

PCE Working Group  
Internet-Draft  
Intended Status: Standards Track  
Expires: March 2, 2018

R. Gandhi  
Cisco Systems, Inc.  
B. Wen  
Comcast  
C. Barth  
Juniper Networks  
D. Dhody  
Huawei Technologies  
August 29, 2017

PCEP Extensions for MPLS-TE LSP Performance Measurements  
with Stateful PCE  
draft-gandhi-pce-pm-08

## Abstract

The Path Computation Element Communication Protocol (PCEP) provides mechanisms for Path Computation Elements (PCEs) to perform path computations in response to Path Computation Clients (PCCs) requests.

The Stateful PCE extensions allow Stateful control of Multi-Protocol Label Switching (MPLS) Traffic Engineering Label Switched Paths (TE LSPs) using PCEP.

In certain networks, network performance data such as packet loss, delay and delay variation (jitter) as well as bandwidth utilization is a critical measure for Traffic Engineering (TE). This data provides operators the characteristics of their networks for performance evaluation that is required to ensure the Service Level Agreements (SLAs). Performance Measurement (PM) mechanisms can be employed to monitor these metrics end-to-end for TE Label Switched Paths (LSPs). This document describes Path Computation Element Protocol (PCEP) extensions for enabling and reporting such performance measurements to an Active Stateful PCE for both PCE-Initiated and PCC-Initiated LSPs.

## Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months

and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at  
<http://www.ietf.org/lid-abstracts.html>

The list of Internet-Draft Shadow Directories can be accessed at  
<http://www.ietf.org/shadow.html>

#### Copyright and License Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

1. Introduction . . . . .	5
1.1. Use-cases . . . . .	6
1.2. Dependencies and Considerations . . . . .	8
1.3. Auto-bandwidth Considerations . . . . .	9
2. Conventions Used in This Document . . . . .	9
2.1. Requirements Language . . . . .	9
2.2. Terminology . . . . .	9
2.3. Measurement Units . . . . .	10
3. Overview of the PCEP Extensions . . . . .	10
3.1. Report Thresholds . . . . .	11
4. Sub-TLVs for Measurements Attributes . . . . .	12
4.1. Measurement-Enable sub-TLV . . . . .	12
4.2. Measurement-Interval sub-TLV . . . . .	13
4.3. Report-Threshold sub-TLV . . . . .	13
4.4. Report-Threshold-Percentage sub-TLV . . . . .	14
4.5. Report-Interval sub-TLV . . . . .	14
4.6. Report-Upper-Bound sub-TLV . . . . .	15
5. PCEP Extensions for Reporting Delay Measurement . . . . .	16
5.1. Delay Measurement Capability Advertisement . . . . .	16
5.1.1. DELAY-MEASUREMENT-CAPABILITY TLV . . . . .	16
5.2. DELAY-MEASUREMENT-ATTRIBUTES TLV . . . . .	17
5.2.1. Delay Measurement Enable . . . . .	18
5.2.2. Delay Measurement Interval . . . . .	18
5.2.3. Delay Measurement Report Threshold . . . . .	18
5.2.4. Delay Measurement Report Threshold Percentage . . . . .	18
5.2.5. Delay Measurement Report Interval . . . . .	19
5.2.6. Delay Measurement Upper Bound . . . . .	19
5.3. DELAY-MEASUREMENT Object . . . . .	19
6. PCEP Extensions for Reporting Loss Measurement . . . . .	21
6.1. Loss Measurement Capability Advertisement . . . . .	21
6.1.1. LOSS-MEASUREMENT-CAPABILITY TLV . . . . .	22
6.2. LOSS-MEASUREMENT-ATTRIBUTES TLV . . . . .	23
6.2.1. Loss Measurement Enable . . . . .	24
6.2.2. Loss Measurement Interval . . . . .	24
6.2.3. Loss Measurement Report Threshold . . . . .	24
6.2.4. Loss Measurement Report Threshold Percentage . . . . .	24
6.2.5. Loss Measurement Report Interval . . . . .	25
6.2.6. Loss Measurement Upper Bound . . . . .	25
6.3. LOSS-MEASUREMENT Object . . . . .	25
7. PCEP Extensions for Reporting Bandwidth Utilization . . . . .	26
7.1. Bandwidth Utilization Capability Advertisement . . . . .	26
7.1.1. BANDWIDTH-UTILIZATION-CAPABILITY TLV . . . . .	27
7.2. BW-UTILIZATION-MEASUREMENT-ATTRIBUTES TLV . . . . .	27
7.2.1. Bandwidth Utilization Measurement Enable . . . . .	28
7.2.2. Bandwidth Utilization Measurement Interval . . . . .	28
7.2.3. Bandwidth Utilization Report Threshold . . . . .	28

7.2.4. Bandwidth Utilization Report Threshold Percentage . .	28
7.2.5. Bandwidth Utilization Report Interval . . . . .	29
7.2.6. Bandwidth Utilization Upper Bound . . . . .	29
7.3. BANDWIDTH Object . . . . .	29
8. PCEP Procedure . . . . .	29
8.1. Various MEASUREMENT-ATTRIBUTES TLVs . . . . .	30
8.2. The MEASUREMENT Objects . . . . .	30
9. Scaling Considerations . . . . .	30
9.1. The PCNtf Message . . . . .	31
10. Security Considerations . . . . .	31
11. Manageability Considerations . . . . .	31
11.1. Control of Function and Policy . . . . .	32
11.2. Information and Data Models . . . . .	32
11.3. Liveness Detection and Monitoring . . . . .	32
11.4. Verify Correct Operations . . . . .	32
11.5. Requirements On Other Protocols . . . . .	32
11.6. Impact On Network Operations . . . . .	32
12. IANA Considerations . . . . .	32
12.1. Measurement Capability TLV Types . . . . .	33
12.1.1. Flag Fields for MEASUREMENT-CAPABILITY TLVs . . . .	33
12.2. MEASUREMENT-ATTRIBUTES TLVs . . . . .	34
12.2.1. The Sub-TLVs For MEASUREMENT-ATTRIBUTES TLVs . . . .	34
12.2.1.1. Flag Fields in Measurement-Enable sub-TLV . . . .	34
12.3. Measurement Object-Class . . . . .	35
12.3.1. DELAY-MEASUREMENT Object-Types . . . . .	35
12.3.2. LOSS-MEASUREMENT Object-Types . . . . .	35
12.3.3. BANDWIDTH Object-Type . . . . .	36
12.4. PCE Error Codes . . . . .	36
12.5. Notification Object-Type . . . . .	36
13. References . . . . .	38
13.1. Normative References . . . . .	38
13.2. Informative References . . . . .	38
Acknowledgments . . . . .	40
Authors' Addresses . . . . .	40

## 1. Introduction

[RFC5440] describes the Path Computation Element Protocol (PCEP) as a communication mechanism between a Path Computation Client (PCC) and a Path Control Element (PCE), or between PCE and PCE, that enables computation of Multi-Protocol Label Switching (MPLS) Traffic Engineering Label Switched Paths (TE LSPs).

[DRAFT-PCE-STATEFUL] specifies extensions to PCEP to enable Stateful control of TE LSPs. It describes two mode of operations - Passive Stateful PCE and Active Stateful PCE. Further [DRAFT-PCE-INITIATED-LSP] describes the setup, maintenance and teardown of PCE-Initiated LSPs for the Stateful PCE model. In this document, the focus is on Active Stateful PCE where the LSPs are controlled by the PCE, for both PCE-Initiated and PCC-Initiated LSPs.

In certain networks, such as financial information networks, network performance data (e.g. packet loss, delay and delay variation (jitter), bandwidth utilization) is a critical measure for traffic engineering. The protocol extensions have been defined to advertise link performance metrics, see [RFC7471], [RFC7810], [RFC7823] and [DRAFT-IDR-TE-PM-BGP]. This data provides operators the characteristics of their networks for performance evaluation that is required to ensure the Service Level Agreements (SLAs).

[DRAFT-PCE-SERVICE-AWARE] defines the PCEP extensions for TE LSP path computation using packet loss, delay and delay variation as path selection metrics. Such path computations use link metrics for packet loss and delay and do not provide end-to-end metrics of the TE LSPs. The end-to-end metrics of a TE LSP may be very different than the path computation results due to many factors, such as queuing, etc. There is a need to verify and monitor that the traffic sent over the established TE LSPs does not exceed the requested metric bounds (e.g. total end-to-end delay/loss). The Stateful PCE may need to take some action (such as tear-down or re-optimize the LSP) when the performance requirement is not met. [RFC6374], [RFC6375] and [RFC7876] define protocol extensions needed for measuring end-to-end packet loss, delay and delay variation (jitter) for bidirectional and unidirectional TE LSPs.

