PCE Working Group Internet-Draft Intended status: Standards Track Expires: April 11, 2016

Optimizations of PCEP Link-State(LS) Synchronization Procedures draft-kondreddy-pce-pcep-ls-sync-optimizations-00

Abstract

For a Path Computation Element (PCE) to perform its computations, it is important that Link-State (and TE) information be complete and accurate each time. This requires a reliable Link-State Synchronization mechanism between the PCE and path computation clients (PCCs), and between cooperating PCEs. The basic mechanism for Link-State Synchronization is part of the PCEP Link-State (and TE) draft. This draft presents motivations for optimizations to the base PCEP Link-State (and TE) procedure and specifies the required Path Computation Element Communication Protocol (PCEP) extensions.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of <u>BCP 78</u> and <u>BCP 79</u>.

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 <u>http://datatracker.ietf.org/drafts/current/</u>.

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 April 11, 2016.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to <u>BCP 78</u> and the IETF Trust's Legal Provisions Relating to IETF Documents (<u>http://trustee.ietf.org/license-info</u>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect

Kondreddy & Negi

Expires April 11, 2016

[Page 1]

Internet-Draft PCEP LS SYNC OPT

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

$\underline{1}$. Introduction
<u>1.1</u> . Requirements Language
<u>2</u> . Terminology
$\underline{3}$. LS Synchronization Avoidance
3.1. Motivation
3.2. LS Synchronization Avoidance Procedure
4. Incremental LS Synchronization
<u>4.1</u> . Motivation
4.2. Incremental Synchronization Procedure
5. PCE-triggered Initial Synchronization
<u>5.1</u> . Motivation
5.2. PCE-triggered Initial LS Synchronization Procedure <u>11</u>
<u>6</u> . PCE-triggered Re-synchronization
<u>6.1</u> . Motivation
6.2. PCE-triggered LS Re-synchronization Procedure <u>12</u>
<u>7</u> . PCEP Extensions
<u>7.1</u> . Link-State(LS) Report Message
<u>7.2</u> . Capability Advertisement
7.3. Advertising Support of Synchronization Optimizations <u>14</u>
<u>8</u> . IANA Considerations
<u>8.1</u> . PCEP-Error Object
8.2. PCEP TLV Type Indicators
<u>8.3</u> . LS-CAPABILITY Flags
$\underline{9}$. Manageability Considerations
<u>9.1</u> . Control of Function and Policy \ldots \ldots \ldots \ldots $\frac{17}{2}$
<u>9.2</u> . Information and Data Models
<u>9.3</u> . Liveness Detection and Monitoring
<u>9.4</u> . Verify Correct Operations
9.5. Requirements On Other Protocols
9.6. Impact On Network Operations
<u>10</u> . Security Considerations
<u>11</u> . Ackwonoledgement
<u>12</u> . References
<u>12.1</u> . Normative References
<u>12.2</u> . Informative References
Appendix A. Contributor Addresses
Authors' Addresses

<u>1</u>. Introduction

The Path Computation Element Communication Protocol (PCEP) provides mechanisms for Path Computation Elements (PCEs) to perform path computations in response to Path Computation Clients (PCCs) requests.

[I-D.dhodylee-pce-pcep-ls] describes a set of extensions to PCEP to provide Link-State (and TE) distribution. This draft presents motivations for optimizations to the base PCEP Link-State transport procedure and specifies the required Path Computation Element Communication Protocol (PCEP) extensions. This draft specifies following optimizations for Link-State Synchronization and the corresponding PCEP procedures and extensions:

- Link-State Synchronization Avoidance: To skip Link-State (and TE) synchronization if the state has survived and not changed during session restart.(See <u>Section 3</u>)
- o Incremental Link-State Synchronization: To do incremental (delta) Link-State (and TE) Synchronization when possible.(See Section 4)
- o PCE-triggered Initial Synchronization: To let PCE control the timing of the initial Link-State (and TE) Synchronization.(See Section 5)
- o PCE-triggered Re-synchronization: To let PCE re-synchronize the Link-State (and TE) information for sanity check.(See Section 6)

<u>1.1</u>. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [<u>RFC2119</u>].

