Technical Specification
MEF 26.1

External Network Network Interface (ENNI) – Phase 2

January 2012

Disclaimer

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

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

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

a) any express or implied license or right to or under any patent, copyright, trademark or trade secret rights held or claimed by any MEF member company which are or may be associated with the ideas, techniques, concepts or expressions contained herein; nor

b) any warranty or representation that any MEF member companies will announce any product(s) and/or service(s) related thereto, or if such announcements are made, that such announced product(s) and/or service(s) embody any or all of the ideas, technologies, or concepts contained herein; nor

c) any form of relationship between any MEF member companies and the recipient or user of this document.

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

Table of Contents

List of Figures

List of Tables

1 Abstract

The Metro Ethernet Network Architecture Framework specifies a reference point that is the interface between two Metro Ethernet Networks (MENs), where each Operator MEN is under the control of a distinct administrative authority. This reference point is termed the External Network Network Interface or ENNI.[1] The ENNI is intended to support the extension of Ethernet services across multiple Operator MENs. This Technical Specification specifies:

- The requirements at the ENNI reference point as well as the interface functionality in sufficient detail to ensure interoperability between two Operator MENs including Link OAM.

- The connectivity attributes UNI to UNI, UNI to ENNI, and ENNI to ENNI such that multiple Operator MENs can be interconnected and the Ethernet services and attributes in MEF 6.1 [9] and MEF 10.2 [5] can be realized.

2 Terminology

Term	Definition	Source
Bundled OVC	A non-Rooted Multipoint OVC that associates an OVC End Point that has more than one S-VLAN ID value that maps to it.	This document
CE	Customer Edge	MEF 10.2[5]
CHLI	Consecutive High Loss Interval	This document
Color Forwarding	An OVC attribute defining the relationship between the Color of an egress ENNI Frame and the Color of the corresponding ingress ENNI Frame or Service Frame	This document
Consecutive High Loss Interval	A sequence of small time intervals, each with a high frame loss ratio	This document
C-Tag	Subscriber VLAN Tag	IEEE Std 802.1ad[4]
DSCP	Diff-Serve Code Point	RFC 2474[14]
End Point Map	A mapping of specified S-Tag VLAN ID values to specified OVC End Point Identifiers	This document
End Point Map Bundling	When multiple S-VLAN ID values map to a single OVC End Point in the End Point Map, and the OVC associating that OVC End Point is not a Rooted-Multipoint OVC	This document
End Point Type	A parameter in the End Point Map (In this specification the End Point Type is always OVC End Point.)	This document
ENNI	A reference point representing the boundary between two Operator MENs that are operated as separate administrative domains	MEF 4[1]

[1] MEF 4 [1] hyphenates the acronym but this document does not.

Term	Definition	Source
ENNI Frame	The first bit of the Destination Address to the last bit of the Frame Check Sequence of the Ethernet Frame transmitted across the ENNI	This document
ENNI MTU	MTU of an ENNI frame at the ENNI	This document
ENNI-N_i	The functional element administered by the Operator of the ith Operator MEN that supports the protocols and procedures required in this document	This document
EVC	An association of two or more UNIs	MEF 10.2[5]
E-WAN	An MEF defined ETH services aware network that provides connectivity between two or more MENs via ENNIs	MEF 4[1]
External Interface	Either a UNI or an ENNI[2]	MEF 4[1]
High Loss Interval	A small time interval with a high frame loss ratio	This document
HLI	High Loss Interval	This document
L2CP Tunneling	The process by which a frame containing a Layer 2 Control Protocol is transferred between External Interfaces.	This document
Leaf OVC End Point	An OVC End Point that has the role of Leaf	This document
MEN	A Metro Ethernet Network comprising a single administrative domain	MEF 10.2[5]
MTU	Maximum Transmission Unit	This document
Multipoint-to-Multipoint OVC	An OVC that can associate two or more Root OVC End Points	This document
Network Operator	The Administrative Entity of a MEN	MEF 4[1]
Operator Virtual Connection	An association of OVC End Points	This document
OVC	Operator Virtual Connection	This document
OVC End Point	An association of an OVC with a specific External Interface i.e., UNI, ENNI	This document
OVC End Point Role	A property of an OVC End Point that determines the forwarding behavior between it and other OVC End Points that are associated with the OVC End Point by an OVC	This document

[2] MEF 4 considers several types of External Interfaces. This document is limited to consideration of the UNI and ENNI.

Term	Definition	Source
OVC Identifier	string that is unique among all OVCs in the Operator MEN	This document
Point-to-Point OVC	An OVC that associates exactly two Root OVC End Points	This document
Resiliency Performance	The number of High Loss Intervals and Consecutive High Loss Intervals in a time interval *T*	This document
Root OVC End Point	An OVC End Point that has the role of Root	This document
Rooted-Multipoint OVC	An OVC that can associate at least one Leaf or Trunk OVC End Point	This document
Service Frame	An Ethernet frame transmitted across the UNI toward the Service Provider or an Ethernet frame transmitted across the UNI toward the Subscriber	MEF 10.2[5]
Service Provider	The organization responsible for the UNI to UNI Ethernet service	MEF 10.2[5]
S-Tag	Service VLAN Tag.	IEEE Std 802.1ad[4]
Subscriber	The organization purchasing and/or using Ethernet Services	MEF 10.2[5]
S-VLAN ID	The 12 bit VLAN ID field in the S-Tag of an ENNI Frame	This document
Tag	An optional field in a frame header. In this document it is the 4-byte field that, when present in an Ethernet frame, appears immediately after the Source Address, or another tag in an Ethernet frame header and which consists of the 2-byte Tag Protocol Identification Field (TPID) which indicates S-Tag or C-Tag, and the 2-byte Tag Control Information field (TCI) which contains the 3-bit Priority Code Point, and the 12-bit VLAN ID field	IEEE Std 802.1ad[4]
Trunk OVC End Point	An OVC End Point that has the role of Trunk	This document
UNI	The physical demarcation point between the responsibility of the Service Provider and the responsibility of the Subscriber	MEF 10.2[5]

Table 1 – Acronyms and definitions

Note that throughout this specification, UNI means a demarcation point and ENNI means a demarcation point. Functionality associated with an interface at the ENNI is denoted by ENNI-N_i.

3 Scope

This document is a revision of MEF 26 [16]. MEF 26 was the first phase of specifications for interconnecting Operator MENs in order to support MEF Ethernet Services. MEF 26 includes:

- Support for Point-to-Point and Multipoint-to-Multipoint EVCs spanning an arbitrary number of Operator MENs and ENNIs.

- Ethernet frames at the ENNI with formats according to the Provider Bridges specification IEEE Std 802.1ad [4].
- Gigabit Ethernet or 10-Gigabit physical links according to IEEE Std 802.3 [3].
- Color aware Bandwidth Profiles at the ENNI.
- Hairpin switching, where ENNI Frames associated with an EVC may be sent back across an ENNI from which they were received by the Operator.
- Link protection based on IEEE Std 802.3-2005 Clause 43, Link Aggregation [3].
- Link OAM based on IEEE Std 802.3 [3].

This document represents the second phase. It consolidates MEF 26.0.1 [17], MEF 26.0.2 [18], and MEF 26.0.3 [19]. In addition it introduces specifications for the support of Rooted-Multipoint EVCs as defined in MEF 10.2 [5]. As such it adds the following to MEF 26:

- The definition and requirements for tunneling frames containing a Layer 2 Control Protocol on an Operator Virtual Connection.
- Service Level Specification definitions and related requirements.
- Support for Rooted Multipoint EVCs spanning an arbitrary number of Operator MENs.

This document supersedes MEF 26.

4 Key Concepts

4.1 Motivation and Service Model

It is likely that a potential Subscriber for Ethernet Services will have locations that are not all served by a single MEN Operator. Put another way, in order for such a Subscriber to obtain services, multiple MEN Operators will need to be used to support all of the Subscriber's User Network Interfaces (UNIs). A further potential complication is that the MEN Operators supporting the UNIs may not all interconnect with each other necessitating the use of transit MEN Operators. Figure 1 shows an example where there are four Subscriber UNIs supported by three MEN Operators where Operator A does not directly connect with Operator C necessitating the use of Operator D as an intermediary. The goal of this Technical Specification is to enable configurations like that of Figure 1 to support the service attributes defined in MEF 10.2 [5] and service definitions in MEF 6.1 [9].

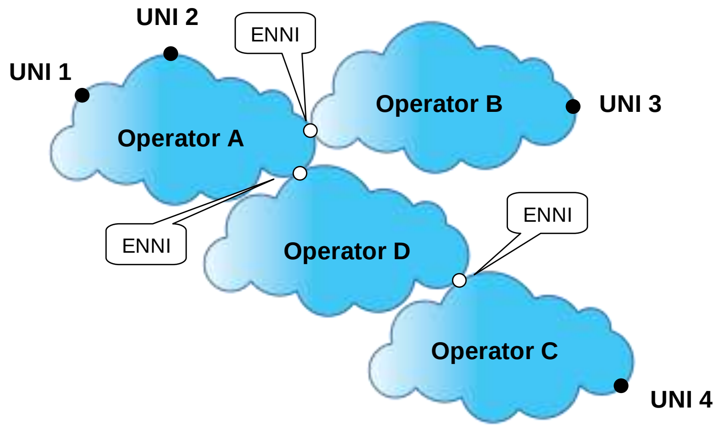

Figure 1 – Example of Multiple MEN Operators Supporting Ethernet Services

This document uses the following Service Model. For a given EVC, the Subscriber contracts with a Service Provider to be responsible for delivering Ethernet Services among the UNIs in the EVC. The Service Provider, in turn, selects and contracts with various MEN Operators to deliver the UNI-to-UNI services. It is the responsibility of the Service Provider to ensure that the appropriate service and interface attribute values from each Operator are such that the UNI to UNI service features purchased by the Subscriber can be delivered. There is no constraint on the type of organization that can act as the Service Provider. Examples include:

- One of the Operators involved in instantiating the Services, e.g., Operator D in Figure 1
- A third party such as a systems integrator
- An enterprise IT department (acting as both Service Provider and Subscriber)

Note that the role of an organization can be different for different EVC instances. For example, an organization can act as an Operator for one EVC and as Service Provider and an Operator for another EVC.

There are two types of technical requirements needed to support this Service Model and covered in this Technical Specification:

Interconnection Interface: These requirements detail the method of interconnection between two Operator MENs including the protocols that support the exchange of the information needed to support the UNI to UNI Ethernet Services. This interface is called the External Network Network Interface (ENNI). The Protocol Data Units exchanged at the ENNI are called ENNI Frames. In this Technical Specification these Protocol Data Units are Ethernet Frames as specified in IEEE Std 802.1ad [4].

Operator Services Attributes: These requirements detail the connectivity attributes that are supported by an Operator MEN. Such attributes can exist between UNIs as described in MEF

10.2 (Operator A in Figure 1), between ENNIs (Operators A and D in Figure 1), and between a UNI and an ENNI (Operators A, B, and C in Figure 1). These attributes can be thought of as the menu from which the Service Provider purchases, from each Operator, what is needed to instantiate the UNI to UNI services purchased by the Subscriber.

It is highly desirable that the UNI to UNI service observed by the Subscriber, when multiple Operator MENs are involved, be indistinguishable from the services that are obtained via a single Operator MEN. However, practical considerations might prevent this.

MEF 10.2 specifies multiple UNI to UNI service attributes, e.g., EVC type, service multiplexing. The Operator Service Attributes detailed in this document are intended to support most of these attributes (see Section 3). Furthermore it is desirable that a given instance of an ENNI simultaneously supports all of the Operator Service Attributes but this capability is not mandated.

4.2 Design Goals

This specification is driven by two main goals:

Rapid deployment: In order to meet Subscriber demand for interconnection of sites served by more than one Operator MEN, it is important that these requirements be amenable to rapid development and deployment. In practical terms, this means restricting the specification to existing Ethernet technology to the extent possible. It is desirable that the ENNI and Operator Service Attributes become available as quickly as possible to satisfy the overall market for Metro Ethernet Services, both on the demand and supply sides. Note that this does not preclude the subsequent development of extensions to this specification to support, for example, additional physical layers and protocol encapsulations, as well as automated provisioning and management functionality.

Minimization of global knowledge: The Service Provider has global knowledge of the sites associated with a particular service, what services have been subscribed to, etc. Coordination will have to occur with all of the MEN Operators involved. However, there is considerable motivation to limit the information that a particular Operator MEN requires to that needed to deploy services within that Operator MEN, and to information amenable to bilateral agreement at each ENNI. Partitioning the required information in this manner will greatly expedite the process of deploying services via multiple Operator MENs.

4.3 ENNI Reference Model

Formally, the Metro Ethernet Network Architecture Framework [1] specifies a reference point that is the interface between two Metro Ethernet Networks (MENs), where each MEN is under the control of a distinct administrative authority. This reference point is termed the External Network Network Interface or ENNI. The MEF external reference point model is displayed in Figure 2 which is derived from Figure 3 in [1]. An Ethernet-aware Wide Area Network (E-WAN) is functionally equivalent to a MEN but is given a distinct name to suggest that an E-WAN may have a larger geographical extent than a typical MEN. In this specification, MEN includes E-WAN.

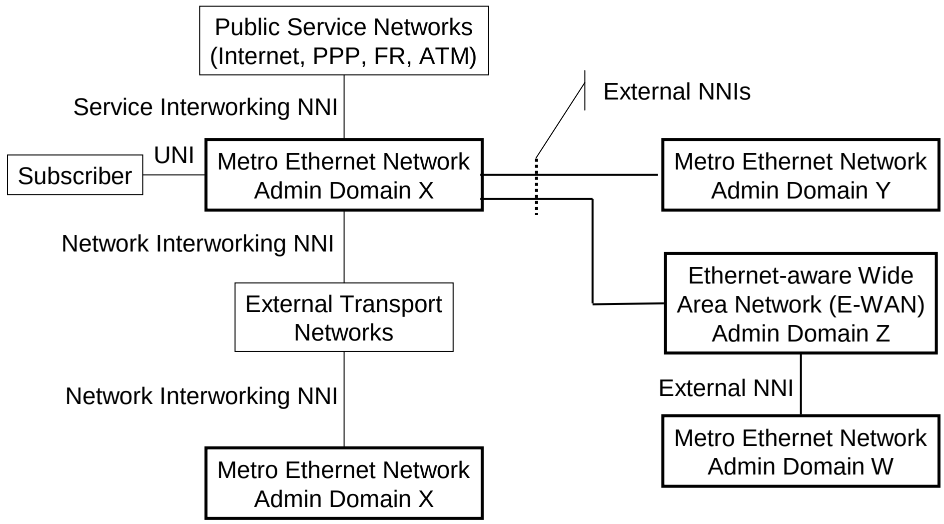

Figure 2 – MEF external reference model

4.4 ENNI Frame

Two Operator MENs exchange ENNI Frames across the ENNI. An ENNI Frame is an Ethernet [3] frame transmitted from an Operator MEN, across the ENNI toward the other Operator MEN. When an ENNI Frame is transmitted by an Operator MEN, from the perspective of this Operator MEN, it is called an egress ENNI Frame. When the ENNI Frame is received by an Operator MEN, from the perspective of this MEN, it is called an ingress ENNI Frame. The ENNI Frame consists of the first bit of the Destination MAC Address through the last bit of the Frame Check Sequence. The protocol, as seen by the network elements operating at the ENNI, complies to the standard Ethernet [3] frame with the exception that may have a length greater than that specified in [3] (see Section 7.1.6). There are no assumptions about the details of the Operator Metro Ethernet Network. It could consist of a single switch or an agglomeration of networks based on many different technologies.

4.5 Operator Virtual Connection

An Operator Virtual Connection (OVC) is the building block for constructing an EVC spanning multiple Operator MENs.

An OVC can informally be thought of as an association of “External Interfaces” within the same Operator MEN. This association constrains the delivery of frames in the sense that an egress Service Frame or ENNI Frame mapped to a given OVC can only be the result of an ingress Service Frame or ENNI Frame mapped to the given OVC.

In the case of an ENNI, an egress ENNI Frame with identical MAC and payload information can result from an ingress ENNI Frame at the same interface. (This behavior is not allowed at a UNI as specified in MEF 10.2 [5].) To describe this behavior, the OVC End Point is used which allows multiple, mutually exclusive ways that an ENNI Frame can be mapped to a single OVC at an ENNI. Section 7.2 defines the OVC as an association of OVC End Points. In turn each OVC End Point is associated with either a UNI or an ENNI. For the scope of this document, at least one OVC End Point associated by an OVC is at an ENNI.

Figure 3 shows an example of two OVCs. One OVC connects UNI A, UNI B, and the ENNI. The other OVC connects UNI A and the ENNI but allows what is called Hairpin Switching (see Section 7.2.3) via OVC Endpoints a and b at the ENNI.

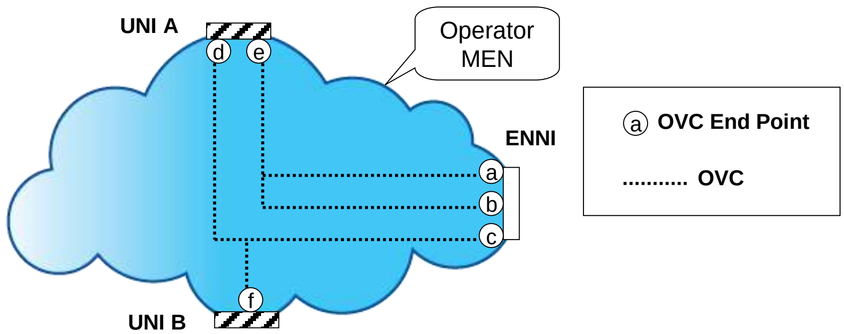

Figure 3 – Examples of OVCs

The formal definition and requirements that apply to OVCs are detailed in Section 7.2.

4.6 Relationship between OVCs and EVC

An Ethernet Virtual Connection (EVC) is an association of two or more UNIs. [5] When an EVC associates UNIs attached to more than one Operator MEN, the EVC is realized by concatenating OVCs. Figure 4 illustrates how this is done with an example.

In this example, there is an EVC associating UNI Q and UNI S and it is constructed by concatenating OVC A2 in MEN A with OVC B2 in MEN B. The concatenation is achieved by properly configuring the End Point Map attribute for ENNI AB in MEN A and the End Point Map attribute for ENNI AB in MEN B. (See Section 7.1.7 for the definitions and requirements for the End Point Map.) These map attributes are configured such that an ENNI Frame at ENNI AB that is mapped to OVC End Point x by MEN A is mapped to OVC End Point y by MEN B and vice versa. An ingress Service Frame at UNI Q that is destined for UNI S will result in an egress ENNI Frame at ENNI AB mapped to OVC End Point x in MEN A. When this frame is received by MEN B as an ingress ENNI Frame, it will be mapped to OVC End Point y and then result in an egress Service Frame at UNI S. The other EVCs in the example can be similarly instantiated by configuring the End Point Maps as shown in the table. It is the responsibility of the Service Provider responsible for an EVC to insure that the OVCs are correctly concatenated for the corresponding EVC.

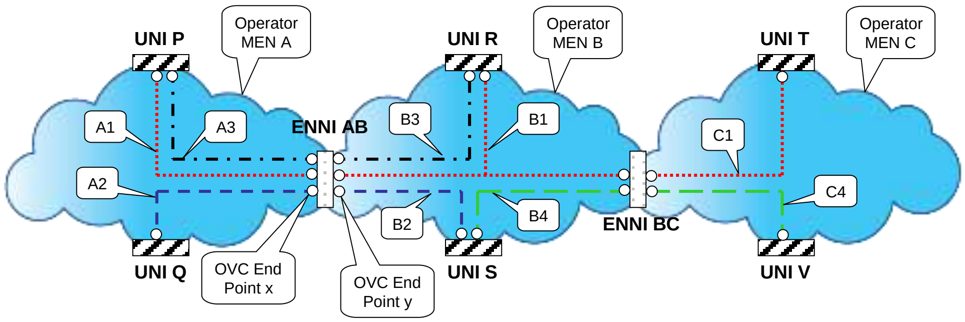

EVC	UNIs	OVCs
1 (red)	UNI P, UNI R, UNI T	A1, B1, C1
2 (blue)	UNI Q, UNI S	A2, B2
3 (black)	UNI P, UNI R	A3, B3
4 (green)	UNI S, UNI V	B4, C4

Figure 4 – Example Relationship of OVCs to EVCs

The example of Figure 4 illustrates one possible configuration. In this example, each Operator MEN has only one OVC for each EVC that it is supporting. However, it is possible to use multiple OVCs in a single Operator MEN to build an EVC. Section 10.4 contains an example. It is also possible that a single OVC in an Operator MEN can support more than one EVC. This can occur when an End Point Map attribute has Bundling as described in Section 7.1.7.

The Operator Service Attributes contained in this document are sufficient to support Point-to-Point, Multipoint-to-Multipoint, and Rooted-Multipoint EVCs.

5 Compliance Levels

The keywords **MUST**, **MUST NOT**, **REQUIRED**, **SHALL**, **SHALL NOT**, **SHOULD**, **SHOULD NOT**, **RECOMMENDED**, **MAY**, and **OPTIONAL**, when they appear in this document, are to be interpreted as described in RFC 2119 [2].

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

6 Requirements for the ENNI

The ENNI is defined as a reference point representing the boundary between two Operator MENs that are operated as separate administrative domains.

Similar to the concept of the UNI-C and UNI-N functional components of the UNI [7], it is useful to identify ENNI-N_1 and ENNI-N_2 as the separately administered functional components that support the ENNI between MEN 1 and MEN 2. Figure 5 illustrates this concept. Each ENNI-Ni is an entity that represents those functions necessary to implement the protocols and procedures specified in this document.

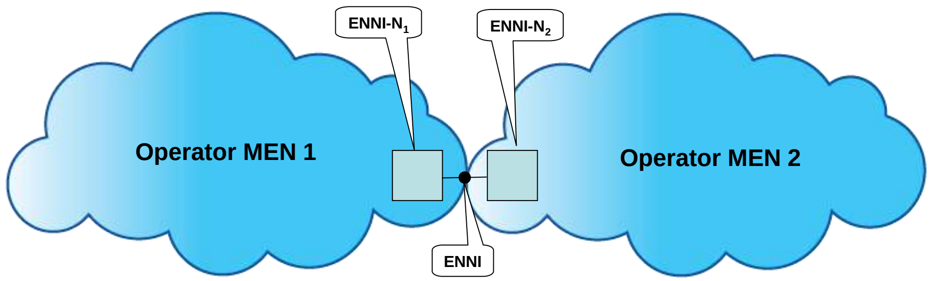

Figure 5 – Representation of ENNI-N_i

An ENNI can be implemented with one or more physical links. However, when there is no protection mechanism among multiple links connecting two Operator MENs, each link represents a distinct ENNI. The following requirements apply.

[R1] When there are two physical links in the ENNI, an ENNI-Ni **MUST** be capable of implementing Link Aggregation as in Clause 43.6.1 of [3] with one Link Aggregation Group (LAG) across the ports supporting an instance of ENNI and with one link in active mode and the other in standby mode.