This document provides mechanisms to enable and report the performance measurements (PM) such as packet loss, delay, delay variation (jitter) and bandwidth utilization for a TE LSP to an Active Stateful PCE, for both PCE-Initiated and PCC-Initiated LSPs.

### 1.1. Use-cases

This section describes a non-exhaustive list of deployment use-cases of PM for TE LSPs when deployed in a network with PCE. A controller may also be deployed in the network capable of "streaming telemetry" for receiving PM metrics and may interact with PCC and PCE for the PM as described in use-cases 3, 4 and 5.

#### Use-case 1: PCE Enables PM on PCC and PCC Takes Action

PCE -----> PCC

In this use-case, the PCE sets the upper-bound threshold condition for TE LSPs at the PCC. The PCC takes a local action when the condition is met. The action could be based on a local policy or policy set by the PCE. The steps involved are -

- o PCE sends the PM attributes as part of PCE-initiated LSPs including upper-bound threshold (Section 4.6 in this document) for the PM metrics using the PCEP extensions defined in this document.
- o PCC takes actions when PM metrics exceed the upper-bound threshold, actions could be to bring down the LSP, trigger protection switchover, remove tunnel from IGP for some prefixes, or request a new path from PCE (based on local policies which may be set by the PCE). PCC may take these actions even when LSPs are delegated to PCE as the upper-bound is set by the PCE.
- o PCC does not report the PM metrics to PCE.
- o PCC may install the new LSP in routing table only if the PM metric is below the upper-bound, otherwise, the PCC may reject the LSP request and send an error to the PCE.
- o The report interval should be set to 0 to disable reporting to PCE. Only the upper-bound threshold should be set.

#### Use-case 2: PCC Reports PM Metrics to PCE, PCE Takes Action

PCE <----- PCC

In this use-case, the PCC reports the PM metrics and parameters to the PCE and the PCE may take an immediate local (reactive) action based on the PM metrics. The steps involved are -

- o PCC sends the PM metrics and parameters to PCE using the PCEP extensions defined in this document and PCE takes an action; action could be to correlate faults, invalidate LSP path, send new

LSP path to PCC (trigger re-optimization), etc.

- o If upper-bound threshold is set, PCC only reports the PM metrics to PCE when upper-bound is crossed. Otherwise the PCC reports the PM metrics to PCE every report-interval.
- o Optionally, PCC may take an immediate local (reactive) action such as trigger path protection switch-over when PM metrics exceed upper-bound.
- o PCE has a global view due to PM metric reports received from various PCCs and hence can make a better decision about LSP placement in the network.
- o PCE can make pro-active decisions based on PM metrics when metrics are reported before crossing of the upper-bound as opposed to reactive action that PCC could make.
- o The report interval should be set to enable reporting by the PCC. Optionally, the upper-bound threshold may also be set.

Use-case 3: PCE Enables PM on PCC, PCC Sends PM Metrics to Controller

PCE -----> PCC -----> Controller

The steps involved are -

- o A controller may be used in a network that is capable of "streaming telemetry" for receiving data and Yang or XML based provisioning using non-PCEP channel. The controller may interact with a PCE for LSP path computation using the PCEP channel.
- o PCE sends the PM attributes as part of PCE-initiated LSPs using the PCEP extensions defined in this document.
- o PCC reports the PM metrics to controller via "streaming telemetry".
- o Controller may request PCE to take an action based on the PM metrics.
- o The report interval should be set to 0 to disable reporting to PCE. The other PM attributes may be set and used for "streaming telemetry".

Use-case 4: Controller Enables PM on PCC, PCC Sends PM Metrics to PCE

PCE <----- PCC <----- Controller

The steps involved are -

- o Controller enables PM on PCC using a non-PCEP channel.
- o PCC then reports the PM metrics to PCE using the PCEP extensions defined in this document.
- o PCE may take an action based on the PM metrics received from PCC.

Use-case 5: Controller Enables PM on PCC, PCC Sends PM Metrics to Controller

PCE ----> PCC <-----> Controller ----> PCE

The steps involved are -

- o Controller enables PM on PCC using a non-PCEP channel.
- o PCC reports the PM metrics to the controller via "streaming telemetry".
- o Controller may request PCE to take an action based on the PM metrics.
- o The PCEP extensions defined in this document are not used in this use-case.

## 1.2. Dependencies and Considerations

[RFC6374] describes several reasons why PMs are valuable to operators. Note that the specification of the use of the reported packet loss, delay, delay variation and bandwidth utilization measurements by a Stateful PCE is outside the scope of this document.

Furthermore, [RFC6374] describes many types of MPLS channels that may leverage PMs and some may have bidirectional dependencies. Defining a mechanism for the verification and/or provisioning of bidirectional or associated bidirectional LSPs within the Stateful PCE architecture is outside the scope of this document.

In all cases, delay and loss PM messages are carried over the MPLS Generic Associated Channel (G-ACh) as described in [RFC5586]. MPLS LSPs that carry the G-ACh can be referred to as MPLS Transport Profile (MPLS-TP) LSPs [RFC5921]. MPLS-TP LSPs require Ultimate Hop Popping (UHP) where LSPs are assigned Non-NULL labels by tail-end nodes. It is beyond the scope of this document to define the

mechanism by which a Stateful PCE verifies and/or provisions an LSP for UHP. Note that for both unidirectional and bidirectional LSPs, packet loss measurement requires UHP.

### 1.3. Auto-bandwidth Considerations

Auto-Bandwidth feature allows a head-end LSR (PCC) to automatically adjust the LSP bandwidth reservation based on the traffic demand of a TE LSP. Auto-bandwidth requested bandwidth computation can be implemented on a PCC or a Stateful PCE.

[DRAFT-IETF-PCE-AUTOBW] describes the PCEP extensions for auto-bandwidth, where the requested bandwidth for the LSP is computed by the PCC and reported to the Stateful PCE. There is a benefit in pushing the responsibility for deciding when auto-bandwidth adjustments are needed to the PCC as this distributes the load of monitoring the bandwidth utilization of the LSPs down to the PCCs and frees up the PCE for path computations. In addition, it reduces the load on PCEP communications for reporting the bandwidth utilization of the LSP.

However, exactly when to adjust an LSP bandwidth could be better left to a Stateful PCE. That is, a PCE could be flexible in its interpretation of thresholds enabling it to trigger auto-bandwidth adjustment early if it believes there is a good reason (for example, doing a set of parallel path re-computations) or late (for example, when it knows that an adjustment would be disruptive to the network).

When the auto-bandwidth computation is delegated to the PCC, the PCC cannot see the impact on other LSPs in the network, and the PCE cannot tell whether the request to adjust the LSP bandwidth is critical or not. The bandwidth utilization reporting defined in this document can be used by the PCE to do computations to determine whether auto-bandwidth adjustments are needed and/or desirable before performing the path computations.

## 2. Conventions Used in This Document

### 2.1. Requirements Language

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 [RFC2119].

### 2.2. Terminology

The reader is assumed to be familiar with the terminology defined in [RFC5440], [RFC6374] and [RFC7471].

### 2.3. Measurement Units

The measurement unit for delay value is defined in [RFC7471], Section 4.1.5.

The measurement unit for loss value is defined in [RFC7471], Section 4.4.5.

The utilized bandwidth [RFC7471] is encoded in IEEE floating point format in bytes per second (see [IEEE.754.1985]).

All average values are calculated as rolling averages.

### 3. Overview of the PCEP Extensions

The high-level overview of the PCEP extensions defined in this document for requesting and reporting end-to-end performance measurements and bandwidth utilization for TE LSPs are shown in Figure 1.

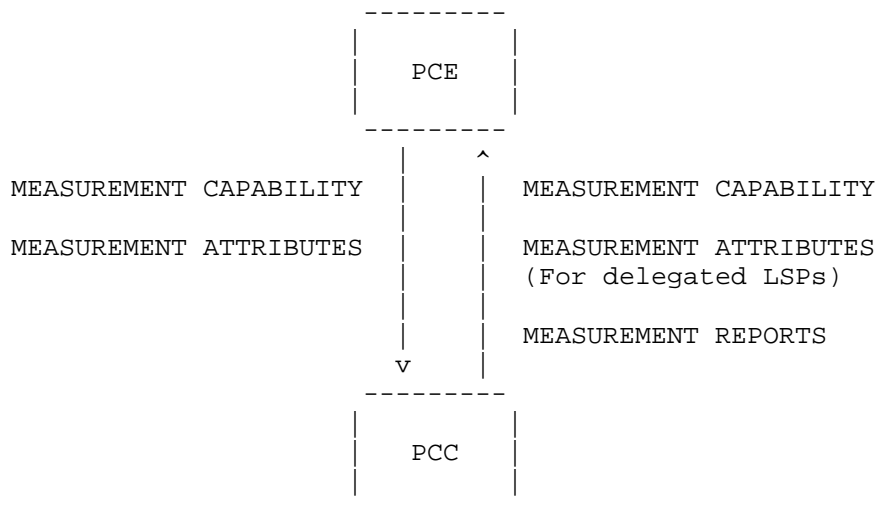


Figure 1: Overview of PCEP Extensions

The following steps describe the PCEP extensions for reporting performance measurements and bandwidth utilization of TE LSPs:

- o The Stateful PCE and PCC (head-end of the LSP) advertise the capability of their support for the delay, loss and bandwidth-utilization measurements and reporting in the PCEP Open message

(in the OPEN Object).

