

PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: February 12, 2015

D. Dhody
F. Zhang
X. Zhang
Huawei Technologies
V. Lopez
O. Gonzalez de Dios
Telefonica I+D
August 11, 2014

PCEP Extensions for Receiving SRLG Information
draft-dhody-pce-recv-srlg-02

Abstract

The Path Computation Element (PCE) provides functions of path computation in support of traffic engineering (TE) in networks controlled by Multi-Protocol Label Switching (MPLS) and Generalized MPLS (GMPLS).

This document provides extensions for the Path Computation Element Protocol (PCEP) to receive Shared Risk Link Group (SRLG) information during path computation via encoding this information in the path computation reply message.

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 February 12, 2015.

Copyright Notice

Copyright (c) 2014 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	3
2.	Terminology	3
3.	Usage of SRLG	4
4.	PCEP Requirements	4
5.	Extension to PCEP	4
5.1.	The Extension of the RP Object	5
5.2.	SRLG Subobject in ERO	5
6.	Other Considerations	6
6.1.	Backward Compatibility	6
6.2.	Confidentiality via PathKey	6
6.3.	Coherent SRLG IDs	6
7.	Security Considerations	6
8.	Manageability Considerations	7
8.1.	Control of Function and Policy	7
8.2.	Information and Data Models	7
8.3.	Liveness Detection and Monitoring	7
8.4.	Verify Correct Operations	7
8.5.	Requirements On Other Protocols	7
8.6.	Impact On Network Operations	7
9.	IANA Considerations	7
9.1.	New Subobjects for the ERO Object	8
10.	Acknowledgments	8
11.	References	8
11.1.	Normative References	8
11.2.	Informative References	8
Appendix A.	Contributor Addresses	10

[1.](#) Introduction

As per [\[RFC4655\]](#), PCE based path computation model is deployed in large, multi-domain, multi-region, or multi-layer networks. In such case PCEs may cooperate with each other to provide end to end optimal path.

It is important to understand which TE links in the network might be at risk from the same failures. In this sense, a set of links may constitute a 'shared risk link group' (SRLG) if they share a resource whose failure may affect all links in the set [[RFC4202](#)]. H-LSP (Hierarchical LSP) or S-LSP (Stitched LSP) can be used for carrying one or more other LSPs as described in [[RFC4206](#)] and [[RFC6107](#)]. H-LSP and S-LSP may be computed by PCE(s) and further form as a TE link. The SRLG information of such LSPs can be obtained during path computation itself and encoded in the PCEP Path Computation Reply (PCRep) message. [[I-D.zhang-ccamp-gmpls-uni-app](#)] describes the use of a PCE for end to end User-Network Interface (UNI) path computation.

Note that [[I-D.ietf-ccamp-rsvp-te-srlg-collect](#)] specifies a extension to Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) where SRLG information is collected at the time of signaling. But in case a PCE or cooperating PCEs are used for path computation it is recommended that SRLG information is provided by the PCE(s) during the path computation itself.

[[I-D.farrel-interconnected-te-info-exchange](#)] describes a scaling problem with SRLGs in multi-layer environment and introduce a concept of Macro SRLG (MSRLG). Lower layer SRLG are abstracted at the time of path computation and can be the basis to generate such a Macro SRLG at the PCE.

[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.

CPS: Confidential Path Segment. A segment of a path that contains nodes and links that the policy requires not to be disclosed outside the domain.

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.

SRLG: Shared Risk Link Group.

UNI: User-Network Interface.

3. Usage of SRLG

[RFC4202] states that a set of links may constitute a 'shared risk link group' (SRLG) if they share a resource whose failure may affect all links in the set. For example, two fibers in the same conduit would be in the same SRLG. If an LSR is required to have multiple diversely routed LSPs to another LSR, the path computation should attempt to route the paths so that they do not have any links in common, and such that the path SRLGs are disjoint.

In case a PCE or cooperating PCEs are used for path computation, the SRLG information is provided by the PCE(s). For example, disjoint paths for inter-domain or inter-layer LSPs. In order to achieve path computation for a secondary (backup) path, a PCC may request the PCE for a route that must be SRLG disjoint from the primary (working) path. The Exclude Route Object (XRO) [[RFC5521](#)] is used to specify SRLG information to be explicitly excluded.

4. PCEP Requirements

Following key requirements are identified for PCEP to receive SRLG information during path computation:

SRLG Indication: The PCEP speaker must be capable of indicating whether the SRLG information of the path should be received during the path computation procedure.

SRLG: If requested, the SRLG information should be received during the path computation and encoded in the PCRep message.

Cooperating PCEs [[RFC4655](#)] with inter-PCE Communication work together to provide the end to end optimal path as well as the SRLG information of this path. During inter-domain or inter-layer path computation, the aggregating PCE (Parent PCE [[RFC6805](#)] or Ingress PCE(1) [[RFC5441](#)] or Higher-Layer PCE [[RFC5623](#)]) should receive the SRLG information of path segments from other PCEs and provide the end to end SRLG information of the optimal path to the Path Computation Client (PCC).