Note that an Operator that is capable of supporting LAG as described in [R1] but elects to use an alternative protection mechanism at a given ENNI by mutual agreement with the connecting Operator is compliant with [R1].

[R2] When Link Aggregation is used at the ENNI, LACP **MUST** be used by each ENNI-Ni per [3].

The choice of which link to use for active or standby is to be decided on an Operator-to-Operator basis.

Note that the above requirements mean that if one link becomes inactive, Link Aggregation is to continue to operate on the remaining active link.

Requirements for a two-link LAG in active/active mode or for a LAG with more than two physical links in the ENNI may be addressed in a later phase of this specification.

7 Operator Services Attributes

The Service Model for the use of the ENNI involves the purchase of services from one or more Operators. These services are the exchange of traffic among ENNIs and UNIs that are supported by each Operator MEN. The purchaser of these services from the Operators is referred to as the Service Provider. Therefore it is important that the attributes of these services be described precisely. The basic model is shown in Figure 6. Operator Service Attributes describe the possible behaviors seen by an observer (the Service Provider) external to the Operator MEN at and between the external interfaces (UNI and ENNI).

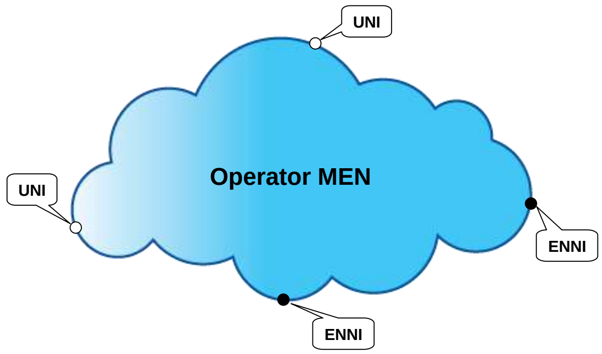

Figure 6 – ENNI Ethernet Services Model

The implementation of the Operator Metro Ethernet Network is opaque to the Service Provider. What is important is the observed behavior among the UNIs and ENNIs in Figure 6. These behaviors can be described by the following sets of attributes:

- ENNI Service Attributes are presented in Section 7.1.
- OVC Service Attributes are presented in Section 7.2.
- OVC End Point per ENNI Service Attributes are presented in Section 7.3.
- UNI Service Attributes are presented in Section 7.4.
- OVC per UNI Service Attributes are presented in Section 7.5.

In the following sections, the term “External Interface” is used to denote either an ENNI or a UNI.

7.1 ENNI Service Attributes

The ENNI is the point of demarcation between the responsibilities of two Operators. For each instance of an ENNI, there are two sets of ENNI Service Attributes, one for each Operator. A given attribute in the set can have an identical value for each Operator while another attribute can have a different value for each Operator. It is expected that many if not all of the ENNI Service Attribute values for each Operator will be known to the Service Provider. In some cases, some of the ENNI Service Attribute values of one Operator will not be known to the other Operator and vice versa.

The ENNI Service Attributes are summarized in Table 2 and described in detail in the following sub-sections.

Attribute Name	Summary Description	Possible Values
Operator ENNI Identifier	An identifier for the ENNI intended for management purposes	A string that is unique across the Operator MEN.
Physical Layer	The physical layer of the links supporting the ENNI	One of the PHYs listed in [R5]
Frame Format	The format of the PDUs at the ENNI	Frame formats as specified in Section 7.1.3
Number of Links	The number of physical links in the ENNI	An integer with value 1 or 2
Protection Mechanism	The method for protection, if any, against a failure	Link Aggregation, none, or other
ENNI Maximum Transmission Unit Size	The maximum length ENNI Frame in bytes allowed at the ENNI	An integer number of bytes greater than or equal to 1526
End Point Map	The map that associates each S-Tagged ENNI Frame with an End Point	A table with rows of the form <S-VLAN ID value, End Point Identifier, End Point Type>
Maximum Number of OVCs	The maximum number of OVCs that the Operator can support at the ENNI	An integer greater than or equal to 1
Maximum Number of OVC End Points per OVC	The maximum number of OVC End Points that the Operator can support at the ENNI for an OVC.	An integer greater than or equal to 1

Table 2 – ENNI Service Attributes

7.1.1 Operator ENNI Identifier

The Operator ENNI Identifier is a string administered by the Operator. It is intended for management and control purposes. An Operator can manage multiple ENNIs.

[R3] The Operator ENNI Identifier **MUST** be unique among all such identifiers for ENNIs supported by the Operator MEN.

[R4] The Operator ENNI Identifier **MUST** contain no more than 45 bytes.

7.1.2 Physical Layer

[R5] Each link in an ENNI **MUST** be one of the following physical layers in full duplex mode as defined in IEEE Std 802.3 – 2005[3]: 1000Base-SX, 1000Base-LX, 1000Base T, 10GBASE-SR, 10GBASE-LX4, 10GBASE-LR, 10GBASE-ER, 10GBASE-SW, 10GBASE-LW, 10GBASE-EW.

Note that the physical layer at one ENNI supported by the Operator MEN can be different than the physical layer at another ENNI supported by the Operator MEN.

7.1.3 Frame Format

The ENNI Frame is an Ethernet frame and is defined to consist of the first bit of the Destination MAC Address through the last bit of the Frame Check Sequence. ENNI Frames use Service VLAN tags (S-tags), as defined in IEEE Std 802.1ad-2005[4], to map frames to End Points as described in Section 7.1.7. An ENNI Frame can have zero or more VLAN tags. When there is a single tag, that tag is an S-Tag. When there are two tags, the outer tag is an S-Tag and the next tag is a C-Tag as defined in IEEE Std 802.1ad-2005[4].

[R6] Each ENNI Frame **MUST** have the standard Ethernet format with one of the tag configurations specified in Table 3. [DA = Destination Address, SA = Source Address, ET = Ethertype/Length, S-Tag with Tag Protocol Identification Field (TPID) = 0x88A8, C-Tag with TPID = 0x8100.]

DA (6 bytes) : SA (6 bytes) : ET (2 bytes): payload and FCS (no VLAN tags)
DA(6) : SA(6) : S-Tag (4) : ET (2) : payload and FCS
DA(6) : SA(6) : S-Tag (4) : C-Tag (4) : ET (2) : payload and FCS

Table 3 – ENNI Frame Formats

[R7] An S-Tag **MUST** have the format specified in Sections 9.5 and 9.7 of [IEEE Std 802.1ad]. [4]

[R8] A C-Tag **MUST** have the format specified in Sections 9.5 and 9.6 of [IEEE Std 802.1ad]. [4]

[R9] An ingress ENNI Frame that is invalid as defined in Clause 3.4 of [3] **MUST** be discarded by the receiving Operator MEN.

The length of an ENNI frame is defined as the number of bytes beginning with the first bit of the Destination Address through the last bit of the Frame Check Sequence.

[R10] An ingress ENNI Frame whose length is less than 64 octets **MUST** be discarded by the receiving Operator MEN as per Clause 4.2.4.2.2 of [3].

Note that this specification provides for ENNI Frames that are longer than the maximum specified in [3]. See Section 7.1.6.

When an ENNI Frame contains an S-Tag, the value of the 12 bit VID field in the S-Tag is defined as the S-VLAN ID Value.

7.1.4 Number of Links

An ENNI can be implemented with one or more physical links. This attribute specifies the number of such links. When there are two links, protection mechanisms are required, see Section 6. Protection mechanisms for more than two links are beyond the scope of this specification.

7.1.5 ENNI Resiliency Mechanisms

[R11] If the Number of Links is one, then the Protection Mechanism attribute **MUST** be set to "none."

[R12] If the Number of Links is 2 and LAG as specified in [R1] is implemented, then the Protection Mechanism attribute **MUST** be set to "Link Aggregation."

[R13] If the conditions specified [R11] and [R12] are not met, then the Protection Mechanism attribute **MUST** be set to "other."

As a consequence of these requirements, unprotected links between two Operator MENs represent distinct ENNIs.

Note that MEN Operators are allowed to decide whether to configure Link Aggregation by agreement between themselves.

The current scope of this document is restricted to the cases of the Gigabit Ethernet or the 10 Gigabit PHY layer, and therefore the discussions of protection and fault recovery mechanisms are directed at these PHYs. It is fully expected that later versions of this document will discuss other PHY layers, and they may contain their own protection and fault recovery mechanisms. Also, this phase of the specification addresses link or port protection only. The protection mechanism is Link Aggregation Protocol (802.3-2005). Similarly, later versions may specify other aspects of protection mechanisms from MEF 2. [13]

7.1.6 ENNI Maximum Transmission Unit Size

The ENNI Maximum Transmission Unit Size specifies the maximum length ENNI Frame in bytes allowed at the ENNI.

[R14] The ENNI Maximum Transmission Unit Size **MUST** be at least 1526 bytes.

[D1] The ENNI Maximum Transmission Unit Size **SHOULD** be at least 2000 bytes.

[R15] When an ENNI Frame is larger than the MTU Size, the receiving Operator MEN for this frame **MUST** discard it, and the operation of a Bandwidth Profile that applies to this is not defined.

The undefined operation of the Bandwidth Profile referred to in [R15] means that an ENNI Frame discarded because it is larger than the OVC MTU Size can result in either a change or no change in the state of the Bandwidth Profile algorithm.

The MTU is part of several attribute specifications. For example, UNI, ENNI, and OVC will have MTU attributes. Please refer to Section 9 for global MTU requirements.

7.1.7 End Point Map

The End Point Map specifies how each S-Tagged ENNI Frame is associated with an OVC End Point within an Operator MEN. (See Section 7.2.1 for the definition of OVC End Point.) As described in the following sub-sections, an End Point Map may or may not have what is called Bundling.

7.1.7.1 Basic End Point Map

The End Point Map can be represented by a three column table. Column 1 contains S-VLAN ID values. Column 2 contains End Point Identifiers. Column 3 contains End Point types. Each row in this table maps the S-VLAN ID value to the End Point Identifier and End Point Type.

Such a table is illustrated by the example in Figure 7. In this example, it is assumed that each End Point Identifier is formed by appending a four digit number to the ENNI Identifier (Gotham Central Exchange_12) and there are two types of End Point, OVC and Type X. Per this example, an S-Tagged ENNI Frame with S-VLAN ID value 189 is mapped to the OVC End Point, Gotham Central Exchange_12-1589. Note that the End Point Map applies to both ingress and egress S-Tagged ENNI Frames.

S-VLAN ID Value	End Point Identifier	End Point Type
158	Gotham Central Exchange_12-1224	OVC End Point
166	Gotham Central Exchange_12-1224	OVC End Point
189	Gotham Central Exchange_12-1589	OVC End Point
3502	Gotham Central Exchange_12-0587	Type X

Figure 7 – End Point Map Example

The following requirements apply to the End Point Map.

[R16] For a given S-Tagged ENNI Frame, the End Point to which it is mapped **MUST** be determined by the S-VLAN ID value in the S-Tag.

[R17] An S-VLAN ID value **MUST** be used in no more than one row of the map.

[R18] The End Point Type **MUST** be OVC End Point or VUNI End Point.[3]

As per the following requirement, an ingress S-Tagged ENNI Frame whose S-VLAN ID value is not in the map is not to be forwarded by the receiving Operator MEN.

[R19] An ingress S-Tagged ENNI Frame that is not mapped to an existing End Point **MUST NOT** be forwarded to an External Interface by the receiving Operator MEN.

[3] The definition of VUNI End Point and related requirements are in MEF 28 [20].

Note that [R19] does not preclude the receiving Operator MEN from acting as a peer for a L2CP protocol carried in an S-Tagged ENNI Frame. For example, S-Tagged ENNI Frames could be used for the ENNI Maintenance Entity Service OAM protocol.

In general the S-VLAN ID values in the End Point Map are local to an ENNI, e.g., OVC End Points associated by the same OVC at different ENNIs within an Operator MEN may be identified by different S-VLAN ID values. In cases where it is desirable to constrain each OVC End Point to be identified by the same S-VLAN ID value for a given OVC, the OVC is specified to have S-VLAN ID Preservation with value “yes” (see Section 7.2.13).

[R20] An ENNI Frame without an S-Tag **MUST NOT** be mapped to an OVC End Point.

[R20] does not necessarily mean that an untagged ENNI is to be discarded by the receiving Operator MEN. For example, an untagged ENNI Frame carrying a Layer 2 Control Protocol might be processed.

Note that at a given ENNI, each Operator MEN will have an End Point Map and these maps will typically have differences. Figure 8 shows an example for two Operator MENs, A and B. In this example, each Operator MEN is using a different scheme for identifying OVC End Points and thus the maps differ in column 2. More examples of End Point Maps are in Section 10.

Operator MEN A End Point Map

S-VLAN ID Value	End Point Identifier	End Point Type
158	Gotham Central Exchange_12-1224	OVC
189	Gotham Central Exchange_12-1589	OVC

Operator MEN B End Point Map

S-VLAN ID Value	End Point Identifier	End Point Type
158	Switch-103-Port-4-038	OVC
189	Switch-103-Port-4-344	OVC

Figure 8 – Example of the two End Point Maps for a Given ENNI

As described in Section 7.2.2, an OVC End Point always has one of three roles; Root, Leaf, or Trunk. When an OVC End Point at an ENNI has the role of Trunk, two S-VLAN ID values map to that OVC End Point in the End Point Map. One S-VLAN ID value identifies ENNI frames that result from Service Frames that originated at a Root UNI, and the other S-VLAN ID value identifies ENNI frames that result from Service Frames that originated at a Leaf UNI. (See Section 6.1.2.2 of MEF 10.2 [5] for descriptions of Root UNI and Leaf UNI.) The Trunk Identifiers attribute (see Section 7.3.2) specifies these S-VLAN ID values.

[R21] When the End Point Map contains an OVC End Point that has the OVC End Point Role of Trunk, the End Point MAP **MUST** contain exactly one Root S-VLAN ID value and one Leaf S-VLAN ID value that map to the OVC End Point.

Note that [R21] is only relevant at an ENNI.

7.1.7.2 End Point Map Bundling

When multiple S-VLAN ID values map to a single OVC End Point in the End Point Map, and the OVC associating that OVC End Point is not a Rooted-Multipoint OVC, the End Point Map is said to have Bundling and the OVC is said to be Bundled. (See Section 7.2.6 for the definition of Rooted-Multipoint OVC.) Note that these definitions are only relevant at an ENNI. When the End Point Map has Bundling, any OVC (that is not a Rooted-Multipoint OVC) that associates an OVC End Point for which the bundling applies has to have S-VLAN ID Preservation = Yes, cannot have Hairpin Switching (see Section 7.2.2), and can only associate OVC End Points that are at two ENNIs. The following requirements detail these properties.

[R22] A Bundled OVC **MUST** associate exactly two OVC End Points.

When there is Bundling, it is possible that frames originated by more than one Subscriber will be carried by the OVC and thus there may be duplicate MAC addresses being used by multiple Subscribers. To avoid the problems in the Operator MEN that can result from this duplication, MAC Address learning can be disabled on the OVC. However, disabling MAC Address learning can lead to poor efficiency when the OVC associates OVC End Points at more than two ENNIs. This is the motivation for [R22]. Future phases of this specification may relax [R22].

[R23] A Bundled OVC **MUST** have its S-VLAN ID Preservation attribute set to Yes. (See Section 7.2.13.)

Note that [R23] and [R47] mean that a Bundled OVC can associate at most one OVC End Point at an ENNI.

[R24] A Bundled OVC **MUST** have its CE-VLAN ID Preservation attribute set to Yes. (See Section 7.2.11.)

[R25] A Bundled OVC **MUST** have its CE-VLAN CoS Preservation attribute set to Yes. (See Section 7.2.12.)

[R26] Each OVC End Point associated by a Bundled OVC **MUST** be at an ENNI.

[R27] Each End Point Map at the ENNIs where there is an OVC End Point associated by a Bundled OVC **MUST** map the same list of S-VLAN ID values to the OVC End Point associated by the Bundled OVC.

As an example of [R27], consider the End Point Map shown in Figure 7. S-VLAN ID values 158 and 166 both map to the OVC End Point, “Gotham Central Exchange_12-1224.” If the OVC associating this OVC End Point also associates an OVC End Point at another ENNI, call it “Metropolis East Exchange_08-1328,” then the End Point Map at this other ENNI must map exactly 158 and 166 to “Metropolis East Exchange_08-1328.”

7.1.8 Maximum Number of OVCs

The Maximum Number of OVCs provides an upper bound on the number of OVCs that the Operator will support at the ENNI.

7.1.9 Maximum Number of OVC End Points per OVC

The Maximum Number of OVC End Points per OVC provides an upper bound on the number of OVC End Points that are associated by an OVC that the Operator can support at the ENNI.

Note that if the Maximum Number of OVC End Points per OVC is one, then hairpin switching cannot be supported at the ENNI. See Section 7.2.3.

7.2 OVC Service Attributes

7.2.1 OVC End Points

In the same way that an EVC defines an association of UNIs, an OVC is an association of OVC End Points. An OVC End Point represents the logical attachment of an OVC to an External Interface (a UNI or ENNI in the context of this document). At each External Interface, there must be a way to map each frame to at most one OVC End Point. Sections 7.1.7 and 7.4 describe the mapping method for an ENNI and a UNI respectively.

[R28] A given OVC **MUST** associate at most one OVC End Point at a given UNI.

[O1] A given OVC **MAY** associate more than one OVC End Point at a given ENNI.

[R29] If an egress frame mapped to an OVC End Point results from an ingress frame mapped to an OVC End Point, there **MUST** be an OVC that associates the two OVC End Points. And, the two OVC End Points **MUST** be different from each other.

[R29] means that, at a given ENNI, an ingress ENNI Frame mapped to a OVC End Point cannot result in an egress ENNI Frame at the given ENNI that is also mapped to that OVC End Point.

[R30] At least one of the OVC End Points associated by an OVC **MUST** be at an ENNI.

Note that if an OVC was allowed to associate End Points that were only at UNIs, then the OVC would not be distinguishable from an EVC as defined in MEF 10.2 [5].

7.2.2 OVC End Point Roles and Forwarding Constraints

An OVC End Point has one of three possible Roles; Root, Leaf, or Trunk.

[R31] The OVC End Point Role of an OVC End Point at a UNI **MUST** have the value either Root or Leaf.

[R32] The OVC End Point Role of an OVC End Point at an ENNI **MUST** have the value Root, Trunk, or Leaf.

Note that the OVC End Role will always have the value Root when the associating OVC is not of the type Rooted-Multipoint. See Section 7.2.6 for the definition of OVC Type.

For ease of exposition:

- An OVC End Point with the role of Root is called a Root OVC End Point,
- An OVC End Point with the role of Leaf is called a Leaf OVC End Point, and
- An OVC End Point with the role of Trunk is called a Trunk OVC End Point.

The following requirements constrain the forwarding behavior of an OVC based on the roles of the OVC End Points associated by the OVC. (See Section 7.3.2 for the definition of Root S-VLAN ID value and Leaf S-VLAN ID value.)

[R33] An egress frame at an EI that is mapped to a Root OVC End Point **MUST** be the result of an ingress frame at an EI that was mapped to a Root, Trunk, or Leaf OVC End Point.

[R34] An egress frame at an EI that is mapped to a Leaf OVC End Point **MUST** be the result of an ingress frame at an EI that was mapped to a Root OVC End Point or mapped to a Trunk OVC End Point via the Root S-VLAN ID value.

[R35] An egress frame at an EI that is mapped to a Trunk OVC End Point **MUST** contain the Root S-VLAN ID value when it is the result of an ingress frame at an EI that was mapped to a Root OVC End Point or mapped to a Trunk OVC End Point via the Root S-VLAN ID value.

[R36] An egress frame at an EI that is mapped to a Trunk OVC End Point **MUST** contain the Leaf S-VLAN ID value when it is the result of an ingress frame at an EI that was mapped to a Leaf OVC End Point or mapped to a Trunk OVC End Point via the Leaf S-VLAN ID value.

These forwarding requirements are summarized in Table 4.

		Ingress OVC End Point Role			
		Root	Leaf	Trunk (Leaf S-VID)	Trunk (Root S-VID)
Egress OVC End Point Role	Root	☑	☑	☑	☑
	Leaf	☑	☐	☐	☑
	Trunk (Leaf S-VID)	☐	☑	☑	☐
	Trunk (Root S-VID)	☑	☐	☐	☑

Table 4 – Allowed Connectivity Between OVC End Point Roles

By correct use of OVC End Point Roles and S-VLAN ID values, each Operator MEN can determine for each ingress frame that it receives whether the frame is the result of an ingress Service Frame at a Root UNI or a Leaf UNI. This information is necessary to implement a Rooted-Multipoint EVC that spans multiple Operator MENs. Section 11 contains examples of the support of Rooted-Multipoint EVCs.

When doing MAC address learning it is useful to do shared VLAN learning. This means that the source address of an ingress ENNI Frame mapped to a Trunk OVC End Point should be learned for both the Root S-VLAN ID value and the Leaf S-VLAN ID value. See Annex F of IEEE Std 802.1Q-2011 [21] for information on shared VLAN learning, and specifically F.1.3.2 for Rooted-Multipoint.

7.2.3 Hairpin Switching

Hairpin Switching occurs when an ingress S-Tagged ENNI Frame at a given ENNI results in an egress S-Tagged ENNI Frame with a different S-VLAN ID value at the given ENNI. This behavior is possible when an OVC associates two or more OVC End Points at a given ENNI. Sections 10.4 and 10.6 show examples of the use of Hairpin Switching.

Note that this configuration of OVC End Points is allowed by [O1]. Also note that [R28] precludes Hairpin Switching at a UNI.

Improper use of Hairpin Switching can result in a data loop between two Operator MENs at a single ENNI. Section 10.5 shows an example of how this can happen. It is up to the Service Provider and Operators to ensure that such loops do not occur. Automated methods for detecting and/or preventing such loops are beyond the scope of this document.

7.2.4 OVC Service Attributes

The OVC Service Attributes are summarized in Table 5 and described in detail in the following sub-sections.

