



MEF 35.1

Service OAM Performance Monitoring Implementation Agreement

May 2015

Disclaimer

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

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

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

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

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

© The Metro Ethernet Forum 2012-2015. All Rights Reserved.

Table of Contents

1.	List of Contributing Members	1
2.	Abstract.....	2
3.	Terminology.....	3
4.	Scope	10
4.1	Change History	10
5.	Compliance Levels.....	12
6.	Introduction.....	13
6.1	OAM Domains	13
6.2	Maintenance Entities	13
6.3	OAM Domains and Performance Metrics	15
6.4	Default Behavior	16
7.	PM Source Documents	17
7.1	MEF 7.2.....	17
7.2	MEF 10.3.....	17
7.3	MEF 15.....	18
7.4	MEF 17.....	18
7.5	MEF 20.....	18
7.6	MEF 23.1.....	18
7.7	MEF 26.1.....	19
7.8	MEF 30.1.....	19
8.	PM Considerations	20
8.1	Frame Delay Measurements	20
8.2	Frame Loss Measurements	20
8.3	Packet-Count-Based versus Time-Based Measurements	23
8.4	CoS Considerations	25
9.	PM Solutions.....	28
9.1	PM-1: Single-Ended Point-to-Point or Multipoint Delay and Synthetic Loss Measurements.....	31
9.2	PM-2: Dual-Ended Point-to-Point or Multipoint Delay.....	32
9.3	PM-3: Single-Ended Service Loss Measurements.....	33
9.4	PM-4: Dual-Ended Synthetic Loss Measurements	34
10.	Common Requirements.....	35
10.1	Life Cycle	35
10.2	Storage.....	41
10.3	OAM Domains	50
10.4	MEP Placement.....	51
10.5	Threshold Crossing Alerts	51

11. PM-1 Requirements	63
11.1 Single-Ended Delay Function for Delay, Frame Delay Range, and Inter-Frame Delay Variation	63
11.2 Single-Ended Synthetic Loss Function for Frame Loss Ratio (FLR) and Availability	68
12. PM-2 Requirements	76
12.1 Dual-Ended Delay Function for Frame Delay, Frame Delay Range, and Inter-Frame Delay Variation	76
13. PM-3 Requirements	81
13.1 Single-Ended Service Loss Function for FLR	81
14. PM-4 Requirements	84
14.1 Dual-Ended Synthetic Loss Function for FLR, Availability, HLI, CHLI	84
15. References	90
16. Appendix A - Performance Management Functions (Informative)	92
16.1 Dual-Ended Delay PM Function	92
16.2 Single-Ended Delay PM Function	93
16.3 Single-Ended Service Loss PM Function	95
16.4 Single-Ended Synthetic Loss PM Function	97
16.5 Dual-Ended Service Loss PM Function	99
16.6 Dual-Ended Synthetic Loss PM Function	100
16.7 PM Session Identifiers	102
17. Appendix B – Life Cycle Terminology (Informative)	103
17.1 Proactive PM Sessions	103
17.2 On-Demand PM Sessions	104
17.3 PM Sessions With Clock-Aligned Measurement Intervals and Repetition Time of “None”	106
17.4 PM Sessions With Clock-Aligned Measurement Intervals and Repetition Times Not Equal To “None”	107
18. Appendix C – Measurement Bins (Informative)	111
18.1 Description of Measurement Bins	111
18.2 One-way Frame Delay Performance	112
18.3 One-way Inter-Frame Delay Variation Performance	112
18.4 One-way Frame Delay Range Performance	112
19. Appendix D - Statistical Considerations for Synthetic Loss Measurement (Informative)	114
19.1 Synthetic Frames and Statistical Methods	114
20. Appendix E – Notes on the Applicability of PM-3 Solutions (Informative)	120
20.1 Summary of Loss Measurement	120
20.2 PM-3 in Multipoint MEGs	120
20.3 PM-3 Considerations in Point-to-Point MEGs	121

21. Appendix F - Frame Loss Count Accuracy (Informative)	123
21.1 Review of the placement of the Down MEPs (VID Aware) to Queuing entities	123
22. Appendix G: Normalizing Measurements for FDR (Informative)	124
22.1 Topology Shifts.....	125
22.2 Impact of Lack of ToD Synchronization.....	126
23. Appendix H – Notes on Dual-Ended PM Functions (Informative)	128
24. Appendix I – Calculation of SLS Performance Metrics (Informative)	131
24.1 One-way Frame Delay.....	131
24.2 One-way Mean Frame Delay.....	132
24.3 One-way Frame Loss Ratio	133
24.4 One-way Availability	133
24.5 One-way High Loss Interval.....	135
24.6 One-way Consecutive High Loss Interval.....	135
24.7 One-way Group Availability	135
25. Appendix J – Statistical Considerations for Availability (Informative)	138
25.1 Availability flr Accuracy	138
25.2 Availability flr Precision	139
25.3 State Change Probability	139
25.4 State Change Probability over a Period of High Loss.....	140
25.5 Accuracy and Precision of Availability.....	141
25.6 Example Parameter Values.....	142
25.7 Relationship between Availability and CHLI.....	152
25.8 Conclusions.....	152

List of Figures

Figure 1 – Maintenance Entities (see MEF 30.1 [19]).....	14
Figure 2 – OAM Domain	15
Figure 3 – MEP Placement.....	23
Figure 4 – Ideal Packet-Count-Based Measurement Samples.....	24
Figure 5 – Possible Actual Packet-Count-Based Measurement Samples	24
Figure 6 – Time-Based Measurement Samples	25
Figure 7 – PM Solution Components.....	28
Figure 8 – Performance Metrics that can be collected with Single-Ended Loss and Delay	30
Figure 9 – Performance Metrics that can be collected with Dual-Ended Loss and Delay.....	31
Figure 10 – Example of Measurement Intervals and Bins	41
Figure 11 – Example of FLR Measurements.....	42
Figure 12 – Example of adjusting Measurement Counters during transition from Available to Unavailable	49
Figure 13 – Example of adjusting Measurement Counters during transition from Unavailable to Available.....	50
Figure 14 – Stateless TCA Reporting Example.....	53
Figure 15 – Stateful TCA Reporting Example	54
Figure 16 – Upper Bin Count for Threshold Crossing.....	55
Figure 17 – Dual-Ended Delay Processes	92
Figure 18 – Single-Ended Delay Processes.....	94
Figure 19 – Single-Ended Loss Processes.....	96
Figure 20 – Single-Ended Synthetic Loss Processes	98
Figure 21 – Dual-Ended Loss Processes	100
Figure 22 – Dual-Ended Synthetic Loss Processes.....	101
Figure 23 – Measurement Interval Terminology.....	104
Figure 24 – Illustration of non-Repetitive, On-Demand PM Session.....	105
Figure 25 – Example of Repetitive On-Demand PM Session.....	105
Figure 26 – Example Proactive PM Session with Clock-Aligned Measurement Interval.....	107
Figure 27 – Example On-Demand PM Session with Clock-Aligned Measurement Interval	108
Figure 28 – Second Example of On-Demand PM Session with Clock-Aligned Measurement Interval.....	109
Figure 29 – Hypothesis Test for Synthetic Frame Loss Measurements.....	114
Figure 30 – Density Curve and Probability of Exceeding the Objective	115
Figure 31 – Synthetic Loss Performance Example 1	116
Figure 32 – Synthetic Loss Performance Example 2.....	116
Figure 33 – Synthetic Loss Performance Example 3.....	117
Figure 34 – Synthetic Loss Performance Example 4.....	117
Figure 35 – 802.1Q Bridge Port.....	123
Figure 36 – Example FDR Distribution (normalized), and Bins.....	124
Figure 37 – Reduction in Minimum Delay, due to Network Topology Change	125
Figure 38 – Increase in Minimum Delay, due to Network Topology Change	126
Figure 39 – Lack of ToD Synchronization.....	126
Figure 40 – Example of a Multipoint MEG and MEPs	128
Figure 41 – Single-Ended Synthetic Loss Example	129

Figure 42 – Dual-Ended Synthetic Loss Example 130
Figure 43 – Example of Availability State Transitions during *T*..... 133
Figure 44 – Example showing a CHLI crossing a Measurement Interval boundary 135
Figure 45 – Group Availability..... 136
Figure 46 – Group Availability..... 137

List of Tables

Table 1 – Terminology and Definitions	9
Table 2 – PM Solutions Summary	29
Table 3 – PM Functions Summary	30
Table 4 – Time Parameters	41
Table 5 – Example Measurement Bin Configuration	45
Table 6 – SOAM Performance Metrics TCA	56
Table 7 – TCA Notification Message Fields	61
Table 8 – Comparison of TCA fields in MEF 35 and ITU-T X.73x	62
Table 9 – Mandatory Single-Ended Delay Data Set	67
Table 10 – Mandatory Single-Ended Delay Data Set with Clock Synchronization	68
Table 11 – Availability State Transition Event Data	71
Table 12 – Mandatory Single-Ended Synthetic Loss Data Set	74
Table 13 – Optional Single-Ended Synthetic Loss Data Set	75
Table 14 – Mandatory Dual-Ended Delay Data Set	79
Table 15 – Mandatory Dual-Ended Delay Data Set with Clock Synchronization	80
Table 16 – Mandatory Single-Ended Service Loss Data Set	83
Table 17 – Availability State Transition Event Data	87
Table 18 – Mandatory Dual-Ended Synthetic Loss Data Set	89
Table 19 – Optional Dual-Ended Synthetic Loss Data Set	89
Table 20 – 1DM Source Control Process Signals	93
Table 21 – 1DM Sink Control Process Signals	93
Table 22 – DM Control Process Signals	95
Table 23 – LM Control Process Signals	97
Table 24 – SL Control Process Signals	98
Table 25 – CCM Generation Process Signals	100
Table 26 – Loss Measurement Process Signals	100
Table 27 – 1SL Source Control Process Signals	101
Table 28 – 1SL Sink Control Process Signals	102
Table 29 – CoV Calculations with Message Period 1s	119
Table 30 – CoV Calculations with Message Period 100ms	119
Table 31 – Examples with $P = 100\text{ms}$, $n = 5$, $\Delta t = 1\text{s}$	143
Table 32 – Examples with $P = 100\text{ms}$, $n = 5$, $\Delta t = 10\text{s}$	144
Table 33 – Examples with $P = 100\text{ms}$, $n = 5$, $\Delta t = 30\text{s}$	145
Table 34 – Examples with $P = 100\text{ms}$, $n = 5$, $\Delta t = 60\text{s}$	146
Table 35 – Examples with varying loss for 300s	148
Table 36 – Examples with $P = 100\text{ms}$, $C = 0.01$, $\Delta t = 10\text{s}$, $n = 10$	149
Table 37 – Examples with $P = 100\text{ms}$, $C = 0.01$, $\Delta t = 10\text{s}$, $n = 5$	150
Table 38 – Examples with $P = 100\text{ms}$, $C = 0.01$, $\Delta t = 10\text{s}$, $n = 3$	151

1. List of Contributing Members

The following members of the MEF participated in the development of this revision of this document and have requested to be included in this list.

ALU	EXFO
AT&T	Huawei
Bell Canada	InfoVista
Ciena	PLDT
Cisco	RAD
Comcast	Siana Systems
Ericsson	

2. 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 [15].

Existing PM Functions are defined by ITU-T G.8013/Y.1731 [1], and ITU-T G.8021 [3] as amended [4] [5]. This document details how to use these functions in order to achieve the requirements of MEF SOAM PM.

3. Terminology

Term	Definition	Reference
IDM	One-way Delay Measurement Message.	ITU-T G.8013/Y.1731 [1]
ISL	One-way Synthetic Loss Measurement Message.	ITU-T G.8013/Y.1731 [1]
Available Time	The set of small time intervals Δt contained in T that do not intersect a Maintenance Interval and are evaluated as Available.	MEF 10.3 [12]
Availability	A measure of the percentage of time that a service is useable.	MEF 10.3 [12]
Availability flr	The Availability flr (in contrast with FLR) is the ratio of lost frames to sent frames over a small interval of time Δt (e.g. 1 sec).	MEF 10.3 [12]
Availability Window	A period of n consecutive intervals of Δt , used to determine whether the Available state has been entered or exited.	This Document
Backward	The direction of performance measurements from the Responder MEP towards the Controller MEP, when One-way measurements are taken using a Single-Ended PM Function. Note: this term is not applicable when Dual-Ended PM Functions are used.	This Document
CEN	Carrier Ethernet Network	MEF 12.2 [13]
CHLI	Consecutive High Loss Interval	MEF 10.3 [12]
Controller MEP	The MEP that initiates SOAM PDUs. Term is applicable to both Dual-Ended and Single-Ended PM Functions. In a Single-Ended PM Function, the Controller MEP also receives responses from the Responder MEP.	This Document
CoS	Class of Service	MEF 23.1 [17]
CoS ID	Class of Service Identifier	MEF 23.1 [17]
CoS ID for SOAM PM Frames	Class of Service Identifier for SOAM PM Frames. One of the mechanisms, and/or the values of the parameters in the mechanisms, available for use by SOAM PM Frames such that their performance is representative of the performance of the Qualified Service Frames being monitored.	This Document
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 [17]
CoS FS	Class of Service Frame Set	MEF 23.1 [17]

Term	Definition	Reference
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 [17]
CoV	Coefficient of Variation	This Document
DEI	Discard Eligible Indicator	IEEE 802.1Q-2014 [22]
DMM	Delay Measurement Message	ITU-T G.8013/Y.1731 [1]
DMR	Delay Measurement Reply	ITU-T G.8013/Y.1731 [1]
Dual-Ended	A process whereby a Controller MEP sends measurement information to a peer Sink MEP that will perform the calculations. Dual-Ended processes can only be used to make One-way measurements.	This Document
EI	External Interface – Either a UNI or an ENNI	MEF 12.2 [13]
EMS	Element Management System	MEF 15 [14]
ENNI	External Network-to-Network Interface	MEF 4 [9]
ETH-DM	Ethernet Frame Delay Measurement Function (term is only used to reference the ITU-T definition)	ITU-T G.8013/Y.1731 [1]
ETH-LM	Ethernet Frame Loss Measurement Function (term is only used to reference the ITU-T definition)	ITU-T G.8013/Y.1731 [1]
ETH-SLM	Ethernet Synthetic Loss Measurement Function (term is only used to reference the ITU-T definition)	ITU-T G.8013/Y.1731 [1]
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 10.3 [12]
FD	Frame Delay	MEF 10.3 [12]
FDR	Frame Delay Range	MEF 10.3 [12]
FLR	Frame Loss Ratio	MEF 10.3 [12]
Forward	The direction of performance measurements from the Controller MEP towards the Responder or Sink MEP, when One-way measurements are taken using a Single-Ended or Dual-Ended PM Function.	This Document
Group Availability	A measure of the percentage of time that at least K subsets of ordered UNI pairs within an EVC are Available	MEF 10.3 [12]
HLI	High Loss Interval	MEF 10.3 [12]
IFDV	Inter-Frame Delay Variation	MEF 10.3 [12]

Term	Definition	Reference
LM	Loss Measurement	This Document
LMM	Loss Measurement Message	ITU-T G.8013/Y.1731 [1]
LMR	Loss Measurement Reply	ITU-T G.8013/Y.1731 [1]
MA	Maintenance Association This term is equivalent to a Maintenance Entity Group, or MEG, as defined by ITU-T G.8013/Y.1731 [1], which is the term used in this IA.	IEEE 802.1Q-2014 [22]
Maintenance Interval	A time interval agreed to by the Service Provider and Subscriber during which the service may not perform well or at all.	MEF 10.3 [12]
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 [15] and used in MEF 30.1 [19] (which is the term used in this IA).	IEEE 802.1Q-2014 [22]
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 Frame'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 G.8013/Y.1731 [1]), which is the term used in this IA.	IEEE 802.1Q-2014 [22]
ME	Maintenance Entity. A point-to-point relationship between two MEPs. This term is used by both IEEE and ITU-T.	IEEE 802.1Q-2014 [22] ITU-T G.8013/Y.1731 [1]
Measurement Bin	A counter that stores the number of FD, IFDV or FDR measurements falling within a specified range, during a Measurement Interval.	This Document
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 <i>T</i> .	This Document

Term	Definition	Reference
Measurement Interval Data Set	The collection of completed measurements that were initiated during a given Measurement Interval.	This Document
MEG	Maintenance Entity Group A set of MEPs, each configured with the same MEG ID and MEG Level, established to verify the integrity of a single service instance. A MEG 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 Association, or MA, as defined by IEEE 802.1Q-2014 [22]. MEG is the term used in this IA.	ITU-T G.8013/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 MEG among those associated with the SOAM Frame'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 G.8013/Y.1731 [1]
MEP	Maintenance Association End Point (IEEE 802.1Q-2014 [22]), or equivalently MEG End Point (ITU-T G.8013/Y.1731 [1] or MEF 17 [15]). 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-2014 [22] ITU-T G.8013/Y.1731 [1] MEF 17 [15]
MFD	Mean Frame Delay	MEF 10.3 [12]
MIP	Maintenance Domain Intermediate Point (IEEE 802.1Q-2014 [22]) or equivalently a MEG Intermediate Point (ITU-T G.8013/Y.1731 [1] or MEF 17 [15]). An intermediate point in a MEG that is capable of reacting to some SOAM PDUs, but does not initiate SOAM PDUs.	IEEE 802.1Q-2014 [22] ITU-T G.8013/Y.1731 [1] MEF 17 [15]
NE	Network Element	MEF 15 [14]
NMS	Network Management System	MEF 15 [14]
OAM	Operations, Administration, and Maintenance	MEF 17 [15]

Term	Definition	Reference
OAM Domain	See MD (Maintenance Domain)	MEF 30.1 [19]
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 [9]
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. One-way measurements can be performed using either Single-Ended or Dual-Ended PM Functions.	This Document
OVC	Operator Virtual Connection	MEF 26.1 [18]
PCP	Priority Code Point	IEEE 802.1Q-2014 [22]
PDU	Protocol Data Unit	This Document
PM	Performance Monitoring 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)	This Document
PM Session	The application of a given PM Function between a given pair of MEPs on a given SOAM PM CoS ID over some (possibly indefinite) period of time.	This Document
PM Solution	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.	This Document
Proactive	OAM actions that are carried on continuously to permit timely reporting of fault and/or performance status.	RFC 5951 [9]
Qualified Service Frames	The set of frames that comply with specific criteria, such as the arrival time at the Ingress UNI and Bandwidth Profile compliance, on which a performance attribute is based.	MEF 10.3 [12]
Responder MEP	In a Single-Ended PM Session, the MEP that receives SOAM PM PDUs from the Controller MEP, and transmits responses to the Controller MEP.	This Document
S	A non-empty subset of ordered UNI pairs within a MEG	MEF 10.3 [12]
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 10.3 [12]

Term	Definition	Reference
Single-Ended	A process whereby a Controller MEP sends a measurement request and a peer Responder MEP replies with the requested information so that the originating MEP can calculate the measurement. Single-Ended processes can be used to make One-way and Two-way measurements.	This Document
Sink MEP	In a Dual-Ended PM Session, the MEP that receives SOAM PM PDUs from the Controller MEP and performs the performance calculations.	This Document
SLM	Synthetic Loss Message	ITU-T G.8013/Y.1731 [1]
SLR	Synthetic Loss Reply	ITU-T G.8013/Y.1731 [1]
SLS	Service Level Specification	MEF 10.3 [12]
SOAM	Service Operations, Administration, and Maintenance	MEF 17 [15]
SOAM PM CoS ID	See CoS ID for SOAM PM Frames	This Document
SOAM PM Frame	An Ethernet frame containing a SOAM PM PDU in the Data field.	This Document
SOAM PM Implementation	Capabilities of an NE that are required to support SOAM Performance Monitoring.	This Document
SOAM PDU	Service OAM Protocol Data Unit. Specifically, those PDUs defined in IEEE 802.1Q-2014 [22], ITU-T G.8013/Y.1731 [1], or MEF specifications. In ITU-T documents the equivalent term OAM PDU is used.	This Document
SOAM PM PDU	Service OAM Protocol Data Unit specifically for Performance Measurement. Examples are LMM/LMR, DMM/DMR/1DM, SLM/SLR/1SL.	This Document
Suspect Flag	A flag included in each Measurement Interval Data Set indicating whether a discontinuity (as described in section 10.2.4) occurred in the measurements taken during the Measurement Interval.	ITU-T X.738 [25]
Synthetic Frame	An Ethernet frame created to emulate service traffic, carry additional information necessary to support calculating performance metrics (e.g. delay or loss) and that is treated the same way as a Qualified Service Frame.	This Document

Term	Definition	Reference
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. <i>T</i> is at least as large as the Measurement Interval, and generally consists of multiple Measurement Intervals. Also note that <i>T</i> can have different values for different performance metrics.	MEF 10.3 [12]
TCA	Threshold Crossing Alert	GR-253 [24]
ToD	Time-of-day	This Document
Two-way	A measurement of the performance of frames that flow from the Controller MEP to Responder MEP and back again. Two-way measurements can only be performed using Single-Ended PM Functions.	This Document
UBC(k)	Upper Bin Count (k)	This Document
Upper Bin Count (k)	The total count of Measurement Bin k and above, i.e., Count of Bin(k) + Count of Bin(k+1) +...+ Count of Bin(n)	This Document
UTC	Coordinated Universal Time	ISO 8601 [23]
Unavailable Time	The set of small time intervals Δt contained in <i>T</i> that do not intersect a Maintenance Interval and are evaluated as Unavailable.	MEF 10.3 [12]
UNI	User-to-Network Interface	MEF 10.3 [12]
VID	Virtual Local Area Network Identifier	IEEE 802.1Q-2014 [22]
VLAN	Virtual Local Area Network	IEEE 802.1Q-2014 [22]

Table 1 – Terminology and Definitions

4. 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 [15]. The goal of this IA is to define specific performance measurement procedures and specify solutions for collecting performance measurements, for informational purposes or to compute the performance metrics defined by MEF 10.3 [12] and MEF 26.1 [18] (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 G.8013/Y.1731 [1], and ITU-T G.8021 [3] as amended [4] [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.

4.1 Change History

The following changes were made between MEF 35 and this revision. Note: references to section or requirement numbers are the numbers used in this revision unless otherwise specified.

- Amendment 1, MEF 35.0.1, was incorporated. This added a new PM Solution, PM-4, for Dual-Ended Synthetic Frame Loss using 1SL PDUs. The changes include new sections 9.4, 14 and 23. The text incorporated in section 23 was updated to apply generally to Dual-Ended PM Functions.
- Amendment 2, MEF 35.0.2, was incorporated. This added Threshold Crossing Alerts. The changes include new section 10.5. Text incorporated from the Amendment was clarified by distinguishing between a threshold and a TCA Function.
- References to the following documents were updated, including the text in section 7 where appropriate: MEF 7.2, MEF 10.3, MEF 12.2, MEF 23.1, MEF 26.1, MEF 30.1, MEF 47, ITU-T G.8013/Y.1731, ITU-T G.8021, ITU-T X.738, IEEE 802, IEEE 802.1Q. In the case of MEF 10.3, text was also added throughout the document relating to the Group Availability performance metric defined in the new revision. The references to MEF 6.1, MEF 36, and ITU-T G.8010 were removed.
- The following definitions were added or clarified in section 3: 1SL, Available Time, Backward, CEN, Controller MEP, CoS ID for SOAM PM Frames, Dual Ended, Forward, Group Availability, MA, ME, MEG, MIP, One-way, Qualified Service Frames, Single-Ended, SOAM PM Frame, Suspect Flag, Two-way, Unavailable Time.
- Changes were made in the scope and elsewhere, and a new section, 6.3, was added, to clarify that performance measurements may be made for informational purposes, and not just for evaluating an SLS.
- Conditional Requirements format was used where appropriate, and the explanatory boilerplate was added in section 5.

- References to EVCs throughout the document were updated to ensure they also covered OVCs, or were otherwise made more generic.
- In section 8.2, the effect of bandwidth profiles on Frame Loss measurements using Service Frames in various MEGs was described in more detail.
- A new section, 8.3, was added describing the difference between packet-count-based and time-based sampling.
- Section 8.4, describing CoS considerations, was completely rewritten, and corresponding updates were made elsewhere in the document.
- A paragraph was added in section 10.1.1 describing how performance measurements can be made when elastic Ethernet services are used.
- In addition to those added from incorporating the amendments, the following requirements were added or modified: [R1], [D1], [R9], [R10], [R11], [R12], [R13], [R16], [R17], [D3], [D5], [D6], [D8], [R22], [R26], [D12], [R43], [R44], [D20], [R45], [R66], [D28], [R78], [R79], [D31], [R82], [D33], [R87], [CR25], [CR26], [CR35], [CR38]. In the sections incorporated from the amendments, the following requirements were added or modified: [CR1], [CR2], [CR5], [CR7], [CR8], [CR9], [CR49], [CD13], [CR58], [CR59], [CD16], [CR62], [CD18], [CR67].
- The following requirements were removed (numbers refer to MEF 35): R2, R13, D21, R57, R69, D28, R72, D38, R91, R93, D45, R106. In the sections incorporated from MEF 35.0.1, the following requirements were removed: A1-D1, A1-D2, A1-R5.
- Informative text was added in section 10.2.2, 11.1 and 22.2 to clarify the treatment of negative delay measurements. Text and figures in section 10.2.5 were updated to clarify the handling of measurements during periods of Unavailable Time.
- Additional text was added in sections 11.2 and 14.1, and a new Appendix was added in section 25, to provide guidance on selecting parameters for Availability measurements.
- In each of the output data sets, it was clarified whether the counters are applicable at the Controller MEP, Responder MEP or Sink MEP. In addition, measurements of minimum IFDV were removed from the required data sets.
- Informative text was added in section 13.1 to explain the differences between MEF 10.3, G.8013/Y.1731 and IEEE 802.1Q with respect to LMMs and FLR.
- A new Appendix was added in section 24 describing the calculation of SLS performance metrics.
- Numerous editorial and typographical corrections were made.

5. 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 are in upper case, bold text.

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

A paragraph preceded by [CRa]< specifies a conditional mandatory requirement that **MUST** be followed if the condition(s) following the "<" have been met. For example, "[CR1]<[D38]" indicates that Conditional Mandatory Requirement 1 must be followed if Desirable Requirement 38 has been met. A paragraph preceded by [CDb]< specifies a Conditional Desirable Requirement that **SHOULD** be followed if the condition(s) following the "<" have been met. A paragraph preceded by [COc]< specifies a Conditional Optional Requirement that **MAY** be followed if the condition(s) following the "<" have been met.

6. 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 [15] (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.

6.1 OAM Domains

SOAM allows a network to be partitioned into a set of hierarchical OAM Domains (see MEF 30.1 [19] 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 10-14 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

6.2 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 30.1 [19] identifies the default MEs and the Maintenance Entity Group (MEG) levels.

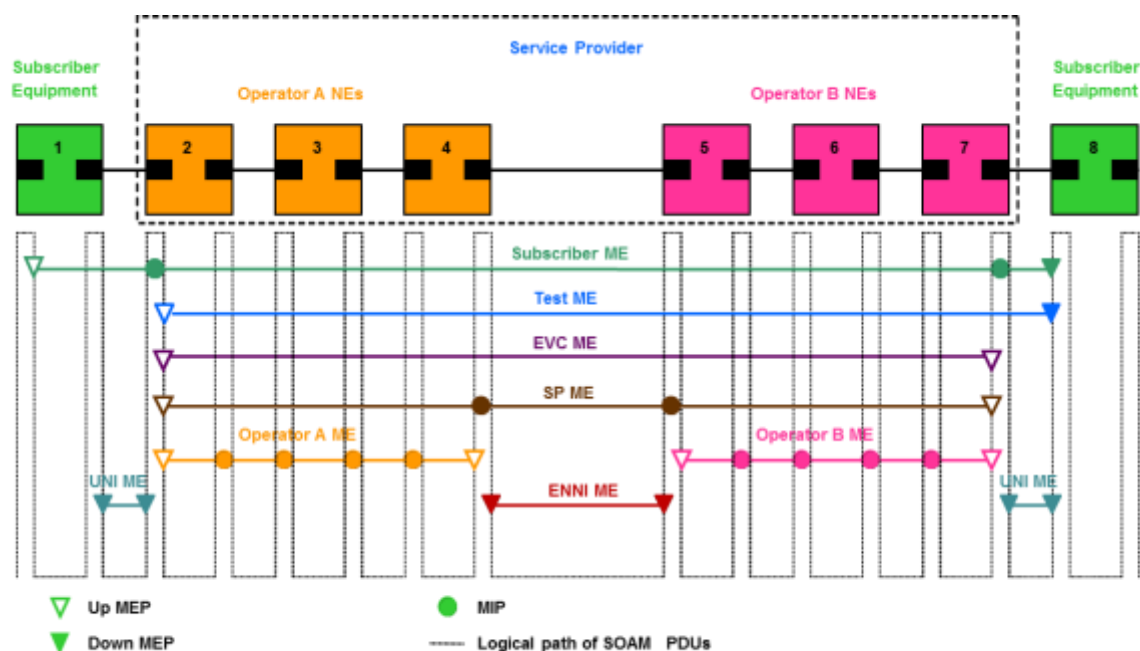


Figure 1 – Maintenance Entities (see MEF 30.1 [19])

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. In this case there is a separate set of 8 MEG Levels for each stacked VLAN tag, as described in MEF 30.1 [19] Appendix B. 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.

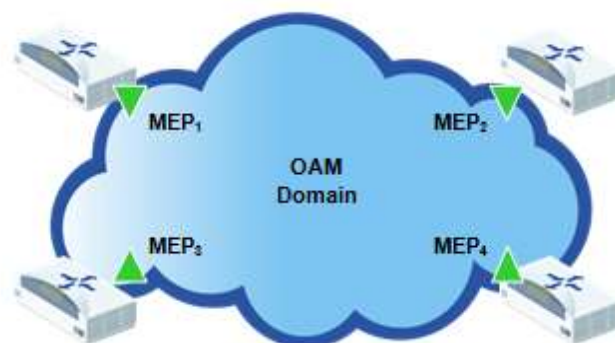


Figure 2 – OAM Domain

6.3 OAM Domains and Performance Metrics

Various performance metrics are defined in MEF 10.3 [12] for EVCs, and equivalently in MEF 26.1 [18] for OVCs (see sections 7.2 and 7.7), for the purpose of evaluating conformance to a Service Level Specification (SLS). EVCs and OVCs correspond respectively to the EVC and Operator OAM Domains described in section 6.1 above. However, performance measurements may also be carried out for informational reasons in other OAM Domains. In this case, an analogous definition of the performance metrics may be assumed as described below.

The performance metrics are defined over a set of ordered pairs of UNIs (in the EVC case according to MEF 10.3 [12]) or a set of ordered pairs of OVC End Points (in the OVC case according to MEF 26.1 [18]). For each performance metric, the performance for each ordered pair in the set is defined, and these values are combined to obtain the value for the EVC or OVC as a whole.

The specific performance measurements defined in this document are always between a single pair of MEPs. The MEPs may be located at UNIs (in the EVC OAM Domain), OVC End Points (in the Operator OAM Domain) or other constructs (in the other cases). In each case, the definition of a performance metric in MEF 10.3 [12] or MEF 26.1 [18] as it applies to a single pair of UNIs or OVC End Points can be applied to any pair of MEPs in the same MEG, and hence performance measurements can be made for any of the OAM Domains described in section 6.1 above.

The performance metrics as defined in MEF 10.3 [12] and MEF 26.1 [18] apply to Qualified Service Frames or to Qualified Frames respectively – i.e., Service Frames or ENNI Frames that meet certain criteria including that they have Ingress Bandwidth Profile compliance of green. However, there are certain OAM Domains – e.g. in the UNI MEG, ENNI MEG or Subscriber MEG – where the Ingress Bandwidth Profile compliance of Service Frames and/or SOAM PM Frames may not be applicable.

Note: in this document the terms “Service Frame” and “Qualified Service Frame” are used, as defined in MEF 10.3 [12]. These may be taken to also apply to ENNI Frames and to Qualified Frames (as defined in MEF 26.1 [18]) generally, as appropriate.

