PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: December 20, 2018

C. Barth
Juniper Networks, Inc.
M. Koldychev
Cisco Systems, Inc.
D. Dhody
Huawei Technology
S. Sivabalan
Cisco Systems, Inc.
June 18, 2018

PCEP extension to support Segment Routing Policy Candidate Paths
draft-barth-pce-segment-routing-policy-cp-00

Abstract

This document introduces a mechanism to specify an SR policy, as a collection of SR candidate paths. An SR policy is identified by <headend, color, destination end-point> tuple. An SR policy can be associated with one or more candidate paths where each candidate path is represented in PCEP by an LSP. This document proposes extension to PCEP to support association among candidate paths of a given SR policy. The mechanism proposed in this document is applicable to both MPLS and IPv6 data plane.

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 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 20, 2018.

Table of Contents

1. Introduction

Path Computation Element (PCE) Communication Protocol (PCEP) [RFC5440] describes the Path Computation Element communication Protocol (PCEP) which enables the communication between a Path Computation Client (PCC) and a Path Control Element (PCE), or between two PCEs based on the PCE architecture [RFC4655].

PCEP Extensions for Stateful PCE Model [RFC8231] describes a set of extensions to PCEP to enable active control of Multiprotocol Label Switching Traffic Engineering (MPLS-TE) and Generalized MPLS (GMPLS) tunnels. [RFC8281] describes the setup and teardown of PCE-initiated

LSPs under the active stateful PCE model, without the need for local configuration on the PCC, thus allowing for a dynamic network.

PCEP Extensions for Segment Routing [I-D.ietf-pce-segment-routing] specifies extensions to the Path Computation Element Protocol (PCEP) that allow a stateful PCE to compute and initiate Traffic Engineering (TE) paths, as well as a PCC to request a path subject to certain constraint(s) and optimization criteria in SR networks.

PCEP Extensions for Establishing Relationships Between Sets of LSPs [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 behaviors) and is equally applicable to stateful PCE (active and passive modes) and stateless PCE.

Segment Routing Policy for Traffic Engineering [I-D.ietf-spring-segment-routing-policy] details the concepts of SR Policy and approaches to steering traffic into an SR Policy.

An SR policy can be associated with one or more candidate paths where one or more such paths can be computed via PCE. This document specifies PCEP extensions to specify a candidate path of an SR policy. Each candidate path being represented by a PCEP LSP. By associating multiple candidate paths together, a PCE becomes aware of the hierarchical structure of an SR policy. Thus the PCE can take computation and control decisions about the candidate paths, with the additional knowledge that these candidate paths belong to the same SR policy. This is accomplished via the use of the existing PCEP Association object, by defining a new association type specifically for associating SR candidate paths into a single SR policy.

2. Terminology

The following terminologies are used in this document:

Destination Endpoint: The destination IPv4 or IPv6 endpoint address of the SR policy in question, as described in [I-D.ietf-spring-segment-routing-policy].

Association parameters: As described in [I-D.ietf-pce-association-group], the combination of the mandatory fields Association type, Association ID and Association Source in the ASSOCIATION object uniquely identify the association group. If the optional TLVs - Global Association Source or Extended Association ID are included, then they MUST be included in combination with mandatory fields to uniquely identifying the association group.

Association information: As described in [I-D.ietf-pce-association-group], the ASSOCIATION object MAY include other optional TLVs based on the association types, that provides 'information' related to the association type.

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 Protocol.

3. Motivation

The new Association Type (SR policy candidate path) and the new TLVs of the ASSOCIATION object, defined in this document, allow PCEP peers to exchange additional parameters of SR policy, defined in [I-D.ietf-spring-segment-routing-policy]. These additional parameters include the color, endpoint, and preference. More parameters may be added in the future.

The motivation for signaling these parameters is summarized in the following subsections.

