PCE Working Group D. Dhody
Internet-Draft Y. Lee
Intended status: Experimental Huawei Technologies
Expires: September 5, 2015 D. Ceccarelli
Ericsson
March 4, 2015

PCEP Extension for Transporting TE Data
draft-dhodylee-pce-pcep-te-data-extn-02

Abstract

In order to compute and provide optimal paths, Path Computation Elements (PCEs) require an accurate and timely Traffic Engineering Database (TED). Traditionally this TED has been obtained from a link state routing protocol supporting traffic engineering extensions.

This document extends the Path Computation Element Communication Protocol (PCEP) with TED population capability.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

In Multiprotocol Label Switching (MPLS) and Generalized MPLS (GMPLS), a Traffic Engineering Database (TED) is used in computing paths for connection oriented packet services and for circuits. The TED contains all relevant information that a Path Computation Element (PCE) needs to perform its computations. It is important that the TED be complete and accurate each time, the PCE performs a path computation.

In MPLS and GMPLS, interior gateway routing protocols (IGPs) have been used to create and maintain a copy of the TED at each node running the IGP. One of the benefits of the PCE architecture [RFC4655] is the use of computationally more sophisticated path computation algorithms and the realization that these may need enhanced processing power not necessarily available at each node participating in an IGP.

Section 4.3 of [RFC4655] describes the potential load of the TED on a network node and proposes an architecture where the TED is maintained by the PCE rather than the network nodes. However, it does not describe how a PCE would obtain the information needed to populate its TED. PCE may construct its TED by participating in the IGP ([RFC3630] and [RFC5305] for MPLS-TE; [RFC4203] and [RFC5307] for GMPLS). An alternative is offered by BGP-LS [I-D.ietf-idr-ls-distribution] .

[I-D.lee-pce-transporting-te-data] proposes some other approaches for creating and maintaining the TED directly on a PCE as an alternative to IGPs and BGP flooding and investigate the impact from the PCE, routing protocol, and node perspectives.

[RFC5440] describes the specifications for the Path Computation Element Communication Protocol (PCEP). PCEP specifies the communication between a Path Computation Client (PCC) and a Path Computation Element (PCE), or between two PCEs based on the PCE architecture [RFC4655].

This document specifies a PCEP extension for TED population capability to support functionalities described in [I-D.lee-pce-transporting-te-data].

1.1. 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 [RFC2119].

2. Terminology

The terminology is as per [RFC4655] and [RFC5440].

3. Applicability

As per [I-D.lee-pce-transporting-te-data], the mechanism specified in this draft is applicable to:

- o Where there is no IGP-TE or BGP-LS running at the PCE to learn TED.
- o Where there is IGP-TE or BGP-LS running but with a need for a faster TED population and convergence at the PCE.
  - * A PCE may receive partial information (say basic TE) from IGP-TE and other information (optical and impairment) from PCEP.
  - * A PCE may receive full information from both IGP-TE and PCEP.

A PCC may further choose to send only local TE information or both local and remote learned TED information.

How a PCE manages the TED information is implementation specific and thus out of scope of this document.

4. Requirements for PCEP extension

Following key requirements associated with TED population are identified for PCEP:

1. The PCEP speaker supporting this draft MUST be a mechanism to advertise the TED capability.

2. PCC supporting this draft MUST have the capability to report the TED to the PCE. This includes self originated TE information and remote TE information learned via routing protocols. PCC MUST be capable to do the initial bulk sync at the time of session initialization as well as changes to TED after.

3. A PCE MAY learn TED from PCEP as well as from existing mechanism like IGP-TE/BGP-LS. PCEP extension MUST have a mechanism to link

the TED information learned via other means. There MUST NOT be any changes to the existing TED population mechanism via IGP-TE/ BGP-LS. PCEP extension SHOULD keep the TE properties in a routing protocol (IGP-TE or BGP-LS) neutral way, such that an implementation which do want to learn about a Link-state topology do not need to know about any OSPF or IS-IS or BGP protocol specifics.

4. It SHOULD be possible to encode only the changes in TED properties (after the initial sync) in PCEP messages.

5. The same mechanism should be used for both MPLS TE as well as GMPLS, optical and impairment aware properties.

6. The extension in this draft SHOULD be extensible to support various architecture options listed in [I-D.lee-pce-transporting-te-data].

