

PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: February 17, 2017

D. Dhody
U. Palle
Huawei Technologies
R. Singh
Juniper Networks
R. Gandhi
Individual Contributor
L. Fang
eBay
August 16, 2016

**PCEP Extensions for MPLS-TE LSP Automatic Bandwidth Adjustment with
Stateful PCE
draft-dhody-pce-stateful-pce-auto-bandwidth-08**

Abstract

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

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

This document describes PCEP extensions for automatic bandwidth adjustment when employing an Active Stateful PCE for both PCE-initiated and PCC-initiated LSPs. 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 bandwidth usage to a PCE and the PCE computes the adjustment bandwidth.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

Copyright 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	5
2.1.	Requirements Language	5
2.2.	Terminology	5
3.	Requirements for PCEP Extensions	6
4.	Architectural Overview	8
4.1.	Auto-Bandwidth Overview	8
4.2.	Theory of Operation	10
4.3.	Scaling Considerations	11
5.	Extensions to the PCEP	12
5.1.	Capability Advertisement	12
5.1.1	AUTO-BANDWIDTH-CAPABILITY TLV	12
5.2.	AUTO-BANDWIDTH-ATTRIBUTE TLV	13
5.2.1	Sample-Interval sub-TLV	14
5.2.2	Adjustment-Interval sub-TLV	15
5.2.3	Adjustment Threshold	15
5.2.3.1	Adjustment-Threshold sub-TLV	15
5.2.3.2	Adjustment-Threshold-Percentage sub-TLV	16
5.2.4	Minimum and Maximum Bandwidth Values	16
5.2.4.1	Minimum-Bandwidth sub-TLV	16
5.2.4.2	Maximum-Bandwidth sub-TLV	17
5.2.5	Overflow and Underflow Condition	17
5.2.5.1	Overflow-Threshold sub-TLV	17
5.2.5.2	Overflow-Threshold-Percentage sub-TLV	18
5.2.5.3	Underflow-Threshold sub-TLV	19
5.2.5.4	Underflow-Threshold-Percentage sub-TLV	20
5.3.	BANDWIDTH-USAGE-ATTRIBUTE TLV	20
5.3.1	Bandwidth-Usage-Report-Interval sub-TLV	22
5.3.2	Bandwidth-Usage-Report-Threshold sub-TLV	22

5.3.3.	Bandwidth-Usage-Report-Threshold-Percentage sub-TLV	23
5.3.4.	Bandwidth-Usage-Report-Flow-Threshold sub-TLV	23
5.3.5.	Bandwidth-Usage-Report-Flow-Threshold-Percent sub-TLV	24
5.4.	BANDWIDTH Object	25
5.4.1.	Auto-Bandwidth Adjusted Bandwidth Report	25
5.4.2.	Bandwidth-Usage Report	25
5.5.	The PCInitiate Message	26
5.6.	The PCRpt Message	26
5.7.	The PCNtf Message	27
6.	Security Considerations	27
7.	Manageability Considerations	28
7.1.	Control of Function and Policy	28
7.2.	Information and Data Models	28
7.3.	Liveness Detection and Monitoring	28
7.4.	Verify Correct Operations	28
7.5.	Requirements On Other Protocols	28
7.6.	Impact On Network Operations	28
8.	IANA Considerations	29
8.1.	PCEP TLV Type Indicators	29
8.2.	AUTO-BANDWIDTH-CAPABILITY TLV Flag Field	29
8.3.	AUTO-BANDWIDTH-ATTRIBUTE Sub-TLV	29
8.4.	BANDWIDTH-USAGE-ATTRIBUTE Sub-TLV	30
8.5.	BANDWIDTH Object	30
8.6.	Error Object	30
8.7.	Notification Object	31
9.	References	31
9.1.	Normative References	31
9.2.	Informative References	32
	Acknowledgments	33
	Contributors' Addresses	33
	Authors' Addresses	34

[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 for 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 reservation based on the actual traffic flowing through the LSP. The initial LSP bandwidth can be set to an arbitrary value (including zero), in practice, it can be operator expected value based on design and planning. Once the LSP is set-up, the LSP 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 the traffic flow. The Auto-Bandwidth 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.

