



Technical Specification

MEF 10.2.1

Performance Attributes Amendment to MEF 10.2

25 January 2011

Disclaimer

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

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

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

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

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

© The Metro Ethernet Forum 2010. All Rights Reserved.

Table of Contents

<i>Table of Contents</i>	<i>i</i>
2. Terminology	1
6.9 EVC Related Performance Service Attributes.....	2
6.9.5 One-way Frame Loss Ratio Performance for a Point-to-Point EVC	2
6.9.6 One-way Frame Loss Ratio Performance for an EVC	3
6.9.7 One-way Frame Loss Ratio Performance for an EVC	3
6.9.8 One-way Availability Performance for an EVC	3
6.9.9 One-way Resiliency Performance for an EVC.....	8

List of Figures

Figure A – Flowchart Definition of $A_{\langle i,j \rangle}(\Delta t_k)$	5
Figure B – Example of the Determination of $A_{\langle i,j \rangle}(\Delta t_k)$	6
Figure C – Example of the Impact of a Maintenance Interval.....	7
Figure D – Hierarchy of Time Showing the Resiliency Attributes	9
Figure E – Determining and Counting Consecutive High Loss Intervals over T	10
Figure F – Example of Counting High Loss Intervals and Consecutive High Loss Intervals.....	11

List of Tables

Table A – Availability Performance Parameters for an EVC.....	8
Table B – Resiliency Performance Parameters for an EVC	11

INTRODUCTION

This amendment makes the following changes to MEF 10.2 [1]:

1. The definition of Qualified Service Frames in Section 6.9 is modified to clarify that Frame Delay, Inter-Frame Delay Variation and Frame Loss Ratio performance commitments only apply during available time.
2. Section 6.9.5 has been deleted since the definition for frame loss ratio needed for availability is now contained in the availability section.
3. Section 6.9.6 has been modified to fix an inconsistency in MEF 10.2, specifically to also relate the *FLR* metric to Qualified Service Frames.
4. Section 6.9.7 has been deleted.
5. Section 6.9.8 has been replaced to modify the definition of Availability that now uses frame loss during a sliding window to determine the availability or unavailability for each short time interval, Δt . In addition, a single loss threshold replaces two loss thresholds in the definition. Text has been changed related to Scheduled Downtime, including the replacement of the term with Maintenance Interval. Text has been removed related to Unscheduled Downtime.
6. Section 6.9.9 has been added to define two new performance attributes for resiliency performance.
7. The new terms used in this amendment are defined in table form for inclusion into section 2 of a future roll-up.

2. Terminology

The following terms used in this amendment are defined here.

CHLI	Consecutive High Loss Interval
Consecutive High Loss Interval	A sequence of small time intervals contained in T , each with a high frame loss ratio.
HLI	High Loss Interval
High Loss Interval	A small time interval contained in T with a high frame loss ratio.
Resiliency Performance	The number of High Loss Intervals and Consecutive High Loss Intervals in T
T	The time interval over which the SLS applies

6.9 EVC Related Performance Service Attributes

This amendment modifies the first paragraph of section 6.9, as shown below in underlined text:

The EVC Related Performance Service Attributes specify the Service Frame delivery performance. Five ~~Four~~ performance attributes are considered in this specification. These are Frame Delay Performance, Inter-Frame Delay Variation Performance, Frame Loss Ratio Performance, Availability Performance and Resiliency Performance. For any given SLS, performance objectives related to these attributes may or may not be specified. If an SLS contains an objective for a given performance attribute, then the SLS MUST specify the related parameters for that objective.

This amendment modifies the second bullet in the Qualified Service Frames definition, by adding a phrase at the end, as shown below in underlined text:

- The first bit of 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 (see Section 6.9.8)

This amendment replaces the third and fourth bullets in the Qualified Service Frames definition, with the following three bullets:

- Each Service Frame **MUST** have the Class of Service Identifier for the Class of Service instance in question,
- Each ingress Service Frame that is subject to an Ingress Bandwidth Profile **MUST** have an Ingress Bandwidth Profile compliance of Green, and
- Each ingress Service Frame that is not subject to an Ingress Bandwidth Profile **MUST** either have no color identifier or a color identifier indicating Green per the color indication requirements of [2].