Attribute Name	Summary Description	Possible Values

Attribute Name	Summary Description	Possible Values
OVC Identifier	An identifier for the OVC intended for management purposes	A string that is unique across the Operator MEN
OVC Type	An indication of the number and roles of the OVC End Points associated by the OVC.	Point-to-Point, Multipoint-to-Multipoint, or Rooted-Multipoint
OVC End Point List	A list of OVC End Points associated by the OVC and the OVC End Point Role of each OVC End Point associated by the OVC	A list of <OVC End Point Identifier, OVC End Point Role> pairs
Maximum Number of UNI OVC End Points	The bound on the number of OVC End Points at different UNIs that can be associated by the OVC	An integer greater than or equal to 0[4]
Maximum Number ENNI OVC End Points	The bound on the number of OVC End Points at ENNIs that can be associated by the OVC	An integer greater than or equal to 1[4]
OVC Maximum Transmission Unit Size	The maximum length in bytes allowed in a frame mapped to the an OVC End Point that is associated by the OVC	An integer number of bytes greater than or equal to 1526
CE-VLAN ID Preservation	A relationship between the format and certain field values of the frame at one External Interface and the format and certain field values of the corresponding frame at another External Interface that allows the CE-VLAN ID value of the UNI Service Frame to be derived from the ENNI Frame and vice versa	Yes or No
CE-VLAN CoS Preservation	A relationship between the format and certain field values of the frame at one External Interface and the format and certain field values of the corresponding frame at another External Interface that allows the CE-VLAN CoS value of the UNI Service Frame to be derived from the ENNI Frame and vice versa	Yes or No
S-VLAN ID Preservation	A relationship between the S-VLAN ID value of a frame at one ENNI and the S-VLAN ID value of the corresponding frame at another ENNI	Yes or No
S-VLAN CoS Preservation	A relationship between the S-VLAN PCP value of a frame at one ENNI and the S-VLAN PCP value of the corresponding frame at another ENNI	Yes or No
Color Forwarding	The relationship between the Color of an egress ENNI Frame and the Color of the corresponding ingress ENNI Frame or Service Frame	Yes or No
Service Level Specification	Frame delivery performance definitions and objectives for frames between External Interfaces	See Section 7.2.16

[4] Note that if the "Maximum Number of UNI OVC End Points" plus the "Maximum Number ENNI OVC End Points" is less than 2, then the OVC is not capable of conveying frames across the Operator MEN.

Attribute Name	Summary Description	Possible Values
Unicast Service Frame Delivery	This attribute describes how ingress frames mapped to an OVC End Point with a unicast destination MAC address are delivered to the other External Interfaces with OVC End Points associated by the OVC	Deliver Unconditionally or Deliver Conditionally
Multicast Service Frame Delivery	This attribute describes how ingress frames mapped to an OVC End Point with a multicast destination MAC address are delivered to the other External Interfaces with OVC End Points associated by the OVC	Deliver Unconditionally or Deliver Conditionally
Broadcast Service Frame Delivery	This attribute describes how ingress frames mapped to an OVC End Point with the broadcast destination MAC address are delivered to the other External Interfaces with OVC End Points associated by the OVC	Deliver Unconditionally or Deliver Conditionally
Layer 2 Control Protocol Tunneling	The Layer 2 Control Protocols that are tunneled by the OVC	A list on Layer 2 Control Protocols

Table 5 – OVC Service Attributes

7.2.5 OVC Identifier

The OVC Identifier is a string administered by the Operator that is used to identify an OVC within the Operator MEN. It is intended for management and control purposes. The OVC Identifier is not carried in any field in the ENNI Frame.

[R37] The OVC Identifier **MUST** be unique among all such identifiers for OVCs supported by the Operator MEN.

[R38] The OVC Identifier **MUST** contain no more than 45 bytes.

Note that the EVC Identifier described in MEF 10.2 [5] is known to the Subscriber and Service Provider. The degree to which the EVC Identifier is made known to the Operators is up to the Service Provider.

7.2.6 OVC Type

There are three types of OVC: Point-to-Point, Multipoint-to-Multipoint, and Rooted-Multipoint.

An OVC that can only associate two or more Root OVC End Points is defined to have OVC Type of Multipoint-to-Multipoint.

An OVC that associates exactly two Root OVC End Points is defined to have OVC Type of Point-to-Point, and can be considered a special case of the Multipoint-to-Multipoint OVC Type.

Note that a Multipoint-to-Multipoint OVC that associates two Root OVC End Points differs from a Point-to-Point OVC in that additional Root OVC End Points can be added to the OVC.

An OVC that can associate at least one Leaf or Trunk OVC End Point is defined to have OVC Type of Rooted-Multipoint. Note that an OVC that associates only Leaf OVC End Points is not useful since it cannot forward frames between External Interfaces. See Section 7.2.2.

The distinction between a Point-to-Point OVC or Multipoint-to-Multipoint OVC and a Rooted-Multipoint OVC with only Root OVC End Points is that a Leaf or Trunk OVC End Point can be added to such a Rooted-Multipoint OVC.

Note that because an OVC can associate more than one OVC End Point at a given ENNI, the type of an OVC is not determined by the number of External Interfaces supported by the OVC.

7.2.7 OVC End Point List

The OVC End Point List is a list of pairs of the form <OVC Point Identifier, OVC End Point Role>. Note that when the OVC type is not Rooted-Multipoint, the list can be simplified to a list of OVC End Point Identifiers since the OVC End Role is always Root. The list contains one item for each OVC End Point associated by the OVC. Section 7.3.1 describes OVC End Point Identifier. Section 7.2.2 describes the OVC End Point Role.

7.2.8 Maximum Number of UNI OVC End Points

The Maximum Number of UNI OVC End Points is the upper bound on the number of OVC End Points that are at different UNIs that can be associated by an OVC.

7.2.9 Maximum Number of ENNI OVC End Points

The Maximum Number of ENNI OVC End Points is the upper bound on the number of OVC End Points that are at an ENNI that can be associated by an OVC.

7.2.10 OVC Maximum Transmission Unit Size

The OVC Maximum Transmission Unit Size specifies the maximum length frame in bytes allowed on the OVC.

[R39] The OVC Maximum Transmission Unit Size **MUST** be at least 1526 bytes.

[D2] The OVC Maximum Transmission Unit Size SHOULD be at least 2000 bytes.

[R40] When an ENNI Frame or a Service Frame is larger than the OVC MTU Size for the OVC associating the OVC End Point to which it is mapped, the receiving Operator for this frame **MUST** discard it, and the operation of a Bandwidth Profile, if any, that applies to this frame is not defined.

The undefined operation of the Bandwidth Profile referred to in [R40] means that frame discarded because it is larger than the OVC MTU Size can result in either a change or no change in the state of the Bandwidth Profile algorithm.

[R41] The OVC Maximum Transmission Unit Size **MUST** be less than or equal to the MTU Size of each External Interface where an OVC End Point exists that is associated by the OVC.

The MTU is part of several attribute specifications. For example, UNI, ENNI, and OVC will have MTU attributes. Please refer to Section 9 for global MTU requirements.

7.2.11 CE-VLAN ID Preservation

CE-VLAN ID Preservation describes a relationship between the format and certain field values of the frame at one External Interface and the format and certain field values of the corresponding frame at another External Interface. This relationship allows the CE-VLAN ID value of the UNI Service Frame to be derived from the ENNI Frame and vice versa. OVC CE-VLAN ID Preservation is used to achieve EVC CE-VLAN ID Preservation that is a key property of the EPL and EP-LAN Service Types specified in MEF 6.1 [9]. See Sections 10.3 and 10.4 for examples of its use.

[R42] When an OVC has the CE-VLAN ID Preservation attribute with a value of Yes and all of the UNIs with an OVC End Point associated by the OVC are such that all CE-VLAN IDs map to the OVC End Point (see Section 7.5.2), then the relationship between the format of the frame at the ingress External Interface and the corresponding frame at the egress External Interface **MUST** be as specified in Table 6.

Ingress Interface	Ingress Frame Format	Egress Interface	Egress Frame Format
UNI	Untagged	UNI	Untagged
UNI	Untagged	ENNI	S-Tag only
UNI	C-Tagged	UNI	C-Tagged with VLAN ID value equal to that of ingress Service Frame at the UNI
UNI	C-Tagged	ENNI	S-Tag and C-Tag with the VLAN ID value in the C-Tag equal to the VLAN ID value in the C-Tag at the UNI
ENNI	S-Tag and C-Tag	UNI	C-Tagged with the VLAN ID value of the C-Tag equal to that of the VLAN ID value in the C-Tag of the ingress frame at the ENNI
ENNI	S-Tag only	UNI	Untagged
ENNI	S-Tag and C-Tag	ENNI	S-Tag and C-Tag with the VLAN ID value in the C-Tag equal to the VLAN ID value in the C-Tag at the ingress ENNI.
ENNI	S-Tag only	ENNI	S-Tag only

Table 6 – OVC CE-VLAN ID Preservation when All CE-VLAN IDs Map to the OVC at all of the UNIs Associated by the OVC

[R43] When an OVC has the CE-VLAN ID Preservation attribute with a value of Yes and not all of the UNIs with an OVC End Point associated by the OVC are such that all CE-VLAN IDs map to the OVC End Point (see Section 7.5.2), then the relationships between the format of the frame at the ingress External Interface and the corresponding frame at the egress External Interface **MUST** be as specified in Table 7.

Ingress Interface	Ingress Frame Format	Egress Interface	Egress Frame Format
UNI	C-Tagged with VLAN ID value in the range 1, … , 4094	UNI	C-Tagged with VLAN ID value equal to that of the ingress Service Frame at the UNI
UNI	C-Tagged with VLAN ID value in the range 1, … , 4094	ENNI	S-Tag and C-Tag with the VLAN ID value in the C-Tag equal to the VLAN ID value in the C-Tag at the UNI
ENNI	S-Tag and C-Tag with VLAN ID value in the C-Tag in the range 1, … , 4094	UNI	Tagged with the VLAN ID value of the C-Tag equal to that of the C-Tag of the ingress frame at the ENNI
ENNI	S-Tag and C-Tag with VLAN ID value in the C-Tag in the range 1, … , 4094	ENNI	S-Tag and C-Tag with the VLAN ID value in the C-Tag equal to the VLAN ID value in the C-Tag at the ingress ENNI.

Table 7 – OVC CE-VLAN ID Preservation when not All CE-VLAN IDs Map to the OVC at all of the UNIs Associated by the OVC

[R44] When an OVC has the CE-VLAN ID Preservation attribute with a value of Yes and none of the End Points associated by the OVC are at UNIs, then the relationships between the format of the frame at the ingress ENNI and the corresponding frame at the egress ENNI **MUST** be as specified in Table 8.

Ingress Frame Format	Egress Frame Format
S-Tag and C-Tag	S and C-Tag with the VLAN ID value in the C-Tag equal to the VLAN ID value in the C-Tag at the ingress ENNI.
S-Tag only	S-Tag only

Table 8 – OVC CE-VLAN ID Preservation when none of the OVC End Points are at UNIs

Note that when a Service Frame is delivered from a UNI in one Operator MEN to a UNI in another Operator MEN via an EVC supported by two or more OVCs with all OVCs having CE-VLAN ID Preservation attribute = Yes, then the Service Frame will have CE-VLAN ID Preservation as defined in Section 6.6.1 in MEF 10.2 [5]. Thus, the EVC that associates these two UNIs will have the CE-VLAN ID Preservation Service Attribute = Yes as defined in Section 6.6.1 in [5]. Also note that CE_VLAN ID Preservation as defined in Section 6.6.1 in MEF 10.2 [5] only applies to tagged Service Frames when there is not All to One Bundling at the UNI and thus Table 7 does not include the cases for untagged Service Frames at the UNI. See Table 2 and Table 3 in Section 6.6.1 of MEF 10.2 [5]. Section 10 shows some examples of the use of CE-VLAN ID Preservation in the construction of Ethernet Services.

7.2.12 CE-VLAN CoS Preservation

CE-VLAN CoS Preservation describes a relationship between the format and certain field values of the frame at one External Interface and the format and certain field values of the corresponding frame at another External Interface. This relationship allows the CE-VLAN CoS value of the UNI Service Frame to be derived from the ENNI Frame and vice versa. OVC CE-VLAN CoS Preservation is used to achieve EVC CE-VLAN CoS Preservation that is a key property of the EPL and EP-LAN Service Types specified in MEF 6.1 [9] See Sections 10.3 and 10.4 for examples of its use.

[R45] When an OVC has the CE-VLAN CoS Preservation attribute with a value of Yes the relationship between the format of the frame at the ingress External Interface and the corresponding frame at the egress External Interface **MUST** be as specified in Table 9.

Ingress Interface	Ingress Frame Format	Egress Interface	Egress Frame Format
UNI	C-Tagged	UNI	C-Tagged with PCP value equal to that of ingress Service Frame
UNI	C-Tagged	ENNI	S-Tag and C-Tag with the PCP value in the C-Tag equal to the PCP value in the C-Tag at the UNI
ENNI	S-Tag and C-Tag	UNI	C-Tagged with the PCP value of the tag equal to that of the C-Tag of the ingress frame at the ENNI
ENNI	S-Tag and C-Tag	ENNI	S-Tag and C-Tag with the PCP value in the C-Tag equal to the PCP value in the C-Tag of the ingress frame at the ingress ENNI

Table 9 – OVC CE-VLAN CoS Preservation

Note that when a Service Frame is delivered from a UNI in one Operator MEN to a UNI in another Operator MEN via two or more OVCs with CE-VLAN CoS Preservation, then the EVC that associates these two UNIs will have the CE-VLAN CoS Preservation Service Attribute as defined in Section 6.6.2 in [5].

7.2.13 S-VLAN ID Preservation

S-VLAN ID Preservation describes a relationship between the S-VLAN ID value of a frame at one ENNI and the S-VLAN ID value of the corresponding frame at another ENNI supported by the Operator MEN where each ENNI has an OVC End Point that is associated by the OVC. The possible values of the S-VLAN ID Preservation attribute are Yes or No.

[R46] When an OVC has the S-VLAN ID Preservation attribute with a value of Yes, an egress ENNI Frame at an ENNI resulting from an ingress ENNI Frame at a different ENNI **MUST** have an S-VLAN ID value identical to the S-VLAN ID value of the ingress ENNI Frame.

[R47] When an OVC has the S-VLAN ID Preservation attribute with a value of Yes, it **MUST** associate at most one OVC End Point located at a given ENNI.

[R48] When an OVC has the S-VLAN ID Preservation attribute with a value of No, an egress ENNI Frame mapped to an OVC End Point resulting from an ingress ENNI Frame mapped to a different OVC End Point **MUST** have an S-VLAN ID value that has a one-to-one association with the S-VLAN ID of the ingress service frame.

Note that the S-VLAN ID Preservation attribute does not apply to frames exchanged between an ENNI and a UNI.

7.2.14 S-VLAN CoS Preservation

S-VLAN CoS Preservation describes a relationship between the S-VLAN PCP value of a frame at one ENNI and the S-VLAN ID of the corresponding frame at another ENNI supported by the Operator MEN where each ENNI has an OVC End Point that is associated by the OVC. The possible values of the S-VLAN CoS Preservation attribute are Yes or No.

[R49] When an OVC has the S-VLAN CoS Preservation attribute with a value of Yes, an egress ENNI Frame at an ENNI resulting from an ingress ENNI Frame at a different ENNI **MUST** have an S-VLAN PCP value identical to the S-VLAN PCP value of the ingress ENNI Frame.

Note that when the S-VLAN PCP is used to indicate the color for an ENNI Frame, see [R85] and [R86], it could be undesirable to have S-VLAN CoS Preservation = Yes because an ingress ENNI Frame marked Green could never be marked Yellow on egress. An attribute that preserves Class of Service indication between ENNIs while allowing for changing the color marking using the S-Tag PCP may be addressed in a future phase of this document.

7.2.15 Color Forwarding

Color Forwarding describes the relationship between the color on an ingress frame into the Operator Network and the color of the resulting egress ENNI Frame. When Color Forwarding is Yes, the OVC cannot “promote” a frame from Yellow to Green. Promoting a frame from Yellow to Green could have an undesired impact on the EVC performance. The newly promoted Green frames are now competing with equal rights for resources as frames marked Green at the ingress UNI. For this reason, this attribute is useful to prevent this behavior.

[R50] When the Color Forwarding attribute is Yes for an OVC, each egress ENNI Frame mapped to an OVC End Point that is associated by the OVC **MUST** be marked Yellow using one of the formats specified in Section 7.3.3 if the corresponding ingress frame into the Operator MEN satisfied one or more of the following:

- The corresponding ingress frame was a Service Frame that was declared Yellow by an Ingress Bandwidth Profile and the Service Frame was mapped to an OVC End Point at the UNI that is associated by the OVC,

- The corresponding ingress frame was a Service Frame with a frame header indicating Yellow as specified in MEF 23.1 [10] and the Service Frame was mapped to an OVC End Point at the UNI that is associated by the OVC,

- The corresponding ingress frame was an ENNI Frame that was declared Yellow by an Ingress Bandwidth Profile and the ENNI Frame was mapped to an OVC End Point at the ENNI that is associated by the OVC,

- The corresponding ingress frame was an ENNI Frame with a frame header indicating Yellow using one of the formats specified in Section 7.3.3 and the ENNI Frame was mapped to an OVC End Point at the ENNI that is associated by the OVC.

[O2] When the Color Forwarding attribute is No, the Color marking of each egress ENNI Frame mapped to an OVC End Point that is associated by the OVC **MAY** be related to the Color of the corresponding ingress frame into the Operator Network in any way.

Note that this attribute does not describe Color marking of an egress Service Frame at a UNI because a method for such marking is not specified in MEF 23.1 [10].

7.2.16 Service Level Specification

The OVC Related Performance Service Attributes specify the frame delivery performance between External Interfaces (EI). Eight performance metrics are detailed in this specification:

1. One-way Frame Delay,

2. One-way Frame Delay Range,

3. One-way Mean Frame Delay,

4. Inter-Frame Delay Variation,

5. One-way Frame Loss Ratio,

6. One-way Availability,

7. One-way High Loss Intervals, and

8. One-way Consecutive High Loss Intervals.

These are similar to the performance attributes for an EVC described in MEF 10.2 [5] and MEF 10.2.1[22] but more general because both the UNI and ENNI are covered in this document.

These Performance Service Attributes describe the performance experienced by the Service Provider who is the user of the OVC. Methods for the Operator and the Service Provider to monitor the OVC performance to estimate this user experience are beyond the scope of this document.

An SLS can specify objectives for all or any subset of the OVC Performance Service Attributes.

[R51] If an SLS contains an objective for a given OVC Related Performance Service Attribute, then the SLS **MUST** specify the related parameters for that objective.

OVC Related Performance Service Attributes, with the exception of Availability[5], apply to "Qualified" Service Frames and "Qualified" ENNI Frames, called Qualified Frames.

[R52] An SLS **MUST** define Qualified Frames as follows for a given ordered pair of OVC End Points <i,j>, a given time interval T, and a given Class of Service Identifier:

1. Each frame **MUST** ingress at the EI where OVC End Point *i* is located.

2. Each frame **MUST** map to OVC End Point *i* via the End Point Map. (See Section 7.1.7 for the descriptions of the End Point Map at the ENNI and Section 7.5.2 for the description of the OVC End Point Map at the UNI.)

3. The first bit of each frame **MUST** arrive at the ingress EI within the time interval *T*, and within a time interval Δt smaller than *T* that has been designated as part of Available time (see Section 7.2.16.4),

4. Each frame **MUST** have the given Class of Service Identifier,

5. Each ingress frame that is subject to an ingress Bandwidth Profile **MUST** have an Ingress Bandwidth Profile compliance of Green, and

6. Each ingress frame that is not subject to an ingress Bandwidth Profile **MUST** meet one of the following two conditions:

    - The frame has no color identifier[6], or

    - The frame has a color identifier that indicates Green per the color indication requirements of MEF 23[10].

Such frames are referred to as **Qualified Frames**.

[5] Availability is used to define Qualified Frames via item 3 in the list.

[6] When there is no color identifier, this bullet means that the Service Frame is treated as if it were declared Green. An example is the use of the H Label as defined in MEF 23 [10] where a color identifier is not specified per Table 2 of [10] .

Recall that both OVC End Points in the ordered pair can be located at the same ENNI. See [O1] and Section 7.2.3.

The assessment of all performance attributes needs to account for unexpected arrival phenomena, such as frame duplication, or frames arriving in a different order from that observed on ingress, and the presence of these phenomena alone do not necessarily exclude a frame from the set of Qualified Frames. Details on how to monitor performance in the face of unexpected arrival phenomena are beyond the scope of this document.

7.2.16.1 One-way Frame Delay Performance for an OVC

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

The One-way Frame delay for a frame that ingresses at EI_1 and results in a frame that egresses at EI_2 is defined as the time elapsed from the reception of the first bit of the ingress frame at EI_1 until the transmission of the last bit of the first corresponding egress frame at EI_2. If the frame is erroneously duplicated in the Operator MEN and multiple copies delivered to EI_2, the delay is based on the first such copy delivered.

Note that this definition of One-way Frame Delay for a frame is the one-way[7] delay that includes the delays encountered as a result of transmission across the ingress and egress EIs as well as that introduced by the Operator MEN.

Note that when the path between two UNIs crosses one or more ENNIs, the UNI to UNI One-way Frame Delay, as defined in MEF 10.2 does not equal the sum of the One-way Frame Delay between each pair of EIs. This is because the sum will double count the time to transmit a frame across each ENNI. To see this, note that the definition of delay, D_O, between a UNI and an ENNI on single Operator MEN can be expressed as $D_O = d_U + d_E + d_M$ were d_U is the time to transmit the frame across the UNI, d_E is the time to transmit the frame across the ENNI, and d_M is the queuing and transmission delay introduced by the Operator MEN. Now consider the case where Operator MEN 1 and Operator MEN 2 are connected via an ENNI to affect an EVC between two UNIs, one on each Operator MEN. The delay between the UNIs is $d_{U1} + d_{M1} + d_E + d_{M2} + d_{U2}$. But

$$D_{O1} + D_{O2} = d_{U1} + d_{M1} + d_{U2} + d_{M2} + 2d_E \neq d_{U1} + d_{M1} + d_E + d_{M2} + d_{U2}.$$

Adding the two OVC delays overstates the UNI to UNI delay by d_E.

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

This effect will need to be taken into account when deriving the UNI to UNI delay performance from the delay performance of each Operator MEN in the path of the frame. This derivation is beyond the scope of this document.

[R53] For a given non-empty set of ordered pairs of OVC End Points, a time interval *T*, and a given Class of Service Identifier, the SLS **MUST** define each Frame Delay Performance metric as follows:

- Let the OVC End Points associated by the OVC be numbered from 1 to *m*. And let S be a non-empty subset of the ordered pairs of OVC End Points associated by the OVC. That is $S \subseteq \{\langle i, j\rangle \mid i = 1,...,m, j = 1,...,m, i \neq j\}, S \neq \varnothing$.

- Let $\overline{d}_{Td}^{\langle i,j\rangle}$ represent the P_d-Percentile of one-way delay for all Qualified Frames delivered to the egress EI where OVC End Point *j* is located resulting from an ingress frame at the EI where OVC End Point *i* is located. If there are no such egress frames, then $\overline{d}_{Td}^{\langle i,j\rangle} =$ Undefined.

- Then the One-way Frame Delay Performance metric **MUST** be defined as the maximum value of all of the defined values $\overline{d}_{Td}^{\langle i,j\rangle}$ for $\langle i, j\rangle \in S$, unless all $\overline{d}_{Td}^{\langle i,j\rangle}$ are Undefined in which case the performance is Undefined.

- Let $\overline{d}_{TR}^{\langle i,j\rangle} = \overline{d}_{Tr}^{\langle i,j\rangle} - d_{\min}^{\langle i,j\rangle}$. $\overline{d}_{Tr}^{\langle i,j\rangle}$ represents the P_r-Percentile of the one-way delay for all Qualified Frames delivered to the egress EI where OVC End Point *j* is located resulting from an ingress frame at the EI where OVC End Point *i* is located. $d_{\min}^{\langle i,j\rangle}$ is the minimum of the one-way delays for all Qualified Frames delivered to the EI where OVC End Point *j* is located resulting from an ingress frame at the EI where OVC End Point *i* is located. If there are no such egress frames, then $\overline{d}_{TR}^{\langle i,j\rangle} =$ Undefined.