- o 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.
- o 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 bandwidth usage 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 bandwidth usage 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

The PCEP speaker supporting this document MUST have a mechanism to advertise the automatic bandwidth adjustment capability.

As discussed in [Section 1](#), there are two deployment models considered in this document for automatic bandwidth adjustments in case of active stateful PCE. In the model 1, where PCC decides the adjusted bandwidth, PCC reports the new adjusted bandwidth and an active stateful PCE updates the bandwidth of a delegated LSP via existing mechanisms defined in [\[I-D.ietf-pce-stateful-pce\]](#). PCEP extensions required for both models are summarized in the following table.

+-----+	
Model 1: PCC decides adjusted BW	
+-----+	
PCC Initiated	PCE Initiated
+-----+	
PCC monitors the traffic	At the time of initiation,
and reports the calculated	PCE request PCC to monitor
bandwidth to be adjusted	the traffic and report the
to the PCE.	calculated bandwidth to be
	adjusted to the PCE.
No new extensions are needed.	Extension is needed for PCE
	to pass on the adjustment
	parameters at the time of
	Initiation.
+-----+	
+-----+	
Model 2: PCC reports bandwidth-usage, PCE decides adjusted BW	
+-----+	
PCC Initiated	PCE Initiated
+-----+	
PCC monitors bandwidth usage	At the time of initiation,
and reports the real-time	PCE request PCC to monitor

bandwidth usage to the PCE.	the traffic and reports the	
It is PCE that decides the	real-time bandwidth usage to	
calculated bandwidth to be	the PCE. It is PCE that decides	
adjusted and updates the	the calculated bandwidth to	
LSP accordingly.	be adjusted and updates the	
	LSP accordingly.	
Extension is needed for	Extension is needed for PCE	
PCC to pass on the auto-bw	to pass on the real-time	
adjustment parameters at	bandwidth usage reporting	
the time of delegation to	parameters at the time of	
PCE.	Initiation.	
Further extension to report	Further extension to report	
the bandwidth-usage to	the bandwidth-usage to	
PCE are also needed	PCE are also needed.	
+-----+	+-----+	+-----+

Table 1: Auto-Bandwidth PCEP extensions

Further Auto-Bandwidth deployment considerations are summarized below:

- o 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 bandwidth usage but be constant as set by the operator.
- o 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 bandwidth usage instead and the PCE decides the adjusted bandwidth.
 - * Note that PCEP extension for reporting real-time bandwidth usage, 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.
- o 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 bandwidth demand during a time interval. This is the maximum value of the averaged traffic bandwidth 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-Count: This value is used to decide when the bandwidth should be adjusted when there is a sudden increase in traffic demand. This value indicates how many times consecutively, the percentage or absolute difference between the current MaxAvgBw and the current bandwidth reservation is greater than or equal to the Overflow-Threshold value.

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-Count.

Underflow-Count: This value is used to decide when the bandwidth should be adjusted when there is a sudden decrease in traffic demand. This value indicates how many times consecutively, the percentage or absolute difference between the current MaxAvgBw and the current bandwidth reservation is greater than or equal to the Underflow-Threshold value.

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-Count.

Report-Interval: This value indicates the periodic interval when the collected real-time bandwidth-usage 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 bandwidth-usage 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, the bandwidth usage report condition is said to be met. The collected bandwidth usage information is reported bypassing the report-interval if the bandwidth usage report condition is met consecutively for the Report-Flow-Count.

Report-Flow-Count: This value is used to decide when the bandwidth usage collected so far should be reported when there is a sudden change in bandwidth usage. This value indicates how many times consecutively, the percentage or absolute difference between the current bandwidth sample and the current bandwidth reservation is greater than or equal to the Report-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). Some LSPs are less eventful while other LSPs may encounter a lot of changes in the traffic pattern. PCE sets the intervals for reporting and adjustment based on the traffic pattern of the LSP.