This amendment deletes the paragraph “Values of the Service Frame delay, delay variation, and loss performance during periods of unavailable time **MUST NOT** be used to determine Service Frame delivery compliance. A process **MUST** be established to exclude all performance during unavailable periods from comparison with Service Frame performance objectives.” This paragraph is not needed given the changes to the Qualified Service Frames definition.

6.9.5 One-way Frame Loss Ratio Performance for a Point-to-Point EVC

The contents of Section 6.9.5 is deleted leaving an empty section. The old Section 6.9.5 consisted of a definition for *FLR* which was used in the Availability definition. The new text for Section 6.9.8 now contains the definition for *flr*.

6.9.6 One-way Frame Loss Ratio Performance for an EVC

The second bullet is replaced to clarify the dependency on Qualified Service Frames. The new second bullet follows:

- Let $I_T^{(i,j)}$ denote the number of ingress Qualified Service Frames at ingress UNI i that should have been delivered to UNI j according to the Service Frame Delivery service attributes (see Section 6.5). Each Service Frame can be a Unicast (see section 6.5.1.1), Multicast (see section 6.5.1.2), Broadcast (see section 6.5.1.3) or Layer 2 Control Protocol (see section 6.5.1.4) Service Frame.

6.9.7 Availability Performance for a Point-to-Point EVC

The contents of section 6.9.7 are deleted leaving an empty section. The old Section 6.9.7 defined Availability for a Point-to-Point EVC and the old Section 6.9.8 defined Availability for a Multipoint EVC. The new Section 6.9.8 covers all types of EVCs in the same way that MEF 10.2 does for delay, delay variation, and frame loss ratio.

6.9.8 One-way Availability Performance for an EVC

The contents of Section 6.9.8 of MEF 10.2 [1] are replaced with the following text.

Availability Performance is the percentage of time within a specified time interval during which the frame loss is small. (The precise definition is presented in the following paragraphs.) As an example, a service provider can define the Availability performance to be measured over a month and the value for the Availability Performance objective to be 99.9%. In a month with 30 days and no Maintenance Interval this objective will allow the service to be unavailable for approximately 43 minutes out of the whole month.

Informally, Availability Performance is based on Service Frame loss during a sequence of consecutive small time intervals and the availability state during the previous small time interval. When the previous sequence was defined as available, if the frame loss is high for each small time interval in the current sequence, then the small time interval at the beginning of the current sequence is defined as unavailable; otherwise it is defined as available. On the other hand, when the previous sequence was defined as unavailable, if frame loss is low for each small time interval in the current sequence, then the small time interval at the beginning of the current sequence is defined as available; otherwise, it is defined as unavailable. The formal definition follows.

The Availability for a particular Class of Service instance from ingress UNI i to egress UNI j for a time interval T is based on the following three parameters:

- Δt , a time interval much smaller than T ,
- C , a frame loss ratio threshold which if exceeded suggests unavailability,

- n , the number of consecutive small time intervals, Δt , over which to assess availability.

Each Δt_k in T is defined to be either “Available” or “Unavailable” and this is represented by $A_{\langle i,j \rangle}(\Delta t_k)$ where $A_{\langle i,j \rangle}(\Delta t_k)=1$ means that the service during Δt_k is Available and $A_{\langle i,j \rangle}(\Delta t_k)=0$ means that the service during Δt_k is Unavailable. The definition of $A_{\langle i,j \rangle}(\Delta t_k)$ is based on the following frame loss ratio function, $flr_{\langle i,j \rangle}(\Delta t_k)$, which is defined as follows.

Let $I_{\Delta t}^{(i,j)}$ be the number of ingress Service Frames that meet the following conditions:

- The first bit of each Service Frame **MUST** arrive at UNI i within the time interval Δt ,
- Each Service Frame **MUST** be one that is to be delivered to the UNI j according to the Service Frame Delivery service attributes (see Section 6.5.1). Each Service Frame can be a Unicast, Multicast, or Broadcast Service Frame.
- Each Service Frame **MUST** belong to the Class of Service instance for the SLS,
- Each Service Frame that is subject to an Ingress Bandwidth Profile **MUST** have an Ingress Bandwidth Profile compliance of Green, and
- Each Service Frame that is not subject to an Ingress Bandwidth Profile **MUST** either have no color identifier or a color identifier that corresponds to Green, as per the color indication requirements of MEF 23 [2].

