ï»¿

Internet Engineering Task Force (IETF)
Request for Comments: 8733
Category: Standards Track
ISSN: 2070-1721

D. Dhody, Ed.
Huawei Technologies
R. Gandhi, Ed.
Cisco Systems, Inc.
U. Palle
R. Singh
Individual Contributor
L. Fang
Expedia Group, Inc.

February 2020

Path Computation Element Communication Protocol (PCEP) Extensions for MPLS-TE Label Switched Path (LSP) Auto-Bandwidth Adjustment with Stateful PCE

Abstract

The Path Computation Element Communication Protocol (PCEP) provides mechanisms for Path Computation Elements (PCEs) to perform path computations in response to Path Computation Client (PCC) requests. Stateful PCE extensions allow stateful control of MPLS-TE Label Switched Paths (LSPs) using PCEP.

The auto-bandwidth feature 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 auto-bandwidth adjustment when employing an active stateful PCE for both PCE-initiated and PCC-initiated LSPs.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

[RFC5440] describes the Path Computation Element Protocol (PCEP) as a communication mechanism between a Path Computation Client (PCC) and a Path Computation Element (PCE), or between a PCE and a PCE, that enables computation of MPLS-TE Label Switched Paths (LSPs).

[RFC8231] specifies extensions to PCEP to enable stateful control of MPLS-TE LSPs. It describes two modes of operation: passive stateful PCE and active stateful PCE. Further, [RFC8281] describes the setup, maintenance, and teardown of PCE-initiated LSPs for the stateful PCE model. In this document, the focus is on the active stateful PCE, where the LSPs are controlled by the PCE.

Over time, based on the varying traffic pattern, an LSP established with a certain bandwidth may require adjustment of the bandwidth reserved in the network dynamically. The head-end Label Switching Router (LSR) monitors the actual bandwidth demand of the established LSP and periodically computes new bandwidth. The head-end LSR automatically adjusts the bandwidth reservation of the LSP based on the computed bandwidth. This feature, when available in the head-end LSR implementation, is commonly referred to as auto-bandwidth. The auto-bandwidth feature is described in detail in Section 4 of this document.

In the model considered in 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 the passive stateful PCE model: while the passive stateful PCE uses a path request/reply mechanism, the active stateful PCE uses a report/update mechanism. With a 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 PCEs.

Another approach would be to send the measured values themselves to the PCE, which is considered out of scope for this document.

This document defines the PCEP extensions needed to support an auto-bandwidth feature in an active stateful PCE model where the LSP bandwidth to be adjusted is calculated on the PCC (head-end of the LSP). The use of PCE to calculate the bandwidth to be adjusted is out of scope of this document.

2. Conventions Used in This Document

2.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2.2. Abbreviations

PCC: Path Computation Client

PCE: Path Computation Element

PCEP: Path Computation Element Communication Protocol

TE: Traffic Engineering

LSP: Label Switched Path

2.3. Terminology

The reader is assumed to be familiar with the terminology defined in [RFC5440], [RFC8231], and [RFC8281].

In this document, the PCC is considered to be the head-end LSR of the LSP. Other types of PCCs are not in scope.

The following auto-bandwidth terminology is defined in this document.

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

Adjusted Bandwidth: This is the auto-bandwidth 'computed' bandwidth that is used to adjust the bandwidth reservation of the LSP.

Sample-Interval: The periodic time interval at which the measured traffic rate of the LSP is collected as a Bandwidth-Sample.

Bandwidth-Sample: The Bandwidth-Sample of the measured traffic rate of the LSP collected at every Sample-Interval.

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-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 of the LSP.

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

Up-Adjustment-Threshold: This parameter is used to decide when the LSP 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 parameter is used to decide when the LSP 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 parameter is used to decide when the LSP bandwidth should be adjusted when there is a sudden increase in traffic demand. This value indicates how many times, consecutively, that the percentage or absolute difference between the current MaxAvgBw and the current bandwidth reservation of the LSP needs to be greater than or equal to the Overflow-Threshold value in order to meet the overflow condition.

Overflow-Threshold: This parameter is used to decide when the LSP 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 of the LSP is greater than or equal to the threshold value, the overflow condition is said to be met. The LSP bandwidth is adjusted to the current bandwidth demand, bypassing the Up-Adjustment-Interval if the overflow condition is met consecutively for the Overflow-Count. The Overflow-Threshold needs to be greater than or equal to the Up-Adjustment-Threshold.

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

Underflow-Threshold: This parameter is used to decide when the LSP 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 of the LSP is greater than or equal to the threshold value, the underflow condition is said to be met. The LSP bandwidth is adjusted to the current bandwidth demand, bypassing the Down-Adjustment-Interval if the underflow condition is met consecutively for the Underflow-Count. The Underflow-Threshold needs to be greater than or equal to the Down-Adjustment-Threshold.

Minimum-Threshold: When percentage-based thresholds are in use, they are accompanied by this Minimum-Threshold, which is used to ensure that the magnitude of deviation of the calculated LSP bandwidth to be adjusted from the current bandwidth reservations exceeds a specific non-percentage-based criterion (represented as an absolute bandwidth value) before any adjustments are made. This serves to suppress unnecessary auto-bandwidth adjustments and resignaling of the LSP at low bandwidth values.

3. Requirements for PCEP Extensions

The PCEP extensions required for auto-bandwidth are summarized in the following table as well as in Figure 1.

