Implementation Agreement

MEF 35

Service OAM Performance Monitoring Implementation Agreement

April 2012

Table of Contents

List of Figures

List of Tables

1. Abstract

This document specifies an Implementation Agreement (IA) for Service Operations, Administration, and Maintenance (SOAM) that satisfies and extends the Performance Monitoring (PM) framework and requirements described in MEF 17 [16].

Existing PM functions are defined by ITU-T Y.1731 [1], ITU-T G.8021 [4] and ITU-T G.8021 Amendment 1 [5]. This document details how to use these functions in order to achieve the requirements of MEF SOAM PM.

2. Terminology

Term	Definition	Reference
1DM	One-way Delay Measurement Message.	ITU-T Y.1731 [1]
Availability	A measure of the percentage of time that a service is useable.	MEF 10.2.1 [13]
Availability Indicator	A binary indication of whether an interval Δt is available or not.	MEF 10.2.1 [13]
Availability flr	The Availability flr (in contrast with FLR) is the ratio of lost frames over a small interval of time Δt (e.g. 1 sec).	MEF 10.2.1 [13]
Availability Window	A period of n consecutive intervals of Δt, used to determine whether the Availability state has been entered or exited.	
Backward	OAM information sent from the Responder MEP to the Controller MEP. For LM it indicates the frame counts/loss from the Responder MEP to the Controller MEP. For DM is indicates the delay and delay variation from the Responder MEP to the Controller MEP.	
CHLI	Consecutive High Loss Interval	MEF 10.2.1 [13]
Controller MEP	The Controller MEP initiates SOAM PM PDUs, and in a single-ended session receives responses from the Responder MEP.	
CoS	Class of Service	MEF 23.1 [19]
CoS ID	Class of Service Identifier	MEF 23.1 [19]
CoS ID for SOAM PM frames	Class of Service Identifier for SOAM PM frames. The mechanism and/or values of the parameters in the mechanism to be used to identify the CoS Name (H, M, L) that applies to a given SOAM frame.	

CoS Frame Set	Class of Service Frame Set A set of Service or ENNI Frames that have a commitment from the Operator or Service Provider subject to a particular set of performance objectives.	MEF 23.1 [19]
CoS FS	Class of Service Frame Set	MEF 23.1 [19]
CoS Name	Class of Service Name A designation given to one or more sets of performance objectives and associated parameters by the Service Provider or Operator.	MEF 23.1 [19]
CoV	Coefficient of Variation	
DEI	Discard Eligible Indicator	IEEE 802.1Q-2011 [24]
DM	Delay Measurement	
DMM	Delay Measurement Message	ITU-T Y.1731 [1]
DMR	Delay Measurement Response	ITU-T Y.1731 [1]
Dual-Ended	A type of process where a MEP sends measurement information to a peer MEP that will perform the calculations.	
EEAF	Ethernet ECS Adaptation Function	MEF 12.1[14]
EEIF	Ethernet EC Interworking Function	MEF 12.1[14]
EI	External Interface – Either a UNI or an ENNI	MEF 12.1 [14]
E-LAN	An Ethernet service type that is based on a Multipoint-to-Multipoint EVC.	MEF 6.1 [10]
EMS	Element Management System	MEF 15 [15]
ENNI	External Network-to-Network Interface	MEF 4 [9]
EPCF	Ethernet Provider Conditioning Function	MEF 12.1[14]
ESCF	Ethernet Subscriber Conditioning Function	MEF 12.1[14]
ETH-DM	Ethernet Frame Delay Measurement Function (Term is only used to reference the ITU-T definition)	ITU-T Y.1731 [1]
ETH-LM	Ethernet Frame Loss Measurement Function (Term is only used to reference the ITU-T definition)	ITU-T Y.1731 [1]
ETH-SLM	Ethernet Synthetic Loss Measurement Function (Term is only used to reference the ITU-T definition)	ITU-T Y.1731 [1]
Ethernet Frame	A data packet on a wire from preamble to FCS	IEEE 802-2001 [23]
EVC	Ethernet Virtual Connection An association of two or more UNIs that limits the exchange of Service Frames to UNIs in the Ethernet Virtual Connection.	MEF 10.2 [12]
FCS	Frame Check Sequence	
FD	Frame Delay	MEF 10.2 [12]

FDR	Frame Delay Range	MEF 10.2 [12]
FLR	Frame Loss Ratio	MEF 10.2 [12]
Forward	OAM information sent from the Controller MEP to the Responder MEP. For LM it indicates the frame counts/loss from the Controller MEP to the Responder MEP. For DM is indicates the delay and delay variation from the Controller MEP to the Responder MEP.	
HLI	High Loss Interval	MEF 10.2.1 [13]
IA	Implementation Agreement	
IEEE	Institute of Electrical and Electronics Engineers	
IFDV	Inter-Frame Delay Variation	MEF 10.2 [12]
ITU-T	International Telecommunication Union - Telecommunication Standardization Bureau	
LAN	Local Area Network	
LM	Loss Measurement	
LMM	Loss Measurement Message	ITU-T Y.1731 [1]
LMR	Loss Measurement Reply	ITU-T Y.1731 [1]
MAC	Media Access Control	
MA	Maintenance Association A set of MEPs, each configured with the same MAID and MD Level, established to verify the integrity of a single service instance. An MA can also be thought of as a full mesh of Maintenance Entities among a set of MEPs so configured. This term is equivalent to a Maintenance Entity Group, or MEG, as defined by ITU-T Y.1731 [1], which is the term used in this IA.	IEEE 802.1Q-2011 [24]
MAID	Maintenance Association Identifier. An identifier for a Maintenance Association, unique over the OAM Domain. The MAID has two parts: the MD Name and the Short MA Name. These are discussed in the SOAM FM IA. A MAID is equivalent to a MEG ID (as defined by ITU-T Y.1731 [1]), which is the term used in this IA.	IEEE 802.1Q-2011 [24]
Maintenance Interval	A Maintenance Interval is a time interval agreed to by the Service Provider and Subscriber during which the service may not perform well or at all.	MEF 10.2.1 [13]

MD	Maintenance Domain. The network or the part of the network for which faults in connectivity can be managed. This term is equivalent to an OAM Domain, as defined by MEF 17 [16] and used in MEF 30 [22] (which is the term used in this IA).	IEEE 802.1Q-2011 [24]
MD Level	Maintenance Domain Level. An integer in a field in a SOAM PDU with a value in the range (0..7) that is used, along with the VID in the VLAN tag, to identify to which Maintenance Domain among those associated with the SOAM PDU's VID, and thus to which MEG, a SOAM PDU belongs. The MD Level determines the MPs a) that are interested in the contents of a SOAM PDU, and b) through which the frame carrying that SOAM PDU is allowed to pass. This term is equivalent to MEG Level (defined in ITU-T Y.1731 [1]), which is the term used in this IA.	IEEE 802.1Q-2011 [24]
ME	Maintenance Entity. A point-to-point relationship between two MEPs within a single MA. This term is equivalent to a Maintenance Entity, or ME, as defined by ITU-T Y.1731 [1].	IEEE 802.1Q-2011 [24]
Measurement Bin	A Measurement Bin is a counter that stores the number of delay measurements falling within a specified range, during a Measurement Interval.	
Measurement Interval	A period of time during which measurements are taken. Measurements initiated during one Measurement Interval are kept separate from measurements taken during other Measurement Intervals. It is important to note that this is different from T.	
Measurement Interval Data Set	The collection of completed measurements that were initiated during a Measurement Interval.	
MEF	Metro Ethernet Forum	
MEG	Maintenance Entity Group Note that IEEE 802.1Q-2011 [24] uses the term Maintenance Association, while ITU-T Y.1731 [1] uses “Maintenance Entity Group”. These two terms are equivalent. This is the term used in this IA.	ITU-T Y.1731 [1]

MEG Level	Maintenance Entity Group Level A small integer in a field in a SOAM PDU that is used, along with the VID in the VLAN tag, to identify to which Maintenance Domain among those associated with the SOAM PDU's VID, and thus to which ME, a SOAM PDU belongs. The MEG Level determines the MPs a) that are interested in the contents of a SOAM PDU, and b) through which the frame carrying that SOAM PDU is allowed to pass. Note that IEEE uses the term "MD Level", but MEG Level is the term used in this IA.	ITU-T Y.1731 [1]
MEN	Metro Ethernet Network	MEF 4 [9]
MEP	Maintenance association End Point (IEEE 802.1Q-2011 [24]), or equivalently MEG End Point (ITU-T Y.1731 [1] or MEF 17 [16]). An actively managed SOAM entity associated with a specific service instance that can generate and receive SOAM PDUs and track any responses. It is an end point of a single MEG, and is an end-point of a separate Maintenance Entity for each of the other MEPs in the same MEG.	IEEE 802.1Q-2011 [24] ITU-T Y.1731 [1] MEF 17 [16]
MFD	Mean Frame Delay	MEF 10.2[12]
MIP	Maintenance domain Intermediate Point (IEEE 802.1Q-2011 [24]) or equivalently a MEG Intermediate Point (ITU-T Y.1731 [1] or MEF 17 [16]).	IEEE 802.1Q-2011 [24] ITU-T Y.1731 [1] MEF 17 [16]
MP	Maintenance Point One of either a MEP or a MIP.	IEEE 802.1Q-2011 [24]
MTU	Maximum Transmission Unit	MEF 10.2[12]
NE	Network Element	MEF 15 [15]
NMS	Network Management System	MEF 15 [15]
OAM	Operations, Administration, and Maintenance	MEF 17 [16]
OAM Domain	See MD (Maintenance Domain)	MEF 30 [22]
On-Demand	OAM actions that are initiated via manual intervention for a limited time to carry out diagnostics. On-Demand OAM can result in singular or periodic OAM actions during the diagnostic time interval.	RFC 5951 [26]
One-way	A measurement performed in the forward or backward direction. For example from MEP A to MEP B or from MEP B to MEP A.	
OVC	Operator Virtual Connection	MEF 26 [20]
P2P	Point-to-Point	

PCP	Priority Code Point	IEEE 802.1Q-2011 [24]
PDU	Protocol Data Unit	
PM	Performance Monitoring involves the collection of data concerning the performance of the network.	ITU-T M.3400 [7]
PM Function	A MEP capability specified for performance monitoring purposes (e.g., Single-Ended Delay, Single-Ended Synthetic Loss)	
PM Session	A PM Session is the application of a given PM Function between a given pair of MEPs and using a given CoS Frame Set over some (possibly indefinite) period of time.	
PM Solution	A PM Solution is a set of related requirements that when implemented allow a given set of performance metrics to be measured using a given set of PM functions.	
PM Tool	A generic term used to discuss the application of a PM Function.	
PT	Performance Tier	MEF 23.1 [19]
Proactive	OAM actions that are carried on continuously to permit timely reporting of fault and/or performance status.	RFC 5951 [26]
Responder MEP	In a single-ended session, the Responder MEP receives SOAM PM PDUs, from the Controller MEP, and transmits a response to the Controller MEP.	
RFC	Request For Comment	
S	A non-empty subset of ordered UNI pairs within a MEG	MEF 10.2.1 [13]
Service Frame	An Ethernet frame transmitted across the UNI toward the Service Provider or an Ethernet frame transmitted across the UNI toward the Subscriber.	MEF 10.2 [12]
Single-Ended	A type of process where a MEP sends a measurement request and the peer MEP replies with the requested information so the originating MEP can calculate the measurement.	
Sink MEP	In a dual-ended session, the Sink MEP receives SOAM PM PDUs, from the Controller MEP and performs the performance calculations.	
SLM	Synthetic Loss Message	ITU-T Y.1731 [1]
SLR	Synthetic Loss Reply	ITU-T Y.1731 [1]
SLS	Service Level Specification	MEF 10.2 [12]
SOAM	Service Operations, Administration, and Maintenance	MEF 17 [16]
SOAM PM CoS ID	See CoS ID for SOAM PM frames	

SOAM PM Implementation	Capabilities of an NE that are required to support SOAM Performance Monitoring.	
SOAM PDU	Service OAM Protocol Data Unit. Specifically, those PDUs defined in IEEE 802.1Q-2011 [24], ITU-T Y.1731 [1], or MEF specifications. In ITU-T documents the equivalent term OAM PDU is used.	
SOAM PM PDU	Service OAM Protocol Data Unit specifically for Performance Measurement. Examples are LMM/LMR, DMM/DMR/1DM, SLM/SLR.	
Synthetic Frame	An Ethernet frame created to emulate service traffic, carry additional information necessary to support calculating delay or loss and that is treated the same way as a Service Frame.	
Synthetic Traffic	SOAM traffic that emulates service traffic in order to measure the performance experience. Delay measurements must use synthetic traffic, because user traffic does not contain standardized timestamp fields. Other measurements, such as Frame Loss, may also use synthetic frames for certain advantages (e.g., ability to measure loss in multipoint services).	
T	Time Interval for SLS Metrics. The time over which a Performance Metric is defined. It is important to note that this is different from Measurement Interval. T is at least as large as the Measurement Interval, and generally consists of multiple Measurement Intervals. Also note that T can have different values for different performance metrics.	MEF 10.2 [12]
ToD	Time-of-day	
Traffic Conditioning	A process that classifies the traffic units according to configured rules and ensures traffic is conformant before forwarding the traffic.	ITU-T G.8010 [3]
Two-way	A measurement of the performance of frames that flow from the Controller MEP to Responder MEP and back again.	
UTC	Coordinated Universal Time	ISO 8601 [25]
UNI	User-to-Network Interface	MEF 10.2 [12]
VID	Virtual Local Area Network Identifier	IEEE 802.1Q-2011 [24]
VLAN	Virtual Local Area Network	IEEE 802.1Q-2011 [24]

Table 1 - Terminology and Definitions

3. Scope

The scope of this document is to define an Implementation Agreement (IA) for MEF Service Operations, Administration, and Maintenance (SOAM) Performance Monitoring (PM). These requirements are primarily driven by, but not limited to, MEF 17 [16]. The goal of this IA is to define specific performance measurement procedures and specify solutions for collecting the information needed to compute the performance metrics defined by MEF 10.2 [12] as enhanced by MEF 10.2.1 [13] (and in section 8 and section 9 of Y.1563 [2] as well), that may be included in Service Level Specifications (SLSs) over a typical SLS interval. The solutions use the PM functions defined by ITU-T Y.1731 [1], ITU-T G.8021 [4] and ITU-T G.8021 Amendment 1 [5]. When and if necessary, this document may include enhancements to the protocols and/or procedures of existing PM functions in order to satisfy MEF SOAM PM requirements.

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 [8]. All key words **MUST** be in upper case, bold text.

5. Introduction

Among other things, SOAM provides the protocols, mechanisms, and procedures for monitoring the performance of an Ethernet Virtual Connection (EVC) or an Operator Virtual Connection (OVC) across a defined Maintenance Domain (MD). The term used in MEF 17 [16] (and in this document) for an MD is OAM Domain.

While PM measurements can be used for troubleshooting, this document does not attempt to provide a comprehensive treatment of troubleshooting.

5.1 OAM Domains

SOAM allows a network to be partitioned into a set of hierarchical OAM Domains (see MEF 30 [22] Section 7), where an OAM Domain is a contiguous (sub)-network, and may be further partitioned into additional (sub)-domains.

The OAM Domains relevant to this document, and to which the requirements in sections 9-12 apply are:

- EVC (The span of provided service to a subscriber from UNI to UNI)
- Service Provider (The span of the service viewed by the Service Provider)

- Operator (The span of a portion of the service monitored by a Network Operator)
- ENNI (The span of a portion of a service monitored between Network Operators at the ENNI)

However, the following OAM Domains are not precluded (they are allowed but are out of scope for this IA):

- Subscriber (The span of the provided service from subscriber equipment to subscriber equipment)
- UNI (The span of a portion of the service monitored between the UNI-C and UNI-N)

The following domain is not supported for performance monitoring (and is out of scope for this IA):

- Test (used by service providers to test the connectivity to UNI-C, see MEF 20 [17] Appendix A)

5.2 Maintenance Entities

The following figure illustrates the OAM Domains and Maintenance Entities (MEs) defined by the MEF. The figure illustrates pairs of MEPs (thus MEs) that are communicating across various OAM Domains, and also illustrates the hierarchical relationship between these OAM Domains. MEF 30 [22] identifies the default MEs and the Maintenance Entity Group (MEG) levels.

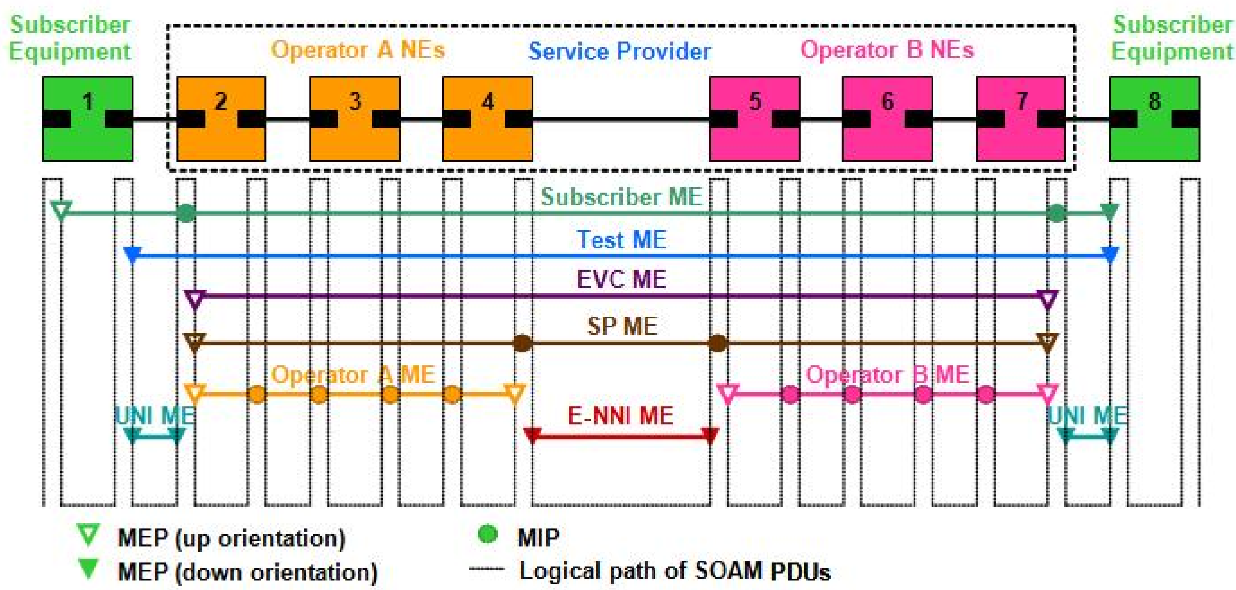

Figure 1 - Maintenance Entities (See MEF 30 [22])

Note that the given MEP and MIP locations, and MEP orientations, are for example purposes only. There are cases where the locations and orientations may differ, and where orientation is not applicable.

In addition, the hierarchical relationship between OAM Domains is also for example purposes only. The scope of an OAM Domain is restricted to its associated VLAN, which has implications when VLAN identifiers are stacked. Service Frames with a C-tag are stacked with a S-tag at the ENNI. The chosen model is sharing the 8 MEG Levels of the OAM Domain space between the C-VLANs and the S-VLANs ensuring no conflicts will occur on MIPs at intermediate nodes, especially at the ENNI interface. MIPs are not involved in performance monitoring so they are not further discussed in this document.

The following figure looks more closely at one example OAM Domain and its MEs. The OAM Domain consists of {MEP1, MEP2, MEP3, MEP4}, where each unique MEP pair (i.e., {{MEP1, MEP2}, {MEP1, MEP3}, {MEP1, MEP4}, {MEP2, MEP3}, {MEP2, MEP4}, {MEP3, MEP4}}) constitutes a ME.

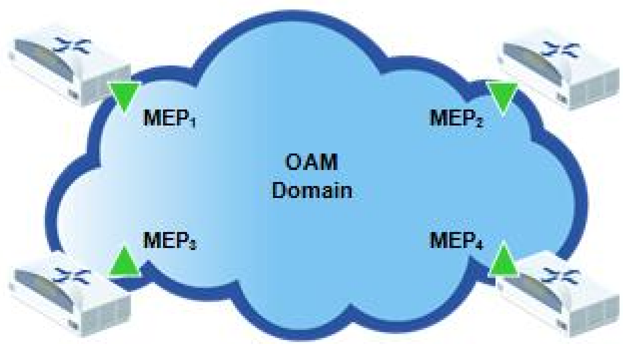

Figure 2 - OAM Domain

5.3 Default Behavior

One of the important functions of this document is to simplify the provisioning of SOAM across the Metro Ethernet Network (MEN). To this end, a default value for an attribute of a maintenance object is defined as the recommended value to be used for that attribute when no other value has been specified during the creation of that object. The use of default values aids interoperability.

Note that the specification of default values does not relieve carriers / equipment of being capable of using a different value if one of the parties has an issue. In other words, specification of a default value assumes that the value is settable and that other values could be used. The default value is suggested as a value to shorten or obviate the need for negotiations in most cases, however other values should also be available for those cases where the default may not be suitable to one of the parties.

6. PM Source Documents

The following sections provide a brief summary of existing MEF specifications having SOAM requirements relating directly (or indirectly) to PM. This discussion is not intended to be complete or exhaustive. For additional information, refer to the corresponding MEF specification.

6.1 MEF 7.1

MEF 7.1 [11] defines the EMS-NMS Information Model that can be used to create interoperable management systems for a Carrier Ethernet network based on MEF specifications.

6.2 MEF 10.2

MEF 10.2 [12] defines service metrics to create MEF compliant services, with some of these being related to performance. The following one-way performance metrics have objectives defined on a per-EVC/OVC per CoS Name basis:

- Frame Delay (FD)
- Frame Delay Range (FDR)
- Mean Frame Delay (MFD)
- Inter-Frame Delay Variation (IFDV)
- Frame Loss Ratio (FLR)
- Availability (see MEF 10.2.1 [13])

The performance metrics encompass service frames flowing in one direction over a subset of ordered EI pairs (i.e., some or all) of an EVC or OVC.

The objectives are uni-directional (specified in MEF 10.2 [12] section 6.9), however, the measurement can be done using bi-directional means. Also see section 6.9.2 in MEF 10.2 [12].

6.3 MEF 10.2.1

MEF 10.2.1 [13] changes MEF 10.2 [12] in a number of important ways. MEF 10.2.1 [13] changes the definition of “qualified” service frame, provides a new definition of Availability, and defines performance metrics for Resiliency such as counts of High Loss Intervals (HLI) and Consecutive High Loss Intervals (CHLI) during the interval T.

Qualified service frames include the following requirements. Each service frame must:

- arrive at the ingress UNI within the time interval T, and within a small time interval Δt that has been designated as part of Available time t

- have a valid Class of Service Identifier for the Class of Service Name in question
- have an Ingress Bandwidth Profile compliance of Green (if it is subject to an Ingress Bandwidth Profile)
- either have no color identifier or a color identifier indicating Green if it is not subject to an Ingress Bandwidth Profile

6.4 MEF 15

MEF 15 [15] defines a number of statistics that NEs should maintain related to the performance of individual services, and the behavior NEs should exhibit related to maintaining and making these statistics available.

6.5 MEF 17

MEF 17 [16] provides a high level overview of SOAM architecture and capabilities, and discusses some of the requirements for a SOAM PM implementation. Included in this specification are the definitions of Connectivity Status for a ME and MA.

According to MEF 17 [16], SOAM must provide the ability to determine Connectivity Status, and measure one-way FLR, two-way FD, and one-way IFDV for point-to-point EVCs. One-way FD and two-way IFDV are listed as optional measurements.

MEF 10.2 [12] and MEF 10.2.1 [13] have evolved over time to include performance metrics based on one-way measurements whereas MEF 17 [16] also reflects considerations for two way measures. MEF 10.2 [12] also defines the metrics FDR and MFD.

6.6 MEF 20

MEF 20 [17] defines SOAM requirements for UNI Type II interfaces or NEs with UNI Type II interfaces, and its scope includes the following OAM Domains:

- Subscriber
- Test (only used by SOAM FM)
- UNI

6.7 MEF 23

MEF 23 [18] defines a common CoS model for all EVC types. See MEF23.1 [19] for updates.