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 bandwidth-usage 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 bandwidth-usage, 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 bandwidth usage, 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

It should be noted that with both deployment models, any bandwidth change would require re-signaling of an LSP in a make-before-break fashion, which can further trigger preemption of lower priority LSPs in the network. When deployed under scale, this can lead to a signaling churn in the network. The Auto-bandwidth application algorithm is thus advised to take this into consideration before adjusting the LSP bandwidth. Operators are advised to set the values of various auto-bandwidth adjustment and reporting parameters appropriate for the deployed LSP scale.

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

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.

If a PCE gets overwhelmed, it can notify the PCC to temporarily suspend its auto-bandwidth reporting (see [Section 5.7](#)). Similarly if a PCC gets overwhelmed due to signaling churn, it can notify the PCE to temporarily suspend the LSP bandwidth adjustment.

5. Extensions to the PCEP


5.1. Capability Advertisement

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

The PCEP protocol extensions for Auto-Bandwidth adjustments MUST NOT be used if one or both PCEP Speakers have not included the "Auto-Bandwidth 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 AUTO-BANDWIDTH-ATTRIBUTE TLV in LSPA object, it SHOULD generate a PCErr with error-type 19 (Invalid Operation), error-value TBD4 (Auto-Bandwidth capability was not advertised) and it will terminate the PCEP session.

5.1.1 AUTO-BANDWIDTH-CAPABILITY TLV

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



AUTO-BANDWIDTH-CAPABILITY TLV format

The type of the TLV is [TBD5] and it has a fixed length of 4 octets.

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

- o B (bandwidth usage - 1 bit): if set to 1 by a PCC, the B Flag indicates that the PCC allows reporting of bandwidth usage information; if set to 1 by a PCE, the B Flag indicates that the PCE is capable of receiving bandwidth usage information from the

PCC. The BANDWIDTH-USAGE-ATTRIBUTE TLV MUST be encoded when both PCEP speakers have the B Flag set.

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

Advertisement of the Auto-Bandwidth capability TLV implies support of auto-bandwidth adjustment, as well as the objects, TLVs and procedures defined in this document.


5.2. AUTO-BANDWIDTH-ATTRIBUTE TLV

The AUTO-BANDWIDTH-ATTRIBUTE TLV provides the 'configurable knobs' of the feature and it can be included as an optional TLV in the LSPA Object (as described in [RFC5440]).

For PCE-Initiated LSP ([I-D.ietf-pce-pce-initiated-lsp]), this TLV is included in the LSPA Object with PCInitiate message. For delegated LSPs, this TLV is carried in PCRpt message in LSPA Object.

The TLV is encoded in all PCEP messages for the LSP till the auto bandwidth adjustment feature is enabled, the absence of the TLV indicate the PCEP speaker wish to disable the feature.

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



AUTO-BANDWIDTH-ATTRIBUTE TLV format

Type: TBD1

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

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 all others MUST be ignored.

The AUTO-BANDWIDTH-ATTRIBUTE TLV can also be carried in PCUpd message in LSPA Object in order to make updates to auto-bandwidth attributes such as Adjustment-Interval.

If sub-TLVs are not present, the default values based on the local policy are assumed.

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

- o For PCE-Initiated LSPs, in the Deployment Model 1, inform the PCC of the various sampling and adjustment parameters.
- o For PCC-Initiated LSPs, in the Deployment Model 2, inform the PCE of the various sampling and adjustment parameters.

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

5.2.1. 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 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=1               |               Length=4               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Sample-Interval               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Sample-Interval sub-TLV format

5.2.2. 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 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               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Adjustment-Interval               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Adjustment-Interval sub-TLV format

5.2.3. 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.