2. Terminology

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

This document uses the following terms defined in [<u>I-D.dhodylee-pce-pcep-ls</u>]: LSRpt, Link-State (and TE), LS

Link-State (and TE): "LS" is interchangably used for the keyword "Link-State (and TE)".

Within this document, when describing PCE-PCE communications, the requesting PCE fills the role of a PCC. This provides a saving in documentation without loss of function.

The following terms are defined in this document:

LS-DB: Link-State (and TE) Database.

LS Sync: LS Synchronization, an operation to send LS information synchronization request to PCC by LSRpt message with LS-ID 0.

LS-Info: One LS information (i.e. LS Node/Link/Prefix defined in I-D.dhodylee-pce-pcep-ls).

3. LS Synchronization Avoidance

<u>3.1</u>. Motivation

The purpose of LS Synchronization is to provide a checkpoint-in-time state replica of a PCC's Link-State (and TE) information in a PCE. LS Synchronization is performed immediately after the initialization phase ([<u>RFC5440</u>]). [<u>I-D.dhodylee-pce-pcep-ls</u>] describes the basic mechanism for LS synchronization.

LS Synchronization is not always necessary following a PCEP session restart. If the Link-State (and TE) information of both PCEP peers did not change, the synchronization phase may be skipped. This can result in significant savings in both control-plane data exchanges and the time it takes for the PCE to become fully operational.

3.2. LS Synchronization Avoidance Procedure

LS Synchronization MAY be skipped following a PCEP session restart if the Link-State (and TE) information of both PCEP peers did not change during the period prior to session re-initialization. To be able to make this determination, LS-DB must be exchanged and maintained by both PCE and PCC during normal operation. This is accomplished by keeping track of the changes to the Link-State (and TE) Database (LS-DB), using a version tracking field called the LS-DB Version Number.

The LS-DB Version Number, carried in LS-DB-VERSION TLV (see <u>Section 7.2</u>), is owned by a PCC and it MUST be incremented by 1 for each successive change in the PCC's LS-DB. The LS-DB Version Number MUST start at 1 and may wrap around. Values 0 and 0xFFFFFFFFFFFFFFFFF are reserved. If either of the two values are used during LS resynchronization, the PCE speaker receiving this node should send back a PCErr with Error-type TBD1 Error-value 1 'Received an invalid LS-DB Version Number', and close the PCEP session. Operations that trigger a change to the local LS-DB include but not limited to -

- a change in the link status or attributes(i.e. bandwidth, metric etc), addition or deletion of link.

- a change in the node attributes, addition or deletion of node.

- a change in the prefix attributes, addition or deletion of prefix.

LS Synchronization avoidance is advertised on a PCEP session during session startup using the LS-INCLUDE-DB-VERSION (S) bit in the LS-CAPABILITY TLV (see <u>Section 7.3</u>). The peer may move in the network, either physically or logically, which may cause its connectivity details and transport-level identity (such as IP address) to change. To ensure that a PCEP peer can recognize a previously connected peer even in case of such mobility, each PCEP peer includes the SPEAKER-ENTITY-ID TLV in the OPEN message. SPEAKER-ENTITY-ID TLV is described in [I-D.ietf-pce-stateful-sync-optimizations]

If both PCEP speakers set the S flag in the OPEN object's LS-CAPABILITY TLV to 1, the PCC MUST include the LS-DB-VERSION TLV in each LS object of the LSRpt message. If the LS-DB-VERSION TLV is missing in a LSRpt message, the PCE will generate an error with Error-Type 6 (mandatory object missing) and Error-Value TBD2 'LS-DB-VERSION TLV missing' and close the session. If LS Synchronization avoidance has not been enabled on a PCEP session, the PCC SHOULD NOT include the LS-DB-VERSION TLV in the LS object and the PCE SHOULD ignore it, if it were to receive one.

