

Internet Engineering Task Force (IETF)
Request for Comments: 9604
Category: Standards Track
ISSN: 2070-1721

S. Sivabalan
Ciena Corporation
C. Filsfils
Cisco Systems, Inc.
J. Tantsura
Nvidia
S. Previdi
C. Li, Ed.
Huawei Technologies
August 2024

Carrying Binding Label/SID in PCE-Based Networks

Abstract

In order to provide greater scalability, network confidentiality, and service independence, Segment Routing (SR) utilizes a Binding Segment Identifier (BSID), as described in RFC 8402. It is possible to associate a BSID to an RSVP-TE-signaled Traffic Engineering (TE) Label Switched Path (LSP) or an SR TE path. The BSID can be used by an upstream node for steering traffic into the appropriate TE path to enforce SR policies. This document specifies the concept of binding value, which can be either an MPLS label or a Segment Identifier (SID). It further specifies an extension to Path Computation Element Communication Protocol (PCEP) for reporting the binding value by a Path Computation Client (PCC) to the Path Computation Element (PCE) to support PCE-based TE policies.

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/rfc9604>.

Copyright Notice

Copyright (c) 2024 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
 - 1.1. Motivation and Example

1.2.	Summary of the Extension
2.	Requirements Language
3.	Terminology
4.	Path Binding TLV
4.1.	SRv6 Endpoint Behavior and SID Structure
5.	Operation
6.	Binding SID in SR-ERO
7.	Binding SID in SRv6-ERO
8.	PCE Allocation of Binding Label/SID
9.	Security Considerations
10.	Manageability Considerations
10.1.	Control of Function and Policy
10.2.	Information and Data Models
10.3.	Liveness Detection and Monitoring
10.4.	Verify Correct Operations
10.5.	Requirements on Other Protocols
10.6.	Impact on Network Operations
11.	IANA Considerations
11.1.	PCEP TLV Type Indicators
11.1.1.	TE-PATH-BINDING TLV
11.2.	LSP Object
11.3.	PCEP Error Type and Value
12.	References
12.1.	Normative References
12.2.	Informative References
	Acknowledgements
	Contributors
	Authors' Addresses

1. Introduction

A Path Computation Element (PCE) can compute Traffic Engineering (TE) paths through a network where those paths are subject to various constraints. Currently, TE paths are set up using either 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 head-end node to steer a packet flow along a given path via an SR Policy. As per [RFC9256], an SR Policy is a framework that enables the 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], a Binding SID (BSID) is bound to an SR Policy, instantiation of which may involve a list of Segment Identifiers (SIDs). Any packets received with an active segment equal to a BSID are steered onto the bound SR Policy. A BSID may be either a local (SR Local Block (SRLB)) or a global (SR Global Block (SRGB)) SID. As per Section 6.4 of [RFC9256], a BSID can also be associated with any type of interface or tunnel to enable the use of a non-SR interface or tunnel as a segment in a SID list. In this document, the term "binding label/SID" is used to generalize the allocation of a binding value for both SR and non-SR paths.

[RFC5440] describes the PCEP for communication between a Path Computation Client (PCC) and a PCE or between a pair of PCEs as per [RFC4655]. [RFC8231] specifies extensions to PCEP that allow a PCC to delegate its Label Switched Paths (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. This document specifies an extension to PCEP to manage the binding of label/SID that can be applied to SR, RSVP-TE, and other path setup types.

[RFC8664] provides a mechanism for a PCE (acting as a network

controller) to instantiate SR-TE paths (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 [RFC9256].

1.1. Motivation and Example

A 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, a binding label/SID reported from the PCC to the stateful PCE is useful for enforcing an end-to-end TE/SR policy. A sample Data Center (DC) and IP/MPLS WAN use case is illustrated in Figure 1 with a multi-domain PCE. In the IP/MPLS WAN, an SR-TE LSP is set up using the PCE. The list of SIDs of the SR-TE LSP is {A, B, C, D}. The gateway Node-1 (which is the PCC) allocates a binding SID X and reports it to the PCE. In the MPLS DC network, an end-to-end SR-TE LSP is established. In order for the access node to steer the traffic towards Node-1 and over the SR-TE path in WAN, the PCE passes the SID stack {Y, X} where Y is the node SID of the gateway Node-1 to the access node and X is the BSID. In the absence of the BSID X, the PCE would need to pass the SID stack {Y, A, B, C, D} to the access node. This example also illustrates the additional benefit of using the binding label/SID to reduce the number of SIDs imposed by the access nodes with a limited forwarding capacity.

