

Internet Engineering Task Force (IETF)
Request for Comments: 9488
Updates: 5440
Category: Standards Track
ISSN: 2070-1721

A. Stone
M. Aissaoui
Nokia
S. Sidor
Cisco Systems, Inc.
S. Sivabalan
Ciena Corporation
October 2023

Local Protection Enforcement in the Path Computation Element Communication Protocol (PCEP)

Abstract

This document updates RFC 5440 to clarify usage of the Local Protection Desired bit signaled in the Path Computation Element Communication Protocol (PCEP). This document also introduces a new flag for signaling protection enforcement in PCEP.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9488>.

Copyright Notice

Copyright (c) 2023 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction
2. Requirements Language
3. Terminology
4. Motivation
 - 4.1. Implementation Differences
 - 4.2. SLA Enforcement
5. Protection Enforcement Flag (E Flag)
 - 5.1. Backwards Compatibility
6. Security Considerations
7. IANA Considerations
8. References

8.1. Normative References

8.2. Informative References

Acknowledgements

Authors' Addresses

1. Introduction

The Path Computation Element Communication Protocol (PCEP) [RFC5440] enables the communication between a Path Computation Client (PCC) and a 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 LSP Attributes (LSPA) object in [RFC5440], which was originally defined in the Session Attribute object in [RFC3209]. In RSVP, this flag signals to downstream routers that they may use a local repair mechanism. The headend router calculating the path does not know if a downstream router will or will not protect a hop during its calculation. Therefore, the L flag does not require the transit router to satisfy protection in order to establish the RSVP-signaled path. This flag is signaled in PCEP as an attribute of the Label Switched Path (LSP) via the LSPA object.

PCEP Extensions for Segment Routing [RFC8664] extends support in PCEP for Segment Routing paths. The path list is encoded with Segment Identifiers (SIDs), each of which might offer local protection. The PCE may discover the protection eligibility for a SID via the Border Gateway Protocol - Link State (BGP-LS) [RFC9085] and take the protection into consideration as a path constraint.

It is desirable for an operator to be able to define the enforcement of the protection requirement.

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

The document contains descriptions in the context of Segment Routing; however, the content described is agnostic in regard to path setup type and data plane technology.

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Terminology

This document uses the following terminology:

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

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

PROTECTION PREFERRED: The path should have protection eligibility on all links but might contain links that do not have protection eligibility.

UNPROTECTED PREFERRED: The path should not have protection

eligibility on all links but might contain links that 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 Communication Protocol

LSPA: LSP Attributes object [RFC5440]

4. Motivation

4.1. Implementation Differences

As defined in [RFC5440], the mechanism to signal protection enforcement in PCEP is 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". The use case for this flag originated in RSVP. For RSVP-signaled paths, local protection is not within control of the PCE. However, [RFC5440] does state that "[w]hen set, this means that the computed path must include links protected with Fast Reroute as defined in [RFC4090]." Implementations that use PCEP [RFC5440] have interpreted the L flag as either PROTECTION MANDATORY or PROTECTION PREFERRED, leading to operational differences.

4.2. SLA Enforcement

The L flag is a boolean bit and thus unable to distinguish between the different options of PROTECTION MANDATORY, UNPROTECTED MANDATORY, PROTECTION PREFERRED, and UNPROTECTED PREFERRED. Selecting one of these options is typically dependent on the Service Level Agreement (SLA) the operator wishes to impose on the LSP. A network may be providing transit to multiple SLA definitions against the same base topology network, whose behavior could vary, such as wanting local protection to be invoked on some LSPs and not wanting local protection on others. When enforcement is used, the resulting shortest path calculation is impacted.