5.2.3.1. Adjustment-Threshold sub-TLV

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

```

      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=3               |               Length=4               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Adjustment Threshold               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```



```

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

```

Adjustment-Threshold sub-TLV format

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

- o 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.

[5.2.3.2.](#) Adjustment-Threshold-Percentage sub-TLV

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

```

      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=4               |               Length=4               |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|               Reserved               |               Percentage |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Adjustment-Threshold-Percentage sub-TLV format

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

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

[5.2.4.](#) Minimum and Maximum Bandwidth Values

[5.2.4.1.](#) 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.

```

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=5           |           Length=4           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Minimum-Bandwidth           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Minimum-Bandwidth sub-TLV format

[5.2.4.2.](#) 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.

```

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=6           |           Length=4           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Maximum-Bandwidth           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Maximum-Bandwidth sub-TLV format

[5.2.5.](#) 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.

[5.2.5.1.](#) Overflow-Threshold sub-TLV

[illegible]

Overflow-Threshold sub-TLV format

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


5.2.5.2. Overflow-Threshold-Percentage sub-TLV

[illegible]

Overflow-Threshold-Percentage sub-TLV format

- 0 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.
- 0 Reserved: SHOULD be set to zero on transmission and MUST be ignored on receipt.
- 0 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.

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




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

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

The TLV is encoded in all PCEP messages for the LSP till the bandwidth usage reporting feature is enabled, the absence of the TLV

indicate the PCEP speaker wish to disable this feature.

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



BANDWIDTH-USAGE-ATTRIBUTE TLV format

Type: TBD2

Length: Variable

Value: This comprises one or more sub-TLVs.

Following sub-TLVs are defined in this document:

Type	Len	Name

1	4	Bandwidth-Usage-Report-Interval sub-TLV
2	4	Bandwidth-Usage-Report-Threshold sub-TLV
3	4	Bandwidth-Usage-Report-Threshold-Percentage sub-TLV
4	4	Bandwidth-Usage-Report-Flow-Threshold sub-TLV
5	4	Bandwidth-Usage-Report-Flow-Threshold-Percentage sub-TLV

Future specification can define additional sub-TLVs.

The presence of BANDWIDTH-USAGE-ATTRIBUTE TLV in LSPA Object means that the bandwidth usage reporting to PCE 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 all others MUST be ignored.

The BANDWIDTH-USAGE-ATTRIBUTE TLV can also be carried in PCUpd message in LSPA Object in order to make updates to the attributes such as Bandwidth-Usage-Report-Interval.

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 the BANDWIDTH-USAGE-ATTRIBUTE TLV, and serves the following purpose -

- o For PCE-Initiated LSPs, in the Deployment Model 2, inform the PCC of the various sampling and reporting parameters.

5.3.1. Bandwidth-Usage-Report-Interval sub-TLV

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

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. Default value is 3600 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=1                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Bandwidth-Usage-Report-Interval       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Bandwidth-Usage-Report-Interval sub-TLV format

5.3.2. Bandwidth-Usage-Report-Threshold sub-TLV

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

```

      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                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Bandwidth-Usage-Report Threshold       |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Bandwidth-Usage-Report-Threshold sub-TLV format

The Type is 2, Length is 4, and the value comprises of -

- o **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 (BwSamples) 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.

5.3.3. Bandwidth-Usage-Report-Threshold-Percentage sub-TLV

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

```

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=3                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Reserved                               | Percentage |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Bandwidth-Usage-Report-Threshold-Percentage sub-TLV format

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

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

5.3.4. Bandwidth-Usage-Report-Flow-Threshold sub-TLV

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

```

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=4                               | Length=8   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```



```

|           Reserved           |      Count      |
+-+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|      Bandwidth-Usage-Report-Flow Threshold      |
+-+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Bandwidth-Usage-Report-Flow-Threshold sub-TLV format

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

- o Reserved: SHOULD be set to zero on transmission and MUST be ignored on receipt.
- o Count: The Report-Flow-Count value, encoded in integer. The value 0 is considered to be invalid. The number of consecutive samples for which the bandwidth usage report condition MUST be met for the the bandwidth samples collected so far are reported immediately, bypassing the report-interval.
- o 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, bandwidth usage report condition is said to be met.