The majority of performance measurements defined in this document actually measure the performance of Synthetic SOAM PM Frames, rather than of Qualified Service Frames; since SOAM PM Frames always have Ingress Bandwidth Profile compliance of green within the CEN, they always

measure the performance of green frames, as required. An exception to this is when Single-Ended Service Frame Loss Measurement is used as part of PM Solution PM-3 (see section 9), as in this case it is the loss of Qualified Service Frames that is measured. As explained in section 8.2, the application of the PM-3 Solution is not recommended for OAM Domains other than the EVC or Operator OAM Domains.

With the above points in mind, references in the remainder of this document to the performance metrics defined in MEF 10.3 [12] and MEF 26.1 [18] can be taken as applying to any OAM Domain.

Note: details of how the performance measurements defined in this document can be used to calculate the performance metrics defined in MEF 10.3 [12] and MEF 26.1 [18] can be found in Appendix I – Calculation of SLS Performance Metrics (Informative).

6.4 Default Behavior

One of the important functions of this document is to simplify the provisioning of SOAM across the Carrier Ethernet Network (CEN). 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 desires it. 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.

7. 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.

7.1 MEF 7.2

MEF 7.2 [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.

7.2 MEF 10.3

MEF 10.3 [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 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
- High Loss Intervals (HLI)
- Consecutive High Loss Intervals (CHLI)
- Group Availability

The performance metrics encompass Qualified Service Frames flowing in one direction over a subset of ordered UNI pairs (i.e., some or all) of an EVC, over a time period, 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
- 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

The objectives are uni-directional (specified in MEF 10.3 [12] section 8.8), however, the measurement can be done using bi-directional means. Also see section 8.8.1 in MEF 10.3 [12].

7.3 MEF 15

MEF 15 [14] 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.

7.4 MEF 17

MEF 17 [15] provides a high level overview of SOAM architecture and capabilities, and discusses some of the requirements for a SOAM PM Implementation.

According to MEF 17 [15], SOAM must provide the ability to 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.

Note: The definition of performance metrics has evolved over time and the performance metrics defined in MEF 10.3 [12] differ from the measurement requirements in MEF 17 [15], and include additional performance metrics not envisaged in MEF 17, such as FDR and Availability. This document is aligned with MEF 10.3.

7.5 MEF 20

MEF 20 [16] 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

Note: The SOAM requirements in MEF 20 [16] have been incorporated in MEF 30.1 [19].

7.6 MEF 23.1

MEF 23.1 [17] updates how the term "Class of Service" (CoS) is used in MEF specifications. To avoid ambiguity, the terms "CoS" and "CoS ID" should never be used on their own, but always with additional context. MEF 23.1 also introduces the concept of Performance Tiers, and specifies performance objectives for these Performance Tiers for three standardized CoS Names designated as 'CoS Labels'. The MEF 23.1 framework can be used for additional CoS Names beyond the three standardized CoS Labels.

7.7 MEF 26.1

MEF 26.1 [18] defines the requirements for the External Network Network Interface (ENNI). The document specifies a reference point that is the interface between two Carrier Ethernet Networks. The term Operator Virtual Connection (OVC) is defined in that document. MEF 26.1 also defines Service Level Specification performance metrics and related requirements for OVCs, i.e. where the set S contains ordered pairs of OVC End Points. These definitions are equivalent to the performance metrics specified in MEF 10.3 for EVCs and UNIs (see section 7.2); hence wherever MEF 10.3 is referred to within this document, it can be read as also being a reference to MEF 26.1 in the case of OVCs.

7.8 MEF 30.1

MEF 30.1 [19] and MEF 30.1.1 [20] (SOAM FM IA) provide 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.

8. PM Considerations

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

8.1 Frame Delay Measurements

Measuring the One-way FD of a Qualified 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 Qualified 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 of Service 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 10.3 [12] for Qualified Service 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.

8.2 Frame Loss Measurements

Measuring the One-way FLR of Qualified 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 Qualified Service Frames, as specified in G.8013/Y.1731 [1] and ITU-T G.8021 [3] as amended [4] [5] using the LM process.
- Measuring the loss of Synthetic Frames (SOAM PM PDUs using SLM/SLR or 1SL), as specified in G.8013/Y.1731 [1], and ITU-T G.8021 [3] as amended [4] [5].

A One-way FLR measurement that measures loss of Qualified 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.
- As only frames with an Ingress Bandwidth Profile compliance of green are counted at each MEP, the measurement can be affected if the Service Frames flow through an Ingress or Egress Bandwidth Profile between the two MEPs which changes the color marking of some frames. The details depend on the MEG, as follows:
 - SP MEG, EVC MEG and Operator MEG: MEPs are located at EIs within the CEN (after the Ingress Bandwidth Profile and before the Egress Bandwidth Profile), and hence the counters can correctly determine the color of Service Frames.
 - If color marking is preserved across the EVC or OVC, then both MEPs count the same green frames.
 - If color marking is not preserved, then yellow frames may be promoted to green, or green frames may be demoted to yellow, as they flow across the EVC or OVC; hence different frames might be counted at each MEP. Service Frame Loss Measurement may be inaccurate in this case.
 - Subscriber MEG, UTA SP MEG: the MEPs do not coincide with the ingress and egress of the CEN, and hence the coloring of Service Frames may be altered by an Ingress or Egress Bandwidth Profile. This may result in different frames being counted by each MEP, leading to inaccuracies in the measurement. Note in the case of the UTA SP MEG there may be an Egress BWP per VUNI, which is implemented in the Ethernet Provider Conditioning Function (EPCF) after the UTA SP MEP. There may also be an ingress BWP on the access provider side of the ENNI.
 - UNI MEG, ENNI MEG: As these are not per-service MEGs, the counters could treat all frames as green, i.e. all frames flowing over the interface are counted. For the UNI MEG, the Service Frames are counted at the MEP before the determination of the color, which occurs after the MEP at the ESCF; therefore the only choice within the UNI MEG is to treat all frames as green, i.e. to count all frames. Service Frames flowing over the ENNI may belong to different services, with different color markings; in this case implementing aggregate counters that count only green frames over all services may be impractical.

For the reasons given above, measuring the loss of Qualified Service Frames is not recommended for MEGs other than the EVC and Operator MEGs, and then only when color marking is preserved in the corresponding EVC or OVC.

Another limitation of the LM process is that in a multipoint MEG, counters of Qualified 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, the limitations on counter accuracy and with multipoint MEGs can be overcome by using directed and periodic Synthetic Frames. By counting and measuring the One-way FLR of uniform Synthetic Frames, statistical methods can be used to estimate the One-way FLR of service traffic. This can be achieved by inserting periodic Synthetic SOAM PM Frames into an EVC or OVC, 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 for different SOAM PM CoS IDs in a straightforward 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 number of Synthetic SOAM PM PDUs used to make the measurement, which in turn depends on the rate at which they are sent, and the time over which the measurement is made. In general, more frequent transmission and/or longer timeframes are needed to obtain estimates with the required accuracy.

8.2.1 Location of PM Measurement Points (for Loss)

As discussed in sections 7.2 and 4, MEF 10.3 [12] 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 point¹.

Figure 3 shows the location of MEPs within a UNI, in relation to the traffic conditioning points and the Ethernet ECS Adaptation Function (EEAF), as specified in MEF 12.2 [13]. 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 12.2 [13], where the MEP is between the Ethernet Subscriber Conditioning Function (ESCF) and the 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.

¹ Note that in MEF 12.2 [13], the ESCF is the traffic conditioning point for the UNI-N, 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).




Figure 3 – MEP Placement

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 or OVCs 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 EI 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.

8.3 Packet-Count-Based versus Time-Based Measurements

The ideal performance metrics specified in MEF 10.3 [12] are all based on the actual performance experienced by Qualified Service Frames. However, measuring this actual Qualified Service Frame performance would require modifying Service Frames to include fields such as timestamps, sequence numbers, etc. that are required in order to measure performance. As noted above, an alternative is to insert Synthetic Frames into the Service Frame traffic stream, and then use fields in these Synthetic Frames to obtain performance measurements. Using Synthetic Frames to obtain performance measurements overcomes the limitations of obtaining performance measurements from actual Service Frames (e.g., the process of modifying Service Frames to include necessary fields would, among other issues, affect the performance of the Service Frames). Each Synthetic Frame can be considered to be one statistical measurement point.

MEF 10.3 [12] defines several different performance metrics, including maximums, minimums, percentiles, and averages. The averages defined in MEF 10.3 are all packet-count-based averages, that is, for an average metric \bar{M} over a time interval T:

$$\bar{M} = \frac{\sum \text{measurement of Qualified Service Frames in } T}{\sum \text{number of Qualified Service Frames in } T}$$

Similarly, percentiles are based on percentages of all Qualified Service Frames (for a given CoS FS) received by the CEN during T.

Obtaining these measurements through statistical sampling would require the measurement points to be placed based upon received Service Frame counts, that is, it would require a Synthetic Frame to be inserted after every x Qualified Service Frames of the given CoS FS. While such a packet-count-based sampling method might be theoretically ideal, it faces several practical implementation problems, especially in a high packet rate environment. The received Service Frames may consist of a mixture of multiple CoS Frame Sets (or other similar sets of frames requiring separate performance measurements). For each such set, the implementation needs to count the number of appropriate frames (e.g., only Qualified Service Frames of only that CoS FS), create an appropriate Synthetic Frame, and insert it immediately after the correct number of appropriate Service Frames have passed. The Synthetic Frame must have the correct data for measurements, in particular, delay measurements require an accurate timestamp. A slow implementation may not be able to generate and insert Synthetic Frames quickly enough during a Service Frame burst, possibly leading to a backup of measurements being taken after the burst has passed and therefore not accurately measuring the performance experienced by Qualified Service Frames.

Figure 4 below shows an example of theoretically ideal packet-count-based measurement points, and Figure 5 shows what the actual measurement points might look like in a practical implementation. In these figures “f” represents a Qualified Service Frame and “m” represents a measurement point (i.e., an inserted Synthetic Frame). For illustration, in these figures the sampling rate is one measurement point after every 3 Qualified Service Frames. Time flows from left to right, so frames on the left occur earlier than frames on the right.




Figure 4 – Ideal Packet-Count-Based Measurement Samples




Figure 5 – Possible Actual Packet-Count-Based Measurement Samples

Another implementation is to insert Synthetic Frames on a time basis, rather than on a packet-count basis. In a time-based implementation, a Synthetic Frame is inserted once every y time units

(e.g., milliseconds), regardless of how many or how few Qualified Service Frames have been received. Figure 6 shows an example of a time-based implementation for the same Service Frame flow used in Figure 4 and Figure 5.



Figure 6 – Time-Based Measurement Samples

A time-based implementation has two practical advantages over a packet-count-based implementation. First, a time-based implementation frees the Synthetic Frame generator from having to monitor the Service Frames. Second, the processor load in a time-based implementation is constant, whereas a packet-count-based implementation leads to increased processor load as the rate of Qualified Service Frames (of the given CoS FS) increases (because more Service Frames results in more Synthetic Frames being generated).

Besides the practical implementation advantages, time-based implementations offer some statistical advantages over packet-count-based implementations. Time-based implementations measure network performance at all times, whereas packet-count-based implementations only measure network performance when the network is being used, and provide no information when the network is not being used. Network issues that occur when no Qualified Service Frames are being sent would be spotted by time-based measurements, but would be missed by packet-count-based measurements. In the extreme case, a standby link or path on which there are no Service Frames flowing can be monitored using time-based implementations, but cannot be monitored at all using packet-count-based implementations.

In addition, time-based measurements result in a constant number of measurement samples per Measurement Interval, whereas under a packet-count-based implementation the number of measurement samples depends on the number of Qualified Service Frames (of the given CoS FS) received during the Measurement Interval and thus can vary from one Measurement Interval to another. Having a constant number of measurement samples in each Measurement Interval provides considerable advantages to statistical analysis of the measurement results.

8.4 CoS Considerations

A single Ethernet service might encompass multiple Classes of Service, and therefore it may be desirable to take performance measurements for each Class of Service between the same two end points. The Class of Service Identifiers (CoS IDs) that are available for use by SOAM PM Frames are called CoS IDs for SOAM PM Frames, or "SOAM PM CoS IDs." A SOAM PM CoS ID is limited to mechanisms that can be carried by a SOAM PM Frame. For example, SOAM PM CoS ID cannot be based on IP DSCP because SOAM PM Frames do not carry any IP information. Similarly, SOAM PM CoS ID cannot be based on L2CP since, as defined in MEF 10.3 [12], SOAM Frames are not considered to be L2CP Frames. SOAM PM CoS IDs are defined as one of the following:

- VLAN ID
- combination of VLAN ID and PCP value

Note: SOAM PM Frames can be untagged (e.g. on the UNI and ENNI MEGs); in this case they do not have a SOAM PM CoS ID.

SOAM PM Frames measure the performance experienced by frames with a particular SOAM PM CoS ID. This measurement is applicable to a set of Qualified Service Frames to the extent that the performance of frames with a SOAM PM CoS ID reflects the performance of that set of Qualified Service Frames. For example, if the set of Service Frames is the CoS Frame Set (CoS FS) for some CoS Name, and if SOAM PM Frames with a SOAM PM CoS ID of {VLAN ID = 42, PCP = 3} are used to measure the performance of this CoS FS, then the measurement is only accurate if the network provides the same performance to SOAM PM Frames with {VLAN ID = 42, PCP = 3} as the network does to the CoS FS.

A SOAM PM CoS ID is used for Performance Monitoring in different ways by Synthetic SOAM PM Frames and by LMM / LMR frames. Synthetic frames use a SOAM PM CoS ID to directly measure the performance experienced by frames that have that SOAM PM CoS ID. While LMM/LMR frames are used to transfer measurement information rather than being used to directly experience performance, nonetheless it is essential for accurate loss measurements that the LMM and LMR frames do not get misordered with respect to the Service Frames they are counting. The easiest way to avoid such misordering is to ensure that the CoS ID for LMM/LMR frames is treated the same as the CoS ID for the Service Frames being measured. In addition, LMM / LMR frames use the SOAM PM CoS ID as a means to identify the correct set of service traffic counters whose values should be carried by these frames.

Note that, if Synthetic SOAM PM Frames are used, then multiple sets of performance objectives (e.g., CoS Names) could be measured using the same SOAM PM CoS ID between a given pair of MEPs provided the different sets of Service Frames (e.g., CoS Frame Sets) experience the same performance over the network spanned by that particular ME. For example, suppose that CoS Name Platinum uses a CoS ID of PCP = 6 and CoS Name Gold uses a CoS ID of PCP = 5. Suppose also that both CoS Name Gold and CoS Name Platinum use the same UTA Service and that both use the same UTA Service CoS. As a result, frames with PCP = 5 and frames with PCP = 6 get mapped to the same queues and get treated identically within the network spanned by the UTA SP ME. In this case, performance for both CoS Name Gold and CoS Name Platinum could be monitored by one PM Session, e.g., using SOAM PM CoS ID of PCP = 5. The interpretation of the measurements could be different for each CoS Name (e.g., the performance data could indicate a Pass for CoS Name Gold but a Fail for CoS Name Platinum) but the same performance data could be used for both CoS Names. In contrast, if LMM/LMR frames are used, then each CoS Name (or other set of performance objectives) to be monitored must use a unique SOAM PM CoS ID due to the requirement for LMM/LMR frames to identify separate service traffic counters for each CoS Name.

To provide some examples, consider the Performance Monitoring of Qualified Service Frames between two UNI-Ns on an EVC, which corresponds to two MEPs on an EVC ME. Under MEF 10.3, the CoS ID for Service Frames may be based on EVC, PCP, or IP. If the CoS ID for Service

Frames is based on EVC or PCP, then the mapping of CoS Name to SOAM PM CoS ID is straightforward. If multiple Service Frame VLAN IDs and/or PCP values are mapped to the same CoS Name, then any one of these may be used as the SOAM PM CoS ID. In this case, since at least some of the Service Frames in the CoS FS use the SOAM PM CoS ID themselves, the performance of SOAM PM Frames will reflect the performance of the CoS FS.

When CoS ID for Service Frames is based on IP, then a different basis must be used to map CoS Names to SOAM PM CoS IDs. As noted earlier, to obtain relevant measurements using SOAM PM Frames, the performance experienced by SOAM PM Frames with a given SOAM PM CoS ID must be equivalent to the performance experienced by the CoS FS. In particular, in order to use measurements from SOAM PM Frames, then any decisions affecting performance inside the CEN cannot be made on the basis of IP DSCP values.

CoS Frame Sets, and corresponding CPOs, are defined across EVCs and OVCs, and over an ENNI; these correspond with the EVC, Operator and ENNI MEGs, and are the network sections over which an SLS may be specified. Performance Measurements may also be taken in other MEGs for informational purposes, e.g. in the SP MEG, SP UTA MEG or Subscriber MEG. In this case the same considerations apply, i.e. the SOAM PM CoS ID chosen must be chosen such that the performance of the SOAM PM frames reflects the performance of the Service Frames whose performance is being monitored, across the network section spanned by the MEG.

9. 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 a specific mechanism that is described in ITU-T G.8013/Y.1731 (e.g. Single-Ended ETH-SLM). A PM Session is an instantiation of a particular PM Function within a PM Solution between a given pair of MEPs using a given SOAM PM CoS ID 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 7 – PM Solution Components below.




Figure 7 – PM Solution Components

This implementation agreement covers requirements on the components in the Network Element Layer of Figure 7 – 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, which are outside the scope of this IA but nonetheless part of the overall PM Solution, 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 Element Management Layer and the Network Element Layer are documented in MEF 7.2. [11].

A conforming implementation of a PM Solution provides the SOAM PM and Management mechanisms necessary to meet the goals identified in section 4, including measurement of the performance metrics defined in MEF 10.3 [12]. The SOAM mechanisms covered in this IA are realized, in part, through the maintenance association architecture of IEEE 802.1Q-2014 [22], the PM Functions of ITU-T G.8013/Y.1731 [1], and the (network element based) atomic functions and processes of ITU-T G.8021 [3] as amended [4] [5].

A PM Solution can be categorized as to the types of MEG that it can be applied to and the PM Functions used. 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
PM-4	point-to-point multipoint	Synthetic Testing	Dual-Ended Synthetic Loss	Optional

Table 2 – PM Solutions Summary


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

PM Function	ITU-T PM Mechanism	ITU-T PDU(s)
Single-Ended Delay	ITU-T Single-Ended ² ETH-DM	DMM/DMR
Dual-Ended Delay	ITU-T Dual-Ended ³ 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
Dual-Ended Synthetic Loss	ITU-T Dual-Ended ETH-SLM	1SL

Table 3 – PM Functions Summary

An overview of the PM Functions is provided in Appendix A - Performance Management Functions (Informative). Note that use of Dual-Ended Service Loss PM Function is not recommended as part of any of the PM Solutions described in this document, as it cannot be used to measure loss for more than one Class of Service.

The following figures describe the performance metrics that can be calculated from the measurements collected with each PM Function. Note that calculating One-way FD requires Time-of-Day synchronization.




This can be used to produce measurements for the following performance metrics:

- One-way FD, IFDV, MFD and FDR (Forward and Backward), using DMM/DMR
- Two-way FD, IFDV, MFD and FDR, using DMM/DMR
- One-way FLR (Forward and Backward), using LMM/LMR
- One-way FLR (Forward and Backward), using SLM/SLR
- One-way Availability and Group Availability (Forward and Backward), using SLM/SLR
- One way HLI and CHLI (Forward and Backward), using SLM/SLR

Figure 8 – Performance Metrics that can be collected with Single-Ended Loss and Delay

² In older revisions of the ITU-T Recommendations, Single-Ended ETH-DM was known as Two-way ETH-DM.

³ In older revisions of the ITU-T Recommendations, Dual-Ended ETH-DM was known as One-way ETH-DM.



This can be used to produce measurements for the following performance metrics:

- One-way FD, IFDV, MFD and FDR (Forward), using 1DM
- One-way FLR (Forward), using 1SL
- One-way Availability and Group Availability (Forward), using 1SL
- One-way HLI and CHLI (Forward), using 1SL

Figure 9 – Performance Metrics that can be collected with Dual-Ended Loss and Delay

The following sections serve to briefly describe the individual PM Solutions, which are realized through the NE requirements specified in section 10 (Common Requirements) and sections 11 (PM-1), 12 (PM-2), 13 (PM-3), and 14 (PM-4) that follow.

9.1 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). Single-Ended Synthetic Loss measurement is used to measure Frame Loss (FLR), Availability, Group Availability, and count of High Loss Intervals (HLI, CHLI).

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 and MFD measurements can be taken. Two-way FD, MFD, FDR, and IFDV measurements and One-way FDR and 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 10.3 [12] can be made with this solution. For FD and MFD, if ToD synchronization is not sufficiently accurate for performance measurement purposes, the One-way performance metrics of MEF 10.3 [12] can be estimated by dividing the Two-way measurement by 2, although this introduces considerable statistical bias. 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. This mechanism can be used to take One-way measurements from which FLR, Availability and Group Availability can be calculated. FLR, Availability and Group Availability are defined in MEF 10.3 [12].

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 10.3 [12] 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 frames are 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 PM Frames to be statistically valid. Appendix D - Statistical Considerations for Synthetic Loss Measurement (Informative) contains further information with respect to FLR, and Appendix J – Statistical Considerations for Availability contains further information with respect to Availability.

All Synthetic SOAM PM Frames need to be similar to the Qualified Service Frames carried by the EVC or OVC, 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 Qualified Service Frames are treated. In addition, it is important that Synthetic SOAM PM Frames be inserted irrespective of 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 10.3 [12] that can be calculated for each ordered EI pair in the set S using the PM-1 Solution:

- One-way Frame Delay Performance (MEF 10.3 [12] section 8.8.1)
- One-way Mean Frame Delay (MEF 10.3 [12] section 8.8.1)
- One-way Frame Delay Range (MEF 10.3 [12] section 8.8.1)
- One-way Inter-Frame Delay Variation Performance (MEF 10.3 [12] section 8.8.2)
- One-way Frame Loss Ratio Performance (MEF 10.3 [12] section 8.8.3)
- One-way Availability Performance (MEF 10.3 [12] section 8.8.4)
- One-way Resiliency Performance (HLI and CHLI) (MEF 10.3 [12] section 8.8.5)

One-way Group Availability Performance (MEF 10.3 [12] section 8.8.6) for a group of EI pairs can also be calculated.

9.2 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. Dual-Ended PM Sessions can be 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 10.3 [12] are made with the PM-2 Solution.

The PM-2 Solution can be applied to either point-to-point or multipoint MEGs. 1DM frames can be unicast or multicast. In multipoint MEGs, use of multicast 1DM frames can help to simplify PM Session configuration and reduce SOAM traffic. This is further described in Appendix H – Notes on Dual-Ended PM Functions (Informative).

For One-way FD 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 the Dual-Ended Delay PM Function needs to generate enough SOAM PM Frames to be statistically valid. All Synthetic SOAM PM Frames need to be similar to the Qualified Service Frames carried by the EVC or OVC, 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 Qualified 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 10.3 [12] that can be calculated for each ordered EI pair in the set S using the PM-2 Solution:

- One-way Frame Delay Performance (MEF 10.3 [12] section 8.8.1)
- One-way Mean Frame Delay (MEF 10.3 [12] section 8.8.1)
- One-way Frame Delay Range (MEF 10.3 [12] section 8.8.1)
- One-way Inter-Frame Delay Variation Performance (MEF 10.3 [12] section 8.8.2)

9.3 PM-3: Single-Ended Service Loss Measurements

The PM-3 Solution is an optional solution that uses Service Frame counters to measure performance. This solution uses Single-Ended Service Loss measurement to measure Frame Loss Ratio (FLR). The PM-3 Solution is not applicable to multipoint MEGs.

LMM/LMR PDUs are used for FLR measurements. These collect the counts of the number of Qualified 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 frames 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 10.3 [12] that can be calculated using the PM-3 Solution:

- One-way Frame Loss Ratio Performance (FLR) (MEF 10.3 [12] section 8.8.3)

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

9.4 PM-4: Dual-Ended Synthetic Loss Measurements

The PM-4 Solution is an optional solution that uses 1SL PDUs to measure performance. This solution uses Dual-Ended Synthetic Loss measurement to measure the Frame Loss Ratio (FLR), Availability, Group Availability and the count of High Loss Intervals (HLI, CHLI).

For the PM-4 Solution using 1SL PDUs, One-way measurements from a Controller MEP to a Sink MEP (in the Forward direction) are taken. Dual-Ended PM Sessions can be configured so that one runs from MEP i to MEP j and another runs from MEP j to MEP i.

1SL frames can be unicast or multicast. In multipoint MEGs, use of multicast 1SL frames can help to simplify PM Session configuration and reduce SOAM traffic. This is further described in Appendix H – Notes on Dual-Ended PM Functions (Informative).

Like any synthetic measurement approach, a PM Session using the Dual-Ended Synthetic Loss PM Function needs to generate enough SOAM PM Frames to be statistically valid. Appendix D - Statistical Considerations for Synthetic Loss Measurement (Informative) contains further information with respect to FLR, and Appendix J – Statistical Considerations for Availability contains further information with respect to Availability.

All Synthetic SOAM PM Frames need to be similar to the Qualified Service Frames carried by the EVC or OVC. 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 Qualified Service Frames are treated. In addition, it is important that Synthetic SOAM PM Frames be inserted irrespective of 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 10.3 [12] that can be calculated for each ordered EI pair in the set S using the PM-4 Solution:

- One-way Frame Loss Ratio Performance (MEF 10.3 [12] section 8.8.3)
- One-way Availability Performance (MEF 10.3 [12] section 8.8.4)
- One-way Resiliency Performance (HLI, CHLI) (MEF 10.3 [12] section 8.8.5)

One-way Group Availability Performance (MEF 10.3 [12] section 8.8.6) for a group of EI pairs can also be calculated.

10. Common Requirements

This section provides requirements that are applicable to all of the PM Solutions that follow in sections 11 (PM-1), 12 (PM-2), 13 (PM-3), and 14 (PM-4). 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.

10.1 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 measurements will be taken for a PM Session.

10.1.1 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 SOAM PM CoS ID and a given Responder/Sink MEP within the same MEG. A PM Session can be used for either Loss Measurement or Delay Measurement, depending on the PM Function applied.

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 end points are the Controller MEP and a Responder/Sink MEP.
- The SOAM PM CoS ID for the PM Session is chosen such that the performance of SOAM PM Frames is representative of the performance of the Qualified Service Frames being monitored. See section 8.4 for further details.
- The PM Function is any of the functions described in section 9 (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 Frame transmission frequency (the time between SOAM PM Frame 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).

Elastic Ethernet services (as defined in MEF 47 [21]) allow subscribers to change attributes of their services. They are able to change the CE-VLAN map, the CoS Frame Set, and the CIR/EIR. In the event that service attributes such as CE-VLAN ID or CoS Frame Set are changed, and there is an active PM Session on that service, the active PM Session is stopped and deleted and a new PM Session is created with the appropriate attributes.

10.1.2 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 MEPs for the given SOAM PM CoS ID. 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.

10.1.3 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 11, 12, 13, and 14.

- [R1] A SOAM PM Implementation **MUST** support multiple concurrent PM Sessions to the same destination, regardless of the setting of other parameters for the PM Sessions, and regardless of whether the PM Sessions use the same or different PM Functions.

Multiple PM Sessions using the same PM Function could be used, for example, to monitor different SOAM PM CoS IDs (and hence measure performance for different CoS FSs), different frame lengths, or to support both Proactive and On-Demand sessions. Multiple PM Sessions using different PM Functions could be used, for example, to monitor both loss- and delay-related performance metrics concurrently.

- [R2] A SOAM PM Implementation **MUST** provide a way to indicate to the EMS/NMS whether a PM Session is Proactive or On-Demand.

10.1.4 Delete

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

- [R3] A SOAM PM Implementation **MUST** support the capability to delete a PM Session.
- [R4] After a PM Session is deleted, further SOAM PM Frames relating to the session **MUST NOT** be sent.
- [R5] After a PM Session is deleted, further measurements associated with the deleted PM Session **MUST NOT** be made.
- [O1] Before the data from a deleted PM Session is lost, a SOAM PM Implementation **MAY** issue a report (similar to the report that would happen when Stop Time is reached).
- [R6] After a PM Session is deleted, all the stored measurement data relating to the deleted PM Session **MUST** be deleted.

Note: a PM Session may be deleted at any point in its lifecycle, including before it has started.

10.1.5 Start and Stop

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

- [R7] For Proactive PM Sessions, the Start Time **MUST** be “immediate”.
- [R8] 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 an offset from the current time, or as a fixed absolute time in the future.

An offset from the current time (i.e., a "relative" time) could be specified as a given number of hours, minutes, and seconds from the current time. A fixed absolute time could be specified as a given UTC date and time.

- [D1] For On-Demand PM Sessions, the default Start Time **SHOULD** be “immediate”.

The following requirements apply to stopping of a PM Session.

- [R9] For Proactive PM Sessions, the Stop Time **MUST** be “forever”.
- [R10] For On-demand PM Sessions, a SOAM PM Implementation **MUST** support a configurable Stop Time per PM Session. The Stop Time can be specified as “forever” or as an offset from the current time.

An offset from the current time (i.e., a “relative” time) could be specified as a given number of hours, minutes, and seconds from the Start Time.

- [R11] For On-demand PM Sessions, if the Stop Time is specified as an offset from the Start Time, then the Stop Time **MUST** be equal to or greater than the Message Period of the PM Session.
- [D2] For On-demand PM Sessions, the default Stop Time **SHOULD** be "forever".
- [R12] A SOAM PM Implementation **MUST** support stopping a PM Session by management action, prior to the Stop Time being reached.
- [R13] After a PM Session is stopped, whether by reaching the scheduled Stop Time or by other means, further SOAM PM Frames relating to the session **MUST NOT** be sent.
- [R14] After a PM Session is stopped, the stored measurements relating to the PM Session **MUST NOT** be deleted.

Note: a PM Session cannot be restarted once it has been stopped, as this would make it difficult to interpret the results. Instead, a new PM Session can be started.

10.1.6 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 10.2).

- [R15] A SOAM PM Implementation **MUST** support a configurable duration for Measurement Intervals.
- [R16] A SOAM PM Implementation **MUST** support a Measurement Interval with duration of 15 minutes for Proactive PM Sessions.
- [R17] A SOAM PM Implementation **MUST** support Measurement Intervals with a duration of between 1 minute and 15 minutes (in 1 minute increments) for On-Demand PM Sessions.
- [D3] The default Measurement Interval duration for On-Demand PM Sessions **SHOULD** be 5 minutes.

10.1.7 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 Frames are sent by the Controller MEP relating to the PM Session, and no measurements are initiated. Note that Responder MEPs may send SOAM PDUs during the time between two Measurement Intervals in response to SOAM PDUs that may have previously been sent by the Controller MEP.

- [R18] A SOAM PM Implementation **MUST** support a configurable Repetition Time per PM Session. The Repetition Time can be specified as “none” or as a repeating time interval.

A repeating time interval (i.e., a relative time) could be specified as every given number of hours, minutes, and seconds from the Start Time.

- [D4] The default 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 Frames relating to the PM Session **MUST NOT** be sent by the Controller MEP.

10.1.8 Alignment of Measurement Intervals

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

- [D5] A SOAM PM Implementation **SHOULD** by default align the start of each Measurement Interval, other than the first 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”).
- [D6] A SOAM PM Implementation **SHOULD** by default align the start of each Measurement Interval, other than the first 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”).

When Measurement Intervals are aligned with the ToD clock, the Start Time of a PM Session might not correspond with the alignment boundary. In this case, the first Measurement Interval could be truncated. Further examples can be found in Appendix B – Life Cycle Terminology (Informative).

- [D7] A SOAM PM Implementation **SHOULD** allow for no alignment to the time-of-day clock.
- [D8] A SOAM PM Implementation **SHOULD** support a configurable (in minutes) offset from ToD time for alignment of the start of Measurement Intervals other than the first Measurement Interval.

