

PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 9, 2020

S. Sivabalan
C. Filsfils
Cisco Systems, Inc.
J. Tantsura
Apstra, Inc.
J. Hardwick
Metaswitch Networks
S. Previdi
C. Li
Huawei Technologies
September 6, 2019

**Carrying Binding Label/Segment-ID in PCE-based Networks.
draft-ietf-pce-binding-label-sid-00**

Abstract

In order to provide greater scalability, network opacity, and service independence, SR utilizes a Binding Segment Identifier (BSID). It is possible to associate a BSID to RSVP-TE signaled Traffic Engineering Label Switching Path or binding Segment-ID (SID) to Segment Routed (SR) Traffic Engineering path. Such a binding label/SID can be used by an upstream node for steering traffic into the appropriate TE path to enforce SR policies. This document proposes an approach for reporting binding label/SID to Path Computation Element (PCE) for supporting PCE-based Traffic Engineering policies.

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.

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 March 9, 2020.

Copyright Notice

Copyright (c) 2019 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 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	3
2.	Terminology	5
3.	Path Binding TLV	6
4.	Operation	7
5.	Binding SID in SR-ERO	8
6.	Binding SID in SRv6-ERO/	8
7.	Implementation Status	8
7.1.	Huawei	9
8.	Security Considerations	9
9.	Manageability Considerations	10
9.1.	Control of Function and Policy	10
9.2.	Information and Data Models	10
9.3.	Liveness Detection and Monitoring	10
9.4.	Verify Correct Operations	10
9.5.	Requirements On Other Protocols	10
9.6.	Impact On Network Operations	10
10.	IANA Considerations	11
10.1.	PCEP TLV Type Indicators	11
10.1.1.	TE-PATH-BINDING TLV	11
10.2.	PCEP Error Type and Value	11
11.	Acknowledgements	11
12.	References	12
12.1.	Normative References	12
12.2.	Informative References	13
Appendix A.	Contributor Addresses	15
	Authors' Addresses	15

1. Introduction

A PCE can compute Traffic Engineering paths (TE paths) through a network that are subject to various constraints. Currently, TE paths are either set up using the RSVP-TE signaling protocol or Segment Routing (SR). We refer to such paths as RSVP-TE paths and SR-TE paths respectively in this document.

As per [RFC8402] SR allows a headend node to steer a packet flow along any path. The headend node is said to steer a flow into an Segment Routing Policy (SR Policy). Further, as per [I-D.ietf-spring-segment-routing-policy], an SR Policy is a framework that enables instantiation of an ordered list of segments on a node for implementing a source routing policy with a specific intent for traffic steering from that node.

As described in [RFC8402], Binding Segment Identifier (BSID) is bound to an Segment Routed (SR) Policy, instantiation of which may involve a list of SIDs. Any packets received with an active segment equal to BSID are steered onto the bound SR Policy. A BSID may be either a local (SRLB) or a global (SRGB) SID. As per Section 6.4 of [I-D.ietf-spring-segment-routing-policy] a BSID can also be associated with any type of interfaces or tunnel to enable the use of a non-SR interface or tunnels as segments in a SID-list.

[RFC5440] describes the Path Computation Element Protocol (PCEP) for communication between a Path Computation Client (PCC) and a PCE or between a pair of PCEs as per [RFC4655]. [RFC8231] specifies extension to PCEP that allows a PCC to delegate its LSPs to a stateful PCE. A stateful PCE can then update the state of LSPs delegated to it. [RFC8281] specifies a mechanism allowing a PCE to dynamically instantiate an LSP on a PCC by sending the path and characteristics. The PCEP extension to setup and maintain SR-TE paths is specified in [I-D.ietf-pce-segment-routing].

[I-D.ietf-pce-segment-routing] provides a mechanism for a network controller (acting as a PCE) to instantiate candidate paths for an SR Policy onto a head-end node (acting as a PCC) using PCEP. For more information on the SR Policy Architecture, see [I-D.ietf-spring-segment-routing-policy].

Binding label/SID has local significance to the ingress node of the corresponding TE path. When a stateful PCE is deployed for setting up TE paths, it may be desirable to report the binding label or SID to the stateful PCE for the purpose of enforcing end-to-end TE/SR policy. A sample Data Center (DC) use-case is illustrated in the following diagram. In the MPLS DC network, an SR LSP (without traffic engineering) is established using a prefix SID advertised by

In this document, we introduce a new OPTIONAL TLV that a PCC can use in order to report the binding label/SID associated with a TE LSP, or

a PCE to request a PCC to allocate a specific binding label/SID value. This TLV is intended for TE LSPs established using RSVP-TE, SR, or any other future method. Also, in the case of SR-TE LSPs, the TLV can carry a binding MPLS label (for SR-TE path with MPLS data-plane) or a binding IPv6 SID (e.g., IPv6 address for SR-TE paths with IPv6 data-plane). Binding value means either MPLS label or SID throughout this document.

