PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 18, 2015

D. Dhody
U. Palle
Huawei Technologies
R. Singh
Juniper Networks
R. Gandhi
Cisco Systems, Inc.

June 16, 2015

PCEP Extensions for MPLS-TE LSP Automatic Bandwidth Adjustment with Stateful PCE
draft-dhody-pce-stateful-pce-auto-bandwidth-05

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 provide stateful control of Multi-Protocol Label Switching (MPLS) Traffic Engineering Label Switched Paths (TE LSPs) via PCEP, for the case where PCC delegates control over one or more locally configured LSPs to the PCE.

This document describes automatic bandwidth adjustment of such LSPs when employing an Active Stateful PCE. In one of the models described, PCC computes the bandwidth to be adjusted and informs the PCE whereas in the second model, PCC reports the real-time traffic to a PCE and the PCE computes the adjustment bandwidth.

This document also describes automatic bandwidth adjustment for stateful PCE-initiated LSPs.

This Internet-Draft will expire on December 18, 2015.

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).

[I-D.ietf-pce-stateful-pce] specifies extensions to 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, the focus is on Active Stateful PCE where LSPs are configured at the PCC and control over them is delegated to the PCE. Further [I-D.ietf-pce-pce-initiated-lsp] describes the setup, maintenance and teardown of PCE-initiated LSPs under the stateful PCE model.

Over time, based on the varying traffic pattern, an LSP established with certain bandwidth may require to adjust the bandwidth, reserved in the network automatically. Ingress Label Switch Router (LSR) collects the traffic rate at each sample interval to determine the bandwidth demand of the LSP. This bandwidth information is then used to adjust the LSP bandwidth periodically. This feature is commonly referred to as Auto-Bandwidth.

Enabling Auto-Bandwidth feature on an LSP results in the LSP automatically adjusting its bandwidth based on the actual traffic flowing through the LSP. An LSP set-up with some arbitrary

(including zero) bandwidth value, automatically monitors the traffic flow and adjusts its bandwidth every adjustment-interval period. The bandwidth adjustment uses the make-before-break signaling method so that there is no interruption to traffic flow. This is described in detail in Section 4.1. [I-D.ietf-pce-stateful-pce-app] describes the use-case for Auto-Bandwidth adjustment for passive and active stateful PCE.

In this document, following deployment models are considered for employing Auto-Bandwidth feature with active stateful PCE.

- Deployment model 1: PCC to decide adjusted bandwidth:
  - In this model, the PCC (head-end of the LSP) monitors and calculates the new adjusted bandwidth. The PCC reports the calculated bandwidth to be adjusted to the PCE.
  - This approach would be similar to passive stateful PCE model, while the passive stateful PCE uses path request/reply mechanism, the active stateful PCE uses report/update mechanism to adjust the LSP bandwidth.
  - For PCE-initiated LSP, the PCC is requested during the LSP initiation to monitor and calculate the new adjusted bandwidth.
- Deployment model 2: PCE to decide adjusted bandwidth:
  - In this model, the PCE calculates the new adjusted bandwidth for the LSP.
  - Active stateful PCE can use information such as historical trending data, application-specific information about expected demands and central policy information along with real-time actual flow volumes to make smarter bandwidth adjustment to the delegated LSPs. Since the LSP has delegated control to the PCE, it is inherently suited that it should be the stateful PCE that decides the bandwidth adjustments.
  - For PCE-initiated LSP, the PCC is requested during initiation, to monitor and report the real-time bandwidth usage.
  - This model does not exclude use of any other mechanism employed by stateful PCE to learn real-time traffic information. But at the same time, using the same protocol (PCEP in this case) for updating and reporting the adjustment parameters as well as to learn real-time bandwidth usage is operationally beneficial.

This document defines extensions needed to support Auto-Bandwidth feature on the LSPs in a active stateful PCE model using PCEP.

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.

Delegation: An operation to grant a PCE temporary rights to modify a subset of tunnel parameters on one or more PCC's tunnels. Tunnels are delegated from a PCC to a 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.

Note the Auto-Bandwidth feature specific terms defined in Section 4.1.

3. Requirements for PCEP Extensions

There are two deployment models considered in this document for automatic bandwidth adjustments in case of active stateful PCE. In the model where PCC decides the adjusted bandwidth, PCC can report the new requested bandwidth and an active stateful PCE can update the bandwidth for a delegated LSP via existing mechanisms defined in [I-D.ietf-pce-stateful-pce]. Additional PCEP extensions required are summarized in the following table.

+--------+---------------------------+---------------------------+

Model	PCC Initiated	PCE Initiated
PCC to decide adjusted bandwidth	PCC monitors the traffic and reports the calculated bandwidth to be adjusted to the PCE. No new extensions are needed. Optionally AUTO-BANDWIDTH-ATTRIBUTE TLV can be used to identify the LSP with Auto-Bandwidth Feature enabled.	At the time of initiation, PCE request PCC to monitor the traffic and reports the calculated bandwidth to be adjusted to the PCE. Extension is needed for PCE to pass on the adjustment parameters at the time of Initiation. Refer the AUTO-BANDWIDTH-ATTRIBUTE TLV (and sub-TLVs e.g. Adjustment-Interval, Minimum-Bandwidth) in Section 5.1.
------	----------	----------
PCC reports real-time traffic and PCE to decide adjusted bandwidth	PCC monitors the traffic and reports the real-time traffic to the PCE. It is PCE that decides the calculated bandwidth to be adjusted and updates the LSP accordingly. Extension is needed for PCC to pass on the adjustment parameters at the time of delegation to PCE. Refer the AUTO-BANDWIDTH-ATTRIBUTE TLV (and sub-TLVs e.g. Adjustment-Threshold, Real-time-Traffic-Report-Interval) in Section 5.1. Further extension to	At the time of initiation, PCE request PCC to monitor the traffic and reports the real-time traffic to the PCE. It is PCE that decides the calculated bandwidth to be adjusted and updates the LSP accordingly. Extension is needed for PCE to pass on the real-time traffic reporting parameters at the time of Initiation. Refer the Real-time Traffic Reporting (e.g. Real-time-Traffic-Report-Interval, Real-time-Traffic-Report-Threshold) in Section 5.1.2. Further extension to report

		report the real-time traffic to PCE are also needed (Refer Bandwidth-Usage type in Section 5.2.2).	the real-time traffic to PCE are also needed (Refer Bandwidth-Usage type in Section 5.2.2).

