Technical Specification

MEF 29

Ethernet Services Constructs

January 2010

Table of Contents

List of Figures

List of Tables

1 Abstract

There are many ways in which Service Provider technology can be deployed to support the UNI-to-UNI Ethernet Services Attributes specified in MEF 10.2 [10]. The network between these UNIs can be broken up into smaller Ethernet Sub-Networks with defined interfaces between them. This document defines an abstract set of Ethernet Service Constructs that are widely applicable regardless of how the UNI-to-UNI Network is broken down. Follow-on MEF specifications that wish to define specific types of Ethernet Sub-Networks or interfaces can do so by referring to and building upon this specification.

2 Terminology

Term	Definition	Source
All to One Bundling	A UNI attribute in which all CE-VLAN IDs are associated with a single EVC	MEF 10.2 [10]
Bandwidth Profile	A characterization of EI Frame arrival times and lengths at a reference point. In this document the reference point is the EI.	This document
Bundling	A UNI attribute in which more than one CE-VLAN ID can be associated with an EVC	MEF 10.2 [10]
CBS	Committed Burst Size	This document
CE	Customer Edge	MEF 10.2 [10]
CE-VLAN CoS	Customer Edge VLAN CoS	MEF 10.2 [10]
CE-VLAN ID	Customer Edge VLAN ID	MEF 10.2 [10]
CE-VLAN ID Preservation	An EVC attribute in which the CE-VLAN ID of an egress Service Frame is identical in value to the CE-VLAN ID of the corresponding ingress Service Frame	MEF 10.2 [10]
CE-VLAN Tag	Customer Edge VLAN Tag	MEF 10.2 [10]
CF	Coupling Flag	MEF 10.2 [10]
Child Document	A MEF Specification that uses the Ethernet Service Constructs model (defined in this specification) to specify a particular type of EI or ESN	This document
CIR	Committed Information Rate	This document
Class of Service	A set of Service Frames that have a commitment from the Service Provider to receive a particular level of performance	MEF 10.2 [10]

Term	Definition	Source
Class of Service Identifier	Information derivable from a) the EVC to which the Service Frame is mapped, b) the combination of the EVC to which the Service Frame is mapped and a set of one or more CE-VLAN CoS values, c) the combination of the EVC to which the Service Frame is mapped and a set of one or more DSCP values, or d) the combination of the EVC to which the Service Frame is mapped and a set of one or more tunneled Layer 2 Control Protocols	MEF 10.2 [10]
CM	Color Mode	This document
Color Mode	CM is a Bandwidth Profile parameter. The Color Mode parameter indicates whether the color-aware or color-blind property is employed by the Bandwidth Profile. It takes a value of "color-blind" or "color-aware" only	This document
Color-aware	A Bandwidth Profile property where a pre-determined level of Bandwidth Profile compliance for each EI Frame is taken into account when determining the level of compliance for each Service Frame	This document
Color-blind	A Bandwidth Profile property where a pre-determined level of Bandwidth Profile compliance for each EI Frame, if present, is ignored when determining the level of compliance for each EI Frame	This document
Committed Burst Size	CBS is a Bandwidth Profile parameter. It limits the maximum number of bytes available for a burst of EI Frames sent at the EI speed to remain CIR-conformant	This document
Committed Information Rate	CIR is a Bandwidth Profile parameter. It defines the average rate in bits/s of EI Frames up to which the network delivers Service Frames and meets the performance objectives defined by the Class of Service	This document
Connected Sub-Network	A Sub-Network that is connected by physical External Interfaces to the Ethernet Sub-Network under examination (see Figure 4)	This document
Coupling Flag	CF is a Bandwidth Profile parameter. The Coupling Flag allows the choice between two modes of operations of the rate enforcement algorithm. It takes a value of 0 or 1 only	This document

Term	Definition	Source
CSN	Connected Sub-Network	This document
Customer Edge	Equipment on the Subscriber side of the UNI	MEF 10.2 [10]
Customer Edge VLAN CoS	The user priority bits in the IEEE 802.1Q Tag in a Service Frame that is either tagged or priority tagged	MEF 10.2 [10]
Customer Edge VLAN ID	The identifier derivable from the content of a Service Frame that allows the Service Frame to be associated with an EVC at the UNI	MEF 10.2 [10]
Customer Edge VLAN Tag	The IEEE 802.1Q Tag in a tagged Service Frame	MEF 10.2 [10]
DSCP	Diff-Serve Code Point	RFC 2474 [13]
EBS	Excess Burst Size	This document
Egress Bandwidth Profile	An attribute that specifies the length and arrival time characteristics of egress EI Frames at the egress EI	This document
Egress EI Frame	An Ethernet Frame transmitted across the EI from the ESN toward the CSN	This document
Egress Service Frame	A Service Frame sent from the Service Provider network to the CE	MEF 10.2 [10]
EI	External Interface	This document
EI Frame	An Ethernet frame transmitted across the EI toward the ESN or an Ethernet frame transmitted across the EI toward the CSN	This document
EIR	Excess Information Rate	This document
ESC	Ethernet Service Construct	This document
ESN	Ethernet Sub-Network	This document
ESNC	Ethernet Sub-Network Connection	This document
ESNC End Point	A logical reference point located at a physical External Interface (EI)	This document

Term	Definition	Source
Ethernet Service Construct	An abstract entity used to facilitate describing the externally visible behavior of an Ethernet Sub-Network (ESN) as seen by looking into the physical demarcation points of this ESN. The four basic types of Ethernet Service Constructs introduced in this document are EIs, ESNCs, ESNC End Points, and TEs. Each of these entities represents a collection of defined attributes that are used to describe this behavior	This document
Ethernet Sub-Network	The Sub-Network that is under examination by the Ethernet Service Constructs model (see Figure 4)	This document
Ethernet Sub-Network Connection	An association of two or more ESNC End Points that limits the exchange of EI frames to EIs in the ESNC	This document
Ethernet Virtual Connection	An association of two or more UNIs that limits the exchange of Service Frames to UNIs in the Ethernet Virtual Connection	MEF 10.2 [10]
EVC	Ethernet Virtual Connection	MEF 10.2 [10]
Excess Burst Size	EBS is a Bandwidth Profile parameter. It limits the maximum number of bytes available for a burst of Service Frames sent at the UNI speed to remain EIR-conformant	This document
Excess Information Rate	EIR is a Bandwidth Profile parameter. It defines the average rate in bits/s of Service Frames up to which the network may deliver Service Frames without any performance objectives	This document
External Interface	A physical point of demarcation between the Ethernet Sub-Network under examination and the Connected Sub-Network	This document
Frame	Short for Ethernet frame	MEF 10.2 [10]
Ingress Bandwidth Profile	An attribute that specifies the length and arrival time characteristics of ingress EI Frames at the ingress EI	This document
Ingress EI Frame	An Ethernet frame transmitted across the EI toward the ESN	This document
Ingress Service Frame	A Service Frame sent from the CE into the Service Provider network	MEF 10.2 [10]
LAG	Link Aggregation Group	[15]

Term	Definition	Source
Layer 2 Control Protocol Service Frame	A Service Frame that is used for Layer 2 control, e.g., Spanning Tree Protocol	MEF 10.2 [10]
Layer 2 Control Protocol Tunneling	The process by which a Layer 2 Control Protocol Service Frame is passed through the Service Provider network without being processed and is delivered unchanged to the proper UNI(s)	MEF 10.2 [10]
Link Protection Mechanism	Any mechanism (e.g. LAG) used to protect traffic in the event of link failure across a multi-link EI	This document
Link Protection Mechanism	An attribute of an EI that specifies the mechanism used to protect traffic across an EI with multiple links	This document
LPM	Link Protection Mechanism	This document
L2CP	Layer 2 Control Protocol	MEF 10.2 [10]
Multipoint-to-Multipoint EVC	An EVC with two or more UNIs. A Multipoint-to-Multipoint EVC with two UNIs is different from a Point-to-Point EVC because one or more additional UNIs can be added to it	MEF 10.2 [10]
Operator	The administrative organization responsible for a particular ESN	This document
Point-to-Point EVC	An EVC with exactly 2 UNIs	MEF 10.2 [10]
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 [10]
Service Multiplexing	A UNI service attribute in which the UNI can be in more than one EVC instance	MEF 10.2 [10]
Service Provider	The organization providing Ethernet Service(s)	MEF 10.2 [10]
SNI	Services Node Interface	MEF 4 [8]
Subscriber	The organization purchasing and/or using Ethernet Services	MEF 10.2 [10]
TE	Tunnel End Point	This document
Tunnel End Point	An ESC that represents a collection of ESNC End Points at an EI. It is used to facilitate the description of the externally observable behavior of the case where multiple ESNCs are encapsulated into a single service stream within an ESN	This document
UNI	User Network Interface	MEF 10.2 [10]

Term	Definition	Source
Unicast Service Frame	A Service Frame that has a unicast destination MAC address	MEF 10.2 [10]
User Network Interface	The physical demarcation point between the responsibility of the Service Provider and the responsibility of the Subscriber	MEF 10.2 [10]

3 Scope

This document describes Ethernet Service Constructs. It is not intended to be a standalone specification. The intent is to define a base set of Ethernet Service Constructs that allow multiple follow-on MEF specifications to more easily and consistently specify the externally observable behavior of specific Ethernet Sub-Networks and/or specific External Interfaces within the UNI-to-UNI Ethernet network.

Follow-on MEF specifications that refer to and build on the Ethernet Service Constructs defined in this document are referred to as “Child Documents”. The goal of this specification is to define abstract Ethernet Service Constructs that will be widely applicable to multiple Child Documents, reducing the workload of describing specific types of Ethernet Sub-Networks or External Interfaces and improving the consistency between them.

An Ethernet Service Construct (ESC) is defined to be an abstract entity used to facilitate describing the externally visible behavior of an Ethernet Sub-Network (ESN) as seen by looking into the physical demarcation points of the ESN. The four basic types of Ethernet Service Constructs introduced in this document are External Interfaces (EIs), Ethernet Sub-Network Connections (ESNCs), ESNC End Points, and Tunnel End Points (TEs). Each of these entities represents a collection of defined attributes that are used to describe this behavior.

Ethernet Service Constructs are distinct from the functional logical entities used to describe network architecture in other MEF specifications such as MEF 12 [14]. These architectural models are intended to facilitate designing and building carrier Ethernet networks and network elements as opposed to describing the externally visible behavior of Ethernet services. While these entities can be related (e.g., The “Ethernet Subscriber Conditioning Function (ESCF)” in MEF 12 [14] is related to the “Ingress Bandwidth Profile” of an “ESNC” in this specification), describing these relationships is beyond the scope of this specification.

It is not the intention of this document to describe all constructs (or their attributes) that are universally applicable. The intention is to provide a base on which Child Documents can build, either by constraining the constructs and attributes used in this specification or by extending them. In the course of time, as more Child Documents build on this base, further versions of this specification will likely be developed to describe common constructs or attributes that emerge from those documents.

4 Compliance Levels

The key words "**MUST**", "**MUST NOT**", "**REQUIRED**", "**SHALL**", "**SHALL NOT**", "**SHOULD**", "**SHOULD NOT**", "**RECOMMENDED**", "**MAY**", and "**OPTIONAL**" in this document are to be interpreted as described in RFC 2119 [1]. All key words must be in upper case, bold text.

5 Introduction

There are many ways in which Service Provider technology can be deployed to support the UNI-to-UNI Ethernet Services Attributes specified in MEF 10.2 [10]. In many cases, business and/or technical reasons lead to the deployment of various networks based on Ethernet technology interconnected in such a way that Service Frames can be exchanged among the UNIs in question.

Figure 1 shows an example of three such networks interconnected to implement Ethernet Services among six UNIs. An Ethernet Sub-Network (ESN) is defined to be one or more interconnected physical network elements under the control of a particular administrative domain. The administrative authority for a particular ESN is known as an Operator.

As shown in the example, there are two interfaces used to connect ESNs together, i.e., Interface A/B and Interface B/C. The nature of these interfaces can vary greatly depending on the technology, scale, and ownership of the networks being connected. There are a large number of permutations.

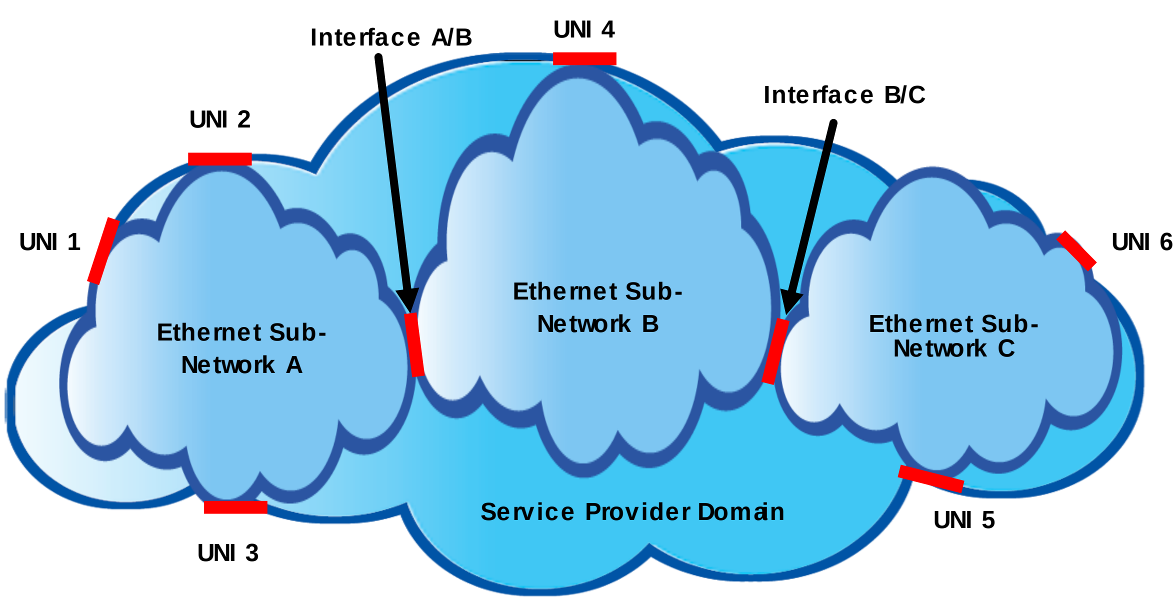

Figure 1 – Example of Multiple Ethernet Sub-Networks

Even though there are many ways of interconnecting Ethernet Sub-Networks, there are some characteristics that are in force in all cases:

- The boundaries of any particular Ethernet Sub-Network are physical demarcation points, across which Ethernet Frames traverse. These interfaces are denoted as External Interfaces (EIs) and are shown as thick red lines in Figure 1. Each ESN must allow the exchange of Ethernet frames between its External Interfaces in an appropriate way to support its part of the service being provided. In Figure 1, for example, “UNI 5”, “UNI 6” and “Interface B/C” are all “External Interfaces” of Ethernet Sub-Network C.

- Each Ethernet Sub-Network must provide connectivity among the External Interfaces. An instance of this connectivity is denoted as an “Ethernet Sub-Network Connection” (ESNC). The Service Provider must ensure that an ESNC has the appropriate attributes so that the end-to-end service being provided (of which this ESNC is a part) has the Ethernet features as agreed upon by the Service Provider and the Subscriber.

These two observations allow the formulation of the basic Ethernet Service Construct (ESC) model as shown in Figure 2.