6.8 MEF 23.1

MEF 23.1 [19] defines performance objectives across different performance tiers (PTs) for the 3 CoS model from MEF 23 [18]. It changes how the term "Class of Service" (CoS) is used in MEF specifications. To avoid ambiguity, the terms "CoS" and "CoS ID" are never used on their own, but always with additional context. In this document the term "CoS ID for SOAM PM Frames" (or "SOAM PM CoS ID") is used to describe how a CoS Frame Set is identified for SOAM PM frames.

6.9 MEF 26

MEF 26 [20] defines the requirements for the External Network Network Interface (ENNI). The document specifies a reference point that is the interface between two Metro Ethernet Networks. The term Operator Virtual Connection (OVC) is defined in that document.

6.10 MEF 26.0.3

MEF 26.0.3 [21] adds Service Level Specification performance metric definitions and related requirements to MEF 26 [20].

6.11 MEF 30

MEF 30 (SOAM FM IA) [22] provides the basis for the SOAM terminology used in this document. The SOAM FM IA defines the default configuration for different MEGs. The document has the fault management aspects of SOAM.

7. PM Considerations

The following sections describe specific considerations relating to Delay Measurement, Loss Measurement and handling of multiple Classes of Service.

7.1 Frame Delay Measurements

Measuring the one-way FD of a service frame between two measurement points requires transmission and reception timestamps, where the difference between them corresponds to the one-way FD.

Independent of whether the service frames contain timestamps and sequence numbers, a synthetic frame that does carry that information can be used. This synthetic frame is an Ethernet frame that is created specifically to carry the information necessary to accurately calculate frame delay. If a sufficiently large number of synthetic frames are included in a Measurement Interval, we can assume that the collective experience of these synthetic frames is representative of the performance experience that would be measured during the same Measurement Interval for service frames on the same path. To achieve this, the synthetic frames must be marked so they are treated by the network as belonging to the same class as the service traffic being monitored.

A one-way FD measurement is affected by the accuracy of the transmission and reception timestamps:

- One-way FD is defined in MEF 10.2.1 [13] for qualified frames as the time elapsed from reception at the ingress UNI of the first bit of the service frame until the transmission of the last bit at the egress UNI. However, timestamps are not always taken precisely at these moments.
- To accurately measure one-way FD requires synchronized clocks between the two measurement points, which are impacted by the synchronization method and clock frequency drift. In the absence of clock synchronization, one-way FD can be estimated from the two-way FD.

7.2 Frame Loss Measurements

Measuring the one-way FLR of service frames between two measurement points requires transmission and reception counters, where the one-way FLR can be determined as the ratio of the difference of these quantities to the number of frames transmitted.

Two categories of measurement are possible:

- Measuring the loss of service frames, as specified in Y.1731 [1] and ITU-T G.8021 [4] using the LM process.
- Measuring the loss of synthetic frames (SOAM PM PDUs using SLM/SLR), as specified in Y.1731 [1] and ITU-T G.8021 Amendment 1 [5].

A one-way FLR measurement that measures loss of service frames using the LM process is affected by the accuracy of the transmission and reception counters:

- To accurately measure one-way FLR requires coordinated collection of the counters. Specifically, the reception counter should not be collected until after the last service frame (i.e., the last service frame transmitted prior to collecting the transmission counter) would have been received.
- The counters do not differentiate between green and yellow frames when the traffic conditioning point is not located at the ingress point to the EVC MEG. Rather, the counters reflect only compliant frames.

A limitation of the LM process is that in a multipoint EVC such counters of service frames may not be directly comparable since there are multiple ingress and egress points as well as the potential for frame replication. Similar to delay measurements, directed and periodic synthetic frames can be used. By counting and measuring the one-way FLR of uniform synthetic traffic, statistical methods can be used to estimate the one-way FLR of service traffic. This can be achieved by inserting periodic Synthetic SOAM PM PDUs belonging to a specific CoS Frame Set into an EVC, ensuring that they are treated as green frames by the device inserting them, and measuring the losses of those frames. Advantages of this approach include the ability to measure loss on multipoint connections, the ability to measure loss on a per-CoS Frame Set basis in a straight-forward manner, and the guarantee that there will be traffic to measure. On the other hand, a major challenge of the approach is that the accuracy depends on the rate of Synthetic SOAM PM PDUs; and so, in general longer timeframes are needed to obtain estimates with required accuracy.

7.2.1 Location of PM Measurement Points (for Loss)

As discussed in sections 6.2 and 3, MEF 10.2 [12] (along with MEF 10.2.1 [13]) specifies that the performance metrics are applicable to qualified service frames, which have a level of bandwidth profile conformance determined to be green. This is determined at the traffic conditioning point[1].

Consider an upward facing MEP at an interface, and its placement relative to the traffic conditioning point. Ingress service frame traffic from the customer should encounter the traffic conditioning point before it encounters the performance measurement point. This is consistent with MEF 12.1[14], where the MEP is between the Ethernet Subscriber Conditioning Function (ESCF) and the Ethernet ECS Adaptation Function (EEAF) on a UNI, and between the Ethernet Provider Conditioning Function (EPCF) and the Ethernet EC Interworking Function (EEIF) at an ENNI. This placement also implies that synthetic frames inserted in the upstream direction must be inserted after the traffic conditioning point.

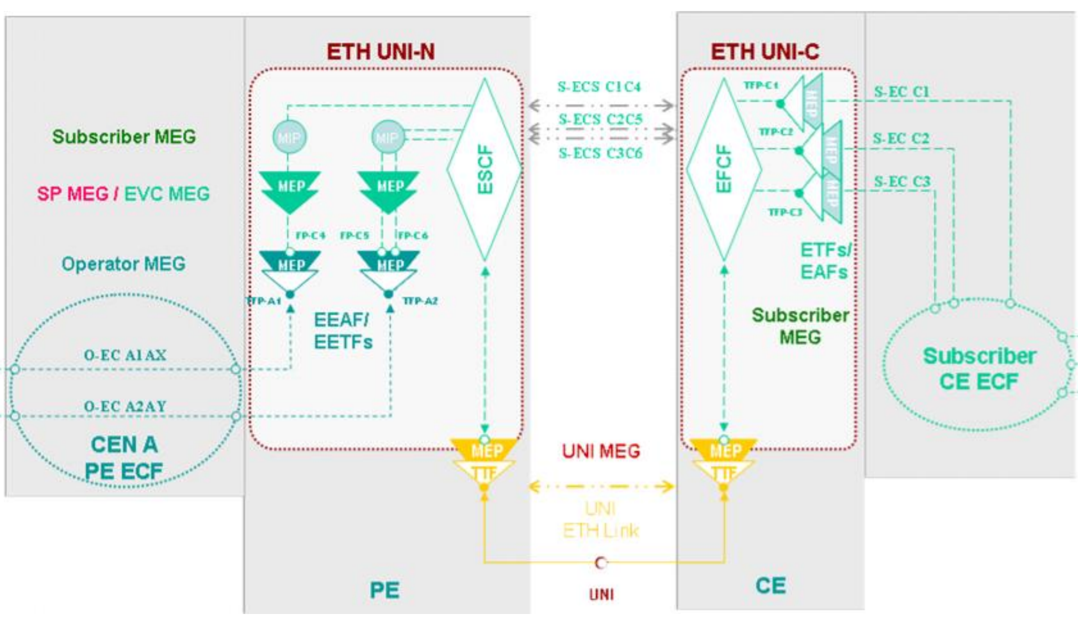

Figure 3 - MEG Placement

[1] Note that in MEF 12.1 [14], the ESCF is the traffic conditioning point for the UNI, and the EPCF is the traffic conditioning point in the ENNI. Also note that both are defined as applying to both ingress and egress traffic conditioning (although egress conditioning is not always applied).

Egress service traffic toward the customer would then encounter the traffic conditioning point after it encounters the MEP. This is reasonable, especially for cases involving multipoint EVCs that can experience focused overloads due to customer behavior (i.e., irrespective of network problems). Such arrangements are likely to use an Egress Bandwidth Profile at the egress UNI that discards frames in the focused overload scenario, and such discards are not indications of network performance problems.

Note that for certain cases, the closer the MEP can be located to the egress link (including the queuing buffers), the more accurate the performance measurements will be. For example, when the UNI link speed is relatively slow and the burst size value is restrictive, the egress buffer at the UNI could be a key contributor for delay and loss impairments.

7.3 CoS Considerations

In Ethernet services, there can be one or more CoS Frame Sets (CoS FS). Performance objectives are defined per CoS FS, {S, PT, CoS ID}, where S is the subset of ordered UNI Pairs or OVC EPs, PT is a certain Performance Tier, and CoS ID is the mechanism to identify the CoS Name at the EI. SOAM PM measurements are taken to measure the performance of traffic belonging to a particular CoS FS, and hence determine whether it meets the performance objectives for that CoS FS.

MEF 23.1 (CoS IA) [19] recognizes that valid CoS IDs differ depending on the type of frame that is being referred to, and defines different types of CoS IDs for Service Frames, ENNI frames mapped to an OVC, and ENNI frames mapped to a VUNI. However, this document is concerned with the types of CoS IDs that can be used by SOAM PM frames as they emulate those frame types. This is called the CoS ID for SOAM PM, or "SOAM PM CoS ID".

The SOAM PM CoS ID can be defined by the following at MEPs:

- VLAN ID
- Combination of the VLAN ID and the PCP value (whether the PCP is an S-tag PCP or a C-tag PCP is determined by the type of MEG)

Frames arriving at an EI can be classified to a CoS Name based on a variety of parameters, for example the VLAN ID, PCP, or (at a UNI) the DSCP or L2CP. However, within the ETH layer, the SOAM PM CoS ID can consist only of untagged or combinations of the VLAN ID and PCP values in one or more VLAN or priority tags.

In order to separately measure the performance of different streams of incoming frames at the EI, those streams can be mapped to different CoS Names (and hence different CoS FSs) within the ETH layer, where each CoS FS uses a SOAM PM CoS ID (as described above) that distinguishes it from other CoS FSs for a given PT and Set S.

To measure the performance of a given CoS FS, SOAM PM PDUs are inserted by the MEPs of the MEG using a SOAM PM CoS ID (that is, either untagged or a combination of VLAN ID and PCP values in one or more tags) that indicates the same CoS FS as the traffic whose performance is being measured. In addition, the SOAM PM PDUs have a Color ID indicating green, i.e., not drop-eligible.

8. PM Solutions

In the context of this specification, a PM Solution is a collection of interdependent and related requirements on the components that implement that solution. A PM Solution uses PM Functions which are capabilities that are specified for performance monitoring purposes (e.g. Single-Ended Delay, Single-Ended Synthetic Loss). A PM Function is associated with an ITU-T PM Tool which is a specific tool that is described in ITU-T Y.1731 (e.g. Single-Ended ETH-SLM). A PM Session is an instantiation of a PM Solution between a given pair of MEPs using a given CoS FS over a given (possibly indefinite) period of time.

The NE is responsible for conducting performance measurements, while the EMS/NMS components are responsible for configuring, collecting, and processing these performance measurements to determine one or more performance metrics for the MEG. An implementation of a PM solution consists of a MEG, supported by NEs in which the MEPs of that MEG are implemented, and the management functionality supported by the EMS and NMS system(s) that typically manage them as shown in Figure 4 - PM Solution Components) below.

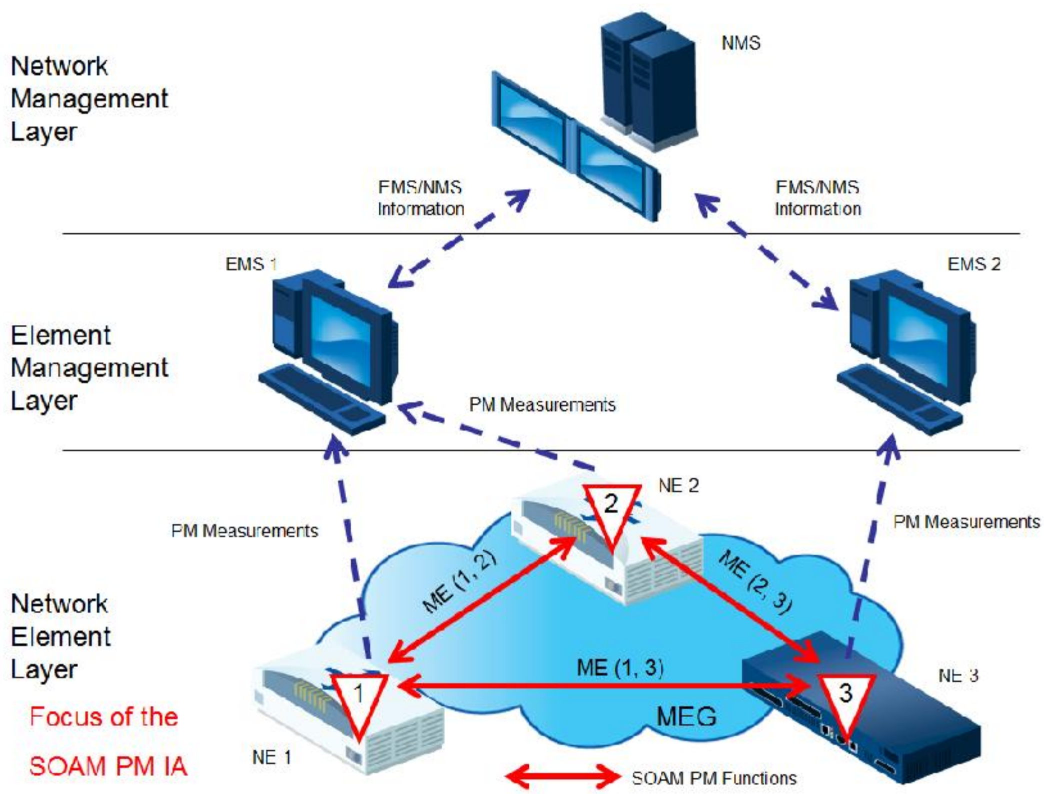

Figure 4 - PM Solution Components

This implementation agreement covers requirements on the components in the Network Element Layer of Figure 4 - PM Solution Components) which shows examples of the network equipment (switches, routers, end stations, or test equipment) that implement the MEPs that make up the MEG. The management systems include the Element Management Systems (EMS) and/or Network Management Systems (NMS) that are responsible for managing the NEs, MEPs and the MEG that is being measured. Requirements on the interface between the EML and the NML are documented in MEF 7.1. [11].

A conforming implementation of a PM Solution provides the SOAM PM and Management mechanisms necessary to meet the goals identified in section 3, including measurement of the performance metrics defined in MEF 10.2 [12] and MEF 10.2.1 [13]. The SOAM mechanisms covered in this IA are realized, in part, through the maintenance association architecture of IEEE 802.1Q-2011 [24] and the PM functions of ITU-T Y.1731 [1] and the (network element based) atomic functions and processes of ITU-T G.8021 [4] as amended by amendment 1 of G.8021[5].

A PM Solution can be categorized as to the types of MEG that they can be applied to and the performance metrics they can measure. A PM solution that can be applied to a MEG with 2 MEPs is a point-to-point solution. A PM Solution that can be applied to a MEG with 2 or more MEPs is a multipoint solution. Note that all multipoint solutions are also point-to-point solutions.

This specification specifies the following PM Solutions:

PM Solution	MEG Type(s)	Measurement Technique for Loss	PM Function(s)	Mandatory or Optional
PM-1	point-to-point multipoint	Synthetic Testing	Single-Ended Delay Single-Ended Synthetic Loss	Mandatory
PM-2	point-to-point multipoint	n/a	Dual-Ended Delay	Optional
PM-3	point-to-point	Counting Service Frames	Single-Ended Service Loss	Optional

Table 2 - PM Solutions Summary

Each PM Session uses a PM Function. Each PM Function uses a specific ITU-T PM Tool which in turn uses specific ITU-T PDU(s), as shown below.

PM Function	ITU-T PM Tool	ITU-T PDU(s)
Single-Ended Delay	ITU-T Two-way ETH-DM	DMM/DMR
Dual-Ended Delay	ITU-T One-way ETH-DM	1DM
Single-Ended Service Loss	ITU-T Single-Ended ETH-LM	LMM/LMR
Single-Ended Synthetic Loss	ITU-T Single-Ended ETH-SLM	SLM/SLR

Table 3 - PM Functions Summary

An overview of the PM Functions (Single-Ended Delay, Dual-Ended Delay, Single-Ended Service Loss Measurement, Single-Ended Synthetic Loss Measurement, and Dual-Ended Loss Measurement) is provided in Appendix A - Performance Management Functions (Informative).

The following figures describe the metrics that can be collected with each PM tool.

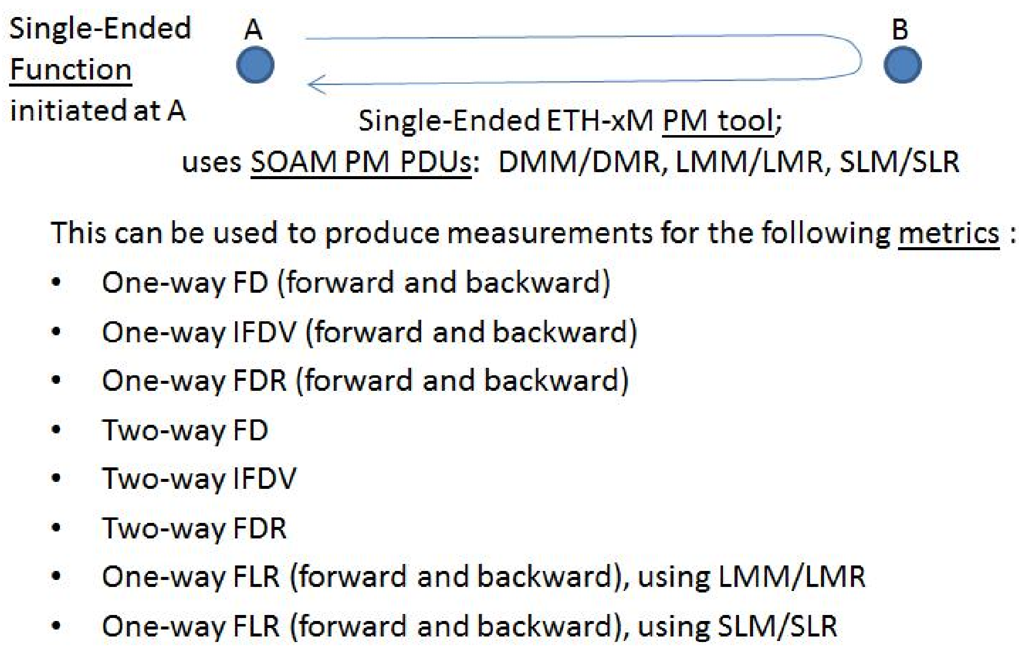

Figure 5 - Metrics that can be collected with Single-Ended Loss and Delay

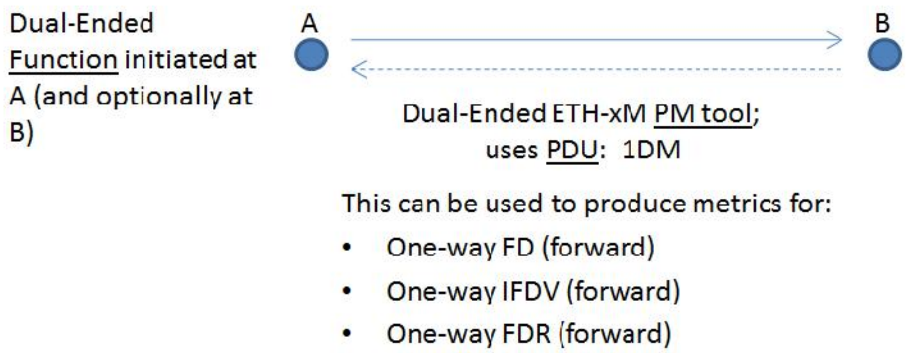

Figure 6 - Metrics that can be collected with Dual-Ended Delay

The following sections serve to briefly describe the individual PM Solutions, which are realized through the NE requirements specified in Section 9 (Common Requirements) and Sections 10 (PM-1), 11 (PM-2), and 12 (PM-3) that follow.

8.1 PM-1: Single-Ended Point-to-Point or Multipoint Delay and Synthetic Loss Measurements

The PM-1 Solution uses synthetic SOAM PM PDUs to measure performance. This solution uses Single-Ended Delay measurement for Frame Delay (FD), Mean Frame Delay (MFD), Frame Delay Range (FDR), and Inter-Frame Delay Variation (IFDV). To measure Frame Loss (FLR), Availability and count of High Loss Intervals (HLI, CHLI), Single-Ended Synthetic Loss measurement is used.

When using DMM/DMR PDUs, DMM frames are sent from a Controller MEP to a Responder MEP which in turn responds with DMR frames. Controller to Responder measurements and Responder to Controller measurements are also known as forward and backward measurements, respectively. With optional time-of-day (ToD) clock synchronization one-way FD, MFD, and FDR measurements can be taken. Two-way FD, MFD, FDR, and IFDV measurements and one-way IFDV measurements can always be taken and do not require ToD clock synchronization. The FD, MFD, FDR, and IFDV delay-related performance metrics as defined in MEF 10.2 [12] and MEF 10.2.1 [13] are made with this solution. For FD, MFD, and FDR, if ToD synchronization is not accurate enough for PM functions, the one-way metrics of MEF 10.2 [12] and MEF 10.2.1 [13] can be estimated by dividing the two-way measurement by 2, although this introduces considerable statistical bias for delay metrics other than MFD. Also note that when measuring one-way FDR, it is necessary to normalize measurements by subtracting the minimum delay. This allows one-way FDR to be measured even if ToD synchronization is not present.

When using SLM/SLR PDUs, SLM frames are sent from a Controller MEP to a Responder MEP which in turn responds with SLR frames. One-way measurements of FLR and Availability are always taken with this mechanism. FLR and Availability are defined in MEF 10.2 [12], and MEF 10.2.1 [13]. See Appendix D - Statistical Considerations for Synthetic Loss Measurement (Informative) for more information about frame loss count accuracy using Synthetic Loss Measurement.

The PM-1 solution using both Single-Ended Delay and Single-Ended Synthetic Loss PM Functions allows all of the performance metrics defined in MEF 10.2 [12] and MEF 10.2.1 [13] to be collected. The PM-1 solution can be applied to point-to-point and multipoint MEGs. Multiple PM Sessions can be run simultaneously between the MEPs, allowing for multiple classes of service to be tested.

DMM and SLM PDUs can be sent to the unicast address of the Responder MEP at the MEG Level of the MEG.

Like any synthetic measurement approach, a PM Session using Single-Ended synthetic loss needs to generate enough SOAM frames to be statistically valid (see Appendix D - Statistical Considerations for Synthetic Loss Measurement (Informative)). All synthetic SOAM PM frames need to be similar to the service frames carried by the EVC, in particular, such SOAM PM frames must have representative frame length and be treated by the network elements between the MEPs in the same way that service frames are treated. In addition, it is important that synthetic SOAM PM frames be inserted irrespective to the load / congestion at the insertion point. To do otherwise would bias measurements away from instances of poor network performance.

The following is a list of the performance metrics defined in MEF 10.2 [12] and MEF 10.2.1 [13] that can be measured for each ordered EI pair in the set S using the PM-1 Solution:

- One-way Frame Delay (MEF 10.2 [12] section 6.9.2)
- One-way Mean Frame Delay (MEF 10.2 [12] section 6.9.2)
- One-way Frame Delay Range (MEF 10.2 [12] section 6.9.2)
- Inter-Frame Delay Variation (MEF 10.2 [12] section 6.9.4)
- One-way Frame Loss Ratio (MEF 10.2 [12] and MEF 10.2.1 [13] section 6.9.6)
- Availability for a EVC or OVC (MEF 10.2.1 [13] section 6.9.8)
- Resiliency-related metrics for EVC or OVC (MEF 10.2.1 [13] section 6.9.9).

8.2 PM-2: Dual-Ended Point-to-Point or Multipoint Delay

The PM-2 solution is an optional solution that uses 1DM PDUs to measure performance. For one-way Frame Delay (FD), Mean Frame Delay (MFD), Frame Delay Range (FDR), and Inter-Frame Delay Variation (IFDV) measurements, Dual-Ended Delay measurement is used.

