Technical Specification

MEF 26.0.3

Service Level Specification Amendment to MEF 26

October 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.

Table of Contents

List of Figures

List of Tables

The purpose of this document is to add Service Level Specification performance attribute definitions and related requirements to MEF 26.[1] The following performance attributes are covered:

- One-way Frame Delay (Section 7.2.15.1),
- One-way Frame Delay Range (Section 7.2.15.1),
- One-way Mean Frame Delay (Section 7.2.15.1),
- One-way Inter-Frame Delay Variation (Section 7.2.15.2),
- One-way Frame Loss Ratio (Section 7.2.15.3),
- One-way Availability (Section 7.2.15.4),
- One-way High Loss Intervals (Section 7.2.15.5), and
- One-way Consecutive High Loss Intervals (Section 7.2.15.5).

The terminology items shown in Table A are added to Table 1 in MEF 26[1].

Term	**Definition**	**Source**
CHLI	Consecutive High Loss Interval	This document
Consecutive High Loss Interval	A sequence of small time intervals contained in *T*, each with a high frame loss ratio.	This document
HLI	High Loss Interval	This document
High Loss Interval	A small time interval contained in *T* with a high frame loss ratio.	This document
Resiliency Performance	The number of High Loss Intervals and Consecutive High Loss Intervals in a time interval *T*	This document

Table A – New Terminology

Section 7.2.15 of MEF 26, *External Network Network Interface – Phase 1*[1] is replaced with the following text. The requirement numbers in this document contain the suffix "A" to distinguish them from requirements in MEF 26. References to requirement numbers without the suffix "A" refer to existing requirements in MEF 26.

7.2.15 Service Level Specification (SLS)

The OVC Related Performance Service Attributes specify the frame delivery performance between External Interfaces (EI). Eight performance metrics are detailed in this specification:

[1] Metro Ethernet Forum, MEF 26, *External Network Network Interface – Phase* 1, January 2010.

1. One-way Frame Delay,
2. One-way Frame Delay Range,
3. One-way Mean Frame Delay,
4. Inter-Frame Delay Variation,
5. One-way Frame Loss Ratio,
6. One-way Availability,
7. One-way High Loss Intervals, and
8. One-way Consecutive High Loss Intervals.

These are similar to the performance attributes for an EVC described in MEF 10.2 [5] and MEF 10.2.1[2] but more general because both the UNI and ENNI are covered in this document.

These Performance Service Attributes describe the performance experienced by the Service Provider who is the user of the OVC. Methods for the Operator and the Service Provider to monitor the OVC performance to estimate this user experience are beyond the scope of this document.

An SLS can specify objectives for all or any subset of the OVC Performance Service Attributes.

[R1A] If an SLS contains an objective for a given OVC Related Performance Service Attribute, then the SLS **MUST** specify the related parameters for that objective.

OVC Related Performance Service Attributes, with the exception of Availability[3], apply to "Qualified" Service Frames and "Qualified" ENNI Frames, called Qualified Frames.

[R2A] An SLS **MUST** define Qualified Frames as follows for a given ordered pair of OVC End Points $<i,j>$, a given time interval T, and a given Class of Service Identifier:

1. Each frame **MUST** ingress at the EI where OVC End Point *i* is located.
2. Each frame **MUST** map to OVC End Point *i* via the End Point Map. (See Section 7.1.1 for the descriptions of the End Point Map at the ENNI and Section 7.5.2 for the description of the OVC End Point Map at the UNI.)
3. The first bit of each frame **MUST** arrive at the ingress EI within the time interval *T*, and within a time interval Δt smaller than *T* that has been designated as part of Available time (see Section 7.2.15.4),
4. Each frame **MUST** have the given Class of Service Identifier,

[2] Metro Ethernet Forum MEF 10.2.1, *Performance Attributes Amendment to MEF 10.2*, January, 2011.
[3] Availability is used to define Qualified Frames via item 3 in the list.

5. Each ingress frame that is subject to an ingress Bandwidth Profile **MUST** have an Ingress Bandwidth Profile compliance of Green, and

6. Each ingress frame that is not subject to an ingress Bandwidth Profile **MUST** meet one of the following two conditions:

   - The frame has no color identifier[4], or

   - The frame has a color identifier that indicates Green per the color indication requirements of MEF 23[10].

Such frames are referred to as **Qualified Frames**.

Recall that both OVC End Points in the ordered pair can be located at the same ENNI. See **[O1]** and Section 7.2.2.

The assessment of all performance attributes needs to account for unexpected arrival phenomena, such as frame duplication, or frames arriving in a different order from that observed on ingress, and the presence of these phenomena alone do not necessarily exclude a frame from the set of Qualified Frames. Details on how to monitor performance in the face of unexpected arrival phenomena are beyond the scope of this document.

7.2.15.1 One-way Frame Delay Performance for an OVC

This section defines three performance attributes: the One-way Frame Delay Performance corresponding to a percentile of the distribution, the One-way Mean Frame Delay, and the One-way Frame Delay Range.

The One-way Frame delay for a frame that ingresses at EI_1 and results in a frame that egresses at EI_2 is defined as the time elapsed from the reception of the first bit of the ingress frame at EI_1 until the transmission of the last bit of the first corresponding egress frame at EI_2. If the frame is erroneously duplicated in the Operator MEN and multiple copies delivered to EI_2, the delay is based on the first such copy delivered.

Note that this definition of One-way Frame Delay for a frame is the one-way[5] delay that includes the delays encountered as a result of transmission across the ingress and egress EIs as well as that introduced by the Operator MEN.

Note that when the path between two UNIs crosses one or more ENNIs, the UNI to UNI One-way Frame Delay, as defined in MEF 10.2 does not equal the sum of the One-way Frame Delay between each pair of EIs. This is because the sum will double count the time to transmit a frame across each ENNI. To see this, note that the definition of delay, D_O, between a UNI and an

[4] When there is no color identifier, this bullet means that the Service Frame is treated as if it were declared Green. An example is the use of the H Label as defined in MEF 23 [10] where acolor identifier is not specified per Table 2 of [10].

[5] One-way delay is difficult to measure and therefore one way delay may be approximated from two way measurements. However these techniques are beyond the scope of this document.