Table 1: Auto-Bandwidth Deployment Models

Additional Auto-Bandwidth deployment considerations are summarized below:

- It is required to identify and inform the PCEP peer, the LSP that are enabled with Auto-Bandwidth feature. Not all LSPs in some deployments would like their bandwidth to be dependent on the real-time traffic but be constant as set by the operator.

- It is also required to identify and inform the PCEP peer the model of operation i.e. if PCC decides the adjusted bandwidth, or PCC reports the real-time traffic instead and the PCE decides the adjusted bandwidth.

  * Note that PCEP extension for reporting real-time traffic, as specified in this document, is one of the ways for a PCE to learn this information. But at the same time a stateful PCE may choose to learn this information from other means like management, performance tools, which are beyond the scope of this document.

- Further for the LSP with Auto-Bandwidth feature enabled, an operator should be able to specify the adjustment parameters (i.e. configuration knobs) to control this feature (e.g. minimum/ maximum bandwidth range) and PCEP peer should be informed.

4. Architectural Overview

4.1. Auto-Bandwidth Overview

Auto-Bandwidth feature allows an LSP to automatically and dynamically adjust its reserved bandwidth over time, i.e. without network operator intervention. The bandwidth adjustment uses the make-before-break signaling method so that there is no interruption to the traffic flow.

The new bandwidth reservation is determined by sampling the actual traffic flowing through the LSP. If the traffic flowing through the LSP is lower than the configured or current bandwidth of the LSP, the

extra bandwidth is being reserved needlessly and being wasted. Conversely, if the actual traffic flowing through the LSP is higher than the configured or current bandwidth of the LSP, it can potentially cause congestion or packet loss in the network. With Auto-Bandwidth feature, the LSP bandwidth can be set to some arbitrary value (including zero) during initial setup time, and it will be periodically adjusted over time based on the actual bandwidth requirement.

Note the following definitions of the Auto-Bandwidth terms:

Maximum Average Bandwidth (MaxAvgBw): The maximum average bandwidth represents the current traffic demand during a time interval. This is the maximum value of the averaged traffic rate in a given adjustment-interval.

Adjusted Bandwidth: This is the Auto-Bandwidth computed bandwidth that needs to be adjusted for the LSP.

Sample-Interval: The periodic time interval at which the traffic rate is collected as a sample.

Bandwidth-Sample (BwSample): The bandwidth sample collected at every sample interval to measure the traffic rate.

Adjustment-Interval: The periodic time interval at which the bandwidth adjustment should be made using the MaxAvgBw.

Maximum-Bandwidth: The maximum bandwidth that can be reserved for the LSP.

Minimum-Bandwidth: The minimum bandwidth that can be reserved for the LSP.

Adjustment-Threshold: This value is used to decide when the bandwidth should be adjusted. If the percentage or absolute difference between the current MaxAvgBw and the current bandwidth reservation is greater than or equal to the threshold value, the LSP bandwidth is adjusted to the current bandwidth demand (Adjusted Bandwidth) at the adjustment-interval expiry.

Overflow-Threshold: This value is used to decide when the bandwidth should be adjusted when there is a sudden increase in traffic demand. If the percentage or absolute difference between the current MaxAvgBw and the current bandwidth reservation is greater than or equal to the threshold value, the overflow-condition is set to be met. The LSP bandwidth is adjusted to the current

bandwidth demand bypassing the adjustment- interval if the overflow-condition is met consecutively for the overflow-counts.

Underflow-Threshold: This value is used to decide when the bandwidth should be adjusted when there is a sudden decrease in traffic demand. If the percentage or absolute difference between the current MaxAvgBw and the current bandwidth reservation is greater than or equal to the threshold value, the underflow-condition is set to be met. The LSP bandwidth is adjusted to the current bandwidth demand bypassing the adjustment- interval if the underflow-condition is met consecutively for the underflow-counts.

Report-Interval: This value indicates the periodic interval when the collected real-time traffic bandwidth samples (BwSample) should be reported to the stateful PCE via the PCRpt message.

Report-Threshold: This value is used to decide if the real-time traffic bandwidth samples collected should be reported. Only if the percentage or the absolute difference between at least one of the bandwidth samples collected and the current bandwidth reservation is greater than or equal to the threshold value, the bandwidth samples collected during the Report-Interval are reported otherwise the bandwidth sample(s) are skipped.

Report-Flow-Threshold: This value is used to decide when the real-time traffic bandwidth samples should be reported immediately when there is a sudden change in traffic demand. If the percentage or absolute difference between the current bandwidth sample and the current bandwidth reservation is greater than or equal to the flow threshold value, all the bandwidth samples collected so far are reported to the PCE immediately.

4.2. Theory of Operation

The traffic rate is periodically sampled at each sample-interval (which can be configured by the user and the default value as 5 minutes) by the head-end node of the LSP. The sampled traffic rates are accumulated over the adjustment-interval period (which can be configured by the user and the default value as 24 hours). The PCEP peer which is in-charge of calculating the bandwidth to be adjusted, will adjust the bandwidth of the LSP to the highest sampled traffic rate (MaxAvgBw) amongst the set of bandwidth samples collected over the adjustment-interval.

Note that the highest sampled traffic rate could be higher or lower than the current LSP bandwidth. Only if the difference between the current bandwidth demand (MaxAvgBw) and the current bandwidth reservation is greater than or equal to the Adjustment-Threshold

(percentage or absolute value), the LSP bandwidth is adjusted to the current bandwidth demand (MaxAvgBw).

In order to avoid frequent re-signaling, an operator may set a longer adjustment-interval value. However, longer adjustment-interval can result in an undesirable effect of masking sudden changes in traffic demands of an LSP. To avoid this, the Auto-Bandwidth feature may pre-maturely expire the adjustment-interval and adjust the LSP bandwidth to accommodate the sudden bursts of increase in traffic demand as an overflow condition or decrease in traffic demand as an underflow condition.

In case of Deployment model 2, the PCC reports the real-time traffic information and the PCE decides the adjusted bandwidth. Multiple bandwidth samples are collected every report-interval, and reported together to the PCE. To avoid reporting minor changes in real-time traffic, report-threshold is used, to suppress the sending of the collected samples during the report-interval. The collected samples are reported if at least one sample crosses the Report-Threshold (percentage or absolute value). In order to accommodate sudden changes in the real-time traffic, report flow threshold is employed by pre-maturely expiry of the report-interval to report the unreported bandwidth samples collected so far.