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

Table 1: Requirements for Auto-Bandwidth PCEP Extensions

```
                          ----------
                         |          |
                         |   PCE    |
                         |          |
                          ----------
                            |     ^
AUTO-BANDWIDTH CAPABILITY   |     |  AUTO-BANDWIDTH CAPABILITY
                            |     |
AUTO-BANDWIDTH ATTRIBUTES   |     |  AUTO-BANDWIDTH ATTRIBUTES
                            |     |  (For Delegated LSPs)
                            |     |
                            |     |  REQUESTED BANDWIDTH
                            v     |
                          ----------
                         |          |
                         |   PCC    |
                         |          |
                          ----------
```

Figure 1: Overview of Auto-Bandwidth PCEP Extensions

A PCEP speaker supporting this document must have a mechanism to advertise the auto-bandwidth adjustment capability for both PCC-initiated and PCE-initiated LSPs.

Auto-bandwidth deployment considerations for PCEP extensions are summarized below:

* It is necessary to identify and inform the PCC which LSPs have enabled the auto-bandwidth feature. Not all LSPs in some deployments would like their bandwidth to be dependent on real-time bandwidth usage; for some LSPs, leaving the bandwidth constant as set by the operator is preferred.

* In addition, an operator should be able to specify the auto-bandwidth adjustment parameters (i.e., configuration knobs) to control this feature (e.g., Minimum/Maximum-Bandwidth range). The PCC should be informed about these adjustment parameters.

4. Architectural Overview

4.1. Auto-Bandwidth Overview

The auto-bandwidth feature allows automatic and dynamic adjustment of the reserved bandwidth of an LSP over time (i.e., without network operator intervention) to accommodate the varying traffic demand of the LSP. If the traffic flowing through the LSP is lower than the configured or current reserved bandwidth of the LSP, the extra bandwidth is being reserved needlessly and is being wasted. Conversely, if the actual traffic flowing through the LSP is higher than the configured or current reserved bandwidth of the LSP, it can potentially cause congestion or packet loss in the network. The initial LSP bandwidth can be set to an arbitrary value (including zero). In practice, it can be set to an expected value based on design and planning. The head-end LSR monitors the actual traffic flowing through the LSP and uses that information to adjust the bandwidth reservation of the LSP in the network.

Bandwidth adjustment must not cause disruption to the traffic flow carried by the LSP. One way to achieve this is to use the make-before-break signaling method [RFC3209].

4.2. Auto-Bandwidth Theory of Operation

This section describes the auto-bandwidth feature in a general way. When the auto-bandwidth feature is enabled, the measured traffic rate is periodically sampled at each Sample-Interval by the PCC when the PCC is the head-end node of the LSP. The Sample-Interval can be configured by an operator, with a default value of 5 minutes. A very low Sample-Interval could have some undesirable interactions with transport protocols (see Section 6.6).

The traffic rate samples are accumulated over the Adjustment-Interval period (in the Up or Down direction). The period can be configured by an operator, with a default value of 24 hours. The PCC in charge of calculating the bandwidth to be adjusted can decide to adjust the bandwidth of the LSP to the highest traffic rate sample (MaxAvgBw) amongst the set of Bandwidth-Samples collected over the Adjustment-Interval period (in the Up or Down direction) depending on the operator policy.

Note that the highest traffic rate sample could be higher or lower than the current LSP bandwidth. The LSP is adjusted (upsized) to the current bandwidth demand (MaxAvgBW) only if the difference between the current bandwidth demand (MaxAvgBw) and the current bandwidth reservation is greater than or equal to the Adjustment-Threshold. The Adjustment-Threshold could be an absolute value or a percentage. The threshold can be configured by an operator, with a default value of 5 percent. 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 thresholds and intervals for bandwidth adjustment are configured based on the traffic pattern of the LSP.

In order to avoid frequent resignaling, an operator may set a longer Adjustment-Interval value (Up and/or Down). However, a longer Adjustment-Interval can result in the undesirable effect of masking sudden changes in the traffic demands of an LSP. To avoid this, the auto-bandwidth feature may force the Adjustment-Interval to prematurely expire 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. An operator needs to configure appropriate values for the Overflow-Threshold and/or Underflow-Threshold parameters, and they do not have default values defined in this document.

All thresholds in this document could be represented in both absolute value and percentage and could be used together. This is provided to accommodate cases where the LSP bandwidth reservation may become very large or very small over time. For example, an operator may use the percentage threshold to handle small to large bandwidth values and absolute values to handle very large bandwidth values. The auto-bandwidth adjustment is made when either one of the two thresholds, the absolute or percentage, is crossed.

When using the (adjustment/overflow/underflow) percentage thresholds, if the LSP bandwidth changes rapidly at very low values, it may trigger frequent auto-bandwidth adjustments due to the crossing of the percentage thresholds. This can lead to unnecessary resignaling of the LSPs in the network. This is suppressed by setting the Minimum-Threshold parameters along with the percentage thresholds. The auto-bandwidth adjustment is only made if the LSP bandwidth crosses both the percentage threshold and the Minimum-Threshold parameters.

4.3. Scaling Considerations

It should be noted that any bandwidth change requires resignaling of an LSP, 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. Similarly, if a PCC gets overwhelmed due to signaling churn, it can notify the PCE to temporarily suspend new LSP setup requests. See Section 5.7 of this document.