For Dual-Ended Delay measurement, one-way measurements from a Controller MEP to a Sink MEP (in the forward direction) are taken. Typically, Dual-Ended PM Sessions are configured so that one runs from MEP i to MEP j and another runs from MEP j to MEP i. Only delay-related performance metrics defined in MEF 10.2 [12] and MEF 10.2.1 [13] are made with the PM-2 solution. The PM-2 solution can be applied to either point-to-point or multipoint MEGs. 1DM PDUs can be sent to the unicast address of the Sink MEP. For one-way FD, FDR and MFD, ToD synchronization is required and the considerations described for PM-1 in the previous section also apply to PM-2.

Like any synthetic measurement approach, a PM Session using Dual-Ended Delay needs to generate enough SOAM frames to be statistically valid (see Appendix D - Statistical Considerations for Synthetic Loss Measurement (Informative)). All synthetic SOAM PM frames need to be similar to the service frames carried by the EVC, in particular, such SOAM PM frames must have representative frame length and be treated by the network elements between the MEPs in the same way that service frames are treated. In addition, it is important that synthetic SOAM PM frames be inserted irrespective to the load / congestion at the insertion point. To do otherwise would bias measurements away from instances of poor network performance.

The following is a list of the performance metrics defined in MEF 10.2 [12] that can be measured for each ordered EI pair in the set S using the PM-2 Solution:

- One-way Frame Delay Performance (MEF 10.2 [12] section 6.9.2)
- One-way Mean Frame Delay (MEF 10.2 [12] section 6.9.2)
- One-way Frame Delay Range (MEF 10.2 [12] section 6.9.2)
- Inter-Frame Delay Variation Performance (MEF 10.2 [12] section 6.9.4)

8.3 PM-3: Single-Ended Service Loss Measurements

The PM-3 solution is an optional solution that uses Single-Ended Service Loss tools to collect FLR measurements for each ordered EI pair in the set S.

LMM/LMR PDUs are used for FLR measurements. These collect the counts of the number of service frames transmitted and received by the two MEPs in a point-to-point MEG. When using LMM/LMR PDUs, LMM frames are sent from a Controller MEP to a Responder MEP, which in turn responds with LMR frames. LMM PDUs can be sent to the unicast address of the Responder MEP at the MEG Level of the MEG.

The following is a list of the performance metrics defined in MEF 10.2 [12] and MEF 10.2.1 [13] that can be measured using the PM-3 Solution:

One-way Frame Loss Ratio Performance (FLR) (MEF 10.2.1 [13] section 6.9.6)

See Appendix E – Notes on the Applicability of PM-3 Solutions (Informative) for considerations on the use of PM-3 to measure loss.

9. Common Requirements

This section provides requirements that are applicable to all of the PM Solutions that follow in sections 10 (PM-1), 11 (PM-2), and 12 (PM-3). The requirements below provide for the Life Cycle (starting, stopping etc.), Storage, OAM Domains, and MEP Placement.

Many requirements apply to a “SOAM PM Implementation”, which refers to the capabilities of an NE that are required to support SOAM Performance Monitoring.

9.1 Life Cycle

The requirements of this section apply to the life cycle of a PM Session, and to the scheduling of performance measurements conducted as part of a PM Session. Specifically, scheduling controls when, how long, and how often a PM Session will run.

9.1.1 General Overview of Parameters

The Performance Monitoring process is made up of a number of Performance Monitoring instances, known as PM Sessions. A PM Session is initiated on a Controller MEP to take performance measurements for a given CoS Frame Set and a given Responder/Sink MEP within the same MEG. A PM Session can be classified as either a Loss Measurement Session (LM session) or a Delay Measurement Session (DM session) depending on the PM Function used.

The PM Session is specified by several direct and indirect parameters. A general description of these parameters is listed below, with more detailed requirements provided elsewhere in the document. Note that not every parameter applies to every type of PM Session

- The Endpoints are the Controller MEP and a Responder/Sink MEP.
- The SOAM PM CoS ID for the PM Session identifies the CoS Frame Set for which performance is being measured.
- The PM Function is any of the functions described in Section 8 (for example loss measurement, delay measurement, or synthetic frame loss measurement). A discussion of the PM Functions is provided in Appendix A - Performance Management Functions (Informative).
- The Message Period is the SOAM PM PDU transmission frequency (the time between SOAM PM PDU transmissions).
- The Start Time is the time that the PM Session begins.
- The Stop Time is the time that the PM Session ends.
- The Measurement Intervals are discrete, non-overlapping periods of time during which the PM Session measurements are performed and results are gathered. SOAM PM PDUs

for a PM Session are transmitted only during a Measurement Interval. Key characteristics of Measurement Intervals are the alignment to the clock and the duration of the Measurement Interval. Measurement Intervals can be aligned to either the PM Session Start Time or to a clock, such as the local time-of-day clock. The duration of a Measurement Interval is the length of time spanned by a non-truncated Measurement Interval.

- The Repetition Time is the time between the start times of the Measurement Intervals.

For more details on the interaction between these parameters, refer to Appendix B – Life Cycle Terminology (Informative).

9.1.2 Proactive and On-Demand PM Sessions

A PM Session can be classified as either a Proactive or an On-Demand session. A Proactive session is intended to perpetually measure the performance between the endpoints for the given CoS Frame Set. An On-Demand session is intended to monitor the performance for some finite period of time.

A Proactive session runs all the time once it has been created and started. Since the intent is to provide perpetual performance measurement, Proactive sessions use a Start Time of "immediate" and a Stop Time of "forever". Measurements are collected into multiple fixed length Measurement Intervals covering different periods of time. Measurement Intervals for Proactive sessions are generally aligned to a clock, rather than the Session Start Time. Data is collected and a history of data is stored for a number of Measurement Intervals. Monitoring continues until the PM Session is deleted.

On-Demand sessions are run when needed, and a report is provided at the end. Since On-Demand sessions are intended to cover some finite period of time, absolute or relative Start and Stop Times may be used if those values are known. Alternatively, a Start Time of "immediate" and/or a Stop Time of "forever" may be used (with the intention of manually ending the session when no longer needed), especially if the monitoring period is of unknown duration (e.g., "until troubleshooting is completed".) Measurements may be gathered into one Measurement Interval spanning the entire session duration, or multiple Measurement Intervals covering different periods of time. When multiple Measurement Intervals are used, then historical data from past Measurement Intervals may or may not be stored on the device. In addition, Measurement Intervals may be aligned with the session Start Time or aligned with a clock.

9.1.3 Create

A PM Session has to be created before it can be started. This applies for both On-Demand and Proactive PM Sessions. In order to create a PM Session, a PM Function must be assigned to the PM Session. Requirements relating to specific PM Functions are found in Sections 10, 11, and 12.

[R1] A SOAM PM implementation **MUST** support the capability to simultaneously monitor multiple CoS Frame Sets (CoS FS) for a given EVC.

[R2] If multiple CoS FSs are being monitored at the same time, each **MUST** be monitored by a separate PM session.

[R3] A SOAM PM implementation **MUST** provide a way to indicate whether a PM Session is Proactive or On-Demand.

9.1.4 Delete

The requirements of this section apply to the deletion of a PM Session.

[R4] A SOAM PM implementation **MUST** support the capability to delete a PM Session.

[R5] After the PM Session is deleted, further SOAM PM PDUs relating to the session **MUST NOT** be sent.

[R6] Further measurements associated with the deleted PM Session, **MUST NOT** be made.

[O1] Before the data from the delete PM Session is lost, a SOAM PM implementation **MAY** issue a report (similar to the report that would happen when Stop Time is reached).

[R7] All the stored measurement data relating to the deleted PM Session **MUST** be deleted.

9.1.5 Start and Stop

When a PM Session is started, it can be specified to start immediately, or be scheduled to start in the future.

[R8] For Proactive PM Sessions, the Start Time **MUST** be immediate.

[R9] For On-Demand PM Sessions, a SOAM PM implementation **MUST** support a configurable Start Time per PM Session. The Start Time can be specified as "immediate", as relative time (e.g., a given number of hours, minutes, and seconds from the current time), or as fixed time (e.g., a given UTC date and time).

[D1] The default configured start time **SHOULD** be "immediate".

The following requirements apply to stopping of a PM Session.

[R10] A SOAM PM implementation **MUST** support a configurable Stop Time per PM Session. The Stop Time can be specified as forever or as relative time (e.g., a given number of hours, minutes, and seconds from the start time).

[R11] If the Stop Time is relative time, then the Stop Time **MUST** be equal to or greater than the message period of the PM function(s) comprising the PM Session.

[D2] The default configured Stop Time **SHOULD** be "forever".

[R12] A SOAM PM implementation **MUST** support stopping a PM Session.

[R13] When a PM Session is scheduled to be stopped, it **MUST** cease running.

[R14] Once a PM Session has been stopped, further measurements relating to the session **MUST NOT** be initiated.

[R15] When a PM Session is stopped, the stored measurements relating to the PM Session **MUST NOT** be deleted.

9.1.6 Measurement Intervals

For the duration of a PM Session, measurements are partitioned into fixed-length Measurement Intervals. The length of the period of time associated with a Measurement Interval is called the duration of the Measurement Interval. The results of the measurements are captured in a Measurement Interval Data Set. The results in a Measurement Interval Data Set are stored separately from the results of measurements performed during other Measurement Intervals. This section contains requirements pertaining to Measurement Intervals in the Life Cycle of the PM Session. Requirements pertaining to storage of Measurement Interval data Sets are found under Storage (Section 9.2).

[R16] A SOAM PM implementation **MUST** support a configurable duration for Measurement Intervals.

[R17] A SOAM PM implementation **MUST** support a Measurement Interval with duration of 15 minutes.

9.1.7 Repetition time

For each PM Session, a Repetition Time can be specified if it is not desirable to perform measurements continuously. If the Repetition time is none, then a new Measurement Interval is started immediately after the previous one finishes, and hence performance measurements are made continuously. If a Repetition time is specified, a new Measurement Interval is not started until after Repetition Time has passed since the previous Measurement Interval started. During the time between the end of the previous Measurement Interval and the start of the next one, no SOAM PM PDUs are sent relating to the PM Session, and no measurements are initiated.

[R18] A SOAM PM implementation **MUST** support a configurable Repetition Time per PM Session. The Repetition Time can be specified as none or in relative time (e.g., every given number of hours, minutes, and seconds from the start time).

[D3] The default configured Repetition Time **SHOULD** be none.

[R19] If the Repetition Time is a relative time, the time specified **MUST** be greater than the duration of the Measurement Interval.

[R20] During the time between two Measurement Intervals, SOAM PM PDUs relating to the Session **MUST NOT** be sent.

9.1.8 Alignment of Measurement Intervals

The following requirements pertain to the alignment of Measurement Intervals with time-of-day clock or PM Session start time.

[D4] A SOAM PM implementation **SHOULD** by default align the start of a Measurement Interval on a boundary of the local time-of-day clock that is divisible by the duration of the Measurement Interval (when Repetition Time is "none").

[D5] A SOAM PM implementation **SHOULD** by default align the start of a Measurement Interval on a boundary of the local time-of-day clock that is divisible by the Repetition Time (when Repetition Time is not "none").

[D6] A SOAM PM implementation **SHOULD** allow for no alignment to the time-of-day clock.

[D7] A SOAM PM implementation **SHOULD** support a configurable (in minutes) boundary (offset from ToD time) for alignment of the start of a Measurement Interval.

For example, if the Measurement Interval is 15 minutes and the Repetition Time is none and if ToD offset is 5 mins, the Measurement Intervals would start at 5, 20, 35, 50 minutes past each hour.

9.1.9 Summary of Time Parameters

Possible values for the time parameters are summarized in the table below:

Attribute	Possible Values	PM Session Type
Start Time	Immediately (default)	Proactive or On-Demand
	Relative Time	On-Demand
	Fixed Time	On-Demand

Stop Time	Forever (default) Relative Time	Proactive or On-Demand On-Demand
Repetition Time	"none" Relative Time	Proactive or On-Demand Proactive or On-Demand

Table 4 - Time Parameters

9.2 Storage

The requirements of this section apply to storage of performance measurement results taken during Measurement Intervals, counters or Measurement Bins (for some delay-related parameters). Performance measurements are stored for each Measurement Interval. A Measurement Bin is a counter, and records the number of performance measurements falling within a specified range. Figure 7 - Example of Measurement Intervals and Bins (below) is an example that illustrates the relationship between Measurement Intervals and Measurement Bins:

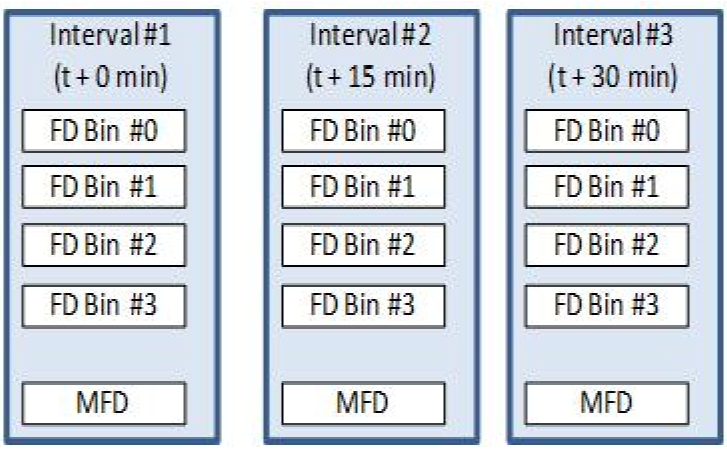

Figure 7 - Example of Measurement Intervals and Bins

Figure 8 shows an example of a MEP running a Single Ended Synthetic Loss PM Function using SLM/SLR. It measures loss, separately for each direction.

Figure 8 - Example of FLR Measurements

9.2.1 Measurement Interval Data Sets

The following requirements apply to the storage of the results of FD, FDR, MFD, IFDV, FLR, Availability or Resiliency performance measurements conducted between a given source and destination pair of MEPs (i.e., ME), for a given PM Session during a given Measurement Interval.

Note that specific requirements relating to the performance parameters that must be stored in a Measurement Interval are enumerated on a per PM Function basis in sections titled PM-1 Requirements, PM-2 Requirements, and PM-3 Requirements.

[R21] A SOAM PM implementation **MUST** store measurement data for a current Measurement Interval and at least 8 hours of historic measurement data (captured per Measurement Interval) for a given data set of a Proactive PM Session.

[D8] A SOAM PM implementation **SHOULD** store measurement data for a current Measurement Interval and at least 24 hours of historic measurement data (captured per Measurement Interval) for a given data set of a Proactive PM Session.

[D9] A SOAM PM implementation **SHOULD** store measurement data for a current Measurement Interval and at least 8 hours of historic measurement data (captured per Measurement Interval) for a given data set of an On-Demand PM Session.

[R22] A SOAM PM implementation **MUST** record the time when the Measurement Interval begins.

[R23] The implementation **MUST** record the value of the local time-of-day clock in UTC at the end of the Measurement Interval.

[R24] A SOAM PM implementation **MUST** support an elapsed time counter per Measurement Interval, which records the number of seconds that have elapsed since the Measurement Interval began.

[D10] A SOAM PM implementation **SHOULD** support synchronization of the local time-of-day clock with UTC to within one second of accuracy.

[R25] A SOAM PM implementation **MUST** record the results of a completed performance measurement as belonging to the Measurement Interval Data Set for the Measurement Interval in which the performance measurement was initiated.

[R26] A SOAM PM PDU arriving more than 5 seconds after the end of the Measurement Interval in which it was initiated **MUST** be discarded and considered lost.

9.2.2 Measurement Bins

The following requirements apply to the use of Measurement Bins for recording the results of delay performance measurements which can be used to determine conformance to FD, IFDV, and FDR objectives conducted between a given source and destination NE for a given PM Session during a Measurement Interval.

FD can be measured over a two-way direction, and/or one-way in the forward and backward directions. The particular FD measurements supported in a SOAM PM implementation depend on the PM Solutions supported and on NE capabilities (e.g., time-of-say clock synchronization between Controller and Responder.) The following requirements apply to each FD measurement supported in a SOAM PM implementation.

[R27] A SOAM PM implementation **MUST** support a configurable number of FD Measurement Bins per Measurement Interval.

[D11] For a SOAM PM implementation, the default number of FD Measurement Bins per Measurement Interval **SHOULD** be 3.

[R28] For a SOAM PM implementation, the minimum number of FD Measurement Bins per Measurement Interval supported **MUST** be 2.

[D12] For a SOAM PM implementation, the maximum number of FD Measurement Bins per Measurement Interval supported **SHOULD** be 10.

IFDV and FDR are measured in both the forward direction and backward direction. The following requirements apply to each direction.

[R29] A SOAM PM implementation **MUST** support a configurable number of IFDV Measurement Bins per Measurement Interval.

[D13] For a SOAM PM implementation, the default number of IFDV Measurement Bins per Measurement Interval supported **SHOULD** be 2.

[R30] For a SOAM PM implementation, the minimum number of IFDV Measurement Bins per Measurement Interval supported **MUST** be 2.

[D14] For a SOAM PM implementation, the maximum number of IFDV Measurement Bins per Measurement Interval supported **SHOULD** be 10.

[R31] A SOAM PM implementation **MUST** support a configurable number of FDR Measurement Bins per Measurement Interval.

[D15] For a SOAM PM implementation, the default number of FDR Measurement Bins per Measurement Interval supported **SHOULD** be 2.

[R32] For a SOAM PM implementation, the minimum number of FDR Measurement Bins per Measurement Interval supported **MUST** be 2.

[D16] For a SOAM PM implementation, the maximum number of FDR Measurement Bins per Measurement Interval supported **SHOULD** be 10.

Note that to support binning, each FDR measurement is normalized by subtracting the estimated minimum of each Measurement Interval (see Appendix G: Normalizing Measurements for FDR (Informative))

The following general Measurement Bin requirements apply. Each bin is associated with a specific range of observed delay, IFDV or FDR. Bins are defined to be contiguous, and each is configured with its lower threshold. Because the bins are contiguous, it is only necessary to configure the lower threshold of each bin. Furthermore, the lowest bin is assumed to always have a threshold of 0, and the highest bin is assumed to have an upper threshold of ∞.

A Measurement Bin is associated with a single counter that can take on non-negative integer values. The counter records the number of measurements whose value falls within the range represented by that bin.

[R33] A SOAM PM implementation **MUST** support a configurable measurement threshold for all but the first Measurement Bin.

[R34] The measurement threshold for each Measurement Bin **MUST** be larger than the measurement threshold of the preceding Measurement Bin.

[R35] The unit for a measurement threshold **MUST** be in microseconds (µs).

[R36] The measurement threshold of the first Measurement Bin **MUST** be fixed to 0µs.

[R37] Received SOAM PM frames with delay values that are greater than or equal to the measurement threshold of a given bin and strictly less than the measurement threshold of the next bin (if any), **MUST** be counted in that, and only that bin.

[D17] The default configured measurement threshold for a Measurement Bin **SHOULD** be an increment of 5000 µs larger than the measurement threshold of the preceding Measurement Bin.

For example, four Measurement Bins gives the following:

Bin	Threshold	Range
bin 0	0 µs	0 µs ≤ measurement < 5,000 µs
bin 1	5,000 µs	5,000 µs ≤ measurement < 10,000 µs
bin 2	10,000 µs	10,000 µs ≤ measurement < 15,000 µs
bin 3	15,000 µs	15,000 µs ≤ measurement < ∞

Table 5 - Example Measurement Bin Configuration

[R38] Each Measurement Bin counter **MUST** be initialized to 0 at the start of the Measurement Interval.

9.2.3 Volatility

The following requirement applies to the volatility of a Measurement Interval.

[D18] A SOAM PM implementation in an NE **SHOULD** store the data for each completed Measurement Interval in local non-volatile memory.

The set of completed Measurement Intervals whose data is stored represents a contiguous and moving window over time, where the data from the oldest historical Measurement Interval is aged out at the completion of the current Measurement Interval.

9.2.4 Measurement Interval Status

The following requirements apply to a discontinuity within a Measurement Interval. Conditions for discontinuity include, but are not limited to, the following:

- Loss of connectivity
- Per section 10.1.6.1 of ITU-T G.7710 [6], the local time-of-day clock is adjusted by at least 10 seconds.
- The conducting of performance measurements is started part way through a Measurement Interval (in the case that Measurement Intervals are not aligned with the start time of the PM Session).
- The conducting of performance measurements is stopped before the current Measurement Interval is completed.

- A local test, failure, or reconfiguration disrupts service on the EVC.
- Maintenance Interval (See MEF 10.2.1 [13])

[R39] A SOAM PM implementation **MUST** support a Suspect Flag per Measurement Interval.

[R40] The Suspect Flag **MUST** be set to false (0) at the start of the current Measurement Interval.

[R41] A SOAM PM implementation **MUST** set the Suspect Flag to true (1) when there is a discontinuity in the performance measurements conducted during the Measurement Interval.

[R42] The value of the Suspect Flag for a Measurement Interval **MUST** always be stored along with the other results for that Measurement Interval when that Measurement Interval's data is moved to history.

9.2.5 Measurement Behavior During Periods of Unavailability and Maintenance Intervals

Measurements of Performance do not apply during Maintenance Intervals. By definition (see MEF 10.2.1), measurements that occur within a Maintenance Interval must not be included in performance metric calculations. When a Measurement Interval lies completely within a Maintenance Interval, its data must be ignored. If a Measurement Interval lies partly within and partly outside of a Maintenance Interval, its data must be marked suspect. Whether this is done by the NE or by an EMS is not specified by this document.

During non-Maintenance Interval time, measurements of Performance apply during periods of availability. This means that if Availability is measured for a given SOAM PM CoS ID on an ME, during periods of unavailability for that SOAM PM CoS ID, measurements of metrics for that same SOAM PM CoS ID (other than Availability) are to be excluded, so such impairments are not double counted. Availability will be evaluated per Maintenance Entity (ME), because a single NE does not necessarily have visibility of all MEs within the MEG.

However, a Maintenance Entity does not know whether it is in a period of unavailability or availability for a SOAM PM CoS ID until a period of n Δt (the Availability Window) has passed, where Δt is a small time interval (e.g., 1 second), and n is the number of consecutive Δt intervals over which Availability transitions are assessed, as defined in Section 6.9 of MEF 10.2.1[2]. Therefore, a PM implementation that is measuring Availability for a SOAM PM CoS ID must store not only the running count of measurements and measurement bins, but also must store information for each Δt within the Availability Window, so the information used in calculating performance metrics can be included/ excluded as dictated by the ME's Availability state for that SOAM PM CoS ID.

[2] *n* consecutive intervals of loss > C are required to transition from the Available to the Unavailable state, and *n* consecutive intervals of loss < C are required to transition from the Unavailable to the Available state. See Section 6.9.8 of MEF10.2.1 for the authoritative discussion.

Correcting the FLR metric for Unavailability periods is of primary importance. Correcting for delay-related metrics is secondary.

[R43] When an Unavailable state is determined to have been entered for a given SOAM PM CoS ID on an ME, a SOAM PM implementation **MUST NOT** include measurements for the previous n Δt intervals in performance metrics for FLR, in any PM session for the same SOAM PM CoS ID and ME.

[D19] When an Unavailable state is determined to have been entered for a given SOAM PM CoS ID on an ME, a SOAM PM implementation **SHOULD NOT** include measurements of the previous n Δt intervals in performance metrics other than FLR and Availability, in any PM sessions for the same SOAM PM CoS ID and ME.

[R44] When an Available state is determined to have been entered for a given SOAM PM CoS ID on an ME, a SOAM PM implementation **MUST** include measurements for the previous n Δt intervals in performance metrics for FLR, in any PM session for the same SOAM PM CoS ID and ME.

[D20] When an Available state is determined to have been entered for a given SOAM PM CoS ID on an ME, a SOAM PM implementation **MUST** include measurements of the previous n Δt intervals in performance metrics other than FLR and Availability, in any PM sessions for the same SOAM PM CoS ID and ME.

A direct consequence of these requirements is that the current counts of a measurement interval cannot be moved into history until an interval of up to $n\,\Delta t$ has passed.

