



Technical Specification

MEF 12.1

Carrier Ethernet Network Architecture Framework

Part 2: Ethernet Services Layer - Base Elements

April 15, 2010

Disclaimer

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

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

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

- a) any express or implied license or right to or under any patent, copyright, trademark or trade secret rights held or claimed by any MEF member company which are or may be associated with the ideas, techniques, concepts or expressions contained herein; nor
- b) any warranty or representation that any MEF member companies will announce any product(s) and/or service(s) related thereto, or if such announcements are made, that such announced product(s) and/or service(s) embody any or all of the ideas, technologies, or concepts contained herein; nor
- c) any form of relationship between any MEF member companies and the recipient or user of this document.

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

© The Metro Ethernet Forum 2010. All Rights Reserved.

Table of Contents

1. Abstract	1
2. Terminology	1
3. Scope.....	3
4. Compliance Levels.....	3
5. The ETH Layer Model.....	4
5.1 ETH Layer Characteristic Information	5
5.1.1 Connection Indications and VLAN IDs.....	6
5.1.2 CoS Indications, UP/PCP and DEI.....	6
5.1.3 EtherType, User Data and ETH_CI.....	6
5.2 ETH Layer Functional Model: Topological Components & Transport Entities.....	6
5.2.1 ETH Layer Network	7
5.2.2 ETH Flow Domain.....	7
5.2.3 ETH Link	8
5.2.4 ETH Access Group	9
5.3 ETH Sublayers, Aggregation and the IEEE Bridging Model	9
5.3.1 Topological Representation of IEEE 802.1D Bridged Network	10
5.3.2 Topological Representation of IEEE 802.1Q Bridged Networks	10
5.3.3 ETH Layer and the MEF Service Model	10
5.4 ETH Connection and ETH Connection Segments.....	11
5.4.1 EC Aggregation and Sublayering in the ETH Layer	12
5.4.2 EC Roles: Subscriber EC, Operator ECs and Service Provider EC	13
5.4.3 Relationship between ECs and EVCs/OVCs.....	15
6. Reference Models for the ETH Layer Connections.....	18
6.1 EC Types	18
6.1.1 Point-to-Point ECs.....	19
6.1.2 Multi-Point ECs.....	19
6.1.3 Rooted Multi-Point EC	20
6.1.3.1 Trunk EC and Branch EC	21
6.2 Hairpin EC	23
6.3 Tunnel ECs	25
7. ETH Layer Processing Functions.....	25
7.1 APP to ETH Adaptation Function (EAF)	26
7.2 ETH Flow Adaptation Function (EFAF).....	27
7.3 ETH Flow Termination Function (EFTF).....	27
7.4 ETH Conditioning Functional Elements.....	27
7.4.1 ETH Flow Conditioning Function (EFCF).....	28
7.4.2 ETH Subscriber Conditioning Function (ESCF).....	28
7.4.3 ETH Provider Conditioning Function (EPCF).....	29
7.5 ETH EC Adaptation Function (EEAF).....	29
7.6 ETH EC Termination Function (EETF).....	30
7.7 ETH Connection function (ECF).....	30
7.8 ETH EC Interconnect Function (EEIF)	31
7.9 ETH to TRAN Adaptation function (TAF).....	31
8. Base ETH Layer Functional Element Models.....	32
8.1 ETH UNI	33
8.2 ETH ENNI.....	34

9. ETH Layer Interface Extensions and Their Functional Elements	34
10. References.....	35
Appendix I: ETH Layer Model for Service OAM and the Base EIs.....	36
I.1 Definitions and Modeling Conventions.....	36
I.1.1 SOAM Reference Model	36
I.1.2 Mapping between IEEE and ITU Terminology	37
I.1.3 Representing MEPs and MIPs	37
I.2 SOAM Functional Model for the ETH Layer	38
I.2.1 OAM Components at the UNI	38
I.2.1.1 UNI MEG.....	38
I.2.1.2 Subscriber MEG	38
I.2.1.3 EVC MEG.....	39
I.2.1.4 Operator MEGs at the UNI.....	39
I.2.2 OAM Components at the ENNI	39
I.2.2.1 ENNI MEG	40
I.2.2.2 Operator MEG.....	40
I.2.2.3 EVC MEG at the ENNI.....	40

List of Figures

Figure 1: Base ETH Layer Interfaces and Reference Points.....	4
Figure 2: ETH Layer Characteristic Information (ETH_CI)	5
Figure 3: Representation of an ETH Layer Network and its topological components	7
Figure 4: Topological representation of selective broadcast within an ETH Flow Domain	8
Figure 5: MEF Service view from the ETH Layer perspective	11
Figure 6: Example of a Subscriber Centric (Topological) Representation of ECs	12
Figure 7: Example of a Service Provider/Network Operator Centric (Topological) Representation of ECs.....	12
Figure 8: Topological representation of ECSs aggregation into an EC	13
Figure 9: Example of Subscriber EC (S-EC), Operator EC (O-EC) and Service Provider EC (SP-EC) in a multi-CEN environment	14
Figure 10: Relationship between EVC, OVCs, Subscriber EC, Operator EC(s) and Service Provider EC	15
Figure 11: Example of Topological Representations of Point-to-point SP-ECs and associated O-ECs	19
Figure 12: Example of Topological Representations of a Multipoint SP-EC and Associated O-ECs.....	20
Figure 13: Topological representation of a portion of a RMP EC in a single CEN.....	22
Figure 14: Root/Trunk O-EC.....	23
Figure 15: Point-to-Point and Multipoint Hairpin SP-ECs.....	24
Figure 16: Diagrammatic representation of ETH Layer processing entities.....	26
Figure 17: Functional representation of an ETH UNI-N and ETH UNI-C.....	33
Figure 18: Functional representation of an ETH ENNI-N.....	34
Figure I.1: SOAM Framework (derived from MEF 17, Figure 5).....	36
Figure I.2: Diagrammatic representation of OAM processing entities in the ETH Layer	37
Figure I.3: Functional representation of MEPs & MIPs Placement at an UNI.....	38
Figure I.4: Functional representation of MEPs & MIPs Placement at an ENNI.....	40

List of Tables

Table 1: Acronyms and Definitions.....	2
Table 2: Summary of ETH Layer Topological Components, Transport Entities and Reference Points.....	6
Table 3: Relationship between Services and Architecture constructs.....	17
Table 4: Allowed Connectivity between TFP types	21
Table I.1: Mapping between IEEE and ITU Terminology	37

1. Abstract

This document provides the architectural framework to model the Ethernet Services Layer of MEF compliant Carrier Ethernet Networks. It also introduces the base functional elements of the MEF architecture. The Ethernet Service Layer architecture framework describes the high-level topological and functional constructs used to model the various architectural components of subscriber and provider networks, their associated functional elements, and their interconnect relationships. The architecture framework also describes the relationship between Ethernet Services Layer functional elements and their reference points, and other architectural elements in the Transport Layer (TRAN) and Application (APP) Layers of the MEF Generic Architecture Framework (MEF 4 [17]).

2. Terminology

This section summarizes terms and acronyms used in this document. In many cases, the normative definitions to terms are found in other documents. In these cases, the third column is used to provide the reference that is controlling. In cases of conflict with other documents, the controlling document is shown in the reference column

Terms	Definitions	Reference
AF	Adaptation Function	MEF4
APP	Application Layer	MEF4
B-FP	Branch Flow Point	This Document
BR-FP	Branch-Root Flow Point	This Document
BWP	Bandwidth Profile	MEF10.2
C-VLAN	Customer VLAN	IEEE 802.1
CE	Customer Edge	MEF10.2
CEN	Carrier Ethernet Network	This Document
CI	Characteristic Information	MEF4
CoS	Class of Service	MEF10.2
CoS ID	Class of Service Identifier	MEF10.2
DEI	Discard Eligibility Indicator	IEEE 802.3
E-NNI	External Network-to-Network Interface ¹	MEF4
EAF	ETH Adaptation Function	MEF4
EC	ETH Connection	This Document
ECF	ETH Connection Function	This Document
ECS	EC Segment	This Document
ECT	Ethernet Connectionless Trail/Ethernet Connection-oriented Trail	MEF4
EETF	ETH EVC Termination Function	This Document
EFCF	ETH Flow Conditioning Function	This Document
EFD	ETH Flow Domain	This Document
EFTF	ETH Flow Termination Function	This Document
EI	External Interface	MEF4
ENNI	External Network-to-Network Interface ¹	MEF26
ENNI-N	ENNI – Network (Functional Element)	MEF26
EPCF	ETH Provider Conditioning Function	This Document
ESCF	ETH Subscriber Conditioning Function	This Document
ETF	ETH Termination Function	This Document
ETH	Ethernet Services Layer	MEF4
EtherType	Ethernet Length/Type	IEEE 802.3

¹ Note that MEF4 uses “E-NNI” as the acronym for the same interface.

EVC	Ethernet Virtual Connection	MEF10.2
FCS	Frame Check Sequence	IEEE 802.3
FP/FPP	Flow Point/Flow Point Pool	ITU-T G.809
H-FP	Hairpin Flow Point	This Document
II	Internal Interface	MEF4
L2CP	Layer Two Control Protocols	MEF10.2
L2DP	Layer Two Data Plane	This Document
LAN	Local Area Network	IEEE 802.1
LLC	Logical Link Control	IEEE 802.1
MAC	Media Access Control	IEEE 802.1
MCF	MAC Convergence Function	IEEE 802.1
MEG	Maintenance Entity Group	ITU-T Y.1731
MEL	Maintenance Entity Level	ITU-T Y.1731
MEN	Metro Ethernet Network	MEF4
MTP	Multipoint	This Document
NE	Network Element	MEF4
NNI	Network-Network Interface	MEF4
O-EC	Operator EC	This Document
OOF	Out of Franchise	This Document
P2P	Point-to-point	MEF10.2
PDU	Protocol Data Unit	This Document
PE	Provider Edge	This Document
RMP	Rooted Multipoint	MEF10.2
S-EC	Subscriber EC	This Document
S-VLAN	Service VLAN (also referred to as Provider VLAN)	IEEE 802.1
SLS	Service Level Specification	MEF10.2
SP-EC	Service Provider EC	This Document
T-FP	Trunk Flow Point	This Document
TAF	Transport Adaptation Function	MEF4
TF	Termination Function	MEF4
TFP	Termination Flow Point	MEF4
TL-FP	Trunk Leaf Flow Point	This Document
TRAN	Transport Layer	MEF4
UNI	User-Network Interface	MEF4
UNI-C	UNI Client	MEF4
UNI-N	UNI Network	MEF4
UP/PCP	User Priority/Priority Code Point	IEEE 802.3
VLAN	Virtual LAN	IEEE 802.1
VLAN ID	VLAN Identifier	IEEE 802.3
VPLS	Virtual Private LAN Service	RFC 4761
WAN	Wide Area Network	MEF4
WEN	Wide Area Ethernet Network	MEF4

Table 1: Acronyms and Definitions

3. Scope

The Ethernet Services Layer, also referred to as the ETH Layer, is the networking layer in the MEF multi-layer architecture framework responsible for the realization of service-aware transport capabilities in support of MEF specified services (see MEF4 [17]).

The ETH Layer architecture framework provides the descriptive model to represent Ethernet service constructs in terms of topological and functional constructs. These architectural and functional constructs are intended to provide a technology neutral decomposition with respect to the TRAN Layer, of the logical and physical connectivity between Subscribers, the Service Provider(s), and any intermediate Network Operators. The connectivity model is described in terms of the ETH Layer functional elements, their relationships to ETH Layer processing and transport entities used to implement the MEF interface functions, and the interconnect rules among them. The networks from Service Providers and Network Operators supporting the MEF service and architecture models are referred to as Carrier Ethernet Networks (CENs).

The architecture framework also describes the interactions of the ETH Layer with the TRAN and APP Layers across MEF specified External Interfaces (EIs), such as the User-Network Interface (UNI) and the External Network-Network Interface (ENNI), and Internal Interfaces (IIs), such as the Internal Network-to-Network Interfaces (I-NNIs), and their associated reference points. The ETH Layer architecture framework is not intended to require or preclude any particular networking technology from being used on any given implementation of a CEN but it does presume a minimal set of Ethernet networking functions. The framework provides a generic interconnect model and information transfer guidelines to facilitate the specification of interoperable ETH Layer components at relevant EIs and IIs conforming to this architecture model. Detailed Technical Specifications and Implementation Agreements for specific EIs/IIs are outside the scope of this document.

The initial MEF ETH Layer specification, MEF12 [22], focused on architecture constructs in support of point-to-point (E-Line) and multipoint (E-LAN) Ethernet Virtual Connection (EVC) services within a single Service Provider network as specified in the base Ethernet Service Model (MEF6 [18]), its Service Attributes (MEF10 [20]), and subscriber service interface (MEF13 [23]). This update introduces topological and functional constructs, and their interconnect rules, in support of rooted multi-point EVC services as specified in MEF6.1 [19] and MEF10.2 [21]. It also introduces architecture constructs in support of inter CEN Operator services. New definitions and constructs introduced in this document are:

1. Carrier Ethernet Network (CEN) as a generalization of the Metro Ethernet Network (MEN) and Wide-Area Ethernet Network (WEN) concepts
2. ETH Connection (EC) and EC Segment (ECS), and associated Termination Flow Points (TFPs) and Flow Points (FPs), as architecture constructs to represent the connectivity between functional elements in the ETH Layer
3. Topological and functional models in support of the External Network-Network-Interface (ENNI) Phase I (MEF26 [27]). Multi-CEN capabilities introduced in the ENNI Phase I include:
 - support E-Line and E-LAN and E-Tree services,
 - support “hairpin” switching (i.e., forwarding of Ethernet frames across separate connections on the same physical link) for E-Line and E-LAN and services,
 - support “EVC tunneling” (i.e., service unaware forwarding of frames) across another CEN.

This document also:

- Clarifies the relationship between the EVC construct as defined in MEF10.2, the Operator Virtual Connection (OVC) construct defined in MEF26 and the EC/ECS constructs defined in this document.
- Expands on any other impacts of this document to existing MEF specifications.

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. The ETH Layer Model

The Ethernet Services Layer, or ETH Layer, refers to the Ethernet networking layer defined by the MEF to specify Ethernet oriented connectivity services. The ETH Layer is responsible for the service view presented by the Service Provider to its Subscribers and CENs participating in the Ethernet service. The ETH Layer is also responsible for all service-aware processing aspects associated with the treatment of the Service Frames, including operations, administration, maintenance and provisioning capabilities required in support of the Ethernet connectivity services. This document refers generically as a CEN to any MEN or WEN as defined in MEF4.

The service interface presented by the ETH Layer at the EI reference points is expected to conform to the frame format for unicast, multicast or broadcast Ethernet frame as specified in IEEE 802.3 [2]. Figure 1 illustrates the administrative relationship between the base service-interfaces, or Base EIs, defined in the MEF Generic Architecture Framework (MEF 4) at the ETH Layer: the User Network Interface (UNI) and the External Network-to-Network Interface (ENNI). The ETH Layer processing functions associated with the instantiation of a UNI within a Subscriber network is referred to as an ETH UNI-C. The ETH Layer processing functions associated with the instantiation of a UNI within a CEN Operator is referred to as an ETH UNI-N. The ETH Layer processing functions associated with the instantiation of an ENNI within a CEN Operator is referred to as an ETH ENNI-N.

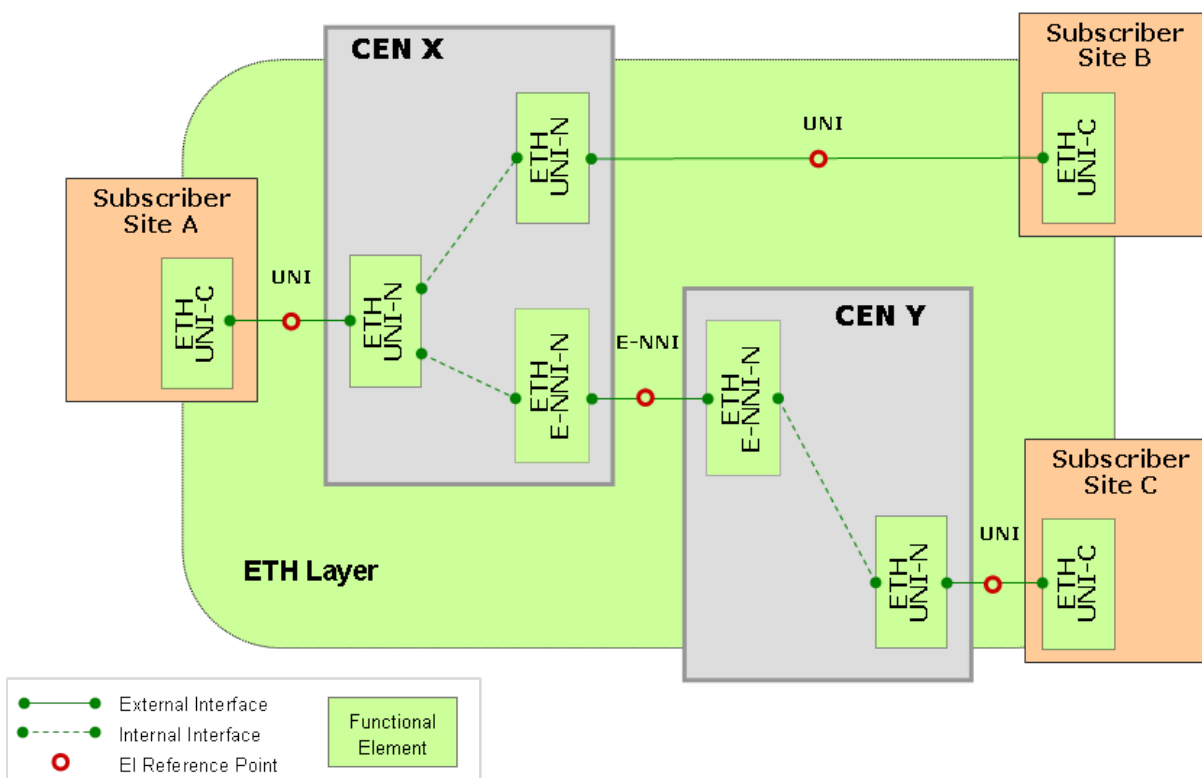


Figure 1: Base ETH Layer Interfaces and Reference Points

The rest of this section relates the architectural model of the ETH Layer to its service model specification in terms of topological and transport entities at relevant reference points, such as those illustrated in Figure 1. The functional view of the ETH Layer described in this document is derived from the functional modeling approach described in MEF4 (which is derived from ITU-T Recommendations G.805 [9] and G.809 [10]). The traffic units exchanged among these reference points are discussed in Section 5.1. From a functional modeling viewpoint, the ETH Layer consists of topological, transport and processing entities. The relevant ETH Layer topological entities and associated transport entities are introduced in Section 5.2. Their relationship with the various Ethernet bridging models and MEF defined connection types are discussed in Section 5.3 and Section 5.4, respectively.