Let $E_{\Delta t}^{(i,j)}$ be the number of unique (not duplicate) unerrored egress Service Frames where each Service Frame is the first egress Service Frame at UNI j that results from a Service Frame counted in $I_{\Delta t}^{(i,j)}$.

$$\text{Then, } flr_{\langle i,j \rangle}(\Delta t) = \begin{cases} \left(\frac{I_{\Delta t}^{(i,j)} - E_{\Delta t}^{(i,j)}}{I_{\Delta t}^{(i,j)}} \right) & \text{if } I_{\Delta t}^{(i,j)} \geq 1 \\ 0 & \text{otherwise} \end{cases}$$

In the case of a Multipoint-to-Multipoint or a Rooted-Multipoint EVC, the Service Provider and the Subscriber **MAY** agree to define $flr_{\langle i,j \rangle}(\Delta t)$ as

$$flr_{\langle i,j \rangle}(\Delta t) = \begin{cases} \left(\frac{\tilde{I}_{\Delta t}^{(i,j)} - E_{\Delta t}^{(i,j)}}{\tilde{I}_{\Delta t}^{(i,j)}} \right) & \text{if } \tilde{I}_{\Delta t}^{(i,j)} \geq 1 \\ 0 & \text{otherwise} \end{cases}$$

where $\tilde{I}_{\Delta t}^{(i,j)} = I_{\Delta t}^{(i,j)} -$ the number of Service Frames discarded by the Service Provider, in order to conform to either the line rate of UNI j or an Egress Bandwidth Profile (if one is used) at UNI

j . Such Service Frame drops may occur anywhere in the network, not just close to UNI j . One example of this could be where an Egress Bandwidth Profile is applied on a link within the network. Another example of this could be where excessive Service Frames exceed the line rate on a link within the network. Good traffic engineering principles would suggest dropping such excessive Service Frames as close to the ingress as possible. This adjustment is meant to account for a focused overload of traffic sent to UNI j from multiple ingress UNIs. The details of such an adjustment are beyond the scope of this document.

Δt_0 is the first short time interval agreed by the Service Provider and Subscriber at or after turn-up of the EVC. $A_{\langle i,j \rangle}(\Delta t_k)$ is defined by the flow diagram in Figure A for $k = 0, 1, 2, \dots$

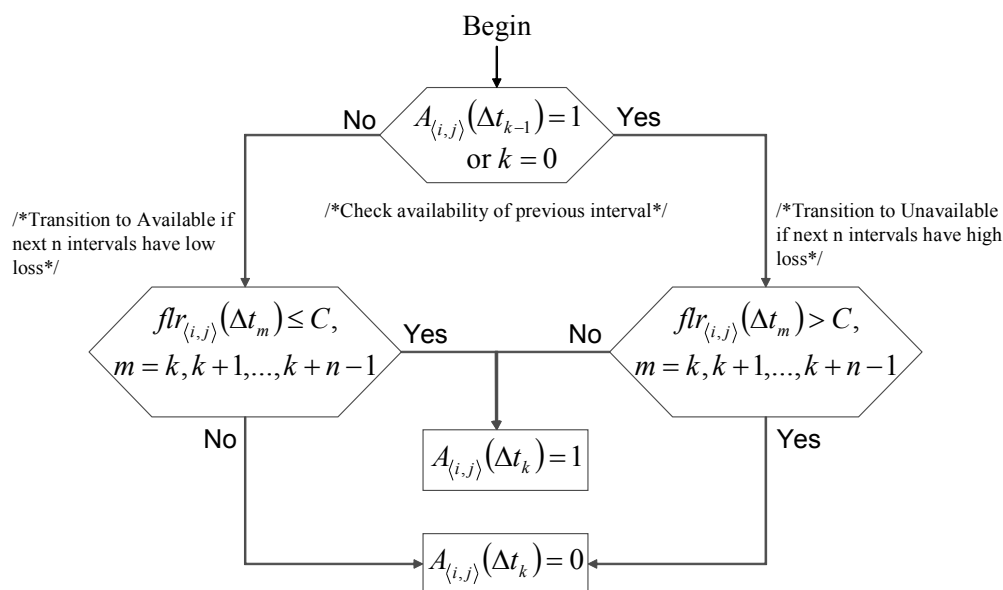


Figure A – Flowchart Definition of $A_{\langle i,j \rangle}(\Delta t_k)$

An alternative way of expressing $A_{\langle i,j \rangle}(\Delta t_k)$ for $k = 0$ is