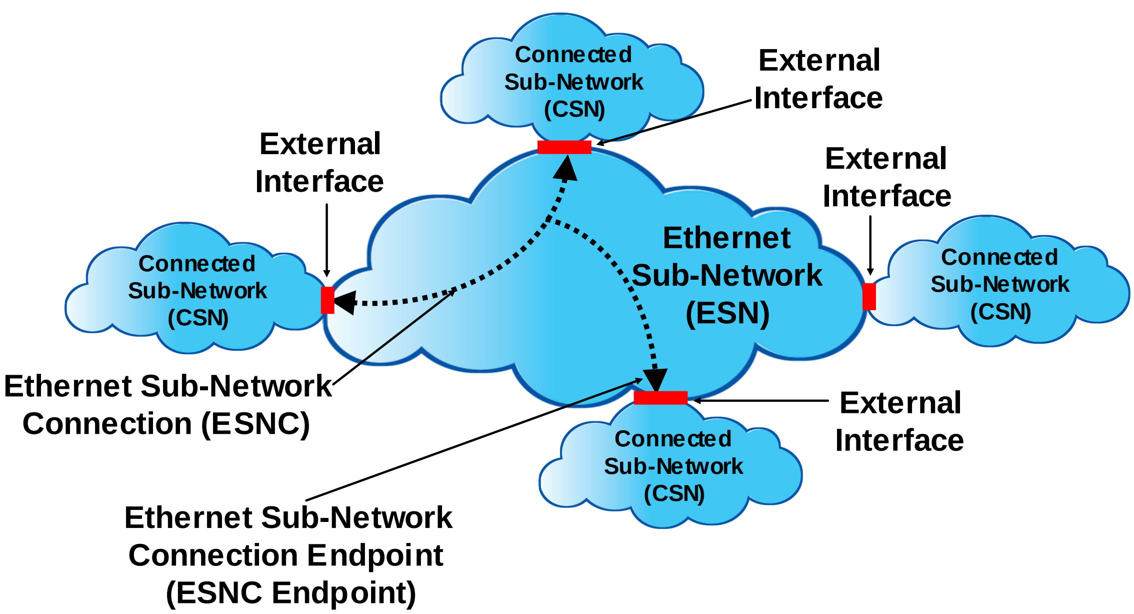

Figure 2 – Basic Ethernet Service Constructs Model

An External Interface (EI) and an Ethernet Sub-Network Connection (ESNC) are said to be “Ethernet Service Constructs” associated with a particular Ethernet Sub-Network (ESN). The Ethernet Service Constructs (ESC) Model described in this document specifies the attributes and behaviors of these constructs that apply generically to a number of different types of EI (e.g., UNIs, **ENNI**s, SNIs, etc.) across a number of different types of ESN (e.g., an Access Network, an Operator Network, etc.). A Connected Sub-Network (CSN) could be any type of network. For example, it could be a Subscriber's CE (making this EI a UNI) or it could be another independent sub-network. (Note that in the simplest case, the ESN can be one single network element, and the ESC model can be used to describe the external behavior of this network element without specifying how it is implemented internally.)

This model is sufficient for "flat" implementations, where end-to-end EVCs are composed of a set of ESNCs concatenated together with a one-to-one association of ESNCs at each EI. There are many cases, however, where EVCs are implemented using more complex topologies. For example, Figure 3 shows the case of two point-to-point EVCs, one from UNI 2 to UNI 5, the other from UNI 1 to UNI 6. Inside the Service Provider network, however, both of these EVCs are "tunneled" through a single service component through ESN B. In general, the intent of such tunnels is that the Operator of ESN B need not be aware of the EVCs inside the service being provided through ESN B.

The behavior of ESN B (in Figure 3) can be readily described using the basic Ethernet Service Constructs model shown in Figure 2. ESN B is the ESN under examination. ESN A and ESN C are two CSNs connected to ESN B. Interface A/B and Interface B/C are EIs representing the physical demarcation between the sub-networks. The line marked "B1" in Figure 3 is an ESNC with two ESNC End Points, one at Interface A/B, the other at Interface B/C.

Much of the behavior of ESN A (in Figure 3) can also be described using the basic Ethernet Service Constructs model shown in Figure 2. In this case, ESN A is the ESN under examination. The CEs connected to UNI 1 and UNI 2 as well as ESN B are the CSNs. UNI 1, UNI 2 and Interface A/B are the EIs. A1 and A2 are ESNCs, both of which have an ESNC End Point at Interface A/B.

But it is clear that ESNC A1 and A2 are not at the same level as ESNC B1. They are all ESNCs since ESNCs are always at the level of the service being provided by the Operator of the ESN under examination. But when looking at ESN A and ESN B together, ESNC B1 is a "tunnel" through which the EVCs associated with A1 and A2 are carried. In the case shown, the End Point of this "tunnel" is contained in ESN A. ESN A is performing the functions associated with this tunnel End Point such as frame encapsulation and de-capsulation and processing of service OAM at the tunnel level when such frames are sent from ESN B into ESN A. Attempting to describe this externally observable behavior of ESN A using only the objects in Figure 2 would be awkward.

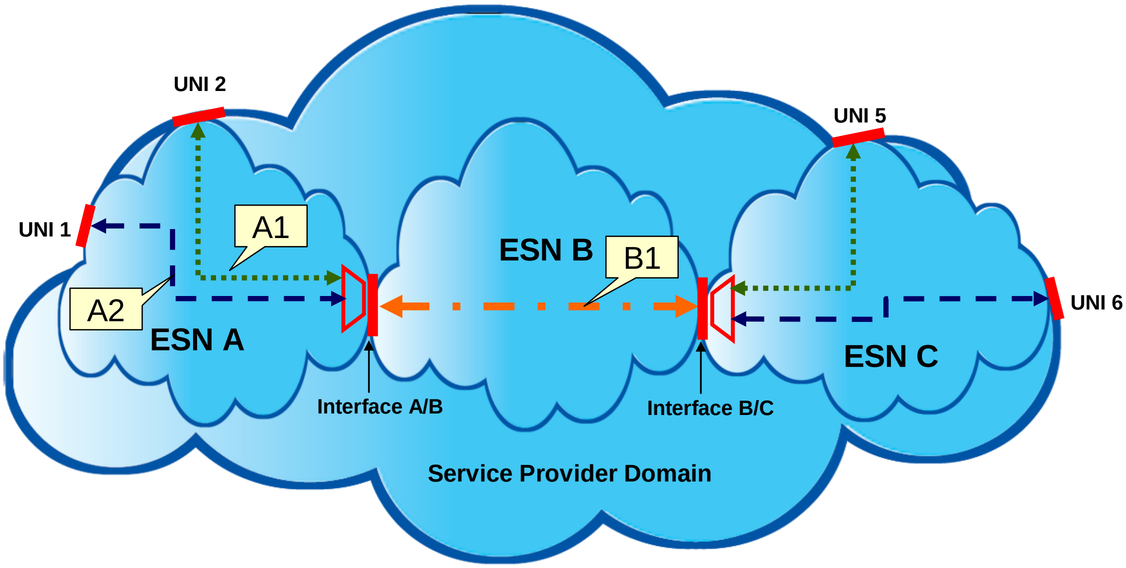

Figure 3 – Example of more complex construction of EVCs

A more complete model, which allows for hierarchy is shown in Figure 4. The model extends that shown in Figure 2 by adding a new Ethernet Service Construct called a Tunnel End Point (TE). A TE is an ESC that represents a collection of ESNC End Points at an EI. It is used to facilitate the description of the externally observable behavior of the case where multiple ESNCs are encapsulated into a single service stream within an ESN.

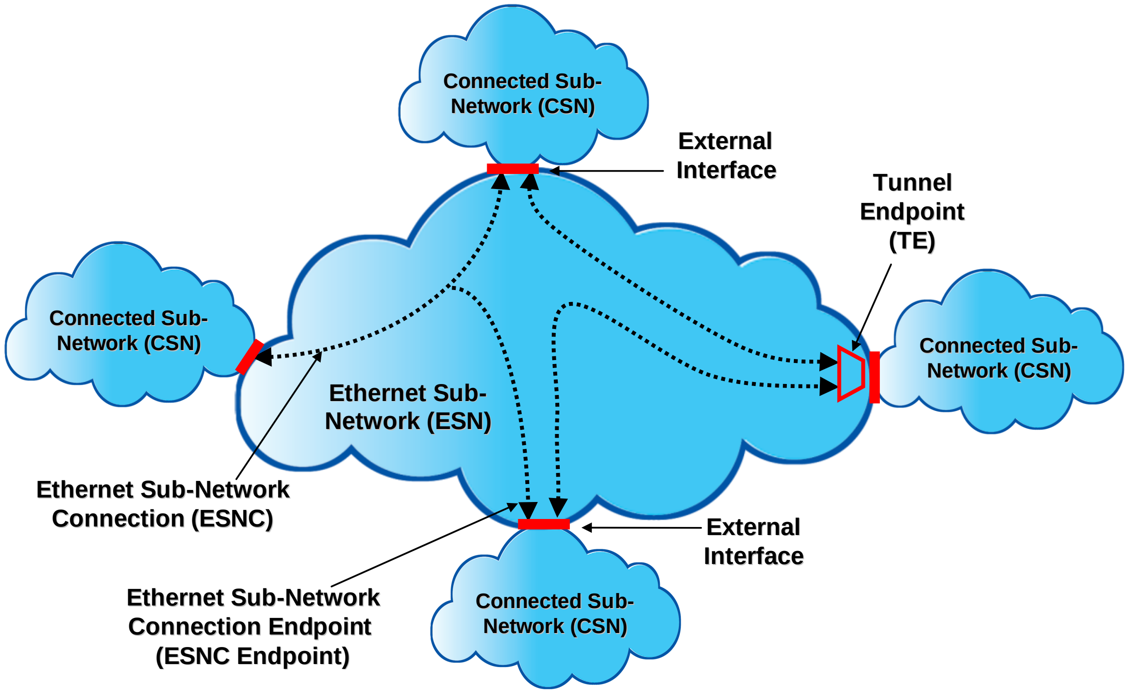

Figure 4 – Ethernet Service Construct Model

Just as there can be multiple ESNC End Points at an EI, there can also be multiple TEs at an EI (not shown in the Figure 4). Note that in Figure 4, the ESNC End Points are depicted as being on the TE as opposed to the EI. This is for graphical simplicity. To an external observer, the ESNC End Points are at the EI.

The model is expected to be used as a reference point for follow-on MEF specifications. For example:

- Subsequent documents specifying the implementation of specific External Interfaces will, where appropriate, build on these basic service constructs rather than starting from the beginning. This reduces the documentation workload and improves the consistency of terminology and behavior across different types of EI specifications.

- Subsequent documents specifying the implementation of specific service types (e.g., A Standardized Transit Tunnel service) will, where appropriate, build on these basic service constructs rather than starting from the beginning. This also reduces the documentation workload and improves the consistency of terminology and behavior of services specifications across different types of ESN.

The ESC model (Figure 4) provides a technical definition of the attributes of the portion of an Ethernet service provided by this particular ESN in terms of what is seen by devices on the other side of each External Interface (i.e., the Connected Sub-Networks). This includes the External

Interface (EI) which is the physical demarcation point between the responsibility of the Operator of this ESN and the Operator of the connected CSN.

Connectivity between EIs is specified by the Ethernet Sub-Network Connection (ESNC).

The CSNs and the ESN exchange EI Frames across the EI. An EI Frame is an Ethernet frame transmitted across the EI toward the ESN (called an Ingress EI Frame) or an Ethernet frame transmitted across the EI toward the CSN (called an Egress EI Frame). The EI Frame consists of the first bit of the destination MAC address through the last bit of the Frame Check Sequence.

EI attributes are described in Section 6.

TE attributes are described in Section 7.

ESNC End Point attributes are described in Section 8.

ESNC attributes are described in Section 9.

Section 10 defines the bandwidth profile algorithm that is used both in EI attributes and ESNC End Point attributes.

Section 11 provides an informative description of the different ESNC End Point Types and the rationale behind them.

Section 12 provides an informative description of the generalized concept of CE VLAN preservation.

Section 13 provides an informative description of the concept of “tunneling”, the rationale for the model chosen, and the power of this approach to describe a broad range of behaviors.

6 External Interface (EI) Attributes

There are attributes for each EI in the ESN. These attributes are summarized in Table 1 and described in detail in the sub-sections that follow this table.

Attribute Name	Summary Description
Operator EI Identifier	An arbitrary string that uniquely identifies the EI within the Ethernet Sub-Network
Physical Layer	A specification of the physical layer over which EI Frames traverse
Physical-to-Frame Format Mapping	A specification of how the EI Frame can be derived from the bit stream on this physical link.
Link Protection Mechanism	Describes the method and nature of protection against a link failure in the EI. It can be LAG <LACP-ENABLED <periodic-timer-value> | LACP-DISABLED>. Child documents can define other types of Link Protection Mechanisms.
Frame Format	The format of the Ethernet Frames at the EI
EI Maximum Transmission Unit Size	The maximum length EI Frame in bytes allowed at the EI
Maximum Number of ESNCs	The maximum number of ESNCs that the ESN can support at the EI

Attribute Name	Summary Description
Maximum Number of ESNC End Points	The maximum total number of ESNC End Points that the ESN can support at the EI.
Maximum Number of ESNC End Points per ESNC	The maximum number of ESNC End Points associated with one ESNC that the ESN can support at the EI.
Maximum number of Tunnel End Points	The maximum number of Tunnel End Points that the ESN can support at the EI
Ingress Bandwidth Profile per EI	The Ingress Bandwidth Profile enforced by the ESN on all ingress EI Frames
Egress Bandwidth Profile per EI	The Egress Bandwidth Profile enforced by the ESN on all egress EI Frames
EI Frame Map	The way that EI Frames are mapped to either ESNC End Points or TEs

Table 1 – EI Attributes

6.1 Operator EI Identifier

The Operator EI Identifier is an arbitrary string assigned to the EI by the Operator of the Ethernet Sub-Network (The entity that is administratively responsible to manage this ESN).

[R1] The EI Identifier **MUST** be no more than 45 bytes in length.

[R2] The EI Identifier **MUST** be unique among all EIs for the ESN.

A child document could place additional requirements on this identifier. For example, the identifier could be mandated to be a Character String as defined in Table 21-19 of IEEE 802.1ag – 2007 [17].

As an example, the Operator might use “SCPOP1-Node3-Interface to ACME" as an EI Identifier and this could signify an EI handoff between Node 3 in Santa Clara POP1 (part of this ESN) and ACME Service Provider (the connected network).

6.2 Physical Layer

The Physical Layer is a list of physical (bidirectional) links that make up the EI along with the attributes of each link. Links that are associated together to form one EI are said to be “in the EI”.

[R3] The Physical Layer is a list of <Link Identifier, Link Type, Speed, Mode, Physical Medium> associations. The list **MUST** have exactly one such association for each link that is in the EI.

The Link Identifier is an arbitrary string assigned to the physical link by the Operator of the Ethernet Sub-Network.

[R4] The Link Identifier **MUST** be no more than 45 bytes in length.

[R5] The Link Identifier **MUST** be unique for all the links in this EI.

[R6] For every EI, agreement **MUST** be reached between the Operator of the ESN and the Operator of the CSN connected by this EI on the EI Identifier and the Link Identifier(s) to allow network management across this demarcation. Agreement can mean the use of the same identifiers by both parties, or it could be agreement on how to map between different identifiers used by the two Operators.