A detailed reference model for MEF ETH Layer in terms of these architecture constructs is provided in Sections 6, 7 and 8. The detailed reference models are based on a functional modeling description of the various CEN components. In particular, the relationship between the Ethernet Service components, such as EVCs and OVCs, and the transport entities are described in Section 6. The ETH Layer processing entities used to instantiate these transport entities are described in Section 7. Functional Element representations for key service interfaces defined by the MEF are introduced in Section 8 and Section 9. The detailed descriptions of functional elements are intended to be consistent with ITU-T Recommendations G.8010/Y.1306 [11] and G.8021 [12].

5.1 ETH Layer Characteristic Information

Functional models in ITU-T Recommendations G.805/G.809 refer to the intrinsic information elements exchanged over a layer network as the Characteristic Information (CI) of the layer network.

For the ETH Layer, the characteristic information exchanged over the ETH Layer links, the ETH_CI, consists of the following information elements from the IEEE 802.3 Ethernet MAC frame:

- Destination MAC Address (DA)
- Source MAC Address (SA)
- Optional Connection Indication via the IEEE 802.1Q VLAN Tag
- Optional CoS Indication via the IEEE 802.1Q VLAN Tag
- Optional Discard Eligibility Indication via the IEEE 802.1Q S-VLAN Tag
- Ethernet Length/Type (EtherType)
- User Data

The optional 802.1Q VLAN Tag is a four octet field composed of a two-octet 802.1Q VLAN Tag Type and the Tag Control Information, which contains 3-bits of User Priority/Priority Code Point (UP/PCP) information, the single-bit Canonical Format Indicator/Discard Eligibility Indicator (DEI) and the 12-bit VLAN Identifier (VLAN ID). Note that the VLAN ID, DEI and UP/PCP are also optional information elements in any IEEE 802.3-2005 Ethernet MAC frame. Figure 2 illustrates the ETH_CI as conveyed in IEEE 802.3 compliant Ethernet MAC frames.

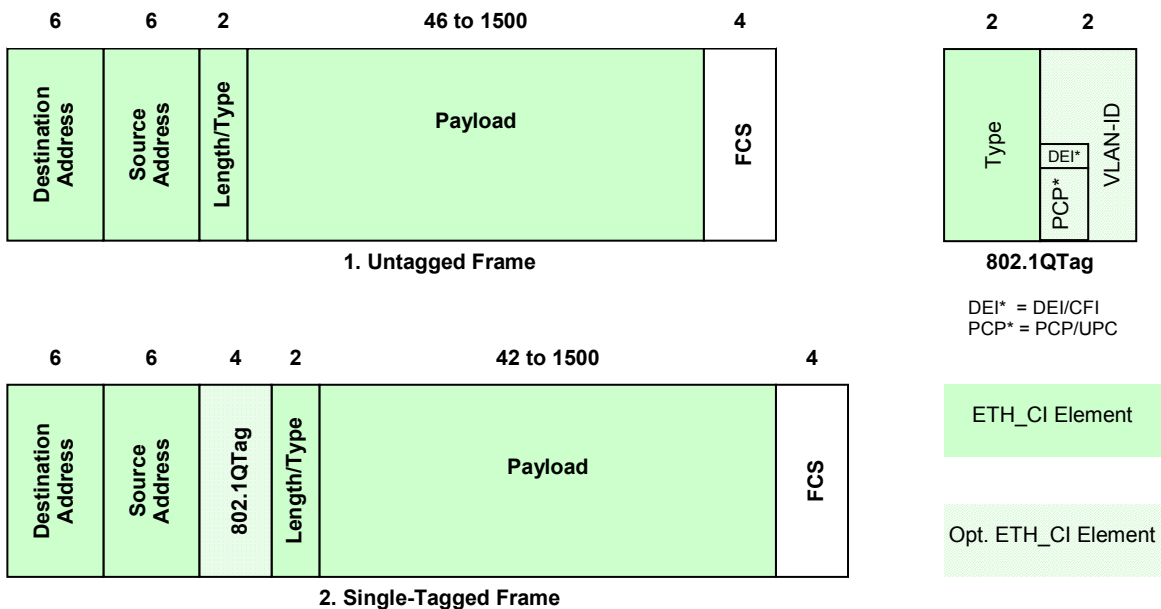


Figure 2: ETH Layer Characteristic Information (ETH_CI)

The term ETH PDUs is used in this document to refer to the frames used to exchange the ETH_CI across standardized ETH Layer interfaces and associated reference points. In particular, the Service Frame refers to the ETH PDU exchanged across the UNI. Operations on Service Frames to support the ETH Layer service model are specified in

MEF 10.2 [21]. Note, as illustrated in Figure 2, that the information elements conveyed by ETH PDU is a super set of the information elements of the ETH_CI.

5.1.1 Connection Indications and VLAN IDs

For the purposes of the MEF4/G.809 functional model, identifiers associated with layer network flows/connections are network specific, and thus, not part of the CI. VLAN IDs specifically are viewed as internal constructs of the ETH Layer that assist with network partitioning (connections) and sub-layering (aggregation) functions but not as formal elements of the ETH_CI. Nonetheless, many networking scenarios, including the connection services specified by the MEF, require VLAN IDs to be preserved. In those scenarios, VLAN IDs **MUST** be treated as if part of the ETH_CI by intermediate CENs.

5.1.2 CoS Indications, UP/PCP and DEI

CoS is deemed to be as an inherent characteristic of a flow/connection by the MEF4/G.809 functional model; hence, it is part of the ETH_CI. Yet, CoS indication via UP/PCP and DEI, or any other CoS ID propagation mechanism, is considered optional. Nonetheless, many networking scenarios, including the connection services specified by the MEF, require CoS-IDs to be preserved. In those scenarios, CoS IDs, **MUST** be treated as if part of the ETH_CI by intermediate CENs.

The CoS associated with a given ETH PDU is service specific. As specified by MEF10.2 and MEF26, the CoS of an ETH PDU at an EI may be conveyed implicitly via information derived from the fields in the ETH Layer (e.g., VLAN IDs), or APP Layer (e.g., DSCP) or TRAN Layer link identifiers. CoS may also be conveyed explicitly as in the UP/UPC field of Service Frames at UNIs and ENNIs. MEF23 specifies the MEF CoS model [26].

5.1.3 EtherType, User Data and ETH_CI

The Length/Type field, or EtherType field, in conjunction with SA/DA, is part of the mechanism specified by IEEE 802.3 to indicate the type of payload conveyed as User Data. User Data PDUs are typically conveyed in their format, or less common, as a Logical Link Control (LLC) encapsulated PDU. The presence of VLAN tags is indicated via the EtherType field. For functional modeling purposes, each VLAN tag corresponds to an ETH Sublayer within the ETH Layer (see Section 5.3).

5.2 ETH Layer Functional Model: Topological Components & Transport Entities

As described in MEF4, topological components are the abstract entities used by the model to represent networked (connectivity) entities. Transport entities represent the physical or logical entities created by functional elements for the purpose of transferring the ETH_CI as implied by the topological constructs. Table 2 summarizes the defined ETH Layer topological components, transport entities, and reference points in the ITU-T Recommendation G.8010 and the MEF Architecture Framework (MEF 4).

Topological Components:	Transport Entities:	Associated Reference Points
ETH Layer Network	ETH Network Flow / Connection	ETH Access Point
ETH Flow Domain	ETH Flow Domain Flow / Subnetwork Connection	ETH Flow Point ETH Flow Point Pool
ETH Link ²	ETH Link Flow / Connection	ETH Flow Point ETH Flow Point Pool
ETH Access Group	ETH Connectionless Trail	ETH Termination Flow Point ETH Termination Flow Point Pool

Table 2: Summary of ETH Layer Topological Components, Transport Entities and Reference Points

In this document the terms flow point and flow point pool are used interchangeably to refer to either one flow point or a set of flow points treated as a single unit.

² This document uses ETH Link as a short hand for ETH Flow Point/Flow Point Pool Link in MEF4.

Figure 3 illustrates the relationships among ETH Layer topological components from an Ethernet Service Subscriber viewpoint. There are three LAN sites (LAN Site A, LAN Site B and LAN Site C) each represented as an ETH Flow Domain. The three LAN sites are fully interconnected via ETH Links (ETH Link AB, ETH Link BC and ETH Link AC). In this particular example there are several end-systems (access points) on each LAN (5 on Site A, 3 on Site B and 4 on Site C).

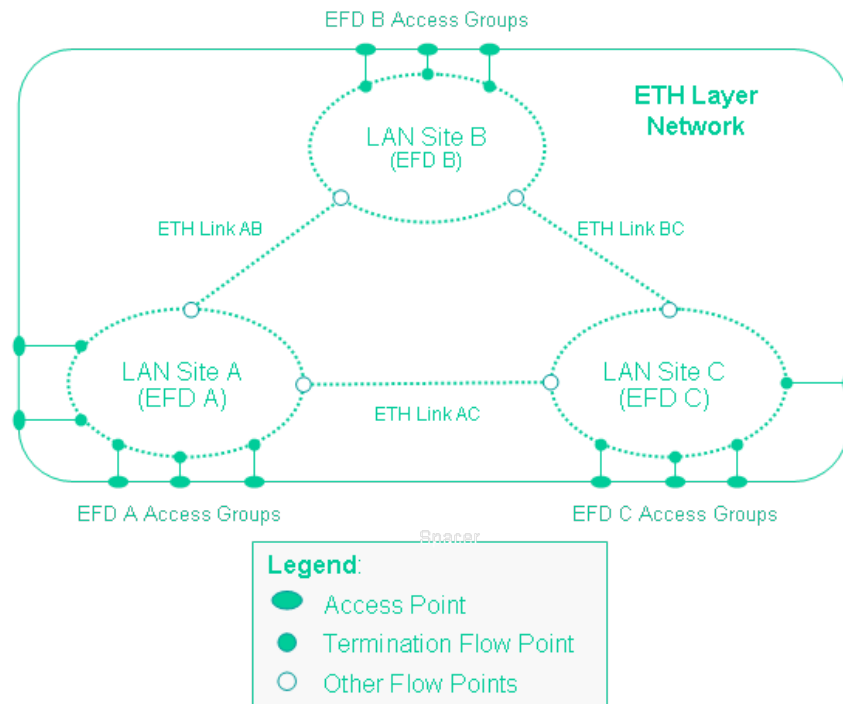


Figure 3: Representation of an ETH Layer Network and its topological components

5.2.1 ETH Layer Network

The ETH Layer network is bounded by the complete set of ETH access groups that may be associated for the purpose of transferring information among access points. The scope of the ETH Layer network is the broadcast domain of all access groups that can be addressed by IEEE 802.1D/802.1Q MAC forwarding entities. The unit of information transferred within the ETH Layer is the ETH_CI. The provisioning of a transport entity in the ETH Layer creates an association between two or more access points and enables the exchange of ETH_CI. The transport entity can be either a connectionless transport entity referred to as the Ethernet Connectionless Trail or connection oriented transport entity referred to as an Ethernet Connection-oriented Trail (ECT) depending on the expected forwarding behavior for said transport entity. The resulting topology of the ETH Layer network can be described in terms of the set of all ETH access groups, the ETH flow domains and the ETH Links³ or ETH Server Subnetworks between them. This ETH Layer network view is illustrated in Figure 3. As illustrated there, this document uses a “rounded rectangle” symbol as a topological representation of an Ethernet Layer Network and an “ellipse” symbol as a topological representation of an EFD within an Ethernet Layer Network.

5.2.2 ETH Flow Domain

An ETH Flow Domain (EFD) is a topological component of the ETH Layer formally defined by the set of ETH flow points, including termination flow points, made available for the purpose of transferring information within a given administrative portion of the ETH Layer network. The transport entities associated with flow points on a given EFD are referred to as EFD Flows (see G.809) or ETH subnetwork connections (see G.805). EFDs may be partitioned

³ Also referred to as a Flow Point Pool Link in G.809

into sets of non-overlapping EFDs interconnected by ETH Links. A switching “matrix” represents the smallest instance of an EFD. An example of a switching matrix is the ETH Layer function associated with the representation of the switching components of the IEEE 802.1D/802.1Q MAC Relay function.

The scope of the EFD transport entity is the selective broadcast domain of the set of associated ETH (termination) flow points as illustrated in Figure 4. Before MAC address learning occurs, ETH PDUs received via an input port (e.g., Link XA) of EFD X and associated with the selective broadcast domain H are forwarded to all output ports on EFD X associated with this domain. In the example, FPs named FP XH* (e.g., FP XH1, TFP XH2, FP XH3 and FP XH4) denote the FPs in the selective broadcast domain H of EFD X. The one exception in this forwarding rule is with the FPs in the same bi-directional Link as the input port (e.g., FP XH1). The connectivity among ETH (termination) flow points in an EFD can be restricted by means of ETH network management or control plane actions. ETH transport entity is commonly referred to as a VLAN.

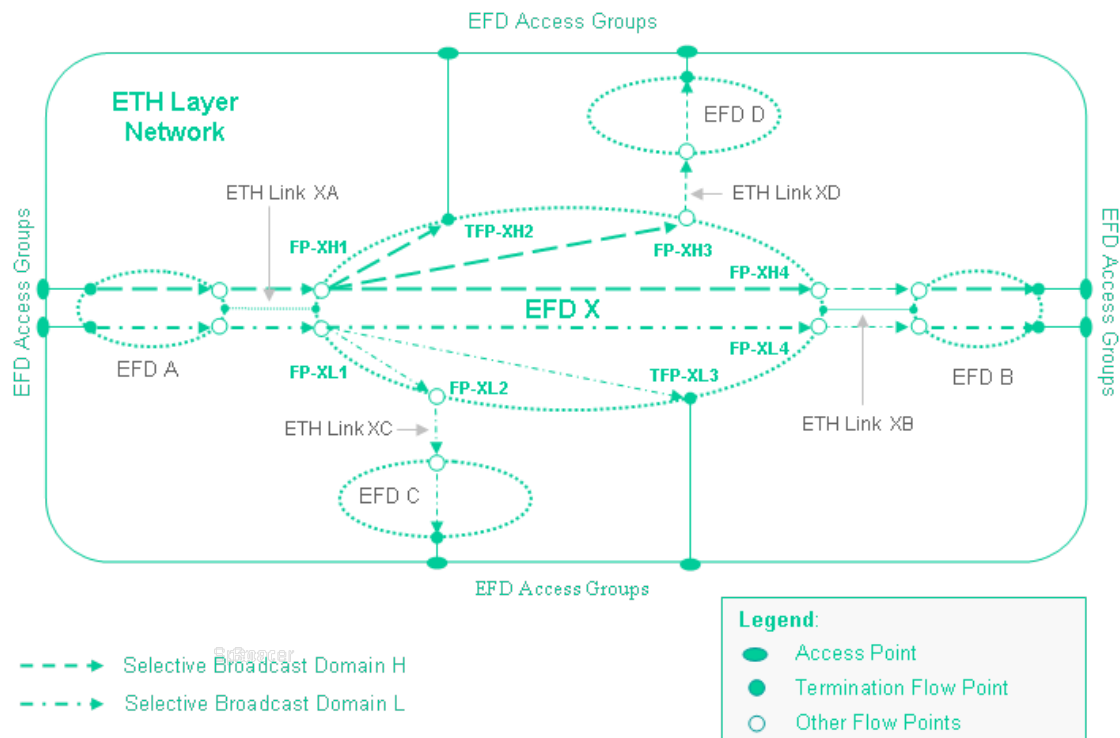


Figure 4: Topological representation of selective broadcast within an ETH Flow Domain

There may be multiple EFDs within the ETH Layer of a particular CEN. At the coarsest level of granularity an EFD represents a particular administrative domain within the ETH Layer. Examples include:

- EFDs associated with the Subscriber networks
- EFDs associated with multiple CENs when an EVC traverses one or more ENNI.
- EFDs associated with different administrative/operational/maintenance boundaries within the same CEN.

5.2.3 ETH Link

An ETH Link, also referred to as ETH flow point pool link (ETH FPP link) in ITU-T G.809, represents a fixed topological relationship and available capacity between a set of ETH reference points. Thus, an ETH Link may interconnect:

- two sets of ETH flow points (i.e., FPPs) at the edge of two EFDs, or
- a set of ETH termination flow points (ETH access groups) and a set of ETH flow points (i.e., FPPs) at the edge of an EFD, or
- two ETH access groups

The transport entity associated with two flow points on a given ETH Link is also referred to as an ETH Link connection. In some situations it is useful to call out ETH Links used between EIs and ETH Links used between IIs. In MEF4 the ETH Links at an EI are also referred to as access links⁴ while ETH Links at the IIs are referred to as trunk links.

ETH Links between EFDs are typically long lived and established at the timescale of the server Layer Network connections rather than the timescale of the individual flows transported at the ETH Layer. Hence, an ETH Link represents a non-switchable entity within the particular Layer Network such as a fiber or any other “connection” created by mechanisms outside the given Layer Network (such as MPLS, ATM or FR connections). As illustrated in Figure 3 and Figure 4, this document uses the symbol of a line between (T)FPs or between a (T)FP and Access Groups to represent an ETH link.

It is also possible to create a point-to-point transport entity from a serial concatenation of two or more ETH Links. Such concatenation of ETH Links is also referred to as a Serial Compound ETH Link.

5.2.4 ETH Access Group

An ETH Access Group is a group of co-located ETH flow termination functions that are connected to the same EFD or EFD link. The ETH Access Group demarcates the point of access into the ETH Layer network.

Hence, an ETH access group represents the “service access point” for applications into the ETH Layer Network. As illustrated in Figure 3 and Figure 4, this document uses the symbol of a solid ellipse to represent an Access Group.

5.3 ETH Sublayers, Aggregation and the IEEE Bridging Model

Packet switching technology employs packet encapsulation techniques to create “partitions” of network resources for the purposes of flow identification, traffic multiplexing, aggregation and overall network engineering. For packet switching technologies that allow multiple levels of encapsulation each new encapsulation tag is said to represent a sublayer of the given layer network within the context of the MEF4/G.809 Layer Network model. Although each sublayer has independent OAM there are interdependences between sublayers as they share the Access Groups’ address space and potentially control and management functions.

An encapsulation sublayer can be viewed as a distinct domain of the layer network, from the viewpoint of the ISO reference model. In addition, each individual flow/connection identifier value conveyed in the encapsulation header (or tag) represents a transport entity of the given sublayer and provides the means to identify individual logical/virtual transport entities of the sublayer for the purposes of traffic forwarding and network resource allocation. These transport entities⁵ of the Layer Network also provide restricted forwarding domains for the Layer Network PDUs. These domains and transport entities of a Layer Network constitute the primary mechanisms for logical connectivity management and traffic engineering within a given Layer Network.

Note that each encapsulation sublayer is intended to represent a new domain of the same Layer Network technology, not a new Layer Network in its own right. From the viewpoint of the architecture framework in MEF4 an encapsulation sublayer is modeled in the same manner as its corresponding Layer Network technology except that its access points represent sublayer encapsulation flow points (also represented as Termination Flow Points), not access groups. This document uses the term “aggregation” to refer to the process associated with the instantiation of a new sublayer and the term “multiplexing” to refer to the process associated with the instantiation of a new layer network.