Other direct consequences are that:

- A SOAM PM implementation that is measuring Availability and FLR for a given SOAM PM CoS ID on an ME will need to support the ability to store FLR-related counters for that SOAM PM CoS ID and ME for n previous Δt intervals.

- A SOAM PM implementation that is measuring Availability and metrics other than FLR or Availability for a SOAM PM CoS ID will need to support the ability to store measurements for that SOAM PM CoS ID and ME for n previous Δt intervals.

Note that it is not specified how a SOAM PM implementation stores measurements; e.g., it may store all raw measurements, store a separate set of counters for each Δt, or use other approaches.

Figure 9 shows an example where n sets of counters are maintained for the Availability Window, so counts of lost frames during Availability / Unavailability state transitions can be included / excluded as required. Again, note that this is just an example, and other implementations are possible.

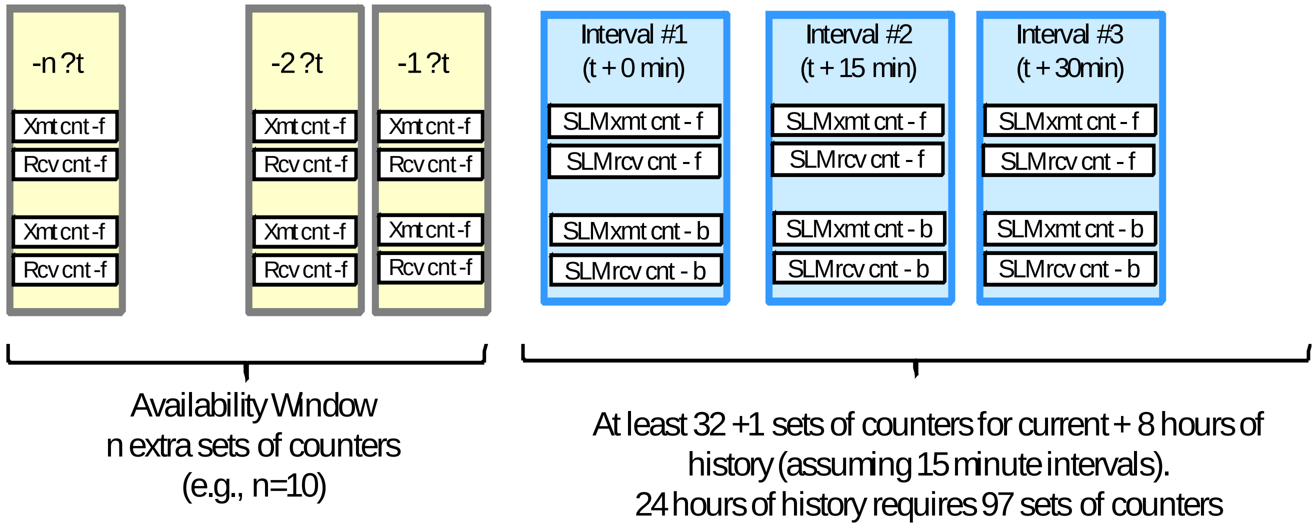

Figure 9 - Example of Measurement Counters to Adjust for Availability

Note that the information stored for each Δt is not reported to the EMS. The MEP PM implementation just uses it locally to perform any necessary adjustments to the counters during transitions.

9.3 OAM Domains

The following requirements provide information about OAM Domains.

[R45] A SOAM PM implementation **MUST** support EVC MEG.

[R46] A SOAM PM implementation **MUST** support Service Provider MEG.

[R47] A SOAM PM implementation **MUST** support Operator MEG.

[R48] A SOAM PM implementation **MUST** support ENNI MEG.

[O2] A SOAM PM implementation **MAY** support Subscriber MEG.

[O3] A SOAM PM implementation **MAY** support UNI MEG.

9.4 MEP Placement

Section 7.2.1 describes the location of measurement points for loss measurement. The following requirements are provided to point out where the MEPs need to be placed in order to support accurate loss measurement.

[R49] On a UNI, the MEP **MUST** be placed between the ESCF and the EEAF.

[R50] On an ENNI, the MEP **MUST** be placed between the EPCF and the EEIF.

10. PM-1 Requirements

The PM-1 solution uses Single-Ended Delay functions for Frame Delay (FD), Frame Delay Range (FDR), Mean Frame Delay (MFD), and Inter-Frame Delay Variation (IFDV) measurements and Single-Ended Synthetic Loss functions for Frame Loss Ratio (FLR) and Availability measurements. The mechanisms support both point-to-point and multipoint connections.

[R51] A SOAM PM implementation **MUST** support the Single-Ended Delay Function as described in Section 10.1.

[R52] A SOAM PM implementation **MUST** support the Single-Ended Synthetic Loss Function as described in Section 10.2.

Section 10.1 lists the requirements for performing Frame Delay and Inter-Frame Delay Variation measurements using the DMM/DMR PDUs. Section 10.2 lists the requirements for performing Frame Loss Ratio and Availability measurements.

Both the Single-Ended Delay and the Single-Ended Synthetic Loss functions can be configured for multiple classes of service per pair of MEPs. Each unique pair of MEPs and Class of Service being monitored forms a unique PM Session. The tools support both point-to-point and multipoint configurations. The two tools can be run at different frequencies.

On multipoint EVCs any subset of the pairs of MEPs can be used and it is not required that measurement be configured for every pair of MEPs. A set of results data will be collected for each pair of MEPs in the configured subset, per class of service. The EMS/NMS will use the data collected for each pair of MEPs in the configured subset and compute a single value for the EVC and Class of Service as specified in MEF 10.2 [12] and MEF 10.2.1 [13].

10.1 Single-Ended Delay Function for Delay, Frame Delay Range, and Inter Delay Variation

The following requirements apply to a SOAM PM implementation of the Single-Ended Delay function and its client application. Each PM Session applies to one ME (i.e., pair of MEPs).

[R53] A SOAM PM implementation **MUST** support the ITU-T Two-way ETH-DM function protocol and the procedures as specified by ITU-T Y.1731 [1], ITU-T G.8021 [4] and ITU-T G.8021 Amendment 1 [5]. Any exceptions to the requirements, behavior, and default characteristics as defined in those specifications are called out in this section.

[R54] A SOAM PM implementation **MUST** support the receive timestamp in the forward direction (***RxTimeStampf***), and the transmit timestamp in the backward direction (***TxTimeStampb***) in the DMR frame. The Controller MEP receives and processes these timestamps and the Responder MEP generates and sends them.

The following requirements specify the ***input parameters*** that are to be supported by the client application for each PM Session.

[R55] A SOAM PM implementation **MUST** support a configurable unicast destination MAC address.

[R56] A SOAM PM implementation **MUST** support a configurable SOAM PM CoS ID for DMM PDU transmission. This requirement is not applicable if the SOAM PM PDUs are untagged.

[D21] The default configured SOAM PM CoS ID **SHOULD** correspond to the CoS Name which yields the best frame delay performance for this MEG.

[R57] An implementation of a Controller MEP **MUST** support configuring a PM Session per class of service with the appropriate SOAM PM CoS IDs for the MEG used across the interface it is associated with.

[R58] The SOAM PM CoS IDs that can be configured **MUST** support at least the following configuration for SOAM PM CoS IDs:

- VLAN ID
- A combination of the PCP and VLAN ID

[R59] An implementation of a Responder MEP **MUST** accept the SOAM PM CoS ID received in a DMM PDU and copy the class of service identifier to the associated DMR response it sends.

[R60] If the MEG is tagged and the VLAN DEI is supported, then a SOAM implementation of a Controller MEP **MUST** use a VLAN DEI of 0 (discard ineligible) for DMM PDU transmission.

[R61] A SOAM PM implementation **MUST** support a configurable period for DMM PDU transmission.

[R62] The periods of {100 ms, 1 sec, 10 sec} **MUST** be supported for DMM PDU transmission.

[D22] The default configured period **SHOULD** be {1 sec}.

[R63] A SOAM PM implementation on the Controller MEP **MUST** support a configurable frame size for DMM PDU transmission.

The frame size corresponds to a valid MEF Service Ethernet frame and is inclusive of the Ethernet header, the DMM PDU with any required PDU padding, and the FCS. This parameter excludes preamble and minimum Interframe Gap. A Data TLV can be used as padding within the DMM PDU based on the Ethernet header applicable at the Controller MEP and the configured frame size.

[R64] The range of Ethernet frame sizes from 64 through 2000 octets **MUST** be supported.

[D23] The range of Ethernet frame sizes from 2001 through 9600 octets **SHOULD** be supported.

[D24] The default configured frame size **SHOULD** be 64 octets, which is the minimum valid Ethernet frame size.

[O4] A SOAM PM implementation **MAY** support the configurable selection of DMR frame pairs for IFDV measurement purposes. A parameter, *n*, is used to control DMR PDU pair selection, where *n* is the ***selection offset***. Given a sequence of received periodic DMR frames, the set of DMR frame pairs can be expressed as $\{ \{f_1, f_{1+n}\}, \{f_2, f_{2+n}\}, \{f_3, f_{3+n}\}, \ldots\}$.

[D25] The default configured selection offset for IFDV **SHOULD** be 1. This parameter, when multiplied by the period parameter of [R61], is equivalent to the IFDV parameter of Δt as specified by MEF 10.2 [12].

[R65] A SOAM PM implementation **MUST** support, for FDR measurement purposes, normalizing delays by subtracting the estimated minimum delay of the interval.

[D26] A SOAM PM implementation **SHOULD** use the observed minimum of the previous Measurement Interval as the estimated minimum delay to normalize FDR measurements at the beginning of a Measurement Interval.

[D27] During the Measurement Interval a SOAM PM implementation **SHOULD** set the estimated minimum to the lower of the previous estimate and the minimum for the current Measurement Interval.

A shift of the minimum may be significant, or it may be minor. The NE relies on the NMS/EMS to determine whether the change in the minimum is such that the FDR measurements for the Measurement Interval should be invalidated. In the case where the minimum has increased, the FDR measurements for the previous Measurement Interval may also need to be invalidated. This is discussed in Appendix G: Normalizing Measurements for FDR (Informative).

The following requirements specify the ***output data set*** that is sent by the Controller MEP to the client application per Measurement Interval.

[R66] A SOAM PM implementation of the Single-Ended Delay PM Function **MUST** support the following counters per Proactive or On-Demand Measurement Interval per PM Session:

Data	Description
Start Time-of-day timestamp	A 64-bit timestamp of the time-of-day at the start of the Measurement Interval.
End Time-of-day timestamp	A 64-bit timestamp of the time-of-day at the end of the Measurement Interval.
Measurement interval elapsed time	A 32-bit counter of the number of seconds of the Measurement Interval as calculated by the NE.
SOAM PM PDUs Sent	A 32-bit counter reflecting the number of SOAM PM PDUs sent.
SOAM PM PDUs Received	A 32-bit counter reflecting the number of SOAM PM PDUs received.
Two-way FD counter per configured FD Measurement Bin	A 32-bit counter per Measurement Bin that counts the number of FD measurements that fall within the configured range.
Mean two-way FD	A 32-bit integer reflecting the average (arithmetic mean) two-way FD measurement in microseconds.
Minimum two-way FD	A 32-bit integer reflecting the minimum two-way FD measurement in microseconds.
Maximum two-way FD	A 32-bit integer reflecting the maximum two-way FD measurement in microseconds.
One-way IFDV counter in the forward direction per configured IFDV Measurement Bin	A 32-bit counter per Measurement Bin that counts the number of IFDV measurements (i.e., each instance of $\|D_i – D_j\|$ in the forward direction that fall within a configured bin.
Mean one-way IFDV in the forward direction	A 32-bit integer reflecting the average (arithmetic mean) one-way IFDV measurement in the forward direction in microseconds.
Minimum one-way IFDV in the forward direction	A 32-bit integer reflecting the minimum one-way IFDV measurement in the forward direction in microseconds.
Maximum one-way IFDV in the forward direction	A 32-bit integer reflecting the maximum one-way IFDV measurement in the forward direction in microseconds.
One-way IFDV counter in the backward direction per configured IFDV Measurement Bin	A 32-bit counter per Measurement Bin that counts the number of IFDV measurements in the backward direction that fall within a configured bin.
Mean one-way IFDV in the backward direction	A 32-bit integer reflecting the average (arithmetic mean) one-way IFDV measurement in the backward direction in microseconds.
Minimum one-way IFDV in the backward direction	A 32-bit integer reflecting the minimum one-way IFDV measurement in the backward direction in microseconds.

Maximum one-way IFDV in the backward direction	A 32-bit integer reflecting the maximum one-way IFDV measurement in the backward direction in microseconds.
One-way FDR counter in the forward direction per configured FDR Measurement Bin	A 32-bit counter per Measurement Bin that counts the number of FDR measurements in the forward direction that fall within a configured bin.
Mean one-way FDR in the forward direction	A 32-bit integer reflecting the average (arithmetic mean) one-way FDR measurement in the forward direction in microseconds.
Maximum one-way FDR in the forward direction	A 32-bit integer reflecting the maximum one-way FDR measurement in the forward direction in microseconds.
One-way FDR counter in the backward direction per configured FDR Measurement Bin	A 32-bit counter per Measurement Bin that counts the number of FDR measurements in the backward direction that fall within a configured bin.
Mean one-way FDR in the backward direction	A 32-bit integer reflecting the average (arithmetic mean) one-way FDR measurement in the backward direction in microseconds.
Maximum one-way FDR in the backward direction	A 32-bit integer reflecting the maximum one-way FDR measurement in the backward direction in microseconds.
Minimum one-way FD in the forward direction	A 32-bit integer reflecting the minimum one-way FD measurement in the forward direction in microseconds.
Minimum one-way FD in the backward direction	A 32-bit integer reflecting the minimum one-way FD measurement in the backward direction in microseconds.

Table 6 - Mandatory Single-Ended Delay Data Set

The minimum one-way FD measurements do not provide intrinsic information about the Frame Delay when time-of-day clock synchronization is not in effect, but are needed to detect changes in the minimum that may invalidate FDR measurements.

[R67] If time-of-day clock synchronization is in effect for both MEPs in the ME, a SOAM PM implementation **MUST** be able to support the following additional data per Measurement Interval per PM Session:

Data	Description
One-way FD counter in the forward direction per configured FD Measurement Bin	A 32-bit counter per Measurement Bin that counts the number of [one-way] FD measurements in the forward direction that fall within the configured bin.
Mean one-way FD in the forward direction	A 32-bit integer reflecting the average (arithmetic mean) one-way FD measurement in the forward direction in microseconds.

Maximum one-way FD in the forward direction	A 32-bit integer reflecting the maximum one-way FD measurement in the forward direction in microseconds.
One-way FD counter in the backward direction per configured FD Measurement Bin	A 32-bit counter per Measurement Bin that counts the number of [one-way] FD measurements in the backward direction that fall within the configured bin
Mean one-way FD in the backward direction	A 32-bit integer reflecting the average (arithmetic mean) one-way FD measurement in the backward direction in microseconds.
Maximum one-way FD in the backward direction	A 32-bit integer reflecting the maximum one-way FD measurement in the backward direction in microseconds.

Table 7 - Mandatory Single-Ended Delay Data Set with Clock Synchronization

10.2 Single-Ended Synthetic Loss Function for Frame Loss Ratio (FLR) and Availability

Single-Ended Synthetic Loss can be configured for multiple classes of service per pair of MEPs. Each unique pair of MEPs and class of service requires a unique PM Session. Single-Ended Synthetic Loss supports both point-to-point and multipoint configurations.

The following requirements apply to a SOAM PM implementation of the Single-Ended Synthetic Loss function and its client application.

[R68] A SOAM PM implementation **MUST** support the ITU-T ETH-SLM protocol and procedures as specified by ITU-T Y.1731 [1] and ITU-T G.8021 Amendment 1 [5]. Any exceptions to the requirements, behavior, and default characteristics as defined in those specifications are called out in this section.

[R69] A SOAM PM implementation **MUST** collect ITU-T ETH-SLM sequence number values applicable for ingress and egress synthetic frames where the sequence numbers maintain a count of transmitted and received synthetic frames between a set of MEPs.

The requirements of this section apply to an instance of the PM-1 solution, as summarized in section 8.1, operating in both directions of a MEP pair for a given ME.

The following requirements specify the ***input parameters*** that are to be supported by the client application for each PM Session.

[R70] A SOAM PM implementation **MUST** support a configurable unicast destination MAC address.

[R71] A SOAM PM implementation **MUST** support a configurable SOAM PM CoS ID for SLM frame transmission. This requirement is not applicable if the SOAM PM PDUs are untagged.

[D28] The default configured SOAM PM CoS ID **SHOULD** correspond to the CoS Name which yields the best frame loss performance for this MEG.

[R72] An implementation of a Controller MEP **MUST** support configuring a PM Session per class of service with the appropriate SOAM PM CoS IDs for the MEG used across the interface it is associated with.

[R73] The SOAM PM CoS IDs that can be configured **MUST** support at least the following configuration for SOAM PM CoS IDs:

- VLAN ID
- A combination of the PCP and VLAN ID

[R74] An implementation of a Responder MEP **MUST** accept SOAM PM CoS IDs received in SLM PDUs and copy the class of service identifier to the associated SLR response it sends.

[R75] If the MEG is tagged and the VLAN DEI is supported, then a SOAM implementation of a Controller MEP **MUST** use a VLAN DEI of 0 (discard ineligible) for SLM PDU transmission.

[R76] A SOAM PM implementation **MUST** support a configurable period for SLM PDU transmission.

[R77] The periods of {100 ms, 1 sec, 10 sec} **MUST** be supported for SLM PDU transmission.

[D29] The period of 10ms **SHOULD** be supported for SLM PDU transmission.

[D30] The default configured period **SHOULD** be {1 sec}.

[R78] A SOAM PM implementation of the Controller MEP must support a configurable frame size for SLM frame transmission.

The frame size corresponds to a valid MEF Service Ethernet frame and is inclusive of the Ethernet header, the SLM PDU with any required PDU padding, and the FCS. This parameter excludes preamble and minimum interframe Gap. A Data TLV can be used as padding within the SLM PDU based on the Ethernet header applicable at the Controller MEP and the configured frame size.

[R79] The range of Ethernet frame sizes from 64 through 2000 octets **MUST** be supported.

[D31] The range of Ethernet frame sizes from 2001 through 9600 octets **SHOULD** be supported.

[D32] The default configured frame size **SHOULD** be 64 octets, which is the minimum valid Ethernet frame size.

When the Single-ended Synthetic Loss Function is used, each transmitted SLM has three possible outcomes: a corresponding SLR is received; the SLM is lost in the forwards direction; or the SLM is lost in the backwards direction. To calculate the forwards or backwards FLR, a number of SLMs must be transmitted, and the corresponding number lost in each direction must be measured. The FLR can then be calculated in the normal way. Note: the more SLMs used for each FLR calculation, the more precise the resulting FLR value will be.

The following requirements apply to the calculation of Availability, which is explained in detail in MEF10.2.1 [13]. A brief summary is that Availability is determined by first calculating the "Availability flr" over a small interval of time Δt and comparing it to a frame loss threshold. If a sufficient number of consecutive Δt intervals exceed the threshold, an Unavailability state is entered. Note that Availability flr is different from FLR, which is calculated over the much larger interval T.

[R80] The number range of 1 through 10 **MUST** be supported for the configurable number of consecutive availability flr measurements to be used to determine Availability/Unavailability transitions. This parameter is equivalent to the Availability parameter of *n* as specified by MEF 10.2.1 [13].

[D33] The default configured number of *n* for Availability **SHOULD** be 10.

The availability flr measurements are the basis to evaluate Availability. Within each small time period *Δt* (e.g., one second), the loss ratio "availability flr" is calculated and compared with a threshold *C*. If a window of consecutive seconds all have loss ratio exceeding the threshold, then an Unavailable state has been entered and all seconds within that window will be designated as having Availability state = 0. Details are in MEF 10.2.1 [13].

[R81] A SOAM PM implementation **MUST** support a configurable availability frame loss ratio threshold to be used in evaluating the availability/unavailability status of an availability indicator per MEF 10.2.1 [13].

[R82] The availability frame loss ratio threshold range of 0.00 through 1.00 **MUST** be supported.

[D34] The default configured availability frame loss ratio threshold **SHOULD** be 0.50.

[R83] A SOAM PM implementation **MUST** report to the managing system whenever a transition between available and unavailable occurs in the status of an adjacent pair of availability indicators per MEF 10.2.1 [13].

[R84] The availability transition report **MUST** include the following data:

Data	Description
Source	Controller MEP
Destination	Responder MEP

Cos ID	SOAM PM CoS ID
Direction	Forward or Backward
Timestamp	Reflects the value of the local time-of-day clock in UTC at the time of transition.
Status	Reflects whether the transition was from available to unavailable, or unavailable to available.

Table 8 - Availability Transition Event Data

[R85] If the NE maintains a time-stamped log, an entry **MUST** also be generated with the same data as the report.

[R86] A SOAM PM implementation **MUST** support a configurable parameter to indicate the number of HLIs that constitute a CHLI. This is equivalent to *p* in MEF 10.2.1 [13].

[D35] The default value for the number of HLIs that constitute a CHLI **SHOULD** be 5.

[D36] The range of values for the number of HLIs that constitute a CHLI **SHOULD** be 1 to (*n* - 1). Where *n* is the Availability parameter as specified in [R80].

The following requirements specify the ***output data set*** that are to be supported by the client application per Measurement Interval.

[R87] A SOAM PM implementation **MUST** support the following additional data per Measurement Interval per PM Session:

Data	Description
Start Time-of-day timestamp	A 64-bit timestamp of the time-of-day at the start of the Measurement Interval.
End Time-of-day timestamp	A 64-bit timestamp of the time-of-day at the end of the Measurement Interval.
Measurement interval elapsed time	A 32-bit counter of the number of seconds of the Measurement Interval as calculated by the NE.
SOAM PM PDUs Sent	A 32-bit counter reflecting the number of SOAM PM PDUs sent.
SOAM PM PDUs Received	A 32-bit counter reflecting the number of SOAM PM PDUs received.
Tx frame count in the forward direction	A 32-bit counter reflecting the number of SLM frames transmitted in the forward direction.
Rx frame count in the forward direction	A 32-bit counter reflecting the number of SLM frames received in the forward direction.
Tx frame count in the backward direction	A 32-bit counter reflecting the number of SLR frames transmitted in the backward direction.
Rx frame count in the backward direction	A 32-bit counter reflecting the number of SLR frames received in the backward direction.

Count of Availability indicators in the forward direction	A 32-bit counter reflecting the number of availability indicators evaluated as available in the forward direction.
Count of Availability indicators in the backward direction	A 32-bit counter reflecting the number of availability indicators evaluated as available in the backward direction.
Count of Unavailable indicators in the forward direction	A 32-bit counter reflecting the number of availability indicators evaluated as unavailable in the forward direction.
Count of Unavailable indicators in the backward direction	A 32-bit counter reflecting the number of availability indicators evaluated as unavailable in the backward direction.
Count of HLIs in the forward direction	Count of HLIs in the forward direction during the Measurement Interval.
Count of HLIs in the backward direction	Count of HLIs in the backward direction during the Measurement Interval.
Count of CHLIs in the forward direction	Count of CHLIs in the forward direction during the Measurement Interval.
Count of CHLIs in the backward direction	Count of CHLIs in the backward direction during the Measurement Interval.

Table 9 - Mandatory Single-Ended Synthetic Loss Data Set

[D37] A SOAM PM implementation **SHOULD** support the following additional availability related data per Measurement Interval per PM Session:

Data	Description
Minimum one-way flr in the forward direction	The minimum one-way flr measurement during this Measurement Interval.
Maximum one-way flr in the forward direction	The maximum one-way flr measurement during this Measurement Interval.
Average one-way flr in the forward direction	The average (arithmetic mean) one-way flr measurement during this Measurement Interval.
Minimum one-way flr in the backward direction	The minimum one-way flr measurement during this Measurement Interval.
Maximum one-way flr in the backward direction	The maximum one-way flr measurement during this Measurement Interval.
Average one-way flr in the backward direction	The average (arithmetic mean) one-way flr measurement during this Measurement Interval.