ENNI on single Operator MEN can be expressed as $D_O = d_U + d_E + d_M$ were d_U is the time to transmit the frame across the UNI, d_E is the time to transmit the frame across the ENNI, and d_M is the queuing and transmission delay introduced by the Operator MEN. Now consider the case where Operator MEN 1 and Operator MEN 2 are connected via an ENNI to affect an EVC between two UNIs, one on each Operator MEN. The delay between the UNIs is $d_{U1} + d_{M1} + d_E + d_{M2} + d_{U2}$. But

$$D_{O1} + D_{O2} = d_{U1} + d_{M1} + d_{U2} + d_{M2} + 2d_E \neq d_{U1} + d_{M1} + d_E + d_{M2} + d_{U2}.$$

Adding the two OVC delays overstates the UNI to UNI delay by d_E.

This effect will need to be taken into account when deriving the UNI to UNI delay performance from the delay performance of each Operator MEN in the path of the frame. This derivation is beyond the scope of this document.

[R3A] For a given non-empty set of ordered pairs of OVC End Points, a time interval *T*, and a given Class of Service Identifier, the SLS **MUST** define each Frame Delay Performance metric as follows:

- Let the OVC End Points associated by the OVC be numbered from 1 to *m*. And let S be a non-empty subset of the ordered pairs of OVC End Points associated by the OVC. That is $S \subseteq \{\langle i, j \rangle \mid i = 1,...,m, j = 1,...,m, i \neq j\}, S \neq \varnothing$.

- Let $\overline{d}_{Td}^{\langle i,j \rangle}$ represent the P_d-Percentile of one-way delay for all Qualified Frames delivered to the egress EI where OVC End Point *j* is located resulting from an ingress frame at the EI where OVC End Point *i* is located. If there are no such egress frames, then $\overline{d}_{Td}^{\langle i,j \rangle}$ = Undefined.

- Then the One-way Frame Delay Performance metric **MUST** be defined as the maximum value of all of the defined values $\overline{d}_{Td}^{\langle i,j \rangle}$ for $\langle i, j \rangle \in S$, unless all $\overline{d}_{Td}^{\langle i,j \rangle}$ are Undefined in which case the performance is Undefined.

- Let $\overline{d}_{TR}^{\langle i,j \rangle} = \overline{d}_{Tr}^{\langle i,j \rangle} - d_{\min}^{\langle i,j \rangle}$. $\overline{d}_{Tr}^{\langle i,j \rangle}$ represents the P_r-Percentile of the one-way delay for all Qualified Frames delivered to the egress EI where OVC End Point *j* is located resulting from an ingress frame at the EI where OVC End Point *i* is located. $d_{\min}^{\langle i,j \rangle}$ is the minimum of the one-way delays for all Qualified Frames delivered to the EI where OVC End Point *j* is located resulting from an ingress frame at the EI where OVC End Point *i* is located. If there are no such egress frames, then $\overline{d}_{TR}^{\langle i,j \rangle}$ = Undefined.

- Then the One-way Frame Delay Range Performance metric **MUST** be defined as the maximum value of all of the defined values of $\overline{d}_{TR}^{\langle i,j\rangle}$ for $\langle i, j\rangle \in S$, unless all $\overline{d}_{TR}^{\langle i,j\rangle}$ are Undefined in which case the performance is Undefined.

- Let $\overline{\mu}_{T}^{\langle i,j\rangle}$ represent the arithmetic mean of one-way delay for all Qualified Frames delivered to the EI where OVC End Point *j* is located resulting from an ingress frame at the EI where OVC End Point *i* is located. If there are no such egress frames, then $\overline{\mu}_{T}^{\langle i,j\rangle}$ = Undefined.

- Then the One-way Mean Frame Delay Performance metric **MUST** be defined as the maximum value of all of the defined values $\overline{\mu}_{T}^{\langle i,j\rangle}$ for $\langle i, j\rangle \in S$, unless all $\overline{\mu}_{T}^{\langle i,j\rangle}$ are Undefined in which case the performance is Undefined.

To restate the Frame Delay definition mathematically, let the OVC End Points associated by the OVC be numbered from 1 to *m* and let $D_{T}^{\langle i,j\rangle}$ be the set of one-way Frame Delay values for all Qualified Frames at the EI where OVC End Point *j* is located resulting from an ingress frame at the EI where OVC End Point *i* is located. $D_{T}^{\langle i,j\rangle}$ can be expressed as $D_{T}^{\langle i,j\rangle} = \left\{d_{1}^{\langle i,j\rangle}, d_{2}^{\langle i,j\rangle}, \ldots, d_{N_{\langle i,j\rangle}}^{\langle i,j\rangle}\right\}$, where $d_{k}^{\langle i,j\rangle}$ is the one-way Frame Delay of the k^{th} frame. Define $\overline{d}_{Td}^{\langle i,j\rangle}$ for $P_d > 0$ as

$$\overline{d}_{Td}^{\langle i,j\rangle} = \begin{cases} \min\left\{ d \mid P_d \leq \dfrac{100}{N_{\langle i,j\rangle}} \displaystyle\sum_{k=1}^{N_{\langle i,j\rangle}} I\left(d, d_{k}^{\langle i,j\rangle}\right)\right\} \text{if } N_{\langle i,j\rangle} \geq 1 \\ \mathit{Undefined} \text{ otherwise} \end{cases}$$

where $I(\cdot, d_k)$ is an indicator function defined by

$$I(d, d_k) \equiv \begin{cases} 1 \text{ if } d \geq d_k \\ 0 \text{ otherwise} \end{cases}.$$

$\overline{d}_{Td}^{\langle i,j\rangle}$ is the minimal delay during the time internal *T* that P_d percent of the frames do not exceed.

Then a One-way Frame Delay Performance metric for an OVC can be expressed as

$$\overline{d}_{T,S} = \begin{cases} \max\left\{\overline{d}_{Td}^{\langle i,j\rangle} \mid \langle i, j\rangle \in S \text{ and where } N_{\langle i,j\rangle} > 0\right\} \\ \mathit{Undefined} \text{ when all } N_{\langle i,j\rangle} = 0 \text{ for all } \langle i, j\rangle \in S \end{cases}.$$

The One-way Frame Delay attribute permits specification of multiple values for P_d, (P_0, P_1, P_2, …) and corresponding objectives ($\hat{d}_0^{\langle i,j\rangle}$, $\hat{d}_1^{\langle i,j\rangle}$, $\hat{d}_2^{\langle i,j\rangle}$, …).

