Technical Specification

MEF 32

Requirements for Service Protection Across External Interfaces

July 2011

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:

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

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

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.

Table of Contents

List of Figures

List of Tables

1. Abstract

This document contains requirements for protecting Ethernet Services at an External Interface against link or a Network Element (NE) failure.

This specification identifies resiliency requirements at external interfaces applicable to MEF Service and associated Service End Points, which associate services with External Interfaces. MEF would like these requirements to be considered when designing a mechanism for Service Protection across an External Interface.

The current MEF standardized External Interfaces are the UNI and ENNI. These requirements aim to cover a wide range of MEF Ethernet service types as well as a wide range of network deployments in which the EIs are MEF defined EIs.

2. Abbreviations

Abbreviation	Definition	Reference
BW	Bandwidth	
CE	Customer Edge	[4]
CoS	Class of Service	[4]
EI	External Interface	[6]
ENNI	External Network Network Interface	[6]
SEP	Service End Point	[6]
EPL	Ethernet Private Line	[3]
EVC	Ethernet Virtual Connection	[3]
EVPL	Ethernet Virtual Private Line	[3]
EP-LAN	Ethernet Private LAN	[3]
EVP-LAN	Ethernet Virtual Private LAN	[3]
EP-Tree	Ethernet Private Tree	[3]
EVP-Tree	Ethernet Virtual Private Tree	[3]
FD	Frame Delay	[6]
FDV	Frame Delay Variation	[6]
FLR	Frame Loss Ratio	[6]
MEN	Metro Ethernet Network	[4]
NE	Network Element	[10][10]
OVC	Operator Virtual Connection	[6]
SLS	Service Level Specification	[4]
SDO	Standards Development Organization	
SP	Service Provider	[4]
UNI	User Network Interface	[5]
VLAN	Virtual LAN	[4]

Table 1: Abbreviations

3. Introduction

Reliability, in terms of availability, is a key attribute of a Carrier Ethernet service. High Availability commitments in SLSs require a resilient network that can rapidly detect interface failure, NE failure and performance degradation, and can rapidly restore service operation. Network survivability plays a critical factor in the delivery of reliable services.

Many resiliency schemes are deployed to date. Those resilient schemes include linear schemes such as 1:1, 1+1, 1:N as well as protection schemes that operate over Mesh and Ring topologies. Each of the schemes has its pros and cons. The requirements specify behaviors that any protection mechanism should achieve.

4. Scope and Definitions

This specification specifies requirements for Service protection across External Interfaces. The requirements for protection address only the Interconnection Zone (see definition below).

4.1 Definitions

Networks are connected to each other at demarcation points. In many cases the resources supporting the connections, i.e., nodes and Link Connections (see definition below) are redundant, providing improved protection for Service.

This specification defines requirements for protection of Ethernet Services that would be performed only at the External Interfaces.

4.1.1 Service End Point (SEP)

In this specification, a Service End Point is an association of a service (EVC), a service construct (OVC), a UNI Tunnel Access (UTA) or a Virtual UNI (VUNI) to an External Interface (a UNI or ENNI in the context of this specification). Note that the End Point normative definition, as specified in the ENNI specification [6][6], allows the association of End Points only with OVCs, and to VUNI or UTA, as specified in [7][7]. This specification extends the definition of an End Point, as specified in the ENNI specification [6], to implicitly assume association of a single End Point with an EVC at a given UNI [4][5], as well. Hence the term Service End Point is inclusive of OVC End Points at an ENNI and OVC End Point and EVC End Point at a UNI. This concept is introduced in this document for the purpose of describing requirements related to protecting connectivity across an EI.

4.1.2 Link Connection

A Link Connection as defined in MEF 4 (Ethernet Network Architecture Framework - Part 1: Generic Framework specification) [2] denotes the connectivity supporting the exchange of Ethernet Service Frames or ENNI Frames as defined in MEF 10.2 (Ethernet Services Attributes - Phase 2) [4] and MEF 26 (External Network Network Interface (ENNI) – Phase 1) [6], respectively, across an EI. The transport layer is an Ethernet phy sub-layer as specified in [4] and [6]. Both [4] and [4] describe Link Aggregation [8][8] as a mechanism protecting a MEF defined EI against a Link Connection failure. The requirements described in this specification are intended to work in conjunction with, or instead of, Link Aggregation. The Ethernet phy sub-layer must be supported while other transport layers such as Sonet/SDH or MPLS/PW are not precluded.

