 DSCP Values	IPv6 DSCP Values	Class of Service Name
11,37,45	11,37,45	Platinum
8,10,12		Diamond
	38,213	Ruby
All other values	All other values	Discard

Table 15 – Internet Protocol Class of Service Identifier Example for an EVC at the UNI

Table 15 is an example where IPv4 and IPv6 are simultaneously supported. Table 16 shows an example where only IPv4 is supported. In this example, the Class of Service Name for ingress Data Service Frames not carrying an IP packet is "Good Enough" per [R97]. This fact and the last row means that "Good Enough" applies to any ingress Data Service Frame not carrying an IPv4 packet. Consequently, if IPv6 is not recognized but instead is treated as non-IP, the mapping shown in Table 16 is achieved.

IPv4 DSCP Values	IPv6 DSCP Values	Class of Service Name
11,37,45		Superior
8,10,12		Near Superior
All other values		Discard
	All values	Good Enough

Table 16 – IPv4 Class of Service Identifier Example for an EVC at the UNI

Table 17 shows an example where only IPv6 is supported. In this example, the Class of Service Name for ingress Data Service Frames not carrying an IP packet is "Good Enough" per [R97]. This fact and the last row means that "Good Enough" applies to any ingress Data Service Frame not carrying an IPv6 packet. Consequently, if IPv4 is not recognized but instead is treated as non-IP, the mapping shown in Table 17 is achieved.

IPv4 DSCP Values	IPv6 DSCP Values	Class of Service Name
	11,37,45	Superior
	8,10,12	Near Superior
	All other values	Discard
All values		Good Enough

Table 17 – IPv6 Class of Service Identifier Example for an EVC at the UNI

The three previous examples show how IPv4 and IPv6, only IPv4, and only IPv6 can be supported by appropriate structuring of the mapping from DSCP value to Class of Service Name.

[R95], [R96], and [R97] along with the fact that the mapping of Class of Service Identifier to Class of Service Names is not necessarily the same for each EVC at given UNI, mean that to determine the Class of Service Name for an ingress Data Service Frame, both the EVC to which the frame is mapped and the DSCP of the IP packet need to be considered. Stated differently, the Class of Service Name for an ingress Data Service Frame is determined from the EVC plus the value of the DSCP.²¹

10.2.2 Class of Service Identifier for Layer 2 Control Protocol Service Frames

Irrespective of which approach is used to determine the Class of Service Identifier for ingress Data Service Frames for the EVC as detailed in Sections 10.2.1.1, 10.2.1.2, and 10.2.1.3, a list of Layer 2 Control Protocols can be specified with each entry in the list being a Class of Service Identifier. Thus each entry in the list maps to a single Class of Service Name. The entries in this list are referred to as the L2CP Class of Service Identifiers.

Note that since the mapping of Class of Service Identifier to Class of Service Names is not necessarily the same for each EVC at given UNI, the L2CP Class of Service Identifiers can be different for different EVCs that associate a given UNI.

- [R98] If the Layer 2 Control Protocol carried by an ingress Layer 2 Control Protocol Service Frame²² is contained in the L2CP Class of Service Identifiers, then the Class of Service Name for this frame **MUST** be the Class of Service Name mapped to by the Layer 2 Control Protocol.
- [R99] If the Layer 2 Control Protocol carried by an ingress Layer 2 Control Protocol Service Frame²² is not contained in the L2CP Class of Service Identifiers, then the Class of Service Name for this frame **MUST** be determined as if it is a Data Service Frame per Section 10.2.1.

[R98] and [R99] along with the fact that the mapping of Class of Service Identifier to Class of Service Names is not necessarily the same for each EVC at given UNI, mean that to determine the Class of Service Name for an ingress Layer 2 Control Protocol Service Frame, both the EVC to which the frame is mapped and the type of Layer 2 Control Protocol need to be considered.

10.2.3 Class of Service Identifier for SOAM Service Frames

- [R100] For a given EVC at a given UNI, the basis for the Class of Service Identifier for ingress SOAM Service Frames **MUST** be the same as that for ingress Data Service Frames.
- [R101] For a given EVC at a given UNI, the mapping from Class of Service Identifier to Class of Service Name for ingress SOAM Service Frames **MUST** be the same as that for ingress Data Service Frames.

²¹ Per [R96], Service Frames not containing an IP packet can be considered to contain a special DSCP value for this purpose.

²² The method by which the type of Layer 2 Control Protocol is determined is beyond the scope of this document.

10.3 Color Identifier Service Attribute

The Color Identifier is the means by which the Color for a Service Frame is indicated by the content in the Service Frame header.

[R102] The Color for a Service Frame **MUST** be either Green or Yellow.

Note that Color Identifier applies to both ingress Service Frames and egress Service Frames.

When an Ingress Bandwidth Profile is color-aware, the value of the Color for a Service Frame influences the behavior of the Bandwidth Profile Algorithm (see Section 12.1). When no Ingress Bandwidth Profile is applied to an ingress Service Frame, the Color for the Service Frame helps determine if the Service Level Specification applies to the frame (see Section 8.8).

The Color for an egress Service Frame is important when there is a color-aware Egress Bandwidth Profile (see Section 12.4).

A Service Frame need not have a Color Identifier.

[R103] When a Service Frame has no Color Identifier it **MUST** have the Color Green.

The Color Identifier can be based on a specific field in the Service Frame or it can be implicit. An example of an implicit Color Identifier is to specify that all Service Frames mapped to a given EVC have the Color Yellow. The Color Identifier for Service Frames is one of:

1. EVC (all Service Frames mapped to the EVC have the same Color),
2. CE-VLAN CoS value,
3. CE-VLAN Tag DEI, or
4. DSCP value if the Service Frame contains an IPv4 packet or an IPv6 packet.

[O4] At a given UNI, the Color Identifiers for Service Frames mapped to one EVC **MAY** be different than the Color Identifiers for Service Frames mapped to a different EVC at the given UNI.

[O5] The Color Identifiers for Service Frames mapped to an EVC at a given UNI **MAY** be different from the Color Identifiers for Service Frames mapped to the EVC at other UNIs in the EVC.

The following sections describe the various ways that Color Identifier can be specified. Note that the requirements in Sections 10.3.1, 10.3.2, 10.3.3, and 10.3.4 apply to both ingress and egress Service Frames.

10.3.1 Color Identifier Based on EVC

- [R104] When the Color Identifier is based on EVC, all Service Frames mapped to the EVC **MUST** have the same Color.

10.3.2 Color Identifier Based on CE-VLAN Tag DEI

- [R105] When the Color Identifier is based on CE-VLAN Tag DEI, the Color for a Tagged Service Frame **MUST** be Green if DEI = 0 and Yellow if DEI = 1.

10.3.3 Color Identifier Based on CE-VLAN CoS

When the Color Identifier is based on the CE-VLAN CoS, each CE-VLAN CoS value maps to a single Color.

- [R106] When the Color Identifier is based on the CE-VLAN CoS, the CE-VLAN CoS **MUST** determine the Color of the Service Frame.
- [R107] When the Color Identifier is based on the CE-VLAN CoS, each possible CE-VLAN CoS value **MUST** map to exactly one Color.

[R107] means that the sets of CE-VLAN CoS values that map to each Color are disjoint sets and the union of these sets is the set of all possible CE-VLAN CoS values.

Table 18 shows an example of Color Identifiers based on CE-VLAN CoS for EVC 1 at the UNI. In this example, an ingress Tagged Service Frame with CE-VLAN CoS = 2 has the Color Yellow. Table 19 shows an example of Color Identifiers based on CE-VLAN CoS for EVC 2 at the UNI. These examples illustrate that the Color Identifiers can be different for different EVCs at the UNI.

CE-VLAN CoS Values	Color
0,2	Yellow
1,3,4,5,6,7	Green

Table 18 – CE-VLAN CoS Color Identifier Example for EVC 1 at the UNI

CE-VLAN CoS Values	Color
0,2,4,6	Yellow
1,3,5,7	Green

Table 19 – CE-VLAN CoS Color Identifier Example for EVC 2 at the UNI

Note that Table 4 in MEF 23.1 [22] combines the mappings for Color Identifiers and Class of Service Identifiers into a single table. Table 20 contains the mappings from Table 4 of MEF 23.1 [22] and it is easy to see how Table 18 is constructed from Table 20.

CE-VLAN CoS Value	Color	CoS Label
0	Yellow	L
1	Green	L
2	Yellow	M
3	Green	M
4	Green	Discard
5	Green	H
6	Green	Discard
7	Green	Discard

Table 20 – Combined Mappings Consistent with Table 4 in MEF 23.1 [22]

10.3.4 Color Identifier Based on Internet Protocol

When the Color Identifier is based on Internet Protocol, the Color Identifier is determined from the DSCP for a Service Frame carrying an IPv4 packet or an IPv6 packet.

- [R108]** When the Color Identifier is based on Internet Protocol, each possible DSCP value **MUST** map to exactly one Color for Data Service Frames carrying an IPv4 packet.
- [R109]** When the Color Identifier is based on Internet Protocol, each possible DSCP value **MUST** map to exactly one Color for Data Service Frames carrying an IPv6 packet.
- [R110]** When the Color Identifier is based on Internet Protocol, all Service Frames mapped to the given EVC that do not contain either an IPv4 or an IPv6 packet **MUST** have a Color that is agreed upon by the Subscriber and the Service Provider.

Per [R103], when the Color Identifier is based on Internet Protocol, all Service Frames mapped to the given EVC that do not contain either an IPv4 or an IPv6 packet have the Color Green.

Table 21 shows an example of using Color Identifier based on Internet Protocol. In this example a Service Frame carrying an IPv4 packet with DSCP = 30 would have the Color Yellow. Similarly, a Service Frame carrying an IPv6 packet with DSCP = 14 would have the Color Yellow. Note that the mappings in Table 21 are consistent with the mapping in Table 4 of MEF 23.1. [22]

IPv4 DSCP Values	IPv6 DSCP Values	Color
0,12,14,28,30	0,12,14,28,30	Yellow
All other values	All other values	Green

Table 21 – Internet Protocol Color Identifier Example

Table 22 shows an example where only IPv4 is supported. In this example, the Color for a Service Frame not carrying an IPv4 or IPv6 packet is Green per [R110]. In other words, a Service Frame carrying an IPv6 packet is treated as a Service Frame not carrying an IPv4 packet. Conse-

quently, if only Service Frames carrying IPv4 are recognized as carrying IP (IPv6 is not recognized), the mapping in Table 22 is achieved.

IPv4 DSCP Values	IPv6 DSCP Values	Color
0,12,14,28,30		Yellow
All other values	All values	Green

Table 22 – IPv4 Only Color Identifier Example

An analogous approach can be used for the case where only IPv6 is supported.

10.3.5 Color Identifier Constraints

The basis for the Color Identifier for an EVC at a given UNI is constrained by the following requirements.

- [R111] When the Class of Service Identifier is based on the Priority Code Point Field (see Section 10.2.1.2) for a given EVC at a given UNI, the Color Identifier **MUST** be based on either the DEI or the Priority Code Point Field.
- [R112] When the Class of Service Identifier is based on Internet Protocol (see Section 10.2.1.3) for a given EVC at a given UNI, the Color Identifier **MUST** be based on Internet Protocol.

10.4 Egress Equivalence Class Identifier Service Attribute

Each egress Service Frame mapped to the given EVC has a single Egress Equivalence Class Identifier. The Egress Equivalence Class Identifier can be determined from inspection of the content of the egress Service Frame. Each Egress Equivalence Class Identifier maps to a single Egress Equivalence Class. The Egress Equivalence Class is used to specify Egress Bandwidth Profiles as described in Section 12.

- [O6] At a given UNI, the Egress Equivalence Class Identifiers for egress Service Frames mapped to one EVC **MAY** be different than the Egress Equivalence Class Identifiers for egress Service Frames mapped to a different EVC.

10.4.1 Egress Equivalence Class Identifier for Data Service Frames

- [R113] A given EVC at a given UNI **MUST** use the method described in exactly one of the Sections 10.4.1.1 or 10.4.1.2 to determine the Egress Equivalence Class from the content of Data Service Frames.

[R113] means that a given EVC at a given UNI either uses the method in 10.4.1.1, or uses the method in 10.4.1.2. Different EVCs at the same UNI can use different methods. Also, different methods can be used by the same EVC at different UNIs.

10.4.1.1 Egress Equivalence Class Identifier Based on CE-VLAN CoS

- [R114] When the Egress Equivalence Class Identifier is based on the CE-VLAN CoS, the CE-VLAN CoS **MUST** determine the Egress Equivalence Class.
- [R115] When the Egress Equivalence Class Identifier is based on the CE-VLAN CoS, each CE-VLAN CoS value **MUST** map to exactly one Egress Equivalence Class.

[R115] means that the sets of CE-VLAN CoS values that each map to a different Egress Equivalence Class are disjoint sets and the union of all such sets is the set of all possible CE-VLAN CoS values. [R115] also means that each egress Service Frame has a single Egress Equivalence Class that applies to it.

- [R116] When the Egress Equivalence Class Identifier is based on the CE-VLAN CoS, the Egress Equivalence Class for all egress Untagged Data Service Frames mapped to the given EVC **MUST** be agreed to by the Subscriber and Service Provider.
- [D4] When the Egress Equivalence Class of Service Identifier is based on the CE-VLAN CoS, the Egress Equivalence Class for egress Untagged Data Service Frames mapped to the given EVC (see [R116]) **SHOULD** be the same Egress Equivalence Class as that for an egress Tagged Data Service Frame with CE-VLAN CoS= 0.

10.4.1.2 Egress Equivalence Class Identifier Based on Internet Protocol

When the Egress Equivalence Class Identifier is based on Internet Protocol, the Egress Equivalence Class Identifier is determined from the DSCP for a Data Service Frame carrying an IPv4 or an IPv6 packet.

- [R117] When the Egress Equivalence Class Identifier is based on Internet Protocol, each possible DSCP value **MUST** map to exactly one Egress Equivalence Class for egress Data Service Frames carrying an IPv4 packet.
- [R118] When the Egress Equivalence Class Identifier is based on Internet Protocol, each possible DSCP value **MUST** map to exactly one Egress Equivalence Class for egress Data Service Frames carrying an IPv6 packet.
- [R119] When the Egress Equivalence Class Identifier is based on Internet Protocol, all egress Data Service Frames mapped to the given EVC that do not contain either an IPv4 or an IPv6 packet **MUST** have an Egress Equivalence Class that is agreed upon by the Subscriber and the Service Provider.

Table 23 shows an example of using Egress Equivalence Class Identifier based on Internet Protocol. In this example an egress Data Service Frame carrying an IPv4 packet with DSCP = 37 would be in Egress Equivalence Class Sophomore. Similarly, an egress Data Service Frame carrying an IPv6 packet with DSCP = 37 would be in Egress Equivalence Class Sophomore. In this

example Senior can only apply to egress Data Service Frames carrying an IPv4 packet and Junior can only apply to egress Data Service Frames carrying an IPv6 packet.

IPv4 DSCP Values	IPv6 DSCP Values	Egress Equivalence Class
11,37,45	11,37,45	Sophomore
8,10,12		Senior
	38,213	Junior
All other values	All other values	Freshman

Table 23 – Internet Protocol Egress Equivalence Class Identifier Example for an EVC at the UNI

Table 24 shows an example where only IPv4 is supported. In this example, the Egress Equivalence Class for egress Data Service Frames not carrying an IP packet is "Baggage" per [R119]. This fact and the last row means that "Baggage" applies to any egress Data Service Frame not carrying an IPv4 packet.