Table 10 - Optional Single-Ended Synthetic Loss Data Set

11. PM-2 Requirements

The PM-2 solution uses Dual-Ended Delay functions for Frame Delay (FD) and Inter-Frame Delay Variation (IFDV) measurements. The mechanisms support both point-to-point and multipoint connections.

Section 11.1 lists the requirements for performing Frame Delay and Inter-Frame Delay Variation measurements using the Dual-Ended Delay functions.

[O5] A SOAM PM implementation **MAY** support the Dual-Ended Delay Function as described in Section 11.1.

11.1 Dual-Ended Delay Function for Frame Delay, Frame Delay Range, and Inter-Frame Delay Variation

Dual-Ended Delay can be configured for multiple classes of service for each direction in a pair of MEPs. Each unique pair of MEPs, direction and class of service forms a unique PM Session. Dual-Ended Delay supports both point-to-point and multipoint configurations.

On multipoint EVCs any subset of the ordered pairs of MEPs can be used and it is not required to configure measurement for every ordered pair of MEPs, nor for both orders (directions) of any given pair of MEPs. A set of results data will be collected for each ordered pair of MEPs in the configured subset, per class of service. The EMS/NMS will use the data collected for each ordered pair of MEPs in the configured subset and compute a single value for the EVC and Class of Service as specified in MEF 10.2 [12] and MEF 10.2.1 [13].

When using Dual-Ended Delay, a single direction (A->B or B->A) can be measured or both directions can be measured (A->B and B->A.) depending on configuration.

[R88] A SOAM PM implementation **MUST** support the ITU-T One-way ETH-DM Function protocol and procedures as specified by ITU-T Y.1731, [1] ITU-T G.8021 [4] and ITU-T G.8021 Amendment 1 [5]. Any exceptions to the requirements, behavior, and default characteristics as defined in those specifications are called out in this section.

The following requirements specify the ***input parameters*** that are to be supported by the client application for each PM Session.

[R89] A SOAM PM implementation **MUST** support a configurable unicast destination MAC address per MEP being measured to.

[R90] A SOAM PM implementation **MUST** support a configurable SOAM PM CoS ID for 1DM frame transmission. This requirement is not applicable if the SOAM PM PDUs are untagged.

[D38] The default configured SOAM PM CoS ID **SHOULD** correspond to the CoS Name which yields the lowest frame delay performance for this MEG.

[R91] An implementation of a Controller MEP **MUST** support configuring a PM Session per class of service with the appropriate SOAM PM CoS IDs for the MEG used across the interface it is associated with.

[R92] The SOAM PM CoS IDs that can be configured **MUST** support at least the following configuration for SOAM PM CoS IDs:

- VLAN ID
- A combination of the PCP and VLAN ID

[R93] An implementation of a Sink MEP **MUST** accept SOAM PM CoS IDs in received 1DM frames.

[R94] If the MEG is tagged and the VLAN DEI is supported, then a SOAM PM implementation on the Controller MEP **MUST** use a VLAN DEI of 0 (discard ineligible) for 1DM PDU transmission.

[R95] A SOAM PM implementation **MUST** support a configurable period for 1DM PDU transmission.

[R96] The periods of {100 ms, 1 sec, 10 sec} **MUST** be supported for 1DM PDU transmission.

[D39] The default configured period **SHOULD** be {1 sec}.

[R97] A SOAM PM implementation on the Controller MEP **MUST** support a configurable frame size for 1DM PDU transmission.

The frame size corresponds to a valid MEF Service Ethernet frame and is inclusive of the Ethernet header, the 1DM PDU with any required PDU padding, and the FCS. This parameter excludes preamble and minimum interframe Gap. A Data TLV can be used as padding within the 1DM PDU based on the Ethernet header applicable at the Controller MEP and the configured frame size.

[R98] The range of Ethernet frame sizes from 64 through 2000 octets **MUST** be supported.

[D40] The range of Ethernet frame sizes from 2001 through 9600 octets **SHOULD** be supported.

[D41] The default configured frame size **SHOULD** be 64 octets, which is the minimum valid Ethernet frame size.

[O6] A SOAM PM implementation **MAY** support the configurable selection of received 1DM PDU pairs for IFDV measurement purposes. A parameter, *n*, is used to control 1DM PDU pair selection, where *n* is the ***selection offset***. Given a sequence of received periodic 1DM frames, the set of 1DM frame pairs can be expressed as { $\{f_1, f_{1+n}\}$, $\{f_2, f_{2+n}\}$, $\{f_3, f_{3+n}\}$, …}.

[D42] The default configured selection offset **SHOULD** be 1. This parameter, in combination with the period parameter of [R95], is equivalent to the IFDV parameter of Δt as specified by MEF 10.2 [12].

[R99] A SOAM PM implementation **MUST** support, for FDR measurement purposes, normalizing delays by subtracting the estimated minimum delay of the interval.

[D43] A SOAM PM implementation **SHOULD** use the observed minimum of the previous Measurement Interval as the estimated minimum delay to normalize FDR measurements at the beginning of a Measurement Interval.

[D44] During the Measurement Interval a SOAM PM implementation **SHOULD** set the estimated minimum to the lower of the previous estimate and the minimum for the current Measurement Interval.

A shift of the minimum may be significant, or it may be minor. The NE relies on the NMS/EMS to determine whether the change in the minimum is such that the FDR measurements for the Measurement Interval should be invalidated. In the case where the minimum has increased, the FDR measurements for the previous Measurement Interval may also need to be invalidated. This is discussed in Appendix G: Normalizing Measurements for FDR (Informative).

The following requirements specify the process ***output data set*** that are to be supported by the client application per Measurement Interval.

[R100] A SOAM PM implementation **MUST** support the following counters per Measurement Interval per PM Session:

Data	Description
Start Time-of-day timestamp	A 64-bit timestamp of the time-of-day at the start of the Measurement Interval.
End Time-of-day timestamp	A 64-bit timestamp of the time-of-day at the end of the Measurement Interval.
Measurement interval elapsed time	A 32-bit counter of the number of seconds of the Measurement Interval as calculated by the NE.
SOAM PM PDUs Sent	A 32-bit counter reflecting the number of SOAM PM PDUs sent.
SOAM PM PDUs Received	A 32-bit counter reflecting the number of SOAM PM PDUs received.
One-way IFDV counter per configured IFDV Measurement Bin	A 32-bit counter per Measurement Bin that counts the number of IFDV measurements that fall within the configured bin.

Mean one-way IFDV	A 32-bit integer reflecting the average (arithmetic mean) one-way IFDV measurement in microseconds.
Minimum one-way IFDV	A 32-bit integer reflecting the minimum one-way IFDV measurement in microseconds.
Maximum one-way IFDV	A 32-bit integer reflecting the maximum one-way IFDV measurement in microseconds.
One-way FDR counter per configured FDR Measurement Bin	A 32-bit counter per Measurement Bin that counts the number of FDR measurements that fall within a configured bin.
Mean one-way FDR	A 32-bit integer reflecting the average (arithmetic mean) one-way FDR measurement in microseconds.
Maximum one-way FDR	A 32-bit integer reflecting the maximum one-way FDR measurement in microseconds.
Minimum one-way FD	A 32-bit integer reflecting the minimum one-way FD measurement in microseconds.

Table 11 - Mandatory Dual-Ended Delay Data Set

The minimum one-way FD measurement does not provide intrinsic information about the Frame Delay when time-of-day clock synchronization is not in effect, but is needed to detect changes in the minimum that may invalidate FDR measurements.

[R101] If clock synchronization is in effect a SOAM PM implementation **MUST** support the following additional data at the Sink MEP per Measurement Interval per PM Session:

Data	Description
One-way FD counter per configured FD Measurement Bin	A 32-bit counter per Measurement Bin that counts the number of [one-way] FD measurements that fall within the configured bin.
Mean one-way FD	A 32-bit integer reflecting the average (arithmetic mean) one-way FD measurement in microseconds.
Maximum one-way FD	A 32-bit integer reflecting the maximum one-way FD measurement in microseconds.

Table 12 - Mandatory Dual-Ended Delay Data Set with Clock Synchronization

12. PM-3 Requirements

PM-3 uses the Single-Ended Service Loss function for service traffic to measure FLR. The Single-Ended Service Loss function can be configured for multiple classes of service per pair of MEPs. Each unique pair of MEPs and class of service forms a unique PM Session. The tools support point-to-point configurations only. The requirements for the Single-Ended Service Loss function are described below.

[O7] A SOAM PM implementation **MAY** support the Single-Ended Service Loss Function as described in Section 12.1.

Note that Availability cannot be measured with PM-3 because it is not possible to distinguish between times when the customer is not sending traffic versus when the ME is broken by looking at service frame counts.

12.1 Single-Ended Service Loss Function

The following requirements apply to a SOAM PM implementation of the Single-Ended Service Loss function and its client application.

[R102] A SOAM PM implementation **MUST** support the ITU-T Single-Ended ETH-LM protocol and procedures as specified by ITU-T Y.1731 [1], ITU-T G.8021 [4] and ITU-T G.8021 Amendment 1 [5]. Any exceptions to the requirements, behavior, and default characteristics as defined in those specifications are called out in this section.

The following requirements specify the ***input parameters*** that are to be supported by the client application for each PM Session.

[R103] A SOAM PM implementation **MUST** convey the frame counts of only those service frames that have a level of bandwidth profile conformance determined to be *green*.

[R104] A SOAM PM implementation **MUST** support a configurable unicast destination MAC address.

[R105] A SOAM PM implementation **MUST** support a configurable SOAM PM CoS ID for LMM PDU transmission. This requirement is not applicable if the SOAM PM PDUs are untagged.

[D45] The default configured SOAM PM CoS ID **SHOULD** correspond to the CoS Name which yields the lowest frame loss performance for this MEG.

[R106] An implementation of a Controller MEP **MUST** support configuring a PM Session per class of service with the appropriate SOAM PM CoS IDs for the MEG used across the interface it is associated with.

[R107] The SOAM PM CoS IDs that can be configured **MUST** support at least the following configuration for SOAM PM CoS IDs:

- VLAN ID

- A combination of the PCP and VLAN ID

[R108] An implementation of a Responder MEP **MUST** accept the SOAM PM CoS ID received in a LMM PDU and copy the class of service identifier to the associated LMR response it sends.

[R109] If the MEG is tagged and the VLAN DEI is supported, then a SOAM PM implementation on the Controller MEP **MUST** use a VLAN DEI of 0 (discard ineligible) for LMM PDU transmission.

[R110] A SOAM PM implementation **MUST** support a configurable period for LMM PDU transmission.

[R111] For the LMM PDU transmission, periods of {100 ms, 1 sec, 10 sec} **MUST** be supported.

[D46] The default configured period **SHOULD** be {1 sec}.

The following requirements specify the ***output data set*** that are to be supported by the client application per Measurement Interval.

[R112] A SOAM PM implementation **MUST** support the following additional data per Measurement Interval per PM Session:

Data	Description
Start Time-of-day timestamp	A 64-bit timestamp of the time-of-day at the start of the Measurement Interval.
End Time-of-day timestamp	A 64-bit timestamp of the time-of-day at the end of the Measurement Interval.
Measurement interval elapsed time	A 32-bit counter of the number of seconds of the Measurement Interval as calculated by the NE.
SOAM PM PDUs Sent	A 32-bit counter reflecting the number of SOAM PM PDUs sent (i.e., LMM frames transmitted).
SOAM PM PDUs Received	A 32-bit counter reflecting the number of SOAM PM PDUs received (i.e., LMR frames received).
Tx frame count in the forward direction	A 64-bit counter reflecting the number of frames transmitted in the forward direction.
Rx frame count in the forward direction	A 64-bit counter reflecting the number of frames received in the forward direction.
Tx frame count in the backward direction	A 64-bit counter reflecting the number of frames transmitted in the backward direction.
Rx frame count in the backward direction	A 64-bit counter reflecting the number of frames received in the backward direction.

Table 13 - Mandatory Single-Ended Service Loss Data Set

13. References

[1] International Telecommunication Union, Recommendation Y.1731 Revision 1 (02/2011), "OAM functions and mechanisms for Ethernet based Networks".

[2] International Telecommunication Union, Recommendation Y.1563 (01/2009), "Ethernet frame transfer and availability performance".

[3] International Telecommunication Union, Recommendation G.8010 (02/2004), "Architecture of Ethernet layer networks".

[4] International Telecommunication Union, Recommendation G.8021 (10/2010), "Characteristics of Ethernet transport network equipment functional blocks".

[5] International Telecommunication Union, Recommendation G.8021 Amendment 1 (02/2011), "Characteristics of Ethernet transport network equipment functional blocks".

[6] International Telecommunication Union, Recommendation G.7710 (07/2007), "Common Equipment Management Function Requirements".

[7] International Telecommunication Union, Recommendation M.3400 (02/2000), "TMN management functions".

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

[9] MEF Technical Specification MEF 4, "Metro Ethernet Network Architecture Framework - Part 1: Generic Framework", May 2004.

[10] MEF Technical Specification MEF 6.1, "Ethernet Services Definitions - Phase 2", April 2008.

[11] MEF Technical Specification MEF 7.1, "EMS-NMS Information Model Phase 2", October 2009.

[12] MEF Technical Specification MEF 10.2, "Ethernet Services Attributes Phase 2", October 2009.

[13] MEF Technical Specification MEF 10.2.1, "Performance Attributes Amendment to MEF 10.2", January 2011.

[14] MEF Technical Specification MEF 12.1, "Carrier Ethernet Network Architecture Framework, Part 2: Ethernet Services Layer – Base Elements", April 2010.

[15] MEF Technical Specification MEF 15, "Requirements for Management of Metro Ethernet Phase 1 Network Elements", November 2005.

[16] MEF Technical Specification MEF 17, "Service OAM Requirements & Framework - Phase 1", April 2007.

[17] MEF Technical Specification MEF 20, "User Network Interface (UNI) Type 2 Implementation Agreement", July 2008.

[18] MEF Technical Specification MEF 23, "Carrier Ethernet Class of Service - Phase 1", June 2009.

[19] MEF Technical Specification MEF 23.1, “Carrier Ethernet Class of Service – Phase 2”, January 2012.

[20] MEF Technical Specification MEF 26, "External Network Network Interface (ENNI) - Phase 1", January 2010.

[21] MEF Technical Specification MEF 26.0.3, “Service Level Specification amendment to MEF 26”, October 2011.

[22] MEF Technical Specification MEF 30, "Service OAM Fault Management Implementation Agreement", January 2011.

[23] IEEE 802-2001, “IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture”, 2001

[24] IEEE 802.1Q-2011, "Virtual Bridged Local Area Networks", 2011

[25] ISO 8601, "Data elements and interchange formats - Information interchange - Representation of dates and times", Third Edition, December 1, 2004.

[26] RFC 5951, "Network Management Requirements for MPLS-based Transport Networks", September 2010.

14. Appendix A - Performance Management Functions (Informative)

The following sections provide an overview of the PM functions specified by ITU-T Y.1731 [1], ITU-T G.8021 [4] and ITU-T G.8021 Amendment 1 [5].

14.1 Dual-Ended Delay PM Function

The Dual-Ended Delay PM Function is intended to measure one-way synthetic FD, and is specified for use in a point-to-point service but is not precluded from use in a multipoint service.

One message is defined to enable a uni-directional mechanism, or dual-ended process, to exchange timestamps. The One-Way Delay Message (1DM) conveys the transmit timestamp at the Controller MEP at the time of 1DM transmission.

The Sink MEP can estimate one-way synthetic FD by comparing the transmit timestamp in the 1DM and the receive timestamp at the time of 1DM reception. Successive measurements can be used to determine one-way synthetic IFDV. With an adjustment to account for the minimum Delay, one-way FDR can also be estimated.

Frame generation and reception processes are defined for 1DM. In addition, a single 1DM Source Control Process and a single 1DM Sink Control Process are defined. The 1DM Source Control Process coordinates 1DM generation to a given destination at a given SOAM PM CoS ID and periodicity. The 1DM Sink Control Process coordinates 1DM reception from a given source. A FD measurement is generated for each successful 1DM exchange. The following figure illustrates these processes:

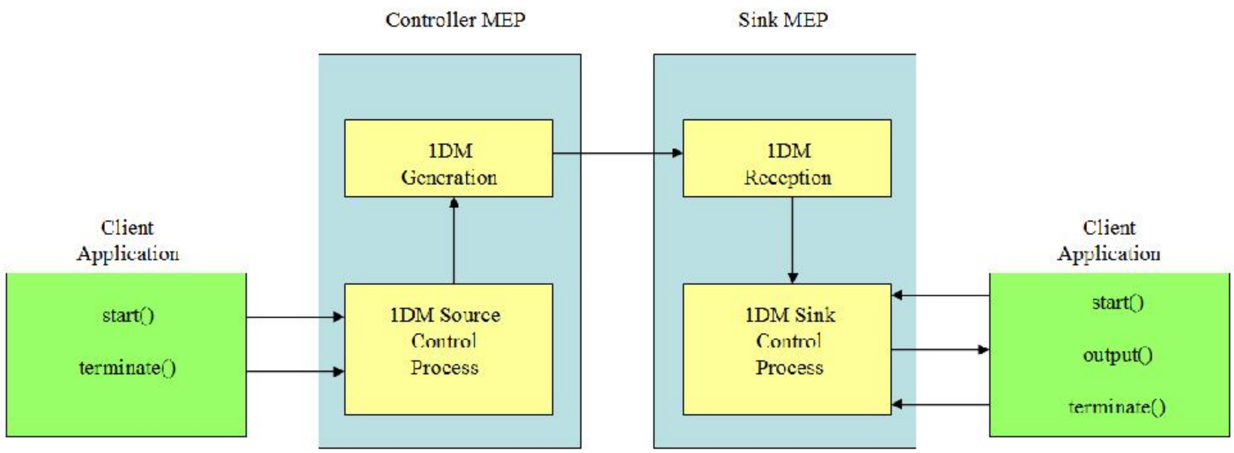

Figure 10 - Dual-Ended Delay Processes

The parameters of the signals generated and received by the 1DM Source Control Process are as follows:

Signal	Parameters
start()	DA (destination unicast MAC address)
	VLAN PCP (0..7, not applicable if untagged)
	Period for 1DM generation (ms)
	TestID
	Length
terminate()	None

Table 14 - 1DM Source Control Process Signals

The parameters of the signals generated and received by the 1DM Sink Control Process are as follows:

Signal	Parameters
start()	SA (source unicast MAC address)
	TestID
output()	One-way FD of last successful 1DM exchange
terminate()	None

Table 15 - 1DM Sink Control Process Signals

Clock synchronization is required in order for the one-way synthetic FD measurement to be accurate.

Since this function is a dual-ended process, administrative access to both measurement points is required.

14.2 Single-Ended Delay PM Function

The Single-Ended Delay PM Function is intended to measure two-way synthetic FD (i.e., Round Trip Time), and is specified for use in a point-to-point service but is not precluded from use in a multipoint service.

Two messages are defined to enable a bi-directional mechanism, or Single-Ended process, to exchange timestamps. The first is a Delay Measure Message (DMM) which conveys the transmit timestamp at the Controller MEP at the time of DMM transmission. The second is a Delay Measure Reply (DMR) which conveys the receive timestamp at the Responder MEP at the time of DMM reception and the transmit timestamp at the Responder MEP at the time of DMR transmission. The transmit timestamp in the DMM is also conveyed in the DMR.

The Controller MEP can estimate two-way synthetic FD using the DMM transmit, DMM receive, and DMR transmit timestamps returned in the DMR, and the receive timestamp at the time of DMR reception. The difference between the DMM receive timestamp and DMR transmit timestamp is processing overhead at the Responder MEP that is removed from the measurement. Successive measurements can be used to determine two-way synthetic IFDV. With an adjustment to account for the minimum Delay, two-way FDR can also be estimated.

The Controller MEP can also estimate one-way synthetic frame delay in each direction, by comparing the DMM transmit and DMM receive timestamps (for forward measurements) and the DMR transmit timestamp and the receive timestamp at the time of DMR reception (for backward measurements). Successive measurements can be used to determine one-way synthetic IFDV. With an adjustment to account for the minimum Delay, the one-way FDR can also be estimated. Clock synchronization is required in order for the one-way synthetic FD measurement to be accurate.

Frame generation and reception processes are defined for DMM and DMR. In addition, a single DM control process is defined to coordinate DMM generation to a given destination at a given SOAM PM CoS ID and periodicity. The following figure illustrates these processes:

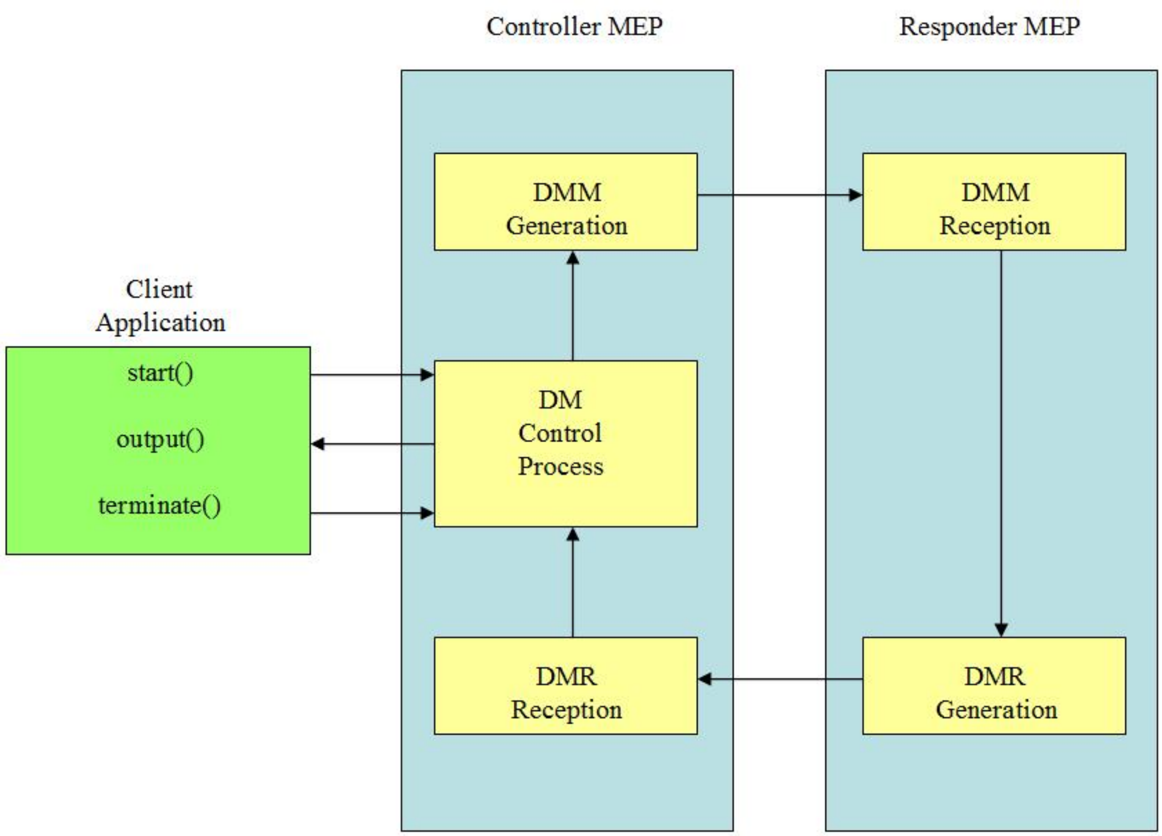

Figure 11 - Single-Ended Delay Processes

The parameters of the signals generated and received by the DM Control Process are as follows:

Signal	Parameters
start()	DA (destination unicast MAC address)
	VLAN PCP (0..7, not applicable if untagged)
	Period of DMM generation (ms)
	TestID
	Length
output()	Two-way FD of last successful DMM/DMR exchange
	One-way FD in each direction
terminate()	None

Table 16 - DM Control Process Signals

Since this function is a single-ended process, administrative access to both measurement points may not be required.

14.3 Single-Ended Service Loss PM Function

The Single-Ended Service Loss PM Function is intended to measure one-way service frame loss, and is specified for use in a point-to-point service only. Ethernet behaviors such as flooding and replication of multicast service frames may limit its application in a multipoint service.