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

10.1.9 Summary of Time Parameters

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

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

10.2 Storage

The requirements of this section apply to storage of performance measurement results taken during Measurement Intervals, using counters or Measurement Bins (for some delay-related parameters). Performance measurements are stored separately for each Measurement Interval. A Measurement Bin is a counter, and records the number of performance measurements falling within a specified range. Figure 10 – Example of Measurement Intervals and Bins (below) is an example that illustrates the relationship between Measurement Intervals and Measurement Bins:

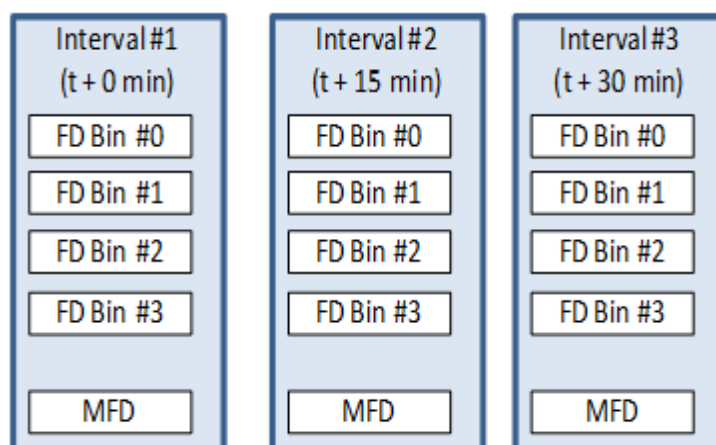


Figure 10 – Example of Measurement Intervals and Bins

Figure 11 shows an example of a MEP running a Single-Ended Synthetic Loss PM Function using SLM/SLR. It measures loss, separately for each direction.

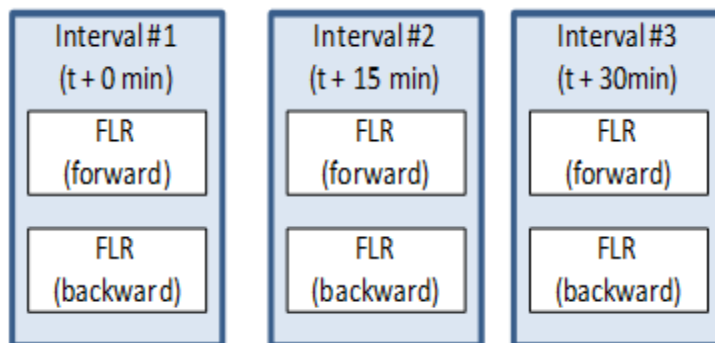


Figure 11 – Example of FLR Measurements

10.2.1 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 12, 13, 13 and 14.

- [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.
- [D9] 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.
- [D10] 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 value of the local time-of-day clock in UTC at the scheduled start of the Measurement Interval.
- [R23] A SOAM PM Implementation **MUST** record the value of the local time-of-day clock in UTC at the scheduled 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.
- [D11] 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] For Single-Ended measurement, a SOAM PM response frame received by the Controller MEP more than 5 seconds after the end of the Measurement Interval in which the corresponding SOAM PM request frame was transmitted **MUST** be discarded and considered lost.

Note: For Dual-Ended measurements, in some cases the Sink MEP cannot determine reliably the Measurement Interval in which a received frame was initiated by the Controller MEP.

10.2.2 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 MEP for a given PM Session during a Measurement Interval.

When using Single-Ended Delay Measurement, FD, IFDV and FDR can be monitored using Two-way measurements, and/or using One-way measurements in the Forward and/or Backward direction. When using Dual-Ended Delay Measurement, FD, IFDV and FDR can be monitored using One-way measurements in the Forward direction only. The particular FD measurements supported in a SOAM PM Implementation depend on the PM Solutions supported and on NE capabilities (e.g., time-of-day 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.
- [D12] For a SOAM PM Implementation, the default number of FD Measurement Bins per Measurement Interval **SHOULD** be 2.
- [R28] A SOAM PM Implementation **MUST** support at least 2 FD Measurement Bins per Measurement Interval.
- [D13] A SOAM PM Implementation **SHOULD** support at least 10 FD Measurement Bins per Measurement Interval.

The following requirements apply to each IFDV or FDR measurement supported in a SOAM PM Implementation.

- [R29] A SOAM PM Implementation **MUST** support a configurable number of IFDV Measurement Bins per Measurement Interval.
- [D14] For a SOAM PM Implementation, the default number of IFDV Measurement Bins per Measurement Interval supported **SHOULD** be 2.

- [R30] A SOAM PM Implementation **MUST** support at least 2 IFDV Measurement Bins per Measurement Interval.
- [D15] A SOAM PM Implementation **SHOULD** support at least 10 IFDV Measurement Bins per Measurement Interval.
- [R31] A SOAM PM Implementation **MUST** support a configurable number of FDR Measurement Bins per Measurement Interval.
- [D16] For a SOAM PM Implementation, the default number of FDR Measurement Bins per Measurement Interval supported **SHOULD** be 2.
- [R32] A SOAM PM Implementation **MUST** support at least 2 FDR Measurement Bins per Measurement Interval.
- [D17] A SOAM PM Implementation **SHOULD** support at least 10 FDR Measurement Bins per Measurement Interval.

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 bound. Because the bins are contiguous, it is only necessary to configure the lower bound of each bin. Furthermore, the lowest bin is assumed to always have a lower bound of 0, and the highest bin is assumed to have an upper bound of ∞ .

Note: All values for IFDV, FDR and Two-way FD are positive by definition. Values for One-way FD can be negative if there is no ToD synchronization, and such measurements would not match any Measurement Bin as defined above; however, in this case taking One-way FD measurements is not recommended except for the purpose of finding the minimum FD for normalization of FDR, and finding the minimum FD does not require Measurement Bins.

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 lower bound for all but the first Measurement Bin.
- [R34] The lower bound for each Measurement Bin **MUST** be larger than the lower bound of the preceding Measurement Bin.
- [R35] The unit for a lower bound **MUST** be in microseconds (μs).
- [R36] The lower bound of the first Measurement Bin **MUST** be fixed to $0\mu\text{s}$.

- [R37] Measured performance values that are greater than or equal to the lower bound of a given bin and strictly less than the lower bound of the next bin (if any), **MUST** be counted in that, and only that bin.
- [D18] The default lower bound for a Measurement Bin **SHOULD** be an increment of 5000 μ s larger than the lower bound of the preceding Measurement Bin.

For example, four Measurement Bins gives the following:

Bin	Lower Bound	Range
bin 0	0 μ s	$0 \mu\text{s} \leq \text{measurement} < 5,000 \mu\text{s}$
bin 1	5,000 μ s	$5,000 \mu\text{s} \leq \text{measurement} < 10,000 \mu\text{s}$
bin 2	10,000 μ s	$10,000 \mu\text{s} \leq \text{measurement} < 15,000 \mu\text{s}$
bin 3	15,000 μ s	$15,000 \mu\text{s} \leq \text{measurement} < \infty$

Table 5 – Example Measurement Bin Configuration

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

10.2.3 Volatility

The following requirement applies to the volatility of storage for Measurement Interval data.

- [D19] 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.

10.2.4 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 between the Controller MEP and the Responder or Sink MEP.
- 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 or OVC.
- Maintenance Interval (see MEF 10.3 [12])

[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.

Note: Loss of measurement frames does not affect whether the Suspect Flag is set.

[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.

10.2.5 Measurement Behavior During Unavailable Time and Maintenance Intervals

Measurements of Performance do not apply during Maintenance Intervals. By definition (see MEF 10.3 [12]), 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 Available Time. This means that if Availability is measured for a given SOAM PM CoS ID on an ME, during Unavailable Time for that SOAM PM CoS ID, measurements of performance metrics for that same SOAM PM CoS ID (other than Availability) are to be excluded, so such impairments are not double counted. Availability is evaluated per Maintenance Entity (ME), because a single NE does not necessarily have visibility of all MEs within the MEG.

However, whether a Maintenance Entity is in Available Time or Unavailable Time for a given SOAM PM CoS ID cannot be determined until a period of $n \Delta t$ (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 8.8.4 of MEF 10.3

[12]⁴. 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.

Correcting the FLR performance metric to account for Unavailable Time is of primary importance. Correcting for delay-related performance metrics is secondary.

- [R43] For all Δt intervals that are determined to be Available for a given SOAM PM Cos ID on a given ME, a SOAM PM Implementation **MUST** include measurements for those Δt intervals in all performance metrics, in any PM Session for the same SOAM PM CoS ID and ME.
- [R44] For all Δt intervals that are determined to be Unavailable for a given SOAM PM Cos ID on a given ME, a SOAM PM Implementation **MUST** exclude measurements for those Δt intervals from performance metrics for FLR, in any PM Session for the same SOAM PM CoS ID and ME.
- [D20] For all Δt intervals that are determined to be Unavailable for a given SOAM PM Cos ID on a given ME, a SOAM PM Implementation **SHOULD** exclude measurements for those Δt intervals from performance metrics other than FLR and Availability, in any PM Session for the same SOAM PM CoS ID and ME.
- [R45] If a SOAM PM Implementation does not conform to [D20], then for all Δt intervals that are determined to be Unavailable for a given SOAM PM CoS ID on a given ME, it **MUST** include measurements for those Δt intervals in performance metrics other than FLR and Availability, in any PM Session for the same SOAM PM CoS ID and ME.

Measurements are always collected during Available Time for all performance metrics, and measurements are always excluded during Unavailable Time for FLR. Excluding measurements during Unavailable time for other performance metrics is also recommended; however, if measurements are not excluded then all measurements (except FLR) are included. In this last case, the SOAM PM Implementation need not distinguish between Available Time and Unavailable Time when taking measurements for metrics other than FLR, since all the measurements are included either way. This eliminates possible issues with aligning MIs between PM Sessions for FLR and for other metrics.

When correcting for Unavailable Time, the correction also applies when transitioning between Available and Unavailable. That is, if the current state for a given SOAM PM CoS ID on an ME is Available, and n Δt intervals occur in which the Availability threshold is crossed, then the state is changed to Unavailable. In this case, those n Δt intervals are determined to be Unavailable and the measurements for them are excluded from the performance metrics for that SOAM PM CoS

⁴ 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 8.8.4 of MEF 10.3 [12] for the authoritative discussion.

ID and ME. Similarly, if the current state is Unavailable, and $n \Delta t$ intervals occur in which the Availability threshold is not crossed, then the state is changed to Available. In this case, those $n \Delta t$ intervals are determined to be Available and the measurements for them are included in the performance metrics.

A direct consequence of the above 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 performance 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 12 shows one example of the impact of an Availability state change on the Measurement Interval counters. In this example, $n=10$, $\Delta t=1s$ and the frame interval is 100ms. The figure shows counts of SLMs sent and received in the forwards and backwards directions (indicated by “-f” and “-b” respectively) in the last 10 Δt intervals, as well as the counts of SLMs sent and received and the counts of Δt intervals evaluated as Available or Unavailable in the current Measurement Interval. Note again that storing separate counters for the last $n \Delta t$ intervals, as shown in the figure, is just an example, and other implementations are possible.

In this example, the last 10 consecutive Δt intervals (between 30s and 40s after the start of the Measurement Interval) experience sufficient frame loss that the Availability state changes from Available to Unavailable in both the forward and backward directions. The count of frames transmitted and received is decremented by the number of frames in last 10 Δt intervals (for instance by 100 in the case of SLMs transmitted in the forward direction). In addition, the count of Unavailable Δt intervals increases from 0 to 10 and the count of Available Δt intervals decreases from 40 to 30.

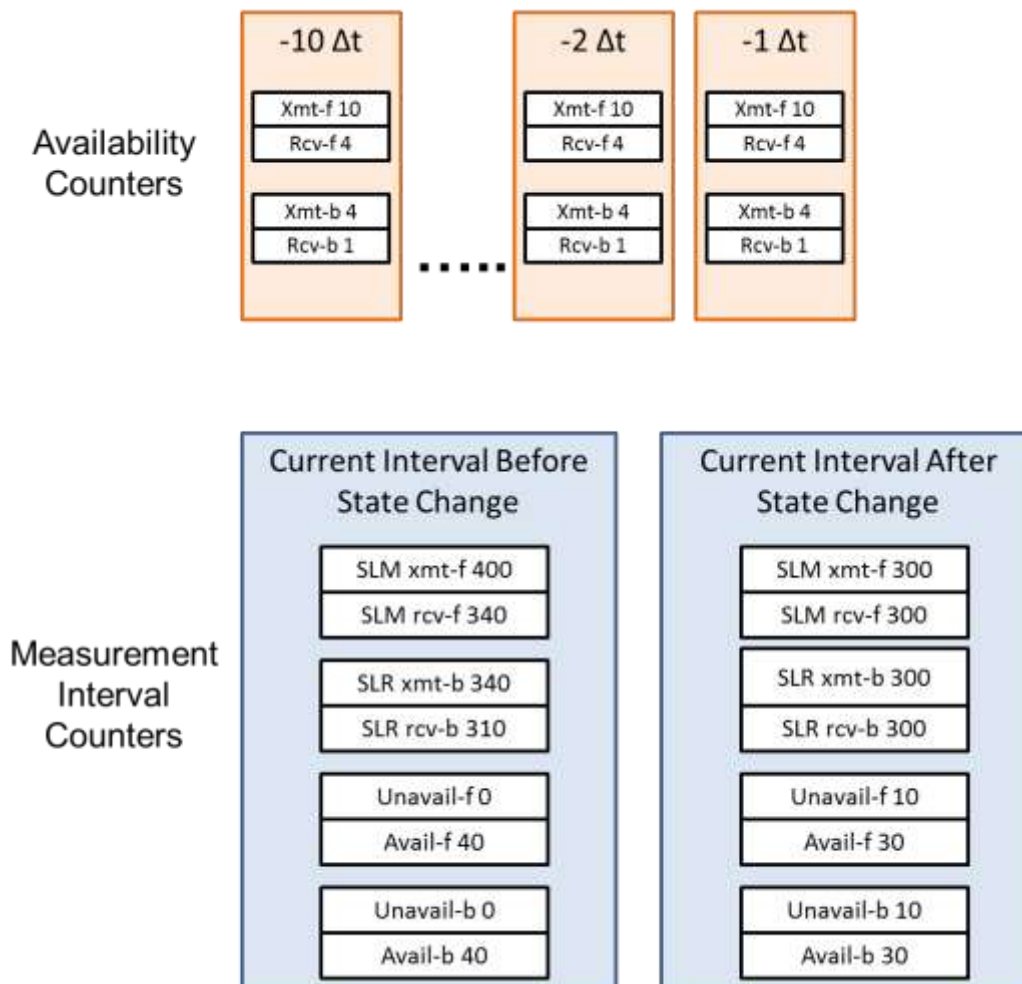


Figure 12 – Example of adjusting Measurement Counters during transition from Available to Unavailable

Figure 13 shows an example of the change from Unavailable to Available state. In this case, the Unavailable state lasted for 40 seconds. After 10 consecutive Δt intervals with the number of lost frames below the Availability threshold, the Availability state transitions from Unavailable to Available. When this occurs the transmitted and received frame counters are incremented by the number of frames sent or received in the 10 consecutive Δt intervals (for instance by 100 in the case of SLMs transmitted in the forward direction). In addition, the count of Unavailable Δt intervals decreases from 50 to 40 and the count of Available Δt intervals decreases from 30 to 40.

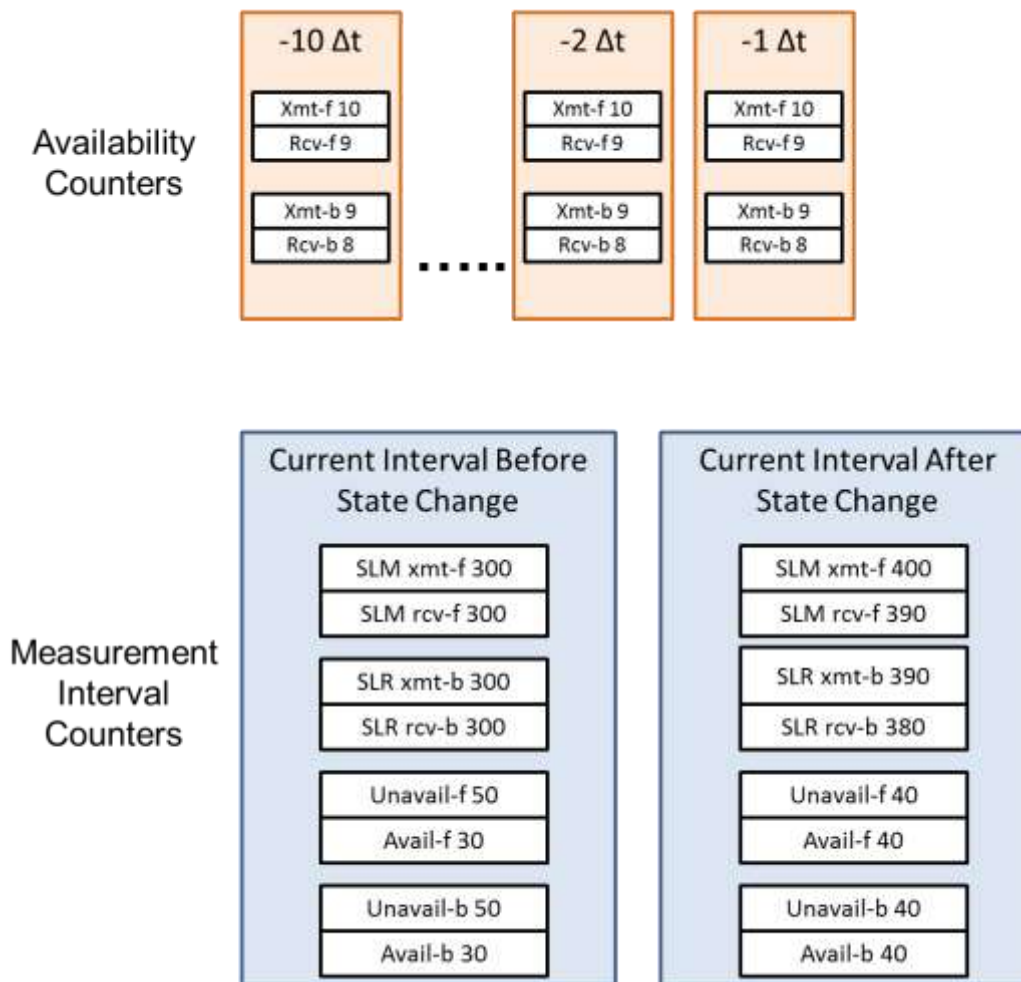


Figure 13 – Example of adjusting Measurement Counters during transition from Unavailable to Available

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.

10.3 OAM Domains

The following requirements provide information about OAM Domains.

- [R46] A SOAM PM Implementation **MUST** support EVC MEG.
- [R47] A SOAM PM Implementation **MUST** support Service Provider MEG.
- [R48] A SOAM PM Implementation **MUST** support Operator MEG.
- [R49] 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.

Note: SOAM PM using the EVC MEG or the Operator MEG may be used to evaluate conformance to an SLS for an EVC or OVC respectively, as defined in MEF 10.3 [12] and MEF 26.1 [18]. SOAM PM using these or other MEGs can also be used for informational purposes, as described in section 6.3.

10.4 MEP Placement

Section 8.2.1 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.

- [R50] On a UNI-N, the MEP **MUST** be placed between the Ethernet Subscriber Conditioning Function (ESCF) and the Ethernet ECS Adaptation Function (EEAF).
- [R51] On an ENNI, the MEP **MUST** be placed between the Ethernet Provider Conditioning Function (EPCF) and the Ethernet EC Interworking Function (EEIF).

10.5 Threshold Crossing Alerts

Performance thresholds, and corresponding Threshold Crossing Alerts (TCAs), can be configured for certain performance metrics, and used to detect when service performance is degraded beyond a given pre-configured level. Thresholds are always specific to a particular performance metric and a particular PM Session. When the measured performance in a Measurement Interval for that session reaches or exceeds the configured threshold level, a TCA can be generated and sent to an Element Management System (EMS) or Network Management System (NMS).

In normal operation, performance data is collected from an NE by the EMS/NMS either periodically (e.g. once an hour) or on-demand. TCAs can be used as warning notifications to the EMS/NMS of possible service degradation, thus allowing more timely action to further investigate or address the problem. For example, if the maximum One-way FD threshold was set to 10ms, and a One-way FD value was measured at more than 10ms, a TCA would be generated.

- [O4] A SOAM PM Implementation **MAY** support Threshold Crossing Alert functionality as described in section 10.5.1, 10.5.2 and 10.5.3.

The requirements in the following subsections only apply if TCA functionality is supported.

10.5.1 TCA Reporting

Thresholds and associated TCAs are specific to a particular performance metric in a given PM Session. There are two types of TCA reporting: stateless and stateful. With stateless reporting, a TCA is generated in each Measurement Interval in which the threshold is crossed. With stateful reporting, a SET TCA is generated in the first Measurement Interval in which the threshold is

crossed, and a CLEAR TCA is subsequently generated at the end of the first Measurement Interval in which the threshold is not crossed.

Note: In ITU-T G.7710 [6] terminology, stateless TCA reporting corresponds to a transient condition, and stateful TCA reporting corresponds to a standing condition.

Regardless of the type of TCA reporting (stateless or stateful), it is not desirable to generate more than one TCA for a given threshold during each Measurement Interval, as to do otherwise could cause unnecessary load both on the NE and on the EMS/NMS receiving the TCAs.

Thresholds and TCAs are only defined for certain performance metrics, as described in section 10.5.2. Note that all of these performance metrics have the property that the value cannot decrease during a given Measurement Interval.

The process that takes a given threshold configuration for a given performance metric in a given PM Session and generates corresponding TCAs is termed a TCA Function. Multiple TCA Functions with different threshold values can be configured for the same PM Session and performance metric, so that TCAs can be generated for different degrees of service degradation. Where multiple TCA Functions are configured, corresponding TCAs are generated independently for each TCA Function.

10.5.1.1 Stateless TCA Reporting

The stateless TCA reporting treats each Measurement Interval separately. When using stateless TCA reporting, each TCA Function has a single configured threshold. As soon as the threshold is reached or crossed in a Measurement Interval for a given performance metric, a TCA is generated.

The following figure illustrates the behavior of stateless TCA reporting.

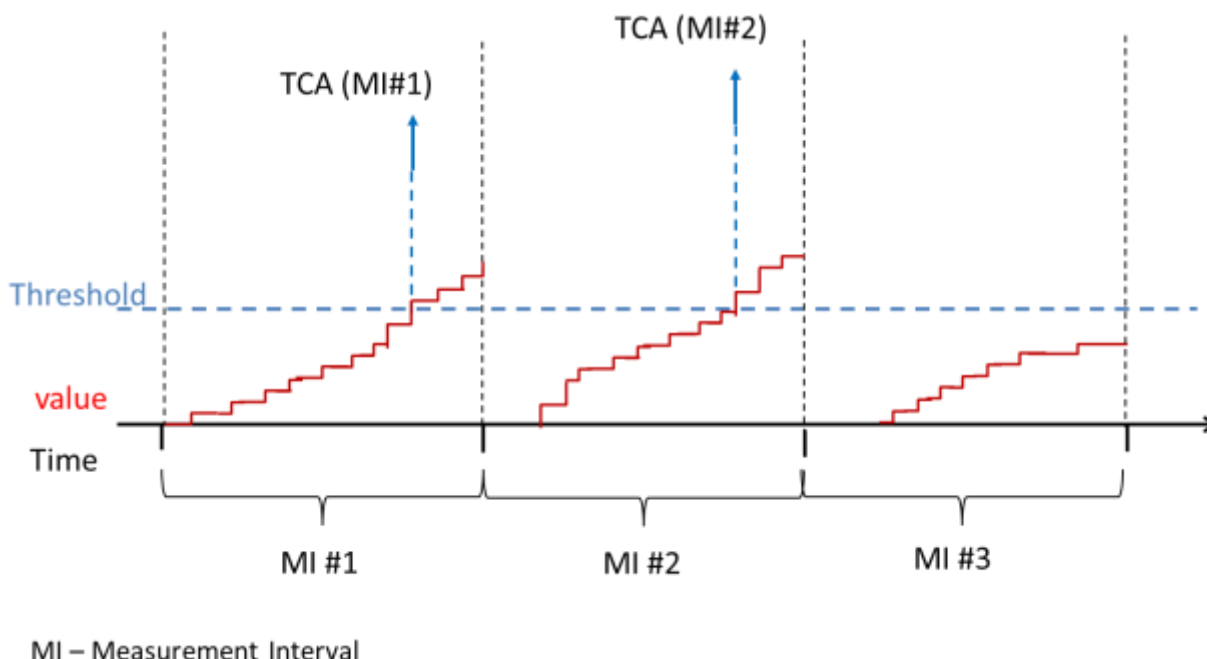


Figure 14 – Stateless TCA Reporting Example

As shown in the example in Figure 14, in MI #1, the measured performance value (e.g., Maximum Frame Delay) crosses the corresponding threshold. Therefore a TCA is generated for MI #1. In MI #2, this threshold is crossed again. Another TCA is generated for MI #2. In MI #3, the measured performance value doesn't reach the threshold. There is no TCA for that performance metric for MI #3.

10.5.1.2 Stateful TCA Reporting

Stateful TCA reporting is another option for how TCAs are generated, that can reduce the total number of TCAs. The intent is to provide a notification when a degradation is first encountered, followed by another when the problem is resolved. This contrasts with stateless TCA reporting, in which TCAs are generated continuously for as long as the degradation lasts.

When using stateful TCA reporting, each TCA Function has two configured thresholds: a SET threshold and a CLEAR threshold. These may be the same, or the CLEAR threshold may be lower than the SET threshold. The TCA Function also has an internal state, which may be 'set' or 'clear'.

The TCA Function begins in the 'clear' state. A SET TCA is generated in the first Measurement Interval as soon as the SET threshold is reached or exceeded. The TCA Function is then considered to be in a 'set' state, and no further SET TCAs are generated in this state. In each subsequent Measurement Interval in which the CLEAR threshold is reached or exceeded, no TCA is generated. At the end of the first Measurement Interval in which the CLEAR threshold is not reached or exceeded, a CLEAR TCA is generated, and the TCA Function returns to the 'clear' state. Thus, each SET TCA is followed by a single CLEAR TCA.

The following figure shows an example of stateful TCA reporting. In this example, the CLEAR threshold is equal to the SET threshold.

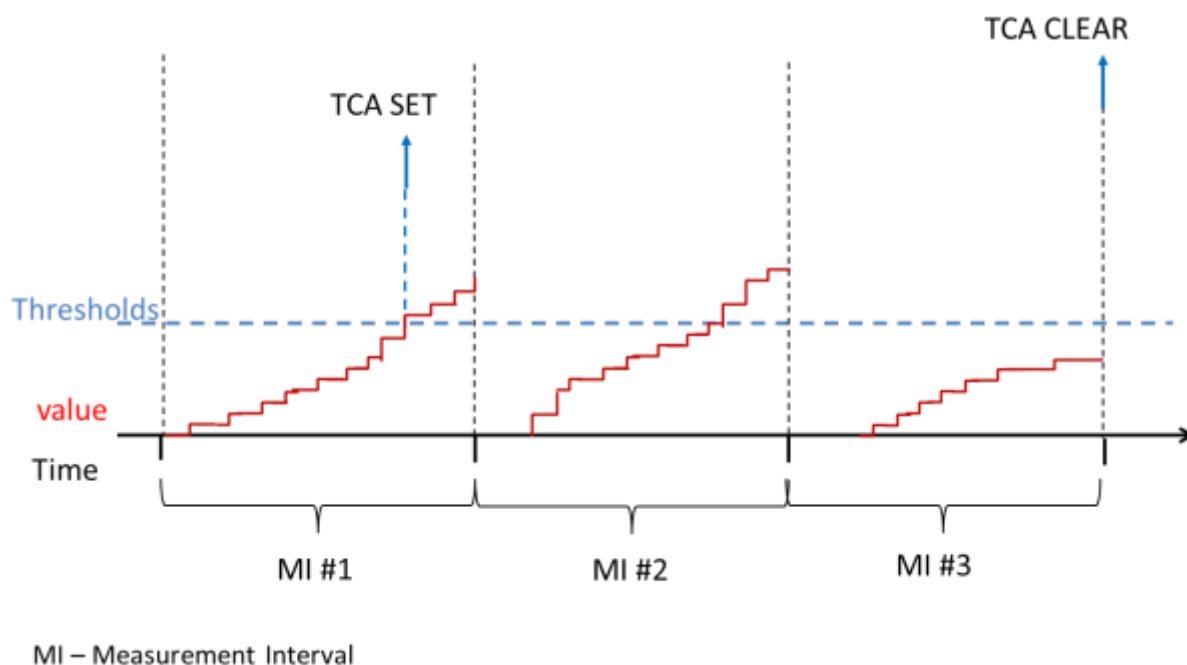


Figure 15 – Stateful TCA Reporting Example

In the example, a SET TCA is generated in MI #1. In MI #2, the threshold is crossed again but no SET TCA is generated because a SET TCA had been generated in MI #1. MI #3 is the first subsequent Measurement Interval that the measured performance value is below the CLEAR threshold. A CLEAR TCA is generated at the end of MI #3.

10.5.2 SOAM PM Thresholds for TCA

TCAs are useful for some performance metrics but may not be meaningful for others. This section describes which performance metrics are required and how to support TCAs.

For performance metrics that use Measurement Bins, thresholds are defined in terms of an Upper Bin Count (UBC). The Upper Bin Count of bin k is the total of the counts for bins k and above, i.e. $UBC(k) = \text{count of bin } (k) + \text{count of bin } (k+1) + \dots + \text{count of bin } (n)$, where n is the last bin.

To configure a threshold, both the bin number, k , and the total count, N , need to be specified - this is represented as (N, k) . A threshold (N, k) is considered to have been crossed when $UBC(k) \geq N$. Figure 16 illustrates how a threshold is configured using bins.

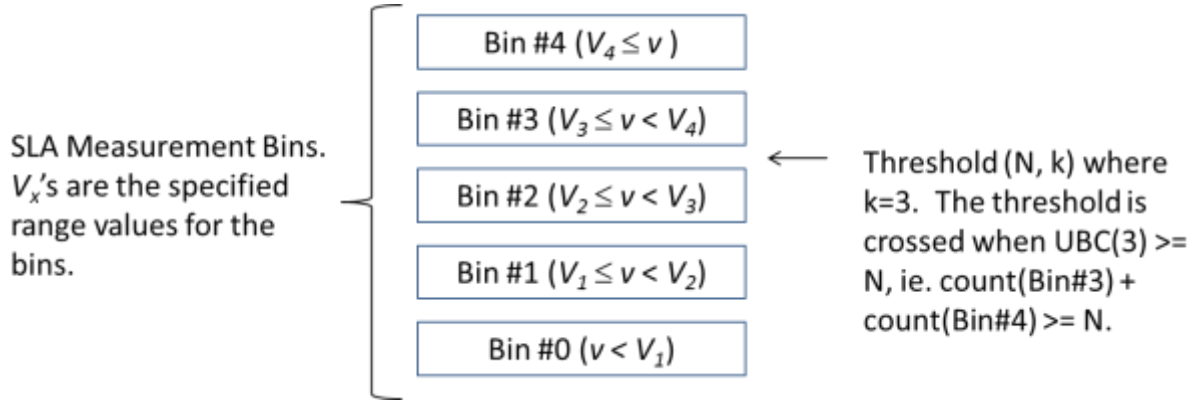


Figure 16 – Upper Bin Count for Threshold Crossing

Most performance metrics, such as Frame Delay, are defined only during Available Time. TCAs are, by definition, alerts – that is, they alert the user to something unexpected, but don't necessarily indicate that a fault has occurred. Therefore, it is acceptable to keep generating them even if the service becomes Unavailable; in particular, it is not necessary to wait for $n \Delta t$ after the threshold crossing is detected before generating a TCA in case the service becomes Unavailable (where Δt is a small time interval, and n is the number of consecutive ' Δt 's required to have high loss before the service is declared Unavailable – see MEF 10.3 [12]). The receivers of the TCAs, e.g., EMS/NMS, may use the combined information of TCAs and Availability state change notifications to decide what actions to take.