3.1. Instantiation of SR policy candidate paths

A PCE may want to instantiate one or more candidate paths on the PCC, as specified in [RFC8281]. In this scenario, the PCE needs to signal to a PCC <headend, color, destination end-point, preference> tuple using which the PCC can instantiate a candidate path for the SR policy identified. Current PCEP standards (as of the time of this writing) do not provide a way to signal color and preference. Although destination end-point can be signaled via the PCEP END-POINTS object, this object may not be suitable because the destination end-point to which the path is computed is not required to be the same IPv4/IPv6 address as the actual endpoint of the SR policy. Thus, a separate way to specify SR policy's destination end-point is provided in this draft.

3.2. Color based path computation

A PCE can make use of the color of the SR policy, when computing the path. For example, the PCE can map a certain set of colors to a specific set of constraints or optimization objectives based on its local path computation policy.

3.3. Avoid computing lower preference candidate paths

When a PCE knows that a given set of LSPs all belong to the same SR policy, then path computation can be done on only the highest preference candidate-path(s). Path computation for lower preference paths is not necessary if one or two higher preference paths are already computed. Since computing their paths will not affect traffic steering, it may be postponed until the higher preference paths become invalid, thus saving a lot of computation resources on the PCE.

A PCE MAY take candidate path preference into consideration in path computation subject to its local path computation policy. For example, it MAY choose to compute path with higher preference first before considering path with lower order preference in the context of a given SR policy. It is also possible for the PCE to compute only the path with the highest preference for a given SR policy.

3.4. Minimal signaling overhead

When an SR policy contains multiple candidate paths computed by PCE, such candidate paths can be created, updated and deleted independently of each other. This is achieved by making each candidate path correspond to a unique LSP. For example, if an SR policy has 4 candidate paths, then if the PCE wants to update one of those candidate paths, only one set of PCUpd and PCRpt messages needs to be exchanged.

4. Overview

As per [I-D.ietf-pce-association-group], LSPs are placed into an association group. In our case, each LSP corresponds to a candidate path of an SR policy, and the association group corresponds to the SR policy itself.

A new Association Type is defined in this document, based on the generic ASSOCIATION object. Association type = TBD1 ("SR Candidate Path Association Type") for SR Candidate Path Association Group (SRCPAG).

The SRCPAG Association is only meant to be used for SR LSPs and with PCEP peers which advertise SR capability.

An Association object of SRCPAG group contains the following information:

o Color of SR policy, represented by a 32-bit unsigned integer.

- Destination end-point of SR policy, represented by IPv4 or IPv6 address.

- Preference of SR policy, represented by an unsigned integer.

The values of Color, Destination Endpoint and Preference are encoded as a TLV in the SRCPAG ASSOCIATION object. This TLV is part of Association Information.

The Association Parameters (see Section 2) are derived from the Color and Destination End-point. The mapping from Color and Destination End-point to Association Parameters is left up to the implementation. An implementation MAY choose Association Parameters in such a way that every possible Color and Destination End-point maps to a unique value of Association Parameters, thus requiring the use of Extended Association ID TLV. Alternatively, an implementation MAY implement a lookup table to generate Association Parameters incrementally as new Color and Destination End-point values are created, thus not requiring the use of Extended Association ID TLV.

As per the processing rules specified in section 5.4 of [I-D.ietf-pce-association-group], if a PCEP speaker does not support the SRCPAG association type, it MUST return a PCErr message with Error-Type 26 (Early allocation by IANA) "Association Error" and Error-Value 1 "Association-type is not supported". Please note that the corresponding PCEP session is not reset.

5. SR Candidate Path Association Group

Two object types for IPv4 and IPv6 are defined. The ASSOCIATION object includes "Association type" indicating the type of the association group. This document adds a new Association type.

Association type = TBD1 ("SR Candidate Path Association Type") for SR Candidate Path Association Group (SRCPAG).

The operator configured Association Range SHOULD NOT be set for this association type and MUST be ignored.

SRCPAG may carry additional TLVs to communicate Association Information. This document specifies a new such TLV called "SRCPAG-INFO-TLV" which is used to carry Color, Destination End-point and Preference.

