

PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: August 26, 2017

D. Dhody
U. Palle
Huawei Technologies
R. Singh
Juniper Networks
R. Gandhi
Individual Contributor
L. Fang
eBay
February 22, 2017

**PCEP Extensions for MPLS-TE LSP Automatic Bandwidth Adjustment with
Stateful PCE
draft-ietf-pce-stateful-pce-auto-bandwidth-02**

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.

Automatic bandwidth allows automatic and dynamic adjustment of the TE LSP bandwidth reservation based on the volume of traffic flowing through the LSP. This document describes PCEP extensions for automatic bandwidth adjustment when employing an Active Stateful PCE for both PCE-Initiated and PCC-Initiated LSPs.

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

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

Table of Contents

1.	Introduction	3
2.	Conventions Used in This Document	4
2.1.	Requirements Language	4
2.2.	Terminology	4
3.	Requirements for PCEP Extensions	5
4.	Architectural Overview	6
4.1.	Auto-Bandwidth Overview	6
4.2.	Auto-bandwidth Theory of Operation	8
4.3.	Scaling Considerations	9
5.	Extensions to the PCEP	9
5.1.	Capability Advertisement	9
5.1.1	AUTO-BANDWIDTH-CAPABILITY TLV	10
5.2.	AUTO-BANDWIDTH-ATTRIBUTE TLV	10
5.2.1.	Sample-Interval sub-TLV	12
5.2.2.	Adjustment Interval	12
5.2.2.1.	Up-Adjustment-Interval sub-TLV	12
5.2.2.2.	Down-Adjustment-Interval sub-TLV	13
5.2.3.	Adjustment Threshold	13
5.2.3.1.	Up-Adjustment-Threshold sub-TLV	13
5.2.3.2.	Up-Adjustment-Threshold-Percentage sub-TLV	14
5.2.3.3.	Down-Adjustment-Threshold sub-TLV	15
5.2.3.4.	Down-Adjustment-Threshold-Percentage sub-TLV	15
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	16
5.2.5.	Overflow and Underflow Conditions	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	18
5.2.5.4.	Underflow-Threshold-Percentage sub-TLV	19

5.3.	BANDWIDTH Object	20
5.4.	The PCInitiate Message	20
5.5.	The PCUpd Message	20
5.6.	The PCRpt Message	21
5.7.	The PCNtf Message	21
6.	Security Considerations	22
7.	Manageability Considerations	22
7.1.	Control of Function and Policy	22
7.2.	Information and Data Models	22
7.3.	Liveness Detection and Monitoring	23
7.4.	Verify Correct Operations	23
7.5.	Requirements On Other Protocols	23
7.6.	Impact On Network Operations	23
8.	IANA Considerations	24
8.1.	PCEP TLV Type Indicators	24
8.2.	AUTO-BANDWIDTH-CAPABILITY TLV Flag Field	24
8.3.	AUTO-BANDWIDTH-ATTRIBUTE Sub-TLV	24
8.4.	Error Object	25
8.5.	Notification Object	25
9.	References	26
9.1.	Normative References	26
9.2.	Informative References	26
	Acknowledgments	28
	Contributors' Addresses	28
	Authors' Addresses	29

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. Further [I-D.ietf-pce-pce-initiated-lsp] describes the setup, maintenance and teardown of PCE-Initiated LSPs for the stateful PCE model. In this document, the focus is on Active stateful PCE where the LSPs are controlled by the PCE.

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. The head-end Label Switch Router (LSR)

needs to monitor the actual bandwidth demand of the LSP and adjust the LSP bandwidth reservation periodically. This feature is commonly referred to as Auto-Bandwidth.

Enabling Auto-Bandwidth feature on an LSP results in the head-end LSR automatically adjusting the LSP bandwidth reservation based on the 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 head-end monitors the traffic flow on the LSP and adjusts the bandwidth every adjustment-interval. The bandwidth adjustment uses the make-before-break signaling method so that there is no disruption to the traffic flow. The Auto-Bandwidth feature is described in detail in [Section 4.1](#) of this document.

