PCE Working Group
Internet-Draft
Intended Status: Standards Track
Expires: April 30, 2017

R. Gandhi, Ed.
Individual Contributor
B. Wen
Comcast
C. Barth
Juniper Networks
October 27, 2016

PCEP Extensions for Reporting MPLS-TE LSP Performance Measurements draft-gandhi-pce-pm-03

Abstract

In certain networks, network performance data such as packet loss, delay and delay variation (jitter) 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 for TE Label Switched Paths (LSPs) in real-time. This document describes Path Computation Element Protocol (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/1id-abstracts.html

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

Table of Contents

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], [RFC7810] 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. However, there is a need to verify that the traffic sent over the TE LSP does not exceed requested metric bounds. [RFC6374], [RFC6375] and [RFC7876] define protocol extensions needed for measuring packet loss, delay and delay variation (jitter) for bidirectional and unidirectional TE LSPs in real-time.

1.1. Dependencies and Considerations

This document provides mechanisms to report the performance measurements (PM) such as packet loss, delay and delay variation (jitter) for a TE LSP to an Active Stateful PCE. [RFC6374] describes several reasons why PMs are valuable to operators. Note that the specification of the use of the reported packet loss, delay and delay variation measurement 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, PMs 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. MPLS-TP LSPs require Ultimate Hop Popping (UHP). 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.

1.2. Overview

High-level overview of the PCEP extensions defined in this document for reporting performance measurements is shown in Figure 1.

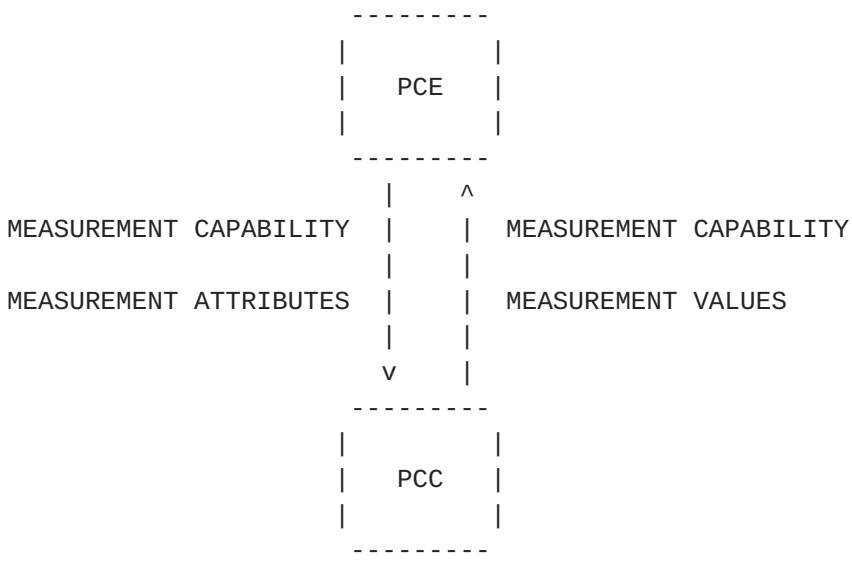

Figure 1: Overview of PCEP Extensions for Performance Measurement

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 Capability Advertisement

During PCEP Initialization Phase, PCEP Speakers (PCE or PCC) advertise their support of DELAY-MEASUREMENT. A PCEP Speaker (PCE or PCC) includes the "DELAY-MEASUREMENT Capability" TLV, in the OPEN Object to advertise its support for PCEP Delay-Measurement extensions. The presence of the "Delay-Measurement Capability" TLV in the OPEN Object indicates that the Delay Measurement capability is supported as described in this document.

The PCEP protocol extensions for Delay Measurement MUST NOT be used if one or both PCEP Speakers have not included the "Delay-Measurement Capability" TLV in their respective OPEN message. If the PCEP speaker that supports the extensions of this draft but did not advertise this capability, then upon receipt of DELAY-MEASUREMENT-ATTRIBUTE TLV in LSPA object, it SHOULD generate a PCErr with error-type 19 (Invalid Operation), error-value TBD7 (Delay-Measurement capability was not advertised) and it will terminate the PCEP session. 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 DELAY-MEASUREMENT-ATTRIBUTE TLV in LSPA object with two-way measurement request and one-way measurement request, respectively, when it did not advertise this capability. Further, the PCEP speaker SHOULD generate error-value TBD11 (Inferred Mode Measurement capability was not advertised) and TBD12 (Direct Mode Measurement capability was not advertised) upon receipt of DELAY-MEASUREMENT-ATTRIBUTE TLV in LSPA object with Inferred Mode delay measurement request and Direct Mode delay measurement request, respectively, when it did not advertise this capability.

