MEF Standard

MEF 10.4

Subscriber Ethernet Service Attributes

December 2018

Disclaimer

Table of Contents

List of Figures

List of Tables

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.

albis-elcon
AT&T
Bell Canada
Ciena Corporation
Cisco Systems, Inc.
Colt
Coriant
Nokia
Verizon
ZTE Corporation

2. Abstract

This document describes Service Attributes for Subscriber Ethernet Services provided to an Ethernet Subscriber by an Ethernet Service Provider. The Service Attributes describe behaviors observable at an Ethernet User Network Interface and from Ethernet User Network Interface to Ethernet User Network Interface. In addition, a framework for defining specific instances of Subscriber Ethernet Services is described. Finally, extensive examples and explanatory material is included in appendices.

This document supersedes and replaces MEF 10.3, *Ethernet Services Attributes Phase 3* [24], MEF 10.3.1, *Composite Performance Metric (CPM) Amendment to MEF 10.3* [25], and MEF 10.3.2, *Amendment to MEF 10.3 - UNI Resiliency Enhancement* [26].

3. Terminology and Abbreviations

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 of the following table is used to provide the reference that is controlling, in other MEF or external documents.

Term	Definition	Reference
All to One Bundled UNI	A UNI with a single EVC EP to which all Service Frames are mapped.	This document
Available Time	A set of time intervals, contained in some longer time interval, when the service is considered available for use.	This document
Bandwidth Profile	A specification of the temporal properties of a sequence of Service Frames at a UNI, along with rules for determining the level of conformance to the specification for each Service Frame in the sequence.	This document
Bandwidth Profile Flow	A set of Service Frames at a UNI that meet specific criteria.	This document
Behavioral Service Attribute	A Service Attribute that directly affects the behavior of the service.	This document
Broadcast Data Service Frame	A Data Service Frame that has the broadcast Destination MAC Address.	This document
Bundling	(1) The condition when there is more than one C-Tag VLAN ID value mapped to an EVC EP at a UNI, or (2) The capability of the Service Provider to map more than one C-Tag VLAN ID to an EVC EP at a UNI.	This document
BWP Flow Parameters	Shorthand for the 10-tuple $\langle CIR, CIR_{max}, CBS, EIR, EIR_{max}, EBS, CF, CM, ER, F \rangle$	This document
Capability Service Attribute	A Service Attribute that does not directly affect the behavior of the service.	This document
Carrier Ethernet Network	Alternative term for Service Provider Network.	This document
CEN	Alternative term for Service Provider Network.	This document
Class of Service Name	An administrative name assigned to a particular set of Performance Objectives and related Bandwidth Profiles.	This document
CoS	Class of Service.	This document
C-Tag VLAN ID	The VID field of a C-Tagged Service Frame as defined in IEEE Std 802.1Q – 2018 [4].	This document

Term	Definition	Reference
C-Tag DEI	The DEI field of a C-Tagged Service Frame as defined in IEEE Std 802.1Q – 2018 [4].	This document
C-Tag PCP	The PCP field of a C-Tagged Service Frame as defined in IEEE Std 802.1Q – 2018 [4].	This document
C-Tagged Service Frame	A Service Frame that is either a VLAN Tagged Service Frame or a Priority Tagged Service Frame.	This document
Data Service Frame	A Service Frame that is intended to be delivered from ingress UNI to one or more egress UNIs.	This document
DEI	Drop Eligible Indicator.	IEEE Std 802.1Q – 2018 [4]
DSCP	Differentiated Services Code Point.	RFC 3260 [16]
Egress Bandwidth Profile	A Bandwidth Profile that applies to Egress Service Frames.	This document
Egress Service Frame	A Service Frame transmitted across the UNI toward the SN.	This document
Envelope	A set that contains an integer $n \geq 1$ number of Bandwidth Profile Flows that can share bandwidth resources that are represented by tokens.	This document
Ethernet Service	A connectivity service that carries Ethernet Frames irrespective of the underlying technology and that is specified using Service Attributes as defined in an MEF Specification.	This document
Ethernet Service Provider	An organization that provides Ethernet Services to Ethernet Subscribers.	This document
Ethernet Subscriber	The end-user of an Ethernet Service.	This document
Ethernet User Network Interface	The demarcation point between the responsibility of the Ethernet Service Provider and the Ethernet Service Subscriber.	This document
Ethernet Virtual Connection	An association of EVC EPs.	This document
EVC	Ethernet Virtual Connection.	This document
EVC End Point	A construct at a UNI that selects a subset of the Service Frames that pass over the UNI.	This document
EVC EP	EVC End Point.	This document

Term	Definition	Reference
Information Rate	The average bit rate of Ethernet frames at the measurement point, where each Ethernet frame is measured as starting with the first MAC address bit and ending with the last FCS bit.	Adapted from ITU Y.1564 [21]
Ingress Bandwidth Profile	A Bandwidth Profile that applies to Ingress Service Frames.	This document
Ingress Service Frame	A Service Frame transmitted across the UNI toward the SP Network.	This document
L2CP Service Frame	Layer 2 Control Protocol Service Frame.	This document
LAG	Link Aggregation Group.	IEEE Std 802.1AX – 2014 [3]
Layer 2 Control Protocol Service Frame	A Service Frame that could be used in a recognized Layer 2 Control Protocol.	This document
Leaf EVC EP	An EVC EP that can only exchange Service Frames with Root EVC EPs in an EVC.	This document
Link Number ID	An integer ≥ 1 that is uniquely assigned to each physical link at a given UNI.	Adapted from IEEE Std 802.1 AX – 2014 [3]
Link Selection Priority List	A sequence of Link Number IDs, in the order of usage preference, highest to lowest, for the link that is to carry the Service Frames corresponding to a given Port Conversation ID.	Adapted from IEEE Std 802.1 AX – 2014 [3]
Maintenance Interval	A time interval agreed to by both the Subscriber and Service Provider during which the service may not perform well or at all.	This document
Multicast Data Service Frame	A Data Service Frame with a multicast Destination MAC Address.	This document
Multipoint-to-Multipoint EVC	An EVC that can associate more than two EVC EPs and can only associate Root EVC EPs.	This document
One-way Frame Delay	The time elapsed from the reception of the first bit of the Ingress Service Frame at the ingress UNI until the transmission of the last bit of the first corresponding Egress Service Frame at the egress UNI.	This document

Term	Definition	Reference
PCP	Priority Code Point.	IEEE Std 802.1Q – 2018 [4]
Performance Metric	A quantitative characterization of Service Frame delivery quality experienced by the Subscriber.	This document
Performance Objective	A value associated with a Performance Metric that reflects the promised Service Frame delivery quality.	This document
Point-to-Point EVC	An EVC that can only associate exactly two Root EVC EPs.	This document
Port Conversation ID	An identifier for a set of Service Frames that are selected to pass over a physical link at a given UNI.	Adapted from IEEE Std 802.1 AX – 2014 [3]
Priority Tagged Service Frame	A Service Frame with a TPID = 0x8100 following the Source Address and a corresponding C-Tag VLAN ID value of 0x000 in the tag following the TPID.	This document
Qualified Service Frames	A set of Ingress Service Frames that comply with specific criteria and to which most of the Performance Metrics apply.	This document
Root EVC EP	An EVC EP that can exchange Service Frames with any other EVC EP in an EVC.	This document
Rooted-Multipoint EVC	An EVC that can associate more than two EVC EPs, at least one of which is a Root EVC EP and any number of which are Leaf EVC EPs.	This document
Service Attribute	Specific information that is agreed between the provider and the user of the service, that describes some aspect of the service behavior or capability.	This document
Service Frame	The IEEE Std 802.3 – 2015 Ethernet Frame.	This document
Service Frame Arrival Time	The time at which the first bit of the Service Frame arrives at the UNI.	This document
Service Level Agreement	The contract between the Subscriber and Service Provider specifying the service level commitments and related business agreements for a Service.	This document
Service Level Specification	The technical details of the service level, in terms of Performance Objectives, agreed between the Service Provider and the Subscriber as part of the SLA.	This document

Term	Definition	Reference
Service Multiplexing	(1) The condition when there is more than one EVC End Point located at a UNI. (2) The capability of the Service Provider to locate more than one EVC EP at a UNI.	This document
Service Provider	Short for Ethernet Service Provider in this document.	This document
Service Provider Network	An interconnected network used by the Service Provider to provide services to one or more Subscribers.	This document
SLA	Service Level Agreement.	This document
SLS	Service Level Specification.	This document
SN	Subscriber Network	This document
SOAM	Service Operations Administration and Maintenance.	MEF 30.1 [32]
SOAM Service Frame	A Service Frame, with or without a C-Tag, whose MAC Destination Address does not indicate it to be an L2CP Service Frame and whose Ethertype = 0x8902.	This document
SP Network	Short for Service Provider Network in this document	This document
Subscriber	Short for Ethernet Subscriber in this document.	This document
Subscriber Ethernet Service	An Ethernet Service that is provided by a Service Provider to Subscribers between two or more UNIs.	This document
Subscriber Network	A network belonging to a given Subscriber that is connected to the Service Provider at one or more UNIs.	This document
TPID	Tag Protocol Identifier.	IEEE Std 802.1Q – 2018 [4]
Unavailable Time	A set of time intervals, contained in some longer time interval, when the service is considered not available for use.	This document
UNI	Short for Ethernet User Network Interface in this document.	This document
UNI MAC Data Rate	The sum of the MAC Data Rate of each physical link 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

Term	Definition	Reference
Untagged UNI	A UNI with a single EVC EP to which only ingress Untagged Service Frames, ingress Priority Tagged Service Frames, and egress Untagged Service Frames are mapped.	This document
Virtual Frame	The information passed across the UNI when there are no physical links.	This document
VLAN Tagged Service Frame	A Service Frame with a TPID = 0x8100 following the Source Address and the corresponding C-Tag VLAN ID value is not 0x000 in the tag following the TPID.	This document
VLAN Tagged UNI	A UNI where only VLAN Tagged Service Frames are mapped to one or more EVC EPs.	This document

Table 1 – Terminology and Abbreviations

4. Compliance Levels

The key words "**MUST**", "**MUST NOT**", "**REQUIRED**", "**SHALL**", "**SHALL NOT**", "**SHOULD**", "**SHOULD NOT**", "**RECOMMENDED**", "**NOT RECOMMENDED**", "**MAY**", and "**OPTIONAL**" in this document are to be interpreted as described in BCP 14 (RFC 2119 [14], RFC 8174 [17]) when, and only when, they appear in all capitals, as shown here. All key words must be in 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.

5. 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

6. Introduction

The goals of this Technical Specification are three-fold. The first goal is to provide sufficient technical specificity to allow an Ethernet Subscriber to successfully plan and integrate Subscriber Ethernet Services into its overall networking infrastructure. The second goal is to provide enough detail so that vendors of equipment for the Subscriber Network can implement capabilities into their products such that they can be used to successfully access Subscriber 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 Subscriber Network. The third goal is to provide the Ethernet Service Provider with 1) the technical information that needs to be agreed upon with the Ethernet Subscriber, and 2) the details of the Subscriber Ethernet Service behaviors mandated by this technical information.

The model for Subscriber Ethernet Services is described as follows:

- Key concepts and definitions are detailed in Section 7.
- Service Attributes for the Ethernet Virtual Connection (Section 7.8) are described in Section 8.
- Service Attributes for the Ethernet User Network Interface (Section 7.2) are described in Section 9.
- Service Attributes for the EVC End Point (Section 7.8) are described in Section 10.
- Performance Service Attributes that apply to multiple Ethernet Virtual Connections are described in Section 11.
- The details of the Bandwidth Profile parameters and algorithm are described in Section 7.10 and Section 12.
- The Subscriber Ethernet Service Framework, which summarizes the possible values for each of the Service Attributes, is contained in Section 13.
- Extensive examples and supplementary explanations are contained in Appendix A through Appendix G.
- Appendix H summarizes the changes in going from MEF 10.3 [24] to this document.

7. Key Concepts and Definitions

This section introduces concepts and definitions that are used throughout this document.

7.1 Ethernet Subscriber and Ethernet Service Provider

This document deals with two types of organizations, the Ethernet Subscriber and the Ethernet Service Provider. The Ethernet Subscriber is the end-user of services described using the Service Attributes specified in this document and the Ethernet Service Provider is the provider of these services.

In the interest of brevity, the remainder of this document uses "Service Provider" as short for "Ethernet Service Provider" and "Subscriber" as short for "Ethernet Subscriber".

7.2 Ethernet User Network Interface, Subscriber Network, and Service Provider Network

The model underlying this document consists of three components as shown in Figure 1. The Ethernet User Network Interface is the demarcation point between the responsibility of the Service Provider and the responsibility of the Subscriber. Throughout this document, "UNI" refers to the "Ethernet User Network Interface".

The "Subscriber Network" (SN) is defined as a network belonging to a given Subscriber that is connected to the Service Provider at one or more UNIs. There are no assumptions about the details of the SN.

The Service Provider Network (SP Network) is defined as an interconnected network used by the Service Provider to provide services to one or more Subscribers. There are no assumptions about the details of the SP Network. It could consist of a single switch or an agglomeration of networks based on many different technologies.

Note that MEF 10.3 [24] uses the term "Carrier Ethernet Network" and the acronym "CEN". This document replaces those terms with "Service Provider Network" and "SP Network" respectively.

For an architectural perspective of the SP Network see "Carrier Ethernet Network", in MEF 12.2. [27].

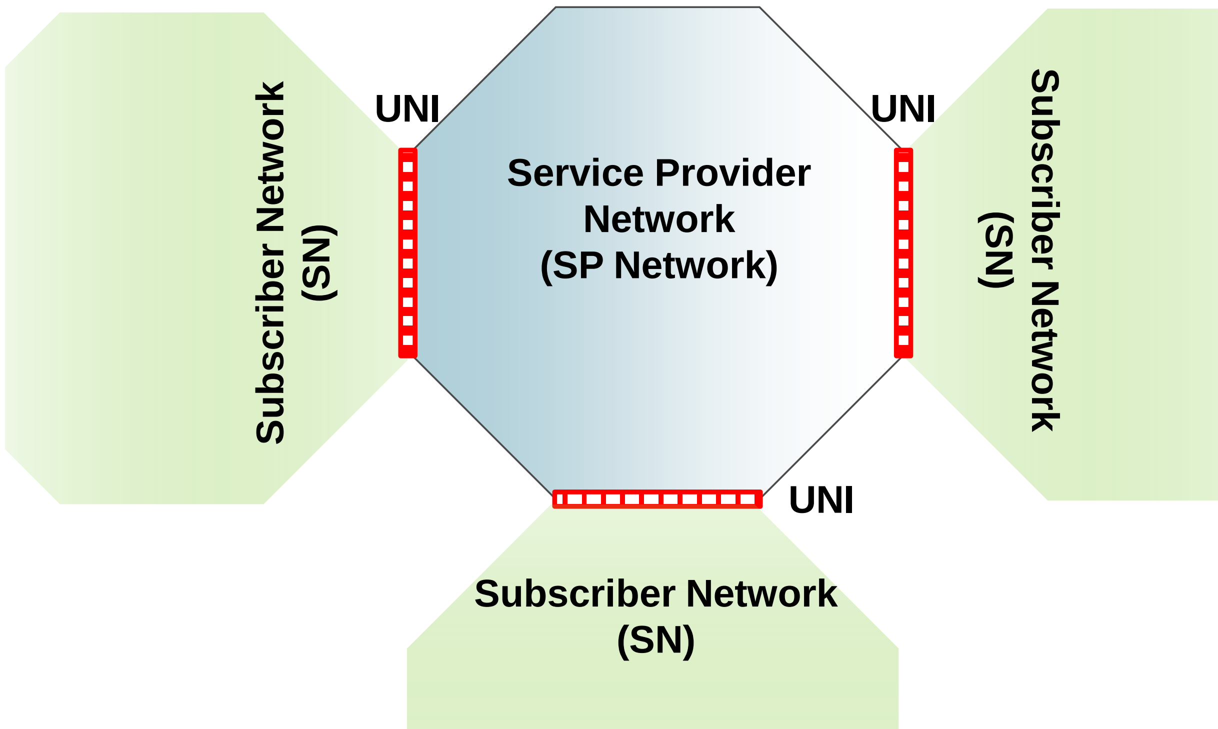

Figure 1 – Fundamental Model

This document is based on the two following requirements.

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

[R2] Subscriber Ethernet Services provided at a UNI **MUST** be provided by a single Service Provider.

Although this document is focused on Subscriber Ethernet Service, [R2] does not preclude the Service Provider from offering other kinds of service at the UNI.

7.3 Service Attributes

Ethernet Services are specified using Service Attributes. A Service Attribute is specific information that is agreed between the provider and the user of the service that describes some aspect of the service behavior or capability. Values for all of the Service Attributes in this document need to be agreed to by the Subscriber and Service Provider for an Ethernet Service.

How the agreement is reached, and the specific values agreed, may have an impact on the price of the service or on other business or commercial aspects of the relationship between the Subscriber and the Service Provider; this is outside the scope of this document. Some examples are given below, but this is not an exhaustive list.

- The Service Provider mandates a particular value.
- The Subscriber selects from a set of options specified by the Service Provider.

- The Subscriber requests a particular value, and the Service Provider accepts the value.
- The Subscriber and Service Provider negotiate to reach a mutually acceptable value.

Service Attributes describe the externally visible behaviors at a UNI and between UNIs; they do not constrain how the service is implemented in the SP Network, or how the Subscriber implements the SN.

There are two types of Service Attributes: Behavioral Service Attributes and Capability Service Attributes.

Behavioral Service Attributes directly affect the behavior of the service as experienced by the Subscriber. As soon as the Service Provider has configured a particular value, the Subscriber can test this to ensure the observed behavior is as expected, for example by sending appropriate traffic over the service.

In contrast, Capability Service Attributes do not directly affect the behavior of the service; instead, they serve as hints to the Service Provider as to what changes to the service the Subscriber might request in the future, and as hints to the Subscriber as to the likely response to such requests. Particular values of a Capability Service Attribute can constrain the acceptable values for other Service Attributes, but do not directly affect the behavior of the service.

For example, the EVC List of EVC EPs Service Attribute (Section 8.2) is a Behavioral Service Attribute that specifies, in effect, the UNIs that are connected by the service. A Subscriber can test that the correct UNIs are connected by testing to see if data can be sent between each pair of UNIs.

In contrast, the EVC Connection Type Service Attribute (Section 8.3) is a Capability Service Attribute that specifies whether the EVC is a Point-To-Point EVC, Multipoint-To-Multipoint EVC, or Rooted-Multipoint EVC. The value of this attribute could constrain the value of other Service Attributes – for example, it could limit the number of EVC EPs that can be present in the EVC List of EVC EPs Service Attribute (Section 8.2), and hence the number of UNIs that can be connected together by the service. However, it does not have any direct bearing on the behavior of the service. For example, if the EVC List of EVC EPs contains exactly two EVC EPs (i.e., the service connects two UNIs together), it is not possible to test whether the EVC Connection Type would allow additional UNIs to be added by sending data between the existing UNIs.

Capability Service Attributes can be used by the Service Provider to guide how they implement the service. This might affect the ease with which they can support future changes to the service requested by the Subscriber. For example, if the Subscriber has a service connecting two UNIs and the service is multipoint-to-multipoint, this might lead the Service Provider to implement the service using a technology that can support multipoint. The Subscriber could add additional UNIs by agreeing a new value for the EVC List of EVC EPs Service Attribute (Section 8.2). On the other hand, if the Subscriber has a service with two UNIs and the service is point-to-point, this might lead the Service Provider to implement the service using a technology that can only support point-to-point. The Subscriber could add additional UNIs by agreeing a change to the

EVC Connection Type to multipoint-to-multipoint, and a further change to the EVC List of EVC EPs Service Attribute. However, the Service Provider might have pricing policies that make such a change more expensive than in the former case, as they might need to change how the service is implemented. For the same reason, the change might have a greater operational impact on the existing UNIs – for example, requiring them to be taken out of service while the change is enacted.

As can be seen from the above example, Capability Service Attributes can also be used by the Subscriber to guide how future requests for changes might be treated by the Service Provider.

It should be stressed that this document does not constrain which Service Attribute values can be changed or when, or what impact particular changes might have on the cost of the service or on other terms and conditions. The commercial arrangements between the Service Provider and the Subscriber are outside the scope of this document.

Note that Capability Service Attributes are not binding – that is, they do not commit the Service Provider to accepting or rejecting particular changes requested by the Subscriber. They serve only as hints about what changes the Subscriber might request in future, and the likely response to such requests.

In this document, Capability Service Attributes are identified as such. Service Attributes not identified as Capability Service Attributes are Behavioral Service Attributes.

7.4 "Support" in Normative Language

When the term "support" is used in a normative context in this document and the normative language applies to the Service Provider, it means that the Service Provider must/should/may be capable of meeting the requirement upon agreement between the Subscriber and Service Provider.

7.5 Service Frames

Ethernet Frames exchanged between the SN and the SP Network across the UNI are referred to as Service Frames. A Service Frame transmitted across the UNI toward the SP Network is called an Ingress Service Frame. A Service Frame transmitted across the UNI toward the SN is called an Egress Service Frame. Service Frames are exchanged via a standard Physical Layer (Section 9.4) or by other means from which an IEEE Std 802.3 – 2015 [5] Ethernet frame can be inferred.

The Service Frame is defined as the IEEE Std 802.3 – 2015 [5] Frame as shown in Figure 2. It consists of the first bit of the Destination MAC Address through the last bit of the Frame Check Sequence. The allowed formats for the Service Frame are detailed in Section 9.7.

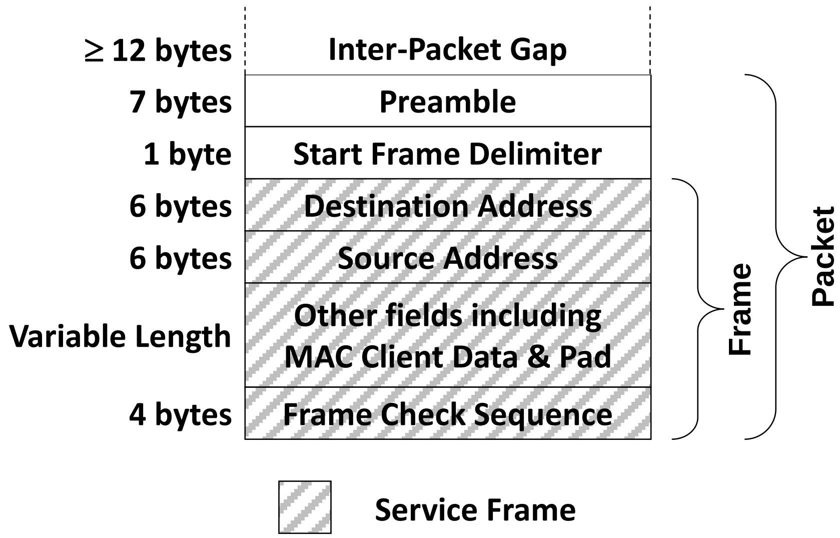

Figure 2 – IEEE Std 802.3 – 2015 Packet Format (Extension Field not shown)

Throughout this document:

- The Service Frame Arrival Time is defined as the time at which the first bit of the Service Frame arrives at the UNI.
- When the field following the Source Address field is a TPID[1] with the value 0x8100 and the corresponding VID 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 TPID[1] with the value 0x8100 and a corresponding VID value of 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 C-Tagged Service Frame.
- When the field following the Source Address field is a TPID[1] 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.
- The C-Tag VLAN ID is defined as the VID field of a C-Tagged Service Frame as defined in IEEE Std 802.1Q – 2018 [4].
- The C-Tag PCP is defined as the PCP field of a C-Tagged Service Frame as defined in IEEE Std 802.1Q – 2018 [4].

[1] TPID, Tag Protocol Identifier, is defined in IEEE Std 802.1Q – 2018 [4].

- The C-Tag DEI is defined as the DEI field of a C-Tagged Service Frame as defined in IEEE Std 802.1Q – 2018 [4].

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 – 2018 [4]) is considered to be 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 the Service Provider to determine the behavior for such Service Frames.

7.6 Types of Service Frame

There are three types of Service Frame as detailed in Sections 7.6.1, 7.6.2, and 7.6.3.

7.6.1 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 SP Networks be able to process such information effectively.[2]

[R3] 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 address.

MAC Destination Addresses[3]	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 and Subscriber **MAY** agree to define additional fields and/or field values for identifying Layer 2 Control Protocol Service Frames in addition to those in Table 3.

Note that [O1] allows Service Frames carrying proprietary protocols to be classified as Layer 2 Control Protocol Service Frames.

[2] This capability will be especially important for Subscribers who choose to deploy IEEE Std 802.1Q – 2018 [4] bridges (as opposed to routers) at UNIs.

[3] Hexadecimal canonical format

Sections 9.15 and 9.17 describe how L2CP Service Frames carrying standardized Layer 2 Control Protocols can be handled by the SP Network. In addition, the Subscriber and Service Provider are free to agree on how proprietary multicast Layer 2 Control Protocols are handled. Such agreements are beyond the scope of this document.

Note that proprietary Layer 2 Control Protocols not identified via agreement per [O1] are treated as Data Service Frames (Section 7.6.3), but could have special delivery behavior agreed using the value of the EVC Service Frame Delivery Service Attribute (Section 8.4).

7.6.2 SOAM Service Frame

A Service Frame, with or without a C-Tag, that is not an L2CP Service Frame per [R3] or [O1] and with Ethertype = 0x8902 is defined to be a SOAM Service Frame. (See Table 21-1 of IEEE Std 802.1Q – 2018 [4].) Section 9.14 and Section 10.13 contain requirements for the formatting and handling of certain SOAM Service Frames. Other requirements for formatting and handling SOAM Service Frames can be found in MEF 17 [28], MEF 30.1 [32], MEF 30.1.1 [33], and MEF 35.1 [34].

7.6.3 Data Service Frame

A Data Service Frame is defined as a Service Frame that meets one of the following three conditions:

1. It is not a Layer 2 Control Protocol Service Frame and not a SOAM Service Frame,
2. It is an ingress Layer 2 Control Protocol Service Frame whose proper handling requires it to be passed at the ingress and egress UNIs per the requirements in MEF 45.1 [36], or
3. It is an ingress SOAM Service Frame at or above the value of the EVC Available MEG Level Service Attribute (Section 8.11) and of a type that is passed transparently by any Subscriber MEG MIP.

Note that a Layer 2 Control Protocol Service Frame is also a Data Service Frame if condition 2 is met and a SOAM Service Frame is also a Data Service Frame if condition 3 is met.

A Data Service Frame with a unicast Destination MAC Address is defined to be a Unicast Data Service Frame. A Data Service Frame with a multicast Destination MAC Address is defined to be a Multicast Data Service Frame. A Data Service Frame with a broadcast Destination MAC Address is defined to be a Broadcast Data Service Frame.

7.7 Subscriber Ethernet Service

An Ethernet Service is a connectivity service that carries Ethernet Frames irrespective of the underlying technology and that is specified using Service Attributes as defined in an MEF Specification. A Subscriber Ethernet Service is an Ethernet Service that is provided by a Service Provider to Subscribers between two or more UNIs.

An ingress Data Service Frame at UNI A is said to be delivered to UNI B when the ingress Data Service Frame at UNI A results in an egress Data Service Frame at UNI B with Destination Address through MAC Client Data & Pad derived from the Ingress Service Frame as specified below.

The following requirements mandate which fields of a Data Service Frame cannot be changed from ingress to egress. The requirements on the format of Service Frames are in Section 9.7.

[R4] If the egress Data Service Frame resulting from an ingress C-Tagged Data Service Frame is also a C-Tagged Service Frame then fields in the egress Data Service Frame **MUST** be unchanged as indicated in Figure 3.

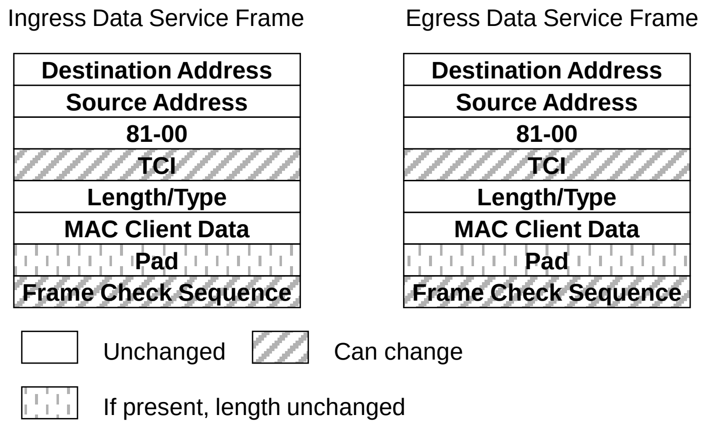

Figure 3 – C-Tagged In to C-Tagged Out Data Service Frame Transparency[4]

[R5] If the egress Data Service Frame resulting from an ingress C-Tagged Data Service Frame is an Untagged Service Frame then fields in the egress Data Service Frame **MUST** be unchanged as indicated in Figure 4.

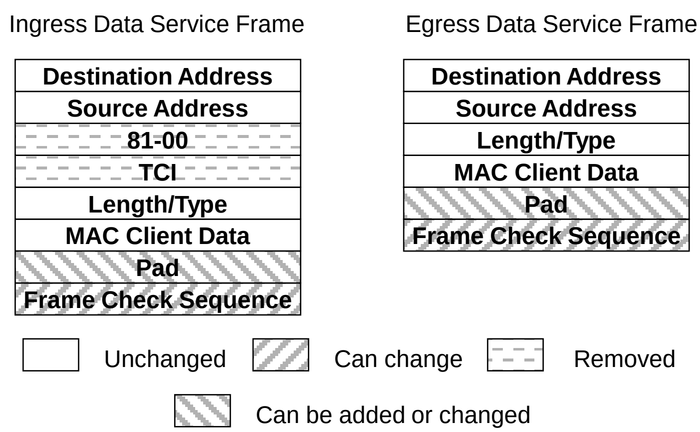

Figure 4 – C-Tagged In to Untagged Out Data Service Frame Transparency[4]

[R6] If the egress Data Service Frame resulting from an ingress Untagged Data Service Frame is a C-Tagged Data Service Frame then fields in the egress Data Service Frame **MUST** be unchanged as indicated in Figure 5.

[4] Clause 4.2.8 of IEEE Std 802.3 – 2015 [5] states that the Pad Field contains arbitrary bits. This document does not constrain the SP Network to preserve these bits.

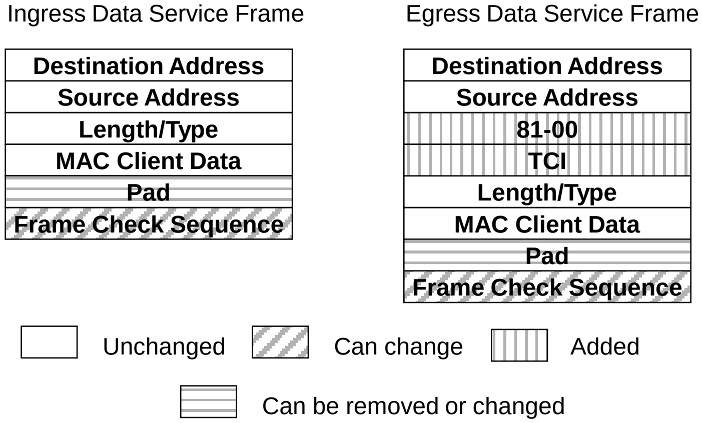

Figure 5 – Untagged In to C-Tagged Out Data Service Frame Transparency[4]

[R7] 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 6.

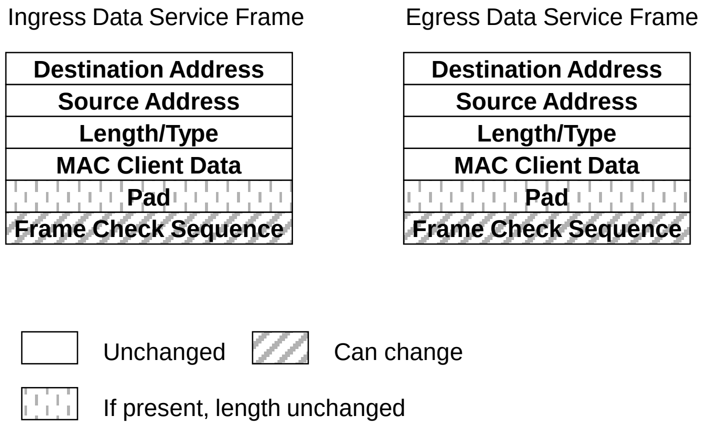

Figure 6 – Untagged In to Untagged Out Data Service Frame Transparency[4]

Under certain conditions, additional fields are mandated to be identical at ingress and egress. See [R118] and Sections 8.5 and 8.6.

7.8 Ethernet Virtual Connection, EVC End Point and EVC EP Map Service Attribute

A fundamental aspect of Subscriber Ethernet Services is the Ethernet Virtual Connection (EVC). An EVC is an association of two or more EVC End Points (EVC EPs). An EVC EP is a construct at a UNI that selects a subset of the Service Frames that pass over the UNI. The subset of Service Frames is specified via the value of the EVC EP Map Service Attribute as described in Section 10.4 and, per [R108], is disjoint from the subsets selected by other EVC EPs at the same UNI. An EVC EP represents the logical attachment of an EVC to a UNI. The EVC EPs associated by an EVC are said to be "in the EVC." A given UNI can support more than one EVC EP but each EVC EP has to be in a different EVC.

The EVC EP Map Service Attribute replaces the CE-VLAN ID/EVC Map Service Attribute that is in MEF 10.3 [24]. As a result, the concept of the CE-VLAN ID in MEF 10.3 is not included in this document. In MEF 10.3, it is possible to map Untagged Service Frames, Priority Tagged Service Frames, and VLAN Tagged Service Frames with a specific C-Tag VLAN ID value to an EVC using the CE-VLAN ID/EVC Map Service Attribute. Such a configuration was included in MEF 10.3 and earlier versions to accommodate an SN that attaches to the UNI with a shared media network and a mix of devices that are compatible and not compatible with IEEE Std 802.1Q – 2018 [4]. Such SN configurations are no longer important in the industry. Consequently, the EVC EP Map Service Attribute (Section 10.4) does not support this SN configuration.

An Ingress Service Frame that is mapped to an EVC EP associated by an EVC can be delivered to zero or more UNIs that have EVC EPs in that EVC; and only to UNIs other than the ingress UNI.

For ease of discourse, "frame is delivered to an EVC EP" means "frame is an Egress Service Frame at the UNI where the EVC EP is located that is mapped to that EVC EP via the value of the EVC EP Map Service Attribute (Section 10.4)". And "UNIs that are in the EVC" means the set of UNIs such that each UNI has an EVC EP that is associated by the EVC.

[R8] If an Egress Service Frame mapped to an EVC EP results from an Ingress Service Frame mapped to an EVC EP then the following two conditions **MUST** hold:

- There is an EVC that associates the two EVC EPs
- The two EVC EPs are different from each other

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

In the context of this document, a Subscriber Ethernet Service consists of a single EVC, associated UNIs and EVC EPs, that is provided to a Subscriber by a Service Provider.

7.9 Service Level Specification

The Subscriber and Service Provider can agree to a Service Level Agreement (SLA) that is a contract specifying the service level commitments and related business agreements for a service.

This document describes the Service Level Specification (SLS) which contains the technical details of the service level, in terms of Performance Objectives, agreed between the Service Provider and the Subscriber as part of the SLA. The business agreements in an SLA, e.g., penalties for failure to meet the SLS, are beyond the scope of this document.

The SLS deals with what the Subscriber experiences. It does not address measurement techniques to estimate the Subscriber experience. Methods for the Service Provider and the Subscriber to monitor the EVC performance to estimate this user experience are beyond the scope of this document. Methods for monitoring performance can be found in MEF 35.1 [34].

7.10 Bandwidth Profiles

A Bandwidth Profile is a specification of the temporal properties of a sequence of Service Frames at a UNI, along with rules for determining the level of conformance to the specification for each Service Frame in the sequence. The details of how a Bandwidth Profile is specified and the details of the rules for determining the level of conformance are in Section 12.

A Bandwidth Profile that applies to Ingress Service Frames is called an Ingress Bandwidth Profile. A Bandwidth Profile that applies to Egress Service Frames is called an Egress Bandwidth Profile. The purpose of each type of Bandwidth Profile is discussed in the following subsections.

7.10.1 Use of Ingress Bandwidth Profiles

The available UNI MAC Data Rates for physical Ethernet interfaces are few in number, e.g., 10 Mbps, 100 Mbps, 1 Gbps, 10 Gbps. However the Subscriber's bandwidth need for an EVC may not be close to one of these MAC Data Rates. Thus it is desirable to enable the Subscriber and Service Provider to agree on EVCs that use less than a UNI MAC Data Rate to

- Reduce cost by consuming less SP Network bandwidth than the UNI MAC Data Rate and/or
- Allow multiple EVCs to share a UNI.

In addition, if an EVC supports multiple classes of service, it is desirable to control how the available bandwidth is apportioned between the different classes.

An Ingress Bandwidth Profile is the vehicle for the agreement between the Subscriber and Service Provider regarding limits on the amount of traffic contained in Ingress Service Frames. It can be specified using the EVC EP Ingress Bandwidth Profile Service Attribute (Section 10.8) or the EVC EP Class of Service Name Ingress Bandwidth Profile Service Attribute (Section 10.9). The agreement is executed by using the level of conformance to the Ingress Bandwidth Profile for each Ingress Service Frame as the basis for the Service Provider to deliver the frame. There are three possible levels of conformance: Green, Yellow, and Red. The action for each conformance level is:

- Green: Deliver the Ingress Service Frame with performance per the EVC Service Level Specification Service Attribute (Section 8.8),
- Yellow: Deliver the Ingress Service Frame on a "Best Effort" basis, i.e., the EVC Service Level Specification Service Attribute (Section 8.8) does not apply, and
- Red: Discard the Ingress Service Frame.

7.10.2 Use of Egress Bandwidth Profiles

An Egress Bandwidth Profile is the vehicle for the agreement between the Subscriber and Service Provider regarding possible limits on the amount of data from an EVC that can egress at a UNI. It can be specified using the EVC EP Egress Bandwidth Profile Service Attribute (Section

10.10) or the EVC EP Class of Service Name Egress Bandwidth Profile Service Attribute (Section 10.11).

Such an agreement can be important for an EVC that associates more than two EVC EPs. With such an EVC, there can be periods of focused load, i.e., traffic is flowing from several UNIs to one target UNI. A focused load can overwhelm SP Network resources on the paths toward the target UNI and/or cause an EVC to crowd out other EVCs on the target UNI. This leads to a need to limit the amount of egress traffic at a UNI.

How the Service Provider accomplishes the limit on the amount of egress traffic is typically some combination of delay and discard within the SP Network. Details of such a SP Network implementation are beyond the scope of this document.

8. Ethernet Virtual Connection Service Attributes

This section contains Service Attributes that apply to an EVC (Section 7.8). Table 33 contains a list of the Ethernet Virtual Connection Service Attributes and their possible values.

8.1 EVC ID Service Attribute

The value of the EVC ID Service Attribute is a string that is used to identify an EVC within the SP Network.

[R9] The EVC ID **MUST** be unique across all EVCs in the SP Network.

[R10] The EVC ID **MUST** contain no more than 45 characters.

[R11] The EVC ID **MUST** be a non-null RFC 2579 [15] DisplayString but not contain the characters 0x00 through 0x1f.

The value of the EVC ID Service Attribute is intended for joint Subscriber/Service Provider management and control purposes.[5] As an example, the Acme Service Provider might use "EVC-0001898-ACME-MEGAMART" to represent the 1898th EVC in the SP Network with the customer for the EVC being MegaMart.

8.2 EVC List of EVC EPs Service Attribute[6]

The value of the EVC List of EVC EPs Service Attribute is a list of EVC EP ID Service Attribute (Section 10.1) values. The list contains one EVC EP ID Service Attribute value for each EVC EP associated by the EVC.

[R12] The value of the EVC List of EVC EPs Service Attribute **MUST** have at least two entries.

[R13] An EVC **MUST** associate at most one EVC EP at a given UNI.

[R13] means that an Ingress Service Frame at a UNI that is mapped to an EVC EP cannot result in an Egress Service Frame at that UNI.

[R14] A given EVC EP ID Service Attribute value (Section 10.1) **MUST NOT** be included in the value of the EVC List of EVC EPs Service Attribute for more than one EVC.

8.3 EVC Type Service Attribute

The value of the EVC Type Service Attribute is one of *Point-to-Point*, *Multipoint-to-Multipoint*, or *Rooted-Multipoint*. The value of the EVC Type Service Attribute indicates the number of EVC EPs that can be in the EVC and which EVC EPs can communicate with each other. This is

[5] For example, see Section 8.1 of MEF 30.1. [32]
[6] This Service Attribute replaces the UNI List Service Attribute in MEF 10.3 [24].

a Capability Service Attribute (Section 7.3). The behaviors for each of these values are as described in Sections 8.3.1, 8.3.2 and 8.3.3.

8.3.1 Point-to-Point EVC

A Point-to-Point EVC is defined as an EVC that can only associate exactly two Root EVC EPs. (See Section 10.3 for the definition of a Root EVC EP.)

[R15] When the value of the EVC Type Service Attribute = *Point-to-Point*, the value of the EVC List of EVC EPs Service Attribute (Section 8.2) **MUST** consist of exactly two Root EVC EPs (Section 10.3).

Figure 7 illustrates two Point-to-Point EVCs.

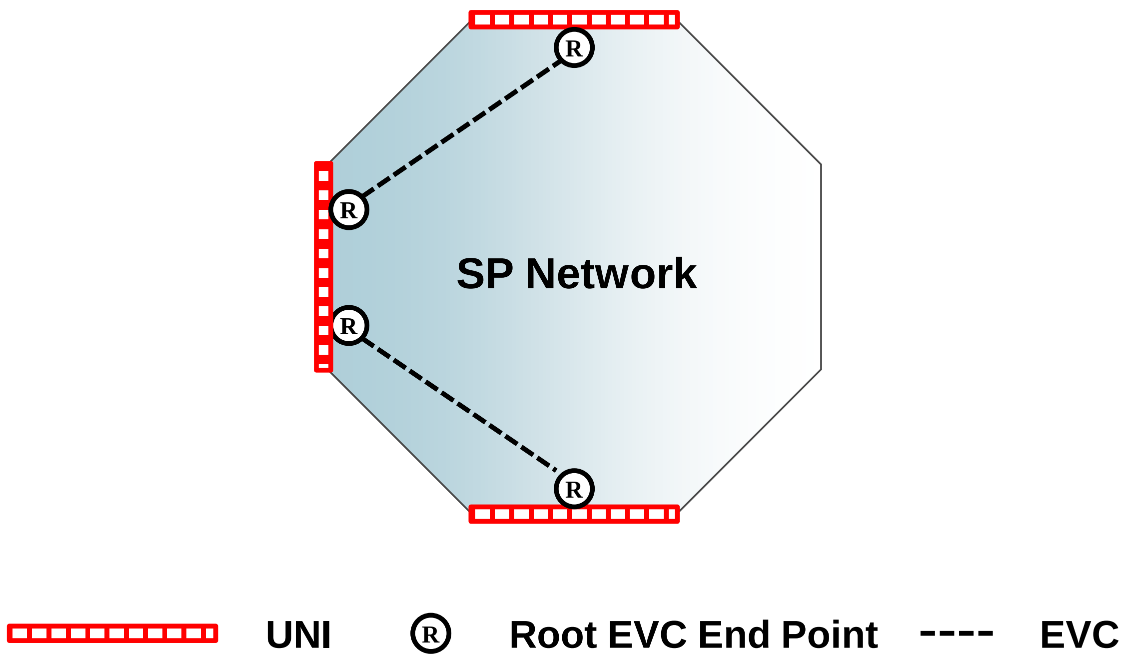

Figure 7 – Point-to-Point EVCs

8.3.2 Multipoint-to-Multipoint EVC

A Multipoint-to-Multipoint EVC is defined as an EVC that can associate more than two EVC EPs and can only associate Root EVC EPs. (See Section 10.3 for the definition of a Root EVC EP.) A Multipoint-to-Multipoint EVC that associates two Root EVC EPs is different from a Point-to-Point EVC because additional Root EVC EPs can be added to the value of EVC List of EVC EPs Service Attribute (Section 8.2) without changing the value of the EVC Type Service Attribute.

[R16] When the value of the EVC Type Service Attribute = *Multipoint-to-Multipoint*, the value of the EVC List of EVC EPs Service Attribute (Section 8.2) **MUST** contain only Root EVC EPs (Section 10.3).

[R17] When the value of the EVC Type Service Attribute = *Multipoint-to-Multipoint*, the Service Provider **MUST** support at least three entries in the value of the EVC List of EVC EPs Service Attribute (Section 8.2).

In a Multipoint-to-Multipoint EVC, the rules under which a frame is delivered to an EVC EP in the EVC are specific to the particular service definition. Typically, a single Ingress Service Frame with a broadcast or multicast MAC Destination Address at a given EVC EP would be replicated in the SP Network and a single copy would be delivered to each of the other EVC EPs in the EVC. This kind of delivery would also typically apply to a Service Frame for which the SP Network has not yet learned an association of the destination MAC address with an EVC EP. Figure 8 illustrates a Multipoint-to-Multipoint EVC.

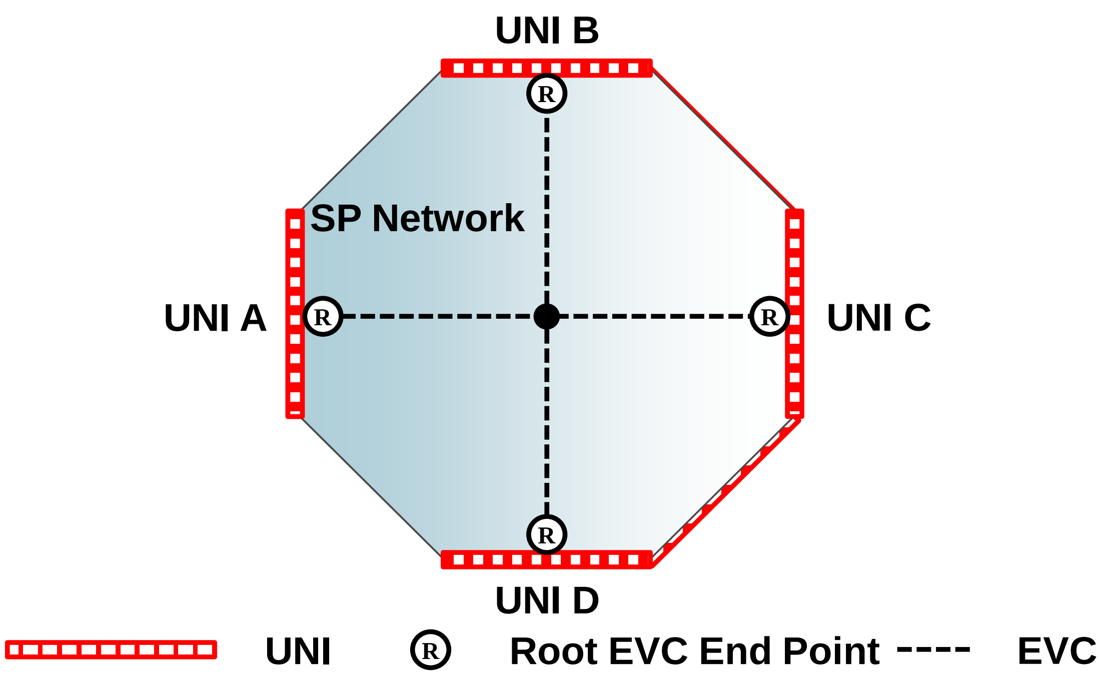

Figure 8 – Multipoint-to-Multipoint EVC

8.3.3 Rooted-Multipoint EVC

A Rooted-Multipoint EVC is defined as an EVC that can associate more than two EVC EPs, at least one of which is a Root EVC EP and any number of which are Leaf EVC EPs. (See Section 10.3 for the definitions of a Root EVC EP and Leaf EVC EP.) A Rooted-Multipoint EVC that associates only Root EVC EPs is different from a Multipoint-to-Multipoint EVC and a Point-to-Point EVC because Leaf EVC EPs can be added to the value of EVC EP List Service Attribute (Section 8.2) without changing the value of the EVC Type Service Attribute.