Sublayering mechanisms are technology specific and not all packet switching technologies support sublayering. At an EI the transport entities of a layer network may be aggregated into a sublayer transport entity or multiplexed into a server layer transport entity. In Provider Bridged Ethernet, the C-VLAN is specified as the transport entity distinguishing customer traffic. It is also the basis for the transport entities at the UNI. Similarly, the S-VLAN is specified as the transport entity distinguishing aggregated C-VLANs. It is also the basis for the transport entities at the ENNI. Other type of transport entities may be specified in future EI specifications. Specification of encapsulation sublayering technology for any specific EI is outside the scope of this specification.

⁴ Not to be confused with links of an Access Network.

⁵ Referred to as a “flow domain fragment” in ITU-T G.8010.

5.3.1 Topological Representation of IEEE 802.1D Bridged Network

An IEEE 802.1D bridged network is represented as a sublayer of IEEE 802.1 bridged network model consisting of an “arbitrary” collection of EFDs and links as illustrated in Figure 3 and Figure 4. Each EFD in the topological model represents a collection of one or more IEEE 802.1D Bridges interconnected via ETH Links (i.e., links to other IEEE 802.1D based LANs) and ETH Access Groups (end-systems) in the ETH Layer Network. This is the simplest topological representation of an Ethernet Network, and hence, a Subscriber Network at the ETH Layer.

For topological modeling purposes an IEEE 802.1D bridged network is considered a single multipoint transport entity supporting a number of flows identified by their Source and Destination MAC addresses. There is a single transport entity allowed on any topological component of IEEE 802.1D networks. It is still possible to create a point-to-point transport link as a concatenation of one or more point-to-point ETH Links, e.g., a serial compound ETH Link (see 5.2.3).

5.3.2 Topological Representation of IEEE 802.1Q Bridged Networks

An IEEE 802.1Q bridged network is also represented as an “arbitrary” collection of EFDs and ETH Links within each sublayer in the IEEE 802.1Q bridge network model. Two sublayers can be identified, one associated with the C-component (or C-VLANs) of IEEE 802.1Q, and one associated with the S-Component (or S-VLANs) of IEEE 802.1Q. For instance, in the C-component sublayer, each EFD in the topological model can represent a collection of IEEE 802.1Q Bridges interconnected via ETH Links (i.e., links to other IEEE 802.1Q LANs) and ETH T-FPs (end-systems) in the C-VLAN based ETH Layer Network.

Similarly, an IEEE 802.1 Provider Bridge network [5] is represented as an “arbitrary” collection of EFDs and ETH Links in the IEEE802.1 Provider Bridge network (i.e., S-Component) sublayer of the IEEE 802.1Q Layer Network. Each EFD can represent a collection of IEEE 802.1 Provider Bridges interconnected via ETH Links (i.e., links to other IEEE 802.1 Provider Bridge LANs) and ETH Access Groups/T-FPs (end-systems) in the S-VLAN based ETH Sublayer. Note that any two VLAN IDs may be associated with separate physical entities, or with the same physical entity.

The C-Component of IEEE 802.1Q (a.k.a. VLAN Bridging) and the S-Component of IEEE 802.1 Provider Bridge are the base sublayering technologies used by the UNI and ENNI. A number of other technologies are available to create ETH Sublayers, including IEEE 802.1 Provider Backbone Bridge [7] and IETF VPLS[30][31]. Although the topological representations of these ETH Sublayers are equivalent to those in Figure 3 and Figure 4 their processing entities are specific to each ETH Sublayer technology. Specification of the applicable frame format on MEF specified IIs and EIs are outside the scope of this document.

5.3.3 ETH Layer and the MEF Service Model

From the ETH Layer model perspective the MEF Service Model, such as in MEF10.2 for UNI services and MEF26 for ENNI services, defines the ETH (Sub)Layer used to described the behavior of the exchanged ETH PDUs at a specified EI/II. The Service Model behaviors are described as a set of Service Attributes from the viewpoint of an external observer able to monitor the exchange of ETH PDUs among the involved parties (subscriber or CENs). The abstract vantage point for the specification of Ethernet Services at a given EI/II is referred to as the EI Reference Point.

The flow points/flow point pools of the ETH Layer model are related to reference points on *functional components* associated with the topological representation of a network. The functional representation of the ETH Layer consists of a set of processing functions in the abstract network equipment that implement the desired topological configuration. This contrasting view of the Carrier Ethernet Network is depicted in Figure 5.




Figure 5: MEF Service view from the ETH Layer perspective

Transport entities are created and interconnected by the functional entities of the ETH Layer for the purposes of handing-off the ETH_CI specified in the Ethernet Service definition and creating an end-to-end connection among the networks participating in the ETH Layer service. The next section discusses the representation of the generic transport entities from Section 5.2 as applied to the ETH Layer.

5.4 ETH Connection and ETH Connection Segments

The ETH (Sub)Layer transport entity created in order to convey ETH_CI between a set of two or more ETH Layer termination flow points in a given ETH Sublayer is generically referred to in this document as an ETH Connection (EC). An EC can be further decomposed as an interconnected (concatenated) set of segments delimited by two or more flow points. These segments are generically referred to in this document as ETH Connection Segments (ECS). ECs can be either unidirectional or bidirectional,

An EC can correspond to either:

- an end-to-end network connection in the ETH Layer (also referred to as a network flow), or
- an edge-to-edge subnetwork connection across an EFD (also referred to as FP/FPP connection) or
- an edge-to-edge link connection between two EFDs (also referred to as a FP/FPP link connection)

An ECS can correspond to either:

- an edge-to-edge link connection, or
- an edge-to-edge subnetwork connection.

As in the case of ETH Links in Section 5.2.3, it is also possible to create an EC from a serial concatenation of two or more ECs. Such concatenation of ECs is more formally referred to in this document as a compound EC. The Link EC between any two CENs is also an O-EC instance.

Figure 6 illustrates a generic representation of a C-VLAN based ETH Sublayer from a Subscriber centric perspective (i.e., when representing the CEN as a “transparent cloud”). It shows two C-VLAN tagged flows among two Subscriber LANs (Site A and Site B). Here, EC A1B1 and EC A2B2 represent the two Subscriber flows, indicated by two non-overlapping C-VLAN IDs, crossing two ETH flow domains, EFD A and EFD B, representing IEEE 802.1Q capable Subscriber LANs. EC A1B1 consists of three serially compound ECSs: 1) ECS A1A3 between the TFP-A1 and FP-A3 on EFD A, 2) ECS B3B1 between the FP-B3 and TFP-B1 on EFD B, and 3) ECS A3B3 between FP-A3 and FP-B3 on the ETH Link between EFD A and EFD B. Note that the TFs depicted the figure are not part of the topological view of the network, they are presented there to highlight the point where the ECs were created.

A similar representation would be used to represent an S-VLAN based ETH Sublayer when the Subscribers are two CENs that contract an Ethernet Service via an intermediate Ethernet Service Provider. In this case the ECs would correspond to the S-VLAN tagged flows among the two CENs. The EFDs would represent the CEN Operators when assuming their networks are implemented via IEEE 802.1 Provider Bridge.

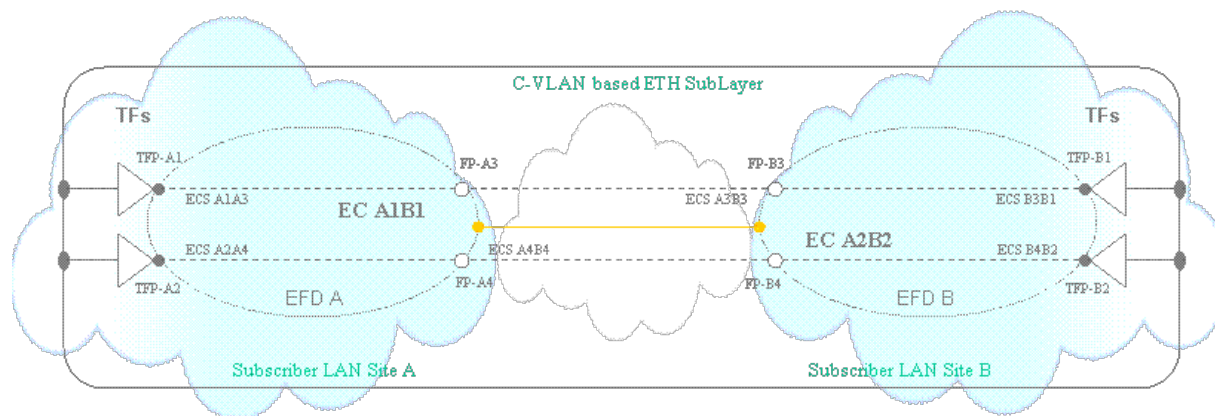


Figure 6: Example of a Subscriber Centric (Topological) Representation of ECs

Figure 7 illustrates a generic representation of a C-VLAN based ETH Sublayer from a Service Provider/Network Operator centric perspective (i.e., when representing the CEN as an “opaque cloud”). It shows the same two C-VLAN tagged flows among two Subscriber LANs (Site A and Site B). Here, the Service Provider depicts the Subscriber ECS A3B3 as actually consisting of three serial segments, ECS A3A5, ECS A5B5 and ECS B5B3 related to the connection between Site A and the SP, the connection segment across the SP and the connection between the SP and Site B, respectively.

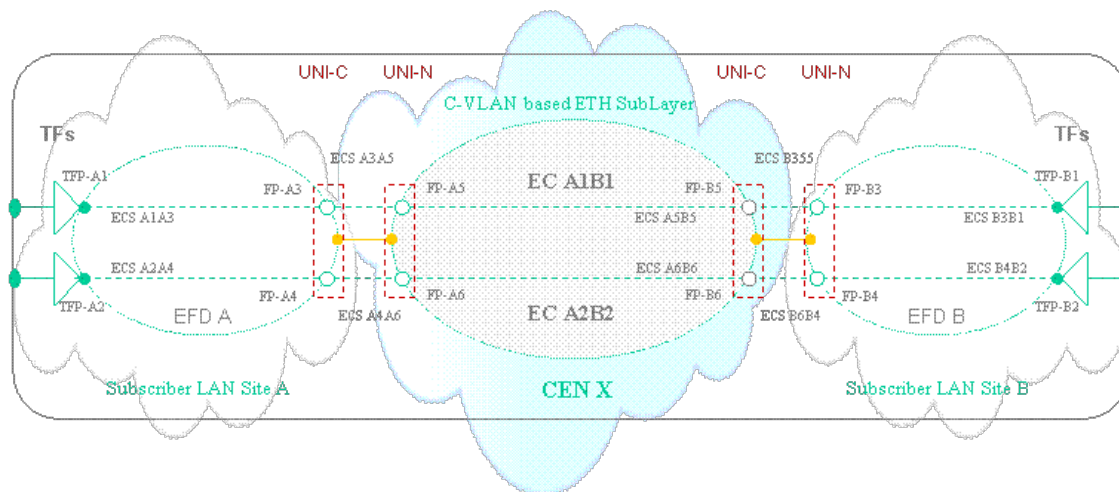


Figure 7: Example of a Service Provider/Network Operator Centric (Topological) Representation of ECs

5.4.1 EC Aggregation and Sublayering in the ETH Layer

An EC in a given ETH Sublayer can be used to convey two or more aggregated ECSs from ECs in another ETH Sublayer. The aggregated ECSs are said to be in the *client* ETH Sublayer while the new EC conveying the aggregated flow is said to be in the *server* ETH Sublayer. ECS aggregation is achieved via an ETH Adaptation Function (AF). Figure 8 illustrates a generic representation of ECS multiplexing into an EC when incorporating both the Subscriber and the Service Provider/Network Operator view of the involved ECs. Here, the representation of EC A1B1 and EC A2B2 are expanded to include the Service Provider/Network Operator specific segments of the EC. ECS

A1A3 and ECS A2A4 on EFD A (correspondingly ECS B1B3 and ECS B2B4 on EFD B) are aggregated into EC X1X2 on the Service Provider EFD X. Note that the TFs and AFs depicted the figure are not part of the topological components of the network. They are shown here to highlight the points where the ECs are created and aggregated.

Note that in practice AFs are specific to the server layer technology used to carry the subscriber flows and the encapsulation format of any intermediate ETH Sublayer. As an example, the AF used to aggregate C-VLAN tagged flows onto an S-VLAN based EC is different from the AF used to aggregate C-VLAN tagged flows onto a non-Ethernet (e.g., MPLS LSP, SDH VC, OTN ODU) connection. Similarly, the AF used to aggregate a C-VLAN tagged flow onto an S-VLAN based EC is different from the AF that uses an I-SID tag, and it is different from the AF that uses any other type of tag plus a MAC header with Type, SA and DA fields (e.g., VPLS). Relevant distinctions will be noted in the EI and I-NNI specific implementation agreements and models.

Also note that in different implementation scenarios the same Ethernet technology (e.g., IEEE 802.1 Provider Bridge), may be used to create transport entities in the MEF ETH Layer or transport entities in the MEF TRAN Layer. Although the high-level topological representation from the data plane viewpoint may appear to be the same the distinction in the usage becomes more evident when the associated operational, administrative and maintenance entities are explicitly shown in the functional representation.

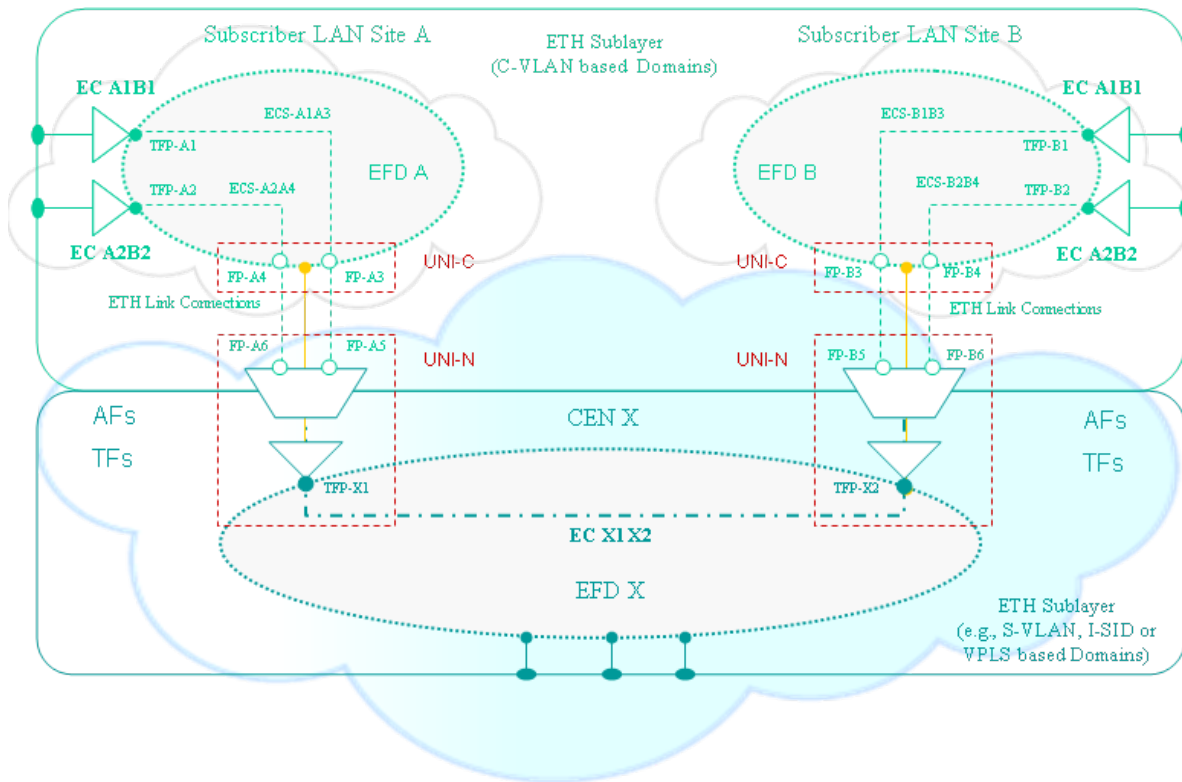


Figure 8: Topological representation of ECSs aggregation into an EC

An EC, and hence, any associated ECS', may convey multiple CoS instances. For example, at a UNI the PCP field in a C-tagged or priority tagged Service Frame may be used to indicate multiple CoS instances associated with a C-tagged flow as per MEF10.2. Subscribers and CENs would typically invoke traffic engineering considerations when aggregating ECS'. Such considerations are Subscriber and CEN (and SLS) specific and outside the scope of this document.

5.4.2 EC Roles: Subscriber EC, Operator ECs and Service Provider EC

When describing MEF services delivered to Subscribers and the mechanisms used by CENs to support those services it is useful to be able to refer generically to either the Subscriber or CEN flows irrespective of the particular tech-

nology used to instantiate the relevant Ethernet frame flows. In this specification the term Subscriber EC (S-EC) is used to refer generically to an ETH Layer flow instantiated by a Subscriber. The term Operator EC is used to refer generically to an ETH Layer flow instantiated by a Network Operator. These flows can correspond to link flow/connections (e.g., between EFDs), flow domain/subnetwork connections (e.g., across an EFD) or network connections (e.g., across a network) introduced in Section 5.2.

In a multi-CEN situation, a Subscriber EC is typically conveyed over a number of Operator ECs, at least one for each traversed CEN. There is also an Operator EC used to interconnect the two CENs, also referred here to as a Link EC. Such multi-CEN scenario arises in situations where a SP contracts OVC from multiple Network Operators to offer a multi-CEN EVC, or when an Operator EC is carried over one or more Operator ECs, e.g., a carrier's carrier service. The concatenation of the O-ECs used to support a single Subscriber or CEN service instance is referred to in this document as a Service Provider EC (SP-EC). Note that a SP-EC can be from any EI to any EI as long as there are 1 or more O-ECs and 1 or more Link ECs. Thus, a SP-EC is not necessarily UNI-N to UNI-N. An example of this is given in Section 6.3 (Tunnel EC).