IPv4 DSCP Values	IPv6 DSCP Values	Egress Equivalence Class
11,37,45		First
8,10,12		Business
All other values		Economy
	All values	Baggage

Table 24 – IPv4 Only Egress Equivalence Class Identifier Example for an EVC at the UNI

An analogous approach can be used for the case where only IPv6 is supported.

10.4.2 Egress Equivalence Class Identifier for Layer 2 Control Protocol Service Frames

Irrespective of which approach is used to determine the Egress Equivalence Class Identifier for egress Data Service Frames for the EVC as detailed in Sections 10.4.1.1 and 10.4.1.2, a list of Layer 2 Control Protocols can be specified with each entry in the list being an Egress Equivalence Class Identifier. The entries in this list are referred to as the L2CP Egress Equivalence Class Identifier.

Note that [O6] means that the L2CP Egress Equivalence Class Identifiers can be different for different EVCs that associate a given UNI.

- [R120]** If the Layer 2 Control Protocol carried by an egress Layer 2 Control Protocol Service Frame²³ is contained in the L2CP Egress Equivalence Class Identifiers, then the Egress Equivalence Class for this frame **MUST** be the matching entry in L2CP Egress Equivalence Class Identifiers.

²³ The method by which the type of Layer 2 Control Protocol is determined is beyond the scope of this document.

- [R121] If the Layer 2 Control Protocol carried by an egress Layer 2 Control Protocol Service Frame²³ is not contained in the L2CP Egress Equivalence Class Identifiers, then the Egress Equivalence Class for this frame **MUST** be determined as if it was an egress Data Service Frame per Section 10.4.1.

10.4.3 Egress Equivalence Class Identifier for SOAM Service Frames

- [R122] For a given EVC at a given UNI, the basis for the Egress Equivalence Identifier for egress SOAM Service Frames **MUST** be the same as that for egress Data Service Frames.
- [R123] For a given EVC at a given UNI, the mapping from Egress Equivalence Class Identifier to Egress Equivalence Class for egress SOAM Service Frames **MUST** be the same as that for egress Data Service Frames.

10.5 Ingress Bandwidth Profile per EVC Service Attribute

- [R124] For an Ingress Bandwidth Profile per EVC, a single Bandwidth Profile Flow based on criterion 2 of [R134] for this EVC **MUST** be contained in an Envelope that contains only Bandwidth Profile Flows based on criterion 2 of [R134]. (See Section 12.1 for the definitions of Bandwidth Profile Flow and Envelope.)

Thus, if a UNI has 3 Ethernet Virtual Connections, there could one Envelope containing three Bandwidth Profile Flows, one for each EVC. In such a case, each set of parameters for a Bandwidth Profile Flow described in Section 12.1 would be associated with a different EVC at the UNI.

When an Envelope contains a single Bandwidth Profile Flow based on criterion 2 of [R134] and both CIR_{\max}^1 and EIR_{\max}^1 are sufficiently large, then the behavior is exactly that of the per EVC Ingress Bandwidth Profile described in MEF 10.2 [15]. (See Section 12.2.)

Figure 24 illustrates an example of the application of Ingress Bandwidth Profiles per EVC. In this example, EVC₁ corresponds to Bandwidth Profile Flow 1 with $CIR^1 = CIR_{\max}^1 = 15$ Mbps, EVC₂ corresponds to Bandwidth Profile Flow 2 with $CIR^2 = CIR_{\max}^2 = 10$ Mbps, and EVC₃ corresponds to Bandwidth Profile Flow 3 with $CIR^3 = CIR_{\max}^3 = 20$ Mbps. The superscripts on the parameters correspond to the rank within the Envelope. The values for the remaining parameters are not addressed in order to keep the example simple.




Figure 24 – Ingress Bandwidth Profile per EVC

10.6 Ingress Bandwidth Profile per Class of Service Identifier Service Attribute

[R125] For an Ingress Bandwidth Profile per Class of Service Identifier, a single Bandwidth Profile Flow based on criterion 3 of [R134] for the Class of Service Name mapped to by this Class of Service Identifier **MUST** be contained in an Envelope that contains only Bandwidth Profile Flows based on criterion 3 of [R134]. (See Section 12.1 for the definitions of Bandwidth Profile Flow and Envelope.)

Class of Service Identifier is described in Section 10.2.

When an Envelope contains a single Bandwidth Profile Flow based on criterion 3 of [R134] and both CIR_{max}^1 and EIR_{max}^1 are sufficiently large, then the behavior is exactly that of the per Class of Service Identifier Ingress Bandwidth Profile described in MEF 10.2 [15]. (See Section 12.2.)

In contrast, an example of multiple Bandwidth Profile Flows in an Envelope is shown in Figure 25. In this example, there are three Class of Service Names within EVC_1 , each corresponding to a different Bandwidth Profile Flow. The Envelope in this example contains these three Bandwidth Profile Flows. In such a case, the three sets of parameters, one for each of the three Bandwidth Profile Flows described in Section 12.1 would be associated with EVC_1 at the UNI.




Figure 25 – Ingress Bandwidth Profile per Class of Service Identifier

[D5] When there are multiple CoS Names supported on an EVC, Ingress Bandwidth Profile per Class of Service Identifier **SHOULD** be used.

10.7 Egress Bandwidth Profile per EVC Service Attribute

[R126] For an Egress Bandwidth Profile per EVC, a single Bandwidth Profile Flow based on criterion 5 of [R134] for this EVC **MUST** be contained in an Envelope that contains only Bandwidth Profile Flows based on criterion 5 of [R134]. (See Section 12.1 for the definitions of Bandwidth Profile Flow and Envelope.)

When an Envelope contains a single Bandwidth Profile Flow based on criterion 5 of [R134] and both CIR_{max}^1 and EIR_{max}^1 are sufficiently large, then the behavior is exactly that of the per EVC Egress Bandwidth Profile described in MEF 10.2 [15]. (See Section 12.2.)

10.8 Egress Bandwidth Profile per Egress Equivalence Class Identifier Service Attribute

[R127] For an Egress Bandwidth Profile per Egress Equivalence Class, a single Bandwidth Profile Flow based on criterion 6 of [R134] for the Egress Equivalence Class mapped to by this Egress Equivalence Class Identifier **MUST** be contained in an Envelope that contains only Bandwidth Profile Flows based on criterion 6 of [R134]. (See Section 12.1 for the definitions of Bandwidth Profile Flow and Envelope.)

When an Envelope contains a single Bandwidth Profile Flow of egress Service Frames with a single Egress Equivalence Class, both CIR_{max}^1 and EIR_{max}^1 are sufficiently large, and the set of egress Service Frames with this Egress Equivalence Class is identical to the set of egress Service Frames with a single Class of Service Name, then the behavior is exactly that of the per Class of Service Identifier Egress Bandwidth Profile described in MEF 10.2 [15]. For example, suppose that:

- The Egress Equivalence Class Identifier is based on Priority Code Point Field,

- PCP values 0 and 1 in the CE-VLAN CoS map to Egress Equivalence Class A,
- An Egress Bandwidth Profile Flow is specified for Egress Equivalence Class A,
- This Bandwidth Profile Flow is in an Envelope by itself,
- The Class of Service Identifier is based on Priority Code Point Field, and
- PCP values 0 and 1 in the CE-VLAN CoS map to Class of Service Name Platinum.

For the same traffic stream at the egress UNI, this Egress Bandwidth Profile per Egress Equivalence Class Identifier has identical behavior to that of the Egress Bandwidth Profile per Class of Service Identifier in MEF 10.2 [15].

As an example, consider an Egress UNI with two EVCs with each EVC having 3 Egress Equivalence Classes. With this model, there can be up to six Egress Bandwidth Profiles.

10.9 Source MAC Address Limit Service Attribute

Source MAC Address Limit can be either Enabled or Disabled.

This attribute limits the number of source MAC Addresses that can be used in ingress Service Frames of all types over a time interval. Two independent parameters control the behavior of this attribute:

- N : A positive integer and
- τ : A time interval.

This attribute operates by maintaining a list of maximum length N of source MAC addresses which are aged-out of the list if not seen in a time interval τ . If an ingress Service Frame arrives with a new source MAC address when the list is full, the Service Frame is discarded.

In algorithmic terms, this can be stated as maintaining a list L where

$$L = \{ \langle A_i, t_i \rangle \mid i = 1, 2, \dots, q \leq N, A_i = \text{unicast MACAddress}, t_i = \text{a time} \}.$$

The t_i in each $\langle A_i, t_i \rangle$ is the most recent time that an ingress Service Frame arrived at the UNI that was mapped to the EVC and contained the source MAC address A_i .

[R128] If Source MAC Address Limit is Enabled, then for a sequence of ingress Service Frames mapped to the EVC with source MAC Addresses A_j and arrival times at the UNI t_j for $j = 0, 1, 2, \dots$, the frames **MUST** be discarded per the logic of Figure 26 where $L = \emptyset$ at time t_0 .

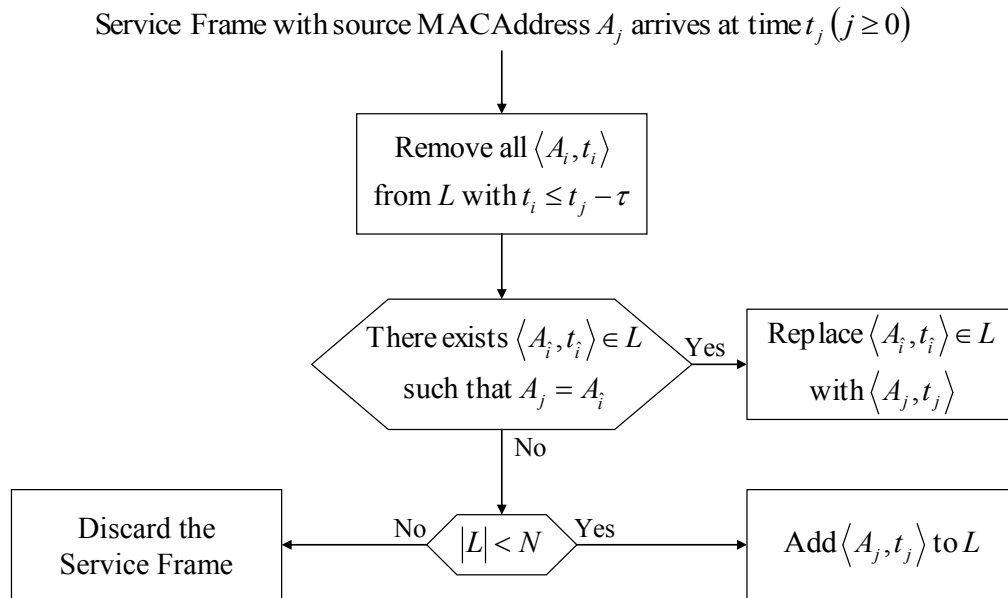


Figure 26 – Frame Discard for Source MAC Address Limit

Note that [R128] does not mandate a specific implementation in the CEN. Any implementation that yields the same behavior as that of Figure 26 meets the requirement. For example, an implementation that removes an $\langle A_i, t_i \rangle$ from the list at time $t = t_i + \tau$ yields the same Service Frame discard behavior as that of Figure 26.

10.10 Test MEG Service Attribute

Test MEG can be either Enabled or Disabled.

- [R129]** When the value of the Test MEG Service Attribute is Enabled, the CEN **MUST** meet the mandatory requirements in Section 7.5 of MEF 30.1 [25] that apply to the Test MEG.

When the value of the Test MEG Service Attribute is Enabled, several parameter values need to be determined as described in Section 7.5 of MEF 30.1 [25]. An Implementation Agreement can specify such values.

10.11 Subscriber MEG MIP Service Attribute

The Subscriber MEG MIP can be either Enabled or Disabled.

- [R130]** When the value of the Subscriber MEG MIP Service Attribute is Enabled, the CEN **MUST** instantiate a Subscriber Level MIP as described in MEF 30.1 [25].

When the value of the Subscriber MEG MIP Service Attribute is Enabled, several parameter values need to be determined as described in MEF 30.1 [25]. An Implementation Agreement can specify such values.

11. Multiple EVC Related Performance Service Attributes

Section 8.8 defines several Performance Metrics that apply to a single EVC. In some situations it is useful to specify performance involving more than one EVC. An example is where the Subscriber is using one EVC as a backup for another EVC.

The SLS can include the Multiple EVC Related Performance Metric defined in the following sub-section. This Performance Metric applies to Qualified Service Frames as detailed in [R31].

11.1 One-way Multiple EVC Group Availability

For a given Class of Service Name, Multiple EVC Group Availability performance is defined for a collection $G = \{S_1, \dots, S_m\}$ of two or more non-empty sets S_k of ordered UNI pairs such that:

- Each set S_k consists of ordered UNI pairs that are associated by one EVC, and
- Two or more EVCs are used to specify G .

The Multiple EVC Group Availability performance metric depends on the following parameters:

- T , a time interval,
- Δt , a time interval smaller than T ,
- C , a frame loss ratio threshold,
- K , a positive integer, and
- n , a positive integer.

The Multiple EVC Group Availability for the time interval Δt is defined as

$$A_G^M(\Delta t) = \begin{cases} 1 & \text{if } \sum_{S \in G} A^g(\Delta t, S) \geq K \\ 0 & \text{otherwise} \end{cases}$$

where $A^g(\Delta t, S) = \min \{A_{\langle i, j \rangle}(\Delta t) \mid \langle i, j \rangle \in S\}$ and $A_{\langle i, j \rangle}(\Delta t)$ is as defined in Section 8.8.4. In words, $A^g(\Delta t, S) = 1$ when all ordered UNI pairs in S have been experiencing low loss. Otherwise $A^g(\Delta t, S) = 0$.

[R131] The SLS **MUST** define the One-way Multiple EVC Group Availability Per-

$$\text{formance Metric as } A_G^M = \begin{cases} \frac{100}{|W_T|} \sum_{\Delta t_k \in W_T} A_G^M(\Delta t_k) & \text{if } |W_T| > 0 \\ 100 & \text{otherwise} \end{cases} \quad \text{where } W_T \text{ is the}$$

set of Δt time intervals contained in T that don't intersect a Maintenance Interval as defined in Section 8.8.4.

The One-way Multiple EVC Group Availability Performance Metric can be viewed as the percentage of time that at least K S_k 's are Available. For example, consider two Point-to-Point EVCs where one EVC associates UNI_A and UNI_B and the other EVC associates UNI_A and UNI_C . Let $G = \{S_1, S_2\}$, $K = 1$, $S_1 = \{\langle UNI_A, UNI_B \rangle, \langle UNI_B, UNI_A \rangle\}$ and $S_2 = \{\langle UNI_A, UNI_C \rangle, \langle UNI_C, UNI_A \rangle\}$. Then A_G^M represents the percentage of time that UNI_A has good connectivity with at least one of UNI_B or UNI_C .

The parameters of a One-way Multiple EVC Group Availability Performance Metric are given in Table 25. Note that there are constraints on Δt , C and n per [R28], [R29], and [R30].

Parameter	Description
T	The time interval
S_1, \dots, S_m	Non-empty subsets of the ordered UNI pairs based on two or more EVCs
CoS_Name	The Class of Service Name
Δt	A time interval much smaller than T
C	Unavailability frame loss ratio threshold
K	Minimum number of S 's that are to be Available
n	Number of consecutive small time intervals for assessing availability
\hat{A}_G^M	Group Availability Performance Objective expressed as a percentage

Table 25 – One-way Multiple EVC Group Availability Performance Parameters

[R132] Given T , S_1, \dots, S_m , CoS_Name , Δt , C , K , n , and \hat{A}_G^M , the SLS **MUST** define the One-way Multiple EVC Group Availability Performance Objective as being met if and only if $A_G^M \geq \hat{A}_G^M$.