[R18] When the value of the EVC Type Service Attribute = *Rooted-Multipoint*, the Service Provider **MUST** support at least three entries in the value of the EVC EP List Service Attribute (Section 8.2).

[R19] When the value of the EVC Type Service Attribute = *Rooted-Multipoint*, the Service Provider **MUST** support a value of the EVC EP List Service Attribute

(Section 8.2) that contains at least one Root EVC EP (Section 10.3) and at least one Leaf EVC EP (Section 10.3).

Per [R103], an Ingress Service Frame mapped to a Leaf EVC EP at a UNI can only be delivered to Root EVC EPs in the EVC.

The rules under which a frame is delivered to a UNI in the EVC are specific to the particular service definition. Typically, a single Ingress Service Frame with a broadcast or multicast MAC Destination Address mapped to a Root EVC EP would be replicated in the SP Network and a single copy would be delivered to each of the other UNIs where there is an EVC EP that is in the EVC. This kind of delivery would also typically apply to a Service Frame for which the SP Network has not yet learned an association of the destination MAC address with an EVC EP. Figure 9 illustrates a Rooted-Multipoint EVC with one Root EVC EP.

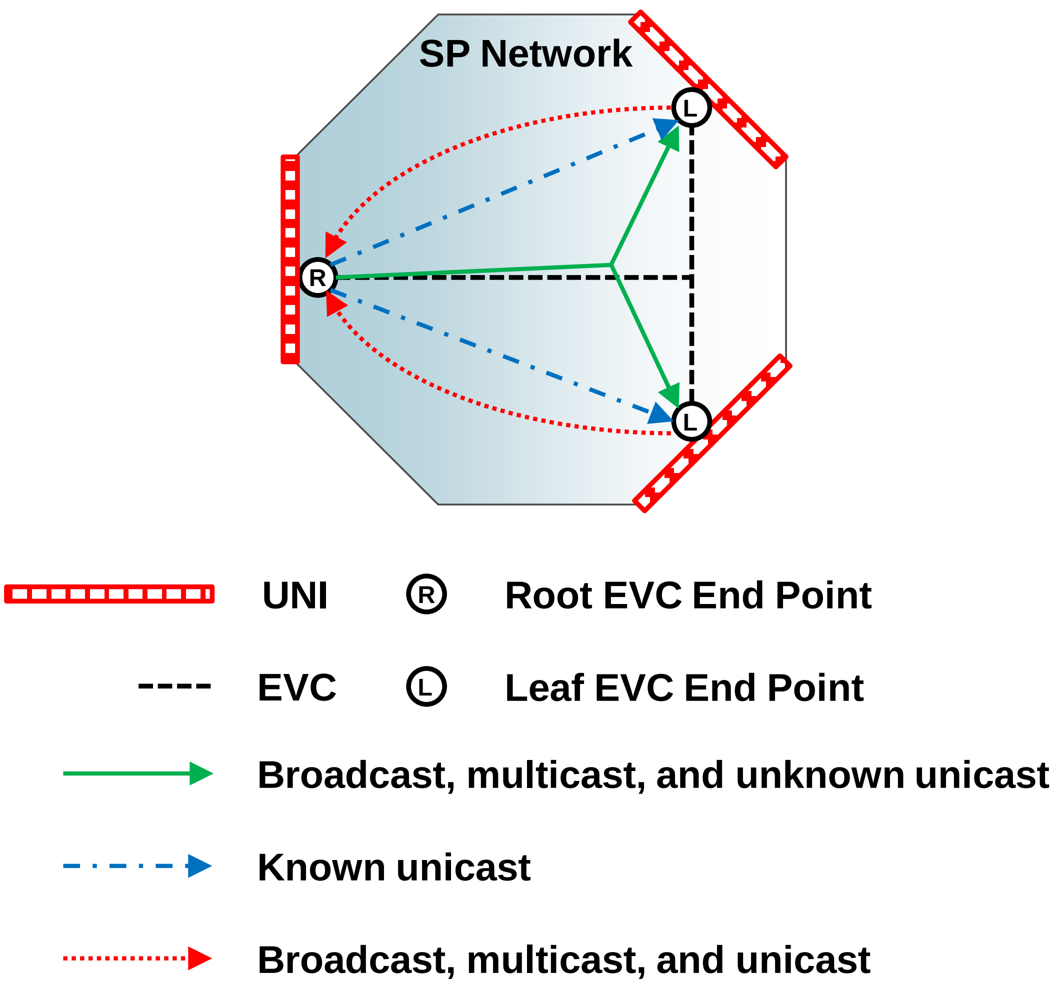

Figure 9 – Rooted-Multipoint EVC Example

8.4 EVC Data Service Frame Disposition Service Attribute

The EVC Data Service Frame Disposition Service Attribute indicates whether different types of Service Frame are to be delivered by the EVC. The value of the EVC Data Service Frame Disposition Service Attribute is a 3-tuple of the form $\langle u, m, b \rangle$ where each element in the 3-tuple has the value *Discard*, *Deliver Unconditionally*, or *Deliver Conditionally*. The value of u in the 3-tuple applies to ingress Unicast Data Service Frames. The value m in the 3-tuple applies to ingress Multicast Data Service Frames. The value of b in the 3-tuple applies to ingress Broadcast Data Service Frames. The EVC Data Service Frame Disposition Service Attribute applies to any ingress Data Service Frame that is not discarded per [R82], [R103], [R119], [R141], [R145], [R149], [D1], [D2], [D6], or [O4]. The dispositions corresponding to each element value are:

- *Discard*: The Data Service Frame is discarded.
- *Deliver Unconditionally*: The Data Service Frame is delivered to all EVC EPs other than the ingress EVC EP..
- *Deliver Conditionally*: The Data Service Frame is delivered to some, none, or all of the EVC EPs other than the ingress EVC EP depending on certain specified conditions.

Examples of *Deliver Conditionally* include:

- MAC Address learning where the destination MAC Address is known by the SP Network to be at a given EVC EP,
- Broadcast throttling where some Broadcast Data Service Frames are dropped in order to limit the amount of such traffic,
- Multicast pruning where some Multicast Data Service Frames are dropped at places where there are known to be no downstream listeners, and
- MAC Address filtering where Data Service Frames with certain MAC Addresses are discarded.

[R20] When an element in the 3-tuple = *Deliver Conditionally*, the conditions that determine to which other EVC EPs in the EVC, if any, an ingress Data Service Frame at a given EVC EP is delivered **MUST** be agreed to by the Subscriber and Service Provider.

Note that the disposition of ingress Unicast Data Service Frames can be different from the disposition of ingress Multicast Data Service Frames which in turn can be different from the disposition of ingress Broadcast Data Service Frames.

Recall that, per Section 7.2, this document makes no assumptions about the implementation of the SP Network. Thus the discarding of a Service Frame when the disposition value is *Discard* or *Deliver Conditionally* can occur anywhere within the SP Network.

Note that this is a description of the ideal service. Data Service Frames that are intended to be delivered might be discarded due to network failure or congestion conditions. See Section 8.8 and Section 11.

8.5 EVC C-Tag PCP Preservation Service Attribute[7]

The EVC C-Tag PCP Preservation Service Attribute can be used to preserve the value of the PCP field in C-Tagged Service Frames across an EVC. The value of the EVC C-Tag PCP Preservation Service Attribute is either *Enabled* or *Disabled*.

[R21] In an EVC with the value of the EVC C-Tag PCP Preservation Service Attribute = *Enabled*, if an ingress C-Tagged Service Frame that is mapped to an EVC EP in the EVC results in an egress C-Tagged Service Frame, the C-Tag PCP value in the egress C-Tagged Service Frame **MUST** be equal to the C-Tag PCP value in the ingress C-Tagged Service Frame.

When the value of the EVC C-Tag PCP Preservation Service Attribute = *Disabled*, the value of the C-Tag PCP can change or stay the same across the EVC. See the EVC EP Egress Map Service Attribute (Section 10.7).

8.6 EVC C-Tag DEI Preservation Service Attribute

The EVC C-Tag DEI Preservation Service Attribute can be used to preserve the value of the DEI field in C-Tagged Service Frames across an EVC. The value of the EVC C-Tag DEI Preservation Service Attribute is either *Enabled* or *Disabled*.

[R22] In an EVC with the value of the EVC C-Tag DEI Preservation Service Attribute = *Enabled*, if an ingress C-Tagged Service Frame that is mapped to an EVC EP in the EVC results in an egress C-Tagged Service Frame, the C-Tag DEI value in the egress C-Tagged Service Frame **MUST** be equal to the C-Tag DEI value in the ingress C-Tagged Service Frame.

When the value of the EVC C-Tag DEI Preservation Service Attribute = *Disabled*, the value of the C-Tag DEI can change or stay the same across the EVC. See the EVC EP Egress Map Service Attribute (Section 10.7).

8.7 EVC List of Class of Service Names Service Attribute

The EVC List of Class of Service Names Service Attribute is used to specify all of the Class of Service Names supported by an EVC. The value of the EVC List of Class of Service Names Service Attribute is a non-empty list of Class of Service Names (which may include one or more of the CoS Labels defined in MEF 23.2 [29]). Each Ingress Service Frame is assigned a Class of Service Name from this list via the EVC EP Ingress Class of Service Map Service Attribute as described in Section 10.5. The Class of Service Name that is assigned to a frame indicates the Performance Objectives that apply to the frame under appropriate conditions as detailed in Section 8.8, and is used to determine how to set certain fields in an Egress Service Frame that results

[7] In MEF 10.3 [24] this Service Attribute is called CE-VLAN CoS Preservation Service Attribute.

from this Ingress Service Frame, as described by the EVC EP Egress Map Service Attribute (Section 10.7). It can also be used to indicate a Bandwidth Profile Flow (Section 12).

Discard is a reserved Class of Service Name that can be included in the value of the EVC List of Class of Service Names Service Attribute. An Ingress Service Frame assigned the Class of Service Name *Discard*, is mandated to be discarded per [R119].

Table 4 shows an example of the value of the EVC List of Class of Service Names Service Attribute.

Platinum
Gold
Silver
Discard

Table 4 – Example Value for the EVC List of Class of Service Names Service Attribute

Note that the value of the EVC EP Ingress Class of Service Map Service Attribute (Section 10.5) can be such that the Class of Service Names for Service Frames mapped to an EVC EP includes only a strict subset of the entries in the value of the EVC List of Class of Service Names Service Attribute for the EVC that associates the EVC EP. In other words, not all entries in the value of the EVC List of Class of Service Names Service Attribute for an EVC need to be available at every EVC EP associated by the EVC.

8.8 EVC Service Level Specification Service Attribute

The EVC Service Level Specification Service Attribute (SLS) is the technical details of the service level, in terms of Performance Objectives, agreed between the Service Provider and the Subscriber as part of the Service Level Agreement. A given SLS might contain 0, 1, or more Performance Objectives for each Performance Metric.

Throughout this document, numeric ranges are denoted as follows: z in $[x, y]$ means $x \leq z \leq y$, z in $[x, y)$ means $x \leq z < y$, z in $(x, y]$ means $x < z \leq y$, and z in (x, y) means $x < z < y$.

The value of the EVC Service Level Specification Service Attribute (SLS) is either *None* or a 3-tuple of the form $\langle ts, T, CN \rangle$ where:

- ts is a time that represents the date and time for the start of the SLS.[8]
- T is a time duration, e.g., 1 month, 2 weeks, that is used in conjunction with ts to specify time intervals for determining when Performance Objectives are met. Note that the units for T are not constrained; in particular, 1 month is an allowable value for T, corresponding to a calendar month, e.g. from midnight on the 10th of one month up to but not including midnight the 10th of the following month.
- CN is a non-empty list of 5-tuples of the form $\langle CoS\_Name, \Delta t, C, n, PM \rangle$ where

[8] Note that ts is the start of the SLS and might not be the time that a service is first turned up.

– CoS_Name is a Class of Service Name contained in the value of the EVC List of Class of Service Names Service Attribute (Section 8.7) and is not *Discard*.

– Δt is a time duration much smaller than T, e.g., 10 seconds.

– C is a real number in the range $[0,1]$ used as a threshold to determine whether a given time interval Δt_k (Section 8.8.1.3) has high loss.

– n is an integer ≥ 1, used to identify how many consecutive Δt_k intervals must have high loss to trigger a change in Availability.

– PM is a non-empty list where each element in the list consists of a Performance Metric Name, a list of parameter values specific to the definition of the Performance Metric, and a Performance Objective.

[R23] A Class of Service Name **MUST** appear at most once in the value of CN.

Figure 10 shows an example of the value for the EVC Service Level Specification Service Attribute. Some observations:

- Performance Metrics and Objectives are specified for the *Gold* and *Silver* Class of Service Names for the EVC. If *Straw* is also in the value of the EVC List of Class of Service Names Service Attribute, this SLS example does not specify Performance Objectives for *Straw*.
- A single value of ts and a single value of T apply to all Class of Service Names.
- Different values of $\Delta t, C,$ and n apply to *Gold* and *Silver*.
- Two instances of the One-way Delay Performance Metric apply to *Gold*. There is a different set of ordered EVC EP pairs (Section 8.8.1.1), a different percentile (Section 8.8.2), and a different objective for each instance. This could be due to EVC EP *c* being geographically distant from EVC EPs *a* and *b*.
- The One-way Frame Delay Range Performance Metric and the One-way Availability Performance Metric along with corresponding objectives apply to *Gold*.
- The One-way Availability Performance Metric applies to *Silver*.

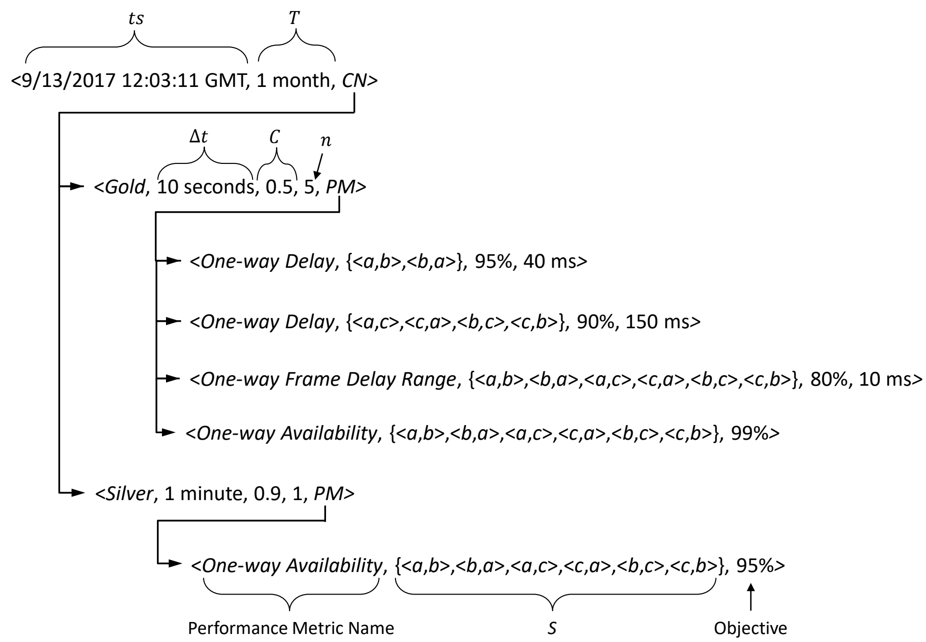

Figure 10 – EVC Service Level Specification Service Attribute Value Example

A Performance Metric is a quantitative characterization of Service Frame delivery quality experienced by the Subscriber. This section specifies the following Performance Metrics all of which apply to a single EVC:

- One-way Frame Delay Performance Metric (Section 8.8.2),
- One-way Mean Frame Delay Performance Metric (Section 8.8.3),
- One-way Frame Delay Range Performance Metric (Section 8.8.4),
- One-way Inter-Frame Delay Variation Performance Metric (Section 8.8.5),
- One-way Frame Loss Ratio Performance Metric (Section 8.8.6),
- One-way Availability Performance Metric (Section 8.8.7),
- One-way High Loss Intervals Performance Metric (Section 8.8.8),
- One-way Consecutive High Loss Intervals Performance Metric (Section 8.8.9),
- One-way Composite Performance Metric (Section 8.8.10), and
- One-way Group Availability Performance Metric (Section 8.8.11).

Section 11 defines the Performance Metric for the Multiple EVC Service Level Specification Service Attribute.

Each Performance Metric has an associated Performance Objective whose value reflects the agreed Service Frame delivery quality. For example, for the One-way Mean Frame Delay Performance Metric (Section 8.8.3), a Performance Objective of 20 ms means that the Service Provider has committed to a mean Service Frame delay, as experienced by the Subscriber, of less than or equal to 20 ms.

[R24] If PM contains an entry with a given Performance Metric Name, then the entry **MUST** specify the related parameters and the Performance Objective for that Performance Metric as specified in the following sections.

These Performance Metrics describe the performance experienced by the Subscriber who is the user of the EVC. Methods for the Service Provider and the Subscriber to monitor the EVC performance to estimate this user experience are beyond the scope of this document. Methods for monitoring performance can be found in MEF 35.1 [34].

[R25] If PM contains multiple entries with a given Performance Metric Name, then for any two such entries, at least one of the parameter values **MUST** be different.

PM can contain multiple entries with a given Performance Metric Name, but as per [R25], one or more of the parameter values associated with each objective for a given Performance Metric Name need to be different from each other. For example, PM could contain two objectives for the One-way Frame Delay Performance Metric each corresponding to a different value of the percentile Pd (see Section 8.8.2).

8.8.1 Key SLS Definitions, Concepts, and Notation

This section describes definitions, concepts, and notation used throughout Section 8.8.

8.8.1.1 Parameters S and CoS_Name

The parameter S is used in most of the Performance Metric definitions. S is a non-empty set of ordered pairs of EVC EP Identifiers taken from the value of the EVC List of EVC EPs Service Attribute (Section 8.2). S is represented by

$$\{\langle i,j\rangle |\ i \text{ and } j \text{ are in the value of the EVC List of EPs Service Attribute and } i \neq j\}$$

Note that for a given entry in PM, S can contain one, some, or all of the possible ordered EVC EP pairs in the EVC.

[R26] For each ordered EVC EP pair $\langle x,y\rangle \in S$, at least one of the EVC EPs **MUST** be a Root EVC EP (Section 10.3).

Consider $CN = \langle Gold, 10\ seconds, 0.5, 4, PM\rangle$, Figure 11 shows an example value for PM in this 5-tuple.

$$PM = \begin{Bmatrix} \text{One-way Frame Loss Ratio}, \{\langle x,y\rangle\}, 0.02\% \\ \text{One-way Frame Loss Ratio}, \{\langle y,x\rangle\}, 0.01\% \end{Bmatrix}$$

Figure 11 – Example of a Value for *PM*

In this example, the first item in each element is the Performance Metric Name, i.e., One-way Frame Loss Ratio (Section 8.8.6). The second item is S. The third item is the Performance Objective. In this example, S varies from entry to entry. Note that the Class of Service Name for this PM is *Gold* per the value of CN.

Figure 12 shows an example value for PM in the case of $CN = \langle Silver, 10\ seconds, 0.5,4, PM\rangle$. Note that the objective (0.05%) for *Silver* is different from that for *Gold* (0.01%) in Figure 11.

$$PM = \{\text{One-way Frame Loss Ratio}, \{\langle x,y\rangle, \langle y,x\rangle\}, 0.05\%\}$$

Figure 12 – Another Example of a Value for *PM*

8.8.1.2 SLS Notation

In Sections 8.8.2 through 8.8.9, Section 8.8.11, and Section 11, numerous mathematical functions are introduced. The following notational conventions are used:

- A subscript always represents an index, e.g., $T_l, l = 0,1,2, \ldots$ (Section 8.8.1.3).
- Dependencies are indicated via the arguments in each function, e.g., $AT(\langle i,j\rangle, T_l)$ means that the value of the function depends on $\langle i,j\rangle$ and T_l (Section 8.8.1.3).
- All functions depend on $CoS\_Name, \Delta t, C,$ and n but, for simplicity, these dependencies are not explicitly indicated in the functional notation.
- Performance Objectives are represented by symbols with a "hat", e.g., $\hat{d}$ (Section 8.8.2).
- A function value that is to be compared to a Performance Objective is indicated by a "bar", e.g., $\bar{d}(S, Pd, T_l)$ (Section 8.8.2).

The above notation is not used in Section 8.8.10 because it would cause the equations and figures to become unwieldy.

8.8.1.3 Time Interval Sequences, Available Time, and Maintenance Interval

For each *CoS_Name* in CN, the time interval sequence $\{\Delta t_k, k = 0,1,2, \ldots\}$ is used where

$$\Delta t_k = [ts + k\Delta t, ts + (k+1)\Delta t)$$

Δt is from the item in CN with *CoS_Name*. Note that Δt can be different for different Class of Service Names.

For the SLS, the sequence $\{T_l, l = 0,1,2, \ldots\}$ is used where

$$T_l = [ts + lT, ts + (l+1)T)$$

Each element of $\{T_l\}$ is used for assessing the success of the EVC in meeting the Performance Objectives of the SLS.

For most of the Performance Metrics in the SLS, the Performance Objective only applies to Ingress Service Frames that arrive at the Service during Available Time. Informally, Available Time is a set of time intervals, contained in some longer time interval, when the service is considered available for use. Formally, Available Time is defined for each *CoS_Name* in CN and for each possible ordered pair of EVC EPs $\langle i,j\rangle, i \neq j$ in the EVC as follows.

For a given ordered EVC EP pair $\langle i,j\rangle$ and a given *CoS_Name* in CN, each Δt_k 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.

Informally, Available Time is based on 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 beginning of the current sequence is defined as available; otherwise, it is defined as unavailable. The formal definition follows. 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.

For a given ordered EVC EP pair $\langle i,j\rangle$, a given *CoS_Name*, and a Δt_k, let $I(\langle i,j\rangle, \Delta t_k)$ be the number of Ingress Service Frames that meet the following conditions:

- The Service Frame arrives at the UNI where EVC EP i is located and its Service Frame Arrival Time is within the time interval Δt_k,
- Each Service Frame maps to EVC EP i via the value of the EVC EP Map Service Attribute (Section 10.4).
- Each Service Frame is intended to be delivered to the EVC EP j according to the value of the EVC Data Service Frame Disposition Service Attribute (Section 8.4),
- Each Service Frame is not discarded per [R82], [R119], [R141], [D1], [D2], [D6], or [O4],
- Each Service Frame has the given *CoS_Name* via the value of the EVC EP Ingress Class of Service Map Service Attribute (Section 10.5),
- Each Service Frame that is subject to an Ingress Bandwidth Profile has an Ingress Bandwidth Profile color declaration equal to Green, and
- Each Ingress Service Frame that is not subject to an Ingress Bandwidth Profile has the Color Green per the value of the EVC EP Color Map Service Attribute (Section 10.6).

Let $E(\langle i,j\rangle, \Delta t_k)$ be the number of unique (not duplicate) Egress Service Frames where each Service Frame is the first, unerrored Egress Service Frame that is mapped to the EVC EP j at the UNI where EVC EP j is located that results from a frame counted in $I(\langle i,j\rangle, \Delta t_k)$. Then

$$flr(\langle i,j\rangle,\Delta t_k)=\begin{cases}\left(\frac{I(\langle i,j\rangle,\Delta t_k)-E(\langle i,j\rangle,\Delta t_k)}{I(\langle i,j\rangle,\Delta t_k)}\right) \text{ if } I(\langle i,j\rangle,\Delta t_k)>0\\ 0 \text{ otherwise}\end{cases}$$

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

$$flr(\langle i,j\rangle,\Delta t_k)=\begin{cases}\left(\frac{\tilde{I}(\langle i,j\rangle,\Delta t_k)-E(\langle i,j\rangle,\Delta t_k)}{\tilde{I}(\langle i,j\rangle,\Delta t_k)}\right) \text{ if } \tilde{I}(\langle i,j\rangle,\Delta t_k)>0\\ 0 \text{ otherwise}\end{cases}$$

where $\tilde{I}(\langle i,j\rangle,\Delta t_k)=I(\langle i,j\rangle,\Delta t_k)-D(\langle i,j\rangle,\Delta t_k)$ and $D(\langle i,j\rangle,\Delta t_k)$ is the number of frames discarded by the Service Provider in order to conform to either the MAC Data Rate of the UNI where EVC EP j is located or an Egress Bandwidth Profile (if one is used) at the UNI where EVC EP j is located. Such frame drops may occur anywhere in the network, not just near the egress UNI. One example of this could be where an Egress Bandwidth Profile is implemented by applying a policer or shaper on a link within the SP Network. Another example of this could be where Green frames for this EVC and Class of Service Name that are intended to be delivered to the UNI with EVC EP j exceed the MAC Data Rate on a link within the network, provided the MAC Data Rate of that link is greater than or equal to the MAC Data Rate of the UNI. Good traffic engineering principles would 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 the UNI where EVC EP j is located. The details of such an adjustment are beyond the scope of this document.

$A(\langle i,j\rangle,\Delta t_k)$ is defined in Figure 13 for $k=0,1,2,\ldots$.

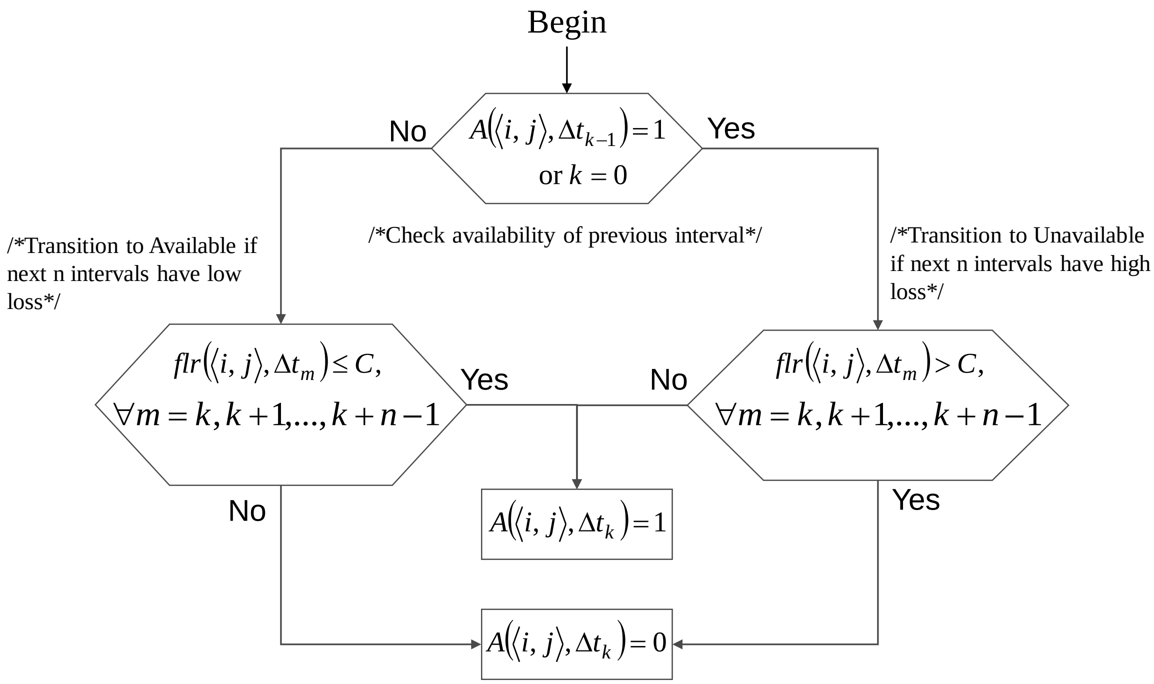

Figure 13 – Flowchart Definition of $A(\langle i,j\rangle,\Delta t_k)$ for a Given Class of Service Name

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

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

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

$$A(\langle i,j\rangle,\Delta t_k)=\begin{cases}0\ if\ A(\langle i,j\rangle,\Delta t_{k-1})=1 \text{ and } flr(\langle i,j\rangle,\Delta t_m)>C \text{ for all } m=k,k+1,\ldots,k+n-1\\ 1\ if\ A(\langle i,j\rangle,\Delta t_{k-1})=0 \text{ and } flr(\langle i,j\rangle,\Delta t_m)\le C \text{ for all } m=k,k+1,\ldots,k+n-1\\ A(\langle i,j\rangle,\Delta t_{k-1}) \text{ otherwise}\end{cases}$$

In the event of a conflict between the above equations and Figure 13, the content of Figure 13 is controlling.

Recall that $\{\Delta t_k, k=0,1,\ldots\}, C,$ and n can be different for different Class of Service Names.

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 define availability. This use of a sliding window is similar to that of ITU-T Y.1563. [20]

Figure 14 presents an example of the determination of the availability for the small time intervals with a sliding window of $n=10$ small time intervals.

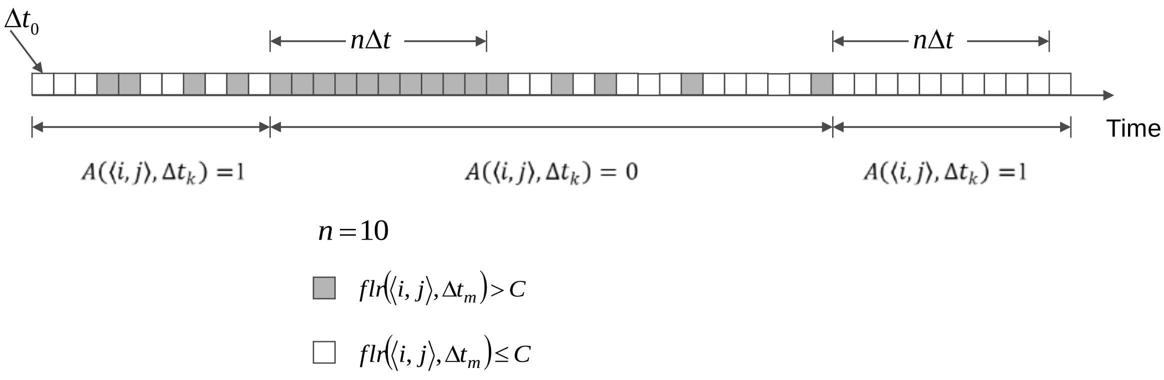

Figure 14 – Example of the Determination of $A(\langle i,j\rangle,\Delta t_k)$ for a Given Class of Service Name

As can be seen below, the definition of Available Time excludes small time intervals that intersect a Maintenance Interval. A Maintenance Interval is a time interval agreed to by the Subscriber and Service Provider during which the EVC may not perform well or at all. Examples of a Maintenance Interval include:

- A time interval during which the Service Provider may disable the EVC for network maintenance such as equipment replacement,
- A time interval during which the Subscriber and Service Provider 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 EVC.

Figure 15 shows an example of a Maintenance Interval.

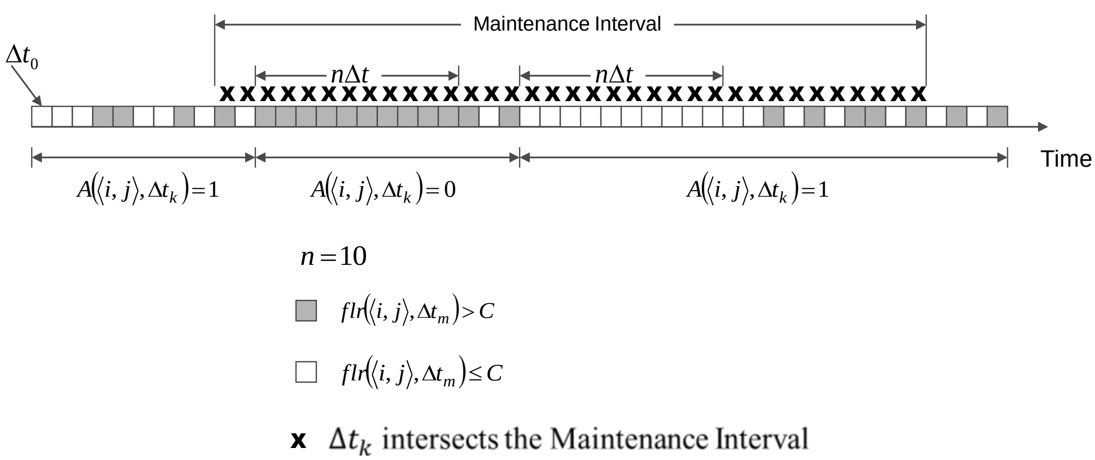

Figure 15 – Example of a Maintenance Interval

Note that the definition of $A(\langle i,j\rangle,\Delta t_k)$ is independent of the presence of any Maintenance Interval.

Let $W(T_l)$ be the set of all Δt_k's contained in T_l that do not intersect a Maintenance Interval.

Available Time for the ordered EVC EP pair $\langle i,j\rangle$, *CoS_Name*, and T_l denoted by $AT(\langle i,j\rangle,T_l)$, is defined as

$$AT(\langle i,j\rangle,T_l) = \{\Delta t_k | \Delta t_k \in W(T_l), A(\langle i,j\rangle,\Delta t_k) = 1\}$$

Unavailable Time, for the ordered EVC EP pair $\langle i,j\rangle$, *CoS_Name*, and T_l denoted by $UT(\langle i,j\rangle,T_l)$, is defined as

$$UT(\langle i,j\rangle,T_l) = \{\Delta t_k | \Delta t_k \in W(T_l), A(\langle i,j\rangle,\Delta t_k) = 0\}$$

Note that a Δt_k that intersects a Maintenance Interval is not included in Available Time and not included in Unavailable Time. Note that the definition of $W(T_l)$ means that the boundaries of T_l and the boundaries of a Maintenance Interval do not have to align with the boundary of a Δt_k.

Figure 16 shows the time classifications for the small time intervals Δt_k.

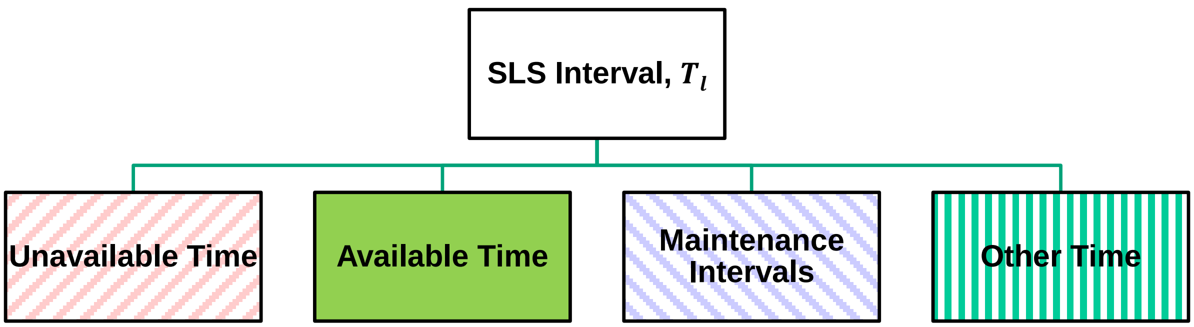

Figure 16 – Time Classifications

In Figure 16, the box labeled Other Time represents the time in Δt_k's with the following properties:

- Δt_k intersects but is not contained within T_l, i.e., it intersects with the start or end of T_l, and
- Δt_k does not intersect a Maintenance Interval.

These small time intervals are not in Unavailable Time, not in Available Time and not in a Maintenance Interval. Other Time arises because the boundaries of T_l do not have to align with a boundary of a small time interval. Figure 17 shows an example of three small time intervals that are in Other Time when there are no Maintenance Intervals. Note that a Δt_k that is in Other Time is not included in the calculation of any Performance Metric.

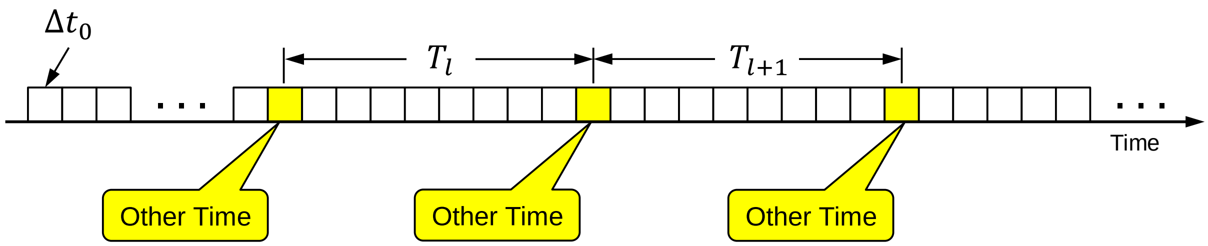

Figure 17 – Example of Small Time Intervals in Other Time

8.8.1.4 Qualified Service Frames

Informally, Qualified Service Frames are defined as a set of Ingress Service Frames that comply with specific criteria and to which most of the Performance Metrics apply.

Qualified Service Frames are defined as every Service Frame that satisfies the following criteria for a given ordered EVC EP pair $\langle i, j \rangle$, a given *CoS_Name*, and a given T_l:

- Each Service Frame ingresses at the UNI where EVC EP i is located,
- Each Service Frame maps to EVC EP i via the value of the EVC EP Map Service Attribute (Section 10.4).
- Each Service Frame is intended to be delivered to the EVC EP j according to the value of the EVC Data Service Frame Disposition Service Attribute (Section 8.4),
- Each Service Frame is not discarded per [R82], [R119], [R141], [D1], [D2], [D6], or [O4],
- Each Service Frame arrives at the ingress UNI within the time interval T_l, and within a time interval Δt_k that is in Available Time, i.e., $\Delta t_k \in AT(\langle i, j \rangle, T_l)$ (Section 8.8.1.3),
- Each Service Frame has the given *CoS_Name* (Section 10.5),
- Each Service Frame that is subject to an Ingress Bandwidth Profile has an Ingress Bandwidth Profile color declaration equal to Green, and
- Each Service Frame that is not subject to an Ingress Bandwidth Profile has the Color Green per the EVC EP Color Map Service Attribute (Section 10.6).

8.8.1.5 One-way Frame Delay

The One-way Frame Delay for a Service Frame that ingresses at UNI_1 and results in a Service Frame that egresses at UNI_2 is defined as the time elapsed from the reception by the SP Network of the first bit of the Ingress Service Frame at the ingress UNI until the transmission by the SP Network of the last bit of the first corresponding Egress Service Frame at the egress UNI. If the Service Frame is erroneously duplicated in the SP Network and multiple copies delivered to

UNI2, the delay is based on the first such copy delivered. The components of the One-way Frame Delay are illustrated in Figure 18. Note that the figure shows $t_1 < t_2$ but $t_1 > t_2$ is also possible.

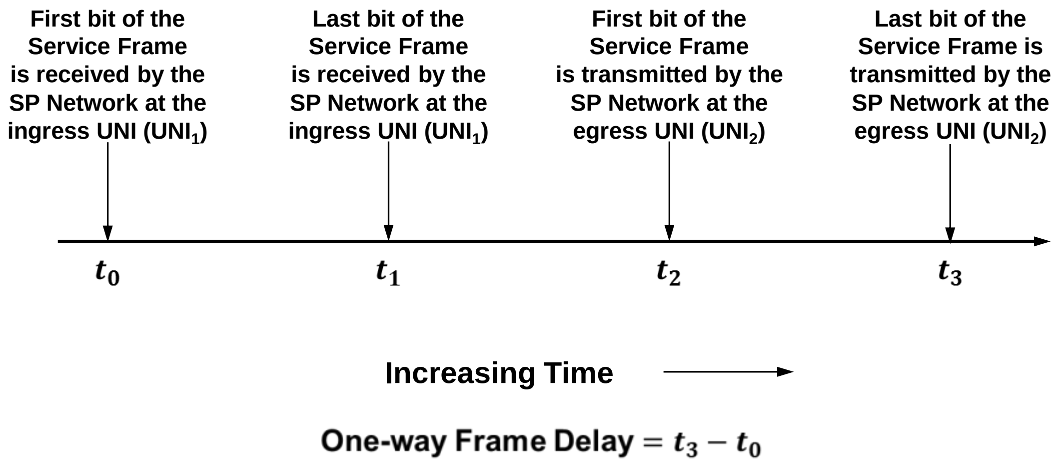

Figure 18 – Components of the One-way Frame Delay for Service Frame

The definition of One-way Frame Delay for a Service Frame is the one-way[9] delay $(t_3 - t_0)$ that includes the delays encountered as a result of transmission of the Service Frame across the ingress UNI $(t_1 - t_0)$ and egress UNI $(t_3 - t_2)$. Note that the One-way Frame Delay does not include delays on the SN side of the UNI, e.g., propagation across a cable attaching an SN router to the UNI is not included.

When the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Physical* for the ingress UNI, $(t_1 - t_0) =$ the length of the Service Frame divided by the MAC Data Rate of the physical link carrying the Service Frame. When the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Virtual* for the ingress UNI, the value of $(t_1 - t_0) = 0$ because [R56] mandates that $t_0 = t_1$. Similar comments apply to the value of $(t_3 - t_2)$ at the egress UNI.

8.8.2 One-way Frame Delay Performance Metric

This section defines the One-way Frame Delay Performance Metric.

[R27] The SLS **MUST** define the One-way Frame Delay Performance Metric as follows:

- Let $d(\langle i,j\rangle, Pd, T_l)$ represent the Pd-Percentile, expressed as a percentage in $(0, 100]$, of the One-way Frame Delay for all Qualified Service Frames[10] mapped to EVC EP i at the UNI where EVC EP i is located that result in an Egress Service Frame at the UNI where EVC EP j is located. If there are no such egress frames, then $d(\langle i,j\rangle, Pd, T_l) = 0$.

[9] 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.

[10] The Class of Service Name used in the definition of Qualified Service Frames is the CoS_Name used in the *CN* 5-tuple in which the Performance Metric resides.

- Then the One-way Frame Delay Performance Metric, denoted by $\bar{d}(S, Pd, T_l)$, is defined as the maximum value of all of the values $d(\langle i,j\rangle, Pd, T_l)$ for $\langle i,j\rangle \in S$.

To restate the definition mathematically, let $D(\langle i,j\rangle, T_l)$ be the set of one-way Frame Delay values for all Qualified Service Frames mapped to EVC EP i at the UNI where EVC EP i is located that result in an Egress Service Frame at the UNI where EVC EP j is located. $D(\langle i,j\rangle, T_l)$ can be expressed as $D(\langle i,j\rangle, T_l) = \{d_1(\langle i,j\rangle, T_l), d_2(\langle i,j\rangle, T_l), \dots, d_N(\langle i,j\rangle, T_l)\}$ where $d_k(\langle i,j\rangle, T_l)$ is the one-way Frame Delay of the k^{th} frame and where N is the number of elements in $D(\langle i,j\rangle, T_l)$. Define $d(\langle i,j\rangle, Pd, T_l)$ for $Pd > 0$ as

$$d(\langle i,j\rangle, Pd, T_l) = \begin{cases} min\left\{\delta \in D(\langle i,j\rangle, T_l) | Pd \le \frac{100}{N} \sum_{k=1}^{N} I\left(\delta, d_k(\langle i,j\rangle, T_l)\right)\right\} & \text{if } N > 0 \\ 0 \text{ otherwise} \end{cases}$$

where $I(\delta, \varepsilon)$ is an indicator function defined by

$$I(\delta, \varepsilon) = \begin{cases} 1 \text{ if } \delta \ge \varepsilon \\ 0 \text{ otherwise} \end{cases}$$

$d(\langle i,j\rangle, Pd, T_l)$ is the minimal delay during the time internal T_l that Pd percent of the frames do not exceed.

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

$$\bar{d}(S, Pd, T_l) = max\{d(\langle i,j\rangle, Pd, T_l) | \langle i,j\rangle \in S\}$$

Table 5 shows what is contained in a PM entry for the One-way Frame Delay Performance Metric.

Item	Value
Performance Metric Name	One-Way Frame Delay Performance Metric
S	Non-empty subset of the ordered pairs of EVC EPs
Pd	A percentage in $(0, 100]$
$\hat{d}$	Performance Objective in time units > 0

Table 5 – *PM* Entry for the One-way Frame Delay Performance Metric

[R28] The SLS **MUST** define the One-way Frame Delay Performance Objective as met over T_l for a PM entry of the form in Table 5 if and only if $\bar{d}(S, Pd, T_l) \le \hat{d}$.

8.8.3 One-way Mean Frame Delay Performance Metric

This section defines the One-way Mean Frame Delay Performance Metric.

[R29] The SLS **MUST** define the One-way Mean Frame Delay Performance Metric as follows:

- Let $\mu(\langle i,j\rangle,T_l)$ represent the arithmetic mean of One-way Frame Delay for all Qualified Service Frames[11] mapped to EVC EP i at the UNI where EVC EP i is located that result in an Egress Service Frame at the UNI where EVC EP j is located. If there are no such egress frames, then $\mu(\langle i,j\rangle,T_l) = 0$.

- Then the One-way Mean Frame Delay Performance Metric, denoted by $\bar{\mu}(S,T_l)$, is defined as the maximum value of all of the values $\mu(\langle i,j\rangle,T_l)$ for $\langle i,j\rangle \in S$.

To restate the definition mathematically, let $D(\langle i,j\rangle,T_l)$ be the set of one-way Frame Delay values for all Qualified Service Frames mapped to EVC EP i at the UNI where EVC EP i is located that result in an Egress Service Frame at the UNI where EVC EP j is located. $D(\langle i,j\rangle,T_l)$ can be expressed as $D(\langle i,j\rangle,T_l) = \{d_1(\langle i,j\rangle,T_l), d_2(\langle i,j\rangle,T_l), \ldots, d_N(\langle i,j\rangle,T_l)\}$ where $d_k(\langle i,j\rangle,T_l)$ is the One-way Frame Delay of the k^{th} frame and where N is the number of elements in $D(\langle i,j\rangle,T_l)$. Define $\mu(\langle i,j\rangle,T_l)$ as

$$\mu(\langle i,j\rangle,T_l) = \begin{cases} \frac{1}{N}\sum_{k=1}^{N} d_k(\langle i,j\rangle,T_l) \text{ if } N > 0 \\ 0 \; if \; N = 0 \end{cases}$$

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

$$\bar{\mu}(S,T_l) = max\{\mu(\langle i,j\rangle,T_l) | \langle i,j\rangle \in S\}$$

Table 6 shows what is contained in a PM entry for the One-way Mean Frame Delay Performance Metric.

Item	Value
Performance Metric Name	One-way Mean Frame Delay Performance Metric
S	Non-empty subset of the ordered pairs of EVC EPs
$\hat{\mu}$	Performance Objective in time units > 0

Table 6 – *PM* Entry for the One-way Mean Frame Delay Performance Metric

[R30] The SLS **MUST** define the One-way Mean Frame Delay Performance Objective as met over T_l for a PM entry of the form in Table 6 if and only if $\bar{\mu}(S,T_l) \le \hat{\mu}$.

8.8.4 One-way Frame Delay Range Performance Metric

This section defines the One-way Frame Delay Range Performance Metric.

[R31] The SLS **MUST** define the One-way Frame Delay Range Performance Metric as follows:

[11] The Class of Service Name used in the definition of Qualified Service Frames is the CoS_Name used in the CN 5-tuple in which the Performance Metric resides.