5. Extension to PCEP

This document extends the existing RP (Request Parameters) object [[RFC5440](#)] so that a PCEP speaker can request SRLG information during path computation. The SRLG subobject maybe carried inside the Explicit Route Object (ERO) in the PCRep message.

5.1. The Extension of the RP Object

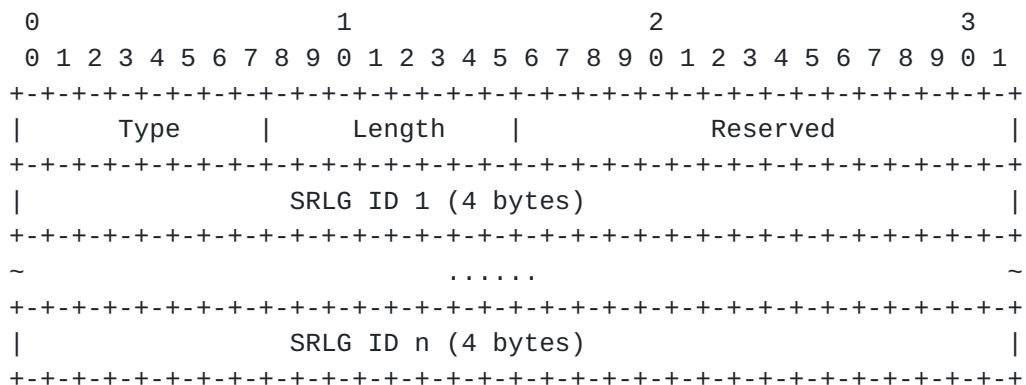
This document adds the following flags to the RP Object:

S (SRLG - 1 bit): when set, in a PCReq message, this indicates that the SRLG information of the path should be provided in the PCRep message. Otherwise, when cleared, this indicates that the SRLG information should not be included in the PCRep message. In a PCRep message, when the S bit is set this indicates that the returned path in ERO also carry the SRLG information; otherwise (when the S bit is cleared), the returned path does not carry SRLG information.

5.2. SRLG Subobject in ERO

As per [RFC5440], ERO is used to encode the path and is carried within a PCRep message to provide the computed path when computation was successful.

The SRLG of a path is the union of the SRLGs of the links in the path [RFC4202]. The SRLG subobject is defined in [I-D.ietf-ccamp-rsvp-te-srlg-collect] for ROUTE_RECORD object (RRO). The same subobject format (as shown below) can be encoded inside the ERO object in the PCRep message.



The meaning and description of Type, Length and SRLG ID can be found in [[I-D.ietf-ccamp-rsvp-te-srlg-collect](#)]. Reserved field SHOULD be set to zero on transmission and MUST be ignored on receipt.

The SRLG subobject should be encoded inside the ERO object in the PCRep message when the S-Bit (SRLG) is set in the PCReq message.

6. Other Considerations

6.1. Backward Compatibility

If a PCE receives a request and the PCE does not understand the new SRLG flag in the RP object, then the PCE SHOULD reject the request.

If PCEP speaker receives a PCRep message with SRLG subobject that it does not support or recognize, it must act according to the existing processing rules of ERO.

6.2. Confidentiality via PathKey

[RFC5520] defines a mechanism to hide the contents of a segment of a path, called the Confidential Path Segment (CPS). The CPS may be replaced by a path-key that can be conveyed in the PCEP and signaled within in a RSVP-TE ERO.

When path-key confidentiality is used, encoding SRLG information in PCRep along with the path-key could be useful to compute a SRLG disjoint backup path at the later instance.

The path segment that needs to be hidden (that is, CPS) MAY be replaced in the ERO with a PKS. The PCE MAY use the SRLG Sub-objects in the ERO along with the PKS sub-object.

6.3. Coherent SRLG IDs

In a multi-layer multi-domain scenario, SRLG ids may be configured by different management entities in each layer/domain. In such scenarios, maintaining a coherent set of SRLG IDs is a key requirement in order to be able to use the SRLG information properly. Thus, SRLG IDs must be unique. Note that current procedure is targeted towards a scenario where the different layers and domains belong to the same operator, or to several coordinated administrative groups. Ensuring the aforementioned coherence of SRLG IDs is beyond the scope of this document. Further scenarios, where coherence in the SRLG IDs cannot be guaranteed are out of the scope of the present document and are left for further study.

7. Security Considerations

This document does not add any new security concerns beyond those discussed in [[RFC5440](#)].

8. Manageability Considerations

8.1. Control of Function and Policy

A PCE involved in inter-domain or inter-layer path computation should be capable of being configured with a SRLG processing policy to specify if the SRLG IDs of the domain or specific layer network can be exposed to the PCC outside the domain or layer network, or whether they should be summarized, mapped to values that are comprehensible to PCC outside the domain or layer network, or removed entirely.

8.2. Information and Data Models

[I-D.ietf-pce-pcep-mib] describes the PCEP MIB, there are no new MIB Objects for this document.

8.3. 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](#)].

8.4. Verify Correct Operations

Mechanisms defined in this document do not imply any new operation verification requirements in addition to those already listed in [[RFC5440](#)].

8.5. Requirements On Other Protocols

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

8.6. 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](#)].