3.1.1. DELAY-MEASUREMENT-CAPABILITY TLV

The DELAY-MEASUREMENT-CAPABILITY TLV is an optional TLV for use in the OPEN Object for Delay Measurement via PCEP capability

advertisement. Its format is shown in the following figure:

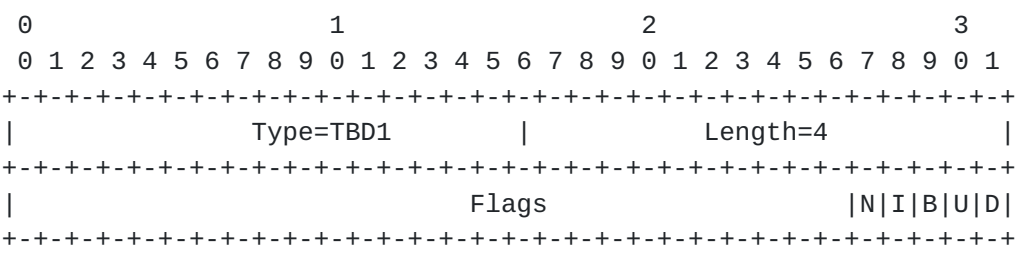

DELAY-MEASUREMENT-CAPABILITY TLV format

The type of the TLV is TBD1 and it has a fixed length of 4 octets.

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

- D (Delay Measurement - 1 bit): if set to 1 by a PCC, the D Flag indicates that the PCC allows reporting of delay measurement information; if set to 1 by a PCE, the D Flag indicates that the PCE is capable of receiving delay measurement information from the PCC. The DELAY-MEASUREMENTE-ATTRIBUTE TLV MUST be encoded when both PCEP speakers have the D Flag set.

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

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

- I (Inferred Mode - 1 bit): if set to 1 by a PCC, the I Flag indicates that the PCC allows reporting of inferred mode delay measurement information; if set to 1 by a PCE, the I Flag indicates that the PCE is capable of receiving inferred mode delay measurement information from the PCC.

- N (Direct Mode - 1 bit): if set to 1 by a PCC, the D Flag indicates that the PCC allows reporting of direct mode delay measurement information; if set to 1 by a PCE, the D Flag indicates that the PCE is capable of receiving direct mode delay measurement information from the PCC.

Unassigned bits are considered reserved. They MUST be set to 0 on transmission and MUST be ignored on receipt.

Advertisement of the Delay-Measurement capability TLV implies support of delay measurement, as well as the objects, TLVs and procedures defined in this document.

3.2. 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:

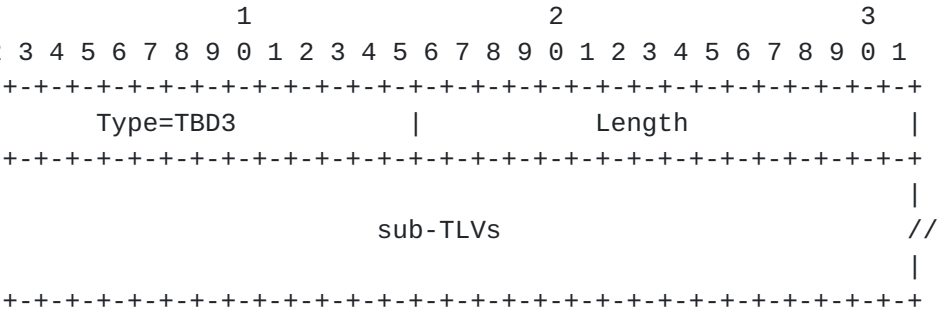

DELAY-MEASUREMENT-ATTRIBUTE TLV format

Type: TBD3

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
3	4	Measurement-Mode sub-TLV
4	4	Report-Threshold sub-TLV
5	4	Report-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.