Another parameter is the objective for the difference between the delay P_r percentile delay and $d_{\min}^{\langle i,j\rangle} = \min\left\{d \in D_T^{\langle i,j\rangle}\right\}$, expressed as

$$\bar{d}_{TR}^{\langle i,j\rangle} = \begin{cases} (\bar{d}_{Tr}^{\langle i,j\rangle} - d_{\min}^{\langle i,j\rangle}) \text{ if } N_{\langle i,j\rangle} > 0 \\ \mathit{Undefined} \text{ if } N_{\langle i,j\rangle} = 0 \end{cases}$$

where

$$\bar{d}_{Tr}^{\langle i,j\rangle} = \begin{cases} \min\left\{ d \mid P_r \le \dfrac{100}{N_{\langle i,j\rangle}} \sum_{k=1}^{N_{\langle i,j\rangle}} I\left(d, d_k^{\langle i,j\rangle}\right) \right\} \text{ if } N_{\langle i,j\rangle} \ge 1 \\ \mathit{Undefined} \text{ otherwise} \end{cases}.$$

Then a One-way Frame Delay Range Performance metric for an OVC can be expressed as

$$\bar{d}_{TR,S} = \begin{cases} \max\left\{ \bar{d}_{TR}^{\langle i,j\rangle} \mid \langle i,j\rangle \in S \text{ and where } N_{\langle i,j\rangle} > 0 \right\} \\ \mathit{Undefined} \text{ when all } N_{\langle i,j\rangle} = 0 \text{ for all } \langle i,j\rangle \in S \end{cases}.$$

Another One-way Frame Delay attribute is the arithmetic mean of $D_T^{\langle i,j\rangle}$, which can be expressed as

$$\bar{\mu}_T^{\langle i,j\rangle} = \begin{cases} \dfrac{1}{N_{\langle i,j\rangle}} \sum_{k=1}^{N_{\langle i,j\rangle}} \left(d_k^{\langle i,j\rangle} \right) \text{ if } N_{\langle i,j\rangle} > 0 \\ \mathit{Undefined} \text{ if } N_{\langle i,j\rangle} = 0 \end{cases}$$

Then a One-way Mean Frame Delay Performance metric for an OVC can be expressed as

$$\bar{\mu}_{T,S} = \begin{cases} \max\left\{ \bar{\mu}_T^{\langle i,j\rangle} \mid \langle i,j\rangle \in S \text{ and where } N_{\langle i,j\rangle} > 0 \right\} \\ \mathit{Undefined} \text{ when all } N_{\langle i,j\rangle} = 0 \text{ for all } \langle i,j\rangle \in S \end{cases}.$$

The parameters of a One-way Frame Delay Performance metric are given in Table B.

Parameter	Description
T	The time interval
S	Non-empty subset of the ordered pairs of OVC End Points associated by the OVC
COS ID	The Class of Service Identifier
P_d	A specific percentile for the Frame Delay Performance, $P_d > 0$
P_r	A specific percentile for the Frame Delay Range Performance, $P_r > 0$
$\hat{d}$	One-way Frame Delay Performance Objective corresponding to P_d
$\hat{d}_R$	One-way Frame Delay Range Objective corresponding to P_r
$\hat{\mu}$	One-way Mean Frame Delay Performance Objective

Table B – One-way Frame Delay Performance Parameters

[R4A] Given *T*, *S*, *COS ID*, P_d, and a one-way Frame Delay Performance objective $\hat{d}$, expressed in time units, an SLS **MUST** define the one-way Frame Delay Performance objective as met over the time interval *T* for the subset *S* if and only if $\bar{d}_{T,S} \leq \hat{d}$ or $\bar{d}_{T,S}$ is Undefined.

[R5A] Given *T*, *S*, *COS ID*, P_r, and a One-way Frame Delay Range Performance objective $\hat{d}_R$, expressed in time units, an SLS **MUST** define the one-way Frame Delay Range Performance objective as met over the time interval *T* for the subset *S* if and only if $\bar{d}_{TR,S} \leq \hat{d}_R$ or $\bar{d}_{TR,S}$ is Undefined.

[R6A] Given *T*, *S*, *COS ID*, a One-way Mean Frame Delay Performance objective $\hat{\mu}$, expressed in time units, the Frame Delay Performance **MUST** be defined by an SLS as met over the time interval *T* if and only if $\bar{\mu}_{T,S} \leq \hat{\mu}$ and $\bar{\mu}_{T,S}$ is Undefined.

Recall that if any of the above performance attributes are Undefined for time interval *T* and ordered pair $\langle i, j \rangle$, then the performance for that ordered pair is to be excluded from calculations on the performance of pairs in *S*.

[O1A] For a Point-to-Point OVC, *S* **MAY** include one or both of the ordered pairs of OVC End Points associated by the OVC.

[O2A] For a Multipoint-to-Multipoint OVC, *S* **MAY** be any subset of the ordered pairs of OVC End Points associated by the OVC so long as it is non-empty.

7.2.15.2 Inter-Frame Delay Variation Performance for an OVC

Inter-Frame Delay Variation (IFDV) is the difference between the one-way delays of a pair of selected frames. This definition is borrowed from RFC3393[6] where IP packet delay variation is defined.

Let a_i be the time of the arrival of the first bit of the i^{th} frame at the ingress EI, then the two frames *i* and *j* are selected according to the selection criterion:

$$\{a_j - a_i = \Delta\tau \quad and \quad j > i\}$$

Let r_i be the time frame *i* is successfully received (last bit of the frame) at the egress EI, then the difference in the delays encountered by frame *i* and frame *j* is given by $d_i - d_j$. Define

$$\Delta d_{ij} = \left|d_i - d_j\right| = \left|(r_i - a_i) - (r_j - a_j)\right| = \left|(a_j - a_i) - (r_j - r_i)\right|$$

With d_j being the delay of the j^{th} frame, a positive value for $d_i - d_j$ implies that the two frames are closer together at the egress EI while a negative value implies that the two frames are further apart at the egress EI. If either or both frames are lost or not delivered due to, for example, FCS violation, then the value Δd_{ij} is not defined and does not contribute to the evaluation of the Inter-Frame Delay Variation.

Figure A shows a depiction of the different times that are related to Inter-Frame Delay Variation Performance.