A given LSP MUST belong to at most one SRCPAG, since a candidate path cannot belong to multiple SR policies. If a PCEP speaker receives a PCEP message with more than one SRCPAG for an LSP, then the PCEP speaker MUST send a PCErr message with Error-Type 26 "Association

Error" and Error-Value TBD "Multiple SRCPAG for one LSP". If the message is a PCRpt message, then the PCEP speaker MUST close the PCEP connection. Closing the PCEP connection is necessary because ignoring PCRpt messages may lead to inconsistent LSP DB state between the two PCEP peers.

5.1. SR Candidate Path Association Group Information TLV

The SRCPAG-INFO-TLV is a mandatory TLV for the SRCPAG Association. Only one SRCPAG-INFO-TLV can be carried and only the first occurrence is processed and any others MUST be ignored. SRCPAG-INFO-TLV is part of the Association Information (see Section 2).

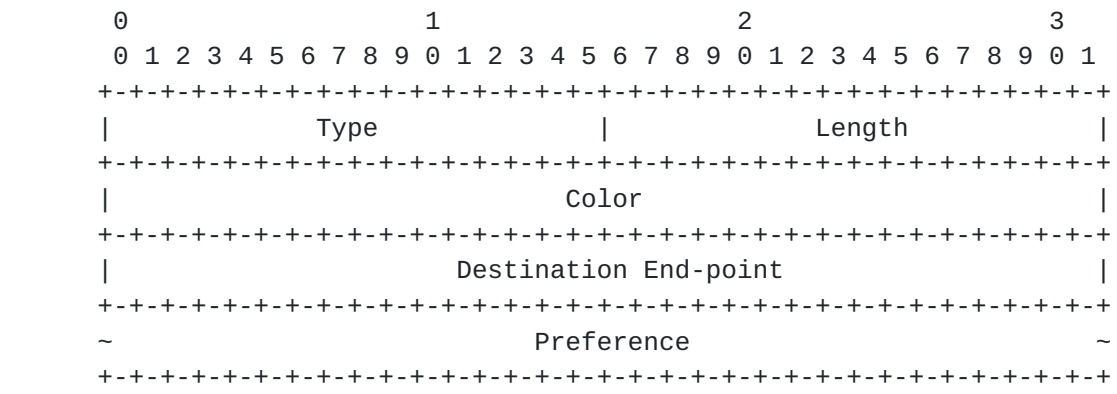

Figure 1: The SRCPAG-INFO-TLV format

Type: TBD2.

Length: variable, depending on length of Destination End-point (IPv4 or IPv6) and the length of the Preference.

Color: any unsigned 32-bit number.

Destination End-point: can be either IPv4 or IPv6, depending on whether the policy endpoint IPv4 or IPv6. This value may be different from the one contained in the existing END-POINTS object, or in the LSP IDENTIFIERS TLV of the LSP object. Destination Endpoint is meant to strictly correspond to the endpoint of the SR policy, as it is defined in [I-D.ietf-spring-segment-routing-policy].

Preference: value of the candidate path preference.

6. Examples

6.1. PCC Initiated SR Policy with single candidate-path

PCReq and PCRep messages are exchanged in the following sequence:

1. PCC sends PCReq message to the PCE, encoding the SRCPAG ASSOCIATION object and TLVs in the PCReq message.

2. PCE returns the path in PCRep message, and echoes back the SRCPAG object that was used in the computation.

PCRpt and PCUpd messages are exchanged in the following sequenece:

1. PCC sends PCRpt message to the PCE, including the LSP object and the SRCPAG ASSOCIATION object.

2. PCE computes path, possibly making use of the Color, Endpoint and Preference that were extracted from the SRCPAG ASSOCIATION object.

3. PCE updates the SR policy candidate path's ERO using PCUpd message.

6.2. PCC Initiated SR Policy with multiple candidate-paths

PCReq and PCRep messages are exchanged using the sequence specified in section 6.1 with individual query for each candidate-path.