$$A_{\langle i,j \rangle}(\Delta t_0) = \begin{cases} 0 & \text{if } flr_{\langle i,j \rangle}(\Delta t_m) > C, \forall m = 0, 1, \dots, n-1 \\ 1 & \text{otherwise} \end{cases}$$

and for $k = 1, 2, \dots$ is

$$A_{\langle i,j \rangle}(\Delta t_k) = \begin{cases} 0 & \text{if } A_{\langle i,j \rangle}(\Delta t_{k-1}) = 1 \text{ and } flr_{\langle i,j \rangle}(\Delta t_m) > C, \forall m = k, k+1, \dots, k+n-1 \\ 1 & \text{if } A_{\langle i,j \rangle}(\Delta t_{k-1}) = 0 \text{ and } flr_{\langle i,j \rangle}(\Delta t_m) \leq C, \forall m = k, k+1, \dots, k+n-1. \\ A_{\langle i,j \rangle}(\Delta t_{k-1}) & \text{otherwise} \end{cases}$$

In the event of a conflict between the above equations and Figure A, the content of Figure A is controlling.

The availability for Δt_k is based on the frame loss ratio during the short interval and each of the following $n-1$ short intervals and the availability of the previous short time interval. In other words, a sliding window of width $n\Delta t$ is used to determine availability. This use of a sliding window is similar to that of ITU-T Y.1563 [3].

Figure B presents an example of the determination of the availability for the small time intervals with a sliding window of 10 small time intervals.




Figure B – Example of the Determination of $A_{(i,j)}(\Delta t_k)$

The Availability for a particular Class of Service instance from UNI i to UNI j for a time interval T excludes the small time intervals that occur during a Maintenance Interval (MI). An MI is a time interval agreed to by the Service Provider and Subscriber during which the service may not perform well or at all. Examples of a Maintenance Interval include:

- A time interval during which the Service Provider may disable the service for network maintenance such as equipment replacement,
- A time interval during which the Service Provider and Subscriber may perform joint fault isolation testing, and
- A time interval during which the Service Provider may change service features and making the changes may disrupt the service.

Figure C shows an example of the elimination of short time intervals for a Maintenance Interval.

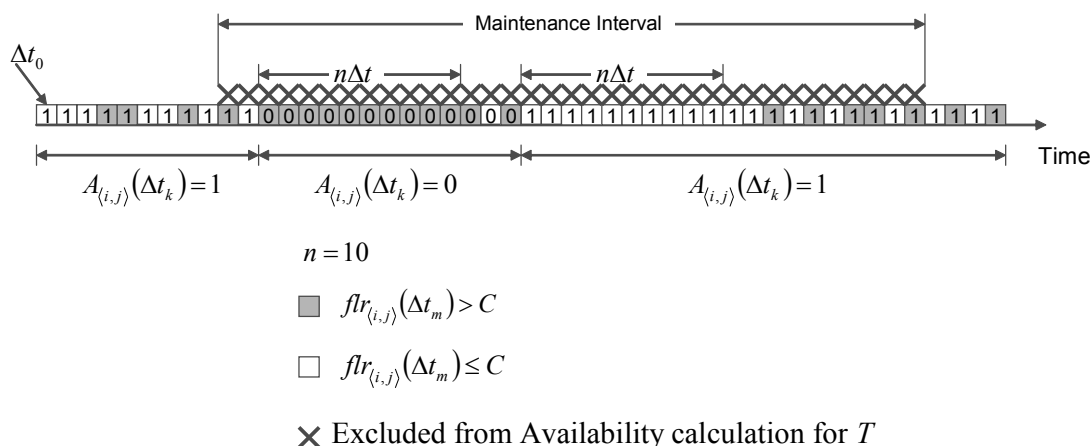


Figure C – Example of the Impact of a Maintenance Interval

Note - the $n\Delta t$ window for Δt_0 straddles a boundary between normal service time and a maintenance interval. This is a point of inconsistency with the Service Availability definition in Y.1563, where Service time alone determines Availability.

