Network Working Group
Internet-Draft
Updates: 5440 (if approved)
Intended status: Standards Track
Expires: July 17, 2021

A. Stone
M. Aissaoui
Nokia
S. Sidor
Cisco Systems, Inc.
S. Sivabalan
Ciena Coroporation

January 13, 2021

Local Protection Enforcement in PCEP
draft-ietf-pce-local-protection-enforcement-01

Abstract

This document updates [RFC5440] to clarify usage of the local protection desired bit signalled in Path Computation Element Protocol (PCEP). This document also introduces a new flag for signalling protection strictness in 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 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 July 17, 2021.

Table of Contents

1. Introduction

Path Computation Element (PCE) Communication Protocol (PCEP) [RFC5440] 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 [RFC5440] utilizes flags, values and concepts previously defined in RSVP-TE Extensions [RFC3209] and Fast Reroute Extensions to RSVP-TE [RFC4090]. One such concept in PCEP is the 'Local Protection Desired' (L-flag in the LSPA Object in RFC5440), which was originally defined in the SESSION-ATTRIBUTE Object in RFC3209. In RSVP, this flag signals to downstream routers that local protection is desired, which indicates to transit routers that they may use a local repair mechanism. The headend router calculating the path does not know whether a downstream router will or will not protect a hop during it's calculation. Therefore, a local protection desired does not require the transit router to satisfy protection in order to establish the RSVP signalled path. This flag is signalled in PCEP as an attribute of the LSP via the LSP Attributes object.

PCEP Extensions for Segment Routing (draft-ietf-pce-segment-routing) extends support in PCEP for Segment Routed LSPs (SR-LSPs) as defined in the Segment Routing Architecture [RFC8402]. As per the Segment Routing Architecture, Adjacency Segment Identifiers(Adj-SID) may be eligible for protection (using IPFRR or MPLS-FRR). The protection eligibility is advertised into IGP (draft-ietf-ospf-segment-routing-extensions and draft-ietf-isis-segment-routing-extensions) as the B-Flag part of the Adjacency SID sub-tlv and can be discovered by a PCE via BGP-LS [RFC7752] using the BGP-LS Segment Routing Extensions (draft-ietf-idr-bgp-ls-segment-routing-ext). An Adjacency SID may or may not have protection eligibility and for a given adjacency between two routers there may be multiple Adjacency SIDs, some of which are protected and some which are not.

A Segment Routed path calculated by PCE may contain various types of segments, as defined in [RFC8402] such as Adjacency, Node or Binding. The protection eligibility for Adjacency SIDs can be discovered by PCE, so therefore the PCE can take the protection eligibility into consideration as a path constraint. If a path is calculated to include other segment identifiers which are not applicable to having their protection state advertised, as they may only be locally significant for each router processing the SID such as Node SIDs, it may not be possible for PCE to include the protection constraint as part of the path calculation.

It is desirable for an operator to define the enforcement, or strictness of the protection requirement when it can be applied.

This document updates [RFC5440] by further describing the behaviour with Local Protection Desired Flag (L-Flag) and extends on it with the introduction of Enforcement Flag (E-Flag).

2. Requirements Language

In this document, the key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" are to be interpreted as described in BCP 14, [RFC2119].

3. Terminology

This document uses the following terminology:

PROTECTION MANDATORY: path MUST have protection eligibility on all links.

UNPROTECTED MANDATORY: path MUST NOT have protection eligibility on all links.

PROTECTION PREFERRED: path SHOULD have protection eligibility on all links but MAY contain links which do not have protection eligibility.

UNPROTECTED PREFERRED: path SHOULD NOT have protection eligibility on all links but MAY contain links which have protection eligibility.

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.