5. PCEP Extensions

5.1. Capability Advertisement

During the PCEP initialization phase, PCEP speakers (PCE or PCC) advertise their support of the auto-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 auto-bandwidth feature 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.

* A PCEP speaker that does not recognize the extensions defined in this document would simply ignore the TLVs as per [RFC5440].

* If a PCEP speaker supports the extensions defined in this document but did not advertise this capability, then upon receipt of AUTO-BANDWIDTH-ATTRIBUTES TLV in the LSP Attributes (LSPA) object, it SHOULD generate a PCErr with Error-Type 19 (Invalid Operation) and Error-value 14 (Auto-Bandwidth capability was not advertised) and ignore the AUTO-BANDWIDTH-ATTRIBUTES TLV.

5.1.1. AUTO-BANDWIDTH-CAPABILITY TLV

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

```
 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=36         |             Length=4          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                             Flag                              |
```

```
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 2: AUTO-BANDWIDTH-CAPABILITY TLV Format

The TLV Type is 36, and it has a fixed Length of 4 octets.

The value comprises a single field: Flag (32 bits). No flags are defined for this TLV in this document.

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-ATTRIBUTES TLV

The AUTO-BANDWIDTH-ATTRIBUTES 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 a PCE-initiated LSP [RFC8281], this TLV is included in the LSPA object with the PCInitiate message. For the PCC-initiated delegated LSPs, this TLV is carried in the Path Computation State Report (PCRpt) message in the LSPA object. This TLV is also carried in the LSPA object with the Path Computation Update Request (PCUpd) message to direct the PCC (LSP head-end) to make updates to auto-bandwidth attributes such as Adjustment-Interval.

The TLV is encoded in all PCEP messages for the LSP while the auto-bandwidth adjustment feature is enabled. The absence of the TLV indicates the PCEP speaker wishes to disable the feature. This TLV includes multiple AUTO-BANDWIDTH-ATTRIBUTES sub-TLVs. The AUTO-BANDWIDTH-ATTRIBUTES sub-TLVs are included if there is a change since the last information sent in the PCEP message. The default values for missing sub-TLVs apply for the first PCEP message for the LSP.

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

```
    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=37             |             Length            |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
   |                                                               |
   //                          sub-TLVs                           //
   |                                                               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 3: AUTO-BANDWIDTH-ATTRIBUTES TLV Format

Type: 37

Length: The Length field defines the length of the value portion in bytes as per [RFC5440].

Value: This comprises one or more sub-TLVs.

The following sub-TLVs are defined in this document:

Type	Len	Name
1	4	Sample-Interval
2	4	Adjustment-Interval
3	4	Down-Adjustment-Interval

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

Table 2: Sub-TLV Types of the AUTO-BANDWIDTH-ATTRIBUTES TLV

Future specifications can define additional sub-TLVs.

The sub-TLVs are encoded to inform the PCEP peer of the various sampling and adjustment parameters. In the case of a missing sub-TLV, as per the local policy, either the default value (as specified in this document) or some other operator-configured value is used.

All sub-TLVs are optional, and any unrecognized sub-TLV MUST be ignored. If a sub-TLV of the same type appears more than once, only the first occurrence is processed, and all others MUST be ignored.

The following subsections describe the sub-TLVs that are currently defined as being carried within the AUTO-BANDWIDTH-ATTRIBUTES TLV.

5.2.1. Sample-Interval Sub-TLV

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

```
 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                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 4: Sample-Interval Sub-TLV Format

The Type is 1, the Length is 4 octets, and the value comprises the following:

Sample-Interval: The 4-octet time interval for the Bandwidth-Sample collection. The valid range is from 1 to 604800 (7 days), in seconds. The default value is 300 seconds. Due care needs to be taken in a case with a very low Sample-Interval, as it can have some undesirable interactions with transport protocols (see Section 6.6). The Sample-Interval parameter MUST NOT be greater than the (down) Adjustment-Interval. In the case in which an invalid value is present, the sub-TLV MUST be ignored and the previous value will be maintained.

5.2.2. Adjustment-Intervals

The sub-TLVs in this section are encoded to inform the PCEP peer of the Adjustment-Interval parameters. The Adjustment-Interval sub-TLV specifies the time interval for both upward (Up-Adjustment-Interval) and downward (Down-Adjustment-Interval) trends. An implementation MAY require that a different Adjustment-Interval value be set when the bandwidth usage trend is moving downwards from the one used when it is moving upwards. In that case, the operator could use the Down-Adjustment-Interval sub-TLV, which overrides the Adjustment-Interval value for Down-Adjustment-Interval.

5.2.2.1. Adjustment-Interval Sub-TLV

The Adjustment-Interval sub-TLV specifies a time interval in seconds in which a bandwidth adjustment should be made in an upward or downward direction. This sub-TLV specifies the value for Up-Adjustment-Interval and Down-Adjustment-Interval when they are the same and when the Down-Adjustment-Interval sub-TLV is not included.

```
 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                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 5: Adjustment-Interval Sub-TLV Format

The Type is 2, the Length is 4 octets, and the value comprises the following:

Adjustment-Interval: The 4-octet time interval for bandwidth adjustments. The valid range is from 1 to 604800 (7 days), in seconds. The default value is 86400 seconds (1 day). The Adjustment-Interval parameter MUST NOT be less than the Sample-Interval; otherwise, the sub-TLV MUST be ignored, and the previous value will be maintained.