All thresholds in this document could be represented in both absolute value and percentage, and could be used together.

4.3. Scaling Considerations

There are potential scaling concerns for the model where PCC (ingress LSR) reports real-time traffic information to the stateful PCE for a large number of LSPs. It is recommended to combine multiple bandwidth samples (BwSample) using larger report-interval and report them together to the PCE, thus reducing the number of PCRpt messages. Further Report-Threshold can be use to skip reporting the bandwidth samples for small changes in the bandwidth.

The processing cost of monitoring a large number of LSPs at the PCC and handling bandwidth change requests at PCE should be taken into consideration. Note that, this will be implementation dependent.

5. Extensions to the PCEP

5.1. AUTO-BANDWIDTH-ATTRIBUTE TLV

The AUTO-BANDWIDTH-ATTRIBUTE TLV can be included as an optional TLV in the LSPA object (as described in [RFC5440]). Whenever the LSP with Auto-Bandwidth feature enabled is delegated, AUTO-BANDWIDTH-

ATTRIBUTE TLV is carried in PCRpt message in LSPA object. The TLV provides PCE with the 'configurable knobs' of this feature. In case of PCE-Initiated LSP ([I-D.ietf-pce-pce-initiated-lsp]) with Auto-Bandwidth feature enabled, this TLV is included in LSPA object with PCInitiate message.

The format of the AUTO-BANDWIDTH-ATTRIBUTE TLV is shown in the following figure:

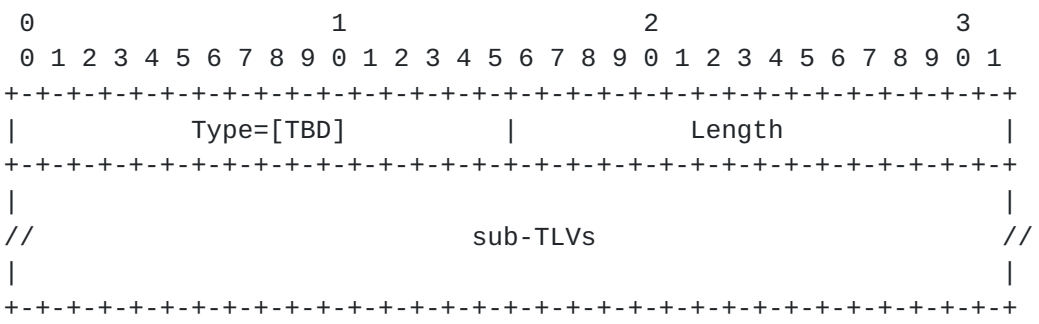

AUTO-BANDWIDTH-ATTRIBUTE TLV format

Type: TBD

Length: Variable

Value: This comprises one or more sub-TLVs.

Following sub-TLVs are defined in this document:

Type	Len	Name
1	4	Sample-Interval sub-TLV
2	4	Adjustment-Interval sub-TLV
3	4	Adjustment-Threshold sub-TLV
4	4	Adjustment-Threshold-Percentage sub-TLV
5	4	Minimum-Bandwidth sub-TLV
6	4	Maximum-Bandwidth sub-TLV
7	8	Overflow-Threshold sub-TLV
8	4	Overflow-Threshold-Percentage sub-TLV
9	8	Underflow-Threshold sub-TLV
10	4	Underflow-Threshold-Percentage sub-TLV
11	4	Real-time-Traffic-Report-Interval sub-TLV
12	4	Real-time-Traffic-Report-Threshold sub-TLV
13	4	Real-time-Traffic-Report-Threshold-Percentage sub-TLV
14	4	Real-time-Traffic-Report-Flow-Threshold sub-TLV
15	4	Real-time-Traffic-Report-Flow-Threshold-Percentage sub-TLV

Future specification can define additional sub-TLVs.

The presence of AUTO-BANDWIDTH-ATTRIBUTE TLV in LSPA object means that the automatic bandwidth adjustment feature is enabled. All 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 any others MUST be ignored.

If the sub-TLV are not encoded, the defaults based on the local policy are assumed.

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

5.1.1. Adjustment Parameters

The sub-TLVs in this section are encoded to inform the PCEP peer the various sampling and adjustment parameters, and serves the following purpose -

- For PCE-Initiated LSPs inform the PCC of the various sampling and adjustment parameters.
- For PCC-Initiated LSPs in the Deployment Model 2 (where PCE decides the adjusted bandwidth), inform the PCE of the various sampling and adjustment parameters.

5.1.1.1. Sample-Interval sub-TLV

The Sample-Interval sub-TLV specifies a time interval in seconds at which traffic samples are collected at the PCC.

The Type is 1, 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.

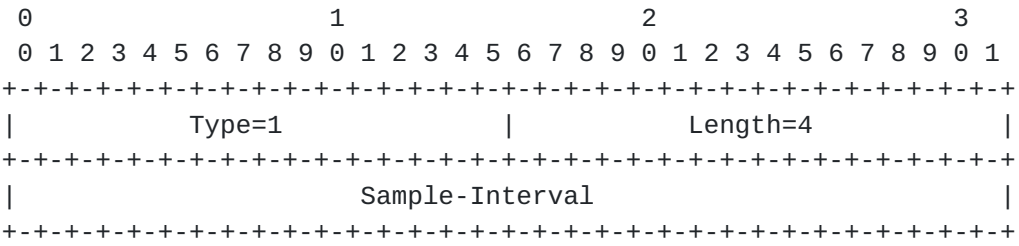

Sample-Interval sub-TLV format

5.1.1.2. Adjustment-Interval sub-TLV

The Adjustment-Interval sub-TLV specifies a time interval in seconds at which bandwidth adjustment should be made.

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.

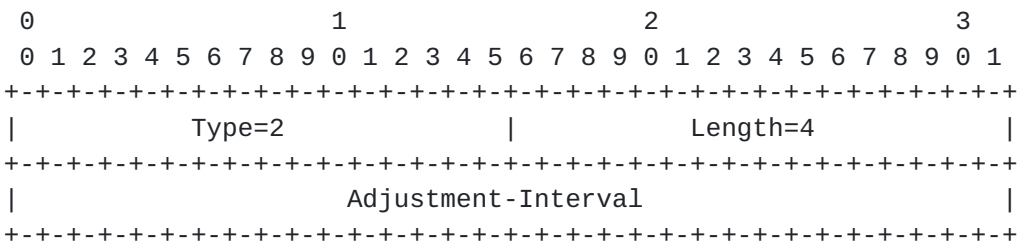

Adjustment-Interval sub-TLV format

5.1.1.3. Adjustment Threshold

The sub-TLVs in this section are encoded to inform the PCEP peer the adjustment threshold parameters. An implementation MAY include both sub-TLVs for the absolute value and the percentage, in which case the bandwidth is adjusted when either of the adjustment threshold conditions are met.