- o The Stateful PCE enables the measurement of a feature and sends and updates the configuration parameters of the feature using the LSPA object to PCC in PCInitiate and PCUpd messages respectively.
- o The PCC reports the measured values of the feature to the Stateful PCE at the end of the specified report-interval or when measured values cross the specified report-threshold. The periodic reporting can be used at the PCE to monitor the TE LSP metrics whereas report-threshold can be used to trigger an immediate action at the PCE on the TE LSP.
- o In some cases, the periodic reporting of the measurements may be disabled and only an upper-bound threshold is set, which when exceeded, a local or PCE-set action may be taken.
- o The PCE and PCC notify each other of their entering and clearing the overwhelmed state when operating under high LSP scale.

### 3.1. Report Thresholds

When explicitly configured, report threshold (absolute or percentage) parameter (along with the configured number of counts) is used to trigger an immediate reporting of the delay and loss measurements and bandwidth utilization, bypassing the report interval. Threshold is used to detect a sudden change in the performance measurement metric of a TE LSP. In order to detect that a measured value has crossed the threshold, the measured (delay/loss) metric is compared with the last reported value. If the change (increase or decrease) in the value is above the threshold (absolute or percentage) consecutively for the given number of counts, the measurement from the current interval is reported immediately. In case of bandwidth utilization, the last reported MaxSampleBw (see [DRAFT-IETF-PCE-AUTOBW]) value is compared with the MaxSampleBW from the current interval to detect the threshold crossing. The delay and loss measurements are still reported at the end of the report interval even if they were reported due to the crossing of the threshold. Refer to [RFC7471], Section 5, for additional considerations.

All thresholds in this document could be represented in both absolute value and percentage, and could be used together. This is provided to accommodate the cases where the metric values may become very large or very small over time. For example, an operator may use the percentage threshold to handle small to large metric values and absolute values to handle very large metric values. The metrics are reported when either one of the two thresholds, the absolute or

percentage, is crossed.

When using the percentage threshold, if the metric changes rapidly at very low values, it may trigger frequent reporting due to the crossing of the percentage threshold. This can lead to unnecessary scale issues in the network. This is suppressed by setting the minimum-threshold parameter along with the percentage threshold. The metric value is only reported if the value crosses both the percentage threshold and the minimum-threshold parameters.

#### 4. Sub-TLVs for Measurements Attributes

This section specifies the generic sub-TLVs those provide various configurable parameters for reporting measurements to a Stateful PCE. These sub-TLVs are carried in various measurement attributes TLVs defined in this document.

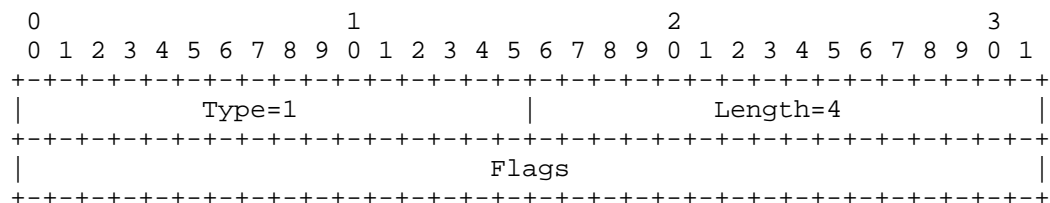
Following sub-TLVs are defined:

Type	Len	Name
1	4	Measurement-Enable sub-TLV
2	4	Measurement-Interval sub-TLV
3	8	Report-Threshold sub-TLV
4	8	Report-Threshold-Percentage sub-TLV
5	4	Report-Interval sub-TLV
6	8	Report-Upper-Bound sub-TLV

The Measurement-Enable sub-TLV MUST be added in the LSPA Object when the measurement feature is enabled for the LSP. All other sub-TLVs are optional and any unrecognized sub-TLV MUST be silently ignored. If a sub-TLV of same type appears more than once, only the first occurrence is processed and all others MUST be ignored. If sub-TLVs are not present, the default values based on the local policy are assumed.

##### 4.1. Measurement-Enable sub-TLV

The Measurement-Enable sub-TLV specifies that the given measurement feature is enabled.



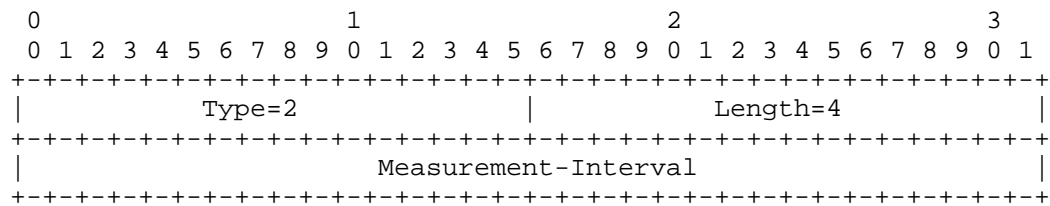
## Measurement-Enable sub-TLV Format

The Type is 1, Length is 4 bytes, and the value comprises of flags (32 bits) for enabling various measurement features.

Unassigned flags are considered reserved, they MUST be set to 0 when sent and MUST be ignored when received.

## 4.2. Measurement-Interval sub-TLV

The Measurement-Interval sub-TLV specifies a time interval in seconds for the measurement.

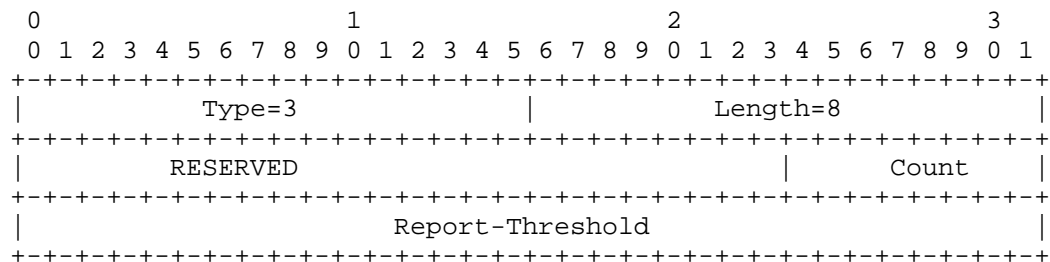


## Measurement-Interval sub-TLV Format

The Type is 2, Length is 4 bytes, and the value comprises of 4-byte time interval, the valid range is from 1 to 604800, in seconds. The default value is 300 seconds. The Measurement-Interval MUST NOT be greater than Report-Interval.

## 4.3. Report-Threshold sub-TLV

The Report-Threshold sub-TLV specifies the threshold value used to trigger an immediate reporting of the measurements bypassing the report-interval.



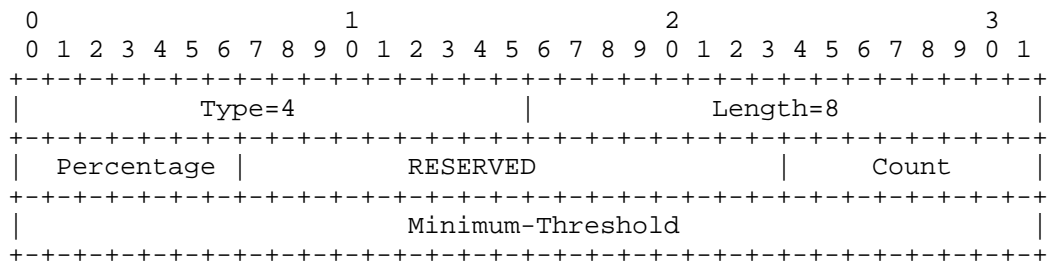
## Report-Threshold sub-TLV Format

The Type is 3, Length is 8 bytes, and the value comprises of -

- o Report-Threshold: 32-bit absolute threshold value.
- o Count: 8-bit integer counter. The number of consecutive measurement values MUST be above the threshold before reporting the measurement. The value 0 is considered to be invalid. By default, report-threshold is not set and threshold check based reporting is disabled.
- o RESERVED: It MUST be set to zero when sent and MUST be ignored when received.

#### 4.4. Report-Threshold-Percentage sub-TLV

The Report-Threshold-Percentage sub-TLV specifies the threshold value used to trigger an immediate reporting of the measurements bypassing the report-interval.