5. New Functions to Support TED via PCEP

Several new functions are required in PCEP to support TED population. A function can be initiated either from a PCC towards a PCE (C-E) or from a PCE towards a PCC (E-C). The new functions are:

- o Capability advertisement (E-C,C-E): both the PCC and the PCE must announce during PCEP session establishment that they support PCEP extensions for TED population defined in this document.

- o TE synchronization (C-E): after the session between the PCC and a PCE is initialized, the PCE must learn PCC's TED before it can perform path computations. In case of stateful PCE it is RECOMENDED that this operation be done before LSP state synchronization.

- o TE Report (C-E): a PCC sends a TE report to a PCE whenever the TED changes.

6. Overview of Extension to PCEP

6.1. New Messages

In this document, we define a new PCEP messages called TE Report (TERpt), a PCEP message sent by a PCC to a PCE to report TED. Each TE Report in a TERpt message can contain the TE node or TE Link properties. An unique PCEP specific TE identifier (TE-ID) is also carried in the message to identify the TE node or link and that remains constant for the lifetime of a PCEP session. This identifier on its own is sufficient when no IGP-TE or BGP-LS running in the

network for PCE to learn TED. Incase PCE learns some information from PCEP and some from the existing mechanism, the PCC SHOULD include the mapping of IGP-TE or BGP-LS identifier to map the TED information populated via PCEP with IGP-TE/BGP-LS. See Section 8.1 for details.

6.2. Capability Advertisement

During PCEP Initialization Phase, PCEP Speakers (PCE or PCC) advertise their support of TED population PCEP extensions. A PCEP Speaker includes the "TED Capability" TLV, described in Section 9, in the OPEN Object to advertise its support for PCEP TED extensions. The presence of the TED Capability TLV in PCC's OPEN Object indicates that the PCC is willing to send TE Reports whenever local TE information changes. The presence of the TED Capability TLV in PCE's

OPEN message indicates that the PCE is interested in receiving TE Reports whenever local TE changes.

The PCEP protocol extensions for TED population MUST NOT be used if one or both PCEP Speakers have not included the TED Capability TLV in their respective OPEN message. If the PCE that supports the extensions of this draft but did not advertise this capability, then upon receipt of a PCRpt message from the PCC, it SHOULD generate a PCErr with error-type 19 (Invalid Operation), error-value TBD1 (Attempted TE Report if TED capability was not advertised) and it will terminate the PCEP session.

The TE reports sent by PCC MAY carry the remote TE information learned via existing means like IGP-TE and BGP-LS only if both PCEP Speakers set the R (remote) Flag in the "TED Capability" TLV to 'Remote Allowed (R Flag = 1)'. If this is not the case and TE reports carry remote TE information, then a PCErr with error-type 19 (Invalid Operation) and error-value TBD1 (Attempted TE Report if TED capability was not advertised) and it will terminate the PCEP session.

6.3. Initial TED Synchronization

The purpose of TED Synchronization is to provide a checkpoint-in-time state replica of a PCC's TED in a PCE. State Synchronization is performed immediately after the Initialization phase (see [RFC5440]]). In case of stateful PCE ([I-D.ietf-pce-stateful-pce]) it is RECOMENDED that the TED synchronization should be done before LSP state synchronization.

During TED Synchronization, a PCC first takes a snapshot of the state of its TED, then sends the snapshot to a PCE in a sequence of TE Reports. Each TE Report sent during TE Synchronization has the SYNC

Flag in the TE Object set to 1. The end of synchronization marker is a TERpt message with the SYNC Flag set to 0 for an TE Object with TED-ID equal to the reserved value 0. If the PCC has no TED state to synchronize, it will only send the end of synchronization marker.

Either the PCE or the PCC MAY terminate the session using the PCEP session termination procedures during the synchronization phase. If the session is terminated, the PCE MUST clean up state it received from this PCC. The session re-establishment MUST be re-attempted per the procedures defined in [RFC5440], including use of a back-off timer.

