PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 30, 2018

C. Barth
R. Torvi
Juniper Networks

June 28, 2018

PCEP Extensions for RSVP-TE Local-Protection with PCE-Stateful draft-cbrt-pce-stateful-local-protection-01

Abstract

Stateful PCE [RFC8231] can apply global concurrent optimizations to optimize LSP placement. In a deployment where a PCE is used to compute all the paths, it may be beneficial for the local protection paths to also be computed by the PCE. This document defines extensions needed for the setup and management of RSVP-TE protection paths by the PCE.

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 December 30, 2018.

Copyright Notice

Table of Contents

1. Introduction 2
2. Terminology 3
3. Architectural Overview 3
3.1. Local Protection Overview 3
4. Extensions for the LSPA object 4
4.1. The Preference TLV 4
4.2. The Bypass TLV 5
4.3. The LOCALLY-PROTECTED-LSPS TLV 6
5. IANA considerations 8
5.1. PCEP-Error Object 8
5.2. PCEP TLV Type Indicators 8
6. Security Considerations 8
7. Contributors 8
8. References 9
8.1. Normative References 9
8.2. Informative References 9
Appendix A. Additional Stuff 10
Authors' Addresses 10

1. Introduction

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

Stateful PCE [RFC8231] specifies a set of extensions to PCEP to enable stateful control of paths such as MPLS TE LSPs between and across PCEP sessions in compliance with [RFC4657]. It includes mechanisms to effect LSP state synchronization between PCCs and PCEs and allow delegation of control of LSPs to PCEs.

In a network where all LSPs have control delegated to a PCE, the PCE can apply global concurrent optimization to optimize LSP placement. The PCE can also control the timing and sequence of path computation and applying path changes. In a deployment where a PCE is used to compute all the paths, it may be beneficial for the protection paths

to also be controlled through the PCE. This document defines extensions needed for the setup and management of protection paths by the PCE.

Benefits of stateful synchronization and control of the protection paths include:

o Better control over traffic after a failure and more deterministic path computation of protection paths. The PCE can optimize the protection path based on data not available to the PCC, for instance the PCE can make sure the protection path will not violate the delay specified by [I-D.ietf-pce-pcep-service-aware].

o Satisfy more complex constraints and diversity requirements, such as maintaining diverse paths for LSPs as well as their local protection paths.

o Given the PCE's global view of network resources, act as a form of LSP admission control into a protection path to ensure links are not overloaded during failure events.

o On a PLR with multiple available protection routes, allows the PCE to map LSPs to all available protection routes versus a single best protection route.

o Most of the benefits stated in the stateful PCE applicability draft [I-D.ietf-pce-stateful-pce-app-04] apply equally to protection paths.

2. Terminology

This document uses the following terms defined in [RFC5440] PCC PCE, PCEP Peer.

This document uses the following terms defined in [RFC8231] Stateful PCE, Delegation, Delegation Timeout Interval, LSP State Report, LSP Update Request.

The message formats in this document are specified using Routing Backus-Naur Format (RBNF) encoding as specified in RFC5511.

3. Architectural Overview

3.1. Local Protection Overview

Local protection refers to the ability to locally route around failure of an LSP. Two types of local protection are possible:

(1) 1:1 protection - the protection path protects a single LSP.

(2) N:1 protection - the protection path protects multiple LSPs traversing the protected resource.

It is assumed that the PCE knows what resources require protection through mechanisms outside the scope of this document. In a PCE controlled deployment, support of 1:1 protection has limited applicability, and can be achieved as a degenerate case of 1:N protection. For this reason, local protection will be discussed only for the N:1 case.

Local protection requires the setup of a bypass at the PLR. This bypass can be PCC-initiated and delegated, or PCE-initiated. In either case, the PLR MUST maintain a PCEP session to the PCE. A bypass identifier (the name of the bypass) is required for disambiguation as multiple bypasses are possible at the PLR. There two types Bypass LSPs mappings:

(1) Independent Bypass LSP Mapping: In this case Bypass LSP mapping is handled by a local policy on PCC and the PCC reports all mappings to the PCE. In other words, bypass LSP(s) are mapped to any protected LSP(s) that satisfy PCC local policy.

(2) Dependent Bypass LSP mapping: Mapping of LSPs to bypass is done through a new TLV, the LOCALLY-PROTECTED-LSPS TLV in the LSP Update message from PCE to PLR. See section Section 4.3. When an LSP requiring protection is set up through the PLR, the PLR checks if it has a mapping to a bypass and only provides protection if such a mapping exists. The status of bypasses and what LSPs are protected by them is communicated to the PCE via LSP Status Report messages.