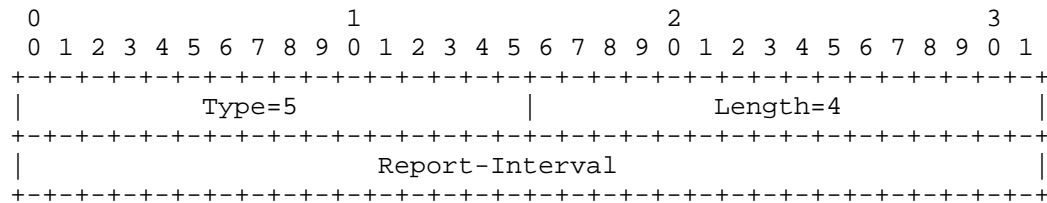
Report-Threshold-Percentage sub-TLV Format

The Type is 4, Length is 8 bytes, and the value comprises of -

- o Percentage: 7-bit threshold value, encoded in percentage (an integer from 1 to 100).
- o Count: 8-bit integer counter. The number of consecutive measurement values MUST be above threshold before reporting the measurement. The value 0 is considered to be invalid. By default, report-threshold-percentage is not set and threshold check based reporting is disabled.
- o RESERVED: It MUST be set to zero when sent and MUST be ignored when received.
- o Minimum-Threshold: The 32-bit absolute Minimum-Threshold value. The increase or decrease should be at least or above this value.

#### 4.5. Report-Interval sub-TLV

The Report-Interval sub-TLV specifies the time interval in seconds when measured values are to be reported.

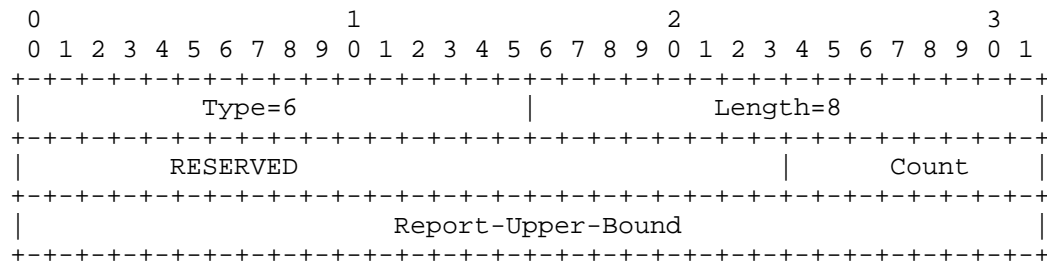


Report-Interval sub-TLV Format

The Type is 5, Length is 4 bytes, and the value comprises of 4-byte time interval, the valid range is from 0 to 604800, in seconds. The default value is 3600 seconds. The value 0 is used to disable the periodic reporting of the measurements.

#### 4.6. Report-Upper-Bound sub-TLV

The Report-Upper-Bound sub-TLV specifies the upper-bound value used to trigger an immediate reporting of the measurements when crossed. This may also result in PCC taking an immediate local action on the LSP.



Report-Upper-Bound sub-TLV Format

The Type is 6, Length is 8 bytes, and the value comprises of -

- o Report-Upper-Bound: 32-bit absolute value.
- o Count: 8-bit integer counter. The number of consecutive measurement values MUST be above the upper-bound before reporting the measurement. The value 0 is considered to be invalid. By default, upper-bound is not set.
- o RESERVED: It MUST be set to zero when sent and MUST be ignored when received.



```

+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                                           Flags                                           |T|O|
+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

#### DELAY-MEASUREMENT-CAPABILITY TLV Format

The Type of the TLV is TBD1 and it has a fixed length of 4 bytes.

The value comprises a single field - Flags (32 bits):

- o O (One-way Delay Measurement - 1 bit): if set to 1 by a PCC, the O Flag indicates that the PCC allows reporting of one-way delay measurement information; if set to 1 by a PCE, the O Flag indicates that the PCE is capable of receiving one-way delay measurement information from the PCC.
- o T (Two-way Delay Measurement - 1 bit): if set to 1 by a PCC, the T Flag indicates that the PCC allows reporting of two-way delay measurement information; if set to 1 by a PCE, the T Flag indicates that the PCE is capable of receiving two-way delay measurement information from the PCC. Two-way measurement is only applicable to the bidirectional LSPs (e.g. MPLS-TP LSPs [RFC5921]).

Unassigned bits are considered reserved. They MUST be set to 0 when sent and MUST be ignored when received.

Advertisement of the DELAY-MEASUREMENT-CAPABILITY TLV implies support of delay measurement, as well as the objects, TLVs and procedures defined in this document. Either O or T flag MUST be set in the TLV.

#### 5.2. DELAY-MEASUREMENT-ATTRIBUTES TLV

The DELAY-MEASUREMENT-ATTRIBUTES TLV provides the configurable parameters of the delay measurement feature.

The format of the DELAY-MEASUREMENT-ATTRIBUTES TLV is shown in the following figure:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|           PCEP TLV Type TBD4           |           Length           |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+
|                                           |
//                               sub-TLVs                               //
|                                           |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+

```

## DELAY-MEASUREMENT-ATTRIBUTES TLV Format

PCEP TLV Type is defined as following:

Type	Name
-----	
TBD4	DELAY-MEASUREMENT-ATTRIBUTES

Length: The Length field defines the length of the value portion in bytes as per [RFC5440].

Value: Comprises of one or more sub-TLVs as described in Section 4 of this document.

The following sub-sections describe the parameters which are currently defined to be carried within this TLV.

#### 5.2.1. Delay Measurement Enable

The Measurement-Enable sub-TLV specifies the delay measurement mode enabled using following flags:

Bit	Description
-----	
31	One-Way Delay Measurement Enabled
30	Two-Way Delay Measurement Enabled

#### 5.2.2. Delay Measurement Interval

The Measurement-Interval sub-TLV specifies a time interval in seconds for the delay measurement.

#### 5.2.3. Delay Measurement Report Threshold

The Report-Threshold sub-TLV specifies the threshold value used to trigger an immediate reporting of the delay measurements bypassing the report-interval.

- o Report-Threshold: Delay in microseconds, encoded as 24-bit integer, as defined in [RFC7471].

Same report-threshold is used for all delay measurement values.

#### 5.2.4. Delay Measurement Report Threshold Percentage

The Report-Threshold-Percentage sub-TLV specifies the threshold value used to trigger an immediate reporting of the measurements bypassing the report-interval.

Same report-threshold-percentage is used for all delay measurement values.

#### 5.2.5. Delay Measurement Report Interval

The Report-Interval sub-TLV specifies the time interval in seconds when measured delay values are to be reported.

#### 5.2.6. Delay Measurement Upper Bound

The Report-Upper-Bound sub-TLV specifies the upper-bound value in microseconds, and is used to trigger an immediate reporting of the delay values when crossed. This may also result in PCC taking an immediate local action on the LSP.

### 5.3. DELAY-MEASUREMENT Object

A new object, DELAY-MEASUREMENT with Object-Class (Value TBD5) is defined in this document to report the delay measurement of a TE LSP.

When the LSP is enabled with the delay measurement feature, the PCC SHOULD include the DELAY-MEASUREMENT Object to report the measured delay values to the PCE in the PCRpt message. The PCC SHOULD report (either one-way or two-way) average delay, min/max delay and delay variations to the PCE in the PCRpt message.