Each MEP is required to maintain a pair of transmit and receive counters per monitored SOAM PM CoS ID. These counters reflect all service frames (i.e., unicast, multicast, and broadcast) which transit the MEP.

Two messages are defined to enable a bi-directional mechanism, or single-ended process, to exchange counters. The first is a Loss Measure Message (LMM) which conveys the service frame transmit count at the Controller MEP at the time of LMM transmission. The second is a Loss Measure Reply (LMR) which conveys the service frame transmit and receive counts at the Responder MEP at the time of LMM reception. The service frame transmit count in the LMM is also conveyed in the LMR.

The Controller MEP can estimate one-way service frame loss in both directions using the service frame transmit and receive counts contained in the LMR and the service frame receive count at the time of LMR reception. These measurements reflect service frame loss since the counters were activated. To determine service frame loss over a given interval of time, it is necessary to take a measurement at the beginning and end of the interval where the difference reflects service frame loss over that period.

Note that the interval of time at the Controller MEP and the Responder MEP are not precisely aligned due to the forwarding delay of the messages. If more precision is desired, an alternative approach is to run an independent measurement process at both points and only use the results of each in the forward direction.

Frame generation and reception processes are defined for LMM and LMR. In addition, a single LM Control Process is defined to coordinate LMM generation to a given destination at a given SOAM PM CoS ID and periodicity. On termination of the LM control process, measures are returned that reflect one-way service frame loss in both directions over the lifetime of the LM control process. The following figure illustrates these processes:

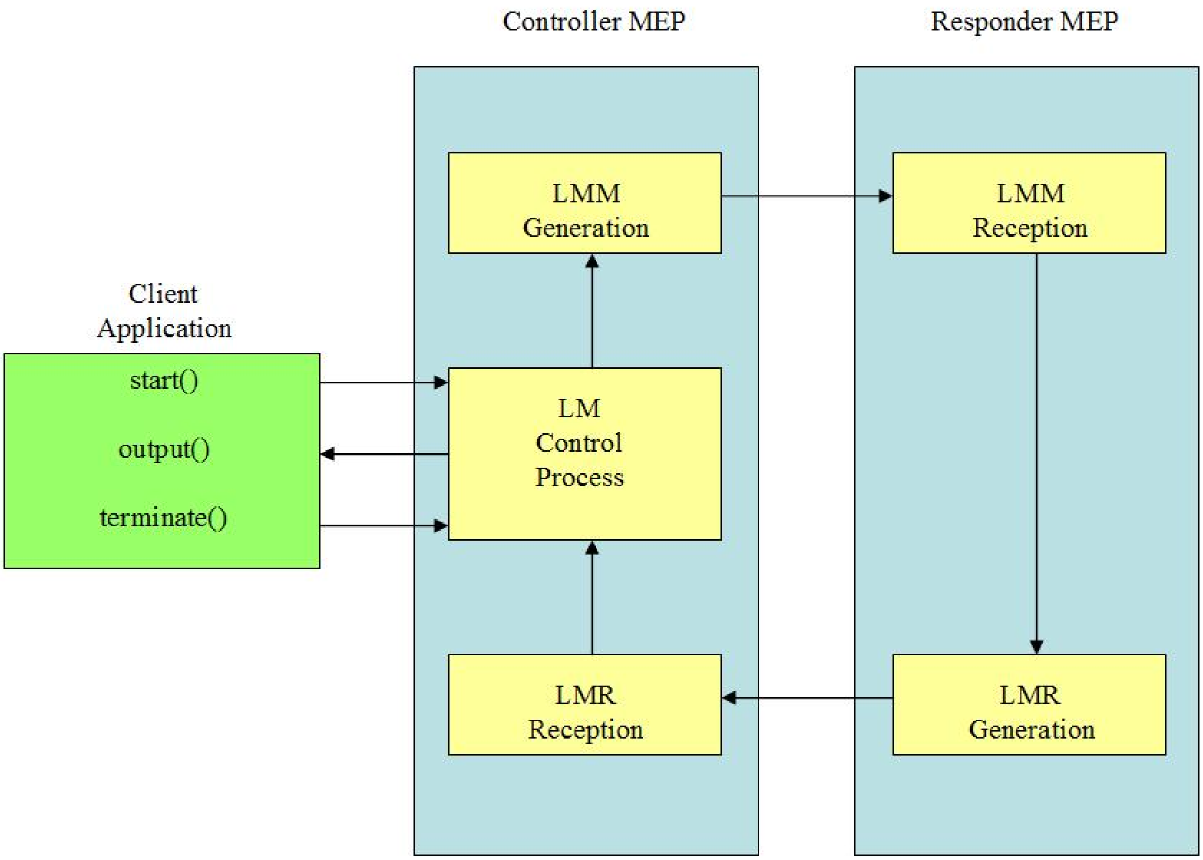

Figure 12 - Single-Ended Loss Processes

The parameters of the signals generated and received by the LM Control Process are as follows:

Signal	Parameters
start()	DA (destination unicast MAC address)
	VLAN PCP (0..7, not applicable if untagged)
	Period of LMM generation (ms)
terminate()	None
result()	Near-end total frames transmitted (NTF)
	Near-end lost frames not received (NLF)
	Far-end total frames transmitted (FTF)
	Far-end lost frames not received (FLF)

Table 17 - LM Control Process Signals

Since this function is a single-ended process, administrative access to both measurement points may not be required.

14.4 Single-Ended Synthetic Loss PM Function

The Single-Ended Synthetic Loss PM Function is intended to measure one-way synthetic frame loss, and is specified for use in a point-to-point service but is not precluded from use in a multipoint service.

Two messages are defined to enable a bi-directional mechanism, or single-ended process, to exchange sequence numbers. The first is a Synthetic Loss Measurement Message (SLM) which conveys a sequence number from the Controller MEP to the Responder MEP. The second is a Synthetic Loss Measurement Reply (SLR) which adds a sequence number from the Responder MEP to the Controller MEP. The original sequence number from the SLM is also conveyed in the SLR.

The Controller MEP can estimate one-way service frame loss in each direction by calculating the loss of the synthetic SLM and SLR frames, using the sequence numbers in a series of received SLR frames. Gaps in one or both sequence numbers indicate frames lost in the forward or backward direction. To determine synthetic frame loss over a given interval of time, it is necessary to send a number of SLM frames over that period, and monitor the received SLRs. The accuracy of the measurement depends on the number of SLM frames sent, as described in Appendix D - Statistical Considerations for Synthetic Loss Measurement (Informative).

Frame generation and reception processes are defined for SLM and SLR. In addition, a single SL Control Process is defined to coordinate SLM generation to a given destination at a given SOAM PM CoS ID and periodicity. On termination of the SL control process, measures are returned that reflect one-way synthetic frame loss in each direction over the lifetime of the SL control process. The following figure illustrates these processes:

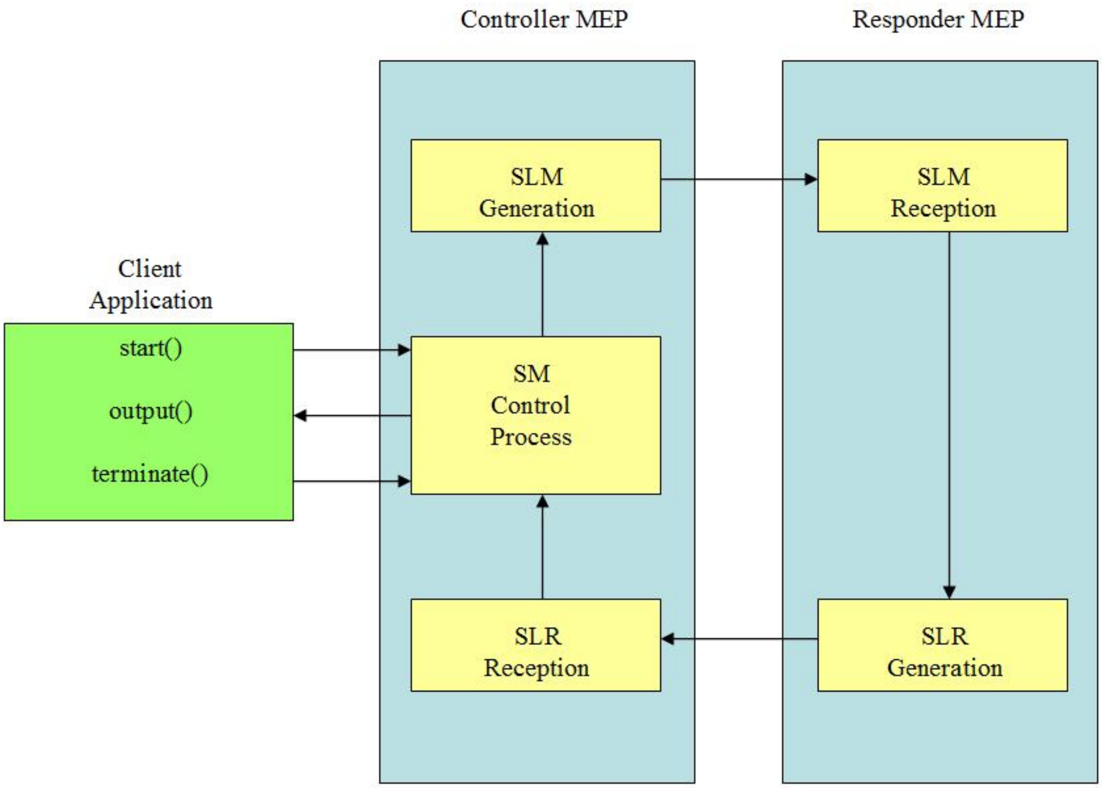

Figure 13 - Single-Ended Synthetic Loss Processes

The parameters of the signals generated and received by the SL Control Process are as follows:

Signal	Parameters
start()	DA (destination unicast MAC address)
	VLAN PCP (0..7, not applicable if untagged)
	Period of SLM generation (ms)
	Test ID
	Length
terminate()	None
result()	Near-end total frames transmitted (NTF)
	Near-end lost frames not received (NLF)
	Far-end total frames transmitted (FTF)
	Far-end lost frames not received (FLF)

Table 18 - SL Control Process Signals

Since this function is a single-ended process, administrative access to both measurement points may not be required.

14.5 Dual-Ended Service Loss PM Function

The Dual-Ended Service Loss PM Function is intended to measure one-way service frame loss, and is specified for use in a point-to-point service only. Ethernet behaviors such as flooding and replication of multicast service frames may limit its application in a multipoint service. The Duel-Ended Service Loss PM Function is not recommended for use as part of any of the PM Solutions described in this document.

Each MEP is required to maintain a pair of transmit and receive counters per monitored SOAM PM CoS ID. These counters reflect all services frames (i.e., unicast, multicast, and broadcast) which transit the MEP.

One message is defined to enable a uni-directional mechanism, or dual-ended process, to exchange counters. The Continuity Check Message (CCM) conveys the service frame transmit count at the Controller MEP at the time of CCM transmission, the service frame transmit count in the last CCM frame received from the Responder MEP, and the service frame receive count at the Controller MEP at the time of CCM reception.

The Responder MEP can estimate one-way service frame loss using the service frame transmit and receive counts contained in the CCM and the service frame receive count at the time of CCM reception. These measurements reflect service frame loss since the counters were activated. To determine service frame loss over a given interval of time, it is necessary to take a measurement at the beginning and end of the interval where the difference reflects service frame loss over that period.

Note that the interval of time at the Controller MEP and the Responder MEP are not precisely aligned due to the forwarding delay of the messages. If more precision is desired, an alternative approach is to run an independent measurement process at both points and only use the results of each in the forward direction.

Frame generation and reception processes are defined for CCM. In addition, a single LM Control Process is defined to calculate the one-way service frame loss over the lifetime of the process, and return it when the process is terminated. The following figure illustrates these processes:

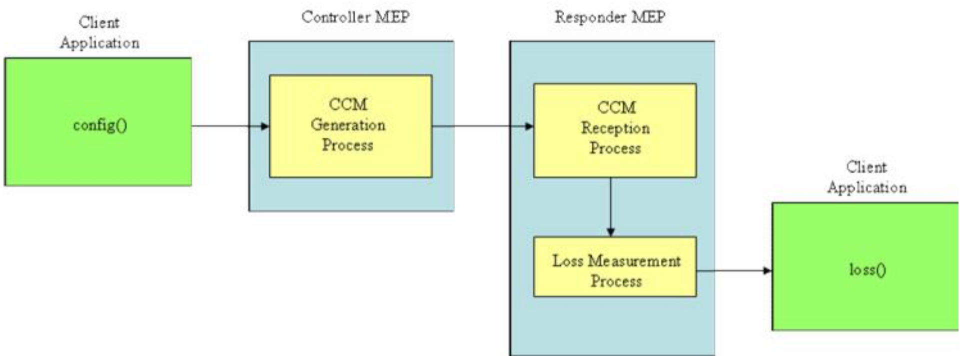

Figure 14 - Dual-Ended Loss Processes

The parameters of the signals received by the CCM Generation Process are as follows:

Signal	Parameters
config()	Continuity Check enable or disable
	Loss Measurement enable or disable
	MEP ID
	MEG ID
	VLAN PCP (0..7, not applicable if untagged)
	Period of CCM generation (ms)

Table 19 - CCM Generation Process Signals

The parameters of the signals generated by the Loss Measurement Process are as follows:

Signal	Parameters
loss()	Near-end total frames transmitted (NTF) for the last second
	Near-end lost frames not received (NLF) for the last second
	Far-end total frames transmitted (FTF) for the last second
	Far-end lost frames not received (FLF) for the last second

Table 20 - Loss Measurement Process Signals

Since this function is a dual-ended process, administrative access to both measurement points is required.

14.6 PM Session Identifiers

In the architecture of a PM Function, there is typically a Control Process that interfaces with a Client Application, and PDU Generation and Reception processes.

In supporting independent PM Sessions, one implementation approach is to extend the interfaces of the Control Process to include an identifier for the session. In this way, an instance of a Control Process can be associated with a specific Session identifier. The session identifier could be a Test ID, SOAM PM CoS ID, or specific VLAN PCP value.

Thus, signals into the Control Process (including received SOAM PM PDUs) would contain a session ID parameter in order to identify the target instance of the Control Process. Similarly, signals out of the Control Process (including transmitted SOAM PM PDUs) would contain a session ID to identify the source instance of the Control Process.

15. Appendix B – Life Cycle Terminology (Informative)

The following diagrams show how the life cycle terminology (see Section 9) for a PM Session is used in this document. While a PM Session is running, the Message Period specifies the time interval between SOAM PDUs, and therefore how often the SOAM PDUs are being sent. The Measurement Interval is the amount of time over which the statistics are collected and stored separately from statistics of other time intervals.

Each PM Session supports a specific PM Function (e.g., Single-ended Delay, Single-ended Synthetic Loss) for a specific CoS Frame Set on a specific MEG on a specific MEG Level.

A PM Session can be Proactive or On-Demand. While there are similarities, there are important differences and different attributes for each. Each is discussed below in turn.

15.1 Proactive PM Sessions

For a Proactive PM Session, there is a time at which the session is created, and the session may be deleted later. Other attributes include the Message Period, Message Interval, and a Start Time (that can be later than the time that the session is created).

The SOAM PM PDUs associated with the PM Session are transmitted every “Message Period”. Data in the form of counters are collected during a Measurement Interval (nominally 15 minutes) and stored in a Current data set. When time progresses past the Measurement Interval, the former Current data set is identified as a History data set. There are multiple History data sets, and the oldest is overwritten.

The EMS/NMS will combine the counters retrieved from NEs to calculate estimates over the SLS period T.

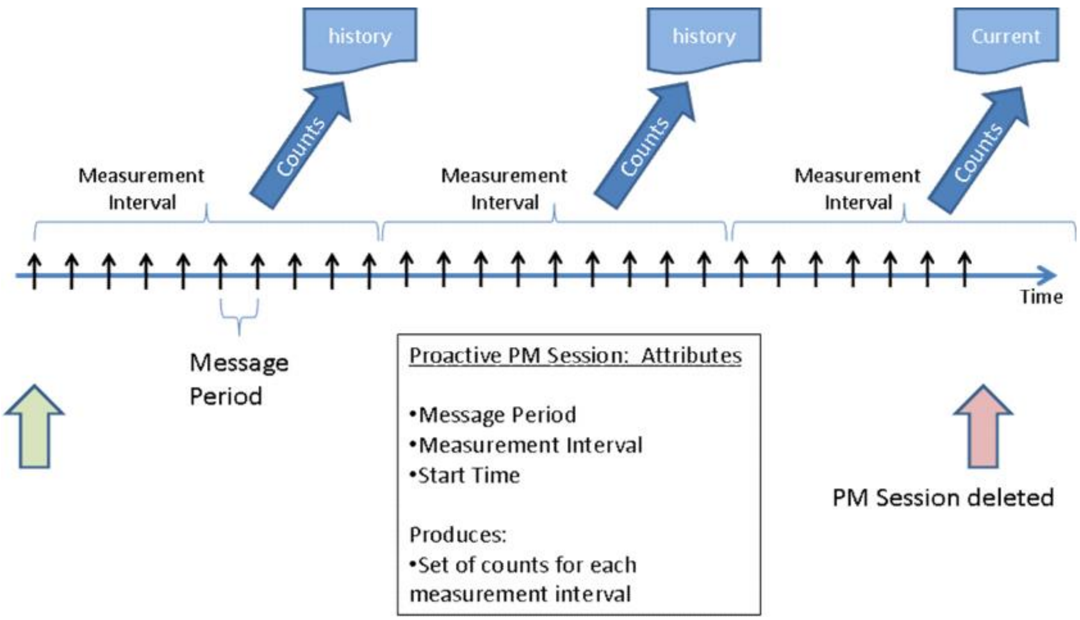

Figure 15 - Measurement Interval Terminology

15.2 On-Demand PM Sessions

For On-Demand PM Sessions, there is a Start Time and a Stop Time. Other attributes can include Message Period, Measurement Interval, and Repetition Time, depending on the type of session that is requested. Different examples are shown in the subsequent diagrams.

Note, in all examples it is assumed that during the interval data is being collected for a report, the counters of the report do not wrap. This is affected by the frequency SOAM frames are sent, the length of time they are sent, and the size of the report counters; the details are not addressed in this specification. At least one report is assumed to be saved after the Measurement Interval is complete.

In the first example, the On-Demand session is run and one set of data is collected. That is, in this example, multiple Measurement Intervals are not used.

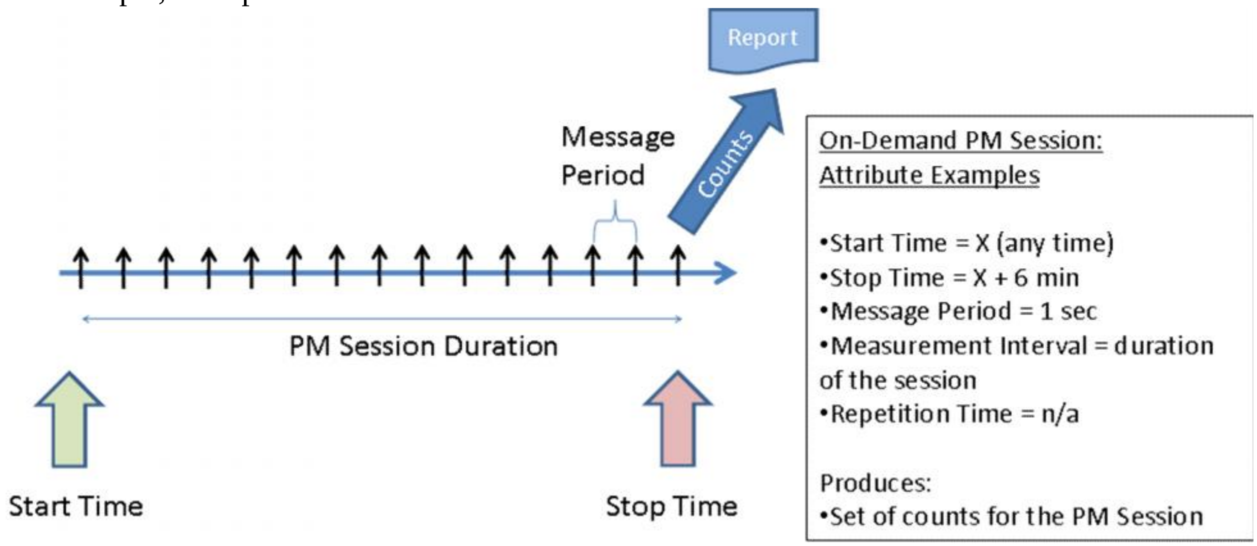

Figure 16 - Illustration of non-Repetitive, On-Demand PM Session

On-Demand PM Sessions can be specified so that Repetitions are specified. This is shown below. Note that a report is created at the end of each Measurement Interval (or Stop Time, if that occurs before the end of the Measurement Interval).

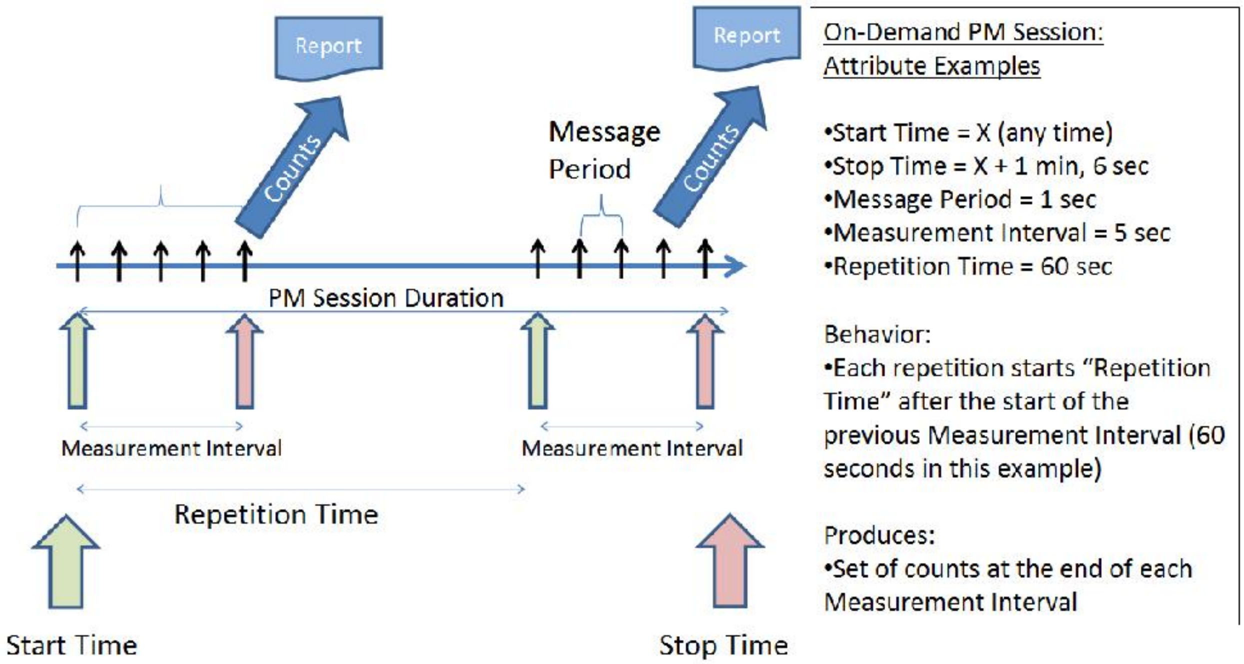

Figure 17 - Example of Repetitive On-Demand PM Session

15.3 PM Sessions With Clock-Aligned Measurement Intervals and Repetition Time of "None"

In all of the previous examples, Measurement Intervals were aligned with the PM Session, so that a PM Session Start Time always occurred at the beginning of a Measurement Interval. Measurement Intervals can instead be aligned to a clock, such as a local time-of-day clock. When Measurement Intervals are aligned to a clock, then in general the PM Session Start Time will not coincide with the beginning of a Measurement Interval.

