

PCE Working Group
Internet-Draft
Updates: 8733 (if approved)
Intended status: Standards Track
Expires: 4 November 2026

S. Peng
Huawei Technologies
D. Dhody
Huawei
R. Gandhi
Cisco Systems, Inc.
3 May 2026

Update to Automatic Bandwidth Adjustment Procedure of Stateful PCE for
MPLS-TE and SR-TE LSPs
draft-ietf-pce-stateful-pce-autobw-update-04

Abstract

The Stateful PCE extensions allow Stateful control of Traffic Engineering (TE) LSPs using PCEP for RSVP-TE and Segment Routing (SR) (for both MPLS and IPv6 Data planes) for both PCE-Initiated and PCC-Initiated LSPs.

Extensions to the Path Computation Element Communication Protocol (PCEP) for MPLS-TE Label Switched Path (LSP) Automatic Bandwidth Adjustments with Stateful PCE are defined in RFC 8733. It defines the AUTO-BANDWIDTH-ATTRIBUTES TLV and a set of sub-TLVs for each of the attributes. The sub-TLVs are included if there is a change since the last information sent in the PCEP message. However, it lacks a mechanism to remove an attribute identified by the sub-TLV explicitly.

This document updates RFC 8733 by defining the behaviour to remove an attribute explicitly. In addition, this updates allow applying this PCEP extensions to SR-TE LSPs (for both MPLS and IPv6 Data planes), in addition to MPLS-TE 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 <https://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."

This Internet-Draft will expire on 4 November 2026.

Copyright Notice

Copyright (c) 2026 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Conventions	5
3. Updated Procedures	5
4. PCEP Extensions	5
5. Examples	6
5.1. Example 1	6
5.2. Example 2	7
5.3. Example 3	7
6. Auto-Bandwidth for Segment Routing LSPs	7
7. Backward Compatibility	7
8. Security Considerations	8
9. IANA Considerations	8
9.1. AUTO-BANDWIDTH-CAPABILITY TLV Flag Field	8
10. References	8
10.1. Normative References	8
10.2. Informative References	9
Appendix A. Acknowledgments	9
Authors' Addresses	10

1. Introduction

[RFC5440] describes the Path Computation Element Communication Protocol (PCEP). PCEP defines the communication between a Path Computation Client (PCC) and a Path Computation Element (PCE), or between PCEs, enabling computation of Multiprotocol Label Switching (MPLS) for Traffic Engineering Label Switched Path (TE LSP) characteristics.

[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.

PCEP Extensions for Segment Routing (SR) [RFC8664] specifies extensions to PCEP that allow a stateful PCE to compute and initiate TE paths, as well as a PCC to request a path subject to certain constraint(s) and optimization criteria for Segment Routing. [RFC9603] extends PCEP for Segment Routing for IPv6 data plane. As specified in [RFC8664], an LSP can be MPLS-TE LSP or SR-TE LSP based on the path setup type. As specified in [RFC9603], the term "LSP" used in the PCEP specifications for SRv6 would be equivalent to an SRv6 path (represented as a list of SRv6 segments).

[RFC8733] describes the auto-bandwidth feature that allows automatic and dynamic adjustment of the TE LSP bandwidth reservation based on the volume of traffic flowing through the LSP. It describes PCEP extensions for auto-bandwidth adjustment when employing an active stateful PCE for both PCE-initiated [RFC8281] and PCC-initiated LSPs. It defines the AUTO-BANDWIDTH-ATTRIBUTES TLV that provides the 'configurable knobs' of the feature, and it can be included as an optional TLV in the LSPA object. 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. The TLV includes multiple AUTO-BANDWIDTH-ATTRIBUTES sub-TLVs defined in [RFC8733]. The AUTO-BANDWIDTH-ATTRIBUTES sub-TLVs are included if there is a change since the last information was sent in the PCEP message. It also states that in the case of a missing sub-TLV, as per the local policy, either the default value or some other operator-configured value is used.

Since the missing sub-TLV in a subsequent PCEP message is considered to indicate no change, there is no mechanism to remove a particular attribute encoded in the sub-TLV. This document updates [RFC8733] to define such a procedure.

Note that for the attributes that have an associated default value, they could simply encode the default value in the sub-TLV, but this cannot be used for the attributes that do not have a default value.

