

PCE Working Group
Internet-Draft
Intended Status: Standards Track
Expires: July 25, 2016

R. Gandhi
Cisco Systems
B. Wen
Comcast
January 22, 2016

**PCEP Extensions for Reporting MPLS-TE LSP Performance Measurements
draft-gandhi-pce-pm-00**

Abstract

In certain networks, network performance data such as packet loss, delay and delay variation (jitter) is a critical measure for traffic engineering. This data provides operators the characteristics of their networks for performance evaluation that is required to ensure the Service Level Agreements (SLAs). Performance measurement mechanisms can be employed to measure these metrics for TE LSPs in real-time. This document describes PCEP extensions for reporting such performance measurements to an Active Stateful PCE.

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) 2016 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	3
2.	Conventions Used in This Document	3
2.1.	Requirements Language	3
2.2.	Terminology	3
3.	PCEP Extensions for Reporting Delay Measurement	4
3.1.	DELAY-MEASUREMENT-ATTRIBUTE TLV	4
3.1.1.	Measurement-Enable sub-TLV	5
3.1.2.	Measurement-Interval sub-TLV	6
3.2.	DELAY-MEASUREMENT Object	6
3.3.	DELAY-MEASUREMENT Capability Advertisement	7
4.	PCEP Extensions for Reporting Loss Measurement	7
4.1.	LOSS-MEASUREMENT-ATTRIBUTE TLV	7
4.1.1.	Measurement-Enable sub-TLV	8
4.1.2.	Measurement-Interval sub-TLV	9
4.2.	LOSS-MEASUREMENT Object	9
4.3.	LOSS-MEASUREMENT Capability Advertisement	10
5.	Security Considerations	10
6.	IANA Considerations	10
6.1.	PCEP TLV Type Indicators	11
6.1.1.	DELAY-MEASUREMENT-ATTRIBUTE Sub-TLVs	11
6.1.2.	LOSS-MEASUREMENT-ATTRIBUTE Sub-TLVs	11
6.2.	DELAY-MEASUREMENT Object	11
6.2.1.	DELAY-MEASUREMENT Object-Types	12
6.3.	LOSS-MEASUREMENT Object	12
6.3.1.	LOSS-MEASUREMENT Object-Types	12
7.	References	13
7.1.	Normative References	13
7.2.	Informative References	13
	Acknowledgments	15
	Authors' Addresses	15

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 for PCEP to enable stateful control of MPLS TE LSPs. It describes two mode of operations - Passive Stateful PCE and Active Stateful PCE. In this document, Active Stateful PCE is considered.

In certain networks, such as financial information networks, network performance data (e.g. packet loss, delay and delay variation (jitter)) is a critical measure for traffic engineering [[RFC7471](#)], [[DRAFT-OSPF-TE-METRIC](#)] and [[DRAFT-ISIS-TE-METRIC](#)]. 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. However, there is a need to verify that the traffic sent over the TE LSP does not exceed requested metric bounds. [[RFC6374](#)], [[RFC6375](#)] and [[DRAFT-RFC6374-UDP-RETURN-PATH](#)] define protocol extensions needed for measuring packet loss, delay and delay variation (jitter) for bidirectional and unidirectional TE LSPs in real-time.

This document provides mechanisms to report the packet loss, delay and delay variation (jitter) measurements of a TE LSP to an Active Stateful PCE. Note that specification of the use of the reported packet loss, delay and delay variation measurement is outside the scope of this document.

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 following terminology is used in this document.

Active Stateful PCE: PCE that uses tunnel state information learned from PCCs to optimize path computations. Additionally, it actively updates tunnel parameters in those PCCs that delegated control over their tunnels to the PCE.

PCC: Path Computation Client. Any client application requesting a path computation to be performed by a Path Computation Element.

PCE: Path Computation Element. An entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints.

TE LSP: Traffic Engineering Label Switched Path.