The following table lists the applicable performance metrics that support TCAs. In each case, both One-way, and where applicable, Two-way performance metrics can be used. The table describes in each case the parameters that must be configured for the threshold, and the definition of when the threshold is crossed. For stateful TCA reporting, the "SET" thresholds and "CLEAR" thresholds are defined in the same way (although the configured values may be different).

Performance Metric	Configured Threshold	Threshold Crossing Detection	Notes
One-way FD in the Forward direction	Forward One-way (N_{FD}, k)	$UBC(k) \geq$ Forward One-way N_{FD}	Using Measurement Bins. Requires ToD sync
One-way Maximum FD in the Forward direction	Forward One-way V_{maxFD}	Max FD \geq Forward One-way V_{maxFD}	Requires ToD sync
One-way FDR in the Forward direction	Forward One-way (N_{FDR}, k)	$UBC(k) \geq$ Forward One-way N_{FDR}	Using Measurement Bins
One-way Maximum FDR in the Forward direction	Forward One-way V_{maxFDR}	Max FDR \geq Forward One-way V_{maxFDR}	
One-way IFDV in the Forward direction	Forward One-way (N_{IFDV}, k)	$UBC(k) \geq$ Forward One-way N_{IFDV}	Using Measurement Bins

Performance Metric	Configured Threshold	Threshold Crossing Detection	Notes
One-way Maximum IFDV in the Forward direction	Forward One-way $V_{\max\text{IFDV}}$	Max IFDV \geq Forward One-way $V_{\max\text{IFDV}}$	
One-way HLI in the Forward direction	Forward One-way N_{HLI}	HLI count \geq Forward One-way N_{HLI}	
One-way CHLI in the Forward direction	Forward One-way N_{CHLI}	CHLI count \geq Forward One-way N_{CHLI}	
One-way FD in the Backward direction	Backward One-way (N_{FD}, k)	UBC(k) \geq Backward One-way N_{FD}	Using Measurement Bins. Requires ToD sync
One-way Maximum FD in the Backward direction	Backward One-way $V_{\max\text{FD}}$	Max FD \geq Backward One-way $V_{\max\text{FD}}$	Requires ToD sync
One-way FDR in the Backward direction	Backward One-way (N_{FDR}, k)	UBC(k) \geq Backward One-way N_{FDR}	Using Measurement Bins
One-way Maximum FDR in the Backward direction	Backward One-way $V_{\max\text{FDR}}$	Max FDR \geq Backward One-way $V_{\max\text{FDR}}$	
One-way IFDV in the Backward direction	Backward One-way (N_{IFDV}, k)	UBC(k) \geq Backward One-way N_{IFDV}	Using Measurement Bins
One-way Maximum IFDV in the Backward direction	Backward One-way $V_{\max\text{IFDV}}$	Max IFDV \geq Backward One-way $V_{\max\text{IFDV}}$	
One-way HLI in the Backward direction	Backward One-way N_{HLI}	HLI count \geq Backward One-way N_{HLI}	
One-way CHLI in the Backward direction	Backward One-way N_{CHLI}	CHLI count \geq Backward One-way N_{CHLI}	
Two-way FD	Two-way (N_{FD}, k)	UBC(k) \geq Two-way N_{FD}	Using Measurement Bins.
Two-way Maximum FD	Two-way $V_{\max\text{FD}}$	Max FD \geq Two-way $V_{\max\text{FD}}$	

Table 6 – SOAM Performance Metrics TCA

Note that not all performance metrics are listed in Table 6. They are either not suitable or not necessary. For example:

- MFD – MFD is a performance metric measuring an average and thus a poor metric for immediate attention, compared to FD, FDR and IFDV.

- FLR – FLR is a performance metric for long time period and thus not suitable for immediate action, compared to HLI and CHLI.
- Availability and Group Availability – Since Availability state transition reporting is required ([R83] and [CR63]), having a TCA would be redundant.

If TCA functionality is supported, the following requirements are applicable for a SOAM PM Implementation:

[CR1]< [O4] A SOAM PM Implementation **MUST** support per performance metric, per PM Session configuration of TCA Functions and associated thresholds, using the parameters described in Table 6, for the following performance metrics:

- One-way FDR in the Forward direction
- One-way maximum FDR in the Forward direction
- One-way IFDV in the Forward direction
- One-way maximum IFDV in the Forward direction
- One-way HLI in the Forward direction
- One-way CHLI in the Forward direction
- One-way FDR in the Backward direction
- One-way maximum FDR in the Backward direction
- One-way IFDV in the Backward direction
- One-way maximum IFDV in the Backward direction
- One-way HLI in the Backward direction
- One-way CHLI in the Backward direction
- Two-way FD
- Two-way maximum FD

[CR2]< [O4] If time-of-day synchronization is supported, a SOAM PM Implementation **MUST** support per performance metric, per PM Session configuration of TCA Functions and associated thresholds, using the parameters described in Table 6, for the following performance metrics:

- One-way FD in the Forward direction

- One-way maximum FD in the Forward direction
- One-way FD in the Backward direction
- One-way maximum FD in the Backward direction

- [CR3]< [O4] A SOAM PM Implementation **MUST** support stateless TCA reporting.
- [CD1]< [O4] A SOAM PM Implementation **SHOULD** support stateful TCA reporting.
- [CR4]< [O4], [CD1] If a SOAM PM Implementation supports stateful TCA reporting, it **MUST** support a configurable parameter per TCA Function to indicate whether the TCA Function uses stateful or stateless TCA reporting.
- [CR5]< [O4] A SOAM PM implementation **MUST** support a single configurable parameter for the threshold value for each TCA Function that uses stateless TCA reporting.
- [CR6]< [O4], [CD1] If a SOAM PM Implementation supports stateful TCA reporting, it **MUST** support the CLEAR threshold being equal to the SET threshold..
- [CO1]< [O4], [CD1] If a SOAM PM Implementation supports stateful TCA reporting, it **MAY** support the CLEAR threshold being different to the SET threshold.

For thresholds defined using bins, a CLEAR threshold (N_c, k_c) is defined to be less than or equal to a SET threshold (N_s, k_s) if $k_c = k_s$ and $N_c \leq N_s$.

- [CR7]< [O4], [CD1], [CO1] If a SOAM PM Implementation supports stateful TCA reporting with different SET and CLEAR thresholds, the CLEAR threshold **MUST** be less than or equal to the SET threshold
- [CR8]< [O4], [CD1] If a SOAM PM Implementation supports stateful TCA reporting, it **MUST** support a configurable parameter for the SET threshold for each TCA Function that uses stateful TCA reporting.
- [CR9]< [O4], [CD1], [CO1] If a SOAM PM Implementation supports stateful TCA reporting with different SET and CLEAR thresholds, it **MUST** support a configurable parameter for the CLEAR threshold for each TCA Function that uses stateful TCA reporting.

If different SET and CLEAR thresholds are not used, the value configured for the SET threshold is also used for the CLEAR threshold.

- [CR10]< [O4] If a TCA Function is configured to use stateless TCA reporting, a TCA **MUST** be generated for each Measurement Interval in which the threshold is crossed as defined in Table 6.

- [CD2]< [O4] If a TCA Function is configured to use stateless TCA reporting, the TCA for a given Measurement Interval **SHOULD** be generated as soon as the threshold is crossed.
- [CR11]< [O4] If a TCA Function is configured to use stateless TCA reporting, the TCA for a given Measurement Interval **MUST** be generated within 1 minute of the end of the Measurement Interval.
- [CR12]< [O4], [CD1] If a TCA Function is configured to use stateful TCA reporting, in the 'clear' state a SET TCA **MUST** be generated for a given Measurement Interval if the SET threshold is crossed as defined in Table 6 during that Measurement Interval.
- [CR13]< [O4], [CD1] If a TCA Function is configured to use stateful TCA reporting, in the 'clear' state, if the SET threshold is crossed during a given Measurement Interval, the state **MUST** be changed to 'set' by the end of that Measurement Interval.
- [CD3]< [O4], [CD1] If a TCA Function is configured to use stateful TCA reporting, the SET TCA for a given Measurement Interval **SHOULD** be generated as soon as the SET threshold is crossed.
- [CR14]< [O4], [CD1] If a TCA Function is configured to use stateful TCA reporting, the SET TCA for a given Measurement Interval **MUST** be generated within 1 minute of the end of the Measurement Interval.
- [CR15]< [O4], [CD1] If a TCA Function is configured to use stateful TCA reporting, SET TCAs **MUST NOT** be generated when in the 'set' state.
- [CR16]< [O4], [CD1] If a TCA Function is configured to use stateful TCA reporting, in the 'set' state a CLEAR TCA **MUST** be generated for a given Measurement Interval if the CLEAR threshold is not crossed as defined in Table 6 during that Measurement Interval.
- [CR17]< [O4], [CD1] If a TCA Function is configured to use stateful TCA reporting, in the 'set' state, if the CLEAR threshold is not crossed during a given Measurement Interval, the state **MUST** be changed to 'clear' at the end of that Measurement Interval.
- [CD4]< [O4], [CD1] If a TCA Function is configured to use stateful TCA reporting, the CLEAR TCA for a given Measurement Interval **SHOULD** be generated immediately at the end of the Measurement Interval.
- [CR18]< [O4], [CD1] If a TCA Function is configured to use stateful TCA reporting, the CLEAR TCA for a given Measurement Interval **MUST** be generated within 1 minute of the end of the Measurement Interval.

- [CR19]< [O4], [CD1] If a TCA Function is configured to use stateful TCA reporting, CLEAR TCAs **MUST NOT** be generated when in the 'clear' state.
- [CR20]< [O4] For a given TCA Function applying to a given performance metric and a given PM Session, a SOAM PM Implementation **MUST NOT** generate more than one TCA for each Measurement Interval.
- [CR21]< [O4] A SOAM PM Implementation **MUST** support the configuration of at least one TCA Function for each performance metric listed in Table 6, for each PM Session.

Note: this does not require that a SOAM PM Implementation is able to support configuration of a TCA Function for every performance metric for every PM Session simultaneously.

- [CO2]< [O4] A SOAM PM Implementation **MAY** support the configuration of more than one TCA Function for a performance metric, for each PM Session.

10.5.3 SOAM PM TCA Notification Messages

Table 7 lists the SOAM PM TCA Notification message attributes used when sending a TCA to an EMS/NMS.

Field Name	Field description
Date and Time	Time of the event, in UTC. For stateless TCAs, and stateful SET TCAs, this is the time the threshold was crossed; for stateful CLEAR TCAs, it is the time at the end of the Measurement Interval for which the CLEAR TCA is being generated.
PM Session	Identification of the PM Session for which the TCA Function was configured. The specific parameters needed to uniquely identify a PM Session are implementation-specific.
Measurement Interval	The time, in UTC, at the start of the Measurement Interval for which the TCA was generated.
Performance Metric Name	Performance Metric for which the TCA Function was configured, i.e., one of those listed in Table 6.
Configured Threshold	The configured threshold parameters. For bin-based thresholds, this includes the bin number and the total count, i.e., (N, k).

Field Name	Field description
Measured Performance Metric	Measured value that caused the TCA to be generated. For bin-based thresholds configured as (N, k), this is always equal to N for stateless TCAs and stateful SET TCAs; for stateful CLEAR TCAs, it is the value of UBC(k) at the end of the Measurement Interval. For "maximum" performance metrics, for stateless TCAs and stateful SET TCAs, this is the first value in the Measurement Interval that reaches or exceeds the configured threshold; for stateful CLEAR TCAs it is the maximum value at the end of the Measurement Interval. For HLI and CHLI thresholds, this is always equal to the configured threshold value for stateless TCAs and stateful SET TCAs; for stateful CLEAR TCAs it is the total count at the end of the Measurement Interval.
Suspect Flag	Value of the Suspect Flag for the Measurement Interval for which the TCA was generated. Suspect Flag is true when there is a discontinuity in the performance measurements conducted during the Measurement Interval.
TCA Type	The type of TCA, i.e. one of STATELESS (if stateless TCA reporting was configured for the TCA Function), STATEFUL-SET (if stateful TCA reporting was configured and this is a SET TCA) or STATEFUL-CLEAR (if stateful TCA reporting was configured and this is a CLEAR TCA).
Severity	WARNING (for STATELESS or STATEFUL-SET) or INFO (for STATEFUL-CLEAR)

Table 7 –TCA Notification Message Fields

[CR22]< [O4] A SOAM PM Implementation **MUST** include the fields in the TCA notification messages listed in Table 7

Table 8 shows the correlation between the general alarm and event notification parameters described in ITU-T X.733 and X.734, and the notification attributes considered in this document.

ITU-T X.733, X.734	Consideration for MEF 35
Event time	Date and time
Managed Obj Class	PM Session
Managed Obj Instance	Included in PM Session
Monitored Attribute	Performance Metric Name, Measurement Interval
Threshold Info	Configured Threshold, Measured Performance Metric
<i>No equivalent</i>	Suspect Flag
Event type (service degraded)	TCA Type
Severity	Severity

ITU-T X.733, X.734	Consideration for MEF 35
Probable cause	Not applicable

Table 8 – Comparison of TCA fields in MEF 35 and ITU-T X.73x

11. PM-1 Requirements

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

- [R52] A SOAM PM Implementation **MUST** support the Single-Ended Delay Function as described in section 11.1.
- [R53] A SOAM PM Implementation **MUST** support the Single-Ended Synthetic Loss Function as described in section 11.2.

Section 11.1 lists the requirements for performing Frame Delay, Inter-Frame Delay Variation, Mean Frame Delay and Frame Delay Range measurements using the DMM/DMR PDUs. Section 11.2 lists the requirements for performing Frame Loss Ratio, Resiliency (HLI and CHLI), and Availability measurements.

Both the Single-Ended Delay and the Single-Ended Synthetic Loss functions can be configured for multiple SOAM PM CoS IDs per pair of MEPs. Each unique pair of MEPs and SOAM PM CoS ID being measured results in one or more distinct PM Sessions. The functions support both point-to-point and multipoint configurations.

On multipoint MEGs 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 SOAM PM CoS ID. If the measurements are being used to evaluate conformance to an SLS, the EMS/NMS can use the data collected for each pair of MEPs in the configured subset and compute a single value for the EVC or OVC and Class of Service as specified in MEF 10.3 [12] or MEF 26.1 [18] – see Appendix I – Calculation of SLS Performance Metrics (Informative).

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

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

- [R54] A SOAM PM Implementation **MUST** support the ITU-T Single-Ended ETH-DM function protocol and the procedures as specified by ITU-T G.8013/Y.1731 [1], ITU-T G.8021 [3] as amended [4] [5]. Any exceptions to the requirements, behavior, and default characteristics as defined in those specifications are called out in this section.
- [R55] A SOAM PM Implementation **MUST** support the receive timestamp in the Forward direction (*RxTimeStamp*), and the transmit timestamp in the Backward

direction (*TxTimeStamp*) 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 for each PM Session.

- [R56] A SOAM PM Implementation **MUST** support a configurable unicast destination MAC address for DMM frames.
- [R57] A SOAM PM Implementation **MUST** support a configurable SOAM PM CoS ID for DMM frame transmission. This requirement is not applicable if the SOAM PM Frames are untagged.
- [R58] A SOAM PM Implementation **MUST** support the following options for the configuration of SOAM PM CoS IDs:
- VLAN ID
 - A combination of the PCP and VLAN ID
- This requirement is not applicable if the SOAM PM Frames are untagged.
- [R59] An implementation of a Responder MEP **MUST** accept the SOAM PM CoS ID received in a DMM frame and copy the CoS ID to the associated DMR response it sends. This requirement is not applicable if the DMM frames are untagged.
- [R60] If the DMM frames are 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 frame transmission.
- [R61] A SOAM PM Implementation **MUST** support a configurable period for DMM frame transmission.
- [R62] The periods of {100 ms, 1 sec, 10 sec} **MUST** be supported for DMM frame transmission.
- [D21] The default period **SHOULD** be {1 sec}.
- [R63] A SOAM PM Implementation on the Controller MEP **MUST** support a configurable frame size for DMM frame 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.

- [R64] The range of Ethernet frame sizes from 64 through 2000 octets **MUST** be supported.
- [D22] The range of Ethernet frame sizes from 2001 through 9600 octets **SHOULD** be supported.
- [D23] The default frame size **SHOULD** be 64 octets, which is the minimum valid Ethernet frame size.

Measurement accuracy can be improved by ensuring the DMM frame size closely represents the average size of data frames for the class of service for which the measurement is being taken.

- [O5] 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 $\{f_1, f_{1+n}\}, \{f_2, f_{2+n}\}, \{f_3, f_{3+n}\}, \dots\}$.

- [D24] The default 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 10.3 [12].

- [R65] A SOAM PM Implementation **MUST** support, for FDR measurement purposes, normalizing delays by subtracting the estimated minimum delay of the interval.
- [D25] A SOAM PM Implementation **SHOULD** use the observed minimum delay of the previous Measurement Interval as the estimated minimum delay to normalize FDR measurements at the beginning of a Measurement Interval.
- [D26] 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 EMS/NMS 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 recorded by the Controller MEP per Measurement Interval.

- [R66] A SOAM PM Implementation **MUST** support the following data at the Controller MEP per Measurement Interval per PM Session:

Data	Description
Start Time-of-day timestamp	A 64-bit timestamp of the time-of-day in UTC at the scheduled start time of the Measurement Interval.
End Time-of-day timestamp	A 64-bit timestamp of the time-of-day in UTC at the scheduled end time 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. Note: this may differ from the difference between the start and end times if measurements started or stopped part way through the Measurement Interval, or if there was a shift in the time-of-day clock. Some of these conditions will result in the Suspect Flag being set.
SOAM PM Frames Sent	A 32-bit counter reflecting the number of SOAM PM Frames sent.
SOAM PM Frames Received	A 32-bit counter reflecting the number of SOAM PM Frames 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.
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.

Data	Description
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.
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 9 – 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. Note that when time-of-day clock synchronization is not in effect, measurements of One-way FD may result in a negative value for the minimum. This does not impact the ability to monitor changes in the minimum for the purpose of invalidating 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 at the Controller MEP 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 10 – Mandatory Single-Ended Delay Data Set with Clock Synchronization

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

The following requirements apply to a SOAM PM Implementation of the Single-Ended Synthetic Loss function.

- [R68] A SOAM PM Implementation **MUST** support the ITU-T ETH-SLM protocol and procedures as specified by ITU-T G.8013/Y.1731 [1] and ITU-T G.8021 [3] as amended [4] [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 for each PM Session.

- [R69] A SOAM PM Implementation **MUST** support a configurable unicast destination MAC address for SLM frames.
- [R70] 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 Frames are untagged.
- [R71] A SOAM PM Implementation **MUST** support the following options for the configuration of SOAM PM CoS IDs:
- VLAN ID
 - A combination of the PCP and VLAN ID
- This requirement is not applicable if the SOAM PM Frames are untagged.
- [R72] An implementation of a Responder MEP **MUST** accept SOAM PM CoS IDs received in SLM frames and copy the CoS ID to the associated SLR response it sends. This requirement is not applicable if the SLM frames are untagged.
- [R73] If the SLM frames are 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 frame transmission.
- [R74] A SOAM PM Implementation **MUST** support a configurable period for SLM frame transmission.
- [R75] The periods of {100 ms, 1 sec, 10 sec} **MUST** be supported for SLM frame transmission.
- [D27] The period of 10ms **SHOULD** be supported for SLM frame transmission.
- [D28] The default period **SHOULD** be {100 ms}.
- [R76] 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.

- [R77] The range of Ethernet frame sizes from 64 through 2000 octets **MUST** be supported.

- [D29] The range of Ethernet frame sizes from 2001 through 9600 octets **SHOULD** be supported.
- [D30] The default frame size **SHOULD** be 64 octets, which is the minimum valid Ethernet frame size.

Measurement accuracy can be improved by ensuring the SLM frame size closely represents the average size of data frames for the class of service for which the measurement is being taken.

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 Forward direction; or the SLR is lost in the Backward direction. To calculate the Forward or Backward FLR, a number of SLMs are transmitted, and the corresponding number lost in each direction is 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. Conversely, the shorter the period between SLM frames (and the longer the SLM frame size), the more bandwidth will be used for SLM frames, and the higher the load will be on the SOAM PM Implementation. See Appendix D - Statistical Considerations for Synthetic Loss Measurement (Informative).

The following requirements apply to the calculation of Availability and Group Availability, which are explained in detail in MEF 10.3 [12]. 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 Unavailable state is entered. Note that Availability flr is different from FLR, which is calculated over the much larger interval T . Availability measurements can also be used to calculate Group Availability: the Group Availability of a set of EI pairs is Unavailable if and only if the Availability is Unavailable for every EI pair in the set.

- [R78] A SOAM PM Implementation **MUST** support a configurable parameter for the length of time over which each Availability flr value is calculated, with a range of 1s – 300s. This parameter is equivalent to Δt as specified by MEF 10.3 [12].
- [R79] The length of time over which each Availability flr value is calculated (Δt) **MUST** be an integer multiple of the interval between each SLM frame transmission.
- [D31] The default length of time over which each Availability flr value is calculated **SHOULD** be 1s.
- [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 Available/Unavailable state transitions. This parameter is equivalent to the Availability parameter of n as specified by MEF 10.3 [12].
- [D32] The default 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 Δt intervals all have loss ratio exceeding the threshold, then an Unavailable state has been entered and all Δt intervals within that window will be designated as Unavailable. Details are in MEF 10.3 [12].

- [R81] A SOAM PM Implementation **MUST** support a configurable Availability frame loss ratio threshold to be used in evaluating the Available/Unavailable state of each Δt interval per MEF 10.3 [12].
- [R82] The Availability frame loss ratio threshold range of 0.00 through 1.00 **MUST** be supported in increments of 0.01.
- [D33] The default Availability frame loss ratio threshold **SHOULD** be 0.1.
- [R83] A SOAM PM Implementation at a Controller MEP **MUST** report to the managing system whenever a transition between Available and Unavailable occurs in the status of an adjacent pair of Δt intervals per MEF 10.3 [12].
- [R84] The Availability state 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 11 – Availability State Transition Event Data

Note: a transition cannot be detected for $n \Delta t$ seconds after it has occurred. The timestamp in the Availability state transition report should be the time of transition, not the time of detection.

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

A number of parameters relating to Availability are interrelated and have an impact on the precision of Availability flr measurements, the bandwidth consumed by SLM frames, the time taken to detect a change in Availability state and the damping effect on Available/Unavailable state changes due to the sliding window algorithm defined in MEF 10.3 [12]. These parameters are:

- SLM Frame Transmission Period, P ([R74])

- SLM Frame Size, s ([R76])
- Time over which each Availability flr is calculated, Δt ([R78])
- The number of consecutive Availability flr measurements to be used to determine Available/Unavailable state transitions, n ([R80])
- The Availability threshold, C ([R81])

These parameters are related in the following way:

- The precision of Availability flr measurements is determined by the number of SLMs for the calculation, i.e. the number SLM frames transmitted during each Δt interval, $\Delta t/P$. The precision can be improved by increasing Δt or decreasing P . The lower the threshold C , the higher precision is needed.
- The bandwidth used for SLM frames is determined by the size and frequency of SLM frames, i.e. bandwidth in bps is s/P . The bandwidth can be lowered by decreasing s or increasing P .
- The time taken to detect a change in Availability state is determined by the time over which each Availability flr is calculated, and the number of Availability flrs used to determine a state change, i.e. $n \Delta t$. The time can be reduced by decreasing n or Δt .
- The damping effect of the sliding window algorithm is determined by the number of Availability flrs used to determine a state change, n . It can be improved by increasing n .

It can be seen that there are several conflicting considerations in determining the value of the various parameters, and hence tradeoffs are needed; for example, more precise Availability flr measurements can be obtained at the expense of higher bandwidth use or longer detection time. It is therefore impossible to give a single set of values that is suitable for all cases.

Note that the default values of P ([D28]) and Δt ([D31]) are such that 10 SLMs are used for each Availability flr calculation. In general it is important that sufficient SLMs are used for each Availability flr calculation to give a high degree of confidence that the threshold C is crossed only when the actual frame loss ratio over Δt is above the threshold. Appendix J – Statistical Considerations for Availability (Informative) gives more detail on the statistical considerations for Availability calculation.

The following requirements apply to the measurement of HLI and CHLI, which are explained in detail in MEF 10.3 [12]

- [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 10.3 [12].

- [D34] The default value for the number of HLIs that constitute a CHLI **SHOULD** be 5.
- [D35] 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 [R78].

As with Availability, a number of consecutive HLIs that constitute a CHLI could span the end of one Measurement Interval and the start of the following Measurement Interval. In this case, the CHLI is counted in the Measurement Interval in which it ends.

- [R87] A SOAM PM Implementation **MUST** include CHLIs that end during a given Measurement Interval, and only those CHLIs, in the count of CHLIs for that Measurement Interval

The following requirements specify the *output data set* that is recorded by the Controller MEP per Measurement Interval.

- [R88] A SOAM PM Implementation **MUST** support the following data at the Controller MEP per Measurement Interval per PM Session:

Data	Description
Start Time-of-day timestamp	A 64-bit timestamp of the time-of-day in UTC at the scheduled start time of the Measurement Interval.
End Time-of-day timestamp	A 64-bit timestamp of the time-of-day in UTC at the scheduled end time 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. Note: this may differ from the difference between the start and end times if measurements started or stopped part way through the Measurement Interval, or if there was a shift in the time-of-day clock. Some of these conditions will result in the Suspect Flag being set.
SOAM PM Frames Sent ⁵	A 32-bit counter reflecting the number of SOAM PM Frames sent.
SOAM PM Frames Received ⁵	A 32-bit counter reflecting the number of SOAM PM Frames received.
Tx frame count in the Forward direction ⁵	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.

⁵ For Single-Ended Synthetic Loss, SOAM PM Frames Sent is equal to Tx frame count in the Forward Direction and SOAM PM Frames Received is equal to Rx frame count in the Backward Direction. Both fields are specified so as to retain consistency with other PM Functions.

Data	Description
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 ⁵	A 32-bit counter reflecting the number of SLR frames received in the Backward direction.
Count of Δt intervals evaluated as Available in the Forward direction	A 32-bit counter reflecting the number of Δt intervals evaluated as Available in the Forward direction (i.e., for which $A_{\langle \text{Controller}, \text{Responder} \rangle}(\Delta t) = 1$).
Count of Δt intervals evaluated as Available in the Backward direction	A 32-bit counter reflecting the number of Δt intervals evaluated as Available in the Backward direction (i.e., for which $A_{\langle \text{Responder}, \text{Controller} \rangle}(\Delta t) = 1$).
Count of Δt intervals evaluated as Unavailable in the Forward direction	A 32-bit counter reflecting the number of Δt intervals evaluated as Unavailable in the Forward direction (i.e., for which $A_{\langle \text{Controller}, \text{Responder} \rangle}(\Delta t) = 0$).
Count of Δt intervals evaluated as Unavailable in the Backward direction	A 32-bit counter reflecting the number of Δt intervals evaluated as Unavailable in the Backward direction (i.e., for which $A_{\langle \text{Responder}, \text{Controller} \rangle}(\Delta t) = 0$).
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 12 – Mandatory Single-Ended Synthetic Loss Data Set

[D36] A SOAM PM Implementation **SHOULD** support the following additional Availability related data at the Controller MEP per Measurement Interval per PM Session:

Data	Description
Minimum One-way Availability flr in the Forward direction	The minimum One-way Availability flr measurement during this Measurement Interval.
Maximum One-way Availability flr in the Forward direction	The maximum One-way Availability flr measurement during this Measurement Interval.
Mean One-way Availability flr in the Forward direction	The average (arithmetic mean) One-way Availability flr measurement during this Measurement Interval.

Data	Description
Minimum One-way Availability flr in the Backward direction	The minimum One-way Availability flr measurement during this Measurement Interval.
Maximum One-way Availability flr in the Backward direction	The maximum One-way Availability flr measurement during this Measurement Interval.
Mean One-way Availability flr in the Backward direction	The average (arithmetic mean) One-way Availability flr measurement during this Measurement Interval.

Table 13 – Optional Single-Ended Synthetic Loss Data Set

12. PM-2 Requirements

The PM-2 Solution uses the Dual-Ended Delay PM Function for Frame Delay (FD), Inter-Frame Delay Variation (IFDV), Frame Delay Range (FDR) and Mean Frame Delay (MFD) measurements. The mechanisms support both point-to-point and multipoint connections.

Section 12.1 lists the requirements for performing Frame Delay and Inter-Frame Delay Variation measurements using the Dual-Ended Delay functions.

- [O6] A SOAM PM Implementation **MAY** support the Dual-Ended Delay Function as described in section 12.1.

12.1 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 results in one or more distinct PM Sessions. Dual-Ended Delay supports both point-to-point and multipoint configurations.

On multipoint MEGs 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 SOAM PM CoS ID. If the measurements are being used to evaluate conformance to an SLS, the EMS/NMS can use the data collected for each ordered pair of MEPs in the configured subset and compute a single value for the EVC or OVC and Class of Service as specified in MEF 10.3 [12] or MEF 26.1 [18] – see Appendix I – Calculation of SLS Performance Metrics (Informative).

When using Dual-Ended Delay, a single direction (A->B or B->A) can be measured using one PM Session, or both directions can be measured (A->B and B->A.) by using a separate PM Session for each direction.

- [CR23]< [O6] A SOAM PM Implementation **MUST** support the ITU-T Dual-Ended ETH-DM Function protocol and procedures as specified by ITU-T G.8013/Y.1731, [1] ITU-T G.8021 [3] as amended [4] [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 for each PM Session.

- [CR24]< [O6] A SOAM PM Implementation of a Controller MEP **MUST** support a configurable unicast destination MAC address for 1DM frames.
- [CR25]< [O6] A SOAM PM Implementation of a Sink MEP **MUST** support a configurable unicast source MAC address for 1DM frames.

A Sink MEP could also support a mode where 1DM frames from any source MAC are accepted; in this case 1DM frames received from different source MAC addresses are treated as belonging to different PM Sessions.

[CR26]< [O6] A SOAM PM Implementation **MUST** support Class 1 multicast destination MAC address for 1DM frames.

[CR27]< [O6] 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 Frames are untagged.

[CR28]< [O6] A SOAM PM Implementation **MUST** support the following options for the configuration of SOAM PM CoS IDs:

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

This requirement is not applicable if the SOAM PM Frames are untagged.

[CR29]< [O6] If the 1DM frames are 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 frame transmission.

[CR30]< [O6] A SOAM PM Implementation **MUST** support a configurable period for 1DM frame transmission.

[CR31]< [O6] The periods of {100 ms, 1 sec, 10 sec} **MUST** be supported for 1DM frame transmission.

[CD5]< [O6] The default period **SHOULD** be {1 sec}.

[CR32]< [O6] A SOAM PM Implementation on the Controller MEP **MUST** support a configurable frame size for 1DM frame transmission.

Note: The frame size does not need to be configured at the Sink MEP.

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.

[CR33]< [O6] The range of Ethernet frame sizes from 64 through 2000 octets **MUST** be supported.

[CD6]< [O6] The range of Ethernet frame sizes from 2001 through 9600 octets **SHOULD** be supported.

- [CD7]< [O6] The default frame size **SHOULD** be 64 octets, which is the minimum valid Ethernet frame size.

Measurement accuracy can be improved by ensuring the 1DM frame size closely represents the average size of data frames for the class of service for which the measurement is being taken.

- [CO3]< [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 $\{f_1, f_{1+n}\}, \{f_2, f_{2+n}\}, \{f_3, f_{3+n}\}, \dots\}$.

- [CD8]< [O6] The default selection offset for IFDV **SHOULD** be 1.

This parameter, when multiplied by the period parameter of [CR30], is equivalent to the IFDV parameter of Δt as specified by MEF 10.3 [12].