3.2.1. 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

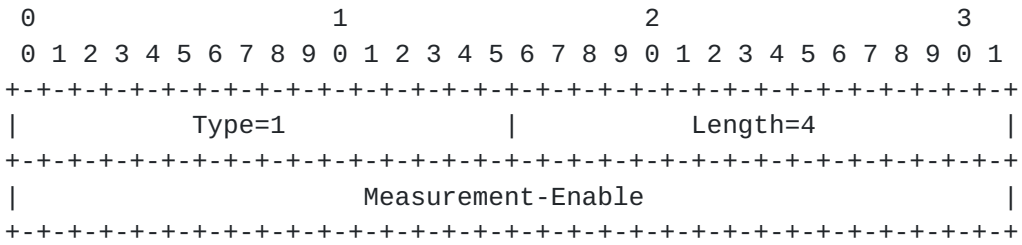

Measurement-Enable sub-TLV format

3.2.2. 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.

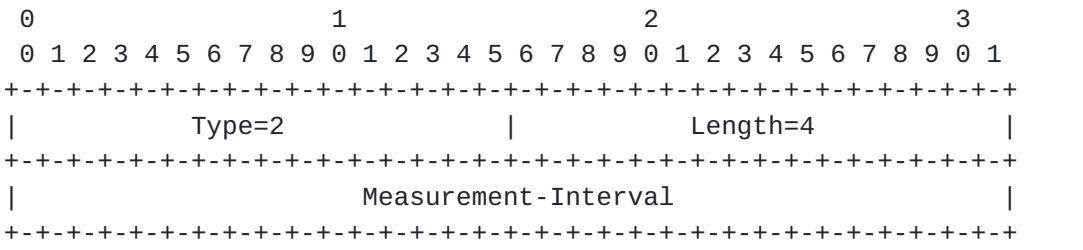

Measurement-Interval sub-TLV format

3.2.3. Measurement-Mode sub-TLV

The Measurement-Mode sub-TLV specifies the delay measurement mode which can be direct or inferred.

The Type is 3, Length is 4, and the value comprises of 4-octet delay measurement mode.

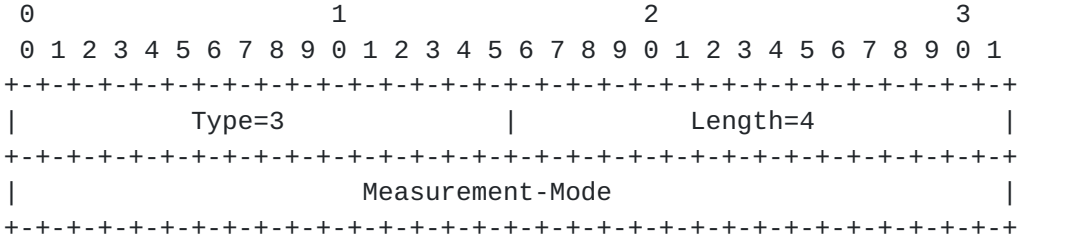

Measurement-Mode sub-TLV format

Value	Name
0	Inferred Mode (using test traffic)
1	Direct Mode (using data traffic)

3.2.4. Report-Threshold sub-TLV

The Report-Threshold sub-TLV specifies the delay threshold value.

The Type is 4, Length is 4, and the value comprises of 4-octet delay threshold value.

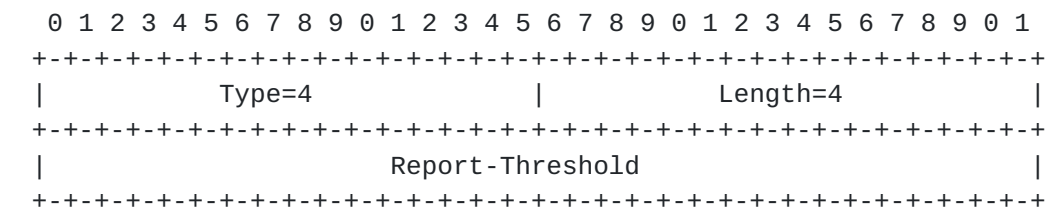

Report-Threshold sub-TLV format

The measured delay value MUST only be reported once the value exceeds the specified threshold value. This reporting can be used to trigger an immediate action at the PCE peer.