This document proposes to use a special value of all zeros to indicate "restore to default", which could mean going back to the default values or removal of the attribute itself.

The following table includes the sub-TLVs and the default values as per [RFC8733].

Type	Len	Name	Default
1	4	Sample-Interval	300 seconds
2	4	Adjustment-Interval	86400 seconds
3	4	Down-Adjustment-Interval	Adjustment-Interval
4	4	Adjustment-Threshold	none
5	8	Adjustment-Threshold-Percentage	5%, 0
6	4	Down-Adjustment-Threshold	Adjustment-Threshold
7	8	Down-Adjustment-Threshold-Percentage	Adjustment-Threshold-Percentage
8	4	Minimum-Bandwidth	0
9	4	Maximum-Bandwidth	none
10	8	Overflow-Threshold	none
11	8	Overflow-Threshold-Percentage	none
12	8	Underflow-Threshold	none
13	8	Underflow-Threshold-Percentage	none

Table 1

Thus, the use of the special value of all zeros in the value portion of the sub-TLV can be used to indicate "restore to default", which could mean :

- * if an explicit default value is set for the sub-TLV:
 - Restore to the default values
- * if the default value is set to another sub-TLV value:
 - Remove the associated attribute

- * if there is no default value for the sub-TLV:
 - Remove the associated attribute

The value portion of the sub-TLV consists of encoded data (of the specified length and type), which is set to zero. In cases where the value portion of the sub-TLV contains multiple fields, all fields are set to zero.

2. Conventions

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.

3. Updated Procedures

Section 5.2 of [RFC8733] defines the AUTO-BANDWIDTH-ATTRIBUTES TLV and its associated sub-TLVs.

This document updates [RFC8733] by adding this text at the end of paragraph 3 in section 5.2:

A special value of all zeros in the value portion of the sub-TLV indicates that the attribute identified by the sub-TLV is restored to the default value. The value of all zeros is not considered an invalid value and MUST be checked before individual fields.

For the attributes that have an associated default value, on receiving such a sub-TLV, the PCEP speaker MUST consider it as an instruction to restore to the default values. Note that, the PCEP speaker could also set the default value in the sub-TLV itself.

For the attributes that do not have an associated default value, on receiving such a sub-TLV, the PCEP speaker MUST consider it as a removal of the specific auto-bandwidth attribute.

4. PCEP Extensions

Section 5.1.1 of [RFC8733] defines the AUTO-BANDWIDTH-CAPABILITY TLV as an optional TLV for use in the OPEN Object for auto-bandwidth adjustment. This document adds a new flag -

Z (TBD): The flag indicates that a PCEP speaker supports the use of the special value of all zeros in the value field as specified in this document.

The presence of the Z flag can give a clear indication to the PCEP peer whether they can use the updated procedures defined in this document.

A PCEP speaker MUST NOT use the all-zero value semantics defined in this document unless the peer has advertised the Z flag in the AUTO-BANDWIDTH-CAPABILITY TLV.

5. Examples

The examples in this section assume that both PCEP speakers have advertised support for the Z flag.

5.1. Example 1

Consider an LSP with the following information in the AUTO-BANDWIDTH-ATTRIBUTES TLV in the PCInitiate message:

- * Sample-Interval: 600 (in sec)
- * Adjustment-Interval: 172800 (2 days in sec)
- * Adjustment-Threshold: 0x49989680 (10 Mbps in bps)

Now, if the PCE would like to not use the Adjustment-Thresholds feature for the LSP and set the Adjustment-Interval to 1 day, it could send the AUTO-BANDWIDTH-ATTRIBUTES TLV in the PCUpd message with the following sub-TLVs:

- * Adjustment-Interval: 86400 (1 day in seconds, the default value)
- * Adjustment-Threshold: 0x0

On receiving the special value of all zeros in the value portion of the Adjustment-Threshold sub-TLV, the PCEP speaker would consider that to be the removal of the Adjustment-Threshold feature.

Note that, the PCE could also set the Adjustment-Interval: 0x0 instead of the default value to trigger the restore to default. The Sample-Interval remains unchanged.

5.2. Example 2

Consider an LSP with the following information in the AUTO-BANDWIDTH-ATTRIBUTES TLV in the PCInitiate message:

- * Sample-Interval = 1000

Now, if the PCC receives an update with Sample-Interval with the special value of all zeros, this will lead to Sample-Interval being set to the default value of 300.