5.2.2.2. Down-Adjustment-Interval Sub-TLV

The Down-Adjustment-Interval sub-TLV specifies a time interval in seconds in which a bandwidth adjustment should be made when MaxAvgBw is less than the current bandwidth reservation of the LSP. This parameter overrides the Adjustment-Interval for the downward trend. This sub-TLV is used only when there is a need for different Adjustment-Intervals in the upward and downward directions.

```
 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                    |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 6: Down-Adjustment-Interval Sub-TLV Format

The Type is 3, the Length is 4 octets, and the value comprises the following:

Down-Adjustment-Interval: The 4-octet time interval for downward bandwidth adjustments. The valid range is from 1 to 604800 (7 days), in seconds. The default value equals the Adjustment-Interval. The Down-Adjustment-Interval parameter MUST NOT be less than the Sample-Interval; otherwise, the sub-TLV MUST be ignored and the previous value will be maintained.

5.2.3. Adjustment-Thresholds

The sub-TLVs in this section are encoded to inform the PCEP peer of 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. The Adjustment-Threshold sub-TLV specifies the threshold for both upward (Up-Adjustment-Threshold) and downward (Down-Adjustment-Threshold) trends. If the operator would like to use a different Adjustment-Threshold during the downward trend, the Down-Adjustment-Threshold sub-TLV is included. Similarly, the Adjustment-Threshold-Percentage sub-TLV specifies the threshold percentage for both upward and downward trends. If the operator would like to use a different Adjustment-Threshold percentage during the downward trend, the Down-Adjustment-Threshold-Percentage sub-TLV is included. It is worth noting that regardless of how the thresholds are set, the adjustment will not be made until at least one Sample-Interval has passed simply because no sample will be made on which to base a comparison with a threshold.

5.2.3.1. Adjustment-Threshold Sub-TLV

The Adjustment-Threshold sub-TLV is used to decide when the LSP bandwidth should be adjusted in an upward or downward direction. This sub-TLV specifies the absolute value for Up-Adjustment-Threshold and Down-Adjustment-Threshold when they are the same and when the Down-Adjustment-Threshold sub-TLV is not included.

```
 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          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Adjustment-Threshold                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 7: Adjustment-Threshold Sub-TLV Format

The Type is 4, the Length is 4 octets, and the value comprises the following:

Adjustment-Threshold: The absolute Adjustment-Threshold bandwidth difference value, encoded in IEEE floating point format (see [IEEE.754.1985]) and expressed in bytes per second. The default Adjustment-Threshold value is not set. Refer to Section 3.1.2 of [RFC3471] for a table of commonly used values.

If the modulus of 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 (MaxAvgBw).

In the case in which an invalid value is present, the sub-TLV MUST be ignored and the previous value will be maintained.

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 in an upward or downward direction. This sub-TLV specifies the percentage value for Up-Adjustment-Threshold and Down-Adjustment-Threshold when they are the same and when the Down-Adjustment-Threshold-Percentage sub-TLV is not included.

```
 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=8          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                   Reserved                    |  Percentage   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Minimum-Threshold                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 8: Adjustment-Threshold-Percentage Sub-TLV Format

The Type is 5, the Length is 8 octets, and the value comprises the following:

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

Percentage: The Adjustment-Threshold value (7 bits), encoded in a percentage (an integer from 1 to 100). The value 0 is considered to be invalid. The default value is 5 percent.

Minimum-Threshold: The absolute Minimum-Threshold bandwidth value, encoded in IEEE floating point format (see [IEEE.754.1985]) and expressed in bytes per second. The increase or decrease of the LSP bandwidth MUST be at or above the Minimum-Threshold before the bandwidth adjustment is made. The default value is 0.

If the percentage absolute difference between the current MaxAvgBw and the current bandwidth reservation is greater than or equal to the threshold percentage and the difference in the bandwidth is at or above the Minimum-Threshold, the LSP bandwidth is adjusted to the current bandwidth demand (MaxAvgBw).

In the case in which an invalid value is present, the sub-TLV MUST be ignored and the previous value will be maintained.

5.2.3.3. Down-Adjustment-Threshold Sub-TLV

The Down-Adjustment-Threshold sub-TLV is used to decide when the LSP bandwidth should be adjusted when MaxAvgBw is less than the current bandwidth reservation. This parameter overrides the Adjustment-Threshold for the downward trend. This sub-TLV is used only when there is a need for a different threshold in the upward and downward directions.

```
 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            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                    Down-Adjustment-Threshold                  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 9: Down-Adjustment-Threshold Sub-TLV Format

The Type is 6, the Length is 4 octets, and the value comprises the following:

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

If the difference between the 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).

In the case in which an invalid value is present, the sub-TLV MUST be ignored and the previous value will be maintained.

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

The Down-Adjustment-Threshold-Percentage sub-TLV is used to decide when the LSP bandwidth should be adjusted when MaxAvgBw is less than the current bandwidth reservation. This parameter overrides the Adjustment-Threshold-Percentage for the downward trend. This sub-TLV is used only when there is a need for a different threshold percentage in the upward and downward directions.

```
 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=7              |           Length=8            |
```

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                   Reserved                    |  Percentage |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         Minimum-Threshold                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 10: Down-Adjustment-Threshold-Percentage Sub-TLV Format

The Type is 7, the Length is 8 octets, and the value comprises the following:

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

Percentage: The Down-Adjustment-Threshold value (7 bits), encoded in a percentage (an integer from 1 to 100). The value 0 is considered to be invalid. The default value equals the Adjustment-Threshold-Percentage.