[R133] Each S_k that is based on a Rooted-Multipoint EVC **MUST** be a non-empty subset of the ordered pairs of UNIs in the EVC, such that all ordered pairs contain at least one UNI that has the Root Role.

12. Bandwidth Profile at the UNI

A Bandwidth Profile is a characterization of the lengths and arrival times for Service Frames at a UNI.

There are two types of Bandwidth Profile. An Ingress Bandwidth Profile is used to regulate the amount of ingress traffic at a particular UNI, while an Egress Bandwidth Profile is used to regulate the amount of egress traffic at a particular UNI. The Ingress Bandwidth Profile is described in Section 12.3. The Egress Bandwidth Profile is described in Section 12.4.

Typically a Bandwidth Profile defines Service Frame traffic that is less than the full bandwidth of the UNI.

Bandwidth Profiles are associated with the UNI. This allows different Bandwidth Profiles at each UNI in an EVC. For example, on a Multipoint-to-Multipoint EVC, a different Bandwidth Profile could apply at each UNI in the EVC.

The Bandwidth Profile defines the set of traffic parameters applicable to a sequence of Service Frames. Associated with the Bandwidth Profile is an algorithm to determine Service Frame compliance with the specified parameters. In the case of an Ingress Bandwidth Profile, rate enforcement is accomplished via the disposition of Service Frames. Rate enforcement for an Egress Bandwidth Profile is a network design issue and beyond the scope of this document.

All Bandwidth Profiles in this Technical Specification are based on the parameters and algorithm described in this Section.

A Bandwidth Profile is specified using the concepts of Bandwidth Profile Flow and Envelope. A Bandwidth Profile Flow is defined as a set of Service Frames arriving at a UNI that meet a specific criterion.

[R134] A Bandwidth Profile Flow **MUST** be specified using one of the following six criteria:

1. All ingress Service Frames at the UNI that are not discarded per requirements [R69], [R70], [R89], [R128], [D1], [D2], or any of the conditions specified per [R16],
2. All ingress Service Frames at the UNI that are mapped to a given EVC and that are not discarded per requirements [R69], [R70], [R89], [R128], [D1], [D2], or any of the conditions specified per [R16],
3. All ingress Service Frames at the UNI that have a given Class of Service Name, that map to a given EVC, and that are not discarded per requirements [R69], [R70], [R89], [R128], [D1], [D2], or any of the conditions specified per [R16],
4. All egress Service Frames at the UNI,
5. All egress Service Frames at the UNI that are mapped to a given EVC, or
6. All egress Service Frames at the UNI that have a given Egress Equivalence Class and that map to a given EVC.

Note that in some cases the criteria in the above requirements can be equivalent, e.g., criterion 1 and criterion 2 are equivalent at a UNI with All to One Bundling Enabled since there is only one EVC at the UNI and all Service Frames map to it (see Section 9.13).

Note that criteria 1, 2, and 3 of [R134] mean that ingress Service Frames discarded per requirements [R69], [R70], [R89], [R128], [D1], [D2], or any of the conditions specified per [R16] cannot be in a Bandwidth Profile Flow and thus they will not consume tokens in the Bandwidth Profile Algorithm described in Section 12.1.

An Envelope is a set of n Bandwidth Profile Flows in which each Bandwidth Profile Flow is assigned a unique rank between 1 (lowest) and n (highest). n can be any positive integer.

[R135] All Bandwidth Profile Flows in an Envelope **MUST** satisfy the same criterion in [R134].

One implication of [R135] is that the Bandwidth Profile Flows in an Envelope are either all Ingress Bandwidth Profile Flows or all Egress Bandwidth Profile Flows. When the Envelope contains Bandwidth Profile Flows meeting either criterion 1 or criterion 4 of [R134], there is just one Bandwidth Profile Flow in the Envelope.

Multiple Envelopes containing Bandwidth Profile Flows meeting criterion 2 of [R134] can co-exist at a UNI, each based on disjoint set of EVCs. Similarly multiple Envelopes containing Bandwidth Profile Flows meeting criterion 5 of [R134] can co-exist at a UNI, each based on disjoint set of EVCs.

Multiple Envelopes containing Bandwidth Profile Flows meeting criterion 3 of [R134] can co-exist at a UNI. Similarly multiple Envelopes containing Bandwidth Profile Flows meeting criterion 6 of [R134] can co-exist at a UNI.

[R136] Each Bandwidth Profile Flow at a UNI **MUST** belong to exactly one Envelope.

[R137] A Service Frame **MUST** be mapped to at most one Bandwidth Profile Flow.

Note that a given Service Frame does not need to be mapped to a Bandwidth Profile Flow. When this is the case, it is said that the Service Frame is not subject to a Bandwidth Profile. When a Service Frame is mapped to a Bandwidth Profile Flow, [R136] and [R137] mean that this Service Frame is subject to exactly one Envelope and thus the Bandwidth Profile algorithm of [R155] is applied to this Service Frame exactly once. As a consequence, each Service Frame mapped to a Bandwidth Profile Flow has exactly one color declaration.

12.1 Bandwidth Profile Parameters and Algorithm

The Bandwidth Profile algorithm can be configured such that bandwidth that is not used by one Bandwidth Profile Flow can be reallocated to other Bandwidth Profile Flows in the same Envelope. The significance of the ranking of the Bandwidth Profile Flows is that it controls how the algorithm reallocates unused bandwidth from a given Bandwidth Profile Flow to another Bandwidth Profile Flow.

Each Envelope at a UNI has two parameters, the **Envelope ID** and the **Envelope Coupling Flag**, CF^0 .

- [R138] The Envelope ID **MUST** contain no more than 45 characters.⁵
- [R139] The Envelope ID **MUST** be unique among all Envelope IDs at a UNI.
- [R140] The Envelope ID **MUST** be a non-null RFC 2579 [10] DisplayString but not contain the characters 0x00 through 0x1f.
- [R141] CF^0 **MUST** have only one of two possible values, 0 or 1.
- [R142] When one Bandwidth Profile Flow is mapped to an Envelope, CF^0 **MUST** equal 0.

Note that the concatenation of the UNI ID and the Envelope ID is unique among all Envelopes for the CEN.

The parameters for each Bandwidth Profile Flow of rank $i = 1, 2, \dots, n$:

- **Committed Information Rate per Bandwidth Profile Flow** (CIR^i) expressed as bits per second. This parameter limits the average rate in bits per second of Service Frames that will be declared Green by the Bandwidth Profile.
- [R143] CIR^i **MUST** be ≥ 0 .
- **Maximum Committed Information Rate per Bandwidth Profile Flow** (CIR_{\max}^i) expressed as bits per second. This parameter limits the rate of tokens added to the committed token bucket.
- [R144] CIR_{\max}^i **MUST** be ≥ 0
- **Committed Burst Size per Bandwidth Profile Flow** (CBS^i) expressed as bytes. This parameter limits the maximum number of bytes available for a burst of Service Frames sent at the UNI line rate that will be declared Green by the Bandwidth Profile.
- [R145] When $CIR^i > 0$, CBS^i **MUST** be greater than or equal to the largest EVC Maximum Service Frame Size value of the EVCs that Service Frames in the Bandwidth Profile Flow are mapped to. See Section 8.9.
- **Excess Information Rate per Bandwidth Profile Flow** (EIR^i) expressed as bits per second. This parameter limits the average rate in bits per second of Service Frames that will be declared Yellow by the Bandwidth Profile.
- [R146] EIR^i **MUST** be ≥ 0 .

- **Maximum Excess Information Rate per Bandwidth Profile Flow** (EIR_{\max}^i) expressed as bits per second. This parameter limits the rate of tokens added to the excess token bucket.

[R147] EIR_{\max}^i **MUST** be ≥ 0 .

- **Excess Burst Size per Bandwidth Profile Flow** (EBS^i) expressed as bytes. This parameter limits the maximum number of bytes available for a burst of Service Frames sent at the UNI line rate that will be declared Yellow by the Bandwidth Profile

[R148] When $EIR^i > 0$, EBS^i **MUST** be greater than or equal to the largest EVC Maximum Service Frame Size value of the EVCs that Service Frames in the Bandwidth Profile Flow are mapped to. See Section 8.9.

- **Coupling Flag per Bandwidth Profile Flow** (CF^i). This parameter determines whether or not overflow tokens not used for Service Frames declared Green can be used as Yellow tokens.

[R149] CF^i **MUST** have only one of two possible values, 0 or 1.

[R150] If $CF^0 = 1$ for an Envelope, then CF^i **MUST** equal 0 for all Bandwidth Profile Flows mapped to the Envelope.

- **Color Mode per Bandwidth Profile Flow** (CM^i). This parameter indicates whether the color-aware or color-blind property is employed by the Bandwidth Profile.

[R151] CM^i **MUST** have only one of two possible values, “color-blind” and “color-aware.”

- **Envelope and Rank** (ER^i) a pair <Envelope ID, rank>. The Envelope ID indicates the Envelope that the Bandwidth Profile Flow belongs to and the rank i is a positive integer that specifies the rank of the Bandwidth Profile Flow within the Envelope.

[R152] The value of the rank in ER^i **MUST** be in the range 1,2,..., n where n is the number of Bandwidth Profile Flows with the same Envelope ID value.

[R153] The value of the rank in ER^i **MUST NOT** equal the rank of any of the other Bandwidth Profile Flows with the same Envelope ID value.

Each Service Frame is classified to determine which, if any, Envelope is applicable to the Service Frame.²⁴ Operation of the Bandwidth Profile algorithm is governed by the parameters, $\langle CIR^i, CIR_{\max}^i, CBS^i, EIR^i, EIR_{\max}^i, EBS^i, CF^i, CM^i, ER^i \rangle$ for $i=1,2,\dots,n$ and CF^0 . The algorithm declares each Service Frame 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 Red.²⁵

The Bandwidth Profile algorithm is said to be in color-aware mode for Bandwidth Profile Flow i when each Service Frame in Bandwidth Profile Flow i that already has a level of compliance (i.e., a color) associated with it via the Color Identifier (see Section 10.3) has that color taken into account in determining the level of compliance by the Bandwidth Profile algorithm. The Bandwidth Profile algorithm is said to be in color-blind mode for Bandwidth Profile Flow i when the color (if any) already associated with each Service Frame in Bandwidth Profile Flow i is ignored by the Bandwidth Profile Algorithm.

[R154] A UNI **MUST** be able to support color-blind mode for Bandwidth Profiles.

[O7] A UNI **MAY** support color-aware mode for Bandwidth Profiles.

The color mode of operation for Bandwidth Profile Flow i is determined using the parameter CM^i .

A token bucket model is used to describe the Bandwidth Profile behavior. Each Bandwidth Profile Flow i has a committed token bucket that can contain up to CBS^i tokens, and an excess token bucket that can contain up to EBS^i tokens. Tokens for each bucket are sourced at a rate of CIR^i and EIR^i respectively. Tokens that overflow each bucket for a Bandwidth Profile Flow may be made available to other buckets depending on the setting of the Coupling Flags (CF^i) and what other Bandwidth Profile Flows are mapped to the Envelope. Tokens flowing into a bucket are limited to CIR_{\max}^i and EIR_{\max}^i respectively. Each Service Frame that maps to Bandwidth Profile Flow i will be declared Green, Yellow, or Red by the Bandwidth Profile algorithm, depending on the arrival time of the Service Frame, the Color of the Service Frame (if any), and the number of tokens in the committed and excess token buckets. When a Service Frame for Bandwidth Profile Flow i is declared Green a number of tokens equal to the length of that Service Frame is removed from the committed bucket for that Bandwidth Profile Flow. When a Service Frame for Bandwidth Profile Flow i is declared Yellow a number of tokens equal to the length of that Service Frame is removed from the excess bucket for that Bandwidth Profile Flow.

To specify the Bandwidth Profile behavior, we define the following functions:

$B_C^i(t)$ and $B_E^i(t)$ represent the number of tokens in the committed and excess token buckets, respectively, for Bandwidth Profile Flow i at time t , for $i=1,2,\dots,n$.

²⁴ 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 and thus can be subject to a Bandwidth Profile. Similarly SOAM Service Frames can be subject to a Bandwidth Profile.

²⁵ 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 CEN but such procedures are beyond the scope of this Technical Specification.

$T_C^i(t_1, t_2)$ and $T_E^i(t_1, t_2)$ represent the maximum number of tokens that might be added to the committed and excess token buckets, respectively, for Bandwidth Profile Flow i over the time interval t_1 to t_2 , for $i = 1, 2, \dots, n$.

$A_C^i(t_1, t_2)$ and $A_E^i(t_1, t_2)$ represent the number of tokens actually added to the committed and excess token buckets, respectively, for Bandwidth Profile Flow i over the time interval t_1 to t_2 for $i = 1, 2, \dots, n$.

$O_C^i(t_1, t_2)$ and $O_E^i(t_1, t_2)$ represent the number of tokens that overflow the committed and excess token buckets, respectively, for Bandwidth Profile Flow i over the time interval t_1 to t_2 for $i = 1, 2, \dots, n$.

$S_C^i(t_1, t_2)$ and $S_E^i(t_1, t_2)$ represent the number of tokens available to be shared to the committed and excess token buckets of Bandwidth Profile Flow $i-1$, respectively, from Bandwidth Profile Flow i over the time interval t_1 to t_2 for $i = 2, \dots, n, n+1$.

The maximum number of tokens that might be added to the committed token bucket, $T_C^i(t_1, t_2)$, includes tokens sourced at the Committed Information Rate (CIR^i) over the time interval and any tokens shared from the next higher rank Bandwidth Profile Flow. The maximum number of tokens that might be added to the excess token bucket, $T_E^i(t_1, t_2)$, includes tokens sourced at the Excess Information Rate (EIR^i) over the time interval, any tokens shared from the higher rank Bandwidth Profile Flow, and any overflow tokens from the committed token bucket that are allowed by the Coupling Flag (CF^i). Note that for the committed token bucket, Bandwidth Profile Flow n has no tokens shared from a higher rank since there are no Bandwidth Profile Flows with a rank higher than n . Therefore $S_C^{n+1}(t_1, t_2) = 0$. For the excess token bucket there are no tokens shared from a rank higher than n , however depending on CF^0 there may be tokens shared from the overflow of the committed token bucket at rank 1. Therefore $S_E^{n+1}(t_1, t_2) = CF^0 \times O_C^1(t_1, t_2)$.

$$T_C^i(t_1, t_2) = \frac{CIR^i}{8} \times (t_2 - t_1) + S_C^{i+1}(t_1, t_2) \text{ for } i = 1, 2, \dots, n$$

$$T_E^i(t_1, t_2) = \frac{EIR^i}{8} \times (t_2 - t_1) + S_E^{i+1}(t_1, t_2) + CF^i \times O_C^i(t_1, t_2) \text{ for } i = 1, 2, \dots, n$$

Note that [R150] mandates that $CF^n = 0$ if $CF^0 = 1$ in the equation for $T_E^n(t_1, t_2)$.