W_T is informally defined as the set of all Δt 's in T that are not included in a Maintenance Interval. Let $W_T = \{k \mid \Delta t_k \subseteq T \text{ and } \Delta t_k \text{ does not intersect a Maintenance Interval}\}$ and let $|W_T|$ represent the number of elements in the set W_T . Then the Availability (%) for a particular Class of Service instance from UNI i to UNI j for a time interval T , such as 1 month, is defined by

$$A_T^{(i,j)} = \begin{cases} \frac{100}{|W_T|} \sum_{k \in W_T} A_{(i,j)}(\Delta t_k) & \text{if } |W_T| > 0 \\ 100 & \text{otherwise} \end{cases}$$

Note that the definition of W_T means that the boundaries of T and the boundaries of a Maintenance Interval do not have to align with the boundary of a Δt_k . A Δt_k that *straddles the boundary between two T 's* is excluded from the definition of Availability Performance for each interval T . And a Δt_k that straddles the boundary of a Maintenance Interval is also excluded from the definition of Availability Performance.

Let the UNIs associated by the EVC be numbered $1, 2, \dots, m$ and let S be a non-empty subset of the ordered pairs of UNIs, i.e., $S \subseteq \{(i, j) \mid i = 1, 2, \dots, m, j = 1, 2, \dots, m, i \neq j\}$, $S \neq \emptyset$. Then the Availability for a particular Class of Service instance for the set S is defined by

$$A_T^S = \min \{A_T^{(i,j)} \mid (i, j) \in S\}.$$

For the SLS, an Availability Performance metric for a particular Class of Service **MUST** specify a set of parameters and an objective as shown in Table A.

Parameter	Description
T	The time interval
S	Non-empty subset of the ordered UNI pairs
Δt	A time interval much smaller than T
C	Unavailability frame loss ratio threshold
n	Number of consecutive small time intervals for assessing availability
\hat{A}	Availability Performance Objective expressed as a percentage

Table A – Availability Performance Parameters for an EVC

Given T , S , Δt , C , n , and \hat{A} , the SLS **SHALL** define the Availability Performance Objective as being met if and only if $A_T^S \geq \hat{A}$.

For a Point-to-Point EVC, S **MAY** include one or both of the ordered pairs of UNIs in the EVC.

For a Multipoint-to-Multipoint EVC, S **MAY** be any non-empty subset of the ordered pairs of UNIs in the EVC.

For a Rooted-Multipoint EVC, S **MUST** be a non-empty subset of the ordered pairs of UNIs in the EVC, such that all ordered pairs in S contain at least one UNI that is designated as a Root.

6.9.9 One-way Resiliency Performance for an EVC

This section defines attributes for the Resiliency performance of an ordered pair of UNIs, $\langle i, j \rangle$. The definitions depend on the availability status of each Δt to determine whether resiliency performance counts toward objectives. The Resiliency attributes are similar to the definitions of Severely Errored Seconds (SES) and Consecutive SES in section 9 and Annex B (respectively) of Y.1563 [3], when $\Delta t = 1$ second.

Figure D illustrates how the two resiliency attributes defined here, counts of High Loss Intervals and counts of Consecutive High Loss Intervals, fit into the hierarchy of time and other attributes.