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-----+-----+-----+-----+-----+-----+-----+-----+
| Class=TBD5 | OT | Res | P | I | Object Length (bytes) |
+-----+-----+-----+-----+-----+-----+-----+
|                                                     |
|//                               (Object body)                               //
|                                                     |
+-----+-----+-----+-----+-----+-----+-----+

```

#### DELAY-MEASUREMENT Object Format

Object Length (16 bits): Specifies the total object length including the header, in bytes as per [RFC5440].

Object-Types (OT) are defined as follows:

Object-Type	Len	Name
1	8	One-Way Delay Measurement Value

2	12	One-Way Delay Measurement Min/Max Values
3	8	One-Way Delay Variation Measurement Value
4	8	Two-Way Delay Measurement Value
5	12	Two-Way Delay Measurement Min/Max Values
6	8	Two-Way Delay Variation Measurement Value

The object body formats are defined as following:

For Object-Types 1 and 4:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   RESERVED   |           Delay Value Average           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

For Object-Types 2 and 5:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   RESERVED   |           Delay Value Minimum           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   RESERVED   |           Delay Value Maximum           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

For Object-Types 3 and 6:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|   RESERVED   |           Delay Variation Value           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

#### DELAY-MEASUREMENT Object Body Formats (One-Way and Two-Way)

All delay values are reported in microseconds, encoded as 24-bit integer, as defined in [RFC7471]. When set to the maximum value 16,777,215 (16.777215 sec), the delay is at least that value and may be larger.

- o One-way Delay Measurement Value: Average Delay of the LSP in one (forward) direction.

- o One-way Delay Measurement Variation Value: Average Delay Variation of the LSP in one (forward) direction.
- o One-Way Delay Measurement Value Minimum/Maximum: Minimum and Maximum values of the Delay of the LSP in one (forward) direction in the last measurement interval.
- o Two-way Delay Measurement Value: Average Delay of the bidirectional LSP in both (forward and reverse) directions.
- o Two-way Delay Measurement Variation Value: Average Delay Variation of the bidirectional LSP in both (forward and reverse) directions.
- o Two-Way Delay Measurement Value Minimum/Maximum: Minimum and Maximum values of the Delay of the bidirectional LSP in both (forward and reverse) directions in the last measurement interval.
- o RESERVED: This field is reserved for future use. It MUST be set to 0 when sent and MUST be ignored when received.

## 6. PCEP Extensions for Reporting Loss Measurement

### 6.1. Loss Measurement Capability Advertisement

During PCEP Initialization Phase, PCEP Speakers (PCE or PCC) advertise their support of LOSS-MEASUREMENT. A PCEP Speaker includes the LOSS-MEASUREMENT-CAPABILITY TLV, in the OPEN Object to advertise its support for PCEP Loss-Measurement extensions. The presence of the LOSS-MEASUREMENT-CAPABILITY TLV in the OPEN Object (in the Open message) indicates that the Loss Measurement capability is supported as described in this document. Additional procedure is defined as following:

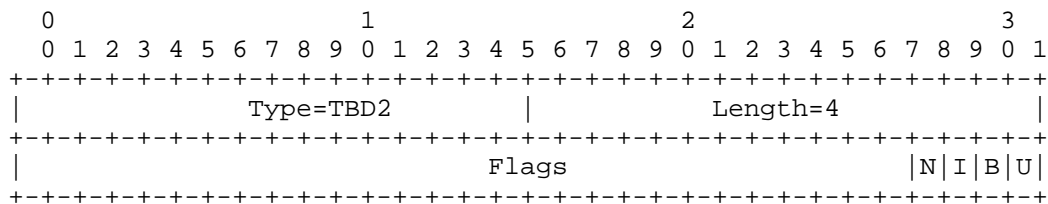
- o The PCEP protocol extensions for Loss Measurement MUST NOT be used if one or both PCEP Speakers have not included the LOSS-MEASUREMENT-CAPABILITY TLV in their respective Open message.
- o If the PCEP speaker that supports the extensions of this document but did not advertise this capability, then upon receipt of LOSS-MEASUREMENT-ATTRIBUTES TLV in LSPA object, it SHOULD generate a PCErr with error-type 19 (Invalid Operation), error-value TBD8 (Loss-Measurement capability was not advertised) and it will terminate the PCEP session.
- o Similarly, the PCEP speaker SHOULD generate error-value TBD9 (Bidirectional Measurement capability was not advertised) and

TBD10 (Unidirectional Measurement capability was not advertised) upon receipt of LOSS-MEASUREMENT-ATTRIBUTES TLV in LSPA object with two-way measurement request and one-way measurement request, respectively, when it did not advertise this capability.

- o Further, the PCEP speaker SHOULD generate error-value TBD11 (Inferred Mode Loss Measurement capability was not advertised) and TBD12 (Direct Mode Loss Measurement capability was not advertised) upon receipt of LOSS-MEASUREMENT-ATTRIBUTES TLV in LSPA object with Inferred Mode loss measurement request and Direct Mode loss measurement request, respectively, when it did not advertise this capability.
- o If the PCEP speaker that supports the extensions of this document but did not advertise this capability, then upon receipt of LOSS-MEASUREMENT object, it SHOULD generate a PCErr with error-type 19 (Invalid Operation), error-value TBD8 (Loss-Measurement capability was not advertised) and it will terminate the PCEP session.

#### 6.1.1. LOSS-MEASUREMENT-CAPABILITY TLV

The LOSS-MEASUREMENT-CAPABILITY TLV is an optional TLV for use in the OPEN Object for Loss Measurement via PCEP capability advertisement. Its format is shown in the following figure:



## LOSS-MEASUREMENT-CAPABILITY TLV Format

The Type of the TLV is TBD2 and it has a fixed length of 4 bytes.

The value comprises a single field - Flags (32 bits):

- o U (Unidirectional Measurement - 1 bit): if set to 1 by a PCC, the U Flag indicates that the PCC allows reporting of unidirectional loss measurement information; if set to 1 by a PCE, the U Flag indicates that the PCE is capable of receiving unidirectional loss measurement information from the PCC.
- o B (Bidirectional Measurement - 1 bit): if set to 1 by a PCC, the B Flag indicates that the PCC allows reporting of bidirectional loss

measurement information; if set to 1 by a PCE, the B Flag indicates that the PCE is capable of receiving bidirectional loss measurement information from the PCC. Bidirectional measurement is only applicable to the bidirectional LSPs (e.g. MPLS-TP LSPs [RFC5921]).

- o I (Inferred Loss Measurement Mode - 1 bit): if set to 1 by a PCC, the I Flag indicates that the PCC allows reporting of inferred mode loss measurement [RFC6374] information; if set to 1 by a PCE, the I Flag indicates that the PCE is capable of receiving inferred mode loss measurement information from the PCC.
- o N (Direct Loss Measurement Mode - 1 bit): if set to 1 by a PCC, the N Flag indicates that the PCC allows reporting of direct mode loss measurement [RFC6374] information; if set to 1 by a PCE, the N Flag indicates that the PCE is capable of receiving direct mode loss measurement information from the PCC.

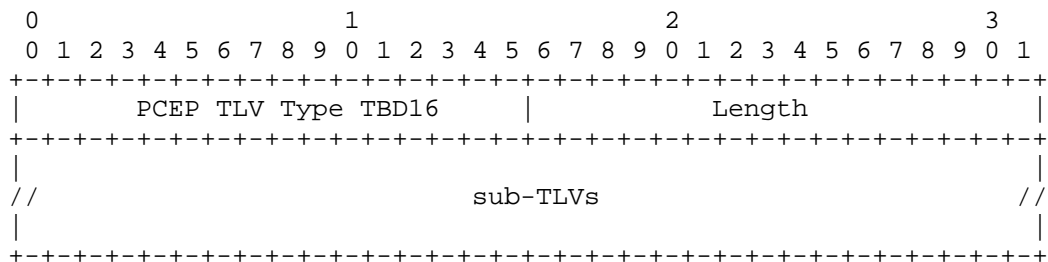
Unassigned bits are considered reserved. They MUST be set to 0 when sent and MUST be ignored when received.

Advertisement of the LOSS-MEASUREMENT-CAPABILITY TLV implies support of loss measurement, as well as the objects, TLVs and procedures defined in this document. Either U or B flag MUST be set in the TLV. Similarly, either I or N flag MUST be set in the TLV.

## 6.2. LOSS-MEASUREMENT-ATTRIBUTES TLV

The LOSS-MEASUREMENT-ATTRIBUTES TLV provides the configurable parameters of the loss measurement feature.

The format of the LOSS-MEASUREMENT-ATTRIBUTES TLV is shown in the following figure:



## LOSS-MEASUREMENT-ATTRIBUTES TLV Format

PCEP TLV Type is defined as following:

Type	Name
-----	
TBD16	LOSS-MEASUREMENT-ATTRIBUTES

Length: The Length field defines the length of the value portion in bytes as per [RFC5440].

Value: Comprises of one or more sub-TLVs as described in Section 4 of this document.

The following sub-sections describe the parameters which are currently defined to be carried within this TLV.

#### 6.2.1. Loss Measurement Enable

The Measurement-Enable sub-TLV specifies the loss measurement mode enabled using following flags:

Bit	Description
-----	
29	Unidirectional Loss Measurement Enabled
28	Bidirectional Loss Measurement Enabled
27	Inferred Loss Measurement Enabled

#### 6.2.2. Loss Measurement Interval

The Measurement-Interval sub-TLV specifies a time interval in seconds for the loss measurement.

#### 6.2.3. Loss Measurement Report Threshold

The Report-Threshold sub-TLV specifies the threshold value used to trigger an immediate reporting of the loss measurements bypassing the report-interval.

- o Report-Threshold: This 24-bit field identifying the packet loss as a percentage of the total packets sent or received. The encoding is as per [RFC7471].

Same report-threshold is used for all loss measurement values.

#### 6.2.4. Loss Measurement Report Threshold Percentage

The Report-Threshold-Percentage sub-TLV specifies the threshold value used to trigger an immediate reporting of the loss measurements bypassing the report-interval.

Same report-threshold-percentage is used for all loss measurement

values.

#### 6.2.5. Loss Measurement Report Interval

The Report-Interval sub-TLV specifies the time interval in seconds when measured loss values are to be reported.

#### 6.2.6. Loss Measurement Upper Bound

The Report-Upper-Bound sub-TLV specifies the upper-bound value in percentage packet loss, and is used to trigger an immediate reporting of the packet loss values when crossed. This may also result in PCC taking an immediate local action on the LSP.

### 6.3. LOSS-MEASUREMENT Object

The LOSS-MEASUREMENT Object with Object-Class (Value TBD6) is defined in this document to report the packet loss measurement of a TE LSP.

When the LSP is enabled with the loss measurement feature, the PCC SHOULD include the LOSS-MEASUREMENT Object to report the measured packet loss to the PCE in the PCRpt message.