3.2.5. Report-Interval sub-TLV

The Report-Interval sub-TLV specifies the time interval in seconds when measured delay values are to be reported. Maximum value of the measured metric (e.g. average delay, delay variation) collected during the report interval is reported.

The Type is 5, 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 equals to the Measurement-Interval.

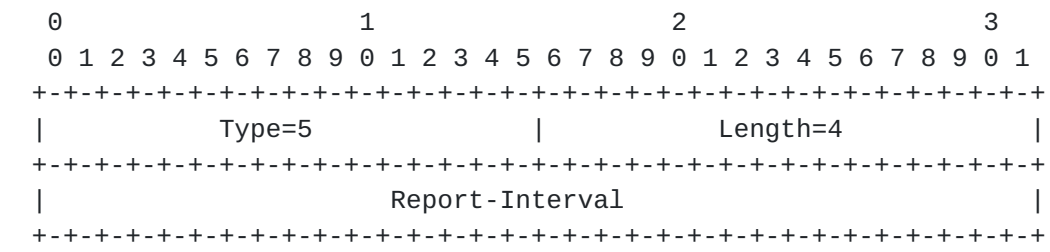

Report-Interval sub-TLV format

The measured delay value MUST be reported if the specified threshold is crossed OR the report-interval expires. The periodic reporting can be used at the PCEP peer to monitor the LSP metrics.

3.3. DELAY-MEASUREMENT Object

The DELAY-MEASUREMENT Object 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, locally or by PCE, the PCC SHOULD include the DELAY-MEASUREMENT Object to report the measured delay values to the PCE in the PCRpt message.

Object-Types are defined as follows:

Object-Type	Len	Name
0	4	Reserved
1	4	One-Way Delay Measurement Value
2	8	One-Way Delay Measurement Min/Max Values
3	4	One-Way Delay Variation Measurement Value
4	4	Two-Way Delay Measurement Value
5	8	Two-Way Delay Measurement Min/Max Values
6	4	Two-Way Delay Variation Measurement Value

The payload format is as follows:

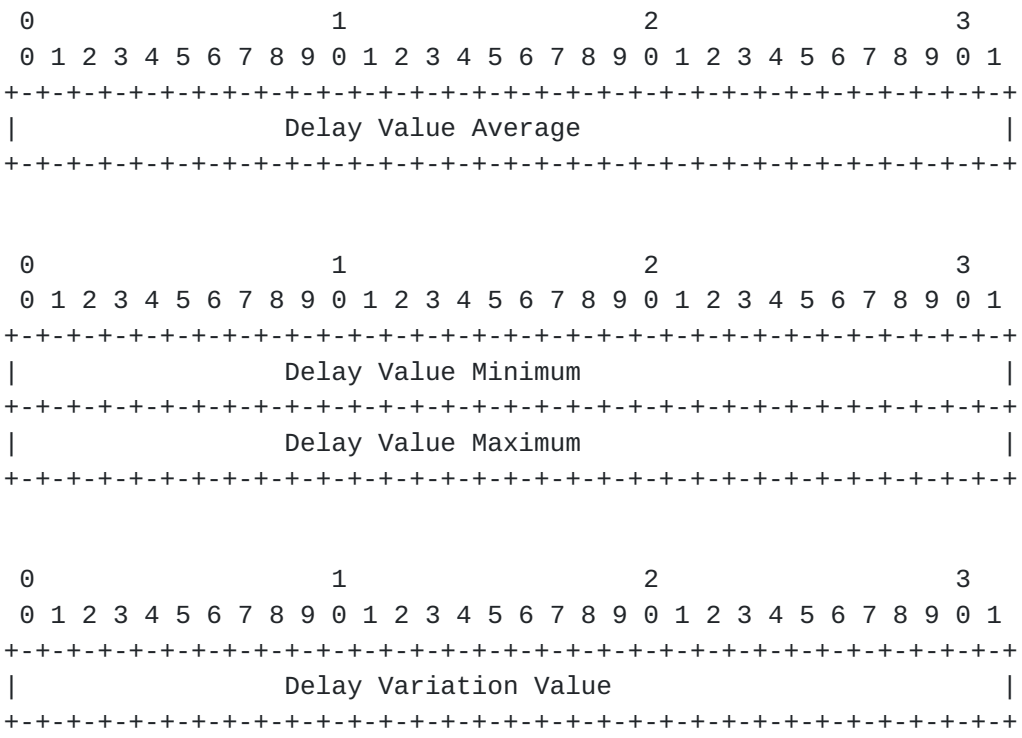

DELAY-MEASUREMENT Object formats (One-Way and Two-Way)

- One-way Delay Measurement 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.

