

Internet Engineering Task Force (IETF)
Request for Comments: 8741
Category: Standards Track
ISSN: 2070-1721

A. Raghuram
A. Goddard
AT&T
J. Karthik
S. Sivabalan
Cisco Systems, Inc.
M. Negi
Huawei Technologies
March 2020

Ability for a Stateful Path Computation Element (PCE) to Request and
Obtain Control of a Label Switched Path (LSP)

Abstract

A stateful Path Computation Element (PCE) retains information about the placement of Multiprotocol Label Switching (MPLS) Traffic Engineering Label Switched Paths (TE LSPs). When a PCE has stateful control over LSPs, it may send indications to LSP head-ends to modify the attributes (especially the paths) of the LSPs. A Path Computation Client (PCC) that has set up LSPs under local configuration may delegate control of those LSPs to a stateful PCE.

There are use cases in which a stateful PCE may wish to obtain control of locally configured LSPs that it is aware of but have not been delegated to the PCE.

This document describes an extension to the Path Computation Element Communication Protocol (PCEP) to enable a PCE to make requests for such control.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8741>.

Copyright Notice

Copyright (c) 2020 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 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
2. Terminology
 - 2.1. Requirements Language
3. LSP Control Request Flag

4. Operation
5. Security Considerations
6. IANA Considerations
7. Manageability Considerations
 - 7.1. Control of Function and Policy
 - 7.2. Information and Data Models
 - 7.3. Liveness Detection and Monitoring
 - 7.4. Verify Correct Operations
 - 7.5. Requirements on Other Protocols
 - 7.6. Impact on Network Operations
8. References
 - 8.1. Normative References
 - 8.2. Informative References

Acknowledgements

Contributors

Authors' Addresses

1. Introduction

"Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE" [RFC8231] specifies a set of extensions to PCEP [RFC5440] to enable stateful control of Traffic Engineering Label Switched Paths (TE LSPs) between and across PCEP sessions in compliance with [RFC4657]. It includes mechanisms to synchronize LSP state between Path Computation Clients (PCCs) and PCEs, delegate control of LSPs to PCEs, and allow PCEs to control the timing and sequence of path computations within and across PCEP sessions. The stateful PCEP defines the following two useful network operations:

Delegation: As per [RFC8051], an operation to grant a PCE temporary rights to modify a subset of LSP parameters on one or more LSPs of a PCC. LSPs are delegated from a PCC to a PCE and are referred to as "delegated" LSPs.

Revocation: As per [RFC8231], an operation performed by a PCC on a previously delegated LSP. Revocation revokes the rights granted to the PCE in the delegation operation.

For redundant stateful PCEs (Section 5.7.4 of [RFC8231]), during a PCE failure, one of the redundant PCEs might want to request to take control over an LSP. The redundant PCEs may use a local policy or a proprietary election mechanism to decide which PCE would take control. In this case, a mechanism is needed for a stateful PCE to request control of one or more LSPs from a PCC so that a newly elected primary PCE can request to take over control.

In case of virtualized PCEs (vPCEs) running in virtual network function (VNF) mode, as the computation load in the network increases, a new instance of vPCE could be instantiated to balance the current load. The PCEs could use a proprietary algorithm to decide which LSPs can be assigned to the new vPCE. Thus, having a mechanism for the PCE to request control of some LSPs is needed.

In some deployments, the operator would like to use stateful PCE for global optimization algorithms but would still like to keep the control of the LSP at the PCC. In such cases, a stateful PCE could request to take control during the global optimization and return the delegation once done.

Note that [RFC8231] specifies a mechanism for a PCC to delegate an orphaned LSP to another PCE. The mechanism defined in this document can be used in conjunction with [RFC8231]. Ultimately, it is the PCC that decides which PCE to delegate the orphaned LSP to.

This specification provides a simple extension that allows a PCE to request control of one or more LSPs from any PCC over the stateful PCEP session. The procedures for granting and relinquishing control of the LSPs are specified in accordance with [RFC8231] unless explicitly set aside in this document.

2. Terminology

This document uses the following terms defined in [RFC5440]:

PCC: Path Computation Client

PCE: Path Computation Element

PCEP: Path Computation Element communication Protocol

This document uses the following terms defined in [RFC8231]:

PCRpt: Path Computation State Report message

PCUpd: Path Computation Update Request message

PLSP-ID: A PCEP-specific identifier for the LSP

SRP: Stateful PCE Request Parameters

Readers of this document are expected to have some familiarity with [RFC8231].