5.3. Example 3

Consider an LSP with the following information in the AUTO-BANDWIDTH-ATTRIBUTES TLV in the PCInitiate message:

- * Adjustment-Threshold: 0x49989680 (10 Mbps in bps)

- * Down-Adjustment-Threshold: 0x93312D00 (20 Mbps in bps)

Now, if the PCC receives an update with Down-Adjustment-Threshold with the special value of all zeros, this will lead to the removal of the Down-Adjustment-Threshold attribute and only Adjustment-Threshold remains.

If the PCC receives an update with Adjustment-Threshold with the special value of all zeros, this will lead to the removal of the Adjustment-Threshold attribute as well.

6. Auto-Bandwidth for Segment Routing LSPs

The PCEP extensions defined in [RFC8733] for the MPLS-TE LSP path setup type, apply equally to the SR-TE LSP path setup types, for both SR-MPLS [RFC8664] and SRv6 [RFC9603] data planes.

7. Backward Compatibility

Note that to achieve the same objective, an [RFC8733] compliant implementation could send a PCEP message without the AUTO-BANDWIDTH-ATTRIBUTES TLV first and then include the AUTO-BANDWIDTH-ATTRIBUTES TLV with the updated sub-TLV. This is the same as "turning it off and on again", but would cause unnecessary path computation churn (compared to targeted removal of the attribute).

An existing implementation of [RFC8733] that does not support this update (where the Z flag is not set) will not recognize or use the special value of all zeros in the sub-TLV. If such a sub-TLV is received, as per [RFC8733], implementations may treat the sub-TLV as malformed and ignore it.

8. Security Considerations

This document does not add any substantial new security concerns beyond those already discussed in [RFC8733].

Incorrect or unauthorized use of the all-zero value semantics could remove or reset auto-bandwidth attributes and affect LSP bandwidth adjustment behavior. Therefore, deployments MUST apply the security mechanisms and access controls described in [RFC8733] and the underlying PCEP security mechanisms.

9. IANA Considerations

9.1. AUTO-BANDWIDTH-CAPABILITY TLV Flag Field

[RFC8733] defines the AUTO-BANDWIDTH-CAPABILITY TLV. IANA created a registry to manage the Flag field of the AUTO-BANDWIDTH-CAPABILITY TLV within the "Path Computation Element Protocol (PCEP) Numbers" registry group. This document requests IANA to allocate a new bit in the AUTO-BANDWIDTH-CAPABILITY TLV Flag Field registry, as follows. IANA is requested to make allocations starting from the least significant bit (31).

+=====+		
Bit	Description	Reference
+=====+		
TBD	Z flag (special value of all zeros)	[This.I-D]
+-----+		

Table 2

10. References

10.1. Normative References

[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/rfc/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/rfc/rfc5440>>.
- [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/rfc/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/rfc/rfc8231>>.
- [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/rfc/rfc8281>>.
- [RFC8733] Dhody, D., Ed., Gandhi, R., Ed., Palle, U., Singh, R., and L. Fang, "Path Computation Element Communication Protocol (PCEP) Extensions for MPLS-TE Label Switched Path (LSP) Auto-Bandwidth Adjustment with Stateful PCE", RFC 8733, DOI 10.17487/RFC8733, February 2020, <<https://www.rfc-editor.org/rfc/rfc8733>>.

10.2. Informative References

- [RFC8664] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing", RFC 8664, DOI 10.17487/RFC8664, December 2019, <<https://www.rfc-editor.org/rfc/rfc8664>>.
- [RFC9603] Li, C., Ed., Kaladharan, P., Sivabalan, S., Koldychev, M., and Y. Zhu, "Path Computation Element Communication Protocol (PCEP) Extensions for IPv6 Segment Routing", RFC 9603, DOI 10.17487/RFC9603, July 2024, <<https://www.rfc-editor.org/rfc/rfc9603>>.

Appendix A. Acknowledgments

Thanks to Aijun Wang, Andrew Stone, and Luis Miguel Contreras Murillo for their review comments.

Authors' Addresses

Shuping Peng
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing
100095
China
Email: pengshuping@huawei.com

Dhruv Dhody
Huawei
India
Email: dhruv.ietf@gmail.com

Rakesh Gandhi
Cisco Systems, Inc.
Canada
Email: rgandhi@cisco.com