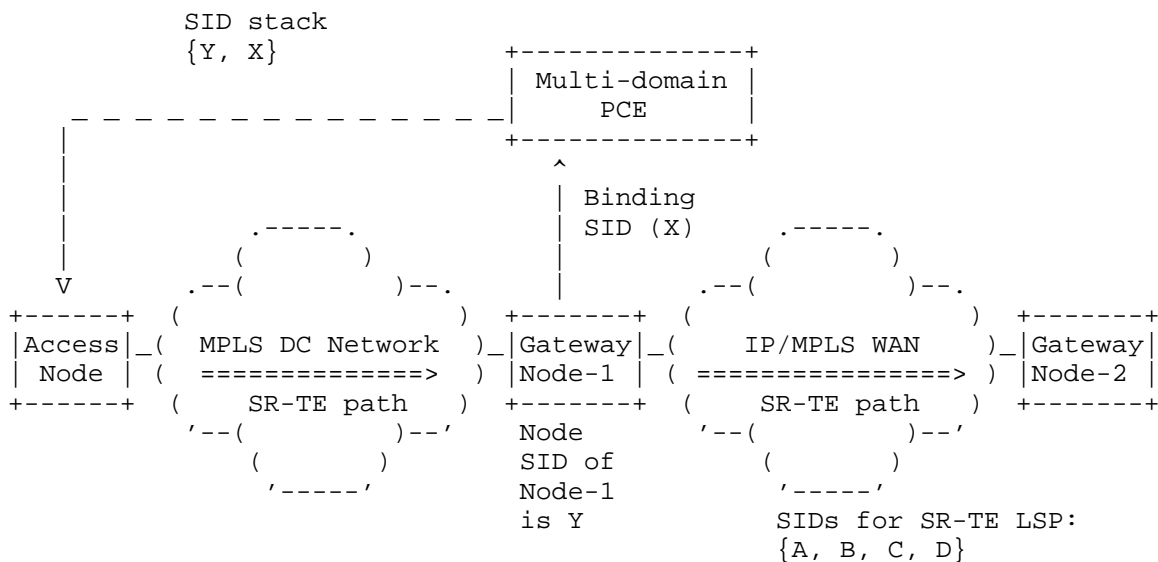


Figure 1: A Sample Use Case of Binding SID

Using the extension defined in this document, a PCC could report to the stateful PCE the binding label/SID it allocated via a Path Computation LSP State Report (PCRpt) message. It is also possible for a stateful PCE to request a PCC to allocate a specific binding label/SID by sending a Path Computation LSP Update Request (PCUpd) message. If the PCC can successfully allocate the specified binding value, it reports the binding value to the PCE. Otherwise, the PCC sends an error message to the PCE indicating the cause of the failure. A local policy or configuration at the PCC SHOULD dictate if the binding label/SID needs to be assigned.

1.2. Summary of the Extension

To implement the needed changes to PCEP, this document introduces a new OPTIONAL TLV that allows a PCC to report the binding label/SID associated with a TE LSP or a PCE to request a PCC to allocate any or a specific binding label/SID value. This TLV is intended for TE LSPs established using RSVP-TE, SR-TE, or any other future method. In the case of SR-TE LSPs, the TLV can carry a binding label (for SR-TE paths with the MPLS data plane) or a binding IPv6 SID (e.g., IPv6

address for SR-TE paths with the IPv6 data plane). Throughout this document, the term "binding value" means either an MPLS label or a SID.

As another way to use the extension specified in this document, 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 Section 8 for details.

In addition to specifying a new TLV, this document specifies how and when a PCC and PCE can use this TLV, how they can allocate a binding label/SID, and the associated error handling.

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

The following terminologies are used in this document:

BSID: Binding SID

binding label/SID: a generic term used for the binding segment for both SR and non-SR paths

binding value: a generic term used for the binding segment as it can be encoded in various formats (as per the Binding Type (BT))

LSP: Label Switched Path

PCC: Path Computation Client

PCEP: Path Computation Element Communication Protocol

RSVP-TE: Resource Reservation Protocol - Traffic Engineering

SID: Segment Identifier

SR: Segment Routing

4. Path Binding TLV

The new optional TLV called "TE-PATH-BINDING TLV" (the format is shown in Figure 2) is defined to carry the binding label/SID for a TE path. This TLV is associated with the LSP object specified in [RFC8231]. This TLV can also be carried in the PCEP-ERROR object [RFC5440] in case of error. Multiple instances of TE-PATH-BINDING TLVs MAY be present in the LSP and PCEP-ERROR object. The type of this TLV is 55. The length is variable.

