PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: 12 October 2026

S. Sidor
Z. Ali
Cisco Systems, Inc.
C. Li
Huawei Technologies
M. Koldychev
Ciena Corporation

10 April 2026

Binding Label/Segment Identifier (SID) Extensions in Path Computation Element Communication Protocol (PCEP)
draft-sidor-pce-binding-label-sid-extensions-02

Abstract

The Path Computation Element Communication Protocol (PCEP) provides mechanisms for Path Computation Elements (PCEs) to instantiate and manage Label Switched Paths (LSPs) on a Path Computation Client (PCC). This includes the ability for a PCE to specify a Binding Segment Identifier (SID) for an LSP.

A binding value specified by a PCE may not be available for use on the PCC. This can lead to LSP instantiation failures or entire PCEP message being rejected.

This document proposes extensions to PCEP to allow a PCC to fall back to allocating a Binding SID from its own dynamic range if the value specified by the PCE is unavailable. It also defines a mechanism for the PCC to report both the requested and the allocated binding values back to the PCE.

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 12 October 2026.

Table of Contents

1. Introduction

This document proposes extensions to the Path Computation Element Communication Protocol (PCEP) to enhance the management of Binding Segment Identifiers (SIDs) for Label Switched Paths (LSPs). Specifically, it defines mechanisms for a Path Computation Client (PCC) to handle situations where a Binding SID (BSID) requested by a Path Computation Element (PCE) is unavailable, allowing for fallback

allocation and subsequent reporting of the allocated values back to the PCE. The ability for a PCE to specify a Binding SID for an LSP is defined in [RFC9604]. These extensions aim to improve the robustness and flexibility of LSP instantiation and management in PCEP-controlled networks.

1.1. 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.

2. Terminology

This document uses the following terms defined in [RFC5440]: PCC, PCE, PCEP Peer, and PCEP speaker.

The base PCEP specification [RFC4655] originally defined the use of the PCE architecture for Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS) networks with Label Switched Paths (LSPs) instantiated using the Resource Reservation Protocol - Traffic Engineering (RSVP-TE) signaling protocol. Over time, support for additional path setup types, such as SRv6, has been introduced [RFC9603]. The term "LSP" is used extensively in PCEP specifications and, in the context of this document, refers to a Candidate Path within an SR Policy, which may be an Segment Routing over IPv6 (SRv6) path (still represented using the LSP Object as specified in [RFC8231].

It also uses the term Binding Segment Identifier (BSID), as defined in [RFC9604], which refers to a local label or SID that represents an SR Policy or an SR-TE LSP.

3. Motivation

The PCEP provides mechanisms for PCEs to instantiate and manage LSPs on a PCC. A Stateful PCE [RFC8231] can instantiate LSPs on a PCC. When instantiating a Segment Routing Traffic Engineering (SR-TE) LSP [RFC8664], the PCE may request a specific BSID to be associated with the LSP using the TE-PATH-BINDING Type-Length-Value (TLV) [RFC9604].

A significant operational challenge arises when the BSID requested by the PCE is already in use, falls outside the valid range, or is otherwise unavailable on the PCC. In the current PCEP specification, such a conflict or unavailability typically results in an LSP instantiation failure. This "hard failure" approach can be

disruptive, requiring manual steps from an operator or complex retry logic at the PCE, and can have negative impact on automated provisioning capabilities that PCEP aims to provide. It can also lead to entire PCEP messages being rejected, forcing the PCE to re-evaluate and re-initiate the entire LSP setup process.

To improve network resilience and operational efficiency, it is desirable to have more flexible mechanisms for handling BSID unavailability scenarios. Instead of failure, a PCC should ideally be able to gracefully handle such situations, for instance, by allocating a Binding SID from its local dynamic range. Furthermore, the PCE needs to be aware of the actual BSID allocated by the PCC to maintain an accurate view of the network state. This document defines extensions to PCEP to address these operational needs.