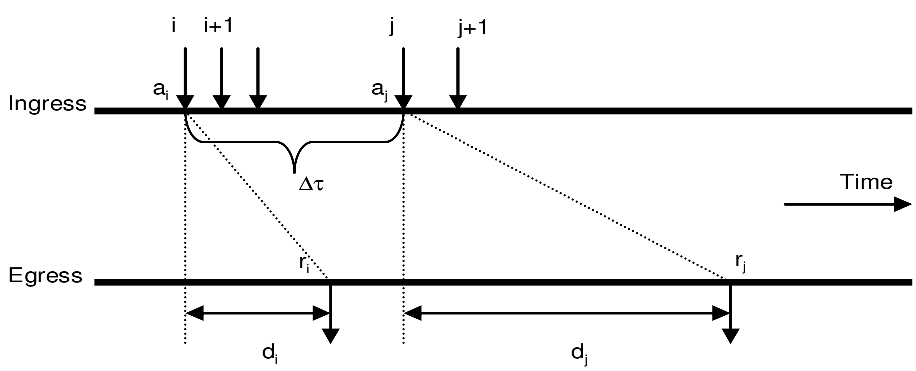

Figure A – Inter-Frame Delay Variation Definition

[R7A] For a given non-empty set of ordered pairs of OVC End Points, a time interval *T*, and a given Class of Service Identifier, the SLS **MUST** define the Inter-Frame Delay Variation Performance metric as follows:

[6] C. Demichelis and P. Chimento, *IP Packet Delay Variation Metric for IP Performance Metric (IPPM)*, RFC 3393, November 2002.

- Let the OVC End Points associated by the OVC be numbered from 1 to m. And let S be a non-empty subset of the ordered pairs of OVC End Points associated by the OVC. That is $S \subseteq \{\langle i, j\rangle \mid i = 1,...,m, j = 1,...,m, i \neq j\}, S \neq \emptyset$.

- Let $\Delta\tilde{d}_T^{\langle i,j\rangle}$ be the P_v-percentile of the absolute value of the difference between the Frame Delays of all Qualified Frame pairs whose difference in the arrival times of the first bit of each frame in the pair at EI where the OVC End Point i is located was exactly $\Delta\tau$.

- If there are no such pairs of frames for OVC End Point i and OVC End Point j, then $\Delta\tilde{d}_T^{\langle i,j\rangle} = Undefined$.

- Then the Inter-Frame Delay Variation Performance metric **MUST** be the maximum of the defined values $\Delta\tilde{d}_T^{\langle i,j\rangle}$ for $\langle i, j\rangle \in S$, unless all $\Delta\tilde{d}_T^{\langle i,j\rangle}$ are Undefined in which case the performance is Undefined.

This definition is in agreement with the IP packet delay variation definition given in RFC3393 where the delay variation is defined as the difference between the one-way delay of two packets selected according to some selection function and are within a given interval [T_1, T_2].

The choice of the value for $\Delta\tau$ can be related to the application timing information. As an example for voice applications where voice frames are generated at regular intervals, $\Delta\tau$ may be chosen to be few multiples of the inter-frame time.

To restate the definition mathematically, let the OVC End Points associated by the OVC be numbered from 1 to m. And let S be a non-empty subset of the ordered pairs of the OVC End Points associated by the OVC. That is

$$S \subseteq \{\langle i, j\rangle \mid i = 1,...,m, j = 1,...,m, i \neq j\}, S \neq \emptyset.$$

Let

$$V_T^{\langle i,j\rangle} = \{\Delta d_1^{\langle i,j\rangle}, \Delta d_2^{\langle i,j\rangle}, ..., \Delta d_{N_{\langle i,j\rangle}}^{\langle i,j\rangle}\}$$

be the set of all absolute value of delay variations for all eligible pairs of Qualified Frames from the EI where OVC End Point i is located to the EI where OVC End Point j is located where the difference in the arrival times of the first bit of each Service Frame at the ingress EI was exactly $\Delta\tau$. Define

$$\Delta\tilde{d}_T^{\langle i,j\rangle} = \begin{cases} \min\left\{d \mid P_v \leq \frac{100}{N_{\langle i,j\rangle}} \sum_{k=1}^{N_{\langle i,j\rangle}} I(d, \Delta d_k^{\langle i,j\rangle})\right\} \text{ if } N_{\langle i,j\rangle} \geq 1 \\ Undefined \text{ otherwise} \end{cases}$$

where $I(\cdot,\Delta d)$ is an indicator function defined by

$$I(d,\Delta d) \equiv \begin{cases} 1 \text{ if } d \geq \Delta d \\ 0 \text{ otherwise} \end{cases}.$$

Then an Inter-Frame Delay Variation Performance metric for an OVC can be expressed as

$$\Delta\tilde{d}_{T,S} = \begin{cases} \max\left\{\Delta\tilde{d}_T^{\langle i,j\rangle} \mid \langle i,j\rangle \in S \text{ and where } N_{\langle i,j\rangle} \geq 1\right\} \\ \mathit{Undefined} \text{ when all } N_{\langle i,j\rangle} = 0 \text{ for all } \langle i,j\rangle \in S \end{cases}.$$

The parameters of an Inter-Frame Delay Variation performance metric are given in Table C.

Parameter	Description
T	The time interval
S	Non-empty subset of the ordered pairs of OVC End Points associated by the OVC
COS ID	The Class of Service Identifier
P_v	Inter-Frame Delay Variation Performance percentile, $P_v > 0$
$\Delta\tau$	The separation between frame pairs for which Inter-Frame Delay Variation Performance is defined
$\breve{d}$	Inter-Frame Delay Variation Performance Objective

Table C – Inter-Frame Delay Variation Parameters

[R8A] Given T, S, *COS ID*, P_v, $\Delta\tau$, and $\breve{d}$, the Inter-Frame Delay Variation Performance objective **MUST** be defined by an SLS as met over the time interval T for the subset S if and only if $\Delta\tilde{d}_{T,S} \leq \breve{d}$ or $\Delta\tilde{d}_{T,S}$ is Undefined.

Recall that if the Inter-Frame Delay Variation is Undefined for time interval T and ordered pair $\langle i,j\rangle$, then the performance for that ordered pair is excluded from calculations on the performance of pairs in S.

[O3A] For a Point-to-Point OVC, S **MAY** include one or both of the ordered pairs of OVC End Points associated by the OVC.