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

If both PCEP speakers include the LS-DB-VERSION TLV in the OPEN Object and the TLV values match, the PCC MAY skip LS Synchronization. Otherwise, the PCC MUST perform complete LS Synchronization. If the PCC attempts to skip LS Synchronization (i.e., the SYNC Flag = 0 on the first LS Report from the PCC, the PCE MUST send back a PCErr with Error-Type TBD1 Error-Value 2 'LS-DB version mismatch', and close the PCEP session.

If LS Synchronization is required, then prior to completing the initialization phase, the PCE MUST mark any LS-Infos in the LS-DB that were previously reported by the PCC as stale. When the PCC reports a LS-Info during LS Synchronization, if the LS-Info already exists in the LS-DB, the PCE MUST update the LS-DB and clear the stale marker from the LS-Info. When it has finished LS

Internet-Draft

Synchronization, the PCC MUST immediately send an end of LS Synchronization marker. The end of synchronization marker is a LS Report (LSRpt) message with an LS object containing a LS-ID of 0 and with the SYNC flag set to 0. The LS-DB-VERSION TLV MUST be included in this LSRpt message. On receiving this LS Report, the PCE MUST purge any LS-Infos from the LS-DB that are still marked as stale. It should be noted that PCE may receive the same Link-state and TE information from multiple PCCs and the purging should take that into account.

Note that a PCE/PCC MAY force LS Synchronization by not including the LS-DB-VERSION TLV in its OPEN object.

Since a PCE does not make changes to the LS-DB Version Number, a PCC should never encounter this TLV in a message from the PCE (other than the OPEN message). A PCC SHOULD ignore the LS-DB-VERSION TLV, were it to receive one from a PCE.

Figure 1 shows an example sequence where the LS synchronization is skipped.



Figure 1: LS Synchronization Skipped

Internet-Draft

Figure 2 shows an example sequence where the LS Synchronization is performed due to LS-DB Version mismatch during the PCEP session setup. Note that the same LS Synchronization sequence would happen if either the PCC or the PCE would not include the LS-DB-VERSION TLV in their respective Open messages.

+	+ - +	+-+-+
P(CCI	PCE
+	+-+	+-+-+
(Do sync)	 Open, LS-DBv=86 \ ,Open S=1 \ /LS-DBv=82 \/ S=1 /\ /\ <`	 > (Expect sync)
. ,	 LSRpt,LS-DBv=86,SYNC=1 	> (Sync start)
	LSRpt,LS-DBv=86,SYNC=0 	> (Sync done) (Purge LS-Info if applicable)
	LSRpt,LS-DBv=87,SYNC=0 LSRpt,LS-DBv=88,SYNC=0 LSRpt,LS-DBv=89,SYNC=0	> (Regular LS Report) > (Regular LS Report) >
	l	I

Figure 2: LS Synchronization Performed

Figure 3 shows an example sequence where the LS Synchronization is skipped, but because one or both PCEP speakers set the S Flag to 0, the PCC does not send LS-DB-VERSION TLVs in subsequent LSRpt messages to the PCE. If the current PCEP session restarts, the PCEP speakers will have to perform LS Synchronization, since the PCE does not know the PCC's latest LS-DB Version Number information.



Figure 3: LS Synchronization Skipped, no LS-DB-VERSION TLVs sent from PCC

<u>4</u>. Incremental LS Synchronization

[I-D.dhodylee-pce-pcep-ls] describes the LS synchronization mechanism during session initialization between PCCs and PCEs. During the LS synchronization, a PCC sends the information of its LS-DB to the PCE based on the local policy. In order to reduce the LS Synchronization overhead when there is a small number of LS-DB change in the network between PCEP session restart, this section defines a mechanism for incremental (Delta) LS synchronization.

<u>4.1</u>. Motivation

According to [I-D.dhodylee-pce-pcep-ls], if a PCEP session restarts, PCCs send snapshot of LS-DB information to the PCE, though LS-DB did not change. And as per <u>Section 3</u> (LS Synchronization Avoidance Procedure), if there is a change in a small number of LS-Infos. PCC yet sends complete snapshot of LS-DB information to the PCE, which takes a long time and consume large communication channel bandwidth.