4. PCEP Extensions

4.1. STATEFUL-PCE-CAPABILITY TLV

A new flag is proposed for the STATEFUL-PCE-CAPABILITY TLV, originally defined in Section 5.4 of [RFC8231].

* E (BSID-FALLBACK-CAPABILITY): If set, indicates that the PCEP peer supports LSP creation and fall back to dynamic binding value allocation if the specific binding value is unavailable, as detailed in Section 5.

4.2. TE-PATH-BINDING TLV

New flags are proposed in the TE-PATH-BINDING TLV, which was originally defined in Section 4 of [RFC9604].

* A (Allocated): If set, indicates that the binding value encoded in the TLV represents an allocated binding value.
* D (Down on BSID Unavailability): If set, indicates that LSP can be created even if specified binding value is unavailable, but LSP will be in down state.
* F (Fallback): If set, indicates that binding value allocation from the dynamic range will be performed if the specified binding value is unavailable.

5. Operation

The PCEP protocol extensions defined in this document MUST NOT be used if one or both PCEP speakers have not indicated support for the extensions by setting the E flag (BSID-FALLBACK-CAPABILITY) in the STATEFUL-PCE-CAPABILITY TLV in their respective OPEN messages.

When a PCE wants to instantiate or update an LSP and suggest a binding value, it includes the TE-PATH-BINDING TLV in the Path Computation LSP Initiate Request (PCInitiate) or Path Computation LSP Update Request (PCUpd) message [RFC8231]. The PCE can set the F flag or the D flag in this TLV to control the PCC's behavior in case the requested binding value is unavailable. The F and D flags are mutually exclusive. If a PCEP speaker receives a TE-PATH-BINDING TLV where both the F flag and the D flag are set, the PCEP speaker MUST send a PCErr message with Error-Type 10 (Reception of an invalid object) and Error-Value TBD5 (Mutually exclusive F and D flags are both set). The LSP instantiation or update request associated with this malformed TLV MUST be rejected.

When both F=0 and D=0, the current behavior as specified in [RFC9604] applies: the LSP instantiation fails if the requested binding value is unavailable.

If a PCEP speaker receives a TE-PATH-BINDING TLV with the A flag set in a PCInitiate or PCUpd message, the PCEP speaker MUST send a PCErr message with Error-Type 10 (Reception of an invalid object) and Error-Value TBD7 (A flag incorrectly set by PCE). The LSP instantiation or update request associated with this malformed TLV MUST be rejected.

If the PCC receives a TE-PATH-BINDING TLV with the F flag set and the requested binding value is unavailable, the PCC MUST attempt to allocate a new binding value from its dynamic pool. If successful, the LSP is brought up with the new binding value.

If the PCC receives a TE-PATH-BINDING TLV with the D flag set and the requested binding value is unavailable, the PCC MUST instantiate the LSP but keep it in a down state.

If the PCC attempts to allocate a binding value from its dynamic pool (when the F flag is set) but the allocation fails due to pool exhaustion or other reasons, the PCC MUST report the LSP in a down state with appropriate error indication in the PCRpt message.

In its Path Computation LSP State Report (PCRpt) message [RFC8231], the PCC reports the status of the binding value allocation. If the originally requested binding value and the allocated binding value differ, two instances of the TE-PATH-BINDING TLV MUST be included in the PCRpt message:

* A TLV instance with the originally requested binding value with the A flag cleared.
* A TLV instance with the actually allocated binding value with the A flag set.

For example, if the PCE requested BSID value 100 with the F flag set, but value 100 was unavailable and the PCC allocated BSID value 200 from its dynamic pool, the PCRpt message would contain:

* TE-PATH-BINDING TLV with binding value 100, A flag = 0, F flag = 1
* TE-PATH-BINDING TLV with binding value 200, A flag = 1, F flag = 1

This allows the PCE to correlate what it requested with what was actually allocated.