[O4A] For a Multipoint-to-Multipoint OVC, S **MAY** be any subset of the ordered pairs of OVC End Points associated by the OVC so long as it is non-empty.

7.2.15.3 One-way Frame Loss Ratio Performance for an OVC

[R9A] For a given non-empty set of ordered pairs of OVC End Points, a time interval T, and a given Class of Service Identifier, the SLS **MUST** define the One-way Frame Loss Ratio Performance metric as follows:

- Let the OVC End Points associated by the OVC be numbered from 1 to *m*. And let S be a non-empty subset of the ordered pairs of OVC End Points associated by the OVC. That is $S \subseteq \{\langle i, j\rangle \mid i = 1,...,m,\ j = 1,...,m, i \neq j\}, S \neq \varnothing$.

- Let $I_T^{\langle i,j\rangle}$ denote the number of ingress Qualified Frames at the EI where OVC End Point *i* is located that should have been delivered to the EI where OVC End Point *j* is located according to the frame Delivery service attributes (see Sections 7.1.16, 7.2.17, and 7.2.18). Each frame can be a Unicast, Multicast, or Broadcast frame.

- Let $E_T^{\langle i,j\rangle}$ be the number of unique (not duplicate) egress frames where each frame is the first egress frame at the EI where OVC End Point *j* is located that results from a frame counted in $I_T^{\langle i,j\rangle}$.

- Define $FLR_T^{\langle i,j\rangle} = \begin{cases} \left(\dfrac{I_T^{\langle i,j\rangle} - E_T^{\langle i,j\rangle}}{I_T^{\langle i,j\rangle}}\right) \times 100 \text{ if } I_T^{\langle i,j\rangle} \geq 1 \\ \mathit{Undefined} \text{ otherwise} \end{cases}$.

- Then the One-way Frame Loss Ratio Performance metric **MUST** be defined as $FLR_{T,S} = \begin{cases} \max\left\{FLR_T^{\langle i,j\rangle} \mid \langle i, j\rangle \in S \text{ and where } I_T^{\langle i,j\rangle} \geq 1\right\} \\ \mathit{Undefined} \text{ when all } I_T^{\langle i,j\rangle} = 0 \text{ for all } \langle i, j\rangle \in S \end{cases}$.

The parameters of a One-way Frame Loss Ratio Performance metric are given in Table D.

Parameter	Description
T	The time interval
S	Non-empty subset of the ordered pairs of OVC End Points associated by the OVC
COS ID	The Class of Service Identifier
$\hat{L}$	One-way Frame Loss Ratio Performance objective

Table D – One-way Frame Loss Ratio Performance Parameters

[R10A] Given *T*, *S*, *COS ID*, and a One-way Frame Loss Ratio Performance objective, the One-way Frame Loss Performance objective **MUST** be defined by an SLS as met over the time interval *T* for the subset *S*if and only if $FLR_{T,S} \leq \hat{L}$ or $FLR_{T,S}$ is Undefined.

Recall that if the One-way Frame Loss Ratio Performance is Undefined for time interval *T* and ordered pair $\langle i, j\rangle$, then the performance for that ordered pair is to be excluded from calculations on the performance of pairs in *S*.

[O5A] For a Point-to-Point OVC, *S* **MAY** include one or both of the ordered pairs of OVC End Points associated by the OVC.

[O6A] For a Multipoint-to-Multipoint OVC, *S* **MAY** be any subset of the ordered pairs of OVC End Points associated by the OVC so long as it is non-empty.

7.2.15.4 One-way Availability Performance for an OVC

Availability Performance is the percentage of time within a specified time interval during which frame loss is small. (The precise definition is presented in the following paragraphs.) As an example, an Operator 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 (defined below) this objective will allow the service to be unavailable for approximately 43 minutes out of the whole month.

Informally, Availability Performance is based on frame loss during a sequence of consecutive small time intervals. 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.

For a time interval *T*, and a given Class of Service Identifier, Availability from OVC End Point *i* to OVC End Point *j* is based on the following three parameters:

- Δt, a time interval much smaller than *T*,
- *C*, a loss ratio threshold which if equaled or 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 Δt_k is Available and $A_{\langle i,j \rangle}(\Delta t_k)=0$ means that Δt_k is Unavailable.

The definition of $A_{\langle i,j \rangle}(\Delta t_k)$ is based on the frame loss ratio function $flr_{\langle i,j \rangle}(\Delta t_k)$ which is defined as follows.

Let $I_{\Delta t}^{\langle i,j \rangle}$ be the number of ingress frames that meet the following conditions:

- The first bit of each frame arrives at the EI where OVC End Point *i* is located within the time interval Δt,

- Each frame should be delivered to the EI where OVC End Point *j* is located according to the frame delivery service attributes (see Sections 7.2.16, 7.2.17, and 7.2.18). Each frame can be a Unicast, Multicast, or Broadcast frame,

- Each frame has the given Class of Service Identifier,

- Each frame that is subject to an ingress Bandwidth Profile has an Ingress Bandwidth Profile compliance of Green, and

- Each frame that is not subject to an ingress Bandwidth Profile either has no color indicator or a color identifier that corresponds to Green per the color indication requirements of MEF 23[10].

Let $E_{\Delta t}^{\langle i,j \rangle}$ be the number of unique (not duplicate) egress frames where each frame is the first, unerrored egress frame at the EI where OVC End Point *j* is located that results from a frame counted in $I_{\Delta t}^{\langle i,j \rangle}$.

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

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

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

where $\tilde{I}_{\Delta t}^{\langle i,j \rangle} = I_{\Delta t}^{\langle i,j \rangle} -$ the number of frames discarded by the Service Provider, in order to conform to either the line rate of the EI where OVC End Point *j* is located or an egress Bandwidth Profile (if one is used) at the EI where OVC End Point *j* is located. Such frame drops may occur anywhere in the network, not just near the egress EI. 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 Green frames for this OVC and Class of Service Identifier that should be delivered to the EI with OVC End Point *j* exceed the line rate on a link within the network, provided the line rate of that link is greater than or equal to the line rate of the EI. Good traffic engineering principles would suggest dropping such excessive frames as close to the ingress as possible. This adjustment is meant to account for a focused overload of traffic sent to the EI where OVC End Point *j* is located. The details of such an adjustment are beyond the scope of this document.