The PCC (head-end of the LSP) collects the traffic rate samples flowing through the LSP and calculates the new adjusted bandwidth. The PCC reports the calculated bandwidth to be adjusted to the PCE. This is similar to a 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. In case of PCE-Initiated LSP, the PCC is requested during the LSP initiation to monitor and calculate the new adjusted bandwidth. [\[RFC8051\]](#) describes the use-case for Auto-Bandwidth adjustment for passive and active stateful PCE.

The document [\[I-D.gandhi-pce-pm\]](#) describes the PCEP extensions for reporting the performance measurements to the PCE, and includes the real-time bandwidth usage information of a TE LSP. Those extensions can be used to implement the auto-bandwidth feature on a stateful PCE, i.e. can be used to calculate the new bandwidth to be adjusted on the stateful PCE.

This document defines the extensions needed to support Auto-Bandwidth features on the LSPs in a active stateful PCE model using PCEP where the bandwidth to be adjusted is calculated on the PCC (the head-end of the LSP).

[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 are 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.

Additional auto-bandwidth PCEP extensions required are summarized in the following table.

PCC Initiated	PCE Initiated
PCC monitors the traffic and reports the calculated bandwidth to be adjusted to the PCE.	At the time of initiation, PCE request PCC to monitor the traffic and report the calculated bandwidth to be adjusted to the PCE.
Extension is needed for PCC to pass on the adjustment parameters at the time of Delegation.	Extension is needed for PCE to pass on the adjustment parameters at the time of Initiation.

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

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 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 automatic and dynamic adjustment of the reserved bandwidth of an LSP over time, i.e. without network operator intervention. 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 demand.

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: The bandwidth sample collected at every sample interval to measure the traffic rate.

Up-Adjustment-Interval: The periodic time interval at which the bandwidth adjustment should be made using the MaxAvgBw, when MaxAvgBw is greater than the current bandwidth reservation.

Down-Adjustment-Interval: The periodic time interval at which the bandwidth adjustment should be made using the MaxAvgBw, when MaxAvgBw is lesser than the current bandwidth reservation.

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

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

Up-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 (upsized) to the current bandwidth demand (Adjusted Bandwidth) at the up-adjustment-interval expiry.

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

4.2. Auto-bandwidth 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 PCC which is 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 PCC (the head-end of the LSP) 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 period (Up or Down).

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 Up-Adjustment-Threshold (percentage or absolute value), the LSP bandwidth is adjusted (upsized) to the current bandwidth demand (MaxAvgBw). Similarly if the difference between the current bandwidth reservation and the current bandwidth demand (MaxAvgBw) is greater than or equal to the Down-Adjustment-Threshold (percentage or absolute value), the LSP bandwidth is adjusted (downsized) to the current bandwidth demand (MaxAvgBw). Some LSPs are less eventful while other LSPs may encounter a lot of changes in the traffic pattern. The intervals for adjustment is configured based on the traffic pattern of the LSP.

In order to avoid frequent re-signaling, an operator may set a longer adjustment-interval value (Up and/or Down). 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.

All thresholds in this document could be represented in both absolute value and percentage, and could be used together. This is done to accommodate case where the LSP bandwidth reservation may become very large or very small over time, the two representations help the operator to provide conditions when the bandwidth usage is too large or too small.

4.3. Scaling Considerations

It should be noted that any bandwidth change requires 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 parameters appropriate for the deployed LSP scale.

If a PCE gets overwhelmed, it can notify the PCC to temporarily suspend the reporting of the new LSP bandwidth to be adjusted (see [Section 5.7](#)). Similarly if a PCC gets overwhelmed due to signaling churn, it can notify the PCE to temporarily suspend the new LSP setup requests.