4.2. Incremental Synchronization Procedure

<u>Section 3</u> describes LS Synchronization avoidance by using LS-DB-VERSION TLV in its OPEN object. This section extends this idea to only synchronize the delta (changes) in case of version mismatch.

If both PCEP speakers include the LS-DB-VERSION TLV in the OPEN object and the LS-DB-VERSION TLV values match, the PCC MAY skip LS Synchronization. Otherwise, the PCC MUST perform LS Synchronization. Incremental LS Synchronization capability is advertised on a PCEP session during session startup using the LS-DELTA-SYNC-CAPABILITY (D) bit in the capabilities TLV (see <u>Section 7.3</u>). Instead of dumping full LS-DB to the PCE again, PCC synchronizes the delta (changes) as described in Figure 4 when D flag and S flag is set to 1 by both PCC and PCE. Other combinations of D and S flags setting by PCC and PCE result in complete LS Synchronization procedure as described in [<u>I-D.dhodylee-pce-pcep-ls</u>]. If a PCC has to force complete LS Synchronization due to reasons including but not limited: (1) local policy configured at the PCC; (2) no sufficient LS-DB caches for incremental update, the PCC can set the D flag to 0. Note a PCC may have to bring down the current session and force a complete LS Synchronization with D flag set to 0 in the subsequent open message.

+-+-+ +-+-+ | PCC | |PCE| +-+-+ + - + - + |---Open--, |LS-DBv=86 \ ,---Open-----| S=1 \ /LS-DBv=82 D=1 \/ S=1 \land D=1 ---->| (Expect Delta sync) (Do |<----` | (Donot purge Delta | | LS-Infos) sync) | |---LSRpt,LS-DBv=86,SYNC=1-->| (Delta Sync start) . |---LSRpt,LS-DBv=86,SYNC=0-->| (Sync done) | LS-ID=0) |---LSRpt,LS-DBv=87,SYNC=0-->| (Regular | LS Report) |---LSRpt,LS-DBv=88,SYNC=0-->| (Regular | LS Report) |---LSRpt,LS-DBv=89,SYNC=0-->|

Figure 4: Incremental Synchronization Procedure

As per <u>Section 3</u>, the LS-DB Version Number is incremented each time a change is made to the PCC's local LS-DB. Each LS-Info is associated with the DB version at the time of change. This is needed to determine which LS-Info needs to be synchronized in incremental LS Synchronization.

PCC MAY store a then history of LS-DB change that happened between the PCEP session(s) restart in order to carry out incremental LS Synchronization. After the synchronization procedure finishes, the PCC can dump this history information. In the example shown in Figure 4, the PCC needs to store the LS-DB changes that happened between DB Version 83 to 86 and synchronizes these changes only when performing incremental LS-DB update. So a PCC needs to remember the LS-DB changes that happened when an existing PCEP session to a PCE goes down in the hope of doing incremental synchronization when the session is re-established.

If a PCC finds out it does not have sufficient information to complete incremental synchronization after advertising incremental LS Synchronization capability, it MUST send a PCErr with Error-Type TBD1 and Error-Value 3 'A PCC indicates to a PCE that it can not complete the LS synchronization' and terminate the session.

<u>5</u>. PCE-triggered Initial Synchronization

5.1. Motivation

In networks such as optical transport networks, the control channel between network nodes can be realized through in-band overhead thus has limited bandwidth. With a PCE connected to the network via one network node, it is desirable to control the timing of PCC LS Synchronization so as not to overload the low communication channel available in the network during the initial synchronization (be it incremental or full) when the session restarts, when there is comparatively large amount of control information needing to be synchronized between the PCE and the network. The method proposed, i.e., allowing PCE to trigger the LS synchronization, is similar to the function proposed in <u>Section 6</u> but is used in different scenarios and for different purposes.

5.2. PCE-triggered Initial LS Synchronization Procedure

Support of PCE-triggered LS Synchronization is advertised during session startup using the LS-TRIGGERED-INITIAL-SYNC (F) bit in the LS-CAPABILITY TLV (see <u>Section 7.3</u>).