Δt_0 is the first short time interval agreed by Service Provider and the Operator at or after turning up the association of OVC End Point *i* and OVC End Point *j*. $A_{\langle i,j\rangle}(\Delta t_k)$ is defined in Figure B for $k = 0,1,2,\ldots$.

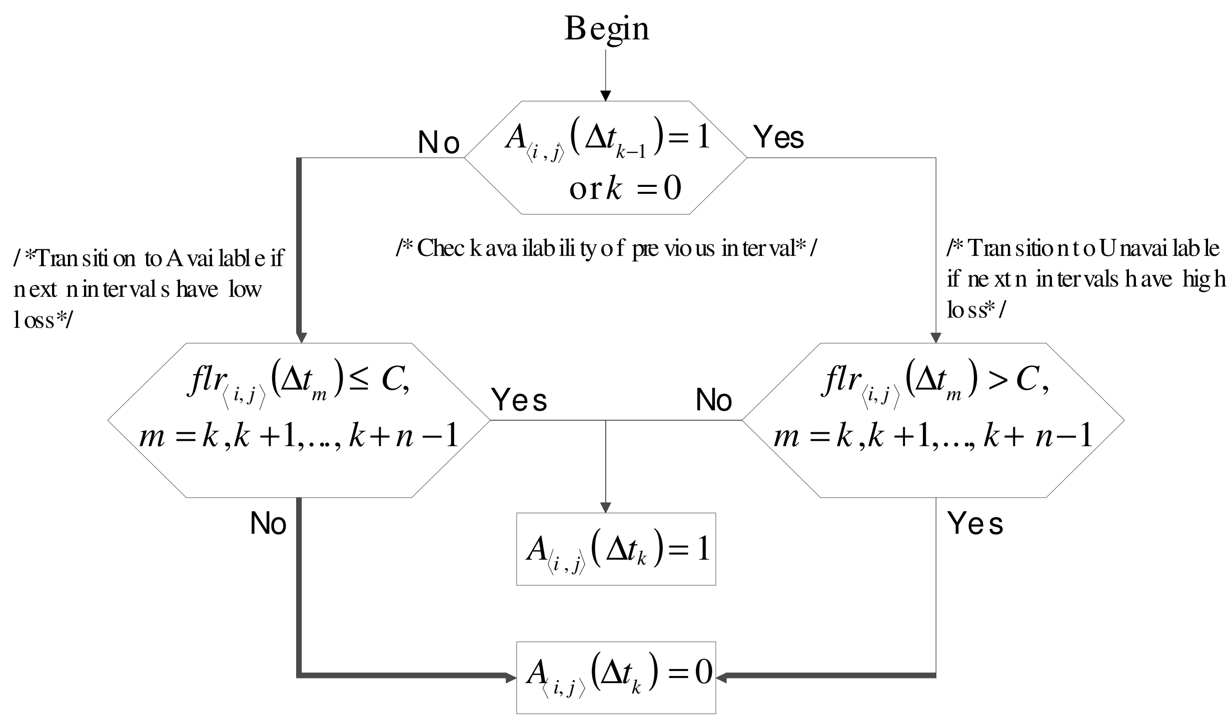

Figure B – Flowchart Definition of $A_{\langle i,j\rangle}(\Delta t_k)$

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

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

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

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

In the event of a conflict between the above equations and Figure B, the content of Figure B 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 define availability. This use of a sliding window is similar to that of ITU-T Y.1563.[7]

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

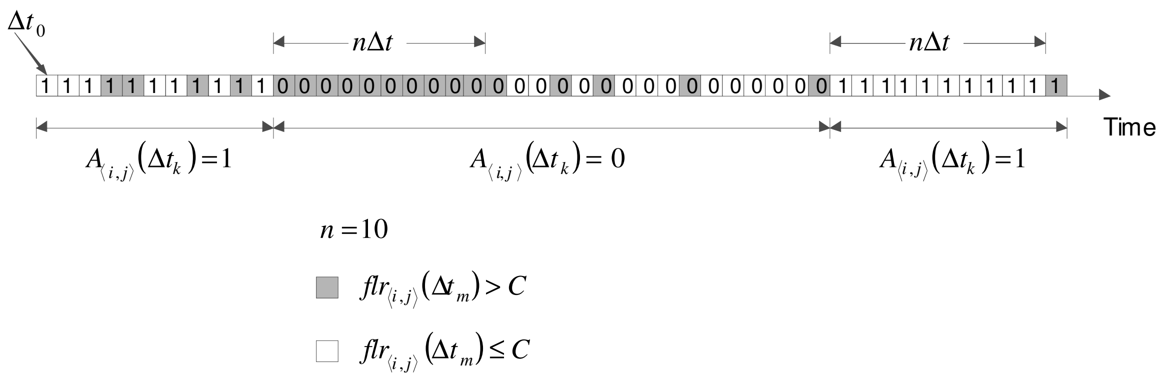

Figure C – Example of the Determination of $A_{\langle i,j \rangle}(\Delta t_k)$

The Availability for a particular Class of Service Identifier from OVC End Point i to OVC End Point j for a time interval T is based on the percentage of small time intervals that are Available. However, the small time intervals that occur during a Maintenance Interval are not included in the calculation of this percentage. A Maintenance Interval is a time interval agreed to by the Service Provider and Operator during which the service may not perform well or at all. Examples of a Maintenance Interval include:

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

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

[7] International Telecommunication Union, Recommendation Y.1563, Global information infrastructure, internet protocol aspects and next-generation networks, Internet protocol aspects – Quality of service and network performance, Ethernet frame transfer and availability performance, 2009.