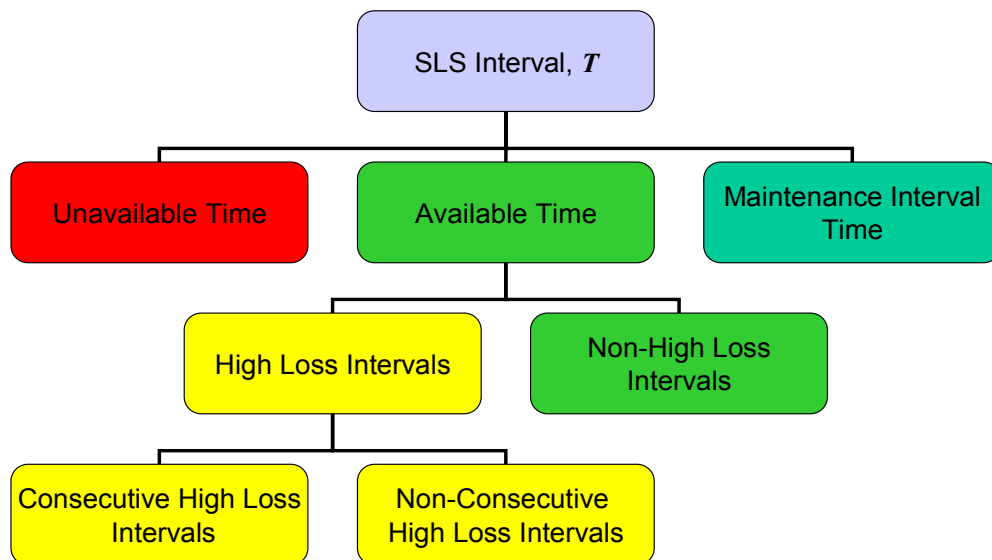


Figure D – Hierarchy of Time Showing the Resiliency Attributes

A High Loss Interval (HLI) is a small time interval contained in T (having the same duration as the interval, Δt , defined in Section 6.9.8) with a high flr . When sufficient HLIs are adjacent, the interval is designated as a Consecutive High Loss Interval (CHLI). Section 6.9.8 defines terminology for Availability. This section re-uses that terminology and defines the following terms:

- $H_{\langle i,j \rangle}(\Delta t_k)$: the high loss state of Δt_k ,
 - equal to 1 when $flr_{\langle i,j \rangle}(\Delta t_k) > C$ and $A_{\langle i,j \rangle}(\Delta t_k) = 1$, equal to 0 otherwise, including any Δt_k that intersects a Maintenance Interval
- $L_T^{\langle i,j \rangle}$: Count of High Loss Intervals (HLIs) over T
- \hat{L} : HLI Count Objective for S , T , and a given Class of Service instance
- p : the minimum integer number of consecutive HLIs in the (sliding) window (with $0 < p < n$) to qualify as a CHLI
- $B_T^{\langle i,j \rangle}$: Count of p or more consecutive HLIs occurring in T
- \hat{B} : CHLI Count Objective for S , T , and a given Class of Service Instance

For every Δt in T that does not intersect a Maintenance Interval, the flr and Availability state determine the value of $H_{\langle i,j \rangle}(\Delta t_k)$, either 1 or 0 as defined above.

For the SLS, the count of HLIs over T **MUST** be determined by

$$L_T^{(i,j)} = \sum_{\Delta t \in T} H_{(i,j)}(\Delta t).$$

Note that the counter for H may be implemented in post processing (e.g., in a Management System), outside the Network Element that is monitoring the frame loss rate of each Δt . This could be necessary to correlate with Δt 's in a Maintenance Interval (MI).

When counting CHLI, the threshold p is used similarly to the variable n for the window size in the Availability attribute, and $p < n$.

For the SLS, the Consecutive High Loss Intervals over T **MUST** be determined according to the flow chart in Figure E.