- Let $dr(\langle i,j\rangle, Pr, T_l) = d(\langle i,j\rangle, Pr, T_l) - dmin(\langle i,j\rangle, T_l)$. The term $d(\langle i,j\rangle, Pr, T_l)$ represents the Pr-Percentile, expressed as a percentage in $(0, 100]$, of the One-way Frame Delay for all Qualified Service Frames[12] mapped to EVC EP *i* at the UNI where EVC EP *i* is located that result in an Egress Service Frame at the UNI where EVC EP j is located. $dmin(\langle i,j\rangle, T_l)$ is the minimum of the One-way Frame Delays for all Qualified Service Frames delivered to the UNI where EVC EP j is located resulting from an ingress frame mapped to EVC EP *i* at the UNI where EVC EP i is located. If there are no such egress frames, then $dr(\langle i,j\rangle, P_r, T_l) = 0$.

- Then the One-way Frame Delay Range Performance Metric, denoted by $\overline{dr}(S, Pr, T_l)$, is defined as the maximum value of all of the values of $dr(\langle i,j\rangle, P_r, T_l)$ for $\langle i,j\rangle \in S$.

To restate the definition mathematically, let $D(\langle i,j\rangle, T_l)$ be the set of one-way Frame Delay values for all Qualified Service Frames[13] mapped to EVC EP *i* at the UNI where EVC EP *i* is located that result in an Egress Service Frame at the UNI where EVC EP *j* is located. $D(\langle i,j\rangle, T_l)$ can be expressed as $D(\langle i,j\rangle, T_l) = \{d_1(\langle i,j\rangle, T_l), d_2(\langle i,j\rangle, T_l), \ldots, d_N(\langle i,j\rangle, T_l)\}$ where $d_k(\langle i,j\rangle, T_l)$ is the one-way Frame Delay of the k^{th} frame and where N is the number of elements in $D(\langle i,j\rangle, T_l)$. With $dmin(\langle i,j\rangle, T_l) = \min\{d \in D(\langle i,j\rangle, T_l)\}$, define $dr(\langle i,j\rangle, P_r, T_l)$ as

$$dr(\langle i,j\rangle, Pr, T_l) = \begin{cases} d(\langle i,j\rangle, Pr, T_l) - dmin(\langle i,j\rangle, T_l) \text{ if } N > 0 \\ 0 \text{ if } N = 0 \end{cases}$$

where

$$d(\langle i,j\rangle, Pr, T_l) = \begin{cases} min\left\{\delta \in D(\langle i,j\rangle, T_l) | Pr \leq \dfrac{100}{N}\displaystyle\sum_{k=1}^{N} I\big(\delta, d_k(\langle i,j\rangle, T_l)\big)\right\} \text{ if } N > 0 \\ 0 \text{ otherwise} \end{cases}$$

and where $I(\delta, \varepsilon)$ is an indicator function defined by

$$I(\delta, \varepsilon) = \begin{cases} 1 \text{ if } \delta \geq \varepsilon \\ 0 \text{ otherwise} \end{cases}$$

Then the One-way Frame Delay Range Performance Metric for an EVC can be expressed as

$$\overline{dr}(S, Pr, T_l) = max\{dr(\langle i,j\rangle, Pr, T_l) | \langle i,j\rangle \in S\}$$

Table 7 shows what is contained in a PM entry for the One-way Frame Delay Range Performance Metric.

[12] The Class of Service Name used in the definition of Qualified Service Frames is the CoS_Name used in the *CN* 5-tuple in which the Performance Metric resides.

[13] The Class of Service Name used in the definition of Qualified Service Frames is the CoS_Name used in the *CN* 5-tuple in which the Performance Metric resides.

Item	Value
Performance Metric Name	One-way Frame Delay Range Performance Metric
S	Non-empty subset of the ordered pairs of EVC EPs
Pr	A percentage in $(0, 100]$
$\widehat{dr}$	Performance Objective in time units > 0

Table 7 – *PM* Entry for the One-way Frame Delay Range Performance Metric

[R32] The SLS **MUST** define the One-way Frame Delay Range Performance Objective as met over T_l for a PM entry of the form in Table 7 if and only if $\overline{dr}(S, Pr, T_l) \leq \widehat{dr}$.

8.8.5 One-way Inter-Frame Delay Variation Performance Metric

The One-way Inter-Frame Delay Variation Performance Metric characterizes the difference between the One-way Frame Delays of pairs of selected Service Frames.

Let a_q be the Service Frame Arrival Time of the q^{th} Service Frame at the ingress UNI and let $\Delta\tau$ be a short time duration > 0, then the two Service Frames k and l are selected according to the selection criterion:

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

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

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

For $k < l$, a positive value for $d_k - d_l$ implies that the two Service Frames are closer together at the egress UNI while a negative value implies that the two Service 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 Performance Metric.

Figure 19 shows the different times that are related to One-way Inter-Frame Delay Variation Performance Metric.

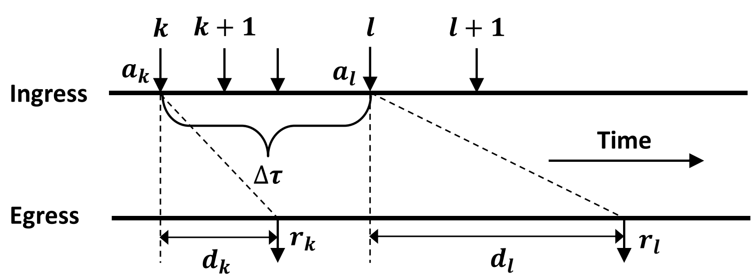

Figure 19 – One-way Inter-Frame Delay Variation Definition

[R33] The SLS **MUST** define the One-way Inter-Frame Delay Variation Performance Metric as follows:

- Let $dv(\langle i,j\rangle, Pv, \Delta\tau, T_l)$ be the Pv-percentile, expressed as a percentage in $(0, 100]$, of the Δd_{kl}'s of all Qualified Service Frame pairs where each Qualified Service Frame[14] ingresses at EVC EP i and results in an egress frame at EVC EP j and whose difference in Service Frame Arrival Times of each frame in the pair at EVC EP i was $\Delta\tau$.

- If there are no such pairs of frames for EVC EP i and EVC EP j, then $dv(\langle i,j\rangle, Pv, \Delta\tau, T_l) = 0$.

- Then the One-way Inter-Frame Delay Variation Performance Metric, denoted by $\overline{dv}(S, Pv, \Delta\tau, T_l)$, is the maximum of the values $dv(\langle i,j\rangle, Pv, \Delta\tau, T_l)$ for $\langle i,j\rangle \in S$.

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 a few multiples of the inter-frame time.

To restate the definition mathematically, let

$$V(\langle i,j\rangle, \Delta\tau, T_l) = \{\Delta d_1(\langle i,j\rangle, \Delta\tau, T_l), \Delta d_2(\langle i,j\rangle, \Delta\tau, T_l), \dots, \Delta d_N(\langle i,j\rangle, \Delta\tau, T_l)\}$$

be the set of the absolute value of delay variations for all eligible pairs of Qualified Service Frames mapped to the EVC EP i from the UNI where EVC EP i is located to the UNI where EVC EP j is located where the difference in the Service Frame Arrival Times of each frame at the ingress UNI was $\Delta\tau$. Define

$$dv(\langle i,j\rangle, Pv, \Delta\tau, T_l) = \begin{cases} min\left\{\delta \in V(\langle i,j\rangle, \Delta\tau, T_l) \middle| Pv \le \frac{100}{N}\sum_{k=1}^{N} I\big(\delta, \Delta d_k(\langle i,j\rangle, \Delta\tau, T_l)\big)\right\} & \text{if } N > 0 \\ 0 \text{ otherwise} \end{cases}$$

where N is the number of elements in $V(\langle i,j\rangle, \Delta\tau, T_l)$ and $I(\delta, \varepsilon)$ is an indicator function defined by

$$I(\delta, \varepsilon) = \begin{cases} 1 \text{ if } \delta \ge \varepsilon \\ 0 \text{ otherwise} \end{cases}$$

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

$$\overline{dv}(S, Pv, \Delta\tau, T_l) = max\{dv(\langle i,j\rangle, Pv, \Delta\tau, T_l) | \langle i,j\rangle \in S\}$$

Table 8 shows what is contained in a PM entry for the One-way Inter-Frame Delay Variation Performance Metric.

[14] The Class of Service Name used in the definition of Qualified Service Frames is the CoS_Name used in the CN 5-tuple in which the Performance Metric resides.

Item	Value
Performance Metric Name	One-way Inter-Frame Delay Variation Performance Metric
S	Non-empty subset of the ordered pairs of EVC EPs
Pv	A percentage in $(0, 100]$
$\Delta\tau$	A time duration in time units
$\widehat{dv}$	Performance Objective in time units > 0

Table 8 – *PM* Entry for the One-way Inter-Frame Delay Variation Performance Metric

[R34] The SLS **MUST** define the One-way Inter-Frame Delay Variation Performance Objective as met over T_l for a PM entry of the form in Table 8 if and only if $\overline{dv}(S, Pv, \Delta\tau, T_l) \leq \widehat{dv}$.

8.8.6 One-way Frame Loss Ratio Performance Metric

The One-way Frame Loss Ratio Performance Metric characterizes frame loss as a percentage.

[R35] The SLS **MUST** define the One-way Frame Loss Ratio Performance Metric as follows:

- Let $I(\langle i,j\rangle, T_l)$ denote the number of ingress Qualified Service Frames[15] mapped to the EVC EP i at the UNI where EVC EP i is located that were intended to have been delivered to the UNI where EVC EP j is located according to the value of the EVC Data Service Frame Disposition Service Attribute (Section 8.4).
- Let $E(\langle i,j\rangle, T_l)$ be the number of unique (not duplicate) Egress Service Frames where each Service Frame is the first Egress Service Frame mapped to the EVC EP j at the UNI where EVC EP j is located that results from a frame counted in $I(\langle i,j\rangle, T_l)$.
- Define $FLR(\langle i,j\rangle, T_l) = \begin{cases} \left(\frac{I(\langle i,j\rangle,T_l) - E(\langle i,j\rangle,T_l)}{I(\langle i,j\rangle,T_l)}\right) \times 100 \text{ if } I(\langle i,j\rangle, T_l) > 0 \\ 0 \text{ otherwise} \end{cases}$
- Then the One-way Frame Loss Ratio Performance Metric is defined as $\overline{FLR}(S, T_l) = max\{FLR(\langle i,j\rangle, T_l) | \langle i,j\rangle \in S\}$.

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

$$FLR(\langle i,j\rangle, T_l) = \begin{cases} \left(\frac{\tilde{I}(\langle i,j\rangle,T_l) - E(\langle i,j\rangle,T_l)}{\tilde{I}(\langle i,j\rangle,T_l)}\right) \times 100 \text{ if } \tilde{I}(\langle i,j\rangle, T_l) > 0 \\ 0 \text{ otherwise} \end{cases}$$

[15] The Class of Service Name used in the definition of Qualified Service Frames is the CoS_Name used in the *CN* 5-tuple in which the Performance Metric resides.

where $\tilde{I}(\langle i,j\rangle,T_l) = I(\langle i,j\rangle,T_l) - \tilde{D}(\langle i,j\rangle,T_l)$ and $\tilde{D}(\langle i,j\rangle,T_l)$ is the number of frames discarded by the Service Provider, in order to conform to either the MAC Data Rate of the UNI where EVC EP j is located or an Egress Bandwidth Profile (if one is used) at the UNI where EVC EP j is located. Such frame drops may occur anywhere in the network, not just near the egress UNI. One example of this could be where an Egress Bandwidth Profile is implemented by applying a policer or shaper on a link within the network. Another example of this could be where Green frames for this EVC and Class of Service Name that are intended to be delivered to the UNI with EVC EP j exceed the MAC Data Rate on a link within the network, provided the MAC Data Rate of that link is greater than or equal to the MAC Data Rate of the UNI. Good traffic engineering principles would 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 the UNI where EVC EP j is located. The details of such an adjustment are beyond the scope of this document.

Table 9 shows what is contained in a PM entry for the One-way Frame Loss Ratio Performance Metric.

Item	Value
Performance Metric Name	One-way Frame Loss Ratio Performance Metric
S	Non-empty subset of the ordered pairs of EVC EPs
$\widehat{FLR}$	Performance Objective expressed as a percentage

Table 9 – *PM* Entry for the One-way Frame Loss Ratio Performance Metric

[R36] The SLS **MUST** define the One-way Frame Loss Performance Objective as met over T_l for a PM entry of the form in Table 9 if and only if $\overline{FLR}(S,T_l) \leq \widehat{FLR}$.

8.8.7 One-way Availability Performance Metric

The One-way Availability Performance Metric is the percentage of time which is Available Time (Section 8.8.1.3). The precise definition is presented in the following paragraphs. As an example, suppose the Performance Objective for the One-way Availability Performance Metric is 99.9% when the duration of $T_l = 30$ days. If there is no Maintenance Interval (Section 8.8.1.3) this objective will allow the service to be unavailable for approximately 43 minutes out of the 30 days.

For a given $\langle i,j\rangle \in S$, a given *CoS_Name*,[16] and a given T_l, define

$$\tilde{A}(\langle i,j\rangle,T_l) = \begin{cases} 100 \times \dfrac{|AT(\langle i,j\rangle,T_l)|}{|W(T_l)|} & \text{if } |W(T_l)| > 0 \\ 0 \text{ otherwise} \end{cases}$$

where the sets $AT(\langle i,j\rangle,T_l)$ and $W(T_l)$ are defined in Section 8.8.1.3 and $|\quad|$ represents the number of elements in a set. Note that by the definitions in Section 8.8.1.3, Δt_k's that intersect a Maintenance Interval are not included in the calculation of $AT(\langle i,j\rangle,T_l)$.

[R37] The SLS **MUST** define the One-way Availability Performance Metric as:

[16] The Class of Service Name used is the CoS_Name used in the CN 5-tuple in which the Performance Metric resides.

$$\bar{A}(S, T_l) = min\{\tilde{A}(\langle i,j\rangle, T_l) | \langle i,j\rangle \in S\}$$

Table 10 shows what is contained in a PM entry for the One-way Availability Performance Metric.

Item	Value
Performance Metric Name	One-way Availability Performance Metric
S	Non-empty subset of the ordered EVC EP pairs
$\hat{A}$	Performance Objective expressed as a percentage

Table 10 – *PM* Entry for the One-way Availability Performance Metric

[R38] The SLS **MUST** define the One-way Availability Performance Objective as met over T_l for a PM entry of the form in Table 10 if and only if $\bar{A}(S, T_l) \geq \hat{A}$.

8.8.8 One-way High Loss Intervals Performance Metric

The One-way High Loss Intervals Performance Metric is a count of the number of Δt_k's that are in Available Time (Section 8.8.1.3) for a given T_l and that have a high frame loss. When $\Delta t = 1$ second, it is analogous to the definition of "Severely Errored Seconds" (SES) defined in Section 9 of ITU Recommendation Y.1563 [20].

Section 8.8.9 defines a related Performance Metric called the One-Way Consecutive High Loss Intervals Performance Metric. Figure 20 illustrates how the One-way High Loss Intervals Performance Metric and the One-way Consecutive High Loss Intervals Performance Metric fit into the hierarchy of time and other attributes.

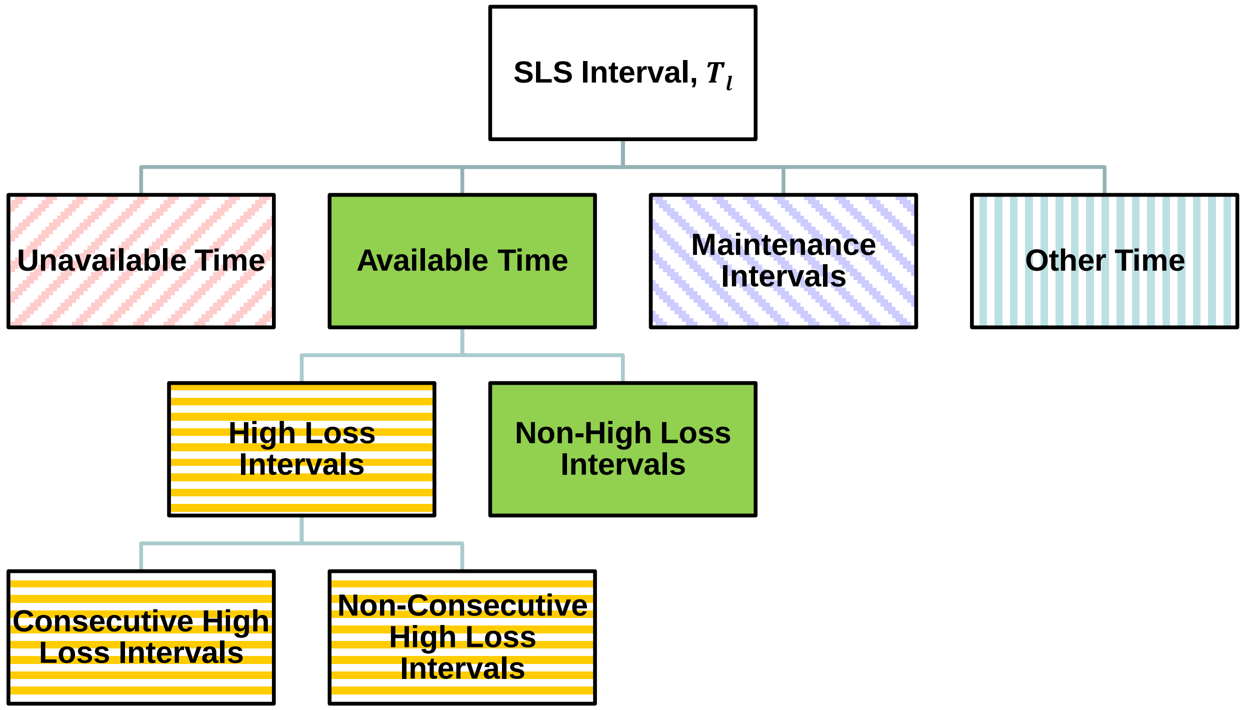

Figure 20 – Hierarchy of Time

For a given $\langle i,j\rangle \in S$, a given *CoS_Name*,[17] and a given C,[18] the high loss state of Δt_k contained in T_l is represented by $H(\langle i,j\rangle, \Delta t_k, T_l)$ that is defined as

$$H(\langle i,j\rangle, \Delta t_k, T_l) = \begin{cases} 1 \text{ if } \Delta t_k \in AT(\langle i,j\rangle, T_l) \text{ and } flr(\langle i,j\rangle, \Delta t_k) > C \\ 0 \text{ otherwise} \end{cases}$$

where $AT(\langle i,j\rangle, T_l)$ and $flr(\langle i,j\rangle, \Delta t_k)$ are defined in Section 8.8.1.3. Recall that $\{\Delta t_k, k = 0,1, \ldots\}$ and C can be different for different Class of Service Names.

For a given $\langle i,j\rangle \in S$ and a given *CoS_Name*, the count of High Loss Intervals, $HL(\langle i,j\rangle, T_l)$, is defined as

$$HL(\langle i,j\rangle, T_l) = \sum_{\Delta t_k \in AT(\langle i,j\rangle, T_l)} H(\langle i,j\rangle, \Delta t_k, T_l)$$

[R39] The SLS **MUST** define the One-way High Loss Intervals Performance Metric as:

$$\overline{HL}(S, T_l) = max\{HL(\langle i,j\rangle, T_l) | \langle i,j\rangle \in S\}$$

Table 11 shows what is contained in a PM entry for the One-way High Loss Intervals Performance Metric.

Item	Value
Performance Metric Name	One-way High Loss Intervals Performance Metric
S	Non-empty subset of the ordered EVC EP pairs
$\widehat{HL}$	Performance Objective expressed as a non-negative integer

Table 11 – *PM* Entry for the One-way High Loss Intervals Performance Metric

[R40] The SLS **MUST** define the One-way High Loss Intervals Performance Objective as met over T_l for a PM entry of the form in Table 11 if and only if $\overline{HL}(S, T_l) \leq \widehat{HL}$.

8.8.9 One-way Consecutive High Loss Intervals Performance Metric

The One-way Consecutive High Loss Intervals Performance Metric is a count of the number of runs of consecutive Δt_k's that are in Available Time (Section 8.8.1.3) and that have a high frame loss. When $\Delta t = 1$ second, it is analogous to the definition of Consecutive Severely Errored Seconds defined in Annex B of ITU Recommendation Y.1563 [20].

For a given $\langle i,j\rangle \in S$, a given *CoS_Name*,[19] a given T_l, a given n,[20] and an integer p such that $1 \leq p < n$, the count of Consecutive High Loss Intervals is denoted by $B(\langle i,j\rangle, p, T_l)$ and defined by the flow chart in Figure 21. Note that $H(\langle i,j\rangle, T_l)$ is defined in Section 8.8.8.

[17] The Class of Service Name used is the CoS_Name used in the *CN* 5-tuple in which the Performance Metric resides.

[18] C is the parameter in the *CN* 5-tuple in which the Performance Metric resides.

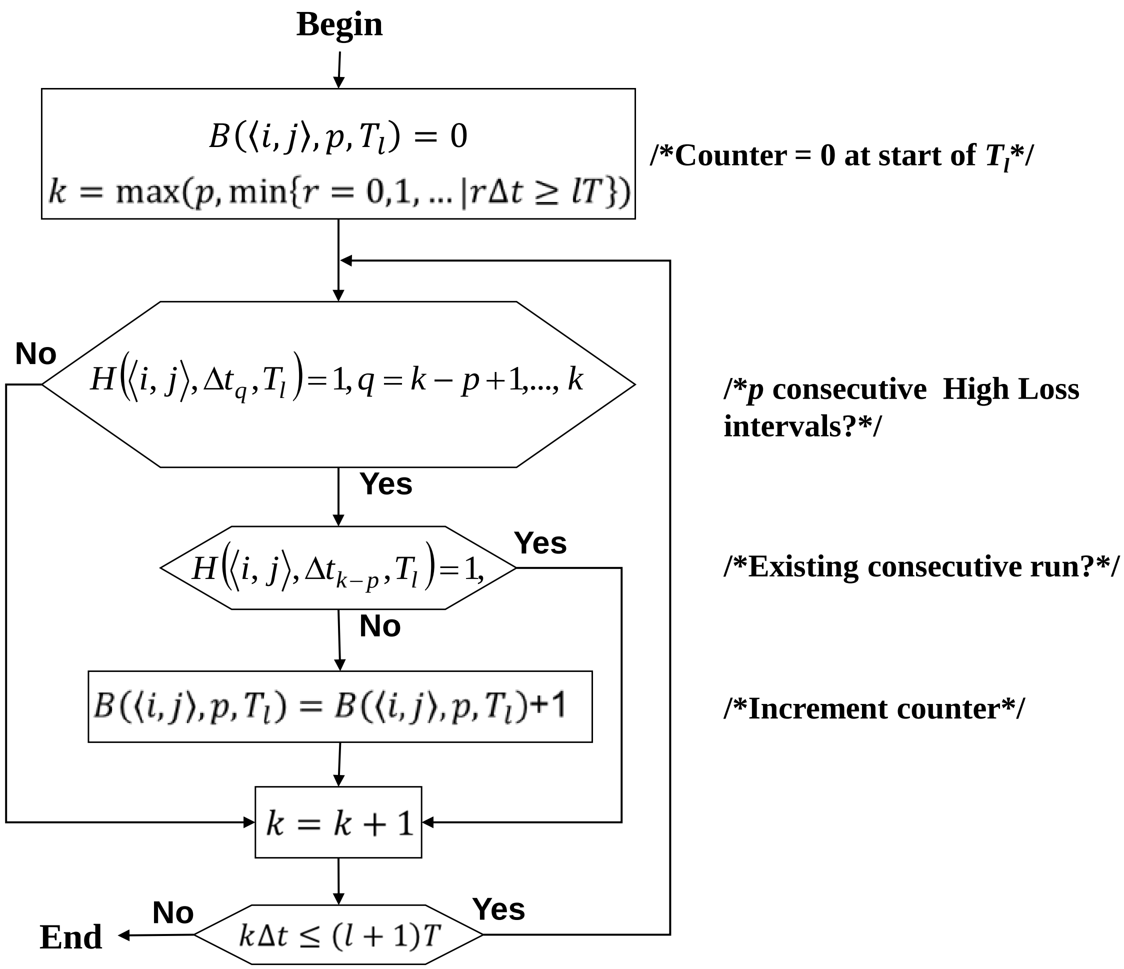

Figure 21 – Determining and Counting Consecutive High Loss Intervals for a Given Class of Service Name

Recall that $\{\Delta t_k, k = 0,1,\ldots\}$, C, and n can be different for different Class of Service Names.

Figure 22 shows an example that depicts the Consecutive High Loss Intervals counting processes when there is no Maintenance Interval. Note that $A(\langle i,j\rangle,\Delta t_k)$ is defined in Section 8.8.1.3 and $H(\langle i,j\rangle,\Delta t_k,T_l)$ is defined in Section 8.8.8.

[19] The Class of Service Name used is the CoS_Name used in the *CN* 5-tuple in which the Performance Metric resides.
[20] n is the parameter in the *CN* 5-tuple in which the Performance Metric resides.

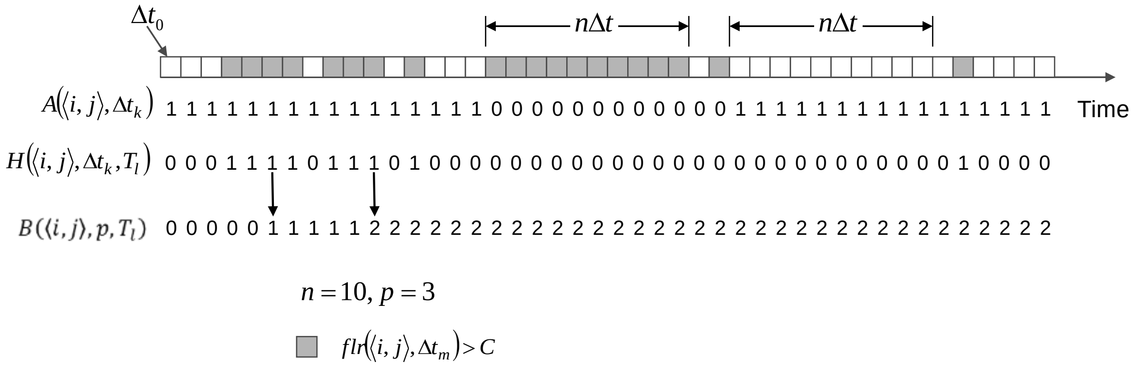

Figure 22 – Example of Counting Consecutive High Loss Intervals

[R41] The SLS **MUST** define the One-way Consecutive High Loss Intervals Performance Metric as:

$$\bar{B}(S,p,T_l) = max\{B(\langle i,j\rangle,p,T_l)|\langle i,j\rangle \in S\}$$

Table 12 shows what is contained in a PM entry for the One-way Consecutive High Loss Intervals Performance Metric.

Item	Value
Performance Metric Name	One-way Consecutive High Loss Intervals Performance Metric
S	Non-empty subset of the ordered EVC EP pairs
p	An integer ≥ 1 and $< n$
$\hat{B}$	Performance Objective expressed as a non-negative integer

Table 12 – *PM* Entry for the One-way Consecutive High Loss Intervals Performance Metric

[R42] The SLS **MUST** define the One-way Consecutive High Loss Intervals Performance Objective as met over T_l for a PM entry of the form in Table 12 if and only if $\bar{B}(S,p,T_l) \leq \hat{B}$.

8.8.10 One-way Composite Performance Metric

The One-way Composite Performance Metric is expressed as the percentage of Δt_k time intervals contained in time interval T_l, that are deemed to have Acceptable Performance. Acceptable Performance for a short time interval, Δt_k , is based on a combination of three frame delivery characteristics (Composite One-way Frame Delay, Composite One-way Frame Delay Variation and Composite One-way Frame Loss). Composite One-way Frame Delay Variation is the difference in delay for Service Frames sent one right after the other. (The precise definition is presented in the following paragraphs.) The combination of three frame delivery characteristics for the short time interval, Δt_k, is denoted as the Composite Performance Indicator, CPI, which is defined in following paragraphs. Note that the definition of Acceptable Performance for a Δt_k is different from the definition of Available Time (Section 8.8.1.3). Also note that:

- The Composite One-way Frame Delay characteristic is different from the One-way Frame Delay Performance Metric (Section 8.8.2) and it is different from the One-way Mean Frame Delay Performance Metric (Section 8.8.3),

- The Composite One-way Frame Delay Variation characteristic is different from the One-way Frame Delay Range Performance Metric (Section 8.8.4) and it is different from the One-way Inter-Frame Delay Variation Performance Metric (Section 8.8.5), and

- The Composite One-way Frame Loss characteristic is different from the One-way Frame Loss Ratio Performance Metric (Section 8.8.6).

Using the SLS Notation (Section 8.8.1.2) would cause the equations and figures in this section to be unwieldy. Thus the SLS notation is not always used. The dependencies on the parameters for the expressions used in specifying the One-way Composite Performance Metric can be deduced by tracing through the derivations.

An example of the use of the One-way Composite Performance Metric for clock synchronization is presented in Appendix E.

Informally, the One-way Composite Performance Metric is based on Composite Performance Indicator values during a sequence of consecutive short time intervals. When the previous sequence was defined as *Acceptable*, if the Composite Performance Indicator exceeds its pre-set threshold for each short time interval in the current sequence, then the short time interval at the beginning of the current sequence is defined as *Unacceptable*; otherwise it is defined as *Acceptable*. On the other hand, when the previous sequence was defined as *Unacceptable*, if the Composite Performance Indicator is within its threshold for each short time interval in the current sequence, then the short time interval at the beginning of the current sequence is defined as *Acceptable*; otherwise, it is defined as *Unacceptable*. The formal definition follows.

The following parameters are specific to the definition of the One-way Composite Performance Metric:

- U in $(0,1)$, the Composite Performance Indicator threshold which if exceeded suggests an unacceptable time interval.

- $Wfl = 0$ or 1, the indicator for Composite One-way Frame Loss.

- $Wfd = 0$ or 1, the indicator for Composite One-way Frame Delay.

- $Wfdv = 0$ or 1, the indicator for Composite One-way Frame Delay Variation.

- $DL > 0$, the Composite One-way Frame Delay threshold in time units.

- $Jt > 0$, the Composite One-way Frame Delay Variation threshold in time units.

For a given $\langle i,j \rangle \in S$, a given T_l, and a given *CoS_Name*, Δt_k is defined to be either Acceptable or Unacceptable. This is represented by $AC(\langle i,j \rangle, \Delta t_k, T_l)$ where $AC(\langle i,j \rangle, \Delta t_k, T_l) = 1$ means that Δt_k is Acceptable and $AC(\langle i,j\rangle, \Delta t_k, T_l) = 0$ means that Δt_k is Unacceptable. $AC(\langle i,j\rangle, \Delta t_k, T_l)$ is based on the Composite Performance Indicator, $CPI(\langle i,j\rangle, \Delta t_k, T_l)$, over a number of consecutive short time intervals.

For a given $\langle i,j\rangle \in S$, a given *CoS_Name*, and a given T_l, the Composite Performance Indicator, $CPI(\langle i,j\rangle, \Delta t_k, T_l)$, is based on the frame delivery characteristics of ingress Qualified Service Frames[21] (Section 8.8.1.4) mapped to the EVC EP i at the ingress UNI where EVC EP i is located that are intended to be delivered to the egress UNI where EVC EP j is located according to the value of the EVC Data Service Frame Disposition Service Attribute (Section 8.4), during the short time interval, Δt_k. The short time interval, Δt_k, is considered as an unacceptable interval if $CPI(\langle i,j\rangle, \Delta t_k, T_l) > U$.

For a given $\langle i,j\rangle \in S$, a given *CoS_Name*, and a given T_l, let $M(\langle i,j\rangle, \Delta t_k, T_l)$ be the number of ingress Qualified Service Frames that are mapped to EVC EP i, are intended to be delivered to EVC EP j according to the value of the EVC Data Service Frame Disposition Service Attribute (Section 8.4), and arrive at the UNI during Δt_k. Let these ingress Qualified Service Frames be numbered $1,2, \ldots, M(\langle i,j\rangle, \Delta t_k, T_l)$, where Qualified Service Frame p arrived at the ingress UNI before Qualified Service Frame q if and only if $p < q$.

Note that by the definition of Qualified Service Frames (Section 8.8.1.4), for a given T_l, ingress Qualified Service Frames can only occur during a $\Delta t_k \in AT(\langle i,j\rangle, T_l)$ which means that Δt_k is contained in T_l. Thus if $\Delta t_k \notin AT(\langle i,j\rangle, T_l)$, then $M(\langle i,j\rangle, \Delta t_k, T_l) = 0$.

Let $N(\langle i,j\rangle, \Delta t_k, T_l)$ be the number of unique (not duplicate) Egress Service Frames where each Service Frame is the first unerrored Egress Service Frame at the UNI where EVC EP j is located that results from a frame counted in $M(\langle i,j\rangle, \Delta t_k, T_l)$. A frame is considered lost if the ingress Qualified Service Frame does not result in an egress frame at the UNI where EVC EP j is located.

Let $m \in \{1,2, \ldots, M(\langle i,j\rangle, \Delta t_k, T_l)\}$ and $M(\langle i,j\rangle, \Delta t_k, T_l) > 0$. Then the Frame Loss Characteristic, $fl(m, \langle i,j\rangle, \Delta t_k, T_l)$, is defined as

$$fl(m, \langle i,j\rangle, \Delta t_k, T_l) = \begin{cases} 1 \text{ if frame } m \text{ is lost} \\ 0 \text{ if frame } m \text{ is not lost} \end{cases}$$

Let $d(m, \langle i,j\rangle, \Delta t_k, T_l)$ equal the One-way Frame Delay for frame m. Then the Frame Delay Characteristic, $fd(m, \langle i,j\rangle, \Delta t_k, T_l)$, is defined as

$$fd(m, \langle i,j\rangle, \Delta t_k, T_l) = \begin{cases} 1 \text{ if } fl(m, \langle i,j\rangle, \Delta t_k, T_l) \neq 1 \text{ and } d(m, \langle i,j\rangle, \Delta t_k, T_l) > DL \\ 0 \text{ otherwise} \end{cases}$$

When both frame m and frame $m-1$ are not lost, let $dv(m, \langle i,j\rangle, \Delta t_k, T_l) = |d(m, \langle i,j\rangle, \Delta t_k, T_l) - d(m-1, \langle i,j\rangle, \Delta t_k, T_l)|$ for $m > 1$. Then the Frame Delay Variation Characteristic, $fdv(m, \langle i,j\rangle, \Delta t_k, T_l)$, is defined as

[21] The Class of Service Name used in the definition of Qualified Service Frames is the CoS_Name used in the *CN* 5-tuple in which the Performance Metric resides.

$$fdv(m,\langle i,j\rangle,\Delta t_k,T_l) = \begin{cases}1 \text{ if } v(\langle i,j\rangle,\Delta t_k,T_l)\text{=1 and } dv(m,\langle i,j\rangle,\Delta t_k,T_l) > Jt \text{ and } m > 1\\ 0 \text{ otherwise}\end{cases}$$

where

$$v(m,\langle i,j\rangle,\Delta t_k,T_l) = \begin{cases}1 \text{ if } fl(m,\langle i,j\rangle,\Delta t_k,T_l) \neq 1 \text{ and } fl(m-1,\langle i,j\rangle,\Delta t_k,T_l) \neq 1 \text{ and } m > 1\\ 0 \text{ otherwise}\end{cases}$$

Note that the above characteristics are undefined when $M(\langle i,j\rangle,\Delta t_k,T_l) = 0$.

The Composite Performance Indicator, $CPI(\langle i,j\rangle,\Delta t_k,T_l)$, is a ratio of the number of characteristics with value of 1 to the number of all characteristics during the short time interval, Δt_k.

In order to fit the expression for $CPI(\langle i,j\rangle,\Delta t_k,T_l)$ on a single line, the following intermediate expressions are defined.

For $M(\langle i,j\rangle,\Delta t_k,T_l) > 0$,

$$Zl(\langle i,j\rangle,\Delta t_k,T_l) = Wfl \times \sum_{m=1}^{M(\langle i,j\rangle,\Delta t_k,T_l)} fl(m,\langle i,j\rangle,\Delta t_k,T_l)$$

$Zl(\langle i,j\rangle,\Delta t_k,T_l)$ is the number of frames that are lost going from EVC EP *i* to EVC EP *j* during Δt_k multiplied by Wfl.

For $M(\langle i,j\rangle,\Delta t_k,T_l) > 0$,

$$Zd(\langle i,j\rangle,\Delta t_k,T_l) = Wfd \times \sum_{m=1}^{M(\langle i,j\rangle,\Delta t_k,T_l)} fd(m,\langle i,j\rangle,\Delta t_k,T_l)$$

$Zd(\langle i,j\rangle,\Delta t_k,T_l)$ is the number of frames going from EVC EP *i* to EVC EP *j* during Δt_k that have unacceptable delay multiplied by Wfd.

$$Zdv(\langle i,j\rangle,\Delta t_k,T_l) = \begin{cases}Wfdv \times \sum_{m=2}^{M(\langle i,j\rangle,\Delta t_k,T_l)} fdv(m,\langle i,j\rangle,\Delta t_k,T_l) \text{ if } M(\langle i,j\rangle,\Delta t_k,T_l) > 1\\ 0 \text{ otherwise}\end{cases}$$

$Zdv(\langle i,j\rangle,\Delta t_k,T_l)$ is the number of frame pairs going from EVC EP *i* to EVC EP *j* during Δt_k that have unacceptable delay variation multiplied by $Wfdv$.

For $M(\langle i,j\rangle,\Delta t_k,T_l) > 0$,

$$Yl(\langle i,j\rangle,\Delta t_k,T_l) = Wfl \times M(\langle i,j\rangle,\Delta t_k,T_l)$$

$$Yd(\langle i,j\rangle,\Delta t_k,T_l) = Wfd \times N(\langle i,j\rangle,\Delta t_k,T_l)$$

$$Ydv(\langle i,j\rangle,\Delta t_k,T_l) = \begin{cases} Wfdv \times \sum_{m=2}^{M(\langle i,j\rangle,\Delta t_k,T_l)} v(m,\langle i,j\rangle,\Delta t_k,T_l) \text{ if } M(\langle i,j\rangle,\Delta t_k,T_l) > 1 \\ 0 \text{ otherwise} \end{cases}$$

$$Num(\langle i,j\rangle,\Delta t_k,T_l) = Zl(\langle i,j\rangle,\Delta t_k,T_l) + Zd(\langle i,j\rangle,\Delta t_k,T_l) + Zdv(\langle i,j\rangle,\Delta t_k,T_l)$$

$$Den(\langle i,j\rangle,\Delta t_k,T_l) = Yl(\langle i,j\rangle,\Delta t_k,T_l) + Yd(\langle i,j\rangle,\Delta t_k,T_l) + Ydv(\langle i,j\rangle,\Delta t_k,T_l)$$

With the above expressions defined, the Composite Performance Indicator, $CPI(\langle i,j\rangle,\Delta t_k,T_l)$, is defined as

$$CPI(\langle i,j\rangle,\Delta t_k,T_l) = \begin{cases} 0 \text{ if } M(\langle i,j\rangle,\Delta t_k,T_l) = 0 \\ \dfrac{Num(\langle i,j\rangle,\Delta t_k,T_l)}{Den(\langle i,j\rangle,\Delta t_k,T_l)} \text{ if } Den(\langle i,j\rangle,\Delta t_k,T_l) > 0 \text{ and } M(\langle i,j\rangle,\Delta t_k,T_l) > 0 \\ 1 \text{ otherwise} \end{cases}$$

Note that it is necessary for at least one of Wfl, Wfd, and $Wfdv$ to be 1 for $CPI(\langle i,j\rangle,\Delta t_k,T_l)$ to reflect performance when $M(\langle i,j\rangle,\Delta t_k,T_l) > 0$.

[R43] At least one of Wfl, Wfd, and $Wfdv$ **MUST** equal 1.

Recall that if $\Delta t_k \notin AT(\langle i,j\rangle,T_l)$, then $M(\langle i,j\rangle,\Delta t_k,T_l) = 0$. Consequently if Δt_k is not contained in T_l or straddles the boundary of T_l, then $CPI(\langle i,j\rangle,\Delta t_k,T_l) = 0$.

Recall that Δt_0 is the first small time interval after the start of the SLS (Section 8.8.1.3). Then $AC(\langle i,j\rangle,\Delta t_k,T_l)$ is defined by the flowchart in Figure 23 for $k = 0,1,\ldots$.

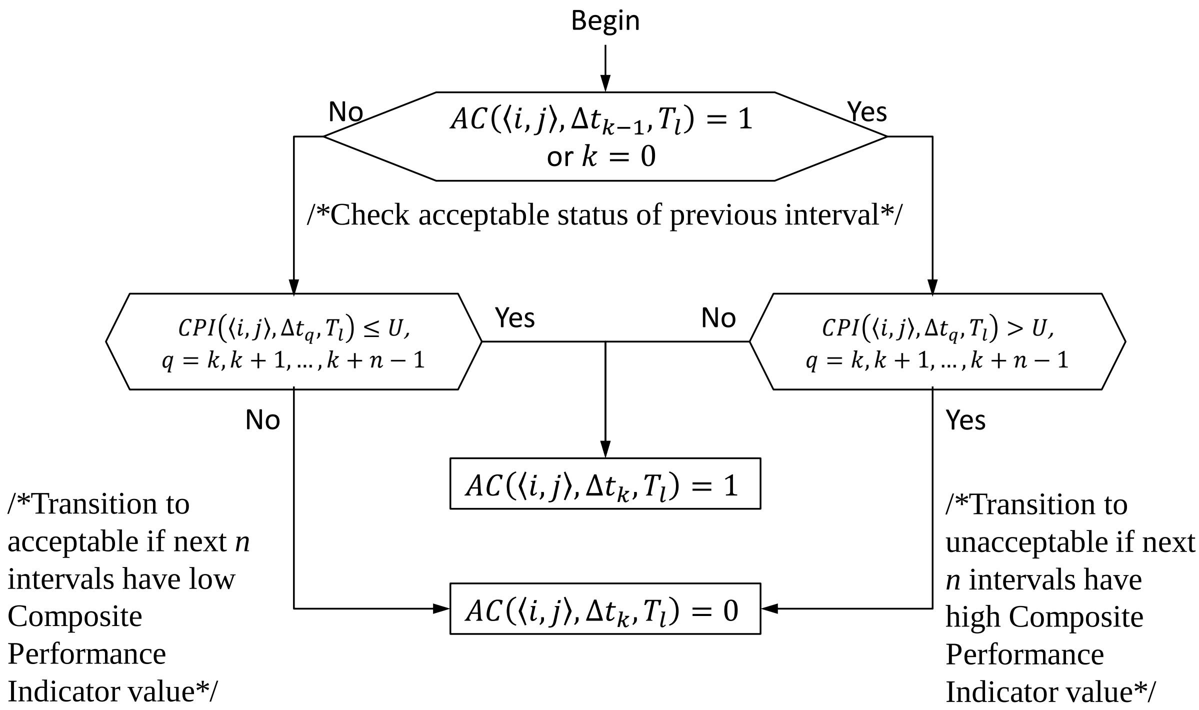

Figure 23 – Flowchart Definition of $AC(\langle i,j\rangle, \Delta t_k, T_l)$

Recall that $\{\Delta t_k, k = 0,1,\ldots\}$ and n can be different for different Class of Service Names.

The Acceptable status for Δt_k is based on the Composite Performance Indicator during Δt_k and each of the following $n-1$ short time intervals and the Acceptable status of the previous short time interval. In other words, a sliding window of width $n\Delta t$ is used to determine the One-way Composite Performance Metric. This use of a sliding window is similar to that of ITU-T Y.1563 [20].

Figure 24 presents an example of the determination of the Acceptable status for the short time intervals with a sliding window of 10 short time intervals where it assumed that all small time intervals belong to $AT(\langle i,j\rangle, T_l)$ (Section 8.8.1.3).

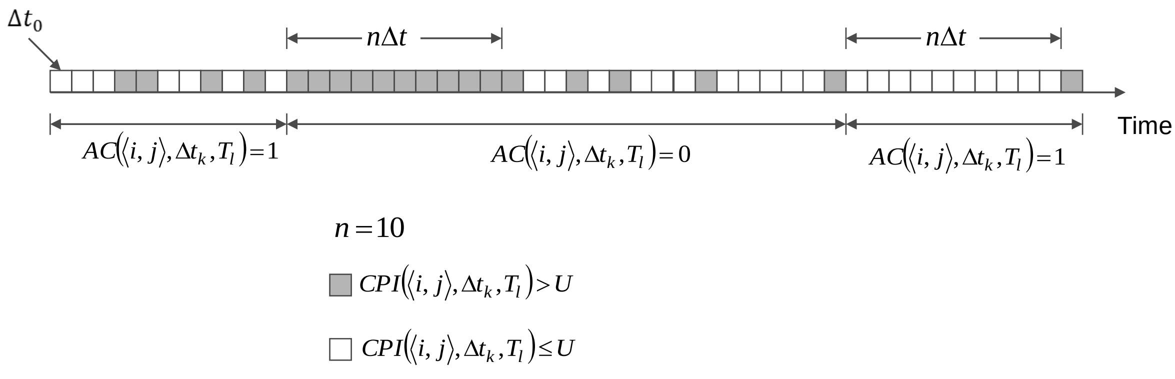

Figure 24 – Example of the Determination of $AC(\langle i,j\rangle, \Delta t_k, T_l)$

For a given $\langle i,j\rangle \in S$, a given *CoS_Name*, and a given T_l, the One-way Composite Performance is defined by

$$ACP(\langle i,j\rangle, U, DL, Jt, Wfl, Wfd, Wfdv, T_l) = \begin{cases} \frac{100}{|W(T_l)|} \sum_{\Delta t_k \in W(T_l)} AC(\langle i,j\rangle, \Delta t_k, T_l) \text{ if } |W(T_l)| > 0 \\ 100 \text{ otherwise} \end{cases}$$

where $W(T_l)$ is defined in Section 8.8.1.3.

[R44] The SLS **MUST** define the One-way Composite Performance Metric as:

$$\overline{AC}(S, U, DL, Jt, Wfl, Wfd, Wfdv, T_l) = min\{ACP(\langle i,j\rangle, U, DL, Jt, Wfl, Wfd, Wfdv, T_l) | \langle i,j\rangle \in S\}$$

Table 13 shows what is contained in a PM entry for the One-way Composite Performance Metric.

Item	Value
Performance Metric Name	One-way Composite Performance Metric
S	Non-empty subset of ordered EVC EP pairs
U	A real number in the range $(0,1)$
DL	A positive real number
Jt	A positive real number
Wfl	An integer $= 0$ or 1
Wfd	An integer $= 0$ or 1
$Wfdv$	An integer $= 0$ or 1
$\widehat{AC}$	Performance Objective expressed as a percentage

Table 13 – *PM* Entry for the One-way Composite Performance Metric

[R45] The SLS **MUST** define the One-way Composite Performance Objective as met over T_l for a PM entry of the form in Table 13 if and only if $\overline{AC}(S, U, DL, Jt, Wfl, Wfd, Wfdv, T_l) \geq \widehat{AC}$.

8.8.11 One-way Group Availability Performance Metric

The One-way Group Availability Performance Metric characterizes the percentage of time that at least a specified number of sets of ordered EVC EP Pairs are available. It is defined for a collection of two or more non-empty sets of ordered pairs of EVC EPs that are associated by the EVC denoted by $G = \{S_1, \ldots, S_m\}$.

Note that the Multiple EVC Service Level Specification Service Attribute (Section 11) contains a Performance Metric similar to the One-way Group Availability Performance Metric but is based on multiple EVCs.