5.3.5. Bandwidth-Usage-Report-Flow-Threshold-Percent sub-TLV

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

```

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=5           |      Length=4      |
+-+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Percentage |      Reserved      |      Count      |
+-+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Bandwidth-Usage-Report-Flow-Threshold-Percentage sub-TLV format

The Type is 5, Length is 4, and the value comprises of -

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

- o 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, bandwidth usage report condition is said to be met.
- o Count: The Report-Flow-Count value, encoded in integer. The value 0 is considered to be invalid. The number of consecutive samples for which the bandwidth usage report condition MUST be met for the the bandwidth samples collected so far are reported immediately, bypassing the report-interval.

5.4. BANDWIDTH Object

5.4.1. Auto-Bandwidth Adjusted Bandwidth Report

As per [RFC5440], the BANDWIDTH object (Object-Class value 5) is defined with two Object-Type values as following:

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

In the first deployment model, where PCC calculates the adjusted bandwidth, PCC 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.

5.4.2. Bandwidth-Usage Report

A new Object-type for BANDWIDTH object is defined to report the real-time bandwidth usage of a TE LSP.

The Object-type is [TBD3], the object length is variable with multiples of 4 bytes. 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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               BwSample1                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               ...                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               BwSampleN                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```


Bandwidth-Usage format

- o 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-type is 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-type 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, details of which are beyond the scope of this document.

5.5. The PCInitiate Message

A PCInitiate message is a PCEP message sent by a PCE to a PCC to trigger LSP instantiation or deletion [[I.D.ietf-pce-pce-initiated-lsp](#)].

For the PCE-initiated LSP [[I-D.ietf-pce-pce-initiated-lsp](#)] with Auto-Bandwidth feature enabled, AUTO-BANDWIDTH-ATTRIBUTE TLV or BANDWIDTH-USAGE-ATTRIBUTE TLV MUST be included in the LSPA object with the PCInitiate message, depending on the deployment model being used. The rest of the processing remains unchanged.

5.6. The PCRpt Message

As specified in [[I.D.ietf-pce-pce-initiated-lsp](#)], the PCC creates the LSP using the attributes communicated by the PCE, and local values for the unspecified parameters. After the successful instantiation of the LSP, PCC automatically delegates the LSP to the PCE and generates an LSP State Report (PCRpt) for the LSP.

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 using 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 using Deployment Model 2, PCC SHOULD include the BANDWIDTH object of type

[TBD3] to report the real-time bandwidth-usage 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.

5.7. The PCNtf Message

As per [[RFC5440](#)], the PCEP Notification message (PCNtf) can be sent by a PCEP speaker to notify its peer of a specific event. As described in [Section 4.3](#) of this document, a PCEP speaker SHOULD notify its PCEP peer that it is overwhelmed, and on receipt of such notification the peer SHOULD NOT send any PCEP messages related to auto-bandwidth adjustment or bandwidth usage reporting. If a PCEP message related to auto-bandwidth adjustment or bandwidth usage reporting is received, it MUST be silently ignored.

When a PCEP speaker is overwhelmed, it SHOULD notify its peer by sending a PCNtf message with Notification Type = TBD6 (Auto-bandwidth Overwhelm State) and Notification Value = 1 (Entering auto-bandwidth overwhelm state). Optionally, OVERLOADED-DURATION TLV [[RFC5440](#)] MAY be included that specifies the time period during which no further PCEP messages related to auto-bandwidth adjustment or bandwidth usage reporting should be sent. When the PCEP speaker is no longer in the overwhelm state and is available to process the auto-bandwidth adjustment or bandwidth usage reporting, it SHOULD notify its peer by sending a PCNtf message with Notification Type = TBD6 (Auto-bandwidth Overwhelm State) and Notification Value = 2 (Clearing auto-bandwidth overwhelm state).