If the PCC encounters a problem which prevents it from completing the TED population, it MUST send a PCErr message with error-type TBD2 (TE Synchronization Error) and error-value 5 (indicating an internal PCC error) to the PCE and terminate the session.

The PCE does not send positive acknowledgements for properly received TED synchronization messages. It MUST respond with a PCErr message with error-type TBD2 (TE Synchronization Error) and error-value 1 (indicating an error in processing the TERpt) if it encounters a problem with the TE Report it received from the PCC and it MUST terminate the session.

The TE reports may carry local as well as remote TED information depending on the R flag in TED capability TLV.

The successful TED Synchronization sequences is shown in Figure 1.


```
                   +-+-+                    +-+-+
```

```
                   |PCC|                    |PCE|
                   +-+-+                    +-+-+
                     |                        |
                     |-----TERpt, SYNC=1----->| (Sync start)
                     |                        |
                     |-----TERpt, SYNC=1----->|
                     |            .           |
                     |            .           |
                     |            .           |
                     |-----TERpt, SYNC=1----->|
                     |            .           |
                     |            .           |
                     |            .           |
                     |                        |
                     |-----TERpt, SYNC=0----->| (End of sync marker
                     |                        |  TE Report
                     |                        |  for TED-ID=0)
                     |                        | (Sync done)
```

Figure 1: Successful state synchronization

The sequence where the PCE fails during the TED Synchronization phase is shown in Figure 2.

```
                   +-+-+                    +-+-+
                   |PCC|                    |PCE|
                   +-+-+                    +-+-+
                     |                        |
                     |-----TERpt, SYNC=1----->|
                     |                        |
                     |-----TERpt, SYNC=1----->|
                     |            .           |
                     |            .           |
                     |            .           |
                     |-----TERpt, SYNC=1----->|
                     |                        |
                     |-TERpt, SYNC=1          |
                     |         \    ,-PCErr---|
                     |          \  /          |
                     |           \/           |
                     |           /\           |
                     |          /   `-------->| (Ignored)
                     |<--------`              |
```

Figure 2: Failed TED synchronization (PCE failure)

The sequence where the PCC fails during the TED Synchronization phase is shown in Figure 3.

```
         +-+-+                    +-+-+
         |PCC|                    |PCE|
         +-+-+                    +-+-+
           |                        |
           |-----TERpt, SYNC=1----->|
           |                        |
           |-----TERpt, SYNC=1----->|
           |            .           |
           |            .           |
           |            .           |
           |-------- PCErr--------->|
           |                        |
```

Figure 3: Failed TED synchronization (PCC failure)

6.3.1. Optimizations for TED Synchronization

TBD

6.4. TE Report

The PCC MUST report any changes in the TEDB to the PCE by sending a TE Report carried on a TERpt message to the PCE, indicating that the TE state. Each TE node and TE Link would be uniquely identified by a PCEP TE identifier (TE-ID). The TE reports may carry local as well as remote TED information depending on the R flag in TED capability TLV. In case R flag is set, It MAY also include the mapping of IGP-TE or BGP-LS identifier to map the TED information populated via PCEP with IGP-TE/BGP-LS.

More details about TERpt message are in Section 8.1.

7. Transport

A permanent PCEP session MUST be established between a PCE and PCC supporting TED population via PCEP. In the case of session failure, session re-establishment MUST be re-attempted per the procedures defined in [RFC5440].

8. PCEP Messages

As defined in [RFC5440], a PCEP message consists of a common header followed by a variable-length body made of a set of objects that can be either mandatory or optional. An object is said to be mandatory

in a PCEP message when the object must be included for the message to be considered valid. For each PCEP message type, a set of rules is defined that specify the set of objects that the message can carry. An implementation MUST form the PCEP messages using the object ordering specified in this document.

8.1. TE Report Message

A PCEP TE Report message (also referred to as TERpt message) is a PCEP message sent by a PCC to a PCE to report the TED state. A TERpt message can carry more than one TE Reports. The Message-Type field of the PCEP common header for the PCRpt message is set to [TBD3].