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

As another example, UNPROTECTED MANDATORY is for use cases in which an operator may intentionally prefer an LSP to not be locally protected and thus would rather local failures cause the LSP to go down. An example scenario is one where an LSP is protected via a secondary diverse LSP. Each LSP is traffic engineered to follow specific traffic-engineered criteria computed by the PCE to satisfy the SLA. Upon a failure, if local protection is invoked on the active LSP traffic, the traffic may temporarily traverse links that violate the TE requirements and could negatively impact the resources being traversed (e.g., insufficient bandwidth). In addition, depending on the network topological scenario, it may not be feasible for the PCE to reroute the LSP while respecting the TE requirements, which include path diversity; this results in the LSP being torn down and switched to the protected path anyways. In such scenarios, it is desirable for the LSP to be simply torn down immediately and not rerouted through local protection, so that traffic may be forwarded through an already-established traffic-engineered secondary path.

Both the UNPROTECTED PREFERRED and PROTECTED PREFERRED options provide a relaxation of the protection constraint. These options can be used when an operator does not require protection enforcement. Regardless of the option selected, the protection status of a resource does not influence whether the link must be pruned during a path calculation. Furthermore, the selection of either option indicates a priority selection to the PCE when there is an option to choose a protected or unprotected instruction associated with a resource, ensuring consistent PCE behavior across different implementations.

When used with Segment Routing, an adjacency may have both a protected SID and an unprotected SID. If the UNPROTECTED PREFERRED option is selected, the PCE chooses the unprotected SID. Alternatively, if the PROTECTED PREFERRED option is selected, the PCE chooses the protected SID.

5. Protection Enforcement Flag (E Flag)

Section 7.11 of [RFC5440] describes the encoding of the Local Protection Desired (L) flag. The Protection Enforcement (E) flag, which extends the L flag, is specified below.

Bit	Description	Reference
6	Protection Enforcement	RFC 9488
7	Local Protection Desired	RFC 5440

Table 1: Codespace of the Flag Field (LSPA Object)

The following shows the format of the LSPA object as defined in [RFC5440] with the addition of the E flag defined in this document:

0										1										2										3									
0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1	2	3	4	5	6	7	8	9	0	1								
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-								
										Exclude-any																													
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-								
										Include-any																													
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-								
										Include-all																													
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-								
	Setup Prio									Holding Prio									Flags E L					Reserved															
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-								
//										Optional TLVs																		//											
+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-	+	-								

Flags (8 bits):

L (Local Protection Desired): This flag is defined in [RFC5440] and further updated by this document. When set to 1, protection is desired. When set to 0, protection is not desired. The enforcement of the protection is identified via the E flag.

E (Protection Enforcement): This flag controls the strictness with which the PCE must apply the L flag. When set to 1, the value of the L flag needs to be respected during resource selection by the PCE. When the E flag is set to 0, an attempt to respect the value of the L flag is made; however, the PCE

could relax or ignore the L flag when computing a path. The statements below indicate preference when the E flag is set to 0 in combination with the L flag value.

When both the L flag and E flag are set to 1, then the PCE MUST consider the protection eligibility as a PROTECTION MANDATORY constraint.

When the L flag is set to 1 and the E flag is set to 0, then the PCE MUST consider the protection eligibility as a PROTECTION PREFERRED constraint.

When both the L flag and E flag are set to 0, then the PCE SHOULD consider the protection eligibility as an UNPROTECTED PREFERRED constraint but MAY consider the protection eligibility as an UNPROTECTED MANDATORY constraint. An example of when the latter behavior might be chosen is if the PCE has some means (outside the scope of this document) to detect that it is interacting with a legacy PCC that expects the legacy behavior.

When the L flag is set to 0 and the E flag is set to 1, then the PCE MUST consider the protection eligibility as an UNPROTECTED MANDATORY constraint.

If a PCE is unable to infer the protection status of a resource, the PCE MAY use local policy to define protected status assumptions. When computing a Segment Routing path, it is RECOMMENDED that a PCE assume a Node SID is protected. It is also RECOMMENDED that a PCE assume an Adjacency SID is protected if the backup flag advertised with the Adjacency SID is set.

5.1. Backwards Compatibility

This section outlines considerations for the E flag bit in the message passing between the PCC and the PCE that are not supported by the entity. The requirements for the PCE and the PCC implementing this document are described at the end.

