

PCE Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: September 5, 2018

U. Palle  
D. Dhody  
S. Karunanithi  
Huawei Technologies  
March 4, 2018

LABEL-DB Synchronization Procedures for a PCE as a central  
controller(PCECC)  
draft-palle-pce-controller-labeldb-sync-03

Abstract

PCE as a central controller specifies the procedures and PCEP protocol extensions where LSPs are calculated/setup/initiated and label forwarding entries are downloaded through a centralized PCE server to each network devices along the LSP path while leveraging the existing PCE technologies as much as possible.

Labels downloaded to forwarding entries requires a reliable synchronization mechanism between the path computation clients (PCCs) and the PCE. The basic mechanism for label database (LABEL-DB) synchronization is part of the PCE as a central controller specification. This document presents motivations for optimizations to the LABEL-DB synchronization and the corresponding PCEP procedures and extensions.

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 September 5, 2018.

## Copyright Notice

Copyright (c) 2018 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
1.1.	Requirements Language . . . . .	3
2.	LABEL-DB Synchronization . . . . .	3
3.	Optimizations for LABEL-DB Synchronization . . . . .	4
3.1.	LABEL-DB Synchronization Avoidance Procedure . . . . .	4
3.2.	Incremental LABEL-DB Synchronization Procedure . . . . .	8
4.	PCEP Extensions . . . . .	10
4.1.	Extension of SRP object . . . . .	10
4.2.	Extension of PCECC Capability sub-TLV . . . . .	10
4.3.	New LABEL-DB-VERSION TLV . . . . .	11
5.	IANA Considerations . . . . .	12
5.1.	PCEP TLV Type Indicators . . . . .	12
5.2.	PCECC-CAPABILITY sub-TLV . . . . .	12
6.	Manageability Considerations . . . . .	12
6.1.	Control of Function and Policy . . . . .	12
6.2.	Information and Data Models . . . . .	12
6.3.	Liveness Detection and Monitoring . . . . .	12
6.4.	Verify Correct Operations . . . . .	12
6.5.	Requirements On Other Protocols . . . . .	13
6.6.	Impact On Network Operations . . . . .	13
7.	Security Considerations . . . . .	13
8.	Acknowledgements . . . . .	13
9.	References . . . . .	13
9.1.	Normative References . . . . .	13
9.2.	Informative References . . . . .	14
	Authors' Addresses . . . . .	14

## 1. Introduction

[I-D.zhao-pce-pcep-extension-for-pce-controller] specify the procedures and PCEP protocol extensions for using the PCE as the central controller [RFC8283] and user cases where LSPs are calculated/setup/initiated/downloaded through extending the existing PCE architectures and PCEP.

[I-D.zhao-pce-pcep-extension-for-pce-controller] and [I-D.zhao-pce-pcep-extension-pce-controller-sr] specifies reliable synchronization mechanism between the path computation clients (PCCs) and the PCECC.

This draft specify the optimizations for LABEL-DB synchronization and the corresponding PCEP procedures and extensions.

[Important Note - [I-D.zhao-pce-pcep-extension-for-pce-controller] defined new messages for label download and cleanup. The authors and WG also debated on the use of existing PCEP messages. The appendix of the document includes the details related to use of existing messages. This document is related to new messages as a new procedure for Label DB sync was defined. In case existing messages are used, some modifications needs to be made to the existing LSP-DB synchronization mechanism to also handle the label synchronization. These details are not present in the document at this stage. ]

### 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. LABEL-DB Synchronization

PCECC MUST maintains the LABEL-DB for each PCEP session separately. The purpose of LABEL-DB synchronization is to make sure that the PCECC's view of LABEL-DB matches with the PCC's LABEL-DB. The LABEL-DB synchronization MUST be performed from PCECC to PCC immediately after the LSP state synchronization. [RFC8231] describes the basic mechanism for LSP state synchronization. [RFC8233] describes the optimizations for LSP state synchronization.

Full LABEL-DB synchronization performed from PCECC to PCC on Initial session UP or every session flap is described in [I-D.zhao-pce-pcep-extension-for-pce-controller] and [I-D.zhao-pce-pcep-extension-pce-controller-sr].