If the requested binding value was successfully allocated, only a single instance of the TE-PATH-BINDING TLV with the A flag set SHOULD be included in the PCEP message.

For PCC-initiated LSPs, the PCC MAY set the F or D flags in the TE-PATH-BINDING TLV included in PCRpt messages to indicate the desired fallback behavior for the binding value. For PCE-initiated LSPs, the PCC MUST reflect the D and F flag values from the PCE's PCInitiate or PCUpd message in all TE-PATH-BINDING TLV instances included in PCRpt messages. This reflection ensures that the binding value allocation policy is propagated to all PCEs in redundant PCE deployments.

The A, D, and F flags in the TE-PATH-BINDING TLV MUST NOT be used if one or both PCEP speakers have not set the BSID-FALLBACK-CAPABILITY in the STATEFUL-PCE-CAPABILITY TLV in their respective OPEN messages. If a PCEP speaker receives a PCEP message containing the A, D, or F flags in the TE-PATH-BINDING TLV, or any other element specific to these extensions, from a peer that has not advertised the BSID-FALLBACK-CAPABILITY in its OPEN message, the receiving PCEP speaker MUST send a PCErr message with Error-Type 10 (Reception of an invalid object) and Error-Value TBD6 (Unsupported Binding SID Extension Flags).

6. Operational Considerations

All operational requirements and considerations listed in [RFC5440], [RFC8231], and [RFC9604] apply to the PCEP extensions defined in this document.

6.1. Control of Function and Policy

A PCE or PCC implementation SHOULD allow the BSID fallback capability to be enabled or disabled through configuration, either globally or on a per-LSP basis. An implementation SHOULD allow the operator to view the advertised and received BSID-FALLBACK-CAPABILITY flags.

Implementations SHOULD provide configuration options to:

* Enable or disable the BSID-FALLBACK-CAPABILITY advertisement
* Configure the range of binding values available for dynamic allocation
* Set policies for when to use fallback allocation (F flag) versus keeping LSP down (D flag)
* Define priority or preference for BSID allocation from the dynamic pool

6.2. Information and Data Models

Implementations SHOULD provide operational state information including:

* Whether BSID-FALLBACK-CAPABILITY is enabled and advertised
* List of LSPs with binding values, showing both requested and allocated values when they differ
* History of binding value allocation failures

A YANG data model for PCEP [RFC9604] MAY be extended to include:

* Capability advertisement of BSID-FALLBACK-CAPABILITY
* Operational state showing requested versus allocated binding values
* Configuration parameters for dynamic binding value pool management

6.3. Liveness Detection and Monitoring

Operators SHOULD monitor binding value allocation events and configure alerts for:

* Binding value allocation failures due to unavailability
* Dynamic binding value pool utilization exceeding configured thresholds
* Frequent fallback allocations indicating potential BSID conflicts
* LSPs in down state due to D flag with unavailable binding values

Implementations SHOULD provide logging for binding value allocation events, including requested values, allocated values, and reasons for any allocation failures.

6.4. Fault Management

As specified in Section 5, when BSID fallback allocation fails (for example, due to dynamic pool exhaustion), the LSP is reported as down with appropriate error indication. Implementations should provide clear diagnostic information to help operators identify the root cause of allocation failures, such as pool exhaustion, configuration errors, or BSID conflicts.

Operators need to be aware that:

* Binding value conflicts can occur due to configuration errors, race conditions, or pool exhaustion
* The D flag behavior (LSP down when requested BSID unavailable) may impact service availability and requires monitoring
* In redundant PCE deployments, binding value allocation state should be synchronized or coordinated to avoid conflicts

7. Security Considerations

The security considerations described in [RFC5440], [RFC8231], and [RFC9604] are applicable to this document.