- Then the One-way Frame Delay Range Performance metric **MUST** be defined as the maximum value of all of the defined values of $\overline{d}_{TR}^{\langle i,j\rangle}$ for $\langle i, j\rangle \in S$, unless all $\overline{d}_{TR}^{\langle i,j\rangle}$ are Undefined in which case the performance is Undefined.

- Let $\overline{\mu}_{T}^{\langle i,j\rangle}$ represent the arithmetic mean of one-way delay for all Qualified Frames delivered to the EI where OVC End Point *j* is located resulting from an ingress frame at the EI where OVC End Point *i* is located. If there are no such egress frames, then $\overline{\mu}_{T}^{\langle i,j\rangle} =$ Undefined.

- Then the One-way Mean Frame Delay Performance metric **MUST** be defined as the maximum value of all of the defined values $\overline{\mu}_T^{\langle i,j\rangle}$ for $\langle i, j\rangle \in S$, unless all $\overline{\mu}_T^{\langle i,j\rangle}$ are Undefined in which case the performance is Undefined.

To restate the Frame Delay definition mathematically, let the OVC End Points associated by the OVC be numbered from 1 to *m* and let $D_T^{\langle i,j\rangle}$ be the set of one-way Frame Delay values for all Qualified Frames at the EI where OVC End Point *j* is located resulting from an ingress frame at the EI where OVC End Point *i* is located. $D_T^{\langle i,j\rangle}$ can be expressed as $D_T^{\langle i,j\rangle} = \left\{d_1^{\langle i,j\rangle}, d_2^{\langle i,j\rangle}, \ldots, d_{N_{\langle i,j\rangle}}^{\langle i,j\rangle}\right\}$, where $d_k^{\langle i,j\rangle}$ is the one-way Frame Delay of the k^{th} frame. Define $\overline{d}_{Td}^{\langle i,j\rangle}$ for $P_d > 0$ as

$$\overline{d}_{Td}^{\langle i,j\rangle} = \begin{cases} \min\left\{d \mid P_d \le \dfrac{100}{N_{\langle i,j\rangle}} \displaystyle\sum_{k=1}^{N_{\langle i,j\rangle}} I\left(d, d_k^{\langle i,j\rangle}\right)\right\} \text{ if } N_{\langle i,j\rangle} \ge 1 \\ \textit{Undefined} \text{ otherwise} \end{cases}$$

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

$$I(d, d_k) \equiv \begin{cases} 1 \text{ if } d \ge d_k \\ 0 \text{ otherwise} \end{cases}.$$

$\overline{d}_{Td}^{\langle i,j\rangle}$ is the minimal delay during the time internal *T* that P_d percent of the frames do not exceed.

Then a One-way Frame Delay Performance metric for an OVC can be expressed as

$$\overline{d}_{T,S} = \begin{cases} \max\left\{\overline{d}_{Td}^{\langle i,j\rangle} \mid \langle i, j\rangle \in S \text{ and where } N_{\langle i,j\rangle} > 0\right\} \\ \textit{Undefined} \quad \text{when all } N_{\langle i,j\rangle} = 0 \text{ for all } \langle i, j\rangle \in S \end{cases}.$$

The One-way Frame Delay attribute permits specification of multiple values for P_d, (P_0, P_1, P_2, ...) and corresponding objectives ($\hat{d}_0^{\langle i,j\rangle}$, $\hat{d}_1^{\langle i,j\rangle}$, $\hat{d}_2^{\langle i,j\rangle}$, …).

Another parameter is the objective for the difference between the delay P_r percentile delay and $d_{\min}^{\langle i,j\rangle} = \min\left\{d \in D_T^{\langle i,j\rangle}\right\}$, expressed as

$$\overline{d}_{TR}^{\langle i,j\rangle} = \begin{cases} (\overline{d}_{Tr}^{\langle i,j\rangle} - d_{\min}^{\langle i,j\rangle}) \text{ if } N_{\langle i,j\rangle} > 0 \\ \textit{Undefined} \text{ if } N_{\langle i,j\rangle} = 0 \end{cases}$$

where

$$\bar{d}_{Tr}^{\langle i,j\rangle} = \begin{cases} \min\left\{ d \mid P_r \le \dfrac{100}{N_{\langle i,j\rangle}} \displaystyle\sum_{k=1}^{N_{\langle i,j\rangle}} I\left(d, d_k^{\langle i,j\rangle}\right) \right\} \text{ if } N_{\langle i,j\rangle} \ge 1 \\ \mathit{Undefined} \text{ otherwise} \end{cases}.$$

Then a One-way Frame Delay Range Performance metric for an OVC can be expressed as

$$\bar{d}_{TR,S} = \begin{cases} \max\left\{ \bar{d}_{TR}^{\langle i,j\rangle} \mid \langle i,j\rangle \in S \text{ and where } N_{\langle i,j\rangle} > 0 \right\} \\ \mathit{Undefined} \text{ when all } N_{\langle i,j\rangle} = 0 \text{ for all } \langle i,j\rangle \in S \end{cases}.$$

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

$$\bar{\mu}_T^{\langle i,j\rangle} = \begin{cases} \dfrac{1}{N_{\langle i,j\rangle}} \displaystyle\sum_{k=1}^{N_{\langle i,j\rangle}} \left(d_k^{\langle i,j\rangle}\right) \text{ if } N_{\langle i,j\rangle} > 0 \\ \mathit{Undefined} \text{ if } N_{\langle i,j\rangle} = 0 \end{cases}$$

Then a One-way Mean Frame Delay Performance metric for an OVC can be expressed as

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

The parameters of a One-way Frame Delay Performance metric are given in Table 10.

Parameter	Description
T	The time interval
S	Non-empty subset of the ordered pairs of OVC End Points associated by the OVC
COS ID	The Class of Service Identifier
P_d	A specific percentile for the Frame Delay Performance, $P_d > 0$
P_r	A specific percentile for the Frame Delay Range Performance, $P_r > 0$
$\hat{d}$	One-way Frame Delay Performance Objective corresponding to P_d
$\hat{d}_R$	One-way Frame Delay Range Objective corresponding to P_r
$\hat{\mu}$	One-way Mean Frame Delay Performance Objective

Table 10 – One-way Frame Delay Performance Parameters

[R54] Given *T*, *S*, *COS ID*, P_d, and a one-way Frame Delay Performance objective $\hat{d}$, expressed in time units, an SLS **MUST** define the one-way Frame Delay Performance objective as met over the time interval *T* for the subset *S* if and only if $\bar{d}_{T,S} \le \hat{d}$ or $\bar{d}_{T,S}$ is Undefined.

[R55] Given *T*, *S*, *COS ID*, P_r, and a One-way Frame Delay Range Performance objective $\hat{d}_R$, expressed in time units, an SLS **MUST** define the one-way Frame Delay Range Performance objective as met over the time interval *T* for the subset *S* if and only if $\bar{d}_{TR,S} \le \hat{d}_R$ or $\bar{d}_{TR,S}$ is Undefined.

[R56] Given *T*, *S*, *COS ID*, a One-way Mean Frame Delay Performance objective $\hat{\mu}$, expressed in time units, the Frame Delay Performance **MUST** be defined by an SLS as met over the time interval *T* if and only if $\bar{\mu}_{T,S} \le \hat{\mu}$ or $\bar{\mu}_{T,S}$ is Undefined.

Recall that if any of the above performance attributes are Undefined for time interval *T* and ordered pair $\langle i, j \rangle$, then the performance for that ordered pair is to be excluded from calculations on the performance of pairs in *S*.

[O3] For a Point-to-Point OVC, *S* **MAY** include one or both of the ordered pairs of OVC End Points associated by the OVC.

[O4] For a Multipoint-to-Multipoint OVC, *S* **MAY** be any non-empty subset of the ordered pairs of OVC End Points associated by the OVC.

[R57] For a Rooted-Multipoint OVC, *S* **MUST** be a non-empty subset of the ordered pairs of OVC End Points associated by the OVC such that each ordered pair in *S* contains at least one Root OVC End Point.

7.2.16.2 Inter-Frame Delay Variation Performance for an OVC

Inter-Frame Delay Variation (IFDV) is the difference between the one-way delays of a pair of selected frames. This definition is borrowed from RFC3393[8] where IP packet delay variation is defined.

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

$$\{a_j - a_i = \Delta\tau \quad and \quad j > i\}$$

Let r_i be the time frame *i* is successfully received (last bit of the frame) at the egress EI, then the difference in the delays encountered by frame *i* and frame *j* is given by $d_i - d_j$. Define

$$\Delta d_{ij} = \left|d_i - d_j\right| = \left|(r_i - a_i) - (r_j - a_j)\right| = \left|(a_j - a_i) - (r_j - r_i)\right|$$

[8] C. Demichelis and P. Chimento, *IP Packet Delay Variation Metric for IP Performance Metric (IPPM)*, RFC 3393, November 2002.

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

Figure 9 shows a depiction of the different times that are related to Inter-Frame Delay Variation Performance.

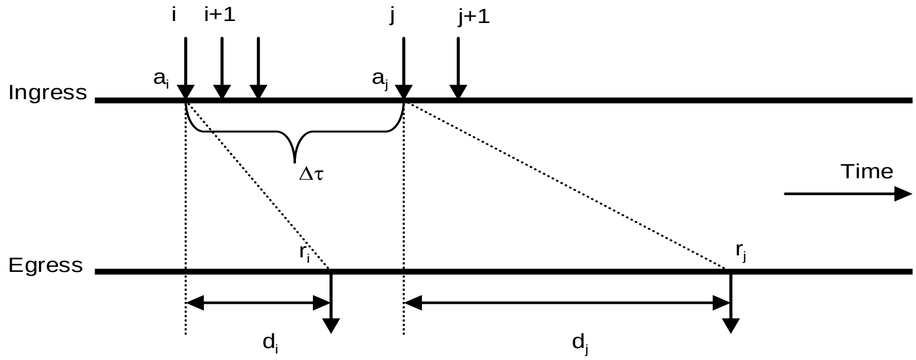

Figure 9 – Inter-Frame Delay Variation Definition

[R58] For a given non-empty set of ordered pairs of OVC End Points, a time interval *T*, and a given Class of Service Identifier, the SLS **MUST** define the Inter-Frame Delay Variation Performance metric as follows:

- Let the OVC End Points associated by the OVC be numbered from 1 to *m*. And let *S* be a non-empty subset of the ordered pairs of OVC End Points associated by the OVC. That is $S \subseteq \{\langle i, j\rangle \mid i = 1,\ldots,m,\ j = 1,\ldots,m, i \neq j\}, S \neq \varnothing$.

- Let $\Delta\tilde{d}_T^{\langle i,j\rangle}$ be the P_v-percentile of the absolute value of the difference between the Frame Delays of all Qualified Frame pairs whose difference in the arrival times of the first bit of each frame in the pair at EI where the OVC End Point *i* is located was exactly $\Delta\tau$.

- If there are no such pairs of frames for OVC End Point *i* and OVC End Point *j*, then $\Delta\tilde{d}_T^{\langle i,j\rangle} = Undefined$.

- Then the Inter-Frame Delay Variation Performance metric **MUST** be the maximum of the defined values $\Delta\tilde{d}_T^{\langle i,j\rangle}$ for $\langle i, j\rangle \in S$, unless all $\Delta\tilde{d}_T^{\langle i,j\rangle}$ are Undefined in which case the performance is Undefined.

This definition is in agreement with the IP packet delay variation definition given in RFC3393 where the delay variation is defined as the difference between the one-way delay of two packets selected according to some selection function and are within a given interval [T_1, T_2].

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

To restate the definition mathematically, let the OVC End Points associated by the OVC be numbered from 1 to *m*. And let *S* be a non-empty subset of the ordered pairs of the OVC End Points associated by the OVC. That is

$$S \subseteq \{\langle i,j\rangle \mid i=1,\ldots,m,\ j=1,\ldots,m, i\neq j\}, S\neq\varnothing.$$

Let

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

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

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

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

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

Then an Inter-Frame Delay Variation Performance metric for an OVC can be expressed as

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

The parameters of an Inter-Frame Delay Variation performance metric are given in Table 11.

Parameter	Description
T	The time interval
S	Non-empty subset of the ordered pairs of OVC End Points associated by the OVC
COS ID	The Class of Service Identifier
P_v	Inter-Frame Delay Variation Performance percentile, $P_v > 0$
$\Delta\tau$	The separation between frame pairs for which Inter-Frame Delay Variation Performance is defined
$\breve{d}$	Inter-Frame Delay Variation Performance Objective

Table 11 – Inter-Frame Delay Variation Parameters

[R59] Given *T, S, COS ID,* P_v, $\Delta\tau$, and $\breve{d}$, the Inter-Frame Delay Variation Performance objective **MUST** be defined by an SLS as met over the time interval *T* for the subset *S* if and only if $\Delta\tilde{d}_{T,S} \le \breve{d}$ or $\Delta\tilde{d}_{T,S}$ is Undefined.

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

[O5] For a Point-to-Point OVC, *S* **MAY** include one or both of the ordered pairs of OVC End Points associated by the OVC.

[O6] For a Multipoint-to-Multipoint OVC, *S* **MAY** be any non-empty subset of the ordered pairs of OVC End Points associated by the OVC.

[R60] For a Rooted-Multipoint OVC, *S* **MUST** be a non-empty subset of the ordered pairs of OVC End Points associated by the OVC such that each ordered pair in *S* contains at least one Root OVC End Point.

7.2.16.3 One-way Frame Loss Ratio Performance for an OVC

[R61] For a given non-empty set of ordered pairs of OVC End Points, a time interval *T*, and a given Class of Service Identifier, the SLS **MUST** define the One-way Frame Loss Ratio Performance metric as follows:

- Let the OVC End Points associated by the OVC be numbered from 1 to *m*. And let S be a non-empty subset of the ordered pairs of OVC End Points associated by the OVC. That is $S \subseteq \{\langle i, j \rangle \mid i = 1,...,m, j = 1,...,m, i \ne j\}, S \ne \emptyset$.
- Let $I_T^{\langle i,j \rangle}$ denote the number of ingress Qualified Frames at the EI where OVC End Point *i* is located that should have been delivered to the EI where OVC End Point *j* is located according to the frame Delivery service attributes (see Sections

7.2.17, 7.2.18, and 7.2.19). Each frame can be a Unicast, Multicast, or Broadcast frame.

- Let $E_T^{\langle i,j \rangle}$ be the number of unique (not duplicate) egress frames where each frame is the first egress frame at the EI where OVC End Point *j* is located that results from a frame counted in $I_T^{\langle i,j \rangle}$.

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

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

The parameters of a One-way Frame Loss Ratio Performance metric are given in Table 12.

Parameter	Description
T	The time interval
S	Non-empty subset of the ordered pairs of OVC End Points associated by the OVC
COS ID	The Class of Service Identifier
$\hat{L}$	One-way Frame Loss Ratio Performance objective

Table 12 – One-way Frame Loss Ratio Performance Parameters

[R62] Given *T*, *S*, *COS ID*, and a One-way Frame Loss Ratio Performance objective, the One-way Frame Loss Performance objective **MUST** be defined by an SLS as met over the time interval *T* for the subset *S* if and only if $FLR_{T,S} \leq \hat{L}$ or $FLR_{T,S}$ is Undefined.

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

[O7] For a Point-to-Point OVC, *S* **MAY** include one or both of the ordered pairs of OVC End Points associated by the OVC.

[O8] For a Multipoint-to-Multipoint OVC, *S* **MAY** be any non-empty subset of the ordered pairs of OVC End Points associated by the OVC.

[R63] For a Rooted-Multipoint OVC, *S* **MUST** be a non-empty subset of the ordered pairs of OVC End Points associated by the OVC such that each ordered pair in *S* contains at least one Root OVC End Point.

7.2.16.4 One-way Availability Performance for an OVC

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

Informally, Availability Performance is based on 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.

For a time interval *T*, and a given Class of Service Identifier, Availability from OVC End Point *i* to OVC End Point *j* is based on the following three parameters:

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

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

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

Let $I_{\Delta t}^{\langle i,j\rangle}$ be the number of ingress frames that meet the following conditions:

- The first bit of each frame arrives at the EI where OVC End Point *i* is located within the time interval Δt,
- Each frame should be delivered to the EI where OVC End Point *j* is located according to the frame delivery service attributes (see Sections 7.2.17, 7.2.18, and 7.2.19). Each frame can be a Unicast, Multicast, or Broadcast frame,

- Each frame has the given Class of Service Identifier,
- Each frame that is subject to an ingress Bandwidth Profile has an Ingress Bandwidth Profile compliance of Green, and
- Each frame that is not subject to an ingress Bandwidth Profile either has no color identifier or a color identifier that corresponds to Green per the color indication requirements of MEF 23 [10]

Let $E_{\Delta t}^{\langle i,j\rangle}$ be the number of unique (not duplicate) egress frames where each frame is the first, unerrored egress frame at the EI where OVC End Point *j* is located that results from a frame counted in $I_{\Delta t}^{\langle i,j\rangle}$.

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

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

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

where $\tilde{I}_{\Delta t}^{\langle i,j\rangle} = I_{\Delta t}^{\langle i,j\rangle} -$ the number of frames discarded by the Service Provider, in order to conform to either the line rate of the EI where OVC End Point *j* is located or an egress Bandwidth Profile (if one is used) at the EI where OVC End Point *j* is located. Such frame drops may occur anywhere in the network, not just near the egress EI. One example of this could be where an egress Bandwidth Profile is applied on a link within the network. Another example of this could be where Green frames for this OVC and Class of Service Identifier that should be delivered to the EI with OVC End Point *j* exceed the line rate on a link within the network, provided the line rate of that link is greater than or equal to the line rate of the EI. 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 EI where OVC End Point *j* is located. The details of such an adjustment are beyond the scope of this document.

Δt_0 is the first short time interval agreed by Service Provider and the Operator at or after turning up the association of OVC End Point *i* and OVC End Point *j*. $A_{\langle i,j\rangle}(\Delta t_k)$ is defined in Figure 10 for $k = 0,1,2,\ldots$.

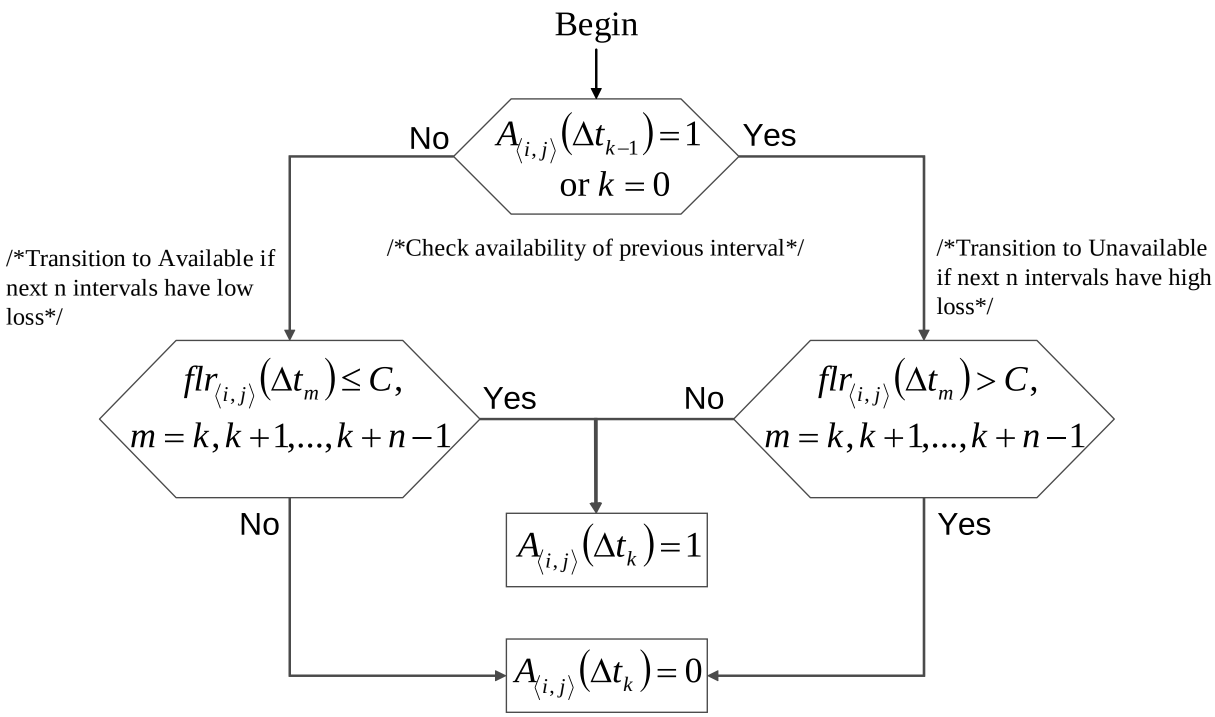

Figure 10 – Flowchart Definition of $A_{\langle i,j\rangle}(\Delta t_k)$

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

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

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

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

In the event of a conflict between the above equations and Figure 10, the content of Figure 10 is controlling.

The availability for Δt_k is based on the frame loss ratio during the short interval and each of the following $n-1$ short intervals and the availability of the previous short time interval. In other words, a sliding window of width $n\Delta t$ is used to define availability. This use of a sliding window is similar to that of ITU-T Y.1563. [23]

Figure 11 presents an example of the determination of the availability for the small time intervals with a sliding window of 10 small time intervals.

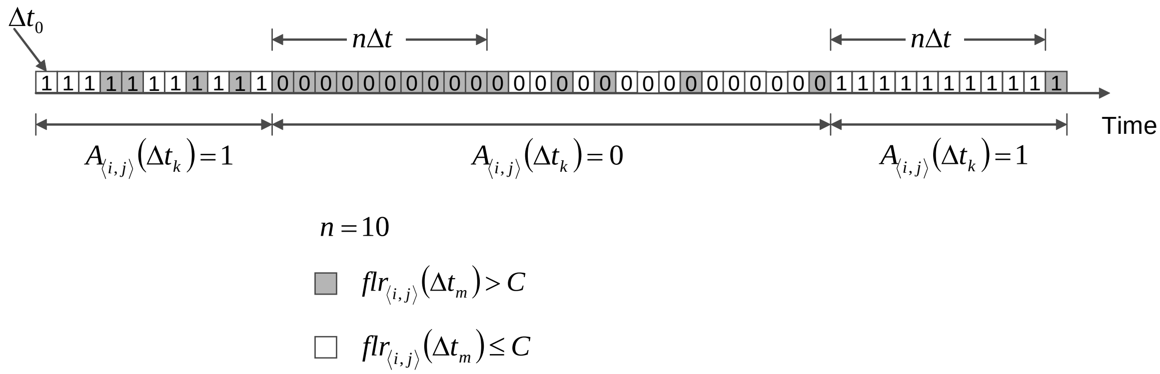

Figure 11 – Example of the Determination of $A_{\langle i,j\rangle}(\Delta t_k)$

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

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

Figure 12 shows an example of the elimination of short time intervals for a Maintenance Interval.