5.1.1.3.1. Adjustment-Threshold sub-TLV

The Adjustment-Threshold sub-TLV is used to decide when the LSP bandwidth should be adjusted.

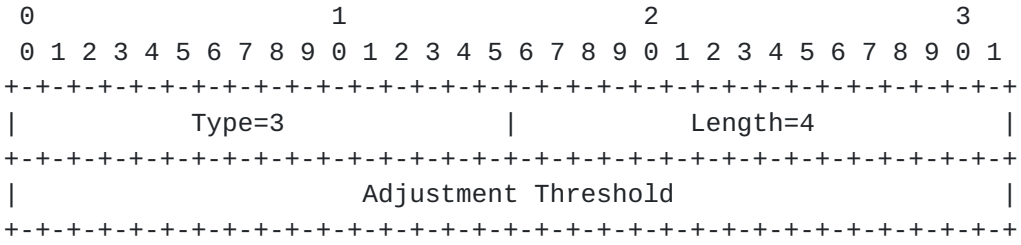

Adjustment-Threshold sub-TLV format

The Type is 3, Length is 4, and the value comprises of -

- Adjustment Threshold: The absolute Adjustment-Threshold bandwidth value, encoded in IEEE floating point format (see [IEEE.754.1985]), expressed in bytes per second. Refer to Section 3.1.2 of [RFC3471] for a table of commonly used values.

If the difference between the current MaxAvgBw and the current bandwidth reservation is greater than or equal to the threshold value, the LSP bandwidth is adjusted to the current bandwidth demand.

5.1.1.3.2. Adjustment-Threshold-Percentage sub-TLV

The Adjustment-Threshold-Percentage sub-TLV is used to decide when the LSP bandwidth should be adjusted.

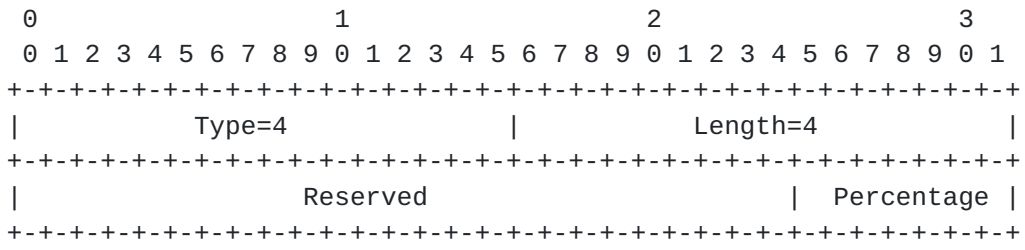

Adjustment-Threshold-Percentage sub-TLV format

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

- Reserved: SHOULD be set to zero on transmission and MUST be ignored on receipt.

- Percentage: The Adjustment-Threshold value, encoded in percentage (an integer from 0 to 100). If the percentage difference between the current MaxAvgBw and the current bandwidth reservation is greater than or equal to the threshold percentage, the LSP bandwidth is adjusted to the current bandwidth demand.

5.1.1.4. Minimum and Maximum Bandwidth

5.1.1.4.1. Minimum-Bandwidth sub-TLV

The Minimum-Bandwidth sub-TLV specify the minimum bandwidth allowed for the LSP, and is expressed in bytes per second. The LSP bandwidth cannot be adjusted below the minimum bandwidth value.

The Type is 5, Length is 4, and the value comprises of 4-octet bandwidth value encoded in IEEE floating point format (see [IEEE.754.1985]), expressed in bytes per second. Refer to Section 3.1.2 of [RFC3471] for a table of commonly used values.

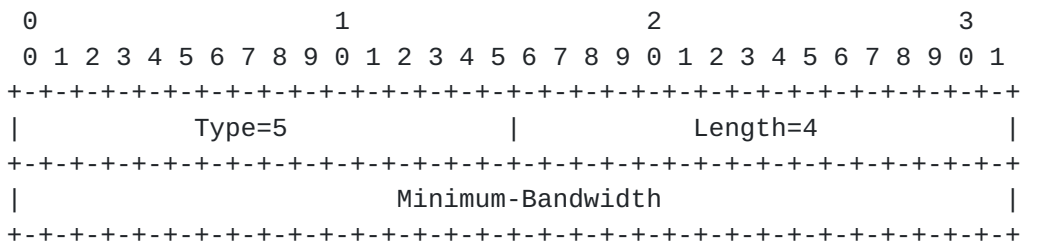

Minimum-Bandwidth sub-TLV format

5.1.1.4.2. Maximum-Bandwidth sub-TLV

The Maximum-Bandwidth sub-TLV specify the maximum bandwidth allowed for the LSP, and is expressed in bytes per second. The LSP bandwidth cannot be adjusted above the maximum bandwidth value.

The Type is 6, Length is 4, and the value comprises of 4-octet bandwidth value encoded in IEEE floating point format (see [IEEE.754.1985]), expressed in bytes per second. Refer to Section 3.1.2 of [RFC3471] for a table of commonly used values.

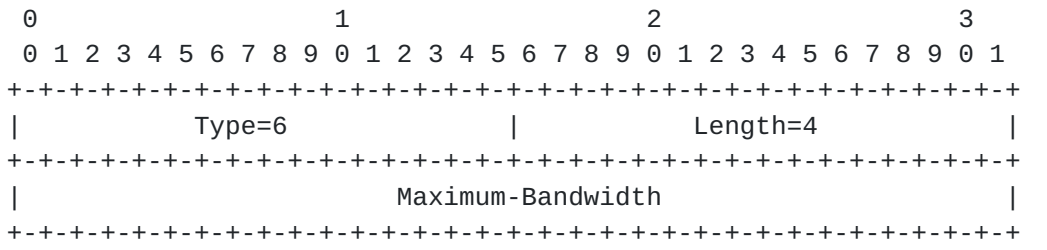

Maximum-Bandwidth sub-TLV format

5.1.1.5. Overflow and Underflow Condition

The sub-TLVs in this section are encoded to inform the PCEP peer the overflow and underflow threshold parameters. An implementation MAY include sub-TLVs for the absolute value and the percentage for the threshold, in which case the bandwidth is immediately adjusted when either of the adjustment threshold conditions are met consecutively for the given count.