- [CR34]< [O6] A SOAM PM Implementation **MUST** support, for FDR measurement purposes, normalizing delays by subtracting the estimated minimum delay of the interval.

- [CD9]< [O6] A SOAM PM Implementation **SHOULD** use the observed minimum delay of the previous Measurement Interval as the estimated minimum delay to normalize FDR measurements at the beginning of a Measurement Interval.

- [CD10]< [O6] 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 EMS/NMS 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 is recorded by the Controller MEP or Sink MEP (as indicated) per Measurement Interval.

- [CR35]< [O6] A SOAM PM Implementation **MUST** support the following data at the Controller or Sink MEP (as indicated) per Measurement Interval per PM Session:

Data	Description	MEP
Start Time-of-day timestamp	A 64-bit timestamp of the time-of-day in UTC at the scheduled start time of the Measurement Interval.	Both
End Time-of-day timestamp	A 64-bit timestamp of the time-of-day in UTC at the scheduled end time of the Measurement Interval.	Both
Measurement Interval elapsed time	A 32-bit counter of the number of seconds of the Measurement Interval as calculated by the NE. Note: this may differ from the difference between the start and end times if measurements started or stopped part way through the Measurement Interval, or if there was a shift in the time-of-day clock. Some of these conditions will result in the Suspect Flag being set.	Both
SOAM PM Frames Sent	A 32-bit counter reflecting the number of SOAM PM Frames sent.	Controller
SOAM PM Frames Received	A 32-bit counter reflecting the number of SOAM PM Frames received.	Sink
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.	Sink
Mean One-way IFDV	A 32-bit integer reflecting the average (arithmetic mean) One-way IFDV measurement in microseconds.	Sink
Maximum One-way IFDV	A 32-bit integer reflecting the maximum One-way IFDV measurement in microseconds.	Sink
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.	Sink
Mean One-way FDR	A 32-bit integer reflecting the average (arithmetic mean) One-way FDR measurement in microseconds.	Sink
Maximum One-way FDR	A 32-bit integer reflecting the maximum One-way FDR measurement in microseconds.	Sink
Minimum One-way FD	A 32-bit integer reflecting the minimum One-way FD measurement in microseconds.	Sink

Table 14 – 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. Note that when time-of-day clock synchronization is not in effect, measurements of One-way FD may result in a negative value for the

minimum. This does not impact the ability to monitor changes in the minimum for the purpose of invalidating FDR measurements.

[CR36]< [O6] 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	MEP
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.	Sink
Mean One-way FD	A 32-bit integer reflecting the average (arithmetic mean) One-way FD measurement in microseconds.	Sink
Maximum One-way FD	A 32-bit integer reflecting the maximum One-way FD measurement in microseconds.	Sink

Table 15 – Mandatory Dual-Ended Delay Data Set with Clock Synchronization

13. PM-3 Requirements

PM-3 uses the Single-Ended Service Loss PM Function 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 results in one or more distinct PM Sessions. The function supports 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 13.1.

Note that Availability cannot be calculated with PM-3 because it cannot measure loss when there are no Service Frames being transmitted.

13.1 Single-Ended Service Loss Function for FLR

The following requirements apply to a SOAM PM Implementation of the Single-Ended Service Loss function.

- [CR37]< [O7] A SOAM PM Implementation **MUST** support the ITU-T Single-Ended ETH-LM protocol and procedures as specified by ITU-T G.8013/Y.1731 [1], ITU-T G.8021 [3] as amended [4] [5]. Any exceptions to the requirements, behavior, and default characteristics as defined in those specifications are called out in this section.
- [CR38]< [O7] A SOAM PM Implementation **MUST** convey frame counts that include only those Service Frames that have a level of bandwidth profile conformance determined to be *green*, and certain Service OAM frames as specified by ITU-T G.8013/Y.1731 [1] and ITU-T G.8021 [3] as amended [4] [5].

The definition of FLR in MEF 10.3 [12] and MEF 26.1 [18] includes only Service Frames that have a level of bandwidth profile conformance determined to be *green* – i.e. it does not include any OAM frames, or service frames marked as *yellow*. However, the Single-Ended Service Loss function defined by ITU-T does include certain OAM frames in the counters. There is hence the potential for a slight discrepancy between the measured value of FLR and the actual value based on only the *green* Service Frames; however, this is expected to be negligible in practice.

Note: IEEE 802.1Q-2014 [22] and ITU-T G.8021 [3] as amended [4] [5] differ as to how Linktrace Messages (LTMs) are sent. In the IEEE model, Linktrace Messages are injected directly at the egress interface (via the “Linktrace Output Multiplexer”), bypassing all other MEPs and MIPs; in the ITU-T model, LTMs are injected by the transmitting MEP or MIP and hence pass through all the lower level MEPs. As specified by ITU-T, LTMs sent at a higher level are included in the counters for LMM; however, if the IEEE model is used for LTM, the transmitted LTMs would bypass the MEP and not be counted. In both cases, received LTMs are counted. Implementations of the IEEE model that do not specifically account for this may cause slight discrepancies in the LMM measurements; however, the impact is expected to be minor and in mitigation, linktrace is an infrequent event.

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

[CR39]< [O7] A SOAM PM Implementation **MUST** support a configurable unicast destination MAC address for LMM frames.

[CR40]< [O7] A SOAM PM Implementation **MUST** support a configurable SOAM PM CoS ID for LMM frame transmission. This requirement is not applicable if the SOAM PM Frames are untagged.

[CR41]< [O7] A SOAM PM Implementation **MUST** support the following options for the configuration of SOAM PM CoS IDs:

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

This requirement is not applicable if the SOAM PM Frames are untagged.

[CR42]< [O7] An implementation of a Responder MEP **MUST** accept the SOAM PM CoS ID received in a LMM frame and copy the CoS ID to the associated LMR response it sends. This requirement is not applicable if the LMM frames are untagged.

[CR43]< [O7] If the LMM frames are 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 frame transmission.

[CR44]< [O7] A SOAM PM Implementation **MUST** support a configurable period for LMM frame transmission.

[CR45]< [O7] For the LMM frame transmission, periods of {100 ms, 1 sec, 10 sec} **MUST** be supported.

[CD11]< [O7] The default period **SHOULD** be {1 sec}.

The following requirements specify the *output data set* that is recorded by the Controller MEP per Measurement Interval.

[CR46]< [O7] A SOAM PM Implementation **MUST** support the following data at the Controller MEP per Measurement Interval per PM Session:

Data	Description
Start Time-of-day timestamp	A 64-bit timestamp of the time-of-day in UTC at the scheduled start time of the Measurement Interval.

Data	Description
End Time-of-day timestamp	A 64-bit timestamp of the time-of-day in UTC at the scheduled end time 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. Note: this may differ from the difference between the start and end times if measurements started or stopped part way through the Measurement Interval, or if there was a shift in the time-of-day clock. Some of these conditions will result in the Suspect Flag being set.
SOAM PM Frames Sent	A 32-bit counter reflecting the number of SOAM PM Frames sent (i.e., LMM frames transmitted).
SOAM PM Frames Received	A 32-bit counter reflecting the number of SOAM PM Frames 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 16 – Mandatory Single-Ended Service Loss Data Set

14. PM-4 Requirements

The PM-4 Solution uses Dual-Ended Synthetic Loss function for Frame Loss Ratio (FLR), Availability, High Loss Intervals (HLI) and Consecutive High Loss Intervals (CHLI) measurements. The mechanisms support both point-to-point and multipoint connections.

Section 14.1 lists the requirements for measuring FLR, Availability, HLI and CHLI using Dual-Ended Synthetic Loss function.

- [O8] A SOAM PM Implementation **MAY** support the Dual-Ended Synthetic Loss function as described in section 14.1.

14.1 Dual-Ended Synthetic Loss Function for FLR, Availability, HLI, CHLI

Dual-Ended Synthetic Loss can be configured for multiple classes of service in each direction between a pair of MEPs. Each unique pair of MEPs, direction, and Class of Service to be measured results in one or more distinct PM Sessions. Dual-Ended Synthetic Loss supports both point-to-point and multipoint configurations.

On multipoint MEGs 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 SOAM PM CoS ID. If the measurements are being used to evaluate conformance to an SLS, 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 or OVC and Class of Service as specified in MEF 10.3 [12] or MEF 26.1 [18] – see Appendix I – Calculation of SLS Performance Metrics (Informative).

When using Dual-Ended Synthetic Loss, a single direction (A->B or B->A) can be measured using one PM Session, or both directions can be measured (A->B and B->A) by using a separate PM Session for each direction.

The following requirements apply to a SOAM PM Implementation of the Dual-Ended Synthetic Loss function.

- [CR47]< [O8] A SOAM PM Implementation **MUST** support the ITU-T Dual-Ended ETH-SLM protocol and procedures as specified by ITU-T G.8013/Y.1731 [1] and ITU-T G.8021 [3] as amended [4] [5].

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

- [CR48]< [O8] A SOAM PM Implementation of a Controller MEP **MUST** support a configurable unicast destination MAC address for 1SL frames.
- [CR49]< [O8] A SOAM PM Implementation of a Sink MEP **MUST** support a configurable unicast source MAC address for 1SL frames.

A Sink MEP could also support a mode where 1SL frames from any source MAC are accepted; in this case 1SL frames received from different source MAC addresses are treated as belonging to different PM Sessions.

[CR50]< [O8] A SOAM PM Implementation **MUST** support Class 1 multicast destination MAC address for 1SL frames.

[CR51]< [O8] A SOAM PM Implementation **MUST** support a configurable SOAM PM CoS ID for 1SL frame transmission. This requirement is not applicable if the SOAM PM Frames are untagged.

[CR52]< [O8] A SOAM PM Implementation **MUST** support the following options for the configuration of SOAM PM CoS IDs:

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

This requirement is not applicable if the SOAM PM Frames are untagged.

[CR53]< [O8] If the 1SL frames are 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 1SL frame transmission.

[CR54]< [O8] A SOAM PM Implementation **MUST** support a configurable period for 1SL frame transmission.

[CR55]< [O8] The periods of {100 ms, 1 sec, 10 sec} **MUST** be supported for 1SL frame transmission.

[CD12]< [O8] The period of 10ms **SHOULD** be supported for 1SL frame transmission.

[CD13]< [O8] The default period **SHOULD** be {100 ms}.

[CR56]< [O8] A SOAM PM Implementation of the Controller MEP **MUST** support a configurable frame size for 1SL frame transmission.

Note: The frame size does not need to be configured at the Sink MEP.

The frame size corresponds to a valid MEF Service Ethernet frame and is inclusive of the Ethernet header, the 1SL 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 1SL PDU.

[CR57]< [O8] The range of Ethernet frame sizes from 64 through 2000 octets **MUST** be supported.

[CD14]< [O8] The range of Ethernet frame sizes from 2001 through 9600 octets **SHOULD** be supported.

[CD15]< [O8] The default frame size **SHOULD** be 64 octets, which is the minimum valid Ethernet frame size.

Measurement accuracy can be improved by ensuring the 1SL frame size closely represents the average size of data frames for the class of service for which the measurement is being taken.

When the Dual-Ended Synthetic Loss Function is used, each transmitted 1SL frame has two possible outcomes: the 1SL is received; or the 1SL is lost. To calculate the FLR, a number of 1SLs must be transmitted, and the corresponding number lost must be measured. The FLR can then be calculated in the normal way. Note: the more 1SLs used for FLR calculation, the more precise the resulting FLR value will be. Conversely, the shorter the period between 1SL frames (and the longer the 1SL frame size), the more bandwidth will be used for 1SL frames, and the higher the load will be on the SOAM PM Implementation. See Appendix D - Statistical Considerations for Synthetic Loss Measurement (Informative).

The following requirements apply to the calculation of Availability and Group Availability, which are explained in detail in MEF 10.3 [12]. 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 Unavailable state is entered. Note that Availability flr is different from FLR, which is calculated over the much larger interval T . Availability measurements can also be used to calculate Group Availability: the Group Availability of a set of EI pairs is Unavailable if and only if the Availability is Unavailable for every EI pair in the set.

[CR58]< [O8] A SOAM PM Implementation **MUST** support a configurable parameter for the length of time over which each Availability flr value is calculated, with a range of 1s – 300s. This parameter is equivalent to Δt as specified by MEF 10.3 [12].

[CR59]< [O8] The length of time over which each Availability flr value is calculated (Δt) **MUST** be an integer multiple of the interval between each 1SL frame transmission.

[CD16]< [O8] The default length of time over which each Availability flr value is calculated **SHOULD** be 1s.

[CR60]< [O8] The number range of 1 through 10 **MUST** be supported for the configurable number of consecutive Availability flr measurements to be used to determine Available/Unavailable state transitions. This parameter is equivalent to the Availability parameter of n as specified by MEF 10.3 [12].

[CD17]< [O8] The default 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 Δt intervals all have loss ratio exceeding the threshold,

then an Unavailable state has been entered and all Δt intervals within that window will be designated as Unavailable. Details are in MEF 10.3 [12].

- [CR61]< [O8] A SOAM PM Implementation **MUST** support a configurable Availability frame loss ratio threshold to be used in evaluating the Available/Unavailable state of each Δt interval per MEF 10.3 [12].
- [CR62]< [O8] The Availability frame loss ratio threshold range of 0.00 through 1.00 **MUST** be supported in increments of 0.01.
- [CD18]< [O8] The default Availability frame loss ratio threshold **SHOULD** be 0.1.
- [CR63]< [O8] A SOAM PM Implementation at a Sink MEP **MUST** report to the managing system whenever a state transition between Available and Unavailable occurs in the status of an adjacent pair of Δt intervals per MEF 10.3 [12].
- [CR64]< [O8] The Availability state transition report from a Sink MEP **MUST** include the following data:

Data	Description
Source	Controller MEP
Destination	Sink MEP
Cos ID	SOAM PM CoS ID
Direction	Always 'Forward'
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 17 – Availability State Transition Event Data

Note: a transition cannot be detected for $n \Delta t$ seconds after it has occurred. The timestamp in the Availability state transition report should be the time of transition, not the time of detection.

- [CR65]< [O8] If the NE maintains a time-stamped log, an entry **MUST** also be generated with the same data as the report by the Sink MEP.

As described in section 11.2, there are a number of Availability parameters that are interdependent, and tradeoffs must be made in setting the values of these.

Note that the default values of P ([CD13]) and Δt ([CD16]) are such that 10 1SLs are used for each Availability flr calculation. In general it is important that sufficient 1SLs are used for each Availability flr calculation to give a high degree of confidence that the threshold C is crossed only when the actual frame loss ratio over Δt is above the threshold. Appendix J – Statistical Considerations

for Availability (Informative) gives more detail on the statistical considerations for Availability calculation.

The following requirements apply to the measurement of HLI and CHLI, which are explained in detail in MEF 10.3 [12]

- [CR66]< [O8] 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 10.3 [12].
- [CD19]< [O8] The default value for the number of HLIs that constitute a CHLI **SHOULD** be 5.
- [CD20]< [O8] 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 [CR60].

As with Availability, a number of consecutive HLIs that constitute a CHLI could span the end of one Measurement Interval and the start of the following Measurement Interval. In this case, the CHLI is counted in the Measurement Interval in which it ends.

- [CR67]< [O8] A SOAM PM Implementation **MUST** include CHLIs that end during a given Measurement Interval, and only those CHLIs, in the count of CHLIs for that Measurement Interval

The following requirements specify the *output data set* that is recorded by the Controller MEP or Sink MEP (as indicated) per Measurement Interval.

- [CR68]< [O8] A SOAM PM Implementation **MUST** support the following data at the Controller MEP or Sink MEP (as indicated) per Measurement Interval per PM Session:

Data	Description	MEP
Start Time-of-day timestamp	A 64-bit timestamp of the time-of-day in UTC at the scheduled start time of the Measurement Interval.	Both
End Time-of-day timestamp	A 64-bit timestamp of the time-of-day in UTC at the scheduled end time of the Measurement Interval.	Both
Measurement Interval elapsed time	A 32-bit counter of the number of seconds of the Measurement Interval as calculated by the NE.	Both

Data	Description	MEP
	Note: this may differ from the difference between the start and end times if measurements started or stopped part way through the Measurement Interval, or if there was a shift in the time-of-day clock. Some of these conditions will result in the Suspect Flag being set.	
Tx frame count	A 32-bit counter reflecting the number of 1SL frames transmitted.	Controller
Rx frame count	A 32-bit counter reflecting the number of 1SL frames received.	Sink
Count of Δt intervals evaluated as Available	A 32-bit counter reflecting the number of Δt intervals evaluated as Available (i.e., for which $A_{\langle \text{Controller, Sink} \rangle}(\Delta t) = 1$).	Sink
Count of Δt intervals evaluated as Unavailable	A 32-bit counter reflecting the number of Δt intervals evaluated as Unavailable (i.e., for which $A_{\langle \text{Controller, Sink} \rangle}(\Delta t) = 0$).	Sink
Count of HLIs	Count of HLIs during the Measurement Interval.	Sink
Count of CHLIs	Count of CHLIs during the Measurement Interval.	Sink

Table 18 – Mandatory Dual-Ended Synthetic Loss Data Set

Note: The data set is based on the assumptions that the Controller MEP consistently transmits the 1SL frames and the performance measurement is processed by the Sink MEP.

[CD21]< [O8] A SOAM PM Implementation **SHOULD** support the following additional Availability related data at the Sink MEP per Measurement Interval per PM Session:

Data	Description	MEP
Minimum One-way Availability flr	The minimum One-way Availability flr measurement during this Measurement Interval.	Sink
Maximum One-way Availability flr	The maximum One-way Availability flr measurement during this Measurement Interval.	Sink
Mean One-way Availability flr	The average (arithmetic mean) One-way Availability flr measurement during this Measurement Interval.	Sink

Table 19 – Optional Dual-Ended Synthetic Loss Data Set

15. References

- [1] International Telecommunication Union, Recommendation G.8013/Y.1731 (11/2013), "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.8021 (05/2012), "Characteristics of Ethernet transport network equipment functional blocks".
- [4] International Telecommunication Union, Recommendation G.8021 Amendment 1 (10/2012), "Characteristics of Ethernet transport network equipment functional blocks".
- [5] International Telecommunication Union, Recommendation G.8021 Amendment 2 (08/2013), "Characteristics of Ethernet transport network equipment functional blocks".
- [6] International Telecommunication Union, Recommendation G.7710 (02/2012), "Common Equipment Management Function Requirements".
- [7] International Telecommunication Union, Recommendation M.3400 (02/2000), "TMN management functions".
- [8] IETF RFC 2119, "Key words for use in RFCs to Indicate Requirement Levels", March 1997.
- [9] IETF RFC 5951, "Network Management Requirements for MPLS-based Transport Networks", September 2010.
- [10] MEF Technical Specification MEF 4, "Metro Ethernet Network Architecture Framework - Part 1: Generic Framework", May 2004.
- [11] MEF Technical Specification MEF 7.2, "Carrier Ethernet Management Information Model", April 2013.
- [12] MEF Technical Specification MEF 10.3, "Ethernet Services Attributes Phase 3", October 2013.
- [13] MEF Technical Specification MEF 12.2, "Carrier Ethernet Network Architecture Framework, Part 2: Ethernet Services Layer", April 2014.
- [14] MEF Technical Specification MEF 15, "Requirements for Management of Metro Ethernet Phase 1 Network Elements", November 2005.
- [15] MEF Technical Specification MEF 17, "Service OAM Requirements & Framework - Phase 1", April 2007.

- [16] MEF Technical Specification MEF 20, "User Network Interface (UNI) Type 2 Implementation Agreement", July 2008.
- [17] MEF Technical Specification MEF 23.1, "Carrier Ethernet Class of Service – Phase 2", January 2012.
- [18] MEF Technical Specification MEF 26.1, "External Network Network Interface (ENNI) - Phase 2", January 2012.
- [19] MEF Technical Specification MEF 30.1, "Service OAM Fault Management Implementation Agreement: Phase 2", April 2013.
- [20] MEF Technical Specification MEF 30.1.1, "Amendment to MEF 30.1 – Correction to Requirement", April 2014.
- [21] MEF Technical Specification MEF 47, "Carrier Ethernet Services for Cloud Implementation Agreement".
- [22] IEEE 802.1Q-2014, "IEEE Standard for Local and Metropolitan Area Networks – Bridges and Bridged Networks", 2014
- [23] ISO 8601, "Data elements and interchange formats - Information interchange – Representation of dates and times", Third Edition, December 1, 2004.
- [24] Telcordia GR-253 (09/2000), "SONET Transport Systems: Common Generic Criteria".
- [25] International Telecommunication Union, Recommendation X.738 (11/1993), " Information technology – Open Systems Interconnection – Systems management: Summarization function".

16. Appendix A - Performance Management Functions (Informative)

The following sections provide an overview of the PM Functions specified by ITU-T G.8013/Y.1731 [1], ITU-T G.8021 [3] as amended [4] [5].

16.1 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:

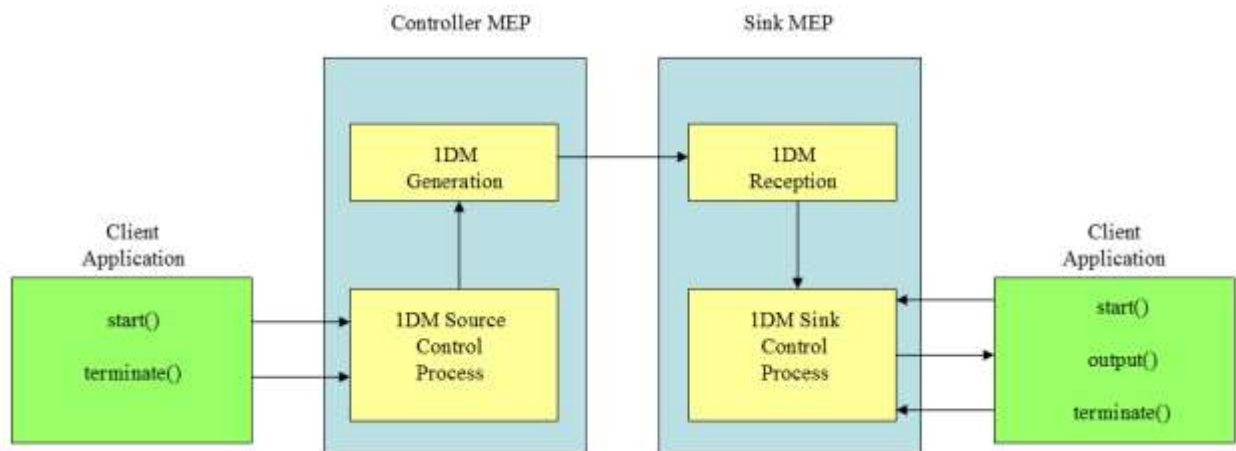


Figure 17 – Dual-Ended Delay Processes

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

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

Table 20 – 1DM Source Control Process Signals

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

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

Table 21 – 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.

16.2 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:

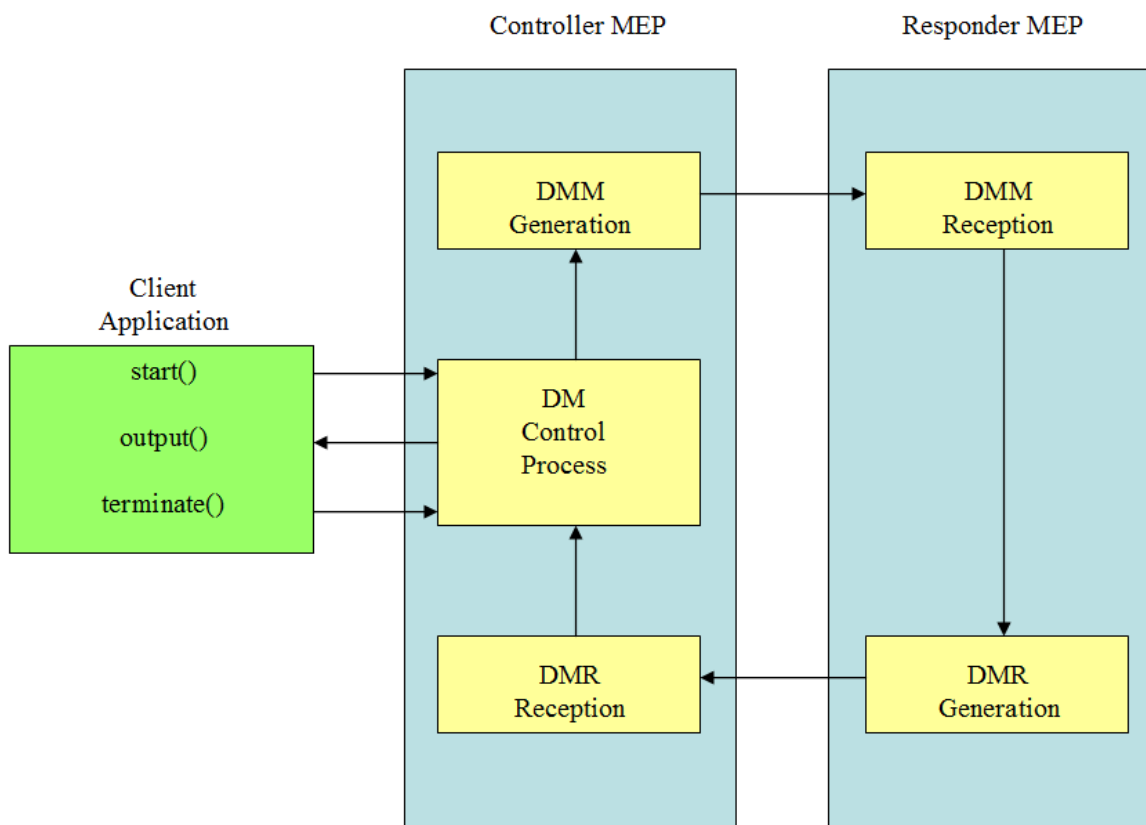


Figure 18 – Single-Ended Delay Processes

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

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

Table 22 – DM Control Process Signals

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

16.3 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 green Service Frames (i.e., unicast, multicast, and broadcast), and some SOAM Frames, 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 frame transmit count at the Controller MEP at the time of LMM transmission. The second is a Loss Measure Reply (LMR) which conveys the frame transmit and receive counts at the Responder MEP at the time of LMM reception. The 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 frame transmit and receive counts contained in the LMR and the 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:

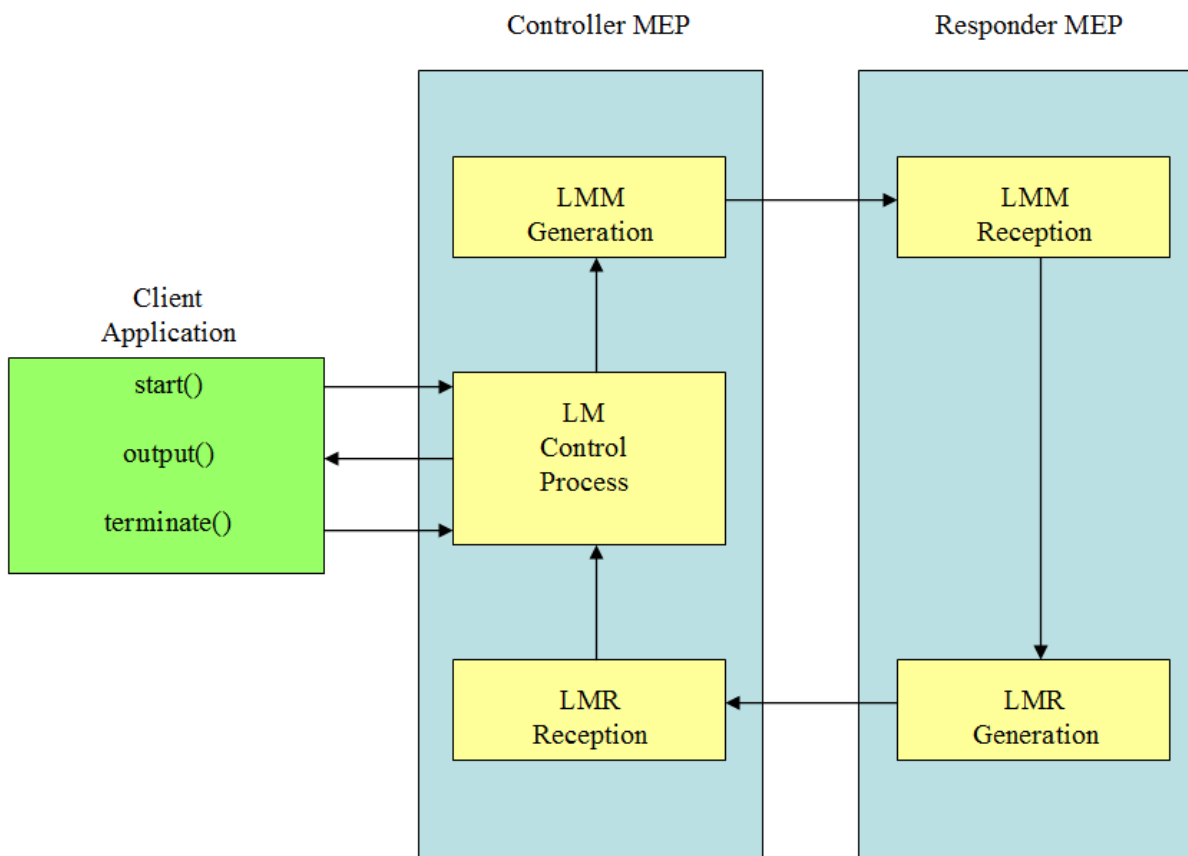


Figure 19 – Single-Ended Loss Processes

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

Signal	Parameters
<i>start()</i>	DA (destination unicast MAC address)
	VLAN PCP (0..7, not applicable if untagged)
	Period of LMM generation (ms)
<i>terminate()</i>	None

<i>result()</i>	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 23 – LM Control Process Signals

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

16.4 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).

Note that the Responder MEP must generate sequence numbers in the SLRs that are specific to each (Controller MEP, Test ID) pair for which it receives SLMs.