The number of tokens actually added to the token bucket is some or all of the tokens available to be added, limited by the maximum number of tokens allowed in the token bucket (CBS^i and EBS^i) and by the maximum rate at which tokens are allowed to be added (CIR_{\max}^i and EIR_{\max}^i).

$$A_C^i(t_1, t_2) = \min \left\{ T_C^i(t_1, t_2), CBS^i - B_C^i(t_1), \frac{CIR_{\max}^i}{8} \times (t_2 - t_1) \right\} \text{ for } i = 1, 2, \dots, n$$

$$A_E^i(t_1, t_2) = \min \left\{ T_E^i(t_1, t_2), EBS^i - B_E^i(t_1), \frac{EIR_{\max}^i}{8} \times (t_2 - t_1) \right\} \text{ for } i = 1, 2, \dots, n$$

The number of tokens that overflow each token bucket is the number of available tokens that are not actually added to the token bucket.

$$O_C^i(t_1, t_2) = T_C^i(t_1, t_2) - A_C^i(t_1, t_2) \text{ for } i = 1, 2, \dots, n$$

$$O_E^i(t_1, t_2) = T_E^i(t_1, t_2) - A_E^i(t_1, t_2) \text{ for } i = 1, 2, \dots, n$$

The number of tokens available to be shared from Bandwidth Profile Flow i to the next lower rank Bandwidth Profile Flow, $i-1$, is the number of tokens that overflow the committed and excess token buckets and, in the case of the committed token bucket, are not made available to the excess token bucket by the Coupling Flag, CF^i .

$$S_C^i(t_1, t_2) = (1 - CF^i) \times O_C^i(t_1, t_2) \text{ for } i = 2, 3, \dots, n$$

$$S_E^i(t_1, t_2) = O_E^i(t_1, t_2) \text{ for } i = 2, 3, \dots, n$$

The Bandwidth Profile algorithm for an Envelope is shown in Figure 27. A visual representation of the algorithm is presented in Appendix C.1.

[R155] For an Envelope and for a sequence of Service Frames, $\{t_j, l_j, i_j\}, j \geq 0, t_{j+1} \geq t_j, i_j \in \{1, 2, \dots, n\}$ with arrival times at the reference point at times t_j , lengths l_j , and Bandwidth Profile Flows i_j that are mapped to the Envelope, the level of compliance assigned to each Service Frame **MUST** be defined according to the algorithm in Figure 27 where $B_C^i(t_0) = CBS^i$ and $B_E^i(t_0) = EBS^i$.

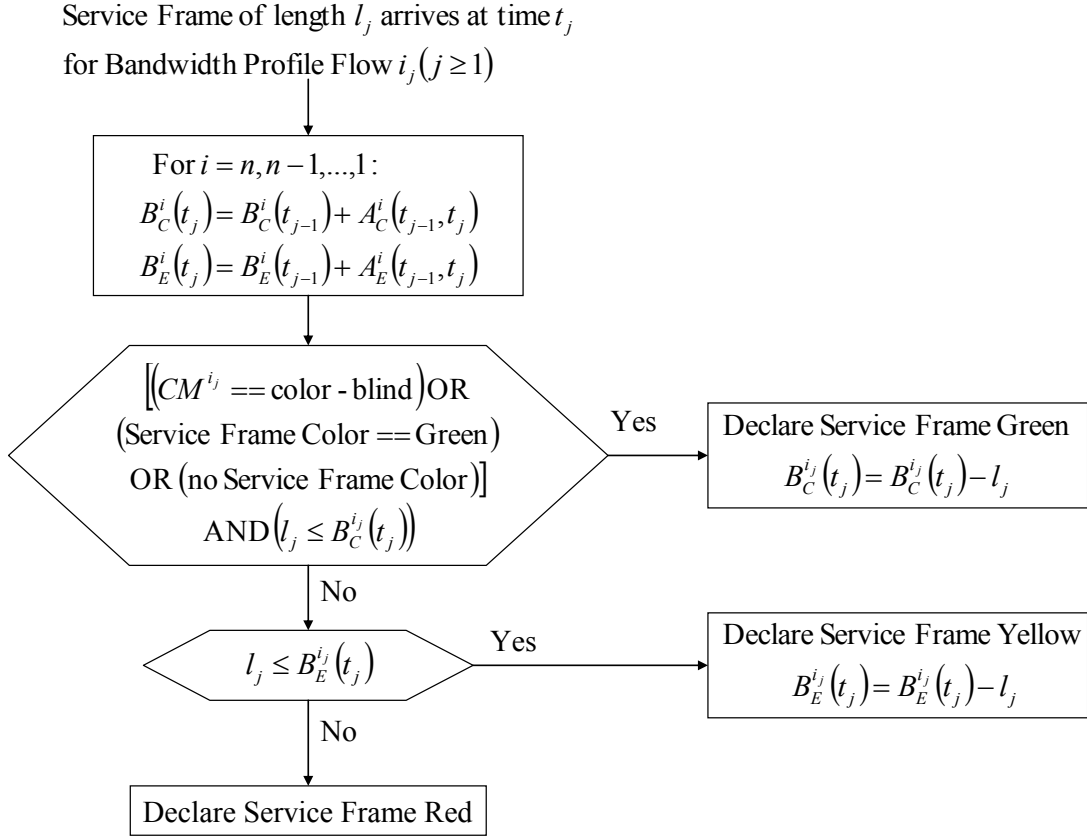


Figure 27 – The Bandwidth Profile Algorithm for an Envelope

Note that this document does not mandate how the behavior specified in Figure 27 is achieved by the CEN. The network functionality could be distributed and configured in any way. Also note that, since the algorithm assumes perfect time precision and zero time required for calculations, it is impossible to achieve the exact behavior described in Figure 27. A practical implementation needs to yield behavior "close" to the ideal behavior. Defining "close" is beyond the scope of this document. See Appendix C.3 for discussions of Bandwidth Profile implementation considerations.

12.2 Bandwidth Profile without Sharing

The algorithm of Figure 27 reduces to the Bandwidth Profile algorithm of Figure 13 of MEF 10.2 [15] when there is only one Bandwidth Profile Flow mapped to an Envelope and both CIR_{\max}^1 and EIR_{\max}^1 are sufficiently large. The following derivation demonstrates this fact.

When $n = 1$, $CIR_{\max}^1 \geq CIR^1$, and $EIR_{\max}^1 \geq EIR^1 + (CF^1 \times CIR^1)$,

$$T_C^1(t_1, t_2) = \frac{CIR^1}{8} \times (t_2 - t_1) + S_C^{n+1}(t_1, t_2) = \frac{CIR^1}{8} \times (t_2 - t_1) \text{ and}$$

$$T_E^1(t_1, t_2) = \frac{EIR^1}{8} \times (t_2 - t_1) + CF^1 \times O_C^1(t_1, t_2) + S_E^{n+1}(t_1, t_2) = \frac{EIR^1}{8} \times (t_2 - t_1) + CF^1 \times O_C^1(t_1, t_2).$$

Recall that [R142] mandates $CF^0 = 0$ when there is only one Bandwidth Profile Flow mapped to the Envelope.

Since $CIR_{\max}^1 \geq CIR^1$, it follows that

$$\begin{aligned} A_C^1(t_1, t_2) &= \min \left\{ \frac{CIR^1}{8} \times (t_2 - t_1), CBS^1 - B_C^1(t_1), \frac{CIR_{\max}^1}{8} \times (t_2 - t_1) \right\} \\ &= \min \left\{ \frac{CIR^1}{8} \times (t_2 - t_1), CBS^1 - B_C^1(t_1) \right\} \end{aligned}$$

and

$$\begin{aligned} O_C^1(t_1, t_2) &= T_C^1(t_1, t_2) - A_C^1(t_1, t_2) = \frac{CIR^1}{8} \times (t_2 - t_1) - \min \left\{ \frac{CIR^1}{8} \times (t_2 - t_1), CBS^1 - B_C^1(t_1) \right\} \\ &= B_C^1(t_1) + \frac{CIR^1}{8} \times (t_2 - t_1) - CBS^1 - \min \left\{ B_C^1(t_1) + \frac{CIR^1}{8} \times (t_2 - t_1) - CBS^1, 0 \right\} \\ &= \max \left\{ B_C^1(t_1) + \frac{CIR^1}{8} \times (t_2 - t_1) - CBS^1, 0 \right\}. \end{aligned}$$

Since $EIR_{\max}^1 \geq EIR^1 + (CF^1 \times CIR^1)$,

$$\begin{aligned} A_E^1(t_1, t_2) &= \min \left\{ \frac{EIR^1}{8} \times (t_2 - t_1) + CF^1 \times O_C^1(t_1, t_2), EBS^1 - B_E^1(t_1), \frac{EIR_{\max}^1}{8} \times (t_2 - t_1) \right\} \\ &= \min \left\{ \frac{EIR^1}{8} \times (t_2 - t_1) + CF^1 \times O_C^1(t_1, t_2), EBS^1 - B_E^1(t_1) \right\}. \end{aligned}$$

Then the computation of the number of tokens in each bucket is

$$\begin{aligned} B_C^1(t_2) &= B_C^1(t_1) + A_C^1(t_1, t_2) = \min \left\{ B_C^1(t_1) + \frac{CIR^1}{8} \times (t_2 - t_1), CBS^1 \right\} \text{ and} \\ B_E^1(t_2) &= B_E^1(t_1) + A_E^1(t_1, t_2) = \min \left\{ B_E^1(t_1) + \frac{EIR^1}{8} \times (t_2 - t_1) + CF^1 \times O_C^1(t_1, t_2), EBS^1 \right\}. \end{aligned}$$

Inspection of these equations reveals that the token counts depend on just the parameters for Bandwidth Profile Flow 1. Thus the level of compliance for each Service Frame in Bandwidth Profile Flow 1 is determined by the algorithm in Figure 28.

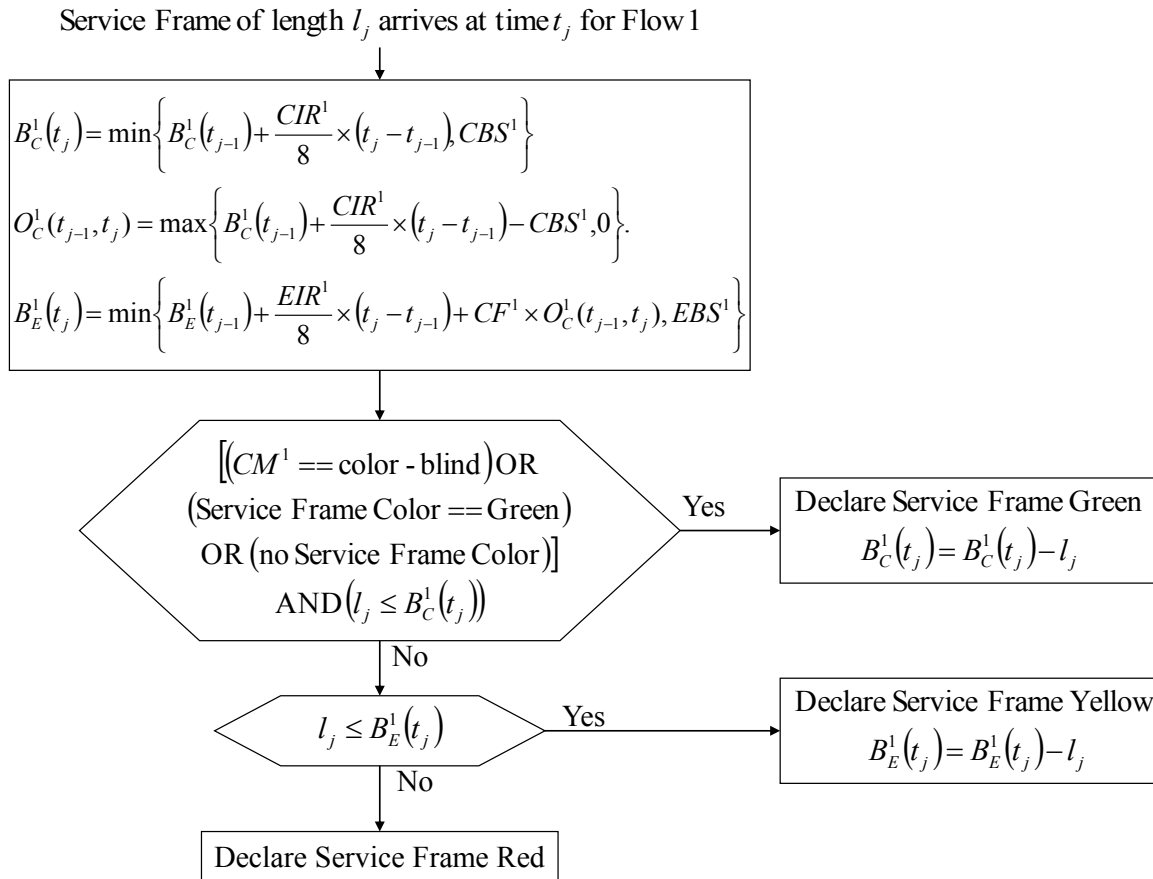


Figure 28 – Bandwidth Profile Algorithm for one Bandwidth Profile Flow in an Envelope

Figure 28 is equivalent to Figure 13 of MEF 10.2 [15] and thus the algorithm of Figure 27 is backward compatible with the Bandwidth Profile algorithm of [15] when a single Bandwidth Profile Flow is in an Envelope, CIR_{\max}^1 is sufficiently large, and EIR_{\max}^1 is sufficiently large.

12.3 Ingress Bandwidth Profiles Service Attributes

The Ingress Bandwidth Profile is used to regulate the amount of ingress traffic at a particular UNI. An Ingress Bandwidth Profile is defined for ingress Service Frames at the particular UNI. In other words, the sequence $\{t_j, l_j, i_j\}, j \geq 0$, to which the algorithm of Section 12.1 is applied is based on ingress Service Frames at a UNI. There are three Ingress Bandwidth Profile models as described in Sections 9.14, 10.5, and 10.6.

The disposition of a given ingress Service Frame with respect to delivery to an egress UNI is dependent on the Service Frame's level of compliance to the Ingress Bandwidth Profile that is applied to it. This is defined to be the Ingress Bandwidth Profile compliance level and it has three possible values: Green, Yellow, or Red.

[R156] An ingress Service Frame that is declared Red by an Ingress Bandwidth Profile **MUST** be discarded.

Note that if an ingress Service Frame is declared Green by an Ingress Bandwidth Profile and it meets other conditions per the definition of Qualified Service Frame [R31] in Section 8.8, then the SLS applies to this Service Frame.

Note that, per the definition of Qualified Service Frame in [R31] in Section 8.8, an ingress Service Frame that is declared Yellow by an Ingress Bandwidth Profile is not subject to the performance objectives of the SLS.

The better the level of Ingress Bandwidth Profile compliance, the fewer Service Frames will be discarded. In order to improve the level of Ingress Bandwidth Profile compliance, a Subscriber may need to shape traffic in the CE (see Appendix A.3).

12.4 Egress Bandwidth Profiles Service Attributes

An Egress Bandwidth Profile is used to regulate the amount of egress traffic at a particular UNI. An Egress Bandwidth Profile is defined for a particular UNI and applies to all or a subset of all egress Service Frames at the UNI in question.

The reference point for an Egress Bandwidth Profile is the egress UNI. An Egress Bandwidth Profile describes arrival times and lengths of Service Frames that will be observed at the egress UNI when an Egress Bandwidth Profile is in operation in the CEN. This description is given in terms of what would happen if an observer at the egress UNI applied the algorithm of Section 12.1 to egress Service Frames. This observer would see traffic after it had been subject to rate limiting and/or shaping in the CEN and thus would have certain characteristics. These characteristics are described in terms of the behavior of the algorithm of Section 12.1 when run by the observer.

[R157] When a sequence of egress Service Frames is subject to an Egress Bandwidth Profile with parameters $\langle CIR^i, CIR_{\max}^i, CBS^i, EIR^i, EIR_{\max}^i, EBS^i, CF^i, CM^i, ER^i \rangle$ for $i = 1, 2, \dots, n$, and CF^0 , the result of the application of the algorithm of Section 12.1 to these egress Service Frames, **MUST** be to declare each Service Frame either Green or Yellow.

The implication is that the regulation of the Service Frames in the CEN is such that all Service Frames that would be determined to be Red by the observer are discarded before reaching the egress UNI. It is important to reiterate that this description of Egress Bandwidth Profile does not mandate or constrain in any way the implementation in the Service Provider network.

There are three Egress Bandwidth Profile models as described in Sections 9.15, 10.7, and 10.8.