4. Motivation

4.1. Implementation differences

As defined in [RFC5440] the mechanism to signal protection enforcement in PCEP is with the previously mentioned L-flag defined in the LSPA Object. The name of the flag uses the term "Desired", which by definition means "strongly wished for or intended" and is rooted in the RSVP use case. For RSVP, this is not within control of the PCE. However, [RFC5440] does state "When set, this means that the computed path must include links protected with Fast Reroute as defined in [RFC4090]." Implementations of [RFC5440] have either interpreted the L-Flag as PROTECTION MANDATORY or PROTECTION PREFERRED, leading to operational differences.

4.2. SLA Enforcement

The boolean bit flag is unable to distinguish between the different options of PROTECTION MANDATORY, UNPROTECTED MANDATORY, PROTECTION PREFERRED and UNPROTECTED PREFERRED. The selection of the options are typically dependent on the service level agreement the operator wishes to impose on the LSP. When enforcement is used, the resulting shortest path calculation is impacted.

For example, PROTECTION MANDATORY is for use cases where an operator may need the LSP to follow a path which has local protection provided along the full path, ensuring that if there is anywhere along the path that traffic will be fast re-routed at the point of failure.

For another example, UNPROTECTED MANDATORY is when an operator may intentionally prefer an LSP to not be locally protected, and thus would rather local failures to cause the LSP to go down and/or rely on other protection mechanisms such as a secondary diverse path.

There are also use cases where there is simply no requirement to enforce protection or no protection along a path. This can be considered as "do not care to enforce". This is a relaxation of the protection constraint. The path calculation is permitted the use of any SID which is available along the calculated path. The SID backup availability does not impact the shortest path computation. Since links may have both protected and unprotected SIDs available, the option PROTECTION PREFERRED or UNPROTECTED PREFERRED is used to instruction PCE a preference on which SID to select, as the behaviour of the LSP would differ during a local failure depending on which SID is selected.

5. Protection Enforcement Flag (E-Flag)

Section 7.11 in Path Computation Element Protocol [RFC5440] describes the encoding of the Local Protection Desired (L-Flag). A new flag is proposed in this document in the LSP Attributes Object which extends the L-Flag to identify the protection enforcement.

The flag bit is to be allocated by IANA following IETF Consensus.

This draft version proposes using the next available bit: {TBD}
//Editor note, next available bit at time of writing is bit 6

Codespace of the Flag field (LSPA Object)

Bit	Description	Reference
7	Local Protection Desired	RFC5440
<TBD>	Local Protection Enforcement	This document

The format of the LSPA Object as defined in [RFC5440] is:

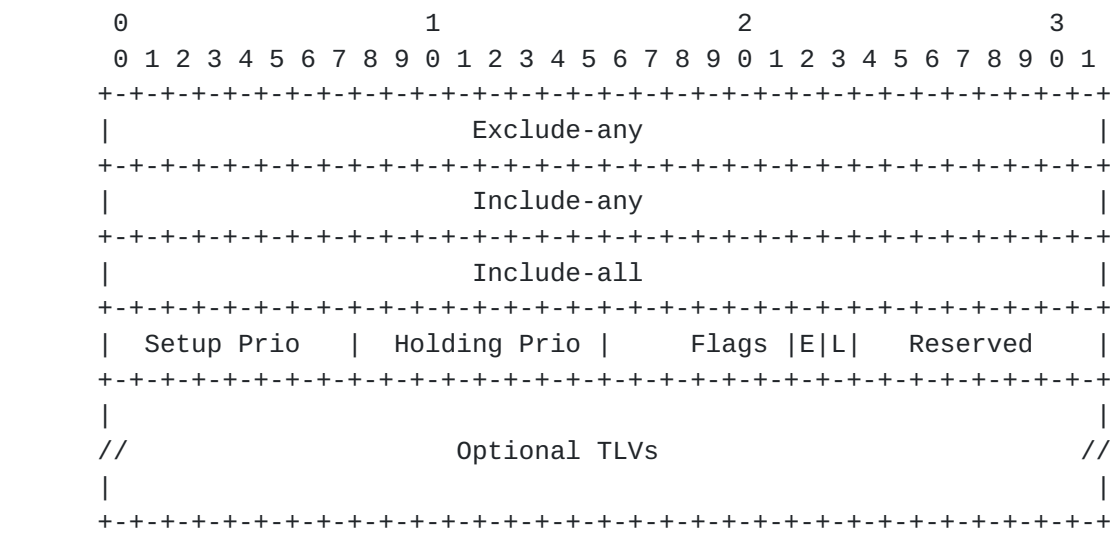

Flags (8 bits)