The format of the PCRpt message is as follows:

```
<TERpt Message> ::= <Common Header>
                    <te-report-list>
Where:

<te-report-list> ::= <TE>[<te-report-list>]
```

The TE object is a mandatory object which carries TE information of a TE node or a TE link. Each TE object has an unique TE-ID as described in Section 9.2. If the TE object is missing, the receiving PCE MUST send a PCErr message with Error-type=6 (Mandatory Object missing) and Error-value=[TBD4] (TE object missing).

A PCE may choose to implement a limit on the TE information a single PCC can populate. If a TERpt is received that causes the PCE to exceed this limit, it MUST send a PCErr message with error-type 19 (invalid operation) and error-value 4 (indicating resource limit exceeded) in response to the TERpt message triggering this condition and MAY terminate the session.

8.2. The PCErr Message

If a PCEP speaker has advertised the TED capability on the PCEP session, the PCErr message MAY include the TE object. If the error reported is the result of an TE report, then the TE-ID number MUST be the one from the TERpt that triggered the error.

The format of a PCErr message from [RFC5440] is extended as follows:

The format of the PCRpt message is as follows:


```
<PCErr Message> ::= <Common Header>
                  ( <error-obj-list> [<Open>] ) | <error>
                  [<error-list>]

<error-obj-list>::=<PCEP-ERROR>[<error-obj-list>]

<error>::=[<request-id-list> | <te-id-list>]
           <error-obj-list>

<request-id-list>::=<RP>[<request-id-list>]

<te-id-list>::=<TE>[<te-id-list>]

<error-list>::=<error>[<error-list>]
```

9. Objects and TLV

The PCEP objects defined in this document are compliant with the PCEP object format defined in [RFC5440]. The P flag and the I flag of the PCEP objects defined in this document MUST always be set to 0 on transmission and MUST be ignored on receipt since these flags are exclusively related to path computation requests.

9.1. Open Object

This document defines a new optional TLV for use in the OPEN Object.

9.1.1. TED Capability TLV

The TED-CAPABILITY TLV is an optional TLV for use in the OPEN Object for TED population via PCEP capability advertisement. Its format is shown in the following figure:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|               Type=[TBD5]       |           Length=4          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                             Flags                           |R|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The type of the TLV is [TBD5] and it has a fixed length of 4 octets.

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

o R (remote - 1 bit): if set to 1 by a PCC, the R Flag indicates that the PCC allows reporting of remote TED information learned via other means like IGP-TE and BGP-LS; if set to 1 by a PCE, the

R Flag indicates that the PCE is capable of receiving remote TED information (from the PCC point of view). The R Flag must be advertised by both a PCC and a PCE for TERpt messages to report remote as well as local TE information on a PCEP session. The TLVs related to IGP-TE/BGP-LS identifier MUST be encoded when both PCEP speakers have the R Flag set.

Unassigned bits are considered reserved. They MUST be set to 0 on transmission and MUST be ignored on receipt.

Advertisement of the TED capability implies support of local TE population, as well as the objects, TLVs and procedures defined in this document.

9.2. TE Object

The TE (traffic engineering) object MUST be carried within TERpt messages and MAY be carried within PCErr messages. The TE object contains a set of fields used to specify the target TE node or link. It also contains a flag indicating to a PCE that the TED synchronization is in progress. The TLVs used with the TE object correlate with the IGP-TE/BGP-LS TE encodings.

TE Object-Class is [TBD6].

Two Object-Type values are defined for the TE object:

o TE Node: TE Object-Type is 1.

o TE Link: TE Object-Type is 2.

The format of the TE object body is as follows:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|  Protocol-ID  |          Flag                             |R|S|
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                          TE-ID                                |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
//                         TLVs                                //
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Protocol-ID (8-bit): The field provide the source information. Incase PCC only provides local information (R flag is not set), it MUST use Protocol-ID as Direct. The following values are defined (same as [I-D.ietf-idr-ls-distribution]):

Protocol-ID	Source protocol
1	IS-IS Level 1
2	IS-IS Level 2
3	OSPFv2
4	Direct
5	Static configuration
6	OSPFv3

Flags (32-bit):

o S (SYNC - 1 bit): the S Flag MUST be set to 1 on each TERpt sent from a PCC during TED Synchronization. The S Flag MUST be set to 0 in other TERpt messages sent from the PCC.

o R (Remove - 1 bit): On TERpt messages the R Flag indicates that the TE node/link has been removed from the PCC and the PCE SHOULD remove from its database. Upon receiving an TE Report with the R Flag set to 1, the PCE SHOULD remove all state for the TE node/link identified by the TE Identifiers from its database.