5.1.1.5.1. Overflow-Threshold sub-TLV

The Overflow-Threshold sub-TLV is used to decide if the bandwidth should be adjusted immediately.

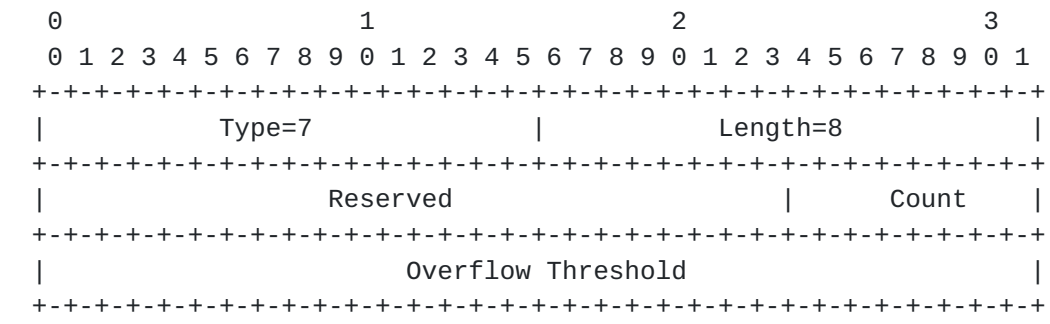

Overflow-Threshold sub-TLV format

The Type is 7, Length is 4, and the value comprises of -

- Reserved: SHOULD be set to zero on transmission and MUST be ignored on receipt.

- Count: The Overflow-Count value, encoded in integer. The value 0 is considered to be invalid. The number of consecutive samples for which the overflow condition MUST be met for the LSP bandwidth to be immediately adjusted to the current bandwidth demand, bypassing the adjustment-interval.

- Overflow Threshold: The absolute Overflow-Threshold bandwidth value, encoded in IEEE floating point format (see [IEEE.754.1985]), expressed in bytes per second. Refer to Section 3.1.2 of [RFC3471] for a table of commonly used values. If the increase of the current MaxAvgBw from the current bandwidth reservation is greater than or equal to the threshold value, the overflow condition is met.

5.1.1.5.2. Overflow-Threshold-Percentage sub-TLV

The Overflow-Threshold-Percentage sub-TLV is used to decide if the bandwidth should be adjusted immediately.

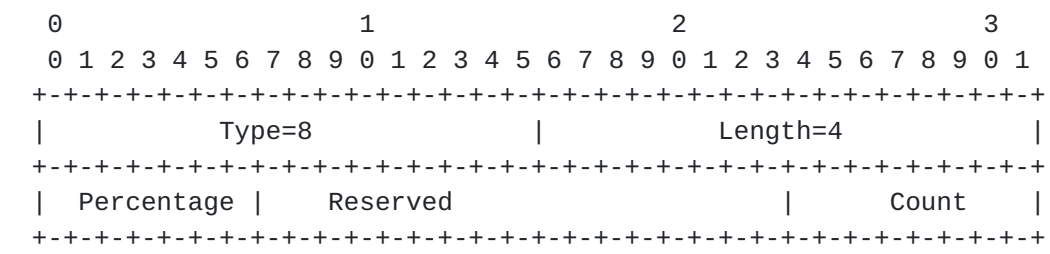

Overflow-Threshold-Percentage sub-TLV format

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

- Percentage: The Overflow-Threshold value, encoded in percentage (an integer from 0 to 100). If the percentage increase of the current MaxAvgBw from the current bandwidth reservation is greater than or equal to the threshold percentage, the overflow condition is met.

- Reserved: SHOULD be set to zero on transmission and MUST be ignored on receipt.

- Count: The Overflow-Count value, encoded in integer. The value 0 is considered to be invalid. The number of consecutive samples for which the overflow condition MUST be met for the LSP bandwidth to be immediately adjusted to the current bandwidth demand, bypassing the adjustment-interval.

5.1.1.5.3. Underflow-Threshold sub-TLV

The Underflow-Threshold sub-TLV is used to decide if the bandwidth should be adjusted immediately.

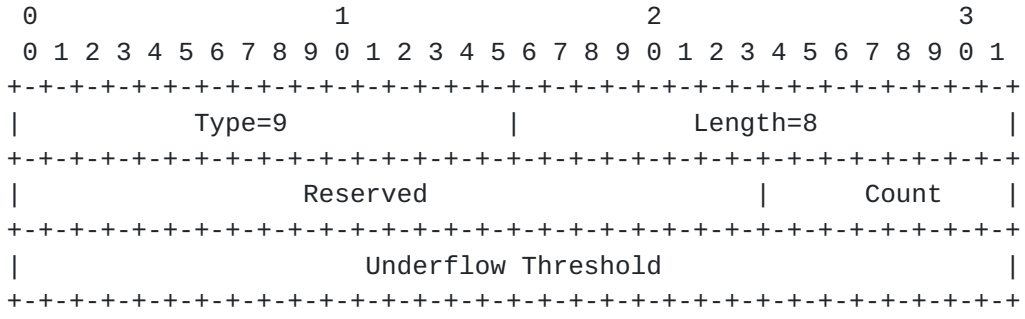

Underflow-Threshold sub-TLV format

The Type is 9, Length is 8, and the value comprises of -

- Reserved: SHOULD be set to zero on transmission and MUST be ignored on receipt.

- Count: The Underflow-Count value, encoded in integer. The value 0 is considered to be invalid. The number of consecutive samples for which the underflow condition MUST be met for the LSP bandwidth to be immediately adjusted to the current bandwidth demand, bypassing the adjustment-interval.

- Underflow Threshold: The absolute Underflow-Threshold bandwidth value, encoded in IEEE floating point format (see [IEEE.754.1985]), expressed in bytes per second. Refer to Section 3.1.2 of [RFC3471] for a table of commonly used values. If the decrease of the current MaxAvgBw from the current bandwidth

reservation is greater than or equal to the threshold value, the underflow condition is met.