[O1] The Link Type **MAY** be ETHERNET.

Other types of physical links can be added by Child Documents. For example, if a Child Document requires a DS3 as a type of physical interface, it can do so, specifying the requirements around the use of that type of physical interface.

[R7] If the Link Type is ETHERNET, then for each link in an EI, the Speed (in bits per second), Mode, and Physical Medium **MUST** be one of the combinations shown in Table 2.

Typically there are no constraints in mixing EIs with different physical media in the same ESNC.[1]

[O2] Constraints on the mix of speeds and modes **MAY** be imposed for some services.

Speed	Mode	Physical Medium
10 Mbps	Full duplex	All Ethernet physical media compatible with Speed and Mode specified in IEEE 802.3 – 2005 [2], IEEE 802.3ae – 2002 [3], or IEEE 802.3ba [16].
100 Mbps	Full duplex	
10/100 Mbps Auto-Negotiation	Full duplex	
1 Gbps	Full duplex	
10 Gbps	Full duplex	
40 Gbps	Full duplex	
100 Gbps	Full duplex	

Table 2 – Possible Physical Layer Characteristics of an Ethernet EI

[R8] All links in a particular Physical Layer for a particular EI **MUST** have the same Link Type, Speed, Mode, and Physical Medium.

6.3 Physical-to-Frame Format

The Physical-to-Frame Format describes how the EI Frame can be derived from the bit stream on the physical links that make up this EI.

[O3] The Physical-to-Frame Format **MAY** be ETHERNET.

[1] An exception might be wireless when the service requires stringent requirements on packet loss.

Other Physical-to-Frame Format types can be added by Child Documents. For example, consider the case where a Child Document requires the use of a Generic Framing Protocol encapsulated DS3 Link. Then it would specify the new DS3 Link Type as described above. It would also specify a new Physical-to-Frame Format Type (e.g., GFP) and would specify the requirements around the use of this mapping (e.g., by referring to the appropriate standards).

[R9] If the Link Type is ETHERNET, then the Physical-to-Frame Format **MUST** be ETHERNET and the mapping **MUST** be as specified in Table 2.

6.4 Link Protection Mechanism

The Link Protection Mechanism (LPM) describes the link protection mechanism used on the EI.

[R10] If there is only one link in the EI, the Link Protection Mechanism **MUST** be NONE.

[R11] If there are two or more links in the EI then there **MUST** be protection between them so that upon failure of a link, traffic that would normally traverse this link is automatically (i.e., without human or management plane intervention) moved to other links in the EI.

6.4.1 Link Aggregation Group (802.1AX – 2008)

[R12] If Link Aggregation Group (LAG) (as specified in 802.1AX – 2008 [15]) is used to protect the links in this EI, all links in the EI **MUST** belong to one Link Aggregation Group.

[R13] If Link Aggregation Group (LAG) (as specified in 802.1AX – 2008 [15]) is used to protect the links in this EI, the LPM **MUST** specify that the protection type is LAG (Link Aggregation Group)

[R14] If Link Aggregation Group (LAG) (as specified in 802.1AX – 2008 [15]) is used to protect the links in this EI, the LPM **MUST** specify if LACP is enabled and if so, what the periodic timer is set to.

[D1] If LACP is enabled, the periodic timer **SHOULD** be set to 1 per second "fast mode".

[R15] If Link Aggregation Group (LAG) (as specified in 802.1AX – 2008 [15]) is used to protect the links in this EI, the Operator **MUST** be able to administratively detach a link from, and add a link to, a LAG[2].

[D2] The mechanism to do so **SHOULD** minimize frame loss, reordering, and impact to other links in the group.

[2] The mechanism that allows the Operator to do this is beyond the scope of this specification but is assumed to be via management plane interaction to the NE housing this EI.

6.4.2 Other Protection Mechanism

Child Documents that refer to this specification can allow other protection mechanisms to be specified. The Child Document would need to specify the protection mechanism and all of its externally visible attributes as seen by the Connected Sub-Network (CSN).

6.5 Frame Format

The Frame Format defines the format of the Frames sufficiently to ensure inter-operability between the ESN and the Connected Sub-Network (typically by referring to another standard such as 802.1ad or 802.1ah).

[R16] All links in an EI **MUST** use this Frame Format for all ingress and egress EI Frames crossing this EI.

6.6 EI Maximum Transmission Unit Size

[R17] The EI Maximum Transmission Unit Size specifies the maximum frame size (in Bytes) allowed at the EI. The EI MTU Size **MUST** be specified to have a value greater than or equal to 1522.

[D3] The EI MTU Size **SHOULD** be specified to have a value greater than or equal to 2000.

[R18] The ESNC MTU Size for each ESNC that has an ESNC End Point (see Section 9.4) at the EI **MUST** be less than or equal to the EI MTU Size.

[R19] The length of each Egress EI Frame **MUST** be smaller or equal to the EI MTU Size.

[R20] An Ingress EI Frame that is larger than the EI MTU Size **MUST** be discarded by the ESN.

[D4] When an ingress EI Frame is discarded due to an MTU size that is larger than the EI MTU Size, it **SHOULD** be counted as part of any ingress EI Bandwidth Profile applied at this EI.

6.7 Maximum Number of ESNCs

[R21] The Maximum Number of ESNCs defines the maximum number of ESNCs that the EI can support. It **MUST** be an integer value of at least one.

6.8 Maximum Number of ESNC End Points

[R22] The Maximum Number of ESNC End Points defines the maximum total number of ESNC End Points that the EI can support. It **MUST** have a value of at least one.

[R23] The Maximum Number of ESNC End Points **MUST** be greater than or equal to the Maximum Number of ESNCs.

[R24] If the Maximum Number of ESNC End Points per ESNC (see Section 6.9) is 1, then the Maximum Number of ESNC End Points **MUST** be equal to the Maximum Number of ESNCs.

6.9 Maximum Number of ESNC End Points per ESNC

Figure 5 shows an example of Multiple ESNC End Points associated with one ESNC at an EI. ESNC B1 in the figure is an ESNC with three ESNC End Points, two of which are at EI 3. In this example, EI Frames associated with ESNC A1 are mapped to ESNC End Point B1.1 at EI 3. Also in this example, EI Frames associated with ESNC A2 are mapped to ESNC End Point B1.2 at EI 3. Consider an EI Frame received at EI 1 with a destination MAC address that was attached to E2. The Frame would be delivered across CSN A on ESNC A1 to EI 3. ESN B would receive this as an Ingress EI Frame via ESNC End Point B1.1 at EI 3 and deliver it via ESNC End Point B1.2 across EI 3 towards CSN A. This can be referred to as a “hairpin”. It would then travel across ESNC A2 towards its destination at EI 2. To support this service EI 3 must be capable of supporting more than one ESNC End Point associated with one ESNC when described as an EI of ESN B. It is important to notice that the diagram could be looked at from a different perspective, where CSN A is the ESN under examination and ESN B is the Connected Sub-Network (CSN). In this case, the same EI 3 does not need to support multiple ESNC End Points associated with one ESNC to enable this service.

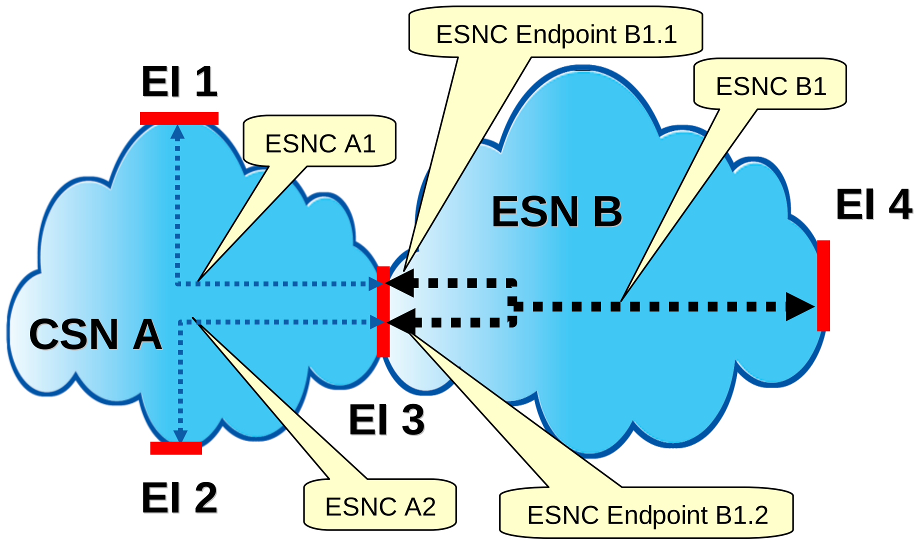

Figure 5 – Example of supporting multiple ESNC End Points for one ESNC at one EI

[R25] The Maximum Number ESNC End Points per ESNC defines the maximum number of ESNC End Points associated with one ESNC that the EI can support. It **MUST** have a value of at least one.

If the Maximum Number ESNC End Points per ESNC is set to be one, then the ESN under examination is not capable of supporting the "hairpin" function (described above) at this EI.

6.10 Maximum Number of Tunnel End Points

A Tunnel End Point is defined in section 7.

[R26] This attribute defines the maximum total number of Tunnel End Points that the EI can support. It **MUST** be an integer greater than or equal to zero.

6.11 Ingress Bandwidth Profile per EI

The Ingress Bandwidth Profile per EI describes ingress policing by the Operator on all ingress EI Frames. The Operator can apply an Ingress Bandwidth profile per EI at a particular EI.

[R27] When the Ingress Bandwidth Profile per EI is in force, suitable parameters <*CIR*, *CBS*, *EIR*, *EBS*, *CF*, *CM*> as defined in Section 9 **SHALL** be specified and the algorithm of Section 10.2 **MUST** be applied to all ingress EI Frames.

Figure 6 illustrates an example of the application of ingress policing with an Ingress Bandwidth Profile per EI. In this example, ingress EI Frames for the three ESNCs would all be subject to a single Ingress Bandwidth Profile. The Ingress Bandwidth Profile per EI manages bandwidth non-discriminately for all ESNC End Points at the EI, i.e., some ESNC End Points can get more bandwidth while others can get less.

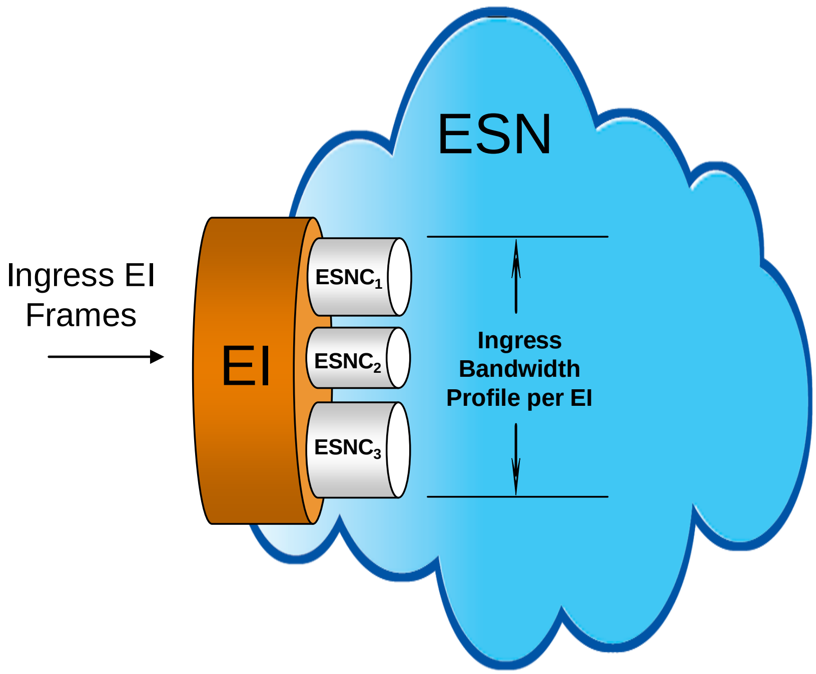

Figure 6 – Example of Ingress Bandwidth Profile per EI

6.12 Egress Bandwidth Profile per EI

The Egress Bandwidth Profile per EI describes egress traffic conditioning by the Operator on all egress EI Frames. The Operator can apply an Egress Bandwidth profile per EI at a particular EI.

[R28] When the Egress Bandwidth Profile per EI is in force, suitable parameters <*CIR*, *CBS*, *EIR*, *EBS*, *CF*, *CM*> as defined in Section 9 **SHALL** be specified and all egress EI Frames **MUST** have the property mandated by [R85] .

Figure 7 illustrates an example of the application of an Egress Bandwidth Profile per EI. In this example, egress EI Frames for the three ESNCs would all be subject to a single Egress Bandwidth Profile. The Egress Bandwidth Profile per EI manages bandwidth non-discriminately for all ESNC End Points at the EI, i.e., some ESNC End Points can get more bandwidth while others can get less.

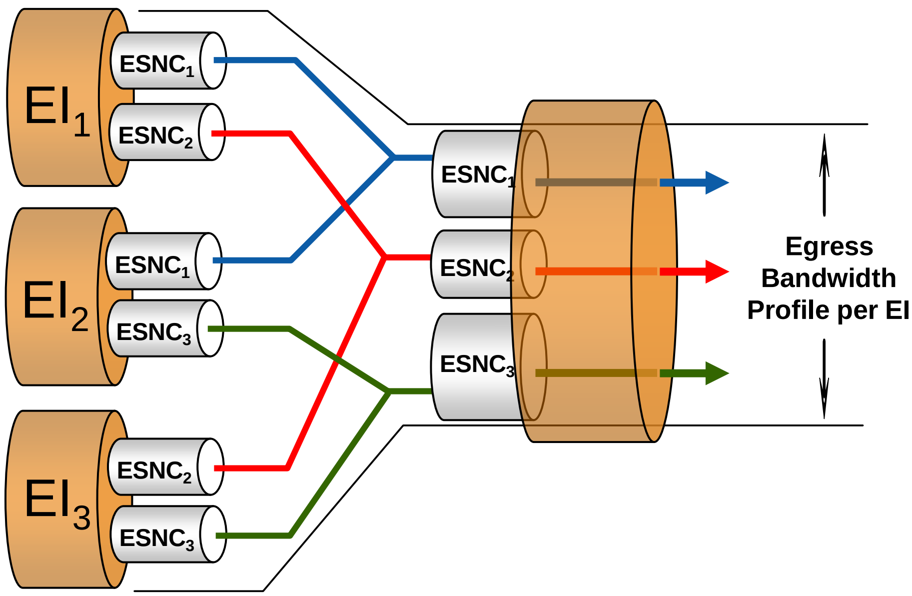

Figure 7 – Example of Egress Bandwidth Profile per EI

6.13 EI Frame Map

The EI Frame Map describes the details of how the content of an EI Frame determines the ESNC End Point or TE with which this EI Frame is associated (if any). It is not possible to define all the details of this mapping in a generic way because the frame format (see Section 6.5) is not specified. A Child Document using the Ethernet Service Constructs to define a specific external interface (e.g., an SNI) is expected to significantly refine this attribute depending on the frame format used on that interface. There are, however, a number of EI Frame Map requirements that are independent of the frame format on the EI.

[R29] For each frame crossing a given EI, the ESNC End Point Map **MUST** unambiguously associate the frame with one ESNC End Point at the given EI, one Tunnel End Point at the EI or neither.