When the Repetition Time is "None", then the PM Session Start Time will always fall inside a Measurement Interval, so measurements will begin to be taken at the Start Time. As Figure 17 illustrates, when Measurement Intervals are aligned with a clock rather than aligned with the PM Session, then the first Measurement Interval could be truncated. The first, truncated Measurement Interval ends when the clock-aligned Measurement Interval boundary is reached. If the PM Session is Proactive, then a report is generated as usual, except that this report will have the Suspect Flag set to indicate the Measurement Interval's truncated status. Figure 17 depicts a Proactive PM Session, but the same principles apply to On-Demand PM Sessions with Repetition Times of "None".

Subsequent Measurement Intervals in the PM Session will be of full length, with Measurement Interval boundaries occurring at regular fixed-length periods, aligned to the clock. The exception may be the last Measurement Interval of the PM Session. When a PM Session is Stopped or Deleted, then the final Measurement Interval could be truncated, and so again the Suspect Flag would be set for this final, truncated Measurement Interval.

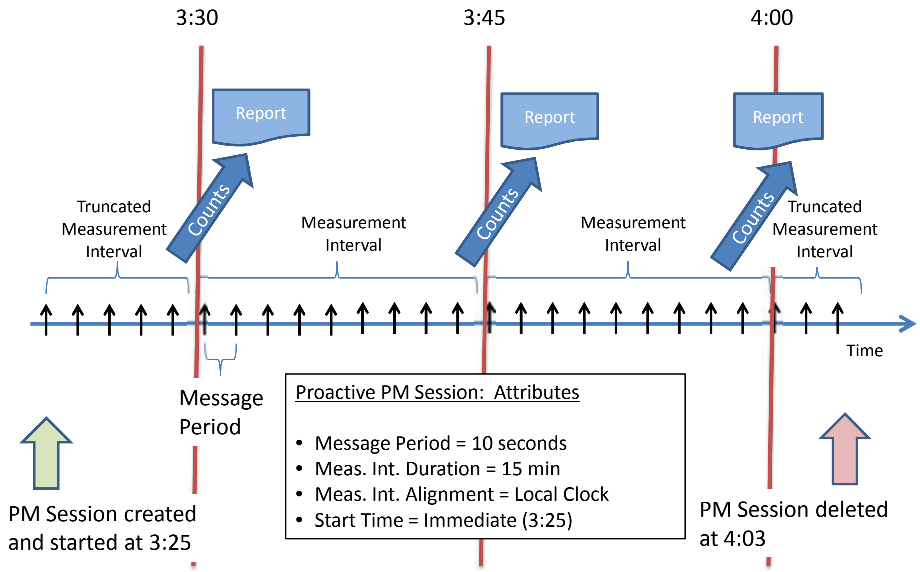

Figure 18 - Example Proactive PM Session with Clock-Aligned Measurement Interval

15.4 PM Sessions With Clock-Aligned Measurement Intervals and Repetition Times Not Equal To "None"

When Measurement Intervals are aligned with a clock and the Repetition Time is not equal to "none", then there are two possibilities for the PM Session Start Time. The first possibility is that the PM Session Start Time is at a time that would fall inside a clock-aligned Measurement Interval. The second possibility when Repetition Times are not equal to "none" is that the PM Session Start Time could fall outside of a clock-aligned Measurement Interval.

If the PM Session Start Time would fall inside a clock-aligned Measurement Interval, then measurements would begin immediately at the PM Session Start Time. In this case, the first Measurement Interval might be truncated (unless PM Session Start Time is also chosen to align with local clock), and thus have its data flagged with a Suspect Flag. An example is illustrated in Figure 18. Figure 18 depicts an On-Demand PM Session, but the same principles apply to a Proactive PM Session whose Repetition Time is not equal to "none".

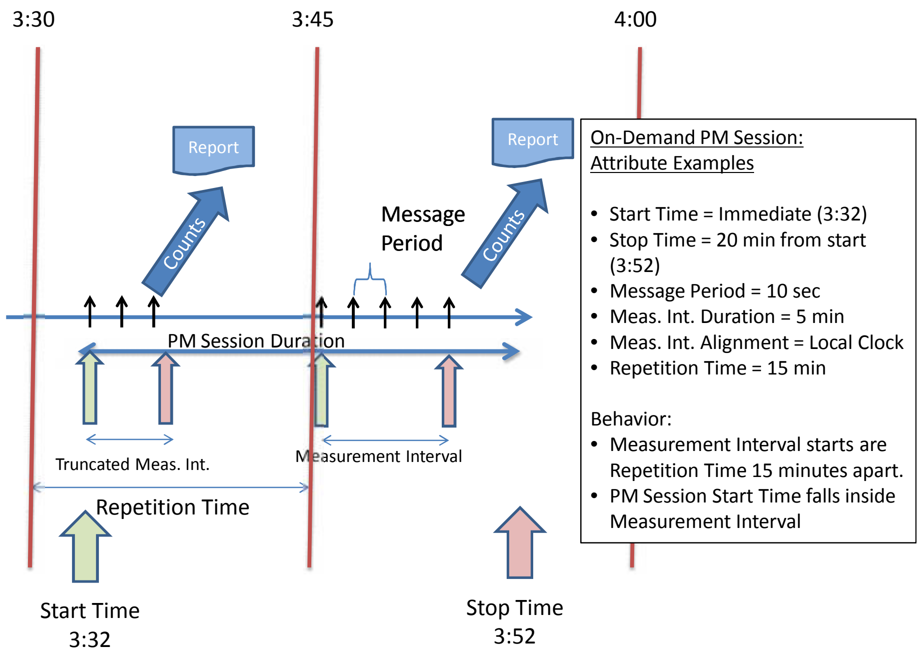

Figure 19 - Example On-Demand PM Session with Clock-Aligned Measurement Interval

In Figure 18, the PM Session starts at 3:32 and has a Stop Time at 3:52. Note that the PM Session might not have been given these explicit times; the PM Session could have had a Start Time of "immediate" and a Stop Time of "20 minutes from start". The Measurement Interval boundary is aligned to the local clock at quadrants of the hour. The next Measurement Interval boundary after the PM Session Start Time is at 3:45. Since the Repetition Time is 15 minutes and the Measurement Interval duration is 5 minutes, the PM Start Time of 3:32 falls inside a Measurement Interval, therefore measurements are begun at the PM Start Time. The first Measurement Interval ends at 3:35 due to its alignment with the local clock. Therefore, the first Measurement Interval is a truncated Measurement Interval (3 minutes long rather than the normal 5 minutes) and its data will be flagged with the Suspect Flag.

The next Measurement Interval begins at 3:45, and runs for its full 5 minute duration, so measurements cease at 3:50. In this example, the PM Session reaches its Stop Time before any more Measurement Intervals can begin. Note that the PM Session Stop Time could fall inside a Measurement Interval, in which case the final Measurement Interval would be truncated; or the PM Session could fall outside a Measurement Interval, in which case the final Measurement Interval would not be truncated. In Figure 18, the data from the second Measurement Interval would not be flagged as suspect.

Figure 18 covered the case where the PM Session Start Time falls inside a clock-aligned Measurement Interval. The second possibility when Repetition Times are not equal to "none" is that the PM Session Start Time could fall outside of a clock-aligned Measurement Interval. In such a case, measurements would not begin immediately at the PM Session Start Time, but rather would be delayed until the next Measurement Interval begins. An example is illustrated in Figure 19. Again, while Figure 19 depicts an On-Demand PM Session, similar principles apply to a Proactive PM Session whose Repetition Time is not equal to "none".

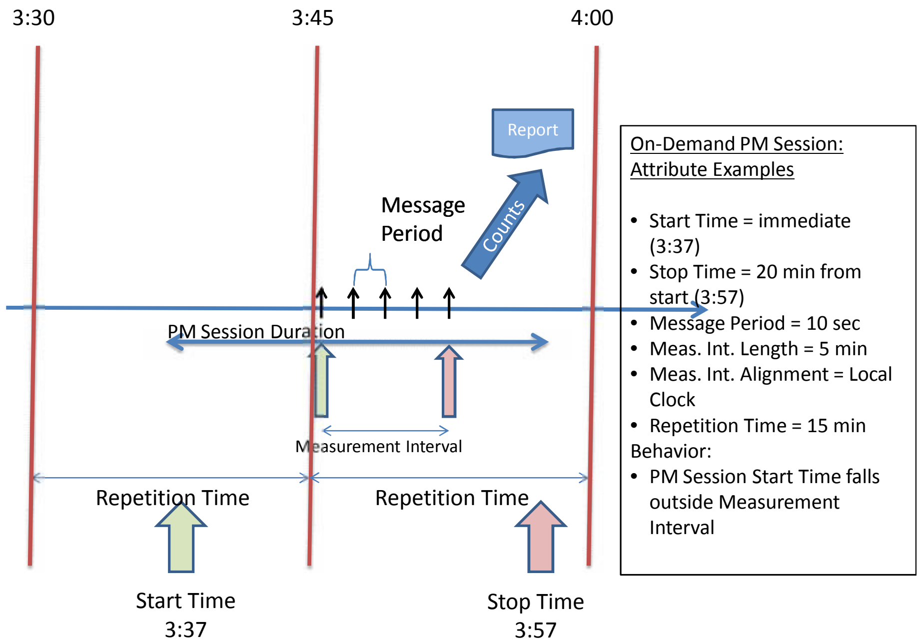

Figure 20 - Second Example of On-Demand PM Session with Clock-Aligned Measurement Interval

In Figure 19, the PM Session starts at 3:37 and has a Stop Time at 3:57. Note that the PM Session might not have been given these explicit times; the PM Session could have had a Start Time of "immediate" and a Stop Time of "20 minutes from start". Note also that in such a case, the parameters given in Figure 19 might be identical to the parameters given in Figure 18, with the only difference being that the "Start button" is pressed 5 minutes later.

The Measurement Interval boundary is aligned to the local clock at quadrants of the hour. The next Measurement Interval boundary after the PM Session Start Time is at 3:45. Since the Repetition Time is 15 minutes and the Measurement Interval duration is 5 minutes, the PM Start Time of 3:37 falls outside a Measurement Interval. Therefore, measurements do not begin at the PM Session Start Time but instead are delayed until the next Measurement Interval boundary.

The first Measurement Interval for this example begins at 3:45, 8 minutes after the PM Session is started. This first Measurement Interval runs for its full 5 minutes, so its data will not have the Suspect Flag set. Measurements cease at 3:50 due to the 5 minute Measurement Interval duration. In this example, the PM Session reaches its Stop Time before any more Measurement Intervals can begin.

Note that, as in the previous case, the PM Session Stop Time could fall either inside or outside a Measurement Interval, and so the final Measurement Interval might or might not be truncated. In general, all Measurement Intervals other than the first and last Measurement Intervals should be full-length.

16. Appendix C – Measurement Bins (Informative)

The MEF 10.2 [12] performance metrics of one-way Frame Delay Performance, one-way Frame Delay Range, and Inter-Frame Delay Variation Performance are all defined in terms of the p-Percentile of frame delay or inter-frame delay variation. Direct computation of percentiles would be resource intensive, requiring significant storage and computation. This informative appendix describes a method for determining whether performance objectives are met using bins for frame delay and inter-frame delay variation.

16.1 Description of Measurement Bins

As described in section 9.2.2, each frame delay bin is one of n counters, B_1, .. B_n, each of which counts the number of frame delay measurements whose measured delay, x, falls into a range. The range for $n+1$ bins (there are n bins, plus Bin 0, so n+1) is determined by n delay thresholds, $D_1, D_2, .. D_n$ such that $0 < D_1 < D_2 < .. < D_n$. Then a frame whose delay is x falls into one of the following delay bin:

Bin 0 if $x < D_1$

Bin i if $D_i \le x < D_{i+1}$

Bin n if $D_n < x$

Note: A Bin 0 (B_0) counter does not need to be implemented, because, B_0 can be determined from R, the total number of frame delay measurement frames received using the following formula:

$$B_0 = R - \sum_{i=1}^{n} B_i$$

Similarly, each inter-frame delay variation (IFDV) bin is one of m counters, $B_1, \ldots, B_m$, each of which counts the number of IFDV measurements whose measured delay, v falls into a range. The range for $m+1$ bins is determined by m IFDV thresholds, $V_1, V_2, .. V_m$ such that $0 < V_1 < V_2 < .. < V_m$. Then a frame whose IFDV v falls into one of the following IFDV bin:

1. Bin 0 if $v < V_1$
2. Bin *i* if $V_i \leq v < V_{i+1}$
3. Bin *m* if $V_m < x$

Note: A Bin 0 (B_0) counter; B_0 can be determined from R_y, the total number of IFDV measurement frame pairs received using the following formula:

$$B_0 = R_y - \sum_{i=1}^{m} B_i$$

16.2 One-way Frame Delay Performance

As defined in MEF 10.2 [12] the one-way Frame Delay Performance is met for an EI pair if $P_p(x) < D$ where $P_p(x)$ is the *p*-th percentile of one way frame delay, *x* and *D* is the one way frame delay performance objective set for that EI pair. To determine if this objective is met, assume that of the *n* delay bins defined for the EI pair bin *j* is defined such that $D_j = D$.

Then we can conclude:

$$P_p(x) < D \text{ if and only if } \sum_{i=j}^{n} B_i < (1-p)R \qquad \text{[Equation 1]}$$

For example, consider an objective for an EI pair that the 95th percentile of one-way delay must be less than 2 milliseconds. If fewer than 5 out of 100 of the received frames have delay greater than 2 milliseconds, then the 95th percentile of delay must be less than 2 milliseconds.

16.3 One-way Inter-Frame Delay Variation Performance

As defined in MEF 10.2 [12] the one-way Inter Frame Delay Variation Performance is met for an EI pair if $P_p(v) < V$ where $P_p(v)$ is the *p*-th percentile of one way IFDV, *v* and *V* is the one way IFDV performance objective set for that EI pair. To determine if this objective is met, assume that of the *m* IFDV bins defined for the EI pair, bin *j* is defined such that $V_j = V$

Then we can conclude:

$$P_p(v) < V \text{ if and only if } \sum_{i=j}^{m} B_i < (1-p)R \qquad \text{[Equation 2]}$$

16.4 One-way Frame Delay Range Performance

As defined in MEF 10.2 [12] the one-way Frame Delay Range Performance is met for an EI pair if $Q_{lh}(x) = P_h(x) - P_l(x) < Q$ where x is the one-way frame delay, *l* and *h* are low and high percentiles such that $0 \leq l < h \leq 1$, $P_l(x)$ is the *l*-th percentile of one way frame delay and the lower bound of the range, $P_h(x)$ is the *h*-th percentile of one way frame delay and the higher bound of the range, and *Q* is the one way frame delay range performance objective for that EI pair. When $l = 0$ then $P_l(x)$ = minimum(x) and when $h = 1$ then $P_h(x)$ = maximum(x).

Note that requirements for measurements of minimum and maximum one-way delay are found in section 10.1. Also note that the minimum delay is lower bounded by *c*, the propagation delay of the shortest path connecting the EI pairs. The constant *c* could be known when the EVC is designed.

There are four cases to consider, depending on the values of *l* and *h*.

16.4.1 Case 1: $Q_{01}(x)$

In the case where $l = 0$ and $h = 1$ then by definition $Q_{01}(x) = \max(x) - \min(x)$ and bins are not required to determine if the range objective is met:

$$Q_{01}(x) < Q \text{ if and only if } \max(x) - \min(x) < Q$$

16.4.2 Case 2: $Q_{0h}(x)$

In the case where $l = 0$ and $h < 1$ then to determine if the objective is met, assume that of the *n* delay bins defined for the EI pair, bin *j* is defined such that $D_j = c+Q$. Then we can transform the range attribute being met into a test that the upper bound on the range $P_h(x)$ is less than a known value, D_j and that the lower bound is above a known value, *c*, then the range will be less than their separation *Q*. Equation 1 gives us a way to determine if the upper bound is less than a known value:

$$P_h(x) < D_j \quad \text{if and only if } \sum_{i=j}^{n} B_i < (1-h)R \qquad \text{[Equation 1]}$$

And so we can conclude:

$$\text{if } \sum_{i=j}^{n} B_i < (1-h)R \text{ and } c \le \min(x) \text{ then } Q_{0h}(x) < Q \qquad \text{[Equation 2]}$$

In other words, the measured range $Q_{01}(x)$ is less than the objective *Q*, and so the range objective is met.

16.4.3 Case 3: $Q_{l1}(x)$

In the case where $l > 0$ and $h = 1$ then an approach is to use bins to determine $P_l(x)$ as follows. Assume that of the *n* bins, bin *g* is defined such that $D_g = L$. Then we can conclude:

$$P_l(x) > L \quad \text{if and only if } \sum_{i=0}^{g-1} B_i < lR \qquad \text{[Equation 3]}$$

By judiciously choosing *L*, then if $\max(x) < Q+L$ and $\sum_{i=0}^{g-1} B_i < lR$ then we can conclude:

$$Q_{l1}(x) = \max(x) - P_l(x) < Q + L - P_l(x) < Q + L - L = Q \qquad \text{[Equation 4]}$$

In other words, the range objective is met.

A reasonable choice for *L* is $L=c$.

16.4.4 Case 4: $Q_{lh}(x)$

In the most general case where $l > 0$ and $h < 1$ then an approach is to combine the techniques in case 2 and case 3 to use bins to determine bounds on $P_l(x)$ and $P_h(x)$ as follows. Assume that of the n bins, bins g and j are defined such that $D_g = L$ and $D_j = L+Q$. Then by Equation 5:

$$P_h(x) < D_j \quad \text{if and only if} \quad \sum_{i=j}^{n} B_i < (1-h)R \qquad \text{[Equation 5]}$$

And by equation 6:

$$P_l(x) > L \quad \text{if and only if} \quad \sum_{i=0}^{g-1} B_i < lR \qquad \text{[Equation 6]}$$

Using the bins, if $\sum_{i=0}^{g-1} B_i < lR$ and $\sum_{i=j}^{n} B_i < (1-h)R$ then we can conclude that:

$$Q_{lh}(x) = P_h(x) - P_l(x) < Q + L - P_l(x) < Q + L - L = Q \qquad \text{[Equation 7]}$$

In other words, the range objective is met.

As with case 3, a reasonable choice for L is $L=c$.

17. Appendix D - Statistical Considerations for Synthetic Loss Measurement (Informative)

This appendix provides considerations on how to configure the Measurement Interval and Measurement Period of the Synthetic Loss Measurement capability.

17.1 Synthetic Traffic and Statistical Methods

One of the first questions of statistical analysis is, “what is the required confidence interval?” This is a central question when one is comparing a null hypothesis against an alternate hypothesis, but for this problem, it is not immediately clear what the null hypothesis is.

The assumption is that if we are promising a loss rate of alpha% to a customer, we have to build the network to a slightly smaller loss rate. (Otherwise, any measurement, no matter how large and accurate the sample size, would yield violations half of the time.) As an example, suppose a carrier promises a network with better than 1% loss, and builds a network to .7% loss. The carrier can then choose a one-tailed confidence interval (say 95%), and then it becomes straightforward to calculate the number of samples that are needed to get the variability of measurements to be as small as needed. This is shown below.

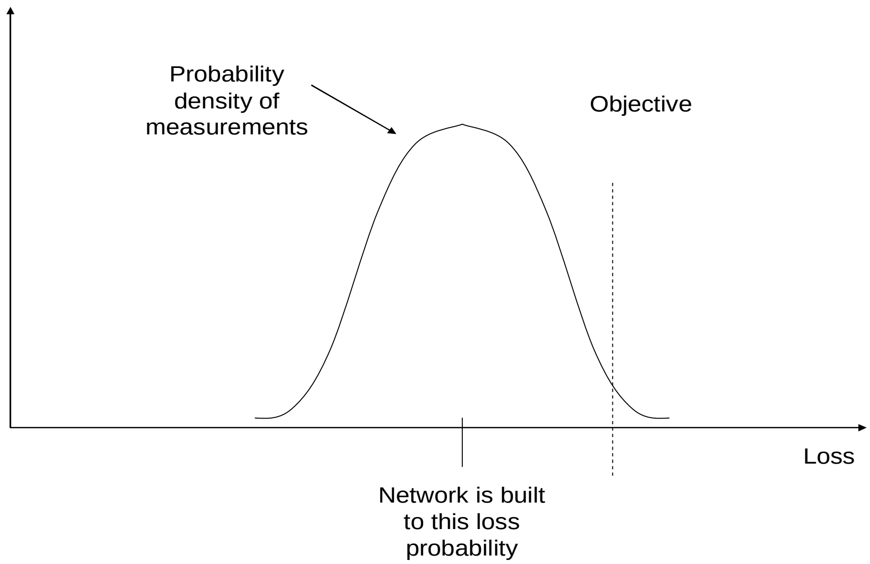

Figure 21 - Hypothesis Test for Synthetic Frame Loss Measurements

Before we specify confidence intervals, or decide how much "better" the network should be built than promised, we can study how the sampling rate and sampling interval relate to the variability of measurements. A useful measure is the Coefficient of Variation (CoV), i.e. the ratio of a probability density's standard deviation to its mean. In the hypothetical diagram above, the value would be roughly 0.2. It should be clear that the smaller the CoV, the more accurate the measurements will be.

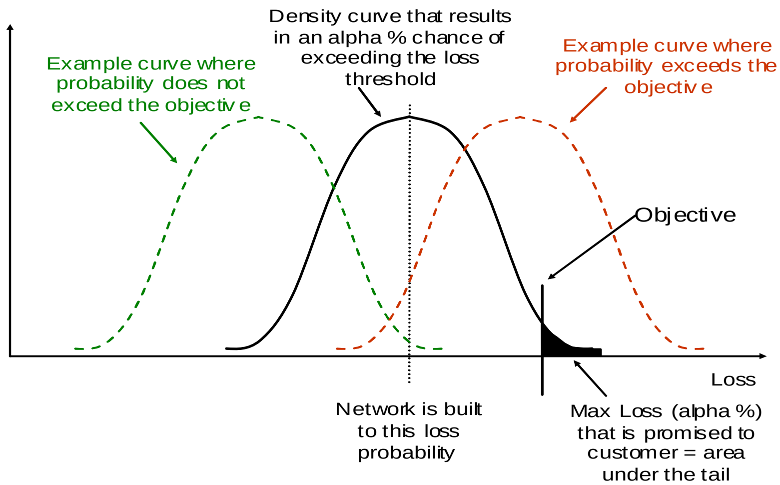

Figure 22 - Density Curve and Probability of Exceeding the Objective

Before getting into the simple equations that are relevant to the analysis, consider what the graphs look like for the synthetic frame approach, with specific examples of different synthetic frame Message Periods, Measurement Intervals, and probabilities of loss (i.e., the true Frame Loss Ratio of the network). These graphs are not hypothetical; they use exact values from the binomial probability density function. The assumption here is that the network is performing at exactly the FLR listed in the title of each graph, and the Y axis shows the probability that a specific percentage of synthetic frames would be lost in practice, i.e., that the measured FLR has the value shown on the X axis. Note that for some combinations of variables, the distribution is quite asymmetric with a long tail to the right, but for many others the distribution is an extremely close approximation to the normal. This, of course, is a well-known property of the binomial density function.

In each example, the number of samples (i.e., the number of synthetic frames) is shown - this is a function of the Message Period and the interval over which the FLR is calculated. For instance, sending one synthetic frame per second for 1 hour yields 3600 samples.

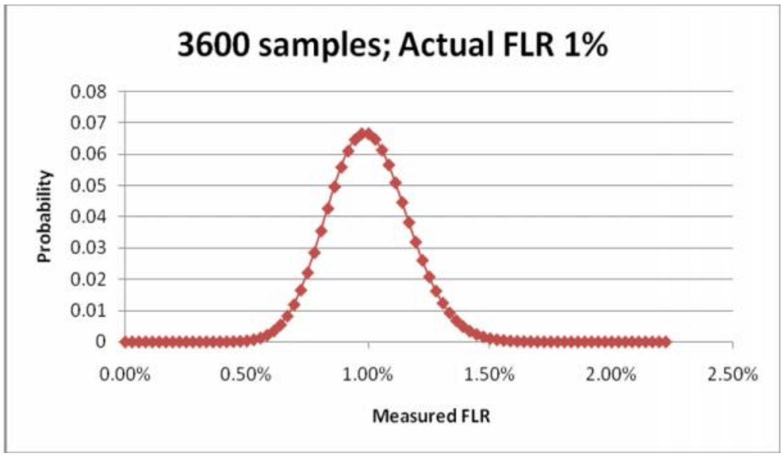

Figure 23 - Synthetic Loss Performance Example 1

The above has a CoV of 0.17. Note how it looks like a normal density.