4. Extensions for the LSPA object

4.1. The Preference TLV

When provisioning a PCC, the PCE can influence primary to bypass LSP association of the PCC using the preference TLV. Bypass LSPs with a higher preference are used first during primary LSP association. Bypass LSPs with identical preferences are used for primary association according to local PCC selection.

The format of the IPv4 Preference TLV is shown in the following figure:

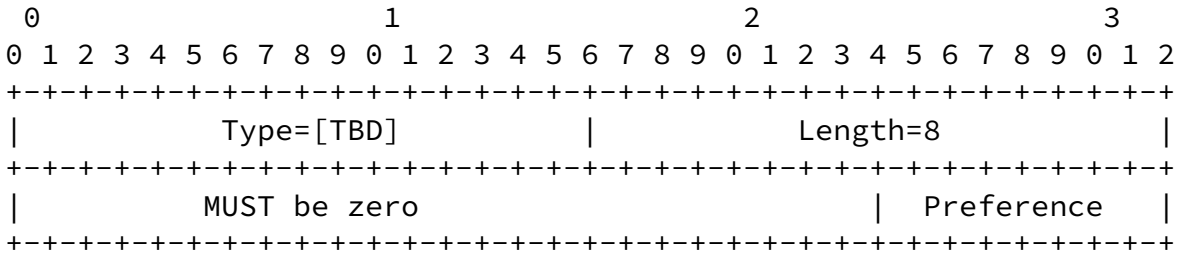

Figure 1: IPv4 Preference TLV format

The type of the TLV is [TBD] and it has a fixed length of 8 octets. The value contains the following fields:

Preference (8 bits): The value indicates the bypass LSP preference during the primary LSP selection process of the PCC. A lower preference value is preferred to a higher value with a default value of 255. A value of 0 would indicate that the bypass is not to be selected for any primary LSP associations.

If the Preference TLV is included, then the LSPA object MUST also carry the SYMBOLIC-PATH-NAME TLV as one of the optional TLVs. Failure to include the mandatory SYMBOLIC-PATH-NAME TLV MUST trigger PCErr of type 6 (Mandatory Object missing) and value TBD (SYMBOLIC-PATH-NAME TLV missing for bypass LSP).

4.2. The Bypass TLV

The facility backup method creates a bypass tunnel to protect a potential failure point. The bypass tunnel protects a set of LSPs with similar backup constraints [RFC4090].

A PCC can delegate a bypass tunnel to PCE control or a PCE can provision the bypass tunnel via a PCC. The procedures for bypass instantiation rely on the extensions defined in [RFC8281] and will be detailed in a future version of this document.

A subscription multiplier can be used to influence the local PCC admission control during primary LSP association. This allows for under subscription or oversubscription policy to be applied to the bandwidth attribute of the bypass LSP.

The Bypass TLV carries information about the bypass tunnel. It is included in the LSPA Object in LSP State Report and LSP Update Request messages.

The format of the IPv4 Bypass TLV is shown in the following figure:

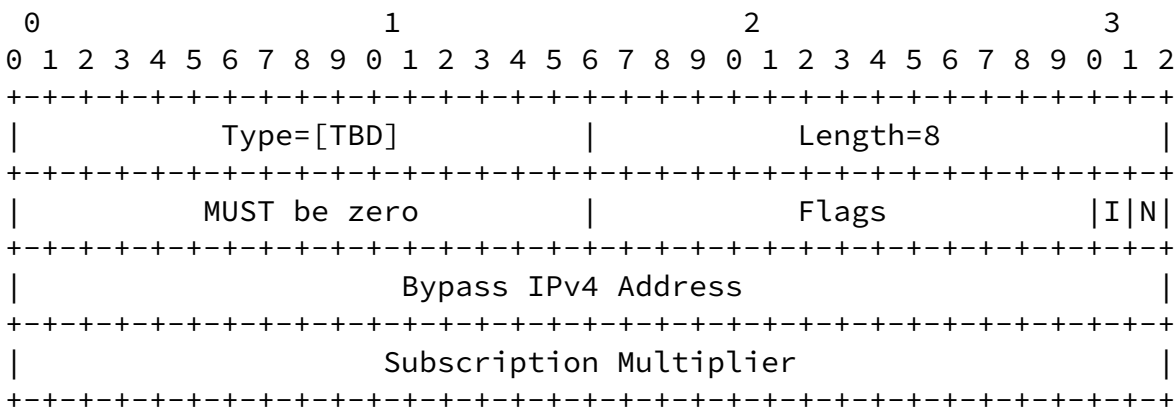

Figure 2: IPv4 Bypass TLV format

The type of the TLV is [TBD] and it has a fixed length of 8 octets.