Additionally, to support the PCE based central controller [[RFC8283](#)] operation where the PCE would take responsibility for managing some part of the MPLS label space for each of the routers that it controls, the PCE could directly make the binding label/SID allocation and inform the PCC. See [[I-D.ietf-pce-pcep-extension-for-pce-controller](#)] for details.

2. Terminology

The following terminologies are used in this document:

BSID: Binding Segment Identifier.

LER: Label Edge Router.

LSP: Label Switched Path.

LSR: Label Switching Router.

PCC: Path Computation Client.

PCE: Path Computation Element

PCEP: Path Computation Element Protocol.

RSVP-TE: Resource ReserVation Protocol-Traffic Engineering.

SID: Segment Identifier.

SR: Segment Routing.

SRGB: Segment Routing Global Block.

SRLB: Segment Routing Local Block.

TLV: Type, Length, and Value.

3. Path Binding TLV

The new optional TLV is called "TE-PATH-BINDING TLV" whose format is shown in the diagram below is defined to carry binding label or SID for a TE path. This TLV is associated with the LSP object specified in ([RFC8231]). The type of this TLV is to be allocated by IANA.




Figure 2: TE-PATH-BINDING TLV

TE-PATH-BINDING TLV is a generic TLV such that it is able to carry MPLS label binding as well as SRv6 Binding SID. It is formatted according to the rules specified in [RFC5440].

Binding Type (BT): A one byte field identifies the type of binding included in the TLV. This document specifies the following BT values:

- o BT = 0: The binding value is an MPLS label carried in the format specified in [RFC5462] where only the label value is valid, and other fields (TC, S, and TTL) fields MUST be considered invalid. The Length MUST be set to 6.
- o BT = 1: Similar to the case where BT is 0 except that all the fields on the MPLS label entry are set on transmission. However, the receiver MAY choose to override TC, S, and TTL values according its local policy.
- o BT = 2: The binding value is a SRv6 SID with a format of an 16 byte IPv6 address, representing the binding SID for SRv6.

Reserved: MUST be set to 0 while sending and ignored on receipt.

Binding Value: A variable length field, padded with trailing zeros to a 4-byte boundary. For the BT as 0, the 20 bits represents the MPLS label. For the BT as 1, the 32-bits represents the label stack entry as per [RFC5462]. For the BT as 2, the 128-bits represent the SRv6 SID.

4. Operation

The binding value is allocated by the PCC and reported to a PCE via PCRpt message. If a PCE does not recognize the TE-PATH-BINDING TLV, it MUST ignore the TLV in accordance with ([RFC5440]). If a PCE recognizes the TLV but does not support the TLV, it MUST send PCErr with Error-Type = 2 (Capability not supported).

If a TE-PATH-BINDING TLV is absent in PCRpt message, PCE MUST assume that the corresponding LSP does not have any binding. If there are more than one TE-PATH-BINDING TLVs, only the first TLV MUST be processed and the rest MUST be silently ignored. If a PCE recognizes an invalid binding value (e.g., label value from the reserved label space when MPLS label binding is used), it MUST send the PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error Value = TBD ("Bad label value") as specified in [I-D.ietf-pce-segment-routing].

If a PCE requires a PCC to allocate a specific binding value, it may do so by sending a PCUpd or PCInitiate message containing a TE-PATH-BINDING TLV. If the value can be successfully allocated, the PCC reports the binding value to the PCE. If the PCC considers the binding value specified by the PCE invalid, it MUST send a PCErr message with Error-Type = TBD ("Binding label/SID failure") and Error Value = TBD ("Invalid SID"). If the binding value is valid, but the PCC is unable to allocate the binding value, it MUST send a PCErr message with Error-Type = TBD ("Binding label/SID failure") and Error Value = TBD ("Unable to allocate the specified label/SID").

If a PCC receives TE-PATH-BINDING TLV in any message other than PCUpd or PCInitiate, it MUST close the corresponding PCEP session with the reason "Reception of a malformed PCEP message" (according to [RFC5440]). Similarly, if a PCE receives a TE-PATH-BINDING TLV in any message other than a PCRpt or if the TE-PATH-BINDING TLV is associated with any object other than LSP object, the PCE MUST close the corresponding PCEP session with the reason "Reception of a malformed PCEP message" (according to [RFC5440]).