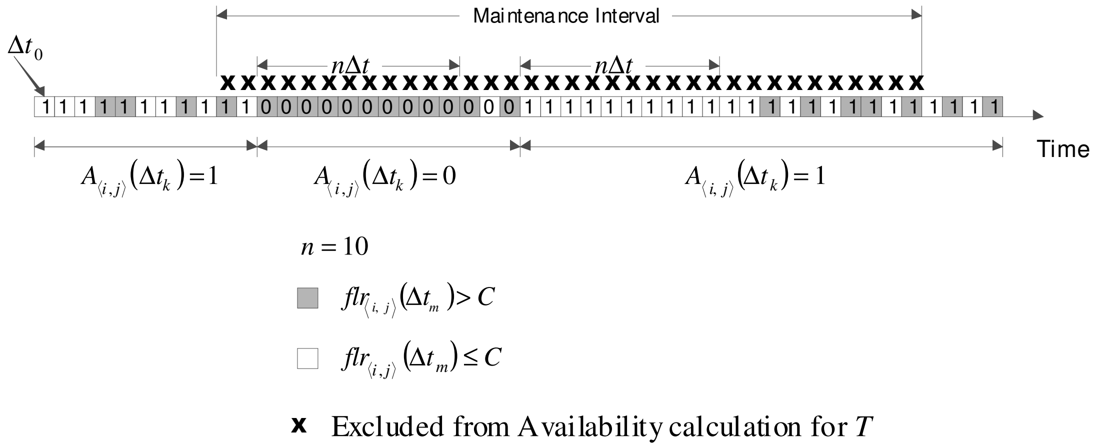

Figure D – Example of the Impact of a Maintenance Interval

Let $W_T = \{k \mid \Delta t_k \subseteq T$ and Δt_k 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 Identifier from OVC End Point *i* to OVC End Point *j* for a time interval *T,* expressed as percentage, is defined by

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

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*. A Δt_k that straddles the boundary of a Maintenance Interval is also excluded from the definition of Availability Performance.

Let the OVC End Points associated by the OVC be numbered $1,2,\ldots,m$ and let *S* be a non-empty subset of the ordered pairs of OVC End Points, i.e.,

$$S \subseteq \{\langle i, j \rangle \mid i = 1,2,\ldots m, j = 1,2,\ldots,m, i \neq j\}, S \neq \emptyset.$$

Then the Availability for a particular Class of Service Identifier for the set *S* is defined by

$$A_T^S = \min\{A_T^{\langle i,j \rangle} \mid \langle i, j \rangle \in S\}.$$

The parameters of a One-Way Availability Performance Metric are given in Table E.

Parameter	Description
T	The time interval
S	Non-empty subset of the ordered OVC End Point pairs
COS ID	The Class of Service Identifier
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 E – Availability Performance Parameters for an OVC

[R11A] Given T, S, *COS ID*, Δt, C, n, and $\hat{A}$, the SLS **MUST** define the Availability Performance objective as being met if and only if $A_T^S \geq \hat{A}$.

[O7A] For a Point-to-Point OVC, S **MAY** include one or both of the ordered pairs of OVC End Points associated by the OVC.

[O8A] For a Multipoint-to-Multipoint OVC, S **MAY** be any subset of the ordered pairs of OVC End Points associated the OVC so long as it is non-empty.

7.2.15.5 One-way Resiliency Performance for an OVC

This section defines attributes for the Resiliency performance of an ordered pair of OVC End Points, <*i*,*j*>. The definitions depend on the availability status of each Δt to determine whether 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 ITU-T Recommendation Y.1563[8], when $\Delta t = 1$ second.

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

[8] International Telecommunication Union, Recommendation Y.1563, *Ethernet frame transfer and availability performance*, 2009.

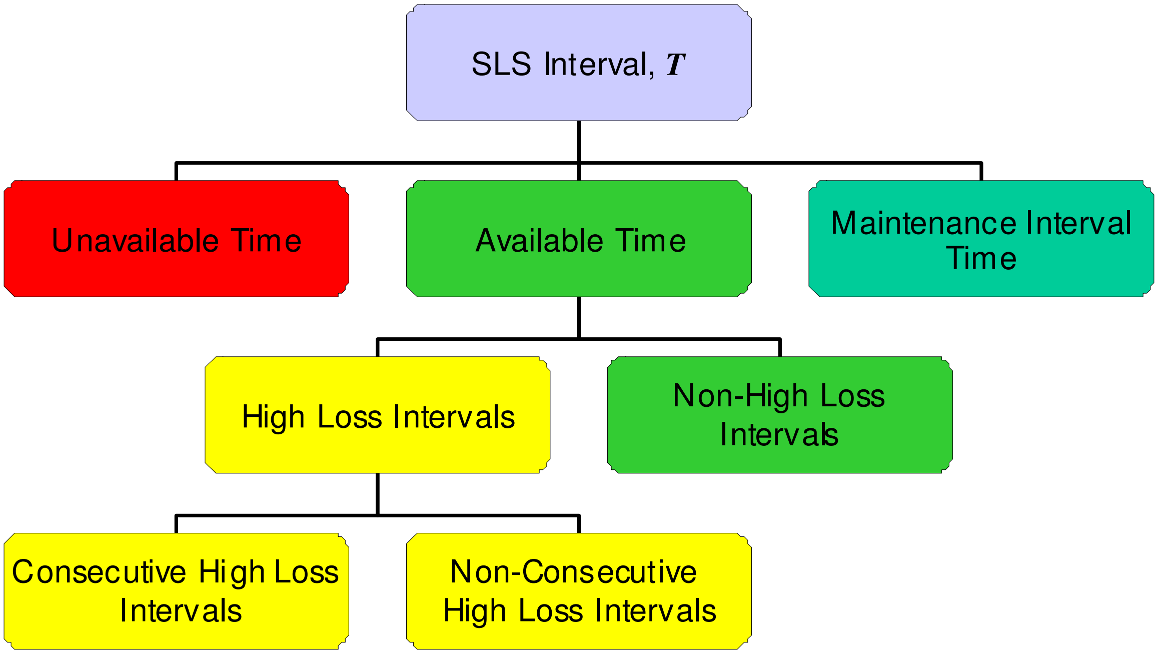

Figure E – Hierarchy of Time Showing the Resiliency Attributes

A High Loss Interval (HLI) is a small time interval (having the same duration as the interval, Δt, defined in Section 7.2.15.4) with a high frame loss ratio. When sufficient HLIs are adjacent, the interval is designated a Consecutive High Loss Interval (CHLI). Section 7.2.15.4 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 $flr_{\langle i,j\rangle}(\Delta t) > C$ 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 over T
- $\hat{L}$: HLI Count Objective for S, T, and a given Class of Service Identifier
- p: the minimum integer number of Δt's in the consecutive (sliding) window (with $0 < p < n$) to qualify as a Consecutive HLIs
- $B_T^{\langle i,j\rangle}$: Count of p or more consecutive High Loss Intervals occurring in T
- $\hat{B}$: Consecutive HLI Count Objective for S, T, and a given Class of Service Identifier

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

[R12A] For the SLS, the count of High Loss Intervals over T **MUST** be determined by $L_T^{\langle i,j \rangle} = \sum_{\Delta t \in T} H_{\langle i,j \rangle}(\Delta t)$.