5.1.1.5.4. Underflow-Threshold-Percentage sub-TLV

The Underflow-Threshold-Percentage sub-TLV is used to decide if the bandwidth should be adjusted immediately.

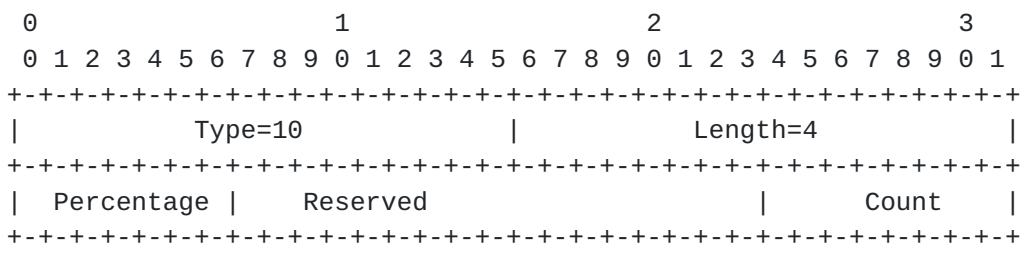

Underflow-Threshold-Percentage sub-TLV format

The Type is 10, Length is 4, and the value comprises of -

- Percentage: The Underflow-Threshold value, encoded in percentage (an integer from 0 to 100). If the percentage decrease of the current MaxAvgBw from the current bandwidth reservation is greater than or equal to the threshold percentage, the underflow condition is met.

- Reserved: SHOULD be set to zero on transmission and MUST be ignored on receipt.

- Count: The Underflow-Count value, encoded in integer. The value 0 is considered to be invalid. The number of consecutive samples for which the underflow condition MUST be met for the LSP bandwidth to be immediately adjusted to the current bandwidth demand, bypassing the adjustment-interval.

5.1.2. Real-time Traffic Reporting

The sub-TLVs in this section are encoded to inform the PCEP peer the various real-time traffic reporting parameters in the Deployment Model 2 (where PCE decides the adjusted bandwidth). In this model, Real-time-Traffic-Report-Interval sub-TLV MUST be included to specify the frequency of reporting.

The report threshold is used to decide if the collected bandwidth samples should be reported or skipped. An implementation MAY include both sub-TLVs for the absolute value and the percentage, in which case the real-time traffic is reported when either of the report threshold conditions are met.

The report flow threshold is used to decide when the collected bandwidth samples should be reported immediately, bypassing the report interval. An implementation MAY include both sub-TLVs for the absolute value and the percentage, in which case the real-time traffic is reported immediately when either of the report flow threshold conditions are met.

5.1.2.1. Real-time-Traffic-Report-Interval sub-TLV

The Real-time-Traffic-Report-Interval sub-TLV specifies a time interval in seconds in which collected bandwidth samples should be reported to PCE.

The Type is 11, Length is 4, and the value comprises of 4-octet time interval, the valid range is from 1 to 604800, in seconds.

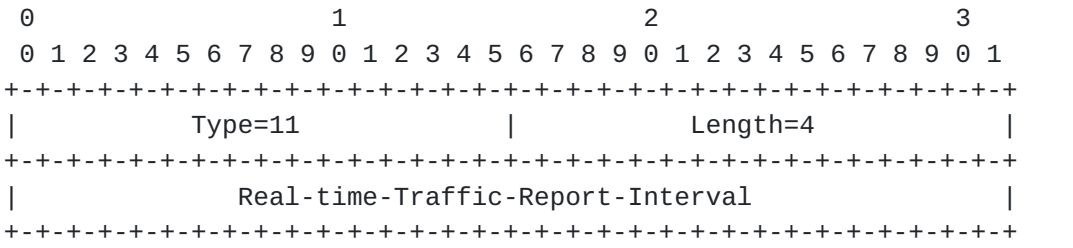

Real-time-Traffic-Report-Interval sub-TLV format

There is no default value. This sub-TLV MUST be included to enable the real-time traffic reporting.

5.1.2.2. Real-time-Traffic-Report-Threshold sub-TLV

The Real-time-Traffic-Report-Threshold sub-TLV is used to decide when the bandwidth samples collected should be reported immediately, bypassing the report-interval.

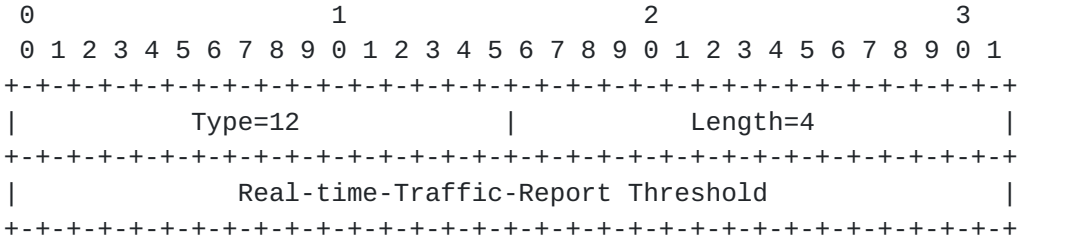

Real-time-Traffic-Report-Threshold sub-TLV format

The Type is 12, Length is 4, and the value comprises of -

- Threshold: The absolute threshold bandwidth value, encoded in IEEE floating point format (see [IEEE.754.1985]), expressed in bytes per second. Refer to Section 3.1.2 of [RFC3471] for a table of commonly used values. If the increase or the decrease of at least one of the bandwidth samples (BwSample) collected so far compared to the current bandwidth reservation is greater than or equal to the threshold value, the bandwidth samples collected so far are reported.

5.1.2.3. Real-time-Traffic-Report-Threshold-Percentage sub-TLV

The Real-time-Traffic-Report-Threshold sub-TLV is used to decide when the bandwidth samples collected should be reported immediately, bypassing the report-interval.

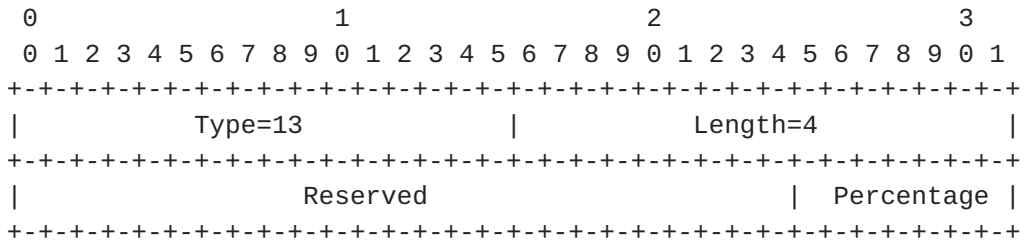

Real-time-Traffic-Report-Threshold-Percentage sub-TLV format

The Type is 13, Length is 4, and the value comprises of -

- Reserved: SHOULD be set to zero on transmission and MUST be ignored on receipt.