```

      0                               1                               2                               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Class=TBD6      |   OT   |Res|P|I|   Object Length (bytes)   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                                           |
//                                                           //
|                                                           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

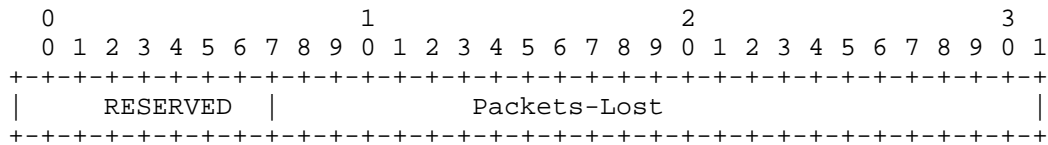
#### LOSS-MEASUREMENT Object Format

Object Length (16 bits): Specifies the total object length including the header, in bytes [RFC5440].

Object-Types (OT) are defined as following:

Object-Type	Len	Name
-----		
1	8	Tx Packets-Lost
2	8	Rx Packets-Lost

The object body format is defined as following:



#### LOSS-MEASUREMENT Object Body Formats (Tx and Rx)

- o Packets-Lost: This 24-bit field identifying the packet loss as a percentage of the total packets sent or received, encoded as per [RFC7471].
- o RESERVED: This field is reserved for future use. It MUST be set to 0 when sent and MUST be ignored when received.

The Packets-Lost in the Rx direction is reported when bidirectional loss measurement is enabled.

## 7. PCEP Extensions for Reporting Bandwidth Utilization

### 7.1. Bandwidth Utilization Capability Advertisement

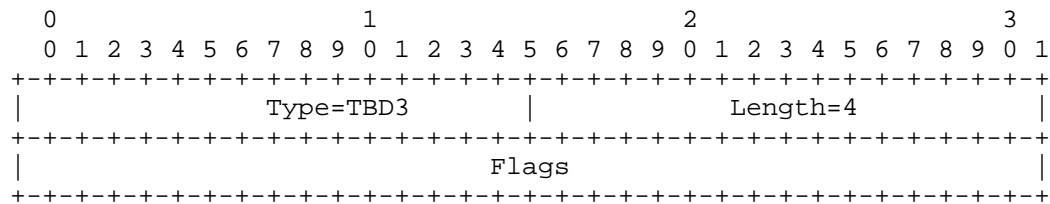
During PCEP Initialization Phase, PCEP Speakers (PCE or PCC) advertise their support of bandwidth utilization reporting. A PCEP Speaker includes the "BANDWIDTH-UTILIZATION-CAPABILITY" TLV, in the OPEN Object to advertise its support for PCEP extension. The presence of the "BANDWIDTH-UTILIZATION-CAPABILITY" TLV in the OPEN Object (in the Open message) indicates that the bandwidth utilization reporting is supported as described in this document. Additional procedure is defined as following:

- o The PCEP protocol extensions for bandwidth utilization MUST NOT be used if one or both PCEP Speakers have not included the "BANDWIDTH-UTILIZATION-CAPABILITY" TLV in their respective Open message.
- o If the PCEP speaker that supports the extensions of this document but did not advertise this capability, then upon receipt of BANDWIDTH-UTILIZATION-ATTRIBUTES TLV in LSPA object, it SHOULD generate a PCerr with error-type 19 (Invalid Operation), error-value TBD13 (Bandwidth utilization capability was not advertised) and it will terminate the PCEP session.
- o If the PCEP speaker that supports the extensions of this document but did not advertise this capability, then upon receipt of BANDWIDTH object of type TBD14, it SHOULD generate a PCerr with error-type 19 (Invalid Operation), error-value TBD13 (Bandwidth

utilization capability was not advertised) and it will terminate the PCEP session.

#### 7.1.1. BANDWIDTH-UTILIZATION-CAPABILITY TLV

The BANDWIDTH-UTILIZATION-CAPABILITY TLV is an optional TLV for use in the OPEN Object for Bandwidth Utilization reporting via PCEP capability advertisement. Its format is shown in the following figure:



BANDWIDTH-UTILIZATION-CAPABILITY TLV Format

The Type of the TLV is TBD3 and it has a fixed length of 4 bytes.

The value comprises a single field - Flags (32 bits). Currently no flags are defined for this TLV.

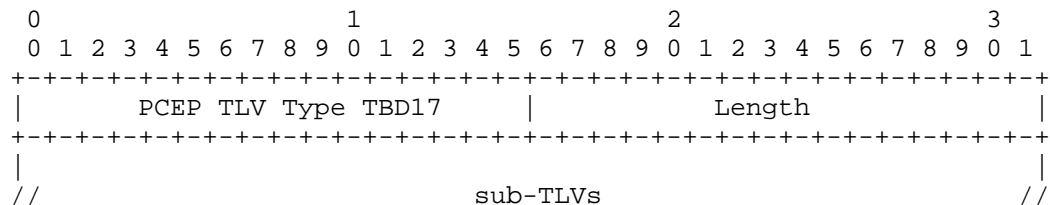
Unassigned bits are considered reserved. They MUST be set to 0 when sent and MUST be ignored when received.

Advertisement of the BANDWIDTH-UTILIZATION-CAPABILITY TLV implies support of bandwidth utilization reporting, as well as the objects, TLVs and procedures defined in this document.

#### 7.2. BW-UTILIZATION-MEASUREMENT-ATTRIBUTES TLV

The BW-UTILIZATION-MEASUREMENT-ATTRIBUTES TLV provides the configurable parameters of bandwidth utilization feature.

The format of the BW-UTILIZATION-MEASUREMENT-ATTRIBUTES TLV is shown in the following figure:



```
|
+-----+
```

#### BW-UTILIZATION-MEASUREMENT-ATTRIBUTES TLV Format

PCEP TLV Type is defined as following:

Type	Name
TBD17	BW-UTILIZATION-MEASUREMENT-ATTRIBUTES

Length: The Length field defines the length of the value portion in bytes as per [RFC5440].

Value: Comprises of one or more sub-TLVs as described in Section 4 of this document.

The following sub-sections describe the parameters which are currently defined to be carried within this TLV.

##### 7.2.1. Bandwidth Utilization Measurement Enable

The Measurement-Enable sub-TLV specifies that the bandwidth utilization reporting is enabled using following flag:

Bit	Description
26	Bandwidth Utilization Reporting Enabled

##### 7.2.2. Bandwidth Utilization Measurement Interval

The Measurement-Interval sub-TLV specifies a time interval in seconds for the bandwidth samples collection interval.

##### 7.2.3. Bandwidth Utilization Report Threshold

The Report-Threshold sub-TLV is used to decide if the bandwidth samples collected so far should be immediately reported bypassing the report-interval.

- o Threshold: The absolute threshold bandwidth value in 32-bits, encoded in IEEE floating point format (see [IEEE.754.1985]), expressed in bytes per second.

##### 7.2.4. Bandwidth Utilization Report Threshold Percentage

The Report-Threshold-Percentage sub-TLV is used to decide if the bandwidth samples collected so far should be immediately reported

bypassing the report-interval.

### 7.2.5. Bandwidth Utilization Report Interval

The Report-Interval sub-TLV specifies a time interval in seconds when the collected bandwidth samples are to be reported to PCE.

### 7.2.6. Bandwidth Utilization Upper Bound

The Report-Upper-Bound sub-TLV specifies the upper-bound bandwidth encoded in IEEE floating point format (see [IEEE.754.1985]), expressed in bytes per second, and is used to trigger an immediate reporting when crossed. This may also result in PCC taking an immediate local action on the LSP.

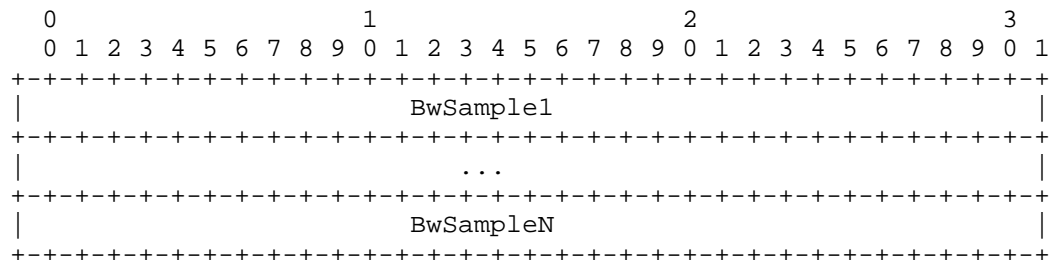
### 7.3. BANDWIDTH Object

A new object-type for BANDWIDTH object (Object-Class 5) is defined to report the bandwidth utilization of a TE LSP.