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:

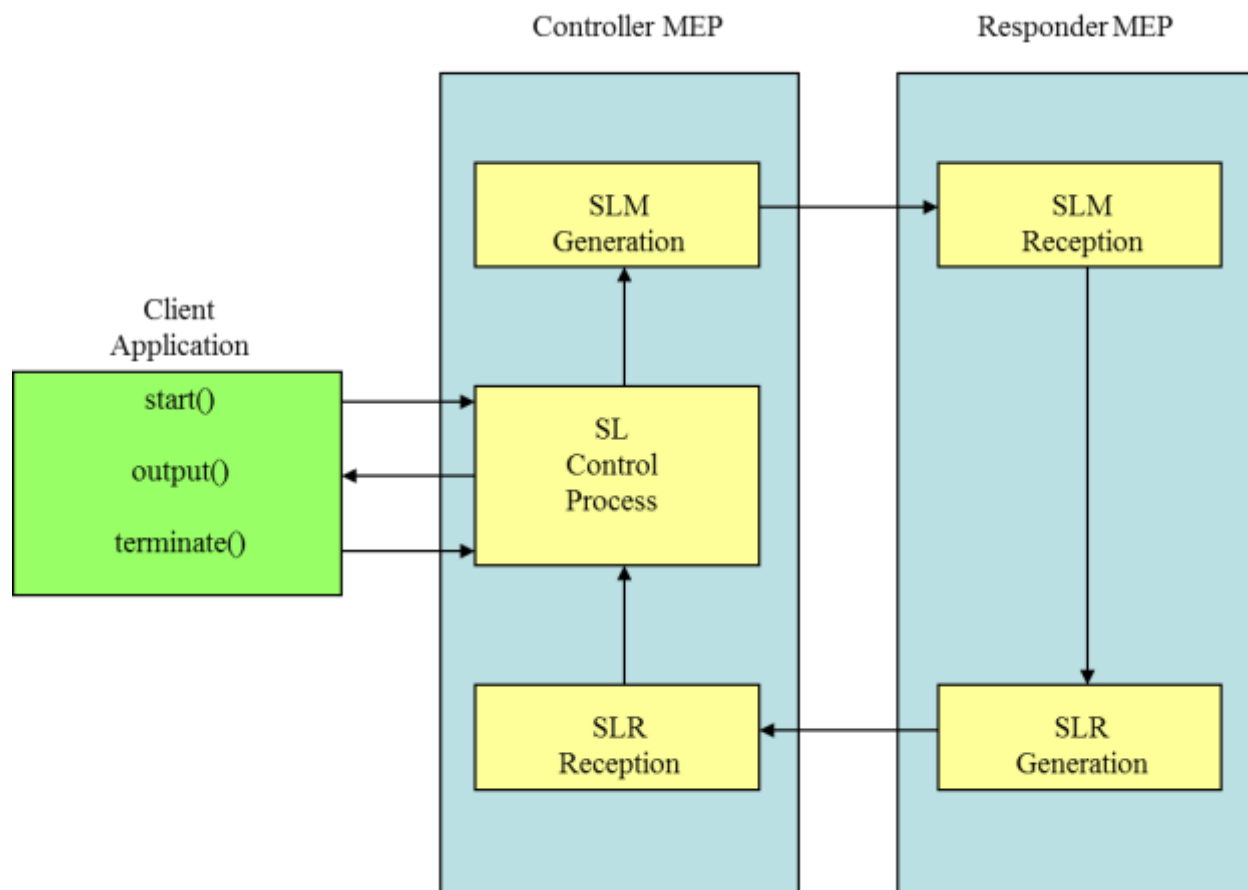


Figure 20 – Single-Ended Synthetic Loss Processes

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

Signal	Parameters
<i>start()</i>	DA (destination unicast MAC address)
	VLAN PCP (0..7, not applicable if untagged)
	Period of SLM generation (ms)
	Test ID
	Length
<i>terminate()</i>	None
<i>result()</i>	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 24 – SL Control Process Signals

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

16.5 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 Dual-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 Qualified Service Frame transmit count at the Controller MEP at the time of CCM transmission, the Qualified Service Frame transmit count in the last CCM frame received from the Responder MEP, and the Qualified 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 Qualified Service Frame transmit and receive counts contained in the CCM and the Qualified 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:

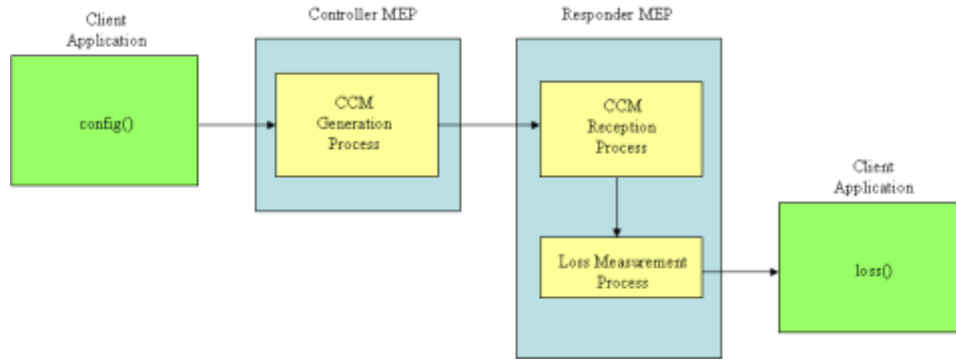


Figure 21 – Dual-Ended Loss Processes

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

Signal	Parameters
<i>config()</i>	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 25 – CCM Generation Process Signals

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

Signal	Parameters
<i>loss()</i>	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 26 – Loss Measurement Process Signals

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

16.6 Dual-Ended Synthetic Loss PM Function

The Dual-Ended Synthetic Loss PM Function is intended to measure One-way synthetic FLR, and is specified for use in a point-to-point or a multipoint service.

One message is defined to enable a uni-directional mechanism, or dual-ended process, to exchange sequence numbers. The One-way Synthetic Loss Message (1SL) conveys a sequence number which is incremented by 1 by the Controller MEP for each 1SL frame transmitted.

The Sink MEP can estimate One-way service frame loss by calculating the loss of the synthetic 1SL frames, using the sequence numbers in a series of received 1SL frames. Gaps in sequence numbers indicate frames lost. To determine synthetic frame loss over a given interval of time, it is necessary to send a number of 1SL frames over that period, and monitor the received 1SL frames. The accuracy of the measurement depends on the number of 1SL frames sent, as described in Appendix D - Statistical Considerations for Synthetic Loss Measurement (Informative).

Frame generation and reception processes are defined for 1SL. In addition, a single 1SL Source Control Process and a single 1SL Sink Control Process are defined. The 1SL Source Control Process coordinates 1SL generation to a given destination at a given SOAM PM CoS ID and periodicity. The 1SL Sink Control Process coordinates 1SL reception from a given source. On termination of the 1SL Sink Control process, measures are returned that reflect the One-way synthetic frame loss over the lifetime of the process. The following figure illustrates these processes:

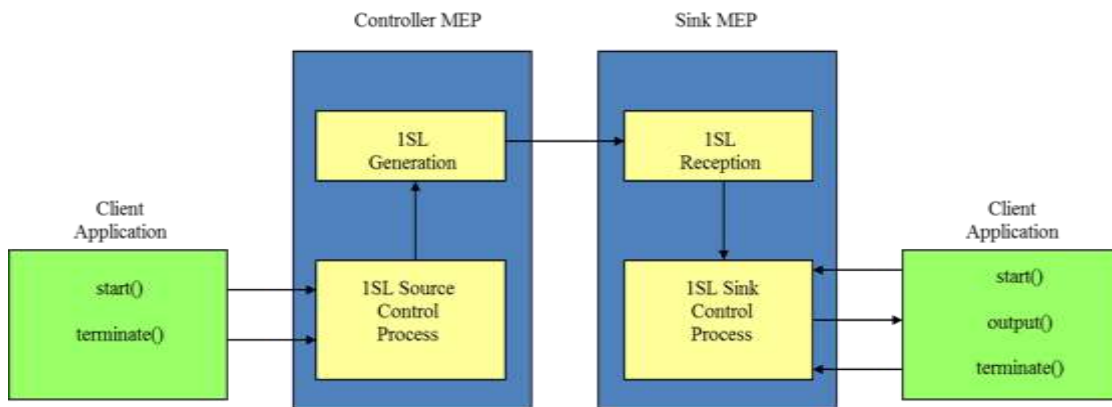


Figure 22 – Dual-Ended Synthetic Loss Processes

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

Signal	Parameters
<i>start()</i>	DA (destination unicast or Class 1 multicast MAC address)
	VLAN PCP (0..7, not applicable if untagged)
	Period for 1SL generation (ms)
	TestID
	Length
<i>terminate()</i>	None

Table 27 – 1SL Source Control Process Signals

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

Signal	Parameters
<i>start()</i>	SA (source unicast MAC address)
	TestID
<i>output()</i>	Total ISL frames transmitted (based on the sequence numbers)
	Total ISL frames not received
<i>terminate()</i>	None

Table 28 – ISL Sink Control Process Signals

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

16.7 PM Session Identifiers

In the architecture of a PM Function, there is typically a Control Process (for a given MEP) that interfaces with an EMS/NMS, 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. In each case, the identifier is specific to a given Controller MEP, so both the MEP and session identifier need to be specified.

Note that not all PDUs contain a Test ID field (e.g., LMM PDUs do not); therefore one of the alternative types of session identifier must be used in these cases.

Thus, signals into the Control Process for a given MEP (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.

17. Appendix B – Life Cycle Terminology (Informative)

The following diagrams show how the life cycle terminology (see section 10) for a PM Session is used in this document. While measurements are being taken for a PM Session, 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 SOAM PM CoS ID on a specific ME.

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.

17.1 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, Measurement Interval, Repetition Period, Start Time (which is always ‘immediate’ for Proactive PM Sessions), and Stop Time (which is always ‘forever’ for Proactive PM Sessions).

The SOAM PM PDUs associated with the PM Session are transmitted every “Message Period”. Data in the form of counters is 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 .

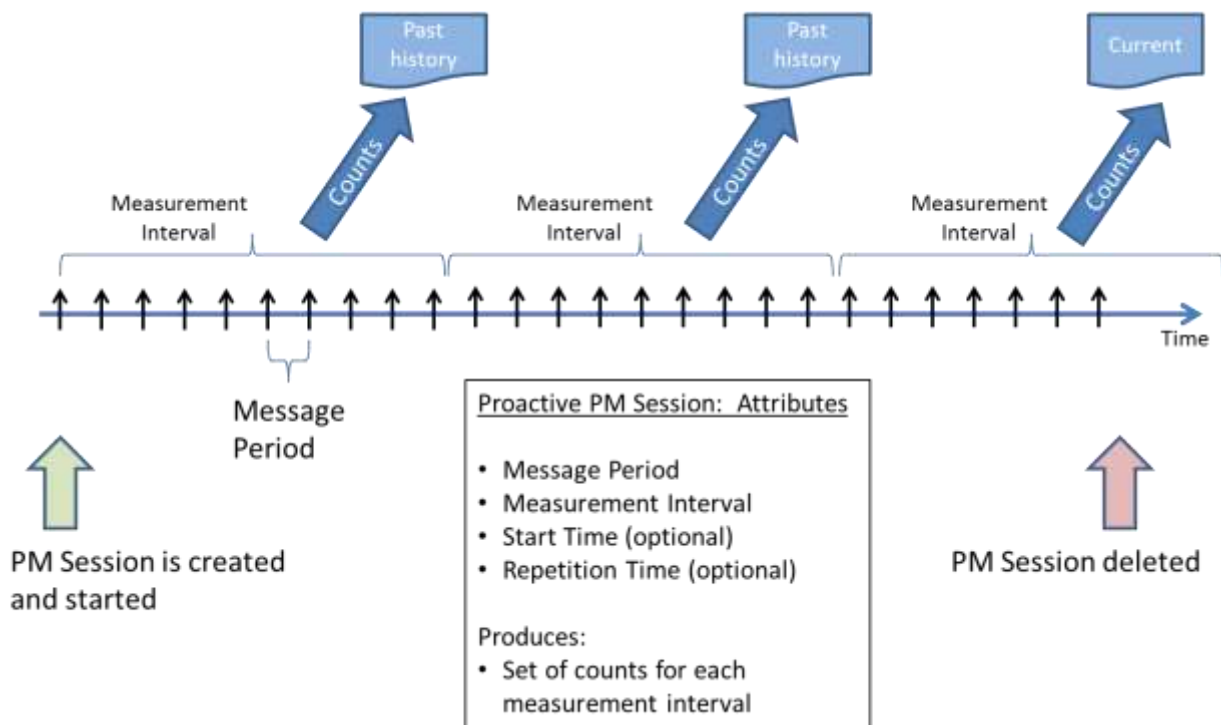


Figure 23 – Measurement Interval Terminology

17.2 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 PM 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.

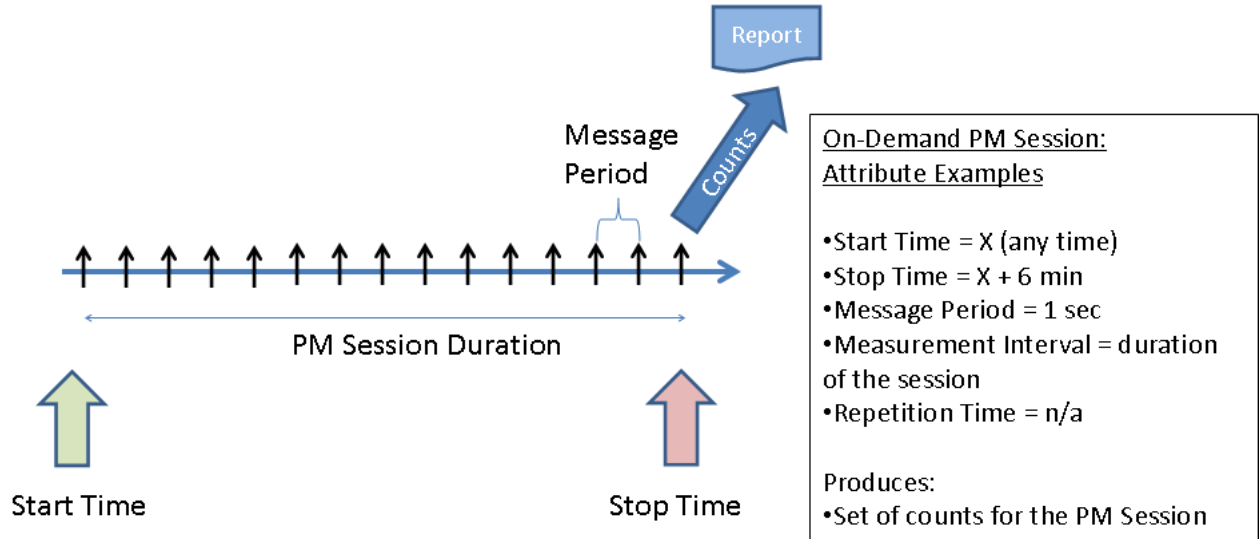


Figure 24 – Illustration of non-Repetitive, On-Demand PM Session

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).

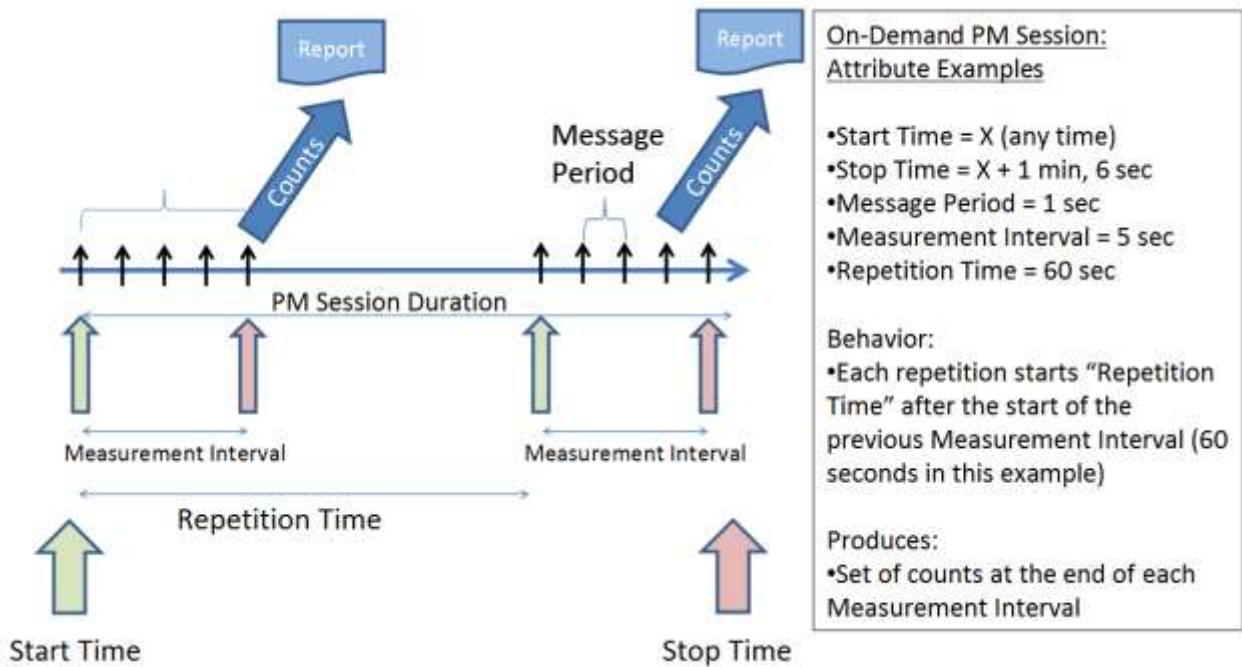


Figure 25 – Example of Repetitive On-Demand PM Session

17.3 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 26 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 26 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.

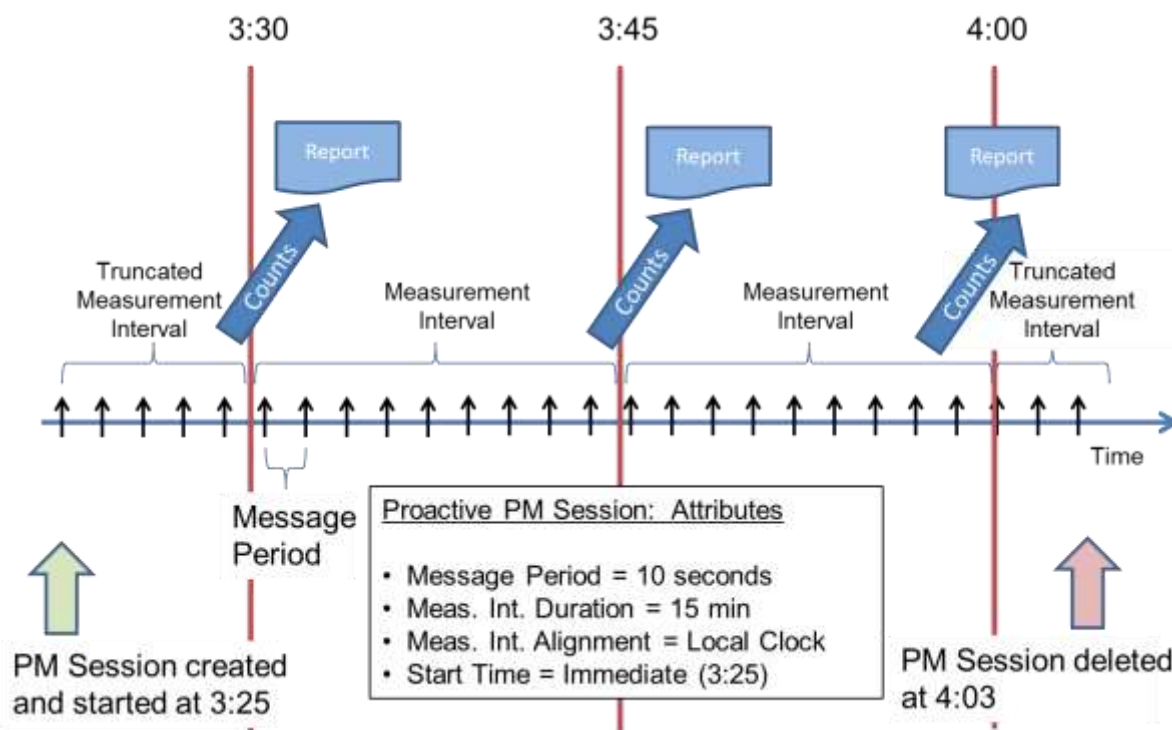


Figure 26 – Example Proactive PM Session with Clock-Aligned Measurement Interval

17.4 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 27. Figure 27 depicts an On-Demand PM Session, but the same principles apply to a Proactive PM Session whose Repetition Time is not equal to “none”.

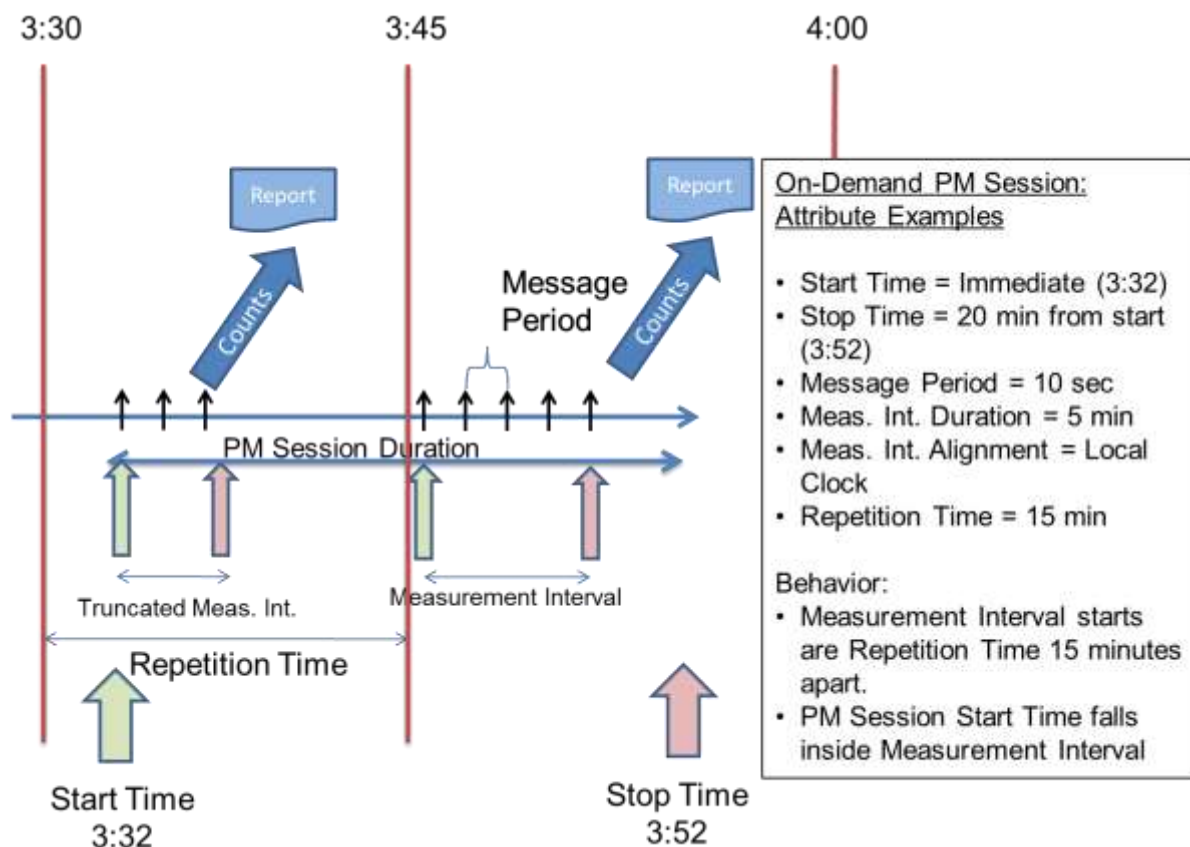


Figure 27 – Example On-Demand PM Session with Clock-Aligned Measurement Interval

In Figure 27, 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 28, the data from the second Measurement Interval would not be flagged as suspect.

Figure 27 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 28. Again, while Figure 28 depicts an On-Demand PM Session, similar principles apply to a Proactive PM Session whose Repetition Time is not equal to “none”.

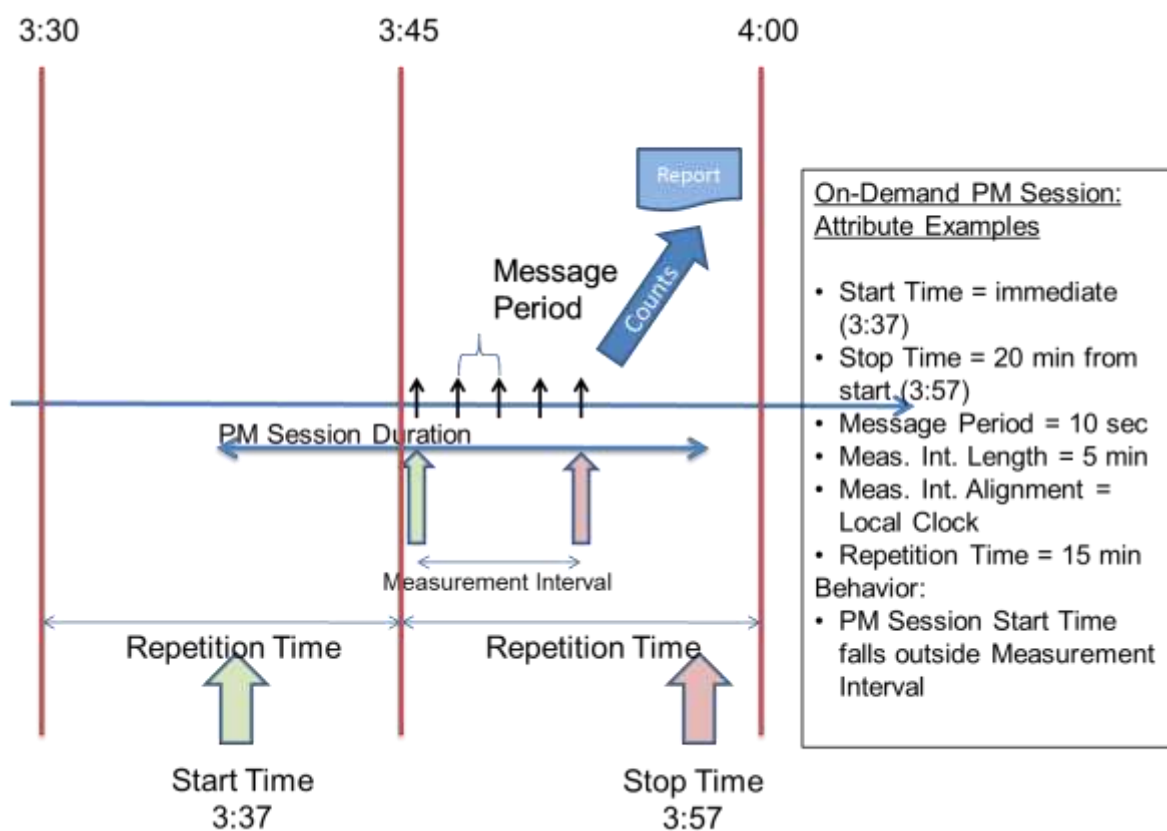


Figure 28 – Second Example of On-Demand PM Session with Clock-Aligned Measurement Interval

In Figure 28, 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 28 might be identical to the parameters given in Figure 27, 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.

18. Appendix C – Measurement Bins (Informative)

The MEF 10.3 [12] and MEF 26.1 [18] 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, inter-frame delay variation, and frame delay range.

18.1 Description of Measurement Bins

As described in section 10.2.2, 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 bins:

Bin 0 if $x < D_1$

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

Bin n if $D_n \leq x$

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, $B_1, .. ,B_m$, 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:

Bin 0 if $v < V_1$

Bin i if $V_i \leq v < V_{i+1}$

Bin m if $V_m \leq x$

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

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

18.2 One-way Frame Delay Performance

As defined in MEF 10.3 [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:

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

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.

18.3 One-way Inter-Frame Delay Variation Performance

As defined in MEF 10.3 [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:

$$P_p(v) < V \text{ if and only if } \sum_{i=j}^m B_i < (1 - p)R_y$$

18.4 One-way Frame Delay Range Performance

As defined in MEF 10.3 [12] the One-way Frame Delay Range Performance is met for an EI pair if $Q_h(x) = P_h(x) - P_0(x) < Q$ where x is the One-way frame delay, h is a high percentile such that $0 < h \leq 1$, $P_0(x)$ is the 0th percentile (i.e., the minimum) 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 $h = 1$ then $P_h(x) = \text{maximum}(x)$.

Note that requirements for measurements of minimum and maximum One-way delay are found in section 11.1. 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 or OVC is designed.

There are two cases to consider, depending on the value of h .

18.4.1 Case 1: $Q_1(x)$

In the case where $h = 1$ then by definition $Q_1(x) = \max(x) - \min(x)$ and bins are not required to determine if the range objective is met:

$$Q_1(x) < Q \text{ if and only if } \max(x) - \min(x) < Q$$

18.4.2 Case 2: $Q_h(x)$

In the case where $h < 1$ 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 . The Equation above for One-way Frame Delay gives us a way to determine if the upper bound is less than a known value:

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

And so we can conclude:

$$\text{if } \sum_{i=j}^n B_i < (1 - h)R \text{ and } c < \min(x) \text{ then } Q_h(x) < Q$$

In other words, the measured range $Q_h(x)$ is less than the objective Q , and so the range objective is met.

19. 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.

19.1 Synthetic Frames 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.

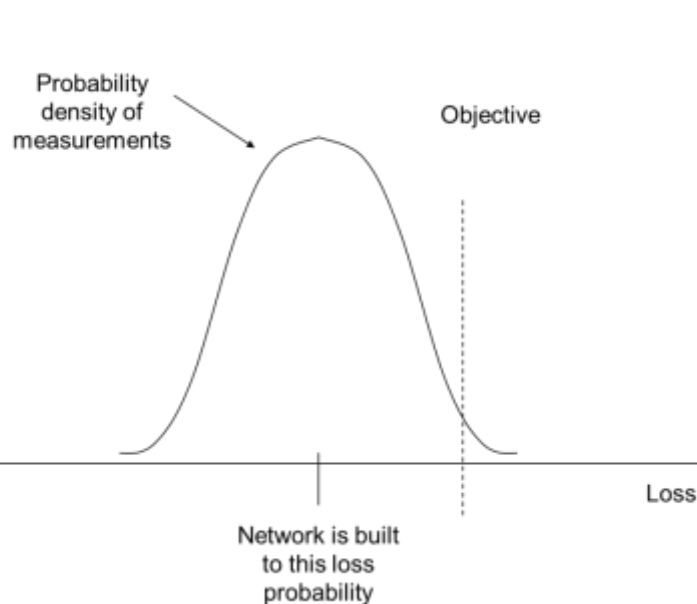


Figure 29 – Hypothesis Test for Synthetic Frame Loss Measurements

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.

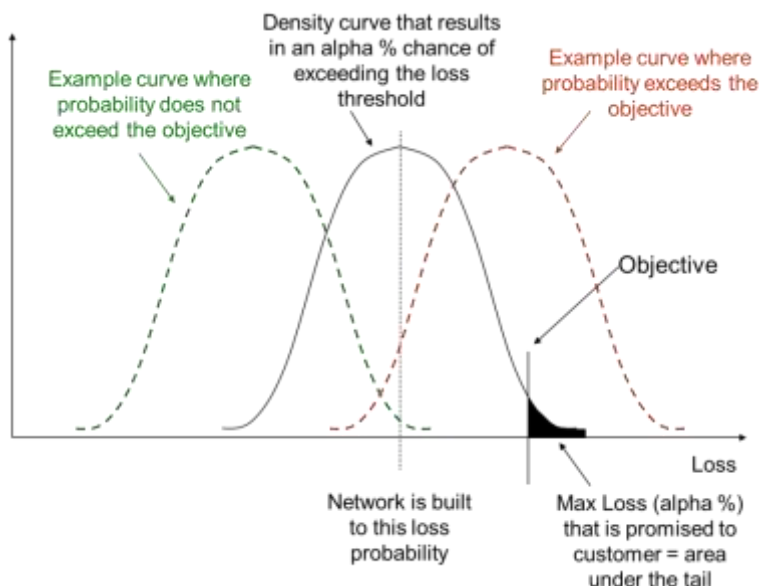


Figure 30 – Density Curve and Probability of Exceeding the Objective

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 31 – Synthetic Loss Performance Example 1

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

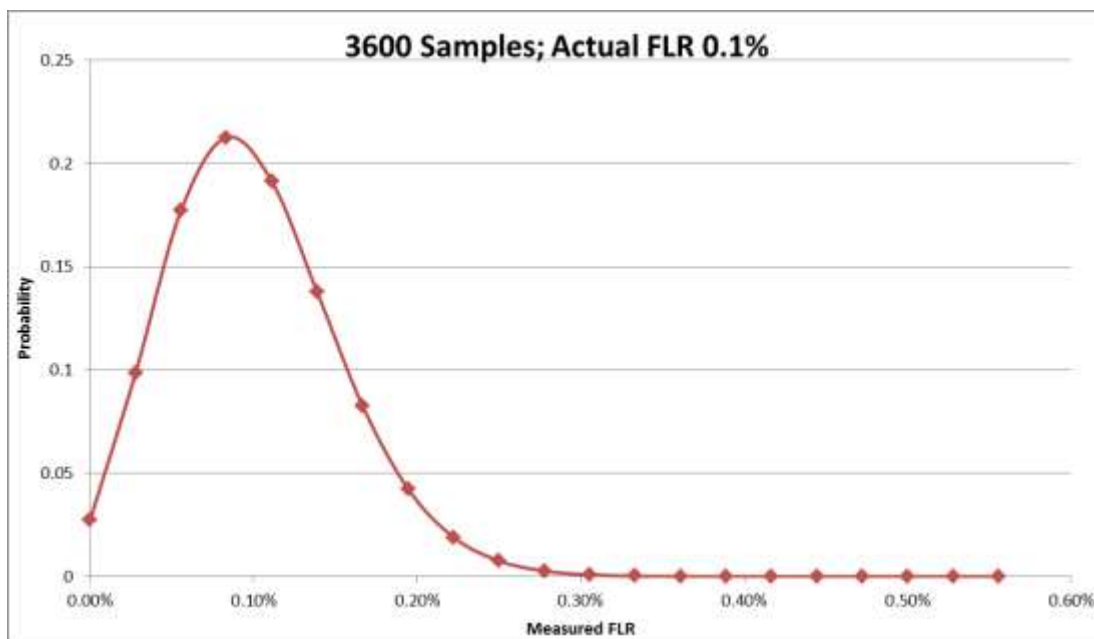


Figure 32 – Synthetic Loss Performance Example 2

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.