Providing an Optimizations for LABEL-DB synchronization can result in significant savings in both control-plane data exchanges and the time it takes for the PCC to become fully operational.

Optimizations for LABEL-DB synchronization describes the need that both PCEP speakers support label database version capability and maintain label database version for each session. See Section 3 for detail procedures.

[Editor's Note: [I-D.zhao-pce-pcep-extension-for-pce-controller] defines new messages PCLabelUpd and PCLabelRpt. Questions where raised on the need for the new messages. Further appendix in [I-D.zhao-pce-pcep-extension-for-pce-controller] and [I-D.zhao-pce-pcep-extension-pce-controller-sr] describes how the existing messages can be extended to add this functionality. WG needs to decide the final direction i.e. new specific messages are needed or existing PCEP messages can be extended. The optimization procedure would need to be modified based on the above decision.]

### 3. Optimizations for LABEL-DB Synchronization

This section add some of the optimization mechanisms for LABEL-DB synchronization. By default, the full LABEL-DB synchronization is performed.

#### 3.1. LABEL-DB Synchronization Avoidance Procedure

The LABEL-DB synchronization MAY be skipped following a PCEP session restart if there is no change in the LABEL-DB of the session at PCECC, during the period prior to session re-initialization. To be able to make this determination, labels must be exchanged and maintained by both PCECC and PCC during normal operation. This is accomplished by keeping track of the changes to the label database, using a version tracking field called the Label Database Version Number.

The Label Database Version Number, carried in LABEL-DB-VERSION TLV (see Section 4.3), is owned by a PCECC and it MUST be incremented by 1 for each successive change in the PCECC's label database. The Label Database Version Number MUST start at 1 and may wrap around. Values 0 and 0xFFFFFFFFFFFFFFFF are reserved. If either of the two values are used during LABEL-DB synchronization, the PCC speaker receiving this node should send back a PCERR with Error-type TBD1 Error-value 3 'Received an invalid Label Database Version Number', and close the PCEP session. Operations that trigger a change to the Label database include an addition or deletion of labels that would trigger a label update to the PCC.

LABEL-DB synchronization avoidance is advertised on a PCEP session during session startup using the INCLUDE-LABEL-DB-VERSION (I) bit in the PCECC capability sub-TLV (see Section 4.2). The PCEP peer MAY include the SPEAKER-ENTITY-ID TLV described in [RFC8233] in the OPEN message to identify the peer in case of IP address change.

If both PCEP speakers set the I flag in the OPEN object's PCECC Capability sub-TLV to 1, the PCECC MUST include the LABEL-DB-VERSION TLV in each LABEL object of the PCLabelUpd message. If the LABEL-DB-VERSION TLV is missing in a PCLabelUpd message, the PCC will generate an error with Error-Type 6 (mandatory object missing) and Error-Value TBD2 'LABEL-DB-VERSION TLV missing' and close the session. If LABEL-DB synchronization avoidance has not been enabled on a PCEP session, the PCECC SHOULD NOT include the LABEL-DB-VERSION TLV in the LABEL Object and the PCC SHOULD ignore it were it to receive one.

If a PCC's label database survived the restart of a PCEP session, the PCC will include the LABEL-DB-VERSION TLV in its OPEN object, and the TLV will contain the last Label Database Version Number received on an Label Update from the PCECC in the previous PCEP session. If a PCECC's Label Database survived the restart of a PCEP session, the PCECC will include the LABEL-DB-VERSION TLV in its OPEN object and the TLV will contain the latest Label Database Version Number. If a PCEP speaker's label database did not survive the restart of a PCEP session, the PCEP speaker MUST NOT include the LABEL-DB-VERSION TLV in the OPEN object.

If both PCEP speakers include the LABEL-DB-VERSION TLV in the OPEN Object and the TLV values match, the PCECC MAY skip LABEL-DB synchronization. Otherwise, the PCECC MUST perform full LABEL-DB synchronization ([I-D.zhao-pce-pcep-extension-for-pce-controller] and [I-D.zhao-pce-pcep-extension-pce-controller-sr]) or incremental LABEL-DB synchronization (see Section 3.2) to the PCC, In case, the PCECC attempts to skip LABEL-DB synchronization, by setting the SYNC Flag to 0 on the first Label Update from the PCECC, the PCC MUST send back a PCErr with Error-type TBD1 (Label Database Synchronization Error) and Error-value 4(Label Database Version mismatch), and close the PCEP session.