- One-way Delay Measurement 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.

- Two-way Delay Measurement 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.

- Two-way Delay Measurement 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.

4. PCEP Extensions for Reporting Loss Measurement

4.1. 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 indicates that the Loss Measurement capability is supported as described in this document.

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. If the PCEP speaker that supports the extensions of this draft but did not advertise this capability, then upon receipt of LOSS-MEASUREMENT-ATTRIBUTE 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. 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-ATTRIBUTE TLV in LSPA object with two-way measurement request and one-way measurement request, respectively, when it did not advertise this capability. Further, the PCEP speaker SHOULD generate error-value TBD11 (Inferred Mode Measurement capability was not advertised) and TBD12 (Direct Mode Measurement capability was not advertised) upon receipt of LOSS-MEASUREMENT-ATTRIBUTE TLV in LSPA object with Inferred Mode loss measurement request and Direct Mode loss measurement request, respectively, when it did not advertise this capability.

4.1.1. 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:

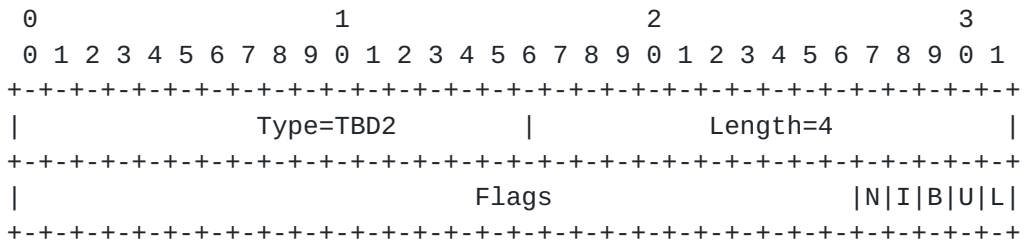

LOSS-MEASUREMENT-CAPABILITY TLV format

The type of the TLV is TBD2 and it has a fixed length of 4 octets.

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

- L (Loss Measurement - 1 bit): if set to 1 by a PCC, the L Flag indicates that the PCC allows reporting of loss measurement information; if set to 1 by a PCE, the L Flag indicates that the PCE is capable of receiving loss measurement information from the PCC. The LOSS-MEASUREMENTE-ATTRIBUTE TLV MUST be encoded when both PCEP speakers have the L Flag set.

- 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.

- 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.

- I (Inferred Mode - 1 bit): if set to 1 by a PCC, the I Flag indicates that the PCC allows reporting of inferred mode loss measurement 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.

- N (Direct Mode - 1 bit): if set to 1 by a PCC, the D Flag indicates that the PCC allows reporting of direct mode loss measurement information; if set to 1 by a PCE, the D 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 on transmission and MUST be ignored on receipt.

Advertisement of the Loss-Measurement capability TLV implies support of loss measurement, as well as the objects, TLVs and procedures defined in this document.

4.2. 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:

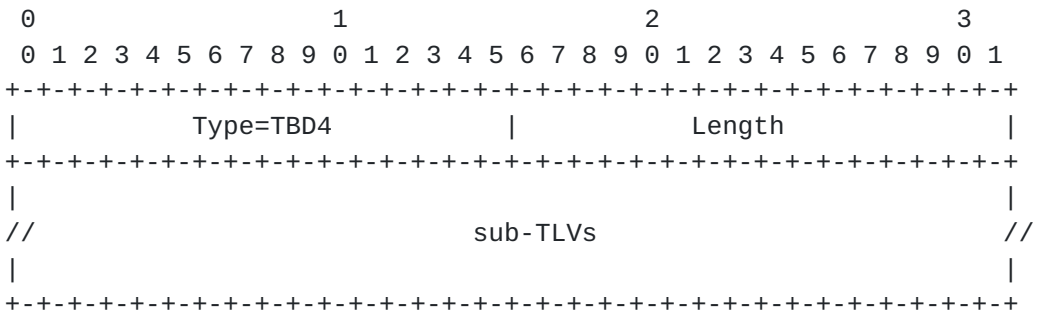

LOSS-MEASUREMENT-ATTRIBUTE TLV format

Type: TBD4

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
3	4	Measurement-Mode sub-TLV
4	4	Report-Threshold-Packets sub-TLV
5	4	Report-Threshold-Bytes sub-TLV
6	4	Report-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.