The following parameters are specific to the definition of the One-way Group Availability Performance Metric:

- $G = \{S_1, \ldots, S_m\}$, where each element in G is a non-empty set of ordered pairs of EVC EPs, and
- K, an integer ≥ 1.

[R46] For $j = 1, \ldots, m$, in each ordered EVC EP pair element of S_j, at least one of the EVC EPs **MUST** have the value of the EVC EP Role Service Attribute (Section 10.3) = *Root*.

For a given *CoS_Name*[22] and a given G, define $Ag(S, \Delta t_k)$ for $S \in G$ as

$$Ag(S, \Delta t_k) = min\left\{A\left(\langle i, j\rangle, \Delta t_k\right) | \langle i, j\rangle \in S\right\}$$

where $A(\langle i, j\rangle, \Delta t_k)$ is defined in Section 8.8.1.3. Note that $Ag(S, \Delta t_k) = 1$ when all ordered EVC EP pairs in S have been experiencing low loss. Otherwise, $Ag(S, \Delta t_k) = 0$.

The One-way Group Availability for the time interval Δt_k is defined as

$$GA(G, K, \Delta t_k) = \begin{cases} 1 \text{ if } \sum_{S \in G} Ag(S, \Delta t_k) \geq K \\ 0 \text{ otherwise} \end{cases}$$

[R47] The SLS **MUST** define the One-way Group Availability Performance Metric as

$$\overline{GA}(G, K, T_l) = \begin{cases} \frac{100}{|W(T_l)|} \sum_{\Delta t_k \in W(T_l)} GA(G, K, \Delta t_k) \text{ if } |W(T_l)| > 0 \\ 100 \text{ otherwise} \end{cases}$$

Note that $W(T_l)$ is defined in Section 8.8.1.3.

The One-way Group Availability Performance Metric can be viewed as the percentage of time that at least K S_j's in G are Available.

Table 14 shows what is contained in a PM entry for the One-way Group Availability Performance Metric.

[22] The Class of Service Name used is the CoS_Name used in the CN 5-tuple in which the Performance Metric resides.

Item	Value
Performance Metric Name	One-way Group Availability Performance Metric
$G = \{S_1, \dots, S_m\}$	Non-empty subsets of ordered EVC EP pairs
K	An integer ≥ 1
$\widehat{GA}$	Performance Objective expressed as percentage

Table 14 – *PM* Entry for the One-way Group Availability Performance Metric

[R48] The SLS **MUST** define the One-way Group Availability Performance Objective as met over T_l for a PM entry of the form in Table 14 if and only if $\overline{GA}(G, K, T_l) \geq \widehat{GA}$.

8.9 EVC Group Membership Service Attribute

The EVC Group Membership Service Attribute is used to specify an instance of the Multiple EVC Service Level Specification Service Attribute (Section 11), if any, in which the EVC is used. The value of the EVC Group Membership Service Attribute is *None* or a non-empty list of 3-tuples of the form $\langle ID, CoS\_Name\_G, SG \rangle$ where:

- ID is a string that is one of the values in an instance of the Multiple EVC Service Level Specification Service Attribute (Section 11).
- CoS_Name_G is an entry in the value of the EVC List of Class of Service Names Service Attribute (Section 8.7) that is not *Discard*.
- SG is a subset of ordered EVC EP pairs constructed from the value of the EVC List of EVC EPs Service Attribute (Section 8.2).

[R49] In each ordered EVC EP pair element of SG, at least one of the EVC EPs **MUST** have the value of the EVC EP Role Service Attribute (Section 10.3) = *Root*.

[R50] When the value of the EVC Group Membership Service Attribute is a non-empty list of 3-tuples, each 3-tuple, $\langle ID, CoS\_Name\_G, SG \rangle$, **MUST** be such that there exists a 5-tuple, $\langle CoS\_Name, \Delta t, C, n, PM \rangle$, in CN in the value of the EVC Service Level Specification (Section 8.8) with $CoS\_Name\_G = CoS\_Name$.

Note that [R50] means that if the value of the EVC Service Level Specification Service Attribute = *None*, then the value of the EVC Group Membership Service Attribute has to be *None*.

[R51] When the value of the EVC Group Membership Service Attribute is a non-empty list of 3-tuples, a given value of ID **MUST** appear in at most one 3-tuple in the list.

[R51] means that an EVC can appear at most once in an instance of the Multiple EVC Service Level Specification Service Attribute (Section 11).

When the value is *None*, the EVC is not used as part of an instance of the Multiple EVC Service Level Specification Service Attribute (Section 11). The EVC can be used in more than one instance of the Multiple EVC Service Level Specification Service Attribute (Section 11) by having more than one value of ID in the list of 3-tuples.

Note that Section 11 constrains the values of the EVC Service Level Specification Service Attribute for all EVCs that have a 3-tuple containing a given ID.

8.10 EVC Maximum Service Frame Size Service Attribute

The value of the EVC Maximum Service Frame Size Service Attribute is an integer ≥ 1522.

[R52] The value of the EVC Maximum Service Frame Size Service Attribute **MUST** be less than or equal to the value of the Subscriber UNI Maximum Service Frame Size Service Attribute (Section 9.8) for all UNIs in the EVC.

[D1] An ingress C-Tagged Service Frame that is mapped to an EVC EP that is in the EVC and whose length in bytes exceeds the value of the EVC Maximum Service Frame Size Service Attribute **SHOULD** be discarded.

[D2] An ingress Untagged Service Frame that is mapped to an EVC EP that is in the EVC and whose length in bytes exceeds the value of the EVC Maximum Service Frame Size Service Attribute minus 4 **SHOULD** be discarded.

8.11 EVC Available MEG Level Service Attribute

The value of the EVC Available MEG Level Service Attribute is an integer from 0 to 7 or *None*. If it is an integer, its value is one MEG Level higher than any MEG level reserved by the Service Provider for MEGs with MEPs that are contained entirely within the SP Network. A value of *None* indicates that MEG level 7 is reserved by the Service Provider for a MEG with MEPs that are contained entirely within the SP Network. If the EVC Available MEG Level is an integer, SOAM Service Frames at or above the EVC Available MEG Level are not peered or discarded by MEPs within the SP Network, but they may be processed by Subscriber MEG MIPs (Section 10.13). SOAM Service Frames at or above the EVC Available MEG Level (if it is an integer) and of a type that is passed transparently by a Subscriber MEG MIP are considered to be Data Service Frames (Section 7.6.3), and consequently the frame delivery requirements of Section 8.4 and the frame transparency requirements of Section 7.7 apply to them. SOAM Service Frames below the EVC Available MEG Level, or all SOAM Service Frames if the value is *None*, might be peered or discarded by MEPs within the SP Network.

9. Subscriber UNI Service Attributes

This section describes Service Attributes at each UNI.[23] Table 32 contains a list of the Subscriber UNI Service Attributes and their possible values.

A UNI can have a number of characteristics that are important to the way that the Subscriber experiences a service. One of the key aspects of a service description is the allowable mix of different characteristics for the UNIs in the EVC.

9.1 Subscriber UNI ID Service Attribute

The value of the Subscriber UNI ID Service Attribute is a string that is used to allow the Subscriber and Service Provider to uniquely identify the UNI for operations purposes.

[R53] The value of the Subscriber UNI ID Service Attribute **MUST** be unique among all UNIs for the SP Network.

[R54] The value of the Subscriber UNI ID Service Attribute **MUST** contain no more than 45 characters.

[R55] The value of the Subscriber UNI ID Service Attribute **MUST** be a non-null RFC 2579 [15] DisplayString but not contain the characters 0x00 through 0x1f.

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

Note that [R53] 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 [R53].

9.2 Subscriber UNI Instantiation Service Attribute

The value of the Subscriber UNI Instantiation Service Attribute is either *Physical* or *Virtual*.

When the value is *Physical*, the UNI is implemented using one or more instances of a standard Physical Layer per [R64].

When the value is *Virtual*, the physical layer is not specified. Figure 25 shows an example where Subscriber and Service Provider UNI-related functions are implemented in two different Virtual Machines inside a Server. The UNI lies between the two Virtual Machines. The connection from the Server to the rest of the Subscriber's network (to the left) and the connection from the Server to the rest of the SP Network (to the right) are beyond the scope of this document.

[23] In MEF 10.3, these Service Attributes are called "UNI Service Attributes". In this document, these have been renamed "Subscriber UNI Service Attributes" whose values are agreed to by the Subscriber and Service Provider to distinguish them from other UNI related Service Attributes such as the Operator UNI Service Attributes in MEF 26.2 [31] whose values are agreed to by the Operator and Service Provider/Super Operator.

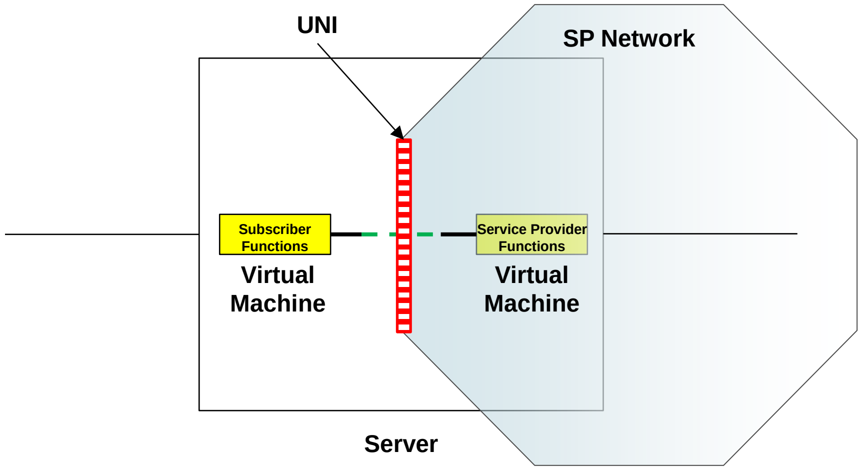

Figure 25 – Example of the Subscriber UNI Instantiation Service Attribute = *Virtual*

When the value of the Subscriber UNI Instantiation Service Attribute = *Virtual*, each instance of the information passed across the UNI is called a Virtual Frame and it is said that the Virtual Frame crosses the UNI. In the case of Figure 25, examples of the Virtual Frame include:

- A pointer to a common memory location that contains information relevant to the Subscriber Ethernet Service, and
- A block of bits transferred over a backplane.

The details of the Virtual Frame are beyond the scope of this document but the following requirements apply.

[R56] When the value of the Subscriber UNI Instantiation Service Attribute = *Virtual*, there **MUST** exist a map that maps the set of Virtual Frames that cross the UNI to a sequence of pairs of the form $\langle s, t \rangle$ where s is a standard Ethernet frame per Clause 3 of IEEE Std 802.3 – 2015 [5] and t is the arrival time at the UNI for all bits in s.

[R57] When the value of the Subscriber UNI Instantiation Service Attribute = *Virtual*, if the result of applying the map referred to in [R56] is $\{\langle s_k, t_k \rangle, k = 0,1,2, \ldots.\}$, then the following **MUST** hold:

$$t_{k+1} \geq t_k, k = 0,1,2, \ldots.$$

In other words, applying the map of [R56] to a set of Virtual Frames yields a corresponding sequence of IEEE Std 802.3 – 2015 [5] Ethernet frames, each with an associated Service Frame Arrival Time and the Arrival Times are monotonically increasing. See the Subscriber UNI Virtual Frame Map Service Attribute (Section 9.3).

9.3 Subscriber UNI Virtual Frame Map Service Attribute

The value of the Subscriber UNI Virtual Frame Map Service Attribute is either *Not Applicable* or a map that meets [R56] and [R57].

[R58] When the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Physical*, the value of the Subscriber UNI Virtual Frame Map Service Attribute **MUST** be *Not Applicable*.

[R59] When the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Virtual*, the value of the Subscriber UNI Virtual Frame Map Service Attribute **MUST** be a map that meets [R56] and [R57].

When the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Physical*, the Service Frame for the UNI is the first bit of the Destination Address through the last bit of the Frame Check Sequence of the IEEE Std 802.3 – 2015 [5] Ethernet frame that passes over a standard Physical Layer per [R64]. When the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Virtual*, the Service Frame is the first bit of the Destination Address through the last bit of the Frame Check Sequence of the IEEE Std 802.3 – 2015 [5] Ethernet frame that results from applying the value of the Subscriber UNI Virtual Frame Map Service Attribute.

The details of possible values of the Subscriber UNI Virtual Frame Map Service Attribute are beyond the scope of this document but, like all Service Attribute values, need to be agreed to by the Subscriber and Service Provider.

9.4 Subscriber UNI List of Physical Links Service Attribute[24]

The value of the Subscriber UNI List of Physical Links Service Attribute is either *Not Applicable* or a non-empty list of 4-tuples of the form $\langle id, pl, fs, pt\rangle$, with one list item for each physical link. The value of id is an identifier for the physical link. The value of pl specifies a physical layer. fs indicates if synchronous Ethernet is used on the physical link corresponding to the 4-tuple and has the value either *Enabled* or *Disabled*. The value of pt indicates if the Precision Time Protocol is used on the physical link corresponding to the 4-tuple and has the value either *Enabled* or *Disabled*.

[R60] The value of the Subscriber UNI List of Physical Links Service Attribute **MUST** be a non-empty list of 4-tuples if and only if the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Physical*.

[R61] When the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Physical*, the value of id in each $\langle id, pl, fs, pt\rangle$ 4-tuple **MUST** be unique among all physical links at UNIs for the SP Network.

[24] In MEF 10.3 [24] this Service Attribute is called Physical Layer Service Attribute.

[R62] When the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Physical*, the value of id in each $\langle id, pl, fs, pt\rangle$ 4-tuple **MUST** contain no more than 45 characters.

[R63] When the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Physical*, the value of id in each $\langle id, pl, fs, pt\rangle$ 4-tuple **MUST** be a non-null RFC 2579 [15] DisplayString but not contain the characters 0x00 through 0x1f.

[R64] When the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Physical*, the value of pl in each $\langle id, pl, fs, pt\rangle$ 4-tuple **MUST** be one of the PHYs specified in:

- IEEE Std 802.3 – 2015 [5] but excluding 1000BASE-PX-D and 1000BASE-PX-U,
- IEEE Std 802.3bp – 2016 [6],
- IEEE Std 802.3bq – 2016 [7],
- IEEE Std 802.3bs – 2017 [8],
- IEEE Std 802.3bu – 2016 [9],
- IEEE Std 802.3bv – 2017 [10],
- IEEE Std 802.3by – 2016 [11],
- IEEE Std 802.3bz – 2016 [12], and
- IEEE Std 802.3cc – 2017 [13].

[R65] When the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Physical*, the Physical Layer specified by each value of pl **MUST** operate in full duplex mode.

Note that when the Subscriber UNI Instantiation Service Attribute = *Physical*, different physical links implementing the UNI can use different physical layers.

When the value of fs in a 4-tuple is *Enabled*, synchronous Ethernet is used on the physical link corresponding to the 4-tuple.

[R66] When the value of fs = *Enabled* in a 4-tuple in the value of Subscriber UNI List of Physical Links Service Attribute, synchronous Ethernet as defined in ITU-T G.8262/Y.1362 [18] and ITU-T G.8264/Y.1364 [19] **MUST** be used on the corresponding physical link.

The accuracy of the frequency reference provided on a physical link with fs = *Enabled* is beyond the scope of this document.

The use of the Ethernet Synchronization Message Channel (ESMC) can be agreed by including it in the value of the L2CP Peering Service Attribute (Section 9.17).

[R67] When the value of pt = *Enabled* in a 4-tuple in the value of Subscriber UNI List of Physical Links Service Attribute, the Precision Time Protocol as specified in IEEE Std 1588-2008 [1] **MUST** be used on the corresponding physical link per Annex F of that specification, such that the SN can derive a Primary Reference Time Clock traceable time synchronization reference from the SP Network.

[R68] When the value of pt = *Disabled* in a 4-tuple in the value of Subscriber UNI List of Physical Links Service Attribute, the Service Provider **MUST NOT** run the Precision Time Protocol as specified in IEEE Std 1588-2008 [1] on the corresponding physical link.

The accuracy of the time reference provided on a physical link with pt = *Enabled* is beyond the scope of this document.

When the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Physical*, a UNI will contain one or more physical links. When multiple physical links are configured at a UNI, the individual links may terminate at the same device in the SP Network and/or in the SN, or at different devices in the SP Network and/or in the SN. The configuration for the termination of the physical link or links that implement the UNI is beyond the scope of this document.

When the value of the Subscriber UNI List of Physical Links Service Attribute is a non-empty list of 4-tuples with more than one 4-tuple in the list, a resiliency mechanism is required and is identified by the Subscriber UNI Link Aggregation Service Attribute specified in Section 9.5.

9.5 Subscriber UNI Link Aggregation Service Attribute[25]

The value of the Subscriber UNI Link Aggregation Service Attribute is one of *2-Link Active/Standby*, *All Active*, *Other*, or *Not Applicable*. The value of this Service Attribute is dependent on the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) and the value of the Subscriber UNI List of Physical Links Service Attribute (Section 9.4). See Appendix F for some configuration examples for multiple physical links.

[R69] If the value of the Subscriber UNI List of Physical Links Service Attribute (Section 9.4) = a non-empty list containing a single 4-tuple, then the value of the Subscriber UNI Link Aggregation Service Attribute **MUST** be set to *Not Applicable*.

[R70] If the value of the Subscriber UNI List of Physical Links Service Attribute (Section 9.4) = a non-empty list containing two 4-tuples, then the value of the Sub-

[25] In MEF 10.3 [24] this is called UNI Resiliency Service Attribute.

scriber UNI Link Aggregation Service Attribute **MUST** be set to one of *2-Link Active/Standby*, *All Active*, or *Other*.

[R71] If the value of the Subscriber UNI List of Physical Links Service Attribute (Section 9.4) = a non-empty list containing three or more 4-tuples, then the value of the Subscriber UNI Link Aggregation Service Attribute **MUST** be set to either *All Active* or *Other*.

Note that [R69], [R70], and [R71] only apply when the value of the Subscriber Instantiation Service Attribute = *Physical*. Table 15 summarizes the above requirements.

Number of 4-tuples*	*Not Applicable*	*Active/Standby*	*All Active*	*Other*
1	☑	☐	☐	☐
2	☐	☑	☑	☑
3 or more	☐	☐	☑	☑

* Applies when the value of the Subscriber UNI Instantiation Service Attribute = *Physical*

Table 15 – Allowed Values of the Subscriber UNI Link Aggregation Service Attribute

[R72] When the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Virtual*, the value of the Subscriber UNI Link Aggregation Service Attribute **MUST** be *Not Applicable*.

The following requirements depend on the value of the Subscriber UNI Link Aggregation Service Attribute.

[R73] When the value of the Subscriber UNI Link Aggregation Service Attribute = *2-Link Active/Standby*, the SP Network **MUST** implement Link Aggregation as in either Clause 5.6.1 of IEEE Std 802.1AX – 2008 [2] or Clause 6.7.1 of IEEE Std 802.1AX – 2014 [3] with one Link Aggregation Group (LAG) across the links supporting the UNI that is configured such that all Service Frames are carried on only one of the two links when both links are operational.

[R74] When the value of the Subscriber UNI Link Aggregation Service Attribute = *All Active* the SP Network **MUST** use Link Aggregation as specified in Clause 5.3 of IEEE Std 802.1AX-2014 [3], including the use of the version 2 LACPDUs as specified in Clause 5.3.1h of IEEE Std 802.1AX-2014 [3], with one Link Aggregation Group (LAG) across the links supporting the UNI.

[R75] When the value of the Subscriber UNI Link Aggregation Service Attribute = *All-Active*, the SP Network **MUST** use “Per-service frame distribution” as specified in Clause 8.2 of IEEE Std 802.1AX-2014 [3], where the Port Conversation ID is equal to the C-Tag VLAN ID for VLAN Tagged Service Frames and equal to 0 for Untagged Service Frames and Priority Tagged Service Frames.

[R75] means that all Service Frames with a given C-Tag will traverse the same link.

[R76] When the value of the Subscriber UNI Link Aggregation Service Attribute = *All-Active*, the SP Network **MUST** be configured such that there is only one aAggActorAdminKey that has the same value as the aAggPortActorAdminKey for the ports terminating the links at the UNI.

The aAggActorAdminKey and aAggPortActorAdminKey are managed objects defined in IEEE Std 802.1AX-2014 [3]. Ensuring that there is only one aAggActorAdminKey with the same value as the aAggPortActorAdminKey for the ports terminating the links at the UNI assures that only a single Link Aggregation Group is formed at the UNI. This eliminates the possibility of any loops potentially arising from multiple UNI links coming up independently or forming separate Link Aggregation Groups. Note that for Link Aggregation to operate correctly at the UNI, the SN needs to be configured so that there is only one aAggActorAdminKey that has the same value as the aAggPortActorAdminKey for the ports terminating the links at the UNI.

[O2] The Subscriber and the Service Provider **MAY** agree to any other resiliency mechanism when the value of the Subscriber UNI Link Aggregation Service Attribute = *Other*.

The other resiliency mechanism referred to in [O2] is beyond the scope of this document.

9.6 Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute

The value of the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute is either a Port Conversation ID to Aggregation Link Map as defined in IEEE Std 802.1AX – 2014 [3] or *Not Applicable*.

[R77] The value of the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute **MUST** be *Not Applicable* when the value of the Subscriber UNI Link Aggregation Service Attribute (Section 9.5) does not equal *All Active* or *Other*.

[R78] When the value of the Subscriber UNI Link Aggregation Service Attribute (Section 9.5) = *All Active*, the value of the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute **MUST** be a Port Conversation ID to Aggregation Link Map as defined in IEEE Std 802.1AX – 2014 [3].

The Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute value is the mapping of each Port Conversation ID (see [R75]) to a Link Selection Priority List at the UNI. The Link Selection Priority List is a sequence of Link Number IDs[26], in the order of usage preference, highest to lowest, for the link that is to carry the Service Frames corresponding to a given Port Conversation ID. The value of a Link Number ID has local significance to the LAG at a given UNI.

[R79] When the value of the Subscriber UNI Link Aggregation Service Attribute (Section 9.5) = *All Active*, the set of Link Number IDs as defined in IEEE Std 802.1AX – 2014 [3] at the UNI **MUST** be $\{1,2,\ldots,m\}$ where m is the number of 4-tuples in the value of the Subscriber UNI List of Physical Links Service Attribute (Section 9.4).

[26] A Link Number ID is an integer ≥ 1 that is uniquely assigned to each physical link at a given UNI.

[R79] mandates the value of Link Number IDs that are used in the value of the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute; this avoids the need to negotiate the values between the Service Provider and Subscriber for a given UNI. The Service Provider and Subscriber do not need to agree on an association of each Link Number ID to a physical link (or the physical port terminating the link) as this association is made during the operation of LACP. However, although not required to do so, the Service Provider and Subscriber can agree on an association of each Link Number ID to a physical link, which could be useful if there is a preference for which physical link carries specific Service Frames.

The Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute is required when the Subscriber UNI Link Aggregation Service Attribute is set to *All Active* and can be used when the Subscriber UNI Link Aggregation Service Attribute is set to *Other*. However the use of the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute for the latter case is beyond the scope of this document.

Note that the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute is equivalent to the aAggConversationAdminLink[] that is defined in Clause 7.3.1 of IEEE Std 802.1AX-2014 [3].

The distribution of Service Frames across the different physical links at a given UNI is based on the agreed value of the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute. If the first link in the Link Selection Priority List for a given Port Conversation ID is operational[27] in the LAG, all of the Service Frames with the corresponding Port Conversation ID are carried on that link in both directions. If the first link is not operational, then the second link in the list is used if the second link is operational, and so on. If all links in the list fail, the Service Frames with the corresponding Port Conversation ID are not carried over the UNI in either direction, even if a link that is not in the list is still operational.

[R80] For a given C-Tag VLAN ID value that is mapped to an EVC EP (Section 10.4), the Link Selection Priority List in the value of the Port Conversation ID to Aggregation Link Map Service Attribute **MUST NOT** be empty.

The number of links in a Link Selection Priority List in the value of the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute is, by definition, less than or equal to the number of physical links for that UNI. A shorter list results in lower resilience for the Service Frames corresponding to the Port Conversation ID. Note that of a Port Conversation ID does not map to an EVC EP, it may have an empty Link Selection Priority List in the value of the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute, in which case Service Frames with the corresponding Port Conversation ID are not carried across the UNI.

If a particular Link Number ID is in a Link Selection Priority List in the value of the Subscriber UNI Port Conversation ID to Link Aggregation Map Service Attribute, but not the first link in any list in the attribute, then the physical link associated with that Link Number ID does not car-

[27] In the interest of brevity, the term "operational" is used as shorthand for MAC_Operational status = TRUE and the term "not operational" is used as shorthand for MAC_Operational status = FALSE. See Clause 6.3.12 in IEEE Std 802.1AX-2014 [3].

ry any Service Frames if all other links at the UNI are operational. In this case, the link can be considered as a “backup link” that is reserved for protection against failure of another link.

Table 16 illustrates an example of a value of the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute that contains three physical links with three Link Number IDs, 1, 2, and 3.

Port Conversation ID	**Link Selection Priority List (decreasing order)**
0, 1, 4	1, 3, 2
5	2, 3, 1
10	2, 1, 3
1000	2, 1
All other values	

Table 16 – Example Value of a Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute

In this example, six Port Conversation IDs have a non-empty Link Selection Priority List while other Port Conversation IDs have an empty Link Selection Priority List at the UNI. As shown in Table 16, the Link Selection Priority List for Port Conversation IDs 0, 1, and 4 contains Link Number IDs 1, 3, and 2 in the sequence. The Link Selection Priority List for Port Conversation ID 5 has 2, 3, and 1; the list for Port Conversation ID 10 has 2, 1, and 3; the list for Port Conversation ID 1000 has 2 and 1. In this example, link 3 is not used when both link 1 and 2 are operational. Thus link 3 is used for protection purposes. The example also indicates that the Service Frames corresponding to Port Conversation IDs 5 and 10 are carried over link 2 when link 2 is operational; when link 2 is not operational and link 1 and 3 are operational, the Service Frames with Port Conversation ID 5 are carried over link 3 and the Service Frames with Port Conversation ID 10 are carried over link 1. The Service Frames with Port Conversation ID 1000 in the example have less resilience than the Service Frames corresponding to Port Conversation IDs 0, 1, 4, 5, and 10. Service Frames with a Port Conversation ID not 0, 1, 4, 5, 10, or 1000 cannot traverse the UNI.

Note that the Table 16 is an abstract description for the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute. This description does not constrain how the contents can be described in a protocol, database, service order form, etc.

The value in the Subscriber UNI Port Conversion ID to Aggregation Link Map Service Attribute is only used for Service Frame distribution at a given UNI. Which EVC EP a Service Frame is mapped to at the UNI is determined by the value of the EVC EP Map Service Attribute (Section 10.4).

[O3] If the value of the Subscriber UNI Link Aggregation Service Attribute (Section 9.5) = *All Active* and if for a given EVC EP the value of the EVC EP Map Service Attribute (Section 10.4) = *All* or *List*, then the Service Provider **MAY** support a value of the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute such that Service Frames with different Port Conversation ID values that are mapped to the given EVC EP can be carried on different physical links.

The value of the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute described in [O3] is useful when the bandwidth needed for the EVC EP exceeds the capacity of a single physical link at the UNI. However, in certain configurations (for example when the links terminate on different devices in the SP Network), supporting such a map could require a Service Provider to make tradeoffs between the Service Frame distribution and the application of MEF SOAM and Bandwidth Profiles.

[R81] When the value of the Subscriber UNI Link Aggregation Service Attribute (Section 9.5) = *All Active* and Ingress Bandwidth Profiles and/or Egress Bandwidth Profiles are used at the UNI, the Service Provider **MUST** support a value in the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute such that all Service Frames that map to a given Envelope (Section 12) are carried on the same physical link.

Note that when Service Frames that map to a given Envelope are carried on different links, it may be difficult to apply the Bandwidth Profile algorithm in the SP Network, and it may be difficult for the Subscriber to apply shaping in the SN, especially if the different links happen to terminate on different devices. The Service Provider can offer a value in the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute where Service Frames that map to a given Envelope are carried on different links, if they have the capability to apply the Bandwidth Profile algorithm to such frames (or if there is no Bandwidth Profile configured at the UNI). However, [R81] requires the Service Provider to also support a map where Service Frames that map to a given Envelope are carried on a single link.

Appendix G contains information on how Service Frames are distributed across multiple links at the UNI based on the value of the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute.

9.7 Subscriber UNI Service Frame Format Service Attribute

The value of the Subscriber UNI Service Frame Format Service Attribute specifies the allowed formats of Service Frames. It has a single value as specified in [R82].

[R82] The format of the Service Frame **MUST** be that of the MAC frame that is specified in Clause 3 of IEEE Std 802.3 – 2015 [5].

[R82] constrains the value of the Subscriber UNI Virtual Frame Map when the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Virtual*. This is because the Service Frame is the result of applying the value of the Subscriber UNI Virtual Frame Map Service Attribute (Section 9.3) to the Virtual Frame when the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Virtual*.

Note that [R82] means that Service Frames will be discarded by the receiving SP Network if they are not properly constructed. For example, a Service Frame with an incorrect Frame Check Sequence will be discarded. However, this document provides for Service Frames that are longer than the maximum specified in IEEE Std 802.3 – 2015 [5]. See Section 9.8.

9.8 Subscriber UNI Maximum Service Frame Size Service Attribute

The value for the Subscriber UNI Maximum Service Frame Size Service Attribute is an integer number of bytes ≥ 1522.

Note that the value of the Subscriber UNI Maximum Service Frame Size Service Attribute constrains the value of the EVC Maximum Service Frame Size Service Attribute (Section 8.10) via [R52].

9.9 Subscriber UNI Maximum Number of EVC EPs Service Attribute

The value of the Subscriber UNI Maximum Number of EVC EPs Service Attribute is an integer ≥ 1 that limits the number of EVC EPs that can be located at the UNI. This is a Capability Service Attribute (Section 7.3).

[R83] The number of EVC EPs for which the value of the EVC EP UNI Service Attribute (Section 10.2) is a given UNI **MUST** be less than or equal to the value of the Subscriber UNI Maximum Number of EVC EPs Service Attribute at that UNI.

Figure 26 shows an example of the value of the Subscriber UNI Maximum Number of EVC EPs Service Attribute. In this example, the value of the Subscriber UNI Maximum Number of EVC EPs = 7 at UNI A, 3 at UNI B, 1 at UNI C, and 1 at UNI D. The number of EVC EPs located at UNI A can be as many as 7 but in this example it is only 3. Similarly the number of EVC EPs located on UNI B is less than the maximum number.

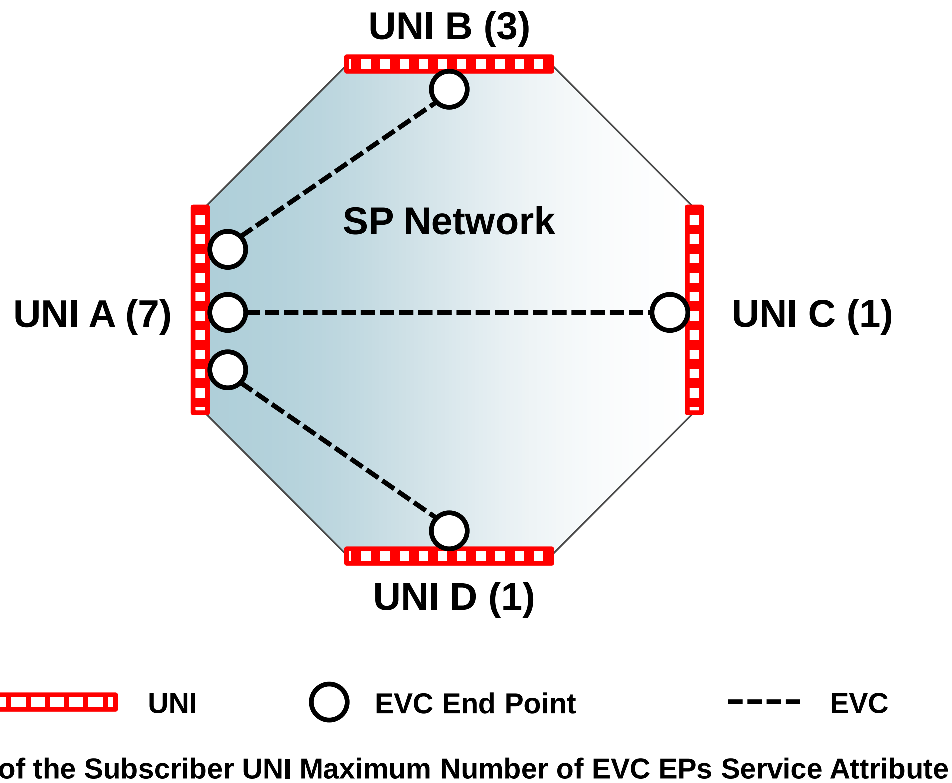

Figure 26 – Example of the Subscriber UNI Maximum Number of EVC EPs Service Attribute

For ease of description of configurations like that in Figure 26, the term "Service Multiplexing" is introduced. There are two definitions for Service Multiplexing:

- **Definition 1:** The condition when there is more than one EVC End Point located at a UNI.
- **Definition 2:** The condition where the value of the Subscriber UNI Maximum Number of EVC End Points Service Attribute is > 1 at a UNI.

In Figure 26, UNI A has Service Multiplexing in the sense of Definition 1. In Figure 26 UNI A and UNI B have Service Multiplexing in the sense of definition 2.

9.10 Subscriber UNI Maximum Number of C-Tag VLAN IDs per EVC EP Service Attribute

The value of the Subscriber UNI Maximum Number of C-Tag VLAN IDs per EVC EP Service Attribute affects behavior only when the value of the EVC EP Map Service Attribute = *List* for EVC EPs located at the UNI (see [R109]). The value of the Subscriber UNI Maximum Number of C-Tag VLAN IDs per EVC EP Service Attribute is an integer ≥ 1. This is a Capability Service Attribute.

For ease of description, the term "Bundling" is introduced. There are two definitions for Bundling:

- **Definition 1**: The condition when there is more than one C-Tag VLAN ID value mapped to an EVC EP at a UNI, or
- **Definition 2**: The condition where the value of the Subscriber UNI Maximum Number of C-Tag VLAN IDs per EVC End Point Service Attribute is > 1 at a UNI.

9.11 Subscriber UNI Token Share Service Attribute

The value of the Subscriber UNI Token Share Service Attribute is either *Enabled* or *Disabled*. This is a Capability Service Attribute (Section 7.3) that indicates whether Bandwidth Profile Envelopes containing more than one Bandwidth Profile Flow (Section 12) are supported by the Service Provider at the UNI (Section 12.3).

[R84] When the value of the Subscriber UNI Token Share Service Attribute = *Enabled*, the Service Provider **MUST** support at least one Envelope with two or more Bandwidth Profile Flows mapped to it.

[D3] When the value of the Subscriber UNI Token Share Service Attribute = *Enabled*, the Service Provider **SHOULD** support mapping two or more Bandwidth Profile Flows to every Envelope.

[R85] When the value of the Subscriber UNI Token Share Service Attribute = *Disabled*, every Envelope at the UNI **MUST** have exactly one Bandwidth Profile Flow mapped to it.

See Section 12 for the descriptions of Envelope and Bandwidth Profile Flow.

9.12 Subscriber UNI Envelopes Service Attribute

The value of the Subscriber UNI Envelopes Service Attribute is *None* or a non-empty list of pairs of the form $\langle x, y \rangle$, where x is an Envelope ID value and y is the Envelope Coupling Flag (CF^0) value.

Table 17 shows the parameters for each pair. Note that the descriptions in the table are informal. The precise role played by a given parameter is specified by the Bandwidth Profile Algorithm (Section 12.2).

Parameter Name	Symbol	Units	Informal Description
Envelope ID	–	RFC 2579 [15] DisplayString	Identifies the Envelope
Envelope Coupling Flag	CF^0	Integer	Determines whether unused committed bandwidth* for the lowest-ranked Bandwidth Profile Flow is made available as excess bandwidth for the highest-ranked Bandwidth Profile Flow or remains unused.

* The unused committed bandwidth for a Bandwidth Profile Flow is the difference between the limit resulting from the CIR for that Bandwidth Profile Flow (taking into account unused bandwidth from higher-ranked Bandwidth Profile Flows if applicable), and the actual information rate of Service Frames for that Bandwidth Profile Flow that are declared Green.

Table 17 – Envelope Parameters

[R86] The Envelope ID **MUST** contain no more than 45 characters.

[R87] The Envelope ID **MUST** be a non-null RFC 2579 [15] DisplayString but not contain the characters 0x00 through 0x1f.

[R88] CF^0 **MUST** have a value of 0 or 1.

[R89] When one Bandwidth Profile Flow (Section 12) is mapped to an Envelope, CF^0 **MUST** equal 0.

[R90] The value of the Subscriber UNI Envelopes Service Attribute **MUST** contain at most one entry with a given Envelope ID value.

Note that the pair <value of the Subscriber UNI ID Service Attribute (Section 9.1), Envelope ID> is unique among all Envelopes for the SP Network.

[R91] The Envelope ID value in the Envelope and Rank Parameter (ER) for each Bandwidth Profile Flow (Section 12) specified at the UNI **MUST** match an Envelope ID value in an entry in the value of the Subscriber UNI Envelopes Service Attribute.[28]

9.13 Subscriber UNI Link OAM Service Attribute

The value of the Subscriber UNI Link OAM Service Attribute is either *Enabled* or *Disabled*. The Subscriber UNI Link OAM Service Attribute controls when and how Link OAM per IEEE Std 802.3-2015 [5] is run on the physical links in the UNI.

[R92] When the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Virtual*, the value of the Subscriber UNI Link OAM Service Attribute **MUST** be *Disabled*.

[28] Note that [R5] of MEF 6.2 [22] conflicts with this requirement.

[R93] When the value of the Subscriber UNI Link OAM Service Attribute = *Enabled*, Link OAM as specified in Clause 57 of IEEE Std 802.3-2015 [5] **MUST** be run on all physical links in the UNI.

[R94] When the value of the Subscriber UNI Link OAM Service Attribute = *Enabled*, the Service Provider **MUST** enable Active DTE (Data Termination Equipment) mode capabilities as specified in clause 57.2.9 of IEEE Std 802.3 – 2015 [5] on each link in the UNI.

[D4] When the value of the Subscriber UNI Link OAM Service Attribute = *Enabled*, Link Events as specified in Clauses 57.2.10 and 57.4.3.2 of IEEE Std 802.3 – 2015 [5] **SHOULD** be enabled on each link in the UNI.

9.14 Subscriber UNI MEG Service Attribute

The Subscriber UNI MEG Service Attribute indicates if the Service Provider has instantiated a MEG End Point at the UNI Maintenance Entity Group (MEG) Level. The value of the Subscriber UNI MEG Service Attribute is either *Enabled* or *Disabled*.

[R95] When the value of the Subscriber UNI MEG Service Attribute = *Enabled*, the Service Provider **MUST** instantiate a UNI MEG MEP.

[R96] When the value of the Subscriber UNI MEG Service Attribute = *Disabled*, the Service Provider **MUST NOT** instantiate a UNI MEG MEP.

When the value of the Subscriber UNI MEG Service Attribute = *Enabled*, several parameter values need to be agreed upon by the Subscriber and the Service Provider.

9.15 Subscriber UNI LAG Link MEG Service Attribute

The value of the Subscriber UNI LAG Link MEG Service Attribute is either *Enabled* or *Disabled*.

[R97] When the value of the Subscriber UNI LAG Link MEG Service Attribute is *Enabled*, the Service Provider **MUST** operate the LAG Link MEG on each link in the UNI.

[R98] If the value of the Subscriber UNI List of Physical Links Service Attribute (Section 9.4) is *Not Applicable* or is a list containing a single 4-tuple, then the value of the Subscriber UNI LAG Link MEG Service Attribute **MUST** be *Disabled*.

[R99] When the value of the Subscriber UNI LAG Link MEG Service Attribute = *Disabled*, the Service Provider **MUST NOT** instantiate a UNI LAG Link MEG MEP.

When the value of the Subscriber UNI LAG Link MEG Service Attribute = *Enabled*, several parameter values need to be agreed upon by the Subscriber and the Service Provider.

9.16 Subscriber UNI L2CP Address Set Service Attribute

The Subscriber UNI L2CP Address Set Service Attribute is defined in MEF 45.1 [36].

9.17 Subscriber UNI L2CP Peering Service Attribute

The Subscriber UNI L2CP Peering Service Attribute is defined in MEF 45.1 [36].

10. EVC EP Service Attributes[29]

This section describes EVC EP Service Attributes. Table 34 contains a list of the EVC EP Service Attributes and their possible values.

An EVC EP is a construct at a UNI that selects a subset of the Service Frames that pass over the UNI. The subset of Service Frames is specified via the EVC EP Map Service Attribute as described in Section 10.4, and per [R108], is disjoint from the subsets selected by other EVC EPs at the same UNI. Per Section 7.8 an EVC is an association of EVC EPs. An EVC EP represents the logical attachment of an EVC to a UNI.

10.1 EVC EP ID Service Attribute

The value of the EVC EP ID Service Attribute is a string that is used to allow the Subscriber and Service Provider to uniquely identify the EVC EP for operations purposes.

[R100] The value of the EVC EP ID Service Attribute **MUST** be unique among all such identifiers for EVC EPs in the SP Network.

[R101] The value of the EVC EP ID Service Attribute **MUST** contain no more than 45 characters.

[R102] The value of the EVC EP ID Service Attribute **MUST** be a non-null RFC 2579 [15] DisplayString but not contain the characters 0x00 through 0x1f.

10.2 EVC EP UNI Service Attribute

The value of the EVC EP UNI Service Attribute is a Subscriber UNI ID Service Attribute value per Section 9.1. The value of the EVC EP UNI Service Attribute serves to specify the UNI where the EVC EP is located. The EVC EP is said to be at this UNI.

10.3 EVC EP Role Service Attribute

The value of the EVC EP Role Service Attribute is one of *Root* or *Leaf*.

For ease of exposition:

- An EVC EP with the value of the EVC EP Role Service Attribute = *Root* is called a Root EVC EP, and
- An EVC EP with the value of the EVC EP Role Service Attribute = *Leaf* is called a Leaf EVC EP.

Per the following requirement, an Ingress Service Frame that is mapped to a Leaf EVC EP can only be delivered to Root EVC EPs. There is no such restriction on a Root EVC EP.

[29] In MEF 10.3 [24] these attributes are called EVC per UNI Service Attributes.

[R103] When the value of the EVC EP Role Service Attribute = *Leaf*, then an Ingress Service Frame mapped to that EVC EP **MUST NOT** result in an Egress Service Frame mapped to any EVC EP which also has the value of the EVC EP Role Service Attribute = *Leaf*.

Note that because of [R15] and [R16], the value of the EVC EP Role Service Attribute (Section 8.3) will always = *Root* when the associating EVC is not of the type *Rooted-Multipoint*.

10.4 EVC EP Map Service Attribute

The value of the EVC EP Map Service Attribute is one of *List*, *All*, or *UT/PT*. When the value of the EVC EP Map Service Attribute = *List*, a non-empty list of integers, each in the range $1,2,\ldots,4094$ is specified.

The EVC EP Map Service Attribute is used to map Service Frames to the EVC EP.

[R104] An Ingress Service Frame that is not mapped to any EVC EP via the values of the EVC EP Map Service Attribute for the EVC EPs located at the ingress UNI **MUST NOT** result in a corresponding Egress Service Frame that is mapped to any EVC EP in the SP Network.

Note that [R104] does not preclude processing ingress SOAM and L2CP Service Frames that do not map to an EVC EP. Also note that [R104] does not preclude frames that do not map to an EVC EP from being mapped to some other service that is not covered by this document.

Requirements arising from each of the possible values of the EVC EP Map Service Attribute are presented in the following subsections.

10.4.1 Value of the EVC EP Map Service Attribute = *List*

[R105] If the value of the EVC EP Map Service Attribute = *List*, then all ingress VLAN Tagged Service Frames at the UNI where the EVC EP is located with a C-Tag VLAN ID value matching an entry in the value of the EVC EP Map Service Attribute **MUST** be mapped to the EVC EP.

[R105] means that Priority Tagged Service Frames are not mapped to the EVC EP when the value of the EVC EP Map Service Attribute = *List* since a Priority Tagged Service Frame is not a VLAN Tagged Service Frame (Section 7.5).

Note that if the value of the EVC EP Map Service Attribute = *List*, then an ingress VLAN Tagged Service Frame with 4095 as the value of the C-Tag VLAN ID field will not be mapped to the EVC EP because 4095 cannot be in the value of the EVC EP Map Service Attribute.

[R106] If the value of the EVC EP Map Service Attribute = *List*, then ingress Untagged Service Frames and ingress Priority Tagged Service Frames at the UNI where the EVC EP is located **MUST NOT** be mapped to the EVC EP.

[R107] If the value of the EVC EP Map Service Attribute = *List*, then all Egress Service Frames that were delivered to the EVC End Point and transmitted across the

UNI **MUST** be VLAN Tagged Service Frames and have a C-Tag VLAN ID value matching an entry in the value of the EVC End Point Map Service Attribute.

[R107] means that an egress VLAN Tagged Service Frame is required to have the C-Tag VLAN ID value in the C-Tag set to one of the values in the value of the EVC End Point Map Service Attribute. Note that the EVC EP Map Service Attribute applies to both Ingress and Egress Service Frames. For an Ingress Service Frame, the C-Tag VLAN ID value for the Service Frame and the value of the EVC EP Map Service Attribute enables the SP Network to know how to deliver the Service Frame. For an Egress Service Frame, the C-Tag VLAN ID value for the Service Frame and the value of the EVC EP Map Service Attribute allow the SN to know which EVC the Service Frame came from.

[R108] At each UNI, a given C-Tag VLAN ID value **MUST** be in at most one value of an EVC EP Map Service Attribute among all EVC EPs located at the UNI.

[R109] If the value of the EVC EP Map Service Attribute = *List*, then the number of C-Tag VLAN ID values in the value of the EVC EP Map Service Attribute for an EVC EP **MUST** be less than or equal to the value of the Subscriber UNI Maximum Number C-Tag of VLAN IDs per EVC EP Service Attribute (Section 9.10) for the UNI where the EVC EP is located.

[R110] For a given EVC EP, if the value of the EVC EP Map Service Attribute = *List* and more than one C-Tag VLAN ID value is in the list, then all of the EVC EPs in the same EVC as the given EVC EP **MUST** have the value of the EVC EP Map Service Attribute = *List* and with the same list entries.

Figure 27 shows some examples of the value of the EVC EP Map Service Attribute = *List*. In this figure, the list contains the same three C-Tag VLAN ID values for both EVC EP a and EVC EP c as mandated by [R110]. Also shown in Figure 27 is that a given C-Tag VLAN ID value can be reused in more than one EVC. In this case, C-Tag VLAN ID 47 is reused.