9. IANA Considerations

IANA assigns values to PCEP parameters in registries defined in [[RFC5440](#)]. IANA is requested to make the following additional assignments.

9.1. New Subobjects for the ERO Object

IANA has previously assigned an Object-Class and Object-Type to the ERO carried in PCEP messages [[RFC5440](#)]. IANA also maintains a list of subobject types valid for inclusion in the ERO.

IANA is requested to assign one new subobject types for inclusion in the ERO as follows:

Subobject	Meaning	Reference
xx (TBA)	SRLG sub-object	This document

10. Acknowledgments

Special thanks to the authors of [[I-D.ietf-ccamp-rsvp-te-srlg-collect](#)]. This document borrows some of text from it.

11. References

11.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", [RFC 5440](#), March 2009.

11.2. Informative References

- [RFC4202] Kompella, K. and Y. Rekhter, "Routing Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", [RFC 4202](#), October 2005.
- [RFC4206] Kompella, K. and Y. Rekhter, "Label Switched Paths (LSP) Hierarchy with Generalized Multi-Protocol Label Switching (GMPLS) Traffic Engineering (TE)", [RFC 4206](#), October 2005.
- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", [RFC 4655](#), August 2006.
- [RFC5441] Vasseur, JP., Zhang, R., Bitar, N., and JL. Le Roux, "A Backward-Recursive PCE-Based Computation (BRPC) Procedure to Compute Shortest Constrained Inter-Domain Traffic Engineering Label Switched Paths", [RFC 5441](#), April 2009.

- [RFC5520] Bradford, R., Vasseur, JP., and A. Farrel, "Preserving Topology Confidentiality in Inter-Domain Path Computation Using a Path-Key-Based Mechanism", [RFC 5520](#), April 2009.
- [RFC5521] Oki, E., Takeda, T., and A. Farrel, "Extensions to the Path Computation Element Communication Protocol (PCEP) for Route Exclusions", [RFC 5521](#), April 2009.
- [RFC5623] Oki, E., Takeda, T., Le Roux, JL., and A. Farrel, "Framework for PCE-Based Inter-Layer MPLS and GMPLS Traffic Engineering", [RFC 5623](#), September 2009.
- [RFC6107] Shiomoto, K. and A. Farrel, "Procedures for Dynamically Signaled Hierarchical Label Switched Paths", [RFC 6107](#), February 2011.
- [RFC6805] King, D. and A. Farrel, "The Application of the Path Computation Element Architecture to the Determination of a Sequence of Domains in MPLS and GMPLS", [RFC 6805](#), November 2012.
- [I-D.ietf-ccamp-rsvp-te-srlg-collect]
Zhang, F., Dios, O., Li, D., Margaria, C., Hartley, M., and Z. Ali, "RSVP-TE Extensions for Collecting SRLG Information", [draft-ietf-ccamp-rsvp-te-srlg-collect-06](#) (work in progress), July 2014.
- [I-D.ietf-pce-pcep-mib]
Koushik, K., Emile, S., Zhao, Q., King, D., and J. Hardwick, "Path Computation Element Protocol (PCEP) Management Information Base", [draft-ietf-pce-pcep-mib-09](#) (work in progress), July 2014.
- [I-D.farrel-interconnected-te-info-exchange]
Farrel, A., Drake, J., Bitar, N., Swallow, G., Ceccarelli, D., and X. Zhang, "Problem Statement and Architecture for Information Exchange Between Interconnected Traffic Engineered Networks", [draft-farrel-interconnected-te-info-exchange-06](#) (work in progress), July 2014.
- [I-D.zhang-ccamp-gmpls-uni-app]
Zhang, F., Dios, O., Farrel, A., Zhang, X., and D. Ceccarelli, "Applicability of Generalized Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI)", [draft-zhang-ccamp-gmpls-uni-app-05](#) (work in progress), February 2014.

Appendix A. Contributor Addresses

Udayasree Palle
Huawei Technologies
Leela Palace
Bangalore, Karnataka 560008
India

EMail: udayasree.palle@huawei.com

Avantika
Huawei Technologies
Leela Palace
Bangalore, Karnataka 560008
India

EMail: avantika.sushilkumar@huawei.com

Authors' Addresses

Dhruv Dhody
Huawei Technologies
Leela Palace
Bangalore, Karnataka 560008
India

EMail: dhruv.ietf@gmail.com

Fatai Zhang
Huawei Technologies
Bantian, Longgang District
Shenzhen, Guangdong 518129
P.R.China

EMail: zhangfatai@huawei.com

Xian Zhang
Huawei Technologies
Bantian, Longgang District
Shenzhen, Guangdong 518129
P.R.China

EMail: zhang.xian@huawei.com

Victor Lopez
Telefonica I+D
Distrito Telefonica
Edificio Sur 3, 3rd floor
Madrid 28050
Spain

E-Mail: victor.lopezalvarez@telefonica.com

Oscar Gonzalez de Dios
Telefonica I+D
Distrito Telefonica
Edificio Sur 3, 3rd floor
Madrid 28050
Spain

E-Mail: oscar.gonzalezdedios@telefonica.com