- Percentage: The threshold value, encoded in percentage (an integer from 0 to 100). If the percentage increase or the decrease of at least one of the bandwidth sample (BwSample) compared to the current bandwidth reservation is greater than or equal to the threshold percentage, the bandwidth samples collected so far are reported.

5.1.2.4. Real-time-Traffic-Report-Flow-Threshold sub-TLV

The Real-time-Traffic-Report-Flow-Threshold sub-TLV is used to decide when the bandwidth samples collected should be reported immediately, bypassing the report-interval.

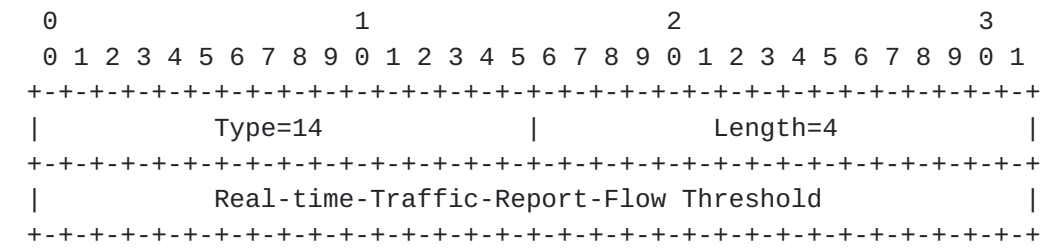

Real-time-Traffic-Report-Flow-Threshold sub-TLV format

The Type is 14, Length is 4, and the value comprises of -

- Threshold: The absolute flow threshold bandwidth value, encoded in IEEE floating point format (see [IEEE.754.1985]), expressed in bytes per second. Refer to Section 3.1.2 of [RFC3471] for a table of commonly used values. If the increase or the decrease of the current bandwidth sample (BwSample) compared to the current bandwidth reservation is greater than or equal to the flow threshold value, all the bandwidth samples collected so far are reported immediately, bypassing the report-interval.

5.1.2.5. Real-time-Traffic-Report-Flow-Threshold-Percentage sub-TLV

The Real-time-Traffic-Report-Flow-Threshold sub-TLV is used to decide when the bandwidth samples collected should be reported immediately, bypassing the report-interval.

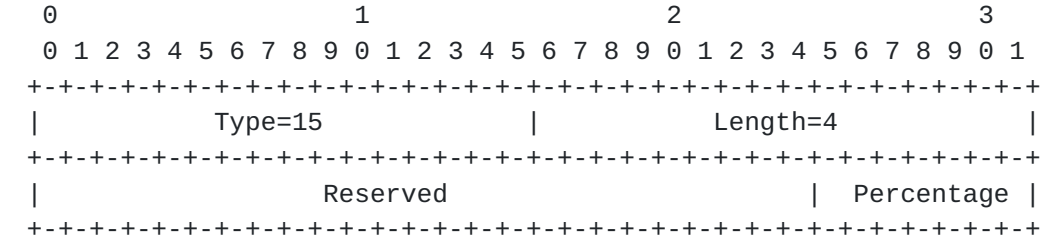

Real-time-Traffic-Report-Flow-Threshold-Percentage sub-TLV format

The Type is 15, Length is 4, and the value comprises of -

- Reserved: SHOULD be set to zero on transmission and MUST be ignored on receipt.

- Percentage: The flow threshold value, encoded in percentage (an integer from 0 to 100). If the percentage increase or the decrease of the current bandwidth sample (BwSample) compared to the current bandwidth reservation is greater than or equal to the threshold percentage, all the bandwidth samples collected so far are reported immediately, bypassing the report-interval.

5.2. BANDWIDTH Object

5.2.1. Auto-Bandwidth Adjusted Bandwidth

As per [RFC5440], the BANDWIDTH object is defined with two Object-Type values as following:

- Requested Bandwidth: BANDWIDTH Object-Type is 1.
- Re-optimization Bandwidth: Bandwidth of an existing TE LSP for which a re-optimization is requested. BANDWIDTH Object-Type is 2.

In the first model, where PCC calculates the adjusted bandwidth, PCC only reports the calculated bandwidth to be adjusted (MaxAvgBw) to the PCE. This is done via the existing 'Requested Bandwidth with BANDWIDTH Object-Type as 1'.

5.2.2. Bandwidth-Usage Report

A new BANDWIDTH object type is defined to report the actual bandwidth usage of a TE LSP.

The Object type is [TBD], the object body has a variable length, multiples of 4 bytes. The payload format is as follows:

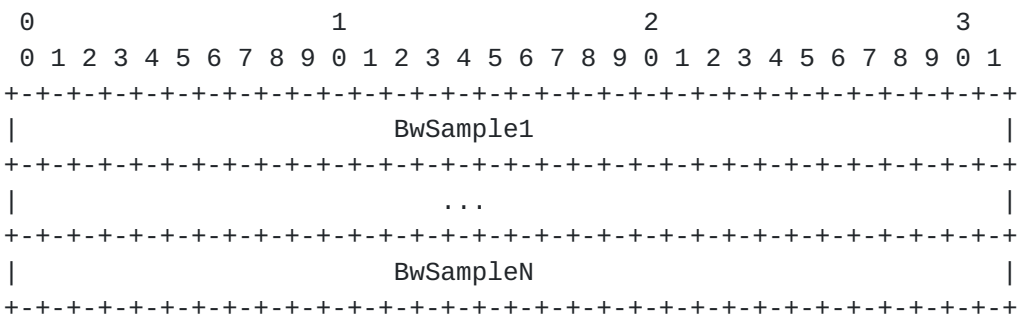

Bandwidth-Usage format

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

The Bandwidth-Usage object can be used in the second deployment model where PCC reports the TE LSP bandwidth usage and the PCE decides the auto-bandwidth adjusted bandwidth.

The Bandwidth-Usage object can also be used for TE LSPs without enabling the auto-bandwidth feature, to learn the actual bandwidth

usage of the LSPs for other applications at the stateful PCE. The details of which are beyond the scope of this document.