4.1.2.1 Working Link Connection

The designated Link Connection that exchanges Ethernet frames between Service End Points under normal condition (i.e., where there is no failure) is either configured or selected automatically. When there is no failure, frames between two Service End Points are carried on a single Working Link.

4.1.2.2 Resilient Link Connection

The Link Connection is either pre-configured or automatically chosen to exchange Ethernet frames between Service End Points when the Working Link Connection fails.

When a resilient Link Connection is pre-configured, it is denoted as a Protection Link

Note that a given Link Connection can serve the role of "working link" for some Service End Points while simultaneously serving the role of "resilient link" for other Service End Points.

4.1.3 Active and Standby Links Connections

Active Link Connection is a dynamic operational status of a pre-configured Link Connection indicating that the Link Connection is currently exchanging Ethernet Service frames or ENNI frames for specific Service End Points.

Standby Link Connection is a dynamic operational status of a pre-configured Link Connection indicating that the Link Connection is currently not forwarding Ethernet Service frames or ENNI frames for specific Service End Points. This definition applies for protection only.

Note that the Working Link Connection and the Resilient Link Connection can each have a status of Active or Standby, but not the same status simultaneously for a given Service End Points.

4.1.4 Resiliency

Resiliency is a generic term covering both Protection and Restoration.

- Protection: re-establishing service delivery using pre-allocated resources. The pre-allocation of resources guarantees the re-establishment of the service.

- Restoration: re-establishing service delivery using resources allocated at the time of need; This scheme does not pre-allocate resources, allowing them to be used during normal operation.

A traffic redirection event occurs when a failure of the Active Link Connection is detected, and as a result, traffic is switched (i.e., redirected) from the failed Link Connection to the Resilient Link. When a Protection Link is pre-configured, the traffic is switched (i.e.,

redirected) from the failed Link Connection to the Standby Link Connection (which now becomes the Active Link Connection).

4.1.5 Failure Event

Throughout this specification the term failure means any event that affects the performance of an Ethernet Service, violating the SLS agreement. Examples of such events are: link failure, NE failure, link degradation.

4.1.6 Interconnection Zone

An Interconnection Zone is an area where External Interfaces are interconnecting two administrative domains. The Interconnection Zone contains a collection of domain border NEs of the two interconnected administrative domains, which are associated with their respective External Interface and Link Connections connecting between the External Interfaces of two administrative domains. In addition, the Interconnection Zone also includes Link Connections connecting between the border domain NEs themselves. An Interconnection Zone can only be assigned to a single instance of the resiliency mechanism. The Interconnection Zones currently supported by MEF Specifications are: UNI Interconnection Zone, between the subscriber domain and the Service Provider domain and ENNI Interconnection Zone, between two Service Providers.

Note that domain border NEs and the Link Connection(s) connecting between the border domain NEs themselves can be associated with one or more Interconnection Zones. Other Link Connections and the same or other domain border NEs associated with External Interfaces may not be assigned to any Interconnection Zone.

4.2 Reference Model

The network reference model is illustrated in Figure 1~~Figure 1~~ below. Several adjacent MENs for which protection requirements may be defined, are illustrated in the figure. The Interconnection Zones currently supported by MEF Specifications are: UNI i.e., between UNI-C and UNI-N and ENNI between adjacent ENNI-Ns. This reference model is used when defining requirements for Service End Point protection and related terminology.