5. Extensions to the PCEP

5.1. Capability Advertisement

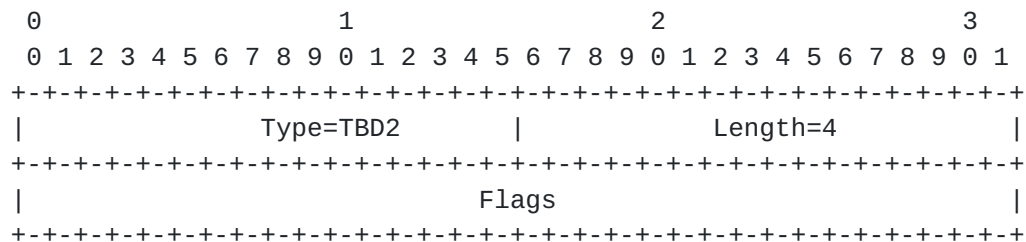
During PCEP Initialization Phase, PCEP Speakers (PCE or PCC) advertise their support of Automatic Bandwidth Adjustment feature. A PCEP Speaker includes the "Auto-Bandwidth Capability" TLV, in the OPEN Object to advertise its support for PCEP Auto-Bandwidth extensions. The presence of the "Auto-Bandwidth 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 (TBD2) and it has a fixed length of 4 octets.

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

Unassigned bits are considered reserved. They MUST be set to 0 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.

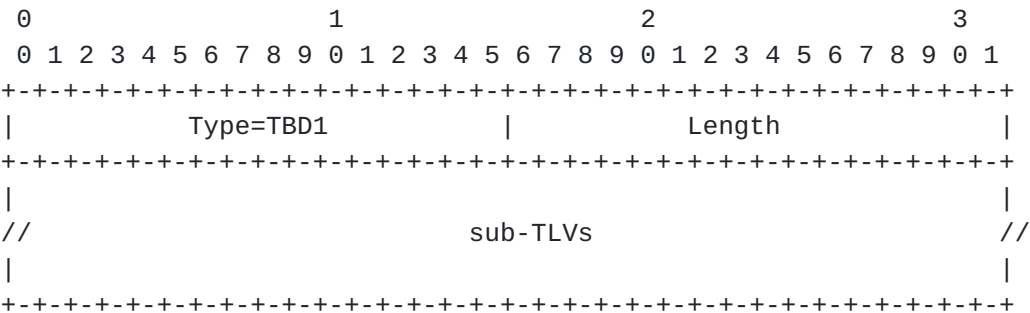
5.2. AUTO-BANDWIDTH-ATTRIBUTE TLV

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 the PCInitiate message. For the PCC-Initiated delegated LSPs, this TLV is carried in the PCRpt message in LSPA Object. This TLV is also included in the LSPA object with the PCUpd message to direct the PCC (the LSP head-end) to use different parameters with the LSP.

The TLV is encoded in all PCEP messages for the LSP while 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: The Length field defines the length of the value portion in bytes as per [RFC5440].

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	Up-Adjustment-Interval sub-TLV
3	4	Down-Adjustment-Interval sub-TLV
4	4	Up-Adjustment-Threshold sub-TLV
5	4	Up-Adjustment-Threshold-Percentage sub-TLV
6	4	Down-Adjustment-Threshold sub-TLV
7	4	Down-Adjustment-Threshold-Percentage sub-TLV
8	4	Minimum-Bandwidth sub-TLV
9	4	Maximum-Bandwidth sub-TLV
10	8	Overflow-Threshold sub-TLV
11	4	Overflow-Threshold-Percentage sub-TLV
12	8	Underflow-Threshold sub-TLV
13	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

The Up-Adjustment-Interval sub-TLV specifies a time interval in seconds at which bandwidth adjustment should be made when MaxAvgBw is

greater than the current bandwidth reservation.

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           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Up-Adjustment-Interval           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Up-Adjustment-Interval sub-TLV format

[5.2.2.2.](#) Down-Adjustment-Interval sub-TLV

The Adjustment-Interval sub-TLV specifies a time interval in seconds at which bandwidth adjustment should be made when MaxAvgBw is less than the current bandwidth reservation.