3. PCEP Extensions for Reporting Delay Measurement

3.1. DELAY-MEASUREMENT-ATTRIBUTE TLV

The DELAY-MEASUREMENT-ATTRIBUTE TLV provides the configurable parameters of the delay measurement feature. This is an optional TLV defined for the LSPA Object.

For PCE-Initiated LSP ([[DRAFT-PCE-INITIATED-LSP](#)]) with delay measurement feature enabled, this TLV is included in the LSPA Object with PCInitiate message. The DELAY-MEASUREMENT-ATTRIBUTE TLV can also be carried in PCUpd message in LSPA Object in order to make updates to delay measurement attributes such as Measurement-Interval.

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



DELAY-MEASUREMENT-ATTRIBUTE TLV format

Type: TBA

Length: Variable

Value: Comprises one or more sub-TLVs.

Following sub-TLVs are defined in this document:

Type	Len	Name
0	4	Reserved
1	4	Measurement-Enable sub-TLV
2	4	Measurement-Interval sub-TLV

The Measurement-Enable sub-TLV MUST be added in LSPA Object when the delay 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.

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

[3.1.1. Measurement-Enable sub-TLV](#)

The Measurement-Enable sub-TLV specifies the delay measurement mode enabled.

The Type is 1, Length is 4, and the value comprises of 4-octet. Value is defined as following:

Value	Name
0	Delay Measurement Disabled
1	One-way Delay Measurement Enabled
2	Two-way Delay Measurement Enabled
3	One-Way and Two-Way Delay Measurements Enabled

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Type=1           |           Length=4           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Measurement-Enable           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Measurement-Enable sub-TLV format

3.1.2. Measurement-Interval sub-TLV

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

The Type is 2, Length is 4, and the value comprises of 4-octet time interval, the valid range is from 1 to 604800, in seconds. The default value is 300 seconds.

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Type=2                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Measurement-Interval                  |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Measurement-Interval sub-TLV format

3.2. DELAY-MEASUREMENT Object

The DELAY-MEASUREMENT Object with Object-Class (Value TBA) 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 value to the PCE in the PCRpt message.

The object length is 4 bytes. Object-Types are defined as follows:

Object-Type	Name

0	Reserved
1	One-Way Delay Value
2	One-Way Delay Variation Value
3	Two-Way Delay Value
4	Two-Way Delay Variation Value

The payload format is as follows:

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Delay Value                          |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```


DELAY-MEASUREMENT Object format

- o One-way Delay Value: Delay of the LSP in one (forward) direction, encoded as 24-bit integer, as defined in [RFC 7471](#) [[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 Variation Value: Delay Variation of the LSP in one (forward) direction, encoded as 24-bit integer, as defined in [RFC 7471](#) [[RFC7471](#)]. When set to the maximum value 16,777,215 (16.777215 sec), the delay variation is at least that value and may be larger.
- o Two-way Delay Value: Delay of the bidirectional LSP in both (forward and reverse) directions, encoded as 24-bit integer, as defined in [RFC 7471](#) [[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 Two-way Delay Variation Value: Delay Variation of the bidirectional LSP in both (forward and reverse) directions, encoded as 24-bit integer, as defined in [RFC 7471](#) [[RFC7471](#)]. When set to the maximum value 16,777,215 (16.777215 sec), the delay variation is at least that value and may be larger.

[3.3.](#) DELAY-MEASUREMENT Capability Advertisement

To be added.


[4.](#) PCEP Extensions for Reporting Loss Measurement

[4.1.](#) LOSS-MEASUREMENT-ATTRIBUTE TLV

The LOSS-MEASUREMENT-ATTRIBUTE TLV provides the configurable parameters of the loss measurement feature. This is an optional TLV defined for the LSPA Object.

For PCE-Initiated LSP ([\[DRAFT-PCE-INITIATED-LSP\]](#)) with loss measurement feature enabled, this TLV is included in the LSPA Object with PCInitiate message. The LOSS-MEASUREMENT-ATTRIBUTE TLV can also be carried in PCUpd message in LSPA Object in order to make updates to loss measurement attributes such as Measurement-Interval.