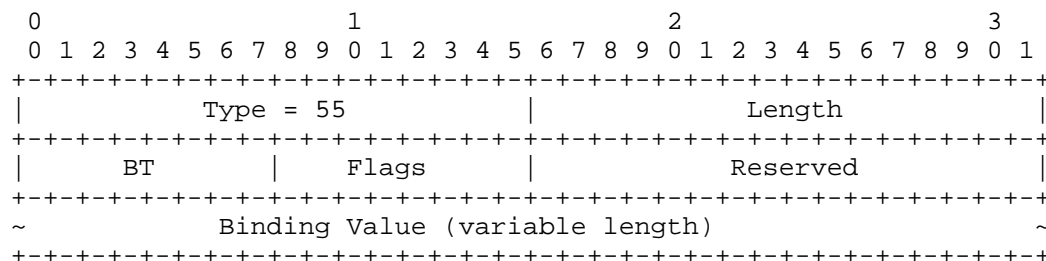


Figure 2: TE-PATH-BINDING TLV

The TE-PATH-BINDING TLV is a generic TLV such that it is able to carry binding label/SID (i.e., MPLS label or SRv6 SID). It is formatted according to the rules specified in [RFC5440]. The value portion of the TLV comprises:

- * Binding Type (BT): A one-octet field that identifies the type of binding included in the TLV. This document specifies the following BT values:
 - BT = 0: The binding value is a 20-bit MPLS label value. The TLV is padded to 4-bytes alignment. The Length MUST be set to 7 (the padding is not included in the length, as per [RFC5440], Section 7.1), and the first 20 bits are used to encode the MPLS label value.
 - BT = 1: The binding value is a 32-bit MPLS Label Stack Entry as per [RFC3032] with Label, Traffic Class (TC) [RFC5462], S, and TTL values encoded. Note that the receiver MAY choose to override TC, S, and TTL values according to its local policy. The Length MUST be set to 8.
 - BT = 2: The binding value is an SRv6 SID with the format of a 16-octet IPv6 address, representing the binding SID for SRv6. The Length MUST be set to 20.
 - BT = 3: The binding value is a 24-octet field, defined in Section 4.1, that contains the SRv6 SID as well as its Behavior and Structure. The Length MUST be set to 28.

Section 11.1.1 defines the IANA registry used to maintain these binding types as well as any future ones. Note that multiple TE-PATH-BINDING TLVs with the same or different binding types MAY be present for the same LSP. A PCEP speaker could allocate multiple TE-PATH-BINDING TLVs (of the same BT) and use different binding values in different domains or use cases based on a local policy.

- * Flags: 1 octet of flags. The following flag is defined in the new "TE-PATH-BINDING TLV Flag field" registry as described in Section 11.1.1:

```
  0 1 2 3 4 5 6 7
+---+---+---+---+
|R|               |
+---+---+---+---+
```

Figure 3: Flags

Where:

- R (Removal - 1 bit): When set, the requesting PCEP peer requires the removal of the binding value for the LSP. When unset, the PCEP peer indicates that the binding value is added or retained for the LSP. This flag is used in the PCRpt and PCUpd messages. It is ignored in other PCEP messages.
- The unassigned flags MUST be set to 0 while sending and ignored on receipt.
- * 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-octet boundary. When the BT is 0, the 20 bits represent the MPLS label. When the BT is 1, the 32 bits represent the MPLS

label stack entry as per [RFC3032]. When the BT is 2, the 128 bits represent the SRv6 SID. When the BT is 3, the binding value also contains the SRv6 Endpoint Behavior and SID Structure, defined in Section 4.1. In this document, the TE-PATH-BINDING TLV is considered to be empty if no binding value is present. Note that the length of the TLV would be 4 in such a case.

4.1. SRv6 Endpoint Behavior and SID Structure

This section specifies the format of the binding value in the TE-PATH-BINDING TLV when the BT is set to 3 for the SRv6 Binding SIDs [RFC8986]. The format is shown in Figure 4.