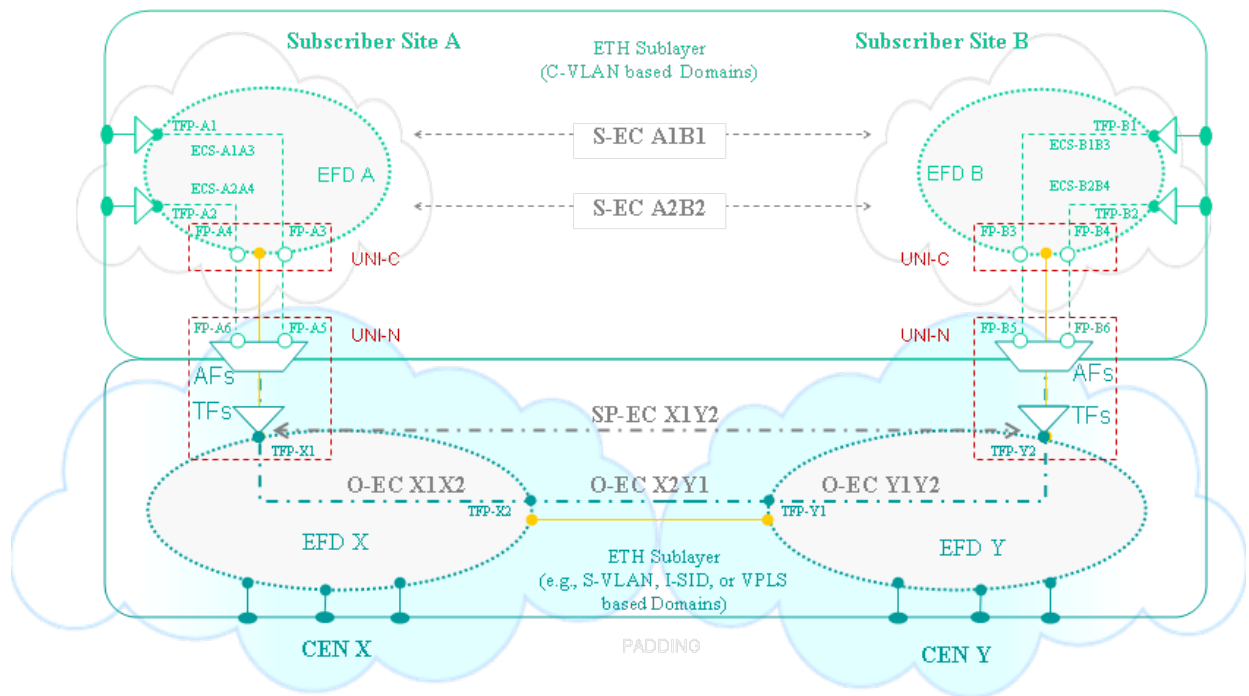


Figure 9: Example of Subscriber EC (S-EC), Operator EC (O-EC) and Service Provider EC (SP-EC) in a multi-CEN environment

Figure 9 illustrates the use of S-ECs, O-ECs and SP-ECs in a multi CEN context. The Subscriber has two S-ECs, S-EC A1B1 and S-EC A2B2, that traverse two CENs, CEN X and CEN Y. In this example the lead Service Provider instantiates an SP-EC, SP-EC X1Y2, to convey the Service Frames from Site A to Site B. This SP-EC is created from three O-ECs:

- an O-EC within CEN X, O-EC X1X2
- an O-EC within CEN Y, O-EC Y1Y2
- an O-EC interconnecting CEN X to CEN Y, O-EC-X2Y1

The EC model here generalizes to any arbitrary number of intervening CENs and Subscriber sites. Note that in some scenarios it may be possible to emulate O-ECs via a TRAns Layer connection. Such an emulation is Network Operator specific and outside the scope of this document.

5.4.3 Relationship between ECs and EVCs/OVCs

An EVC is a service construct used to associate UNI Reference Points, and hence, enable the creation of end-to-end subscriber services instances across one or more CENs as specified by the MEF services definitions (e.g., MEF 6.1). An OVC is a service construct used to associate at least one ENNI with other ENNIs, and UNIs, and hence, enable the creation of Network Operator-oriented services instances across one or more CENs.

From the ETH Layer architecture viewpoint each tagged⁶ Ethernet flow constitutes a separate EC. An EVC indicates the association of one or more instances of Subscriber ECs among a number of UNIs in an untagged or C-VLAN based ETH Sublayer. Within a CEN these Subscriber ECs may be switched as is, or, more typically, may be associated with one or more Operator ECs with a sublayer instance of the ETH Layer. An OVC is represented as an association of one or more instances of Operator ECs among a number of EIs. These O-EC instances include the Link EC between any two CENs. Hence, the concatenation of Operator ECs used to convey an EVC, or an OVC, or any aggregation of EVC or OVCs (but not both), for the purposes of delivering a single instance of an Ethernet Service is referred to as a Service Provider EC. The relationship between an EVC, its associated Subscriber EC(s), Operator EC(s), Service Provider EC and any underlying OVC(s) is illustrated in Figure 10.

Note that a CEN may be composed of two or more internal CENs. For instance, a Service Provider network may consist of a number of independent MENs interconnected over a Wide-Area Ethernet network (WEN) acting as a backbone. Thus, the an O-EC traversing the Service Provider network may consist of a number of concatenated O-ECs across each of the independent domains constructed in the same fashion as Figure 10.

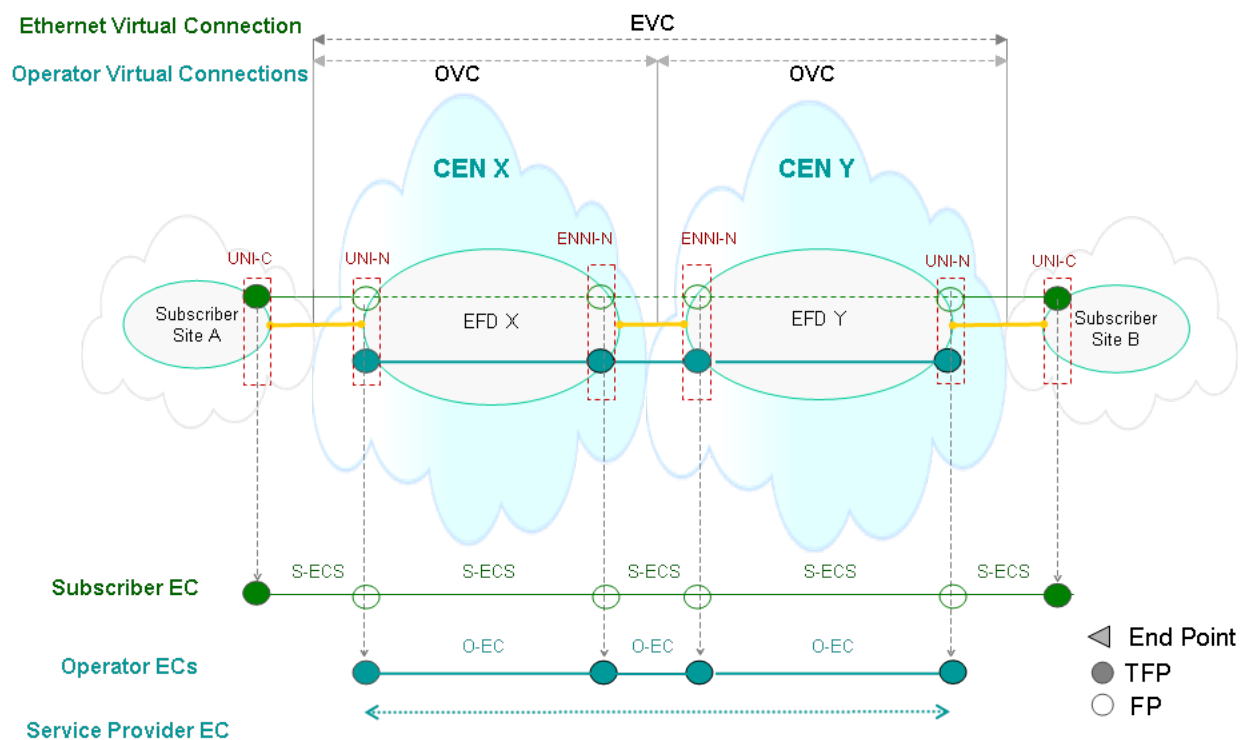


Figure 10: Relationship between EVC, OVCs, Subscriber EC, Operator EC(s) and Service Provider EC

⁶ Untagged and priority tagged Ethernet frames are modeled as being “virtually” tagged.

The following requirements apply to Subscriber ECs:

- a) There **MUST** be at least one Subscriber EC associated with an EVC.
- b) There **MUST** be FPs located at two or more UNI-C in a Subscriber EC⁷.
- c) There **MAY** be more than one Subscriber EC associated with an EVC (see Bundling Service Attribute in MEF 10.2).
- d) There **MAY** be one or more CoS Instances associated with a Subscriber EC (e.g., see BWP per EVC and CoS ID Service Attributes in MEF 10.2).

These requirements underline the 1:1 relationship between a CE VLAN-ID and a S-EC, the 1:n relationship between EVC/OVC and S-ECs (and hence, between EVC/OVC EPs and S-EC FPs), and the 1:n relationship between S-EC and CoS Instances.

The following requirements apply to Operator ECs:

- a) There **MUST** be at least one EVC or OVC associated with an Operator EC.
- b) There **MUST** be FPs located at two or more EIs that delimit an Operator EC⁸.
- c) There **MAY** be one or more CoS Instances associated with an Operator EC.
- d) Service Frames from an EVC or OVC associated with one or more Operator ECs **MUST** be forwarded such that the instantiated Operator EC(s) is consistent with the selective broadcast domain for the given EVC within each Operator domain.

Thus, these requirements underline the n:1 relationship between EVCs/OVCs and O-EC, and between CE-VLAN-IDs and O-ECs. In particular, a separate O-EC could exist for each CE-VLAN-ID bundle. An example of multiple EVCs associated with a single Operator EC is one where the O-EC is used to realize a tunnel service that conveys the ETH PDUs associated with the EVCs between a UNI and an ENNI.

The following requirements apply to Service Provider ECs:

- a) There **MUST** be at least one EVC or OVC associated with a Service Provider EC⁹.
- b) There **MUST** be TFPs located at two or more EIs that delimit a Service Provider EC.
- c) There **MAY** be one or more CoS Instances associated with a Service Provider EC.
- d) Service Frame from an EVC associated with a Service Provider EC **MUST** be forwarded such that the instantiated Service Provider EC supports the selective broadcast domain for the given EVC.

Note that in a case of an UNI-to-UNI Ethernet Service instantiated over a single CEN the Service Provider EC consists of a single Operator EC. For Ethernet Service instantiated over multi-CENs, the Service Provider EC consists of two or more O-ECs and one of these O-ECs is a Link EC.

Mapping EVCs and OVCs to ECs is service specific. Nonetheless, rules used to map Subscriber ECs, and their associated EVCs, to Operator ECs at a UNI need to comply with the MEF Services Model (MEF 10.2). Rules used to map Subscriber ECs or Operator ECs, and their associated OVCs, to Link ECs at ENNIs need to comply with the ENNI Technical Specification (MEF 26). Rules used to map Subscriber ECs or Operator ECs, and their associated EVCs and OVCs with Service OAM MEGs need to comply with the Service OAM Framework Technical Specification – MEF17 [24]. Table 3 shows plausible relationships between MEF Service constructs, such as EVCs, OVCs and Tunnels, and ETH Layer architectures constructs as depicted in Figure 10.

Although there are no requirements on the mechanisms used by CENs to implement Operator ECs, it is expected that at a minimum an S-VLAN based ETH sublayering mechanism is supported at ENNIs as well as any other EIs/Is that may be defined that require support of an EVC tunneling or aggregation mechanism.

⁷ In scenarios where the S-EC boundaries are located at the UNI-C, these FPs actually correspond to TFPs.

⁸ Note that for an Operator EC across an EI, a Link EC, there are exactly two EIs delimiting the O-EC.

⁹ An example of multiple EVCs or OVCs associated with a single SP-EC is one where the SP-EC is used to realize a tunnel O-EC across one or multi-CEN service.

Administrative Entity		Service Entity	Service Entity Reference Points	Associated Architecture Entities	Associated Functional Elements	Associated SOAM MEG	
Subscriber		EVC	UNI	Subscriber ECS ¹	UNI-C	Subscriber	
Service Provider		EVC	UNI	Subscriber ECS ¹	UNI-N	EVC ¹	
		EVC	UNI	SP-EC ²	UNI-N	SP ²	
		Tunnel-1	ENNI & UNI	Tunnel EC ³	ENNI-N & UNI-N	SP ³	
		Tunnel-2	ENNI	Tunnel EC ⁴	ENNI-N	SP ⁴	
Network Operator		Single	EVC	UNI	Subscriber ECS ⁵	UNI-N	EVC ⁵
			EVC	UNI	Operator EC ⁶	UNI-N	SP ⁶
			EVC	UNI	Operator EC ⁶	UNI-N	Operator ⁶
			Tunnel-2	ENNI	Tunnel EC ⁴	ENNI-N	Operator ⁴
		Multi	OVC	ENNI & UNI	Subscriber ECS ⁷	ENNI-N & UNI-N	Operator ⁷
			OVC	ENNI & UNI	Operator EC ⁷	ENNI-N & UNI-N	Operator ⁷
			OVC	ENNI	Operator EC ⁸	ENNI-N	Operator ⁸
		UNI	UNI	ETH Link	UNI-N & UNI-C	UNI	
		ENNI	ENNI	ETH Link	ENNI-N	ENNI	

Table 3: Relationship between Services and Architecture constructs.

The following observations apply to Table 3:

1. The Subscriber ECS is the only transport entity required to traverse all EIs associated by an EVC. As such, it is the only suitable transport entity to be associated with the Subscriber MEG and the EVC MEG when using a Shared MEG Level model (see Y.1731).
2. EVCs are usually instantiated via one or more Operator ECs that create a Service Provider EC across one or more CENs. In use cases where there is a one-to-one mapping between the Service Frames of an EVC and such SP-EC, the Service Provider could use a MEG with the same scope as the SP-EC (referred to here as the Service Provider MEG) to monitor the EVC.
3. A Tunnel-1 based service (see Section 6.3) allows a Service Provider to limit visibility of EVCs at IIs and EIs of intermediate Network Operators via a Tunnel EC between an associated UNI and ENNI. A Service Provider can use a SP MEG within the scope of the Tunnel EC to monitor a Tunnel-1 based service. Note that a Tunnel-1 based service involves more than a single CEN. Each intermediate CEN is responsible for the instantiation of any Operator MEG required to monitor an OVC associated with the Tunnel ECS across a CEN.
4. A Tunnel-2 based service (see Section 6.3) allows Network Operators to limit the visibility of client EVCs or OVCs within their IIs via a Tunnel EC among a number of associated ENNIs (e.g., a carrier's carrier service). The Service Provider can use a Service Provider or Operator MEG (if the Tunnel EC involves a single CEN) to monitor the Tunnel-2 based service. Each intermediate CEN is responsible for the instantiation of any Operator MEG required to monitor an OVC associated with the Tunnel ECS across a CEN.
5. In a single CEN scenario it is always possible to associate the EVC MEG with the Subscriber ECS, irrespective of the presence of any Operator EC.
6. In a single CEN scenario an Operator EC can be used to instantiate an EVC. In use cases where there is a one-to-one mapping between the Service Frames of an EVC and a SP-EC, a Network Operator can use a SP or an Operator MEG associated with the Operator EC to monitor the EVC. Note that in a single operator scenario the SP-EC and Operator EC have the same monitoring scope.
7. In a multi CEN scenario, an OVC is typically monitored via the Operator MEG in each CEN. Either the Subscriber ECS or the Operator EC, if present, may be associated with the Operator MEG for the OVC between an UNI and an ENNI. In use cases where there is a one-to-one mapping between the Service Frames of an EVC and a SP-EC, a Network Operator can use the Operator MEG associated with the Operator EC to monitor the OVC.
8. An Operator EC, and an associated Operator MEG, is always expected to be present in support of OVCs that only associate ENNIs.

A CEN may instantiate an Operator EC to convey Service Frames between UNI-Ns. Similarly, a CEN may instantiate an Operator EC to convey ETH PDUs between a UNI-N and an ENNI-N, or to convey ETH PDUs between ENNI-Ns. Manual (e.g., UNI Type 1 [23]), or management plane (see UNI Type 2 [25]) mechanisms may be invoked to configure the mapping of Service Frames to Operator ECs between the UNI-N and its associated UNI-C. In the future, Control Plane (UNI Type 3) mechanisms may be invoked to configure the mapping of Service Frames to Operator ECs between the UNI-N and its associated UNI-C, or even the end-to-end configurations of an Operator EC among a number of UNI-Ns and associated UNI-Cs.

O-ECs **MAY** be emulated by connections in the TRAN Layer. Such O-EC emulation is outside the scope of this document.

6. Reference Models for the ETH Layer Connections

From a topological perspective the ETH Layer consists of Subscriber EFDs and CEN EFDs, representing Subscriber LANs and CENs interconnected by ETH Links. At the functional level, ETH Layer adaptation functions are used to map the Subscriber ECs to ETH Links and ETH Layer flow termination functions then instantiate the connection/connectionless trail (see MEF4) in the CEN associated with the Subscriber EC(s). The ETH TF also provides access into the ETH Layer functional elements in the Network Operator domains required to instantiate and manage ETH Layer flows within the involved CENs.

The sections below describe the relationship of the EVC and EC concepts and the functional elements associated with the ETH Layer. High-level functional decompositions of the UNI, ENNI and other EI and IIs of interest are discussed in Section 8. External interfaces covered in this specification include:

- User-Network Interface (UNI)
- External Network-to-Network Interface (ENNI)

Other EIs/IIs will be covered in future updates to this document.

6.1 EC Types

The EC type indicates the forwarding behavior and degree of connectivity provided by an EC. An EC can be classified as:

- a) uni-directional or bi-directional point-to-point transport entity,
- b) unidirectional point-to-multipoint transport entity,
- c) bidirectional multipoint transport entity¹⁰, or
- d) a compound concatenation of unidirectional and bi-directional transport entities.

An example of an EC consisting of a concatenation of point-to-point and multipoint ECs is one used to support a rooted multi-point EVC over two or more CENs to instantiate an E-Tree services as specified in MEF 10.2.

The next sections present topological representations of the following EC constructs that may be used by Network Operators and Service Providers to instantiate Operator and Service Provider ECs over CENs:

- Point-to-Point ECs
- Multipoint ECs
- Rooted Multipoint ECs
- Hairpin ECs
- Tunnel ECs

The focus of these representations is on key reference points (TFPs and FPs) in the participating CENs. Although the description is given in terms of generic ECs, the ECs of relevance among the CENs participating in the Ethernet service are Operator ECs in each CEN and the Service Provider EC.

¹⁰ A multipoint transport entity can also be interpreted as an instance of a number of point-to-multipoint transport entities. They are also referred to in ITU-T G.8021 as Flow Domain Fragments.

6.1.1 Point-to-Point ECs

A Point-to-point EC is a transport entity used to convey ETH Layer PDUs among two FP. A typical example is the flow points at the UNI-Ns associated with a point-to-point EVC. Point-to-point ECs, and their associated point-to-point ECSs, may be uni-directional or bi-directional. When bi-directional, each constituent ECS **MUST** be bidirectional.

A multi CEN Point-to-point SP-EC is constructed as a serial compound connection of O-ECs consisting of:

- one Point-to-point O-EC per traversed CEN, and
- one Point-to-point O-EC (i.e., Link EC) per traversed ENNI.