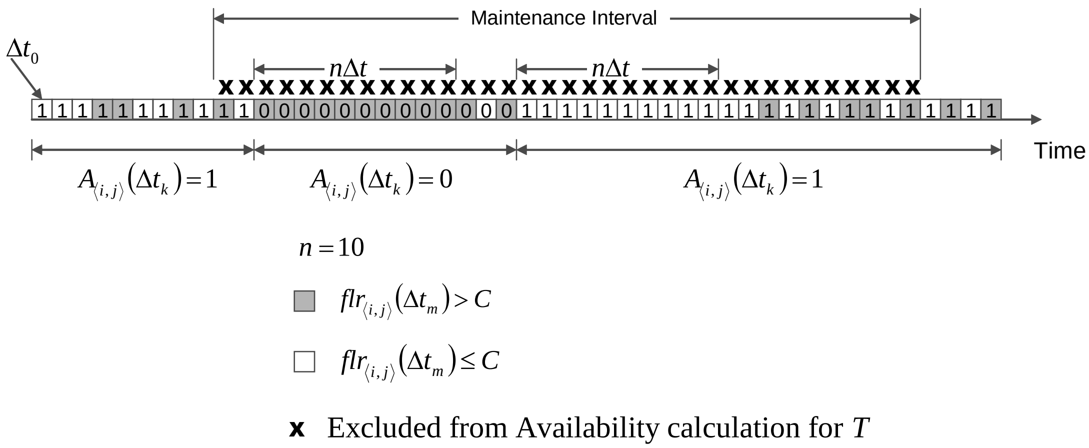

Figure 12 – Example of the Impact of a Maintenance Interval

Let $W_T = \{k \mid \Delta t_k \subseteq T \text{ and } \Delta t_k \text{ does not intersect a Maintenance Interval}\}$ and let $|W_T|$ represent the number of elements in the set W_T. Then the Availability for a particular Class of Service Identifier from OVC End Point *i* to OVC End Point *j* for a time interval *T*, expressed as percentage, is defined by

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

Note that the definition of W_T means that the boundaries of *T* and the boundaries of a Maintenance Interval do not have to align with the boundary of a Δt_k. A Δt_k that straddles the boundary between two *T*'s is excluded from the definition of Availability Performance for each interval *T*. A Δt_k that straddles the boundary of a Maintenance Interval is also excluded from the definition of Availability Performance.

Let the OVC End Points associated by the OVC be numbered $1,2,\ldots,m$ and let *S* be a non-empty subset of the ordered pairs of OVC End Points, i.e.,

$$S \subseteq \{\langle i, j\rangle \mid i = 1,2,\ldots m, j = 1,2,\ldots,m, i \neq j\}, S \neq \varnothing.$$

Then the Availability for a particular Class of Service Identifier for the set *S* is defined by

$$A_T^S = \min\{A_T^{\langle i,j\rangle} \mid \langle i, j\rangle \in S\}.$$

The parameters of a One-Way Availability Performance Metric are given in Table 13.

Parameter	Description
T	The time interval
S	Non-empty subset of the ordered OVC End Point pairs
COS ID	The Class of Service Identifier
Δt	A time interval much smaller than T
C	Unavailability frame loss ratio threshold
n	Number of consecutive small time intervals for assessing availability
$\hat{A}$	Availability Performance Objective expressed as a percentage

Table 13 – Availability Performance Parameters for an OVC

[R64] Given T, S, *COS ID*, Δt, C, n, and $\hat{A}$, the SLS **MUST** define the Availability Performance objective as being met if and only if $A_T^S \geq \hat{A}$.

[O9] For a Point-to-Point OVC, S **MAY** include one or both of the ordered pairs of OVC End Points associated by the OVC.

[O10] For a Multipoint-to-Multipoint OVC, S **MAY** be any non-empty subset of the ordered pairs of OVC End Points associated the OVC.

[R65] For a Rooted-Multipoint OVC, S **MUST** be a non-empty subset of the ordered pairs of OVC End Points associated by the OVC such that each ordered pair in S contains at least one Root OVC End Point.

7.2.16.5 One-way Resiliency Performance for an OVC

This section defines attributes for the Resiliency performance of an ordered pair of OVC End Points, $\langle i,j \rangle$. The definitions depend on the availability status of each Δt to determine whether performance counts toward objectives. The Resiliency attributes are similar to the definitions of Severely Errored Seconds (SES) and Consecutive SES in Section 9 and Annex B (respectively) of ITU-T Recommendation Y.1563 [23], when $\Delta t = 1$ second.

Figure 13 illustrates how the two resiliency attributes, counts of High Loss Intervals and counts of Consecutive High Loss Intervals fit into the hierarchy of time and other attributes.

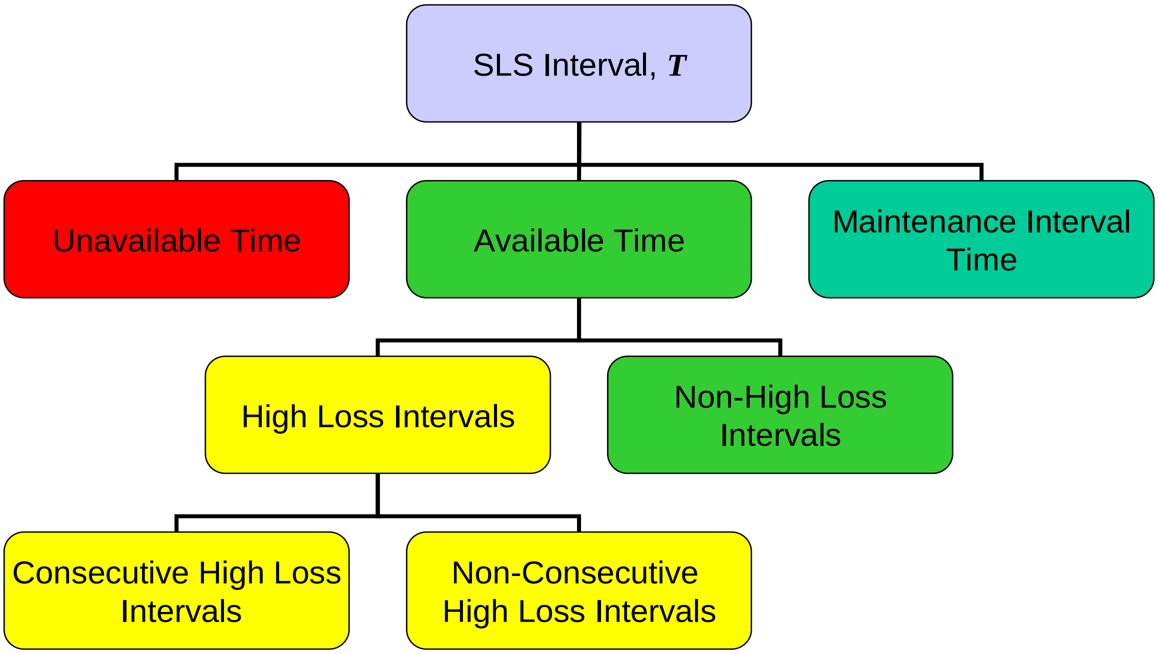

Figure 13 – Hierarchy of Time Showing the Resiliency Attributes

A High Loss Interval (HLI) is a small time interval (having the same duration as the interval, Δt, defined in Section 7.2.16.4) with a high frame loss ratio. When sufficient HLIs are adjacent, the interval is designated a Consecutive High Loss Interval (CHLI). Section 7.2.16.4 defines terminology for Availability. This section re-uses that terminology and defines the following terms:

- $H_{\langle i,j\rangle}(\Delta t_k)$: the high loss state of Δt_k,
  - equal to 1 when $flr_{\langle i,j\rangle}(\Delta t) > C$ and $A_{\langle i,j\rangle}(\Delta t_k) = 1$, equal to 0 otherwise, including any Δt_k that intersects a Maintenance Interval
- $L_T^{\langle i,j\rangle}$: Count of High Loss Intervals over T
- $\hat{L}$: HLI Count Objective for S, T, and a given Class of Service Identifier
- p: the minimum integer number of Δt's in the consecutive (sliding) window (with $0 < p < n$) to qualify as a CHLI
- $B_T^{\langle i,j\rangle}$: Count of p or more consecutive High Loss Intervals occurring in T
- $\hat{B}$: CHLI Count Objective for S, T, and a given Class of Service Identifier

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

[R66] For the SLS, the count of High Loss Intervals over T **MUST** be determined by $L_T^{\langle i,j \rangle} = \sum_{\Delta t \in T} H_{\langle i,j \rangle}(\Delta t)$.

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

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

[R67] For the SLS, the Consecutive High Loss Intervals over T **MUST** be determined according to the flow chart in Figure 14.

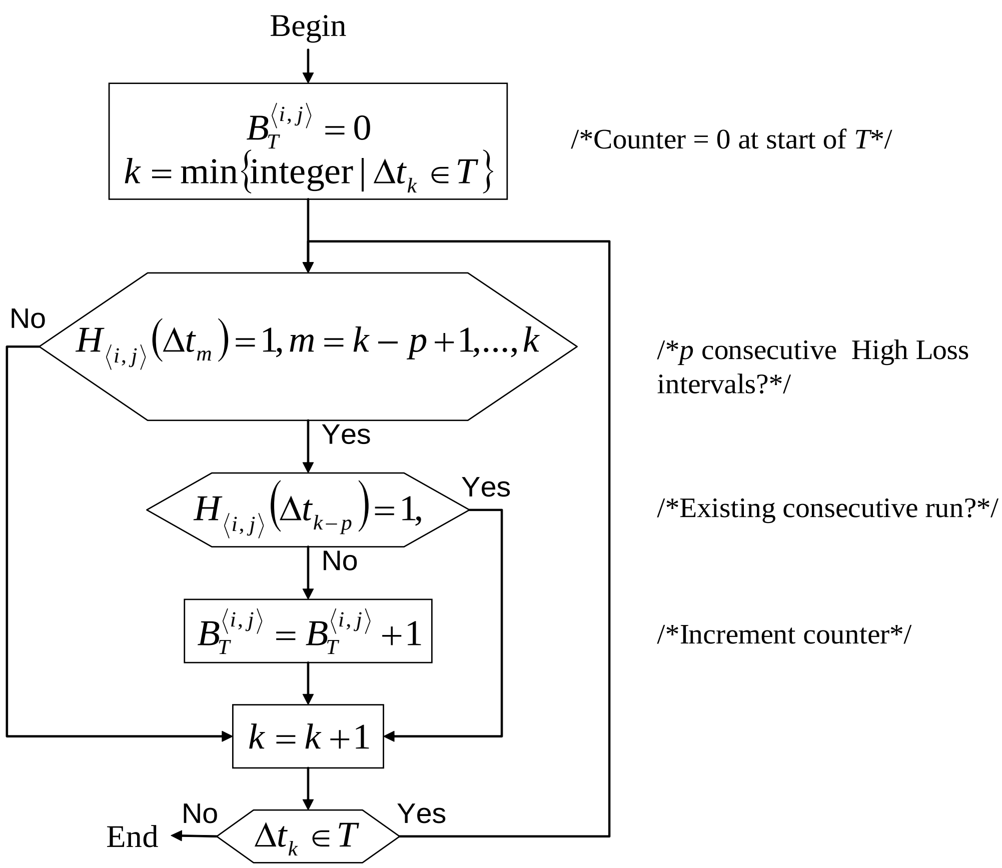

Figure 14 – Determining and Counting Consecutive High Loss Intervals over *T*

Figure 15 shows an example that depicts the HLI and CHLI counting processes.

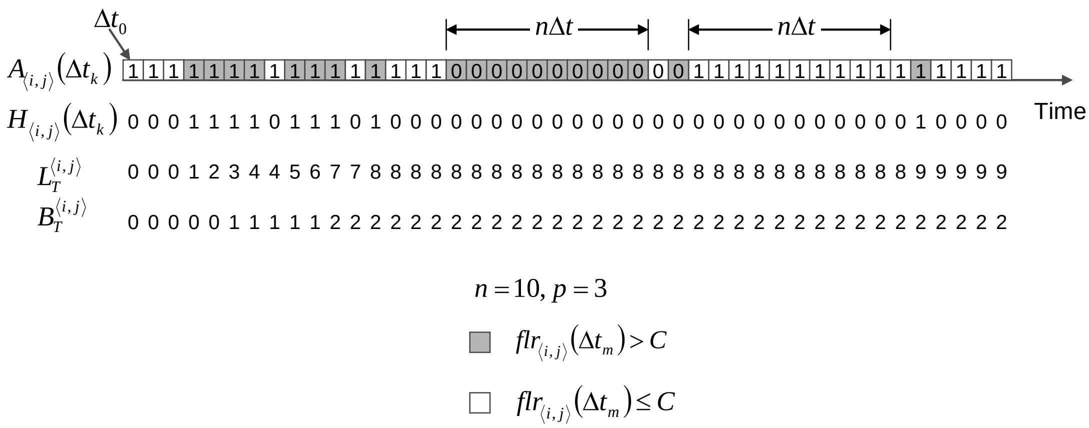

Figure 15 – Example of Counting High Loss Intervals and Consecutive High Loss Intervals

Let the OVC End Points associated by the OVC be numbered $1,2,...,m$ and let S be a non-empty subset of the ordered pairs of OVC End Points, i.e.,

$$S \subseteq \{\langle i,j\rangle \mid i = 1,2,...m,\ j = 1,2,...,m, i \neq j\}, S \neq \emptyset .$$

Then the HLI and CHLI performance attributes for a particular Class of Service Identifier for the set S and time interval T are defined by

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

The parameters of the One-Way Resiliency Performance metrics are given in Table 14.

Parameter	Description
T	The time interval
S	Non-empty subset of the ordered OVC End Point pairs associated by the OVC
COS ID	The Class of Service Identifier
Δt	A time interval much smaller than T
C	Unavailability frame loss ratio threshold
p	Number of consecutive small time intervals for assessing CHLI where $p < n$
$\hat{L}$	HLI Performance Objective expressed as an integer
$\hat{B}$	Consecutive HLI Performance Objective expressed as an integer

Table 14 – Resiliency Performance Parameters for an OVC

[R68] Δt **MUST** have the same value as Δt in Section 7.2.16.4.

[R69] C **MUST** have the same value as C in Section 7.2.16.4.

[R70] Given *T*, *S*, *COS ID*, Δt, *C*, *p*, $\hat{L}$, and $\hat{B}$, the SLS **MUST** define the HLI Performance objective as being met if and only if $L_T^S \leq \hat{L}$, and the CHLI Performance objective as being met if and only if $B_T^S \leq \hat{B}$.

[O11] For a Point-to-Point OVC, *S* **MAY** include one or both of the ordered pairs of OVC End Points associated by the OVC.

[O12] For a Multipoint-to-Multipoint OVC, *S* **MAY** be any non-empty subset of the ordered pairs of OVC End Points associated the OVC so long as it is non-empty.

[R71] For a Rooted-Multipoint OVC, *S* **MUST** be a non-empty subset of the ordered pairs of OVC End Points associated by the OVC such that each ordered pair in *S* contains at least one Root OVC End Point.

7.2.17 Unicast Frame Delivery

This attribute describes how ingress frames mapped to an OVC End Point at an External Interface with a unicast destination MAC address are delivered to the other OVC End Points associated by the OVC[9].

[R72] When the Unicast Frame Delivery is "unconditional," all properly formatted ingress frames mapped to an OVC End Point at an External Interface with a unicast destination MAC address **MUST** be delivered to all of the other OVC End Points associated by the OVC provided [R33],[R34], [R35], and [R36] are satisfied.

[R73] When the Unicast Frame Delivery is "conditional," a properly formatted ingress frame mapped to an OVC End Point at an External Interface with a unicast destination MAC address is delivered to all or a subset of all of the other OVC End Points associated by the OVC depending on certain conditions being met. When "conditional" is in force, the conditions for delivery **MUST** be specified and such conditions **MUST** satisfy [R33],[R34], [R35], and [R36].

An example of such a condition is that the destination MAC address is known by the Operator MEN to be "at" the OVC End Point.

7.2.18 Multicast Frame Delivery

This attribute describes how ingress frames mapped to an OVC End Point at an External Interface with a multicast destination MAC address are delivered to the other OVC End Points associated by the OVC.[9]

[R74] When the Multicast Frame Delivery is "unconditional," all properly formatted ingress frames mapped to an OVC End Point at an External Interface with a mul-

[9] Assuming normal operation, e.g., valid FCS, no network congestion, and assuming that the frames comply with the Bandwidth Profile if any.

ticast destination MAC address **MUST** be delivered to all of the other End Points associated by the OVC provided [R33],[R34], [R35], and [R36] are satisfied.

[R75] When the Multicast Frame Delivery is "conditional," a properly formatted ingress frame mapped to an OVC End Point at an External Interface with a multicast destination MAC address is delivered to all or a subset of all of the other OVC End Points associated by the OVC depending on certain conditions being met. When "conditional" is in force, the conditions for delivery **MUST** be specified and such conditions **MUST** satisfy [R33],[R34], [R35], and [R36].

7.2.19 Broadcast Frame Delivery

This attribute describes how ingress frames mapped to an OVC End Point at an External Interface with the broadcast destination MAC address are delivered to the other End Points associated by the OVC.[9]

[R76] When the Broadcast Frame Delivery is "unconditional," all properly formatted ingress frames mapped to an OVC End Point at an External Interface with the broadcast destination MAC address **MUST** be delivered to all of the other OVC End Points associated by the OVC provided [R33],[R34], [R35], and [R36] are satisfied.

[R77] When the Broadcast Frame Delivery is "conditional," a properly formatted ingress frame mapped to an OVC End Point at an External Interface with the broadcast destination MAC address is delivered to all or a subset of all of the other OVC End Points associated by the OVC depending on certain conditions being met. When "conditional" is in force, the conditions for delivery **MUST** be specified and such conditions **MUST** satisfy [R33],[R34], [R35], and [R36].

7.2.20 Layer 2 Control Protocol Tunneling

The Layer 2 Control Protocol Service Frame is described in MEF 10.2. [5]. An ENNI Frame with a Destination MAC Address shown in Table 15 is defined to be a Layer 2 Control Protocol ENNI Frame. Additional ways of denoting a Layer 2 Control Protocol ENNI Frame at a given ENNI can be agreed to by the two Operators involved in the given ENNI.

MAC Addresses[10]
01-80-C2-00-00-00 through 01-80-C2-00-00-0F
01-80-C2-00-00-20 through 01-80-C2-00-00-2F
01-80-C2-00-00-10

Table 15 – MAC Addresses that Identify a Layer 2 Control Protocol ENNI Frame

[R78] When a L2CP Service Frame or L2CP ENNI Frame is tunneled, the frame **MUST** be delivered to all OVC End Points, other than the ingress OVC End Point, that

[10] Hexadecimal canonical format

are associated by the OVC and the format relationships detailed in Table 16 **MUST** be maintained.

Note that [R98] mandates that an ingress Service Frame that is not mapped to an existing OVC End Point is not to be tunneled.

Note that [R19] mandates that an ingress ENNI Frame that is not mapped to an existing OVC End Point is not to be tunneled. In view of [R20], this means that an ingress L2CP ENNI Frame that does not have an S-Tag is not to be tunneled because the Operator has no information on which OVC to use to tunnel the frame.

Ingress Interface	**Egress Interface**	**Egress Frame Format**[11]
UNI (L2CP Service Frame)	UNI (L2CP Service Frame)	Identical to the ingress frame.
UNI (L2CP Service Frame)	ENNI (L2CP ENNI Frame)	All fields from the Destination Address through the Payload of the ingress Service Frame present and unchanged. S-Tag added in after the Source Address.
ENNI (L2CP ENNI Frame)	UNI (L2CP Service Frame)	All fields from the Destination Address through the Payload except the S-Tag of the ingress ENNI Frame present and unchanged. No S-Tag is present.
ENNI (L2CP ENNI Frame)	ENNI (L2CP ENNI Frame)	All fields from the Destination Address through the Payload of the ingress ENNI Frame present. The content of the S-Tag can be changed while all other fields are unchanged.

Table 16 – Format Relationships for Tunneled L2CP Service and ENNI Frames

7.3 OVC End Point per ENNI Service Attributes

There are service attributes for each instance of an OVC End Point at a given ENNI. These are summarized in Table 17 and described in detail in the following sub-sections.

[11] The Frame Check Sequence in the egress frame might need to be recalculated.

Attribute Name	Summary Description	Possible Values
OVC End Point Identifier	An identifier for the OVC End Point intended for management purposes	A string that is unique across the Operator MEN
Trunk Identifiers	For a Trunk OVC End Point, specifies the S-VLAN ID values used on the ENNI to distinguish frames originating at a Root UNI and frames originating at a Leaf UNI	<Root S-VLAN ID value, Leaf S-VLAN ID value> for Trunk OVC End Points; Not Applicable to Root or Leaf OVC End Points
Class of Service Identifiers	The way that a Class of Service is determined for an ENNI Frame at each ENNI	The OVC associating the End Point and non-overlapping sets of values of the S-Tag Priority Code Point
Ingress Bandwidth Profile Per OVC End Point	Ingress policing by the Operator MEN on all ingress ENNI Frames mapped to the OVC End Point	No or parameters as defined in Section 7.6.1
Ingress Bandwidth Profile Per ENNI Class of Service Identifier	Ingress policing by the Operator MEN on all ingress ENNI Frames with the Class of Service Identifier for the receiving Operator MEN	No or parameters as defined in Section 7.6.1
Egress Bandwidth Profile Per End Point	Egress policing by the Operator MEN on all egress ENNI Frames mapped to the OVC End Point	No or parameters as defined in Section 7.6.1
Egress Bandwidth Profile Per ENNI Class of Service Identifier	Egress policing by the Operator MEN on all egress ENNI Frames with the Class of Service Identifier for the receiving MEN	No or parameters as defined in Section 7.6.1

Table 17 – OVC End Point per ENNI Service Attributes

7.3.1 OVC End Point Identifier

The OVC End Point Identifier is a string administered by the Operator that is used to identify an OVC End Point within the Operator MEN. It is intended for management and control purposes. The OVC End Point Identifier is not carried in any field in the ENNI Frame.

[R79] The OVC End Point Identifier **MUST** be unique among all such identifiers for OVC End Points supported by the Operator MEN.

[R80] The OVC End Point Identifier **MUST** contain no more than 45 bytes.

7.3.2 Trunk Identifiers

When the OVC End Point Role is Trunk, the End Point Map at the ENNI contains two S-VLAN ID values that map to the OVC End Point as mandated by [R21]. One value is the Root S-VLANI ID value and the other is the Leaf S-VLAN ID value.

[R81] The S-VLAN ID field of each ENNI Frame that is the result of an ingress Service Frame at a Root UNI **MUST** contain the Root S-VLAN ID value.

[R82] The S-VLAN ID field of each ENNI Frame that is the result of an ingress Service Frame at a Leaf UNI **MUST** contain the Leaf S-VLAN ID value.