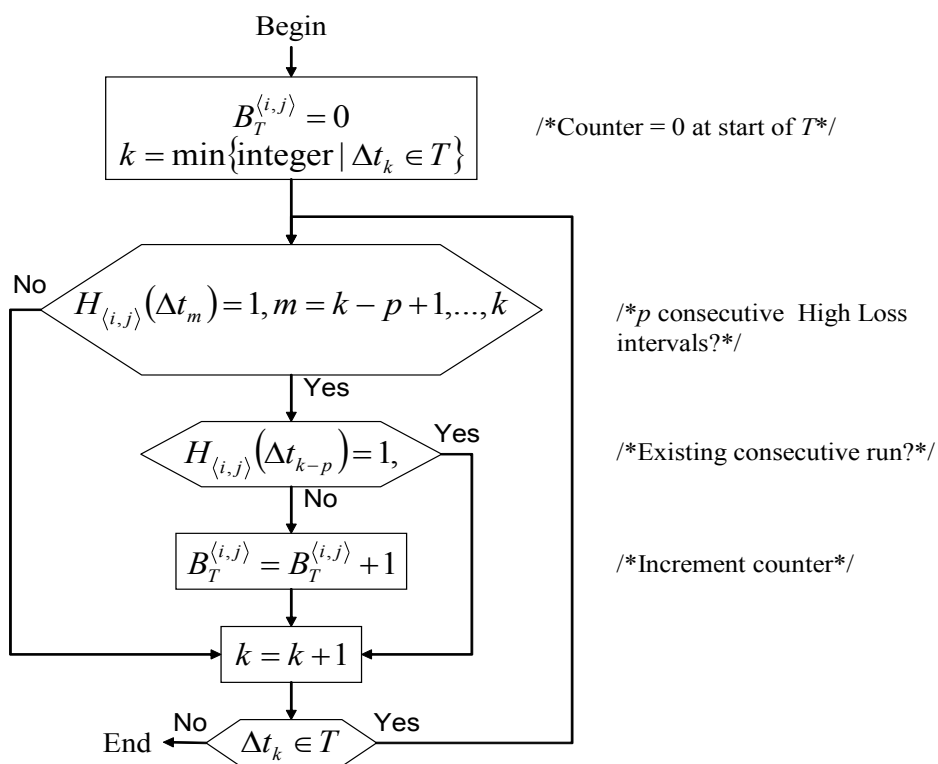


Figure E – Determining and Counting Consecutive High Loss Intervals over T

Figure F shows an example that depicts the HLI and CHLI counting processes.




Figure F – Example of Counting High Loss Intervals and Consecutive High Loss Intervals

Let the UNIs associated by the EVC be numbered $1, 2, \dots, m$ and let S be a non-empty subset of the ordered pairs of UNIs, i.e., $S \subseteq \{\langle i, j \rangle \mid i = 1, 2, \dots, m, j = 1, 2, \dots, m, i \neq j\}, S \neq \emptyset$. Then the HLI and CHLI performance attributes for a particular Class of Service instance for the set S are defined by

$$L_T^S = \max \{L_T^{(i,j)} \mid \langle i, j \rangle \in S\} \text{ and } B_T^S = \max \{B_T^{(i,j)} \mid \langle i, j \rangle \in S\}$$

For the SLS, the Resiliency Performance metrics for a particular Class of Service **MUST** specify a set of parameters and objectives as shown in Table B.

Parameter	Description
T	The time interval
S	Non-empty subset of the ordered UNI pairs associated by the EVC
Δt	A time interval much smaller than T , MUST be the same as in Section 6.9.8
C	Unavailability frame loss ratio threshold, MUST be the same as in Section 6.9.8
p	Number of consecutive small time intervals for assessing CHLI, where $p < n$
\hat{L}	HLI Performance Objective expressed as an integer
\hat{B}	Consecutive HLI Performance Objective expressed as an integer

Table B – Resiliency Performance Parameters for an EVC

Given $T, S, \Delta t, C, p, \hat{L}$, and \hat{B} , the SLS **MUST** define the HLI Performance Objective as being met if and only if $L_T^S \leq \hat{L}$, and the CHLI Performance Objective as being met if and only if $B_T^S \leq \hat{B}$.

For a Point-to-Point EVC, S **MAY** include one or both of the ordered UNI pairs in the EVC.

For a Multipoint-to-Multipoint EVC, S **MAY** be any non-empty subset of the ordered UNI pairs of the EVC.

For a Rooted-Multipoint EVC, S **MAY** be any non-empty subset of the ordered pairs of UNIs in the EVC, and S **MUST** be such that all ordered pairs in S contain at least one UNI that is designated as a Root.

9. REFERENCES

- [1] Metro Ethernet Forum MEF 10.2, *Ethernet Services Attributes Phase 2*, October 2009.
- [2] Metro Ethernet Forum MEF 23, *Carrier Ethernet Class of Service Phase 1*, June 2009.
- [3] ITU-T Y.1563, *Ethernet frame transfer and availability performance*, 2009.