[R30] When an EI Frame is mapped to a TE, there **MUST** be a subsequent mechanism that maps this frame to a particular ESNC End Point.

A child document can define other types of End Points and can map frames to End Points of this type using the EI Frame Map.

[R31] An ingress Data EI Frame (see Section 9.6.1) that is not mapped to any ESNC End Point at this EI, any Tunnel End Point at this EI, or any End Points at this EI of a type defined by the child document **MUST** be discarded.

7 Tunnel End Point Attributes

A TE is an ESC that represents a collection of ESNC End Points at an EI. It is used to facilitate the description of the externally observable behavior of the case where multiple ESNCs are encapsulated into a single service stream within an ESN.

There are attributes for each TE in the ESN. These attributes are summarized in Table 3 and described in detail in the sub-sections that follow this table.

Attribute Name	Summary Description
Operator TE Identifier	A string that uniquely identifies this TE within this ESN.
TE Frame Mapping	The way that EI Frames that have been mapped to this TE are mapped to ESNC End Points at this TE.
Maximum Number of ESNCs in TE	The maximum number of ESNCs that the ESN can support at this TE.
Maximum Number of ESNC End Points in TE	The maximum total number of ESNC End Points that the ESN can support at this TE.
Maximum Number of ESNC End Points per ESNC in TE	The maximum number of ESNC End Points associated with one ESNC that the can be supported at this TE.

Table 3 – TE Attributes

7.1 Operator TE Identifier

The Operator TE Identifier is a string assigned to the TE by the Operator of the Ethernet Sub-Network.

[R32] The Operator TE Identifier **MUST** be no more than 45 bytes in length.

[R33] The Operator TE Identifier **MUST** be unique for each TE in the ESN.

A child document could place additional requirements on this identifier. For example, the identifier could be mandated to be a Character String as defined in Table 21-19 of IEEE 802.1ag – 2007 [17].

7.2 TE Frame Mapping

EI Frames are mapped to a particular TE as described in section 6.13.

The TE Frame Mapping describes how the content of this same EI Frame determines the ESNC End Point within this TE to which this frame is associated. It is not possible to define all the details of this mapping in a generic way because the frame format (see Section 6.5) is not specified. A Child Document using the Ethernet Service Constructs to define a specific external interface is expected to significantly refine this attribute depending on the frame format used on that interface. There are, however, a number of TE Frame Map requirements that are independent of the frame format on the EI.

[R34] For each ingress frame and each egress frame at a given EI that is mapped to this TE, the TE Frame Map **MUST** unambiguously associate the frame with at most one ESNC End Point.

[R35] Ingress Data EI Frames (see Section 9.6.1) that are mapped to this TE but are not mapped to any ESNC End Point **MUST** be discarded.

7.3 Maximum Number of ESNCs at this TE

[R36] The Maximum Number of ESNCs at this TE defines the maximum number of ESNCs that this TE can support. It **MUST** be an integer value of at least one.

7.4 Maximum Number of ESNC End Points at this TE

[R37] The Maximum Number of ESNC End Points at this TE defines the maximum total number of ESNC End Points that the TE can support. It **MUST** be an integer value of at least one.

[R38] The Maximum Number of ESNC End Points at this TE **MUST** be greater than or equal to the Maximum Number of ESNCs at this TE.

[R39] If the Maximum Number of ESNC End Points per ESNC at this TE (see Section 7.5) is 1, then the Maximum Number of ESNC End Points at this TE **MUST** be equal to the Maximum Number of ESNCs at this TE.

7.5 Maximum Number of ESNC End Points per ESNC at this TE

Figure 8 shows an example of multiple ESNC End Points for one ESNC supported at one TE. The service represented across the Service Provider Domain is a Multipoint-to-Multipoint EVC with three End Points, one at UNI 1, one at UNI 2, and one at UNI 3. ESNC A1 is an ESNC with three ESNC End Points (one at UNI 1 and two located at Interface A/B, both of which are associated with the TE at this interface). A Service Frame entering at UNI 2 but with a destination MAC address for UNI 3 would travel through ESNC C1 and be encapsulated into a tunnel before handoff at Interface B/C. It would then be tunneled through ESNC B1 across ESN B. At Interface A/B, it would be mapped to the TE at this interface. The TE would then map the frame to ESNC End Point A1.1. At this point, the frame would then be "hair-pinned" back on ESNC A1 to ESNC End Point A1.2 then encapsulated and handed off at Interface A/B. It would then be tunneled back through ESNC B1 across ESN B. At Interface B/C, it would be mapped to the TE at this interface. The TE would then map the frame to ESNC C2 where it would travel to its final destination on UNI 3.

See section 12 for examples of how a TE can be used to describe the externally observable behavior of different tunneling applications.

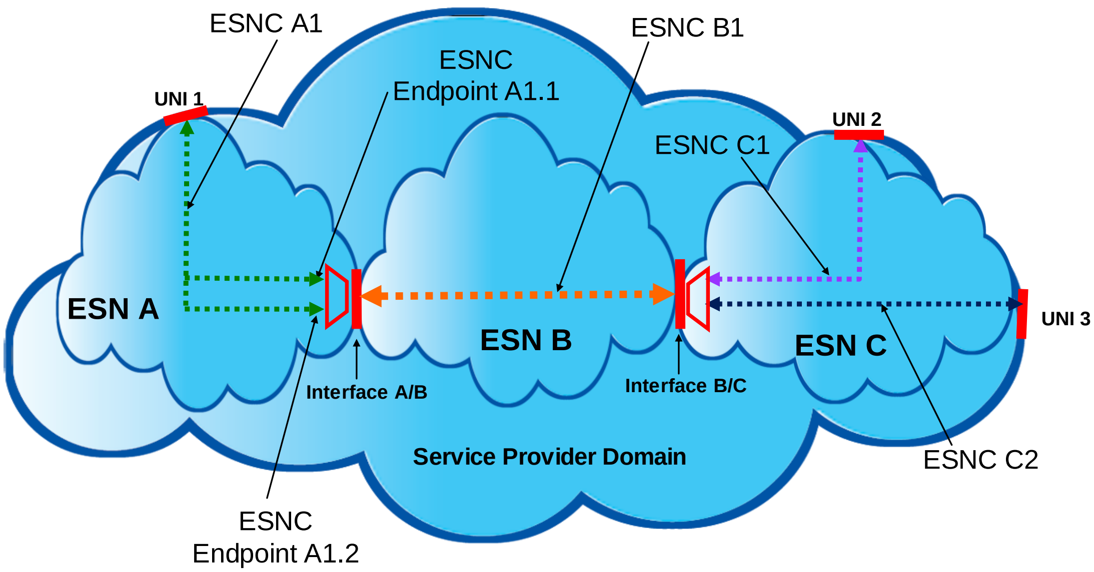

Figure 8 – Example of Supporting multiple ESNC End Points for one ESNC at a TE

[R40] The Maximum Number of ESNC End Points per ESNC at this TE defines the maximum total number of ESNC End Points per ESNC that the TE can support. It **MUST** be an integer of at least one.

8 ESNC End Point Attributes

An ESNC End Point is defined to be a logical reference point located at a physical External Interface (EI).

There are attributes for each ESNC End Point. These are summarized in Table 4 and described in detail in the sub-sections that follow this table.

Attribute Name	Summary Description
ESNC End Point ID	A string that uniquely identifies this ESNC End Point within the ESN.
ESNC End Point Type	The Type of this ESNC End Point. Can be one of Root, Leaf, From-Leaf or From-Root.
ESNC End Point Class of Service Identifier	The way that a Class of Service is determined for a frame at this ESNC End Point
Ingress Bandwidth Profile Per ESNC End Point	Ingress Bandwidth Profile enforced by the ESN on all ingress EI Frames mapped to this ESNC End Point
Ingress Bandwidth Profile Per ESNC End Point Class of Service Identifier	Ingress Bandwidth Profile enforced by the ESN on all ingress EI Frames mapped to this ESNC End Point and associated with each Class of Service Identifier
Egress Bandwidth Profile Per ESNC End Point	Egress Bandwidth Profile enforced by the ESN on all egress EI Frames mapped to this ESNC End Point

Attribute Name	Summary Description
Egress Bandwidth Profile Per ESNC End Point Class of Service Identifier.	Egress Bandwidth Profile enforced by the ESN on all egress EI Frames mapped to this ESNC End Point and associated with each Class of Service Identifier

Table 4 – ESNC End Point Attributes

8.1 ESNC End Point ID

The ESNC End Point ID is a string assigned to the TE by the Operator of the Ethernet Sub-Network.

[R41] The ESNC End Point ID **MUST** be no more than 45 bytes in length.

[R42] The ESNC End Point ID **MUST** be unique within the ESN.

A child document could place additional requirements on this identifier. For example, the identifier could be mandated to be a Character String as defined in Table 21-19 of IEEE 802.1ag – 2007 [17].

8.2 ESNC End Point Type

[O4] The ESNC End Point Type **MAY** have one of the following values: Root, Leaf, From-Root, or From-Leaf.

Delivery of frames between ESNC End Points types are as specified in section 9.

[O5] Child Documents **MAY** define new ESNC End Point Types specifying the delivery criteria between ESNC End Points of all types.

See Section 11 for an informative description of the different ESNC End Point types.

8.3 ESNC End Point Class of Service Identifier

A Child Document can specify EI Frame delivery performance for all EI Frames transported within an ESNC with a particular Class of Service instance. The Class of Service instance for a given EI Frame is derivable from the content of the EI Frame header. For example, suppose that three Classes of Service are offered called silver, gold, and platinum and, at a given EI, there are three instances of silver service, two instances of gold service and one instance of platinum service. Then there would be six Class of Service Identifiers, one for each Class of Service instance.

[O6] An EI Frame delivery performance **MAY** be to discard the EI Frame. Thus a Class of Service Identifier may be specified for EI Frame discard.

[R43] EI Frames mapped to different ESNCs **MUST** have different ESNC End Point Class of Service Identifiers.

[R44] EI Frames mapped to different ESNC End Points of one ESNC **MUST** use the same ESNC End Point Class of Service Identifiers and **MUST** use the same method of determining the ESNC End Point Class of Service identifier at that EI.

[R45] When the ESNC End Point CoS ID is based on the code point values of a field in the EI Frame, each ESNC End Point CoS ID **MUST** consist of a set of such values and each value **MUST** belong to exactly one ESNC End Point CoS ID.

[R45] means that the sets of code point values for the Class of Service Identifiers are disjoint and their union equals the set of all code point values.

8.4 Ingress Bandwidth Profile per ESNC End Point

The Ingress Bandwidth Profile per ESNC End Point describes ingress policing by the Operator on all ingress EI frames mapped to a given ESNC End Point.

[R46] When the Ingress Bandwidth Profile per ESNC End Point is in force for a given ESNC End Point, suitable parameters <CIR, CBS, EIR, EBS, CF, CM> as defined in Section 10 **MUST** be specified and the algorithm of Section 10.2 **MUST** be applied to all ingress EI Frames that are mapped to the given ESNC End Point.

Figure 9 illustrates an example of the application of Ingress Bandwidth Profiles per ESNC End Point. In this example, $ESNC_1$ End Point 1 could have *CIR*=15 Mbps, $ESNC_2$ End Point 1 could have *CIR* = 10 Mbps, and $ESNC_2$ End Point 2 could have *CIR* = 20 Mbps.

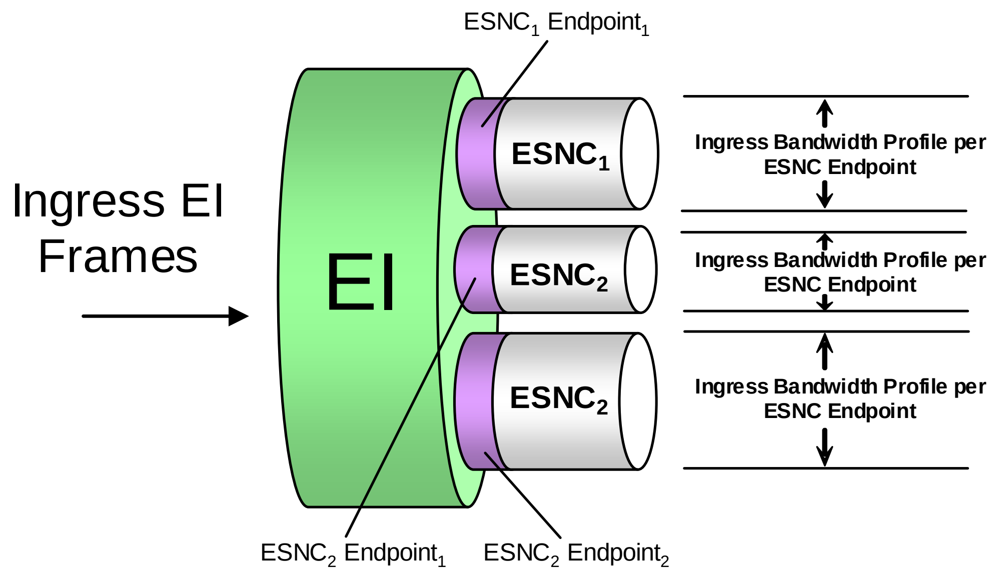

Figure 9 – Example of Ingress Bandwidth Profile per ESNC End Point

8.5 Ingress Bandwidth Profile per Class of Service Identifier

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

[R47] When the Ingress Bandwidth Profile per Class of Service Identifier is in force for a given Class of Service Identifier, suitable parameters <CIR, CBS, EIR, EBS, CF, CM> as defined in Section 10 **MUST** be specified and the algorithm of Section 10.2 **MUST** be applied to all ingress EI Frames that have the given Class of Service Identifier.

Figure 10 shows an example of an Ingress Bandwidth Profile per Class of Service Identifier.

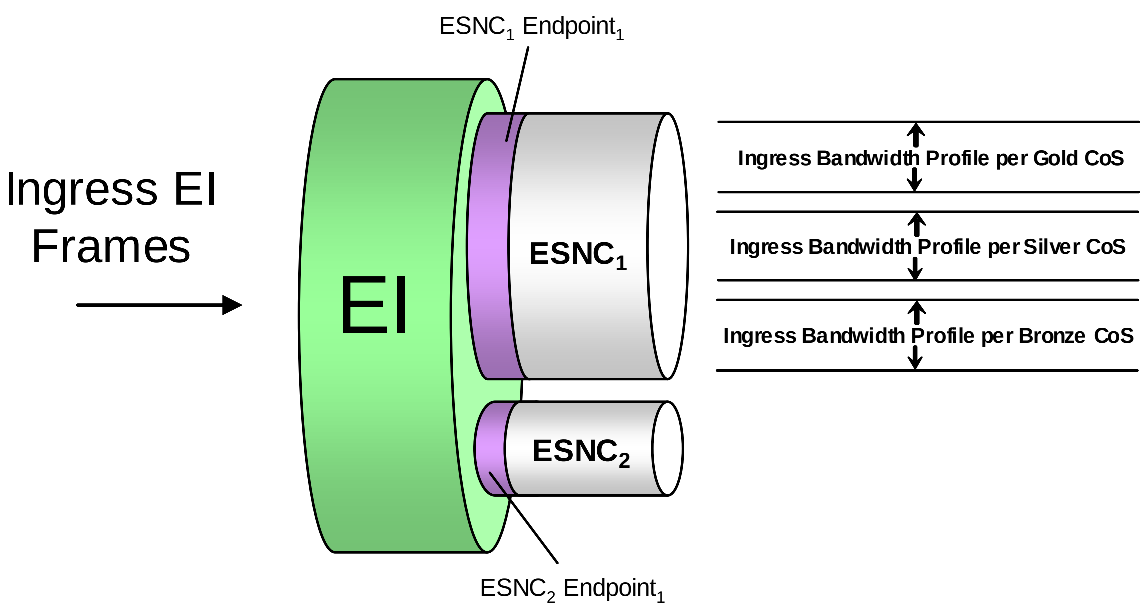

Figure 10 – Example of Ingress Bandwidth Profile per Class of Service Identifier

8.6 Egress Bandwidth Profile per ESNC End Point

The Egress Bandwidth Profile per ESNC End Point describes egress traffic conditioning by the Operator on all egress EI Frames that are mapped to a given ESNC End Point.