If a PCC wishes to withdraw or modify a previously reported binding value, it MUST send a PCRpt message without any TE-PATH-BINDING TLV or with the TE-PATH-BINDING TLV containing the new binding value respectively.

If a PCE wishes to modify a previously requested binding value, it MUST send a PCUpd message with TE-PATH-BINDING TLV containing the new binding value. Absence of TE-PATH-BINDING TLV in PCUpd message means that the PCE does not specify a binding value in which case the binding value allocation is governed by the PCC's local policy.

If a PCC receives a valid binding value from a PCE which is different than the current binding value, it MUST try to allocate the new value. If the new binding value is successfully allocated, the PCC MUST report the new value to the PCE. Otherwise, it MUST send a PCErr message with Error-Type = TBD ("Binding label/SID failure") and Error Value = TBD ("Unable to allocate the specified label/SID").

In some cases, a stateful PCE can request the PCC to allocate a binding value. It may do so by sending a PCUpd message containing an empty TE-PATH-BINDING TLV, i.e., no binding value is specified (making the length field of the TLV as 2). A PCE can also make the request PCC to allocate a binding at the time of initiation by sending a PCInitiate message with an empty TE-PATH-BINDING TLV.

5. Binding SID in SR-ERO

In PCEP messages, LSP route information is carried in the Explicit Route Object (ERO), which consists of a sequence of subobjects. [I-D.ietf-pce-segment-routing] defines a new ERO subobject "SR-ERO subobject" capable of carrying a SID as well as the identity of the node/adjacency (NAI) represented by the SID. The NAI Type (NT) field indicates the type and format of the NAI contained in the SR-ERO. In case of binding SID, the NAI MUST NOT be included and NT MUST be set to zero. So as per Section 5.2.1 of [I-D.ietf-pce-segment-routing], for NT=0, the F bit MUST be 1, the S bit needs to be zero and the Length MUST be 8. Further the M bit MUST be set. If these conditions are not met, the entire ERO MUST be considered invalid and a PCErr message is sent with Error-Type = 10 ("Reception of an invalid object") and Error-Value = 11 ("Malformed object").

6. Binding SID in SRv6-ERO/

[I-D.ietf-pce-segment-routing] defines a new ERO subobject "SRv6-ERO subobject" for SRv6 SID. The NAI MUST NOT be included and NT MUST be set to zero. So as per Section 5.2.1 of [I-D.ietf-pce-segment-routing], for NT=0, the F bit MUST be 1, the S bit needs to be zero and the Length MUST be 24. If these conditions are not met, the entire ERO is considered invalid and a PCErr message is sent with Error-Type = 10 ("Reception of an invalid object") and Error-Value = 11 ("Malformed object") (as per [I-D.ietf-pce-segment-routing]).

7. 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 catalog 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".

[7.1.](#) Huawei

- o Organization: Huawei
- o Implementation: Huawei's Router and Controller
- o Description: An experimental code-point is used and plan to request early code-point allocation from IANA after WG adoption.
- o Maturity Level: Production
- o Coverage: Full
- o Contact: mahendrasingh@huawei.com

[8.](#) Security Considerations

The security considerations described in [\[RFC5440\]](#), [\[RFC8231\]](#), [\[RFC8281\]](#) and [\[I-D.ietf-pce-segment-routing\]](#) are applicable to this specification. No additional security measure is required.

As described [\[I-D.ietf-pce-segment-routing\]](#), SR allows a network controller to instantiate and control paths in the network. A rouge PCE can manipulate binding SID allocations to move traffic around for some other LSPs that uses BSID in its SR-ERO.

Thus, as per [\[RFC8231\]](#), it is RECOMMENDED that these PCEP extensions only be activated on authenticated and encrypted sessions across PCEs

and PCCs belonging to the same administrative authority, using Transport Layer Security (TLS) [RFC8253], as per the recommendations and best current practices in [RFC7525] (unless explicitly set aside in [RFC8253]).

9. Manageability Considerations

All manageability requirements and considerations listed in [RFC5440], [RFC8231], and [I-D.ietf-pce-segment-routing] apply to PCEP protocol extensions defined in this document. In addition, requirements and considerations listed in this section apply.

9.1. Control of Function and Policy

A PCC implementation SHOULD allow the operator to configure the policy based on which PCC needs to allocate the binding label/SID.

9.2. Information and Data Models

The PCEP YANG module [I-D.ietf-pce-pcep-yang] could be extended to include policy configuration for binding label/SID allocation.

9.3. 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].

9.4. Verify Correct Operations

Mechanisms defined in this document do not imply any new operation verification requirements in addition to those already listed in [RFC5440], [RFC8231], and [I-D.ietf-pce-segment-routing].