If LABEL-DB synchronization is required, then prior to completing the initialization phase, the PCC MUST mark any labels in the label database that were previously updated by the PCECC as stale. When the PCECC updates a label during LABEL-DB synchronization, if the label already exists in the label database, the PCC MUST update the label database and clear the stale marker from the label. When it has finished LABEL-DB synchronization, the PCECC MUST immediately send an end of synchronization marker. The end of synchronization marker is a Path Computation Label Update (PCLabelUpd) message with a

SRP object containing the SYNC flag set to 0 (see Section 4.1) and Label as 0 in the LABEL object. The LABEL-DB-VERSION TLV MUST be included in this PCLabelUpd message. On receiving this Label Update, the PCC MUST report all the labels in the label database that are still marked as stale to PCECC.

Note that a PCECC/PCC MAY force LABEL-DB synchronization by not including the LABEL-DB-VERSION TLV in its OPEN object.

Figure 1 shows an example sequence where the LABEL-DB synchronization is skipped.

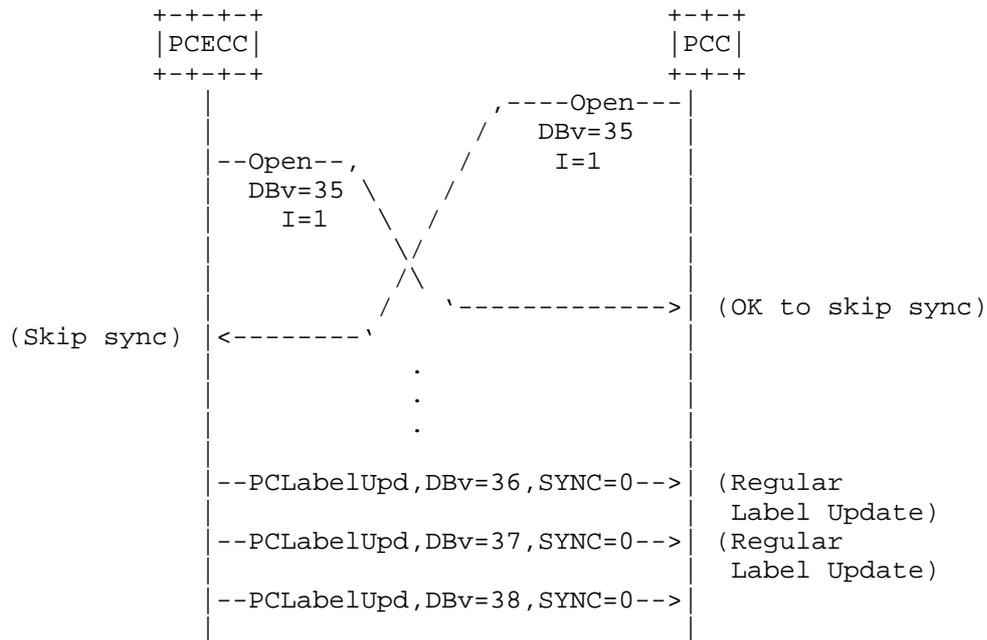


Figure 1: LABEL-DB synchronization Skipped

Figure 2 shows an example sequence where the LABEL-DB synchronization is performed due to label database version mismatch during the PCEP session setup. Note that the same LABEL-DB synchronization sequence would happen if either the PCC or the PCECC would not include the LABEL- DB-VERSION TLV in their respective Open messages.