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



LOSS-MEASUREMENT-ATTRIBUTE TLV format

Type: TBA

Length: Variable

Value: Comprises one or more sub-TLVs.

Following sub-TLVs are defined in this document:

Type	Len	Name
0	4	Reserved
1	4	Measurement-Enable sub-TLV
2	4	Measurement-Interval sub-TLV

The Measurement-Enable sub-TLV MUST be added in the LSPA Object when the loss 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.

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

4.1.1.1. Measurement-Enable sub-TLV

The Measurement-Enable sub-TLV specifies the loss measurement mode enabled.

The Type is 1, Length is 4, and the value comprises of 4-octet. Value is defined as following:

Value	Name
0	Loss Measurement Disabled
1	Loss Measurement Enabled

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Type=1           |           Length=4           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Measurement-Enable           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Measurement-Enable sub-TLV format

4.1.2. Measurement-Interval sub-TLV

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

The Type is 2, Length is 4, and the value comprises of 4-octet time interval, the valid range is from 1 to 604800, in seconds. The default value is 300 seconds.

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Type=2           |           Length=4           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Measurement-Interval           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Measurement-Interval sub-TLV format

4.2. LOSS-MEASUREMENT Object

The LOSS-MEASUREMENT Object with Object-Class (Value TBA) 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.

The object length is 4 bytes. Object-Types are defined as follows:

Object-Type	Name
0	Reserved
1	Packets-Lost
2	Bytes-Lost

The payload format is as follows:

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Packets-Lost                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Bytes-Lost                                |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

LOSS-MEASUREMENT Object formats

- o Packets-Lost: Number of packets sent over the LSP were lost.
- o Bytes-Lost: Number of Bytes sent over the LSP were lost.

4.3. LOSS-MEASUREMENT Capability Advertisement

To be added.

5. Security Considerations

This document defines new MEASUREMENT objects and MEASUREMENT-ATTRIBUTE TLVs for loss and delay 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 information as extra sensitive and thus should employ suitable PCEP security mechanisms like TCP-AO or [\[DRAFT-PCE-PCEPS\]](#).

6. IANA Considerations

6.1. PCEP TLV Type Indicators

This document defines the following new PCEP TLVs; IANA is requested to make the following allocations from this registry.

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

Value	Name	Reference
TBA	DELAY-MEASUREMENT-ATTRIBUTE	[This I.D.]
TBA	LOSS-MEASUREMENT-ATTRIBUTE	[This I.D.]

6.1.1. DELAY-MEASUREMENT-ATTRIBUTE Sub-TLVs

This document specifies the DELAY-MEASUREMENT-ATTRIBUTE Sub-TLVs. IANA is requested to create an "DELAY-MEASUREMENT-ATTRIBUTE Sub-TLV Types" sub-registry in the "PCEP TLV Type Indicators" for the sub-TLVs carried in the DELAY-MEASUREMENT-ATTRIBUTE TLV. This document defines the following types:

Type	Name	Reference
0	Reserved	[This I.D.]
1	Measurement-Enable sub-TLV	[This I.D.]
2	Measurement-Interval sub-TLV	[This I.D.]
3- 65535	Unassigned	[This I.D.]

6.1.2. LOSS-MEASUREMENT-ATTRIBUTE Sub-TLVs

This document specifies the LOSS-MEASUREMENT-ATTRIBUTE Sub-TLVs. IANA is requested to create an "LOSS-MEASUREMENT-ATTRIBUTE Sub-TLV Types" sub-registry in the "PCEP TLV Type Indicators" for the sub-TLVs carried in the LOSS-MEASUREMENT-ATTRIBUTE TLV. This document defines the following types:

Type	Name	Reference
0	Reserved	[This I.D.]
1	Measurement-Enable sub-TLV	[This I.D.]
2	Measurement-Interval sub-TLV	[This I.D.]
3- 65535	Unassigned	[This I.D.]