Minimum-Threshold: The absolute Minimum-Threshold bandwidth value, encoded in IEEE floating point format (see [IEEE.754.1985]) and expressed in bytes per second. The decrease of the LSP bandwidth MUST be at or above the Minimum-Threshold before the bandwidth adjustment is made. The default value equals the Minimum-Threshold for the Adjustment-Threshold-Percentage.

If the percentage difference between the current bandwidth reservation and the current MaxAvgBw is greater than or equal to the threshold percentage and the difference in the bandwidth is at or above the Minimum-Threshold, the LSP bandwidth is adjusted to the current bandwidth demand (MaxAvgBw).

In the case in which an invalid value is present, the sub-TLV MUST be ignored and the previous value will be maintained.

5.2.4. Minimum and Maximum-Bandwidth Values

5.2.4.1. Minimum-Bandwidth Sub-TLV

The Minimum-Bandwidth sub-TLV specifies 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.

```
 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=8              |             Length=4          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                         Minimum-Bandwidth                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 11: Minimum-Bandwidth Sub-TLV Format

The Type is 8, the Length is 4 octets, and the value comprises the following:

Minimum-Bandwidth: The 4-octet bandwidth value encoded in IEEE floating point format (see [IEEE.754.1985]) and expressed in bytes per second. The default Minimum-Bandwidth value is set to 0. Refer to Section 3.1.2 of [RFC3471] for a table of commonly used values.

In the case in which an invalid value is present, the sub-TLV MUST be ignored and the previous value will be maintained.

5.2.4.2. Maximum-Bandwidth Sub-TLV

The Maximum-Bandwidth sub-TLV specifies 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.

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

Figure 12: Maximum-Bandwidth Sub-TLV Format

The Type is 9, the Length is 4 octets, and the value comprises the following:

Maximum-Bandwidth: The 4-octet bandwidth value encoded in IEEE floating point format (see [IEEE.754.1985]) and expressed in bytes per second. The default Maximum-Bandwidth value is not set. Refer to Section 3.1.2 of [RFC3471] for a table of commonly used values.

In the case in which an invalid value is present, the sub-TLV MUST be ignored and the previous value will be maintained.

5.2.5. Overflow and Underflow Conditions

The sub-TLVs in this section are encoded to inform the PCEP peer of the overflow and underflow threshold parameters. An implementation MAY include sub-TLVs for an absolute value and/or a percentage for the threshold, in which case the bandwidth is immediately adjusted when either of the threshold conditions is met consecutively for the given count (as long as the difference in the bandwidth is at or above the Minimum-Threshold). By default, the threshold values for overflow and underflow conditions are not set.

5.2.5.1. Overflow-Threshold Sub-TLV

The Overflow-Threshold sub-TLV is used to decide if the LSP 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                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 13: Overflow-Threshold Sub-TLV Format

The Type is 10, the Length is 8 octets, and the value comprises the following:

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

Count: The Overflow-Count value (5 bits), encoded in an 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 is to be immediately adjusted to the current bandwidth demand, bypassing the (up) Adjustment-Interval.

Overflow-Threshold: The absolute Overflow-Threshold bandwidth value, encoded in IEEE floating point format (see [IEEE.754.1985]) and 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 overflow condition is met.

In the case in which an invalid value is present, the sub-TLV MUST be ignored and the previous value will be maintained.

5.2.5.2. Overflow-Threshold-Percentage Sub-TLV

The Overflow-Threshold-Percentage sub-TLV is used to decide if the LSP 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=11            |           Length=8            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Percentage |                Reserved                 | Count   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Minimum-Threshold                        |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 14: Overflow-Threshold-Percentage Sub-TLV Format

The Type is 11, the Length is 8 octets, and the value comprises the following:

Percentage: The Overflow-Threshold value (7 bits), encoded in a percentage (an integer from 1 to 100). The value 0 is considered to be invalid. If the percentage increase of the current MaxAvgBw from the current bandwidth reservation is greater than or equal to the threshold percentage, the overflow condition is met.

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

Count: The Overflow-Count value (5 bits), encoded in an 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 is to be immediately adjusted to the current bandwidth demand, bypassing the (up) Adjustment-Interval.

Minimum-Threshold: The absolute Minimum-Threshold bandwidth value, encoded in IEEE floating point format (see [IEEE.754.1985]) and expressed in bytes per second. The increase of the LSP bandwidth MUST be at or above the Minimum-Threshold before the bandwidth adjustment is made.

In the case in which an invalid value is present, the sub-TLV MUST be ignored and the previous value will be maintained.

5.2.5.3. Underflow-Threshold Sub-TLV