[R48] When the Egress Bandwidth Profile per ESNC End Point is in force for a given ESNC End Point, suitable parameters <CIR, CBS, EIR, EBS, CF, CM> as defined in Section 10 **MUST** be specified and all egress EI Frames mapped to the given ESNC End Point **MUST** have the property defined in Section 10.3.

Figure 11 illustrates an Egress Bandwidth Profile.

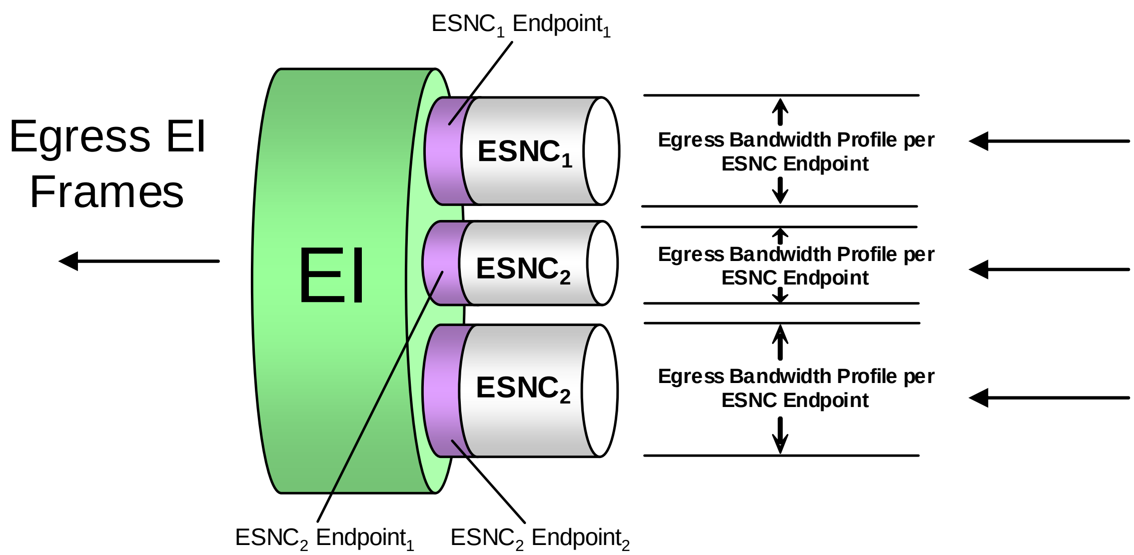

Figure 11 – Example of Egress Bandwidth Profile per ESNC End Point

8.7 Egress Bandwidth Profile per Class of Service Identifier

The Egress Bandwidth Profile per Class of Service Identifier describes egress traffic conditioning by the Operator on all egress EI Frames with a given Class of Service Identifier.

[R49] When the Egress Bandwidth Profile per Class of Service Identifier is in force for a given Class of Service Identifier, suitable parameters <CIR, CBS, EIR, EBS, CF, CM> as defined in Section 10 **MUST** be specified and all egress Frames with the given Class of Service Identifier **MUST** have the property defined in 10.3.

Figure 12 shows an example of an Egress Bandwidth Profile per Class of Service Identifier.

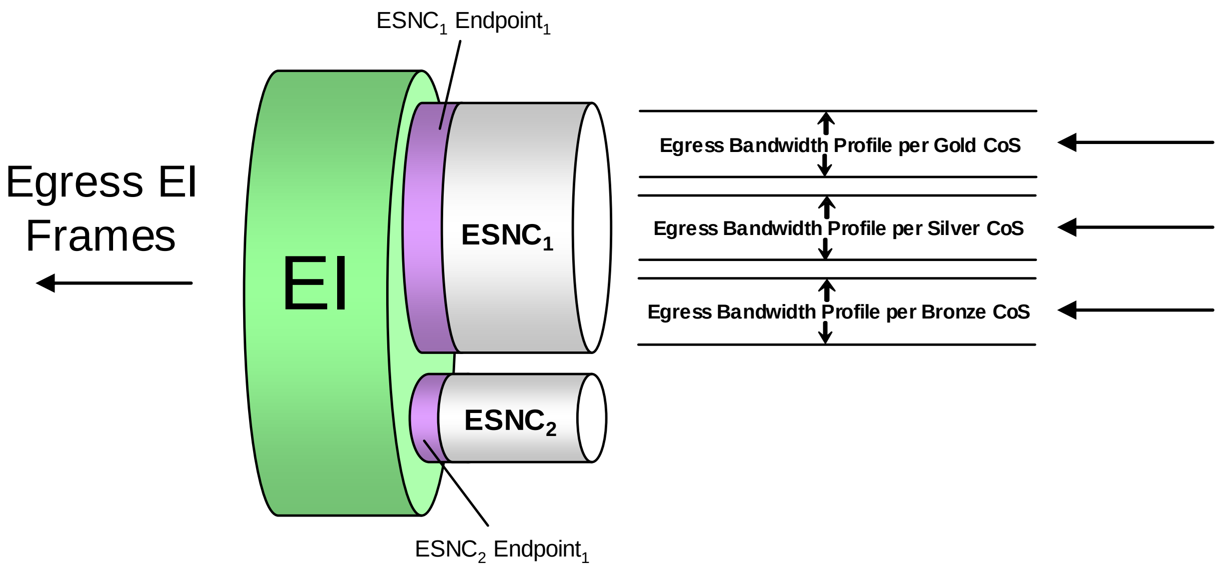

Figure 12 – Egress Bandwidth Profile per EI Class of Service Identifier

9 ESNC Attributes

An ESNC is defined as an association of ESNC End Points, and all of the associated ESNC End Points are said to be "in the ESNC".

An EI Frame that is mapped to the ESNC can be delivered to one or more ESNC End Points of the ESNC.

[R50] An ESNC End Point **MUST** be at one particular EI.

[R51] An ESNC End Point **MUST** be in exactly one ESNC.

[O7] There **MAY** be more than one ESNC End Point for a particular ESNC at a particular EI.

[R52] If an Egress EI Frame Associated with an ESNC End Point results from an Ingress EI Frame associated with an ESNC End Point, the egress ESNC End Point **MUST** be in the same ESNC as that of the ingress ESNC End Point and the egress ESNC End Point **MUST** be different from the ingress ESNC End Point.

There are four pre-defined types of ESNC End Points; Root, Leaf, From-Root, and From-Leaf. See Section 11 for an informative description of the different ESNC End Point types.

[R53] An Ingress EI Frame at a Root ESNC End Point **MUST NOT** result in an Egress EI Frame at a From-Leaf ESNC End Point.

[R54] An Ingress EI Frame at a Leaf ESNC End Point **MUST NOT** result in an Egress EI Frame at a Leaf ESNC End Point or a From-Root ESNC End Point.

[R55] An Ingress EI Frame at a From-Leaf ESNC End Point **MUST NOT** result in an Egress EI Frame at a Leaf ESNC End Point or a From-Root ESNC End Point.

[R56] An Ingress EI Frame at a From-Root ESNC End Point **MUST NOT** result in an Egress EI Frame at a From-Leaf ESNC End Point.

Table 5 summarizes the connectivity constraints based on the type of ESNC End Point.

		Ingress ESNC End Point Type			
		Root	Leaf	From-Leaf	From-Root
Egress ESNC End Point Type	Root	☑	☑	☑	☑
	Leaf	☑	☐	☐	☑
	From-Leaf	☐	☑	☑	☐
	From-Root	☑	☐	☐	☑

Table 5 – Allowed Connectivity between ESNC End Point Types

The ESNC Attributes are summarized in Table 6 and described in detail in the sub-sections that follow this table.

Attribute Name	Summary Description
ESNC ID	An arbitrary string, unique across the ESN, for the ESNC
ESNC End Point List	A list of the ESNC End Point Identifiers that are associated with this ESNC.
Maximum Number of ESNC End Points	The maximum number of ESNC End Points that can be in this ESNC.
ESNC Maximum Transmission Unit Size	The maximum length EI Frame in bytes allowed on the ESNC in this ESN
CE-VLAN ID Preservation	The ability to preserve the CE-VLAN ID of the original Service Frame across this ESNC. Must be ALL, TAGGED, or NO.
CE-VLAN CoS Preservation	The ability to preserve the CE-VLAN CoS of the original Service Frame across this ESNC. Must be YES or NO
Unicast EI Frame Delivery	This attribute describes how Ingress EI Frames at an ESNC End Point with a unicast destination MAC address are delivered to the other ESNC End Points in the ESNC.
Multicast EI Frame Delivery	This attribute describes how Ingress EI Frames at an ESNC End Point with a Multicast destination MAC address are delivered to the other ESNC End Points in the ESNC.
Broadcast EI Frame Delivery	This attribute describes how Ingress EI Frames at an ESNC End Point with the broadcast destination MAC address are delivered to the other ESNC End Points in the ESNC.

Table 6 – ESNC Attributes

9.1 ESNC ID

The ESNC ID is a string administered by the ESN Operator that is used to identify an ESNC within the ESN.

[R57] The ESNC ID **MUST** be unique across all ESNCs in the ESN.

[R58] The ESNC ID **MUST** be no more than 45 bytes in length.

A child document could place additional requirements on this identifier. For example, the identifier could be mandated to be a Character String as defined in Table 21-19 of IEEE 802.1ag – 2007 [17].

The ESNC ID is intended for management and control purposes. The ESNC ID is not carried in any field in the EI Frame. As an example, the Acme Operator might use "ESNC-0001898-ACME-MEGAMART" to represent the 1898th ESNC in the ESN. This ESNC is part of an end-to-end EVC provided to a Subscriber named MegaMart.

9.2 ESNC End Point List

[R59] The ESNC End Point List is a list of all the ESNC End Point Identifiers that are associated with this ESNC. For every ESNC End Point that is in this ENC, there **MUST** be one corresponding entry in this list.

9.3 Maximum Number of ESNC End Points

[R60] The Maximum Number of ESNC End Points specifies the maximum number of ESNC End Points that can be associated by this ESNC. For a Point-to-Point ESNC, the Maximum Number of ESNC End Points **MUST** be two.

9.4 ESNC Maximum Transmission Unit Size

[R61] The ESNC Maximum Transmission Unit Size specifies the maximum EI Frame size (in Bytes) allowed on the ESNC. Every EI with an ESNC End Point in the ESNC **MUST** be capable of supporting this EI Frame size[3] [See section 6.6].

[R62] The MTU **MUST** be specified to have a value greater than or equal to 1522.

[D5] The MTU **SHOULD** be specified to have a value greater than or equal to 2000.

[R63] When an Ingress EI Frame that is mapped to this ESNC has length greater than the ESNC Maximum Transmission Unit Size, the EI Frame **MUST** be discarded.

[3] The MTU size for an ESNC will be constrained by the MTU size of the network equipment used to carry the frame including the network equipment supporting each EI. The method of calculating the MTU Size is beyond the scope of this specification.

9.5 CE-VLAN Tag Preservation

MEF 10.2 [10] defines a CE VLAN Tag and its associated CE VLAN ID and a CE VLAN CoS as it can be derived from a particular Service Frame crossing a UNI. It also defines CE VLAN ID and CE VLAN CoS preservation across an EVC between UNIs.

There is a requirement to define similar preservation attributes across an ESNC from EI to EI. The intent is to ensure that a Service Provider creating a UNI to UNI EVC using a set of ESNCs can infer that as long as all the ESNCs used in the path of the EVC have appropriate preservation, then the end-to-end EVC has the desired preservation behavior.

However, in the context of this document (ESC), the frame format at each of the EIs is not defined (see Section 6.5) making it impossible to describe how to derive the CE VLAN ID or the CE VLAN CoS from the EI Frame. Similarly, it is not possible to specify how to derive whether or not the original Ingress Service Frame was untagged or tagged. Instead, the following are left to Child Documents:

- How to derive from an EI Frame whether the original Ingress Service Frame was tagged or untagged
- If it was tagged, how to derive from an EI Frame the value of the CE VLAN ID
- If it was tagged, how to derive from an EI Frame the value of the CE VLAN CoS

The CE-VLAN ID and CE-VLAN CoS preservation can only be applied in a case where the Child Document specifies the above.

9.5.1 CE-VLAN ID Preservation

[R64] The CE-VLAN ID Preservation **MUST** be either "ALL" "TAGGED" or "NO"

[R65] If CE-VLAN ID preservation is "ALL" for a particular ESNC, 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 7 .

Ingress Frame	Egress Frame
It can be derived that the Ingress EI Frame is Untagged	It can be derived from the Egress EI Frame that the corresponding Ingress EI Frame was Untagged.
The CE VLAN ID can be derived from the Ingress EI Frame	The CE VLAN ID can be derived from the Egress EI Frame and the CE VLAN ID value is equal to that of the corresponding Ingress Frame.

Table 7 – Relationship between Ingress and Egress EI Frames with CE VLAN ID Preservation equal to ALL

[R66] If CE-VLAN ID preservation is "TAGGED" for a particular ESNC, 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 8.

Ingress Frame	Egress Frame
The CE VLAN ID can be derived from the Ingress EI Frame and it is between 1 and 4094.	The CE VLAN ID can be derived from the Egress EI Frame and the CE VLAN ID value is equal to that of the corresponding Ingress Frame.

Table 8 – Relationship between Ingress and Egress EI Frames with CE VLAN ID Preservation equal to TAGGED

9.5.2 CE-VLAN CoS Preservation

[R67] The CE-VLAN CoS Preservation **MUST** have a value of "YES" or "NO".

[R68] If the CE-VLAN CoS Preservation is YES, 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 9.

Ingress Frame	Egress Frame
The CE VLAN CoS can be derived from the Ingress EI Frame.	The CE VLAN CoS can be derived from the Egress EI Frame and the CE VLAN CoS value is equal to that of the corresponding Ingress Frame.

Table 9 – Relationship between Ingress and Egress EI Frames with CE VLAN CoS Preservation equal to YES.

9.6 EI Frame Delivery

9.6.1 Types of EI Frames

There are a number of pre-defined types of EI Frame as follows:

- Layer 2 Control Protocol EI Frames
- Data EI Frames (of three types)
  - Unicast EI Frames
  - Multicast EI Frames
  - Broadcast EI Frames

Child documents can choose to add other types of EI Frames describing how they are identified and how they are to be treated.

9.6.1.1 Layer 2 Control Protocol EI Frames

Given that there are several Layer 2 protocols used for various control purposes, it is important that the ESN be able to process such information effectively.[4]

[R69] An EI Frame whose destination MAC address is one of the addresses listed in Table 1 of MEF 10.2 [10], **MUST** be treated as Layer 2 Control Protocol EI Frame.