The Type is 3, 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=3           |           Length=4           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Down-Adjustment-Interval           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Down-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.](#) Up-Adjustment-Threshold sub-TLV

The Up-Adjustment-Threshold sub-TLV is used to decide when the LSP bandwidth should be adjusted when MaxAvgBw is greater than the

5.2.3.3. Down-Adjustment-Threshold sub-TLV

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	2	3	4	5	6	7	8	9
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-
	Type=6										Length=4																												
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-
	Down-Adjustment-Threshold																																						
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-

The Type is 6, Length is 4, and the value comprises of -

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

5.2.3.4. Down-Adjustment-Threshold-Percentage sub-TLV

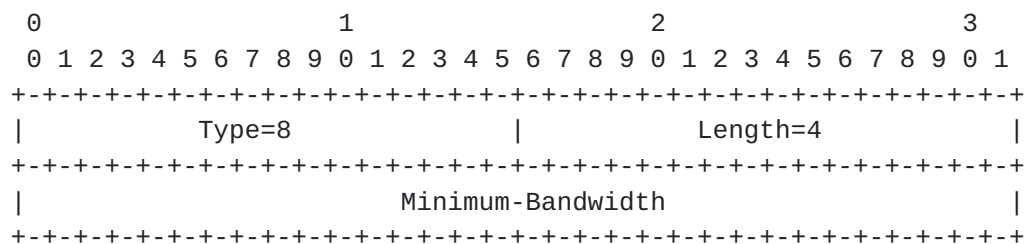
[illegible]

Down-Adjustment-Threshold-Percentage sub-TLV format

- 0 Reserved: SHOULD be set to zero on transmission and MUST be ignored on receipt.
- 0 Percentage: The Down-Adjustment-Threshold value, encoded in percentage (an integer from 0 to 100). If the percentage difference between the current bandwidth reservation and the current MaxAvgBw is greater than or equal to the threshold percentage, the LSP bandwidth is adjusted to the current bandwidth demand (MaxAvgBw).

5.2.4.1. Minimum-Bandwidth sub-TLV

The Type is 8, 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.



5.2.4.2. Maximum-Bandwidth sub-TLV

The Type is 9, 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.




```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Type=9           |           Length=4           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Maximum-Bandwidth           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Maximum-Bandwidth sub-TLV format

5.2.5. Overflow and Underflow Conditions

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

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

```

      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=10           |           Length=8           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Reserved           |           Count           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|           Overflow-Threshold           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Overflow-Threshold sub-TLV format

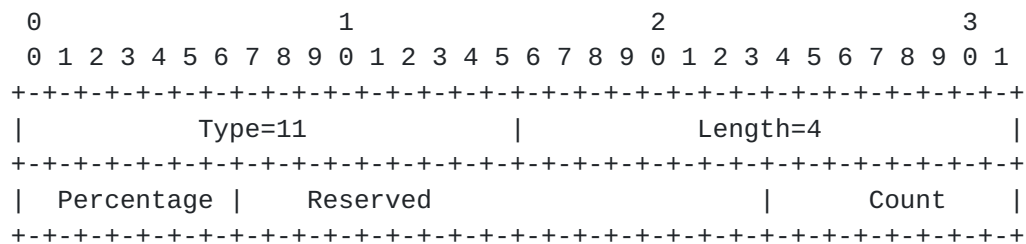
The Type is 10, Length is 8, and the value comprises of -

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

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



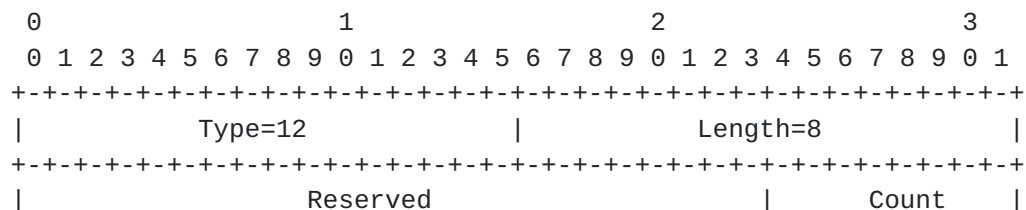
Overflow-Threshold-Percentage sub-TLV format

The Type is 11, Length is 4, and the value comprises of -

- o 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.
- 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.