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.

3. LSP Control Request Flag

The Stateful PCE Request Parameters (SRP) object is defined in Section 7.2 of [RFC8231] and includes a Flags field.

A new "LSP Control Request" flag (30), also called the C flag, is introduced in the SRP object. In a PCUpd message, a PCE sets the C flag to 1 to indicate that it wishes to gain control of LSPs. The LSPs are identified by the PLSP-ID in the LSP object following the SRP object. A PLSP-ID value other than 0 and 0xFFFFFFFF is used to identify the LSP for which the PCE requests control. A PLSP-ID value of 0 indicates that the PCE is requesting control of all LSPs originating from the PCC that it wishes to delegate. The C flag has no meaning in other PCEP messages that carry SRP objects and for which the C flag MUST be set to 0 on transmission and MUST be ignored on receipt.

The C flag is ignored in case the R flag [RFC8281] in the SRP object is set.

4. Operation

During normal operation, a PCC that wishes to delegate the control of an LSP sets the Delegate (D) flag (Section 7.3 of [RFC8231]) to 1 in all PCRpt messages pertaining to the LSP. The PCE confirms the delegation by setting the D flag to 1 in all PCUpd messages pertaining to the LSP. The PCC revokes the control of the LSP from the PCE by setting the D flag to 0 in PCRpt messages pertaining to the LSP. If the PCE wishes to relinquish the control of the LSP, it sets the D flag to 0 in all PCUpd messages pertaining to the LSP.

If a PCE wishes to gain control over an LSP, it sends a PCUpd message with the C flag set to 1 in the SRP object. The LSP for which the PCE requests control is identified by the PLSP-ID in the associated LSP object. A PLSP-ID value of 0 indicates that the PCE wants control over all LSPs originating from the PCC. An implementation of this feature needs to make sure to check for the LSP control feature (C flag set to 1) before any check for PLSP-ID (as per [RFC8231]). The D flag and C flag are mutually exclusive in a PCUpd message. The PCE MUST NOT send a control request for the LSP that is already delegated to the PCE, i.e., if the D flag is set in the PCUpd

message, then the C flag MUST NOT be set. If a PCC receives a PCUpd message with the D flag set in the LSP object (i.e., LSP is already delegated) and the C flag is also set (i.e., PCE is making a control request), the PCC MUST ignore the C flag. A PCC can decide to delegate the control of the LSP at its own discretion. If the PCC grants or denies the control, it sends a PCRpt message with the D flag set to 1 and 0, respectively, in accordance with stateful PCEP [RFC8231]. If the PCC does not grant the control, it MAY choose to not respond, and the PCE MAY choose to retry requesting the control, preferably using an exponentially increasing timer. Note that, if the PCUpd message with the C flag set is received for a currently non-delegated LSP (for which the PCE is requesting delegation), this MUST NOT trigger the error handling as specified in [RFC8231] (a PCErr with Error-type=19 (Invalid Operation) and error-value 1 (Attempted LSP Update Request for a non-delegated LSP)).

As per [RFC8231], a PCC cannot delegate an LSP to more than one PCE at any time. If a PCE requests control of an LSP that has already been delegated by the PCC to another PCE, the PCC MAY ignore the request or MAY revoke the delegation to the first PCE before delegating it to the second. This choice is a matter of local policy.

It should be noted that a legacy implementation of PCC that does not support this extension may receive an LSP control request: a PCUpd message with the C flag set and the D flag unset. The legacy implementation would ignore the C flag and trigger the error condition for the D flag, as specified in [RFC8231] (i.e., a PCErr with Error-type=19 (Invalid Operation) and error-value 1 (Attempted LSP Update Request for a non-delegated LSP)). Further, in case of a PLSP-ID value of 0, the error condition, as specified in [RFC8231], (i.e., a PCErr with Error-type=19 (Invalid Operation) and error-value 3 (Attempted LSP Update Request for an LSP identified by an unknown PLSP-ID)) would be triggered.

[RFC8281] describes the setup, maintenance, and teardown of PCE-initiated LSPs under the stateful PCE model. It also specifies how a PCE may obtain control over an orphaned LSP that was PCE-initiated. A PCE implementation can apply the mechanism described in this document in conjunction with those in [RFC8281].