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

[O8] An Operator **MAY** define additional addresses for identifying Layer 2 Control protocols in addition to those in Table 1 of MEF 10.2 [10].

9.6.1.2 Data EI Frames

Data EI Frames are EI Frames that are not L2CP Frames.

A Unicast EI Frame is a Data EI Frame that has a unicast destination MAC address.

A Multicast EI Frame is a Data EI Frame that has a multicast destination MAC address that has not been designated as carrying a Layer 2 Control Protocol as in Section 9.6.1.1.

A Broadcast EI Frame is a Data EI Frame with the broadcast destination MAC address.

9.6.2 EI Frame Disposition

[R70] An ingress EI Frame that is invalid as defined in Clause 3.4 of IEEE 802.3 – 2005 [2] **MUST** be discarded by the receiving ESN..

The disposition of an Ingress EI Frame at an ESNC End Point is described by one of the following:

- Deliver Unconditionally: No matter what the content (assuming correct FCS) of the EI Frame, it is delivered across the other (egress) ESNC End Points. This might be the behavior of a Point-to-Point ESNC.

- Deliver Conditionally: The EI Frame is delivered across an egress ESNC End Point if certain conditions are met. An example of such a condition is that the destination MAC address is known by the ESN to be "at" the destination ESNC End Point. Another example is broadcast throttling where some EI Frames with the

[4] This capability is especially important for example, when Subscribers choose to deploy IEEE 802.1D [6] or IEEE 802.1Q [7] bridges (as opposed to routers) as CEs.

broadcast destination MAC address are dropped to limit the amount of such traffic.

- Tunnel: This applies only to Layer 2 Control Protocol EI Frames.

Note that this is a description of the ideal service. EI Frames that should be delivered might be discarded due to network failure or congestion conditions.

9.6.3 EI Frame Transparency

Ethernet Frames ingress or egress each EI in the ESN (encapsulated over whatever physical format the EI takes). However, it is not the case that all EIs in a particular ESN use the same Ethernet Frame Format. For example, one EI in a particular ESN could have only single tagged EI frames, while another has 802.1ad based EI frames.

The detailed requirements of converting from a frame format on one EI type (e.g., converting a single tagged Ethernet frame at a UNI to an 802.1ad based Ethernet frame at an **ENNI**) is beyond the scope of this specification. Such requirements would be specified when defining the behavior of a specific type of EI across a specific ESN interfacing (in a pair wise manner) to specific EIs of other types. There are, however, some general requirements that apply in all cases:

[R71] All Egress EI Frames delivered by an ESNC to an egress ESNC End Point **MUST** be delivered using the appropriate frame format specified on the EI of the ESNC End Point (see Section 6.5) regardless of the frame format of the corresponding EI Frame at ingress.

[R72] The payload of all Egress Data EI Frames from a particular ESNC **MUST** be identical to the payload of the corresponding Ingress Data EI Frame. In this case, payload is defined to be Mac Client Data and Pad fields as specified in IEEE 802.3 – 2005 [2].

Specific attributes of an ESNC can enforce the condition that additional fields must be identical at ingress and egress. See Section 9.5.

9.6.4 Service Level Specification

The Service Level Specification (SLS) consists of definitions of delivery performance attributes for frames between External Interfaces, e.g., delay, loss. For each performance attribute, the SLS also contains a quantitative objective. When the frame delivery performance level meets or exceeds the objective for each performance attribute, the SLS is said to be met.

The SLS that applies to a given EI Frame is based on the Class of Service Identifier for the frame. See Section 8.3 for details about the Class of Service Identifier.

When an EI Frame is not sufficiently compliant with an Ingress Bandwidth Profile that applies to it, the SLS does not apply to the frame. See Section 10.2.2.

Details for the SLS might be addressed in a later phase of this specification or in a different future MEF document.

9.6.5 Unicast EI Frame Delivery

This attribute describes the EI Frame Disposition (see Section 9.6.2) of Unicast EI Frames (see Section 9.6.1.2).

[R73] The Unicast EI Frame Delivery **MUST** be "Deliver Unconditionally" or "Deliver Conditionally". If "Deliver Conditionally" is used, then the conditions **MUST** be specified.

9.6.6 Multicast EI Frame Delivery

The Multicast EI Frame Delivery describes the EI Frame Disposition (see Section 9.6.2) of Multicast EI Frames (see Section 9.6.1.2).

[R74] The Multicast EI Frame Delivery **MUST** be "Deliver Unconditionally" or "Deliver Conditionally". If "Deliver Conditionally" is used, then the conditions **MUST** be specified.

9.6.7 Broadcast EI Frame Delivery

The Broadcast EI Frame Delivery describes the EI Frame Disposition (see Section 9.6.2) of Broadcast EI Frames (see Section 9.6.1.2).

[R75] The Broadcast EI Frame Delivery **MUST** be "Deliver Unconditionally" or "Deliver Conditionally". If "Deliver Conditionally" is used, then the conditions **MUST** be specified.

10 Bandwidth profiles

A Bandwidth Profile is a method characterizing EI frames for the purpose of rate enforcement or policing. There are two types of Bandwidth Profile. An Ingress Bandwidth Profile is used to regulate the amount of ingress traffic at a particular EI, while an Egress Bandwidth Profile is used to regulate the amount of egress traffic at a particular EI. The Ingress Bandwidth Profile application is described in Sections 6.11, 8.4, and 8.5. The Egress Bandwidth Profile application described in Sections 6.12, 8.6 and 8.7.

A Bandwidth Profile is a characterization of the lengths and arrival times for EI frames at a reference point, which can be the ingress or egress EI.

Typically a Bandwidth Profile defines EI frame traffic that is less than the full bandwidth of the EI. Thus the Bandwidth Profile can be considered to be analogous to the traffic policing of Frame Relay ITU I.370 [4].

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

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

10.1 Bandwidth Profile Parameters and Algorithm

The parameters comprising the Bandwidth Profile parameters are:

[R76] **Committed Information Rate** (*CIR*) expressed as bits per second. *CIR* **MUST** be ≥ 0.

[R77] **Committed Burst Size** (*CBS*) expressed as bytes. When CIR > 0, *CBS* **MUST** be greater than or equal to the largest Maximum Transmission Unit size among all of the ESNCs that the Bandwidth Profile applies to.

[R78] **Excess Information Rate** (*EIR*) expressed as bits per second. *EIR* **MUST** be ≥ 0.

[R79] **Excess Burst Size** (*EBS*) expressed as bytes. When *EIR* > 0, *EBS* **MUST** be greater than or equal to the largest Maximum Transmission Unit size among all of the ESNCs that the Bandwidth Profile applies to.

[R80] **Coupling Flag** (*CF*) **MUST** have only one of two possible values, 0 or 1.

[R81] **Color Mode** (*CM*) **MUST** have only one of two possible values, “color-blind” and “color-aware.”

Each EI frame is classified to determine which, if any, Bandwidth Profile is applicable to the EI frame. Operation of the Bandwidth Profile algorithm is governed by the six parameters, <*CIR*, *CBS*, *EIR*, *EBS*, *CF*, *CM*>. The algorithm declares each EI 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 said to be in color-aware mode when each EI 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 said to be in color-blind mode when the color (if any) already associated with each EI frame is ignored by the Bandwidth Profile Algorithm.

[O9] Since the Coupling Flag has negligible effect in color-blind mode (*CM* = “color-blind”), a Bandwidth Profile that uses color-blind operation **MAY** be defined without specifying the value of the Coupling Flag.

The Bandwidth Profile algorithm is shown in Figure 13.

[R82] For a sequence of EI 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 EI frame **MUST** be defined according to the algorithm in Figure 13.

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*.

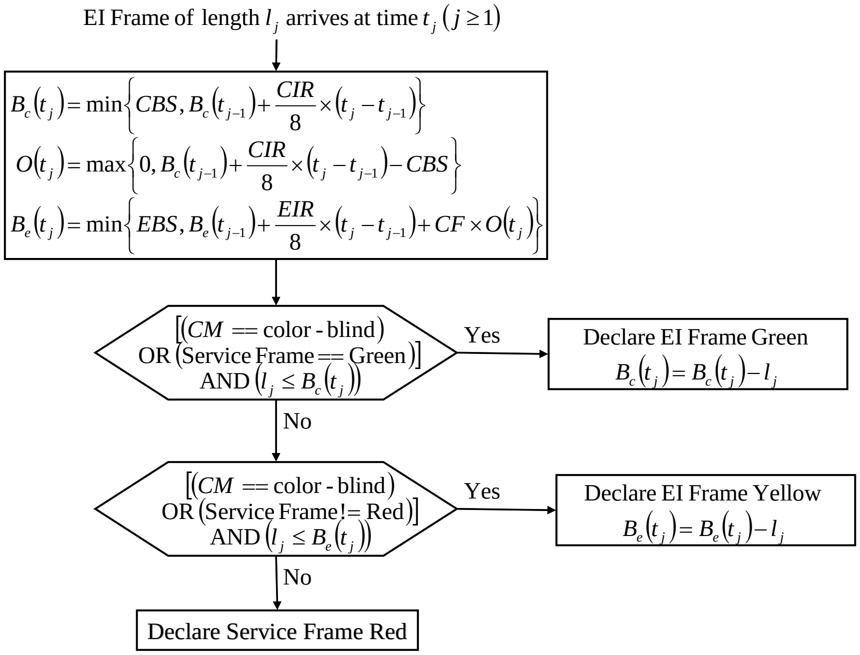

Figure 13 – The Bandwidth Profile Algorithm

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 EI frames that are declared Yellow. When *CF* is set to 0, the long term average bit rate of EI frames that are declared Yellow is bounded by *EIR*. When *CF* is set to 1, the long term average bit rate of EI frames that are declared Yellow is bounded by *CIR* + *EIR* depending on volume of the offered EI frames that are declared Green. In both cases the burst size of the EI frames that are declared Yellow is bounded by *EBS*.

10.2 Ingress Bandwidth Profiles

The Ingress Bandwidth Profile is used to regulate the amount of ingress traffic at a particular EI. An Ingress Bandwidth Profile is defined for ingress EI Frames at the particular EI. In other words, the sequence $\{t_j, l_j\}, j \geq 0,$ to which the algorithm of Section 10.1 is applied is based on ingress EI Frames at an EI. There are three Ingress Bandwidth Profile models as described in Sections 6.11, 8.4, and 8.5.

10.2.1 Simultaneous Application of the Ingress Bandwidth Profile Application Models

[O10] Multiple models of Ingress Bandwidth Profile application **MAY** exist simultaneously at an EI.

[R83] However, an EI **MUST** be configured such that only a single Ingress Bandwidth Profile applies to any given ingress EI Frame.

This means that:

- If there is a per-EI Ingress Bandwidth Profile, then there cannot be any other Ingress Bandwidth Profiles at that EI.
- If there is a per-ESNC End Point Ingress Bandwidth Profile on an ESNC, then there cannot be any per Class of Service Ingress Bandwidth Profiles for instances of CoS on that ESNC End Point.

For example, in the configuration of Figure 14, there cannot be an Ingress Bandwidth Profile for $ESNC_1$. Note also for the configuration in Figure 14, that it is possible to configure a per-End Point Ingress Bandwidth Profile for $ESNC_2$ but there happens to not be an Ingress Bandwidth Profile for $ESNC_2$ in this example.

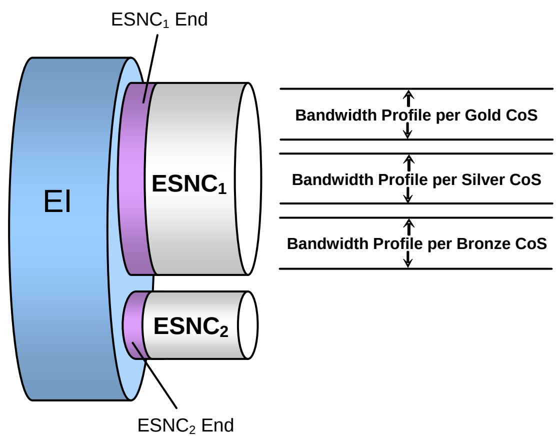

Figure 14 – Ingress Bandwidth Profile per Class of Service Identifier

10.2.2 EI Frame Disposition

The disposition of a given EI frame with respect to delivery to an egress EI is dependent on the EI 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 10 defines how the Ingress Bandwidth Profile compliance is related to the disposition of the EI Frame. In this table, "Not Applied" identifies the case where no Ingress Bandwidth Profile was applied to the EI Frame in question.

[R84] The disposition of each EI Frame for delivery to each egress EI **MUST** be as described in Table 10.

Ingress Bandwidth Profile Compliance	EI frame Disposition
Red	Discard
Yellow	Deliver to the egress EI according to the attributes of the service instance but SLS performance objectives do not apply.
Green	Deliver to the egress EI according to the attributes of the service instance. SLS performance objectives apply.
Ingress Bandwidth Profile Not Applied	

Table 10 – EI frame Disposition for Each Egress EI

The behavior described in Table 10 is based on the arrival of the EI Frame at its ingress EI. It does not mandate or constrain in any way the implementation within the ESN.

From Table 10 it is clear that the better the level of Ingress Bandwidth Profile compliance the better the performance of the service. In order to increase the level of Ingress Bandwidth Profile compliance, a CSN Operator can choose to shape traffic being sent into the ESN.

10.3 Egress Bandwidth Profiles

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

The reference point for an EI Egress Bandwidth Profile is the egress EI. An Egress Bandwidth Profile describes arrival times and lengths of EI Frames that will be observed at the egress EI when an Egress Bandwidth Profile is in operation in the ESN. This description is given in terms of what would happen if an observer at the egress EI applied the algorithm of Section 9 to egress EI Frames. This observer would see traffic after it had been subject to rate limiting and/or shaping in the ESN and thus would have certain characteristics. These characteristics are described in terms of the behavior of the algorithm of Section 10.1 when run by the observer.

[R85] Consider a sequence of egress EI Frames subject to an Egress Bandwidth Profile with parameters <*CIR*, *CBS*, *EIR*, *EBS*, *CF*, *CM*> and with arrival times and lengths at the egress EI, $\{t_j, l_j\}, j \geq 0$. If the algorithm of Section 10.1 is applied to these EI Frames, the result for each EI Frame **SHALL** be to declare the EI Frame either Green or Yellow.

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

There are three Egress Bandwidth Profile models as described in Sections 6.12, 8.6, and 8.7.

10.3.1.1 Simultaneous Application of the Egress Bandwidth Profile Application Models

[O11] Multiple models of Egress Bandwidth Profile application **MAY** exist simultaneously for an egress EI.

[R86] However, an egress EI Frame **MUST** be subject to at most one Egress Bandwidth Profile.

[R86] 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 for instances of CoS on the OVC that associates that OVC End Point.

11 Appendix A: ESNC End Point Types (informative)

MEF 10.2 [10] defines Point-to-Point EVCs as well as two kinds of Multipoint EVCs, Multipoint-to-Multipoint EVCs and Rooted-Multipoint EVCs. In the case of Point-to-Point EVCs and Multipoint-to-Multipoint EVCs, there is no need to define different types of End Points of the EVC (i.e., the different UNIs on which this EVC is terminated) since all such End Points behave the same way.

The Rooted-Multipoint EVC, however, is different. In this case, two different kinds of EVC End Points (attributes of the EVC at the UNI) are defined: Root and Leaf. In this case, Service Frames can be exchanged from Root to Root, Root to Leaf, and Leaf to Root, but not from Leaf to Leaf.