12.5 Simultaneous Application of the Ingress Bandwidth Profile Application Models

[O8] Multiple models of Ingress Bandwidth Profile application **MAY** exist simultaneously at a UNI.

In view of [R137]:

- If there is a per UNI Ingress Bandwidth Profile, then there cannot be any other Ingress Bandwidth Profiles at that UNI.
- If there is a per EVC Ingress Bandwidth Profile on an EVC, then there cannot be any per Class of Service Ingress Bandwidth Profiles on that EVC.

For example, in the configuration of Figure 25, there cannot be an Ingress Bandwidth Profile for EVC₁. Note also that for the configuration in Figure 25, that it is possible to configure a per-EVC Ingress Bandwidth Profile for EVC₂ but there happens to not be an Ingress Bandwidth Profile for EVC₂ in this example.

12.6 Simultaneous Application of the Egress Bandwidth Profile Application Models

- [O9] Multiple models of Egress Bandwidth Profile application **MAY** exist simultaneously for an egress UNI.

In view of [R137]:

- If there is a per UNI Egress Bandwidth Profile, then there cannot be any other Egress Bandwidth Profiles at that UNI.
- If there is a per EVC Egress Bandwidth Profile on an EVC, then there cannot be any per Egress Equivalence Class Identifier Bandwidth Profiles on that EVC.

13. Ethernet Service Framework

The parameter values for the Service Attributes define the capabilities of an Ethernet Service. Some or all of the Service Attributes may apply to an Ethernet Service. Service Attributes are described in Sections 8, 9, 10, 11, and 12.

For a particular Ethernet Service, there are four types of Service Attributes, those that apply to a UNI as described in Section 9, those that apply to an EVC as described in Section 8, those that apply to an EVC per UNI as described in Section 10, and those that apply across the EVC in question and other EVCs as described in Section 11.

The UNI Service Attributes are listed in Table 26 along with their possible parameter values. For a given instance of a service, a table like that of Table 26 needs to be specified for each UNI in the EVC associated with the service.

Attribute	Values
UNI ID (Section 9.1)	A non-null RFC 2579 DisplayString no greater than 45 characters
Physical Layer (Section 9.2)	A Standard Ethernet PHY per IEEE Std 802.3 – 2012 [3] but excluding 1000BASE-PX-D and 1000BASE-PX-U for each physical link implementing the UNI.
Synchronous Mode (Section 9.3)	Enabled or Disabled for each physical link implementing the UNI.
Number of Links (Section 9.4)	≥ 1
UNI Resiliency Attribute (Section 9.5)	None or Link Aggregation
Service Frame Format (Section 9.6)	IEEE Std 802.3 – 2012 [3]
UNI Maximum Service Frame Size (Section 9.7)	Integer ≥ 1522
Service Multiplexing (Section 9.8)	Enabled or Disabled ²⁶
CE-VLAN ID for Untagged and Priority Tagged Service Frames (Section 9.9)	An integer in 1, 2, ..., 4094.
CE-VLAN ID/EVC Map (Section 9.10)	Map as per Section 9.10
Maximum Number of EVCs (Section 9.11)	Integer ≥ 1
Bundling (Section 9.12)	Enabled or Disabled ²⁶
All to One Bundling (Section 9.13)	Enabled or Disabled ²⁷
Ingress Bandwidth Profile per UNI (Section 9.14)	No or parameters as defined in Section 12.1
Egress Bandwidth Profile per UNI (Section 9.15)	No or parameters as defined in Section 12.1
Link OAM (Section 9.16)	Enabled or Disabled
UNI MEG (Section 9.17)	Enabled or Disabled
E-LMI (Section 9.18)	Enabled or Disabled
Layer 2 Control Protocols Processing (Section 9.19)	Requirements for the processing of L2CP Service Frames are contained in MEF 6.1.1.[13]

Table 26 – UNI Service Attributes

The EVC Service Attributes are listed in Table 27 along with their possible parameter values. For a given instance of a service, a table like that of Table 27 needs to be specified for the EVC associated with the service.

²⁶ Must be Disabled if All to One Bundling is Enabled.

²⁷ Must be Disabled if Bundling is Enabled or Service Multiplexing is Enabled.

Attribute	Values
EVC Type (Section 8.1)	Point-to-Point, Multipoint-to-Multipoint, or Rooted-Multipoint
EVC ID (Section 8.2)	A non-null RFC 2579 DisplayString no greater than 45 characters
UNI List (Section 8.3)	A list of <UNI ID, UNI Role> pairs
Maximum Number of UNIs (Section 8.4)	Integer. 2 if EVC Type is Point-to-Point. Greater than or equal to 3 otherwise.
Unicast Data Service Frame Delivery (Section 8.5.2)	Deliver Unconditionally or Deliver Conditionally.
Multicast Data Service Frame Delivery (Section 8.5.2)	Deliver Unconditionally or Deliver Conditionally.
Broadcast Data Service Frame Delivery (Section 8.5.2)	Deliver Unconditionally or Deliver Conditionally.
CE-VLAN ID Preservation (8.6.1)	Enabled or Disabled
CE-VLAN CoS Preservation (Section 8.6.2)	Enabled or Disabled
Layer 2 Control Protocols Processing (Section 8.7)	Requirements for the processing of L2CP Service Frames are contained in MEF 6.1.1.[13]
EVC Performance (Section 8.8)	Performance objectives and parameters as described in Section 8.8. ²⁸
EVC Maximum Service Frame Size (Section 8.9)	Integer ≥ 1522

Table 27 – EVC Service Attributes

The EVC per UNI Service Attributes are listed in Table 28 along with their possible parameter values. For a given instance of a service, a table like that of Table 28 needs to be specified for each UNI in the EVC associated with the service.

²⁸ Note that an SLS can also specify performance objectives spanning more than one EVC as described in Section 11.

Attribute	Values
UNI EVC ID (Section 10.1)	A string formed by the concatenation of the UNI ID and the EVC ID
Class of Service Identifiers (Section 10.2)	Basis and values as described in Section 10.2.
Color Identifier (Section 10.3)	Basis and values as described in Section 10.3.
Egress Equivalence Class Identifiers (Section 10.4)	Basis and values as described in Section 10.4.
Ingress Bandwidth Profile Per EVC (Section 10.5)	No or parameters as defined in Section 12.1 ²⁹
Ingress Bandwidth Profile Per Class of Service Identifier (Section 10.6)	No or parameters as defined in Section 12.1 for each Class of Service Name ³⁰
Egress Bandwidth Profile Per EVC (Section 10.7)	No or parameters as defined in Section 12.1 for each EVC ³¹
Egress Bandwidth Profile Per Egress Equivalence Class Identifier (Section 10.8)	No or parameters as defined in Section 12.1 for each Egress Equivalence Class ³²
Source MAC Address Limit (Section 10.9)	Enabled or Disabled. If Enabled, a positive integer and a positive time interval.
Test MEG (Section 10.10)	Enabled or Disabled
Subscriber MEG MIP (Section 10.11)	Enabled or Disabled

Table 28 – EVC per UNI Service Attributes

14. References

- [1] IEEE Std 802.1AX – 2008, IEEE Standard for Local and metropolitan area networks – Link Aggregation, November 2008.
- [2] IEEE Std 802.1Q – 2011, IEEE Standards for Local and metropolitan area networks—Media Access Control (MAC) Bridges and Virtual Bridge Local Area Networks, 31 August 2011.
- [3] IEEE Std 802.3 – 2012, *Standard for Ethernet*, September 2012.
- [4] International Telecommunication Union, Recommendation G.781, Digital terminal equipments – Principal characteristics of multiplexing equipment for the synchronous digital hierarchy, 2008.
- [5] International Telecommunication Union, Recommendation G.8264, *Packet over Transport aspects – Quality and availability targets*, October 2008.
- [6] International Telecommunication Union, Recommendation Y.1563, *Ethernet frame transfer and availability performance*, 2009.

²⁹ Must be No if Ingress Bandwidth Profile per UNI is not No.

³⁰ Must be No if Ingress Bandwidth Profile per UNI is not No or Ingress Bandwidth Profile per EVC is not No.

³¹ Must be No if Egress Bandwidth Profile per UNI is not No.

³² Must be No if Egress Bandwidth Profile per UNI is not No or Egress Bandwidth Profile per EVC is not No.

- [7] International Telecommunication Union, Recommendation Y.1564, *Ethernet service activation test methodology*, May 2011.
- [8] Internet Engineering Task Force RFC 2119. Key words for use in RFCs to Indicate Requirement Levels, March 1997.
- [9] Internet Engineering Task Force RFC 2474, Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers, December 1998.
- [10] Internet Engineering Task Force RFC 2579. *Textual Conventions for SMIv2*, April 1999.
- [11] Internet Engineering Task Force RFC 3393, IP Packet Delay Variation Metric for IP Performance Metric (IPPM), November 2002.
- [12] MEF Forum MEF 6.1, *Ethernet Services Definitions - Phase 2*, April 2008.
- [13] MEF Forum, MEF 6.1.1, Layer 2 Control Protocol Handling Amendment to MEF 6.1, January 2012.
- [14] MEF Forum MEF 7.2, *Carrier Ethernet Information Model*, April 2013.
- [15] MEF Forum MEF 10.2, *Ethernet Services Attributes Phase 2*, October 2009.
- [16] MEF Forum MEF 10.2.1, *Performance Attributes Amendment to MEF 10.2*, January 2011.
- [17] MEF Forum MEF 12.1, Carrier Ethernet Network Architecture Framework Part 2: Ethernet Services Layer - Base Elements, April 2010.
- [18] MEF Forum MEF 16, *Ethernet Local Management Interface (E-LMI)*, January 2006.
- [19] MEF Forum MEF 17, Service OAM Framework and Requirements – Phase 1, April 2007.
- [20] MEF Forum MEF 20, User Network Interface (UNI) Type 2 Implementation Agreement, July 2008.
- [21] MEF Forum MEF 22.1, *Mobile Backhaul Phase 2*, January 2012.
- [22] MEF Forum MEF 23.1, *Carrier Ethernet Class of Service Phase 2*, January 2012.
- [23] MEF Forum MEF 26.1, External Network Network Interface (ENNI) –Phase 2, January 2012.
- [24] MEF Forum MEF 28, External Network Network Interface (ENNI) Support for UNI Tunnel Access and Virtual UNI, October 2010.

- [25] MEF Forum MEF 30.1, Service OAM Fault Management Implementation Agreement: Phase 2, April 2013.
- [26] MEF Forum MEF 33, *Ethernet Access Services Definition*, January 2012.
- [27] MEF Forum MEF 31, Service OAM Fault Management Definition of Managed Objects, January 2011.
- [28] MEF Forum MEF 35, Service OAM Performance Monitoring Implementation Agreement, April 2012.
- [29] MEF Forum MEF 36, Service OAM SNMP MIB for Performance Monitoring, January 2012.
- [30] MEF Forum MEF 38, *Service OAM Fault Management YANG Modules*, April 2012.
- [31] MEF Forum MEF 40, *UNI and EVC Definition of Managed Objects*, April 2013.

Appendix A Examples (Informative)

This appendix contains examples of some of the attributes specified in this Technical Specification. They are for illustrative purposes only. In the event of a conflict between the material in this appendix and the main body of this text, the material in the main body is controlling.

A.1 CE-VLAN ID Preservation Service Attribute

The following is a list of examples covering the CE-VLAN ID Preservation Service Attribute for both CE-VLAN ID Preservation Enabled and CE-VLAN ID Preservation Disabled. See Section 8.6.1.

A.1.1 CE-VLAN ID Preservation Enabled

Figure 29 shows the notation used for the CE-VLAN ID/EVC Maps in the examples in this subsection.

Ingress Map	Egress Map	Ingress Service Frame Format	Egress Service Frame Format
A*	A*	Untagged Priority Tagged VLAN Tagged	Untagged Priority Tagged VLAN Tagged
B	B	VLAN Tagged	VLAN Tagged
Scenario A*: Untagged/Priority Tagged Service Frames are mapped to the EVC Scenario B: Untagged/Priority Tagged Service Frames are not mapped to the EVC			

Figure 29 – CE-VLAN ID/EVC Map Notation




Figure 30 – Example 1: CE-VLAN ID Preservation Enabled with All to One Bundling Enabled³³

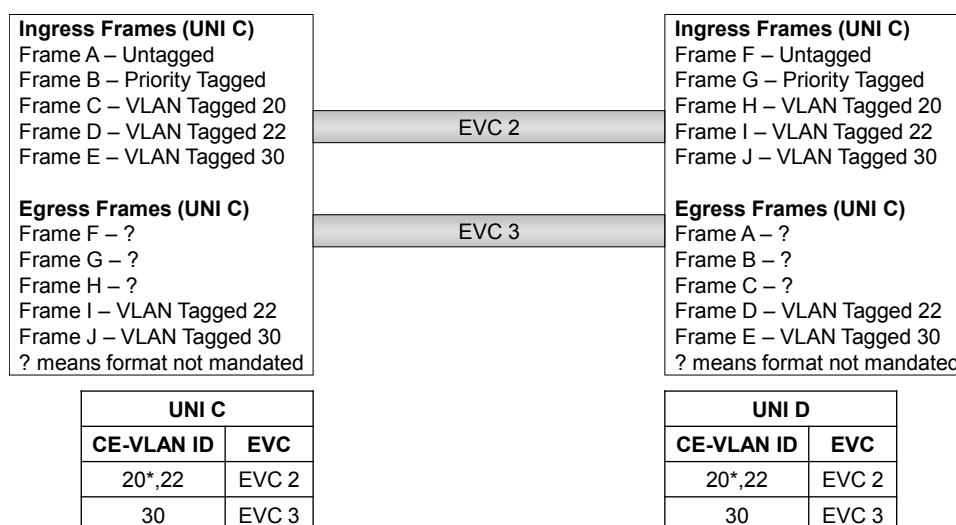


Figure 31 – Example 2: CE-VLAN ID Preservation Enabled with Bundling Enabled and Multiple CE-VLAN IDs mapped to EVC2³³

A.1.2 CE-VLAN ID Preservation Disabled

Figure 32 shows the notation used for the CE-VLAN ID/EVC Maps in the examples in this subsection.

³³ When a UNI has the All to One Bundling Enabled or Bundling Enabled, CE-VLAN ID Preservation is mandated to be Enabled.