When the TE LSP is enabled with the bandwidth utilization reporting, the PCC SHOULD include the BANDWIDTH-UTILIZATION Object to report the bandwidth utilization of the TE LSP to the PCE in the PCRpt message.

The object-type is TBD14, the object length is variable with multiples of 4 bytes.

The object body format is defined as following:



BANDWIDTH-UTILIZATION Object Body Format

- o BwSample: The utilized bandwidth, (the BwSample collected at the end of each measurement-interval) encoded in IEEE floating point format (see [IEEE.754.1985]), expressed in bytes per second.

## 8. PCEP Procedure

Following procedure is defined for the extensions to different PCEP messages for reporting performance measurements.

#### 8.1. Various MEASUREMENT-ATTRIBUTES TLVs

- o For a PCE-Initiated LSP [DRAFT-PCE-INITIATED-LSP] with reporting features enabled, the corresponding MEASUREMENT-ATTRIBUTES TLV for each measurement MUST be included in the LSPA Object with the PCInitiate message.
- o For a PCE-Initiated LSP [DRAFT-PCE-INITIATED-LSP] with reporting features enabled, the corresponding MEASUREMENT-ATTRIBUTES TLV for each measurement is carried in the PCUpd message in the LSPA Object in order to make updates to the attributes such as Report-Interval.
- o For a PCC-Initiated LSP with reporting features enabled, when the LSP is delegated to the PCE, the corresponding MEASUREMENT-ATTRIBUTES TLV for each measurement MUST be included in the LSPA Object of the PCRpt message.
- o The various MEASUREMENT-ATTRIBUTES TLVs are encoded in all PCEP messages for the LSP with reporting features enabled, the absence of the corresponding MEASUREMENT-ATTRIBUTES TLV indicates that the PCEP speaker wishes to disable the feature.

#### 8.2. The MEASUREMENT Objects

When a TE LSP is enabled with a measurement reporting feature, the PCC SHOULD include the corresponding MEASUREMENT Object to report the measured values to the PCE in the PCRpt message [DRAFT-PCE-STATEFUL].

The format of the "actual\_attribute-list" in the PCRpt message is modified as following:

```
<actual_attribute-list>::=[<BANDWIDTH>]
                             [<DELAY-MEASUREMENT>]
                             [<LOSS-MEASUREMENT>]
                             [<metric-list>]
```

#### 9. Scaling Considerations

It should be noted that when measurement reporting is deployed under LSP scale, it can lead to frequent reporting updates to the PCE. Operators are advised to set the values of various measurement reporting parameters appropriate for the deployed LSP scale.

If a PCE gets overwhelmed, it can notify the PCC to temporarily suspend the reporting of the measurements as described below.

#### 9.1. The PCNtf Message

As per [RFC5440], the PCEP Notification message (PCNtf) can be sent by a PCEP speaker to notify its peer of a specific event. A PCEP speaker SHOULD notify its PCEP peer that it is overwhelmed, and on receipt of such notification the peer SHOULD NOT send any PCEP messages related to measurement reporting. If a PCEP message related to measurement reporting is received, it MUST be silently ignored.

- o When a PCEP speaker is overwhelmed, it SHOULD notify its peer by sending a PCNtf message with Notification Type = TBD15 (PM Overwhelm State) and Notification Value = 1 (Entering PM overwhelm state).
- o Optionally, OVERLOADED-DURATION TLV [RFC5440] MAY be included that specifies the time period during which no further PCEP messages related to PM should be sent.
- o When the PCEP speaker is no longer in the overwhelm state and is available to process the PM reporting, it SHOULD notify its peer by sending a PCNtf message with Notification Type = TBD15 (PM Overwhelm State) and Notification Value = 2 (Clearing PM overwhelm state).

#### 10. Security Considerations

This document defines new MEASUREMENT-ATTRIBUTES TLVs, CAPABILITY TLVs and MEASUREMENT Objects for reporting loss, delay measurements and bandwidth utilization which do not add additional security concerns beyond those discussed in [RFC5440] and [DRAFT-PCE-STATEFUL].

Some deployments may find the reporting of the performance measurement and bandwidth utilization information as extra sensitive as it could be used to influence LSP path computation and LSP setup with adverse effect. Additionally, snooping of PCEP messages with such data or using PCEP messages for network reconnaissance, may give an attacker sensitive information about the operations of the network. Thus, such deployment should employ suitable PCEP security mechanisms like TCP Authentication Option (TCP-AO) [RFC5925] or [DRAFT-PCE-PCEPS].

#### 11. Manageability Considerations

### 11.1. Control of Function and Policy

The performance measurement reporting SHOULD be controlled per TE tunnel (at PCC or PCE) and the values for feature attributes e.g. measurement-interval, report-interval, report-threshold SHOULD be configurable by an operator.

### 11.2. Information and Data Models

A Management Information Base (MIB) module for modeling PCEP is described in [RFC7420]. However, one may prefer the mechanism for configuration using YANG data model [DRAFT-PCE-PCEP-YANG]. These SHOULD be enhanced to provide controls and indicators for support of performance measurement reporting feature. Support for various configuration knobs as well as counters of messages sent/received containing the TLVs (defined in this document) SHOULD be added.

### 11.3. Liveness Detection and Monitoring

Mechanisms defined in this document do not imply any new liveness detection and monitoring requirements in addition to those already listed in [RFC5440].

### 11.4. Verify Correct Operations

Mechanisms defined in this document do not imply any new operation verification requirements in addition to those already listed in [RFC5440].

### 11.5. Requirements On Other Protocols

Mechanisms defined in this document do not add any new requirements on other protocols.

### 11.6. Impact On Network Operations

In order to avoid any unacceptable impact on network operations, an implementation SHOULD allow a limit to be placed on the number of LSPs that can be enabled with performance measurement reporting feature. An implementation MAY allow a limit to be placed on the rate of measurement reporting messages sent by a PCEP speaker and received by a peer. An implementation MAY also allow sending a notification when a PCEP speaker is overwhelmed or the rate of messages reach a threshold.

## 12. IANA Considerations

### 12.1. Measurement Capability TLV Types

This document defines the following new PCEP TLVs; IANA is requested to make the following allocations from the "PCEP TLV Type Indicators" registry. <<http://www.iana.org/assignments/pcep/pcep.xhtml#pcep-tlv-type-indicators>>

Type	Name	Reference
TBD1	DELAY-MEASUREMENT-CAPABILITY	[This document]
TBD2	LOSS-MEASUREMENT-CAPABILITY	[This document]
TBD3	BANDWIDTH-UTILIZATION-CAPABILITY	[This document]

#### 12.1.1. Flag Fields for MEASUREMENT-CAPABILITY TLVs

IANA is requested to create a registry to manage the Flag field of the DELAY-MEASUREMENT-CAPABILITY TLV, LOSS-MEASUREMENT-CAPABILITY TLV and BANDWIDTH-UTILIZATION-CAPABILITY TLV.

New bit numbers are allocated only by an IETF Review action [RFC5226]. Each bit should be tracked with the following qualities:

- o Bit number (counting from bit 0 as the most significant bit)
- o Capability description
- o Defining RFC

The following value are defined in this document for the Flag field for -

#### DELAY-MEASUREMENT-CAPABILITY TLV:

Bit	Description	Reference
31	One-way Delay Measurement	[This document]
30	Two-way Delay Measurement	[This document]

#### LOSS-MEASUREMENT-CAPABILITY TLV:

Bit	Description	Reference
31	Unidirectional Loss Measurement	[This document]
30	Bidirectional Loss Measurement	[This document]
29	Inferred Loss Measurement Mode	[This document]
28	Direct Loss Measurement Mode	[This document]

## 12.2. MEASUREMENT-ATTRIBUTES TLVs

This document defines the following new PCEP TLV Types; IANA is requested to make the following TLV type allocations from the "PCEP TLV Type Indicators" registry.

<<http://www.iana.org/assignments/pcep/pcep.xhtml#pcep-tlv-type-indicators>>

Type	Name	Reference
TBD4	DELAY-MEASUREMENT-ATTRIBUTES	[This document]
TBD16	LOSS-MEASUREMENT-ATTRIBUTES	[This document]
TBD17	BW-UTILIZATION-MEASUREMENT-ATTRIBUTES	[This document]

### 12.2.1. The Sub-TLVs For MEASUREMENT-ATTRIBUTES TLVs

IANA is requested to create an "MEASUREMENT-ATTRIBUTES Sub-TLV Types" sub-registry in the "PCEP TLV Type Indicators" registry. New sub-TLV are allocated only by an IETF Review action [RFC5226].