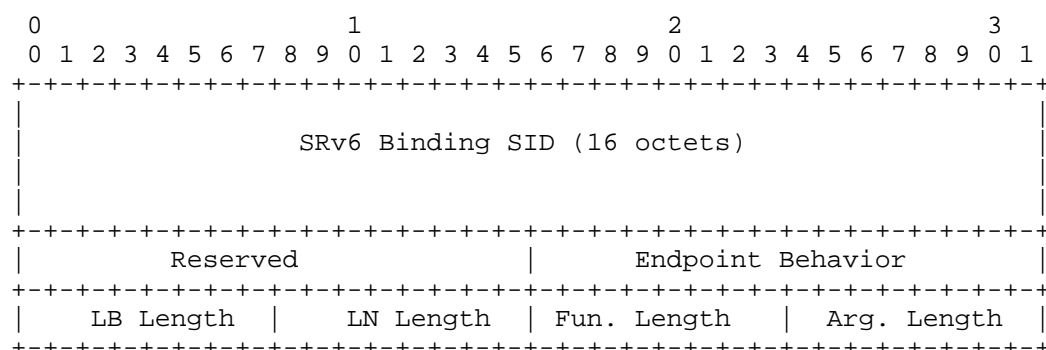


Figure 4: SRv6 Endpoint Behavior and SID Structure

The Binding Value consists of:

- * SRv6 Binding SID: 16 octets. The 128-bit IPv6 address, representing the binding SID for SRv6.
- * Reserved: 2 octets. It MUST be set to 0 on transmit and ignored on receipt.
- * Endpoint Behavior: 2 octets. The Endpoint Behavior code point for this SRv6 SID as defined by the "SRv6 Endpoint Behaviors" registry [RFC8986]. When the field is set with the value 0, the Endpoint Behavior is considered unknown.
- * [RFC8986] defines an SRv6 SID as consisting of LOC:FUNCT:ARG, where a locator (LOC) is encoded in the L most significant bits of the SID, followed by F bits of function (FUNCT) and A bits of arguments (ARG). A locator may be represented as B:N, where B is the SRv6 SID locator block (IPv6 prefix allocated for SRv6 SIDs by the operator) and N is the identifier of the parent node instantiating the SID, called "locator node". The following fields are used to advertise the length of each individual part of the SRv6 SID:
 - LB Length: 1 octet. SRv6 SID Locator Block length in bits.
 - LN Length: 1 octet. SRv6 SID Locator Node length in bits.
 - Function Length: 1 octet. SRv6 SID Function length in bits.
 - Arguments Length: 1 octet. SRv6 SID Arguments length in bits.

The total of the locator block, locator node, function, and arguments lengths MUST be less than or equal to 128 bits. If this condition is not met, the corresponding TE-PATH-BINDING TLV is considered invalid. Also, if the Endpoint Behavior is found to be unknown or inconsistent, it is considered invalid. A PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 37 ("Invalid SRv6 SID Structure") MUST be sent in such cases.

The SRv6 SID Structure could be used by the PCE for ease of operations and monitoring. For example, this information could be used for validation of SRv6 SIDs being instantiated in the network and checked for conformance to the SRv6 SID allocation scheme chosen by the operator as described in Section 3.2 of [RFC8986]. In the future, PCE could also be used for verification and for automatically securing the SRv6 domain by provisioning filtering rules at SR domain boundaries as described in Section 5 of [RFC8754]. The details of these potential applications are outside the scope of this document.

5. Operation

The binding value is usually allocated by the PCC and reported to a PCE via a PCRpt message (see Section 8 where PCE performs the allocation). If a PCE does not recognize the TE-PATH-BINDING TLV, it would ignore the TLV in accordance with [RFC5440]. If a PCE recognizes the TLV but does not support the TLV, it MUST send a PCErr with Error-Type = 2 ("Capability not supported").

Multiple TE-PATH-BINDING TLVs are allowed to be present in the same LSP object. This signifies the presence of multiple binding SIDs for the given LSP. In the case of multiple TE-PATH-BINDING TLVs, the existing instances of TE-PATH-BINDING TLVs MAY be included in the LSP object. In case of an error condition, the whole message is rejected, and the resulting PCErr message MAY include the offending TE-PATH-BINDING TLV in the PCEP-ERROR object.

If a PCE recognizes an invalid binding value (e.g., label value from the reserved MPLS label space), it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-value = 2 ("Bad label value") as specified in [RFC8664].

For SRv6 BSIDs, it is RECOMMENDED to always explicitly specify the SRv6 Endpoint Behavior and SID Structure in the TE-PATH-BINDING TLV by setting BT to 3. This can enable the sender to have control of the SRv6 Endpoint Behavior and SID Structure. A sender MAY choose to set the BT to 2, in which case the receiving implementation chooses how to interpret the SRv6 Endpoint Behavior and SID Structure according to local policy.