An example of where this might be used in a real-world application is for a Subscriber with two data centers and a number of branch offices. The Subscriber wishes to deploy a Multipoint Ethernet Service that allows traffic between the two data centers, and between either data center and any branch office but does not allow traffic to flow between the branch offices. This is achieved by deploying a Rooted-Multipoint Ethernet service where the Data Centers are given Root End Points and the branch offices given Leaf End Points.

Note also that, for consistency, MEF10.2 [10] specifies that all End Points of Point-to-Point and Multipoint-to-Multipoint EVCs must be Root End Points thereby allowing full connectivity.

The above example of a Rooted-Multipoint EVC can be expressed in the form of the ESC model as show in Figure 15. In this case, there is one ESNC (represented as a dashed line) with four different ESNC End Points (represented as arrow-heads). The two ESNC End Points that are located at the data center UNIs are marked with an “R” to represent that these are Root ESNC End Points. The two ESNC End Points that are located at the branch offices are marked with an “L” to represent that these are Leaf ESNC End Points.

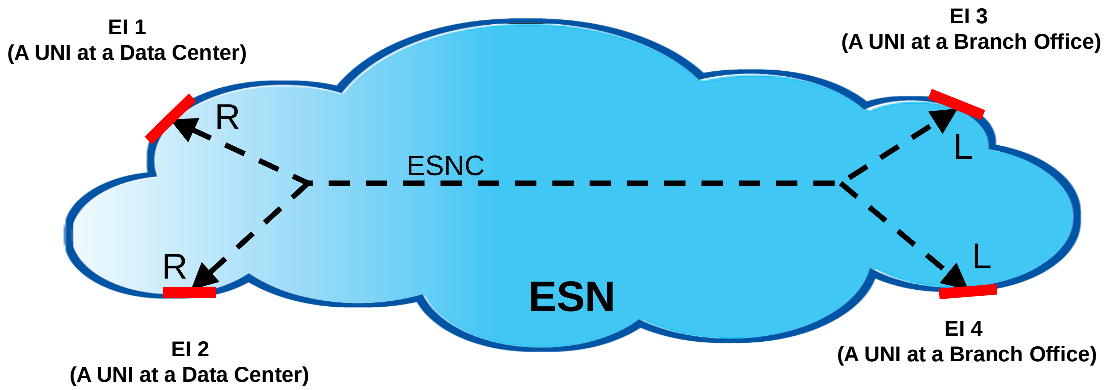

Figure 15 – A Rooted-Multipoint EVC represented in the ESC Model

Connectivity between the different ESNC End Points in Figure 15 is as specified in Table 11.

		Ingress ESNC End Point Type	
		Root	Leaf
Egress ESNC End Point	Root	☑	☑
	Leaf	☑	☐

Table 11 – Allowed Connectivity between ESNC End Point Types (Simple)

Now consider the application of the ESC model where two independent sub-networks exist with an EI between them as shown in Figure 16. Two independent ESNCs (ESNC 1 and ESNC 2) make up the UNI-to-UNI EVC. It is important to note that the ESC model examines one ESN at a time, so if ESN 1 is under examination, then ESN 2 becomes a CSN and vice versa. Treating ESN 1 in isolation, it is necessary to specify the ESNC End Point type of ESNC End Point 1.1. Similarly, treating ESN 2 in isolation, it is necessary to specify the ESNC End Point type of ESNC End Point 2.1. It is be clear that setting the values as shown in the Figure (ESNC End Point 1.1 is set to be a Leaf; ESNC End Point 2.1 is set to be a Root), and applying the connectivity specified in Table 11 within each ESN, then the end-to-end service will perform as expected. In this simple case, where all Root End Points are in one ESN and all Leaf End Points are in another, ESNC End Point 1.1 could also be set as a Root.

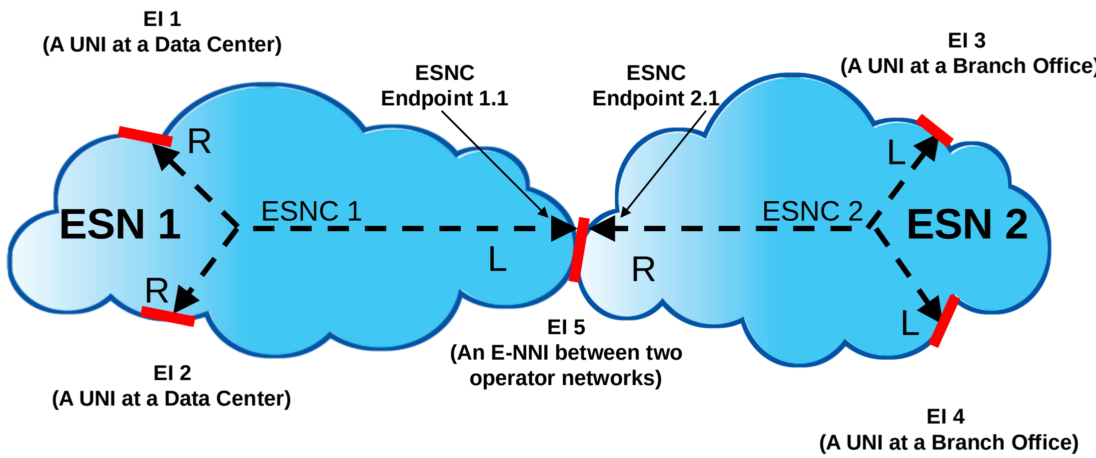

Figure 16 – A Simple Rooted-Multipoint EVC across two sub-networks

Figure 17 gives a view of the much more general case, where both ESNs contain a mixture of Root ESNC End Points and Leaf ESNC End Points. In this case, there is no way of setting the values for ESNC End Points 1.1 and 2.1 in such a way that the end-to-end service performs as expected. Looking at ESN 1, traffic that ingresses at EI 7 (a Leaf) must be deliverable to Root End Points in ESN 2. Therefore ESNC End Point 1.1 must be a Root. Looking at ESN 2 and applying the same argument, ESNC End Point 2.1 must be a Root. But now consider traffic that ingresses at EI 3 (a Leaf). It gets delivered out ESNC End Point 2.1 (a Root) and enters ESNC End Point 1.1 (A root). Since it came from a Root End Point, ESNC 1 is allowed to deliver it out EI 7 (a Leaf). In this way, traffic from a Leaf at EI 3 gets delivered to a Leaf at EI 7. This, of course, is not how the end-to-end service is supposed to behave.

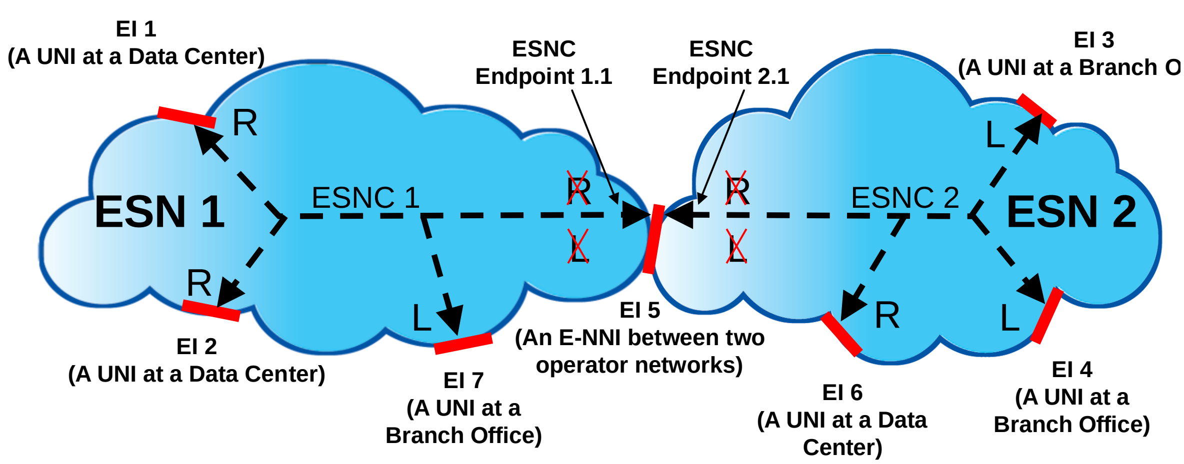

Figure 17 – First Example of Incorrect Rooted-Multipoint EVC Configuration

The fundamental problem is that when either ESNC receives EI Frames across EI 5, it has no way of knowing whether the EI Frame came from a Root ESNC End Point or a Leaf ESNC End Point. The EI Frames are not marked differently. As a result, it cannot know the ESNC End Points to which it is allowed to deliver this Frame.

One way to attempt to solve this problem is to use Multiple ESNC End Points at EI 5 for both ESNC 1 and ESNC 2. The intent is to allow the Frames to be marked differently depending on the originating ESNC End Point type. If, for example, EI 5 had an 802.1ad based frame format, the outer tag (S-tag) would distinguish the traffic. Figure 18 shows an attempt to use this approach.

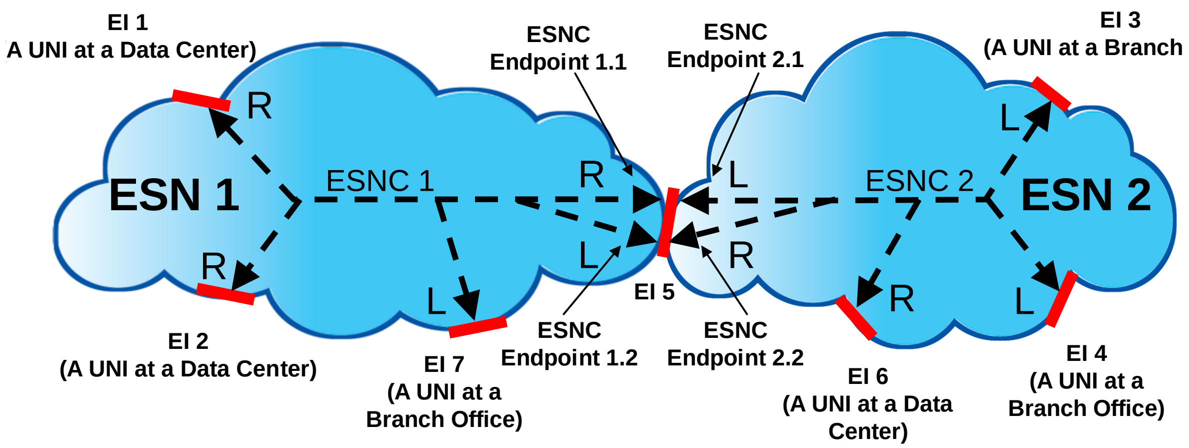

Figure 18 – Second Example of Incorrect Rooted-Multipoint EVC Configuration

This solution is close to solving the problem, but suffers from two further problems. First, traffic that originates at Root UNI End Points (e.g. EI 1) could be delivered to both the Root and Leaf ESNC End Points at EI 5 resulting in two copies of the frame being delivered to ESNC 2. Secondly, loops are a definite concern. An EI Frame arriving at ESNC End Point 1.1 (a Root) could be delivered by ESNC 1 to ESNC End Point 1.2 a Leaf, back across the same EI. The same would happen in ESNC 2 resulting in an endless loop.

This issue is solved by creating two new ESNC End Point Types with different allowed delivery. The new ESNC End Point Types added are From-Root, and From-Leaf to represent the original source of the Service Frame. Figure 19 shows the final solution with “FR” meant to represent a From-Root ESNC End Point and “FL” meant to represent a From-Leaf ESNC End Point. The allowed delivery between ESNC End Point Types is as shown in Table 5 and repeated in Table 12 below.

It is important that the EI Frame Map in each ESN be established to ensure that EI Frames are mapped to ESNC End Points of the same type. For example, in the figure below, the solution would not work if ESNC End Point 1.1 were mapped to ESNC End Point 2.2 across EI 5.

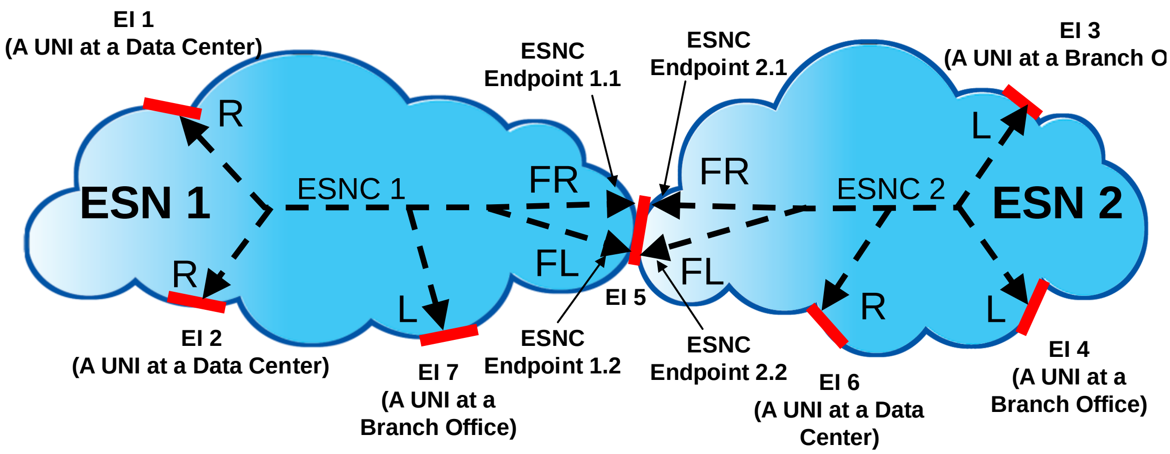

Figure 19 – Final ESC Model in a Rooted-Multipoint Service

		Ingress ESNC End Point Type			
		Root	Leaf	From-Leaf	From-Root
Egress ESNC End Point Type	Root	☑	☑	☑	☑
	Leaf	☑	☐	☐	☑
	From-Leaf	☐	☑	☑	☐
	From-Root	☑	☐	☐	☑

Table 12 – Allowed Connectivity between ESNC End Point Types

Applying this behavior to ESNC 1 in ESN 1 in Isolation, and ESNC 2 in ESN 2 in isolation, it is be clear that the end-to-end service performs as expected.

12 Appendix B: Examples of the use of Tunnel Endpoints (informative)

This section presents several examples of the use of Tunnel Endpoints as described in Section 7 to achieve EI to EI Ethernet Services in complex cases where Operators employ the use of tunneling for business and management reasons. The first sub-section establishes the conventions and notation used in the examples. The remaining sub-sections present the examples. Understanding these examples help support those writing Child Documents to use the Tunnel Endpoint Construct consistently.

12.1 Notation and Conventions

"Tunneling", in the context of this document, can be broadly defined as the bundling of multiple connectivity services into a single connectivity service that can then be managed as a unit without a detailed understanding of the services inside. It is employed by Operators of telecom networks for a number of business reasons. For example:

- Operations Cost Reduction and Scale: The global telecommunications network carries an immense number of UNI to UNI connectivity services. The notion that every Operator be aware of the details of every UNI to UNI connectivity service does not scale operationally. For example, consider a point-to-point EVC that crosses four Operator networks with no tunneling involved in the end-to-end path of the service. To protect their network, each Operator applies a bandwidth profile on this EVC at the contracted service rate. In this case, a change to the rate requires a coordinated configuration change in all four operator networks. In cases where there are thousands of EVCs that follow a similar path through these four Operator networks, the cost of such changes can become prohibitively expensive. By contrast, suppose the Operators on the two ends "bundle" all of these services into a tunnel running through the two middle operators and the size of the tunnel is adjusted on a quarterly basis to stay ahead of the expected aggregate capacity of the EVCs inside. In this case, bandwidth profile changes for individual EVCs can be executed without involving the two middle Operators.