The Underflow-Threshold sub-TLV is used to decide if the LSP 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=12            |           Length=8            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Reserved                     | Count   |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                      Underflow-Threshold                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 15: Underflow-Threshold Sub-TLV Format

The Type is 12, the Length is 8 octets, and the value comprises the following:

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

Count: The Underflow-Count value (5 bits), encoded in an 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 is to be immediately adjusted to the current bandwidth demand, bypassing the Down-Adjustment-Interval.

Underflow-Threshold: The absolute Underflow-Threshold bandwidth value, encoded in IEEE floating point format (see [IEEE.754.1985]) and 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 underflow condition is met.

In the case in which an invalid value is present, the sub-TLV MUST be ignored and the previous value will be maintained.

5.2.5.4. Underflow-Threshold-Percentage Sub-TLV

The Underflow-Threshold-Percentage sub-TLV is used to decide if the LSP 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=8            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
| Percentage  |              Reserved                 |  Count  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                        Minimum-Threshold                      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 16: Underflow-Threshold-Percentage Sub-TLV Format

The Type is 13, the Length is 8 octets, and the value comprises the following:

Percentage: The Underflow-Threshold value (7 bits), encoded in percentage (an integer from 1 to 100). The value 0 is considered to be invalid. If the percentage decrease of the current MaxAvgBw from the current bandwidth reservation is greater than or equal to the threshold percentage, the underflow condition is met.

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

Count: The Underflow-Count value (5 bits), encoded in an 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 is to be immediately adjusted to the current bandwidth demand, bypassing the Down-Adjustment-Interval.

Minimum-Threshold: The absolute Minimum-Threshold bandwidth value, encoded in IEEE floating point format (see [IEEE.754.1985]) and expressed in bytes per second. The decrease of the LSP bandwidth MUST be at or above the Minimum-Threshold before the bandwidth adjustment is made.

In the case in which an invalid value is present, the sub-TLV MUST be ignored and the previous value will be maintained.

5.3. BANDWIDTH Object

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

Requested Bandwidth: The BANDWIDTH Object-Type value is 1.

Reoptimization Bandwidth: The bandwidth of an existing TE LSP for which a reoptimization is requested. The BANDWIDTH Object-Type value is 2.

The PCC reports the calculated bandwidth to be adjusted (MaxAvgBw) to the stateful PCE using the existing 'Requested Bandwidth' with the BANDWIDTH Object-Type as 1. The reporting of the 'reoptimization bandwidth' with BANDWIDTH Object-Type as 2 is not required as the stateful PCE is aware of the existing LSP bandwidth.

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 [RFC8281].

For the PCE-initiated LSP with the auto-bandwidth feature enabled, AUTO-BANDWIDTH-ATTRIBUTES TLV MUST be included in the LSPA object with the PCInitiate message.

The Routing Backus-Naur Form (RBNF) definition of the PCInitiate message [RFC8281] 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 [RFC8231].

For PCE-initiated LSPs with the auto-bandwidth feature enabled, the AUTO-BANDWIDTH-ATTRIBUTES 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 RBNF definition of the PCUpd message [RFC8231] is unchanged by this document.

5.6. The PCRpt Message

The PCRpt message [RFC8231] is a PCEP message sent by a PCC to a PCE to report the status of one or more LSPs.

For PCE-initiated LSPs [RFC8281], the PCC creates the LSP using the attributes communicated by the PCE and the local values for the unspecified parameters. After the successful instantiation of the LSP, the 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 as well as after the successful delegation, the BANDWIDTH object of type 1 is used to specify the requested bandwidth in the PCRpt message.

The RBNF definition of the PCRpt message [RFC8231] 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.

A PCEP speaker (PCE or PCC) SHOULD notify its PCEP peer (PCC or PCE) when it is in an overwhelmed state due to the auto-bandwidth feature. An implementation needs to make an attempt to send this notification (when overwhelmed by auto-bandwidth adjustments) unless sending this notification would only serve to increase the load further. Note that when the notification is not received, the PCEP speaker would continue to request bandwidth adjustments even when they cannot be handled in a timely fashion.

Upon receipt of an auto-bandwidth overwhelm notification, the peer SHOULD NOT send any PCEP messages related to auto-bandwidth adjustment. If a PCEP message related to auto-bandwidth adjustment is received while in an overwhelmed state, it MUST be ignored.

* When a PCEP speaker is overwhelmed, it SHOULD notify its peer by sending a PCNtf message with Notification-type = 5 (Auto-Bandwidth Overwhelm State) and Notification-value = 1 (Entering Auto-

Bandwidth Overwhelm State). Optionally, an OVERLOADED-DURATION TLV [RFC5440] MAY be included to specify 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 adjustments, it SHOULD notify its peers by sending a PCNtf message with Notification-type = 5 (Auto-Bandwidth Overwhelm State) and Notification-value = 2 (Clearing Auto-Bandwidth Overwhelm State). A PCEP speaker SHOULD send such notification to all peers if a Notification message (Notification-type = 5, Notification-value = 1) was sent earlier. This message is not sent if an OVERLOADED-DURATION TLV was included and the PCEP speakers wishes for the peer to wait for the expiration of that period of time before receiving further PCEP messages related to auto-bandwidth adjustment.