9.5. Requirements On Other Protocols

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

9.6. Impact On Network Operations

Mechanisms defined in [RFC5440], [RFC8231], and [I-D.ietf-pce-segment-routing] also apply to PCEP extensions defined in this document. Further, the mechanism described in this document can help the operator to request control of the LSPs at a particular PCE.

[10.](#) IANA Considerations

[10.1.](#) PCEP TLV Type Indicators

This document defines a new PCEP TLV; IANA is requested to make the following allocations from the "PCEP TLV Type Indicators" sub-registry of the PCEP Numbers registry, as follows:

Value	Name	Reference
TBD	TE-PATH-BINDING	This document

[10.1.1.](#) TE-PATH-BINDING TLV

IANA is requested to create a sub-registry to manage the value of the Binding Type field in the TE-PATH-BINDING TLV.

Value	Description	Reference
0	MPLS Label	This document
1	MPLS Label Stack Entry	This document
2	SRV6 SID	This document

[10.2.](#) PCEP Error Type and Value

This document defines a new Error-type and Error-Values for the PCErr message. IANA is requested to allocate new error-type and error-values within the "PCEP-ERROR Object Error Types and Values" subregistry of the PCEP Numbers registry, as follows:

Error-Type	Meaning
-----	-----
TBD	Binding label/SID failure:
	Error-value = TBD: Invalid SID
	Error-value = TBD: Unable to allocate the specified label/SID

[11.](#) Acknowledgements

We like to thank Milos Fabian for his valuable comments.

12. References

12.1. 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>>.
- [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>>.
- [RFC5462] Andersson, L. and R. Asati, "Multiprotocol Label Switching (MPLS) Label Stack Entry: "EXP" Field Renamed to "Traffic Class" Field", [RFC 5462](#), DOI 10.17487/RFC5462, February 2009, <<https://www.rfc-editor.org/info/rfc5462>>.
- [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>>.
- [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>>.
- [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>>.
- [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>>.
- [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-16](#) (work in progress), March 2019.

[12.2. Informative References](#)

- [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>>.
- [RFC8283] Farrel, A., Ed., Zhao, Q., Ed., Li, Z., and C. Zhou, "An Architecture for Use of PCE and the PCE Communication Protocol (PCEP) in a Network with Central Control", [RFC 8283](#), DOI 10.17487/RFC8283, December 2017, <<https://www.rfc-editor.org/info/rfc8283>>.
- [I-D.ietf-spring-segment-routing-policy]
Filsfils, C., Sivabalan, S., daniel.voyer@bell.ca, d., bogdanov@google.com, b., and P. Mattes, "Segment Routing Policy Architecture", [draft-ietf-spring-segment-routing-policy-03](#) (work in progress), May 2019.
- [I-D.ietf-idr-bgp-prefix-sid]
Previdi, S., Filsfils, C., Lindem, A., Sreekantiah, A., and H. Gredler, "Segment Routing Prefix SID extensions for BGP", [draft-ietf-idr-bgp-prefix-sid-27](#) (work in progress), June 2018.
- [I-D.ietf-pce-pcep-extension-for-pce-controller]
Zhao, Q., Li, Z., Negi, M., and C. Zhou, "PCEP Procedures and Protocol Extensions for Using PCE as a Central Controller (PCECC) of LSPs", [draft-ietf-pce-pcep-extension-for-pce-controller-02](#) (work in progress), July 2019.

[I-D.ietf-pce-pcep-yang]

Dhody, D., Hardwick, J., Beeram, V., and J. Tantsura, "A YANG Data Model for Path Computation Element Communications Protocol (PCEP)", [draft-ietf-pce-pcep-yang-12](#) (work in progress), July 2019.

Appendix A. Contributor Addresses

Dhruv Dhody
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore, Karnataka 560066
India

EMail: dhruv.ietf@gmail.com

Mahendra Singh Negi

EMail: mahend.ietf@gmail.com

Authors' Addresses

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

EMail: msiva@cisco.com

Clarence Filsfils
Cisco Systems, Inc.
Pegasus Parc
De kleetlaan 6a, DIEGEM BRABANT 1831
BELGIUM

EMail: cfilsfil@cisco.com

Jeff Tantsura
Apstra, Inc.

EMail: jefftant.ietf@gmail.com

Jonathan Hardwick
Metaswitch Networks
100 Church Street
Enfield, Middlesex
UK

EMail: Jonathan.Hardwick@metaswitch.com

Stefano Previdi
Huawei Technologies

E-Mail: stefano@previdi.net

Cheng Li
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing 100095
China

E-Mail: chengli13@huawei.com