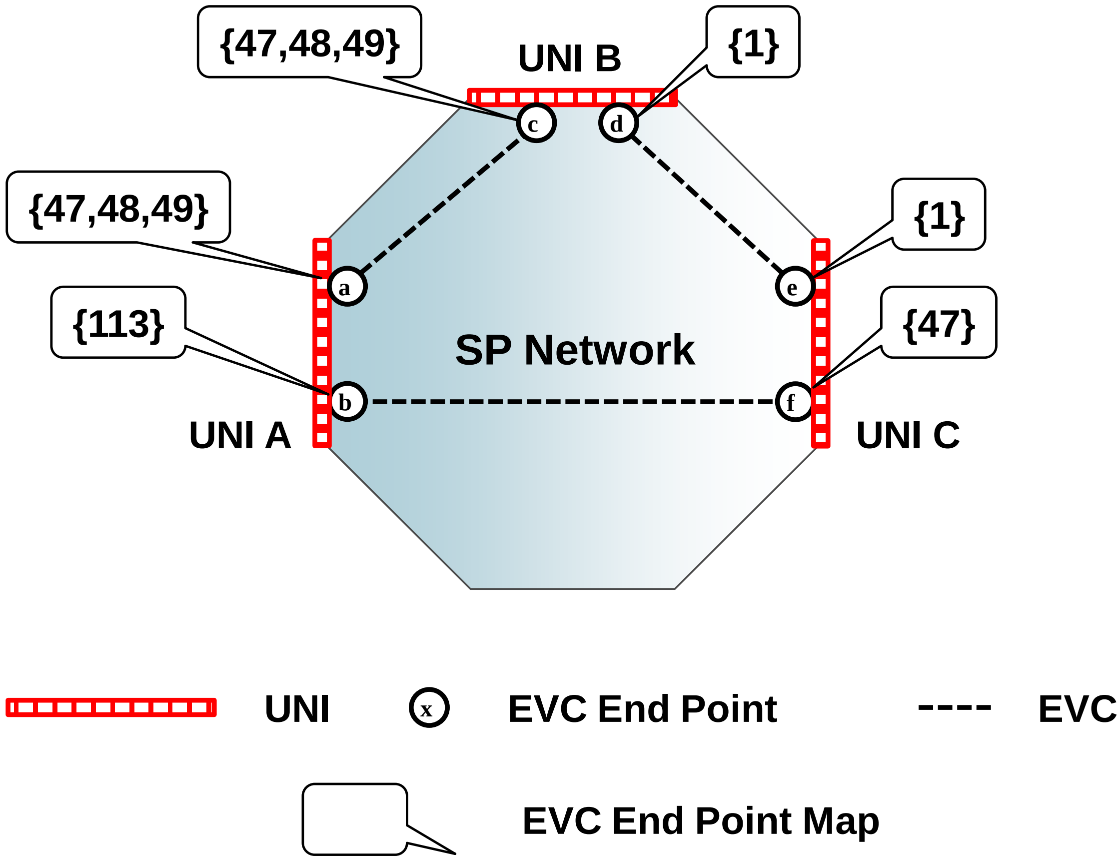

Figure 27 – Examples for the value of the EVC EP Map Service Attribute = *List*

Note that for a given UNI, the value of the EVC EP Map Service Attribute may be constrained by the range of C-Tag VLAN ID values that can be supported by the SN and the range of C-Tag VLAN ID values that can be supported by the SP Network.

10.4.2 Value of the EVC EP Map Service Attribute = *All*

[R111] If the value of the EVC EP Map Service Attribute = *All*, then all Service Frames at the UNI where the EVC EP is located **MUST** be mapped to the EVC EP.

[R111] mandates that an Ingress VLAN Tagged Service Frame with 4095 as the value of the C-Tag VLAN ID be mapped to the EVC EP. However, it is allowed that such an Ingress Service Frame not be delivered per the following optional requirement.

[O4] If the value of the EVC EP Map Service Attribute = *All*, then an Ingress Service Frame with 4095 in C-Tag VLAN ID **MAY** be discarded by the SP Network.

[R112] If the value of the EVC EP Map Service Attribute = *All*, then the EVC EP **MUST** be the only EVC EP at the UNI.

[R113] For a given EVC EP, if the value of the EVC EP Map Service Attribute = *All*, then all of the EVC EPs in the same EVC as the given EVC EP **MUST** have the value of the EVC EP Map Service Attribute = *All*.

10.4.3 Value of the EVC EP Map Service Attribute = *UT/PT*

[R114] If the value of the EVC EP Map Service Attribute = *UT/PT*, then all ingress Untagged Service Frames and all ingress Priority Tagged Service Frames at the UNI **MUST** be mapped to the EVC EP.

[R115] If the value of the EVC EP Map Service Attribute = *UT/PT*, then ingress VLAN Tagged Service Frames at the UNI **MUST NOT** be mapped to the EVC EP.

[R116] If the value of the EVC EP Map Service Attribute = *UT/PT*, then all Egress Service Frames that were delivered to the EVC End Point and transmitted across the UNI **MUST** be Untagged Service Frames.

[R117] If the value of the EVC EP Map Service Attribute = *UT/PT*, then the EVC EP **MUST** be the only EVC EP at the UNI.

10.4.4 Additional Requirements and Observations

For ease of discussion, three terms are defined:

- **VLAN Tagged UNI:** A UNI with one or more EVC EPs where the EVC EP Map Service Attribute value = *List* for all EVC EPs located at the UNI.
- **All to One Bundled UNI:** A UNI with a single EVC EP whose EVC EP Map Service Attribute value = *All*.
- **Untagged UNI:** A UNI with a single EVC EP whose EVC EP Map Service Attribute value = *UT/PT*.

[R112] and [R117] mean that exactly one of the above three terms applies to a given UNI where one or more EVC EPs are located. Figure 28 shows examples of the three terms. Behavior of each EVC in the figure:

- **EVC 1:** An Ingress Service Frame at UNI A with C-Tag VLAN ID = 44 is delivered to UNI B as an Untagged Service Frame. An Ingress Service Frame at UNI B that has no C-Tag is delivered to UNI A as a VLAN Tagged Service Frame with C-Tag VLAN ID = 44.
- **EVC 2:** An Ingress Service Frame at UNI A with C-Tag VLAN ID = 45 is delivered to UNI C as a VLAN Tagged Service Frame with C-Tag VLAN ID = 101. An Ingress Service Frame at UNI C with C-Tag VLAN ID = 101 is delivered to UNI A as a VLAN Tagged Service Frame with C-Tag VLAN ID = 45.

- **EVC 3:** An Ingress Service Frame at UNI A with C-Tag VLAN ID = 7, 8, or 9 is delivered to UNI C as a VLAN Tagged Service Frame with C-Tag VLAN ID = to the value in the Ingress Service Frame at UNI A. An Ingress Service Frame at UNI C with C-Tag VLAN ID = 7, 8, or 9 is delivered to UNI A as a VLAN Tagged Service Frame with C-Tag VLAN ID = to the value in the Ingress Service Frame at UNI C.

- **EVC 4:** An Ingress Service Frame at either UNI D or UNI E that has no C-Tag is delivered to the other UNI as an Untagged Service Frame. An Ingress Service Frame at either UNI D or UNI E that has a C-Tag is delivered to the other UNI as an C-Tagged Service Frame with C-Tag VLAN ID = to the value in the Ingress Service Frame at the ingress UNI.

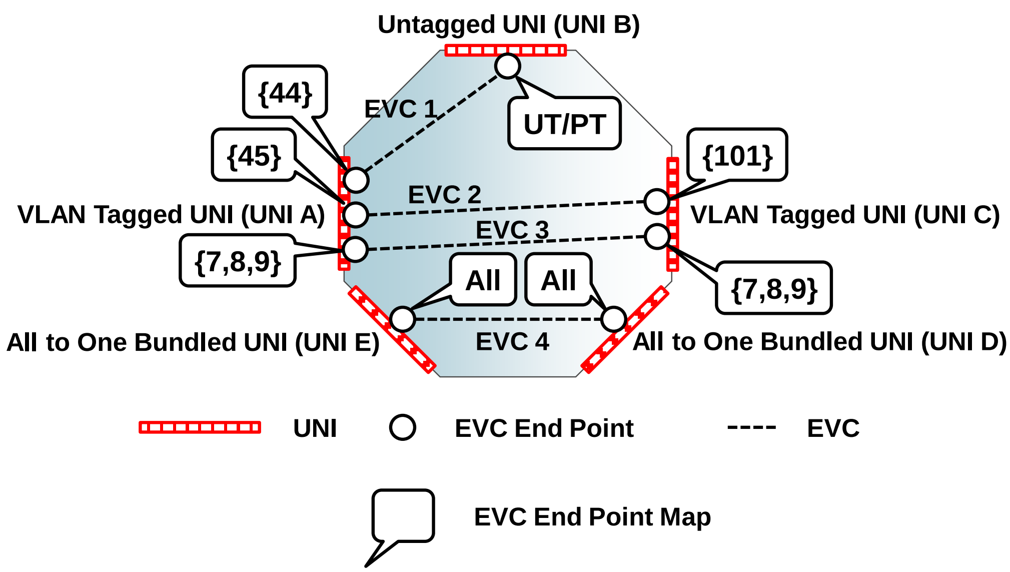

Figure 28 – Examples of VLAN Tagged UNI, All to One Bundled UNI, and Untagged UNI

[R110] and [R113] limit the ways that EVC EP Map Service Attribute values can be mixed in an EVC. Table 18 shows the mixing of values that is allowed.

		EVC EP 2			
		All	***UT/PT***	***List–1***	***List–2+***
EVC EP 1	***All***	✓			
	UT/PT		✓	✓	
	List–1		✓	✓	
	List–2+				✓

List–1: *List* with 1 entry
List–2+: *List* with 2 or more entries

Table 18 – Allowed Mixing of EVC EP Map Service Attribute Values in an EVC

Certain values of the EVC EP Service Attribute require preservation of the presence of the C-Tag and, in some cases, the value of the C-Tag VLAN ID per the following requirement.

[R118] If an Ingress Service Frame results in an Egress Service Frame, then the relationship between the Ingress Service Frame and the Egress Frame Service **MUST** be as specified in Table 19.

EVC EP Map Service Attribute Value at all UNIs in the EVC	Ingress Service Frame	Egress Service Frame
All	Untagged	Untagged
	Priority Tagged	Priority Tagged
	VLAN Tagged	VLAN Tagged with C-Tag VLAN ID value equal to the C-Tag VLAN ID value in the Ingress Service Frame
List with two or more entries	VLAN Tagged	VLAN Tagged with C-Tag VLAN ID value equal to the C-Tag VLAN ID value in the Ingress Service Frame

Table 19 – Required Preservation Behavior

Consider two EVC EPs in an EVC where both have the value of the EVC EP Map Service Attribute = *List* with a single entry in each list. If the Subscriber and Service Provider agree to the same C-Tag VLAN ID value for the two EVC EP Map Service Attribute values, then the value of the C-Tag VLAN ID is preserved for Service Frames traversing the two EVC EPs. The EVC EP pair d and e in Figure 27 is an example of such a configuration.

10.4.5 Describing the Value of the EVC EP Map Service Attribute

This document does not constrain the way that the Service Provider and Subscriber communicate the value of the EVC EP Map Service Attribute when the value of the EVC EP Map Service Attribute = *List*. For example, a Subscriber and a Service Provider could agree to describe the map as shown in Figure 29. Note that the value of the EVC EP Map Service Attribute shown in Figure 29 can only occur when the value of the Subscriber UNI Maximum Number of C-Tag VLAN IDs per EVC EP Service Attribute (Section 9.10) is greater than or equal to 4092.

Description	Actual Value
All C-Tag VLAN IDs but 2000 and 2001	1
	2
	⋮
	1999
	2002
	⋮
	4094

Figure 29 – Example of a Simple Description of the EVC EP Map Service Attribute Value

10.5 EVC EP Ingress Class of Service Map Service Attribute[30]

The value of the EVC EP Ingress Class of Service Map Service Attribute is a 3-tuple of the form $\langle F, M, P \rangle$ where:

- *F* is one of the values *EVC EP*, *C-Tag PCP*, or *DSCP*,
- *M* is a map that can be used to assign Class of Service Names to Service Frames. (The form of *M* depends on the value of *F*.), and
- *P* is a map with entries of the form <Layer 2 Control Protocol type → Class of Service Name> where the Layer 2 Control Protocol type is determined by the Protocol Identifier (see Section 6.2 of MEF 45.1 [36]).

When the value of *F* = *DSCP*, *M* can map a DSCP to two different Class of Service Names, one Class of Service Name for Ingress Service Frames carrying an IPv4 packet and a different Class of Service Name for Ingress Service Frames carrying an IPv6 packet. The map *P* can include mappings for all or a subset of the possible L2CP Protocol Identifiers (per Section 6.2 of MEF 45.1 [36]), or can be empty.

For ease of exposition, a Service Frame that is mapped to a Class of Service Name, either by the values of *F* and *M* or by the value of *P*, is said to have that Class of Service Name. Each Ingress Service Frame mapped to the given EVC EP has a single Class of Service Name taken from the value of the EVC List of Class of Service Names Service Attribute (Section 8.7). The Class of Service Name can be determined from inspection of the content of the Ingress Service Frame.

As described in Section 6.2 of MEF 23.2 [29], the Class of Service Name identifies the Performance Metrics and associated Performance Objectives 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

[30] In MEF 10.3 [24] this is called Class of Service Identifier Service Attribute.

the Class of Service Labels standardized in MEF 23.2 [29]. In addition, per Section 12, the Class of Service Name can also determine the Bandwidth Profile Flow to which the Service Frame is mapped.

[R119] An Ingress Service Frame that has the Class of Service Name *Discard* **MUST** be discarded.

[R120] The values of *M* and *P* in the value of the EVC EP Ingress Class of Service Map Service Attribute for an EVC EP **MUST** only contain Class of Service Names taken from the value of the EVC List of Class of Service Names Service Attribute (Section 8.7) for the EVC that the EVC EP is in.

The values of *F*, *M*, and *P* can be different for each EVC EP that is associated by an EVC.

The requirements regarding the EVC EP Ingress Class of Service Map Service Attribute for Data Service Frames and SOAM Service Frames are different from what they are for Layer 2 Control Protocol Service Frames as detailed below.

10.5.1 EVC EP Ingress Class of Service Map Service Attribute for Data Service Frames

[R121] If an ingress Data Service Frame that is not also an L2CP Service Frame (Section 7.6.3) is mapped to an EVC EP, then the Class of Service Name for this Service Frame **MUST** be the Class of Service Name indicated by the map *M*.

Sections 10.5.1.1, 10.5.1.2, and 10.5.1.3 describe the EVC EP Ingress Class of Service Map Service Attribute for each of the possible values of *F*.

10.5.1.1 EVC EP Ingress Class of Service Map Service Attribute Based on EVC EP

When the value of *F* is *EVC EP*, the map *M* contains a single Class of Service Name, and all ingress Data Service Frames mapped to the EVC EP (Section 10.4) that are not also L2CP Service Frames are assigned this Class of Service Name.

[R122] When the value of *F* = *EVC EP* for an EVC EP, *M* **MUST** contain a single Class of Service Name.

[R123] When the value of *F* = *EVC EP* for an EVC EP, the Class of Service Name contained in *M* **MUST NOT** be *Discard*.

[R122] means that a single Class of Service Name applies to all ingress Data Service Frames that are not also L2CP Service Frames (Section 7.6.3) mapped to the EVC EP.

As an example, consider EVC EP 1 and EVC EP 2 located at a UNI. Ingress Data Service Frames that are not also L2CP Service Frames (Section 7.6.3) mapped to EVC EP 1 have the Class of Service Name *Gold*. ingress Data Service Frames that are not also L2CP Service Frames (Section 7.6.3) mapped to EVC EP 2 have the Class of Service Name *Silver*.

10.5.1.2 EVC EP Ingress Class of Service Map Service Attribute Based on Priority Code Point Field

When the value of *F* = *C-Tag PCP*, the map *M* is based on the C-Tag Priority Code Point.

[R124] When the value of *F* = *C-Tag PCP* for an EVC EP, *M* **MUST** map each possible C-Tag PCP value to exactly one Class of Service Name and absence of a C-Tag to exactly one Class of Service Name.

[R124] means that the sets of C-Tag Priority Code Point 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 C-Tag Priority Code Point values. [R124] also means that each ingress Data Service Frame mapped to the EVC EP has a single Class of Service Name that applies to it.

[D5] When the value of *F* = *C-Tag PCP* for an EVC EP, the Class of Service Name for all ingress Data Service Frames that are not also L2CP Service Frames (Section 7.6.3) mapped to the EVC EP without a C-Tag **SHOULD** be the same Class of Service Name as that for ingress Data Service Frames that are not also L2CP Service Frames (Section 7.6.3) mapped to the EVC EP with a C-Tag that has the Priority Code Point value = 0.

Table 20 shows examples of the mapping of PCP value to Class of Service Name for two different EVC EPs at a UNI.

EVC EP A	
PCP Value	**Class of Service Name**
0	Silver
1	Discard
2	Discard
3	Silver
4	Gold
5	Gold
6	Gold
7	Gold
Untagged	Silver

EVC EP B	
PCP Value	**Class of Service Name**
0	Gold
1	Gold
2	Gold
3	Gold
4	Gold
5	Gold
6	Gold
7	Platinum
Untagged	Gold

Table 20 – Example of *M* when *F* Equals *C-Tag PCP*

Note that a single Class of Service Name for all Ingress Service Frames that are not also L2CP Service Frames (Section 7.6.3) mapped to an EVC EP can be achieved by mapping all possible PCP values and Untagged to the same Class of Service Name. This is the same behavior as when *F* = *EVC EP*.

10.5.1.3 EVC EP Ingress Class of Service Map Service Attribute Based on Internet Protocol

When the value of *F* = *DSCP*, the map *M* is based on Internet Protocol. This means that the Class of Service Name is determined from the DSCP for an ingress Data Service Frame carrying an IPv4 or an IPv6 packet.

[R125] When the value of *F* = *DSCP* for an EVC EP, *M* **MUST** map each possible DSCP to exactly one Class of Service Name for ingress Data Service Frames carrying an IPv4 packet.

[R126] When the value of *F* = *DSCP* for an EVC EP, *M* **MUST** map each possible DSCP to exactly one Class of Service Name for ingress Data Service Frames carrying an IPv6 packet.

[O5] When the value of *F* = *DSCP* for an EVC EP, *M* **MAY** be such that a DSCP maps to one Class of Service Name for an ingress Data Service Frame carrying an IPv4 packet and maps to a different Class of Service Name for an ingress Data Service Frame carrying an IPv6 packet.

[R127] When the value of *F* = *DSCP* for an EVC EP, *M* **MUST** map each ingress Data Service Frame that is not also an L2CP Service Frame (Section 7.6.3), that is not carrying either an IPv4 or IP6 packet, and that is mapped to the EVC EP, to exactly one Class of Service Name.

[R125] and [R126] mean that each possible <IP version, DSCP> pair maps to exactly one Class of Service Name.

Table 21 shows an example of using the EVC EP Ingress Class of Service Map Service Attribute based on Internet Protocol. In this example an ingress Data Service Frame carrying an IPv4 packet with DSCP = 37 would have the Class of Service Name *Platinum*. Similarly, an ingress Data Service Frame carrying an IPv6 packet with DSCP = 37 would have the 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. This example illustrates [O5] that allows a mapping of a DSCP to Class of Service name for IPv4 that is different from the mapping of a DSCP to Class of Service Name for IPv6.

IPv4 DSCPs	IPv6 DSCPs	Class of Service Name
11,37,45	11,37,45	*Platinum*
8,10,12		*Diamond*
	38,40	*Ruby*
No IP Packet	No IP Packet	*Quartz*
All other values	All other values	*Discard*

Table 21 – Example of *M* when *F* Equals *DSCP*

Table 22 shows an example where only IPv4 is supported. In this example, the Class of Service Name for ingress Data Service Frames that are not also L2CP Service Frames (Section 7.6.3) and not carrying an IP packet is *Good Enough*. This fact and the last row in Table 22 means that *Good Enough* applies to any ingress Data Service Frame carrying an IPv6 packet or any ingress Data Service Frame that is not also an L2CP Service Frame (Section 7.6.3) that is not carrying an IP packet. Consequently, IPv6 is not recognized but instead is treated as non-IP. A similar approach can be used to support only IPv6.

IPv4 DSCPs	IPv6 DSCPs	Class of Service Name
11,37,45		*Superior*
8,10,12		*Near Superior*
All other values		*Discard*
No IP Packet	No IP Packet	*Good Enough*
	All values	*Good Enough*

Table 22 – Example of *M* that Only Supports IPv4 when *F* Equals *DSCP*

10.5.2 EVC EP Ingress Class of Service Map Service Attribute for L2CP Service Frames

[R128] If an ingress Layer 2 Control Protocol Service Frame is mapped to an EVC EP and the Layer 2 Control Protocol carried by this ingress Layer 2 Control Protocol Service Frame is contained in the value of the map *P*, then the Class of Service Name for this frame **MUST** be the Class of Service Name mapped to by *P*.

[R129] If an ingress Layer 2 Control Protocol Service Frame is mapped to an EVC EP and the Layer 2 Control Protocol carried by this ingress Layer 2 Control Protocol Service Frame is not contained in the map *P*, then the Class of Service Name for this frame **MUST** be determined as if it were an ingress Data Service Frame that is not also an L2CP Service Frame (Section 7.6.3).

10.5.3 EVC EP Ingress Class of Service Map Service Attribute for SOAM Service Frames

[R130] If an ingress SOAM Service Frame that is not also a Data Service Frame is mapped to an EVC EP, then the Class of Service Name for this ingress SOAM Service Frame **MUST** be the same as if it were an ingress Data Service Frame that is not also an L2CP Service Frame (Section 7.6).

Note that SOAM Service Frames that are also Data Service Frames are covered in Section 10.5.1.

10.6 EVC EP Color Map Service Attribute[31]

The value of the EVC EP Color Map Service Attribute is a pair of the form $\langle F, M \rangle$ where:

- *F* is one of the values *EVC EP*, *C-Tag DEI*, *C-Tag PCP*, or *DSCP* and
- *M* is a map that can be used to assign Color to each Ingress Service Frame. (The form of *M* depends on the value of *F*.).

The EVC EP Color Map Service Attribute is the mechanism by which the Color for an Ingress Service Frame that is mapped to an EVC EP is indicated by the content in the Service Frame header.

[31] In MEF 10.3 [24] this is called Color Identifier Service Attribute.

The color identified for an Ingress Service Frame can affect whether it is a Qualified Service Frame (Section 8.8.1.4) and whether it is subject to the One-way Availability Performance Metric (Section 8.8.7) when:

- None of the Bandwidth Profile Flows that the Service Frame is in are in an Envelope (Section 12) or
- The Service Frame is in an Ingress Bandwidth Profile Flow that is in an Envelope (Section 12) and $CM = \textit{color-aware}$ (Section 12.1.2).

When neither of the two above conditions hold, the value of *M* has no impact on the behavior of the service.

The EVC EP Color Map Service Attribute does not apply to Egress Service Frames. Fields in an Egress Service Frame can be used to reflect the color determination of the corresponding Ingress Service Frame by having an appropriate value of the EVC EP Egress Map Service Attribute (Section 10.7). One use of this information is to allow the Subscriber to know how much delivered traffic is being declared Yellow by the Service Provider.

[R131] The Color for an Ingress Service Frame **MUST** be either Green or Yellow.

Sections 10.6.1 through 10.6.4 contain requirements for each of the above values of *F*.

10.6.1 EVC EP Color Map Service Attribute with *F = EVC EP*

When the value of *F* is *EVC EP*, the map *M* contains a single Color, which is assigned to all Ingress Service Frames mapped to the EVC EP (Section 10.4).

[R132] When the value of *F* = *EVC EP*, *M* **MUST** contain exactly one Color.

[R132] means that a single Color applies to all Ingress Service Frames mapped to the EVC EP.

10.6.2 EVC EP Color Map Service Attribute with *F = C-Tag DEI*

When *F* = *C-Tag DEI*, the Color is determined by the value of the C-Tag DEI when present.

[R133] When the value of *F* = *C-Tag DEI*, *M* **MUST** be such that the C-Tag DEI = 0 maps to Green and C-Tag DEI = 1 maps to Yellow.

[R134] When the value of *F* is *C-Tag DEI*, the Color of an Ingress Service Frame without a C-Tag **MUST** be Green.

10.6.3 EVC EP Color Map Service Attribute with *F = C-Tag PCP*

When *F* = *C-Tag PCP*, the Color is determined by the value of the C-Tag PCP when present.

[R135] When the value of *F* = *C-Tag PCP*, *M* **MUST** be such that each possible value of the C-Tag PCP maps to a Color.

[R136] When the value of *F* = *C-Tag PCP*, the Color of a Ingress Service Frame without a C-Tag **MUST** be Green.

[R135] means that the sets of C-Tag PCP values that map to each Color are disjoint sets and the union of these sets is the set of all possible C-Tag PCP values.

10.6.4 EVC EP Color Map Service Attribute with *F = DSCP*

When *F* = *DSCP*, the Color is determined by the DSCP when the Service Frame is carrying an IPv4 packet or an IPv6 packet.

[R137] When the value of *F* = *DSCP*, *M* **MUST** be such that each possible DSCP maps to exactly one Color for Ingress Service Frames carrying an IPv4 packet.

[R138] When the value of *F* = *DSCP*, *M* **MUST** be such that each possible DSCP maps to exactly one Color for Ingress Service Frames carrying an IPv6 packet.

[O6] When the value of *F* = *DSCP*, *M* **MAY** be such that a DSCP maps to one Color for an Ingress Service Frame carrying an IPv4 packet and maps to a different Color for an Ingress Service Frame carrying an IPv6 packet.

[R139] When the value of *F* = *DSCP*, the Color of an Ingress Service Frame that does not contain either an IPv4 or an IPv6 packet **MUST** be Green.

Note that the mapping of DSCPs to Color can be different for IPv4 and IPv6. Table 23 shows an example of the use of different mappings for IPv4 and IPv6.

IPv4 DSCPs	IPv6 DSCPs	Color
1,2,3,5,7,11,13,17	1,2,3,5,8,13,21,34	Yellow
All other values	All other values	Green

Table 23 – Example of DSCP to Color Mapping

10.7 EVC EP Egress Map Service Attribute

The value of the EVC EP Egress Map Service Attribute is *None* or a map of the form <Corresponding Ingress Service Frame Class of Service Name, Corresponding Ingress Service Frame Color> to either <Egress Service Frame C-Tag PCP value, Egress Service Frame C-Tag DEI value> or *Discard*.

Table 24 contains an example when value of the EVC EP Egress Map Service Attribute is not *None* for the value of the EVC List Class of Service Names Service Attribute value in Table 4.

Ingress Class of Service Name	Ingress Color	Egress C-Tag PCP Value	Egress C-Tag DEI Value
Platinum	*Green*	7	0
Platinum	*Yellow*	7	1
Gold	*Green*	5	0
Gold	*Yellow*	5	1
Silver	*Green*	*Discard*	
Silver	*Yellow*	*Discard*	

Table 24 – Example of the Value of the EVC EP Egress Map Service Attribute

Section 10.5 details the determination of the Class of Service Name for Ingress Service Frames. The Color of an Ingress Service Frame is determined by the Ingress Bandwidth Profile that is applied to it or according to the value of the EVC EP Color Map Service Attribute (Section 10.6) if no Ingress Bandwidth Profile is applied.

Note that the EVC EP Egress Map Service Attribute does not apply to egress SOAM Service Frames and egress Layer 2 Control Protocol Service Frames that are a result of frames generated inside the SP Network. Populating the C-Tag for Service Frames that are a result of frames generated inside the SP Network is beyond the scope of this document.

[R140] The value of the EVC EP Egress Map Service Attribute **MUST NOT** be *None* unless one of the following conditions applies:

- The value of the EVC EP Map Service Attribute (Section 10.4) is *UT/PT*.
- The value of the EVC C-Tag PCP Preservation Service Attribute (Section 8.5) = *Enabled*, the value of the EVC C-Tag DEI Preservation Service Attribute (Section 8.6) = *Enabled*, and the value of the EVC EP Map Service Attribute (Section 10.4) = *List* for all EVC EPs in the EVC.
- The value of the EVC C-Tag PCP Preservation Service Attribute (Section 8.5) = *Enabled*, the value of the EVC C-Tag DEI Preservation Service Attribute (Section 8.6) = *Enabled*, and the value of the EVC EP Map Service Attribute (Section 10.4) = *All*.

[R141] If the value of the EVC EP Egress Map Service Attribute is not *None* and is such that the pair <Class of Service Name, Color> maps to *Discard*, then any Ingress Service Frame with this Class of Service Name and Color **MUST NOT** result in an Egress Service Frame at the UNI where the EVC EP is located.

[R142] If an Egress Service Frame that is mapped to the EVC EP is a C-Tagged Service Frame and if the corresponding Ingress Service Frame is an Untagged Service Frame, then the values of the C-Tag PCP and C-Tag DEI in the Egress Service Frame **MUST** be as specified by the value of the EVC EP Egress Map Service Attribute.

[R143] If an Egress Service Frame that is mapped to the EVC EP is a C-Tagged Service Frame and if the corresponding Ingress Service Frame is a C-Tagged Service Frame, then the values of the C-Tag PCP and C-Tag DEI in the Egress Service Frame **MUST** be as specified in Table 25.

Value of the EVC C-Tag PCP Preservation Service Attribute	Value of the EVC C-Tag DEI Preservation Service Attribute	Value of the egress C-Tag PCP	Value of the egress C-Tag DEI
Enabled	*Enabled*	I	I
Enabled	*Disabled*	I	M
Disabled	*Enabled*	M	I
Disabled	*Disabled*	M	M

I = Value in the corresponding Ingress Service Frame
M = Value in the EVC EP Egress Map Service Attribute

Table 25 – Egress Service Frame C-Tag PCP and C-Tag DEI Determination

Notes on setting the value of the EVC EP Egress Map Service Attribute are contained in Section 10.7.1 and Section 10.7.2.

10.7.1 Conveying Class of Service Name and Color

It can be useful to the Subscriber to know the Class of Service Name and Color used within the SP Network for delivering each Egress Service Frame. For example, the Subscriber can use this information to estimate the percentage of Ingress Service Frames with a given Class of Service Name that were discarded and/or declared Yellow by the SP Network.

If the SN uses the C-Tag DEI to assign Color to Egress Service Frames, then the value the EVC EP Egress Map Service Attribute shown in Table 26 allows the SN to use the C-Tag PCP value to identify the Class of Service Name and the C-Tag DEI to identify the Color for each C-Tagged Egress Service Frame.

Ingress Class of Service Name	Ingress Color	Egress C-Tag PCP Value	Egress C-Tag DEI Value
Platinum	*Green*	7	0
Platinum	*Yellow*	7	1
Gold	*Green*	5	0
Gold	*Yellow*	5	1
Silver	*Green*	3	0
Silver	*Yellow*	3	1

Table 26 – Example of Using the C-Tag PCP to Identify Class of Service Name and C-Tag DEI to Identify Color

If the SN uses the C-Tag PCP to assign Color to Egress Service Frames, then the value the EVC EP Egress Map Service Attribute shown in Table 27 allows the SN to use the C-Tag PCP value to identify both the Class of Service Name and the Color for each C-Tagged Egress Service

Frame. Note that the entries in the fourth column can be any value since the SN does not use the C-Tag DEI to identify Color.

Ingress Class of Service Name	Ingress Color	Egress C-Tag PCP Value	Egress C-Tag DEI Value
Platinum	*Green*	7	0
Platinum	*Yellow*	6	0
Gold	*Green*	5	0
Gold	*Yellow*	4	0
Silver	*Green*	3	0
Silver	*Yellow*	2	0

Table 27 – Example of Using the C-Tag PCP to Identify both Class of Service Name and Color

10.7.2 Use of *Discard* in the Value of the EVC EP Egress Map Service Attribute

As an example of the use of *Discard* in the value of the EVC EP Egress Map Service Attribute, consider an EVC that associates five EVC EPs; one EVC EP at each of UNIs A, B, C, D, and E. If the Class of Service Name *Silver* is only used by the Subscriber at UNIs A and B, then using the value of the EVC EP Egress Map Service Attribute shown in Table 24 at UNIs C, D, and E prevents *Silver* Egress Service Frames at UNIs C, D, and E thus conserving egress bandwidth at these UNIs.

An analogous approach can be used to prevent Service Frames that were deemed to be Yellow at the ingress UNI from consuming egress bandwidth at the egress UNI(s).

10.8 EVC EP Ingress Bandwidth Profile Service Attribute[32]

The EVC EP Ingress Bandwidth Profile Service Attribute is used to limit the rate of Ingress Service Frames mapped to an EVC EP at a UNI (Section 7.10.1). The value of the EVC EP Ingress Bandwidth Profile Service Attribute is either *None* or BWP Flow Parameters as defined in Section 12.1.2.

[R144] When the value of the EVC EP Ingress Bandwidth Profile Service Attribute for a given EVC EP = BWP Flow Parameters (Section 12.1.2), the SP Network **MUST** apply the Bandwidth Profile Algorithm (Section 12.2) using these parameter values to the Ingress EVC EP Bandwidth Profile Flow for that EVC EP (Section 12).

[R145] An Ingress Service Frame that is declared Red by the application of the Bandwidth Profile Algorithm per [R144] **MUST** be discarded.

Note that if an Ingress Service Frame is declared Green and it meets other conditions per the definition of Qualified Service Frame in Section 8.8.1.4, then the SLS applies to this Service Frame.

[32] In MEF 10.3 [24] this is called Ingress Bandwidth Profile per EVC Service Attribute.

Note that, per the definition of Qualified Service Frame in Section 8.8.1.4, an Ingress Service Frame that is declared Yellow is not subject to the Performance Metrics and associated Performance Objectives of the SLS.

There are three levels of Bandwidth Profile conformance in descending order: Green, Yellow and Red. A better conformance level results in fewer discarded frames. In order to improve the level of Bandwidth Profile conformance, a Subscriber may need to shape traffic in the SN (Appendix B contains a shaping example).

Appendix D.1 shows an example of applying the EVC EP Ingress Bandwidth Profile Service Attribute at a UNI.

10.9 EVC EP Class of Service Name Ingress Bandwidth Profile Service Attribute[33]

The EVC EP Class of Service Name Ingress Bandwidth Profile Service Attribute is used to limit the rate of Ingress Service Frames with a given Class of Service Name at a UNI. The value of the EVC EP Class of Service Name Ingress Bandwidth Profile Service Attribute is either *None* or a non-empty list of pairs of the form $\langle x, y\rangle$, where x is a Class of Service Name that is in the value of the EVC List of Class of Service Names Service Attribute (Section 8.7) for the EVC that associates the EVC EP and y = BWP Flow Parameters as defined in Section 12.1.2.

[R146] Each Class of Service Name that is in the value of the EVC List of Class of Service Names Service Attribute (Section 8.7) for the EVC that associates the EVC EP **MUST** appear at most once in the list of pairs $\langle x, y\rangle$ in the value of the EVC EP Class of Service Name Ingress Bandwidth Profile Service Attribute.

[R147] The Class of Service Name *Discard* **MUST NOT** appear in the list of pairs $\langle x, y\rangle$ in the value of the EVC EP Class of Service Name Ingress Bandwidth Profile Service Attribute.

[R148] For each entry $\langle x, y\rangle$ in the value of the EVC EP Class of Service Name Ingress Bandwidth Profile Service Attribute for a given EVC EP, where $y =$ BWP Flow Parameters (Section 12.1.2), the SP Network **MUST** apply the Bandwidth Profile Algorithm (Section 12.2) using these parameter values to the Ingress Class of Service Name Bandwidth Profile Flow (Section 12) for this EVC EP and Class of Service Name x.

[R149] An Ingress Service Frame that is declared Red by the application of the Bandwidth Profile Algorithm per [R148] **MUST** be discarded.

Note that if there is no need to apply the Bandwidth Profile Algorithm to Ingress Service Frames with a given Class of Service Name, e.g., the Class of Service Name is not used at the UNI, then this Class of Service Name can be omitted from the list of $\langle x, y\rangle$'s. This obviates the need to specify parameter values for a Bandwidth Profile Flow that will never be used.

[33] In MEF 10.3 [24] this is called Ingress Bandwidth Profile per Class of Service Identifier Service Attribute.

Note that if an Ingress Service Frame is declared Green and it meets other conditions per the definition of Qualified Service Frame in Section 8.8.1.4, then the SLS applies to this Service Frame.

Note that, per the definition of Qualified Service Frame in Section 8.8.1.4, an Ingress Service Frame that is declared Yellow is not subject to the Performance Metrics and associated Performance Objectives of the SLS.

There are three levels of Bandwidth Profile conformance in descending order: Green, Yellow and Red. A better conformance level results in fewer discarded frames. In order to improve the level of Bandwidth Profile conformance, a Subscriber may need to shape traffic in the SN (Appendix B contains a shaping example).

Appendix D.1 shows an example of applying the EVC EP Class of Service Name Ingress Bandwidth Profile Service Attribute at a UNI.

10.10 EVC EP Egress Bandwidth Profile Service Attribute[34]

The EVC EP Egress Bandwidth Profile Service Attribute is used to limit the rate of all Egress Service Frames mapped to an EVC EP at a UNI (Section 7.10.2). The value of the EVC EP Egress Bandwidth Profile Service Attribute is either *None* or the 3-tuple $\langle CIR, CIR_{max}, ER \rangle$ where the elements of the 3-tuple are defined in Section 12.1.2.

When the value of the EVC EP Egress Bandwidth Profile Service Attribute is a 3-tuple $\langle CIR, CIR_{max}, ER \rangle$, the Service Provider might, from time to time, delay or discard traffic that is mapped to the EVC EP in such a way that applying the Bandwidth Profile Algorithm (Section 12.2) to the Egress EVC EP Bandwidth Profile Flow for the EVC EP would result in all Egress Service frames being declared Green or Yellow, i.e., no frames declared Red.

[R150] If the Envelope identified by ER contains only one Bandwidth Profile Flow (Section 12), then CIR_{max} **MUST** equal CIR.

[R151] If the value of the rank in ER is not the highest rank in the Envelope identified by ER, then CIR **MUST** be 0.

It is important to reiterate that this description does not mandate or constrain in any way the implementation in the SP Network.

When the Bandwidth Profile Algorithm is applied to an Egress EVC EP Bandwidth Profile Flow, some parameter values used by the algorithm are specified by the Service Provider per Table 30 and Table 31. Furthermore, the algorithm uses the Color assigned to the frame at the ingress UNI. Thus, if the Subscriber wants to apply the Bandwidth Profile Algorithm to an Egress EVC EP Bandwidth Profile Flow, it is necessary for the Subscriber to know:

- The Service Provider specified parameter values, and

[34] In MEF 10.3 [24] this is called Egress Bandwidth Profile per EVC Service Attribute.

- The Color assigned to each corresponding Ingress Service Frame at the ingress UNI. (See Section 10.7.1 for how this can be accomplished via the EVC EP Egress Map Service Attribute.)

Note that Service Frames discarded or buffered in order to meet the limits implied by the EVC EP Egress Bandwidth Profile can affect the value of the Performance Metrics in Section 8.8 and Section 11.

10.11 EVC EP Class of Service Name Egress Bandwidth Profile Service Attribute[35]

The EVC EP Class of Service Name Egress Bandwidth Profile Service Attribute is used to limit the rate of all Egress Service Frames with a given Class of Service Name, as determined at the ingress UNI for each frame per the EVC EP Ingress Class of Service Map Service Attribute (Section 10.5). The value of the EVC EP Class of Service Name Egress Bandwidth Profile Service Attribute is either *None* or a non-empty list of pairs of the form $\langle x, y \rangle$ where x is a Class of Service Name contained in the value of the EVC List of Class of Service Names Service Attribute (Section 8.7) but not equal to *Discard* and $y =$ the 3-tuple $\langle CIR, CIR_{max}, ER \rangle$ where the elements of the 3-tuple are defined in Section 12.1.2.

For a pair $\langle x, y \rangle$ in the value of the EVC EP Class of Service Name Egress Bandwidth Profile Service Attribute, the Service Provider might, from time to time, delay or discard traffic in the Egress Class of Service Name Bandwidth Profile Flow (Section 12) for the Class of Service Name x in such a way that applying the Bandwidth Profile Algorithm (Section 12.2), using the value of y, to Egress Service Frames in the Egress Class of Service Name Bandwidth Profile Flow would result in frames being declared Green or Yellow, i.e., no frames declared Red.

[R152] A given Class of Service Name **MUST** appear at most once in the list of pairs $\langle x, y \rangle$ in the value of the EVC EP Class of Service Name Egress Bandwidth Profile Service Attribute.

[R153] For each $\langle x, y \rangle$ in the value of the EVC EP Class of Service Name Egress Bandwidth Profile Service Attribute, if the Envelope identified by ER in y contains only one Bandwidth Profile Flow (Section 12), then CIR_{max} **MUST** equal CIR.

[R154] For each $\langle x, y \rangle$ in the value of the EVC EP Class of Service Name Egress Bandwidth Profile Service Attribute, if the value of the rank in ER in y is not the highest rank in the Envelope identified by ER, then CIR **MUST** be 0.

It is important to reiterate that this description does not mandate or constrain in any way the implementation in the SP Network.

When the Bandwidth Profile Algorithm is applied to an Egress Class of Service Name Bandwidth Profile Flow, some parameter values used by the algorithm are specified by the Service Provider per Table 30 and Table 31. Furthermore, the algorithm uses the Color and Class of Ser-

[35] In MEF 10.3 [24] this is called Egress Bandwidth Profile per Egress Equivalence Class Identifier Service Attribute.

vice Name assigned to the frame at the ingress UNI. Thus, if the Subscriber wants to apply the Bandwidth Profile Algorithm to an Egress Class of Service Name Bandwidth Profile Flow, it is necessary for the Subscriber to know:

- The Service Provider specified parameter values, and
- The Color and Class of Service Name assigned to each corresponding Ingress Service Frame at the ingress UNI. (See Section 10.7.1 for how this can be accomplished via the EVC EP Egress Map Service Attribute.)

Note that Service Frames discarded or buffered in order to meet the limits implied by the EVC EP Class of Service Name Egress Bandwidth Profile Service Attribute can affect the value of the Performance Metrics in Section 8.8 and Section 11.

Note that if there is no need to apply the Bandwidth Profile Algorithm to Egress Service Frames with a given Class of Service Name, then this Class of Service Name can be omitted from the list of $\langle x, y\rangle$'s. This obviates the need to specify parameter values for a Bandwidth Profile Flow that is not subject to an Egress Bandwidth Profile.

10.12 EVC EP Source MAC Address Limit Service Attribute

The value of the EVC EP Source MAC Address Limit Service Attribute is either *None* or the pair $\langle N, \tau\rangle$ where N is an integer ≥ 1 and τ is a time duration.

When the value of the EVC EP Source MAC Address Limit Service Attribute = $\langle N, \tau\rangle$, the number of source MAC Addresses that can be used in Ingress Service Frames is limited.

Within the SP Network, there can be tables for recording Source MAC Addresses. These tables can be shared across multiple EVCs. One motivation for this Service Attribute is to avoid these shared tables from being exhausted by the actions of one Subscriber. For example, if MAC Address learning is being done in the SP Network, exhausting these tables can result in undesirable flooding of frames within the SP Network for other Subscribers.

The limit is applied by maintaining a list of maximum length N of source MAC addresses which are aged-out of the list if not seen for a time duration τ. If an Ingress Service Frame arrives with a new source MAC address when the list is full, it is recommended that the Service Frame be discarded.

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

$$L = \{\langle A_i, t_i\rangle | i = 1,2, \ldots, q \leq N, A_i = \text{unicast MAC Address}, t_i = \text{ a time}\}$$

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 EP and contained the source MAC address A_i.

[D6] If the value of the Source MAC Address Limit Service Attribute = $\langle N, \tau\rangle$, then for a sequence of Ingress Service Frames mapped to the EVC EP with Source MAC Addresses A_j and Service Frame Arrival Times at the UNI t_j for $j =$

0,1, …., the frames **SHOULD** be discarded per the logic of Figure 30 where $L = \emptyset$ at time t_0.

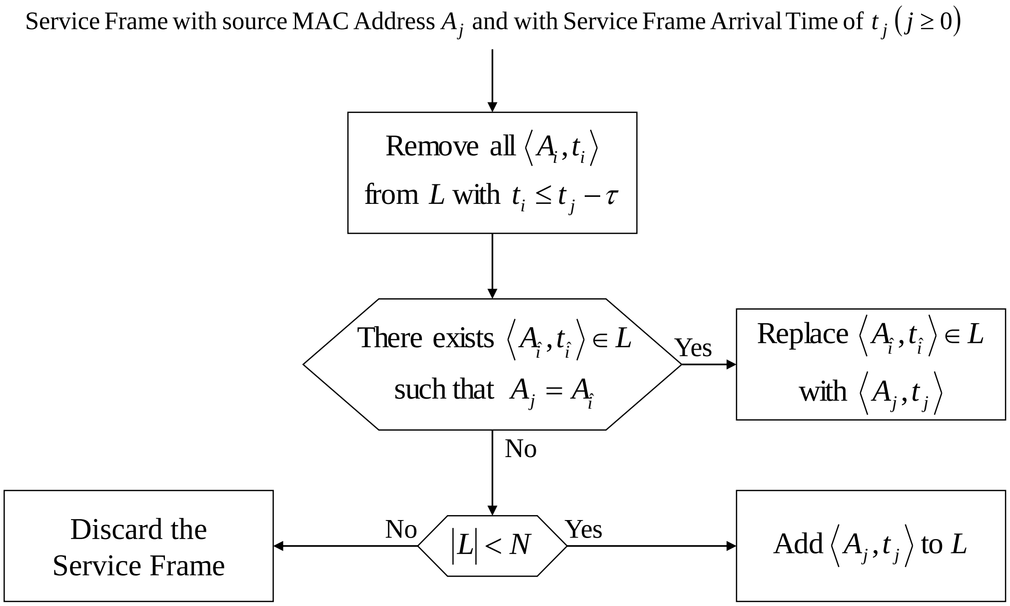

Figure 30 – Frame Discard for Source MAC Address Limit

Note that [D6] does not mandate a specific implementation in the SP Network. Any implementation that yields the same behavior as that of Figure 30 meets the requirement. For example, an implementation that removes $\langle A_i, t_i \rangle$ from the list at time $t = t_i + \tau$ yields the same Service Frame discard behavior as that of Figure 30.

When the value is *None* there is no limit above which Service Frames are discarded.

10.13 EVC EP Subscriber MEG MIP Service Attribute

The value of the EVC EP Subscriber MEG MIP Service Attribute is either *None* or an integer in the range 0 – 7 that indicates the MEG Level of a Subscriber MEG MIP.

[R155] If the value of the EVC EP Subscriber MEG MIP Service Attribute = *None*, a Subscriber MEG MIP **MUST NOT** be instantiated in the SP Network.

[R156] If the value of the EVC EP Subscriber MEG MIP Service Attribute = an integer, then a Subscriber MEG MIP **MUST** be instantiated in the SP Network with the MEG Level set to the value of the EVC EP Subscriber MEG MIP Service Attribute.

[R157] If the value of the EVC EP Subscriber MEG MIP Service Attribute is an integer, then it **MUST** be greater than or equal to the value of the EVC Available MEG Level Service Attribute (Section 8.11).

When the value of the EVC EP Subscriber MEG MIP Service Attribute is not *None*, several parameter values need to be agreed upon by the Subscriber and the Service Provider.

11. Multiple EVC Service Level Specification Service Attribute

The value of the Multiple EVC Service Level Specification Service Attribute is a 3-tuple of the form $\langle ID, KG, \widehat{MGA} \rangle$ where:

- ID is a string that identifies the instance of the Multiple EVC Service Level Specification Service Attribute.
- KG is an integer ≥ 1.
- $\widehat{MGA}$ is a percentage that is the Performance Objective.

The Multiple EVC Service Level Specification Service Attribute specifies an objective for the percentage of time that at least KG sets of ordered EVC EP pairs will be simultaneously available where more than one EVC is involved. It generalizes the One-way Group Availability Performance Metric (Section 8.8.11) by including more than one EVC.

[R158] The value of ID **MUST** contain no more than 45 characters.

[R159] The value of ID **MUST** be a non-null RFC 2579 [15] DisplayString but not contain the characters 0x00 through 0x1f.

For ease of exposition, an EVC is said to belong to an instance of the Multiple EVC Service Level Specification Service Attribute when the value of the EVC Group Membership Service Attribute (Section 8.9) contains a 3-tuple whose first value equals the value of ID.