5.2.5.3. Underflow-Threshold sub-TLV

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




```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Underflow-Threshold                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Underflow-Threshold sub-TLV format

The Type is 12, 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 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.

5.2.5.4. Underflow-Threshold-Percentage sub-TLV

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

```

      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=13                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                                     Length=4                                     |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Percentage |   Reserved   |                                     Count   |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Underflow-Threshold-Percentage sub-TLV format

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

- o 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.

- 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.

5.3. BANDWIDTH Object

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.

PCC reports the calculated bandwidth to be adjusted (MaxAvgBw) to the PCE using the existing 'Requested Bandwidth' with BANDWIDTH Object-Type as 1.

5.4. 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 with Auto-Bandwidth feature enabled, AUTO-BANDWIDTH-ATTRIBUTE TLV MUST be included in the LSPA object with the PCInitiate message.

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

5.5. The PCUpd Message

A PCUpd message is a PCEP message sent by a PCE to a PCC to update the LSP parameters [[I-D.ietf-pce-stateful-pce](#)].

For PCE-Initiated LSPs with Auto-Bandwidth feature enabled, AUTO-BANDWIDTH-ATTRIBUTE TLV MUST be included in the LSPA object with the PCUpd message. The PCE can send this TLV to direct the PCC to change the auto bandwidth parameters.

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

5.6. The PCRpt Message

The PCRpt message ([[I-D.ietf-pce-stateful-pce](#)]) is a PCEP message sent by a PCC to a PCE to report the status of one or more LSPs.

For PCE-Initiated LSPs [[I-D.ietf-pce-pce-initiated-lsp](#)], the PCC creates the LSP using the attributes communicated by the PCE, and using the local values for the unspecified parameters. After the successful instantiation of the LSP, PCC automatically delegates the LSP to the PCE and generates a PCRpt message to provide the status report for the LSP.

For both PCE-Initiated and PCC-Initiated LSPs, when the LSP is delegated to a PCE for the very first time, the BANDWIDTH object of type 1 is used to specify the requested bandwidth in the PCRpt message. After the successful delegation, to specify the LSP bandwidth to be adjusted, the BANDWIDTH object of type 1 is included in the PCRpt message.

For PCC-Initiated LSPs with Auto-Bandwidth feature enabled, AUTO-BANDWIDTH-ATTRIBUTE TLV MUST be included in the LSPA object of 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 (PCE or PCC) SHOULD notify its PCEP peer (PCC or PCE) when it is in overwhelmed state for the auto-bandwidth feature. Upon receipt of such notification, the peer SHOULD NOT send any PCEP messages related to auto-bandwidth adjustment. If a PCEP message related to auto-bandwidth is received during in overwhelmed state, it MUST be silently ignored.

When a PCEP speaker is overwhelmed, it SHOULD notify its peer by sending a PCNtf message with Notification Type = TBD3 (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 should be sent.

When the PCEP speaker is no longer in the overwhelm state and is available to process the auto-bandwidth adjustment, it SHOULD notify its peer by sending a PCNTf message with Notification Type = TBD3 (Auto-bandwidth Overwhelm State) and Notification Value = 2 (Clearing auto-bandwidth overwhelm state).

When Auto-Bandwidth feature is deployed, 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, it can also notify the PCE to not adjust the LSP bandwidth while in overwhelm state.

6. Security Considerations

This document defines AUTO-BANDWIDTH-CAPABILITY TLV 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 auto-bandwidth information as extra sensitive as it could be used to influence LSP path computation and LSP setup with adverse effect. Additionally, snooping of PCEP messages with such data or using PCEP messages for network reconnaissance, may give an attacker sensitive information about the operations of the network. Thus, such deployment should employ suitable PCEP security mechanisms like TCP Authentication Option (TCP-AO) [\[RFC5925\]](#) or [\[I-D.ietf-pce-pceps\]](#).