4.2.1. 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	Unidirectional Loss Measurement Enabled
2	Bidirectional Loss Measurement Enabled

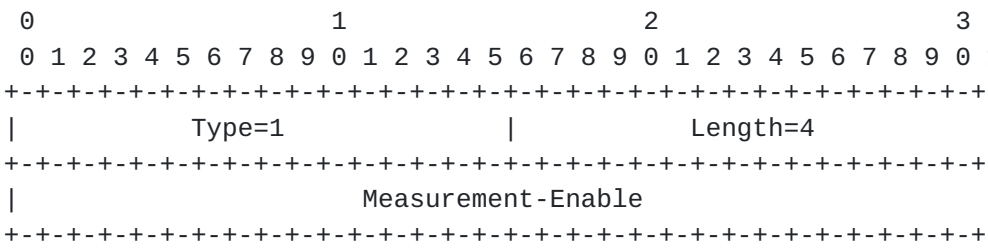

Measurement-Enable sub-TLV format

4.2.2. 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.

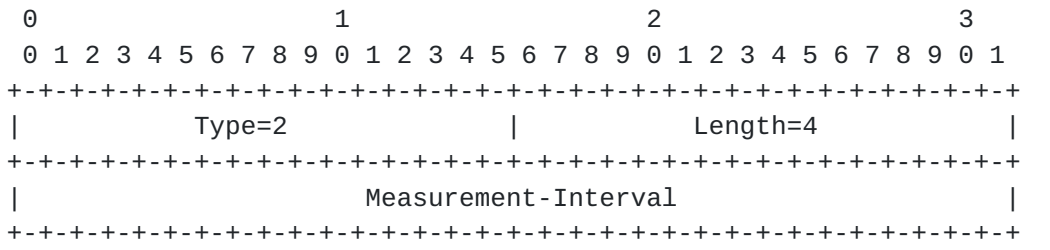

Measurement-Interval sub-TLV format

4.2.3. Measurement-Mode sub-TLV

The Measurement-Mode sub-TLV specifies the loss measurement mode which can be direct or inferred.

The Type is 3, Length is 4, and the value comprises of 4-octet loss measurement mode.

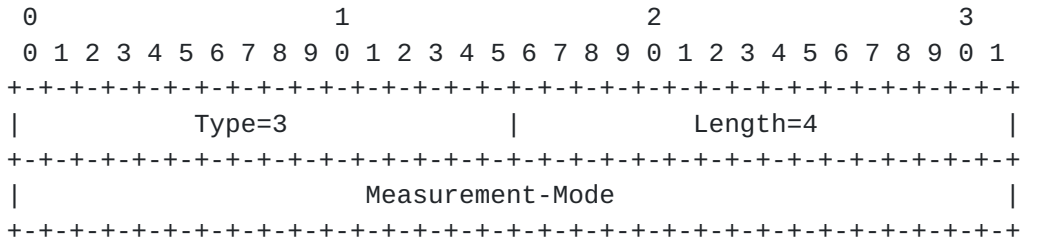

Measurement-Mode sub-TLV format

Value	Name
0	Inferred Mode (using test traffic)
1	Direct Mode (using data traffic)

4.2.4. Report-Threshold-Packets sub-TLV

The Report-Threshold-Packets sub-TLV specifies the loss threshold value.

The Type is 4, Length is 4, and the value comprises of 4-octet loss threshold values.

