Workgroup: Network Working Group
Internet-Draft: draft-chen-pce-mbinding-03
Published: 28 March 2024
Intended Status: Standards Track
Expires: 29 September 2024
Authors: H. Chen B. Decraene G. Mishra A. Wang
Futurewei Orange Verizon China Telecom
X. Liu L. Liu
IBM Corporation Fujitsu

PCE for Mirror Binding

Abstract

PCE is used to distribute a binding to a node. The binding includes a binding SID and a path represented by a list of SIDs. This document describes extensions to PCEP for distributing the information about the binding to a protecting node. For an SR path via the node with the binding SID, when the node fails, the protecting node such as the upstream neighbor on the path uses the information to protect the binding SID of the failed node.

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 RFC 2119 [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 29 September 2024.

Copyright Notice

Table of Contents

1. Introduction
2. Extensions to PCEP
  2.1. Binding Protection Information Distribution Capability
  2.2. Extensions to RP/SRP Object
3. Procedure for Updating Information
4. References
  4.1. Normative References
  4.2. Informative References
Acknowledgments
Authors' Addresses

1. Introduction

[I-D.ietf-pce-binding-label-sid] specifies how PCE may be used to distribute a Segment Routing (SR) Policy to a node in a network. An SR Policy may contain a binding, which includes a binding SID and a path associated with the binding SID. The path is represented by a list of SIDs.

After a PCE as a controller distributes the binding to the node, the node forwards the packet with the binding SID according to the first SID in the list. It replaces the binding SID in the packet with the list of SIDs and forwards the packet using the FIB entry for the top SID (i.e., the first SID) in the packet.

When the node fails, suppose that the upstream neighbor (as PLR) of the node has the corresponding binding protection information for protecting the binding SID of the node. The information includes the binding SID, the list of SIDs and an identifier of the node. After the upstream neighbor as PLR detects the failure of the node, for a packet with the node SID of the failed node received, it protects the binding SID of the failed node. It pops the node SID, replaces the binding SID in the packet with the list of SIDs, forwards the packet without going through the failed node towards the top SID (i.e., the first SID, assuming it is a node SID for simplicity here). How a upstream node protect the binding SID of the failed node is out of scope of this document and described in

[I-D.ietf-spring-segment-protection-sr-te-paths] and [I-D.hu-spring-segment-routing-proxy-forwarding] (Note: the second reference will be removed after the first one includes enough text for protecting binding SIDs of a node).

This document proposes some procedures and extensions to PCEP for distributing the binding protection information to the possible protecting nodes, which are the nodes that may protect the failed node.

2. Extensions to PCEP

A PCC may run on each node in a network. A PCE runs on a server as a controller to communicate with PCCs. The PCE and the PCCs work together to distribute the binding protection information about a binding SID on a node to the possible protecting nodes for protecting the binding SID of the node when the node fails.

2.1. Binding Protection Information Distribution Capability

When a PCE and a PCC running on a network node establish a PCEP session between them, they exchange their capabilities of Binding Protection Information Distribution in Open messages. An Open message includes an Open object. The object contains a PATH_SETUP_TYPE_CAPABILITY TLV with Path Setup Type (PST) = TBD1 and a new sub-TLV, called Binding Protection Information Distribution Capability (BSID-D for short) sub-TLV.

PST = TBD1 indicates Binding Protection Information Distribution.

BSID-D sub-TLV contains parameters used for Binding Protection Information distribution.

The format of BSID-D sub-TLV is shown in Figure 1.

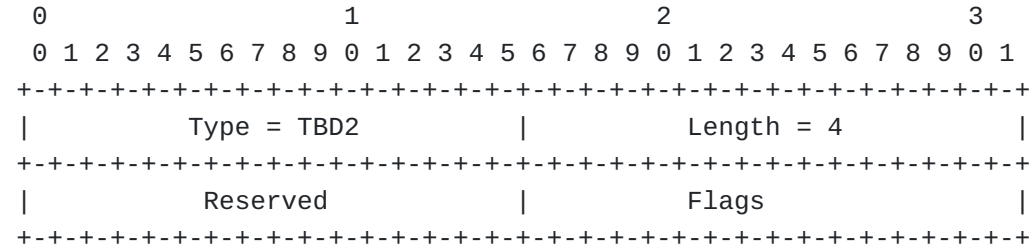

Figure 1: BSID_PROTECTION_DISTRIBUTION_CAPABILITY sub-TLV

Type: TBD2 is to be assigned by IANA.

Length: 4.

Reserved:
2 octets. Must be set to zero in transmission and ignored on reception.

Flags: 2 octets. No flag bits are defined.

A PCC, which supports the capability of Binding Protection Information Distribution, sends a PCE an Open message containing a PATH_SETUP_TYPE_CAPABILITY TLV with Path Setup Type (PST) = TBD1 and a BSID_PROTECTION_DISTRIBUTION_CAPABILITY sub-TLV. PST = TBD1 indicates that the PCC is capable of receiving and processing the binding protection information about a binding SID on a node from the PCE for protecting the binding SID of the node when the node fails.

