

PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 27, 2017

S. Litkowski
Orange
S. Sivabalan
Cisco Systems, Inc.
C. Barth
Juniper Networks
November 23, 2016

Path Computation Element communication Protocol extension for signaling
LSP diversity constraint
[draft-litkowski-pce-association-diversity-01](#)

Abstract

This document introduces a simple mechanism to signal path diversity for a group of Label Switched Paths (LSPs) via an extension to the Path Computation Element Communication Protocol (PCEP).

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 May 27, 2017.

Copyright Notice

Copyright (c) 2016 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 (<http://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
1.1.	Requirements Language	2
2.	Terminology	3
3.	Motivation	3
4.	Protocol extension	6
4.1.	Association group	6
4.2.	Optional TLVs	6
5.	Security Considerations	8
6.	IANA Considerations	8
6.1.	Association object Type Indicators	8
7.	Manageability Considerations	9
7.1.	Control of Function and Policy	9
7.2.	Information and Data Models	9
7.3.	Liveness Detection and Monitoring	9
7.4.	Verify Correct Operations	9
7.5.	Requirements On Other Protocols	9
7.6.	Impact On Network Operations	9
8.	Acknowledgments	9
9.	References	9
9.1.	Normative References	9
9.2.	Informative References	10
	Authors' Addresses	10

[1.](#) Introduction

[I-D.ietf-pce-association-group] introduces a generic mechanism to create a grouping of LSPs which can then be used to define associations between a set of LSPs and a set of attributes (such as configuration parameters or behaviours) and is equally applicable to the active and passive modes of a stateful PCE [I-D.ietf-pce-stateful-pce] or a stateless PCE [RFC5440].

This document specifies a PCEP extension to signal that a particular group of LSPs should use diverse paths including the requested type of diversity.

[1.1.](#) 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].

2. Terminology

The following terminology is used in this document.

LSR: Label Switch Router.

MPLS: Multiprotocol Label Switching.

PCC: Path Computation Client. Any client application requesting a path computation to be performed by a Path Computation Element.

PCE: Path Computation Element. An entity (component, application, or network node) that is capable of computing a network path or route based on a network graph and applying computational constraints.

PCEP: Path Computation Element Communication Protocol.

SRLG: Shared Risk Link Group.

3. Motivation

Path diversity is a very common use case today in IP/MPLS networks especially for layer 2 transport over MPLS. A customer may request that the operator provide two end-to-end disjoint paths across the IP/MPLS core. The customer may use those paths as primary/backup or active/active.

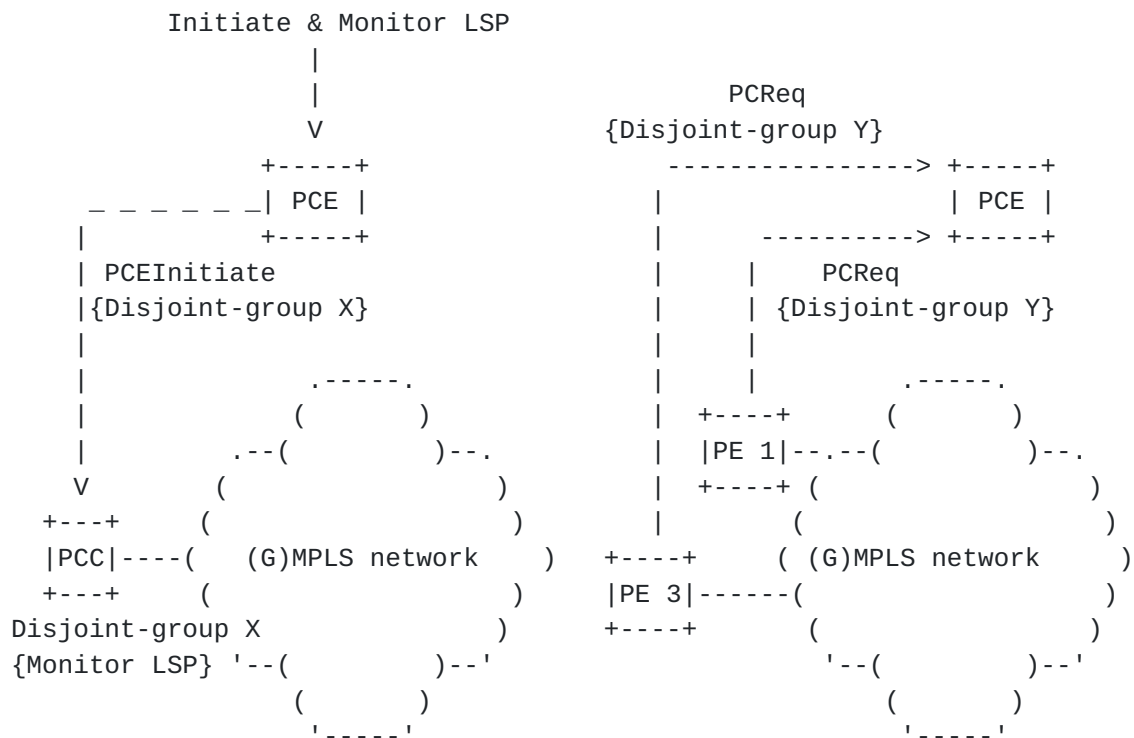
Different level of disjointness may be offered:

- o Link disjointness: the paths of the associated LSPs should transit different links (but may use common nodes or different links that may have some shared fate).
- o Node disjointness: the paths of the associated LSPs should transit different nodes (but may use different links that may have some shared fate).
- o SRLG disjointness: the paths of the associated LSPs should transit different links that do not share fate (but may use common transit nodes).
- o Node+SRLG disjointness: the paths of the associated LSPs should transit different links that do not have any common shared fate and should transit different nodes.

The associated LSPs may originate from the same or from different head-end(s) and may terminate at the same or different tail-end(s).

The management of the disjoint group-ids will be a key point for the operator as the Association ID field is limited to 65535. The local configuration of IPv4/IPv6 association source, or Global Association Source/Extended Association ID should allow to overcome this limitation. For example, when a PCC or PCE initiates all the LSPs in a particular disjoint-group, it can set the IPv4/IPv6 association source can be set to one of its IP address. When disjoint LSPs are initiated from different head-ends, a unique association identifier SHOULD be used for those LSPs: this can be achieved by setting the

IPv4/IPv6 source address to a common value (zero value can be used) as well as the Association ID.



Case 1: Disjointness initiated by PCE and enforced by PCC

Case 2: Disjointness initiated by PCC and enforced by PCE

Figure 2 - Sample use-cases for carrying disjoint-group over PCEP session

Using the disjoint-group within a PCUpdate or PCInit may have two purposes:

- o Information: in case the PCE is performing the path computation, it may communicate to the PCC the locally used configured parameters in the attribute-list of the LSP.
- o Configuration: in case the PCC is performing the path computation but the PCE (without computation engine) is managing the LSP parameters, the PCE should add the disjoint-group within the PCUpdate message to communicate to the PCC the disjointness constraint.

4. Protocol extension

4.1. Association group

As per [[I-D.ietf-pce-association-group](#)], LSPs are associated with other LSPs with which they interact by adding them to a common association group. The Association ID will be used to identify the disjoint group a set of LSPs belongs to. This document defines four new Association types, based on the generic Association object -

- o Association type = TBD1 ("Disjointness Association Type") for link disjoint group.
- o Association type = TBD2 ("Disjointness Association Type") for node disjoint group.
- o Association type = TBD3 ("Disjointness Association Type") for srlg disjoint group.
- o Association type = TBD4 ("Disjointness Association Type") for node+srlg disjoint group.

A disjoint group can have two or more LSPs. But a PCE may be limited in how many LSPs it can take into account when computing disjointness: usually PCEs are able to compute a pair of disjoint paths. If a PCE receives more LSPs in the group than it can handle in its computation algorithm, it SHOULD apply disjointness computation to only a subset of LSPs in the group. The subset of disjoint LSPs will be decided by the implementation.

Local policies on the PCC or PCE MAY define the computational behavior for the other LSPs in the group. For example, the PCE may provide no path, a shortest path, or a constrained path based on relaxing disjointness, etc.