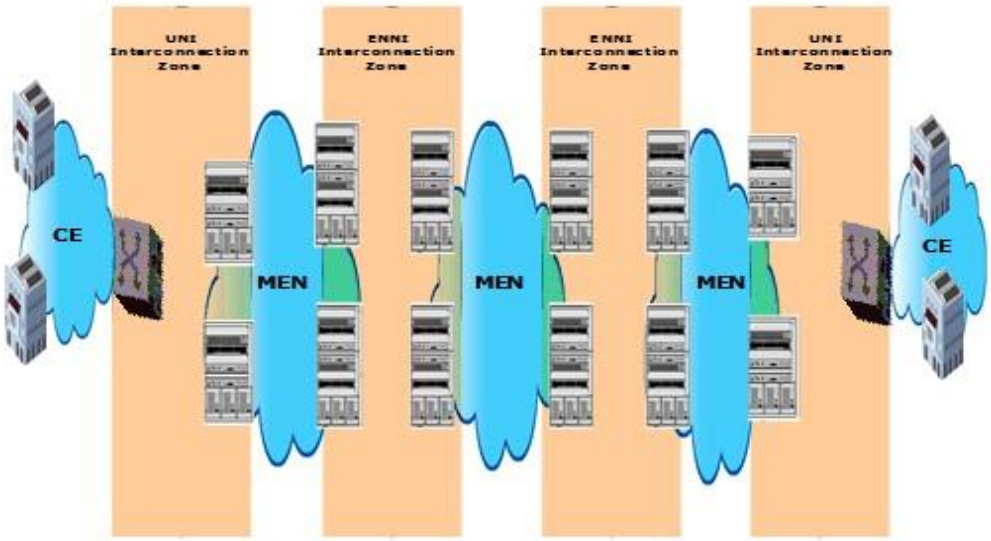

Figure 1: A Reference Model of the Interconnection Zones

Several MENs and subscriber premises are depicted in Figure 1~~Figure 1~~. An Interconnection Zone may be defined between two interconnected networks supporting MEF defined ENNI External Interfaces or between a MEN and a subscriber premise supporting MEF defined UNI External Interfaces.

Appendix A depicts a typical use case where an EVC traverses several networks which are interconnected at Interconnection Zones. Each Interconnection Zone provides protection only within its Zone.

4.3 In Scope for Phase I

The following items highlight the scope of this specification for Phase I:

- ENNI Interconnection Zone.
- UNI Interconnection Zone.

4.4 Out of Scope for Phase I

The following are out of scope of this specification for Phase I, but are candidates for inclusion in this document as future work items:

- Requirements for UTA Service protection across ENNI Interconnection Zone.
- Requirements for NID to MEN Interconnection Zone, when the NID is connected to the MEN over an EI.
- Requirements for 1+1 protection scheme.
- Failure detection time or degradation period is not included in the resiliency switching time (i.e. redirection time).

5. Compliance Levels

The key words "**MUST**", "**MUST NOT**", "**REQUIRED**", "**SHALL**", "**SHALL NOT**", "**SHOULD**", "**SHOULD NOT**", "**RECOMMENDED**", "**MAY**", and "**OPTIONAL**" in this specification are to be interpreted as described in IETF RFC 2119 [1].

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

6. Requirements

When high reliability is required for Carrier Ethernet services at the Interconnection Zone, a resiliency mechanism would enhance reliability by restoration of traffic with minimal impact to EVC and/or OVC SLS provided to the end user. This specification addresses the requirements for a resiliency mechanism. The mechanism should try to avoid a potential single point failure (NE or interface) and react to an SLS violation such as high FLR or FD.

The designed resiliency mechanism used in an Interconnection Zone needs to be robust enough to ensure a service is protected against various types of failures such as:

- An interface failure between two nodes, each residing on a different Administrative Domains.
- A NE failure supporting the service in the Interconnection Zone.
- Service performance degradation, i.e., where the network performance violates the SLS, across the Interconnection Zone.

This specification defines requirements for protecting Service End Points (SEPs) in the Interconnection Zone following the topics depicted below:

6.1 General Requirements

This section details general requirements.

[R1] The Ethernet physical layer as defined in MEF 10.2 [4] and MEF 26 [6] **MUST** be supported by the resiliency mechanism.

[R2] Note that other transport layers such as Sonet/SDH or MPLS/PW are not precluded.

[R3] The resiliency switching mechanism **MUST** be able to operate at a UNI.

[R4] The resiliency mechanism **MUST** be able to operate at an ENNI.

[R5] The resiliency mechanism **MUST** support resiliency per a single Service End Point.

[R6] The resiliency mechanism **MUST** always converge to a state such that the frames for a given Service End Point are carried on a single Link Connection.

Note that this requirement is designed to prevent services from being split into "streams" (or "conversation", as defined in the IEEE 802.1AX [8] specification).

[R7] The resiliency mechanism **MUST** support traffic redirection caused by failure events, from the Working Link Connection to the Resilient Link Connection , in the Interconnection Zone, in case of a failure, without manual intervention.

[R8] Each Service End Point **MUST** be supported by exactly one Working Link Connection and at least one Resilient Link Connection across the Interconnection Zone.

[R9] The resiliency mechanism **MUST** provide indication of a protection state change to a Management System, i.e., which Link Connection has changed its Active or Standby status for each Service End Point.

[R10] When traffic redirection is required, the resiliency mechanism **MUST** perform traffic redirection at each side of the Interconnection Zone such that the traffic for a given Service End Point uses the same Link Connection in each direction.