7. Manageability Considerations

7.1. Control of Function and Policy

The Auto-Bandwidth feature SHOULD be controlled per tunnel (at ingress (PCC) or PCE) and the values for auto-bandwidth parameters e.g. sample-interval, adjustment-interval (up/down), minimum-bandwidth, maximum-bandwidth, adjustment-threshold (up/down) SHOULD be configurable by an operator.

7.2. Information and Data Models

A Management Information Base (MIB) module for modeling PCEP is described in [\[RFC7420\]](#). However, one may prefer the mechanism for configuration using YANG data model [\[I-D.ietf-pce-pcep-yang\]](#). These SHOULD be enhanced to provide controls and indicators for support of auto-bandwidth feature. Support for various configuration knobs as well as counters of messages sent/received containing the TLVs (defined in this document) SHOULD be added.

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

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

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

Value	Name	Reference

TBD2	AUTO-BANDWIDTH-CAPABILITY	[This I.D.]
TBD1	AUTO-BANDWIDTH-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 are allocated only by an IETF Review action [[RFC5226](#)]. Each bit should be tracked with the following qualities:

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

No bit is defined for the AUTO-BANDWIDTH-CAPABILITY TLV Object flag field in this document.

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. New sub-TLV are allocated only by an IETF Review action [[RFC5226](#)].

This document defines the following types:

Type	Name	Reference

0	Reserved	[This I.D.]
1	Sample-Interval sub-TLV	[This I.D.]
2	Up-Adjustment-Interval sub-TLV	[This I.D.]

3	Down-Adjustment-Interval sub-TLV	[This I.D.]
4	Up-Adjustment-Threshold sub-TLV	[This I.D.]
5	Up-Adjustment-Threshold-Percentage sub-TLV	[This I.D.]
6	Down-Adjustment-Threshold sub-TLV	[This I.D.]
7	Down-Adjustment-Threshold-Percentage sub-TLV	[This I.D.]
8	Minimum-Bandwidth sub-TLV	[This I.D.]
9	Maximum-Bandwidth sub-TLV	[This I.D.]
10	Overflow-Threshold sub-TLV	[This I.D.]
11	Overflow-Threshold-Percentage sub-TLV	[This I.D.]
12	Underflow-Threshold sub-TLV	[This I.D.]
13	Underflow-Threshold-Percentage sub-TLV	[This I.D.]
14-	Unassigned	[This I.D.]

65535

8.4. Error Object

This document defines a new Error-Value for PCErr message of type 19 (Invalid Operation) [[I-D.ietf-pce-stateful-pce](https://tools.ietf.org/html/draft-ietf-pce-stateful-pce-04)]; 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.]

8.5. 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

TBD3	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.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 5226](#), May 2008.
- [RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", [RFC 5440](#), March 2009.
- [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).

9.2. Informative References

- [RFC3471] Berger, L., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", [RFC 3471](#), January 2003.
- [RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", [RFC 5925](#), June 2010.
- [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.
- [RFC8051] Zhang, X. and I. Minei, "Applicability of a Stateful Path Computation Element (PCE)", [RFC 8051](#), January 2017.
- [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).
- [I-D.ietf-pce-pcep-yang] Dhody, D., Hardwick, J., Beeram, V., and J. Tantsura, "A YANG Data Model for Path Computation Element Communications Protocol (PCEP)", [draft-ietf-pce-pcep-yang](#)

(work in progress).

[I-D.gandhi-pce-pm] Gandhi, R., Wen, B., Barth, C., and D. Dhody
"PCEP Extensions for Reporting MPLS-TE LSP Performance
Measurements", [draft-gandhi-pce-pm](#) (work in progress).

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

Acknowledgments

Authors would like to thank Robert Varga, Venugopal Reddy, Reeja Paul, Sandeep Boina, Avantika, JP Vasseur, Himanshu Shah and Adrian Farrel 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

EMail: rgandhi.ietf@gmail.com

Luyuan Fang
eBay
USA

EMail: lufang@ebay.com