When the auto-bandwidth feature is deployed, a PCE can send this notification to a PCC when it reports frequent auto-bandwidth adjustments. If a PCC is overwhelmed with resignaling, it can also notify the PCE to not adjust the LSP bandwidth while in the overwhelm state.

Some dampening notification procedure (as per [RFC5440]) to avoid oscillations of the overwhelm state is RECOMMENDED. On receipt of an auto-bandwidth overwhelm notification from the PCE, a PCC should consider the impact on the entire network. Moving the delegations of auto-bandwidth-enabled LSPs to another PCE could cause further overloading.

6. Manageability Considerations

6.1. Control of Function and Policy

The auto-bandwidth feature SHOULD be controlled on a per-LSP basis (at the PCC (head-end of the LSP) or PCE), and the values for auto-bandwidth parameters, e.g., Sample-Interval, Adjustment-Interval (up/down), Minimum-Bandwidth, Maximum-Bandwidth, and Adjustment-Threshold (up/down), SHOULD be configurable by an operator.

The Maximum-Bandwidth (and Minimum-Bandwidth) should be set to an acceptable limit to avoid having an impact on the rest of the MPLS-TE domain.

The operator should make sure that the Overflow-Threshold is greater than or at least equal to the Up-Adjustment-Threshold. And similarly, it is important to ensure that the Underflow-Threshold is greater than or at least equal to the Down-Adjustment-Threshold.

6.2. Information and Data Models

A MIB module for gathering operational information about the PCEP is defined in [RFC7420]. Additionally, the YANG module defined in [PCE-PCEP-YANG] provides both configuration of PCEP as well as operational management. These could be enhanced to provide controls and indicators for support of the auto-bandwidth feature. Support for various configuration knobs as well as counters of messages sent/received containing the TLVs defined in this document could be added.

6.3. Liveness Detection and Monitoring

The mechanisms defined in this document do not imply any new liveness detection and monitoring requirements in addition to those already listed in [RFC5440].

6.4. Verifying Correct Operations

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

In the case in which an invalid value is present, the sub-TLV would get ignored and the previous value will be maintained. In such a case, the implementation SHOULD log the event.

6.5. Requirements for Other Protocols

The mechanisms defined in this document do not add any new requirements for other protocols.

6.6. 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 the auto-bandwidth feature. For each LSP enabled with the auto-bandwidth feature, there is an extra load on the PCC, as it needs to monitor the traffic and report the calculated bandwidth to be adjusted to the PCE. The PCE further recomputes paths based on the requested bandwidth and updates the path to the PCC, which, in turn, triggers the resignaling of the path. All these steps add extra load and churn in the network; thus, the operator needs to take due care while enabling these features on a number of LSPs.

An implementation MAY allow a limit to be placed on the rate of auto-bandwidth-related messages sent by a PCEP speaker and received by a peer. An implementation SHOULD also allow notifications to be sent when a PCEP speaker is overwhelmed or when the rate of messages reaches a threshold.

Due care is required by the operator if a Sample-Interval value significantly smaller than the default (5 minutes) is used, as small Sample-Interval values, e.g., 1 minute or less, could cause undesirable interactions with transport protocols. These undesirable interactions result from providing insufficient time for transport protocol reactions to a prior bandwidth adjustment to settle down before Bandwidth-Samples are taken for the next bandwidth adjustment.

7. Security Considerations

This document defines AUTO-BANDWIDTH-CAPABILITY TLV and AUTO-BANDWIDTH-ATTRIBUTES sub-TLVs, which do not add any substantial new security concerns beyond those already discussed in [RFC8231] and [RFC8281] for stateful PCE operations. As per [RFC8231], it is RECOMMENDED that these PCEP extensions only be activated on authenticated and encrypted sessions across PCEs and PCCs belonging to the same administrative authority, using Transport Layer Security (TLS) [RFC8253], as per the recommendations and best current practices in BCP 195 [RFC7525] (unless explicitly set aside in [RFC8253]).

Incorrect auto-bandwidth parameters in the AUTO-BANDWIDTH-ATTRIBUTES sub-TLVs could have an adverse effect on the LSP as well as on the network.

8. IANA Considerations

8.1. PCEP TLV Type Indicators

This document defines the following new PCEP TLVs; IANA has made the following allocations from the "PCEP TLV Type Indicators" subregistry of the "Path Computation Element Protocol (PCEP) Numbers" registry as follows:

Value	Description	Reference
36	AUTO-BANDWIDTH-CAPABILITY	[RFC8733]
37	AUTO-BANDWIDTH-ATTRIBUTES	[RFC8733]

Table 3: PCEP TLV Type Indicators

8.2. AUTO-BANDWIDTH-CAPABILITY TLV Flag Field

IANA has created a subregistry to manage the Flag field of the AUTO-BANDWIDTH-CAPABILITY TLV within the "Path Computation Element Protocol (PCEP) Numbers" registry.

New bit numbers are to be assigned by Standards Action [RFC8126]. Each bit should be tracked with the following qualities:

* Bit number (counting from bit 0 as the most significant bit)