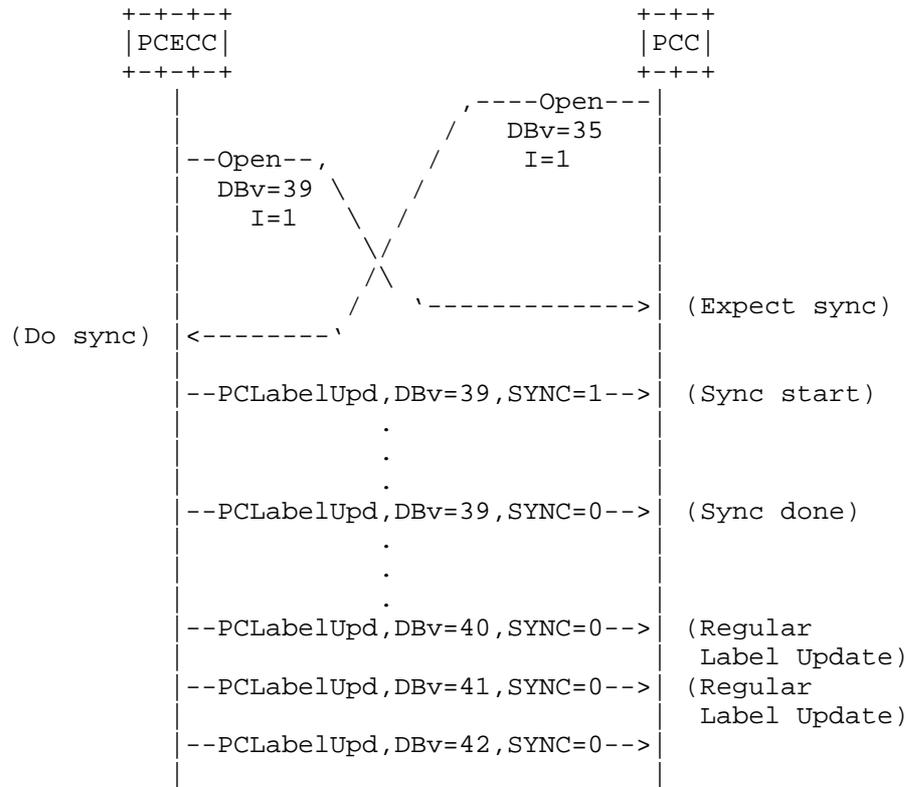


Figure 2: LABEL-DB synchronization Performed

Figure 3 shows an example sequence where the LABEL-DB synchronization is skipped, but because one or both PCEP speakers set the I Flag to 0, the PCECC does not send LABEL-DB-VERSION TLVs in subsequent PCLabelUpd messages to the PCC. If the current PCEP session restarts, the PCEP speakers will have to perform full LABEL-DB synchronization, since the PCC does not know the PCECC's latest Label Database Version Number information.

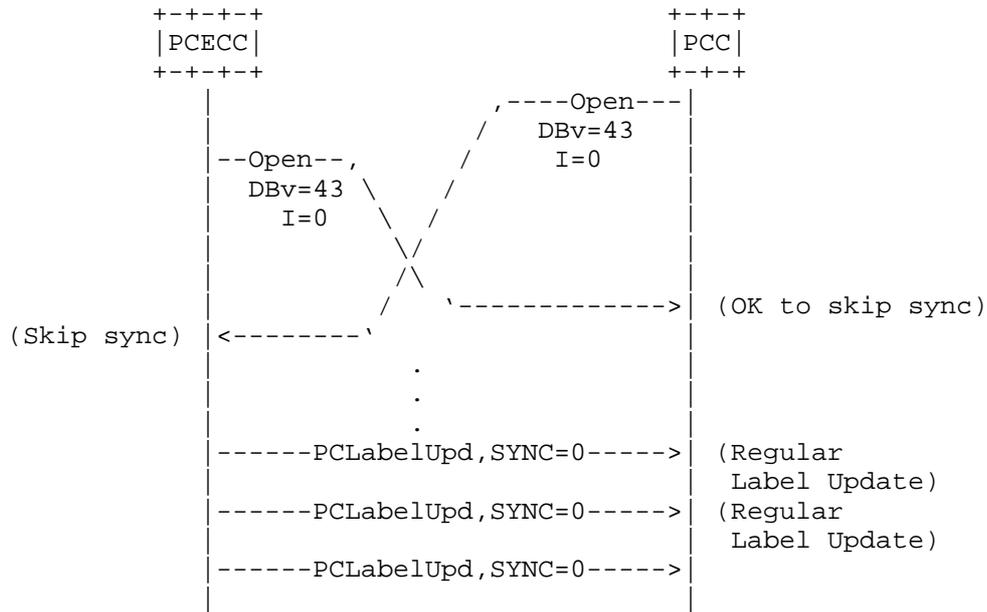


Figure 3: LABEL-DB Synchronization Skipped, no LABEL-DB-VERSION TLVs sent from PCECC