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 $p < n$.

[R13A] For the SLS, the Consecutive High Loss Intervals over T **MUST** be determined according to the flow chart in Figure F.

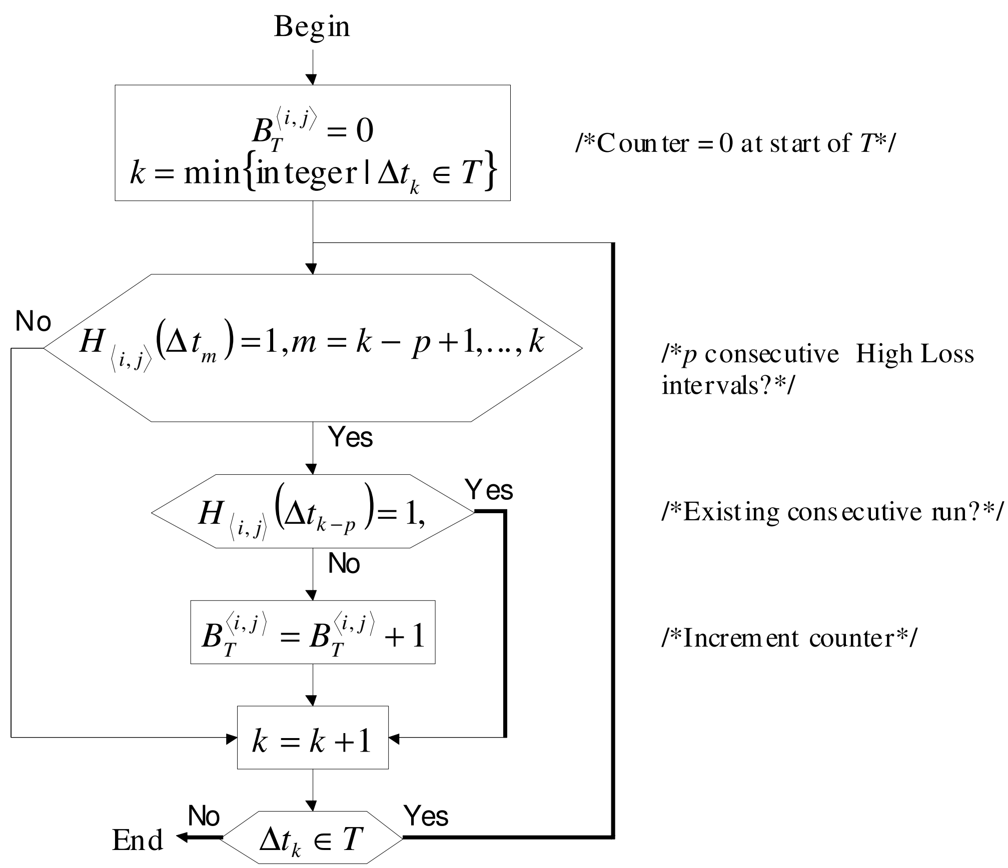

Figure F – Determining and Counting Consecutive High Loss Intervals over *T*

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

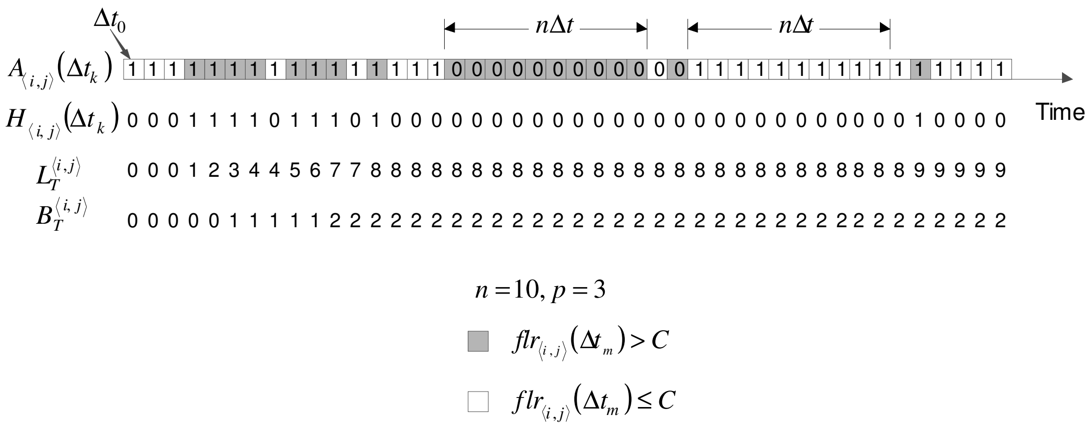

Figure G – Example of Counting High Loss Intervals and Consecutive High Loss Intervals

Let the OVC End Points associated by the OVC be numbered $1,2,\ldots,m$ and let S be a non-empty subset of the ordered pairs of OVC End Points, i.e.,

$$S \subseteq \{\langle i, j\rangle \mid i = 1,2,\ldots m,\ j = 1,2,\ldots,m, i \neq j\}, S \neq \varnothing .$$

Then the HLI and CHLI performance attributes for a particular Class of Service Identifier for the set S and time interval T are defined by

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

The parameters of the One-Way Resiliency Performance metrics are given in Table F.

Parameter	Description
T	The time interval
S	Non-empty subset of the ordered OVC End Point pairs associated by the OVC
COS ID	The Class of Service Identifier
Δt	A time interval much smaller than T
C	Unavailability frame loss ratio threshold
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 F – Resiliency Performance Parameters for an OVC

[R14A] Δt **MUST** have the same value as Δt in Section 7.2.15.4.

[R15A] C **MUST** have the same value as C in Section 7.2.15.4.

[R16A] Given T, S, $COS\ ID$, Δt, 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 $B_T^S \leq \hat{B}$.

[O9A] For a Point-to-Point OVC, S **MAY** include one or both of the ordered pairs of OVC End Points associated by the OVC.

[O10A] For a Multipoint-to-Multipoint OVC, S **MAY** be any subset of the ordered pairs of OVC End Points associated the OVC so long as it is non-empty.