- Cost Optimization: It is very common that the cost per MB/s of bandwidth decreases as the size of a particular circuit grows. For example, the monthly recurring cost for a 150Mb/s circuit will typically be much less than that for fifteen 10 Mb/s circuits. Service Providers looking to buy wholesale circuits from other operators will often take advantage of this by bundling together a set of services with the same source and destination allowing them to purchase one large service instead of many small ones.

- Technology Limitation: A given Operator can have connectivity to areas desired by the Service Provider but the technology used by this Operator does not have the necessary multiplexing capability needed to distinguish the EVCs that need to be provisioned. This technology limitation can be overcome by bundling (multiplexing) and unbundling (de-multiplexing) in the ESNs connected to this Operator.

The following sections examine a number of examples of how tunneling might be used across Operator Networks. In each of these examples, the way that the TE service construct is used to describe the externally observable behavior of the tunnel is shown.

In the diagrams, the following drawing conventions shown in Figure 20 are used. In order to show examples of EI Frame Maps and TE Frame Maps, it is necessary to refer to a specific value of a specific field in the EI Frame. Since this document leaves the specification of the EI Frame format to the child document, a generic notation is used. It has the form of a Greek letter followed by an integer. The Greek letter refers to a value while the integer refers to a field. For example, $\alpha 1$ means that field 1 has value α. To make it clear the extent of the Tunnel in each example, those ESNCs that carry traffic between Tunnel End Points are shown as heavy dashed lines.

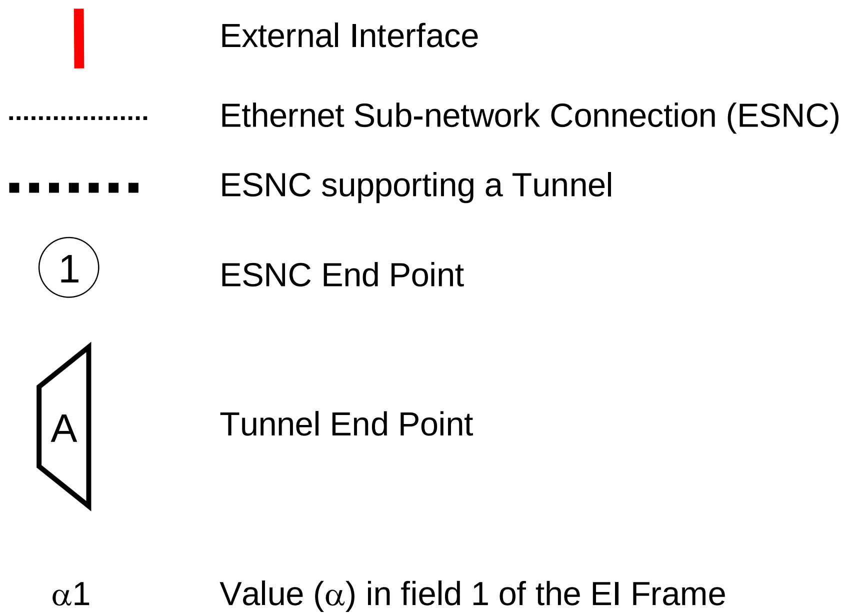

Figure 20 – Drawing Icons used in the Examples

12.2 Example 1: Transit Tunnel

Figure 21 shows an example of a Transit Tunnel. This example is called a Transit Tunnel because there are paired Tunnel End Points outside of ESN B but ESN B connects the two Tunnel End Points. Thus the tunnel can be thought of as transiting ESN B.

The frame maps are shown for interfaces A/B and B/C. The mapping of an EI Frame to an ESNC End Point in ESN B is a one step process using the value in field 1 of the EI Frame. At interface A/B, the value α in field 1 maps to ESNC End Point 5 while at interface B/C, the value β in field 3 maps to ESNC End Point 6. Note that it is not necessary for the EI Frame to have the same format at different interfaces and it is not necessary to use the same fields at different interfaces even if the formats are the same.

The mapping of an EI Frame to an ESNC End Point in ESN A and ESN C is more complex. In ESN A, an EI Frame at interface A/B with value α in field 1 maps to the Tunnel End Point A. The TE Frame Map for Tunnel End Point A is then used to map the EI Frame to an ESNC End Point based on the value in field 2. In ESN C, an EI Frame at interface B/C with value β in field 3 maps to the Tunnel End Point C. The TE Frame Map for Tunnel End Point C is then used to map the EI Frame to an ESNC End Point based on the value in field 4.

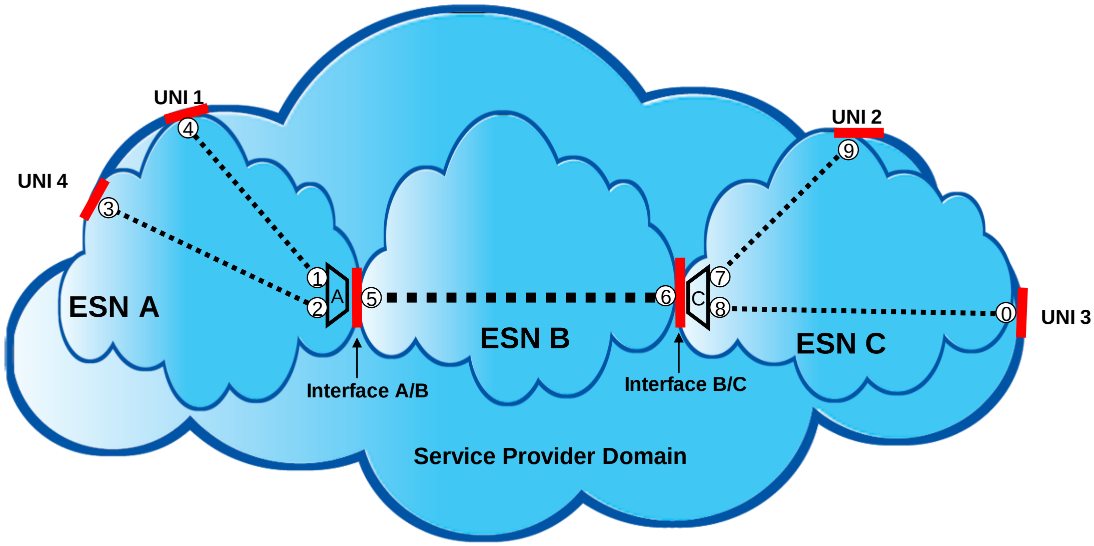

	EI A/B				EI B/C			
	ESN A		ESN B		ESN B		ESN C	
EI Frame Maps	A	α1	**5**	α1	6	β3	C	β3

TE Frame Maps	1	δ2		7	π4
	2	ε2		8	σ4

Figure 21 – Example of a Transit Tunnel

A Transit Tunnel can transit multiple ESNs. An example of two ESNs being transited is shown in Figure 22.

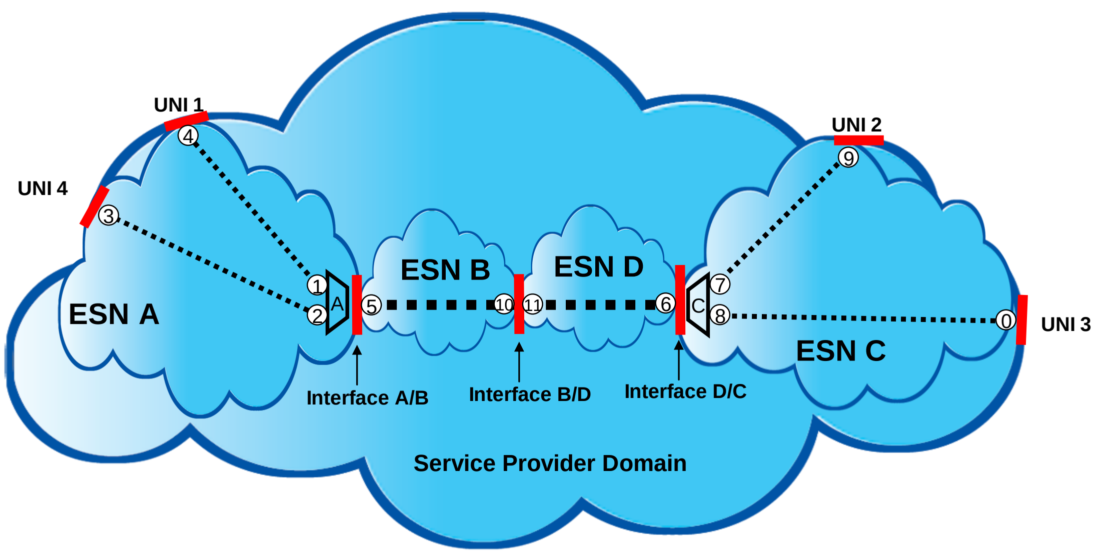

Figure 22 – Example of a Transit Tunnel with Two Intermediate ESNs

12.3 Example 2: Multipoint Tunnel

Figure 23 shows an example of a Multipoint Tunnel. This example is similar to the Transit Tunnel example in Section 12.2 with the difference that the ESNC in ESN B associates 3 ESNC End Points instead of two and there are three Tunnel End Points, one each in ESN A, ESN B, and ESN C.

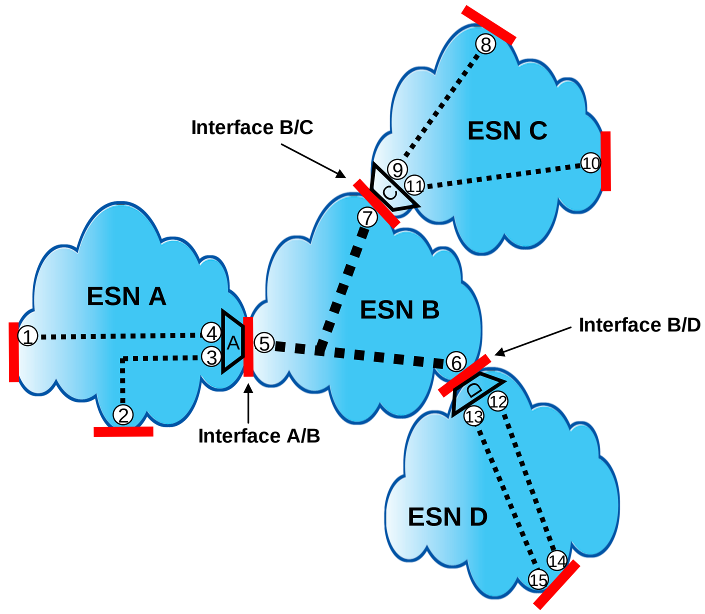

	EI A/B		EI B/C		EI B/D	
	ESN A	ESN B	ESN B	ESN C	ESN B	ESN D
EI Frame Maps	A | $\alpha1$	5 | $\alpha1$	7 | $\beta1$	C | $\beta1$	6 | $\omega1$	D | $\omega1$
TE Frame Maps	3 | $\delta2$ 4 | $\pi2$			9 | $\lambda2$ 11 | $\tau2$		12 | $\rho2$ 13 | $\sigma2$

Figure 23 – Example of a Multipoint Tunnel

As in the case of the Transit Tunnel, the Multipoint Tunnel can span multiple ESNs. An example is shown in Figure 24.

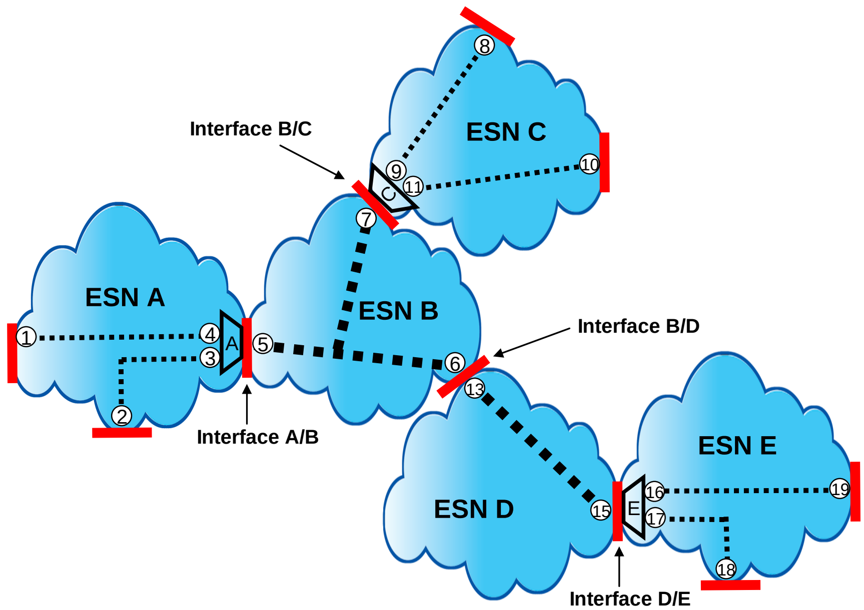

Figure 24 – Example of a Multipoint Tunnel Spanning Two ESNs

13 References

[1] Bradner, S., Key words for use in RFCs to Indicate Requirement Levels, RFC 2119, March 1997. (Normative)

[2] 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. (Normative)

[3] IEEE Std 802.3ae – 2002, IEEE Standard for Information technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific requirements – Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications Amendment: Media Access Control (MAC) Parameters, Physical Layers, and Management Parameters for 10 Gb/s Operation, 13 June 2002. (Normative)

[4] International Telecommunication Union, Recommendation I.370, Integrated Services Digital Network (ISDN), Overall Network Aspects And Functions, ISDN User-Network Interfaces, Congestion Management For The ISDN Frame Relaying Bearer Service, 1991.

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

[6] IEEE Std 802.1D – 2004, IEEE Standard for Local and Metropolitan Area Networks: Media Access Control (MAC) Bridges, June 2004.

[7] IEEE Std 802.1Q – 2005, IEEE Standards for Local and metropolitan area networks—Virtual Bridged Local Area Networks, 19 May 2006.

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

[9] IEEE Standard 802.1ad – Provider Bridges, Draft 6, August 2005.

[10] Metro Ethernet Forum MEF 10.2, Ethernet Service Attributes Phase 2, October 2009.

[11] Metro Ethernet Forum MEF 7, EMS-NMS Information Model, October 2004.

[12] International Telecommunication Union, Recommendation I.370, Integrated Services Digital Network (ISDN), Overall Network Aspects And Functions, ISDN User-Network Interfaces, Congestion Management For The ISDN Frame Relaying Bearer Service, 1991.

[13] 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.

[14] Metro Ethernet Forum MEF 12, Metro Ethernet Network Architecture Framework – Part 2: Ethernet Services Layer, April 2005.

[15] IEEE Std 802.1AX – 2008, IEEE Standard for Local and metropolitan area networks — Link Aggregation, 3 November 2008.

[16] IEEE Std 802.3ba – 2010, IEEE Standard for Information Technology – Telecommunications and information exchange between systems – Local and metropolitan area networks – Specific Requirements – Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications – Amendment 4: Media Access Control Parameters, Physical Layers, and Management Parameters for 40 Gb/s and 100 Gb/s Operation, June 2010.

[17] 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.