PCRpt and PCUpd messages are exchanged in the following sequence:

1. Step 1: For each candidate path of the SR policy, the PCC generates a different PLSP-ID and symbolic-name and sends multiple PCRpt messages (or one message with multiple LSP objects) to the PCE. Each LSP object is followed by SRCPAG ASSOCIATION object with identical Color and Endpoint values.

2. Step 2: PCE takes into account that all the LSPs belong to the same SR policy. PCE prioritizes computation for the highest preference LSP and sends PCUpd message(s) back to the PCC.

3. Step 3: If a new candidate path is added on the PCC by the operator, then a new PLSP-ID and symbolic name is generated for that candidate path and a new PCRpt is sent to the PCE.

4. Step 4: If an existing candidate path is removed from the PCC by the operator, then that PLSP-ID is deleted from the PCE by sending PCRpt with the R-flag set.

6.3. PCE Initiated SR Policy with single candidate-path

A candidate-path is created using the following steps:

1. PCE sends PCInit message, as usual containing the SRCPAG Association object. PCE needs to generate a symbolic-name for this LSP that will not clash with other symbolic names on the same PCC.

2. PCC uses the color, destination endpoint and preference from the SRCPAG object to create a new candidate path. If no SR policy exists to hold the candidate path, then a new SR policy is created to hold the new candidate-path. The Originator of the candidate path is set to be the address of the PCE that is sending the PCInit message.

3. PCC allocates a locally unique PLSP-ID for the newly created candidate path. This PLSP-ID is sent to the PCE in the PCRpt message.

A candidate-path is deleted using the following steps:

1. PCE sends PCInit message, setting the R-flag in the LSP object.

2. PCC uses the PLSP-ID from the LSP object to find the candidate path and delete it. If this is the last candidate path under the SR policy, then the containing SR policy is deleted as well.

6.4. PCE Initiated SR Policy with multiple candidate-paths

A candidate-path is created using the following steps:

1. PCE sends a separate PCInit message for every candidate path that it wants to create, or it sends multiple LSP objects within a single PCInit message. Each candidate-path must get a unique symbolic-name generated on the PCE. SRCPAG object is sent for every LSP in the PCInit message.

2. PCC creates multiple candidate paths under the same SR policy, identified by Color and Endpoint. PCC generates a unique PLSP-ID for every candidate path.

3. PCC allocates a locally unique PLSP-ID for each newly created candidate path. This PLSP-ID is sent to the PCE in the PCRpt message.

A candidate path is deleted using the following steps:

1. PCE sends PCInit message, setting the R-flag in the LSP object.

2. PCC uses the PLSP-ID from the LSP object to find the candidate path and delete it.

7. Acknowledgement

8. Normative References

[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-11 (work in progress), October 2017.

[I-D.ietf-pce-segment-routing]
Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "PCEP Extensions for Segment Routing", draft-ietf-pce-segment-routing-11 (work in progress), November 2017.

[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-21 (work in progress), June 2017.

[I-D.ietf-spring-segment-routing]
Filsfils, C., Previdi, S., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", draft-ietf-spring-segment-routing-15 (work in progress), January 2018.

[I-D.sivabalan-pce-binding-label-sid]
Sivabalan, S., Tantsura, J., Filsfils, C., Previdi, S., Hardwick, J., and D. Dhody, "Carrying Binding Label/ Segment-ID in PCE-based Networks.", draft-sivabalan-pce-binding-label-sid-04 (work in progress), March 2018.

[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>.

Authors' Addresses

Colby Barth
Juniper Networks, Inc.

Email: cbarth@juniper.net

Mike Koldychev
Cisco Systems, Inc.
2000 Innovation Drive
Kanata, Ontario K2K 3E8
Canada

Email: mkoldych@cisco.com

Dhruv Dhody
Huawei Technology

Email: dhruv.dhody@huawei.com

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

Email: msiva@cisco.com