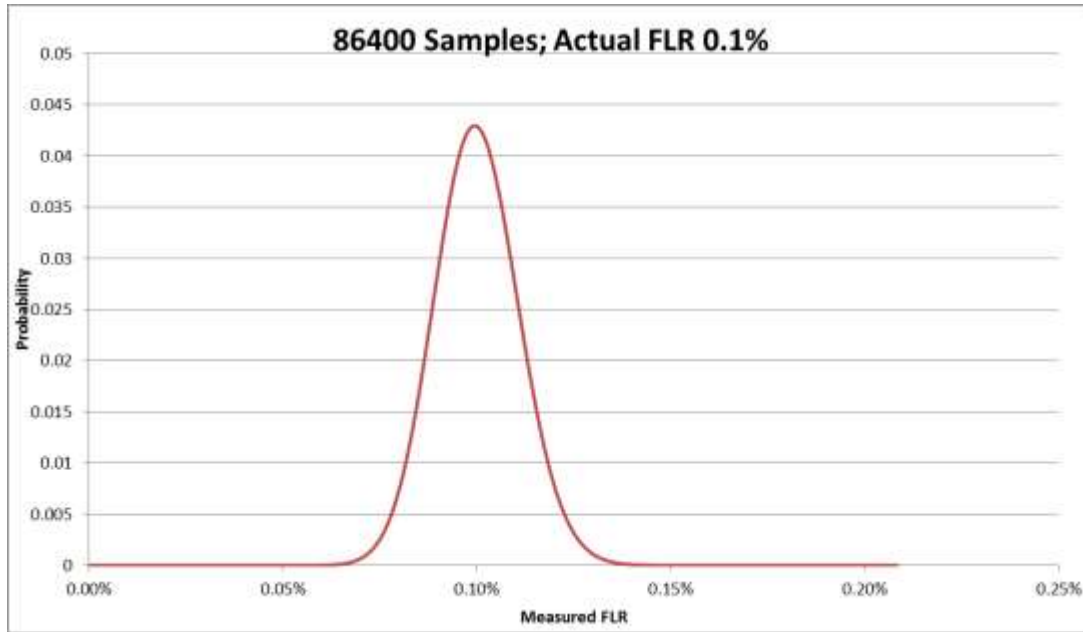


Figure 33 – Synthetic Loss Performance Example 3

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 34 – 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:

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

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

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:

$$\frac{\sigma}{\mu} = \frac{\sqrt{npq}}{np} = \sqrt{\frac{q}{np}} = \sqrt{\frac{q}{p}} * \frac{1}{\sqrt{n}}$$

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 n 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 the Coefficient of Variation. Values are highlighted where the CoV is less than 0.2. This value is proposed as a reasonable bound.

	n	p	μ_{FLR}	σ_{FLR}	CoV
1 hour	3600	0.01	1.000%	0.1658%	0.1658
	3600	0.001	0.100%	0.0527%	0.5268
	3600	0.0001	0.010%	0.0167%	1.6666
	3600	0.00001	0.001%	0.0053%	5.2704
24 hour	86400	0.01	1.000%	0.0339%	0.0339
	86400	0.001	0.100%	0.0108%	0.1075
	86400	0.0001	0.010%	0.0034%	0.3402
	86400	0.00001	0.001%	0.0011%	1.0758
1 month	2592000	0.01	1.000%	0.0062%	0.0062
	2592000	0.001	0.100%	0.0020%	0.0196
	2592000	0.0001	0.010%	0.0006%	0.0621
	2592000	0.00001	0.001%	0.0002%	0.1964

Table 29 – CoV Calculations with Message Period 1s

	n	p	μ_{FLR}	σ_{FLR}	CoV
1 hour	36000	0.01	1.000%	0.0524%	0.0524
	36000	0.001	0.100%	0.0167%	0.1666
	36000	0.0001	0.010%	0.0053%	0.5270
	36000	0.00001	0.001%	0.0017%	1.6667
24 hour	864000	0.01	1.000%	0.0107%	0.0107
	864000	0.001	0.100%	0.0034%	0.0340
	864000	0.0001	0.010%	0.0011%	0.1076
	864000	0.00001	0.001%	0.0003%	0.3402
1 month	25920000	0.01	1.000%	0.0020%	0.0020
	25920000	0.001	0.100%	0.0006%	0.0062
	25920000	0.0001	0.010%	0.0002%	0.0196
	25920000	0.00001	0.001%	0.0001%	0.0621

Table 30 – CoV Calculations with Message Period 100ms

20. 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.

20.1 Summary of Loss Measurement

The ITU-T ETH-LM function is defined in ITU-T G.8013/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 Qualified 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 Qualified 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 Frame⁶ into the flow of Service Frames. The SOAM PM Frame contains the value of $T_i(1)$ in the appropriate field of the SOAM PM PDU. When the SOAM PM Frame 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 Frame into the flow of Service Frames. The SOAM PM Frame contains the value of $T_i(2)$. When the SOAM PM Frame 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 Qualified Service Frames transmitted by the Ingress MEP i between the transmission of the two SOAM PM Frames is $\Delta T_i = T_i(2) - T_i(1)$.

Similarly, the number of Qualified Service Frames received by the Egress MEP j between the receipt of the two SOAM PM Frames 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 Frames which is defined as $\Delta L_{ij} = \Delta T_i - \Delta R_j$ and Frame Loss Ratio is $\Delta FLR_{ij} = \Delta L_{ij} / \Delta T_i$.

20.2 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 frames are sent from MEP 1 to MEP 2. Assume that over the interval of interest, 30 Qualified Service Frames enter the MEG at MEP 1, 20 of the Qualified Service Frames are delivered to MEP 2, and the other 10 Qualified Service Frames are delivered to MEP 3. No Qualified Service Frames are lost in this example. In this example, $\Delta T_1 = 30$ frames, $\Delta R_2 = 20$ frames, $\Delta L_{12} = 10$ frames, and $FLR_{12} = .33$, which is wrong, it should be 0.

⁶ The SOAM PM Frame can be a LMM, LMR, or CCM frame depending on the specific technique being used.

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

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

PM-3 can only work in a point-to-point MEG (that is, one with only 2 MEPs), and only 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 Qualified 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.

20.3.1 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.

20.3.2 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 Qualified 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 Qualified 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 Qualified Service Frame as negative loss.

20.3.3 Frames Consumed by an Internal MAC

If a unicast frame enters the MEG and is counted by the Ingress MEP as a transmitted Qualified 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 Qualified 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.

20.3.4 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 Qualified Service Frame by the Egress MEP then that frame will be incorrectly counted as negative loss.

21. Appendix F - Frame Loss Count Accuracy (Informative)

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.

21.1 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 35, are above the queuing entities (detailed view can be found Figure 22-4 of IEEE 802.1Q-2014 [22]). Hence, in the egress direction, these MPs cannot distinguish between discards (on a per VID basis) in the queuing entities and discards in the CEN 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.

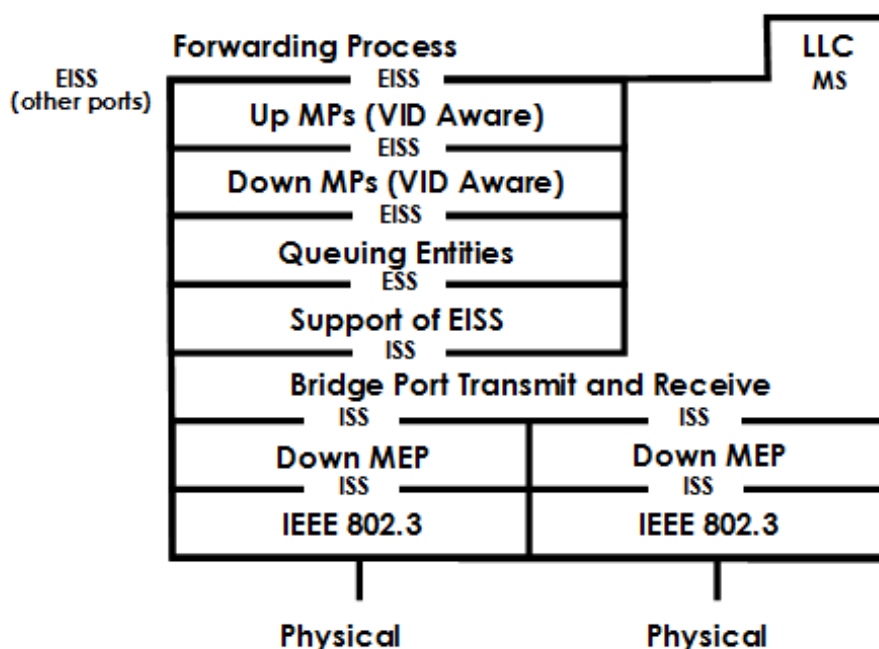


Figure 35 – 802.1Q Bridge Port

22. 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 performance 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 36.




Figure 36 – 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 PM Frame.

One suggested practical approach as shown in Figure 36 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 EMS/NMS 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 EMS/NMS 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 Measurement Interval may be invalid, and the measurements can be ignored by the EMS/NMS. 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.

22.1 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.

22.1.1 Minimum Delay Becomes Significantly Smaller

When the delay becomes significantly smaller, as is shown in MI 2 below in Figure 37, 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 EMS/NMS to simply consider the FDR measurements of that interval as being invalid, and to ignore them.




Figure 37 – Reduction in Minimum Delay, due to Network Topology Change

22.1.2 Minimum Delay Becomes Significantly Larger

When the delay becomes significantly larger, as is shown in MI 6 below in Figure 38, 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.

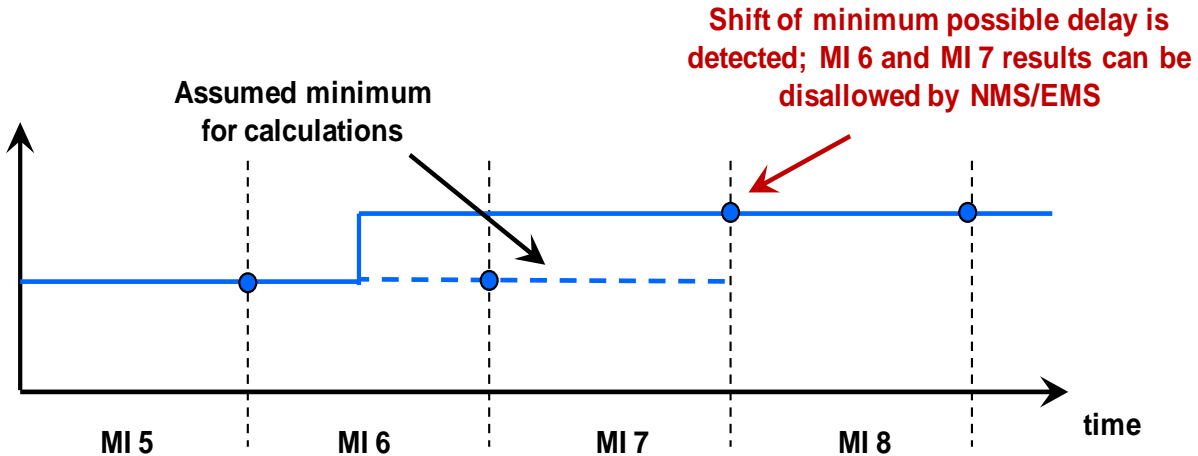


Figure 38 – Increase in Minimum Delay, due to Network Topology Change

22.2 Impact of Lack of ToD Synchronization

When performing One-way measurements using Single-Ended or Dual-Ended Delay Measurement without ToD synchronization between the MEPs, negative frame delay measurements can be seen due to differences in the ToD for each MEP. An example of this is shown in Figure 39.

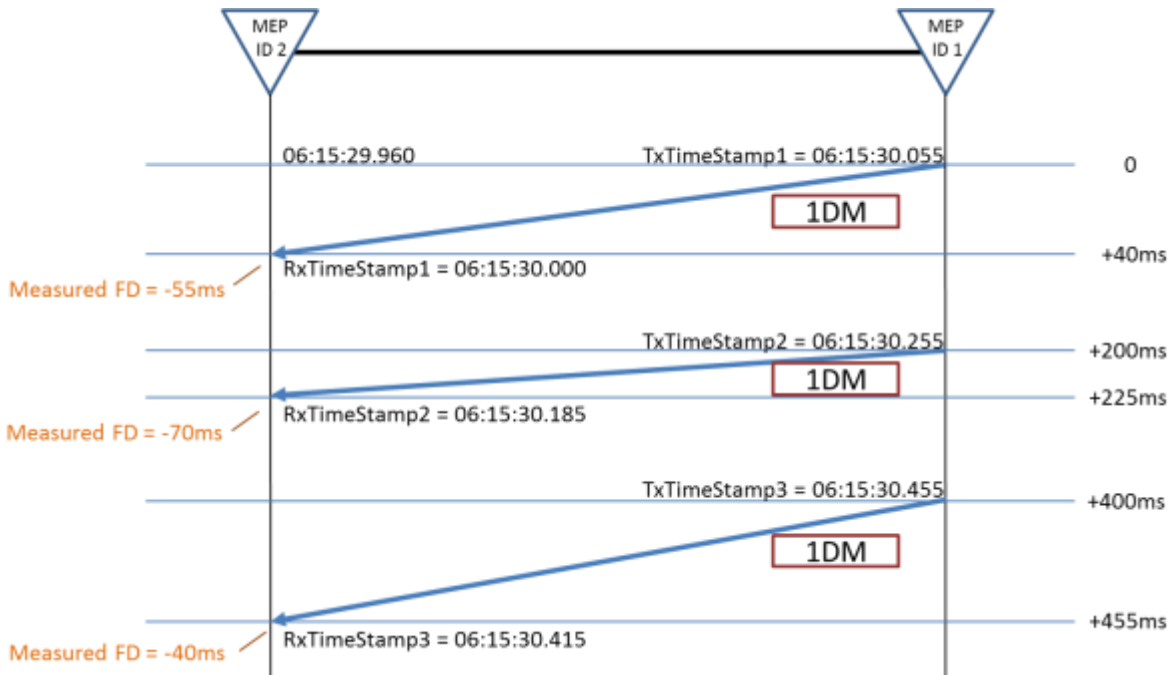


Figure 39 – Lack of ToD Synchronization

In Figure 39, three 1DMs are shown. At the time when the first 1DM is transmitted, the ToD clock at MEP ID 1 reads 06:15:30.055 and the ToD clock at MEP ID 2 reads 06:15:29.960. The FD measured for the first 1DM, using $RxTimestamp1 - TxTimestamp1$, is -55ms since $TxTimestamp1 > RxTimestamp1$. When determining the minimum FD for FDR in this situation,

a “less negative” FD is considered an increase in delay and a “more negative” FD is considered a decrease in delay. Using the example in Figure 39, the FD measured for the second 1DM, $RxTimeStamp2 - TxTimeStamp2$, is -70ms which indicates that the frame arrived 15ms faster than the first 1DM frame. The FD measured for the third 1DM, $RxTimeStamp3 - TxTimeStamp3$, is -40ms which indicates that the frame arrive 15ms slower than the first 1DM frame.

Implementations that are measuring FDR without ToD synchronization are expected to take this into account and react accordingly to negative FD measurements.

23. Appendix H – Notes on Dual-Ended PM Functions (Informative)

When compared to Single-Ended PM Functions, Dual-Ended PM Functions have some advantages for multipoint MEGs, especially those with a large number of MEPs. For point-to-point MEGs, there is little difference between the two mechanisms.

For multipoint MEGs, Dual-Ended PM Functions can help to simplify the SOAM PM configuration, and may reduce the total amount of SOAM traffic in the network if multicast 1SL or 1DM frames are used. These advantages are best illustrated by an example, as shown below.

Note that MEF 35 limits SLM and DMM frames to unicast, due to the additional complexity that would be required in the Controller to handle multiple responses if multicast were used. However, this limitation does not apply to 1SL or 1DM.

An example for a multipoint MEG and MEPs as follows:

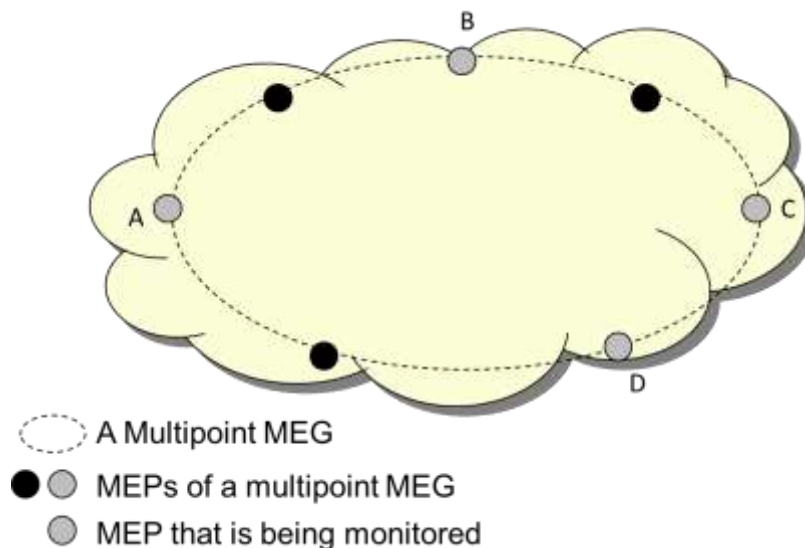


Figure 40 – Example of a Multipoint MEG and MEPs

When a Single-Ended PM Function is used, one example of the configuration can be:

- MEP at A: 3 Controllers + 0 Responder
- MEP at B: 0 Controller + 3 Responders
- MEP at C: 1 Controller + 2 Responders
- MEP at D: 2 Controllers + 1 Responder

The corresponding PDU flows are shown in Figure 41 - this uses the example of Synthetic Loss Measurements with SLM/SLR PDUs, but the same would apply for Delay Measurements with DMM/DMR PDUs.




Figure 41 – Single-Ended Synthetic Loss Example

Since all SLMs are unicast frames, the PDU flows will be 6 SLM flows and 6 SLR flows. In general case, it would be $1/2 * N * (N-1)$ flows of SLM and SLR respectively, where N is the number of MEPs being monitored of the MEG.

Notice that the configuration of each MEP is different. It requires more careful planning in order to cover all directions for all monitored MEPs.

Alternatively, if we use a Dual-Ended PM Function with multicast 1SL or 1DM frames, the configuration can be much simpler and well balanced, compared to using a Single-Ended PM Function:

- MEP at A: 1 Controller + 3 Sinks
- MEP at B: 1 Controller + 3 Sinks
- MEP at C: 1 Controller + 3 Sinks
- MEP at D: 1 Controller + 3 Sinks

The corresponding flows are shown in Figure 42 - again we use the example of Synthetic Loss Measurement with 1SL PDUs, but it applies equally to Delay Measurement with 1DM PDUs.




Figure 42 – Dual-Ended Synthetic Loss Example

Notice that the configuration of each MEP is the same.

Using multicast 1SL frames requires each Controller MEP to send multicast 1SL and each Sink MEP to accept and process the measurement. The number of multicast flows will be N (4 as shown in Figure 42). Depending on where the multicast starts, the actual PDU flows in the network will be at most $N*(N-1)$, where N is the number of MEPs being monitored of the MEG.

24. Appendix I – Calculation of SLS Performance Metrics (Informative)

This document defines the data sets that Network Elements provide to EMS/NMS, while other MEF specifications and applications need to obtain the performance metrics for SLS. This appendix provides some guidelines for how to calculate SLS performance metrics, using data sets as inputs.

The SLS performance metrics are defined in terms of the performance of every Qualified Service Frame; however, the data sets are primarily based on time-based samples, as described in section 8.3. In the remainder of this appendix we assume that time-based sampling is used, and analyze how the data sets can be used to calculate the SLS metrics on that basis.

The data sets are Measurement Interval based. Traditionally, the duration of a Measurement Interval is 15 minutes or 24 hours. This document requires at least that 15 minute Measurement Intervals are supported. When reaching the end of a Measurement Interval, the data set for the current measurement interval is moved to the list of historic Measurement Intervals. The EMS/NMS can retrieve a block of historic data sets from the NEs (or the NEs send the data sets). Usually the performance metrics are measured against the SLS over a much longer time period T , typically one month or so. The processing of performance metrics for an SLS can be done by EMS, NMS, or even OSS. Therefore, the data sets from multiple Measurement Intervals are used for calculating the performance metrics over period T . In the following, we discuss how to obtain the following performance metrics for SLS, using PM-1 defined data sets:

- One-way FD
- One-way FDR
- One-way MFD
- One-way IFDV
- One-way Availability
- One-way HLI
- One-way CHLI
- One-way Group Availability

24.1 One-way Frame Delay

One-way FD can be calculated from the data sets (i.e. counts of each Measurement Bin), when there are n Measurement Intervals in T . If $FD(T) (\%) \leq \text{Objective FD\%}$, then the performance is considered to meet the SLS for time period T . The FD over T can be calculated from:

$$FD(T) = \frac{\sum^n(\text{Total counts of Meas. Bins in the MI that meet the objective})}{\sum^n(\text{Total counts of all Measurement Bins in the MI})}$$

Note that the Measurement Bin thresholds must be chosen such that the FD objective is aligned with the boundary between two bins, as described in Appendix C – Measurement Bins (Informative).

The same calculation applies to all other SLS performance metrics for which Measurement Bins are used, including One-way FDR and One-way IFDV.

24.2 One-way Mean Frame Delay

By MEF 10.3 definition, the MFD shall be:

$$\frac{\sum(\text{Measured frame delays of } T)}{\sum(\text{Total DM PDUs of } T)}$$

Since the MFD will be calculated based on data sets, a possible naïve solution where there are n MIs in T is:

$$\frac{\sum^n(\text{MFD of MI})}{n}$$

This may not give the same results as the MEF 10.3 definition above. For example, suppose there are 3 MIs with delay measurements as follows:

MI(1) has 3 DM: 4,5,6; → MFD(1) = 5

MI(2) has 4 DM: 3,5,3,5; → MFD(2) = 4

MI(3) has 3 DM: 4,6,8; → MFD(3) = 6

Using the naïve solution: $MFD(T) = (5 + 4 + 6)/3 = 5$

However, using the MEF 10.3 definition: $MFD = (4 + 5 + 6 + 3 + 5 + 3 + 5 + 4 + 6 + 8)/10 = 4.9$

One possible solution that gives the same result as the definition would be:

$$\frac{\sum^n(\text{Frame delays of all DM in an MI})}{\sum^n(\text{Total DM PDUs of MI})}$$

The reason that some MIs have fewer delay OAM PDUs may be because those delay PDUs were either lost or had excessively large delays (i.e. such that they are considered as lost at the end of a Measurement Interval). In that case, it is undesirable to account for the missing results in calcu-

lating the One-way FD over T , as that could skew the calculation away from periods of poor performance. Therefore, the naïve solution provides a reasonable method of calculating One-way MFD over T from the Measurement Interval data sets.

24.3 One-way Frame Loss Ratio

Based on the Tx and Rx frame counts of the data sets for n MIs during T , the One-way Frame Loss Ratio over T can be obtained by:

$$FLR(T) = \frac{\sum^n((Tx \text{ frame count for the MI}) - (Rx \text{ frame count for the MI}))}{\sum^n(Tx \text{ frame count for the MI})}$$

24.4 One-way Availability

Based on the count of Δt intervals evaluated as Available and Unavailable in the data sets for n MIs during T , the One-way Availability over T can be obtained by:

$$AV(T) = \frac{\sum^n(\text{Count of Available } \Delta t\text{'s in the MI})}{\sum^n((\text{Count of Available } \Delta t\text{'s in the MI}) + (\text{Count of Unavailable } \Delta t\text{'s in the MI}))}$$

An alternative approach to calculate Availability is to use Availability State Change notifications. The nature of such transitions is that during a time interval T , they must alternate between Available-to-Unavailable transitions and Unavailable-to-Available transitions. To find the total Available Time during interval T , we identify the segments of Available Time during T , calculate the duration of each segment, and add these together. The duration of each segment can be calculated by taking the difference between the timestamp of the Available-to-Unavailable transition at the end of the segment and the timestamp of the Unavailable-to-Available transition at the beginning of the segment.

If the state is Available at the start of T , then there is an additional segment of Available Time starting at the start of T and ending at the timestamp of the first Available-to-Unavailable transition. Similarly, if the state is Available at the end of T , then there is an additional segment of Available Time starting at the timestamp of the last Unavailable-to-Available transition and ending at the end of T .

An example is illustrated in Figure 43 below:




Figure 43 – Example of Availability State Transitions during T

In this example, there are four segments of Available Time during T , starting at times 20, 50, 100 and 130. To find the total Available Time, we must add together the durations of each of these four segments.

More formally:

Let T_{start} be the time in UTC at the start of time period T , and T_{end} be the time in UTC at the end of time period T .

During time period T , there may be Unavailable-to-Available transitions, and there may be Available-to-Unavailable transitions. In addition, if the time period begins in an Available state at T_{start} , then there is considered to be a "virtual" Unavailable-to-Available transition at time T_{start} . If the time period begins in Unavailable state, then there is no "virtual" transition at time T_{start} . Similarly, if the time period ends in Available state, then there is considered to be a "virtual" Available-to-Unavailable transition at time T_{end} . If the time period ends in Unavailable state, then there is no "virtual" transition at time T_{end} .

It can be shown that, including the virtual transitions, over time period T , the first transition is always an Unavailable-to-Available transition, and the last is an Available-to-Unavailable transition. Furthermore, there are exactly the same number of Unavailable-to-Available transitions as Available-to-Unavailable transitions.

It can further be shown that, where the timestamps of the virtual transitions as described above are included, the total Available Time in interval T is:

$$\text{Total Available Time} = \sum(\text{Available} - \text{to} - \text{Unavailable timestamps}) - \sum(\text{Unavailable} - \text{to} - \text{Available timestamps})$$

Therefore the Availability over time period T :

$$AV(T) = \frac{\sum(\text{Available} - \text{to} - \text{Unavailable timestamps}) - \sum(\text{Unavailable} - \text{to} - \text{Available timestamps})}{T}$$

This is best illustrated with reference to the example shown in Figure 43 above. In this example, the T starts at 20 (as UTC timestamp), ends at 200, and the state transitions are shown with their time stamps. Since the state is Available at the start of T , there is a "virtual" transition at time 20. It can be seen that, including this virtual transition, there are four transitions from Unavailable to Available, and four transitions from Available to Unavailable. The total available time over T can be found by taking the difference between each pair of transitions, and summing these:

$$\begin{aligned} \text{Available Time} &= (30-20) + (60-50) + (120-100) + (180-130) \\ &= (30 + 60 + 120 + 180) - (20 + 50 + 100 + 130) \end{aligned}$$

24.5 One-way High Loss Interval

Based on the count of HLI of the data sets for n Measurement Intervals during T , the count of One-way High Loss Intervals in T can be obtained by:

$$HLI(T) = \sum^n (\text{Count of HLI in the MI})$$

24.6 One-way Consecutive High Loss Interval

Based on the count of CHLI of the data sets for n MIs during T , the count of One-way CHLIs during T can be obtained by:

$$CHLI(T) = \sum^n (\text{Count of CHLI in the MI})$$

Note that care is needed when a CHLI crosses the boundary of a Measurement Interval. An example is shown in the figure below.




Figure 44 – Example showing a CHLI crossing a Measurement Interval boundary

In this example, there are two CHLIs ending in MI#2, and therefore the count of CHLIs for MI#2 is 2. One consecutive HLI block is completely within the MI#2 boundary while another consecutive HLI block crosses the boundary between MI#1 and MI#2. The NE must make sure that the count of CHLIs in the data set for MI#2 includes the CHLI block that crosses the Measurement Interval boundary.

24.7 One-way Group Availability

Based on MEF 10.3, the One-way Group Availability is illustrated in the following example. This shows two MEs, from M to B and from M to C. The Availability state is determined for each Δt in each ME, and combined to find the Group Availability state for each Δt . The Group Availability state for each Δt is Available if the Availability state for that Δt is Available for either of the MEs.




Figure 45 – Group Availability

In this example, if the measurements are initiated at B and C, and if B and C are in different NEs, then in order to calculate Group Availability, the Measurement Interval data sets retrieved from B and C would have to carry the Availability state for each Δt . If the Measurement Interval duration is 15 min and Δt is 1 sec, there are 900 Δt intervals in each data set. This is a prohibitively large amount of data to be added into the data sets.

Assuming the Measurement Interval boundaries and the clocks on the NEs are well aligned, the Availability State Change notifications provide the boundaries of Available Time periods and Unavailable Time periods for each ME, as described in section 24.4 above. The EMS/NMS can apply an OR operation over the periods of Available Time for each ME to obtain the Group Availability.

In other words, if t is any interval in T , then we can define:

$$AG(t) = \begin{cases} 1, & \text{if the state is Available during } t \text{ for any ME in the Group} \\ 0, & \text{otherwise} \end{cases}$$

Then the Group Availability over T is:

$$AVG(T) = \frac{\sum_{t \in T} AG(t)}{T}$$

This is illustrated in the figure below.




Figure 46 – Group Availability

In situations where the clocks are not well aligned between the devices transmitting the Availability state change notifications, the NMS can add a positive or negative time factor to the timestamps of one pair's data to align them with the other pair's data. For example, in Figure 46 above, if M-B reports the ToD as x and M-C reports the ToD as $x + 5$ minutes then the NMS can internally adjust the ToD received from M-C to align with the ToD received from M-B (i.e. by subtracting 5 minutes). This adjustment takes place within the NMS and does not impact the clock settings on the elements transmitting the notifications. Instead, the NMS interprets the timestamp on the notification to mean the adjusted time. In order to calculate the necessary offset, the NMS needs to be aware of the ToD on each device; the method used to determine this is outside the scope of this document.

25. Appendix J – Statistical Considerations for Availability (Informative)

There are a number of parameters that must be considered for Availability measurements. This appendix contains statistical analysis and considerations for how suitable combinations can be chosen.

The parameters are:

- P , the SLM transmission period (e.g. 100ms)
- Δt , the short time interval over which the Availability flr is calculated (e.g. 10s)
- C , the Availability flr threshold (e.g. 0.1)
- n , the number of consecutive Δt intervals in which the threshold C must be exceeded to cause a change in the Availability state (e.g. 10)

The subsections below consider various aspects of Availability measurement, followed by some examples:

- Accuracy of Availability flr
- Precision of Availability flr
- Probability of Availability state change
- Probability of Availability state change during a Period of High Loss
- Accuracy and Precision of Availability over T

25.1 Availability flr Accuracy

The Availability flr, used for Availability and Resiliency Performance, is measured using Synthetic Loss Measurement (ETH-SLM) since actual loss measurements based on the service frames (ETH-LM) is not applicable for multipoint EVCs and has a number of other limitations, as described in Appendix E – Notes on the Applicability of PM-3 Solutions (Informative). From the point of view of statistical analysis, one important weakness of traffic-based loss measurement is that there is no control over the sample size, i.e. the number of Qualified Service Frames passing during each Δt . The accuracy of the Availability flr measurement therefore depends on the number of Synthetic Frames, s , transmitted during each Δt interval. This can be easily determined from the SLM frame period, P :

$$s = \frac{\Delta t}{P}$$

As explained in Appendix D - Statistical Considerations for Synthetic Loss Measurement (Informative), the Coefficient of Variation (CoV, c_v) is a good measure for the accuracy of statistical frame loss measurements. If L is the actual probability of frame loss (i.e., the actual FLR), then:

$$c_v = \frac{\sigma}{\mu} = \sqrt{\frac{(1-L)}{sL}}$$

As described in Appendix D, for a given CoV the sample size s needs to go up by the same factor as the loss probability L goes down. This means that to detect a smaller probability of loss, either the period P between the synthetic frames need to be decreased or the sampling interval Δt has to be increased.