[R160] All EVCs that belong to an instance of the Multiple EVC Service Level Specification Service Attribute **MUST** be provided to a single Subscriber by a single Service Provider.

[R161] The values of the EVC Service Level Specification Service Attribute (Section 8.8) for the EVCs that belong to an instance of the Multiple EVC Service Level Specification Service Attribute **MUST** have the same values for the pair $\langle ts, T \rangle$.

[R161] means that that all of the EVCs have an identical series of time intervals $\{T_l\}$, and that it is therefore well-defined to refer to T_l for the instance of the Multiple EVC Service Level Specification Service Attribute.

An example of the application of the Multiple EVC Service Level Specification Service Attribute is the Subscriber using one EVC as a backup for another EVC. Figure 31 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 primary EVC fails, the Mobile Operator routers start sending data on the standby EVC.

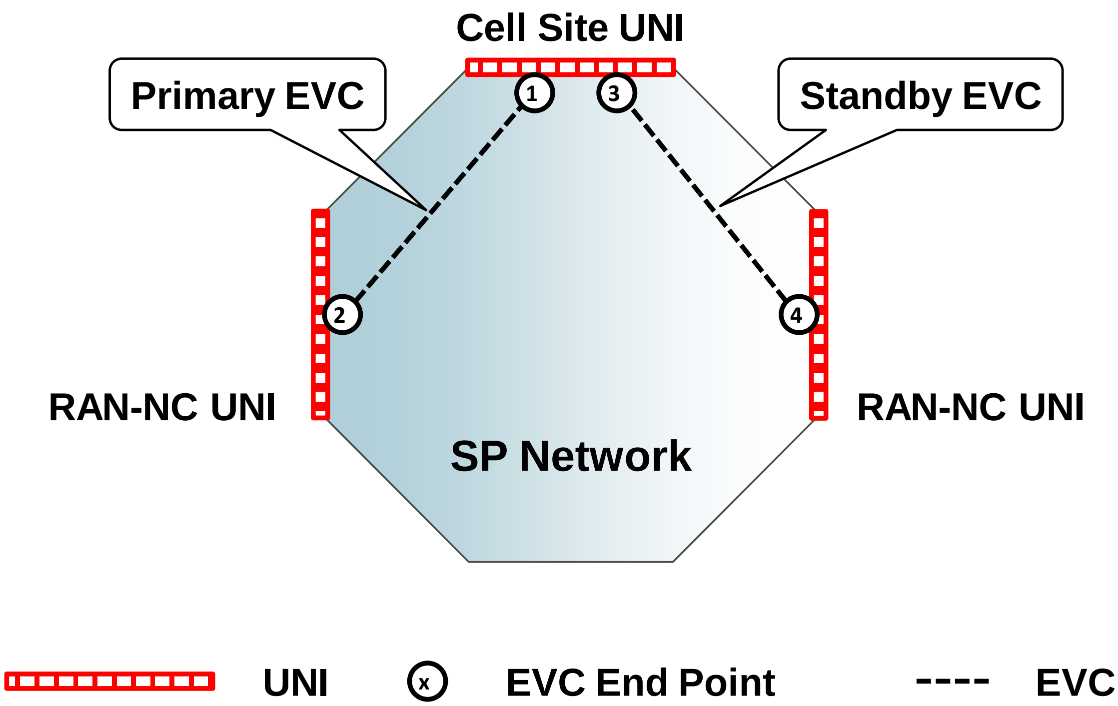

Figure 31 – Example of the Multiple EVC Service Level Specificaton Service Attribute

Consider an instance of the Multiple EVC Service Level Specification Service Attribute whose value is $\langle ID, KG, \widehat{MGA} \rangle$. Let $E_q, q = 1, \ldots, h$ be the values for EVC ID Service Attribute for the EVCs that belong to this instance of the Multiple EVC Service Level Specification Service Attribute. Note that [R51] means that $E_i \neq E_j, i = 1, \ldots h, j = 1, \ldots h, i \neq j$.

In the example in Figure 31:

- The value of the Multiple EVC Service Level Specification Service Attribute = $\langle MobileBackhaulGroup002358, 1, 99.999\% \rangle$,
- The value of the EVC Group Membership Service Attribute (Section 8.9) for the Primary EVC = $\langle MobileBackhaulGroup002358, Platinum, \{\langle 1,2 \rangle, \langle 2,1 \rangle\} \rangle$, and
- The value of the EVC Group Membership Service Attribute (Section 8.9) for the Standby EVC = $\langle MobileBackhaulGroup002358, Gold, \{\langle 3,4 \rangle, \langle 4,3 \rangle\} \rangle$.

Informally, when the Performance Metric for the Multiple EVC Service Level Specification (defined below) is met, 99.999% of the time interval T_l there is at least one EVC that is not experiencing "high loss" or a Maintenance Interval.

[R162] For each instance of the Multiple EVC Service Level Specification Service Attribute, there **MUST** be at least two different EVCs that belong to the Multiple

EVC Service Level Specification via the value of the EVC Group Membership Service Attribute (Section 8.9).

[R162] means that h is at least 2.

[R163] Let the 5-tuple, $\langle CoS\_Name_q, \Delta t_q, C_q, n_q, PM_q \rangle$ be the entry in CN for which CoS_Name_q is equal to CoS_Name_G in the value of the EVC Group Membership Service Attribute (Section 8.9) for EVC E_q. Then all $\Delta t_q, q = 1, \ldots h$ **MUST** have the same value.

[R161] and [R163] mean that all of the EVCs have an identical series of time intervals $\{\Delta t_k\}$ for their respective Class of Service Names, and that it is therefore well-defined to refer to Δt_k for the instance of the Multiple EVC Service Level Specification Service Attribute.

Let Δt_k and T_l be as defined in Section 8.8.1.3. Let $\langle ID_q, CoS\_Name\_G_q, SG_q \rangle$ be the 3-tuple in the value of the EVC Group Membership Service Attribute (Section 8.9) for EVC E_q, and $ID_q = ID$ for ID in $\langle ID, KG, \widehat{MGA} \rangle$. For each E_q, let $AT_q(\langle i,j \rangle, T_l)$ denote Available Time for the ordered EVC EP pair $\langle i,j \rangle \in SG_q$ and $CoS\_Name\_G_q$, as specified in Section 8.8.1.3. Recall that if $\Delta t_k \in AT_q(\langle i,j \rangle, T_l)$, then Δt_k does not intersect a Maintenance Interval and there has been good frame loss performance per the sliding window calculation.

Define

$$ATSG_q\left(SG_{q,} \Delta t_k, T_l\right) = \begin{cases} 1 \text{ if } \Delta t_k \in AT_q(\langle i,j \rangle, T_l) \text{ for all } \langle i,j \rangle \in SG_q \\ 0 \text{ otherwise} \end{cases}$$

Informally, $ATSG_q\left(SG_{q,} \Delta t_k, T_l\right) = 1$ means that during the small time interval, none of the ordered EVC EP pairs in EVC E_q, have been impacted by a Maintenance Interval and all of the ordered EVC EP pairs have had good frame loss performance.

Define

$$MGA(ID, KG, \Delta t_k, T_l) = \begin{cases} 1 \text{ if } \sum_{q=1}^{h} ATSG_q\left(SG_{q,} \Delta t_k, T_l\right) \geq KG \\ 0 \text{ otherwise} \end{cases}$$

Define $\tilde{X}(T_l)$ as the set of Δt_k's contained in T_l. Note that $\tilde{X}(T_l) \neq \emptyset$. The Performance Metric for the Multiple EVC Service Level Specification Service Attribute, expressed as a percentage, is

$$\overline{MGA}(ID, KG, T_l) = \frac{100}{\left|\tilde{X}(T_l)\right|} \sum_{\Delta t_k \in X(T_l)} MGA(ID, KG, \Delta t_k, T_l)$$

where $\left|\tilde{X}(T_l)\right|$ is the number of elements of $\tilde{X}(T_l)$.

Note that when multiple EVCs in an instance of the Multiple EVC Service Level Specification Service Attribute come under maintenance, the value of $\overline{MGA}(ID, KG, T_l)$ can be negatively impacted. This is because a Δt_k that intersects a Maintenance Interval is not an element of Available Time, $AT_q(\langle i,j\rangle, T_l)$ (Section 8.8.1.3). This in turn can mean that $ATSG_q(SG_q, \Delta t_k, T_l) = 0$ for multiple EVCs E_q and hence $MGA(ID, KG, \Delta t_k, T_l) = 0$. An example of such a situation is if both the Primary EVC and the Standby EVC in Figure 31 are put into maintenance in order to repair or reconfigure the Cell Site UNI.

[R164] The Multiple EVC Service Level Specification Service Attribute **MUST** define the Performance Metric for the Multiple EVC Service Level Specification as met over T_l if and only if $\overline{MGA}(ID, KG, T_l) \geq \widehat{MGA}$.

Figure 32 shows the calculation of $MGA(ID, KG, \Delta t_k, T_l)$ for the example of Figure 31. The value of the Performance Metric for this example is $100 \times {}^{15}/_{17}$ =88.24%.

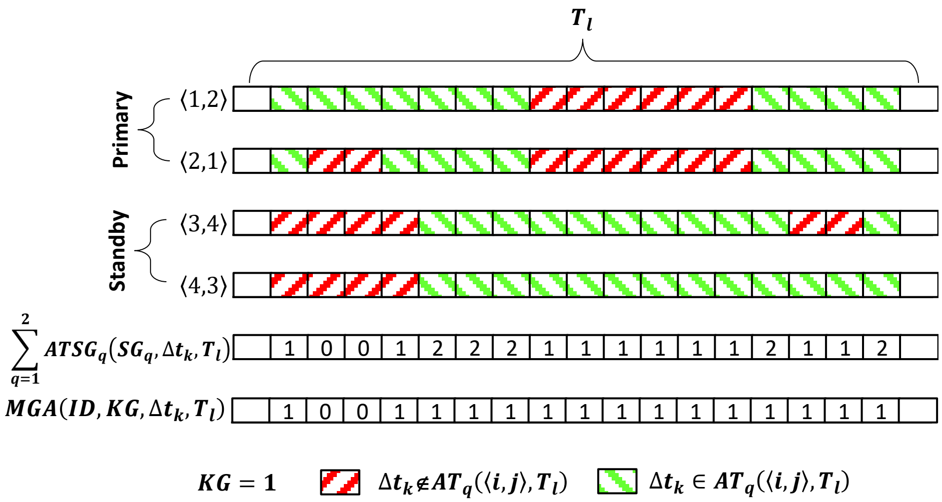

Figure 32 – Example of the Calculation of *MGA*

12. Bandwidth Profiles

A Bandwidth Profile is a specification of the temporal properties of a sequence of Service Frames at a UNI, along with rules for determining the level of conformance to the specification for each Service Frame in the sequence.

A Bandwidth Profile consists of the set of traffic parameters applicable to a sequence of Service Frames. Associated with the Bandwidth Profile is an algorithm to determine Service Frame conformance with the specified parameters.

All Bandwidth Profile Service Attributes (Sections 9.11, 9.12, 10.8, 10.9, 10.10, and 10.11) 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 a set of Service Frames at a UNI that meet specific criteria. There are 4 types of Bandwidth Profile Flow per the following definitions:

- **Ingress EVC EP Bandwidth Profile Flow:** All Ingress Service Frames at a UNI that are mapped to a given EVC EP and that are not discarded per [R82] with the possible exception of Service Frames discarded per [R103], [R119], [D1], [D2], [D6], [O4], or any of the conditions specified per [R20].[36]

- **Ingress Class of Service Name Bandwidth Profile Flow:** All Ingress Service Frames at a UNI that have a given Class of Service Name, that are mapped to a given EVC EP, and that are not discarded per [R82] with the possible exception of Service Frames discarded per [R103], [R119], [D1], [D2], [D6], [O4], or any of the conditions specified per [R20].[36]

- **Egress EVC EP Bandwidth Profile Flow:** All Egress Service Frames at a UNI that are mapped to a given EVC EP.

- **Egress Class of Service Name Bandwidth Profile Flow:** All Egress Service Frames at a UNI that map to a given EVC EP and whose corresponding Ingress Service Frames have a given Class of Service Name .

For ease of exposition, "Ingress Bandwidth Profile Flow" is used to refer to either of the first two bullet items and "Egress Bandwidth Profile Flow" is used to refer to either of the last two bullet items.

Note that in some cases two Bandwidth Profile Flow types can be equivalent, e.g., the Ingress Class of Service Name Profile Flow and the Ingress EVC EP Bandwidth Profile Flow are equivalent when all Service Frames that map to an EVC EP have the same Class of Service Name via the value of the EVC EP Ingress Class of Service Map Service Attribute (Section 10.5).

[36] This document does not specify whether or not a Service Frame discarded per [R103], [R119], [D1], [D2], [D6], [O4], or any of the conditions specified per [R20] is in a Bandwidth Profile Flow that is in an Envelope.

A Service Frame will not be subject to a Bandwidth Profile Algorithm when:

- It is an Ingress Service Frame that is not in an Ingress Bandwidth Profile Flow, or
- None of the Bandwidth Profile Flows that contain the Service Frame are in an Envelope.

An example of the first case is an Ingress Service Frame that is discarded per [R82]. A Service Frame that is not subjected to a Bandwidth Profile Algorithm does not consume any tokens.

An Envelope is a set that contains one or more Bandwidth Profile Flows that can share bandwidth resources that are represented by tokens. In an Envelope that contains n Bandwidth Profile Flows, each Bandwidth Profile Flow is assigned a unique rank between 1 (lowest) and n (highest).

[R165] The Bandwidth Profile Flows in an Envelope **MUST** meet exactly one of the following conditions:

- They are all Ingress EVC EP Bandwidth Profile Flows,
- They are all Ingress Class of Service Name Bandwidth Profile Flows,
- They are all Egress EVC EP Bandwidth Profile Flows, or
- They are all Egress Class of Service Name Bandwidth Profile Flows.

[R165] means that an Envelope contains Bandwidth Profile Flows with only Ingress Service Frames or only Egress Service Frames.

Multiple Envelopes containing Ingress EVC EP Bandwidth Profile Flows can co-exist at a UNI with the EVC EP Bandwidth Profile Flow for a given EVC EP in exactly one Envelope. Similarly multiple Envelopes containing Egress EVC EP Bandwidth Profile Flows can co-exist at a UNI with the EVC EP Bandwidth Profile Flow for a given EVC EP in exactly one Envelope.

Multiple Envelopes containing Ingress Class of Service Name Bandwidth Profile Flows can co-exist at a UNI. Similarly multiple Envelopes containing Egress Class of Service Name Bandwidth Profile Flows can co-exist at a UNI.

An Envelope can contain Ingress Class of Service Name Bandwidth Profile Flows for multiple EVC EPs at the UNI. Similarly an Envelope can contain Egress Class of Service Name Bandwidth Profile Flows for multiple EVC EPs at the UNI. Appendix D.1 shows an example of mapping Bandwidth Profile Flows to Envelopes.

Note that there can be many Bandwidth Profile Flows at a UNI depending on the number of EVC EPs and the number of Class of Service Names for the EVC that each EVC EP is in. An Egress Service Frame is always in the two following Bandwidth Profile Flows:

- Egress EVC EP Bandwidth Profile Flow and
- Egress Class of Service Name Bandwidth Profile Flow.

An Ingress Service Frame that is not discarded is always in the two following Bandwidth Profile Flows:

- Ingress EVC EP Bandwidth Profile Flow and
- Ingress Class of Service Name Bandwidth Profile Flow.

[R166] Any Bandwidth Profile Flow at a UNI **MUST** be in at most 1 Envelope.

[R167] For a given Service Frame at a UNI, consider the set of Bandwidth Profile Flows to which this Service Frame belongs. The number of Bandwidth Profile Flows in this set that are in an Envelope **MUST** be at most 1.

[R167] means that:

- If the value of the EVC EP Ingress Bandwidth Profile Service Attribute (Section 10.8) for an EVC EP is not *None*, then the value of the EVC EP Class of Service Name Ingress Bandwidth Profile Service Attribute (Section 10.9) for that EVP EP is always *None*.
- If the value of the EVC EP Egress Bandwidth Profile Service Attribute (Section 10.10) for an EVC EP is not *None*, then the value of the EVC EP Class of Service Name Egress Bandwidth Profile Service Attribute (Section 10.11) for that EVP EP is always *None*.

For a given Service Frame at a UNI, consider the set of Bandwidth Profile Flows to which this Service Frame belongs. It is allowed that none of the Bandwidth Profile Flows in this set are in an Envelope. When this is the case, it is said that the Service Frame is not subject to a Bandwidth Profile at that UNI. When one of the Bandwidth Profile Flows in this set is in an Envelope per [R167], the given Service Frame is said to be subject to a Bandwidth Profile at that UNI. In this case, [R166] and [R167] mean that none of the other Bandwidth Profile Flows in the set are in an Envelope. Therefore the Bandwidth Profile algorithm of Section 12.2 is applied to this Service Frame exactly once at the UNI. As a consequence, at a given UNI, a Service Frame that is subject to a Bandwidth Profile has exactly one color declaration. Note that a Service Frame might be subject to a Bandwidth Profile at both the ingress UNI and at the egress UNI, and these might result in different color declarations.

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.

12.1 Bandwidth Profile Parameters

The following subsections detail the parameters and related requirements for the parameters used in the Bandwidth Profile Algorithm (Section 12.2).

12.1.1 Envelope Parameters

The details of the envelope parameters are contained in Section 9.12.

12.1.2 Bandwidth Profile Service Attribute Parameters

Table 28 shows the parameters that govern the operation of Bandwidth Profile Algorithm when applied to a Bandwidth Profile Flow that is in an Envelope. All 10 parameter values are agreed on by the Subscriber and Service Provider for the EVC EP Ingress Bandwidth Profile Service Attribute (Section 10.8) and the EVC EP Class of Service Name Ingress Bandwidth Profile Service Attribute (Section 10.9). Only CIR, CIR_{max}, and ER are agreed on by the Subscriber and Service Provider for the EVC EP Egress Bandwidth Profile Service Attribute (Section 10.10) and the EVC EP Class of Service Name Egress Bandwidth Profile Service Attribute (Section 10.11).

Note that the descriptions in the table are informal. The precise role played by a given parameter is determined by the Bandwidth Profile Algorithm (Section 12.2).

Parameter Name	Symbol	Units	Informal Description
Committed Information Rate	CIR	bits per second	When added to unused committed bandwidth* provided from higher-ranked Bandwidth Profile Flows (depending on the value of *CF* for the higher-ranked Bandwidth Profile Flows), limits the average rate in bits per second at which Service Frames for this Bandwidth Profile Flow can be declared Green.
Maximum Committed Information Rate	CIR_{max}	bits per second	Limits the average rate in bits per second at which Service Frames for this Bandwidth Profile Flow can be declared Green (regardless of unused committed bandwidth* from higher-ranked Bandwidth Profile Flows).
Committed Burst Size	CBS	bytes	Limits by how much, and for how long, the amount of traffic declared Green for this Bandwidth Profile Flow in the short term can exceed the committed bandwidth made available to this Bandwidth Profile Flow over the long term.
Excess Information Rate	EIR	bits per second	When added to unused excess bandwidth** from higher-ranked Bandwidth Profile Flows, and to unused committed bandwidth* (depending on the value of *CF* for this Bandwidth Profile Flow and CF^0 for the Envelope), limits the average rate in bits per second at which Service Frames for this Bandwidth Profile Flow can be declared Yellow.
Maximum Excess Information Rate	EIR_{max}	bits per second	Limits the average rate in bits per second at which Service Frames for this Bandwidth Profile Flow can be declared Yellow (regardless of unused excess bandwidth** from higher-ranked Bandwidth Profile Flows or unused committed bandwidth*).
Excess Burst Size	EBS	bytes	Limits by how much, and for how long, the amount of traffic declared Yellow for this Bandwidth Profile Flow in the short term can exceed the excess bandwidth made available to this Bandwidth Profile Flow over the long term.
Coupling Flag per Bandwidth Profile Flow	CF	integer	Determines whether unused committed bandwidth* for this Bandwidth Profile Flow is made available as excess bandwidth for this Bandwidth Profile Flow or as committed bandwidth for the next lower-ranked Bandwidth Profile Flow.
Color Mode	CM	string	Indicates whether Service Frames for this Bandwidth Profile Flow that are identified as Yellow on input to the Bandwidth Profile Algorithm can be declared Green or not.

Parameter Name	Symbol	Units	Informal Description
Envelope and Rank	ER	<Envelope ID, integer ≥ 1>	Identifies the Envelope that the Bandwidth Profile Flow belongs to, and the rank within the Envelope.
Token Request Offset	F	integer	Adjusts the bandwidth consumed by each Service Frame in the Bandwidth Profile Flow relative to the length of the Service Frame.

* The unused committed bandwidth for a Bandwidth Profile Flow is the difference between the limit resulting from the CIR for that Bandwidth Profile Flow (taking into account unused bandwidth from higher-ranked Bandwidth Profile Flows if applicable), and the actual information rate of Service Frames for that Bandwidth Profile Flow that are declared Green.
** The unused excess bandwidth for a Bandwidth Profile Flow is the difference between the limit resulting from the EIR for that Bandwidth Profile Flow (taking into account unused bandwidth from other Bandwidth Profile Flows if applicable), and the actual information rate of Service Frames for that Bandwidth Profile Flow that are declared Yellow.

Table 28 – Bandwidth Profile Flow Parameters

For ease of exposition, “BWP Flow Parameters” is used as shorthand for the 10-tuple $\langle CIR, CIR_{max}, CBS, EIR, EIR_{max}, EBS, CF, CM, ER, F\rangle$.

[R168] CIR **MUST** be ≥ 0.

[R169] CIR_{max} **MUST** be ≥ 0

[R170] If $CIR_{max} > 0$, then CBS **MUST** be greater than or equal to the value of the EVC Maximum Service Frame Size Service Attribute (Section 8.10).

[R171] EIR **MUST** be ≥ 0.

[R172] EIR_{max} **MUST** be ≥ 0.

[R173] If $EIR_{max} > 0$, then EBS **MUST** be greater than or equal to the value of the EVC Maximum Service Frame Size Service Attribute (Section 8.10).

[R174] The value of CF **MUST** be 0 or 1.

[R175] If $CF^0 = 1$ for an Envelope, then CF **MUST** equal 0 for all Bandwidth Profile Flows mapped to the Envelope.

[R176] The value of CM **MUST** be *color-blind* or *color-aware*.

[R177] The value of the rank in ER **MUST** be in the range $1,2,\ldots,n$ where n is the number of Bandwidth Profile Flows with the same Envelope ID value.

[R178] The value of the rank in ER **MUST NOT** equal the rank of any of the other Bandwidth Profile Flows with the same Envelope ID value.

12.2 Bandwidth Profile Algorithm

The Bandwidth Profile Algorithm applies to the Service Frames[37] in the Bandwidth Profile Flows that are mapped to a given Envelope. Operation of the Bandwidth Profile Algorithm is governed by CF^0 and the parameters, $\langle CIR^i, CIR^i_{max}, CBS^i, EIR^i, EIR^i_{max}, EBS^i, CF^i, CM^i, ER^i, F^i \rangle$ for $i = 1,2, \ldots, n$ where i identifies the rank. The algorithm declares a level of conformance relative to the Bandwidth Profile for each Service Frame. The level of conformance is expressed as one of three colors, Green, Yellow, or Red.[38]

A token bucket model is used to describe the Bandwidth Profile Algorithm 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. Committed tokens are sourced for Bandwidth Profile Flow i at a rate of up to CIR^i and excess tokens are sourced for Bandwidth Profile Flow i at a rate of up to EIR^i (either or both of which can be 0 per [R168] and [R171]). Tokens for a Bandwidth Profile Flow may be made available to other buckets depending on the setting of the Coupling Flags ($CF^i, i = 1,2, \ldots, n$) and what other Bandwidth Profile Flows are mapped to the Envelope. The rate of tokens flowing into the committed token bucket or the excess token bucked is limited to CIR^i_{max} and EIR^i_{max} 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 Service Frame Arrival Time, the Color of the Service Frame, and the number of tokens in the committed and excess token buckets. When a Service Frame for Bandwidth Profile Flow *i* is declared Green the number of tokens equal to the length of that Service Frame less the Token Request Offset, F^i, is removed from the committed bucket for that Bandwidth Profile Flow. When a Service Frame for Bandwidth Profile Flow *i* is declared Yellow the number of tokens equal to the length of that Service Frame less the Token Request Offset, F^i, is removed from the excess bucket for that Bandwidth Profile Flow.

An informal description of the Bandwidth Profile algorithm for an Envelope is shown in Figure 33. In the figure, $C^i_G(t)$ represents the number of Green tokens in the committed token bucket and $C^i_Y(t)$ represents the number of Yellow tokens in the excess token bucket for the Bandwidth Profile Flow with rank *i* at time *t*, for $i = 1,2, \ldots, n$. The details of the calculation of $C^i_G(t)$ and $C^i_Y(t)$ are specified in MEF 41 [35]. A visual representation of the token flows within the Envelope is presented in Appendix D.2.

[37] Consequently Data Service Frames, Layer 2 Control Protocol Service Frames and SOAM Service Frames can all be subject to a Bandwidth Profile.

[38] 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 SP Network but such procedures are beyond the scope of this Technical Specification.

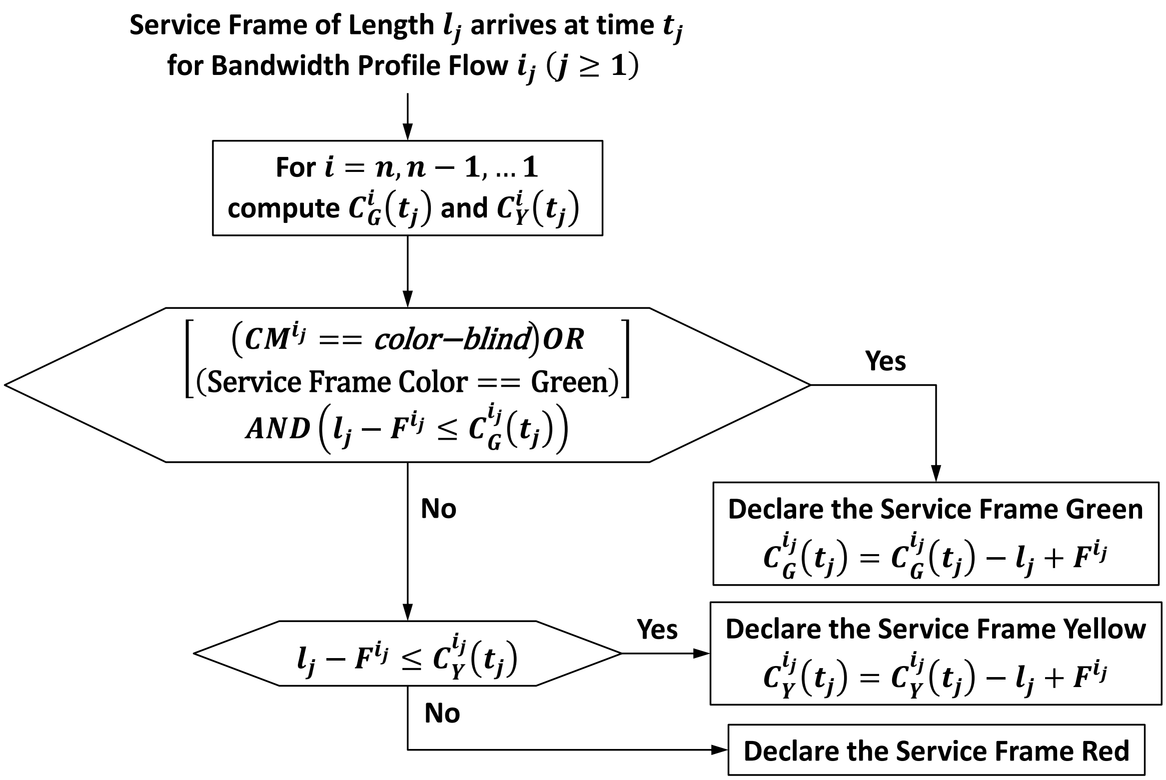

Figure 33 – Informal Description of the Bandwidth Profile Algorithm

The Bandwidth Profile Algorithm is specified by reference to the Generic Token Sharing Algorithm of MEF 41 [35]. Two things are needed to do this:

- The relationship of the parameter values in this document to the values of the parameters in MEF 41 [35], and

- The specification of the Token Requests used in MEF 41.

Table 29 shows the values of the General Parameters from Section 8.1 of MEF 41 [35].

MEF 41 [35] Parameter	Value for an Ingress Bandwidth Profile Flow	Value for an Egress Bandwidth Profile Flow
n	Number of Bandwidth Profile Flows in the Envelope	Number of Bandwidth Profile Flows in the Envelope
CF^0	CF^0	0

Table 29 – Values of MEF 41 General Parameters

Each MEF 41 TRF parameter (Section 8.2 of MEF 41[35]) takes on the value based a corresponding parameter value from this document as shown in Table 30.

MEF 41 [35] **Parameter**	**Value for an Ingress Bandwidth Profile Flow**	**Value for an Egress Bandwidth Profile Flow**
GTR^i	$\frac{CIR^i}{8}$	$\frac{CIR^i}{8}$
GTR^i_{max}	$\frac{CIR^i_{max}}{8}$	$\frac{CIR^i_{max}}{8}$
GTV^i	CBS^i	Service Provider Specified[39]
YTR^i	$\frac{EIR^i}{8}$	Service Provider Specified[39] for $i = n$, 0 for $i < n$
YTR^i_{max}	$\frac{EIR^i_{max}}{8}$	Service Provider Specified[39]
YTV^i	EBS^i	Service Provider Specified[39]
CF^i	CF^i	0

Table 30 – Values of MEF 41 TRF Parameters for Rank $i, i = 1,2,..., n$

A Token Request, as defined in MEF 41 [35], corresponds to each Service Frame that arrives at the UNI and that is contained in Bandwidth Profile Flow with rank i per Table 31.

MEF 41 [35] **Token Request Parameter**	**Value for an Ingress Bandwidth Profile Flow**	**Value for an Egress Bandwidth Profile Flow**
l	$L - F^i$ where L is the length in bytes of the Service Frame	$L - f^i$ where L is the length in bytes of the Service Frame and f^i is Service Provider Specified[39]
t	Ingress Service Frame Arrival Time at the UNI	Egress Service Frame Arrival Time at the UNI
c	If $CM^i =$ *color-aware* the color of the Service Frame per the value of the EVC EP Color Map Service Attribute (Section 10.6). Otherwise, *Green*	If the corresponding Ingress Service Frame was subject to a Bandwidth Profile, the value of the color declaration of the Ingress Bandwidth Profile Algorithm; else, the color of the corresponding Ingress Service Frame as determined per the EVC EP Color Map Service Attribute (Section 10.6)
r	i	i

Table 31 – Values for Token Requests for Bandwidth Profile Flow with Rank $i, i = 1,2,..., n$

Appendix D.4 describes examples of the use of the Token Request Offset, F^i.

[39] Value is specified by the Service Provider. The Subscriber can be told of the value at the discretion of the Service Provider.

[R179] If the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Physical*, let $\langle l_k, t_k, c_k, r_k \rangle, k = 0,1, \ldots$ be the sequence of Token Requests constructed per Table 31 for all Bandwidth Profile Flows in the Envelope such that $t_k < t_{k+1}, k = 0,1, \ldots.$ Then the color declaration for each Service Frame **MUST** equal the Color Determination for the corresponding Token Request using the algorithm of Section 10.2 in MEF 41 [35] with parameter values set as in Table 30.

[R180] If the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Virtual*, let $\langle s_k, t_k \rangle, k = 0,1, \ldots.$ be the subset of the output of the Subscriber UNI Virtual Frame Map Service Attribute (Section 9.3) such that the Service Frames $s_k, k = 0,1, \ldots.$ are in a Bandwidth Profile Flow in the Envelope, and let $\langle l_k, t_k, c_k, r_k \rangle, k = 0,1, \ldots$ be the sequence of Token Requests constructed per Table 31. Then the color declaration for each Service Frame, s_k, **MUST** equal the Color Determination for the corresponding Token Request using the algorithm of Section 10.2 in MEF 41 [35] with parameter values set as in Table 30.

Note that [R57] means that $t_{k+1} \geq t_k, k = 0,1, \ldots.$ in [R180].

Note that this document does not mandate how the behavior specified by [R179] and [R180] is achieved by the SP Network. 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 by [R179] and [R180]. A practical implementation needs to yield behavior “close” to the ideal behavior. Defining “close” is beyond the scope of this document. See Appendix D.5 for discussions of Bandwidth Profile implementation considerations.

12.3 Bandwidth Profile Algorithm With and Without Token Sharing

When there is more than one Bandwidth Profile Flow mapped to an Envelope, it is possible that tokens that are not used by a given Bandwidth Profile Flow can be made available to another Bandwidth Profile Flow. In other words, tokens can be shared. In such a case, the values of the parameters CIR^i_{max} and EIR^i_{max} can be set such that traffic in a higher rank Bandwidth Profile Flow doesn’t starve traffic in a lower rank Bandwidth Profile Flow. Appendix D.3 shows some examples of token sharing.

When there is a single Bandwidth Profile Flow mapped to an Envelope, token sharing is not possible. Furthermore, if $CIR^1_{max} \geq CIR^1$, $EIR^1_{max} \geq EIR^1 + (CF^1 \times CIR^1)$, and $F^1 = 0$, then the algorithm in Section 12.2 reduces to the algorithm of MEF 10.2 [23].

13. Subscriber Ethernet Service Framework

The values for the Service Attributes define the capabilities of a Subscriber Ethernet Service. Service Attributes are described in Sections 8, 9, 10, and 11.

For a particular Subscriber Ethernet Service, there are three types of Service Attributes, those that apply to an EVC as described in Section 8, those that apply to a UNI as described in Section 9, and those that apply to an EVC EP as described in Section 10. In addition, the Multiple EVC Service Level Specification Service Attribute described in Section 11 can be applied across multiple Subscriber Ethernet Services.

The Subscriber UNI Service Attributes are listed in Table 32 along with their possible values. For a given Subscriber Ethernet Service, the Subscriber and Service Provider need to agree on values for each of the Service Attributes in Table 32 for each UNI in the EVC associated with the service.

Service Attribute	Values
Subscriber UNI ID (Section 9.1)	A non-null RFC 2579 DisplayString no greater than 45 characters.
Subscriber UNI Instantiation (Section 9.2)	*Physical* or *Virtual*.
Subscriber UNI Virtual Frame Map (Section 9.3)	A map meeting [R56] and [R57] or *Not Applicable*.
Subscriber UNI List of Physical Links (Section 9.4)	*Not Applicable* or a non-empty list of 4-tuples of the form $\langle id, pl, fs, pt \rangle$.
Subscriber UNI Link Aggregation (Section 9.5)	*2-Link Active/Standby*, *All Active*, *Other*, or *Not Applicable*.
Subscriber UNI Port Conversation ID to Aggregation Link Map (Section 9.6)	Either a Port Conversation ID to Aggregation Link Map as defined in IEEE Std 802.1AX – 2014 [3] or *Not Applicable*.
Subscriber UNI Service Frame Format (Section 9.7)	*Ethernet MAC Frame conforming to Clause 3 of IEEE Std.802.3-2015* [5].
Subscriber UNI Maximum Service Frame Size (Section 9.8)	Integer ≥ 1522.
Subscriber UNI Maximum Number of EVC EPs (Section 9.9)	Integer ≥ 1.
Subscriber UNI Maximum Number of C-Tag VLAN IDs per EVC EP (Section 9.10)	Integer ≥ 1.
Subscriber UNI Token Share (Section 9.11)	*Enabled* or *Disabled*.
Subscriber UNI Envelopes (Section 9.12)	*None* or a non-empty list of pairs of the form <*x*,*y*> where *x* is an Envelope ID value and *y* is the Envelope Coupling Flag value.
Subscriber UNI Link OAM (Section 9.13)	*Enabled* or *Disabled*.
Subscriber UNI MEG (Section 9.14)	*Enabled or Disabled.*
Subscriber UNI LAG Link MEG Service Attribute (Section 9.15)	*Enabled or Disabled.*
Subscriber UNI L2CP Address Set (Section 9.16)	As described in MEF 45.1 [36].
Subscriber UNI L2CP Peering (Section 9.17)	As described in MEF 45.1 [36].

Table 32 – Subscriber UNI Service Attributes

The EVC Service Attributes are listed in Table 33 along with their possible values. For a given Subscriber Ethernet Service, the Subscriber and Service Provider need to agree on values for each of the Service Attributes in Table 33 for the EVC associated with the service.

Service Attribute	Values
EVC ID (Section 8.1)	A non-null RFC 2579 DisplayString no greater than 45 characters.
EVC List of EVC EPs (Section 8.2)	A list of EVC EP ID Service Attribute values.
EVC Type (Section 8.3)	*Point-to-Point*, *Multipoint-to-Multipoint*, or *Rooted-Multipoint.*
EVC Data Service Frame Disposition (Section 8.4)	A 3-tuple of the form $\langle u, m, b \rangle$ where each element in the 3-tuple has the value *Discard*, *Deliver Unconditionally*, or *Deliver Conditionally*.
EVC C-Tag PCP Preservation (Section 8.5)	*Enabled* or *Disabled.*
EVC C-Tag DEI Preservation (Section 8.6)	*Enabled* or *Disabled.*
EVC List of Class of Service Names (Section 8.7)	A non-empty list of Class of Service Names.
EVC Service Level Specification (Section 8.8)	Performance Objectives and parameters as described in Section 8.8.
EVC Group Membership (Section 8.9)	*None* or a non-empty list of 3-tuples of the form $\langle ID, CoS\_Name\_G, SG \rangle$ where: - ID is a string that is one of the values in an instance of the Multiple EVC Service Level Specification Service Attribute (Section 11). - CoS_Name_G is an entry in the value of the EVC List of Class of Service Names Service Attribute (Section 8.7) that is not *Discard*. - SG is a subset of ordered EVC EP pairs constructed from the value of the EVC List of EVC EPs Service Attribute (Section 8.2).
EVC Maximum Service Frame Size (Section 8.10)	Integer ≥ 1522.
EVC Available MEG Level (Section 8.11)	Integer from 0 to 7 or *None.*

Table 33 – EVC Service Attributes

The EVC EP Service Attributes are listed in Table 34 along with their possible parameter values. For a given Subscriber Ethernet Service, the Subscriber and Service Provider need to agree on values for each of the Service Attributes in Table 34 for each EVC EP in the EVC associated with the service.

Service Attribute	Values
EVC EP ID (Section 10.1)	A non-null RFC 2579 DisplayString no greater than 45 characters.
EVC EP UNI (Section 10.2)	Subscriber UNI ID Service Attribute value.
EVC EP Role (Section 10.3)	*Root* or *Leaf*.
EVC EP Map (Section 10.4)	*List* (along with a non-empty list of C-Tag VLAN ID values), *All*, or *UT/PT*.
EVC EP Ingress Class of Service Map (Section 10.5)	A 3-tuple of the form $\langle F, M, P \rangle$ where: • *F* is one of the values *EVC EP*, *C-Tag PCP*, or *DSCP*, • *M* is a map that can be used to assign Class of Service Names to Service Frames. (The form of *M* depends on the value of *F*.), and • *P* is a map with entries of the form <Layer 2 Control Protocol type → Class of Service Name> where the Layer 2 Control Protocol type is determined by the Protocol Identifier.
EVC EP Color Map (Section 10.6)	A pair of the form <*F*, *M*> where: • *F* is one of *EVC EP*, *C-Tag DEI*, *C-Tag PCP*, or *DSCP*, • *M* is a mapping that allows a Color to be assigned to each Ingress Service Frame (The form of *M* depends on the value of *F*.).
EVC EP Egress Map (Section 10.7)	*None* or a map that determines the content of the C-Tag PCP and C-Tag DEI of each Egress Service Frame.
EVC EP Ingress Bandwidth Profile (Section 10.8)	*None* or BWP Flow Parameters (Section 12.1.2).
EVC EP Class of Service Name Ingress Bandwidth Profile (Section 10.9)	*None* or a non-empty list of pairs of the form $\langle x, y \rangle$ where x is a Class of Service Name and y = BWP Flow Parameters (Section 12.1.2).
EVC EP Egress Bandwidth Profile (Section 10.10)	*None* or the 3-tuple $\langle CIR, CIR_{max}, ER \rangle$.
EVC EP Class of Service Name Egress Bandwidth Profile (Section 10.11)	*None* or a non-empty list of pairs of the form $\langle x, y \rangle$ where x is an Class of Service Name and $y = \langle CIR, CIR_{max}, ER \rangle$.
EVC EP Source MAC Address Limit (Section 10.12)	*None* or the pair $\langle N, \tau \rangle$ where N is an integer ≥ 1 and τ is a time duration.
EVC EP Subscriber MEG MIP (Section 10.13)	*None* or an integer in the range $0 - 7$.

Table 34 – EVC EP Service Attributes

The Multiple EVC Group Availability Service Level Specification Service Attribute (Section 11) can be applied to multiple instances of Subscriber Ethernet Service. The value for this Service Attribute is $\langle ID, KG, \widehat{MGA} \rangle$ as described in Section 11.

14. References

[1] IEEE Std 1588 – 2008, *IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems*, 24 July 2008.

[2] IEEE Std 802.1AX – 2008, *IEEE Standard for Local and metropolitan area networks – Link Aggregation*, November 2008.

[3] IEEE Std 802.1AX – 2014, *IEEE Standard for Local and metropolitan area networks – Link Aggregation*, December 2014.

[4] IEEE Std 802.1Q – 2018, *(Revision of IEEE Std 802.1Q-2014) - IEEE Standard for Local and metropolitan area networks--Bridges and Bridged Networks*, May 2018.

[5] IEEE Std 802.3 – 2015, *IEEE Standard for Ethernet*, 3 September 2015.

[6] IEEE Std 802.3bp – 2016, *IEEE Standard for Ethernet Amendment 4: Physical Layer Specifications and Management Parameters for 1 Gb/s Operation over a Single Twisted-Pair Copper Cable*, 30 June 2016.

[7] IEEE Std 802.3bq – 2016, *IEEE Standard for Ethernet Amendment 3: Physical Layer and Management Parameters for 25 Gb/s and 40 Gb/s Operation, Types 25GBASE-T and 40GBASE-T*, 30 June 2016.

[8] IEEE Std 802.3bs – 2017, *IEEE Standard for Ethernet Amendment 10: Media Access Control Parameters, Physical Layers, and Management Parameters for 200 Gb/s and 400 Gb/s Operation*, 12 December 2017.

[9] IEEE Std 802.3bu – 2016, *IEEE Standard for Ethernet Amendment 8: Physical Layer and Management Parameters for Power over Data Lines (PoDL) of Single Balanced Twisted-Pair Ethernet*, 7 December 2016.

[10] IEEE Std 802.3bv – 2017, *IEEE Standard for Ethernet Amendment 9: Physical Layer Specifications and Management Parameters for 1000 Mb/s Operation Over Plastic Optical Fiber*, 14 February 2017.

[11] IEEE Std 802.3by – 2016, *IEEE Standard for Ethernet Amendment 2: Media Access Control Parameters, Physical Layers, and Management Parameters for 25 Gb/s Operation*, 30 June 2016.

[12] IEEE Std 802.3bz – 2016, *IEEE Standard for Ethernet Amendment 7: Media Access Control Parameters, Physical Layers, and Management Parameters for 2.5 Gb/s and 5 Gb/s Operation, Types 2.5GBASE-T and 5GBASE-T*, 22 September 2016.

[13] IEEE Std 802.3cc – 2017, *IEEE Standard for Ethernet Amendment 11: Physical Layer and Management Parameters for Serial 25 Gb/s Ethernet Operation over Single-Mode Fiber*, 15 January 2018.

[14] IETF RFC 2119, *Key words for use in RFCs to Indicate Requirement Levels*, March 1997.

[15] IETF RFC 2579, *Textual Conventions for SMIv2*, April 1999.

[16] IETF RFC 3260, *New Terminology and Clarification for Diffserv*, April 2002.

[17] IETF RFC 8174, *Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words*, May 2017.

[18] ITU-T Recommendation G.8262/Y.1362, *Timing characteristics of synchronous Ethernet equipment slave clock*, January 2015.

[19] ITU-T Recommendation G.8264/Y.1364, *Distribution of timing information through packet networks*, May 2014.

[20] ITU-T Recommendation Y.1563, *Ethernet frame transfer and availability performance*, 2009.

[21] ITU-T Recommendation Y.1564, *Ethernet service activation test methodology*, February 2016.

[22] MEF 6.2, *Ethernet Services Definitions - Phase 3*, August 2014.

[23] MEF 10.2, *Ethernet Services Attributes Phase 2*, October 2009.

[24] MEF 10.3, *Ethernet Services Attributes Phase 3*, October 2013.

[25] MEF 10.3.1, *Composite Performance Metric (CPM) Amendment to MEF 10.3*, February 2015.

[26] MEF 10.3.2, *Amendment to MEF 10.3 - UNI Resiliency Enhancement*, October 2015.

[27] MEF 12.2, *Carrier Ethernet Network Architecture Framework Part 2: Ethernet Services Layer*, May 2014.

[28] MEF 17, *Service OAM Framework and Requirements – Phase 1*, April 2007.

[29] MEF 23.2, *Carrier Ethernet Class of Service – Phase 3*, August 2016.

[30] MEF 23.2.1, *Models for Bandwidth Profiles with Token Sharing Amendment to MEF 23.2*, January 2017.

[31] MEF 26.2, *External Network Network Interfaces (ENNI) and Operator Service Attributes*, August 2016.

[32] MEF 30.1, *Service OAM Fault Management Implementation Agreement: Phase 2*, April 2013.

[33] MEF 30.1.1, *Amendment to MEF 30.1 – Correction to Requirement*, April 2014.

[34] MEF 35.1, *Service OAM Performance Monitoring Implementation Agreement*, May 2015.

[35] MEF 41, *Generic Token Bucket Algorithm*, October 2013.

[36] MEF 45.1, *Layer 2 Control Protocols in Ethernet Services*, December 2018.

Appendix A Examples (Informative)

This appendix contains examples of some of the Service 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 Examples of the Use of the EVC EP Map Service Attribute and EVCs

This section presents examples of the use of EVCs and the EVC EP Map Service Attribute. It is intended to clarify the concepts and present likely deployment scenarios.

A.1.1 Untagged UNI

In connecting branch enterprise locations to a hub enterprise location, it is desirable to make the configuration of the branch SNs simple. A similar objective applies to providing access to higher layer services, e.g., Internet Access, where the configuration of the SN at the sites accessing the service should be kept simple. Figure 34 shows an example of 3 UNIs (A, B, and C) where SNs only capable of handling Untagged Service Frames are attached. The EVC EP Map Service Attribute values are shown for each EVC EP. Note that UNIs A, B, and C are Untagged UNIs and that UNI D is a VLAN Tagged UNI.