Ingress Map	Egress Map	Ingress Service Frame Format	Egress Service Frame Format
A*	A*	Untagged Priority Tagged VLAN Tagged	Untagged Untagged Untagged
A*	B	Untagged Priority Tagged VLAN Tagged	VLAN Tagged VLAN Tagged VLAN Tagged
B	A*	VLAN Tagged	Untagged
B	B	VLAN Tagged	VLAN Tagged
Scenario A*: Untagged/Priority Tagged Service Frames are mapped to the EVC Scenario B: Untagged/Priority Tagged Service Frames are not mapped to the EVC			

Figure 32 – CE-VLAN ID/EVC Map Notation

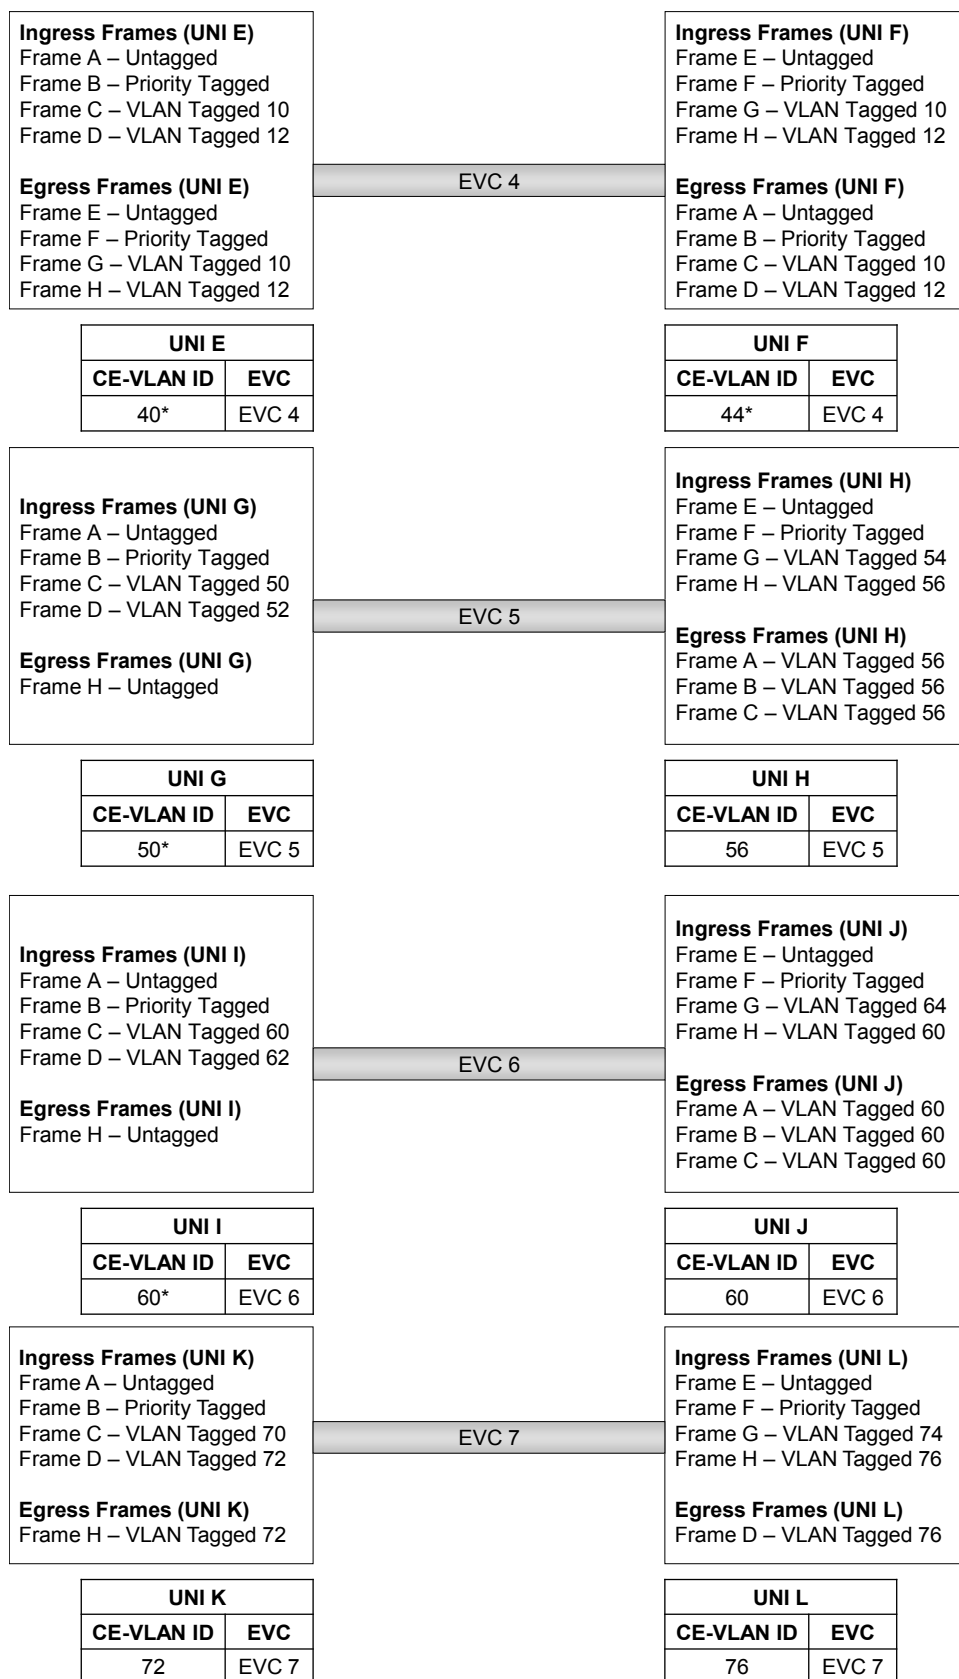


Figure 33 – Example 3: CE-VLAN ID Preservation Disabled

A.2 Examples of the Use of the CE-VLAN ID/EVC Map and EVCs

This section presents examples of the use of EVCs and the CE-VLAN ID/EVC Map. It is intended to clarify the concepts and present likely deployment scenarios.

A.2.1 Untagged UNIs

In connecting branch enterprise locations to a hub enterprise location, it is desirable to make the configuration of the branch CEs simple. A similar objective applies to providing access to higher layer services, e.g., Internet Access, where the configuration of the CE at the sites accessing the service should be kept simple. Figure 34 shows an example of 3 UNIs (A, B, and C) where CEs only capable of handling Untagged Service Frames are attached. The CE-VLAN ID/EVC maps are shown for each UNI. The asterisk indicates the CE-VLAN ID assigned to Untagged/Priority Tagged Service Frames.

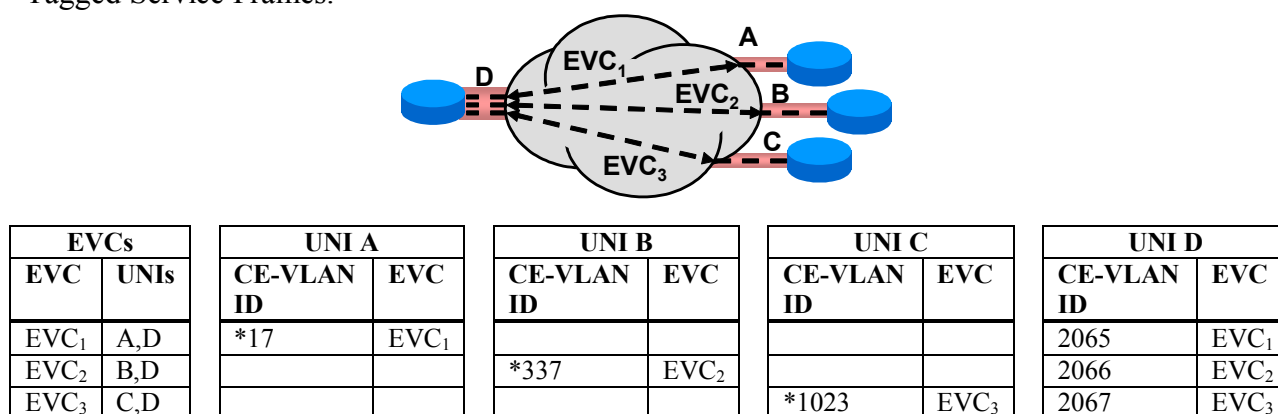


Figure 34 – Untagged UNIs

Consider an ingress Untagged Service Frame at UNI A. It will be mapped to EVC₁ and delivered to UNI D. At UNI D, it will become a VLAN Tagged egress Service Frame with VLAN ID 2065. An ingress VLAN Tagged Service Frame at UNI D with VLAN ID 2065 will be mapped to EVC₁ and delivered to UNI A. At UNI A, it will become an egress Untagged (see Section 9.9) Service Frame.

A.2.2 Use of Rooted-Multipoint EVC

An example of the use of the Rooted-Multipoint EVC is shown in Figure 35. A higher layer service is being provided via UNI D to three different customers at UNIs A, B, and C. By using a Rooted-Multipoint EVC, all three customers can be reached by the higher layer service provider at UNI D using a single EVC. Each customer's CE can only send to the higher layer service CE thus keeping each customer from seeing other customers' sourced traffic. Compared with the example shown in Figure 34, this can save a large number of Point-to-Point EVCs when there are a large number of customers. Note that the CEs do not necessarily have to send and receive Tagged Service Frames. In particular, the CEs at UNIs A and C do not need to send or receive Tagged Service Frames in this example.

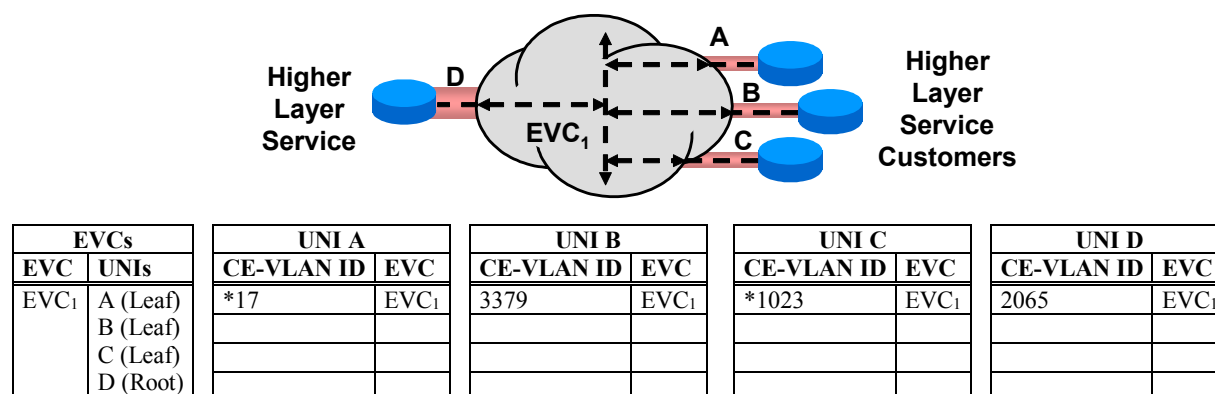


Figure 35 – Use of a Rooted-Multipoint EVC

A.2.3 Redundant Higher Layer Service Access

The example shown in Figure 36 illustrates the use of service multiplexing and Multipoint-to-Multipoint EVCs to provide redundant access to higher layer services. A Multipoint-to-Multipoint EVC is used for each customer of the higher layer service. Higher layer service routers are attached to two UNIs (C and D in the example) in each such EVC. Routing protocols running among the two higher layer service routers and the customer router allow the customer to access the higher layer service in a redundant fashion.

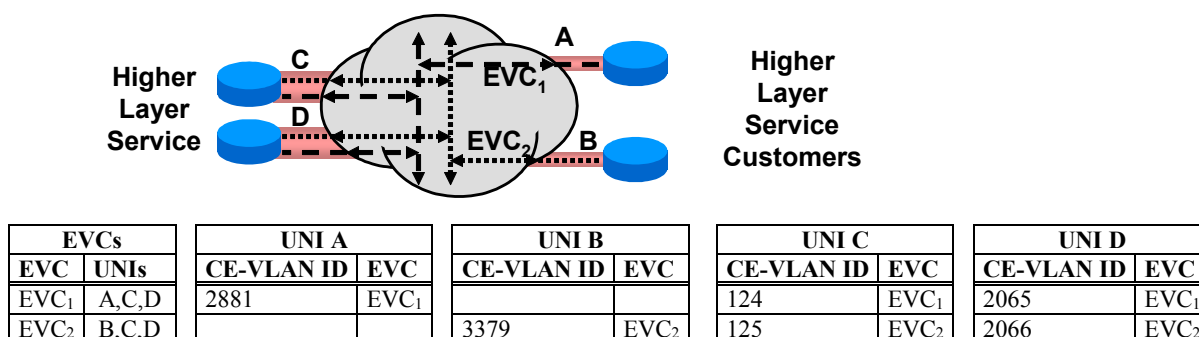


Figure 36 – Redundant Higher Layer Service Access

A.3 Traffic Shaping Example

Shaping is a procedure to reduce the burstiness of traffic. When done in the CE it is meant to increase the level of Ingress Bandwidth Profile compliance (see Section 12.3). A shaper is defined by a set of parameters. Those parameters should be chosen to ensure that the delay introduced by shaping function is bounded within the acceptable limits and that the traffic dropped at the shaper is kept to a minimum.

The example in this appendix assumes a single Bandwidth Profile Flow that is the only Bandwidth Profile Flow mapped to the Envelope. It also assumes the CM = color-blind and CF = 0.

A shaper could be a single rate or a double rate shaper. A single rate shaper could consist of three parameters CIR , CBS^* , and CBS , in which:

- CIR = the shaping rate of Green frames (average output rate of the shaper),

- CBS = the shaping burst of Green frames (maximum output burst of the shaper)
- CBS^* = the accepted burst of Green frames (maximum buffer size for Green frames)
- $CBS^* \geq CBS$, which means the shaper accepts larger bursts at its input and generates smaller bursts at its output.

A double rate shaper could consist of parameters CIR , CBS^* , CBS , EIR , EBS^* , and EBS , in which, CIR , CBS^* , and CBS are as defined above, and EIR , EBS^* , and EBS are:

- EIR = the shaping rate of Yellow frames (average output rate of the shaper),
- EBS = the shaping burst of Yellow frames (maximum output burst of the shaper)
- EBS^* = the accepted burst of Yellow frames (maximum buffer size for Yellow frames)

It is recommended that the CE shape the traffic it sends into the CEN, so that the output of the shaper matches the CIR , CBS , EIR , and EBS parameters of the appropriate Bandwidth Profile.

For example, the CE could shape with a dual rate token bucket shaper using parameters CIR , CBS^* , CBS , EIR , EBS^* , and EBS where $EBS^* = 0$ and CBS^* is the shaper's buffer size. Define the following notation:

- $B(t)$ = the instantaneous buffer occupancy in bytes,
- $C(t)$ = the instantaneous value of the tokens in the Committed token bucket,
- $E(t)$ = the instantaneous value of the tokens in the Excess token bucket,
- L = the length of the frame at the head of the buffer,
- LNF = length of a new frame given to the shaper, and
- THS = a configured buffer threshold such that the difference between THS and the shaper's buffer size, CBS^* , is large enough to hold a maximum sized frame.

Example shaping algorithms are presented below. In these algorithms, "transmit the frame" means transfer the frame to the transmission buffer for the outgoing Ethernet port of the CE. For simplicity, it is assumed that there is always space available in the transmission buffer. The algorithm in Figure 37 is run every Δt seconds where Δt is the period between updating the token bucket values $C(t)$ and $E(t)$, i.e.,

$$C(t) = \min\left(CBS, C(t - \Delta t) + \left(\frac{CIR}{8}\right)\Delta t\right) \text{ and } E(t) = \min\left(EBS, E(t - \Delta t) + \left(\frac{EIR}{8}\right)\Delta t\right).$$

Service Frames sent by the CE using this algorithm should always have an Ingress Bandwidth Profile compliance of Green.

```
while((L <= C(t)) && (B(t) > 0))
{
    C(t) = C(t) - L;
    B(t) = B(t) - L;
    transmit the frame at the head of the buffer; //Should be declared green
}
```

Figure 37 – Periodic Algorithm

The algorithm of Figure 38 is run every time a new frame is given to the shaper to process. This algorithm will send Service Frames with an Ingress Bandwidth Profile compliance of Yellow if necessary to try to make room in the buffer for the new frame.

```
if(B(t) == 0) // If buffer is empty
{
    if(LNF <= C(t))
    {
        C(t) = C(t) - LNF;
        transmit new frame; // Should be declared green
    }
    else
    {
        place new frame in buffer;
        B(t) = B(t) + LNF;
    }
}
else
{
    while(L <= C(t))
    {
        C(t) = C(t) - L;
        B(t) = B(t) - L;
        transmit the frame at the head of the buffer; //Should be declared green
    }
    if(B(t) <= THS)
    {
        place new frame in buffer;
        B(t) = B(t) + LNF;
    }
    else
    {
        while((L <= E(t) && (B(t) > THS))
        {
            E(t) = E(t) - L;
            B(t) = B(t) - L;
            transmit the frame at the head of the buffer; //Should be declared yellow
        }
        if(B(t) <= THS)
        {
            place new frame in buffer;
            B(t) = B(t) + LNF;
        }
        else
        {
            discard new frame;
        }
    }
}
```

Figure 38 – New Frame Algorithm

Shaping can also be used within the CEN to implement a low loss but higher delay SLS and/or to smooth traffic for more efficient use of network buffers.