A PCE, which supports the capability of Binding Protection Information Distribution, sends a PCC an Open message containing a PATH_SETUP_TYPE_CAPABILITY TLV with Path Setup Type (PST) = TBD1 and a BSID_PROTECTION_DISTRIBUTION_CAPABILITY sub-TLV. PST = TBD1 indicates that the PCE supports the capability of Binding Protection Information Distribution.

If both a PCC and a PCE support the capability of Binding Protection Information Distribution, each of the Open messages sent by the PCC and PCE contains a PATH-SETUP-TYPE-CAPABILITY TLV with a PST list containing PST = TBD1 and a BSID_PROTECTION_DISTRIBUTION_CAPABILITY sub-TLV.

If a PCE receives an Open message without a PATH-SETUP-TYPE-CAPABILITY TLV containing PST = TBD1 from a PCC, then the PCE MUST not send the PCC any binding protection information.

If a PCC receives an Open message without a PATH-SETUP-TYPE-CAPABILITY TLV containing PST = TBD1 from a PCE, then the PCC MUST ignore any binding protection information from the PCE.

When PCECC is used, a PCC and PCE exchange capability of binding protection information distribution using PCECC-CAPABILITY Sub-TLV which is included in the PATH_SETUP_TYPE_CAPABILITY TLV in an Open message.

A new flag bit B is defined in the Flags field of the PCECC-CAPABILITY sub-TLV as shown in Figure 2. B flag (for Binding SID Protection): if set to 1 by a PCEP speaker (PCE or PCC), the B flag indicates that the PCEP speaker supports and is willing to handle the PCECC based central controller instructions for Binding SID protection. The bit MUST be set to 1 by both a PCC and a PCE for the PCECC Binding SID protection instruction download/report on a PCEP session.

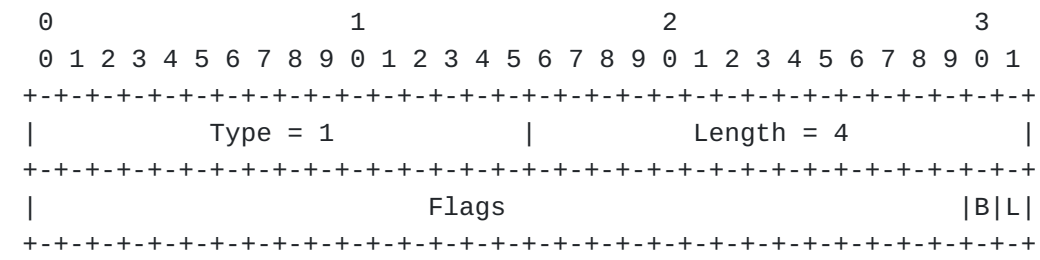

Figure 2: PCECC-Capability sub-TLV with B

2.2. Extensions to RP/SRP Object

After sending the binding to a node (i.e., the PCC running on the node), a PCE sends the corresponding binding protection information to the possible protecting nodes of the node in a PCEP message such as a Path Computation LSP Update Request (PCUpd) message. The message contains a Request Parameters (RP) object or Stateful PCE Request Parameters (SRP) object. The object includes:

- A PATH-SETUP-TYPE TLV with PST = TBD1 indicating binding protection information for a Binding SID of a node.

- A Node ID TLV containing the identifier of the node.

The format of PATH-SETUP-TYPE TLV is shown in Figure 3.

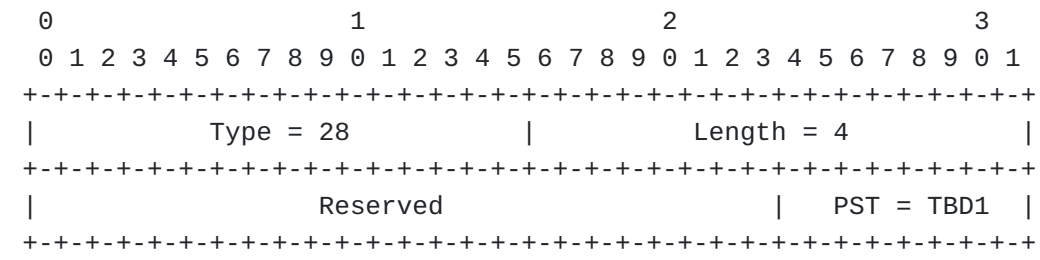

Figure 3: PATH-SETUP-TYPE TLV with PST = TBD1

The format of Node ID TLV is illustrated in Figure 4.

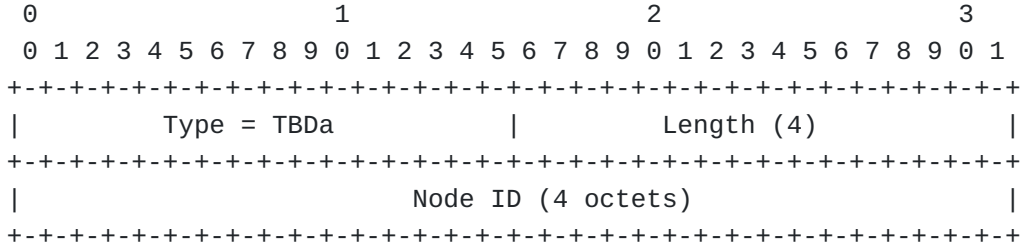

Figure 4: Node ID TLV Format

Type:

Its value (TBDa) is to be assigned by IANA.

