Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 15, 2013

E. Crabbe
Google, Inc.
J. Medved
Cisco Systems, Inc.
I. Minei
R. Torvi
Juniper Networks, Inc.

October 12, 2012

PCEP Extensions for MPLS-TE LSP protection with stateful PCE
draft-crabbe-pce-stateful-pce-protection-00

Abstract

Stateful PCE [I-D.ietf-pce-stateful-pce] 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 protection paths to also be computed by the PCE. This document defines extensions needed for the setup and management of MPLS-TE protection paths by the PCE.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on April 15, 2013.

Copyright Notice

Table of Contents

1. Introduction 3
2. Terminology 3
3. Architectural Overview 3
3.1. Path Protection Overview 3
3.2. Local Protection Overview 4
4. Extensions for the LSPA object 5
4.1. The Standby flag in the LSPA object 5
4.2. The Weight TLV 6
4.3. The Bypass TLV 6
4.4. The LOCALLY-PROTECTED-LSPS TLV 7
5. IANA considerations 9
5.1. PCEP-Error Object 9
5.2. PCEP TLV Type Indicators 9
6. Security Considerations 9
7. Acknowledgements 9
8. References 10
8.1. Normative References 10
8.2. Informative References 11
Authors' Addresses 11

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 Multiprotocol Label Switching (MPLS) for Traffic Engineering Label Switched Path (TE LSP) characteristics.

Stateful pce [I-D.ietf-pce-stateful-pce] 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, delegation of control of LSPs to PCEs, and PCE control of timing and sequence of path computations within and across PCEP sessions and focuses on a model where LSPs are configured on the PCC and control over them is delegated to the PCE.

Stateful PCE 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 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 controlling the protection paths include: better control over traffic after a failure and more deterministic path computation (paths not affected by overload after a failure).

2. Terminology

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

This document uses the following terms defined in [I-D.ietf-pce-stateful-pce]: 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. Path Protection Overview

Path protection refers to switching to a new path on failure. Several cases exist:

(1) MPLS-TE Global Default Restoration - protection paths are computed dynamically by the LSR after the failure. This can be supported without any PCEP protocol changes by specifying a secondary path with an ERO of just the end points of the LSP. Once reestablished, the path is communicated to the PCE via the LSP State Report message.

(2) MPLS-TE Global Path Protection - protection paths are fully specified ahead of the failure. The base Stateful PCE specification [I-D.ietf-pce-stateful-pce] supports sending multiple fully-specified paths in the PCUpd requests.There are 2 further sub-cases:

    (a) Protection paths are pre-signaled ahead of the failure (standby paths).

    (b) Protection paths are set up after the failure.

The protection path setup regimen (standby or not) is specified in the path using a new per-path flag in the LSPA object, the S (standby) flag (see section Section 4.1). Paths for which the S flag is set MUST have a name associated with them, specified using the SYMBOLIC-PATH-NAME TLV in the LSPA object.

Because multiple secondary standby paths are possible, there is also a need for the PCE to be able to specify the relative priorities between the paths (which one to take if there are 3 available). This is done through a weight assigned to each path. See details in Section 4.2.

Reversion from protection paths to the primary path when possible will be controlled by the PCE, by sending a new LSP Update Request. If the primary can be successfully signaled and the secondary does not have the S flag set, then the secondary MUST be torn down. Thus, there is no need to signal the desire for revertive behavior.

3.2. 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) 1:N 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 disccussed only for the 1:N case.

Local protection requires the setup of a bypass at the PLR. This bypass can be locally 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. 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 4.4. 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 Standby flag in the LSPA object

The LSPA object is defined in [RFC5440] and replicated below for easy reference. This document defines a new flag, the S flag in the flags field of the LSPA object.