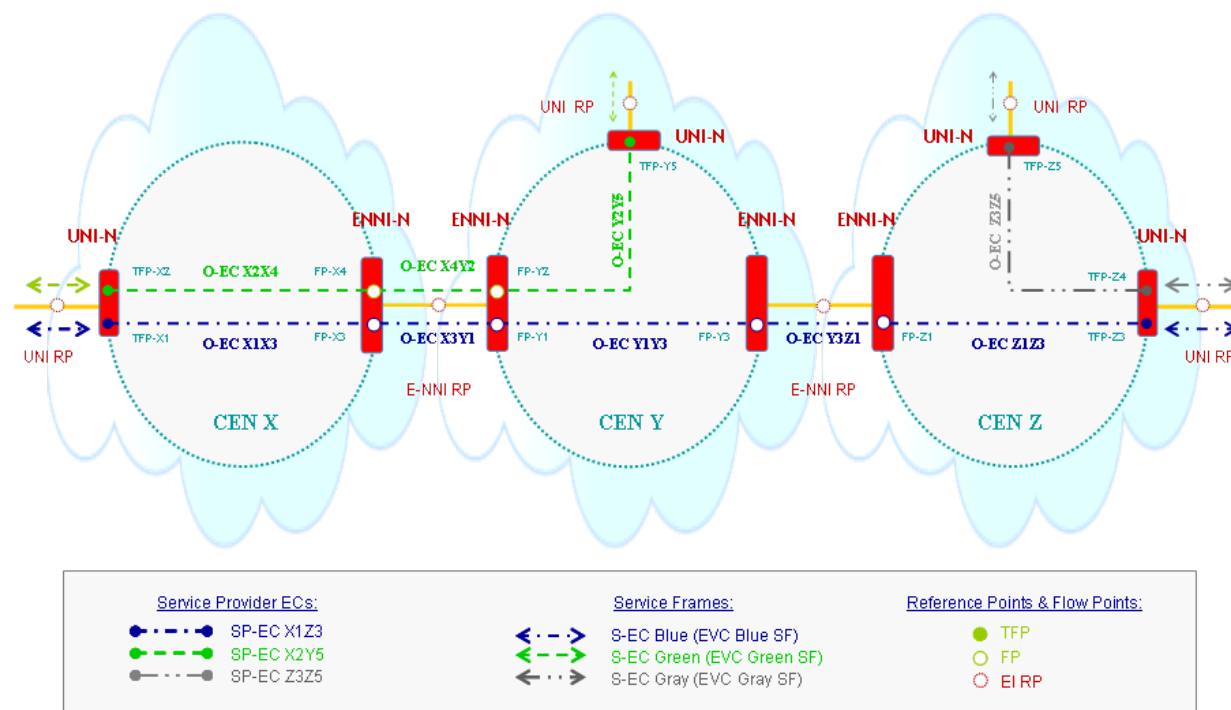


Figure 11: Example of Topological Representations of Point-to-point SP-ECs and associated O-ECs

Figure 11 illustrates examples of topological representations of bidirectional point-to-point SP-ECs traversing one, two or three CENs and identifies relevant reference points. Here, there is a SP-EC (SP-EC X1Z3) traversing CEN X, CEN Y and CEN Z interconnecting TFP X1 and TFP Z3. The SP-EC is constructed from O-EC X1X3 in CEN X, O-EC Y1Y3 in CEN Y and O-EC Z1Z3 in CEN Z. A Link EC, O-EC X3Y1, is used to interconnect O-EC X1X3 to O-EC Y1Y3. Similarly, a Link EC, O-EC Y3Z1, is used to interconnect O-EC Y1Y3 to O-EC Z1Z3. A similar approach is used to instantiate the SP-EC X2Y5 traversing CEN X and CEN Y, and the SP-Z3Z5 traversing CEN Z. The topological model generalizes to any number of participating CENs.

6.1.2 Multi-Point ECs

A Multipoint EC is a transport entity used to convey ETH Layer PDUs among two or more flow points. A typical example is the flow points at the UNI-Ns associated with a multipoint EVC. Multipoint ECs, and their associated ECSs, may be unidirectional or bidirectional. When bidirectional, each constituent ECS **MUST** be bidirectional and the two directions of any connection associated with the ECS are assumed to follow the same path at the given ETH Sublayer.

A Multipoint SP-ECs traversing multiple CENs is constructed as a serial compound connection of O-ECs consisting of:

- one Multipoint O-EC per traversed CEN, and
- one Point-to-point O-EC for each traversed ENNI.

The Multipoint O-EC **MAY** be replaced with a Point-to-point O-EC if there are only two EIs associated with the O-EC at the given CEN.

Figure 12 illustrates examples of topological representations of Multipoint SP-ECs traversing one or multiple CENs. Here, there is a SP-EC (SP-EC X1Y5Z3) traversing CEN X, CEN Y and CEN Z interconnecting TFP X1, TFP Y5 and TFP Z3. The SP-EC is constructed from O-EC X1X3 in CEN X, O-EC Y1Y3Y5 in CEN Y and O-EC Z1Z3 in CEN Z. A Link EC, O-EC X3Y1, is used to interconnect O-EC X1X3 to O-EC Y1Y3. Similarly, a Link EC, O-EC Y3Z1, is used to interconnect O-EC Y1Y3 to O-EC Z1Z3. A similar approach is used to instantiate the SP-EC Z2Z5Z6 traversing CEN Z. Note that the associated O-ECs of a Multipoint EC may be instantiated as point-to-point transport entities within a particular CEN, e.g., it will typically show as a point-to-point transport entity between ENNI components of adjacent CENs. The topological model generalizes to any number of CENs.

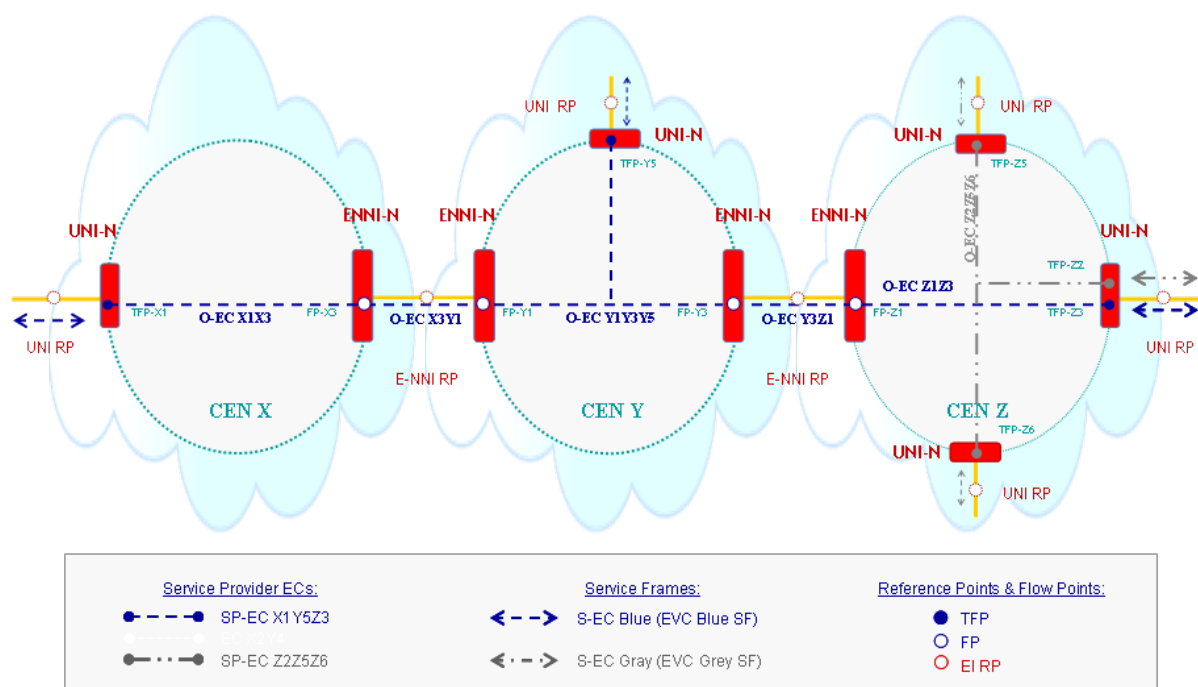


Figure 12: Example of Topological Representations of a Multipoint SP-EC and Associated O-ECs

6.1.3 Rooted Multi-Point EC

A Rooted Multipoint EC is a transport entity used to convey ETH Layer PDUs among FPs used to instantiate a rooted multi-point connection. A typical example is the FPs at Root and Leaf UNI-Ns associated with a rooted multipoint EVC. Rooted multipoint ECs **MAY** be constructed from a set of unidirectional and bidirectional ECSs. For the bidirectional ECSs, the two directions of any connection associated with the ECS are assumed to follow the same path at the given ETH Sublayer. The TFPs associated with the Root and Leaf UNIs are referred to as Root FPs and Leaf FPs, respectively. ETH PDU forwarding rules between Root UNIs (or Root UNI-Ns) and Leaf UNIs (or Leaf UNIN-Ns) and associated FPs described below are intended to be consistent with the RMP service as defined in MEF 10.2.

By the RMP service definition in MEF10.2, service frames received at a Root UNI may be forwarded to other Root UNIs and Leaf UNIs, while service frames received at a Leaf UNI may only be forwarded to Root UNIs. The set of TFPs associated with a RMP connection at EIs are:

- Root FP: associated with a Root UNI-N of a RMP connection
- Leaf FP: associated with a Leaf UNI-N of a RMP connection
- Trunk FP (T-FP): associated with a Trunk O-EC at an ENNI-N
- Branch FP (B-FP): associated with a Branch O-EC at an ENNI-N

Forwarding requirements for Trunk FPs and Branch FPs are:

- a) An egress ETH PDU at a Trunk FP **MUST** be the result of an ingress ETH PDU at either a Root FP or another Trunk FP
- b) An egress ETH PDU at a Branch FP **MUST** be the result of an ingress ETH PDU at either a Leaf FP or another Branch FP

Forwarding requirements for Root FPs and Leaf FPs are:

- a) An egress ETH PDU at a Root FP **MUST** be the result of an ingress ETH PDU at a Root, Leaf, Trunk, or Branch FP
- b) An egress ETH PDU at a Leaf FP **MUST** be the result of an ingress ETH PDU at either a Root or Trunk FPs

Table 4 summarizes the connectivity constraints based on the OVC End Point Subtype.

		Ingress TFP			
		Root	Leaf	Branch	Trunk
Egress TFP	Root	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>
	Leaf	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>
	Branch	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input type="checkbox"/>
	Trunk	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>

Table 4: Allowed Connectivity between TFP types

6.1.3.1 Trunk EC and Branch EC

Two special O-EC types are introduced to show a realization of the RMP TFP forwarding rules within a single CEN:

- 1) A Trunk O-EC, responsible for conveying service frames from a Root UNI to other Root UNIs and Leaf UNIs.
- 2) A Branch O-EC, responsible for conveying service frames from Leaf UNIs to Root UNIs.

To enable a RMP EC to span multiple CENs, it is necessary to interconnect the Trunk O-ECs of each CEN and to interconnect the Branch O-ECs of each CEN. Therefore the Trunk O-ECs and Branch O-ECs have TFPs at each ENNI-N traversed by the RMP EC.

The Trunk O-EC and Branch O-EC are in themselves compound ECs consisting of both bidirectional multipoint and unidirectional point-to-point ECs. In order to specify these ECs it is helpful to define three more types of flow points:

- Trunk-to-Leaf FP (TL-FP): associated with the interconnection of a Trunk O-EC to a Leaf UNI-N
- Leaf-to-Branch FP (LB-FP): associated with the interconnection of a Leaf UNI-N to a Branch O-EC
- Branch-to-Root FP (BR-FP): associated with the interconnection of a Branch O-EC to a Root UNI-N

A Trunk O-EC refers to a compound EC within a CEN consisting of:

- One Multipoint EC interconnecting the Root FPs, Trunk FPs and Trunk-to-Leaf FPs

- One unidirectional Point-to-point EC per Leaf UNI-N connecting from a Trunk-to-Leaf FP to a Leaf FP. One or more of these ECs **MAY** be consolidated into a point-to-multipoint O-EC.

A Branch O-EC refers to a compound EC within a CEN consisting of:

- One Multipoint EC interconnecting Leaf-to-Branch FPs, Branch FPs and Branch-to-Root FPs,
- One unidirectional Point-to-point EC per Leaf UNI-N connecting from a Leaf FP to a Leaf-to-Branch FP
- One unidirectional Point-to-point EC per Root UNI-N connecting from a Branch-to-Root FP to a Root FP. One or more of these ECs **MAY** be consolidated into a point-to-multipoint O-EC

A Rooted Multipoint SP-EC traversing multiple CENs is constructed as a compound serial connection of O-ECs consisting of:

- one Trunk O-EC per traversed CEN,
- one Branch O-EC per traversed CEN,
- a bidirectional Point-to-point O-EC interconnecting T-FPs at each traversed ENNI-N
- a bidirectional Point-to-point O-EC interconnecting B-FPs at each traversed ENNI-N

Figure 13 shows a topological representation in a CEN of a multi-CEN RMP EC interconnecting a Root UNI, a Leaf UNI, and an ENNI. Within the CEN, the Leaf UNI-N forwards ETH PDUs to its Branch EC via the LB-FP and receives ETH PDUs from the Trunk EC via the TL-FP. The Root UNI-N forwards and receives ETH PDUs from its Trunk EC directly. It also receives ETH PDUs from its Branch EC via the BR-FP.

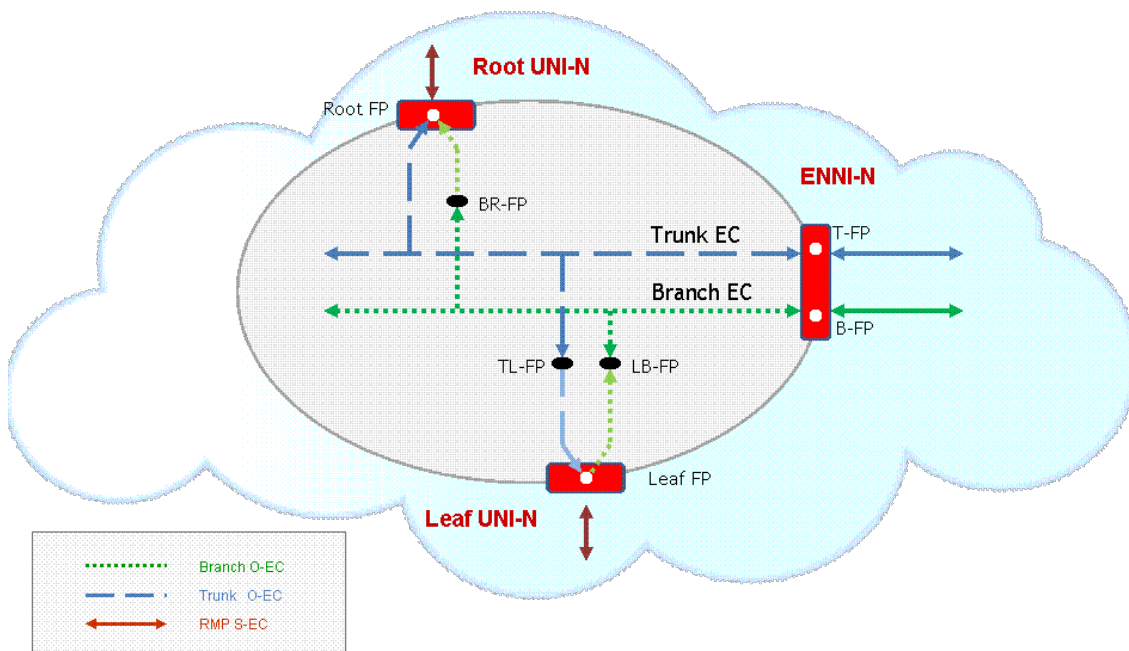


Figure 13: Topological representation of a portion of a RMP EC in a single CEN

Figure 13 does not constrain the location of BR-FP to be at a Root UNI-N and the TL-FP or LB-FP to be at a Leaf UNI-N. While it is possible to place the FPs at locations other than their corresponding UNI-Ns such implementation can impose additional resource management and operation burden on the instantiation and manageability of a RMP EC. These management and operational considerations are Network Operator specific and outside the scope of this specification.

Figure 14 shows a topological representation of a Rooted Multipoint EC, traversing three CENs (CEX X, CEN Y and CEN Z). Within each CEN there is a Trunk EC (ECS TX1X3X4X8 on CEN X, ECS TY4Y6Y8 on CEN Y and TZ2Z3Z4Z6 on CEN Z) and a Branch ECS (ECS BX1X3X4X7 on CEN X, ECS TY4Y5Y7 on CEN Y and

TZ2Z3Z4Z5 on CEN Z) constructed under similar rules as those illustrated in Figure 13. There are two Point-to-point Link ECs used to interconnect the Trunk ECs: ECS X8Y6 between CEN X and CEN Y, and ECS Y8Z6 between CEN Y and CEN Z). There are also two Point-to-point Link ECs used to interconnect the Branch: ECSX7Y5 between CEN X and CEN Y, and ECSY7Z5 between CEN Y and CEN Z.

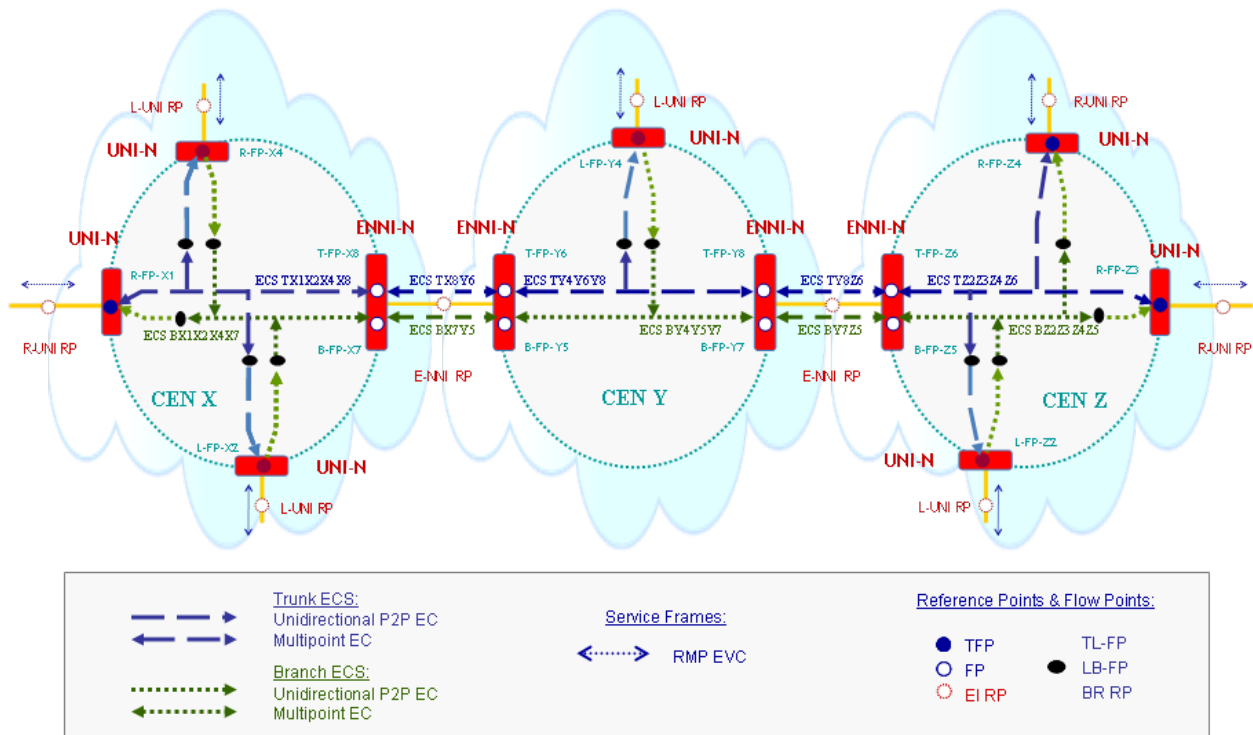


Figure 14: Root/Trunk O-EC

When an RMP EC is implemented via Root/Trunk O-ECs there **SHOULD** be at least one T-FP and one B-FP associated with each ENNI-N participating in a RMP EC. In addition, within any particular CEN there **MAY** be as many BR-FPs as Root UNI-Ns and as many LB-FPs and TL-FPs as the number of Leaf UNI-Ns in the CEN.