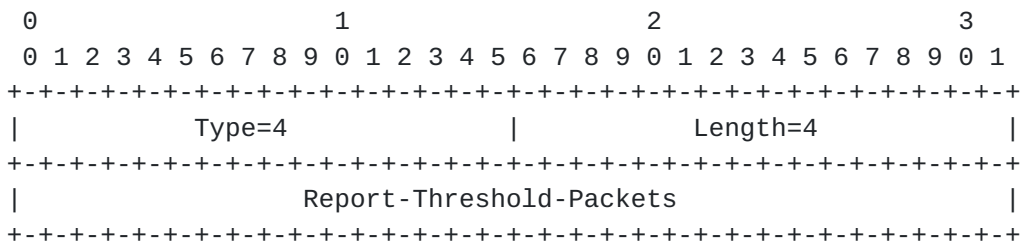

Report-Threshold-Packets sub-TLV format

The measured loss value MUST only be reported once the value exceeds the specified threshold value. This reporting can be used to trigger an immediate action at the PCE peer.

4.2.5. Report-Threshold-Bytes sub-TLV

The Report-Threshold-Bytes sub-TLV specifies the loss threshold value.

The Type is 5, Length is 4, and the value comprises of 4-octet loss threshold values.

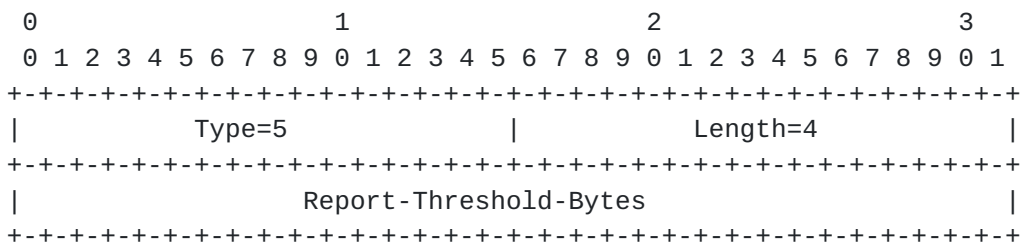

Report-Threshold-Bytes sub-TLV format

The measured loss value MUST only be reported once the value exceeds the specified threshold value. This reporting can be used to trigger an immediate action at the PCE peer.

4.2.6. Report-Interval sub-TLV

The Report-Interval sub-TLV specifies the time interval in seconds when measured loss values are to be reported. Maximum value of the measured metric (e.g. packets-lost, bytes-lost) collected during the report interval is reported.

The Type is 6, 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 equals to the Measurement-Interval.

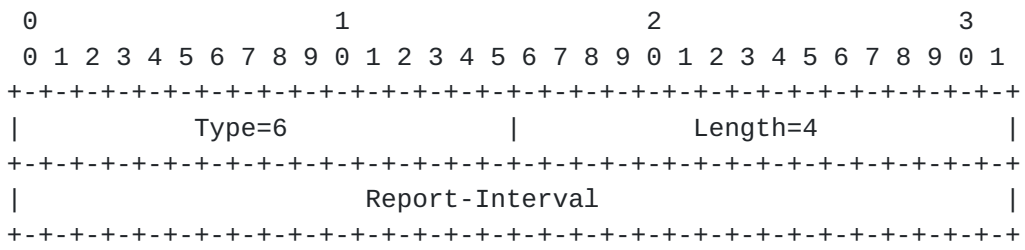

Report-Interval sub-TLV format

The measured loss value MUST be reported if the specified threshold is crossed OR the report-interval expires. The periodic reporting can be used at the PCEP peer to monitor the LSP metrics.

4.3. 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, locally or by PCE, the PCC SHOULD include the LOSS-MEASUREMENT Object to report the measured packet loss to the PCE in the PCRpt message.

Object-Types are defined as follows:

Object-Type	Len	Name
0	4	Reserved
1	4	Tx Packets-Lost
2	4	Tx Bytes-Lost
3	4	Rx Packets-Lost
4	4	Rx Bytes-Lost

The payload format is as follows:

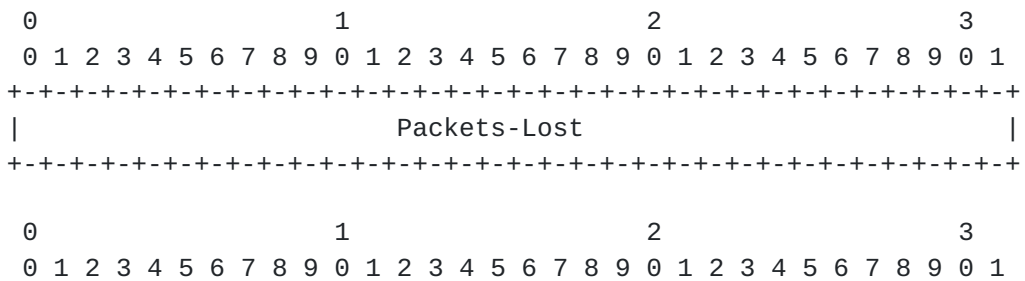

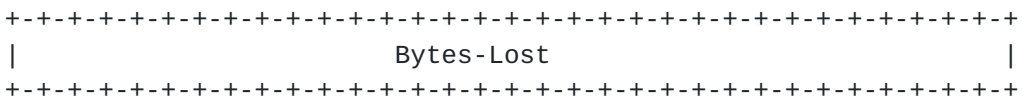

LOSS-MEASUREMENT Object formats (Tx and Rx directions)

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

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

Packets and Bytes lost in the Rx direction is reported only when bidirectional loss measurement is enabled.

5. Security Considerations

This document defines new MEASUREMENT objects and MEASUREMENT-ATTRIBUTE TLVs for reporting loss and delay measurements which do not add additional security concerns beyond those discussed in [RFC5440] and [DRAFT-PCE-STATEFUL].

Some deployments may find the reporting of the network 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. MEASUREMENT-CAPABILITY TLV Flag Field

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

New bit numbers to be allocated only by an IETF Consensus action. Each bit should be tracked with the following qualities:

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

Following types for DELAY-MEASUREMENT-CAPABILITY and LOSS-MEASUREMENT-CAPABILITY TLVs are defined in this document:

Type	Name	Reference
TBD1	DELAY-MEASUREMENT-CAPABILITY	[This I.D.]
TBD2	LOSS-MEASUREMENT-CAPABILITY	[This I.D.]

6.2. PCEP TLV Type Indicators

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

Value	Name	Reference
TBD3	DELAY-MEASUREMENT-ATTRIBUTE	[This I.D.]
TBD4	LOSS-MEASUREMENT-ATTRIBUTE	[This I.D.]

6.2.1. 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	Measurement-Mode sub-TLV	[This I.D.]
4	Report-Threshold sub-TLV	[This I.D.]
5	Report-Interval sub-TLV	[This I.D.]
6-65535	Unassigned	[This I.D.]

6.2.2. 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	Measurement-Mode sub-TLV	[This I.D.]
4	Report-Threshold-Packet sub-TLV	[This I.D.]
5	Report-Threshold-Bytes sub-TLV	[This I.D.]
6	Report-Interval sub-TLV	[This I.D.]
7-65535	Unassigned	[This I.D.]

6.3. 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
TBD5	DELAY-MEASUREMENT Object	[This I.D.]

6.3.1. 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:

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

6.4. 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
TBD6	LOSS-MEASUREMENT Object	[This I.D.]

6.4.1. 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:

Type	Name	Reference
0	Reserved	[This I.D.]
1	Tx Packets-Lost	[This I.D.]
2	Tx Bytes-Lost	[This I.D.]
3	Rx Packets-Lost	[This I.D.]
4	Rx Bytes-Lost	[This I.D.]

6.5. 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 I.D.]
TBD8	Loss-Measurement capability was not advertised	[This I.D.]
TBD9	Bidirectional Measurement capability was not advertised	[This I.D.]

TBD10 Unidirectional Measurement capability
was not advertised [This I.D.]
TBD11 Inferred Mode Measurement capability
was not advertised [This I.D.]
TBD12 Direct Mode Measurement capability
was not advertised [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

[RFC5586] Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", RFC 5586, June 2009.

[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.

[RFC7471] S. Giacalone, D. Ward, J. Drake, A. Atlas, S. Previdi, "OSPF Traffic Engineering (TE) Metric Extensions", RFC 7471, March 2015.

[RFC7810] S. Previdi, S. Giacalone, D. Ward, J. Drake, 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., 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., 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).

Acknowledgments

TBA.

Authors' Addresses

Rakesh Gandhi (editor)
Individual Contributor

EMail: rgandhi.ietf@gmail.com

Bin Wen
Comcast

EMail: Bin_Wen@cable.comcast.com

Colby Barth
Juniper Networks

EMail: cbarth@juniper.net