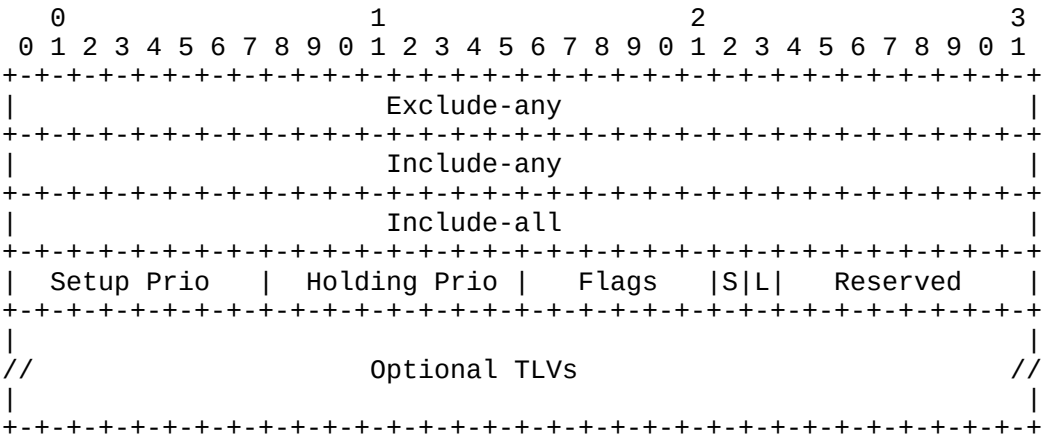

Figure 1: STATEFUL-PCE-CAPABILITY TLV format

The L flag is defined in [RFC5440].

If set to 1, the S Flag indicates this is a standby path.

If the S flag is set, 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 when the S flag is set MUST trigger PCErr of type 6 (Mandatory Object missing) and value TBD (SYMBOLIC-PATH-NAME TLV missing for standby LSP).

4.2. The Weight TLV

This TLV will be discussed in a future version of tihs document.

4.3. 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 [I-D.crabbe-pce-pce-initiated-lsp] and will be detailed in a future version of this document.

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 Bypass TLV is shown in the following figure:

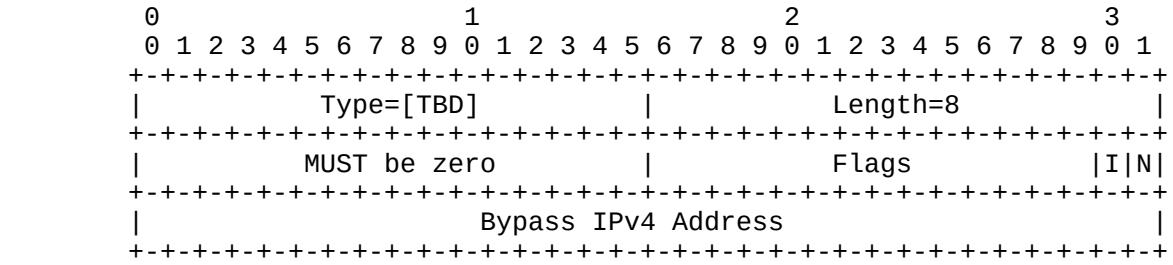

Figure 2: Bypass TLV format

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

Flags

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: For link protection, the Bypass IPv4 Address is the nexthop address of the protected link in the paths of the protected LSPs. For node protection, the Bypass IPv4 Address is the node addresses of the protected node.

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)