A.4 Examples of Availability Metrics for Multipoint EVCs

The Performance Metric definitions for Multipoint EVCs provide a great deal of flexibility. This section provides examples on how the subset of UNIs in the EVC can be used to define UNI-oriented Performance Metrics (Section A.4.1) and EVC-oriented Performance Metrics (Section A.4.2). The Availability Performance Metric is used for these examples.

Both examples use the Multipoint EVC depicted in Figure 39. There are two Class of Service Names supported on the EVC, namely CoS Name 1 and CoS Name 2. The important traffic paths for each Class of Service Name have been agreed to by Subscriber and the Service Provider as shown in the figure.

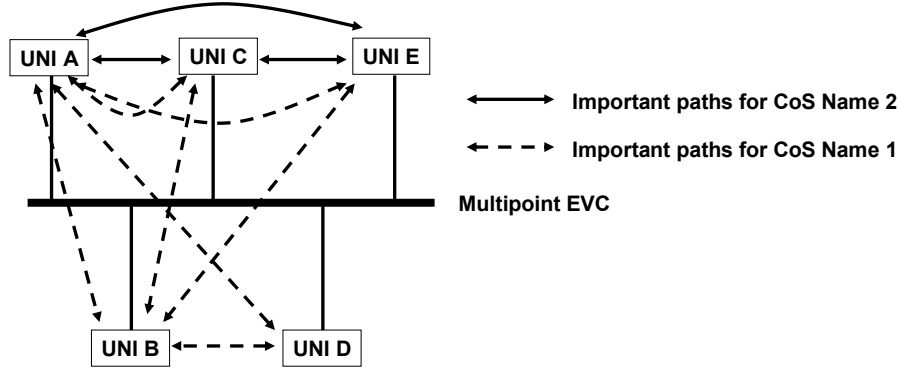


Figure 39 – Multipoint EVC Example

A.4.1 UNI-oriented Availability Example

In this case, an Availability Performance Metric is defined for each UNI for each Class of Service Name. The Performance Metric is based on the ability to communicate between the UNI in question and the other UNIs identified by the important traffic flows. Define the following subsets of ordered UNI pairs:

- $S_{A,1} = \{\langle A, B \rangle, \langle A, C \rangle, \langle A, D \rangle, \langle A, E \rangle\}$
- $S_{B,1} = \{\langle B, A \rangle, \langle B, C \rangle, \langle B, D \rangle, \langle B, E \rangle\}$
- $S_{C,1} = \{\langle C, A \rangle, \langle C, B \rangle\}$
- $S_{D,1} = \{\langle D, A \rangle, \langle D, B \rangle\}$
- $S_{E,1} = \{\langle E, A \rangle, \langle E, B \rangle\}$
- $S_{A,2} = \{\langle A, C \rangle, \langle A, E \rangle\}$
- $S_{C,2} = \{\langle C, A \rangle, \langle C, E \rangle\}$
- $S_{E,2} = \{\langle E, C \rangle, \langle E, A \rangle\}$

For this example, assume that T , Δt , C , and n , are used for all availability Performance Metrics. Then using the definition in Section 8.8.4, $A_T^{S_{A,1}}$ can be viewed as the availability of UNI A for Class of Service Name 1 and this reflects the availability of the important point to point paths

that UNI A is a part of. Similarly, $A_T^{S_{C,2}}$ can be viewed as the availability of UNI C for Class of Service Name 2. Thus, the availability for each UNI for each Class of Service Name can be defined by selecting the appropriate subset of UNI pairs.

A.4.2 EVC-oriented Availability Example

In this case an Availability Performance Metric is defined for each Class of Service Name supported by the EVC. Define the following subsets of ordered UNI pairs:

- $$S_1 = \left\{ \langle A, B \rangle, \langle B, A \rangle, \langle A, C \rangle, \langle C, A \rangle, \langle A, D \rangle, \langle D, A \rangle, \langle A, E \rangle, \langle E, A \rangle, \langle B, C \rangle, \langle C, B \rangle, \langle B, D \rangle, \right. \\ \left. \langle D, B \rangle, \langle B, E \rangle, \langle E, B \rangle \right\}$$
- $$S_2 = \left\{ \langle A, C \rangle, \langle C, A \rangle, \langle A, E \rangle, \langle E, A \rangle, \langle C, E \rangle, \langle E, C \rangle \right\}$$

For this example, assume that T , Δt , C , and n , are used for both availability Performance Metrics. Then using the definition in Section 8.8.4, $A_T^{S_1}$ can be viewed as the availability of Class of Service Name 1 on the EVC and $A_T^{S_2}$ can be viewed as the availability of Class of Service Name 2 on the EVC.

Appendix B Effect of Service Frame Overhead (Informative)

Figure 40 is a simplified version of Figure 3 – 1 in IEEE Std 802.3 – 2012 [3] which shows the format of the IEEE Std 802.3 – 2012 Packet. Because UNIs specified in this document do not run in half duplex mode, the Extension Field is always zero length and not shown in the figure. As can be seen from Figure 40 the Service Frame is exactly the IEEE Std 802.3 – 2012 Frame. IEEE Std 802.3 – 2012 Packets are separated by an Inter-Packet gap of at least 96 bit times (12 bytes). Thus the minimum separation between Service Frames transmitted across a UNI is 20 bytes.

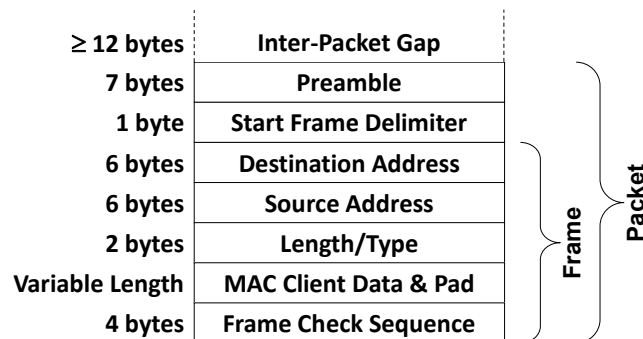


Figure 40 – IEEE Std 802.3 Packet Format (simplified)

ITU Y.1564 [7] defines Information Rate as the average bit rate of Service Frames at the measurement point starting with the first MAC address bit and ending with the last FCS bit. In the context of this document, the measurement point is the UNI.

To calculate the Information Rate of a group of Service Frames, take the total number of Service Frame bits (the sum of the lengths of each Service Frame in bits from Figure 40) divided by the amount of time from the first bit in the first Service Frame to the last bit in the last Service Frame plus the amount of time for 160 more bits (minimum Inter-Packet gap, preamble, and start of frame delimiter). Let S be the sum of the Service Frame bits in bits, let r be the UNI Line Rate in bits per second, let T_1 be the time when the first bit of the first Service Frame arrives at the UNI, and let T_2 be the time when the last bit of the last Service Frame arrives at the UNI. Then

$$\text{Information Rate for a group of Service Frames} = \frac{S}{T_2 - T_1 + \frac{160}{r}}.$$

For given Service Frame size, the Information Rate is maximized when the Service Frames in the group are back to back, i.e., the Service Frames are separated by 160 bits. For n back to back Service Frames each with length x bits,

$$T_2 - T_1 = \frac{nx + (n-1)160}{r} \text{ and}$$

Information Rate for n back to back Service Frames each with length x bits

$$= \frac{\frac{nx}{r} + \frac{(n-1)160}{r}}{\frac{nx + (n-1)160}{r} + \frac{160}{r}} = \left(\frac{x}{x + 160} \right) r.$$

Figure 41 shows the ratio of the Information Rate to UNI Line Rate as a function of Service Frame size for group of fixed length, back to back Service Frames.




Figure 41 – Ratio of Information Rate to UNI Line Rate for Back to Back Service Frames

Figure 41 has implications on the behavior of the Bandwidth Profile Algorithm of Section 12.1. When that algorithm is applied to a sequence of Service Frames, it deducts tokens from a token bucket based on the number of bytes in each Service Frame. Thus it is possible to pick a value of CIR^i that is less than the UNI Line Rate but results (perhaps after a transient period to allow the token bucket to accumulate tokens) in all Service Frames being declared Green. For example, if all Service Frames had lengths of no more than 100 bytes, then choosing CIR^i to be .85 of the UNI Line Rate would yield this behavior.

ITU Y.1564 [7] defines Utilized Line Rate as the average bit rate of the Ethernet line at the measurement point, including the bits a) allocable to the minimum-duration period of each Inter-Packet gap (but not the number of bits allocable to the part of each Inter-Packet gap longer than the minimum), b) in the preamble, c) in the start of frame delimiter and d) in the Ethernet Service Frame starting with the first MAC address bit and ending with the last FCS bit. In the context of this document, the measurement point is the UNI.

To calculate the Utilized Line Rate of a group of n Packets, take the total number of Packet bits (the length of each Packet in bits from Figure 40) plus n times the minimum Inter-Packet gap divided by the amount of time from the first bit in the first Packet to the last bit in the last Packet plus the time for minimum Inter-Packet gap (96 bits). Let P be the sum of the Packet bits in bits, let r be the UNI Line Rate in bits per second, let T_1 be the time when the first bit of the first Packet arrives at the UNI, and let T_2 be the time when the last bit of the last Packet arrives at the UNI.

In equation form:

$$\text{Utilized Line Rate for a group of } n \text{ Packets} = \frac{P + n96}{T_2 - T_1 + \frac{96}{r}}.$$

When the n Packets are back to back, $T_2 - T_1 = \frac{P + (n-1)96}{r}$ and the Utilized Line Rate for a group of back to back Packets is equal to the UNI Line Rate. Consequently,

Information Rate for n back to back Service Frames each with length x bits =

$$\left(\frac{x}{x+160} \right) \times \text{Utilized Line Rate for } n \text{ back to back Packets.}$$

Appendix C Bandwidth Profiles (Informative)

C.1 Visual Representation of the Bandwidth Profile

Figure 42 shows a conceptual diagram of the Bandwidth Profile Algorithm with three Bandwidth Profile Flows.

- The green and yellow trapezoids represent the Committed and Excess token buckets for each Bandwidth Profile Flow.
- The solid arrows labeled CIR^i or EIR^i represent a token source flowing into each bucket at the rate specified by the Bandwidth Profile.
- The dashed arrows show where tokens that overflow the bucket go when the number of tokens in the bucket reaches the capacity (CBS^i or EBS^i) specified by the Bandwidth Profile.
- The circles represent decision points in the path of the overflow tokens that are controlled by the Coupling Flag (CF^i).
- The X's represent points at which tokens are discarded (i.e., not added to the token count for any of the buckets).
- CIR_{\max}^i and EIR_{\max}^i are not shown on the figure.

The Coupling Flag for each Bandwidth Profile Flow controls the path of overflow tokens from the Committed token bucket to the Excess token bucket.

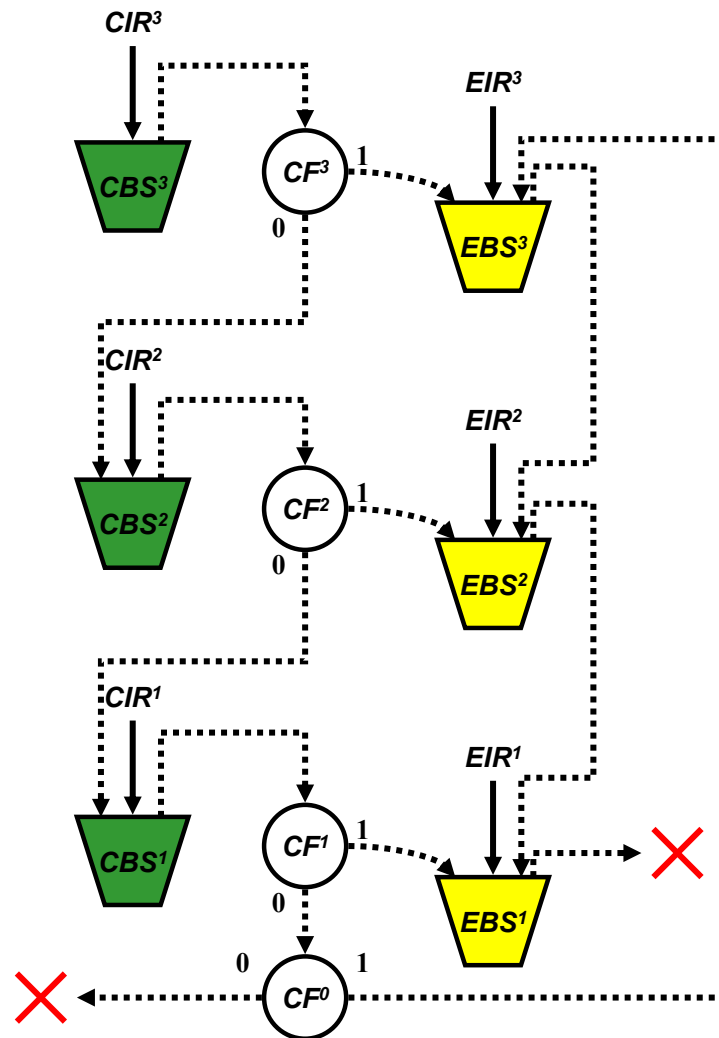


Figure 42 – Behavior of the Bandwidth Profile Algorithm

C.2 Bandwidth Profile Use Cases

C.2.1 Sharing Excess Bandwidth

In this use case, three Bandwidth Profile Flows are included in an Envelope with each Bandwidth Profile Flow having a dedicated CIR^i . Overflow Green tokens for each Bandwidth Profile Flow are sent to the Excess token bucket for that Bandwidth Profile Flow. Overflow Yellow tokens are shared with lower ranked Bandwidth Profile Flows. In this case, the amount of bandwidth of Qualified Service Frames for each Bandwidth Profile Flow is bounded by each CIR^i and overflows of Green tokens are available for frames that have no SLS commitment. The configuration of this example is $CF^0 = 0, CF^i = 1, i = 1, 2, 3$. Figure 43 shows the token paths for this configuration. Note that $EIR_{\max}^i, i = 1, 2, 3$ are not shown in Figure 43 for simplicity but their values can be set to limit the excess bandwidth available to a given Bandwidth Profile Flow.