The requirements regarding OVC End Points associated by a Rooted-Multipoint OVC are specified in Section 7.2.2.

7.3.3 Class of Service Identifiers

The delivery performance for an ingress ENNI Frame is dependent on the particular Class of Service Identifier that applies to it. A Class of Service Identifier is a set of one or more S-Tag PCP values.[12] Each Class of Service Identifier indicates a single Class of Service Name as defined in MEF 23.1[10]. The Class of Service Name that applies to an ingress S-Tagged ENNI Frame that is mapped to the OVC End Point is the Class of Service identified by the Class of Service Identifier that contains the S-Tag PCP value in the frame.

For example, the S-Tag PCP values 0, 1, 2, and 3 could constitute a Class of Service Identifier that indicates the silver service while the S-Tag PCP values 4, 5, 6, and 7 could constitute a different Class of Service Identifier that indicates the gold service. In this example, an S-Tagged ENNI Frame with S-Tag PCP value = 3 would be given the silver service.

Note that when multiple S-VLAN ID values are mapped to the same OVC End Point as with End Point Map Bundling (see Section 7.1.7.2), the CoS Identifier for each ingress S-Tagged ENNI Frame mapped to the given OVC End Point is independent of the S-VLAN ID value in the ENNI Frame.

In general, at a given ENNI, each OVC End Point can have a different Class of Service Identifiers attribute but the following requirements apply.

[R83] For each OVC End Point at an ENNI, each possible S-Tag PCP value **MUST** be included in exactly one Class of Service Identifier.

[R83] means that the sets of S-Tag PCP values for the Class of Service Identifiers are disjoint and their union equals the set of all S-Tag PCP values.

[O13] One of the Class of Service Identifiers **MAY** indicate 100% discard.

[12] MEF 28 [20] introduces additional Class of Service Identifiers at the ENNI..

[O13] means that the Operator MEN can discard all S-Tagged Frames whose S-Tag PCP value belongs to one of the Class of Service Identifiers.

[R84] At a given ENNI, all OVC End Points associated by the same OVC **MUST** have the same Class of Service Identifiers.

[R84] means that if an OVC associates multiple OVC End Points at an ENNI, the Class of Service Identifier for an ENNI Frame at this ENNI is independent of the particular OVC End Point to which it is mapped. Instead the Class of Service Identifier is based on the OVC and the S-Tag PCP value.

[D3] An Operator MEN **SHOULD** support the use of different Class of Service Identifiers attributes for OVC End Points at an ENNI that are associated by different OVCs.

[D3] means that it is recommended that an Operator MEN be able to map a given S-Tag PCP value to a different class of service for OVC End Points at an ENNI that are associated by different OVCs.

Either the Drop Eligible Indicator (DEI) bit or the S-Tag PCP can be used to indicate the ENNI Frame Color and the following requirements apply.

[R85] Color indication for each ENNI Frame **MUST** conform to requirements [R4] and [R5] of MEF 23.1 [10].

[R86] If the S-Tag PCP field is used to indicate Color for the ENNI Frame, then the Class of Service Identifiers **MUST** map S-Tag PCP values to L, M, and H as per [R10] of MEF 23.1 **[10]**.

[R87] If the DEI bit is used to indicate Color for the ENNI Frame, then the Class of Service Identifiers **MUST** map S-Tag PCP values to L, M, and H as per [R10] of MEF 23.1 **[10]**.

7.3.3.1 Class of Service Identifiers for Egress ENNI Frames

An egress ENNI Frame does not specify a Class of Service Identifier for the Operator MEN from which it is being transmitted. The S-Tag PCP value for an egress ENNI Frame only indicates the Class of Service Identifier for the frame with respect to the Operator MEN that is receiving it. Thus it is the responsibility of the transmitting Operator MEN to set the appropriate S-Tag PCP value such that the frame is given the appropriate Class of Service by the receiving Operator MEN. Section 10.8 presents an example of setting Class of Service Identifiers at an ENNI.

Note that MEF 23.1 [10] contains additional requirements on Classes of Service and S-Tag PCP values.

7.3.4 Ingress Bandwidth Profile per OVC End Point

The Ingress Bandwidth Profile per OVC End Point describes ingress policing by the Operator MEN on all ingress ENNI Frames mapped to a given OVC End Point.

[R88] When the Ingress Bandwidth Profile per OVC End Point is in force for a given OVC End Point, suitable parameters <CIR, CBS, EIR, EBS, CF, CM> as defined in Section 7.6.1 **MUST** be specified and the algorithm of Section 7.6.1 **MUST** be applied to all ingress ENNI Frames that are mapped to the given OVC End Point.

[R89] The Color Mode for the Bandwidth Profile Algorithm **MUST** be color-aware.

Figure 16 illustrates an example of the application of Ingress Bandwidth Profiles per OVC End Point. In this example, OVC End Point$_1$ could have *CIR*=15 Mbps, OVC End Point$_2$ could have *CIR* = 10 Mbps, and OVC End Point$_3$ could have *CIR* = 20 Mbps.

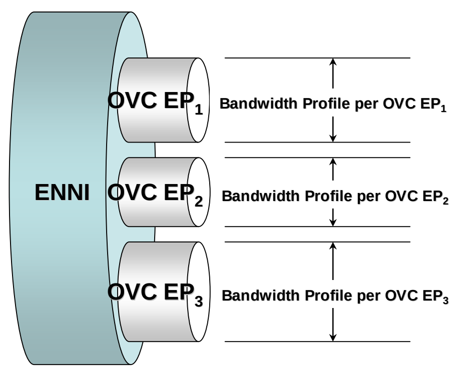

Figure 16 – Ingress Bandwidth Profile per OVC End Point

7.3.5 Egress Bandwidth Profile per OVC End Point

The Egress Bandwidth Profile per OVC End Point describes egress policing by the Operator on all egress ENNI Frames that are mapped to a given OVC End Point.

[R90] When the Egress Bandwidth Profile per OVC End Point is in force for a given OVC End Point, suitable parameters <CIR, CBS, EIR, EBS, CF, CM> as defined in Section 7.6.1 **MUST** be specified and all egress ENNI Frames mapped to the given OVC End Point **MUST** have the property defined in 7.6.3.

[R91] The Color Mode for the Bandwidth Profile Algorithm **MUST** be color-aware.

7.3.6 Ingress Bandwidth Profile per Class of Service Identifier

The Ingress Bandwidth Profile per Class of Service Identifier describes ingress policing by the Operator MEN on all ingress ENNI Frames with a given Class of Service Identifier.

[R92] When the Ingress Bandwidth Profile per Class of Service Identifier is in force for a given ENNI Class of Service Identifier, suitable parameters <CIR, CBS, EIR, EBS, CF, CM> as defined in Section 7.6.1 **MUST** be specified and the algorithm of Section 7.6.1 **MUST** be applied to all ingress ENNI Frames mapped to the OVC End Point that have the given ENNI Class of Service Identifier.

[R93] The Bandwidth Profile Algorithm **MUST** be color-aware.

7.3.7 Egress Bandwidth Profile per Class of Service Identifier

The Egress Bandwidth Profile per Class of Service Identifier describes egress policing by the Operator MEN on all egress ENNI Frames with a given Class of Service Identifier.

[R94] When the Egress Bandwidth Profile per Class of Service Identifier is in force for a given ENNI Class of Service Identifier, suitable parameters <CIR, CBS, EIR, EBS, CF, CM> as defined in Section 7.6.1 **MUST** be specified and all egress ENNI Frames mapped to the OVC End Point with the given Class of Service Identifier **MUST** have the property defined in Section 7.6.3.

[R95] The Color Mode for the Bandwidth Profile Algorithm **MUST** be color-aware.

7.4 UNI Attributes

These are identical to the UNI attributes specified in Sections 7.1, 7.2, 7.3, 7.4, 7.5, 7.6.1, 7.7, 7.8, 7.9, 7.10, 7.11.2.1, 7.11.3.1, and 7.12 of MEF 10.2 [5]. Note that the details of the UNI attributes as agreed by the Service Provider and Operator may be different than the details of the UNI attributes as agreed by the Service Provider and the Subscriber. See the discussion following [R99].

7.5 OVC per UNI Service Attributes

There are service attributes for each instance of an OVC at a specific UNI. Since an OVC can only associate one OVC End Point that it is at a UNI ([R28]), these service attributes can be equivalently viewed as OVC End Point per UNI service attributes. These service attributes are summarized in Table 18 and described in detail in the following sub-sections.

Attribute Name	Summary Description	Possible Values
UNI OVC Identifier	An identifier for the instance of the OVC at a UNI intended for management purposes	A string formed by the concatenation of the UNI Identifier and the OVC Identifier
OVC End Point Map	The CE-VLAN ID(s) that map to the OVC End Point at the UNI	A list of one or more CE-VLAN ID values
Class of Service Identifiers	The way that a Class of Service is determined for a frame at each UNI	Non-overlapping sets of values of C-Tag Priority Code Point values or DSCP values as specified in Section 7.5.3
Ingress Bandwidth Profile Per OVC End Point at a UNI	Ingress policing by the Operator MEN on all ingress Service Frames mapped to the OVC End Point	No or parameters as defined in Section 7.11.1 in MEF 10.2 [5]
Ingress Bandwidth Profile Per Class of Service Identifier at a UNI	Ingress policing by the Operator on all ingress Service Frames with a given Class of Service Identifier	No or parameters as defined in Section 7.11.1 in MEF 10.2 [5]
Egress Bandwidth Profile Per OVC End Point at a UNI	Egress policing by the Operator on all egress Service Frames mapped to the OVC End Point	No or parameters as defined in Section 7.11.1 in MEF 10.2 [5]
Egress Bandwidth Profile Per Class of Service Identifier at a UNI	Egress policing by the Operator on all egress Service Frames with a given Class of Service Identifier	No or parameters as defined in Section 7.11.1 in MEF 10.2 [5]

Table 18 – OVC per UNI Service Attributes

7.5.1 UNI OVC Identifier

The OVC Identifier is a string that is used to identify an OVC at the UNI. It is intended for management and control purposes.

[R96] The UNI OVC Identifier **MUST** be the concatenation of the UNI Identifier and the OVC Identifier.

7.5.2 OVC End Point Map at the UNI

A Service Frame is mapped to the OVC End Point via the CE-VLAN ID of the Service Frame as defined in Section 7.6.1 of MEF 10.2. [5]

[R97] The OVC End Point at the UNI for a Service Frame **MUST** be identified by the value of CE-VLAN ID of the Service Frame.

[R98] An ingress Service Frame that is not mapped to an existing OVC End Point or EVC at the UNI **MUST** be discarded.

[O14] Multiple CE-VLAN IDs **MAY** map to a single OVC End Point.

[R99] Each CE-VLAN ID **MUST** have one of the following mutually exclusive properties; 1) it maps to one OVC End Point, 2) it maps to one EVC that associates UNIs within the Operator MEN, 3) it does not map to either such an EVC or an OVC End Point.

Note that this document is describing the attributes as agreed to by the Service Provider and Operator and therefore there is awareness of an OVC End Point at a UNI. MEF 10.2 [5] describes attributes as agreed to by the Subscriber and Service Provider for which OVC End Points are invisible. From the Subscriber's viewpoint, at a UNI, each CE-VLAN ID either maps to an EVC or it does not map to an EVC.

[R100] When an OVC associating the OVC End Point to which the CE-VLAN ID for untagged and priority tagged Service Frames is mapped does not have the CE-VLAN ID Preservation attribute in force, egress Service Frames mapped to this OVC End Point at the given UNI **MUST** be untagged.

7.5.3 Class of Service Identifiers

[R101] There **MUST** be three mutually exclusive ways to determine the Class of Service Identifier from the content of a given Data Service Frame at UNI as described in Sections 7.5.3.1, 7.5.3.2, and 7.5.3.3.

Note that Sections 7.5.3.1, 7.5.3.2, and 7.5.3.3 describe Class of Service Identifiers for ingress Data Service Frames. A Data Service Frame is a Service Frame that is not carrying a Layer 2 Control Protocol. (See Section 6.5.1 of MEF 10.2[5].)

7.5.3.1 Class of Service Identifier Based on OVC End Point

[R102] When the Class of Service Identifier is based on OVC End Point, all ingress Data Service Frames mapped to the same OVC End Point at the UNI **MUST** have the same Class of Service Identifier.

As an example, consider OVC End Point 1 and OVC End Point 2 at a UNI. Data Service Frames mapped to OVC End Point 1 have a first Class of Service Identifier that indicates gold service. Data Service Frames mapped to OVC End Point 2 have a second Class of Service Identifier that indicates silver service.

7.5.3.2 Class of Service Identifier Based on Priority Code Point Field

[R103] When the Class of Service Identifier is based on Priority Code Point field, the Class of Service Identifier for an ingress Data Service Frame at the UNI **MUST**

be determined by the OVC End Point and non-overlapping sets of values of the PCP field in the C-Tag.

[R104] When the Class of Service Identifier is based on Priority Code Point field, if the ingress frame does not contain a C-Tag, it **MUST** have the same Class of Service Identifier as an ingress frame with Priority Code Point field = 0 in the C-Tag.

[R105] When the Class of Service Identifier is based on Priority Code Point field, the union of the sets of PCP field values **MUST** contain all of the possible values.

As an example, consider OVC End Point 1 and OVC End Point 2 at a UNI. C-Tagged Data Service Frames mapped to OVC End Point 1 with Priority Code Point values 4, 5, 6, and 7 have a first Class of Service Identifier that indicates gold service. C-Tagged Data Service Frames mapped to OVC End Point 1 with Priority Code Point values 0 and 3 have a second Class of Service Identifier that indicates silver service. C-Tagged Data Service Frames mapped to OVC End Point 1 with Priority Code Point values 1 and 2 have a third Class of Service Identifier that indicates Service Frame discard. Data Service Frames without a C-Tag mapped to OVC End Point 1 also have the second Class of Service Identifier that indicates silver service. C-Tagged Data Service Frames mapped to OVC End Point 2 with Priority Code Point value 7 have a fourth Class of Service Identifier that indicates platinum service. All other Data Service Frames mapped to OVC End Point 2 have a fifth Class of Service Identifier that indicates gold service.

7.5.3.3 Class of Service Identifier Based on DSCP

[R106] When the Class of Service Identifier is based on DSCP, the Class of Service Identifier for an ingress Data Service Frame at the UNI containing an IP packet **MUST** be determined by the OVC End Point and non-overlapping sets of values of the DSCP.

[R107] When the Class of Service Identifier is based on DSCP, the union of the sets of DSCP values **MUST** contain all of the possible DSCP values.

[R108] When the Class of Service Identifier is based on DSCP, each ingress Data Service Frame at the UNI not containing an IP packet and mapped to a given OVC End Point **MUST** have the same Class of Service Identifier with a value agreed upon by the Operator and the Service Provider.

7.5.4 Ingress Bandwidth Profile per OVC End Point at a UNI

The Ingress Bandwidth Profile per OVC End Point at a UNI describes ingress policing by the Operator MEN on all ingress Service Frames mapped to a given OVC End Point at a UNI.

[R109] When the Ingress Bandwidth Profile per OVC End Point at a UNI is in force for a given OVC End Point, suitable parameters <*CIR*, *CBS*, *EIR*, *EBS*, *CF*, *CM*> as defined in Section 7.11.1 of [5] **MUST** be specified and the algorithm of Section 7.11.1 of [5] **MUST** be applied to all ingress Service Frames that are mapped to the given OVC End Point.

Note that the algorithm and parameters defined in Section 7.6.1 are identical to the algorithm and parameters defined in 7.11.1 of [5] with the exception of the value of the *CM* parameter. At the ENNI, *CM* is mandated to be "color-aware" by requirements [R89], [R91], [R93], and [R95] while color awareness is optional at the UNI.

7.5.5 Ingress Bandwidth Profile per Class of Service Identifier at a UNI

The Ingress Bandwidth Profile per Class of Service Identifier at a UNI describes ingress policing by the Operator MEN on all ingress Service Frames with a given Class of Service Identifier at a UNI.

[R110] When the Ingress Bandwidth Profile per Class of Service Identifier at a UNI is in force for a given Class of Service Identifier, suitable parameters <*CIR*, *CBS*, *EIR*, *EBS*, *CF*, *CM*> as defined in Section 7.11.1 of [5] **MUST** be specified and the algorithm of Section 7.11.1 of [5] **MUST** be applied to all ingress Service Frames with the given Class of Service Identifier.

Note that the algorithm and parameters defined in Section 7.6.1 are identical to the algorithm and parameters defined in 7.11.1 of [5] with the exception of the value of the *CM* parameter. At the ENNI, *CM* is mandated to be "color-aware" by requirements [R89], [R91], [R93], and [R95] while color awareness is optional at the UNI.

7.5.6 Egress Bandwidth Profile per OVC End Point at a UNI

The Egress Bandwidth Profile per OVC End Point at a UNI describes egress policing by the Operator MEN on all egress Service Frames mapped to a given OVC End Point at a UNI.

[R111] When the Egress Bandwidth Profile per OVC End Point at a UNI is in force for a given OVC End Point, suitable parameters <*CIR*, *CBS*, *EIR*, *EBS*, *CF*, *CM*> as defined in Section 7.11.1 of [5] **MUST** be specified and when the algorithm of Section 7.11.1 of [5] using these parameters is applied to these egress Service Frames, the result for each Service Frame **MUST** be to declare the Service Frame either Green or Yellow.

Note that the algorithm and parameters defined in Section 7.6.1 are identical to the algorithm and parameters defined in 7.11.1 of [5] with the exception of the value of the *CM* parameter. At the ENNI, *CM* is mandated to be "color-aware" by requirements [R89], [R91], [R93], and [R95] while color awareness is optional at the UNI.

7.5.7 Egress Bandwidth Profile per Class of Service Identifier at a UNI

The Egress Bandwidth Profile per Class of Service Identifier at a UNI describes egress policing by the Operator MEN on all egress Service Frames with a given Class of Service Identifier at a UNI.

[R112] When the Egress Bandwidth Profile per Class of Service Identifier at a UNI is in force for a given Class of Service Identifier, suitable parameters <*CIR*, *CBS*, *EIR*, *EBS*, *CF*, *CM*> as defined in Section 7.11.1 of [5] **MUST** be specified and when the algorithm of Section 7.11.1 of [5] using these parameters is applied to these egress Service Frames, the result for each Service Frame **MUST** be to declare the Service Frame either Green or Yellow.

Note that the algorithm and parameters defined in Section 7.6.1 are identical to the algorithm and parameters defined in 7.11.1 of [5] with the exception of the value of the *CM* parameter. At the ENNI, *CM* is mandated to be "color-aware" by requirements [R89], [R91], [R93], and [R95] while this is not mandated at the UNI.

7.6 Bandwidth Profile at the ENNI

The Bandwidth Profile defines the set of traffic parameters applicable to a sequence of ENNI Frames. Associated with the Bandwidth Profile is an algorithm to determine ENNI Frame compliance with the specified parameters. In the case of an Ingress Bandwidth Profile, rate enforcement is accomplished via the disposition of non-compliant ENNI Frames.

Many of the Bandwidth Profiles in this Technical Specification are based on the parameters and algorithm described in Section 7.6.1.

7.6.1 Bandwidth Profile Parameters and Algorithm

The parameters comprising the Bandwidth Profile are:

- **Committed Information Rate** (*CIR*) expressed as bits per second.

[R113] CIR **MUST** be ≥ 0.

- **Committed Burst Size** (*CBS*) expressed as bytes.

[R114] When CIR > 0, CBS **MUST** be greater than or equal to the largest Maximum Transmission Unit size allowed for the ENNI Frames that the Bandwidth Profile applies to.

- **Excess Information Rate** (*EIR*) expressed as bits per second.

[R115] EIR **MUST** be ≥ 0.

- **Excess Burst Size** (*EBS*) expressed as bytes.

[R116] When EIR > 0, EBS **MUST** be greater than or equal to the largest Maximum Transmission Unit size allowed for the ENNI Frames that the Bandwidth Profile applies to.

- **Coupling Flag** (*CF*) has value 0 or 1.
- **Color Mode** (*CM*) has value "color-blind" or "color-aware".

Each ENNI Frame is classified to determine which, if any, Bandwidth Profile is applicable to the ENNI Frame. Operation of the Bandwidth Profile algorithm is governed by the six parameters, <*CIR, CBS, EIR, EBS, CF, CM*>. The algorithm declares each ENNI Frame to be compliant or non-compliant relative to the Bandwidth Profile. The level of compliance is expressed as one of three colors, Green, Yellow, or Red.

The Bandwidth Profile algorithm is in color aware mode when each ENNI Frame already has a level of compliance (i.e., a color) associated with it and that color is taken into account in determining the level of compliance by the Bandwidth Profile algorithm.

The Bandwidth Profile algorithm is shown in Figure 17. For this algorithm, $B_c(t_0) = CBS$ and $B_e(t_0) = EBS$. $B_c(t)$ and $B_e(t)$ can be viewed as the number of bytes in the Committed and Excess token buckets respectively at a given time *t*.

[R117] For a sequence of ENNI Frames, $\{t_j, l_j\}, j \geq 0, t_{j+1} \geq t_j,$ with arrival times at the reference point t_j and lengths l_j, the level of compliance color assigned to each ENNI Frame **MUST** be defined according to the algorithm in Figure 17.

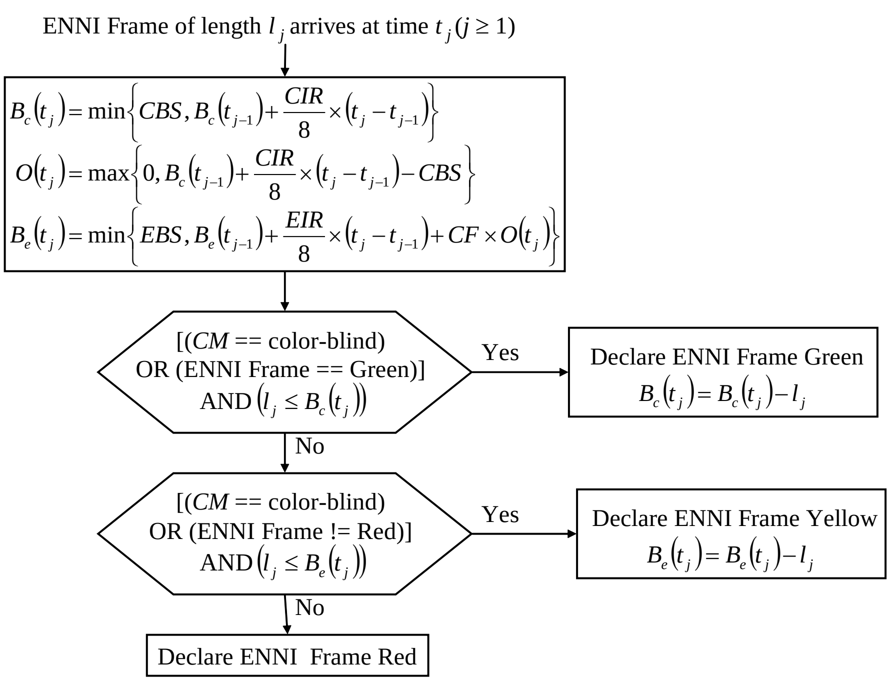

Figure 17 – The Bandwidth Profile Algorithm

The Coupling Flag *CF* is set to either 0 or 1. The choice of the value for CF has the effect of controlling the volume of the ENNI Frames that are declared Yellow. When *CF* is set to 0, the long term average bit rate of ENNI Frames that are declared Yellow is bounded by *EIR*. When *CF* is set to 1, the long term average bit rate of ENNI Frames that are declared Yellow is bounded by

CIR + *EIR* depending on volume of the offered ENNI Frames that are declared Green. In both cases the burst size of the ENNI Frames that are declared Yellow is bounded by *EBS*.