For FLR over time period T , a CoV of 0.2 or lower is proposed. However, for Availability and Resiliency a higher CoV can be accepted since the frame loss is measured for multiple Δt intervals n before a change in the Availability state is triggered or a CHLI is reported.

25.2 Availability flr Precision

Another consideration for Availability is the precision of Availability flr measurements. It is clear that if, for example, 100 samples are taken (i.e. $s = 100$), the only possible values for the measured Availability flr are multiples of 1%. The definition of Availability in MEF 10.3 specifies that a given Δt interval is bad if the measured Availability flr is strictly greater than a threshold, C . Therefore, in this example, if $C = 0.01$, it will be exceeded if the Availability flr is 2% or more, i.e. if at least 2 frames out of 100 were lost. However, if C is any value less than 0.01 – for instance 0.00999, 0.005 or even 0 – then it will be exceeded if the Availability flr is 1% or more, i.e. if at least 1 frame out of 100 was lost – in other words, all these values for C have the same effect.

It can be seen that setting C to different values in between multiples of the precision – that is, multiples of $1/s$ – has no effect. In other words, there is no benefit in setting C to any value other than an integer multiple of $1/s$. Note that setting $C = 0$ may, unintuitively, be a reasonable value in certain scenarios, as illustrated below.

25.3 State Change Probability

For Availability and Resiliency Performance the Availability flr measurements are assessed over multiple time intervals and the state is only changed if it provides the same result for at least n consecutive intervals. The requirement for multiple consecutive intervals with same result (above or below the threshold) means that only periods of continued loss are considered, and short bursts are ignored.

The probability of such a state change depends on the probability of exceeding the threshold C during a single Δt (which in turn depends on the accuracy of the measurement as described above), and the number n of consecutive intervals. Threshold C is exceeded when strictly more than k frames are lost in s samples, where $k = \text{floor}(Cs)$.

The probability p_k of detecting at least k lost frames in s samples, if the actual frame loss probability is L – in other words, the probability of exceeding C during a given Δt interval – is:

$$p_k = 1 - \sum_{i=0}^k \binom{s}{i} L^i (1-L)^{(s-i)} \quad \text{where } \binom{s}{i} \text{ is the binomial coefficient.}^7$$

Clearly, p_k increases as L increases, or as C (and therefore k) decreases. In other words, the probability of exceeding C is higher if the actual loss is significantly above the threshold.

The probability p_n of detecting at least k lost frames in n consecutive intervals – which would trigger a state change to Unavailable – is:

$$p_n = p_k^n$$

If a given probability of a state change, p_n , is desired, then a larger n requires a higher probability p_k of crossing the threshold in a given Δt , which depends amongst other things on the sample size s . If the sample size is small (meaning large CoV), then the actual loss must be significantly higher than the threshold C before a state change is likely.

As an example, for $n = 3$, $C = 0.01$ and $s = 200$, and with an actual probability of loss four times higher than C (i.e. $L = 4C = 0.04$), the CoV is 0.35 and the probability p_n for an Availability state change is about 0.96. For $n = 5$ the probability decreases to 0.94, for $n = 10$ it is 0.88 and for $n=20$ it is only 0.78. If a higher value of n is desired, then the CoV must be lower (higher sample size), e.g. in the above example with $s = 400$, the CoV is 0.25 and even with $n = 20$, there is >0.99 probability of changing state.

25.4 State Change Probability over a Period of High Loss

The formula given above for p_n shows the probability of the threshold being exceeded in a given n intervals – for instance, the first n intervals that coincide with a period of high loss. If the period of high loss lasts for longer than n intervals, then there may be multiple blocks of n consecutive intervals, each of which could trigger a change to Unavailable (if the state change has not been triggered by a previous block).

In other words, we can redefine p_n as the probability of detecting a state change to Unavailable “immediately” when the period of high loss starts, where “immediately” means after the first $n \Delta t$ intervals. If the state change is not detected immediately (i.e., if C is not exceeded in all of the first n intervals), it may still be detected later (i.e. as soon as C is exceeded in n consecutive intervals). The longer the high loss period lasts, the more likely it is that a state change to Unavailable will occur at some point. We can define p_u as the probability that a transition to Unavailable will occur at some point if the high loss lasts for a given time, longer than $n \Delta t$.

For example, suppose $C = 0.01$, $n = 3$, $P = 100\text{ms}$, $\Delta t = 10\text{s}$, and the actual loss probability $L = 0.03$. The probability of detecting a state change after $n \Delta t$, i.e. after 30s, is $p_n = 0.52$. However,

⁷ The binomial coefficient, often expressed as “ s choose i ”, is the number of ways, disregarding order, that i objects can be chosen from among s objects.

after 60s the probability of having changed state to Unavailable rises to $p_u = 0.83$ and after 120s to $p_u = 0.98$.

The probability p_u can be calculated as follows. Let $p_x(i)$ be the probability of detecting a state change in the i^{th} Δt interval after the high loss starts, and let N be the number of Δt intervals during the period of high loss. Then:

$$\begin{aligned}
 p_x(0) &= p_n \\
 p_x(i) &= (1 - p_k)p_n && [0 < i \leq \min(n, N - n)] \\
 p_x(i) &= \left(1 - \sum_{j=0}^{i-(n+1)} p_x(j)\right) (1 - p_k)p_n && [n < i \leq N - n]
 \end{aligned}$$

In other words, the probability of detecting a state change immediately (at $i = 0$) is p_n , as described above. After that, a state change is only detected in interval i if the threshold was not crossed during the previous interval, and is crossed in the subsequent n intervals (if it was crossed in the previous interval and the subsequent n , the state change would have already occurred). For intervals after the first n , we must additionally check that a state change did not occur previously during the period of high loss.

The probability p_u is then given by:

$$p_u = \sum_{i=0}^{N-n} p_x(i)$$

25.5 Accuracy and Precision of Availability

Availability over time period T is defined as the proportion of T that was Available Time – in other words, the number of Δt intervals that were Available divided by the total number of Δt intervals in T .

The accuracy of a measurement may be defined as how close a measured value is to the actual value. The accuracy of the Availability measurement depends on detecting a state change to Unavailable immediately (i.e., after n Δt intervals) when a period of high loss starts (and of course of detecting a change back to Available immediately when the loss returns to a low level). The probability of detecting a change to Unavailable immediately is p_n as described above; so it is desirable to have a high value for p_n . As described above, the probability of detecting a state change at all, p_u , can be very high even when p_n is relatively low, if the period of high loss lasts for long enough; however, obviously the later the state change is detected, the less accurate is the count of Δt intervals evaluated as Unavailable and the less accurate is the calculation of Availability over time T .

Note: it is assumed in this analysis that when a period of high loss ends, the actual probability of loss falls to near zero, resulting in a very high probability of detecting the transition from Unavailable to Available immediately. The analysis therefore focuses on the transition from Available to Unavailable.

The precision of a measurement may be defined as the smallest increment that can be measured (also termed granularity or resolution). The precision of the Availability measurement depends on the shortest period that can be determined to be Unavailable, which is $n \Delta t$. In extreme cases where both n and Δt are large, this may be a significant proportion of T . For instance, with $n = 10$ and $\Delta t = 600\text{s}$, and if $T = 1$ week, if there is any period of Unavailable Time then it must be Unavailable for at least 6000s, so the highest possible Availability (other than 100%) is 99.2%.

It is worthwhile noting that while having a high accuracy is desirable, i.e. a high value for p_n , increasing this by increasing Δt (and hence increasing s) doesn't necessarily help. For example, suppose $C = 0.01$, $n = 5$, $P = 100\text{ms}$ and the actual loss probability, L , is 0.04. With $\Delta t = 10\text{s}$, the probability of an immediate state change, i.e. after 50s, $p_n = 0.63$. After 150s, the probability of having changed to Unavailable is $p_u = 0.98$. If one were to increase Δt to 30s so as to improve the accuracy, then the probability $p_n = 0.99$ but of course $n \Delta t$ is now 150s. In other words, in either case if the loss lasts for 150s or more, the chance of a state change to Unavailable is very high; however in the first case, there is still a reasonable chance of detecting Unavailability if the loss lasts only 50s, whereas in the second case 150s is the shortest time the loss must last in order to detect Unavailability. By increasing Δt , we may gain accuracy but lose precision.

We can conclude that while it is desirable to have a high probability of detecting a state change immediately, so that the overall measure of Availability over T is accurate, a slightly lower probability may be acceptable if this gives higher precision, since the likelihood of a state change occurring at all remains high, provided the period of loss lasts for long enough. In other words, setting Δt to a lower value may be more desirable even though it means a lower sample size s and hence a lower probability of detecting Unavailable state immediately, p_n .

Note that the definition of Availability makes 'false positives' – detecting a state change to Unavailable when the actual loss is significantly less than the threshold – extremely unlikely. For example, with $n = 5$, the probability of changing state to Unavailable when the actual loss probability is $C/10$ (i.e., a tenth of the threshold) is generally less than 0.00002 (and in most cases significantly lower than that), even when the threshold is as high as 0.5.

25.6 Example Parameter Values

The tables below show some example combinations of parameter values, and the effect they have when there is a given level of loss of a given duration. Cells are highlighted green if $p_n > 0.5$ or if $p_u > 0.9$; if both conditions hold, the whole row is highlighted, indicating that this may be a good set of parameters for detecting a period of Unavailable Time if the actual loss is at the given level for the given period.

Note in all the cases, P is set to 100ms; it can be seen that a long value for P requires a long Δt to get the necessary accuracy of Availability flr and hence a sufficiently high value for p_n , but this means a low precision of Availability over T . Conversely, a short value of P , while ideal from a

statistical point of view, may incur too high a penalty in terms of network bandwidth and processing resources on network elements. 100ms is felt to be a reasonable compromise.

In the first set of tables, to cut down the number of variables the value of n is fixed at $n = 5$. The effect of changing n is analyzed in the second set of tables.

25.6.1 Examples with varying C and Δt

In the examples below, P and n are fixed with $P = 100\text{ms}$, $n = 5$.

The first table shows $\Delta t = 1\text{s}$; this gives 10 SLMs per Δt , and $n \Delta t = 5\text{s}$.

Threshold, C	Actual Loss Probability, L	Actual Loss duration (s)	flr CoV	Probability C exceeded in one Δt , p_k	Probability immediate state change, p_n	Probability state change at all, p_u
0.0	0.5	5	0.316	0.999	0.995	0.995
0.0	0.5	10	0.316	0.999	0.995	1.000
0.0	0.4	5	0.387	0.994	0.970	0.970
0.0	0.4	10	0.387	0.994	0.970	0.999
0.0	0.3	5	0.483	0.972	0.867	0.867
0.0	0.3	10	0.483	0.972	0.867	0.989
0.0	0.2	5	0.632	0.893	0.567	0.567
0.0	0.2	10	0.632	0.893	0.567	0.871
0.0	0.1	5	0.949	0.651	0.117	0.117
0.0	0.1	10	0.949	0.651	0.117	0.322
0.1	0.5	5	0.316	0.989	0.947	0.947
0.1	0.5	10	0.316	0.989	0.947	0.998
0.1	0.4	5	0.387	0.954	0.789	0.789
0.1	0.4	10	0.387	0.954	0.789	0.972
0.1	0.3	5	0.483	0.851	0.446	0.446
0.1	0.3	10	0.483	0.851	0.446	0.778
0.1	0.2	5	0.632	0.624	0.095	0.095
0.1	0.2	10	0.632	0.624	0.095	0.273

Table 31 – Examples with $P = 100\text{ms}$, $n = 5$, $\Delta t = 1\text{s}$

This shows that with a threshold of 0, an actual FLR of 30% for 10s or an actual FLR of 40% for 5s is likely to trigger a state change to Unavailable. With a threshold of 0.1 (the next step, bearing

in mind the recommendation that C is a multiple of $1/s$), an actual FLR of 40% for 10s or 50% for 5s is likely to trigger a state change.

The following table shows $\Delta t = 10s$; this gives 100 SLMs per Δt , and $n \Delta t = 50s$.

Threshold, C	Actual Loss Probability, L	Actual Loss duration (s)	flr CoV	Probability C exceeded in one Δt , p_k	Probability immediate state change, p_n	Probability state change at all, p_u
0.00	0.05	60	0.436	0.994	0.971	0.976
0.00	0.05	100	0.436	0.994	0.971	0.999
0.00	0.05	300	0.436	0.994	0.971	1.000
0.00	0.03	60	0.569	0.952	0.784	0.821
0.00	0.03	100	0.569	0.952	0.784	0.970
0.00	0.03	300	0.569	0.952	0.784	1.000
0.00	0.02	60	0.700	0.867	0.491	0.556
0.00	0.02	100	0.700	0.867	0.491	0.817
0.00	0.02	300	0.700	0.867	0.491	0.998
0.00	0.01	60	0.995	0.634	0.102	0.140
0.00	0.01	100	0.995	0.634	0.102	0.290
0.00	0.01	300	0.995	0.634	0.102	0.734
0.01	0.05	60	0.436	0.963	0.828	0.859
0.01	0.05	100	0.436	0.963	0.828	0.981
0.01	0.05	300	0.436	0.963	0.828	1.000
0.01	0.03	60	0.569	0.805	0.339	0.405
0.01	0.03	100	0.569	0.805	0.339	0.669
0.01	0.03	300	0.569	0.805	0.339	0.984
0.01	0.02	60	0.700	0.597	0.076	0.106
0.01	0.02	100	0.700	0.597	0.076	0.228
0.01	0.02	300	0.700	0.597	0.076	0.635

Table 32 – Examples with $P = 100ms$, $n = 5$, $\Delta t = 10s$

This shows that with a threshold of 0, an actual FLR of 3% for 100s or an actual FLR of 5% for 50s is likely to trigger a state change to Unavailable. With a threshold of 0.01 (the next step, bearing in mind the recommendation that C is a multiple of $1/s$), an actual FLR of 5% for 100s is likely to trigger a state change.

The following table shows $\Delta t = 30s$; this gives 300 SLMs per Δt , and $n \Delta t = 150s$.

Threshold, C	Actual Loss Probability, L	Actual Loss duration (s)	flr CoV	Probability C exceeded in one Δt , p_k	Probability immediate state change, p_n	Probability state change at all, p_u
0.0000	0.05000	150	0.252	1.000	1.000	1.000
0.0000	0.05000	300	0.252	1.000	1.000	1.000
0.0000	0.03000	150	0.328	1.000	0.999	0.999
0.0000	0.03000	300	0.328	1.000	0.999	1.000
0.0000	0.02000	150	0.404	0.998	0.988	0.988
0.0000	0.02000	300	0.404	0.998	0.988	1.000
0.0000	0.01000	150	0.574	0.951	0.778	0.778
0.0000	0.01000	300	0.574	0.951	0.778	0.968
0.0000	0.00333	150	0.998	0.633	0.101	0.101
0.0000	0.00333	300	0.998	0.633	0.101	0.288
0.0033	0.05000	150	0.252	1.000	1.000	1.000
0.0033	0.05000	300	0.252	1.000	1.000	1.000
0.0033	0.03000	150	0.328	0.999	0.994	0.994
0.0033	0.03000	300	0.328	0.999	0.994	1.000
0.0033	0.02000	150	0.404	0.983	0.920	0.920
0.0033	0.02000	300	0.404	0.983	0.920	0.996
0.0033	0.01000	150	0.574	0.802	0.333	0.333
0.0033	0.01000	300	0.574	0.802	0.333	0.661
0.0100	0.05000	150	0.252	1.000	0.999	0.999
0.0100	0.05000	300	0.252	1.000	0.999	1.000
0.0100	0.03000	150	0.328	0.980	0.904	0.904
0.0100	0.03000	300	0.328	0.980	0.904	0.994
0.0100	0.02000	150	0.404	0.851	0.448	0.448
0.0100	0.02000	300	0.404	0.851	0.448	0.780

Table 33 – Examples with $P = 100\text{ms}$, $n = 5$, $\Delta t = 30\text{s}$

This shows that with a threshold of 0, an actual FLR of 1% for 300s or an actual FLR of 2% for 150s is likely to trigger a state change to Unavailable. With a threshold of 0.0033 (the next step, bearing in mind the recommendation that C is a multiple of $1/s$), an actual FLR of 2% for 150s remains likely to trigger a state change, but with a threshold of 0.01 this probability decreases, and an actual FLR of 3% is needed to get a high probability of triggering a state change in 150s.

The final table shows $\Delta t = 60\text{s}$; this gives 600 SLMs per Δt , and $n \Delta t = 300\text{s}$.

Threshold, C	Actual Loss Probability, L	Actual Loss duration (s)	flr CoV	Probability C exceeded in one Δt , p_k	Probability immediate state change, p_n	Probability state change at all, p_u
0.0000	0.02000	300	0.286	1.000	1.000	1.000
0.0000	0.02000	600	0.286	1.000	1.000	1.000
0.0000	0.02000	900	0.286	1.000	1.000	1.000
0.0000	0.01000	300	0.406	0.998	0.988	0.988
0.0000	0.01000	600	0.406	0.998	0.988	1.000
0.0000	0.01000	900	0.406	0.998	0.988	1.000
0.0000	0.00500	300	0.576	0.951	0.776	0.776
0.0000	0.00500	600	0.576	0.951	0.776	0.968
0.0000	0.00500	900	0.576	0.951	0.776	0.996
0.0000	0.00167	300	0.999	0.632	0.101	0.101
0.0000	0.00167	600	0.999	0.632	0.101	0.287
0.0000	0.00167	900	0.999	0.632	0.101	0.440
0.0017	0.02000	300	0.286	1.000	1.000	1.000
0.0017	0.02000	600	0.286	1.000	1.000	1.000
0.0017	0.02000	900	0.286	1.000	1.000	1.000
0.0017	0.01000	300	0.406	0.983	0.918	0.918
0.0017	0.01000	600	0.406	0.983	0.918	0.996
0.0017	0.01000	900	0.406	0.983	0.918	1.000
0.0017	0.00500	300	0.576	0.802	0.331	0.331
0.0017	0.00500	600	0.576	0.802	0.331	0.659
0.0017	0.00500	900	0.576	0.802	0.331	0.836
0.0050	0.02000	300	0.286	0.998	0.989	0.989
0.0050	0.02000	600	0.286	0.998	0.989	1.000
0.0050	0.02000	900	0.286	0.998	0.989	1.000
0.0050	0.01000	300	0.406	0.850	0.444	0.444
0.0050	0.01000	600	0.406	0.850	0.444	0.777
0.0050	0.01000	900	0.406	0.850	0.444	0.918

Table 34 – Examples with $P = 100\text{ms}$, $n = 5$, $\Delta t = 60\text{s}$

This shows that with a threshold of 0, an actual FLR of 0.5% for 600s or an actual FLR of 1% for 300s is likely to trigger a state change to Unavailable. With a threshold of 0.0017 (the next step, bearing in mind the recommendation that C is a multiple of $1/s$), an actual FLR of 1% for 300s

remains likely to trigger a state change, but with a threshold of 0.005 this probability decreases, and an actual FLR of 2% is needed to have a high probability of triggering a state change in 300s.

From the above we can draw a number of conclusions. Firstly, this illustrates again that to measure lower loss requires more samples, i.e. a longer Δt . However, even CoV values for Availability flr significantly above the 0.2 suggested for FLR can give acceptable results. Conversely, to have a high likelihood of triggering a state change for a given actual loss, it is better to pick as short a Δt as possible as this means shorter periods of loss can be detected, and improves the precision of the overall Availability calculation over T .

Another, perhaps surprising, corollary to this is that a shorter Δt can be used if the threshold, C , is set to 0 – in other words, triggering state changes if a single frame is lost in each Δt . Intuitively, it seems as though basing a state change off a single lost frame would give too much variability; however, the effect is mitigated by requiring several consecutive Δt intervals to see the lost frame – and this effect increases with larger n .

It may be noted that with $C = 0$, even a very low level of loss has a reasonable chance of triggering a state change if it lasts long enough. In fact this is true regardless of C ; however, for a given duration the probability is much higher with $C = 0$; or to put it another way, for a given duration a lower level of loss will have the same probability of triggering a state change with $C = 0$, as compared with $C > 0$. This is illustrated in the table below, which has $\Delta t = 10$ s, $P = 100$ ms and $n = 5$, giving an Availability flr precision of 1% and $n \Delta t$ of 50s. It shows a variety of actual loss probabilities, each lasting for 300s, with thresholds of 0 or 0.01.

Threshold, C	Actual Loss Probability, L	Actual Loss duration (s)	flr CoV	Probability C exceeded in one Δt , p_k	Probability immediate state change, p_n	Probability state change at all, p_u
0.00	0.030	300	0.569	0.952	0.784	1.000
0.00	0.025	300	0.624	0.920	0.661	1.000
0.00	0.020	300	0.700	0.867	0.491	0.998
0.00	0.015	300	0.810	0.779	0.288	0.970
0.00	0.010	300	0.995	0.634	0.102	0.734
0.00	0.009	300	1.049	0.595	0.075	0.631
0.00	0.005	300	1.411	0.394	0.010	0.146
0.01	0.030	300	0.569	0.805	0.339	0.984
0.01	0.025	300	0.624	0.717	0.189	0.902
0.01	0.020	300	0.700	0.597	0.076	0.635
0.01	0.015	300	0.810	0.443	0.017	0.236
0.01	0.010	300	0.995	0.264	0.001	0.025

Threshold, C	Actual Loss Probability, L	Actual Loss duration (s)	flr CoV	Probability C exceeded in one Δt , p_k	Probability immediate state change, p_n	Probability state change at all, p_u
0.01	0.009	300	1.049	0.227	0.001	0.012
0.01	0.005	300	1.411	0.090	0.000	0.000

Table 35 – Examples with varying loss for 300s

It can be seen that with $C = 0$, there is a high probability of a state change within 300s if the actual FLR is 1.5% or more, and a reasonable probability even down to 0.5%, even though the probability of an immediate change is very low. The same pattern can be seen with $C = 0.01$, except that a higher actual loss is needed. As another illustration, note that in all the combinations of parameters in all of the tables above, there is a significant probability (≥ 0.1) of detecting a state change at some point.

25.6.2 Examples with varying n

The previous section illustrated the effect of changing C and Δt , with a fixed value of $n = 5$. Here, we examine the impact of varying n . For this purpose, we fix $C = 0.01$ and $\Delta t = 10s$, as well as $P = 100ms$ as before.

As previously, cells or rows are highlighted if $p_n > 0.5$ and/or $p_u > 0.9$. In addition, the minimum possible FLR over 1 month is shown, given an actual loss of the given duration (and assuming a uniform rate of traffic). Rows where this exceeds 0.001% (i.e., 10^{-5}) are shown with grey text – in other words, these are rows where an FLR target over T of 10^{-5} is guaranteed to be exceeded.

The first table shows $n = 10$, giving an $n \Delta t$ of 100s.

Actual Loss Probability, L	Actual Loss duration (s)	flr CoV	Probability C exceeded in one Δt , p_k	Probability immediate state change, p_n	Probability state change at all, p_u	FLR over 1 month
0.10	100	0.300	1.000	0.997	0.997	0.0004%
0.10	300	0.300	1.000	0.997	1.000	0.0011%
0.10	600	0.300	1.000	0.997	1.000	0.0023%
0.05	100	0.436	0.963	0.685	0.685	0.0002%
0.05	300	0.436	0.963	0.685	0.990	0.0006%

Actual Loss Probability, L	Actual Loss duration (s)	flr CoV	Probability C exceeded in one Δt , p_k	Probability immediate state change, p_n	Probability state change at all, p_u	FLR over 1 month
0.05	600	0.436	0.963	0.685	1.000	0.0011%
0.04	100	0.490	0.913	0.402	0.402	0.0002%
0.04	300	0.490	0.913	0.402	0.906	0.0005%
0.04	600	0.490	0.913	0.402	0.995	0.0009%
0.03	100	0.569	0.805	0.115	0.115	0.0001%
0.03	300	0.569	0.805	0.115	0.514	0.0003%
0.03	600	0.569	0.805	0.115	0.808	0.0007%
0.02	100	0.700	0.597	0.006	0.006	0.0001%
0.02	300	0.700	0.597	0.006	0.052	0.0002%
0.02	600	0.700	0.597	0.006	0.117	0.0005%
0.01	100	0.995	0.264	0.000	0.000	0.0000%
0.01	300	0.995	0.264	0.000	0.000	0.0001%
0.01	600	0.995	0.264	0.000	0.000	0.0002%

Table 36 – Examples with $P = 100\text{ms}$, $C = 0.01$, $\Delta t = 10\text{s}$, $n = 10$

This shows that with $n = 10$, only actual loss which is significantly above the threshold results in a high probability of detecting a state change to Unavailable – but in most of those cases, the FLR over time T will also be high. In other words, it is likely that if any Availability target in the SLS is broken, the FLR target will also be broken – with $n = 10$, the usefulness of the Availability measurement is marginal. It can also be seen that with lower levels of loss, the probability of a state change decreases rapidly, even when it lasts for a long time. There is a low probability of triggering a state change at an actual FLR of 2%, and an actual FLR of 1% has a near-zero chance of causing a state change even after 600s.

The next table shows $n = 5$, giving an $n \Delta t$ of 50s.

Actual Loss Probability, L	Actual Loss duration (s)	flr CoV	Probability C exceeded in one Δt , p_k	Probability immediate state change, p_n	Probability state change at all, p_u	FLR over 1 month
0.10	50	0.300	1.000	0.998	0.998	0.0002%

Actual Loss Probability, L	Actual Loss duration (s)	flr CoV	Probability C exceeded in one Δt , p_k	Probability immediate state change, p_n	Probability state change at all, p_u	FLR over 1 month
0.10	100	0.300	1.000	0.998	1.000	0.0004%
0.10	150	0.300	1.000	0.998	1.000	0.0006%
0.10	300	0.300	1.000	0.998	1.000	0.0011%
0.05	50	0.436	0.963	0.828	0.828	0.0001%
0.05	100	0.436	0.963	0.828	0.981	0.0002%
0.05	150	0.436	0.963	0.828	0.998	0.0003%
0.05	300	0.436	0.963	0.828	1.000	0.0006%
0.04	50	0.490	0.913	0.634	0.634	0.0001%
0.04	100	0.490	0.913	0.634	0.910	0.0002%
0.04	150	0.490	0.913	0.634	0.981	0.0002%
0.04	300	0.490	0.913	0.634	1.000	0.0005%
0.03	50	0.569	0.805	0.339	0.339	0.0001%
0.03	100	0.569	0.805	0.339	0.669	0.0001%
0.03	150	0.569	0.805	0.339	0.843	0.0002%
0.03	300	0.569	0.805	0.339	0.984	0.0003%
0.02	50	0.700	0.597	0.076	0.076	0.0000%
0.02	100	0.700	0.597	0.076	0.228	0.0001%
0.02	150	0.700	0.597	0.076	0.360	0.0001%
0.02	300	0.700	0.597	0.076	0.635	0.0002%
0.01	50	0.995	0.264	0.001	0.001	0.0000%
0.01	100	0.995	0.264	0.001	0.006	0.0000%
0.01	150	0.995	0.264	0.001	0.011	0.0001%
0.01	300	0.995	0.264	0.001	0.025	0.0001%

Table 37 – Examples with $P = 100\text{ms}$, $C = 0.01$, $\Delta t = 10\text{s}$, $n = 5$

With $n = 5$, a state change is likely to be triggered with high probability at a lower actual loss and in a shorter time, meaning the overall loss is well below that likely to exceed the FLR target. The likelihood of a state change at a low actual loss levels is higher; there is a reasonable chance of an actual FLR of 2% triggering a state change after 300s, but the likelihood of an FLR of 1% triggering a state change remains low. This setting for n appears to offer a reasonable compromise.

The next table shows $n = 3$, giving an $n \Delta t$ of 50s.

Actual Loss Probability, L	Actual Loss duration (s)	flr CoV	Probability C exceeded in one Δt , p_k	Probability immediate state change, p_n	Probability state change at all, p_u	FLR over 1 month
0.050	30	0.436	0.963	0.893	0.893	0.0001%
0.050	60	0.436	0.963	0.893	0.992	0.0001%
0.050	150	0.436	0.963	0.893	1.000	0.0003%
0.050	300	0.436	0.963	0.893	1.000	0.0006%
0.040	30	0.490	0.913	0.761	0.761	0.0000%
0.040	60	0.490	0.913	0.761	0.960	0.0001%
0.040	150	0.490	0.913	0.761	1.000	0.0002%
0.040	300	0.490	0.913	0.761	1.000	0.0005%
0.030	30	0.569	0.805	0.522	0.522	0.0000%
0.030	60	0.569	0.805	0.522	0.827	0.0001%
0.030	150	0.569	0.805	0.522	0.994	0.0002%
0.030	300	0.569	0.805	0.522	1.000	0.0003%
0.020	30	0.700	0.597	0.212	0.212	0.0000%
0.020	60	0.700	0.597	0.212	0.470	0.0000%
0.020	150	0.700	0.597	0.212	0.848	0.0001%
0.020	300	0.700	0.597	0.212	0.981	0.0002%
0.010	30	0.995	0.264	0.018	0.018	0.0000%
0.010	60	0.995	0.264	0.018	0.059	0.0000%
0.010	150	0.995	0.264	0.018	0.173	0.0001%
0.010	300	0.995	0.264	0.018	0.332	0.0001%
0.009	30	1.049	0.227	0.012	0.012	0.0000%
0.009	60	1.049	0.227	0.012	0.039	0.0000%
0.009	150	1.049	0.227	0.012	0.117	0.0001%
0.009	300	1.049	0.227	0.012	0.233	0.0001%

Table 38 – Examples with $P = 100\text{ms}$, $C = 0.01$, $\Delta t = 10\text{s}$, $n = 3$

With $n = 3$, a state change is likely to be triggered with high probability at an even lower actual loss. However, there is also a significant chance of a state change even when the actual loss is below the threshold C , which may be undesirable. This setting for n may be too low.

25.7 Relationship between Availability and CHLI

MEF 10.3 mandates that the number of intervals n used for Availability must be larger than the number of small time intervals p used for CHLI, and CHLI are only counted during Available Time. This means that n must be ≥ 2 if CHLI should be measured. Note that in the case of $n = 2$, p can only be set to 1, i.e. only single Δt intervals with high loss contribute to the count of CHLIs, since 2 consecutive Δt intervals with high loss would cause a transition to Unavailable state. In this case, the count of CHLIs is equal to the count of HLIs, so this is unlikely to be a useful configuration.

25.8 Conclusions

It can be seen from the discussion above – and the length of this Appendix – that there are many factors involved in choosing an appropriate combination of parameters for Availability measurement, and that the relationships between them are complex. For this reason, it is hard to give any single recommended set of parameters, or even simple “rules of thumb”. Much depends on the tradeoff between the level of loss that the operator would like to see trigger a state change to Unavailable, and the length of time that level of loss must persist for to make a state change sufficiently likely. As has been seen, a higher level of loss can be detected more quickly, and therefore even if it lasts for a shorter duration, whereas to trigger a state change given a low level of loss takes much longer.

There are some general conclusions that can be drawn:

- A value of $n = 5$ appears to strike a good balance between mitigating the high CoV for Availability flr measurements, and detecting state changes sufficiently quickly.
- The value of C should be a multiple of $P/\Delta t$ (i.e., a multiple of $1/s$)
- Setting C to 0 may be a reasonable configuration, particularly if there is a desire to measure low levels of loss.
- The defaults recommended in earlier revisions of this document ($n = 10$, $C = 0.5$, $\Delta t = 1s$) are unlikely to be useful for most operators.