If a PCC wishes to withdraw a previously reported binding value, it MUST send a PCRpt message with the specific TE-PATH-BINDING TLV with R flag set to 1. If a PCC wishes to modify a previously reported binding, it MUST withdraw the former binding value (with R flag set in the former TE-PATH-BINDING TLV) and include a new TE-PATH-BINDING TLV containing the new binding value. Note that other instances of TE-PATH-BINDING TLVs that are unchanged MAY also be included. If the unchanged instances are not included, they will remain associated with the LSP.

If a PCE requires a PCC to allocate one (or several) specific binding value(s), it may do so by sending a PCUpd or PCInitiate message containing one or more TE-PATH-BINDING TLVs. If the values can be successfully allocated, the PCC reports the binding values to the PCE. If the PCC considers the binding value specified by the PCE invalid, it MUST send a PCErr message with Error-Type = 32 ("Binding label/SID failure") and Error-value = 1 ("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 = 32 ("Binding label/SID failure") and Error-value = 2 ("Unable to allocate the specified binding value"). Note that, in case of an error, the PCC rejects the PCUpd or PCInitiate message in its entirety and can include the offending TE-PATH-BINDING TLV in the PCEP-ERROR object.

If a PCE wishes to request the withdrawal of a previously reported

binding value, it MUST send a PCUpd message with the specific TE-PATH-BINDING TLV with R flag set to 1. If a PCE wishes to modify a previously requested binding value, it MUST request the withdrawal of the former binding value (with R flag set in the former TE-PATH-BINDING TLV) and include a new TE-PATH-BINDING TLV containing the new binding value. If a PCC receives a PCUpd message with TE-PATH-BINDING TLV where the R flag is set to 1, but either the binding value is missing (empty TE-PATH-BINDING TLV) or the binding value is incorrect, it MUST send a PCErr message with Error-Type = 32 ("Binding label/SID failure") and Error-value = 4 ("Unable to remove the binding value").

In some cases, a stateful PCE may want to request that the PCC allocate a binding value of the PCC's own choosing. It instructs the PCC by sending a PCUpd message containing an empty TE-PATH-BINDING TLV, i.e., no binding value is specified (bringing the Length field of the TLV to 4). A PCE can also request that a PCC allocate a binding value at the time of initiation by sending a PCInitiate message with an empty TE-PATH-BINDING TLV. Only one such instance of empty TE-PATH-BINDING TLV, per BT, SHOULD be included in the LSP object; others should be ignored on receipt. If the PCC is unable to allocate a new binding value as per the specified BT, it MUST send a PCErr message with Error-Type = 32 ("Binding label/SID failure") and Error-value = 3 ("Unable to allocate a new binding label/SID").

As previously noted, if a message contains an invalid TE-PATH-BINDING TLV that leads to an error condition, the whole message is rejected including any other valid instances of TE-PATH-BINDING TLVs, if any. The resulting error message MAY include the offending TE-PATH-BINDING TLV in the PCEP-ERROR object.

If a PCC receives a 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 an LSP or PCEP-ERROR object, the PCE MUST close the corresponding PCEP session with the reason "Reception of a malformed PCEP message" (according to [RFC5440]).

If a TE-PATH-BINDING TLV is absent in the PCRpt message and no binding values were previously reported, the PCE MUST assume that the corresponding LSP does not have any binding. Similarly, if TE-PATH-BINDING TLV is absent in the PCUpd message and no binding values were previously reported, the PCC's local policy dictates how the binding allocations are made for a given LSP.

Note that some binding types have similar information but different binding value formats. For example, BT=(2 or 3) is used for the SRv6 SID, and BT=(0 or 1) is used for the MPLS Label. In case a PCEP speaker receives multiple TE-PATH-BINDING TLVs with the same SRv6 SID or MPLS Label but different BT values, it MUST send a PCErr message with Error-Type = 32 ("Binding label/SID failure") and Error-value = 5 ("Inconsistent binding types").

6. 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. [RFC8664] defines the "SR-ERO subobject" capable of carrying a SID as well as the identity of the Node or Adjacency Identifier (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. Section 5.2.1 of [RFC8664] specifies bit settings and error handling in the case when NT=0.

7. Binding SID in SRv6-ERO

[RFC9603] defines the "SRv6-ERO subobject" for an SRv6 SID. Similarly to SR-ERO (Section 6), the NAI MUST NOT be included and the NT MUST be set to zero. Section 5.2.1 of [RFC8664] specifies bit settings and error handling in the case when NT=0.

8. PCE Allocation of Binding Label/SID

Section 5 already includes the scenario where a PCE requires a PCC to allocate a specified binding value by sending a PCUpd or PCInitiate message containing a TE-PATH-BINDING TLV. This section specifies an OPTIONAL feature for the PCE to allocate the binding label/SID of its own accord in the case where the PCE also controls the label space of the PCC and can make the label allocation on its own as described in [RFC8283]. Note that the act of requesting a specific binding value (Section 5) is different from the act of allocating a binding label/SID as described in this section.