As per [I-D.dhodylee-pce-pcep-ls], LSRpt is sent from PCC to PCE, this document extends the usage of LSRpt to trigger syncronization. Where a PCC can send a LSRpt (for LS Sync) with an LS object containing a LS-ID of 0 and with the SYNC flag set to 1. This LSRpt message is the trigger for the PCC to enter the synchronization phase and start sending LSRpt messages.

If the LS-TRIGGERED-INITIAL-SYNC capability is not advertised and the PCC receives a LSRpt with the SYNC flag set to 1, it MUST send a PCErr for LSRpt(LS Sync from PCE) with Error-Type TBD1 and Error-Value 4 'Attempt to trigger synchronization when the PCE triggered synchronization capability has not been advertised'.

A PCE MAY choose to control the LS Synchronization process. To allow PCE to do so, PCEP speakers MUST set T bit to 1 to indicate this (as described in Section 7.3). If the LS-DB version is mis-matched, it can send a LSRpt message with LS-ID = 0 and SYNC = 1 in order to trigger the LS Synchronization process. In this way, the PCE can control the sequence of LS Synchronization among all the PCCs that

are re-establishing PCEP sessions with it. When the capability of PCE control is enabled, only after a PCC receives this message, it will start sending information to the PCE. The PCC SHOULD NOT send LSRpt messages to the PCE before it triggers the LS Synchronization. This PCE-triggering capability can be applied to both full and incremental LS Synchronization. If applied to the later, the PCCs only send information that PCE does not possess, which is inferred from the LS-DB version information exchanged in the OPEN message (see <u>Section 3.2</u>) for detailed procedure).

6. PCE-triggered Re-synchronization

<u>6.1</u>. Motivation

The accuracy of the computations performed by the PCE is tied to the accuracy of the view the PCE has on the LS-DB. Therefore, it can be beneficial to be able to re-synchronize LS-DB even after the session has been established. The PCE may use this approach to continuously sanity check its LS-DB against the network, or to recover from error conditions without having to tear down sessions.

6.2. PCE-triggered LS Re-synchronization Procedure

Support of PCE-triggered LS Synchronization is advertised during session startup using the LS-TRIGGERED-RESYNC (T) bit in the LS-CAPABILITY TLV (see Section 7.3).

The PCE triggers re-synchronization of the entire LS-DB. The PCE MUST first mark all LS-Infos in the LS-DB that were previously reported by the PCC as stale and then send a LSRpt (for LS Sync) with an LS object containing a LS-ID of 0 and with the SYNC flag set to 1. This LSRpt message is the trigger for the PCC to enter the synchronization phase and start sending LSRpt messages. After the receipt of the end-of-synchronization marker, the PCE will purge LS-Infos which were not refreshed.

If the LS-TRIGGERED-RESYNC capability is not advertised and the PCC receives a LSRpt with the SYNC flag set to 1, it MUST send a PCErr with Error-Type TBD1 and Error-Value 4 'Attempt to trigger synchronization when the TRIGGERED-SYNC capability has not been advertised'.

Once the state re-synchronization is triggered by the PCE, the procedures and error checks remain unchanged from the full state synchronization ([I-D.dhodylee-pce-pcep-ls]). This would also include PCE triggering multiple state re-synchronization requests while synchronization is in progress.

```
+-+-+
                    +-+-+
| PCC |
                    | PCE |
+-+-+
                    +-+-+
 |<----LSRpt, SYNC=1 -----| Trigger</pre>
 | (LSID=0) | re-sync
 |----LSRpt, SYNC=1---->| (Sync start)
                      .
                     .
            .
 |-----LSRpt, SYNC=1---->|
            .
 .
                     .
 |-----LSRpt, SYNC=0---->| (End of sync marker
                     | LS Report
 | for LS-ID=0)
 | (LS Sync done)
```

Figure 5: PCE Triggered Complete LS re-synchronization

7. PCEP Extensions

7.1. Link-State(LS) Report Message

A PCEP LS Report message (also referred to as LSRpt message) is a PCEP message sent by a PCC to a PCE to report the LS information. The definition of the LSRpt message from [I-D.dhodylee-pce-pcep-ls] is extended to use LSRpt message with LS-ID = 0 to request LS Synchronization from PCE to PCC.