5. Security Considerations

The security considerations listed in [RFC8231] and [RFC8281] apply to this document as well. However, this document also introduces a new attack vector. An attacker may flood the PCC with requests to delegate all of its LSPs at a rate that exceeds the PCC's ability to process them, either by spoofing messages or by compromising the PCE itself. The PCC SHOULD be configured with a threshold rate for the delegation requests received from the PCE. If the threshold is reached, it is RECOMMENDED to log the issue.

A PCC is the ultimate arbiter of delegation. As per [RFC8231], a local policy at the PCC is used to influence the delegation. A PCC can also revoke the delegation at any time. A PCC need not blindly trust the control requests and SHOULD take local policy and other factors into consideration before honoring the request.

Note that a PCE may not be sure if a PCC supports this feature. A PCE would try sending a control request to a 'legacy' PCC that would in turn respond with an error, as described in Section 4. So, a PCE would learn this fact only when it wants to take control over an LSP. A PCE might also be susceptible to downgrade attacks by falsifying the error condition.

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 excluded in [RFC8253]).

6. IANA Considerations

IANA has allocated the following code point in the "SRP Object Flag Field" subregistry in the "Path Computation Element Protocol (PCEP) Numbers" registry.

Bit	Description	Reference
30	LSP Control Request	RFC 8741

Table 1

7. Manageability Considerations

All manageability requirements and considerations listed in [RFC5440] and [RFC8231] apply to PCEP extensions defined in this document. In addition, requirements and considerations listed in this section apply.

7.1. Control of Function and Policy

A PCC implementation SHOULD allow the operator to configure the policy rules that specify the conditions under which it honors the request to control the LSPs. This includes the handling of the case where an LSP control request is received for an LSP that is currently delegated to some other PCE. A PCC implementation SHOULD also allow the operator to configure the threshold rate for the delegation requests received from the PCE. Further, the operator MAY be allowed to trigger the LSP control request for a particular LSP at the PCE. A PCE implementation SHOULD also allow the operator to configure an exponentially increasing timer to retry the control requests for which the PCE did not get a response.

7.2. Information and Data Models

The PCEP YANG module [PCEP-YANG] could be extended to include a mechanism to trigger the LSP control request.

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] and [RFC8231].

7.5. Requirements on Other Protocols

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

7.6. Impact on Network Operations

Mechanisms defined in [RFC5440] and [RFC8231] also apply to PCEP extensions defined in this document. Further, the mechanism described in this document can help the operator to request control of the LSPs at a particular PCE.

8. References

8.1. 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/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>>.
- [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>>.
- [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>>.

8.2. Informative References

- [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>>.
- [RFC4657] Ash, J., Ed. and J.L. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol Generic Requirements", RFC 4657, DOI 10.17487/RFC4657, September 2006, <<https://www.rfc-editor.org/info/rfc4657>>.
- [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>>.
- [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>>.
- [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>>.

Acknowledgements

Thanks to Jonathan Hardwick for reminding the authors to not use suggested values in IANA section.

Thanks to Adrian Farrel, Haomian Zheng, and Tomonori Takeda for their valuable comments.

Thanks to Shawn M. Emery for his Security Directorate review.

Thanks to Francesca Palombini for GENART review.

Thanks to Benjamin Kaduk, Martin Vigoureux, Alvaro Retana, and Barry Leiba for IESG reviews.

Contributors

The following people contributed substantially to the content of this document and should be considered coauthors:

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

Email: dhruv.ietf@gmail.com

Jon Parker
Cisco Systems, Inc.
2000 Innovation Drive
Kanata Ontario K2K 3E8
Canada

Email: jdparker@cisco.com

Chaitanya Yadlapalli
AT&T
200 S Laurel Avenue
Middletown, NJ 07748
United States of America

Email: cy098@att.com

Authors' Addresses

Aswatnarayan Raghuram
AT&T
200 S Laurel Avenue
Middletown, NJ 07748
United States of America

Email: ar2521@att.com

Al Goddard
AT&T
200 S Laurel Avenue
Middletown, NJ 07748
United States of America

Email: ag6941@att.com

Jay Karthik
Cisco Systems, Inc.
125 High Street
Boston, Massachusetts 02110
United States of America

Email: jakarthi@cisco.com

Siva Sivabalan
Cisco Systems, Inc.
2000 Innovation Drive
Kanata Ontario K2K 3E8
Canada

Email: msiva@cisco.com

Mahendra Singh Negi
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore 560066
Karnataka
India

Email: mahend.ietf@gmail.com