- L flag: As defined in [RFC5440] and further updated by this document. When set, protection is desired. When not set, protection is not desired. The enforcement of the protection is identified via the E-Flag.

- E flag (Protection Enforcement): When set, the value of the L-Flag MUST be treated as a MUST constraint where applicable, when protection state of a SID is known. When E flag is not set, the value of the L-Flag MUST be treated as a MAY constraint.

When L-flag is set and E-flag is set then PCE MUST consider the protection eligibility as PROTECTION MANDATORY constraint.

When L-flag is set and E-flag is not set then PCE MUST consider the protection eligibility as PROTECTION PREFERRED constraint.

When L-flag is not set and E-flag is not set then PCE SHOULD consider the protection eligibility as UNPROTECTED PREFERRED but MAY consider protection eligibility as UNPROTECTED MANDATORY constraint.

When L-flag is not set and E-flag is set then PCE MUST consider the protection eligibility as UNPROTECTED MANDATORY constraint.

UNPROTECTED PREFERRED and PROTECTED PREFERRED may seem similar but they indicate the preference of selection of a SID if PCE has an option of either protected or unprotected available on a link. When presented with either option, PCE SHOULD select the SID which has a protection state matching the state of the L-Flag.

The protection enforcement constraint can only be applied to resource selection in which the protection state is known to PCE. A PCE calculating a path that includes resources which does not support the protection state being known to PCE (such as Node SID), then the protection state MAY ignore the protection enforcement constraint.

5.1. Backwards Compatibility

Considerations in the message passing between PCC and PCE for the E-Flag bit which are not supported by the entity are outlined in this section, with requirements for PCE and PCC implementing this document described at the end.

For a PCC or PCE which does not yet support this document, the E-flag bit is ignored and set to zero in PCRpt and/or PCUpd as per [RFC5440] for PCC-initiated or as per ([RFC8281]) for PCE-initiated LSPs. It's important to note that [RFC8231] and [RFC8281] permit LSP Attribute Object to be included in PCUpd messages for PCC-initiated and PCE-initiated LSPs. For PCC-initiated LSPs, PCUpd E-Flag (and L-Flag) are an echo from the previous PCRpt however the bit value is ignored on PCE from the previous PCRpt, therefore the E-Flag value set in the PCUpd is zero. A PCE which does not support this document sends PCUpd messages with the E-Flag unset for PCC-initated LSPs even if set in the prior PCReq or PCRpt. A PCC which does not support this document sends PCRpt messages with the E-Flag unset for PCE-initiated LSPs even if set in the prior PCInitiate or PCUpd.

For a PCC which does support this document, it MAY set E-Flag bit depending on local configuration. If communicating with a PCE which does not yet support this document, the PCE follows the behaviour specified in [RFC5440] and will ignore the E-Flag bit thus it will not compute a path respecting the enforcement constraint.

For PCC-initiated LSPs, PCC SHOULD ignore the E-Flag value received from PCE in a PCUpd message.

For PCE-initiated LSPs, PCC MAY process the E-Flag value received from PCE in a PCUpd message. PCE SHOULD ignore the E-Flag value received from PCC in a PCRpt message.

6. Implementation Status

[Note to the RFC Editor - remove this section before publication, as well as remove the reference to RFC 7942.]

This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in [RFC7942].