Note also that there are no requirements to instantiate within a CEN as many BR-FPs as Root UNI-Ns, or as many TL-FPs and LB-FPs as Leaf UNI-Ns. For example, a particular RMP EC realization could choose to associate multiple Root UNI-Ns to the same BR-FP, or associate multiple Leaf UNI-Ns to the same LB-FPs. Moreover, there are no requirements that the Trunk O-EC or the Branch O-EC encompasses multiple NEs. Hence, for example, a particular realization may choose to create any of these O-ECs within a single NE.

6.2 Hairpin EC

EC Hairpin refers to the interconnection, at an ENNI-N, of TFPs of an O-EC to the the TFPs of O-ECs from another CEN. The EC Hairpin enables the Service Provider to deliver Ethernet Services at UNIs in other CENs. The Network Operator instantiating the EC Hairpin is referred to here as the *Hairpin Operator*. The Network Operator delivering the O-ECs to be part of the EC Hairpin is referred to here as the *Partner Operator*. The resulting SP-EC is referred to as a Hairpin EC.

Thus, what differentiates a SP-EC with EC Hairpins from other SP-ECs discussed so far is that each O-EC from the Partner Operator that form part of the SP-EC are treated by the Partner Operator(s) as separate and independent SP-EC instances, not as a segment of the SP-EC from the Hairpin Operator. The Partner Operators need not be concerned that each these O-ECs (SP-ECs from the viewpoint of each of these Partner Operators) are part of another single SP-EC instance.

More specifically, the following requirements apply to an EC Hairpin at an ENNI-N:

- There **MUST** be two or more TFPs associated with the O-EC (Hairpin EC) within the Hairpin Operator (the Network Operator instantiating the EC Hairpin)
- There **MUST** be one TFP from an O-EC from the Partner Operator for each TFP on the Hairpin EC
- There **MUST** be a one-to-one map between a TFP in the Hairpin EC and a TFP from the Partner Operator EC

There are no constraints imposed on the number of O-ECs from the Partner Operator(s) interconnected by the Hairpin Operator to an O-EC at a given ENNI-N in an EC Hairpin other than each one of the O-ECs **MUST** belong to an independent O-EC instance. Otherwise frames will be likely to loop among the interconnected FPs.

From the viewpoint of the involved Partner Operator(s) each of its O-ECs represents an independent end-to-end service instance with its own service OA&M capabilities, not a segment of a longer SP-EC. It is the sole responsibility of the Hairpin Operator to treat the resulting compound EC as a single service entity and handle all service OAM coordination and reporting functions for the resulting SP-EC. A TFP at an ENNI-N interconnecting two O-ECs to help instantiate an EC Hairpin is referred to as a Hairpin Flow Point (H-FP).

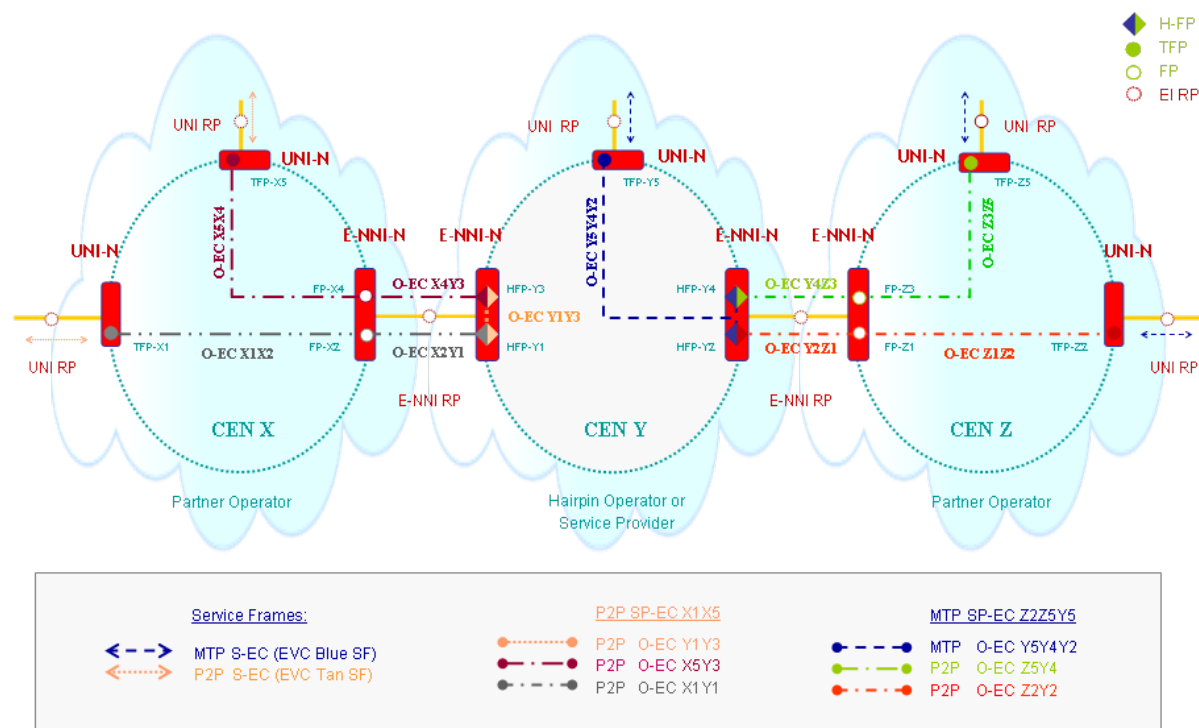


Figure 15: Point-to-Point and Multipoint Hairpin SP-ECs

Figure 15 illustrates various O-ECs used to construct point-to-point and multipoint SP-ECs using EC Hairpins in support of E-Line and E-LAN services. In the E-Line scenario, two Point-to-point O-ECs, represented by O-EC X1X2 and O-EC X5X4 on CEN X, are interconnected via an EC Hairpin to a Point-to-point O-EC, represented by O-EC Y1Y3, at the ENNI-N on CEN Y, to form the SP-EC X1X5. Here, O-EC Y1Y3 exists solely within the scope of the ENNI-N of the Hairpin Operator on CEN Y. In the E-LAN scenario, two Point-to-point O-ECs, represented by O-EC Z1Z2 and O-EC Z3Z5 on CEN Z, are interconnected via an EC Hairpin to a Multipoint O-EC, represented by O-EC Y5Y4Y2, at the ENNI-N on CEN Y, to form the SP-EC Y5Z2Z5. Although O-EC Y5Y4Y2 only associates two EIs, there are three FPs associated by the EC. Note that the topological model here generalizes to incorporate any number of CENs.

6.3 Tunnel ECs

The term “Tunnel EC” is used to describe a particular instance of an EC that transport frames from multiple client ECSs across one or more network domains. The term 'Tunnel' may be applied to an Operator EC within a single network domain or to a Service Provider EC representing the concatenation of multiple such Tunnel O-ECs that have the same set of client ECSs. The client ECS may be related to Subscriber or Operator ECs. One anticipated application for Tunnel ECs, and their associated “Tunnel Services”, is the construction of “All-to-One” bundles of client ECSs such that individual demarcation of the client flow(s) is irrelevant to the transport service provided by the Tunnel EC.

There are no restrictions on the type of EIs that may be associated to create Tunnel ECs. In particular, two usage cases for Tunnel ECs in support of carrier wholesale services have been noted.

Tunnel-1 service: This tunnel service is intended to associate a UNI in one CEN with an ENNI leading to another CEN. The Operator ECs associated with the Tunnel-1 service are intended to transport ETH PDUs between the Subscriber and the Network Operating acting on behalf of the Service Provider. There can be one or more intermediate CEN(s) traversed by the Tunnel-1 service and associated Operator ECs. There can be multiple EVCs associated with the UNI. The Operator EC associating an ENNI-N from the Network Operator acting as the Service Provider with the UNI-N from the Out-of-Franchise (OOF) Network Operator attached to the Subscriber is a Tunnel EC.

Tunnel-2 service: This tunnel service is intended to associate two ENNIs. The Operator EC associated with the Tunnel-2 service is intended to transport ETH PDUs from a given CEN, the client CEN, across one or more CENs, the server CENs. Thus, there can be one or more intermediate server CENs traversed by the Tunnel-2 service and associated Operator ECs. There can be multiple OVCs associated with the ENNIs of the client CEN. The Operator EC associating the two ENNI-Ns from the server CENs across one or more intermediate server CENs is also a Tunnel EC.

7. ETH Layer Processing Functions

The following ETH Layer processing functions are identified:

- ETH Flow Adaptation Function (EFAF)
- ETH Flow Termination Function (EFTF)
- ETH Conditioning Functions:
 - ETH Flow Conditioning Function (EFCF)
 - ETH Subscriber Conditioning Function (ESCF)
 - ETH Provider Conditioning Function (EPCF)
- ETH ECS Adaptation Function (EEAF)
- ETH EC Termination Function (EETF)
- ETH Connection Function (ECF)
- ETH EC Interconnect Function (EEIF)

In addition, the following processing functions also support processing functions related to the ETH Layer, yet they are not fully contained within the ETH Layer:

- APP to ETH Adaptation Functions (EAFs)
- ETH to TRAN Adaptation function (TAFs)

Figure 16 shows the generic representation of the processing entities described in this document. Although not explicitly depicted in Figure 16 the ECF is the smallest instance of an EFD.

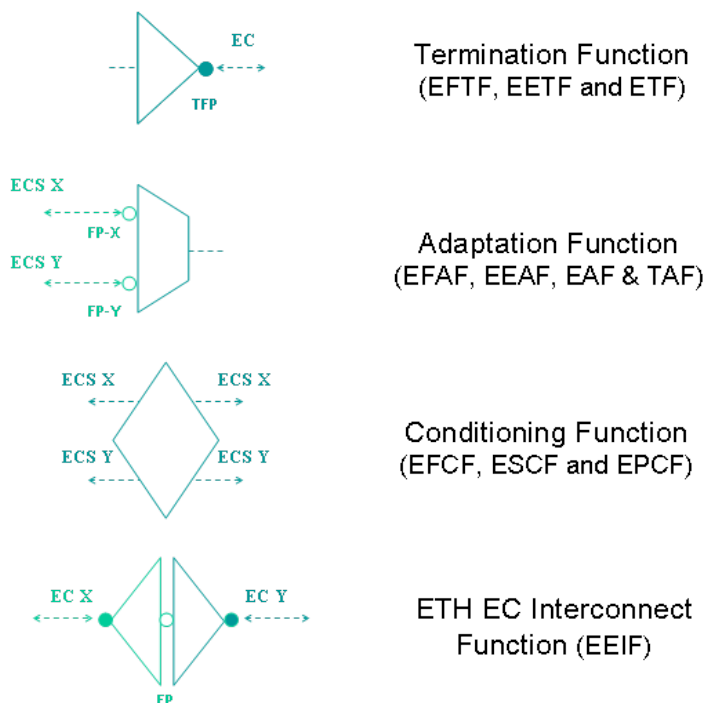


Figure 16: Diagrammatic representation of ETH Layer processing entities

Note that some functional elements such as AFs and TFs are defined as consisting of two separate source and sink components which are paired with counterpart sink and source components. This is the case for functional elements that add, and later remove, information elements to a flow for transport purposes. Other functional elements such as conditioning and classification functions (e.g., EFCF, ESCF and EPCF) are defined as stand-alone elements consisting of separate ingress and egress components. This is the case for functional elements that just modify the information elements being conveyed by a flow. Thus different sets of processing functions apply to different client/server combinations (e.g., the TAF for Ethernet-over-Fiber is different from the TAF for Ethernet over SONET/SDH, see also ITU G.8010 and G.8021). The sections below describe each of these functional elements. Section 8 illustrates the use of these processing functions to represent the ETH Layer base interfaces, the UNI and the ENNI.

7.1 APP to ETH Adaptation Function (EAF)

The APP to ETH Adaptation Function (EAF) refers to a class of processing entities responsible for the adaptation of the PDUs between the APP layer and the ETH Layer. EAFs are application specific as there are multiple application client types (e.g., IP, TDM, encoded voice/G.711, encoded video/BDV-ASI, etc.) that may wish to make use of the ETH Layer services. The EAF consists of separate source and sink functions.

The processes associated with the EAF Source function include:

- Reception of client PDUs
- LLC PDU formation (if LLC is present)
- EtherType allocation as per client application and/or LLC type (if LLC is present)
- Any required padding to minimum transmission unit size
- Multiplexing of adapted client PDUs towards EFTF

The processes associated with the EAF Sink function include:

- Demultiplexing of adapted client PDUs from EFTF
- De-encapsulation and EtherType processing
- LLC PDU extraction (if LLC is present)
- Relaying PDU to client process (as per EtherType)

7.2 ETH Flow Adaptation Function (EFAF)

The EFAF is the processing entity in the ETH Layer responsible for the adaptation of untagged ETH PDUs into C-tagged ETH PDUs as well as imposition/disposition or swapping of C-tags. The EFAF consists of separate source and sink functions.

The processes associated with the EFAF Source function include:

- Reception of untagged ETH PDUs from client ECs
- Mapping of conditioned client ETH PDUs into adapted ETH PDU for the given ETH Sublayer technology
- (Optional) Adaptation of the CoS indications in the EFD into the CoS indication for the tagged EC as per CoS Instance
- (Optional) Ingress traffic management functions including
 - Buffer management as per CoS Instance
 - Scheduling as per CoS Instance
- Aggregation of adapted ETH PDU from various host process
- Relaying the adapted ETH-PDU to the EFTF

The processes associated with the EFAF Sink function include:

- Reception of adapted ETH PDUs from the EFTF
- De-aggregation of adapted ETH PDUs into their corresponding ETH flow instances
- (Optional) Adaptation of the CoS indications from the EFD CoS indications, if applicable
- (Optional) Egress traffic management functions
- Relaying the untagged ETH PDU toward the client EFD

7.3 ETH Flow Termination Function (EFTF)

The ETH Flow Termination Function refers to a class of processing entities in the ETH Layer responsible for the instantiation of untagged Ethernet flows or C-VLAN based ETH Sublayer flows. The EFTF consists of separate source and sink functions.

The processes associated with the EFTF Source function include:

- Reception of adapted client PDUs from the EFAF or EAF
- Ethernet frame preparation including:
 - Values for the Destination and Source MAC Address fields,
 - Optional 802.1Q VLAN ID and UP value (CE-VLAN-CoS value to be populated in the Priority field of the C-VLAN tag), and
 - User Data preparation
- Formatting of ETH PDU (including any required OA&M information elements)
- Relaying of ETH PDU toward target EFD

The processes associated with the EFTF Sink function include:

- Reception of the ETH PDU from EFD
- Extraction of User Data (and potentially OA&M) information elements
- Relaying adapted client PDU to EFAF or EAF

7.4 ETH Conditioning Functional Elements

The ETH conditioning functional elements refer to a class of processing entities responsible for classifying, filtering, metering, marking, policing, shaping and, in general, conditioning of flows into and out of the ETH Layer links typically found between administrative network boundaries. Three types of ETH conditioning functions are identified: i) ETH Flow Conditioning Function, ii) ETH Subscriber Conditioning Function, and iii) ETH Provider Conditioning Function.

7.4.1 ETH Flow Conditioning Function (EFCF)

The ETH Flow Conditioning Function (EFCF) is the processing entity of the ETH Layer responsible for classification, filtering, metering, marking, policing, shaping and, in general, conditioning the Subscriber flow into and out of a Subscriber EFD. The EFCF consists of an ingress function operating on flows from the CEN and an egress function operating on flows to the CEN¹¹.

The processes associated with the EFCF function include:

- Reception of candidate Ingress Service Frame from the Subscriber EFD
- Ingress Service Frame classification into one or more (e.g., Service Multiplexing) ingress flows as per UNI Ingress rules defined in the MEF Services Model (MEF 10.2)
- Ingress Service Frame conditioning towards the Service Provider EFD as applicable to ingress flow(s) as per contracted Ingress BWP including:
 - Metering,
 - Marking, and
 - Policing.
- Relaying the Ingress Service Frame to the TAF at the UNI-C

The processes associated with the EFCF Ingress function include:

- Reception of Egress Service Frames from the TAF at the UNI-C
- Egress Service Frame classification (Optional) as applicable within the Subscriber EFD
- Egress Service Frame conditioning (Optional) as applicable within the Subscriber EFD including:
 - Metering,
 - Marking, and
 - Shaping.
- Relaying the Egress Service Frame towards the Subscriber EFD

7.4.2 ETH Subscriber Conditioning Function (ESCF)

The ETH Subscriber Conditioning Function (ESCF) is the processing entity of the ETH Layer responsible for classification, filtering, metering, marking, policing, scheduling, shaping and, in general, conditioning the subscriber flow into and out of a Service Provider EFD at a UNI-N. The ESCF consists of an ingress function operating on flows toward the CEN and an egress function operating on flows from the CEN.

The processes associated with the ESCF Egress function include:

- Reception of Egress Service Frame from the Service Provider EFD
- Egress Service Frame classification, including filtering, as per Egress UNI rules defined in the MEF Services Model (MEF 10.2)
- Egress Service Frames conditioning towards the Subscriber EFD as per contracted Egress BWP including:
 - Metering,
 - Marking, and
 - Shaping.
- Relaying the Egress Service Frames towards the TAF at the UNI-N

The processes associated with the ESCF Ingress function include:

- Reception of Ingress Service Frame from the TAF at the UNI-N
- Ingress Service Frame classification, including filtering, as per Ingress UNI rules defined in the MEF Services Model (MEF 10.2).
- CoS Instance determination as per contracted SLS

¹¹ Note that the EFCF processes are named after the direction of the data flow at the UNI-C port where it is used. The egress process processes the frames going out at the UNI-C (egress direction) and the ingress process processes the frames coming in at the UNI-C (ingress direction). Service Frames, however, are named in relation to the CEN, hence, irrespective of the port (see MEF 10.2).