Associating a particular LSP to multiple disjoint groups is authorized from a protocol perspective, however there is no insurance that the PCE will be able to compute properly the multi-disjointness constraint.

4.2. Optional TLVs

The disjoint group MAY carry some optional TLVs including but not limited to:

- o VENDOR-INFORMATION-TLV: Used to communicate arbitrary vendor specific behavioral information, described in [[RFC7150](#)].

If a PCEP speaker receives two LSPs with the same disjoint-group but with a different S flag value, it SHOULD apply a strict disjointness path computation for this disjoint-group (it considers S flag set for all LSPs).

5. Security Considerations

This document defines one new type for association, which do not add any new security concerns beyond those discussed in [RFC5440], [I-D.ietf-pce-stateful-pce] and [I-D.ietf-pce-association-group] in itself.

6. IANA Considerations

6.1. Association object Type Indicators

This document defines the following new association type originally defined in [I-D.ietf-pce-association-group].

Value	Name	Reference
TBD1	Link Disjoint-group Association Type	[This I.D.]
TBD2	Node Disjoint-group Association Type	[This I.D.]
TBD3	SRLG Disjoint-group Association Type	[This I.D.]
TBD4	Node+SRLG Disjoint-group Association Type	[This I.D.]

This document defines the following new PCEP TLV:

Value	Name	Reference
TBD5	DISJOINTNESS-INFORMATION-TLV	[This I.D.]

IANA is requested to manage the space of flags carried in the DISJOINTNESS-INFORMATION TLV defined in this document, numbering them from 0 as the least significant bit.

New bit numbers may be allocated in future.

IANA is requested to allocate the following bit numbers in the DISJOINTNESS-INFORMATION-TLV flag space:

Bit Number	Name	Reference
0	Strict disjointness	[This I.D.]
1	Shortest-path	[This I.D.]

7. Manageability Considerations

7.1. Control of Function and Policy

An operator MUST be allowed to configure the disjointness associations and parameters at PCEP peers and associate it with the LSPs.

7.2. Information and Data Models

[RFC7420] describes the PCEP MIB, there are no new MIB Objects for this document.

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](#)].

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 this document do not have any impact on network operations in addition to those already listed in [[RFC5440](#)].

8. Acknowledgments

A special thanks to author of [[I-D.ietf-pce-association-group](#)], this document borrow some of the text from it.

9. References

9.1. 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, <<http://www.rfc-editor.org/info/rfc2119>>.

- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", [RFC 4655](#), DOI 10.17487/RFC4655, August 2006, <<http://www.rfc-editor.org/info/rfc4655>>.
- [RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", [RFC 5440](#), DOI 10.17487/RFC5440, March 2009, <<http://www.rfc-editor.org/info/rfc5440>>.
- [I-D.ietf-pce-association-group]
Minei, I., Crabbe, E., Sivabalan, S., Ananthakrishnan, H., Zhang, X., and Y. Tanaka, "PCEP Extensions for Establishing Relationships Between Sets of LSPs", [draft-ietf-pce-association-group-01](#) (work in progress), July 2016.
- [I-D.ietf-pce-stateful-pce]
Crabbe, E., Minei, I., Medved, J., and R. Varga, "PCEP Extensions for Stateful PCE", [draft-ietf-pce-stateful-pce-17](#) (work in progress), November 2016.

9.2. Informative References

- [RFC7150] Zhang, F. and A. Farrel, "Conveying Vendor-Specific Constraints in the Path Computation Element Communication Protocol", [RFC 7150](#), DOI 10.17487/RFC7150, March 2014, <<http://www.rfc-editor.org/info/rfc7150>>.
- [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, <<http://www.rfc-editor.org/info/rfc7420>>.
- [I-D.ietf-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-ietf-pce-pce-initiated-lsp-07](#) (work in progress), July 2016.

Authors' Addresses

Stephane Litkowski
Orange

EMail: stephane.litkowski@orange.com

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

EMail: msiva@cisco.com

Colby Barth
Juniper Networks

EMail: cbarth@juniper.net