For a PCC or PCE that does not yet support this document, the E flag is ignored and set to 0 in PCRpt and/or PCUpd messages as per [RFC5440] for PCC-initiated LSPs or as per [RFC8281] for PCE-initiated LSPs. It is important to note that [RFC8231] and [RFC8281] permit the LSPA object [RFC5440] to be included in PCUpd messages for PCC-initiated and PCE-initiated LSPs.

For PCC-initiated LSPs, the E flag (and L flag) in a PCUpd message is an echo from the previous PCRpt message; however, the bit value is ignored on the PCE from the previous PCRpt message, so the E flag value set in the PCUpd message is 0. A PCE that does not support this document sends PCUpd messages with the E flag set to 0 for PCC-initiated LSPs even if set to 1 in the prior PCReq or PCRpt message.

A PCC that does not support this document sends PCRpt messages with the E flag set to 0 for PCE-initiated LSPs even if set to 1 in the prior PCInitiate or PCUpd message.

For a PCC that does support this document, the E flag MAY be set to 1 depending on local configuration. If communicating with a PCE that does not yet support this document, the PCE follows the behavior specified in [RFC5440] and ignores the E flag. Thus, a computed path might not respect the enforcement constraint.

For PCC-initiated LSPs, the PCC SHOULD ignore the E flag value received from the PCE in a PCUpd message as it may be communicating with a PCE that does not support this document.

For PCE-initiated LSPs, the PCC MAY process the E flag value received from the PCE in a PCUpd message. The PCE SHOULD ignore the E flag value received from the PCC in a PCRpt message as it may be communicating with a PCC that does not support this document.

6. Security Considerations

This document clarifies the behavior of an existing flag and introduces a new flag to provide further control of that existing behavior. The introduction of this new flag and the behavior clarification do 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 [RFC9325], is RECOMMENDED.

7. IANA Considerations

This document defines a new bit value in the subregistry "LSPA Object Flag Field" in the "Path Computation Element Protocol (PCEP) Numbers" registry. IANA has made the following codepoint allocation.

Bit	Description	Reference
6	Protection Enforcement	RFC 9488

Table 2: Addition to LSPA Object Flag Field Registry

8. References

8.1. 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>>.
- [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>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.
- [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>>.
- [RFC9325] Sheffer, Y., Saint-Andre, P., and T. Fossati, "Recommendations for Secure Use of Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS)", BCP 195, RFC 9325, DOI 10.17487/RFC9325, November 2022, <<https://www.rfc-editor.org/info/rfc9325>>.

8.2. Informative References

- [RFC4655] Farrel, A., Vasseur, J.-P., 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>>.
- [RFC8664] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing", RFC 8664, DOI 10.17487/RFC8664, December 2019, <<https://www.rfc-editor.org/info/rfc8664>>.
- [RFC9085] Previdi, S., Talaulikar, K., Ed., Filsfils, C., Gredler, H., and M. Chen, "Border Gateway Protocol - Link State (BGP-LS) Extensions for Segment Routing", RFC 9085, DOI 10.17487/RFC9085, August 2021, <<https://www.rfc-editor.org/info/rfc9085>>.

Acknowledgements

Thanks to Dhruv Dhody, Mike Koldychev, and John Scudder for reviewing and providing very valuable feedback and discussions on this document.

Thanks to Julien Meuric for shepherding this document.

Authors' Addresses

Andrew Stone
Nokia
600 March Road
Kanata Ontario K2K 2T6
Canada
Email: andrew.stone@nokia.com

Mustapha Aissaoui
Nokia
600 March Road
Kanata Ontario K2K 2T6
Canada
Email: mustapha.aissaoui@nokia.com

Samuel Sidor
Cisco Systems, Inc.
Eurovea Central 3

Pribinova 10
811 09 Bratislava
Slovakia
Email: ssidor@cisco.com

Siva Sivabalan
Ciena Corporation
385 Terry Fox Drive
Kanata Ontario K2K 0L1
Canada
Email: ssivabal@ciena.com