6.2. DELAY-MEASUREMENT Object

This document defines Object-Class for the DELAY-MEASUREMENT Object;

IANA is requested to make the following allocations from this registry. (see <http://www.iana.org/assignments/pcep/pcep.xhtml#pcep-objects>).

Object-Class	Name	Reference

TBA	DELAY-MEASUREMENT Object	[This I.D.]

6.2.1. DELAY-MEASUREMENT Object-Types

This document specifies the DELAY-MEASUREMENT Object-Types. IANA is requested to create an "DELAY-MEASUREMENT Object-Types" %sub-registry for DELAY-MEASUREMENT Object. This document defines the following types:

Object-Type	Name	Reference

0	Reserved	[This I.D.]
1	One-Way Delay Value	[This I.D.]
2	One-Way Delay Variation Value	[This I.D.]
3	Two-Way Delay Value	[This I.D.]
4	Two-Way Delay Variation Value	[This I.D.]

6.3. LOSS-MEASUREMENT Object

This document defines Object-Class for the LOSS-MEASUREMENT Object; IANA is requested to make the following allocations from this registry. (see <http://www.iana.org/assignments/pcep/pcep.xhtml#pcep-objects>).

Object-Class	Name	Reference

TBA	LOSS-MEASUREMENT Object	[This I.D.]

6.3.1. LOSS-MEASUREMENT Object-Types

This document specifies the LOSS-MEASUREMENT Object-Types. IANA is requested to create an "LOSS-MEASUREMENT Object-Types" sub-registry for LOSS-MEASUREMENT Object. This document defines the following types:

Object-Type	Name	Reference

0	Reserved	[This I.D.]
1	Packets-Lost	[This I.D.]
2	Bytes-Lost	[This I.D.]

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [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).

7.2. Informative References

- [RFC3471] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", [RFC 3471](#), January 2003.
- [RFC6373] Andersson, L., Berger, L., Fang, L., Bitar, N., and E. Gray, "MPLS Transport Profile (MPLS-TP) Control Plane Framework", [RFC 6373](#), September 2011.
- [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.
- [RFC7551] Zhang, F., Ed., Jing, R., and Gandhi, R., Ed., "RSVP-TE Extensions for Associated Bidirectional LSPs", [RFC 7551](#), May 2015.
- [DRAFT-PCE-PCEPS] Lopez, D., Dios, O., Wu, W., and D. Dhody, "Secure Transport for PCEP", [draft-ietf-pce-pceps](#), (work in progress).
- [DRAFT-OSPF-TE-METRIC] S. Giacalone, D. Ward, J. Drake, A. Atlas, S.

Previdi, "OSPF Traffic Engineering (TE) Metric Extensions", [draft-ietf-ospf-te-metric-extensions](#), (work in progress).

[DRAFT-ISIS-TE-METRIC] S. Previdi, S. Giacalone, D. Ward, J. Drake, A. Atlas, C. Filsfils, "IS-IS Traffic Engineering (TE) Metric Extensions", [draft-previdi-isis-te-metric-extensions](#), (work in progress).

[DRAFT-EXPRESS-PATH] Atlas, A., Drake, J., Giacalone, S., and S. Previdi, "Performance-based Path Selection for Explicitly Routed LSPs using TE Metric Extensions", [draft-ietf-teas-te-express-path](#), (work in progress).

[DRAFT-RFC6374-UDP-RETURN-PATH] Bryant, S., Sivabalan, S., Soni, S., "RFC6374 UDP Return Path", [draft-ietf-mpls-rfc6374-udp-return-path](#), (work in progress).

Acknowledgments

TBA.

Authors' Addresses

Rakesh Gandhi
Cisco Systems, Inc.

EMail: rgandhi@cisco.com

Bin Wen
Comcast

EMail: Bin_Wen@cable.comcast.com