[R11] In the absence of any other failure in the Interconnection Zone, the resiliency mechanism **MUST** be capable of protecting against any single failure within the Interconnection Zone (NE or Link Connection).

[R12] The resiliency mechanism **SHOULD** be independent of other mechanisms inside the associated MEN and **SHOULD** be able to perform all its functionality independent of the internal functionality of the associated MEN.

[R13] The resiliency mechanism **MUST** provide indication of the protection state to the local associated networks in the Interconnection Zone.

6.2 Requirements Addressing the Ethernet Layer

This section details the requirements addressing frame formats.

[R14] The resiliency mechanism **MUST** support Service Frames (comprising C-Tags, priority tag and untagged Ethernet frames) at the UNI reference point, as defined in the MEF 10.2, “Ethernet Services Attributes - Phase 2” [6].

[R15] The resiliency mechanism **MUST** support ENNI Frames, as well as L2CPs exchanged between the peering ENNI-N functions, as defined in the ENNI specification [6].

[R16] The resiliency mechanism **MUST NOT** modify the Ethernet frames at the EI, however, it MAY add additional fields to the frames. In this case the FCS **MAY** be modified.

Note that this requirement is designed to support the “preservation” Service Attributes. Note also that this requirement does not preclude encapsulation inside the Interconnection Zone, as long as the frames will enter the adjacent network unmodified.

[D1] The resiliency mechanism **SHOULD** minimize the probability to negligible such that Ethernet Frames (unicast, multicast and broadcast frames) of a single Service

End Point are not delivered more than once to the adjacent network beyond the Interconnection Zone.

Note that frames may not be delivered to the adjacent network during the resiliency switching time.

[R17] The resiliency mechanism **MUST** provide a timestamp of the last switchover.

6.3 Requirements Addressing Triggers for Recovery Actions

6.3.1 Requirements addressing operator manual commands

[R18] When a resilient Link Connection is pre-configured, the resiliency mechanism **MUST** support Operator manual commands to switch Service End Points from Active Link Connection to Standby Resilient Link Connection in the Interconnection Zone.

[R19] The resiliency mechanism **MUST** support the Lock command as defined in ITU-T Recommendation Y.1731 [9].

6.3.2 Requirements Addressing Failure Events

[R20] The resiliency mechanism **SHOULD** be able to detect failures causing performance violations of the SLS (e.g., Link Connection failure) and switch the traffic from each affected Service End Point to a resilient Link Connection.

6.4 Requirements Addressing Configuration Aspects

[R21] The resiliency mechanism **MUST** support the ability to manually map Service End Points to specific pre-configured Link Connections in the UNI Interconnection Zone.

[R22] The resiliency mechanism **MUST** support the ability to manually map Service End Points to specific pre-configured Link Connections in the ENNI Interconnection Zone.

Note that there is no requirement to map all Service End Points of the same service to a single Link Connection. This decision is left to the discretion of the Operators. An example of multiple End Points for a single service is the case of Hairpin.

[R23] The resiliency mechanism **MUST** support the ability to configure a Service End Point as 'unprotected'.

Note that in this case when a failure occurs, the Service End Point will not participate in resiliency. This requirement addresses cases where Services are, for example, protected end-to-end by another mechanism and hence do not require local protection at the Interconnection Zone.

[R24] The resiliency mechanism **MUST** support a Management System's ability to retrieve the mapping configuration of Service End Points to Link Connections during normal and failure conditions in the Interconnection Zone.

[R25] The resiliency mechanism **MUST** support a Management System's ability to retrieve the protection state, as defined in [R9], of each Service End Point in the Interconnection Zone.

[R26] The resiliency mechanism **MUST** support the ability to operate in both Revertive and Non-Revertive Modes in the Interconnection Zone.

[R27] The resiliency mechanism **MUST** set Revertive Mode as the default mode per Service End Point in the Interconnection Zone, when one or more Protection Link Connections are configured.

[R28] The resiliency mechanism **MUST** have a configurable time to wait before reverting the traffic back to the repaired Working Connection Link, when one or more Protection Link Connections are configured.

This means that after the failure causing the resiliency switching is repaired, the resiliency mechanism will switch back to the repaired Working Link Connection.

[O1] The resiliency mechanism **MAY** support both Service End Points that operate in Revertive Mode together with other Service End Points that operate in Non-Revertive Mode in the same Link Connection in the Interconnection Zone.