### 3.2. Incremental LABEL-DB Synchronization Procedure

If a PCC restarts and its label database survived, PCECC with mismatched Label Database Version Number will send all their Labels information (full LABEL-DB) to the PCC, even if only a small number of changes happened. It can take a long time and consume large communication channel bandwidth.

This section extends the idea to only synchronize the delta (changes) in case of Label Database Version Number of both PCEP peers is non-zero and mismatch.

If both PCEP speakers include the LABEL-DB-VERSION TLV in the OPEN object and the LABEL-DB-VERSION TLV values match, the PCECC MAY skip LABEL-DB synchronization. Otherwise, the PCECC MUST perform LABEL-DB synchronization. Incremental label database synchronization capability is advertised on a PCEP session during session startup using the DELTA-LABEL-SYNC-CAPABILITY (D) bit in the capabilities TLV (see Section 4.2). Instead of dumping full LABEL-DB to the PCC again, the PCECC synchronizes the delta (changes) as described in Figure 4 when D flag and I flag is set to 1 by both PCC and PCECC. Other combinations of D and I flags setting by PCC and PCECC result in full LABEL-DB synchronization procedure as described in

[I-D.zhao-pce-pcep-extension-for-pce-controller] and [I-D.zhao-pce-pcep-extension-pce-controller-sr]. The PCECC MAY force a full LABEL-DB synchronization by setting the D flag to zero in the OPEN message.