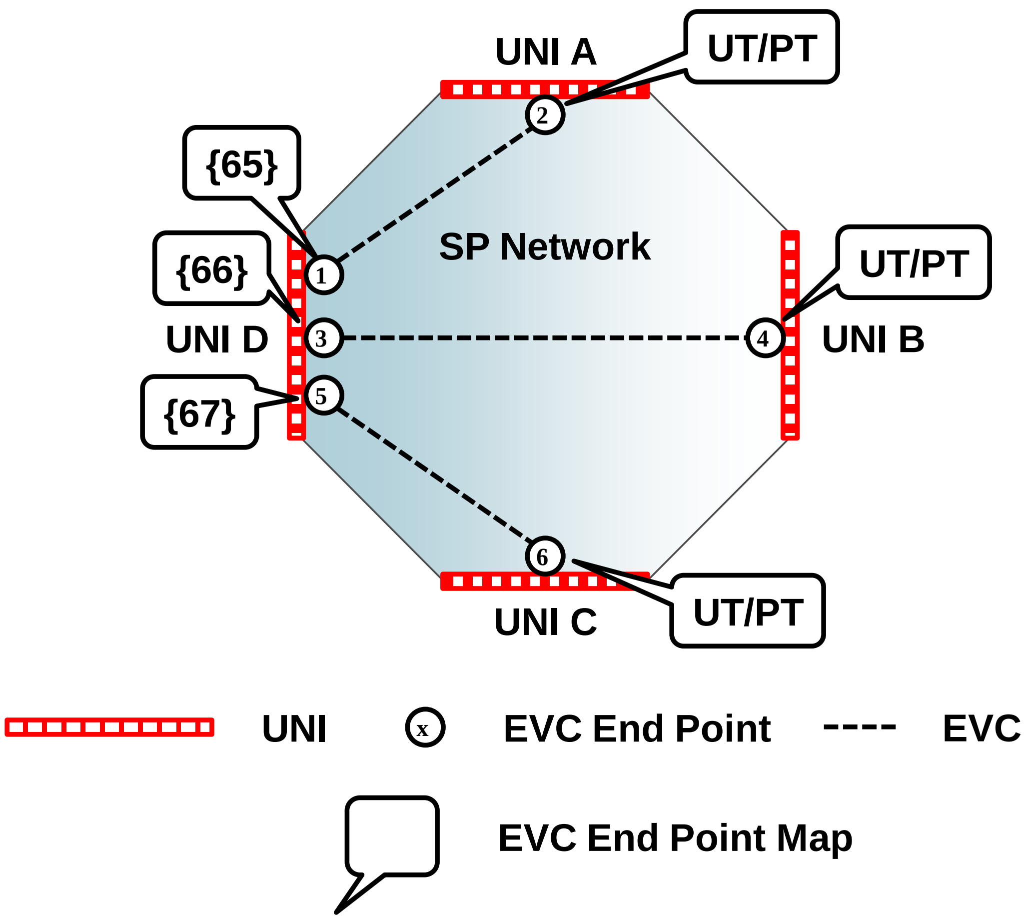

Figure 34 – Example of the Use of the Untagged UNI

Consider an ingress Untagged Service Frame at UNI A. It will be mapped to EVC EP 2 and delivered to EVC EP 1 located at UNI D. At UNI D, it will become a VLAN Tagged Egress Service Frame with C-Tag VLAN ID value 65. An ingress VLAN Tagged Service Frame at UNI D with C-Tag VLAN ID value 65 will be mapped to EVC EP 1 and delivered to EVC EP 2 located at UNI A. At UNI A, it will become an egress Untagged Service Frame per [R116].

A.1.2 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 SN can only send to the higher layer service SN 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 SN does not necessarily have to send and receive C-Tagged Service Frames. In particular, UNIs A and C are Untagged UNIs in this example.

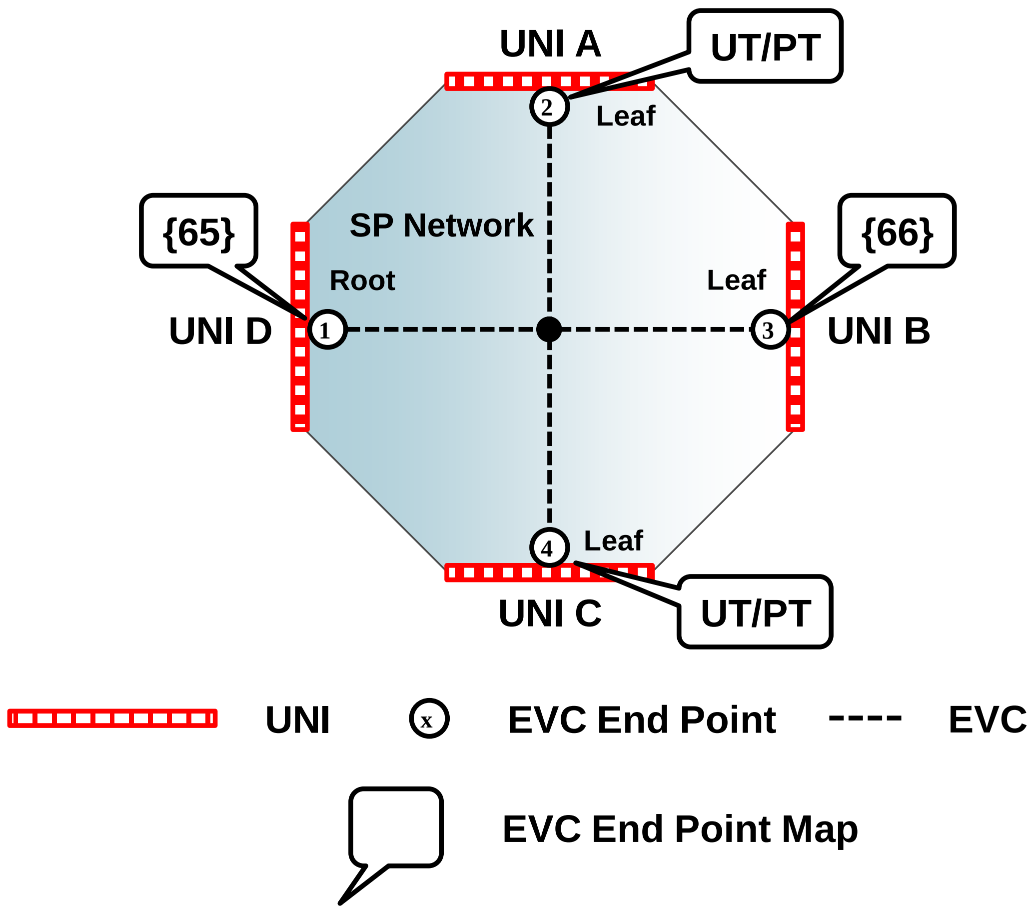

Figure 35 – Example of the Use of a Rooted-Multipoint EVC

A.1.3 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.

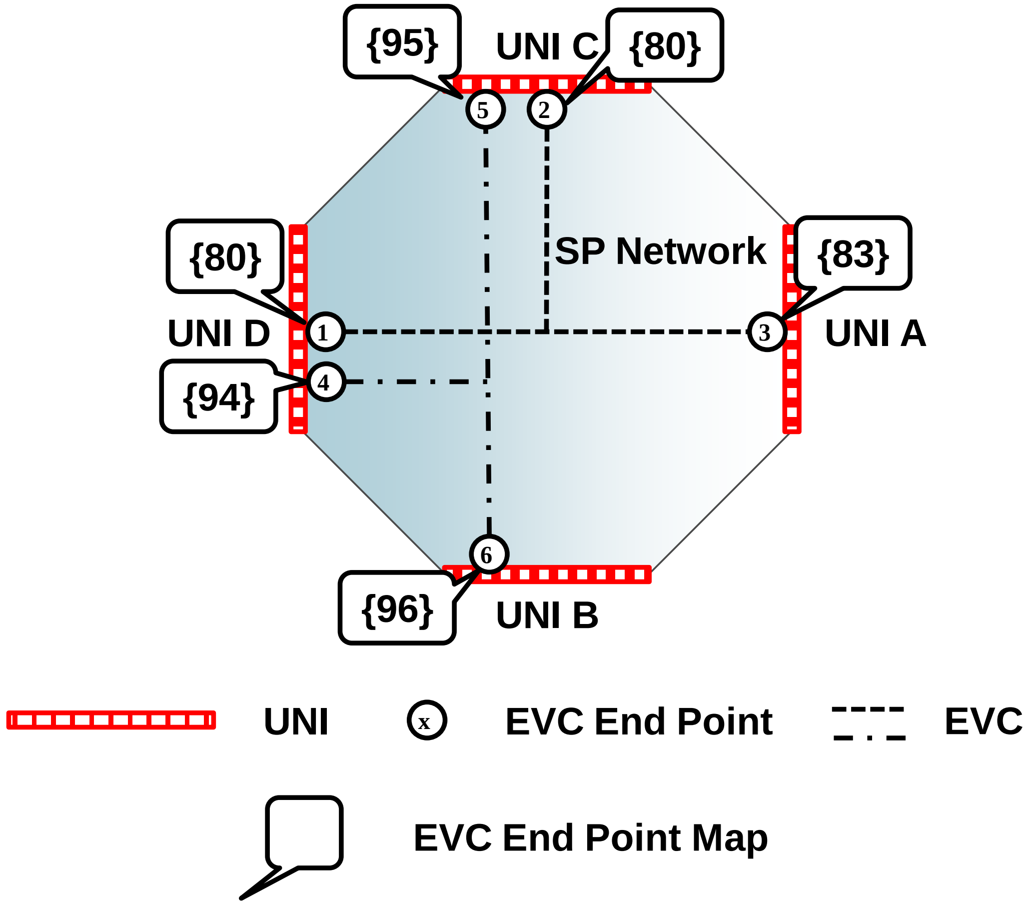

Figure 36 – Example of the Use of Redundant Higher Layer Service Access

A.1.4 EVC EP Map Service Attribute Value = *All*

Figure 37 shows an example when the value of the EVC EP Service Attribute (Section 10.4.2) = *All*.

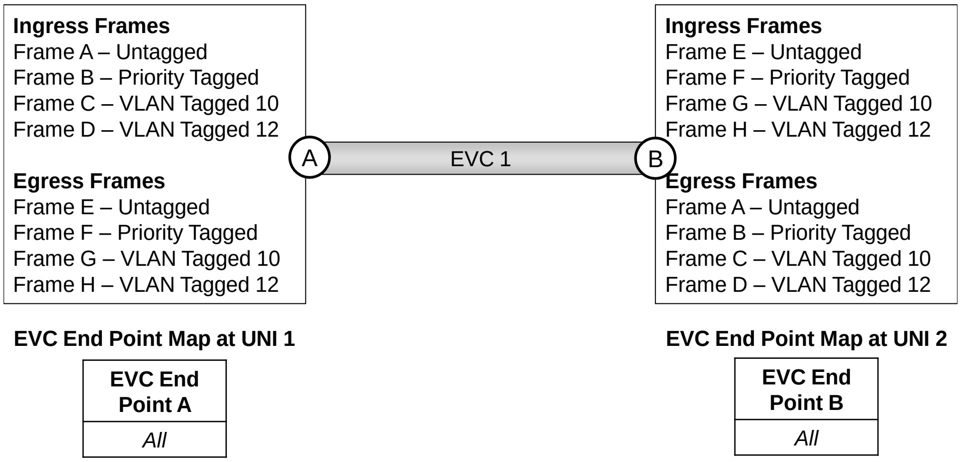

Figure 37 – Example of the EVC EP Map Service Attribute Value = *All*

A.1.5 EVC EP Map Service Attribute = *List*

Figure 38 shows examples of when the value of the EVC EP Map Service Attribute (Section 10.4.1) = *List*. Note that [R110] mandates that EVC EPs C and E have the same value of the EVC EP Map Service Attribute.

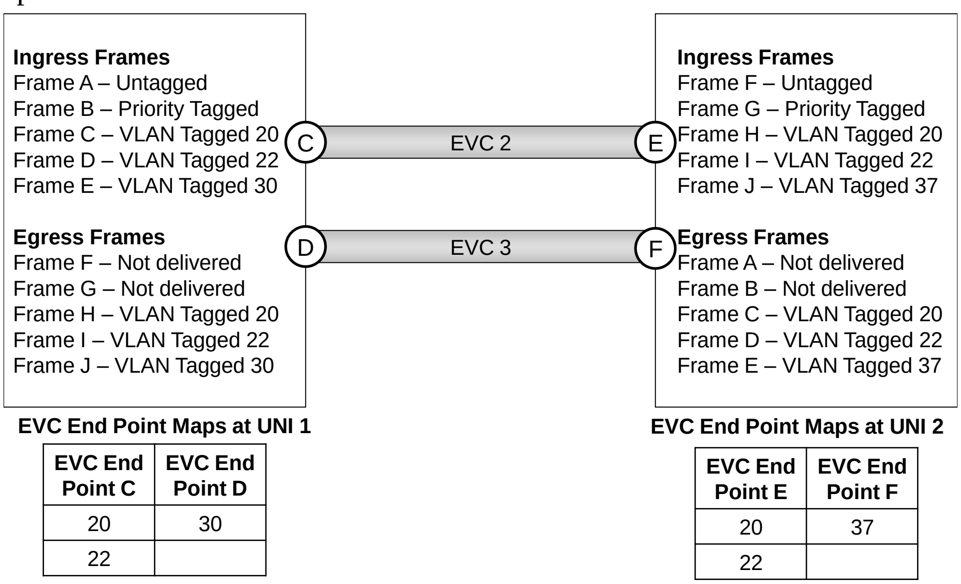

Figure 38 – Examples of the EVC EP Map Service Attribute = *List*

A.1.6 EVC EP Map Service Attribute Value = *UT/PT*

Figure 39 and Figure 40 show examples when the value of the EVC EP Map Service Attribute (Section 10.4.3) = *UT/PT*.

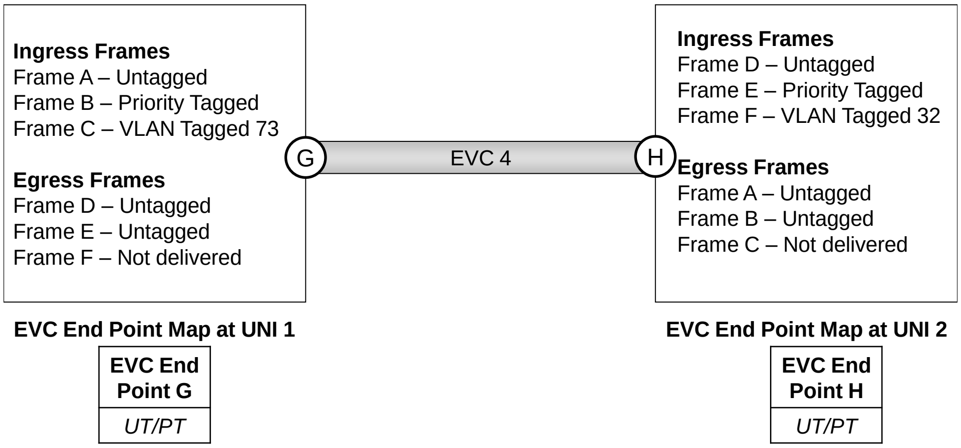

Figure 39 – Example 1 of the EVC EP Service Attribute Value = *UT/PT*

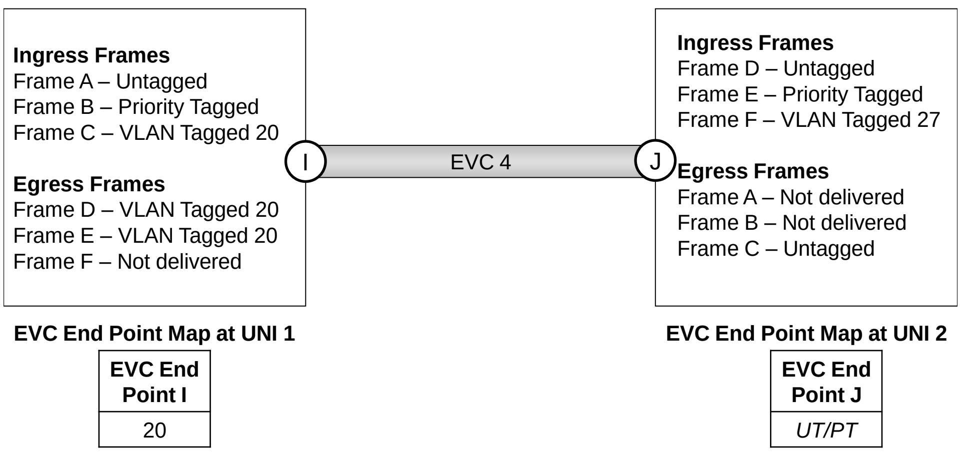

Figure 40 – Example 2 of the EVC EP Service Attribute Value = *UT/PT*

A.2 Examples of Availability Metrics for Multipoint-to-Multipoint EVCs

The Performance Metric definitions for Multipoint-to-Multipoint EVCs provide a great deal of flexibility. This section provides examples on how the subset of EVC EPs in the EVC can be used to define UNI-oriented Performance Metrics (Section A.2.1) and EVC-oriented Performance Metrics (Section A.2.2). The Availability Performance Metric is used for these examples.

Both examples use the Multipoint-to-Multipoint EVC depicted in Figure 41. 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.

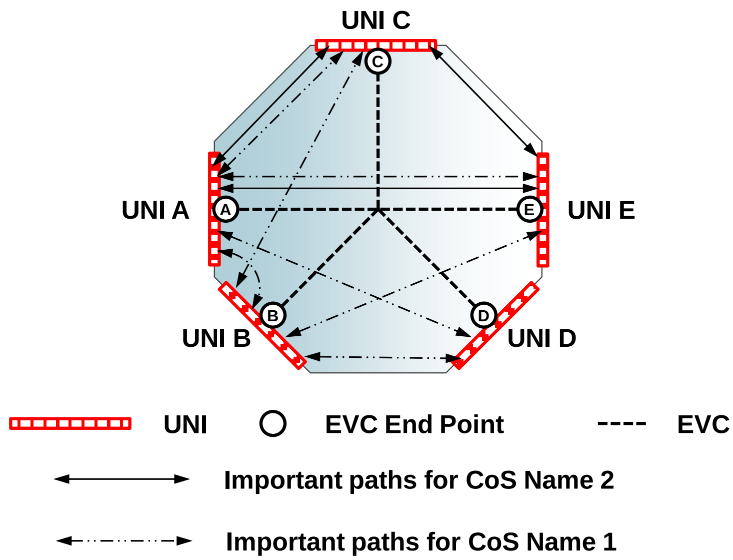

Figure 41 – Multipoint-to-Multipoint EVC Example

A.2.1 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 from the UNI in question and the other UNIs identified by the important traffic flows. Define the subsets of ordered EVC EP pairs $S_{x,y}$ where x represents a UNI identifier and y represents a Class of Service Name:

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

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

$$S_{C,1} = \{\langle C,A\rangle, \langle C,B\rangle\}$$

$$S_{D,1} = \{\langle D,A\rangle, \langle D,B\rangle\}$$

$$S_{E,1} = \{\langle E,A\rangle, \langle E,B\rangle\}$$

$$S_{A,2} = \{\langle A,C\rangle, \langle A,E\rangle\}$$

$$S_{C,2} = \{\langle C,A\rangle, \langle C,E\rangle\}$$

$$S_{E,2} = \{\langle E,A\rangle, \langle E,C\rangle\}$$

For this example, assume that T_l, Δt, C, and n, are used for all availability Performance Metrics. Then using the definition in Section 8.8.7, $\bar{A}(S_{A,1}, T_l)$ can be viewed as the availability of UNI A for Class of Service Name 1 and this reflects the availability from UNI A to the other UNIs identified by the important traffic flows. Similarly, $\bar{A}(S_{C,2}, T_l)$ 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 EVC EP pairs.

A.2.2 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 subsets of ordered EVC EP pairs S_y where y represents a Class of Service Name:

$$S_1 = \left\{ \begin{matrix} \langle A,B\rangle, \langle B,A\rangle, \langle A,C\rangle, \langle C,A\rangle, \langle A,D\rangle, \langle D,A\rangle, \langle A,E\rangle, \langle E,A\rangle, \\ \langle B,C\rangle, \langle C,B\rangle, \langle B,D\rangle, \langle D,B\rangle, \langle B,E\rangle, \langle E,B\rangle \end{matrix} \right\}$$

$$S_2 = \{\langle A,C\rangle, \langle C,A\rangle, \langle A,E\rangle, \langle E,A\rangle, \langle C,E\rangle, \langle E,C\rangle\}$$

For this example, assume that T_l, Δt, C, and n, are used for both availability Performance Metrics. Then using the definition in Section 8.8.7, $\bar{A}(S_1, T_l)$ can be viewed as the availability of Class of Service Name 1 on the EVC and $\bar{A}(S_2, T_l)$ can be viewed as the availability of Class of Service Name 2 on the EVC.

Appendix B Traffic Shaping Example (Informative)

Shaping is a procedure to reduce the burstiness of traffic. When done in the SN it is meant to increase the level of Ingress Bandwidth Profile conformance (Sections 10.8 and 10.9). A shaper is defined by a set of parameters. Those parameters should be chosen to ensure that the delay introduced by the 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 of Ingress Service Frames that is the only Bandwidth Profile Flow mapped to the Envelope.

Two example algorithms are presented that together describe a shaper implementation in the SN. These are the Periodic Algorithm shown in Figure 42 and the New Frame Algorithm shown in Figure 43. In the New Frame Algorithm, a limited number of Yellow frames can be placed in the shaper buffer[40] (and subsequent Yellow frames will be dropped if required). This controls the delay that may be experienced by Green frames due to the presence of Yellow frames. There may be Yellow frames ahead of a Green frame in the transmission buffer, but that will typically not add significant delay.

The following parameters are used in the example algorithms:

- CIR = the shaping rate of Green frames (average output rate of the shaper),
- CBS = the shaping burst size of Green frames (maximum output burst of the shaper),
- EIR = the shaping rate of Yellow frames (average output rate of the shaper),
- EBS = the shaping burst size of Yellow frames (maximum output burst of the shaper),
- CF = the Coupling Flag that controls whether tokens that overflow the Committed token bucket are added to the Excess token bucket,
- CM = the Color Mode can be either *color-blind* or *color-aware* and controls whether the color of a frame is considered in determining when to transfer it from the shaper buffer to the transmission buffer,
- SBL = the Shaper Buffer Limit is the depth of the shaper buffer (in bytes) above which no new frames will be placed in the shaper buffer,
- YBL = the Yellow Buffer Limit is the depth of the shaper buffer (in bytes) above which no new yellow frames will be placed in the shaper buffer,
- Δt = the time duration between the executions of the Periodic Algorithm, and

[40] Note that we differentiate between the shaper buffer and the transmission buffer (outgoing link queue). Frames taken from the head of the shaper buffer are queued in the transmission buffer for transmission at MAC Data Rate.

- $F =$ the token request offset that is subtracted from the length of the frame when tokens are deducted upon sending the frame to the transmission buffer.

The *CM* parameter only affects whether the color of a frame is considered when removing that frame from the shaper buffer. Whether the color of a frame is considered when placing that frame into the shaper buffer is controlled by *YBL* and *SBL*. If *YBL* = *SBL* then yellow and green frames receive identical treatment when determining whether to place the frame in the shaper buffer (i.e. color blind behavior), and otherwise the color affects whether to place the frame in the shaper buffer (i.e. color aware behavior).

YBL is the maximum burst of yellow frames that can be placed in the shaper buffer. *SBL* is the maximum burst of all frames (green and yellow) that can be placed in the shaper buffer. When the shaper buffer is empty, at least *SBL* – *YBL* green frames in a burst can be placed in the shaper buffer. Typically the shaper buffer limits are configured such that *SBL* – *YBL* ≥ *CBS*, which means the shaper accepts larger bursts of green frames at its input and generates smaller bursts of green frames at its output.

The maximum delay that can be experienced by a green frame is *SBL* divided by *CIR*, provided that *CIR* is the minimum rate at which frame are moved from the shaper buffer to the transmission buffer. This will be the case as long as either *CF* = 1 (so tokens overflowing the Committed token bucket are added to the Excess token bucket) or *EIR* ≥ *CIR*.

Yellow frames arriving at the shaper will only be transferred to the transmission buffer using yellow tokens. Green frames arriving at the shaper will be transferred to the transmission buffer using green tokens when they are available, or yellow tokens when no green tokens are available.

The following notation is used in the example algorithms where t represents the time that the algorithm is run:

- $B(t)$ = the instantaneous shaper buffer occupancy in bytes,
- $C(t)$ = the instantaneous value of the tokens in the Committed token bucket with $C(0) = CBS$,
- $E(t) =$ the instantaneous value of the tokens in the Excess token bucket with $E(0) = EBS$,
- $L =$ the length of the frame at the head of the shaper buffer,
- $LNF =$ the length of the newly-arrived frame,
- $CNF =$ the color of the newly arrived frame, and
- $CHF =$ color of the frame at the head of the shaper buffer.

The Periodic Algorithm is shown in Figure 42. It is executed every Δt time units and transfers frames to the transmission buffer in times between the arrival of new frames when sufficient tokens become available.


```
C(t)=min(CBS,(C(t)+(CIR/8)*Δt)); // Update green token count
O(t)=max(0,(C(t)+(CIR/8)*Δt-CBS)); // Compute overflow green tokens
E(t)=min(EBS,(E(t)+(EIR/8)*Δt+CF*O(t))); // Update Yellow token count
while((((L-F)<=C(t))&&(B(t)>0)&&(CHF==Green||CM=Color_Blind))||
       (((L-F)<=E(t))&&(B(t)>0)))
       {
       if(((L-F)<=C(t))&&(B(t)>0)&&(CHF==Green||CM=Color_Blind))
              {
              //Transmit using Green tokens
              C(t)=C(t)-(L-F);
              B(t)=B(t)-L;
              send the frame at the head of the shaper buffer to the
              transmission buffer;
              }
       else
              {
              //Transmit using Yellow tokens
              E(t)=E(t)-(L-F);
              B(t)=B(t)-L;
              Send the frame at the head of the shaper buffer to the
              transmission buffer;
              }
       }
```

Figure 42 – Periodic Algorithm Example

The New Frame Algorithm is shown in Figure 43. It is executed each time a new frame arrives at the shaper. For simplicity it is assumed that $CIR > 0$ and at least one of the following is true:

- $CF = 1$
- $EIR > 0$.

These assumptions mean:

- A Yellow frame with length less than EBS will not get stuck in the shaper buffer because enough Yellow tokens can be accumulated to move it to the transmission buffer via the Periodic Algorithm.
- A Green frame with length less than $max(CBS, EBS)$ will not get stuck in the shaper buffer because enough tokens, either Green or Yellow, can be accumulated to move it to the transmission buffer via the Periodic Algorithm.


```
if(B(t)==0) // If shaper buffer is empty
       {
       if((CNF==Yellow)&&(CM=Color_Aware))
              {
              if((LNF-F)<=E(t))
                     {
                     //Transmit using Yellow tokens
                     E(t)=E(t)-(LNF-F);
                     send new frame to transmission buffer;
                     }
              else if((B(t)+LNF<=YBL)&&((LNF-F)<=EBS))//B(t)=0 so only need LNF
                     {
                     //Room in the shaper buffer for the new frame and it is
                     //possible to accumlate enough Yellow tokens to transmit
                     //the frame in the Periodic Algorithm
                     place new frame in shaper buffer;
                     B(t) = B(t) + LNF; //B(t)=0 so could use B(t)=LNF
                     }
              else
                     {
                     discard new frame;
                     }
              }
       else // New frame is Green or shaper is configured to be Color_Blind
              {
              if((LNF-F)<=C(t))
                     {
                     //Transmit using Green tokens
                     C(t)=C(t)-(LNF-F);
                     send new frame to transmission buffer;
                     }
              else if((LNF-F)<=E(t))
                     {
                     //Transmit using yellow tokens
                     E(t)=E(t)-(LNF-F);
                     send new fame to transmission buffer;
                     }
              else if((B(t)+LNF<=SBL)&&((LNF-F)<=max(CBS,EBS)))
                     {
                     //Room in the shaper buffer for the new frame and it is
                     //possible to accumlate enough tokens to transmit
                     //the frame in the Periodic Algorithm
                     place new frame in shaper buffer;
                     B(t)=B(t)+LNF; //B(t)=0 so could use B(t)=LNF
                     }
              else
                     {
                     discard new frame;
                     }
              }
       }
else // shaper buffer is not empty
       {
       if((CNF==Yellow)&&(CM=Color_Aware))
              {
              if((B(t)+LNF<=YBL)&&((LNF-F)<=EBS))
                     {
```



```
                //Room in the shaper buffer for the new frame and it is
                //possible to accumlate enough Yellow tokens to transmit
                //the frame in the Periodic Algorithm
                place new frame in shaper buffer;
                B(t) = B(t) + LNF;
                }
          else
                {
                discard the new frame;
                }
          }
    else if((B(t)+LNF<=SBL)&&((LNF-F)<=max(CBS,EBS)))
          {
          //Room in the shaper buffer for the new frame and it is
          //possible to accumlate enough tokens to transmit
          //the frame in the Periodic Algorithm
          place new frame in shaper buffer;
          B(t)=B(t)+LNF;
          }
    else
          {
          discard the new frame;
          }
    }
```

Figure 43 – New Frame Algorithm Example

Appendix C Effect of Service Frame Overhead (Informative)

This Appendix assumes that the value of the Subscriber UNI Instantiation Service Attribute (Section 9.2) = *Physical* and the Service Frames are carried on a single link.

Figure 44 is a simplified version of Figure 3 – 1 in IEEE Std 802.3 – 2015 [5] which shows the format of the IEEE Std 802.3 – 2015 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 44 the Service Frame is exactly the IEEE Std 802.3 – 2015 Frame. IEEE Std 802.3 – 2015 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 the UNI is 20 bytes.

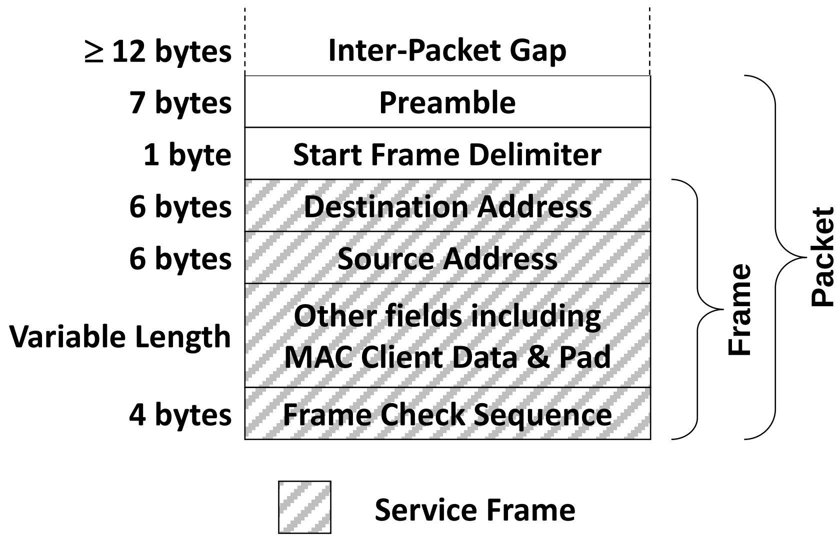

Figure 44 – IEEE Std 802.3 Packet Format (simplified)

Information Rate is defined as the average bit rate of Ethernet frames at the measurement point, where each Ethernet frame is measured as 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 44) 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 MAC Data 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 the Information Rate for a group of Service Frames is

$$\frac{S}{T_2 - T_1 + \frac{160}{r}}$$

For a 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,

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

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

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

Figure 45 shows the ratio of the Information Rate to UNI MAC Data Rate as a function of Service Frame size for a group of fixed length, back to back Service Frames.

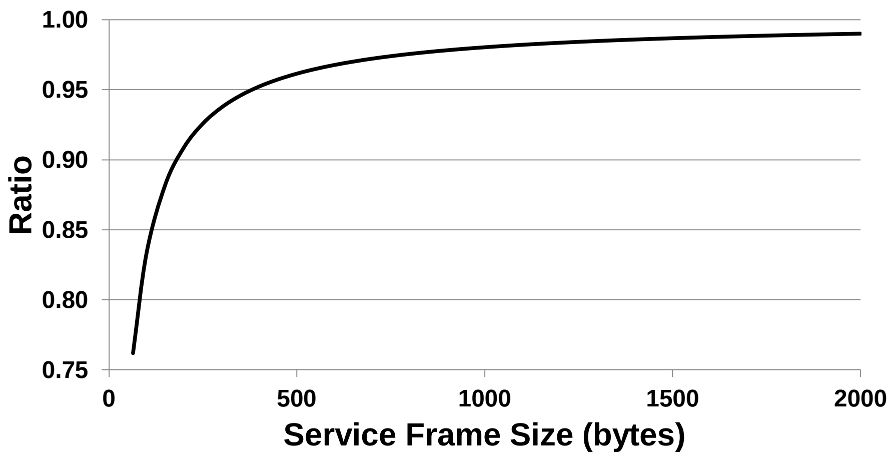

Figure 45 – Ratio of Information Rate to UNI MAC Data Rate for Back to Back Service Frames

Figure 45 has implications on the behavior of the Bandwidth Profile Algorithm of Section 12.2:

- 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* that is less than the UNI MAC Data 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* to be .85 of the UNI MAC Data Rate would yield this behavior.

- When the egress UNI MAC Data Rate is smaller than that of the ingress UNI and an EVC End Point has an ingress *CIR* approaching the egress UNI MAC Data Rate, a burst of back to back small frames may exceed the egress UNI MAC Data Rate and thus could be dropped at the egress UNI even though their Information Rate is within *CIR*. See Appendix D.4.2 for an example.

Appendix D Bandwidth Profiles (Informative)

D.1 Example of Bandwidth Profile Flows and Envelopes

There are many ways that the EVC EP Ingress Bandwidth Profile Service Attribute (Section 10.8), the EVC EP Class of Service Name Ingress Bandwidth Profile Service Attribute (Section 10.9), the EVC EP Egress Bandwidth Profile Service Attribute (Section 10.10), and the EVC EP Class of Service Name Egress Bandwidth Profile Service Attribute (Section 10.11) can co-exist at a UNI.

Figure 46 shows an example for the EVC EP Ingress Bandwidth Profile Service Attribute (Section 10.8) and the EVC EP Class of Service Name Ingress Bandwidth Profile Service Attribute (Section 10.9) co-existing at the UN. There are:

- Four EVC EPs,
- Two instances of the EVC EP Ingress Bandwidth Profile Service Attribute corresponding to the Bandwidth Profile Flows A and B,
- Four instances of the EVC EP Class of Service Name Ingress Bandwidth Profile Service Attribute corresponding to the Bandwidth Profile Flows α, β, γ, and δ and,
- Three Envelopes labeled 1, 2, and 3.

Each of the Envelopes contains Bandwidth Profile Flows of the same type as required by [R165]. All three Envelopes contain Bandwidth Profile Flows from two EVC EPs.

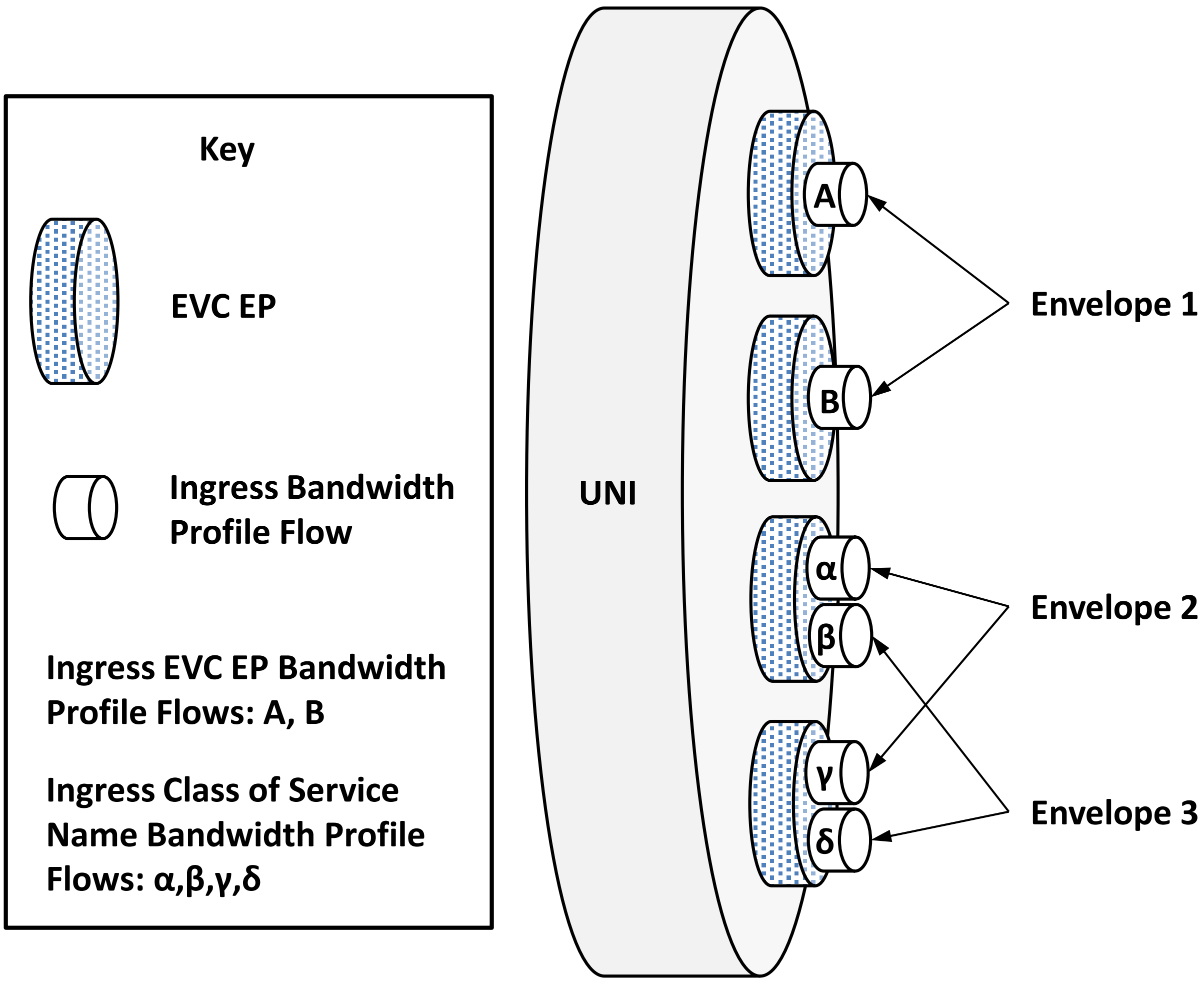

Figure 46 – Example of Bandwidth Profile Flows and Envelopes

D.2 Visual Representation of the Bandwidth Profile Token Flows

Figure 47 shows a conceptual diagram of the Bandwidth Profile Algorithm with three Bandwidth Profile Flows.

- The green and yellow trapezoid icons labeled CBS^i and EBS^i represent the Committed and Excess token buckets for each Bandwidth Profile Flow.
- The green and yellow flag-shaped icons labeled CIR^i or EIR^i represent a token source flowing into each bucket at the rate specified by the Bandwidth Profile.
- The red funnel-shaped icons labeled CIR^i_{max} and EIR^i_{max} represent the limits on how fast tokens can be added to a token bucket.
- The dashed arrows show possible token paths. An arrow that originates at a funnel represents tokens that don't enter a token bucket due to CIR^i_{max} or EIR^i_{max}. An arrow that originates at a trapezoid represent tokens that overflow the bucket when the number of tokens in the bucket reaches the capacity (CBS^i or EBS^i).

- The ellipses labeled CF^i represent decision points in the path of the tokens that are controlled by the Coupling Flag.

- The X's represent points at which tokens are discarded (i.e., not added to the token count for any of the buckets).

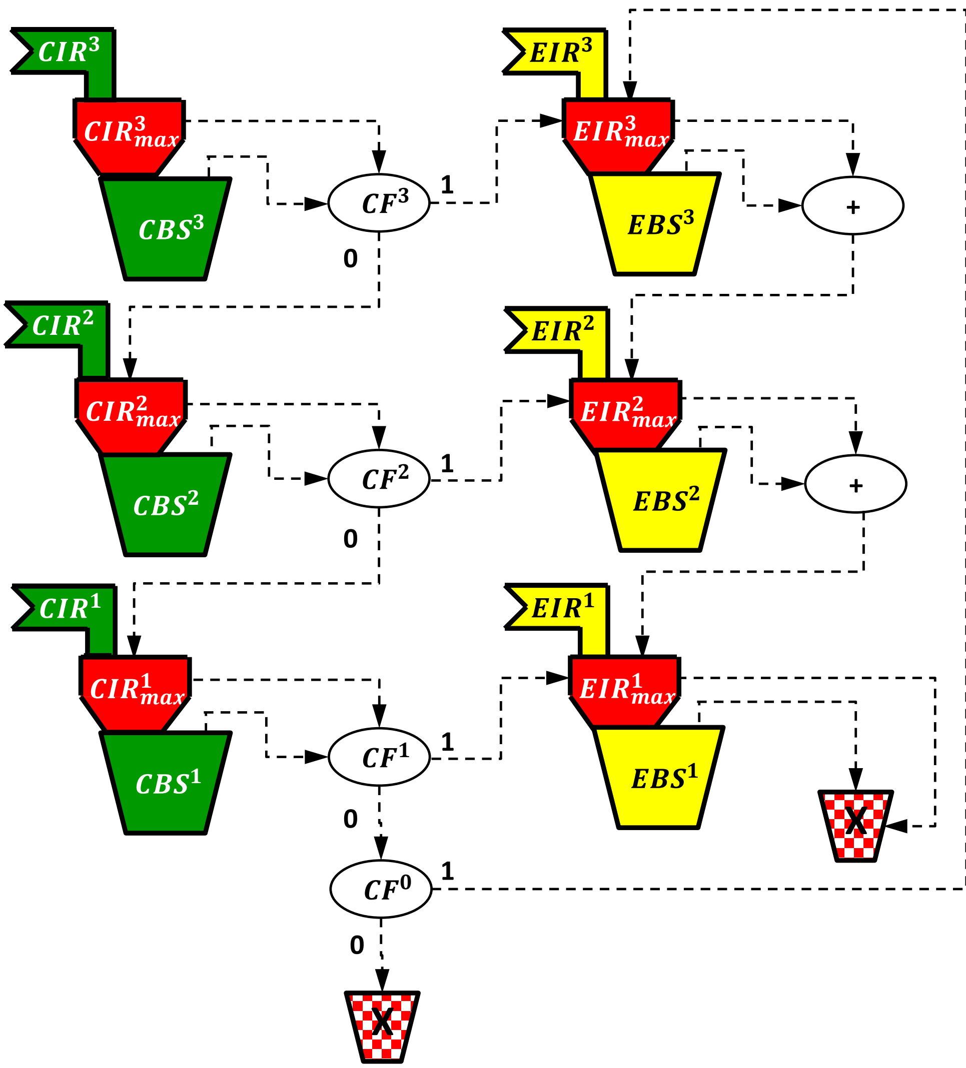

Figure 47 – Bandwidth Profile Token Flows

D.3 Bandwidth Profile Use Cases

This sub-section contains some Bandwidth Profile use cases. Additional Bandwidth Profile use cases can be found in MEF 23.2.1. [30]

D.3.1 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 $min(CIR^i, CIR^i_{max})$ 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 48 shows the token paths for this configuration.

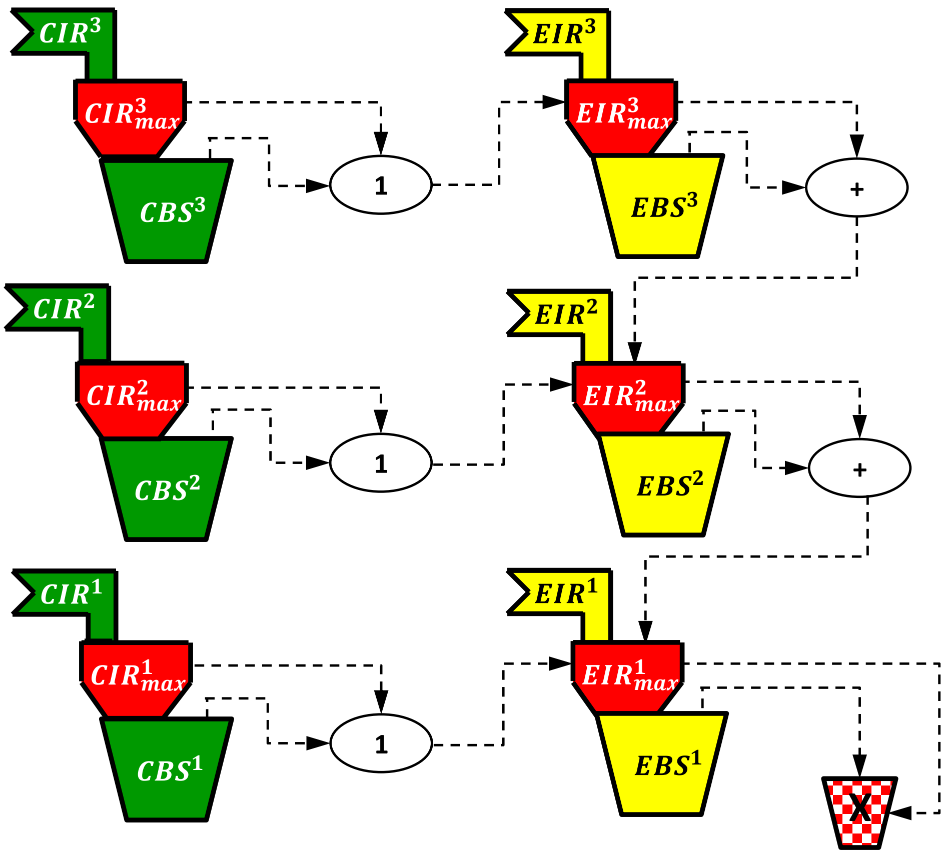

Figure 48 – Bandwidth Profile Algorithm for Sharing Excess Bandwidth

D.3.2 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. CIR_{env} and EIR_{env} are not Bandwidth Profile parameters defined in Section 12. 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. Green tokens are not converted to Yellow tokens. The configuration of this example is $CF^0 = 0, CF^i = 0, i = 1,2,3$, $CIR^1 = CIR^2 = EIR^1 = EIR^2 = 0$, $CIR^3 = CIR_{env}$, and $EIR^3 = EIR_{env}$. Figure 49 shows the token paths for this configuration.

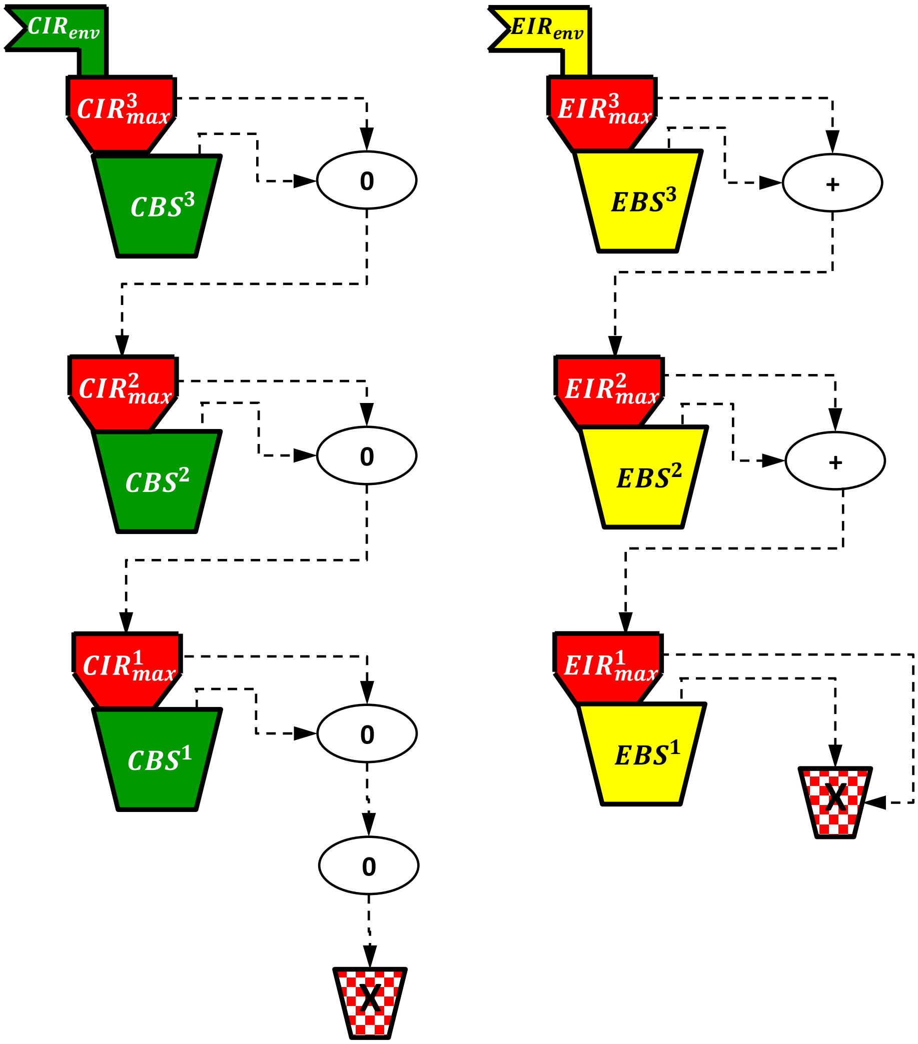

Figure 49 – Bandwidth Profile Algorithm for Uncoupled Bandwidth Sharing

The amount of bandwidth of Qualified Service Frames range from 0 to CIR_{env} for each Bandwidth Profile Flow. CIR_{max}^{3}, CIR_{max}^{2}, and CIR_{max}^{1} can be used to limit committed bandwidth for Bandwidth Profile Flows 3, 2, and 1. Similarly, EIR_{max}^{3}, EIR_{max}^{2}, and EIR_{max}^{1} can be used to limit excess bandwidth for Bandwidth Profile Flows 3, 2, and 1. 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.