* Capability description

* Defining RFC

The initial contents of the subregistry are empty, with all bits marked unassigned.

8.3. AUTO-BANDWIDTH-ATTRIBUTES Sub-TLV

This document specifies the AUTO-BANDWIDTH-ATTRIBUTES sub-TLVs. IANA has created an "AUTO-BANDWIDTH-ATTRIBUTES Sub-TLV Types" subregistry within the "Path Computation Element Protocol (PCEP) Numbers" registry to manage the type indicator space for sub-TLVs of the AUTO-BANDWIDTH-ATTRIBUTES TLV. The valid range of values in the registry is 0-65535. IANA has initialized the registry with the following values. All other values in the registry should be marked as "Unassigned".

IANA has set the Registration Procedure for this registry to read as follows:

Range	Registration Procedure
0-65503	IETF Review
65504-65535	Experimental Use

Table 4: Registration Procedure for the "AUTO-BANDWIDTH-ATTRIBUTES Sub-TLV" Registry

This document defines the following types:

Type	Name	Reference
0	Reserved	[RFC8733]
1	Sample-Interval	[RFC8733]
2	Adjustment-Interval	[RFC8733]
3	Down-Adjustment-Interval	[RFC8733]
4	Adjustment-Threshold	[RFC8733]
5	Adjustment-Threshold-Percentage	[RFC8733]
6	Down-Adjustment-Threshold	[RFC8733]
7	Down-Adjustment-Threshold-Percentage	[RFC8733]
8	Minimum-Bandwidth	[RFC8733]
9	Maximum-Bandwidth	[RFC8733]

10	Overflow-Threshold	[RFC8733]
11	Overflow-Threshold-Percentage	[RFC8733]
12	Underflow-Threshold	[RFC8733]
13	Underflow-Threshold-Percentage	[RFC8733]
14-65503	Unassigned	[RFC8733]

Table 5: Initial Contents of the "AUTO-BANDWIDTH-ATTRIBUTES Sub-TLV" Registry

8.4. Error Object

This document defines a new Error-value for PCErr message of Error-Type 19 (Invalid Operation) [RFC8231]. IANA has allocated a new Error-value within the "PCEP-ERROR Object Error Types and Values" subregistry of the "Path Computation Element Protocol (PCEP) Numbers" registry as follows:

Error-Type	Meaning	Error-value	Reference
19	Invalid Operation	14: Auto-Bandwidth capability was not advertised	[RFC8733]

Table 6: Addition to the "PCEP-ERROR Object Error Types and Values" Registry

8.5. Notification Object

IANA has allocated a new Notification-type and Notification-values within the "Notification Object" subregistry of the "Path Computation Element Protocol (PCEP) Numbers" registry as follows:

Notification-type	Name	Notification-value	Reference
5	Auto-Bandwidth Overwhelm State	0: Unassigned	[RFC8733]
		1: Entering Auto-Bandwidth Overwhelm State	[RFC8733]
		2: Clearing Auto-Bandwidth Overwhelm State	[RFC8733]

Table 7: Additions to the "Notification Object" Registry

9. References

9.1. Normative References

[IEEE.754.1985]
IEEE, "Standard for Binary Floating-Point Arithmetic", DOI 10.1109/IEEESTD.1985.82928, IEEE Standard 754, October 1985, <https://doi.org/10.1109/IEEESTD.1985.82928>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation

Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <https://www.rfc-editor.org/info/rfc5440>.

[RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <https://www.rfc-editor.org/info/rfc7525>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://www.rfc-editor.org/info/rfc8126>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[RFC8231] Crabbe, E., Minei, I., Medved, J., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE", RFC 8231, DOI 10.17487/RFC8231, September 2017, <https://www.rfc-editor.org/info/rfc8231>.

[RFC8253] Lopez, D., Gonzalez de Dios, O., Wu, Q., and D. Dhody, "PCEPS: Usage of TLS to Provide a Secure Transport for the Path Computation Element Communication Protocol (PCEP)", RFC 8253, DOI 10.17487/RFC8253, October 2017, <https://www.rfc-editor.org/info/rfc8253>.

[RFC8281] Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for PCE-Initiated LSP Setup in a Stateful PCE Model", RFC 8281, DOI 10.17487/RFC8281, December 2017, <https://www.rfc-editor.org/info/rfc8281>.

9.2. Informative References

[PCE-PCEP-YANG] Dhody, D., Hardwick, J., Beeram, V., and J. Tantsura, "A YANG Data Model for Path Computation Element Communications Protocol (PCEP)", Work in Progress, Internet-Draft, draft-ietf-pce-pcep-yang-13, 31 October 2019, <https://tools.ietf.org/html/draft-ietf-pce-pcep-yang-13>.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <https://www.rfc-editor.org/info/rfc3209>.

[RFC3471] Berger, L., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, DOI 10.17487/RFC3471, January 2003, <https://www.rfc-editor.org/info/rfc3471>.

[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, DOI 10.17487/RFC7420, December 2014, <https://www.rfc-editor.org/info/rfc7420>.

[RFC8051] Zhang, X., Ed. and I. Minei, Ed., "Applicability of a Stateful Path Computation Element (PCE)", RFC 8051, DOI 10.17487/RFC8051, January 2017, <https://www.rfc-editor.org/info/rfc8051>.

Acknowledgments

The authors would like to thank Robert Varga, Venugopal Reddy, Reeja

Paul, Sandeep Boina, Avantika, JP Vasseur, Himanshu Shah, Jonathan Hardwick, and Adrian Farrel for their useful comments and suggestions.

Thanks to Daniel Franke, Joe Clarke, David Black, and Erik Kline for the directorate reviews.

Thanks to Mirja KÃ¼hlewind, Barry Leiba, Benjamin Kaduk, and Roman Danyliw for the IESG review.

Contributors

He Zekun
Tencent Holdings Ltd.
Shenzhen
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
Samsung

Email: younglee.tx@gmail.com

Authors' Addresses

Dhruv Dhody (editor)
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore 560066
Karnataka
India

Email: dhruv.ietf@gmail.com

Rakesh Gandhi (editor)
Cisco Systems, Inc.
Canada

Email: rgandhi@cisco.com

Udayasree Palle
Individual Contributor

Email: udayasreereddy@gmail.com

Ravi Singh
Individual Contributor

Email: ravi.singh.ietf@gmail.com

Luyuan Fang
Expedia Group, Inc.

United States of America

Email: luyuanf@gmail.com