The description of implementations in this section is intended to assist the IETF in its decision processes in progressing drafts to RFCs. Please note that the listing of any individual implementation here does not imply endorsement by the IETF. Furthermore, no effort has been spent to verify the information presented here that was supplied by IETF contributors. This is not intended as, and must not be construed to be, a catalogue of available implementations or their features. Readers are advised to note that other implementations may exist.

According to [RFC7942], "this will allow reviewers and working groups to assign due consideration to documents that have the benefit of running code, which may serve as evidence of valuable experimentation and feedback that have made the implemented protocols more mature. It is up to the individual working groups to use this information as they see fit".

6.1. Nokia Implementation

- Organization: Nokia
- Implementation: NSP PCE and SROS PCC.
- Description: Implementation for calculation and conveying intention described in this document
- Maturity Level: Demo
- Coverage: Full
- Contact: andrew.stone@nokia.com

6.2. Cisco Implementation

- Organization: Cisco Systems, Inc.
- Implementation: IOS-XR PCE and PCC.
- Description: Implementation for calculation and conveying intention described in this document
- Maturity Level: Demo
- Coverage: Full
- Contact: ssidor@cisco.com

7. Security Considerations

This document clarifies the behaviour of an existing flag and introduces a new flag to provide further control of that existing behaviour. The introduction of this new flag and behaviour clarification does not create any new sensitive information. No additional security measure is required.

Securing the PCEP session using Transport Layer Security (TLS) [RFC8253], as per the recommendations and best current practices in [RFC7525] is RECOMMENDED.

8. IANA Considerations

8.1. LSPA Object

This document defines a new flag in the LSPA Object. IANA is requested to allocate the following code point in the "LSPA Object Flag Field" subregistry in the "Path Computation Element Protocol (PCEP) Numbers" registry.

Value	Name	Reference
TBD	Protection Enforcement	This document

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

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <https://www.rfc-editor.org/info/rfc3209>.

[RFC4090] Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", RFC 4090, DOI 10.17487/RFC4090, May 2005, <https://www.rfc-editor.org/info/rfc4090>.

[RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <https://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, <https://www.rfc-editor.org/info/rfc5440>.

[RFC7525] Sheffer, Y., Holz, R., and P. Saint-Andre, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 7525, DOI 10.17487/RFC7525, May 2015, <https://www.rfc-editor.org/info/rfc7525>.

[RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", RFC 7752, DOI 10.17487/RFC7752, March 2016, <https://www.rfc-editor.org/info/rfc7752>.

[RFC8231] Crabbe, E., Minei, I., Medved, J., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE", RFC 8231, DOI 10.17487/RFC8231, September 2017, <https://www.rfc-editor.org/info/rfc8231>.

[RFC8253] Lopez, D., Gonzalez de Dios, O., Wu, Q., and D. Dhody, "PCEPS: Usage of TLS to Provide a Secure Transport for the Path Computation Element Communication Protocol (PCEP)", RFC 8253, DOI 10.17487/RFC8253, October 2017, <https://www.rfc-editor.org/info/rfc8253>.

[RFC8281] Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for PCE-Initiated LSP Setup in a Stateful PCE Model", RFC 8281, DOI 10.17487/RFC8281, December 2017, <https://www.rfc-editor.org/info/rfc8281>.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <https://www.rfc-editor.org/info/rfc8402>.

9.2. Informative References

[RFC7942] Sheffer, Y. and A. Farrel, "Improving Awareness of Running Code: The Implementation Status Section", BCP 205, RFC 7942, DOI 10.17487/RFC7942, July 2016, <https://www.rfc-editor.org/info/rfc7942>.

Acknowledgements

Thanks to Dhruv Dhody for comments and discussions on this document and Mike Koldychev for reviews.

Authors' Addresses

Andrew Stone
Nokia

Email: andrew.stone@nokia.com

Mustapha Aissaoui
Nokia

Email: mustapha.aissaoui@nokia.com

Samuel Sidor
Cisco Systems, Inc.

Email: ssidor@cisco.com

Siva Sivabalan
Ciena Coroporation

Email: ssivabal@ciena.com