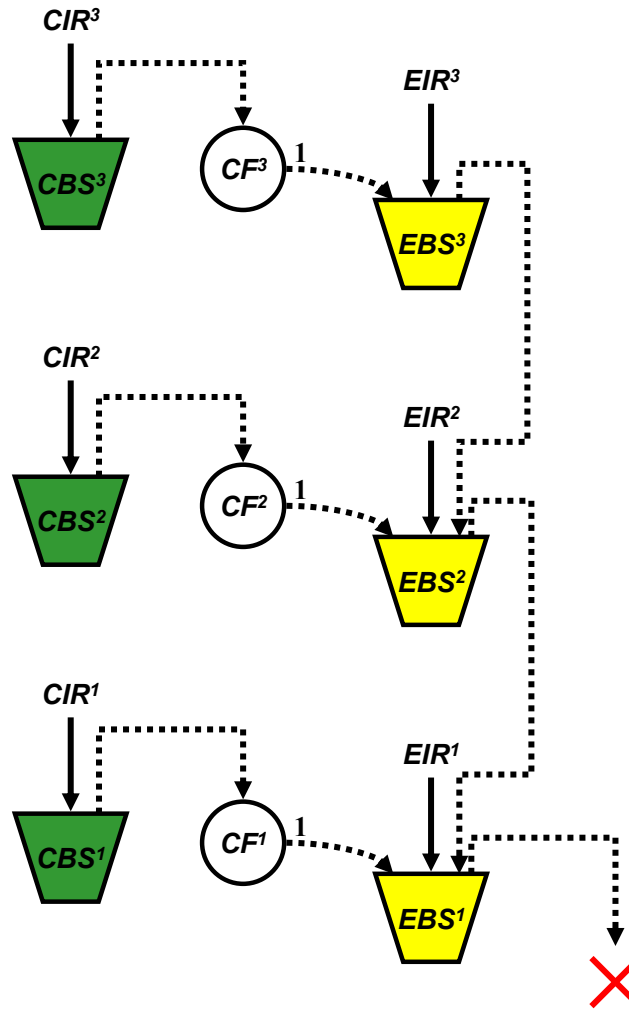


Figure 43 –Bandwidth Profile Algorithm for Sharing Excess Bandwidth

C.2.2 Uncoupled Bandwidth Sharing

In this use case, three Bandwidth Profile Flows are included in an Envelope. The Bandwidth Profile Flows share CIR_{env} committed bandwidth and EIR_{env} excess bandwidth.³⁴ Overflows of Green tokens are not converted to Yellow tokens. The configuration of this example is $CF^0 = 0, CF^i = 0, i = 1, 2, 3$, $CIR^1 = CIR^2 = 0$, $CIR^3 = CIR_{env}$, $EIR^1 = EIR^2 = 0$, and $EIR^3 = EIR_{env}$. Figure 44 shows the token paths for this configuration.

³⁴ CIR_{env} and EIR_{env} are not Bandwidth Profile parameters defined in Section 12.1. CIR_{env} represents the number of bits per second of total committed bandwidth shared by the three Bandwidth Profile Flows. EIR_{env} represents the number of bits per second of total excess bandwidth shared by the three Bandwidth Profile Flows.




Figure 44—Bandwidth Profile Algorithm for Uncoupled Bandwidth Sharing

The amount of bandwidth of Qualified Service Frames range from 0 to CIR_{env} for each Bandwidth Profile Flow. CIR_{max}^1 and CIR_{max}^2 can be used to limit committed bandwidth for Bandwidth Profile Flows 1 and 2 if so desired by the Service Provider and agreed to by the Subscriber. Similarly, EIR_{max}^1 and EIR_{max}^2 can be used to limit excess bandwidth for Bandwidth Profile Flows 1 and 2 if so desired by the Service Provider and agreed to by the Subscriber. A similar but different use case is to make CIR^1 and CIR^2 positive in order to ensure that Bandwidth Profile Flows 1 and 2 are not starved for committed bandwidth by higher ranked Bandwidth Profile Flows. Similarly, EIR^1 and EIR^2 can be used to ensure that Bandwidth Profile Flows 1 and 2 are not starved for excess bandwidth by higher ranked Bandwidth Profile Flows.

C.2.3 Sharing Bandwidth among Class of Service Names

In this case, CoS Labels H, M, and L share CIR_{env} bandwidth within an Envelope. H is rank 3, M is rank 2, and L is rank 1. Frames in the H Bandwidth Profile Flow are declared Green or Red. Frames in the M Bandwidth Profile Flow are declared Green, Yellow, or Red. Frames in the L

Bandwidth Profile Flow are declared Yellow or Red. The Bandwidth Profile Flow parameter values for this configuration are shown in Table 29 and $CF^0 = 0$.

Parameter	Bandwidth Profile Flow 3 (H)	Bandwidth Profile Flow 2 (M)	Bandwidth Profile Flow 1 (L)
CIR	CIR_{env}	0	0
CIR_{max}	$CIR_{max} \geq CIR_{env}$	$CIR_{max} \geq CIR_{env}$	0
CBS	> Maximum Service Frame Size	> Maximum Service Frame Size	0
EIR	0	0	0
EIR_{max}	0	$EIR_{max} \geq CIR_{env}$	$EIR_{max} \geq CIR_{env}$
EBS	0	> Maximum Service Frame Size	> Maximum Service Frame Size
CF	0	1	0

Table 29 – Parameter Values for Sharing Bandwidth among Class of Service Identifiers

Figure 45 shows the token paths for this configuration.

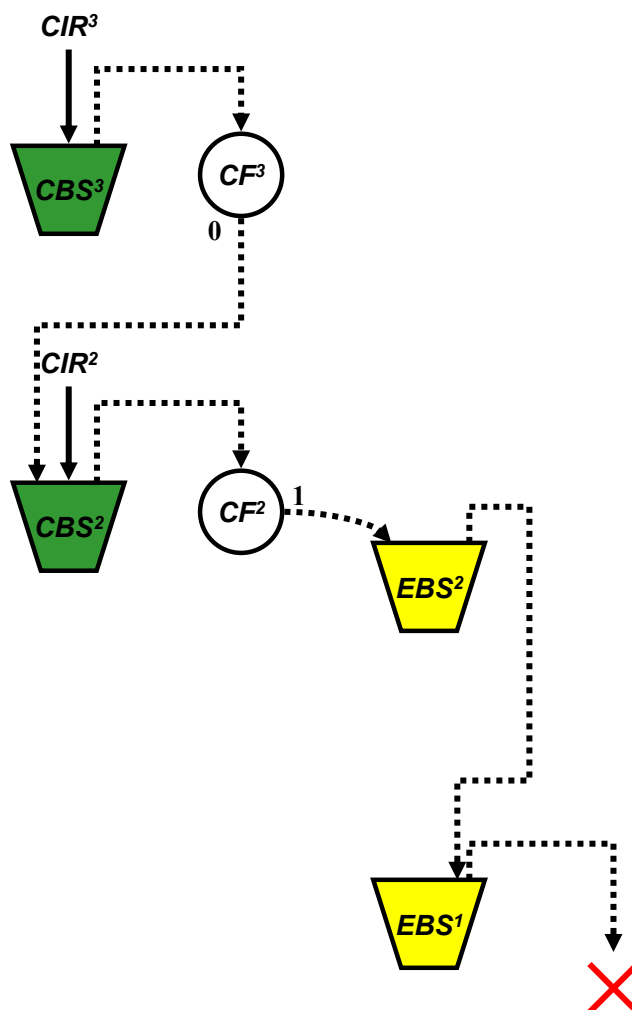


Figure 45–Bandwidth Profile Algorithm for Bandwidth Sharing among Class of Service Identifiers

C.2.4 Active/Standby EVCs

Figure 46 shows an example commonly found in Mobile Backhaul. In this configuration the two EVCs backup each other to provide connectivity from the Cell Site to a Radio Area Network controller. The Mobile Operator uses routing protocols to select the EVC to carry all data other than the routing protocols. If the active EVC fails, the Mobile Operator routers start sending data on the backup EVC.




Figure 46 – Active/Standby EVCs Use Case

The Service Provider constructs the Ingress Bandwidth Profile at the Cell Site UNI to limit the aggregate bandwidth allocated to the two EVCs. For this example, there are three Class of Service Labels supported on each EVC, H, M, and L as specified in MEF 23.1. [22]

To configure the Ingress Bandwidth Profile at the Cell Site UNI, six Bandwidth Profile Flows are specified, one each for each Class of Service Name on an EVC. The six Bandwidth Profile Flows are mapped to three Envelopes as shown in Figure 47. Recall that the Ingress Bandwidth Profile at the Cell Site is applied to traffic moving from right to left.

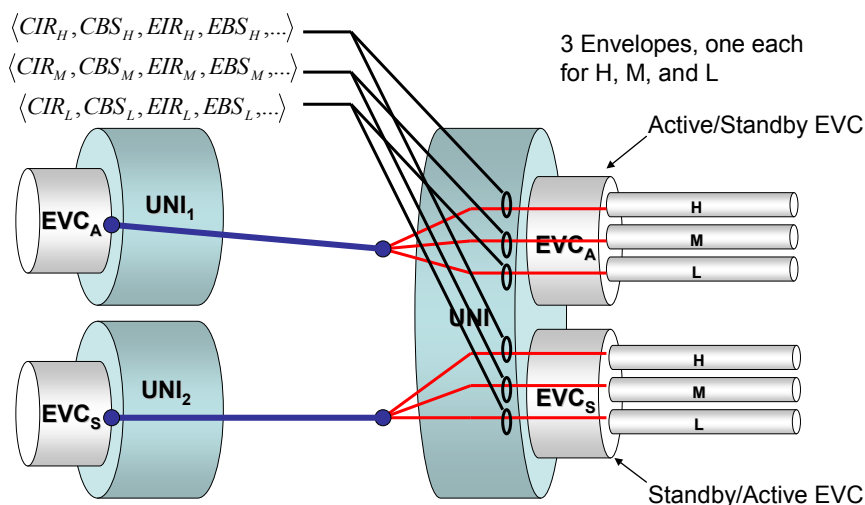


Figure 47 – Mapping of 6 Bandwidth Profile Flows to 3 Envelopes for the Active/Standby Use Case

The parameters shown in Figure 47 for each Envelope are set to accommodate the aggregate bandwidth for the given Class of Service Name. For example, the values for the parameters $\langle CIR_H, CBS_H, EIR_H, EBS_H, \dots \rangle$ are set based on the aggregate H cell user traffic plus the aggregate mobile network control traffic including the routing protocols. Table 30 shows how the Ingress Bandwidth Profile Bandwidth Profile Flow parameter values are set for the H Bandwidth Profile Flows. For the Envelope, $CF^0 = 0$. Note that CIR_r is a small value sufficient to accommodate the routing protocol on the Standby EVC.

Bandwidth Profile Flow 1: H Service Frames Mapped to Standby EVC		Bandwidth Profile Flow 2: H Service Frames Mapped to Active EVC	
Parameter	Value	Parameter	Value
CIR^1	CIR_r	CIR^2	CIR_H
CBS^1	CBS_H	CBS^2	CBS_H
EIR^1	0	EIR^2	EIR_H
EBS^1	EBS_H	EBS^2	EBS_H
CF^1	0	CF^2	0
CIR_{max}^1	∞	CIR_{max}^2	∞
EIR_{max}^1	∞	EIR_{max}^2	∞

Table 30 – Ingress Bandwidth Profile Parameter Values for the H Bandwidth Profile Flows

Figure 48 shows the resulting Ingress Bandwidth Profile algorithm for the Envelope of H Bandwidth Profile Flows. As can be seen, the combined long term average bandwidth of ingress Green Service Frames on the two EVCs is $CIR_H + CIR_r$.




Figure 48 –Bandwidth Profile Algorithm for the Envelope of H Bandwidth Profile Flows

Similar parameter setting and token sharing apply to the Envelope of M Bandwidth Profile Flows and the Envelope of L Bandwidth Profile Flows.

C.3 Bandwidth Profile Implementation

This appendix addresses implementation considerations for a Bandwidth Profile.

C.3.1 Eliminating Bias Toward Short Service Frames

In some circumstances the Bandwidth Profile algorithm specified in Section 12.1 can have a bias in favor of short Service Frames and against long Service Frames. The circumstance where this behavior arises is when the arriving traffic exceeds the specified rate (*CIR* and/or *EIR*) over a time interval significantly greater than can be accommodated by the specified burst tolerance (*CBS* and/or *EBS*). This behavior could be problematic for some applications. For instance, in an environment that typically handles a large number of simultaneous TCP connections, discarding long frames with a higher probability than short frames could conceivably create a situation where new connections are easily established but existing connections are unable to efficiently complete data transfers. In such environments it may be desirable to implement a modification of the Bandwidth Profile algorithm that eliminates the bias toward short Service Frames by making the color declaration of any Service Frame independent of the size of that frame.

The bias toward short Service Frames occurs when the token bucket is near empty. For simplicity, consider an example with a single token bucket (e.g. *CBS* > Maximum Service Frame Size and *EBS* = 0). When the number of tokens in the bucket is less than the maximum frame size, a short frame will be declared Green while a long frame will be declared Red. The arrival of a short frame will deplete the number of tokens in the bucket still further, which decreases the probability that a subsequent long frame would be declared Green. The arrival of several consecutive long frames may allow the bucket to refill sufficiently to declare an occasional long frame Green, but this depletes the bucket again which reinstates the condition for the short frame bias for subsequent frames. The behavior is self-reinforcing as long as the rate of arriving traffic exceeds the rate at which tokens are added to the bucket.

A Bandwidth Profile algorithm implementation can eliminate the bias toward short frames by making the decision to declare an arriving frame Green, Yellow, or Red independent of the length of the frame. This requires two things. First, declare an arriving Service Frame Green if there are any tokens in the bucket, instead of comparing the length of the Service Frame to the number of tokens in the bucket. If the frame length is greater than the number of tokens in the bucket, then declaring the frame Green may result in the token count going negative by up to a maximum frame size (effectively borrowing future tokens). This has a side effect of increasing the effective burst size by up to a maximum frame size. Second, measure the time interval between frames from the first bit of the previous frame to the first bit of the current frame, as opposed to measuring from the last bit of the previous frame to the last bit of the current frame. Measuring last bit to last bit actually introduces a bias toward long frames because of the tokens added to the bucket between the first bit and the last bit of the current frame.

Figure 49, Figure 50, and Figure 51 show simulation results of these variations on the Bandwidth Profile algorithm. The simulation is of an Ingress Bandwidth Profile on a 100Mbps UNI with *CIR* = 10Mbps, *CBS* = 1200B, and *EIR* = *EBS* = 0. The ingress traffic is a mix of 60% short (300 byte) and 40% long (1200 byte) Service Frames arriving in a randomized sequence with minimum inter-frame gap.

Figure 49 is the unmodified Bandwidth Profile algorithm of Section 12.1. After the initial transient when the token bucket gets depleted, short frames are declared Green in a dramatically higher percentage than long frames.




Figure 49 – Service Frame Profiling Results for the Algorithm of Section 12.1

Figure 50 shows the Bandwidth Profile algorithm modified to test for more than zero tokens in the bucket instead of more than the current frame length worth of tokens, but measures the time interval from the last bit of the previous frame to the last bit of the current frame. After the initial transient, the percentage of long frames declared Green is significantly higher than the percentage of long frames in the ingress traffic pattern.




Figure 50 – Service Frame Profiling Results for Modified Algorithm (Time Measured Last Bit to Last Bit)

Figure 51 shows the Bandwidth Profile algorithm modified to test for more than zero tokens in the bucket instead of more than the current frame length worth of tokens, and measures the time interval from the first bit of the previous frame to the first bit of the current frame. After the initial transient, the percentage of short and long frames declared Green is roughly equivalent to the percentage of short and long frames in the ingress traffic pattern.




Figure 51 – Service Frame Profiling Results for Modified Algorithm (Time Measured First Bit to First Bit)

Table 31 shows the numbers of each type of Service Frame declared Green by the three algorithms for the last 250 Service Frames in the simulation which corresponds to the behavior after the initial draining of the tokens. The first row is the offered traffic since all Service Frames not subject to an Ingress Bandwidth Profile are declared Green. The last two columns show the results as a percentage of the total number of Service Frames declared Green.

	Green Short Service Frames	Green Long Service Frames	Percent Green Short Service Frames	Percent Green Long Service Frames
No Bandwidth Profile	157	93	63%	37%
Algorithm of Section 12.1	54	0	100%	0%
Modified Algorithm (Time Measured Last Bit to Last Bit)	4	13	24%	76%
Modified Algorithm (Time Measured First Bit to First Bit)	19	8	70%	30%

Table 31 – Service Frame Profiling Results for the Last 250 Service Frames