The extensions defined in this document introduce new operational behaviors that require careful security consideration:

* Binding Value Allocation: The fallback mechanism allows a PCC to allocate binding values from its dynamic pool when requested values are unavailable. Implementations MUST ensure that the dynamic allocation process includes proper validation and does not allow unauthorized binding value usage. An attacker attempting to exhaust the dynamic pool through repeated requests with unavailable values could cause a denial-of-service condition. Implementations SHOULD implement rate limiting and monitoring of allocation failures.
* State Reporting: The mechanism for reporting both requested and allocated binding values provides visibility into binding value allocation. This information MUST be protected to prevent unauthorized correlation of network state. Implementations MUST validate that reported binding values in PCRpt messages accurately reflect the actual allocated values.
* LSP Down State: The D flag allows LSPs to be created in a down state when binding values are unavailable. Implementations need to ensure that this does not create opportunities for denial-of-service attacks where an attacker forces numerous LSPs into down state by requesting unavailable binding values.
* Flag Manipulation: The A, D, and F flags control critical allocation behavior. Implementations MUST enforce the rules for flag usage, including rejecting messages with the A flag set by a PCE and properly handling mutually exclusive F and D flags, as specified in Section 5.

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 [RFC9325]. This is particularly important given the sensitivity of binding value allocation and the potential for denial-of-service attacks through pool exhaustion.

Operators SHOULD carefully review and configure the dynamic binding value pool ranges to ensure adequate capacity while preventing overlap with statically configured binding values. Regular monitoring of binding value allocation patterns can help detect potential security issues or misconfigurations.

8. IANA Considerations

8.1. STATEFUL-PCE-CAPABILITY TLV Flag

IANA maintains the "STATEFUL-PCE-CAPABILITY TLV Flag Field" registry within the "Path Computation Element Protocol (PCEP) Numbers" registry group. See https://www.iana.org/assignments/pcep/pcep.xhtml#stateful-pce-capability-tlv-flag-field

IANA is requested to make the following assignment:

Bit	Description	Reference
TBA1	E (BSID-FALLBACK-CAPABILITY)	This document

Table 1

8.2. TE-PATH-BINDING TLV Flags

IANA maintains the "TE-PATH-BINDING TLV Flag Field" registry within the "Path Computation Element Protocol (PCEP) Numbers" registry group. See https://www.iana.org/assignments/pcep/pcep.xhtml#te-path-binding-tlv-flag-field

IANA is requested to make the following assignments:

Bit	Description	Reference
TBA2	A (Allocated)	This document
TBA3	D (Down on BSID Unavailability)	This document
TBA4	F (Fallback)	This document

Table 2

8.3. PCEP Errors

IANA maintains the "PCEP-ERROR Object Error Types and Values" registry within the "Path Computation Element Protocol (PCEP) Numbers" registry group. See https://www.iana.org/assignments/pcep/pcep.xhtml#pcep-error-object

IANA is requested to make the following assignments:

Error-Type	Meaning	Error-value	Reference
10	Reception of an invalid object	TBD5: Mutually exclusive F and D flags are both set	This document
		TBD6: Unsupported Binding SID Extension Flags	This document
		TBD7: A flag incorrectly set by PCE	This document

Table 3

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[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>.

[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>.

[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>.

[RFC9604] Sivabalan, S., Filsfils, C., Tantsura, J., Previdi, S., and C. Li, Ed., "Carrying Binding Label/SID in PCE-Based Networks", RFC 9604, DOI 10.17487/RFC9604, August 2024, <https://www.rfc-editor.org/info/rfc9604>.

9.2. 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>.

[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>.

Appendix A. Acknowledgements

The authors would like to thank Rajesh Melarcode Venkateswaran for their contributions to this document.

Authors' Addresses

Samuel Sidor
Cisco Systems, Inc.
Eurovea Central 3
Pribinova 10
811 09 Bratislava
Slovakia
Email: ssidor@cisco.com

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

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

Mike Koldychev
Ciena Corporation
385 Terry Fox Dr.
Kanata Ontario K2K 0L1
Canada
Email: mkoldych@proton.me