7.6.2 Ingress Bandwidth Profiles Service Attributes

The Ingress Bandwidth Profile is used to regulate the amount of ingress traffic at a particular ENNI. An Ingress Bandwidth Profile is defined for ingress ENNI Frames at the particular ENNI. In other words, the sequence $\{t_j, l_j\}, j \geq 0,$ to which the algorithm of Section 7.6.1 is applied is based on ingress ENNI Frames at an ENNI. There are two Ingress Bandwidth Profile attributes as described in Sections 7.3.4, and 7.3.6.

7.6.2.1 Simultaneous Application of the Ingress Bandwidth Profile Application Models

[O15] Multiple models of Ingress Bandwidth Profile application **MAY** exist simultaneously at an ENNI.

[R118] An ENNI **MUST** be configured such that at most a single Ingress Bandwidth Profile applies to any given ingress ENNI Frame.

[R118] means that if there is a per OVC End Point Ingress Bandwidth Profile, then there cannot be any per Class of Service Ingress Bandwidth Profiles on the OVC that associates the OVC End Point.

7.6.2.2 ENNI Frame Disposition

The disposition of a given ENNI Frame with respect to delivery to an egress External Interface is dependent on the ENNI Frame's level of compliance to the Ingress Bandwidth Profile that is applied to it. This is called the Ingress Bandwidth Profile compliance level and it has three possible values: Green, Yellow, or Red. Table 19 defines how the Ingress Bandwidth Profile compliance is related to the disposition of the ENNI Frame. In this table, "Not Applied" identifies the case where no Ingress Bandwidth Profile was applied to the ENNI Frame in question.

[R119] The disposition of each ENNI Frame for delivery to each egress External Interface **MUST** be as described in Table 19.

Ingress Bandwidth Profile Compliance	ENNI Frame Disposition
Red	Discard
Yellow	Deliver to the egress External Interface according to the Service Attributes of the OVC associating the OVC End Point to which the frame is mapped but SLS performance objectives do not apply.
Green	Deliver to the egress External Interface according to the Service Attributes of the OVC associating the OVC End Point to which the frame is mapped and SLS performance objectives apply.
Not Applied	

Table 19 – ENNI Frame Disposition for Each Egress External Interface

The behavior described in Table 19 is based on the arrival of the ENNI Frame at its ingress ENNI. It does not mandate or constrain in any way the implementation within the Operator MEN.

7.6.3 Egress Bandwidth Profiles Service Attributes

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

The reference point for an Egress Bandwidth Profile is the ENNI. An Egress Bandwidth Profile describes arrival times and lengths of ENNI Frames that will be observed at the ENNI when an Egress Bandwidth Profile is in operation in the Operator MEN. This description is given in terms of what would happen if an observer at the ENNI applied the algorithm of Section 7.6.1 to egress ENNI Frames. This observer would see traffic after it had been subject to rate limiting and/or shaping in the Operator network and thus would have certain characteristics.

[R120] When a sequence of egress ENNI Frames with arrival times and lengths at the ENNI, $\{t_j, l_j\}, j \geq 0$ are subjected to an Egress Bandwidth Profile with parameters <CIR, CBS, EIR, EBS, CF, CM>, the result of applying the algorithm of Section 7.6.1 to these frames **MUST** be to declare each ENNI Frame either Green or Yellow.

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

There are two Egress Bandwidth Profile attributes (one per OVC End Point and one per CoS Identifier) as described in Sections 7.3.5 and 7.3.7.

7.6.3.1 Simultaneous Application of the Egress Bandwidth Profile Application Models

[O16] Multiple models of Egress Bandwidth Profile application **MAY** exist simultaneously for an egress ENNI.

[R121] However, an egress ENNI Frame **MUST** be subject to at most one Egress Bandwidth Profile.

[R121] means that if there is a per OVC End Point Egress Bandwidth Profile, then there cannot be any per Class of Service Egress Bandwidth Profiles on the OVC that associates that OVC End Point.

8 Link OAM Function Support for the ENNI

The ENNI has requirements for Link OAM support based on IEEE Std 802.3 [3].

[R122] For each physical link in the ENNI, an ENNI-N **MUST** be capable of supporting Active DTE mode capabilities as specified in clause 57.2.9 of IEEE Std 802.3 [3].

[R123] For each physical link in the ENNI, an ENNI-N **MUST** be capable of supporting Passive DTE mode capabilities as specified in clause 57.2.9 of IEEE Std 802.3 [3].

[R122] and [R123] do not mandate that Link OAM be run on a given link. [R122] and [R123] mean that when the two Operators agree to support Link OAM on a given link, they will also need to negotiate which sides of the ENNI will function in Active DTE mode with at least one side functioning in Active DTE mode.

Operators using Link OAM on the ENNI could receive unwanted loopback messages which could cause an interruption of traffic using the ENNI. The following two requirements are meant to prevent this situation:

[D4] When Link OAM is enabled on an ENNI-N, the loopback capability **SHOULD** be disabled.

[D5] When Link OAM is enabled on an ENNI-N, it **SHOULD** not advertise its loopback capability, as defined in Clause 57.2.11 of IEEE Std 802.3 [3], during the discovery phase if Loopback is not enabled.

9 Maximum Transmission Unit Size

The Maximum Transmission Unit Size specifies the maximum frame length in bytes allowed at an External Interface.

The MTU is part of several attribute specifications. For example, the UNI (defined in [5]), the EVC (defined in [5]), the ENNI (defined in 7.1.6), and the OVC (defined in 7.2.10) all have an MTU attribute.

In order for an EVC to operate properly with respect to its MTU Size, the Service Provider has to ensure that EVC MTU Size is less than or equal to the MTU Size of each OVC used to implement the EVC.

When provisioning a new ENNI or adding an EVC using an existing ENNI, the ENNI Operators and the Service Provider for the new EVC need to agree on MTU sizes which comply with the following requirements:

[D6] The ENNI MTU size **SHOULD** be greater or equal to the MTU size of :

- Each EVC MTU size crossing the ENNI plus 4 bytes (to accommodate the potential addition, at the ENNI, of an S-TAG), and
- Each OVC associating an OVC End Point at this ENNI.

[R124] The OVC MTU size **MUST** be greater or equal to the EVC MTU size of each EVC supported by this OVC plus 4 bytes (to accommodate the potential addition, at the ENNI, of an S-TAG).

10 Appendix A: Examples

This section presents several examples of the use of the Operator Services Attributes described in Section 7 to achieve UNI to UNI Ethernet Services. The first sub-section establishes the conventions and notation used in the examples. The remaining sub-sections present the examples.

10.1 Notation and Conventions

A number of abbreviations are used in the figures to avoid clutter. These are shown in Table 20.

Abbreviation	Object
C-VID	C-VLAN ID value
S-VID	S-VLAN ID value
OEP	OVC End Point Identifier value

Table 20 – Abbreviations Used in the Examples

In addition, the figures accompanying the examples use the icons as shown in Figure 18.

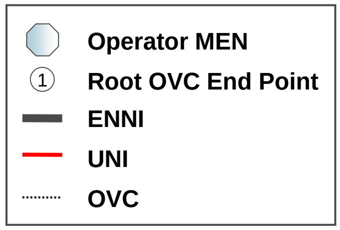

Figure 18 – Key to the Icons Used in the Examples

In the examples, the Service Provider is not explicitly shown. The Service Provider could be any of the Operators or some other organization. The examples are valid no matter who is the Service Provider.

10.2 Example 1: Ethernet Virtual Private Lines to a Hub Location

In this example, four Operator MENs are used by the Service Provider to provide three EVCs, each from a branch location to a hub location. Figure 19 shows the EVCs for this example as perceived by the Subscriber. UNI 1 is the hub location and the other UNIs are the branch locations. The CE-VLAN ID/EVC Maps as agreed to by the Subscriber and the Service Provider for each UNI are included in the figure. From these maps it can be seen that none of the EVCs have CE-VLAN ID Preservation in force as defined in MEF 10.2. [5]

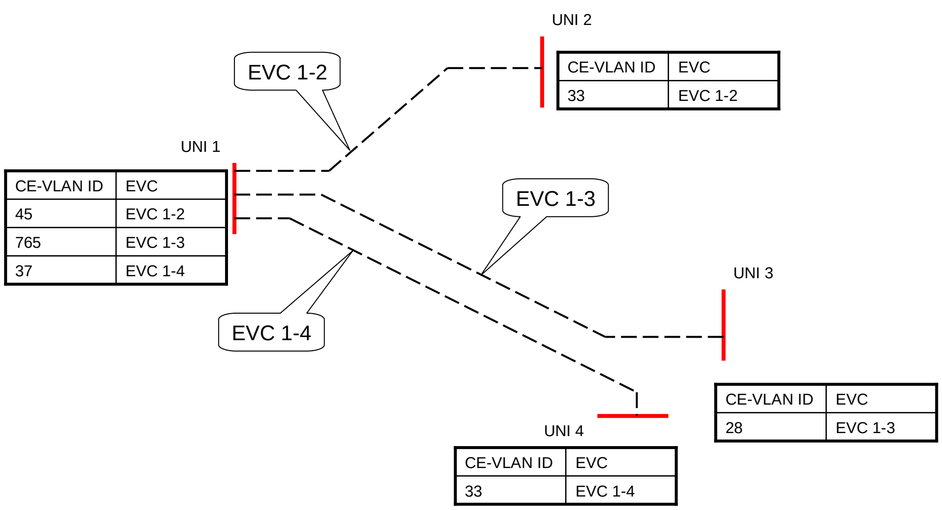

Figure 19 – EVCs to the Hub Location

Figure 20 shows Operator Services Attributes values that instantiate the EVPLs for this example using four Operator MENs. The various ENNI End Point Maps are shown in the figure. To see how this configuration works, consider a Service Frame that ingresses at UNI 1 and is destined for UNI 4 via EVC 1-4. Such a Service Frame will have a CE-VLAN ID value = 37. Operator A maps this frame to OVC End Point 11 and thus transmits a corresponding ENNI Frame with S-VLAN ID value = 1024 across the ENNI with Operator D. Operator D maps this ENNI Frame to

OVC End Point 13 and thus transmits a corresponding ENNI Frame with S-VLAN ID value = 2022 across the ENNI with Operator C. Operator C maps this ENNI Frame to OVC End Point 15 and thus transmits a corresponding Service Frame with CE-VLAN ID value = 33 across UNI 4. A similar sequence of events ensues for the other direction for EVC 1-4 and for the other EVCs in this example.

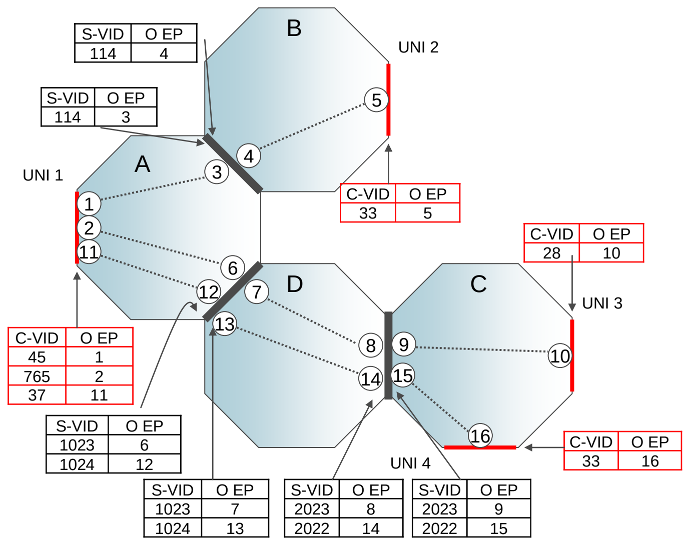

Figure 20 – Details of the Operator Services Attributes for Example 1

This example shows that EVC 1-4 is supported by three OVCs, one in each Operator MEN. The way that these OVCs are "connected" at each ENNI is via the End Point Maps on each side of the ENNI. By using S-VLAN ID value = 1024 to map to OVC End Point 12 in the Operator A MEN and also to map to OVC End Point 13 in the Operator D MEN, the appropriate OVCs in each Operator MEN are connected.

10.3 Example 2: Ethernet Private LAN

In this example, the Service Provider provides a single Ethernet Private LAN connecting four UNIs. Figure 21 shows the EVC for this example as perceived by the Subscriber. Note that EPLAN requires that the EVC have CE-VLAN ID Preservation and CE-CoS Preservation in force as per MEF 6.1. [9] Note also that the CE-VLAN ID/EVC Map at each UNI is All to One as prescribed by [9].

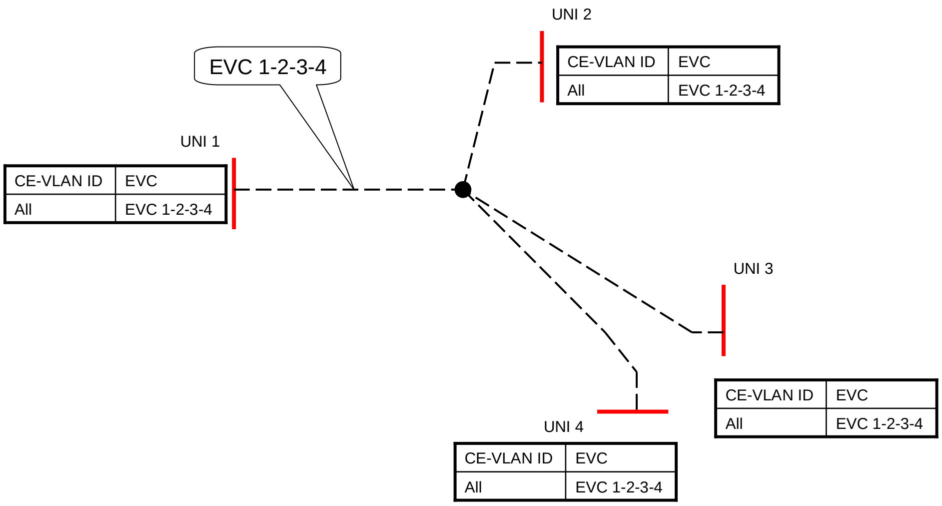

Figure 21 – EPLAN Connecting Four UNIs

Figure 22 shows Operator Services Attributes values that instantiate the EPLAN for this example using four Operator MENs. In order to support All to One Bundling, the OVC End Point Map at each UNI maps all Service Frames to the OVC End Point, e.g., OVC End Point 5 at UNI 2. Each OVC End Point Map at each ENNI maps to the OVC End Point with one S-VLAN ID value as was the case with Example 1.

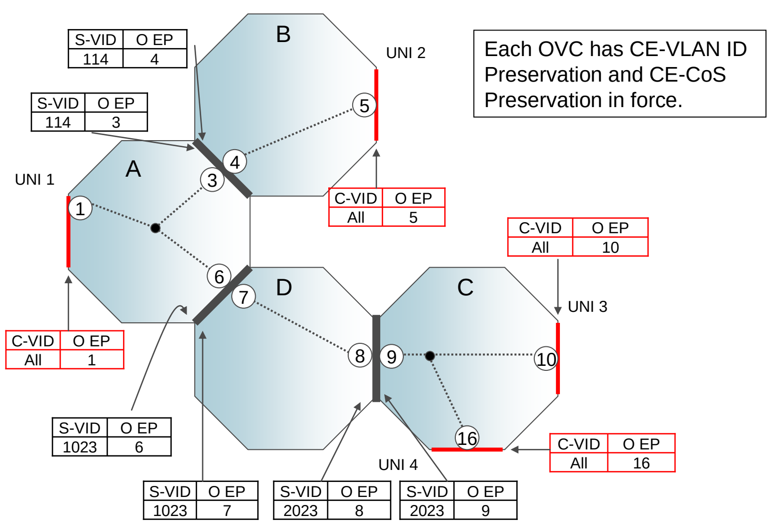

Figure 22 – Details of the Operator Services Attributes for Example 2

The OVC in the Operator MEN A has three OVC End Points as does the OVC in the Operator MEN C. Consequently, if MAC address learning is done in the Operator MEN A, a unicast Service Frame between UNI 1 and UNI 2 need not enter the Operator D and Operator C MENs.

Similarly, if MAC address learning is done in Operator MEN C, a unicast Service Frame between UNI 3 and UNI 4 need not leave the Operator C MEN.

Each OVC has CE-VLAN ID Preservation and CE-CoS Preservation in force as specified in Section 7.2.11 and Section 7.2.12. As an example to see how this works, consider an ingress Service Frame at UNI 1 that has a C-Tag and is destined for UNI 4.

- The resulting ENNI Frame at the ENNI between Operator MENs A and D will have both an S-Tag (with S-VLAN ID value = 1023) and a C-Tag and that C-Tag will have the same C-VLAN ID value and the same C-Tag PCP value as the original C-Tag of the Service Frame.
- The corresponding ENNI Frame at the ENNI between Operator MENs D and C will have both an S-Tag (with S-VLAN ID value = 2023) and a C-Tag and that C-Tag will have the same C-VLAN ID value and the same C-Tag PCP value as the original C-Tag of the ENNI Frame at the ENNI between Operator MENs A and D.
- Finally, the corresponding egress Service Frame at UNI 4 will have just a C-Tag and that C-Tag will have the same C-VLAN ID value and the same C-Tag PCP value as the original C-Tag of the ENNI Frame at the ENNI between Operator MENs C and D.

The result is that the C-VLAN ID values and C-Tag PCP values are the same for both ingress and egress Service Frame. Using the tables in Section 7.2.11, it can be seen that an untagged ingress Service Frame results in an untagged egress Service Frame. (Note that MEF 10.2 [5] specifies that CE-CoS Preservation does not apply to untagged Service Frames.) Consequently, the EVC has CE-VLAN ID Preservation and CE-CoS Preservation in force.

10.4 Example 3: Hairpin Switching

Figure 23 shows an example of the use of OVC End Points for Hairpin Switching. In this example, there is one multipoint EVC that associates UNI Aa, UNI Ab, and UNI B. Operator A has two OVCs. One associates the OVC End Point A4 at UNI Aa and the OVC End Point A1 at the ENNI. The other OVC associates the OVC End Point A3 at UNI Ab and the OVC End Point A2 at the ENNI. Operator B has one OVC that associates the OVC End Points B1 and B2 at the ENNI and the OVC End Point B3 at UNI B. At the ENNI, the End Point Maps, as described in Section 7.1.7, are such that ENNI frames mapped to A1 by Operator A are mapped to B1 by Operator B and similarly for A2 and B2. With this configuration, a Service Frame sent from UNI Aa to UNI Ab will pass through the Operator B MEN where it will be hairpin switched from B1 to B2. And a similar path will be followed by a Service Frame sent from UNI Ab to UNI Aa.

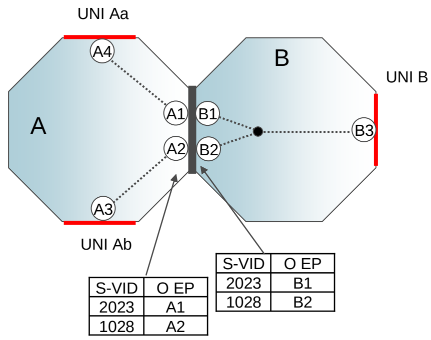

Figure 23 – Example of Hairpin Switching

Note that in this example, two OVCs are used in Operator MEN A to implement the single EVC.

10.5 Example 4: Data Loop at an ENNI with Hairpin Switching

The improper use of Hairpin Switching can lead to data loops at an ENNI. An example of such a improper configuration is shown in Figure 24. In this example, a broadcast frame will pass back and forth across the ENNI following the loop formed by the OVC End Points and Hairpin Switching.

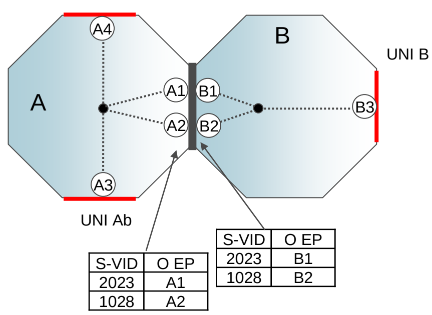

Figure 24 – Example of a Data Loop with Hairpin Switching

10.6 Example 5: Ethernet Private LAN with Hairpin Switching

In this example, the Service Provider provides a single Ethernet Private LAN connecting four UNIs just as was done in the example of Section 10.3. Figure 21 shows the EVC for this example as perceived by the Subscriber.

Figure 25 shows Operator Services Attributes values that instantiate the EPLAN for this example using four Operator MENs. As in the example of Section 10.3, the OVC End Point Map at each UNI maps all Service Frames to the OVC End Point and each OVC has CE-VLAN ID Preservation and CE-CoS Preservation in force as specified in Section 7.2.11 and Section 7.2.12.

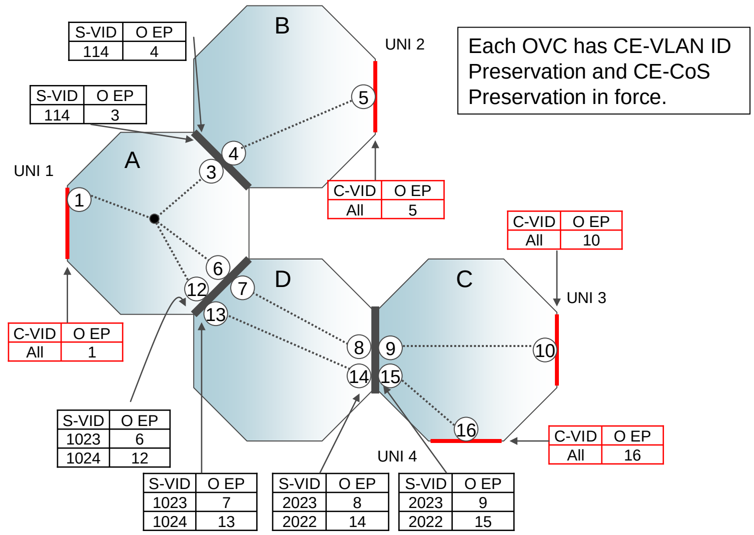

Figure 25 – Details of the Operator Services Attributes for Example 3

The key difference between this example and the example of Section 10.3 is the use of Hairpin Switching in Operator MEN A. The OVC in Operator MEN A has two OVC End Points, 6 and 12, at the ENNI between Operator MENs A and D. In addition, Operator MENs C and D each have two OVCs in support of the EPLAN. As a result, traffic from UNI 3 to UNI 4 will pass through Operator MENs A, C, and D and would follow the path of OVC End Points {10, 9, 8, 7, 6, 12, 13, 14, 15, 16}.