Length: Its value indicates the length of the value field of the TLV is 4.

Node ID: 4-octet field contains the 4-octet TE router identifier (ID) of the (protected) node.

3. Procedure for Updating Information

When a PCE sends a binding to node N (i.e., PCC running on N) in a first Path Computation LSP Update Request (PCUpd) message, the PCE sends the corresponding binding protection information to the possible protecting nodes such as neighbors of node N in a second PCUpd message.

The first message contains an RP/SRP object and an LSP object with a TE-PATH-BINDING TLV. The RP/SRP object does not include any PATH-SETUP-TYPE TLV with PST = TBD1. The TLV includes a binding SID and R bit set to zero.

The second message contains an RP/SRP object and an LSP object with a TE-PATH-BINDING TLV. The RP/SRP object includes a PATH-SETUP-TYPE TLV with PST = TBD1 and a Node ID TLV indicating node N. TE-PATH-BINDING TLV includes the binding SID and R bit set to zero.

After a PCE sends the binding to node N, if the PCE removes the binding from node N through sending a third PCUpd message to node N, the PCE removes the corresponding binding protection information from the nodes through sending a fourth PCUpd message to the nodes.

The third message contains an RP/SRP object and an LSP object with a TE-PATH-BINDING TLV. The RP/SRP object does not include any PATH-SETUP-TYPE TLV with PST = TBD1. The TLV includes a binding SID and R bit set to one (1).

The fourth message contains an RP/SRP object and an LSP object with a TE-PATH-BINDING TLV. The RP/SRP object includes a PATH-SETUP-TYPE TLV with PST = TBD1 and a Node ID TLV indicating node N. TE-PATH-BINDING TLV includes the binding SID and R bit set to one (1).

After a PCE sends the binding to node N, if the PCE changes the binding in node N through sending a fifth PCUpd message to node N, the PCE changes the corresponding binding protection information in the nodes through sending a sixth PCUpd message to the them.

The fifth message contains an RP/SRP object and an LSP object with a TE-PATH-BINDING TLV. The RP/SRP object does not include any PATH-SETUP-TYPE TLV with PST = TBD1. The LSP object includes a (changed) path. The TLV includes a binding SID and R bit set to zero.

The sixth message contains an RP/SRP object and an LSP object with a TE-PATH-BINDING TLV. The RP/SRP object includes a PATH-SETUP-TYPE TLV with PST = TBD1 and a Node ID TLV indicating node N. The LSP object includes the (changed) path. TE-PATH-BINDING TLV includes the binding SID and R bit set to zero.

4. References

4.1. Normative References

[I-D.ietf-spring-segment-protection-sr-te-paths]
Hegde, S., Bowers, C., Litkowski, S., Xu, X., and F. Xu, "Segment Protection for SR-TE Paths", Work in Progress, Internet-Draft, draft-ietf-spring-segment-protection-sr-te-paths-06, 9 February 2024, <https://datatracker.ietf.org/doc/html/draft-ietf-spring-segment-protection-sr-te-paths-06>.

[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>.

4.2. Informative References

[I-D.hu-spring-segment-routing-proxy-forwarding]
Hu, Z., Chen, H., Yao, J., Bowers, C., Zhu, Y., and Y. Liu, "SR-TE Path Midpoint Restoration", Work in Progress, Internet-Draft, draft-hu-spring-segment-routing-proxy-forwarding-24, 21 August 2023, <https://datatracker.ietf.org/doc/html/draft-hu-spring-segment-routing-proxy-forwarding-24>.

[I-D.ietf-pce-binding-label-sid]
Sivabalan, S., Filsfils, C., Tantsura, J., Previdi, S., and C. Li, "Carrying Binding Label/Segment Identifier (SID) in PCE-based Networks.", Work in Progress, Internet-Draft, draft-ietf-pce-binding-label-sid-16, 27 March 2023, <https://datatracker.ietf.org/doc/html/draft-ietf-pce-binding-label-sid-16>.

Acknowledgments

The authors would like to thank Andrew Stone, and Dhruv Dhody for their comments to this work.

Authors' Addresses

Huaimo Chen
Futurewei

Boston, MA,
United States of America

Email: hchen.ietf@gmail.com

Bruno Decraene
Orange
France

Email: bruno.decraene@orange.com

Gyan S. Mishra
Verizon
13101 Columbia Pike
Silver Spring, MD 20904
United States of America

Phone: 301 502-1347
Email: gyan.s.mishra@verizon.com

Aijun Wang
China Telecom
Beiqijia Town, Changping District
Beijing
102209
China

Email: wangaj3@chinatelecom.cn

Xufeng Liu
IBM Corporation
United States of America

Email: xufeng.liu.ietf@gmail.com

Lei Liu
Fujitsu
United States of America

Email: liulei.kddi@gmail.com