- Ingress Service Frame conditioning towards the Service Provider EFD as per contracted Ingress BWP including:
 - Metering,
 - Marking, and
 - Policing.
- Ingress Service Frame Shaping as per Service Provider internal resource management requirements
- Relaying of the Ingress Service Frames towards the Service Provider EFD

7.4.3 ETH Provider Conditioning Function (EPCF)

The ETH Provider Conditioning Function (EPCF) is the processing entity of the ETH Layer responsible for classification, filtering, metering, marking, policing, scheduling, shaping and, in general, conditioning flow(s) between two CENs. The EPCF supports ETH Layer traffic conditioning functions satisfying the flow classification and resource management requirements at the ENNI. The EPCF consists of an ingress function operating on inbound flows to the CEN and an egress function operating on outbound flows from the CEN.

The processes associated with the EPCF Ingress function include:

- Reception of Ingress ETH PDUs from the TAF at the ENNI
- Ingress ETH PDUs classification, including filtering, consistent with the SLS specification for the relevant Ingress ENNI.
- CoS Instance determination as per contracted SLS
- Ingress Service Frame conditioning as per contracted Ingress BWP including:
 - Metering,
 - Marking,
 - Policing.
- Ingress ETH PDU Shaping as per the CEN's internal resource management requirements
- Relaying the Ingress ETH PDU towards the CEN's EFD

The processes associated with the EPCF Egress function include:

- Receiving Egress ETH PDU from the CEN's EFD
- Egress ETH PDU classification, including filtering consistent with ENNI SLS
- Egress ETH PDU conditioning towards the target CEN as per contracted Egress BWP including:
 - Metering,
 - Marking, and
 - Shaping.
- Relaying the Egress ETH PDU to the TAF at the ENNI

7.5 ETH EC Adaptation Function (EEAF)

The EEAF is the processing entity in the ETH Layer responsible for the adaptation of ETH PDUs between the ETH Sublayer technology used in the Subscriber domain and the ETH Sublayer technology used in the Network Operator domain. Thus, the EEAF may be viewed as a specialized "Flow Adaptation Function" for the purposes of handling ETH PDU transitions between two ETH Sublayers. Note that an EEAF is not required to be present. This may be the case when dedicated physical/logical link is used to convey the ETH Layer PDUs. The EEAF consists of separate source and sink functions.

The processes associated with the EEAF Source function include:

- Reception of client ETH PDUs from a ESCF or TAF
- Mapping of conditioned client ETH PDUs into adapted ETH PDU for the given Operator EC technology
- (Optional) Adaptation of the CoS indications in the client ETH PDU into the CEN's CoS indications as per contracted CoS Instance
- Aggregation of adapted ETH PDUs from the various client EC instances according to SLS (Optional)
- Ingress traffic management functions including
 - Buffer management as per CoS Instance
 - Scheduling as per CoS Instance
- Relaying the adapted ETH PDU toward the EETF.

The processes associated with the EEAF Sink function include:

- Receiving adapted ETH PDUs from the EETF
- De-aggregation adapted ETH PDUs from the Operator EC into their corresponding client ETH PDUs for the various Egress ECs according to SLS (Optional)
- (Optional) Adaptation of the CEN's CoS indications into the client EC CoS indications, as applicable
- De-mapping adapted ETH PDUs into ETH PDU for the given client EC technology
- Relaying the client ETH PDU toward the ESCF or TAF.

When the Operator EC is implemented via TRAN layer trails the EEAF may not be present.

7.6 ETH EC Termination Function (EETF)

The EETF is the processing entity in the ETH Layer responsible for the instantiation of an Operator EC. Thus, the EETF may be viewed as a specialized "Flow Termination Function" for the purposes of forwarding ETH PDUs associated with the Operator ECs. The EETF consists of separate source and sink functions.

The processes associated with the EETF Source function include:

- Receiving adapted ETH PDUs from the EEAF
- Aggregation of management (e.g., OA&M), control and data plane PDUs
- Relaying of the ETH PDUs towards the CEN EFD

The processes associated with the EETF Sink function include:

- Reception of the adapted ETH PDU from the CEN EFD
- De-aggregation of management (e.g., OA&M), control and data plane PDUs
- Relaying the adapted ETH PDU to the EEAF

When the Operator EC consists of a single ETH Layer link, the EETF may not be present.

7.7 ETH Connection function (ECF)

The ECF is the processing entity in the ETH Layer that affects the forwarding of the EC PDUs within the CEN. The main role of the ECF is to switch traffic between ETH Layer links to facilitate the creation of point-to-point or multipoint connections¹². ECFs are technology specific. Various connections models may be associated with ECFs. A sample list includes:

- Ethernet Frame Relay¹³
 - The ECF operates as a two port an Ethernet frame relay
 - The ECF emulates a point-to-point link
- IEEE 802.1D Bridge¹⁴
 - The ECF operates as an Ethernet Bridge as per IEEE 802.1D [3]
 - The ECF supports IEEE 802.1D based broadcast domains
- IEEE 802.1Q VLAN Bridge¹⁵
 - The ECF operates as an Ethernet VLAN Bridge as per IEEE 802.1Q [4]
 - The ECF supports IEEE 802.1Q/C-VLAN based broadcast domains
- IEEE 802.1 Provider Bridge
 - The ECF operates as an Ethernet Provider Bridge as per IEEE 802.1 Provider Bridge [7]

¹² The term connection is used here in a loose sense. If the transport mode provided by the underlying layer network were connectionless the connection would be termed a network flow and the ECF would be termed an ETH flow domain function, in the ITU-T Rec. G.809/G.8010 sense.

¹³ The term ECF is also used here in a loose sense as the "relay" function may be emulated via TAFs.

¹⁴ Commonly used to realize enterprise EFDs.

¹⁵ Commonly used to realize CoS-aware enterprise EFDs.

- The ECF supports IEEE 802.1Q/C-VLAN and IEEE 802.1 Provider Bridge/S-VLAN based broadcast domains
 - IEEE 802.1 Provider Backbone Bridge
 - The ECF operates as an Ethernet Provider Backbone Bridge as per IEEE 802.1 Provider Backbone Bridge [8]
 - The ECF supports IEEE 802.1Q/C-VLAN, IEEE 802.1 Provider Bridge/S-VLAN and IEEE 802.1 Provider Backbone Bridge/I-SID based broadcast domains
 - ATM:
 - The ECF operates as an emulate Ethernet Bridge ATM as per ATM Forum LANE specification [16]
 - The ECF supports emulation of multipoint IEEE 802.1Q/C-VLAN based broadcast domains
 - IETF VPLS Split-Horizon bridging function
 - The ECF emulates forwarding aspects of an IEEE 802.1D or an IEEE 802.1Q bridge as per the IETF L2VPN model [30][31]
 - The ECF supports IEEE 802.1Q/C-VLAN and IEEE 802.1 Provider Bridge/S-VLAN based broadcast domain
- Detailed models for each of these ECF implementations are outside the scope of this document.

7.8 ETH EC Interconnect Function (EEIF)

When two CENs are interconnected across an EI there will be common usage scenarios where the ETH Sublayer technology used in one CEN is different from the ETH Sublayer technology used on the EI (e.g., one of the CENs is based on PBB or VPLS and the ETH Sublayer technology specified for the EI is based on IEEE 802.1 Provider Bridge). The ETH EC Interconnect Function (EEIF) refers to the processing entity responsible for the adaptation of the ETH PDUs between two distinct types of ETH Sublayer technologies.

EEIFs are technology specific, as they depend on the type of ETH Sublayer technologies being matched. The EEIF can be realized as a serial compound functional entity consisting of:

- one EETF/EEAF pair used to terminate and expose the payload of the Operator EC traversing the EI,
- one EETF/EEAF pair used to terminate and expose the payload of the Operator EC traversing the CEN,

The EETF/EEAF pairs are interconnected back-to-back with their source and sink processing interfaces facing towards the targeted ECs. The EEIF is represented as depicted in Figure 16. When the same transport technology is used by both Operator ECs interconnected by the EEIF it may be possible to reuse the same Tag ID on both Operator ECs.

Note that the choice of ETH Sublayer technology within a CEN is Network Operator specific and outside the scope of this specification. In addition, the choice of ETH Sublayer technology for an EI is determined by the requirements in the EI specific Implementation Agreement. Hence, the specification of ETH Sublayer technologies for an EI is outside the scope of this document.

7.9 ETH to TRAN Adaptation function (TAF)

The TAF refers to a class of processing entities responsible for the adaptation of the ETH PDUs to its serving TRAN Layer. TAFs are technology specific as there are multiple server layer network types (e.g., Ethernet, SONET/SDH, ATM, FR, MPLS, etc.) that may be used to instantiate the ETH Layer links. In all cases, the Ethernet MAC frame format for the EI is specified as per IEEE 802.3-2005. The TAF consists of separate source and sink processes.

The processes associated with the TAF source include:

- Buffering and scheduling of the ETH_CI information units
- Allocation of VLAN ID field value, if applicable
- Payload padding to meet minimum transmission unit size requirements for the server layer

- Generation of MAC frame FCS
- Encapsulation/encoding (e.g., adaptation) ETH_CI according to TRAN layer specific requirements.
- Multiplexing of EC PDUs into ETH Link
- Rate Adaptation into the TRANs layer
- Insertion of adapted ETH Layer data stream into payload of TRAN layer signal

The processes associated with the TAF sink include:

- Ethernet MAC frame FCS verification
- Ethernet MAC Frame filtering of ETH PDUs from the Subscriber not intended to be forwarded across the EI/NNI
- Extraction of adapted ETH_CI from payload of the TRAN layer signal
- De- multiplexing of encapsulated EC PDUs

For L1 oriented TRANs, applicable adaptations include, among others:

- IEEE 802.3 [2]:
 - 8B/10B encoding
 - 4B/5B encoding
- SDH/SONET:
 - GFP (ITU-T Recommendation G.7041) [13]
 - LAPS (ITU-T Recommendation X.86) [14]

For L2 oriented TRANs, e.g., adaptation applicable adaptations include, among others:

- IEEE 802.1:
 - PBB Traffic Engineering (PBB-TE) [8]
- ATM:
 - Multi-protocol over AAL5 encapsulation as per IETF RFC2684 [28]
- IP/MPLS encapsulation as per IETF PWE3 model [29]

8. Base ETH Layer Functional Element Models

The Generic Architecture Framework (MEF 4) introduced various interfaces and associated functional elements to demark interconnect boundaries between administrative network components in the MEF architecture model. These interfaces, represented as ETH Links in the topological representation of the ETH Layer Network, provide a fixed relationship among functional elements required to implement the demarcation and processing functions at each end of the physical link/media associated with the EI. These set of processing functions are generically referred to as the EI Functional Elements. In particular, this document refers to the UNI and the ENNI as the Base ETH Layer functional elements, or Base EIs, of the MEF architecture model.

From the requirements in Section 5 implicit relationships can be derived among transport entities at the ETH Links of MEF compliant EIs:

- a) It **MUST** be possible to realize a one-to-many relationship between the ETH Link and its underlying TRAN Link(s).
- b) It **MUST** be possible to specify a one-to-many relationship between an EVC/OVC and Subscriber ECs
- c) It **MUST** be possible to define a many-to-one relationship between Subscriber ECs and Operator ECs.

These requirements arise from the peering relationship in the service model between the administrative and functional entities in the architecture model.

Examples of TRAN Links are: SDH VC-n-Xv, SONET STS-n-Xv, IEEE 802.3 PHYs, ATM VCs, and MPLS LSPs, among others. Note that a single ETH Link may consist of multiple TRAN Layer link connections or serial compound link connections. Examples are IEEE 802.3ad (LAG) for Ethernet physical interfaces and ITU-T G.707, ITU-T G.709 and ITU-T G.8040 for virtually concatenated TDM signals. Ethernet-over-SONET and Ethernet-over-OTN trails built on G.7041/Y.1303 are examples of potential multi TRAN Layer link connections.

The rest of this section provides functional representations for ETH Layer components of the Base EIs specified in this document.

8.1 ETH UNI

The UNI is the interface specified as the demarcation point between the responsibilities (and associated processing functions) of a Service Provider and the responsibilities (and associated processing functions) of a Subscriber. The functional element that represents the set of CEN functions is referred to as the UNI-N. The functional element that represents the set of Subscriber functions is referred to as the UNI-C.

From MEF4, there can only be a 1:1 relationship between a CEN and a Subscriber at an ETH Link associated with an UNI.

An ETH UNI-C consists of the following processing functions:

- a TAF instantiating the ETH Link towards the UNI-N
- an (optional) ETH Flow Conditioning Function at the UNI-C end of the ETH Link
- one ETH Flow Termination Function, and associated EFAF, for each Subscriber EC on the ETH Link

An ETH UNI-N consists of the following functions:

- a TAF instantiating the ETH Link toward the UNI-C
- an ETH Subscriber Conditioning Function at the UNI-N end of the ETH Link
- one ETH EC Termination Function, and associated EFAF, for each Operator EC associated with an set of Subscriber ECs.

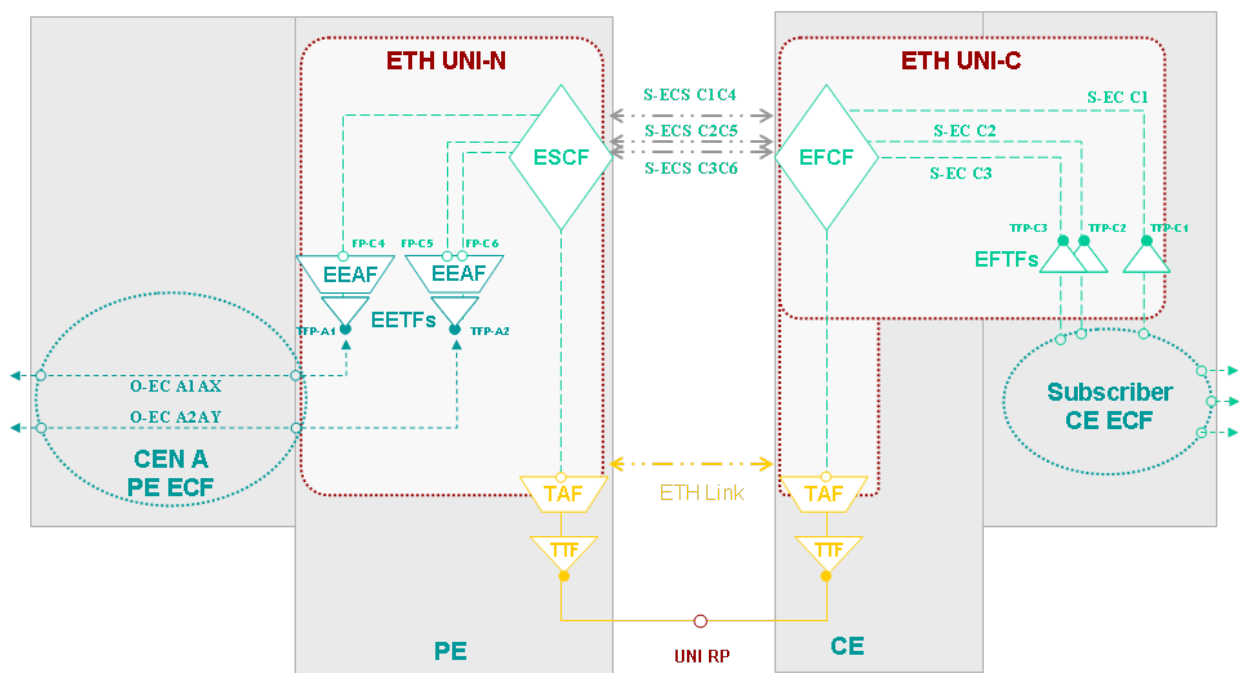


Figure 17: Functional representation of an ETH UNI-N and ETH UNI-C

Figure 17 shows a functional representation of an ETH UNI-C on the Subscriber CE interconnected to an ETH UNI-N on the CEN A PE. In this example the UNI-C shows a number of S-ECs, (S-EC C1, S-EC C2 and S-EC C3) instantiated via their associated EFAFs. Service Frames across the UNI Link associated with these S-ECs are conditioned for transmission/reception across the ETH Link according to the contracted SLS by the EFAF (Subscriber) and EPCF (Network Operator). The Subscriber UNI-C is configured to enable service multiplexing. One EVC corresponds to S-EC C1 and S-EC C2 (e.g., bundling) and another EVC corresponds to S-EC C3. The UNI-N shows two O-ECs, O-EC A1AX and O-EC A2AY, instantiated via their associated EETF in support of the requested EVCs.

The ETH ENNI is the interface specified as the demarcation point between the responsibilities (and associated processing functions) of two interconnected CENs. The functional element that represents the set of CEN functions is referred to as the ENNI-N. From MEF 4, there can only be a 1:1 relationship between two CENs at an ETH Link of an ENNI.

- A TAF instantiates the ETH Link toward the partner CEN
- an ETH Provider Conditioning Function (EPCF) at the ENNI-N end of the ETH Link
- one ETH EEIF for each Operator EC traversing the ENNI

Figure 18 shows a functional representation of interconnected ETH ENNI-Ns between two CENs. In this example, the ENNI-Ns on CEN A and CEN B show two Link ECs (O-EC A1B1 and O-EC A2B2) exchanged over the ENNI ETH Link. ENNI Frames across the ETH Link associated with these O-ECs are mapped to their corresponding O-EC within each CEN by the EEIF and conditioned for transmission/reception according to the contracted SLS by the EPAF.



This section will introduce new ETH Layer interfaces in a future update to this document.

10. References

- [1] IETF RFC 2119, "Key words for use in RFCs to Indicate Requirement Levels." March 1997.
- [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. December 2005.
- [3] IEEE Std 802.1D - 2004, IEEE Standard for Local and Metropolitan Area Networks: Media Access Control (MAC) Bridges. June 2004.
- [4] IEEE Std 802.1Q - 2005, IEEE Standards for Local and metropolitan area networks — Virtual Bridged Local Area Networks. 26 May 2006.
- [5] IEEE Std 802.1ad - 2005, IEEE Standard for Local and Metropolitan Area Networks — Virtual Bridged Local Area Networks — Amendment 4: Provider Bridges. May 2006.
- [6] IEEE Sts 802.1ag - 2007 IEEE Standard for Local and metropolitan area networks — Virtual Bridged Local Area Networks—Amendment 5: Connectivity Fault Management. December 2007.
- [7] IEEE Std 802.1ah - 2008 IEEE Standard for Local and Metropolitan Area Networks — Virtual Bridged Local Area Networks — Amendment 6: Provider Backbone Bridges. August 2008.
- [8] IEEE P802.1Qay - Provider Backbone Bridge Traffic Engineering. June 2009.
- [9] ITU-T Recommendation G.805 "Generic functional architecture of transport networks." March 2002.
- [10] ITU-T Recommendation G.809 "Functional architecture of connectionless layer networks." March 2003.
- [11] ITU-T Recommendation G.8010/Y.1306, "Architecture of Ethernet layer networks." February 2004.
- [12] ITU-T Recommendation G.8021, "Characteristics of Ethernet transport network equipment functional blocks." December 2007.
- [13] ITU-T Recommendation G.7041/Y.1303, "Generic Framing Procedure (GFP)." December 2003.
- [14] ITU-T Recommendation X.86, "Ethernet over LAPS." February 2001.
- [15] ITU-T Recommendation Y.1731, "OAM functions and mechanisms for Ethernet based networks." February 2006.
- [16] ATM Forum, "LAN Emulation over ATM v1.0." January 1995.
- [17] Metro Ethernet Forum, MEF 4, "MEN Architecture Framework – Part 1: Generic Framework." May 2004.
- [18] Metro Ethernet Forum, MEF 6, "Ethernet Services Definition, Phase 1." June 2004.
- [19] Metro Ethernet Forum, MEF 6.1, "Ethernet Services Definition, Phase 2." November 2006.
- [20] Metro Ethernet Forum, MEF 10, "Ethernet Services Attributes, Phase 1." November 2004.
- [21] Metro Ethernet Forum, MEF 10.2, "Ethernet Services Attributes, Phase 2." April 2009.
- [22] Metro Ethernet Forum, MEF 12, "Metro Ethernet Network Architecture Framework – Part 2: Ethernet Services Layer." April 2005.
- [23] Metro Ethernet Forum, MEF 13, "User Network Interface (UNI) Type 1." November 2005.
- [24] Metro Ethernet Forum, MEF 17, "Service OAM Requirements and Framework – Phase I." April 2007.
- [25] Metro Ethernet Forum, MEF 20, "User Network Interface (UNI) Type 2." July 2008.
- [26] Metro Ethernet Forum, MEF 23, "Carrier Ethernet Class of Service". June 2009.
- [27] Metro Ethernet Forum, MEF 26, "External Network-to-Network Interface". January 2009.
- [28] IETF RFC 2684, "Multiprotocol Encapsulation over ATM Adaptation Layer." September 1999.
- [29] IETF RFC 4448, "Encapsulation Methods for Transport of Ethernet over MPLS Networks." April 2006.
- [30] IETF RFC 4761, "Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling." January 2007.
- [31] IETF RFC 4762, "Virtual Private LAN Service (VPLS) Using Label Distribution Protocol (LDP) Signaling." January 2007.