The value contains the following fields:

Flags (16 bit)

N (Node Protection - 1 bit): The N flag indicates whether the Bypass is used for node-protection. If the N flag is set to 1, the Bypass is used for node-protection. If the N flag is 0, the Bypass is used for link-protection.

I (Local Protection In Use - 1 bit): The I Flag indicates that local repair mechanism is in use.

Bypass IPv4 address: The Bypass IPv4 Address is the next-hop address of the protected link in the paths of the protected LSPs.

Subscription Multiplier (32 bits): An optional multiplier represented as a floating point number. The value may be used to influence CAC during primary LSP association. For example, a bypass may reserved 50M but the PCC may want to admit up to (multiplier * reserved bandwidth) to the bypass LSP.

If the Bypass TLV is included, then the LSPA object MUST also carry the SYMBOLIC-PATH-NAME TLV as one of the optional TLVs. Failure to include the mandatory SYMBOLIC-PATH-NAME TLV MUST trigger PCErr of type 6 (Mandatory Object missing) and value TBD (SYMBOLIC-PATH-NAME TLV missing for bypass LSP)

4.3. The LOCALLY-PROTECTED-LSPS TLV

The IPV4-LOCALLY-PROTECTED-LSPS TLV in the LSPA Object contains a list of LSPs protected by the bypass tunnel.

The format of the Locally protected LSPs TLV is shown in the following figure:

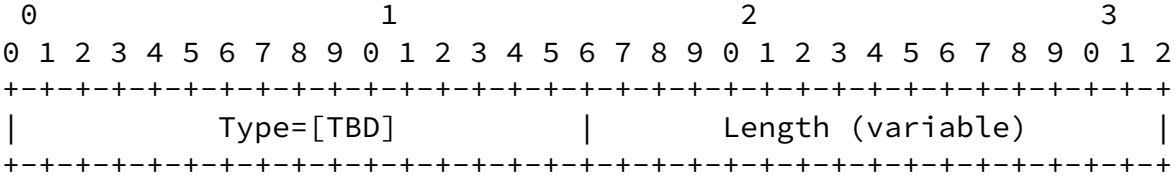

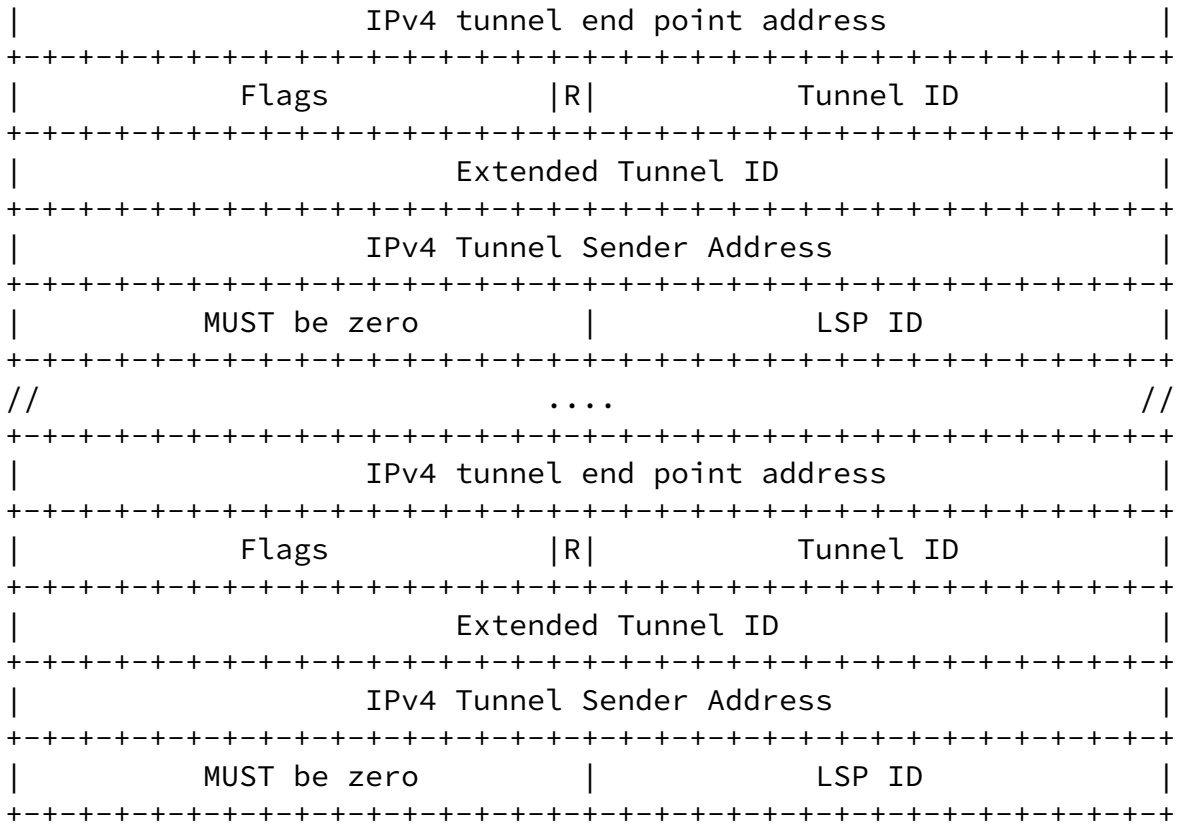

Figure 3: IPv4 Locally protected LSPs TLV format

The type of the TLV is [TBD] and it is of variable length.The value contains one or more LSP descriptors including the following fields filled per [RFC3209]

IPv4 Tunnel end point address: As defined in [RFC3209], Section 4.6.1.1

Flags (16 bit)

R(Remove - 1 bit): The R flag indicates that the LSP has been removed from the list of LSPs protected by the bypass tunnel.

Tunnel ID: As defined in [RFC3209], Section 4.6.1.1

Extended Tunnel ID: As defined in [RFC3209], Section 4.6.2.1

IPv4 Tunnel Sender address: As defined in [RFC3209], Section 4.6.2.1

LSP ID: As defined in RFC 3209

5. IANA considerations

5.1. PCEP-Error Object

This document defines new Error-Type and Error-Value for the following new error conditions:

Error-Type Meaning 6 Mandatory Object missing Error-value=TBD: SYMBOLIC-PATH-NAME TLV missing for a path where the S-bit is set in the LSPA object. Error-value=TBD: SYMBOLIC-PATH-NAME TLV missing for a bypass path.

5.2. PCEP TLV Type Indicators

This document defines the following new PCEP TLVs:

Value #	Meaning	Reference
???	Bypass	This Document
???	Weight	This Document
???	LOCALLY-PROTECTED-LSPS	This Document

Table 1: New PCEP TLVs

6. Security Considerations

The same security considerations apply at the PLR as those describe for the head end in PCE Initiated LSPs [RFC8281].

7. Contributors

The following people have substantially contributed to the editing of this document:

Harish Sitaraman, Juniper Networks, hsitaraman@juniper.net

Vishnu Pavan Beeram, Juniper Networks, vbeeram@juniper.net

Chandrasekar Ramachandran, Juniper Networks, csekar@juniper.net

Ambrose Kwong, Juniper Networks, akwong@juniper.net

Phil Bedard, bedard.phil@gmail.com

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>.

[RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", September 1997.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", December 2001.

[RFC4090] Pan, P., Swallow, G., and A. Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", May 2005.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", May 2008.

[RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", March 2009.

[RFC8231] Crabbe, E., Medved, J., Minie, I., and R. Verga, "PCEP Extensions for Stateful PCE", 2015.

[RFC8281] Crabbe, E., Sivabalan, S., and R. Verga, "PCEP Extensions for PCE-initiated LSP Setup in a Stateful PCE Model", 2014.

8.2. Informative References

[I-D.narten-iana-considerations-rfc2434bis] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", draft-narten-iana-considerations-rfc2434bis-09 (work in progress), March 2008.

[RFC2629] Rose, M., "Writing I-Ds and RFCs using XML", RFC 2629, DOI 10.17487/RFC2629, June 1999, <https://www.rfc-editor.org/info/rfc2629>.

[RFC3552] Rescorla, E. and B. Korver, "Guidelines for Writing RFC Text on Security Considerations", BCP 72, RFC 3552, DOI 10.17487/RFC3552, July 2003, <https://www.rfc-editor.org/info/rfc3552>.

[RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", August 2006.

[RFC4657] Ash, J. and J. Le Roux, "Path Computation Element (PCE) Communication Protocol Generic Requirements", September 2006.

[RFC5394] Bryskin, I., Papadimitriou, D., Berger, L., and J. Ash, "Policy-Enabled Path Computation Framework", December 2008.

[RFC5557] Lee, Y., Le Roux, JL., King, D., and E. Oki, "Path Computation Element Communication Protocol (PCEP) Requirements and Protocol Extensions in Support of Global Concurrent Optimization", July 2009.

Appendix A. Additional Stuff

This becomes an Appendix.

Authors' Addresses

Colby Barth
Juniper Networks
Sunnyvale, CA
USA

Email: cbarth@juniper.net

Raveendra Torvi
Juniper Networks
Sunnyvale, CA
USA

Email: rtorvi@juniper.net