10.7 Example 6: End Point Map Bundling

Consider the EVCs to a hub location in Figure 19. Figure 26 shows the details of the Operator Services Attributes when Bundling is used in Operator MEN D.

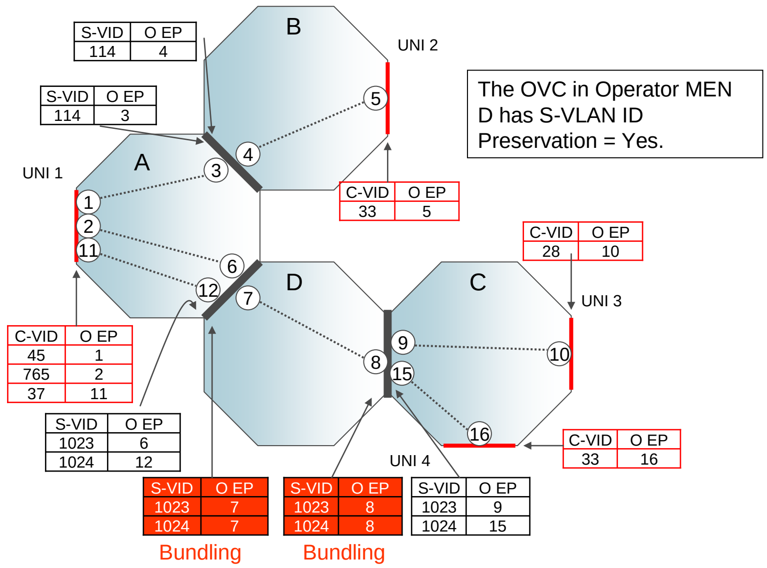

Figure 26 – Example of Using End Point Map Bundling

Figure 26 differs from Figure 20 as follows:

- There is only one OVC in Operator MEN D. Thus, this is a case where more than one EVC is supported by an OVC.
- The End Point Maps in the Operator MEN D have bundling with S-VLAN ID values 1023 and 1024 both mapped to a single OVC End Point at each ENNI.
- The OVC in Operator MEN D has S-VLAN ID Preservation = Yes.
- The End Point Map in Operator MEN C is changed to map each S-VLAN ID value 1023 and 1024 to an OVC End Point.

10.8 Example 7: CoS Identifiers at the ENNI

This example concerns the setting of the Class of Service Identifier in each ENNI Frame at an ENNI. Two scenarios are considered. The first scenario is when both Operator MENs conform to MEF 23.1 [10]. The second scenario is when MEF 23.1 is not conformed to by the Operator MENs.

In the first scenario, both Operator MENs abide by Table 4 in MEF 23.1 [10] which means that they would use the CoS Identifiers shown in Table 21. In this case, H in Operator MEN A would map to H in Operator MEN B and vice versa. In the same way, M would map to M and L would map to L. The result is that ENNI Frames would have the CoS Identifiers shown in Figure 27 where the arrows indicate the direction of flow at the ENNI.

CoS Label	CoS Identifier (S-Tag PCP Value)
H	5
M	3
L	1

Table 21 – MEF 23.1 CoS Identifiers

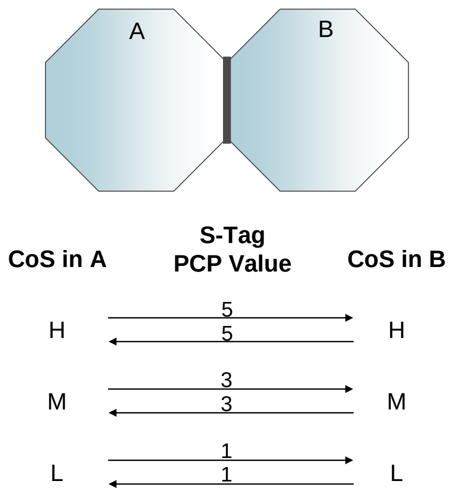

Figure 27 – CoS Identifiers for MEF 23.1 Conforming Operator MENs

In the second scenario, suppose that the Operator MENs had CoS Names and CoS Identifiers as shown in Table 22.

Operator MEN A		Operator MEN B	
CoS Name	CoS ID (S-Tag PCP Value)	CoS Name	CoS ID (S-Tag PCP Value)
Plus	6	Rock	6
Square	4	Paper	3
Heart	2	Scissors	1
Coal	1		

Table 22 – CoS Identifiers for Scenario Two

In this case, the mappings between the CoS instances in each Operator MEN will determine the use of CoS Identifiers at the ENNI. To see this, suppose that Square and Paper are mapped together. Then the ENNI Frames for Square and Paper would have the CoS Identifiers shown in Figure 28. Because the CoS Identifier is set according to the CoS of the receiving Operator MEN, the CoS Identifier in the ENNI Frame is different and depends on the direction of the frame flow.

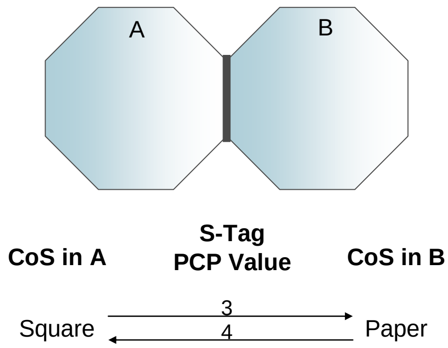

Figure 28 – CoS Identifiers for Square and Paper in Scenario Two

11 Appendix B: Rooted-Multipoint Examples

The forwarding behavior of a Rooted-Multipoint (RMP) EVC is specified in MEF10.2 [9] as any ingress Service Frames at a Root UNI may be forwarded to any other Root UNIs or Leaf UNIs of the same EVC, and any ingress Service Frames at a Leaf UNI may be forwarded to any Root UNIs of the same EVC. This forwarding behavior can be extended to a RMP EVC that spans one or more ENNIs by introducing a RMP OVC and OVC End Point Role designations of Root, Leaf, and Trunk.

The forwarding behavior of a Rooted-Multipoint EVC requires being able to determine whether any given frame originated as an ingress Service Frame at a Root UNI or as an ingress Service Frame at a Leaf UNI. The method of preserving this information for each frame within an Operator MEN depends upon the technology used within the MEN and is beyond the scope of this document. Preserving this information at an ENNI can require using two S-VLAN ID values: the Root S-VLAN ID value identifies frames that originated at a Root UNI, and the Leaf S-VLAN ID value identifies frames that originated at a Leaf UNI. When a Rooted-Multipoint OVC associates a Trunk OVC End Point at an ENNI, both S-VLAN IDs are mapped to the Trunk OVC End Point by the End Point Map.

The examples in this section illustrate the use of these concepts to instantiate a Rooted-Multipoint OVC across multiple Operator MENs.

11.1 Example Using Trunk OVC End Points

Figure 29 shows a Rooted-Multipoint connecting 6 UNIs as seen by the Subscriber.

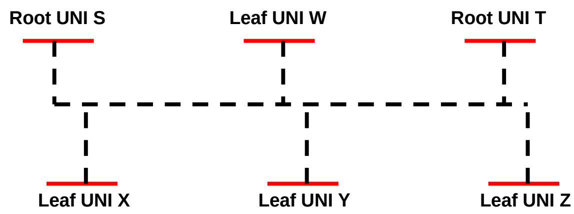

Figure 29 – Subscriber View of the Rooted-Multipoint EVC

Figure 30 shows an example of supporting this Rooted-Multipoint EVC using Trunk OVC End Points at the ENNIs that connect three Operator MENs. Note that each Operator MEN can receive ENNI Frames that originated at a Root UNI or a Leaf UNI. The use of Trunk OVC End Points and Trunk Identifiers at the ENNIs allows the Operator MEN to determine the type of ingress UNI and thus properly forward each ENNI Frame. Figure 31 shows example mappings using Root and Leaf S-VLAN ID values.

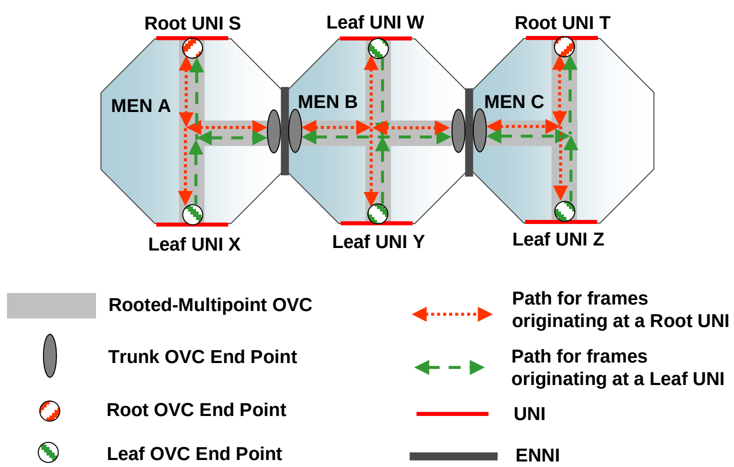

Figure 30 – Rooted-Multipoint EVC using Trunk OVC End Points

MEN A			MEN B		
OVC End Point Identifier	**Root S-VLAN ID value**	**Leaf S-VLAN ID value**	**OVC End Point Identifier**	**Root S-VLAN ID value**	**Leaf S-VLAN ID value**
PAIX A32	1065	1066	PAIX B12	1065	1066

MEN B			MEN C		
OVC End Point Identifier	**Root S-VLAN ID value**	**Leaf S-VLAN ID value**	**OVC End Point Identifier**	**Root S-VLAN ID value**	**Leaf S-VLAN ID value**
ChinaBasin B123	56	57	China Basin C19	56	57

Figure 31 – Example End Point Maps

11.2 Example Using a Rooted-Multipoint OVC in One Operator MEN

Figure 32 shows an example of supporting the Rooted-Multipoint EVC in Figure 29 using a Rooted-Multipoint OVC in just Operator MEN A along with Hairpin Switching in Operator

MEN A. The Hairpin Switching has the effect of including the UNIs in Operator MEN B and Operator MEN C in the Rooted-Multipoint EVC. Note that Operator B and Operator C see only Point-to-Point OVCs with Root OVC End Points at the UNIs and ENNIs. However, the Subscriber and Service Provider recognize the role of each UNI in Operator MEN B and Operator MEN C as either a Root UNI or a Leaf UNI.

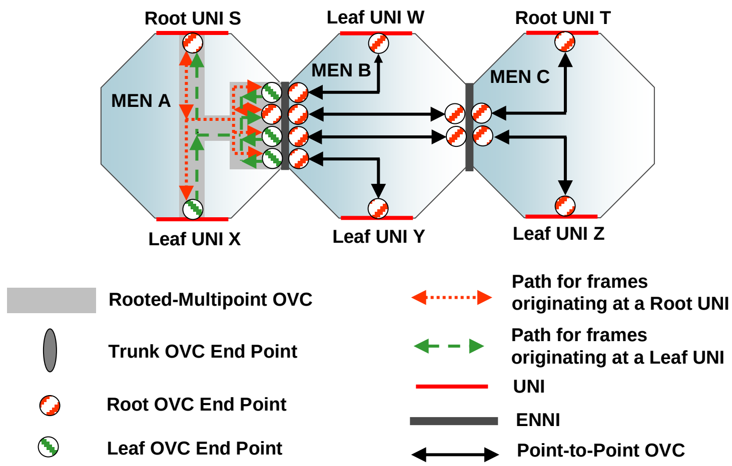

Figure 32 – Rooted-Multipoint EVC with a Rooted-Multipoint OVC in One Operator MEN

11.3 Example with All Root UNIs in One Operator MEN

Trunk OVC End Points are not always necessary to implement a Rooted-Multipoint EVC. Consider the Rooted-Multipoint EVC shown in Figure 33. In this example, Root UNI S and Root UNI T happen to be in Operator MEN A.

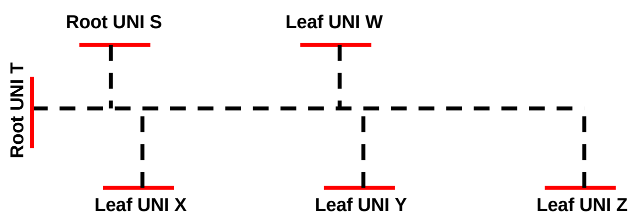

Figure 33 – Subscriber View of the Rooted-Multipoint EVC with all Root UNIs in one Operator MEN

Figure 34 shows an example of supporting the Rooted-Multipoint EVC of Figure 33 using just Root and Leaf OVC End Points. All Root UNIs are in the Operator MEN A, and the other Operator MENs have only Leaf UNIs. At each ENNI, all ENNI Frames going left to right originated at a Root UNI and all ENNI Frames going right to left originated at a Leaf UNI. This allows the use of Root and Leaf OVC End Points at the ENNIs, and a single S-VLAN ID value at the ENNIs. For example at ENNI AB, MEN A maps this single S-VLAN ID value to its Leaf OVC End Point while MEN B maps this same S-VLAN ID value to its Root OVC End Point.

In this figure, MEN A has a Leaf OVC End Point at ENNI AB because any ingress ENNI Frame mapped to this OVC End Point is the result of an ingress Service Frame at a Leaf UNI. For the same reason, MEN B has a Leaf OVC End Point at ENNI BC. MEN B has a Root OVC End Point at ENNI AB because any ingress ENNI Frame mapped to this OVC End Point is the result of an ingress Service Frame at a Root UNI. For the same reason, MEN C has Root OVC End Point at ENNI BC.

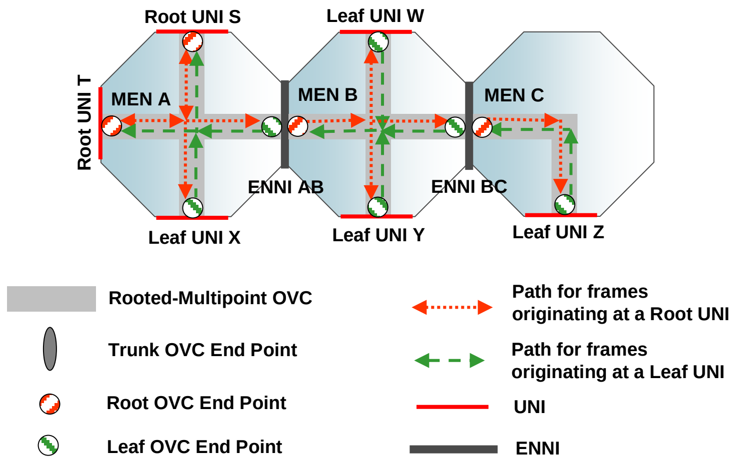

Figure 34 – Rooted-Multipoint EVC with all Root UNIs in one Operator MEN

11.4 Example Using a Bundled OVC

Figure 35 shows a Rooted-Multipoint EVC connecting 4 UNIs as seen by the Subscriber.

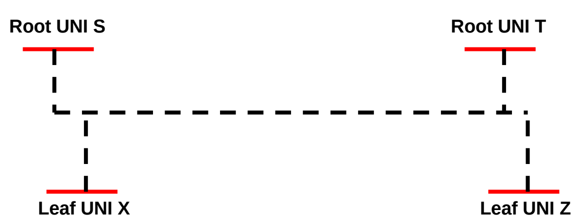

Figure 35 – Subscriber View of the Rooted-Multipoint EVC

Figure 36 shows an example of supporting this Rooted-Multipoint EVC using Trunk OVC End Points at the ENNIs in Operator MEN A and Operator MEN C. Operator MEN B uses a Point-to-Point Bundled OVC in this example. At each ENNI, MEN B maps the two Trunk Identifiers values to the OVC End Point. In this case, Operator MEN B need not be aware of which S-VLAN ID value represents frames originated at a Root UNI and which S-VLAN ID value represents frame originated at a Leaf UNI. Instead, MEN B simply passes the S-VLAN ID values unchanged between the two ENNIs as mandated by [R23]. MEN A and MEN C need to understand the significance of each Trunk Identifiers value.

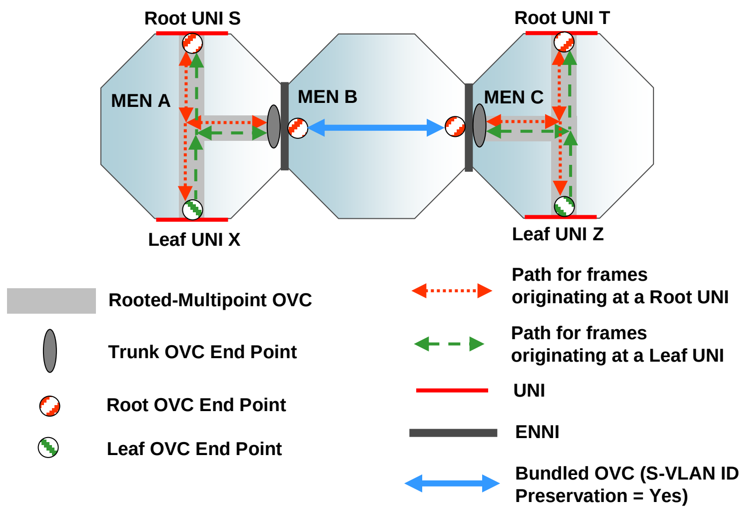

Figure 36 – Rooted-Multipoint EVC using a Bundled OVC

11.5 Example Using Hairpin Switching with Trunk OVC End Points

Figure 37 shows a Rooted-Multipoint EVC with 6 UNIs as seen by the Subscriber.

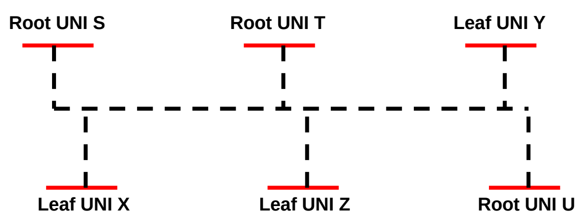

Figure 37 – Subscriber View of the Rooted-Multipoint EVC

Figure 38 shows and example of supporting this Rooted-Multipoint EVC using Hairpin Switching among the Trunk OVC End Points in MEN A. The configuration of this example would be useful if MEN B could only support Point-to-Point OVCs. The Hairpin Switching in MEN A provides the connectivity between the UNIs in MEN C and MEN D that MEN B is unable to provide.

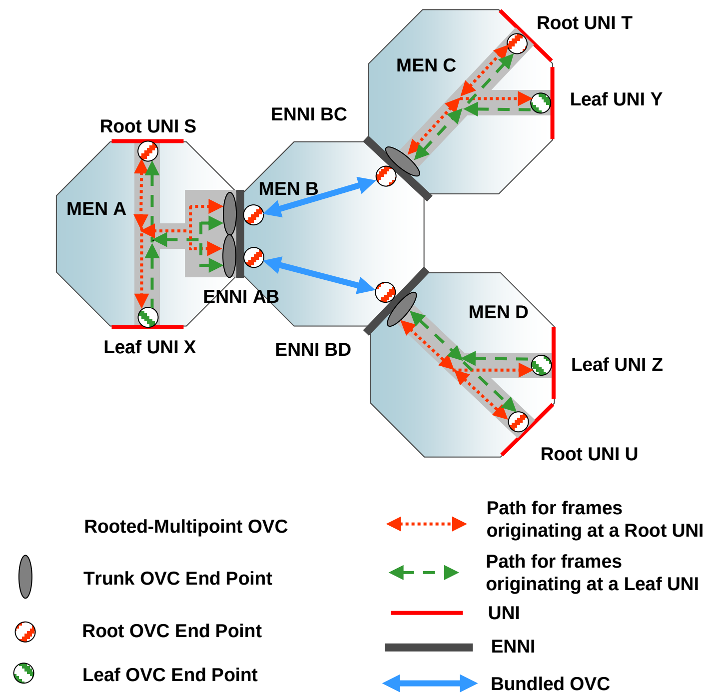

Figure 38 – Hairpin Switching with Trunk OVC End Points

12 References

[1] MEF Forum, MEF 4, *Metro Ethernet Network Architecture Framework – Part 1: Generic Framework*, May 2004.

[2] S. Bradner, RFC 2119, *Key words for use in RFCs to Indicate Requirement Levels*, March 1997.

[3] IEEE Std 802.3 – 2005, *Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications*, 9 December 2005.

[4] IEEE Std 802.1ad – 2005, *IEEE Standard for Local and metropolitan Area Networks – Virtual Bridged Local Area Networks Amendment 4: Provider Bridges*, May 2006.

[5] MEF Forum, MEF 10.2, *Ethernet Services Attributes Phase 2*, October 2009.

[6] International Telecommunication Union, Recommendation Y.1731, *OAM functions and mechanisms for Ethernet based Networks*, February 2008.

[7] MEF Forum, MEF 11, *User Network Interface (UNI) Requirements and Framework*, November 2004.

[8] MEF Forum, MEF 20, *User Network Interface (UNI) Type 2 Implementation Agreement*, July 2008.

[9] MEF Forum, MEF 6.1, *Ethernet Services Definitions - Phase 2*, April 2008.

[10] MEF Forum, MEF 23.1, *Carrier Ethernet Class of Service – Phase 2*, January 2012.

[11] MEF Forum, MEF 17, MEF Forum, *Service OAM Requirements & Framework – Phase 1*, April 2007.

[12] IEEE Std 802.1ag – 2007, *IEEE Standard for Local and metropolitan Area Networks – Virtual Bridged Local Area Networks Amendment 5: Connectivity Fault Management*, December 2007.

[13] MEF Forum, MEF 2, MEF Forum, *Requirements and Framework for Ethernet Service Protection in Metro Ethernet Networks*, February 2004.

[14] K. Nichols, S. Blake, F. Baker, and D. Black, RFC2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*, December 1998.

[15] IEEE Std 802.1ah – 2008, *IEEE Standard for Local and metropolitan Area Networks – Virtual Bridged Local Area Networks Amendment 7: Provider Backbone Bridges*, August 2008.

[16] MEF Forum, MEF 26, *External Network Network Interface (ENNI) – Phase 1*, January 2010.

[17] MEF Forum, MEF 26.0.1, *Corrected Figure 10 for MEF 26*, July 2010.

[18] MEF Forum, MEF 26.0.2, *OVC Layer 2 Control Protocol Tunneling Amendment to MEF 26*, July 2011.

[19] MEF Forum, MEF 26.0.3, *Service Level Specification Amendment to MEF 26*, October 2011.

[20] MEF Forum, MEF 28, *External Network Network Interface (ENNI) Support for UNI Tunnel Access and Virtual UNI*, October 2010.

[21] IEEE Std 802.1Q – 2011, *IEEE Standard for Local and metropolitan area networks--Media Access Control (MAC) Bridges and Virtual Bridged Local Area Networks*, 2011.

[22] MEF Forum, MEF 10.2.1, *Performance Attributes Amendment to MEF 10.2*, January 2011.

[23] International Telecommunication Union, Recommendation Y.1563, *Global information infrastructure, internet protocol aspects and next-generation networks, Internet protocol aspects – Quality of service and network performance, Ethernet frame transfer and availability performance*, 2009.