[R29] The resiliency mechanism **MUST** support the ability to detect miss-configuration of Link Connections across an EI, in the Interconnection Zone.

6.5 Requirements For Scalability And Performance

[R30] The resiliency mechanism **MUST** support 4094 Service End Points in the UNI Interconnection Zone.

[R31] The resiliency mechanism **MUST** support at least 4094 Service End Points in the ENNI Interconnection Zone.

[R32] Automatic and manual traffic redirection between the Working Link Connection and a Resilient Link Connection in the Interconnection Zone **MUST** be performed in not more than 500 ms.

[D2] The Automatic and manual traffic redirection between the Working Link Connection and a Resilient Link Connection in the Interconnection Zone **SHOULD** be performed in not more than 250 ms.

[D3] Manual traffic redirection switching **SHOULD** be hitless (i.e., zero frame loss).

References

[1] RFC 2119, “Key words for use in RFCs to Indicate Requirement Levels”, S. Bradner

[2] MEF Technical Specification MEF 4, Metro Ethernet Network Architecture Framework - Part 1: Generic Framework

[3] MEF Technical Specification MEF 6.1, “Ethernet Services Definitions - Phase 2”

[4] MEF Technical Specification, MEF 10.2, “Ethernet Services Attributes - Phase 2”

[5] MEF Technical Specification, MEF 13 “User Network Interface (UNI) Type 1 Implementation Agreement”

[6] MEF Technical Specification, MEF 26, “External Network Network Interface (ENNI)”

[7] MEF Technical Specification, MEF 28, “External Network Network Interface (ENNI) - Support for UNI Tunnel Access and Virtual UNI”

[8] IEEE802.1AX -2008, "Standard for Local and Metropolitan Area Networks – Link Aggregation"

[9] ITU-T Recommendation Y.1731, OAM functions and mechanisms for Ethernet based networks", 2006.

[10] ITU-T Recommendation Z.341, "Man-Machine Language, Glossary of Terms", 1993.

7. Appendix A – Several Operator Networks And Their Associated Interconnection Zones (Informative)

This Appendix provides an informative description of one model for protecting OVCs across several Interconnection Zones.

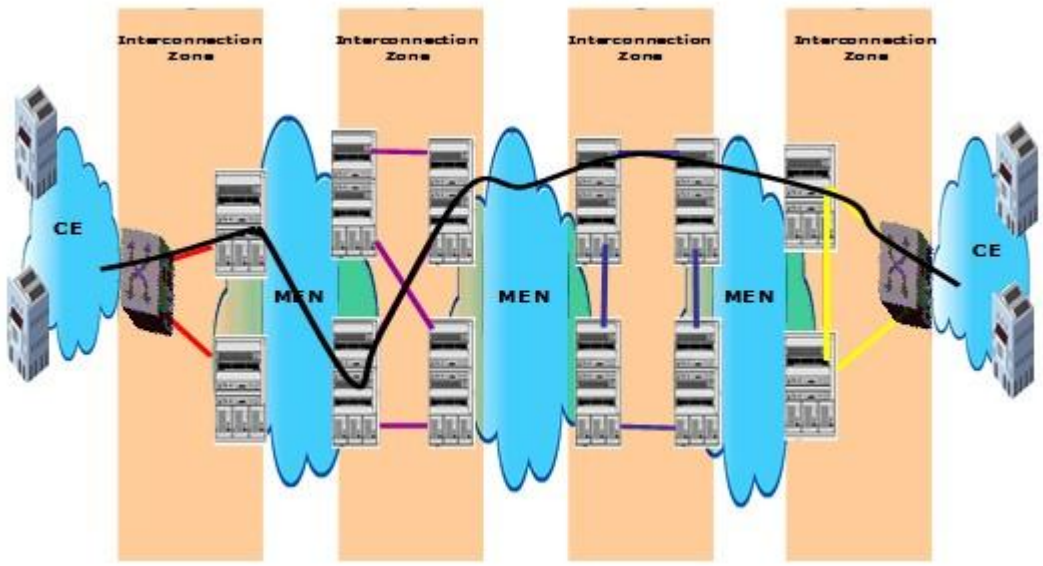

Figure 2: Example of an EVC traversing several Operators' Networks and Interconnection Zones

The figure illustrates various possible interconnection topologies between the Subscriber premises and MENs as well as between the MENs themselves. Note that the interconnection topologies shown in this figure are just for illustration