[RFC8283] introduces the architecture for PCE as a central controller as an extension of the architecture described in [RFC4655] and assumes the continued use of PCEP as the protocol used between PCE and PCC. [RFC9050] specifies the procedures and PCEP extensions for using the PCE as the central controller. It assumes that the exclusive label range to be used by a PCE is known and set on both PCEP peers. A future extension could add the capability to advertise this range via a possible PCEP extension as well (see [PCE-ID-SPACE]).

When PCE as a Central Controller (PCECC) operations are supported as per [RFC9050], the binding label/SID MAY also be allocated by the PCE itself. Both peers need to exchange the PCECC capability as described in [RFC9050] before the PCE can allocate the binding label/SID on its own.

A new P flag in the LSP object [RFC8231] is introduced to indicate that the allocation needs to be made by the PCE. Note that the P flag could be used for other types of allocations (such as path segments [PCEP-SR]) in the future.

P (PCE-allocation): If the bit is set to 1, it indicates that the PCC requests that the PCE make allocations for this LSP. The TE-PATH-BINDING TLV in the LSP object identifies that the allocation is for a binding label/SID. A PCC MUST set this bit to 1 and include a TE-PATH-BINDING TLV in the LSP object if it wishes to request an allocation for a binding label/SID by the PCE in the PCEP message. A PCE MUST also set this bit to 1 and include a TE-PATH-BINDING TLV to indicate that the binding label/SID is allocated by PCE and encoded in the PCEP message towards the PCC. Further, if the binding label/SID is allocated by the PCC, the PCE MUST set this bit to 0 and follow the procedure described in Section 5.

Note that:

- * A PCE could allocate the binding label/SID of its own accord for a PCE-initiated or PCE-delegated LSP and inform the PCC in the PCInitiate message or PCUpd message by setting P=1 and including TE-PATH-BINDING TLV in the LSP object.
- * To let the PCC allocate the binding label/SID, a PCE MUST set P=0 and include an empty TE-PATH-BINDING TLV (i.e., no binding value is specified) in the LSP object in the PCInitiate/PCUpd message.
- * To request that the PCE allocate the binding label/SID, a PCC MUST

set P=1, D=1, and include an empty TE-PATH-BINDING TLV in the PCRpt message. The PCE will attempt to allocate it and respond to the PCC with a PCUpd message that includes the allocated binding label/SID in the TE-PATH-BINDING TLV and P=1 and D=1 in the LSP object. If the PCE is unable to allocate the binding label/SID, it MUST send a PCErr message with Error-Type = 32 ("Binding label/SID failure") and Error-value = 3 ("Unable to allocate a new binding label/SID").

- * If one or both speakers (PCE and PCC) have not indicated support and willingness to use the PCEP extensions for the PCECC as per [RFC9050] and a PCEP peer receives P=1 in the LSP object, they MUST:
 - send a PCErr message with Error-Type = 19 ("Invalid Operation") and Error-value = 16 ("Attempted PCECC operations when PCECC capability was not advertised") and
 - terminate the PCEP session.
- * A legacy PCEP speaker that does not recognize the P flag in the LSP object would ignore it in accordance with [RFC8231].

It is assumed that the label range to be used by a PCE is known and set on both PCEP peers. The exact mechanism is out of the scope of [RFC9050] and this document. Note that the specific BSID could be from the PCE-controlled or the PCC-controlled label space. The PCE can directly allocate the label from the PCE-controlled label space using P=1 as described above, whereas the PCE can request the allocation of a specific BSID from the PCC-controlled label space with P=0 as described in Section 5.

Note that the P flag in the LSP object SHOULD NOT be set to 1 without the presence of TE-PATH-BINDING TLV or any other future TLV defined for PCE allocation. On receipt of such an LSP object, the P flag is ignored. The presence of TE-PATH-BINDING TLV with P=1 indicates the allocation is for the binding label/SID. In the future, some other TLV (such as one defined in [PCEP-SR]) could also be used alongside P=1 to indicate allocation of a different attribute. A future document should not attempt to assign semantics to P=1 without limiting the scope to one that both PCEP peers can agree on.

9. Security Considerations

The security considerations described in [RFC5440], [RFC8231], [RFC8281], [RFC8664], and [RFC9050] are applicable to this specification. No additional security measure is required.