5.3. The PCRpt Message

When LSP is delegated to a PCE for the very first time, BANDWIDTH object of type 1 is used to specify the requested bandwidth in the PCRpt message.

When the LSP is enabled with the Auto-Bandwidth feature, and Real-time-Traffic-Report-Interval sub-TLV is not present (Deployment model 1), PCC SHOULD include the BANDWIDTH object of type 1 to specify the calculated bandwidth to be adjusted to the PCE in the PCRpt message.

When the LSP is enabled with the Auto-Bandwidth feature, and Real-time-Traffic-Report-Interval sub-TLV is present (Deployment model 2), PCC SHOULD include the BANDWIDTH object of type [TBD] to report the real-time traffic to the PCE in the PCRpt message.

The definition of the PCRpt message (see [I-D.ietf-pce-stateful-pce]) is unchanged by this document.

5.4. The PCInitiate Message

For PCE-initiated LSP [I-D.ietf-pce-pce-initiated-lsp] with Auto-Bandwidth feature enabled, AUTO-BANDWIDTH-ATTRIBUTE TLV MUST be included in LSPA object with the PCInitiate message. The rest of the processing remains unchanged.

6. Security Considerations

This document defines a new BANDWIDTH type and AUTO-BANDWIDTH-ATTRIBUTE TLV which do not add any new security concerns beyond those discussed in [RFC5440] and [I-D.ietf-pce-stateful-pce] in itself.

Some deployments may find the reporting of the real-time traffic information as extra sensitive and thus should employ suitable PCEP security mechanisms like TCP-AO or [I-D.ietf-pce-pceps].

7. Manageability Considerations

7.1. Control of Function and Policy

The Auto-Bandwidth feature MUST BE controlled per tunnel (at Ingress (PCC) or PCE), the values for parameters like sample-interval, adjustment- interval, minimum-bandwidth, maximum-bandwidth, adjustment-threshold, report-interval, report-threshold SHOULD be configurable by an operator.

7.2. Information and Data Models

[RFC7420] describes the PCEP MIB, there are no new MIB Objects for this document.

7.3. 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].

7.4. Verify Correct Operations

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

7.5. Requirements On Other Protocols

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

7.6. Impact On Network Operations

Mechanisms defined in this document do not have any impact on network operations in addition to those already listed in [RFC5440].

8. IANA Considerations

8.1. 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
TBD	AUTO-BANDWIDTH-ATTRIBUTE	[This I.D.]

8.2. AUTO-BANDWIDTH-ATTRIBUTE Sub-TLV

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

Type	Name	Reference
0	Reserved	[This I.D.]
1	Sample-Interval sub-TLV	[This I.D.]
2	Adjustment-Interval sub-TLV	[This I.D.]
3	Adjustment-Threshold sub-TLV	[This I.D.]
4	Adjustment-Threshold-Percentage sub-TLV	[This I.D.]
5	Minimum-Bandwidth sub-TLV	[This I.D.]
6	Maximum-Bandwidth sub-TLV	[This I.D.]
7	Overflow-Threshold sub-TLV	[This I.D.]
8	Overflow-Threshold-Percentage sub-TLV	[This I.D.]
9	Underflow-Threshold sub-TLV	[This I.D.]
10	Underflow-Threshold-Percentage sub-TLV	[This I.D.]
11	Real-time-Traffic-Report-Interval sub-TLV	[This I.D.]
12	Real-time-Traffic-Report-Threshold sub-TLV	[This I.D.]
13	Real-time-Traffic-Report-Threshold-Percentage sub-TLV	[This I.D.]
14	Real-time-Traffic-Report-Flow-Threshold sub-TLV	[This I.D.]
15	Real-time-Traffic-Report-Flow-Threshold -Percentage sub-TLV	[This I.D.]
16-65535	Unassigned	[This I.D.]

8.3. BANDWIDTH Object

This document defines new object type for the BANDWIDTH object; IANA is requested to make the following allocations from this registry. http://www.iana.org/assignments/pcep/pcep.xhtml#pcep-objects

Object-Class Value	Name	Reference
5	BANDWIDTH Object-Type	[This I.D.]
	TBD: Bandwidth-Usage Report	

9. Acknowledgments

We would like to thank Venugopal Reddy, Reeja Paul, Sandeep Boina and Avantika for their useful comments and suggestions.

10. References

10.1. 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.

[I-D.ietf-pce-stateful-pce]
Crabbe, E., Minei, I., Medved, J., and R. Varga, "PCEP Extensions for Stateful PCE", draft-ietf-pce-stateful-pce-11 (work in progress), April 2015.

[I-D.ietf-pce-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-04 (work in progress), April 2015.

[IEEE.754.1985]
Institute of Electrical and Electronics Engineers, "Standard for Binary Floating-Point Arithmetic", IEEE Standard 754, August 1985.

10.2. Informative References

[RFC3471] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.

[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.

[I-D.ietf-pce-stateful-pce-app]
Zhang, X. and I. Minei, "Applicability of a Stateful Path Computation Element (PCE)", draft-ietf-pce-stateful-pce-app-04 (work in progress), April 2015.

[I-D.ietf-pce-pceps]
Lopez, D., Dios, O., Wu, W., and D. Dhody, "Secure Transport for PCEP", draft-ietf-pce-pceps-04 (work in progress), May 2015.

Appendix A. Contributor Addresses

He Zekun
Tencent Holdings Ltd,
Shenzhen P.R.China

Email: kinghe@tencent.com

Xian Zhang
Huawei Technologies
Research Area F3-1B,
Huawei Industrial Base,
Shenzhen, 518129, China

Phone: +86-755-28972645
Email: zhang.xian@huawei.com

Young Lee
Huawei Technologies
1700 Alma Drive, Suite 100
Plano, TX 75075
USA

Phone: +1 972 509 5599 x2240
Fax: +1 469 229 5397
EMail: leeyoung@huawei.com

Authors' Addresses

Dhruv Dhody
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore, Karnataka 560037
India

EMail: dhruv.ietf@gmail.com

Udayasree Palle
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore, Karnataka 560037
India

EMail: udayasree.palle@huawei.com

Ravi Singh
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
USA

EMail: ravis@juniper.net

Rakesh Gandhi
Cisco Systems, Inc.

EMail: rgandhi@cisco.com