When Auto-Bandwidth feature is deployed using Deployment Model 1, a PCE can send this notification to PCC when a PCC is reporting frequent auto-bandwidth adjustments. If a PCC is overwhelmed with re-signaling/re-routing, it can also notify the PCE to not adjust the LSP bandwidth while in overwhelm state.

When Auto-Bandwidth feature is deployed using Deployment Model 2, a PCE can send this notification when a PCC is reporting bandwidth-usage too frequently for the PCE to process. In this model as well, if a PCC is overwhelmed with re-signaling/re-routing, it can notify the PCE to not update bandwidth while in overwhelm state.

6. Security Considerations

This document defines a new BANDWIDTH object type, AUTO-BANDWIDTH-CAPABILITY TLV, AUTO-BANDWIDTH-ATTRIBUTE TLV and

BANDWIDTH-USAGE-ATTRIBUTE TLV which do not add any new security concerns beyond those discussed in [[RFC5440](#)] and [[I-D.ietf-pce-stateful-pce](#)].

Some deployments may find the reporting of the real-time bandwidth-usage 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 SHOULD 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 defined in 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 add 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

TBD5	AUTO-BANDWIDTH-CAPABILITY	[This I.D.]
TBD1	AUTO-BANDWIDTH-ATTRIBUTE	[This I.D.]
TBD2	BANDWIDTH-USAGE-ATTRIBUTE	[This I.D.]

8.2. AUTO-BANDWIDTH-CAPABILITY TLV Flag Field

IANA is requested to create a registry to manage the Flag field of the AUTO-BANDWIDTH-CAPABILITY TLV.

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

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

One bit is defined for the AUTO-BANDWIDTH-CAPABILITY TLV Object flag field in this document:

Bit	Name	Reference

31	Bandwidth usage reporting	[This I.D.]

8.3. 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-	Unassigned	[This I.D.]

65535

8.4. BANDWIDTH-USAGE-ATTRIBUTE Sub-TLV

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

Type	Name	Reference
0	Reserved	[This I.D.]
1	Bandwidth-Usage-Report-Interval sub-TLV	[This I.D.]
2	Bandwidth-Usage-Report-Threshold sub-TLV	[This I.D.]
3	Bandwidth-Usage-Report-Threshold-Percentage sub-TLV	[This I.D.]
4	Bandwidth-Usage-Report-Flow-Threshold sub-TLV	[This I.D.]
5	Bandwidth-Usage-Report-Flow-Threshold-Percentage sub-TLV	[This I.D.]
6-	Unassigned	[This I.D.]

65535

8.5. BANDWIDTH Object

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

Object-Type	Name	Reference
TBD3	Bandwidth-Usage Report	[This I.D.]

8.6. Error Object

This document defines a new Error-Value for PCErr message of type 19 (Invalid Operation) [[I-D.ietf-pce-stateful-pce](#)]); IANA is requested to make the following allocation from this registry.

<<http://www.iana.org/assignments/pcep/pcep.xhtml#pcep-error-object>>

Error-Value	Meaning	Reference

TBD4	Auto-Bandwidth Capability was not Advertised	[This I.D.]

8.7. Notification Object

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

Type	Meaning	Reference

TBD6	Auto-Bandwidth Overwhelm State	[This I.D.]
	Notification-value=1: Entering Auto-Bandwidth overwhelm state	
	Notification-value=2: Clearing Auto-Bandwidth overwhelm state	

9. References

9.1. 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](#) (work in progress).
- [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](#) (work in progress).
- [IEEE.754.1985] Institute of Electrical and Electronics Engineers,

"Standard for Binary Floating-Point Arithmetic", IEEE Standard 754, August 1985.

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

Acknowledgments

Authors would like to thank Robert Varga, Venugopal Reddy, Reeja Paul, Sandeep Boina, Avantika, JP Vasseur and Himanshu Shah for their useful comments and suggestions.

Contributors' 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 560066
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
Individual Contributor
USA

EMail: rgandhi.ietf@gmail.com

Luyuan Fang
eBay
USA

EMail: luyuanf@gmail.com