As described in [RFC8402] and [RFC8664], SR intrinsically involves an entity (whether head-end or a central network controller) controlling and instantiating paths in the network without the involvement of (other) nodes along those paths. Binding SIDs are in effect shorthand aliases for longer path representations, and the alias expansion is in principle known only by the node that acts on it. In this document, the expansion of the alias is shared between PCC and PCE, and rogue actions by either PCC or PCE could result in shifting or misdirecting traffic in ways that are hard for other nodes to detect. In particular, when a PCE propagates paths of the form {A, B, BSID} to other entities, the BSID values are opaque, and a rogue PCE can substitute a BSID from a different LSP in such paths to move traffic without the recipient of the path knowing the ultimate destination.

The case of BT=3 provides additional opportunities for malfeasance, as it purports to convey information about internal SRv6 SID Structure. There is no mechanism defined to validate this internal

structure information, and mischaracterizing the division of bits into locator block, locator node, function, and argument can result in different interpretation of the bits by PCC and PCE. Most notably, shifting bits into or out of the "argument" is a direct vector for affecting processing, but other attacks are also possible.

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 RFC 9325 [BCP195] (unless explicitly set aside in [RFC8253]).

10. Manageability Considerations

All manageability requirements and considerations listed in [RFC5440], [RFC8231], and [RFC8664] apply to PCEP protocol extensions defined in this document. In addition, requirements and considerations listed in this section apply.

10.1. Control of Function and Policy

A PCC implementation SHOULD allow the operator to configure the policy the PCC needs to apply when allocating the binding label/SID.

If BT is set to 2, the operator needs to have local policy set to decide the SID structure and the SRv6 Endpoint Behavior of the BSID.

10.2. Information and Data Models

The PCEP YANG module [PCEP-YANG] will be extended to include policy configuration for binding label/SID allocation.

10.3. Liveness Detection and Monitoring

The mechanisms defined in this document do not imply any new liveness detection and monitoring requirements in addition to those already listed in [RFC5440].

10.4. Verify Correct Operations

The mechanisms defined in this document do not imply any new operation verification requirements in addition to those already listed in [RFC5440], [RFC8231], and [RFC8664].

10.5. Requirements on Other Protocols

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

10.6. Impact on Network Operations

The mechanisms defined in [RFC5440], [RFC8231], and [RFC8664] also apply to the PCEP extensions defined in this document.

11. IANA Considerations

IANA has allocated code points for the protocol elements described in this document in the "Path Computation Element Protocol (PCEP) Numbers" registry group.

11.1. PCEP TLV Type Indicators

This document defines a new PCEP TLV. IANA has allocated the following in the "PCEP TLV Type Indicators" registry within the PCEP Numbers registry group:

Value	Description	Reference
55	TE-PATH-BINDING	RFC 9604

Table 1

11.1.1.1. TE-PATH-BINDING TLV

IANA has created the "TE-PATH-BINDING TLV BT Field" registry to manage the values of the binding type field in the TE-PATH-BINDING TLV. Initial values are shown below. New values are assigned by Standards Action [RFC8126].

Value	Description	Reference
0	MPLS Label	RFC 9604
1	MPLS Label Stack Entry	RFC 9604
2	SRv6 SID	RFC 9604
3	SRv6 SID with Behavior and Structure	RFC 9604
4-255	Unassigned	

Table 2

IANA has created a new "TE-PATH-BINDING TLV Flag Field" registry to manage the Flag field in the TE-PATH-BINDING TLV. New values are to be assigned by Standards Action [RFC8126]. Each bit should be tracked with the following qualities:

- * Bit number (count from 0 as the most significant bit)
- * Description
- * Reference

Bit	Description	Reference
0	R (Removal)	RFC 9604
1-7	Unassigned	

Table 3

11.2. LSP Object

IANA has allocated a code point in the "LSP Object Flag Field" registry for the new P flag as follows:

Bit	Description	Reference
0	PCE-allocation	RFC 9604

Table 4

11.3. PCEP Error Type and Value

This document defines a new Error-Type and associated Error-values for the PCErr message. IANA has allocated a new Error-Type and Error-values within the "PCEP-ERROR Object Error Types and Values" registry of the PCEP Numbers registry group, as follows:

Error-Type	Meaning	Error-value
32	Binding label/ SID failure	0: Unassigned
		1: Invalid SID
		2: Unable to allocate the specified binding value
		3: Unable to allocate a new binding label/SID
		4: Unable to remove the binding value
		5: Inconsistent binding types

Table 5

12. References

12.1. Normative References

- [BCP195] Best Current Practice 195,
<<https://www.rfc-editor.org/info/bcp195>>.
At the time of writing, this BCP comprises the following:

Moriarty, K. and S. Farrell, "Deprecating TLS 1.0 and TLS 1.1", BCP 195, RFC 8996, DOI 10.17487/RFC8996, March 2021, <<https://www.rfc-editor.org/info/rfc8996>>.