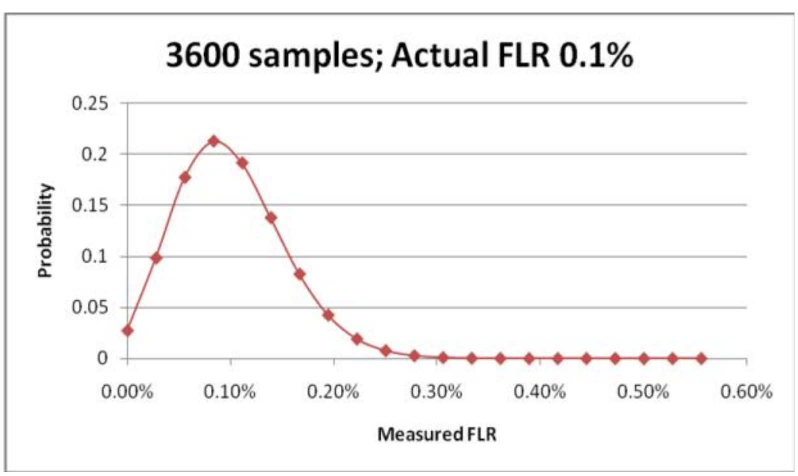

Figure 24 - Synthetic Loss Performance Example 2

In Example 2, the loss rate is smaller, and the CoV is 0.53. This is asymmetric, and variability seems too large for our use.

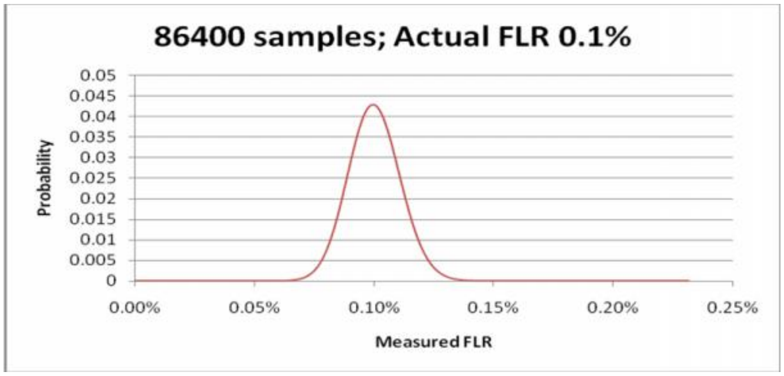

Figure 25 - Synthetic Loss Performance Example 3

Example 3 is the same as Example 2, but with a larger Measurement Interval and hence a higher number of samples. It has a CoV of 0.11 and appears to be precise enough for use.

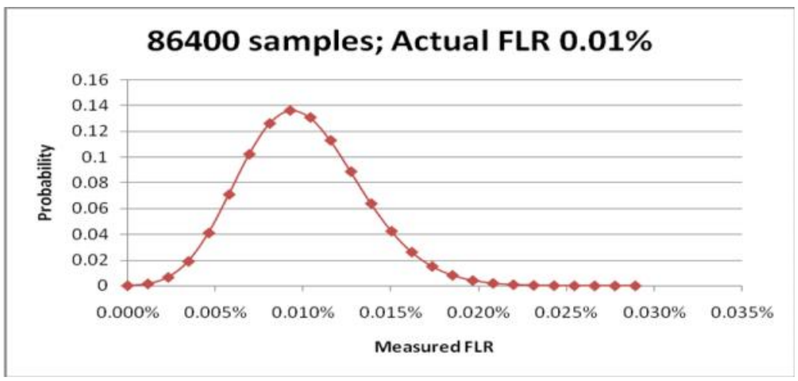

Figure 26 - Synthetic Loss Performance Example 4

In Example 4, the loss rate is even smaller. It has a CoV of 0.34, and may be too variable.

Some similarities in patterns are clear; for example as the probability of frame loss (p) gets smaller, the effects can be mitigated by having a larger number of synthetic loss frames (n). This is predicted by fundamental properties of the density function. The binomial approximates the normal distribution for most of the types of numbers of concern. The exceptions are when the CoV is poor as shown in Examples 2 and 4.

The statistical properties are such that the following equations apply, where p=probability that a frame is lost, $q=1-p$ is the probability that a frame is not lost and n is the sample size:

$$\text{Expected number of frames lost (i.e., mean)} = \mu_n = np$$

$$\text{Standard deviation of number of frames lost} = \sigma_n = \sqrt{npq}$$

These can be easily converted into FLRs:

$$\text{Expected measured FLR (i.e., mean)} = \mu_{FLR} = \frac{100 «_n}{n} = 100p$$

$$\text{Standard deviation of measured FLR} = \sigma_{FLR} = \frac{100 \pm_n}{n} = 100\sqrt{\frac{pq}{n}}$$

Note that the expected value of the measured FLR (μ_{FLR}) is always equal to the probability of loss (p), i.e., the actual FLR of the network.

As introduced above, the coefficient of variation, of the sample statistic is the standard deviation as a fraction of the mean:

$$\frac{\pm}{«} = \frac{\sqrt{n*p*q}}{n*p} = \sqrt{\frac{q}{np}}$$

$$= \sqrt{\frac{q}{p}} * \frac{1}{\sqrt{n}}$$

This is the key result. The smaller CoV is, the better. For a given CoV, we can state the following:

- As n goes up by a factor of 10, the CoV gets smaller (improves) by a factor of $\frac{1}{\sqrt{10}}$, or about 1/3.
- As p goes down by a factor of 10, the CoV gets larger (gets worse) by a factor of $\sqrt{10}$, or about 3.

Furthermore, if p goes down by a certain factor, then n needs to go up by the same factor. That is, if we need to support a loss probability that is 1/100th of what we comfortably support today, we have to either increase the rate of synthetic frames by 100 if we sample over the same interval, increase the interval by a factor of 100, or some combination of the two such as increasing both the rate and the interval by a factor of 10.

Below are example calculations of resolution. Values are highlighted where the resolution is less than 0.2. This value is proposed as a reasonable bound.

Message Period: 1 second

	n	p	μ_{FLR}	σ_{FLR}	CoV
1 hour	3600	0.01	1.000%	0.1658%	0.17
	3600	0.001	0.100%	0.0527%	0.53
	3600	0.0001	0.010%	0.0167%	1.67
	3600	0.00001	0.001%	0.0053%	5.27
24 hour	86400	0.01	1.000%	0.0339%	0.03
	86400	0.001	0.100%	0.0108%	0.11
	86400	0.0001	0.010%	0.0034%	0.34
	86400	0.00001	0.001%	0.0011%	1.08
1 month	2592000	0.01	1.000%	0.0062%	0.01
	2592000	0.001	0.100%	0.0020%	0.02
	2592000	0.0001	0.010%	0.0006%	0.06
	2592000	0.00001	0.001%	0.0002%	0.20

Message Period: 0.1 second

	n	p	$\mu_{\%}$	$\sigma_{\%}$	CoV
1 hour	36000	0.01	1.000%	0.0524%	0.05
	36000	0.001	0.100%	0.0167%	0.17
	36000	0.0001	0.010%	0.0053%	0.53
	36000	0.00001	0.001%	0.0017%	1.67
24 hour	864000	0.01	1.000%	0.0107%	0.01
	864000	0.001	0.100%	0.0034%	0.03
	864000	0.0001	0.010%	0.0011%	0.11
	864000	0.00001	0.001%	0.0003%	0.34
1 month	25920000	0.01	1.000%	0.0020%	0.00
	25920000	0.001	0.100%	0.0006%	0.01
	25920000	0.0001	0.010%	0.0002%	0.02
	25920000	0.00001	0.001%	0.0001%	0.06

Table 21 - CoV Calculations

18. Appendix E – Notes on the Applicability of PM-3 Solutions (Informative)

PM-3 is an optional solution that uses the Loss Measurement function based on LMM/LMR exchanges to measure frame loss and availability within a point-to-point MEG, a MEG with exactly two MEPs. This appendix describes factors which should be considered when deciding which PM solution to apply in a given situation.

18.1 Summary of Loss Measurement (Informative)

The ITU-T ETH-LM function is defined in ITU-T Y.1731 [1] and uses a simple technique for determining loss between a pair of MEPs, which we will denote as the Ingress MEP i and the

Egress MEP j. The ingress MEP maintains a Transmit Counter T_i that counts of all the service frames that pass through it as they enter the network between the MEPs in the MEG. Similarly, the Egress MEP maintains a Receive Counter, R_j that counts all of the service frames that exit the network between the MEPs in the MEG.

At the beginning of a time period we wish to measure loss for, the Ingress MEP inserts a SOAM PM PDU[3] into the flow of service frames. The SOAM PM PDU contains the value of $T_i\ (1)$ in the appropriate field of the SOAM PM PDU. When the SOAM PM PDU is received by the Egress MEP, the current value of the Receive Counter, $R_j\ (1)$ is recorded along with $T_i\ (1)$.

At the end of the time period we wish to measure loss for, the Ingress MEP inserts a second SOAM PM PDU into the flow of service frames. The SOAM PM PDU contains the value of $T_i\ (2)$. When the SOAM PM PDU is received by the Egress MEP, the current value of the Receive Counter, $R_j\ (2)$ is recorded along with $T_i\ (2)$.

The number of service frames transmitted by the Ingress MEP i between the transmission of the two SOAM PM PDUs is $\Delta T_i = T_i(2) – T_i(1)$.

Similarly, the number of service frames received by the Egress MEP j between the receipt of the two SOAM PM PDUs is $\Delta R_j = R_j(2) – R_j(1)$.

The ITU-T ETH-LM function then computes the frames lost between the two SOAM PM PDUs which is defined as $\Delta L_{ij} = \Delta T_i - \Delta R_j$ and Frame Loss Ratio is $\Delta FLR_{ij} = \Delta L_{ij} / \Delta T_i$.

18.2 PM-3 in Multipoint MEGs

PM-3 is not to be used in a MEG with more than 2 MEPs. An example will demonstrate why. Consider a simple three MEP MEG with MEPs 1, 2, and 3. To measure frame loss over a short interval, a pair of LMM OAM frames are sent from MEP 1 to MEP 2. Assume that over the interval of interest, 30 services frames enter the MEG at MEP 1, 20 of the service frames are delivered to MEP 2, and the other 10 service frames are delivered to MEP 3. No service frames are lost in this example. In this example, ΔT_1 = 30 frames, ΔR_2 = 20 frames, ΔL_{12} = 10 frames, and FLR_{12} = .33, which is wrong, it should be 0.

A detailed analysis of why this example fails to give the right answer is beyond the scope of this standard. The quick summary is that to compute loss requires solving series of equations with

[3] The SOAM PM PDU can be a LMM, LMR, or CCM PDU depending on the specific technique being used.

$2N^2$ unknowns and there are only enough counts maintained to solve those equations when $N \leq 2$.

18.3 PM-3 Considerations in Point-to-Point MEGs

- PM-3 will work in only a point-to-point MEG, one with 2 MEPs, if the network between the two MEPs satisfies certain conditions. Those conditions are:
- The network between the MEPs cannot duplicate frames.
- The network between the MEPs cannot deliver frames out of order.
- No frames can be counted as service frames if they enter the MEG through a MEP and are consumed by an Ethernet MAC within the network.
- There cannot be an Ethernet MAC within the network that generates and sends frames that exit the MEG through a MEP.

18.3.1 Duplicate Frames

If a frame counted as a transmitted frame by an Ingress MEP is duplicated within the network then the Egress MEP will receive and count each copy. When Loss is computed, the extra copies will be incorrectly counted as negative loss.

18.3.2 Out of Order Frames

If frames can be delivered out of order then these can affect the loss calculations described above in two ways.

If a frame was received by the Ingress MEP between the two OAM frames, it is possible that it gets delivered before the first OAM frame or it may be delivered after the second OAM frame. In either case, the frame will be counted as a transmitted frame by the Ingress MEP, but not counted as a received frame by the Egress MEP, and incorrectly counted as a lost frame.

Conversely, a service frame that entered the MEG before the first OAM frame of the pair, or after the second OAM frame of the pair could be delivered to the Egress MEP between the two OAM frames. In that case, the service frame would not count as a transmitted frame by the Ingress MEP, but would be counted as a received frame service by the Egress MEP, and the loss formula would incorrectly count this service frame as negative loss.

18.3.3 Frames Consumed by an Internal MAC

If a unicast frame enters the MEG and is counted by the Ingress MEP as a transmitted service frame and that frame is addressed to a MAC within the MEG, that frame should not exit the

MEG and be counted by the Egress MEP as a received frame. It will incorrectly count as a lost frame.

Similarly, if a multicast frame enters the MEG and is counted by the Ingress MEP as a transmitted service frame and that frame is received by a MAC bridge within the network but not forwarded then it will not exit the MEG and will not be counted by the Egress MEP as a received frame. It will incorrectly count as a lost frame.

18.3.4 Frames Transmitted by an Internal MAC

If a MAC within the network that connects the MEPs in the MEG generates and transmits a frame and that frame exits the MEG and that frame is counted as a received service frame by the Egress MEP then that frame will be incorrectly counted as negative loss.

19. Appendix F - Frame Loss Count Accuracy

This appendix provides an overview of the placement of the Down MEPs, VID aware, with respect to the Queuing entities as outlined in IEEE 802.1Q Bridge Port and potential loss of counted In-profile frames.

19.1 Review of the placement of the Down MEPs (VID Aware) to Queuing entities

SOAM-PM can be performed on In-profile frames per CoS ID (e.g., EVC) at the Subscriber, EVC, and Operator MEGs. The MPs (Down or Up MEPs) distinguished by VIDs, as shown in Figure 26, are above the queuing entities (detailed view can be found Figure 22-4 of IEEE 802.1Q-2011 [24]). Hence, in the egress direction, these MPs cannot distinguish between discards (on a per VID basis) in the queuing entities and discards in the MEN cloud. Discards in the egress queue of UNI-Ns can be minimized by setting a higher drop threshold for discard ineligible (green or Qualified Frames) in the queue compared to discard eligible (yellow) frames. Subscriber's UNI-Cs will also need proper configuration (e.g., sufficient queue size) to allow for shaping traffic to contracted Bandwidth Profiles and minimizing discards.

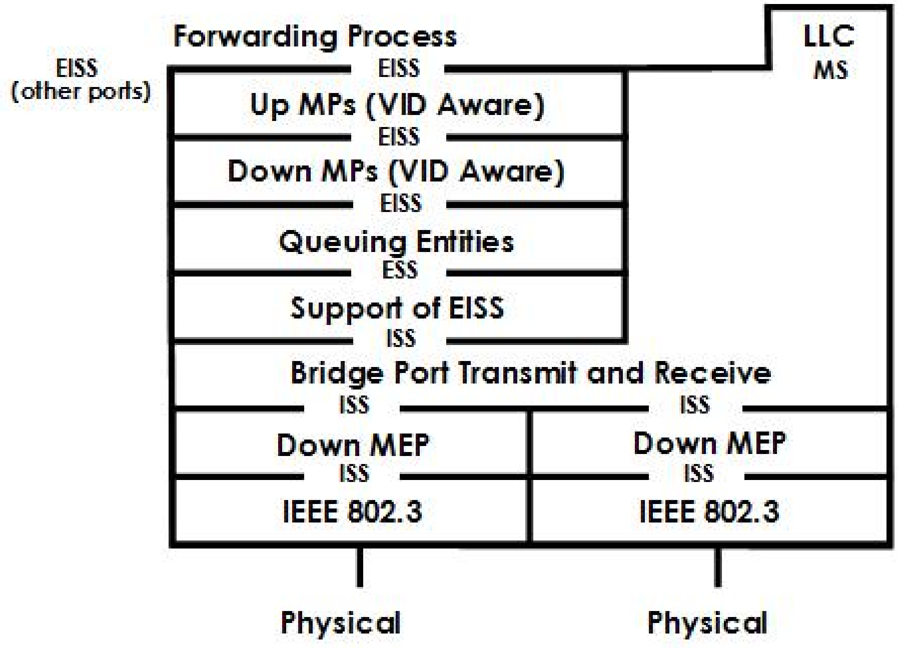

Figure 27 - 802.1Q Bridge Port

20. Appendix G: Normalizing Measurements for FDR (Informative)

This document has specified a binning approach for delay-related measurements. When making measurements of delay variation, normalization is needed.

For the IFDV metric, a pair of delay values are normalized by subtracting one from the other, and taking the absolute value. Thus, the minimum of any IFDV measurement is 0, and as a consequence bins can be set up without any consideration for the actual magnitude of the delay.

A similar normalization is needed for FDR. FDR is defined as the difference between the Y^{th} percentile of delay and the minimum delay, so each delay observation needs to have the estimated minimum subtracted from it, to get a normalized delay. The FDR performance objective *O* is specified relative to a minimum of zero, as shown below in Figure 28.

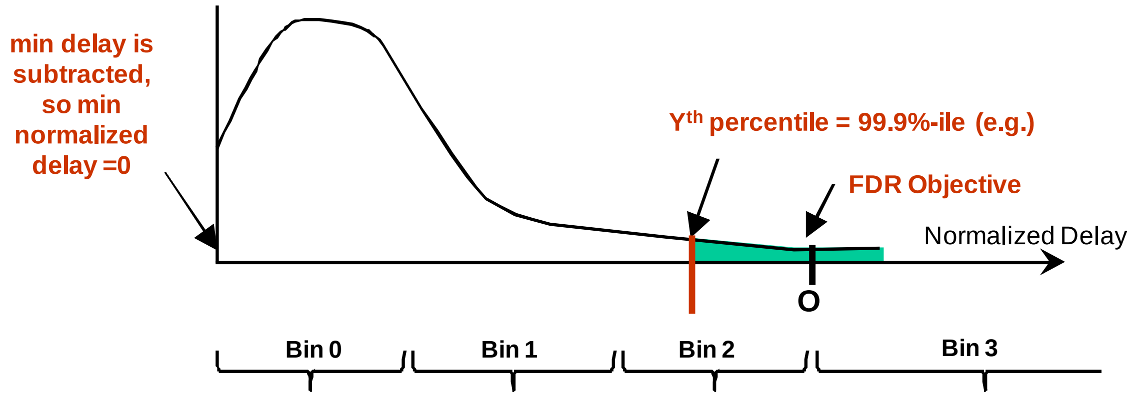

Figure 28 - Example FDR Distribution (Normalized), and Bins

The distribution of delay is generally observed to be skewed to the right; i.e., there would be many measurements at or near the minimum delay, and fewer at higher values. Therefore, a good estimate of the minimum can be determined in a time interval much shorter than a Measurement Interval. Once an estimate of the minimum is available, observed delays can be normalized by subtracting the minimum, and then the appropriate bin counters can be incremented as the normalized delay is processed from each received SOAM frame.

One suggested practical approach as shown in Figure 28 is to record the minimum delay of each Measurement Interval, and to use that value as the estimated minimum at the beginning of the following Measurement Interval. As each delay measurement is received, the estimated minimum can be set to the minimum of the current measured delay and the previous estimate. Then each received delay measurement is normalized by subtracting the estimated minimum. With this approach, there would never be a negative value for a normalized FDR measurement.

Very small shifts in the minimum could be observed that would not be significant. Define ε as the threshold below which a shift is not considered significant (e.g., 10% of the objective). Then the NMS/EMS would not take actions if the shift of the minimum was less than ε. If, on the other hand, the minimum at the end of a Measurement Interval has decreased / increased by a value more than ε, the NMS/EMS is expected to consider as invalid the FDR measurements in the associated Measurement Interval(s).

If there are network changes during the Measurement Interval, then FDR measurements during that MI may be invalid, and the measurements can be ignored by the NMS/EMS. This is discussed next. However, other MIs would still be valid and contribute to the estimate of FDR during the interval *T*.

Note that this approach is presented as an example, and that alternate implementations may improve on it.

20.1 Topology Shifts

For a fixed topology, the minimum delay is essentially fixed. However, network changes (e.g., in response to a network failure) can result in a shift in the minimum delay that can be significant. The minimum delay can of course shift to a lower or to a higher value.

20.1.1 Minimum Delay Becomes Significantly Smaller

When the delay becomes significantly smaller, as is shown in MI 2 below in Figure 29, it will be obvious at the end of MI 2 that the minimum delay is significantly lower than the minimum delay at the end of MI 1. It would be straightforward for an NMS/EMS to simply consider the FDR measurements of that interval as being invalid, and to ignore them.

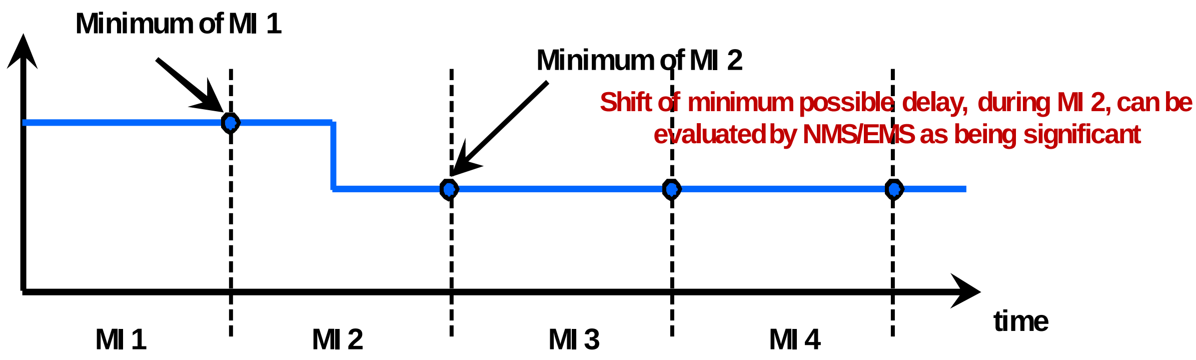

Figure 29 - Reduction in Minimum Delay, due to Network Topology Change

20.1.2 Minimum Delay Becomes Significantly Larger

When the delay becomes significantly larger, as is shown in MI 6 below in Figure 30, it will not be obvious until the end of MI 7 that the minimum delay is significantly higher than the minimum delay observed at the end of MI 5. It would be straightforward for the EMS/NMS to detect that and mark the measurements of MI 6 and MI 7 as being invalid.

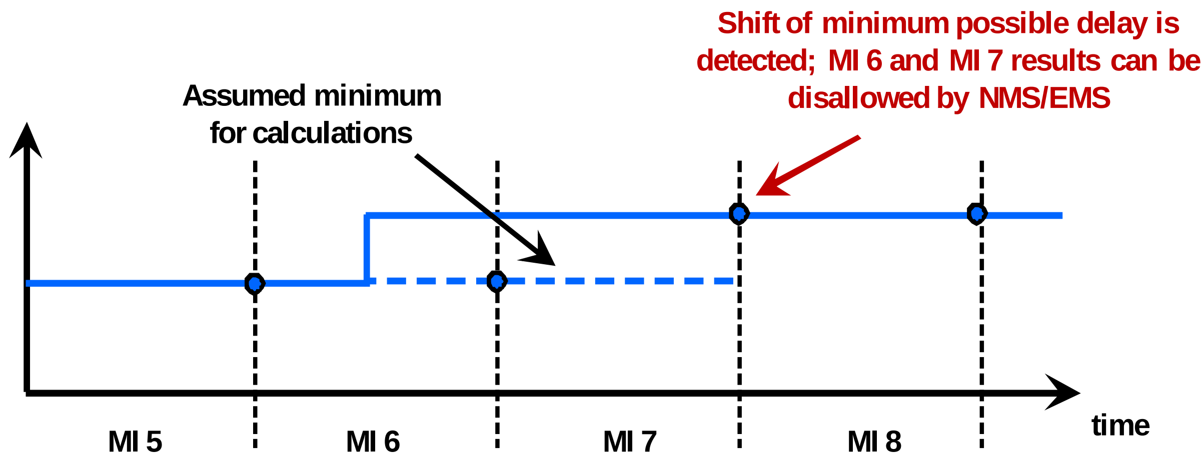

Figure 30 - Increase in Minimum Delay, due to Network Topology Change