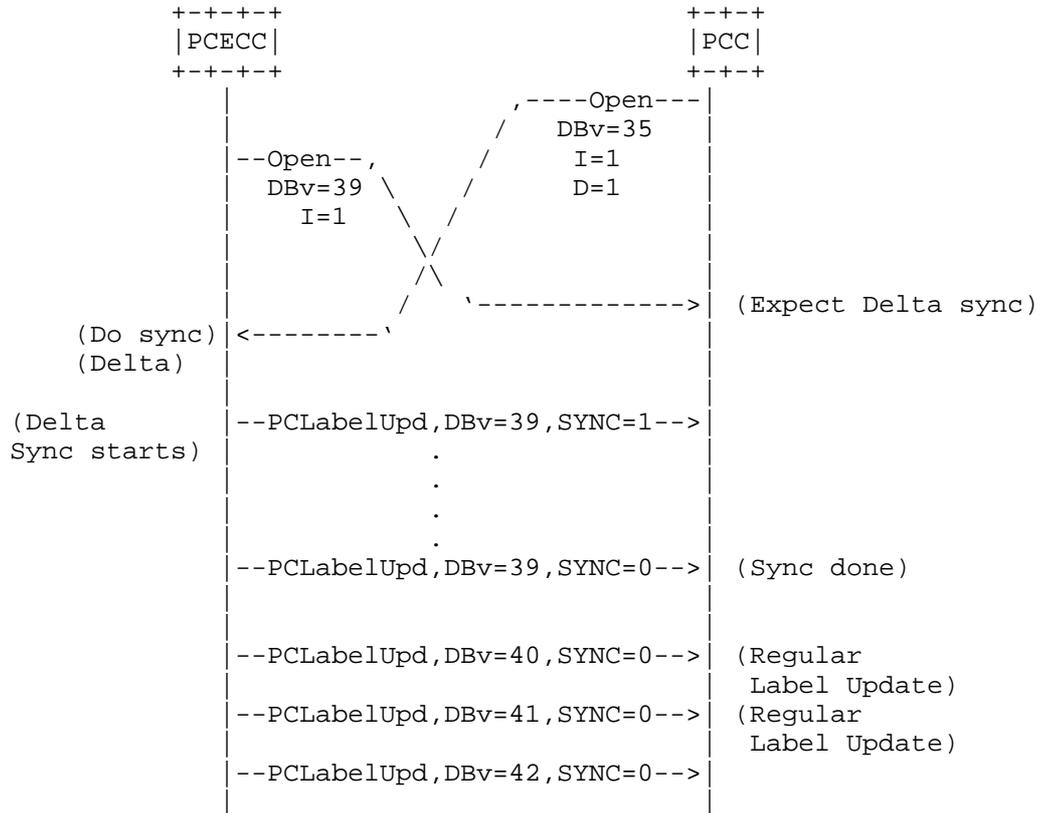


Figure 4: Incremental Synchronization Procedure

As per Section 3.1, the Label Database Version Number is incremented each time a change is made to the PCECC's label database. Each label is associated with the DB version at the time of its addition. This is needed to determine which label and what information needs to be synchronized in incremental LABEL-DB synchronization.

It is not necessary for a PCECC to store a complete history of label database change, but rather remember the labels (including label addition and deletion) that happened between the PCEP session(s) restart in order to carry out incremental LABEL-DB synchronization. After the synchronization procedure finishes, the PCECC can dump this

history information. In the example shown in Figure 4, the PCECC needs to store the label changes that happened between DB Version 35 to 39 and synchronizes these changes only when performing incremental label update. So a PCECC needs to remember at least the label changes that happened after an existing PCEP session with a PCC goes down to have any chance of doing incremental synchronization when the session is re-established.

If a PCECC finds out it does not have sufficient information to complete incremental synchronization after advertising incremental LABEL-DB synchronization capability, it MUST send a PCErr with Error-Type TBD1 and Error-Value 5 'A PCECC indicates to a PCC that it can not complete the LABEL-DB synchronization' and terminate the session. The PCECC SHOULD re-establish the session with the D bit set to 0 in the OPEN message.

The other procedures and error checks remain unchanged from the default LABEL-DB synchronization defined in [I-D.zhao-pce-pcep-extension-for-pce-controller] and [I-D.zhao-pce-pcep-extension-pce-controller-sr].

#### 4. PCEP Extensions

##### 4.1. Extension of SRP object

SRP object extension for SYNC flag to specify the LABEL-DB synchronization operation is defined in [I-D.zhao-pce-pcep-extension-for-pce-controller].

##### 4.2. Extension of PCECC Capability sub-TLV

PCECC Capability sub-TLV is defined in [I-D.zhao-pce-pcep-extension-for-pce-controller]. This draft defines a new 'INCLUDE-LABEL-DB-VERSION' flag (I bit) to specify the label database version capability and 'DELTA-LABEL-SYNC-CAPABILITY' to specify the incremental label database synchronization capability.

The TLV format is as per [RFC5440]. The format of the PCECC Capability sub-TLV is shown Figure 5:

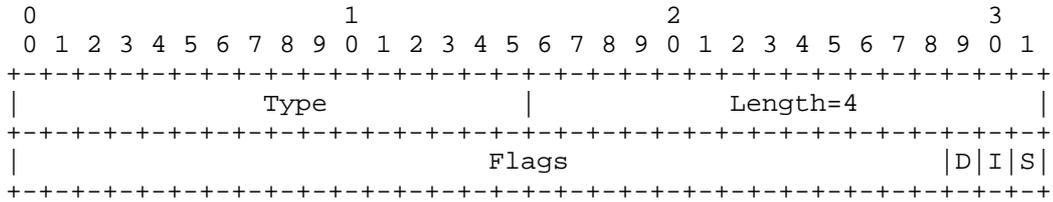


Figure 5: PCECC Capability sub-TLV

I (INCLUDE-LABEL-DB-VERSION - 1 bit): if set to 1 by both PCEP Speakers, the PCECC will include the LABEL-DB-VERSION TLV in each LABEL Object.

D (DELTA-LABEL-SYNC-CAPABILITY - 1 bit): if set to 1 by a PCEP speaker, it indicates that the PCEP speaker allows incremental (delta) LABEL-DB synchronization.

4.3. New LABEL-DB-VERSION TLV

The Label Database Version Number (LABEL-DB-VERSION) TLV is an optional TLV that MAY be included in the OPEN object and the LABEL object.