If a PCC that does not support extention defined in this document receives a LSRpt message, it will act according to existing behavior of receiving invalid message. If a PCC supports the extention, but did not set the flag T or F, and receives the LSRpt message, it sends PCErr message as described earlier in section [x] and [y]. If a PCC supports the extention and set the flag T or F, and receives the LSRpt message without LS-ID as 0 and SYNC flag set, PCC will send an error message with Error-Type TBD1 Error-Value 6 'Invalid LSRpt message'.

7.2. Capability Advertisement

The LS-DB Version Number is an carried in optional LS-DB-VERSION TLV that MAY be included in the OPEN object and the LS object. This TLV MUST NOT be included if LS-INCLUDE-DB-VERSION bit in LS-CAPABILITY TLV is not set.

The format of the LS-DB-VERSION TLV is shown in the following figure:

0 1 2 3 4 5 6 7 8 9 0 1 2

Figure 6: LS-DB-VERSION TLV format

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

7.3. Advertising Support of Synchronization Optimizations

Support for each of the optimizations described in this document requires advertising the corresponding capabilities during session establishment.

New flags are defined for the LS-CAPABILITY TLV defined in [<u>I-D.dhodylee-pce-pcep-ls</u>].



The value comprises a single field - Flags (32 bits):

R (LS-REMOTE-BIT - 1 bit): Defined in [I-D.dhodylee-pce-pcep-ls]

S (LS-INCLUDE-DB-VERSION - 1 bit): If set to 1 by both PCEP speakers, the PCC will include the LS-DB-VERSION TLV in each LS object. See <u>Section 3</u> for details.

T (LS-TRIGGERED-RESYNC - 1 bit): If set to 1 by both PCEP speakers, the PCE can trigger re-synchronization of LS-Infos at any point in the life of the session. See <u>Section 6</u> for details.

D (LS-DELTA-SYNC-CAPABILITY - 1 bit): If set to 1 by a PCEP speaker, it indicates that the PCEP speaker allows incremental (delta) state synchronization. See <u>Section 4</u> for details.

F (LS-TRIGGERED-INITIAL-SYNC - 1 bit): If set to 1 by both PCEP speakers, the PCE SHOULD trigger initial (first) LS synchronization. See <u>Section 5</u> for details.

8. IANA Considerations

This document requests IANA actions to allocate code points for the protocol elements defined in this document.

8.1. PCEP-Error Object

IANA is requested to make the following allocation in the "PCEP-ERROR Object Error Types and Values" registry.

Error-Type	pe Meaning		Reference	
6	Mandatory Object missing		[<u>RFC5440]</u>	
	Error-Value= TBD2	This	document	
	LS-DB-VERSION TLV missing			
TBD1	LS synchronization	This	document	
	error			
	Error-Value= TBD(suggested	This	document	
	value 1):Received an invalid			
	LSDB Version Number			
	Error-Value= TBD(suggested	This	document	
	value 2): LS-DB			
	version mismatch.			
	Error-Value=TBD(suggested	This	document	
	value 3): PCC indicates to a			
	PCE that it cannot complete			
	the LS Synchronization			
	Error-Value=TBD(suggested	This	document	
	value 4): Attempt to trigger a			
	synchronization when the			
	PCE triggered synchronization			
	capability has not been			
	advertised.			
	Error-Value=TBD(suggested	This	document	
	value 5): LS-DB-VERSION			
	ILV Missing when LS			
	synchronization avoidance is			
	enabled.			
	Error-value=IBD(suggested	Ihis	aocument	
	value 6): Invalid LSRpt			
	message.			