This document defines the following sub-TLV types:

Type	Name	Reference
0	Reserved	[This document]
1	Measurement-Enable sub-TLV	[This document]
2	Measurement-Interval sub-TLV	[This document]
3	Report-Threshold sub-TLV	[This document]
4	Report-Threshold-Percentage sub-TLV	[This document]
5	Report-Interval sub-TLV	[This document]
6	Report-Upper-Bound sub-TLV	[This document]
7-65535	Unassigned	[This document]

#### 12.2.1.1. Flag Fields in Measurement-Enable sub-TLV

IANA is requested to create a registry to manage the Flag field of the Measurement-Enable sub-TLV.

New bit numbers are allocated only by an IETF Review action [RFC5226]. Each bit should be tracked with the following qualities:

- o Bit number (counting from bit 0 as the most significant bit)
- o Capability description
- o Defining RFC

The following value are defined in this document for the Flag field.

Bit	Description	Reference
31	One-Way Delay Measurement Enabled	[This document]
30	Two-Way Delay Measurement Enabled	[This document]
29	Unidirectional Loss Measurement Enabled	[This document]
28	Bidirectional Loss Measurement Enabled	[This document]
27	Inferred Loss Measurement Enabled	[This document]
26	Bandwidth Utilization Reporting Enabled	[This document]

### 12.3. Measurement Object-Class

This document defines Object-Class for the following Objects; IANA is requested to make the following allocations from the "PCEP Objects" registry. <<http://www.iana.org/assignments/pcep/pcep.xhtml#pcep-objects>>

Object-Class	Name	Reference
TBD5	DELAY-MEASUREMENT Object	[This document]
TBD6	LOSS-MEASUREMENT Object	[This document]

#### 12.3.1. DELAY-MEASUREMENT Object-Types

IANA is requested to create an "DELAY-MEASUREMENT Object-Types" sub-registry for DELAY-MEASUREMENT Object (Object-class TBD5).

This document defines the following object-types:

Object-Type	Name	Reference
0	Reserved	[This document]
1	One-Way Delay Measurement Value	[This document]
2	One-Way Delay Measurement Min/Max Values	[This document]
3	One-Way Delay Variation Measurement Value	[This document]
4	Two-Way Delay Measurement Value	[This document]
5	Two-Way Delay Measurement Min/Max Values	[This document]
6	Two-Way Delay Variation Measurement Value	[This document]

#### 12.3.2. LOSS-MEASUREMENT Object-Types

IANA is requested to create an "LOSS-MEASUREMENT Object-Types" sub-registry for LOSS-MEASUREMENT Object (Object-class TBD6).

This document defines the following object-types:

Object-Type Name		Reference
-----		
0	Reserved	[This document]
1	Tx Packets-Lost	[This document]
2	Rx Packets-Lost	[This document]

### 12.3.3. BANDWIDTH Object-Type

This document defines new Object-Type for the BANDWIDTH object (Object-Class 5, [RFC5440]); IANA is requested to make the following allocation from the "PCEP Objects" registry.  
[<http://www.iana.org/assignments/pcep/pcep.xhtml#pcep-objects>](http://www.iana.org/assignments/pcep/pcep.xhtml#pcep-objects)

Object-Type Name		Reference
-----		
TBD14	BANDWIDTH-UTILIZATION Object	[This document]

### 12.4. PCE Error Codes

This document defines two new error-values for PCErr with error-code 19 (Invalid Operation). IANA is requested to make the following allocations.

Error-Value	Name	Reference
-----		
TBD7	Delay-Measurement capability	
	was not advertised	[This document]
TBD8	Loss-Measurement capability	
	was not advertised	[This document]
TBD9	Bidirectional Measurement capability	
	was not advertised	[This document]
TBD10	Unidirectional Measurement capability	
	was not advertised	[This document]
TBD11	Inferred Mode Loss Measurement capability	
	was not advertised	[This document]
TBD12	Direct Mode Loss Measurement capability	
	was not advertised	[This document]
TBD13	Bandwidth Utilization capability	
	was not advertised	[This document]

### 12.5. Notification Object-Type

IANA is requested to allocate new Notification Types and Notification Values within the "Notification Object" sub-registry of the PCEP Numbers registry, as follows:

Type	Meaning	Reference
-----		

TBD15

PM Overwhelm State

[This document]

Notification-value=1: Entering PM overwhelm state

Notification-value=2: Clearing PM overwhelm state

## 13. References

### 13.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009.
- [DRAFT-PCE-STATEFUL] Crabbe, E., Minei, I., Medved, J., and R. Varga, "PCEP Extensions for Stateful PCE", draft-ietf-pce-stateful-pce, (work in progress).
- [DRAFT-PCE-INITIATED-LSP] Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "PCEP Extensions for PCE-initiated LSP Setup in a Stateful PCE Model", draft-ietf-pce-pce-initiated-lsp, (work in progress).

### 13.2. Informative References

- [RFC5586] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", RFC 5586, June 2009.
- [RFC5921] Bocci, M., Ed., Bryant, S., Ed., Frost, D., Ed., Levrau, L., and L. Berger, "A Framework for MPLS in Transport Networks", RFC 5921, July 2010.
- [RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", RFC 5925, June 2010.
- [RFC6374] Frost, D. and S. Bryant, "Packet Loss and Delay Measurement for MPLS Networks", DOI 10.17487/RFC6374, RFC 6374, September 2011.
- [RFC6375] Frost, D. and S. Bryant, "A Packet Loss and Delay Measurement Profile for MPLS-Based Transport Networks", RFC 6375, September 2011.
- [RFC7420] Koushik, A., Stephan, E., Zhao, Q., King, D., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Management Information Base (MIB) Module", RFC 7420, December 2014.

- [RFC7471] S. Giacalone, D. Ward, J. Drake, A. Atlas, and S. Previdi, "OSPF Traffic Engineering (TE) Metric Extensions", RFC 7471, March 2015.
- [RFC7810] S. Previdi, S. Giacalone, D. Ward, J. Drake, and Q. Wu, "IS-IS Traffic Engineering (TE) Metric Extensions", RFC 7810, May 2016.
- [RFC7823] Atlas, A., Drake, J., Giacalone, S., and S. Previdi, "Performance-Based Path Selection for Explicitly Routed Label Switched Paths (LSPs) Using TE Metric Extensions", RFC 7823, May 2016.
- [RFC7876] Bryant, S., Sivabalan, S., and Soni, S., "UDP Return Path for Packet Loss and Delay Measurement for MPLS Networks", RFC 7876, July 2016.
- [DRAFT-PCE-PCEPS] Lopez, D., Dios, O., Wu, W., and D. Dhody, "Secure Transport for PCEP", draft-ietf-pce-pceps, (work in progress).
- [DRAFT-PCE-SERVICE-AWARE] Dhody, D., V. Manral, V., Ali, Z., and Kumaki, K., "Extensions to the Path Computation Element Communication Protocol (PCEP) to compute service aware Label Switched Path (LSP)", draft-ietf-pce-pcep-service-aware, (work in progress).
- [DRAFT-IDR-TE-PM-BGP] Wu, Q., Danhua, W., Previdi, S., Gredler, H., and S. Ray, "BGP attribute for North-Bound Distribution of Traffic Engineering (TE) performance Metrics", draft-ietf-idr-te-pm-bgp (work in progress).
- [DRAFT-PCE-PCEP-YANG] Dhody, D., Hardwick, J., Beeram, V., and J. Tantsura, "A YANG Data Model for Path Computation Element Communications Protocol (PCEP)", draft-ietf-pce-pcep-yang (work in progress).
- [DRAFT-IETF-PCE-AUTOBW] Dhody, D., Palle, U., Singh, R., Gandhi, R., and L. Fang, "PCEP Extensions for MPLS-TE LSP Automatic Bandwidth Adjustment with Stateful PCE", draft-ietf-pce-stateful-pce-auto-bandwidth (work in progress).
- [IEEE.754.1985] Institute of Electrical and Electronics Engineers, "Standard for Binary Floating-Point Arithmetic", IEEE Standard 754, August 1985.

Acknowledgments

TBA.

Authors' Addresses

Rakesh Gandhi  
Cisco Systems, Inc.

EMail: rgandhi@cisco.com

Bin Wen  
Comcast

EMail: Bin\_Wen@cable.comcast.com

Colby Barth  
Juniper Networks

EMail: cbarth@juniper.net

Dhruv Dhody  
Huawei Technologies  
India

EMail: dhruv.ietf@gmail.com