The TLV format is as per [RFC5440]. The format of the LABEL-DB-VERSION TLV is shown in the following figure:



Figure 6: LABEL-DB-VERSION TLV format

The type of the TLV is [TBD3] and it has a fixed length of 8 octets. The value contains a 64-bit unsigned integer, representing the Label Database Version Number.

## 5. IANA Considerations

### 5.1. PCEP TLV Type Indicators

IANA is requested to confirm the early allocation of the following TLV Type Indicator values within the "PCEP TLV Type Indicators" sub-registry of the PCEP Numbers registry, and to update the reference in the registry to point to this document, when it is an RFC:

Value	Meaning	Reference
[TBD]	LABEL-DB-VERSION TLV	This document

### 5.2. PCECC-CAPABILITY sub-TLV

[I-D.zhao-pce-pcep-extension-for-pce-controller] defines the PCECC-CAPABILITY sub-TLV and [I-D.zhao-pce-pcep-extension-pce-controller-sr] extends this sub-TLV to add PCECC-SR-CAPABILITY.

Requests that IANA creates a registry to manage the value of the new PCECC-CAPABILITY sub-TLV's Flag field. IANA is requested to allocate a new bits in the PCECC-CAPABILITY sub-TLV Flag Field registry, as follows:

Bit	Description	Reference
TBD	I (INCLUDE-LABEL-DB-VERSION )	This document
TBD	D (DELTA-LABEL-SYNC-CAPABILITY)	This document

## 6. Manageability Considerations

All manageability requirements and considerations listed in [RFC5440], [RFC8231] and [I-D.zhao-pce-pcep-extension-for-pce-controller] apply to PCEP protocol extensions defined in this document. In addition, requirements and considerations listed in this section apply.

### 6.1. Control of Function and Policy

### 6.2. Information and Data Models

### 6.3. Liveness Detection and Monitoring

### 6.4. Verify Correct Operations

## 6.5. Requirements On Other Protocols

## 6.6. Impact On Network Operations

## 7. Security Considerations

The security considerations listed in [RFC8231] and [I-D.zhao-pce-pcep-extension-for-pce-controller] apply to this document as well. Securing the PCEP session using Transport Layer Security (TLS) [RFC8253], as per the recommendations and best current practices in [RFC7525], is RECOMMENDED.

## 8. Acknowledgements

This document borrows some of the structure and text from [RFC8253], and would like to thank the authors and contributors of the document.

## 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>>.
- [I-D.zhao-pce-pcep-extension-for-pce-controller] Zhao, Q., Li, Z., Dhody, D., Karunanithi, S., Farrel, A., and C. Zhou, "PCEP Procedures and Protocol Extensions for Using PCE as a Central Controller (PCECC) of LSPs", draft-zhao-pce-pcep-extension-for-pce-controller-06 (work in progress), October 2017.

## 9.2. Informative References

- [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>>.
- [RFC8233] Dhody, D., Wu, Q., Manral, V., Ali, Z., and K. Kumaki, "Extensions to the Path Computation Element Communication Protocol (PCEP) to Compute Service-Aware Label Switched Paths (LSPs)", RFC 8233, DOI 10.17487/RFC8233, September 2017, <<https://www.rfc-editor.org/info/rfc8233>>.
- [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>>.
- [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.zhao-pce-pcep-extension-pce-controller-sr]  
Zhao, Q., Li, Z., Dhody, D., Karunanithi, S., Farrel, A., and C. Zhou, "PCEP Procedures and Protocol Extensions for Using PCE as a Central Controller (PCECC) of SR-LSPs", draft-zhao-pce-pcep-extension-pce-controller-sr-01 (work in progress), October 2017.

## Authors' Addresses

Udayasree Palle  
Huawei Technologies  
Divyashree Techno Park, Whitefield  
Bangalore, Karnataka 560066  
India

EMail: [udayasreereddy@gmail.com](mailto:udayasreereddy@gmail.com)

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

EMail: dhruv.ietf@gmail.com

Satish Karunanithi  
Huawei Technologies  
Divyashree Techno Park, Whitefield  
Bangalore, Karnataka 560066  
India

EMail: satishk@huawei.com