4.4. The LOCALLY-PROTECTED-LSPS TLV

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

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

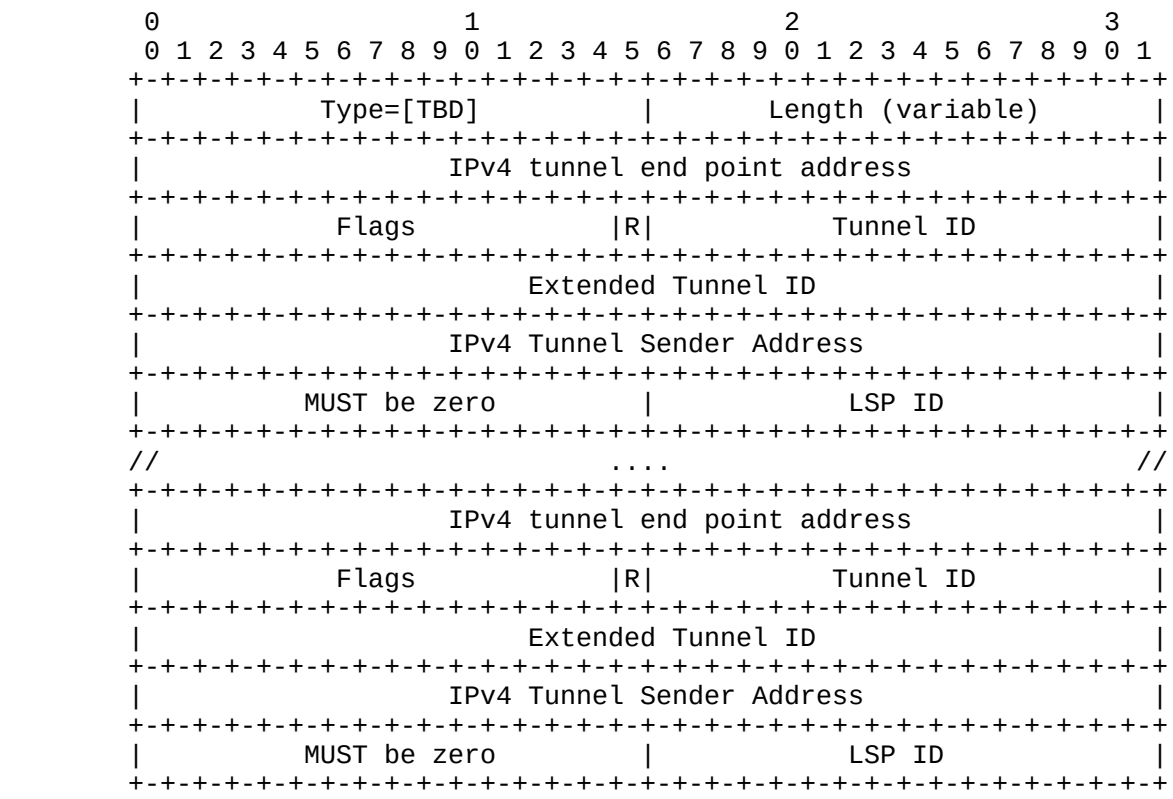

Figure 3: Locally protected LSPs TLV format

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: [RFC3209]

Flags

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: [RFC3209]

Extended Tunnel ID: [RFC3209]

IPv4 Tunnel Sender address: [RFC3209]

LSP ID: [RFC3209]

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

6. Security Considerations

The same security considerations apply at the PLR as those describe for the head end in [I-D.crabbe-pce-pce-initiated-lsp].

7. Acknowledgements

We would like to thank Ambrose Kwong for his contributions to this document.

8. References

8.1. Normative References

[I-D.crabbe-pce-pce-initiated-lsp]
Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "PCEP Extensions for PCE-initiated LSP Setup in a Stateful PCE Model", draft-crabbe-pce-pce-initiated-lsp-00 (work in progress), October 2012.

[I-D.ietf-pce-stateful-pce]
Crabbe, E., Medved, J., Varga, R., and I. Minei, "PCEP Extensions for Stateful PCE", draft-ietf-pce-stateful-pce-01 (work in progress), July 2012.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2205] Braden, B., Zhang, L., Berson, S., Herzog, S., and S. Jamin, "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

[RFC4090] Pan, P., Swallow, G., and A. Atlas, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, May 2005.

[RFC5088] Le Roux, JL., Vasseur, JP., Ikejiri, Y., and R. Zhang, "OSPF Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5088, January 2008.

[RFC5089] Le Roux, JL., Vasseur, JP., Ikejiri, Y., and R. Zhang, "IS-IS Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5089, January 2008.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

[RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009.

[RFC5511] Farrel, A., "Routing Backus-Naur Form (RBNF): A Syntax Used to Form Encoding Rules in Various Routing Protocol Specifications", RFC 5511, April 2009.

8.2. Informative References

[RFC2702] Awduche, D., Malcolm, J., Agogbua, J., O'Dell, M., and J. McManus, "Requirements for Traffic Engineering Over MPLS", RFC 2702, September 1999.

[RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.

[RFC3346] Boyle, J., Gill, V., Hannan, A., Cooper, D., Awduche, D., Christian, B., and W. Lai, "Applicability Statement for Traffic Engineering with MPLS", RFC 3346, August 2002.

[RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003.

[RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.

[RFC4657] Ash, J. and J. Le Roux, "Path Computation Element (PCE) Communication Protocol Generic Requirements", RFC 4657, September 2006.

[RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, October 2008.

[RFC5394] Bryskin, I., Papadimitriou, D., Berger, L., and J. Ash, "Policy-Enabled Path Computation Framework", RFC 5394, 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", RFC 5557, July 2009.

Authors' Addresses

Edward Crabbe
Google, Inc.
1600 Amphitheatre Parkway
Mountain View, CA 94043
US

Email: edc@google.com

Jan Medved
Cisco Systems, Inc.
170 West Tasman Dr.
San Jose, CA 95134
US

Email: jmedved@cisco.com

Ina Minei
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
US

Email: ina@juniper.net

Raveendra Torvi
Juniper Networks, Inc.
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
US

Email: rtorvi@juniper.net