D.3.3 Sharing Bandwidth among Class of Service Names

In this use case, CoS Labels H, M, and L share CIR_{env} bandwidth within an Envelope. CIR_{env} is not a Bandwidth Profile parameter defined in Section 12. CIR_{env} represents the number of bits per second of bandwidth shared by the three Bandwidth Profile Flows. 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 35 and $CF^0 = 0$.

Bandwidth Profile Flow 3 (H)	Bandwidth Profile Flow 2 (M)	Bandwidth Profile Flow 1 (L)
$CIR^3 = CIR_{env}$	$CIR^2 = 0$	$CIR^1 = 0$
$CIR^3_{max} \geq CIR_{env}$	$CIR^2_{max} \geq CIR_{env}$	$CIR^1_{max} = 0$
$CBS^3 \geq$ EVC Maximum Service Frame Size Service Attribute value	$CBS^2 \geq$ EVC Maximum Service Frame Size Service Attribute value	$CBS^1 = 0$
$EIR^3 = 0$	$EIR^2 = 0$	$EIR^1 = 0$
$EIR^3_{max} = 0$	$EIR^2_{max} \geq CIR_{env}$	$EIR^1_{max} \geq CIR_{env}$
$EBS^3 = 0$	$EBS^2 \geq$ EVC Maximum Service Frame Size Service Attribute value	$EBS^1 \geq$ EVC Maximum Service Frame Size Service Attribute value
$CF^3 = 0$	$CF^2 = 1$	$CF^1 = 0$
$ER^3 = \langle$Env-1, 3$\rangle$	$ER^2 = \langle$Env-1, 2$\rangle$	$ER^1 = \langle$Env-1, 1$\rangle$
$F^3 = 0$	$F^2 = 0$	$F^1 = 0$

Table 35 – Parameter Values for Sharing Bandwidth among Class of Service Names

Figure 50 shows the token paths for this configuration.

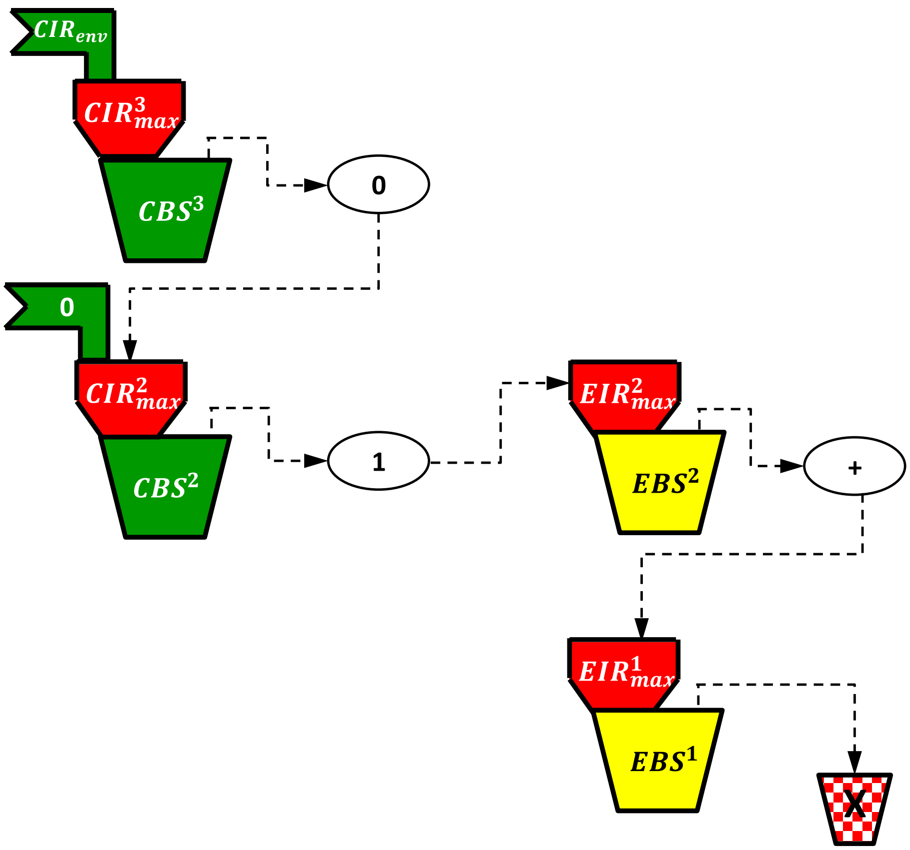

Figure 50 – Bandwidth Profile Algorithm for Bandwidth Sharing among Class of Service Names

D.4 Examples of the Use of the Token Request Offset Parameter

This appendix presents two examples of the use of the Token Request Offset parameter.

D.4.1 Adjusting for C-Tag Removal

Figure 51 shows an example of using the Token Offset Parameter, F^1, to adjust for the case where each ingress C-Tagged Service Frame results in an egress Untagged Service Frame. At UNI 1 the value of the EVC EP Map Service Attribute = *List* with 24 as the only entry (Section 10.4.1). Consequently, Ingress Service Frames mapped to the EVC EP need to be VLAN Tagged Service Frames with the C-Tag VLAN ID value = 24. At UNI 2 the value of the EVC EP Map Service Attribute = *UT/PT* (Section 10.4.3). Because of [R116], the corresponding Egress Service Frames at UNI 2 are Untagged Service Frames.

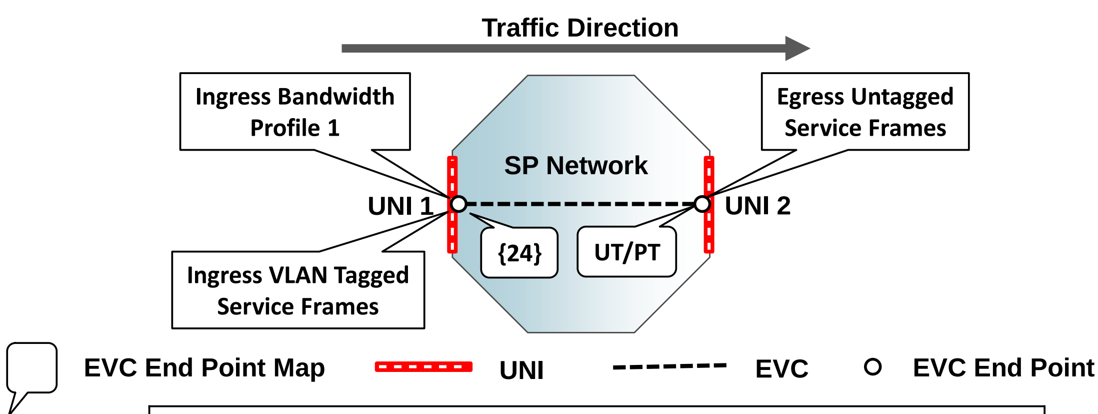

Ingress Bandwidth Profile 1 (Single Bandwidth Profile Flow per Envelope)	
Bandwidth Profile Flow Parameter	**Parameter Value**
CIR^1, CIR^1_{max}	50 Mbps
CBS^1	40,000 Bytes
EIR^1, EIR^1_{max}	0 Mbps
EBS^1	0 Bytes
CF^0, CF^1	0
CM^1	*color-blind*
F^1	4 Bytes

Figure 51 – Use of the Token Request Offset Parameter for C-Tagged to Untagged Service Frames

Now suppose that 1000 byte Service Frames ingress at UNI 1 whenever there are sufficient Green tokens available to declare the frame Green. If $F^1 = 0$, this results in 50 Mbps of Ingress Service Frames that are declared Green.[41] But the corresponding Egress Service Frames are only 996 bytes which means that there is only 49.8 Mbps of Egress Service Frames. If the Ingress Service Frames are 100 bytes, then there are only 48 Mbps of Egress Service Frames. There is

[41] For simplicity of discussion, the transient effects of initially draining the Green token bucket are ignored.

always bandwidth loss and the degree of loss depends on the size of the Ingress Service Frames at UNI 1.

By setting $F^1 = 4$, the C-Tag on the Ingress Service Frames at UNI 1 does not consume tokens. Consequently, the Ingress Service Frame pattern described in the previous paragraph results in 50 Mbps of Egress Service Frames at UNI 2 no matter what the length of the Ingress Service Frames at UNI 1.

D.4.2 Adjusting for Inter-frame Gap, Preamble, and Frame Delimiter

This example assumes that the value of Subscriber UNI Instantiation Service Attribute = *Physical* for the UNIs and the Service Frames are carried on a single link at each UNI.

Figure 52 shows an example of using the Token Offset Parameter, F^1, to adjust for Inter-frame Gap, Preamble and Frame Delimiter. In this case there is a C-Tag on both the Ingress Service Frames at UNI 1 and the Egress Service Frames at UNI 2. In addition, the MAC Data Rate for UNI 1 is 1 Gbps and the MAC Data Rate of UNI 2 is 100 Mbps.

Now suppose that the Subscriber is shaping traffic such that a Service Frame ingresses at UNI 1 whenever there are sufficient Green tokens available to declare the frame as Green. If $F^1 = 0$, this results in 100 Mbps of Ingress Service Frames that are declared Green.[42] At UNI 2, the Egress Service Frames are separated by 20 bytes consisting of the Inter-frame Gap, Preamble, and Frame Delimiter. Thus at UNI 2, Egress Service Frames cannot be generated fast enough to keep up with Ingress Service Frames at UNI 1. The smaller the length of the Service Frames the greater the imbalance because more blocks of 20 bytes are inserted at the egress UNI 2. Consequently, Service Frames will eventually be discarded by the SP Network. Which Service Frames are discarded will depend on the sequence of Ingress Service Frame lengths. (See Appendix C for details regarding these 20 bytes.)

By setting $F^1 = -20$, each Ingress Service Frame at UNI 1 consumes an extra 20 Green tokens which results in less than 100 Mbps of Green Ingress Service Frames in such a way that Egress Service Frames can be generated at UNI 2 that exactly keep up with the Ingress Service Frames at UNI 1 and no Service Frames will be discarded by the SP Network. Note that this is true for all mixes of Service Frame lengths.

[42] For simplicity of discussion, the transient effects of initially draining the Green token bucket are ignored.

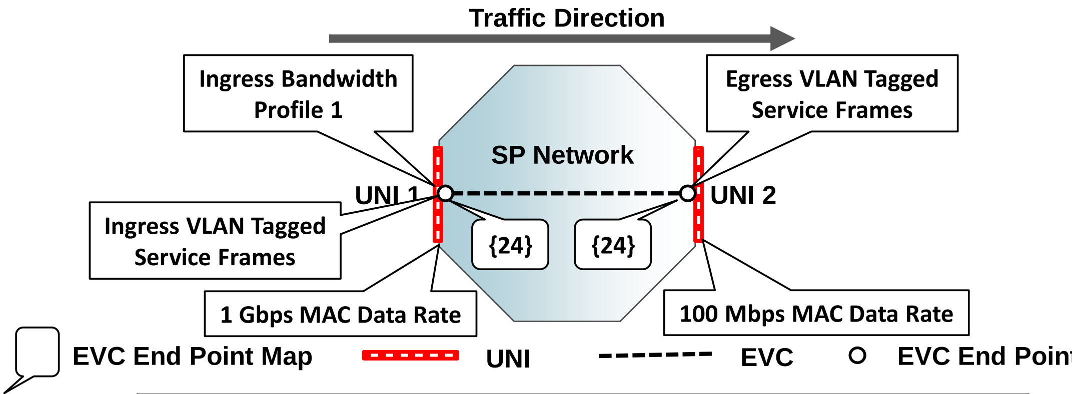

Ingress Bandwidth Profile 1 (Single Bandwidth Profile Flow per Envelope)	
Bandwidth Profile Flow Parameter	**Parameter Value**
CIR^1, CIR^1_{max}	100 Mbps
CBS^1	40,000 Bytes
EIR^1, EIR^1_{max}	0 Mbps
EBS^1	0 Bytes
CF^0, CF^1	0
CM^1	*color-blind*
F^1	-20 Bytes

Figure 52 – Use of the Token Request Offset Parameter to Adjust for Ethernet Frame Overhead

D.5 Eliminating Bias Toward Short Service Frames

In some circumstances the Bandwidth Profile algorithm specified in Section 12.2 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 token refresh rate[43] 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* > the value of the EVC Maximum Service Frame Size Service Attribute 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.

[43] The token refresh rate is bounded by CIR^i_{max} and/or EIR^i_{max}.

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. To do this, 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.

Figure 53 and Figure 54 show simulation results for the existing and the modified Bandwidth Profile algorithms. The simulation is of an Ingress Bandwidth Profile on a 100 Mbps UNI with $CIR = 10$ Mbps, $CIR_{max} = 10$ Mbps, $CBS = 1200$ Bytes, and $EIR = EBS = 0$. The value of the Subscriber UNI Instantiation Service Attribute = *Physical* and the Service Frames are carried on single link. The ingress traffic is a mix of 60% short (100-300 byte) and 40% long (1300-1500 byte) Service Frames arriving in a randomized sequence with minimum inter-frame gap.

Figure 53 is the unmodified Bandwidth Profile algorithm of Section 12.2. After the initial transient when the token bucket gets depleted, short frames are declared Green in a dramatically higher percentage than long frames.

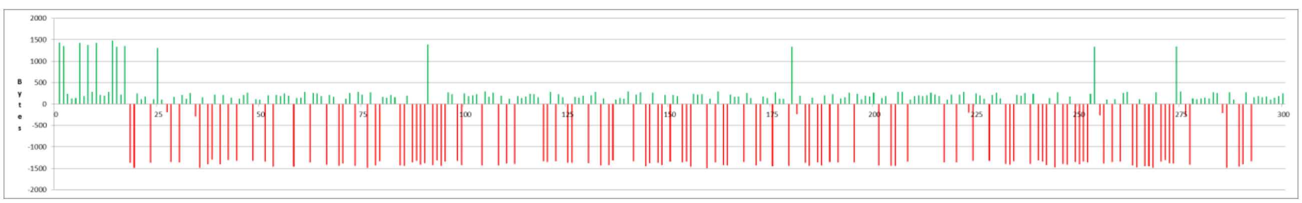

Figure 53 – Service Frame Profiling Results for the Algorithm of Section 12.2

Figure 54 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. 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 54 – Service Frame Profiling Results for Modified Algorithm

Table 36 shows the numbers of each type of Service Frame declared Green by the two 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	147	103	59%	41%
Algorithm of Section 12.2	142	4	97%	3%
Modified Algorithm	27	19	59%	41%

Table 36 – Service Frame Color Declaration Results for the Last 250 Service Frames

Appendix E Example of the Use of the One-way Composite Performance Metric (Informative)

This appendix contains a use case example for the One-way Composite Performance Metric specified in Section 8.8.10. The *PM* entry is for Mobile Backhaul with Packet based methods for synchronization, e.g., IEEE Std 1588 – 2008 [1].

Figure 55 shows a typical example of radio base station synchronization over an EVC. A Master clock is locked to the Primary Reference Clock (PRC) via GPS. This master clock sends timing packets with a configured rate (e.g., 128 packets/s) and packet size (e.g., 80 ~ 90 bytes) to the slave clock deployed at a RAN BS. The slave clock performs synchronization with the master clock. Then the slave clock feeds a clock signal to radio base stations to support air interface radio connections. Alternatively, the slave clock is implemented in radio base stations.

The EVC can have Performance Objectives for the packet synchronization traffic class using metrics such as the One-way Inter-Frame Delay Variation Performance Metric (Section 8.8.5), One-way Frame Delay Performance Metric (Section 8.8.2), and the One-way Frame Loss Ratio Performance Metric (Section 8.8.6).

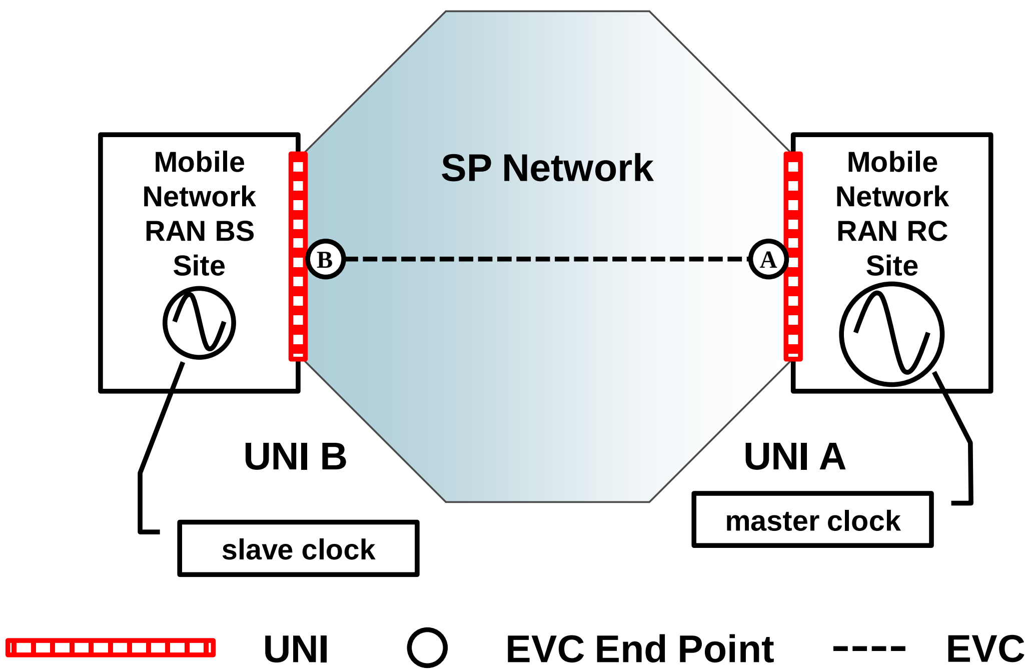

Figure 55 – Mobile Backhaul Base Station Synchronization Over an EVC

The packet synchronization method is vulnerable to network congestions and traffic variation. Therefore, the frequency accuracy, initial synchronization acquiring time, and re-locking time of the slave clock are very sensitive to EVC's delay variation and frame delay, but less sensitive to frame loss. Consecutive intervals of larger inter-frame delay variations in timing packet traffic class may result in the loss of slave clock synchronization with the master clock. Acquiring and re-locking times greatly depend on a sequence of smaller inter-frame delay variations. The slave clock may not be able to get re-locked again if inter-frame delay variations and frame delays become larger after the synchronization is lost. The inter-frame delay variation is actually time error/noise to the slave clock synchronization.

Performance requirements for the EVC in this use case are specified based on: (1) slave clock's Local Oscillator performance, (2) Air interface RF performance requirements, (3) EVC Performance Metrics. From the mobile network operation perspective, Mobile Operators would like to consider the impact of EVC's inter-frame delay variation and frame delay to determine the state of the EVC (i.e., Available or Unavailable) because the performance of radio station synchronizations can be predicted and assured. The One-way Availability Performance Metric specified in Section 8.8.7 is not adequate to address this example of Mobile Backhaul since it does not include inter-frame delay variation and frame delay performance.

The Mobile Operator and Service Provider can agree to include the One-way Composite Performance Metric in the SLS for Mobile Backhaul services for the Class of Service Name carrying the synchronization packets. Example values to use in the PM entry are shown in Table 37.

Item	Example Value
Performance Metric Name	One-way Composite Performance Metric
S	$\{\langle A, B\rangle, \langle B, A\rangle\}$
U	> 1/3
DL	< 8ms
Jt	< 2ms
Wfl	1
Wfd	1
$Wfdv$	1
$\widehat{AC}$	>99.9%

Table 37 – *PM* Entry for the One-way Composite Performance Metric for Packet Synchronization

Appendix F UNI Physical Link Configuration Examples (Informative)

A UNI can contain one or more physical links. When multiple physical links are configured at a UNI, the individual links may terminate at the same device in the SP Network and/or in the SN, or at different devices in the SP Network and/or in the SN. The devices can be located in the same or in different sites in the SP Network and/or in the SN.

The following figures illustrate some configuration examples where a device is shown as a rectangle. The details regarding what constitutes a SP Network device or SN device are beyond the scope of this document.

Figure 56 is configured with three physical links that all terminate at one device in the SP Network as well as in the SN.

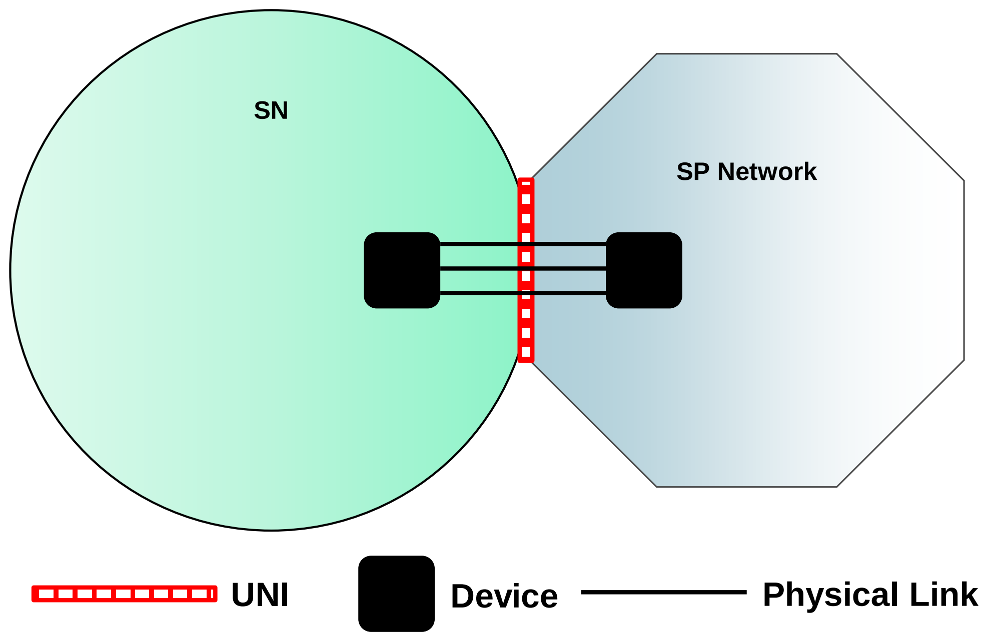

Figure 56 – Single Device in Both the SN and the SP Network

Figure 57 is configured with two links that terminate at different devices in the SP Network but at the same device in the SN.

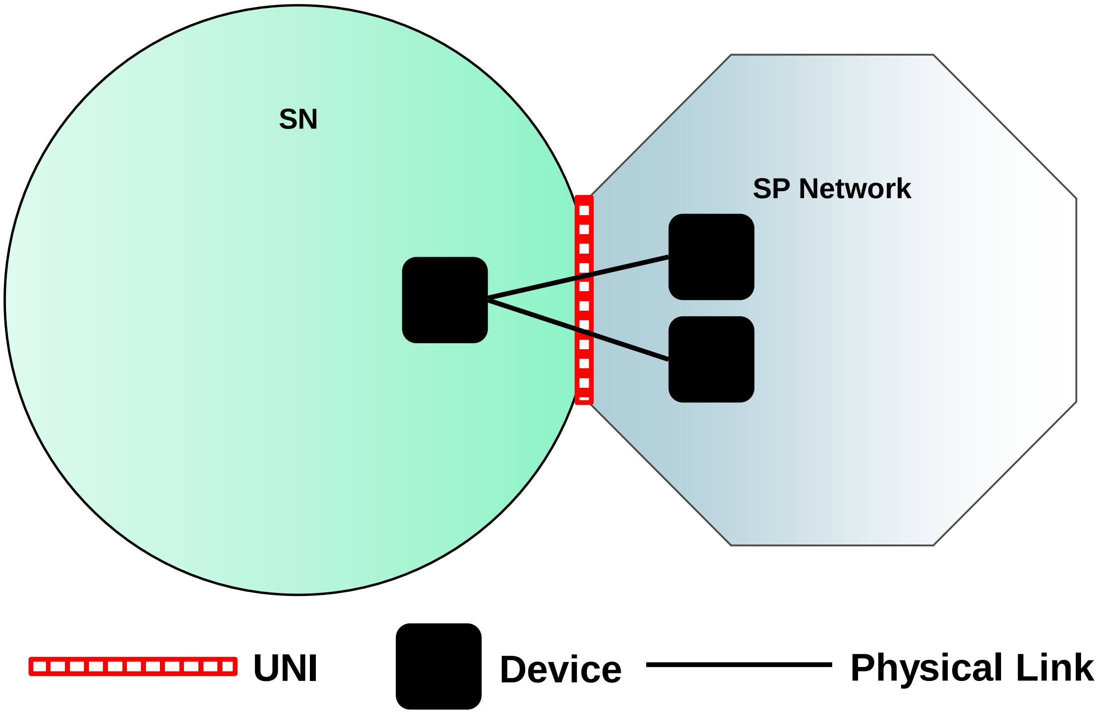

Figure 57 – Single Device in the SN and Two Devices in the SP Network

Figure 58 is configured with four physical links that terminate at two devices in the SP Network and two devices in the SN.

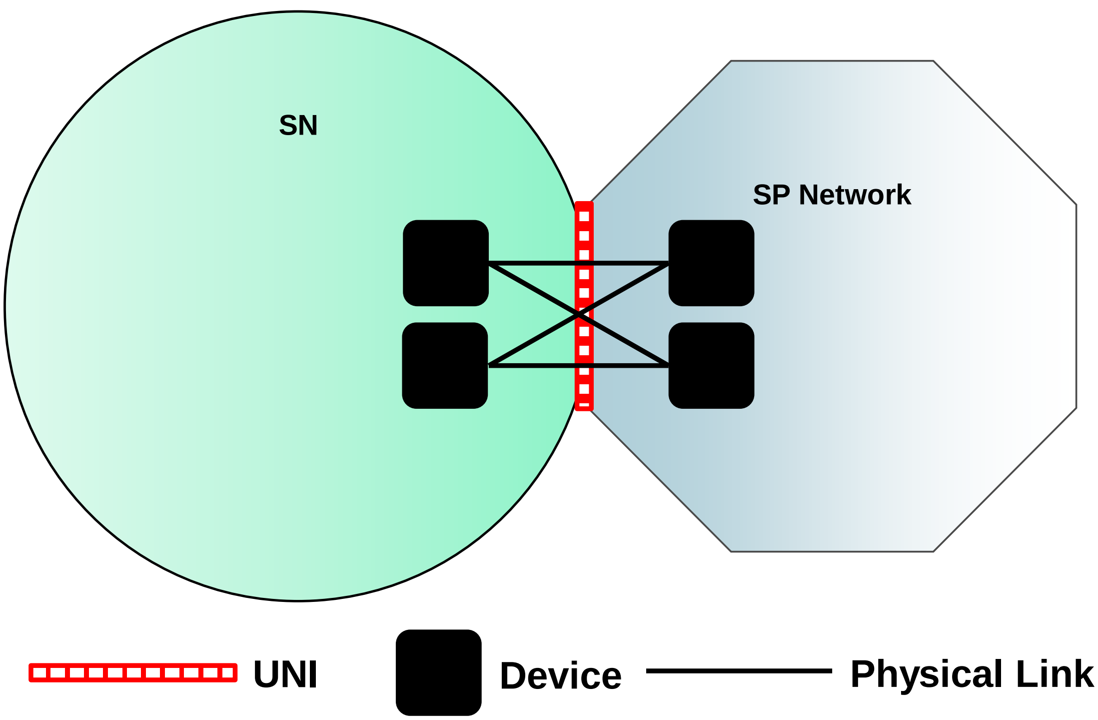

Figure 58 – Two Devices in the SN and Two Devices in the SP Network

Appendix G Subscriber UNI Port Conversation ID to Aggregation Link Map Examples (Informative)

Sections G.1 through G.6 contain examples where the Subscriber UNI Link Aggregation Service Attribute (Section 9.5) is set to *All Active*. A sequence of examples with increasing complexity is presented. In all examples, at UNI-A, the value of the Subscriber UNI List of Physical Links Service Attribute (Section 9.4) = a list of two 4-tuples. The value of Link Number IDs is 1 and 2 per [R79]. For all examples it is assumed that all Service Frames are VLAN Tagged Service Frames and, unless otherwise specified, the value of the EVC EP Map Service Attribute = *List* (Section 10.4.1).

G.1 Single EVC with Single Map Entry

This example has a single EVC at the UNI EVC-A, as shown in Figure 59.

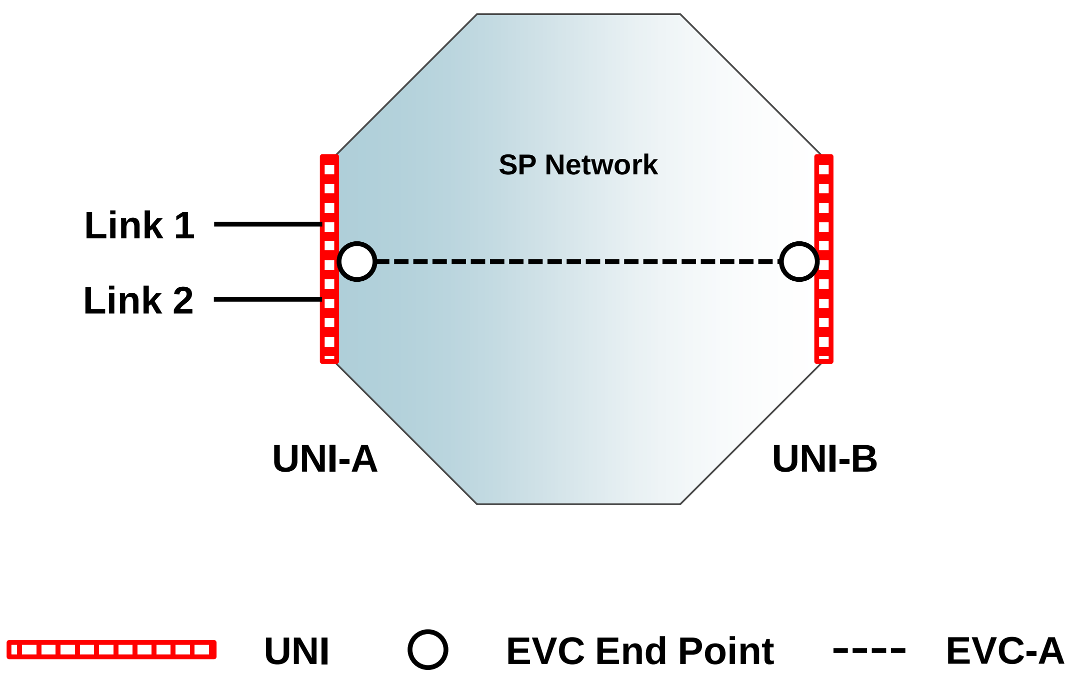

Figure 59 – Single EVC with a Single Port Conversation ID to Aggregation Link Map Entry

In this case, two C-Tag VLAN ID values, 13 and 14, are mapped to the EVC EP at UNI-A. At UNI-A both C-Tag VLAN ID values are mapped using the same row of the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute per Table 38.

Port Conversation ID	Link Selection Priority List
13, 14	1, 2
All other values	

Table 38 – Both C-Tag VLAN ID Values in One Row of the Port Conversation to Link Aggregation Map

In this case, at UNI-A, all Service Frames for EVC-A traverse link 1 as long as it is operational[44] and switch to link 2 when link 1 is not operational and link 2 is operational.

G.2 Single EVC with Multiple Map Entries

This example uses the same EVC as in Section G.1, but the two C-Tag VLAN ID values are mapped using different rows of the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute as shown in Table 39. This example illustrates [O3].

Port Conversation ID	Link Selection Priority List
13	1, 2
14	2, 1
All other values	

Table 39 – C-Tag VLAN IDs in Different Row of the Port Conversation to Link Aggregation Map

In this case, Service Frames with C-Tag VLAN ID value = 13 use link 1 and Service Frames with C-Tag VLAN ID value = 14 use link 2 as long as both links are operational. If either link is not operational, both C-Tag VLAN ID values are carried on the same link.

G.3 Single EVC with the EVC EP Map Service Attribute = *All*

This example has one EVC with the value of the EVC EP Map Service Attribute = *All* (Section 10.4.2). All the Service Frames at UNI-A can be mapped using the same row of the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute as shown in Table 40.

Port Conversation ID	Link Selection Priority List
All C-Tag VLAN ID values	1, 2

Table 40 – All C-Tag VLAN IDs in One Row of the Port Conversation to Link Aggregation Map

In this case, all Service Frames for the EVC traverse link 1 as long as it is operational and switch to link 2 when link 1 is not operational and link 2 is operational.

The Service Frames, at UNI-A, can also be mapped using different rows of the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute as shown in Table 41.

Port Conversation ID	Link Selection Priority List
0,10,13,15,100	1, 2
All other values	2, 1

Table 41 – C-Tag VLAN IDs Spread Across Two Rows of the Port Conversation to Link Aggregation Map

In this case, the Service Frames with C-Tag VLAN ID values = 0, 10, 13, 15, 100 and Untagged Service Frames traverse link 1 and other Service Frames traverse link 2 when both links are op-

[44] In the interest of brevity, the term "operational" is used as shorthand for MAC_Operational status = TRUE and the term "not operational" is used as shorthand for MAC_Operational status = FALSE. See Clause 6.3.12 in IEEE Std 802.1AX-2014 [3].

erational. If one link is not operational and the other link is operational, all Service Frames are carried on the same link.

The Service Frames, at UNI-A, can also be mapped using different rows in the way shown in Table 42.

Port Conversation ID	Link Selection Priority List
0,10,13,15,100	1, 2
All other values	1

Table 42 – Different C-Tag VLAN IDs Have Different Resiliency

In this case, the Service Frames with C-Tag VLAN ID values = 0, 10, 13, 15, 100 and Untagged Service Frames have a higher resilience level than other Service Frames. All other Service Frames are dropped if link 1 is not operational.

G.4 Two EVCs with Multiple Map Entries

This example has two EVCs that are multiplexed at UNI-A. C-Tag VLAN ID values 13 and 14 map to the EVC EP in EVC-A; C-Tag VLAN ID value 23 maps to the EVC EP in EVC-B. This configuration is illustrated in Figure 60.

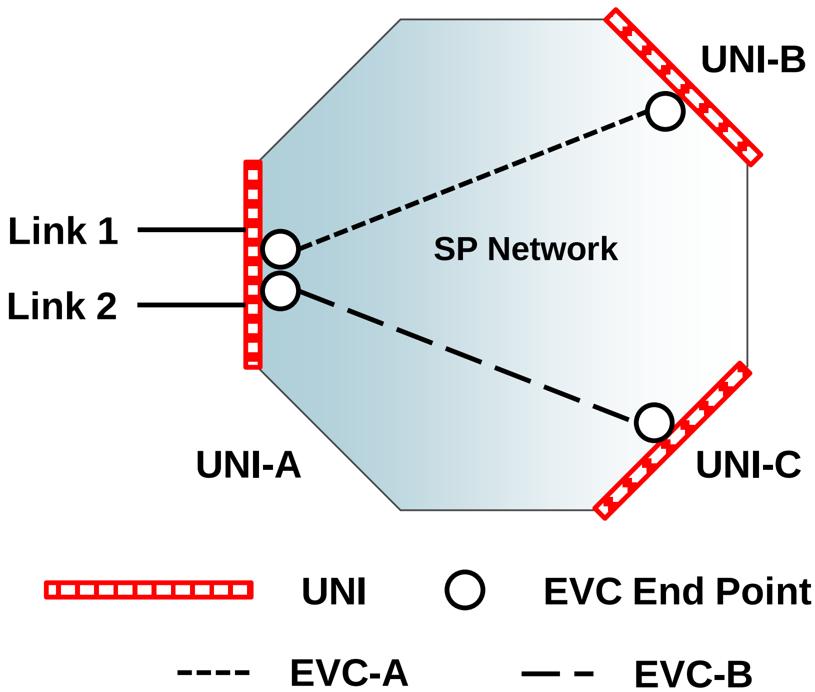

Figure 60 – Two EVCs Multiplexed at UNI-A

The value of the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute for UNI-A could be configured as shown in Table 43.

Port Conversation ID	Link Selection Priority List
13	1, 2
14, 23	2, 1
All other values	

Table 43 – Link Shared by Two EVCs

In this case, Service Frames with C-Tag VLAN ID value = 13 (EVC-A) use link 1 and Service Frames with C-Tag VLAN ID value = 14 (EVC-A) or C-Tag VLAN ID value = 23 (EVC-B) use link 2 as long as both are operational. If either link is not operational, the Service Frames with all three C-Tag VLAN ID values (both EVCs) are carried on the same link provided one link is operational.

G.5 Two EVCs with One Class of Service Name Bandwidth Profile Flow per Envelope

This example has two EVCs with the same C-Tag VLAN ID value to EVC EP mappings as in Section G.4 and has three Envelopes. The EVC-A has two Class of Service Labels, H and L which are differentiated by the C-Tag PCP field (or possibly the IP DSCP field). Service Frames with VLAN values 13 or 14 have PCP values mapping to each Class of Service Label. EVC-B only has Class of Service Label H. There is an EVC EP Class of Service Name Ingress Bandwidth Profile for each Class of Service Label on each EVC. Thus there are three Envelopes with one Bandwidth Profile Flow mapped to each. The configuration of the EVCs is shown in Table 44.

EVC	C-Tag VLAN ID Value	Class of Service Label
EVC-A	13, 14	H, L (based on PCP or DSCP)
EVC-B	23	H

Table 44 – EVC Configurations

If the Subscriber and Service Provider agree to configure the UNI such that all Service Frames that map to a given Envelope are mapped to the same link ([R81] mandates this to be supported), then both C-Tag VLAN ID values for EVC-A have to be carried on the same physical link. This is because EVC-A Service Frames with Class of Service Label H or L could be spread across both C-Tag VLAN ID values since the Class of Service Label is based on PCP value. EVC-B is in a separate Envelope so it can be placed on either of the links. Therefore a value of the Port Conversation to Aggregation Link Map Service Attribute could be configured as shown in Table 45.

Port Conversation ID	Link Selection Priority List
13, 14	1, 2
23	2, 1
All other values	

Table 45 – Each Single Bandwidth Flow Envelope Carried on a Single Link

G.6 Two EVCs with Multiple Class of Service Name Bandwidth Profile Flows per Envelope

This example has two EVCs with the same C-Tag VLAN ID value to EVC EP mappings as in Section G.4 but only has two Envelopes. The first Envelope has two Bandwidth Profile flows that are based on the Class of Service Label H from each of the EVCs, and the second Envelope has a single Bandwidth Profile Flow that is based on the Class of Service Label L. If the Subscriber and Service Provider agree to configure the UNI such that all Service Frames that map to a given Envelope are mapped to the same link ([R81] mandates this to be supported), then all EVC-A and EVC-B Service Frames have to traverse on the same link. Thus, in this case, the value of the Subscriber UNI Port Conversation ID to Aggregation Link Map Service Attribute is as shown in Table 46.

Port Conversation ID	Link Selection Priority List
13, 14, 23	1, 2
All other values	

Table 46 – Multiple Bandwidth Profile Flow Envelope on a Single Link

Appendix H Changes from MEF 10.3 (Informative)

H.1 Terminology Changes

Table 47 shows the terminology changes in this document from MEF 10.3.

MEF 10.3 Term	MEF 10.4 Term
Carrier Ethernet Network	Service Provider Network
CE	SN
CEN	SP Network
CE-VLAN CoS	C-Tag PCP
CE-VLAN CoS Preservation Service Attribute	EVC C-Tag PCP Preservation Service Attribute
Class of Service Identifier Service Attribute	EVC EP Ingress Class of Service Map Service Attribute
Color Identifier Service Attribute	EVC EP Color Map Service Attribute
Customer Edge	Subscriber Network
Data Service Frame Disposition Service Attribute	EVC Data Service Frame Disposition Service Attribute
Egress Bandwidth Profile per Egress Equivalence Class Identifier Service Attribute	EVC EP Class of Service Name Egress Bandwidth Profile Service Attribute
Egress Bandwidth Profile per EVC Service Attribute	EVC EP Egress Bandwidth Profile Service Attribute
EVC per UNI Service Attributes	EVC EP Service Attributes
Ingress Bandwidth Profile per Class of Service Identifier Service Attribute	EVC EP Class of Service Name Ingress Bandwidth Profile Service Attribute
Ingress Bandwidth Profile per EVC Service Attribute	EVC EP Ingress Bandwidth Profile Service Attribute
Maximum Number of EVCs Service Attribute	Subscriber UNI Maximum Number of EVC EPs Service Attribute
Physical Layer Service Attribute	Subscriber UNI List of Physical Links Service Attribute
Service Frame Format Service Attribute	Subscriber UNI Service Frame Format Service Attribute
Service Provider[45]	Ethernet Service Provider
Source MAC Address Limit Service Attribute	EVC EP Source MAC Address Limit Service Attribute
Subscriber[46]	Ethernet Subscriber
Subscriber MEG MIP Service Attribute	EVC EP Subscriber MEG MIP Service Attribute
UNI ID Service Attribute	Subscriber UNI ID Service Attribute
UNI Maximum Service Frame Size Service Attribute	Subscriber UNI Maximum Service Frame Size Service Attribute
UNI MEG Service Attribute	Subscriber UNI MEG Service Attribute
UNI Resiliency Service Attribute	Subscriber UNI Link Aggregation Service Attribute
User Network Interface[47]	Ethernet User Network Interface

Table 47 – Terminology Changes from MEF 10.3

[45] Service Provider is used as short for Ethernet Service Provider in most of the document.
[46] Subscriber is used as short for Ethernet Subscriber in most of the document.
[47] UNI is used as short for Ethernet User Network Interface in most of the document.

H.2 New Material

New material relative to MEF 10.3 includes:

- Addition of the EVC EP concept (Section 7.8).
- Addition of the EVC EP Egress Map Service Attribute (Section 10.7).
- Specification of a virtualized UNI through the Subscriber UNI Instantiation Service Attribute (Section 9.2), the Subscriber UNI Virtual Frame Map Service Attribute (Section 9.3), and various requirements on the values of other Service Attributes.
- Introduction of the Subscriber UNI List of Physical Links Service Attribute (Section 9.4) that specifies a unique identifier and the physical layer for each physical link as well as the frequency synchronization and time synchronization mode for each physical link.
- Addition of informative text making it clear that a frame at a UNI can map to a service other than an EVC at a UNI (Section 10.4).
- Addition of the Token Request Offset parameter to the Bandwidth Profile Algorithm (Section 12.1.2).
- Addition of the Subscriber UNI Token Share Service Attribute (Section 9.11).
- Addition of the Subscriber UNI Envelopes Service Attribute (Section 9.12).
- Addition of the Subscriber UNI L2CP Address Set Service Attribute (Section 9.15).
- Addition of the Subscriber UNI L2CP Peering Service Attribute (Section 9.17).
- Addition of the EVC C-Tag DEI Preservation Service Attribute (Section 8.6).
- Addition of the List of Class of Service Names Service Attribute (Section 8.7).
- Addition of the EVC Available MEG Level Service Attribute (Section 8.11).
- Addition of the Subscriber UNI Maximum Number of C-Tag VLAN IDs per EVC EP Service Attribute (Section 9.10).
- Addition of the Subscriber UNI Maximum Number of EVC EPs Service Attribute (Section 9.9).
- Addition of [O4] to allow discard of an Ingress Service Frame whose value of the C-Tag VLAN ID value = 4095.

- Addition of the content of MEF 10.3.1 [25] (Section 8.8.10 and Appendix E).
- Addition of the content of MEF 10.3.2 [26] (Section 9.5, 9.6, Appendix F, and Appendix G).
- Definition of Service Frame Arrival Time for both physical and virtual UNIs, and use of it for both Performance Metrics and for Bandwidth Profiles.

H.3 Other Changes

Changes relative to MEF 10.3 include:

- Elimination of the per-UNI Ingress Bandwidth Profile Service Attribute.
- Elimination of the per-UNI Egress Bandwidth Profile Service Attribute.
- Removal of the Egress Equivalence Class Service Attribute.
- Changing the Egress Bandwidth Profile algorithm to use the Class of Service Name and Color of the corresponding Ingress Service Frame.
- Simplification of the Egress Bandwidth Profiles by reducing to three the number of parameter values agreed to by the Subscriber and Service Provider.
- Replacement of the requirement that the egress traffic is mandated to always be limited per the Bandwidth Profile Parameters with informative text that the limit can be applied from time to time.
- Removal of the CE-VLAN ID concept (Section 7.8).
- Removal of the CE-VLAN ID for Untagged and Priority Tagged Service Frames Service Attribute.
- Removal of the CE-VLAN ID Preservation Service Attribute and replacing it with [R118].
- Replacement of the CE-VLAN ID/EVC Map Service Attribute with the EVC EP Map Service Attribute (Section 10.4) and this Service Attribute has a value of one of *List*, *All*, or *UT/PT*.
- Restructuring of the Service Level Specification Service Attribute to align with the approach in MEF 26.2 [31] (Section 8.8).
- Removal of the E-LMI Service Attribute.
- Revision of the One-way Multiple EVC Group Availability Performance Metric to simplify and correct a technical flaw (Section 11).

- Removal of the UNI EVC Identifier Service Attribute.
- Revision of the requirement to discard to a recommendation to discard when the Source MAC Address Limit is exceeded (Section 10.12).
- Alignment of the SN shaper example to align with the example in MEF 23.2 [29] (Appendix B).
- Modification of the Bandwidth Profile Algorithm specification to refer to MEF 41 [35] (Section 12.2).
- Removal of the Service Multiplexing Service Attribute and introduction of the term "Service Multiplexing" (Section 9.9).
- Removal of the Bundling Service Attribute and introduction of the term "Bundling" (Section 9.10).
- Restriction of the EVC EP Color Map Service Attribute so that the method of marking Color for Egress Service Frames is no longer part of the Service Attribute. (Section 10.6).
- Change of the notation in the SLS material to use functional notation (Section 8.8.1.2).
- Removal of the Test MEG Service Attribute.
- Removal of the All to One Bundling Service Attribute and introduction of the term "All to One Bundled UNI".
- Addition of the concept of a Behavioral Service Attribute (Section 7.3).
- Addition of the concept of a Capability Service Attribute (Section 7.3).
- Redefinition of several terms.