8.2. PCEP TLV Type Indicators

This document defines the following new PCEP TLVs:

Value	Meaning	Reference
TBD	LS-DB-VERSION	This document

8.3. LS-CAPABILITY Flags

The following values are defined in this document for the Flags field in the LS-CAPABILITY TLV in the OPEN object:

PCEP LS SYNC OPT

Bit			Description	Refer	rence
TBD					
(Suggested	value	30)	LS-INCLUDE-DB-VERSION	This	document
(Suggested	value	29)	LS-TRIGGERED-RESYNC	This	document
(Suggested	value	28)	LS-DELTA-SYNC-CAPABILITY	This	document
(Suggested	value	27)	LS-TRIGGERED-INITIAL-SYNC	This	document

9. Manageability Considerations

All manageability requirements and considerations listed in [<u>RFC5440</u>] and [<u>I-D.dhodylee-pce-pcep-ls</u>] 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 PCE or PCC implementation MUST allow configuring the state synchronization optimization capabilities as described in this document.

9.2. Information and Data Models

PCEP session configuration and information in the PCEP MIB module SHOULD be extended to include advertised LS Capabilities, LS-DB Version Number and LS synchronization status, Message statistics.

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 [<u>RFC5440</u>].

<u>9.4</u>. Verify Correct Operations

Mechanisms defined in <u>section 8.4 of [RFC5440]</u> also apply to PCEP protocol extensions defined in this document. In addition to monitoring parameters defined in [<u>RFC5440</u>], a PCEP implementation with LS-DB SHOULD provide the following parameters:

- o Total number of LSRpt(Synchronization request) requests
- o LS-DB Version Number and synchronization status

<u>9.5</u>. 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 <u>section 8.6 of [RFC5440]</u> also apply to PCEP protocol extensions defined in this document.

Additionally, a PCEP implementation SHOULD allow a limit to be placed on the amount and rate of LSRpt messages sent by a PCEP speaker and processed by the peer. It SHOULD also allow sending a notification when a rate threshold is reached.

<u>10</u>. Security Considerations

The security considerations listed in [<u>I-D.dhodylee-pce-pcep-ls</u>] apply to this document as well. However, because the protocol modifications outlined in this document allow the PCE to control LS Re-synchronization timing and sequence, it also introduces a new attack vector: an attacker may flood the PCC with triggered resynchronization request at a rate which exceeds the PCC's ability to process them, either by spoofing messages or by compromising the PCE itself. The PCC is free to drop any trigger re-synchronization request without additional processing.

<u>11</u>. Ackwonoledgement

The document borrows some of the text and structure from [I-D.ietf-pce-stateful-sync-optimizations].

<u>12</u>. References

<u>12.1</u>. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", <u>BCP 14</u>, <u>RFC 2119</u>, DOI 10.17487/RFC2119, March 1997, <<u>http://www.rfc-editor.org/info/rfc2119</u>>.
- [RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", <u>RFC 5440</u>, DOI 10.17487/RFC5440, March 2009, <<u>http://www.rfc-editor.org/info/rfc5440</u>>.

[I-D.dhodylee-pce-pcep-ls]

Dhody, D., Lee, Y., and D. Ceccarelli, "PCEP Extension for Distribution of Link-State and TE Information.", <u>draft-</u> <u>dhodylee-pce-pcep-ls-00</u> (work in progress), September 2015.

<u>12.2</u>. Informative References

[I-D.ietf-pce-stateful-sync-optimizations]

Crabbe, E., Minei, I., Medved, J., Varga, R., Zhang, X., and D. Dhody, "Optimizations of Label Switched Path State Synchronization Procedures for a Stateful PCE", draftietf-pce-stateful-sync-optimizations-03 (work in progress), October 2015.

<u>Appendix A</u>. Contributor Addresses

Dhruv Dhody Huawei Technologies Divyashree Techno Park, Whitefield Bangalore, Karnataka 560037 India

EMail: dhruv.ietf@gmail.com

Authors' Addresses

Venugopal Reddy Kondreddy Huawei Technologies Divyashree Techno Park, Whitefield Bangalore, Karnataka 560037 India

EMail: venugopalreddyk@huawei.com

Mahendra Singh Negi Huawei Technologies Divyashree Techno Park, Whitefield Bangalore, Karnataka 560037 India

EMail: mahendrasingh@huawei.com