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>>.
- [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>>.
- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, DOI 10.17487/RFC3032, January 2001, <<https://www.rfc-editor.org/info/rfc3032>>.
- [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>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for

Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.

- [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>>.
- [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>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.
- [RFC9050] Li, Z., Peng, S., Negi, M., Zhao, Q., and C. Zhou, "Path Computation Element Communication Protocol (PCEP) Procedures and Extensions for Using the PCE as a Central Controller (PCECC) of LSPs", RFC 9050, DOI 10.17487/RFC9050, July 2021, <<https://www.rfc-editor.org/info/rfc9050>>.
- [RFC9603] Li, C., Ed., Kaladharan, P., Sivabalan, S., Koldychev, M., and Y. Zhu, "Path Computation Element Communication Protocol (PCEP) Extensions for IPv6 Segment Routing", RFC 9603, DOI 10.17487/RFC9603, July 2024, <<https://www.rfc-editor.org/info/rfc9603>>.

12.2. Informative References

- [PCE-ID-SPACE] Li, C., Shi, H., Ed., Wang, A., Cheng, W., and C. Zhou, "Path Computation Element Communication Protocol (PCEP) extension to advertise the PCE Controlled Identifier Space", Work in Progress, Internet-Draft, draft-ietf-pce-controlled-id-space-00, 4 June 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-controlled-id-space-00>>.

- [PCEP-SR] Li, C., Chen, M., Cheng, W., Gandhi, R., and Q. Xiong, "Path Computation Element Communication Protocol (PCEP) Extension for Path Segment in Segment Routing (SR)", Work in Progress, Internet-Draft, draft-ietf-pce-sr-path-segment-09, 26 February 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-sr-path-segment-09>>.
- [PCEP-YANG] Dhody, D., Ed., Beeram, V., Hardwick, J., and J. Tantsura, "A YANG Data Model for Path Computation Element Communications Protocol (PCEP)", Work in Progress, Internet-Draft, draft-ietf-pce-pcep-yang-25, 21 May 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-pce-pcep-yang-25>>.
- [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>>.
- [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>>.
- [RFC8754] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <<https://www.rfc-editor.org/info/rfc8754>>.
- [RFC9256] Filsfils, C., Talaulikar, K., Ed., Voyer, D., Bogdanov, A., and P. Mattes, "Segment Routing Policy Architecture", RFC 9256, DOI 10.17487/RFC9256, July 2022, <<https://www.rfc-editor.org/info/rfc9256>>.

Acknowledgements

We would like to thank Milos Fabian, Mrinmoy Das, Andrew Stone, Tom Petch, Aijun Wang, Olivier Dugeon, and Adrian Farrel for their valuable comments.

Thanks to Julien Meuric for shepherding. Thanks to John Scudder for the AD review.

Thanks to Theresa Enghardt for the GENART review.

Thanks to Martin Vigoureux, Benjamin Kaduk, ric Vyncke, Lars Eggert, Murray Kucherawy, and Erik Kline for the IESG reviews.

Contributors

Jonathan Hardwick
Microsoft
United Kingdom
Email: jonhardwick@microsoft.com

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

Email: dhruv.ietf@gmail.com

Mahendra Singh Negi
RtBrick India
N-17L, Floor-1, 18th Cross Rd, HSR Layout Sector-3
Bangalore 560102
Karnataka
India
Email: mahend.ietf@gmail.com

Mike Koldychev
Cisco Systems, Inc.
2000 Innovation Drive
Kanata Ontario K2K 3E8
Canada
Email: mkoldych@cisco.com

Zafar Ali
Cisco Systems, Inc.
Email: zali@cisco.com

Authors' Addresses

Siva Sivabalan
Ciena Corporation
Email: msiva282@gmail.com

Clarence Filsfils
Cisco Systems, Inc.
Pegasus Parc
De Kleetlaan 6a
1831 Brabant
Belgium
Email: cfilsfil@cisco.com

Jeff Tantsura
Nvidia
Email: jefftant.ietf@gmail.com

Stefano Previdi
Huawei Technologies
Email: stefano@previdi.net

Cheng Li (editor)
Huawei Technologies
Huawei Campus, No. 156 Beiqing Rd.
Beijing
100095
China
Email: c.l@huawei.com