Appendix I: ETH Layer Model for Service OAM and the Base EIs

This appendix provides a functional model for the MEF Service OAM (SOAM) framework at the ETH Layer as specified in MEF17. It identifies suitable placement for OAM functional components (MEPs and MIPs) in the ETH Layer in order to deliver as extensive and complete coverage of OAM maintenance domains as intended by the MEF SOAM framework. This appendix focuses on the functional model of SOAM as applicable to functional elements associated with the UNI and ENNI. Other EIs and IIs will be covered in future updates to this document.

The appendix is structured as follows: Section I.1 overviews the SOAM reference model from MEF17, the OAM terminology and modeling conventions that form the basic assumptions for the OAM model. Section I.2 describes OAM functional model and the placement of OAM components within the EI functional elements.

I.1 Definitions and Modeling Conventions

The Appendix addresses the OAM functional elements associated with the UNI and ENNI. Mapping between MEF Service and Architecture constructs are discussed in Section 5.4.3.

I.1.1 SOAM Reference Model

MEF17 defines the MEF SOAM framework from a Multi CEN perspective. Figure I.1 below is borrowed from MEF 17/Figure 5¹⁶. As noted in MEF17, when an EVC or OVC is provisioned there are a number of SOAM maintenance entity groups (MEGs) that are expected to be in place in order to monitor the proper operations of the intended service. In a single CEN scenario, the minimum set of MEGs traversed by a Subscriber EC includes:

- A Subscriber MEG
- An EVC MEG
- Two (e.g., E-LINE) or more (e.g., E-LAN and E-Tree) UNI MEGs

In addition, two other SOAM maintenance entities are expected to be in place in a multi CEN scenario:

- Two or more Operator MEGs
- One or more ENNI MEGs

Note that the Subscriber ECs need not be visible by all traversed MEGs. An example is any service scenario where Subscriber traffic is tunneled via Tunnel ECs and this Tunnel EC is not terminated at this particular EI.

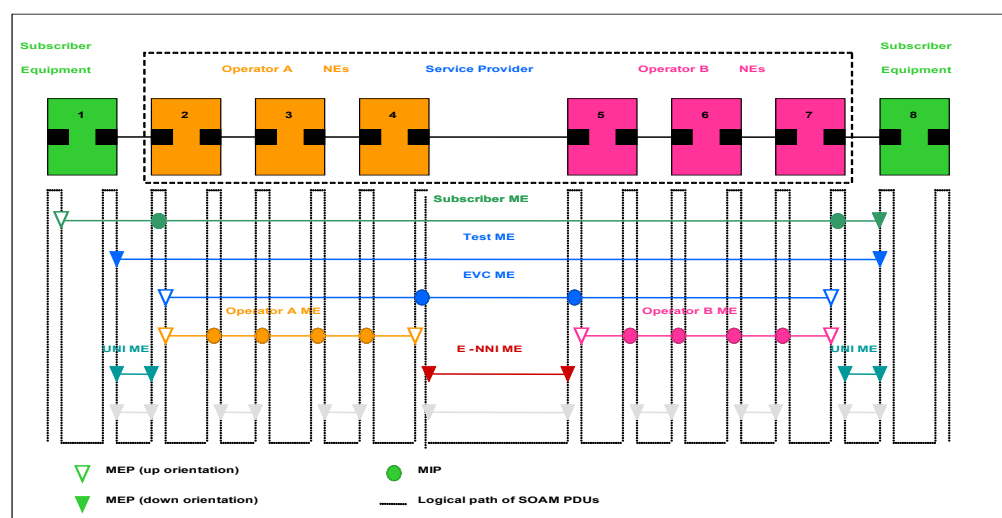


Figure I.1: SOAM Framework (derived from MEF 17, Figure 5)

¹⁶ Note that MEF17/Figure 5 refers to Point-to-point MEG. Thus, the use of MEG vs. ME is equivalent.

1.1.2 Mapping between IEEE and ITU Terminology

The MEF SOAM framework is based on the Ethernet OAM models specified by IEEE 802.1ag [6] and ITU-T Y.1731 [15]. There are some differences in terminology between the two documents. Table I.1 provides a mapping between some common terms used in this appendix.

IEEE 802.1ag	ITU-T Y.1731
Maintenance Association (MA)	Maintenance Entity Group (MEG)
Maintenance Association Identifier (MAID)	Maintenance Entity Group Identifier (MEGID)
Maintenance Domain (MD)	"Domain" ¹⁷
Maintenance Domain Level (MDL)	Maintenance Entity Group Level (MEG Level or MEL)

Table I.1: Mapping between IEEE and ITU Terminology

1.1.3 Representing MEPs and MIPs

ITU-T G.8021 [12] provides a generic functional model for Ethernet Networks, including OAM components. In the ITU-T model, OAM components are represented as a combination of AFs and TFs. More specifically,

- A MEP is represented as an AF/TF pair
- A MIP is represented as "back-to-back" AFs.

In addition, depending on the placement of the OAM functions with respect to the of the monitored transport entity and the use of Shared vs. Independent MELs (see ITU-T Y.1731 Sec 6.2) the TF component of the associated transport entity may be reused in the instantiation of the MEG. This Appendix follows a similar representation convention with the following exceptions:

- A different symbol, a "circle", is used to represent a MIP. This is done to provide a similar representation style as in MEF17.
- A connected TA/AF is used to represent a MEP.
- A shadowed symbol is also used to emphasize the OAM processing entities unique to the functional representation of the OAM components.

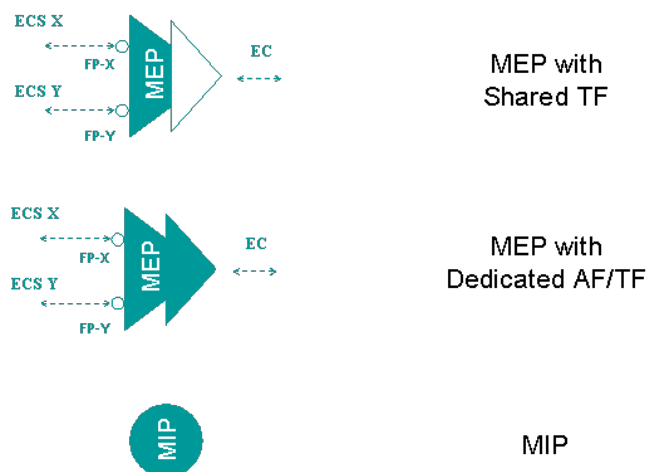


Figure I.2: Diagrammatic representation of OAM processing entities in the ETH Layer

¹⁷ Not a formal term used in the OAM functional descriptions.

Thus, when a shared MEL is used, the AF portion of the MEG is the only new processing entity introduced by the OAM model. The TF portion of the MEG is shared with the associated EC. When an Independent MEL is used, both the TF and AF portions of the MEP are the new processing entities introduced by the OAM model. These representation conventions are illustrated in Figure I.2.

I.2 SOAM Functional Model for the ETH Layer

The following subsections describe a SOAM model for the ETH Layer. It identifies the placement for the OAM processing entities (MEPs and MIPs) that apply to the functional elements associated with UNI and ENNI.

I.2.1 OAM Components at the UNI

As noted in Section I.1.1, the Subscriber, EVC and UNI MEG constitute the minimum set of relevant MEGs at the UNI. Figure I.3 below depicts a functional representation of the UNI and the proposed placement of OAM components at the UNI-N and UNI-C in order to instantiate the Subscriber, EVC and UNI MEGs.

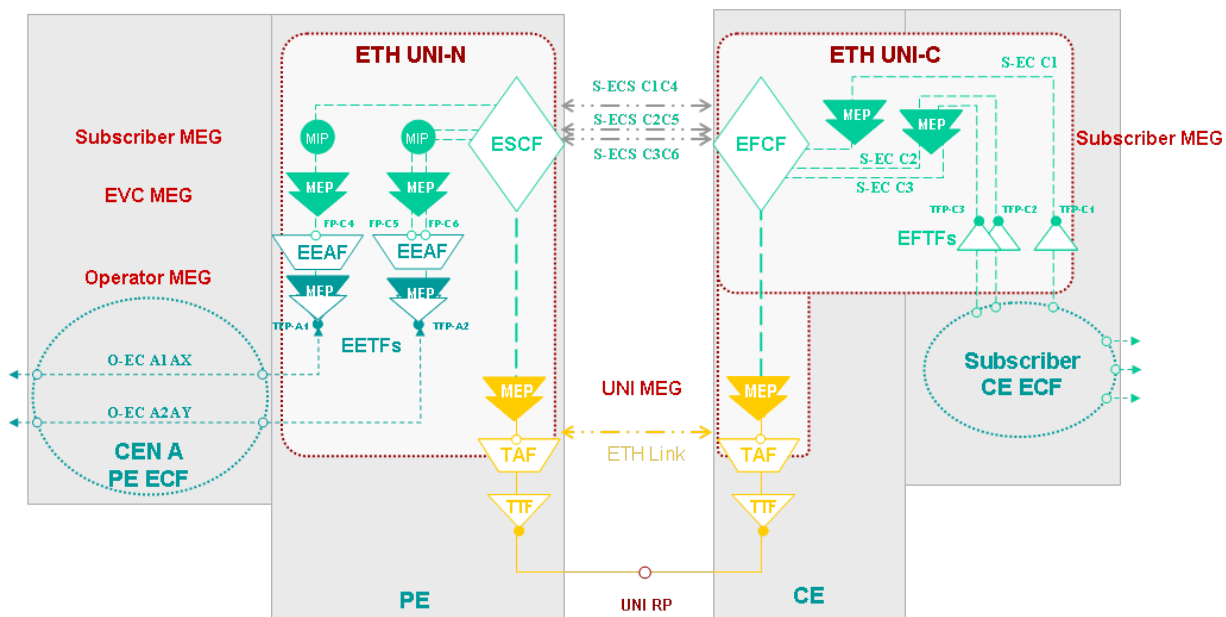


Figure I.3: Functional representation of MEPs & MIPs Placement at an UNI

I.2.1.1 UNI MEG

The UNI MEG is intended to cover the OAM maintenance requirements for the UNI Link. Specifically, a UNI MEG is intended to monitor a link connection between the Network Operator Domain and the Subscriber Domain. Thus, a MEP instantiating a UNI MEG is expected to be placed as close as possible to the monitored ETH Layer link.

I.2.1.2 Subscriber MEG

The Subscriber MEG is intended to cover the OAM maintenance requirements for an EVC from the Subscriber perspective. Specifically, a Subscriber MEG is intended to monitor the Subscriber ECs across the Network Operator Domain(s).

Placement of a Subscriber MEP is primarily dictated by the need to associate incoming Service Frames with an EVC. Thus, a MEP instantiating a Subscriber MEG is expected to be placed between the EFCF responsible for clas-

sification and conditioning of the subscriber flows at the UNI-C and the EETF responsible for instantiating the monitored Subscriber EC(s).

There could be a MIP associated with the Subscriber MEG at the UNI-N. If present, it is expected to be placed between the ESCF responsible for classification and conditioning of the Subscriber ECs into/out off the CEN and the EETF responsible for instantiating the associated Operator EC and before the instantiation of any Operator MEG.

I.2.1.3 EVC MEG

The EVC MEG is intended to cover the OAM maintenance requirements for an EVC from the Service Provider perspective. Specifically, an EVC MEG is intended to monitor the Subscriber ECs across the Network Operator Domain(s).

Placement of an EVC MEP is primarily dictated by the need to associate incoming Service Frames with an EVC. Thus, a MEP instantiating an EVC MEG is expected to be placed between the ESCF responsible for classification and conditioning of the Subscriber ECs at the UNI-N and the EETF instantiating the associated Operator EC¹⁸ and after the instantiation of any Operator MEG.

I.2.1.4 Operator MEGs at the UNI

The Operator MEG is intended to cover the OAM maintenance requirements for segments of client layer ECs within a CEN, such as Subscriber ECs and associated EVCs. In certain scenarios a Network Operator may prefer to instantiate a separate OAM maintenance entity to monitor an ETH Layer transport entity, such as a Subscriber ECs, then to share a MEG with the Subscriber or Service Provider. This can be accomplished via an Operator MEG.

As with EVC MEPs, placement of an Operator MEP at an UNI-N is primarily dictated by the need to associate Service Frames with an EVC. Thus, a MEP instantiating an Operator MEG is expected to be placed between the ESCF responsible for classification and conditioning of the Subscriber ECs at the UNI-N and the EETF associated with the Operator EC.

I.2.2 OAM Components at the ENNI

As noted in Section I.1.1, the EVC, Operator and ENNI MEG constitute the minimum set of relevant MEGs at the ENNI. Figure I.4 below depicts a functional representation of the proposed placement of OAM components at an ENNI-N in order to instantiate the EVC, Operator and ENNI MEGs.

¹⁸ Note that an Operator EC may not always be present in a particular Operator CEN. This may be the case, for instance, in small aggregation MENs interconnected via a larger backbone CEN.

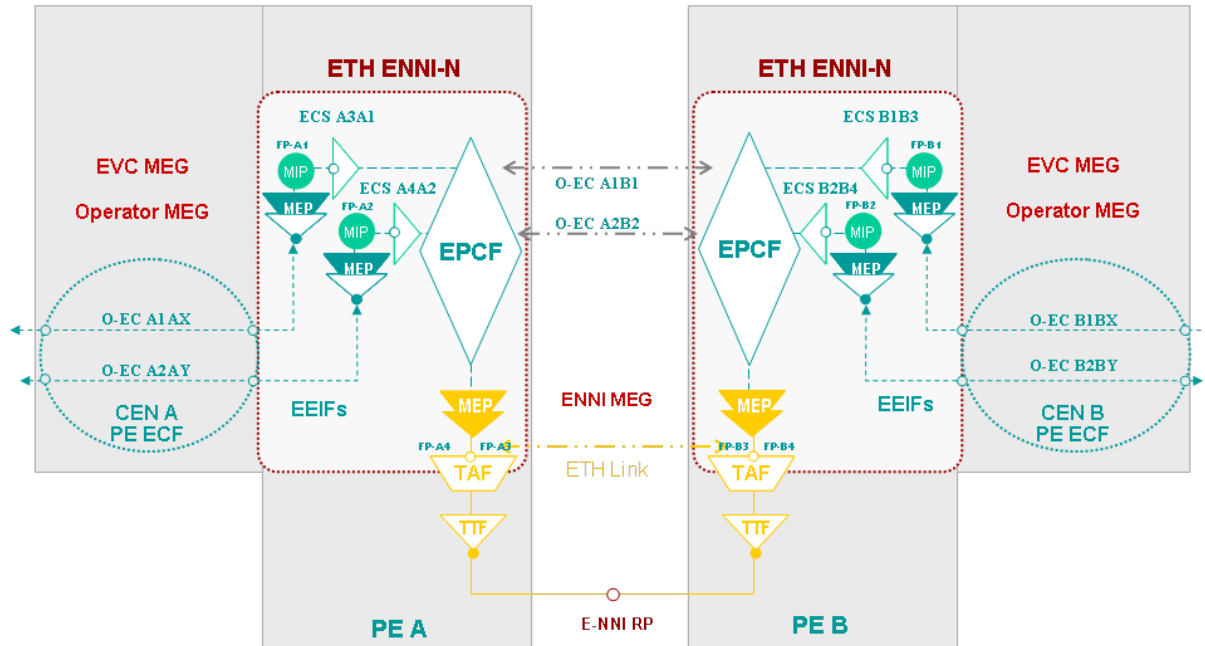


Figure I.4: Functional representation of MEPs & MIPs Placement at an ENNI

I.2.2.1 ENNI MEG

The ENNI MEG is intended to cover the OAM maintenance requirements for the ENNI Link. Specifically, an ENNI MEG is intended to monitor a link connection between CEN domains. Thus, a MEP instantiating an ENNI MEG is expected to be placed as close as possible to the monitored ETH Layer link.

I.2.2.2 Operator MEG

As noted in Section I.2.1.4 the Operator MEG is intended to cover the OAM maintenance requirements for segments of client ECs from a Network Operator perspective. As an example, an Operator MEG can be used to monitor the Operator EC associated with an OVC or a segment of one or more Subscriber EC across the Network Operator Domain.

Placement of an Operator MEPs at an ENNI-N is primarily dictated by the need to associate incoming ENNI Frames, and associated Link ECs, with Operator ECs. Thus, a MEP instantiating an Operator MEG is expected to be placed between the EPCF responsible for classification and conditioning of Link ECs at the ENNI-N and the EETF associated with the Operator EC.

I.2.2.3 EVC MEG at the ENNI

As noted in Section I.2.1.3, the EVC MEG is intended to cover the OAM maintenance requirements for EVCs from a Service Operator perspective. In certain scenarios a Network Operator may wish to allow other Network Operators to monitor a segment of the EVC at the ENNI. Such a capability can be accomplished via a MIP, based on a shared MEG, on the EVC MEG at the ENNI-N.

As with EVC MEPs, placement of an EVC MIPs at an ENNI-N is primarily dictated by the need to associate incoming ENNI Frames, and associated Link ECs, with EVCs and its associated Operator ECs. Thus, a MIP associated with an EVC MEG is expected to be placed between the EPCF responsible for classification and conditioning of the Link ECs at the ENNI-N and the EETF associated with the Operator EC.