TE-ID(32-bit): A PCEP-specific identifier for the TE node or link. A PCC creates a unique TE-ID for each TE node/link that is constant for the lifetime of a PCEP session. The PCC will advertise the same TE-ID on all PCEP sessions it maintains at a given times. All subsequent PCEP messages then address the TE node/link by the TE-ID. The values of 0 and 0xFFFFFFFF are reserved.

Unassigned bits are considered reserved. They MUST be set to 0 on transmission and MUST be ignored on receipt.

TLVs that may be included in the TE Object are described in the following sections.

9.2.1. Routing Universe TLV

In case of remote TED population when existing IGP-TE/BGP-LS are also used, OSPF and IS-IS may run multiple routing protocol instances over the same link as described in [I-D.ietf-idr-ls-distribution]. See [RFC6822] and [RFC6549]. These instances define independent "routing universes". The 64-Bit 'Identifier' field is used to identify the "routing universe" where the TE object belongs. The TE objects representing IGP objects (nodes or links) from the same routing universe MUST have the same 'Identifier' value; TE objects with different 'Identifier' values MUST be considered to be from different routing universes.

The format of the ROUTING-UNIVERSE TLV is shown in the following figure:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           Type=[TBD7]         |             Length=8          |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                             Identifier                        |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Below table lists the 'Identifier' values that are defined as well-known in this draft (same as [I-D.ietf-idr-ls-distribution]).

Identifier	Routing Universe
0	L3 packet topology
1	L1 optical topology

If this TLV is not present the default value 0 is assumed.

9.2.2. Local TE Node Descriptors TLV

As described in [I-D.ietf-idr-ls-distribution], each link is anchored by a pair of Router-IDs that are used by the underlying IGP, namely, 48 Bit ISO System-ID for IS-IS and 32 bit Router-ID for OSPFv2 and OSPFv3. Incase of additional auxiliary Router-IDs used for TE, these MUST also be included in the TE link attribute TLV (see Section 9.2.6).

It is desirable that the Router-ID assignments inside the TE Node Descriptor are globally unique. Autonomous System (AS) Number and PCEP-TED Identifier in order to disambiguate the Router-IDs, as described in [I-D.ietf-idr-ls-distribution].

The Local TE Node Descriptors TLV contains Node Descriptors for the node anchoring the local end of the link. This TLV MUST be included in the TE Report when during a given PCEP session a TE node/link is first reported to a PCE. A PCC sends to a PCE the first TE Report either during State Synchronization, or when a new TE node/link is learned at the PCC. The value contains one or more Node Descriptor Sub-TLVs, which allows specification of a flexible key for any given Node/Link information such that global uniqueness of the TE node/link is ensured.


```
     0                   1                   2                   3
     0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
    +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
    |           Type=[TBD8]           |             Length          |
```

```
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
//              Node Descriptor Sub-TLVs (variable)            //
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The value contains one or more Node Descriptor Sub-TLVs defined in Section 9.2.4.

9.2.3. Remote TE Node Descriptors TLV

The Remote TE Node Descriptors contains Node Descriptors for the node anchoring the remote end of the link. This TLV MUST be included in the TE Report when during a given PCEP session a TE link is first reported to a PCE. A PCC sends to a PCE the first TE Report either during State Synchronization, or when a new TE link is learned at the PCC. The length of this TLV is variable. The value contains one or more Node Descriptor Sub-TLVs defined in Section 9.2.4.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           Type=[TBD9]         |             Length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
//              Node Descriptor Sub-TLVs (variable)            //
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

9.2.4. TE Node Descriptors Sub-TLVs

The Node Descriptor Sub-TLV type Type and lengths are listed in the following table:

Sub-TLV	Description	Length
TBD10	Autonomous System	4
TBD11	BGP-LS Identifier	4
TBD12	OSPF Area-ID	4
TBD13	Router-ID	Variable

The sub-TLV values in Node Descriptor TLVs are defined as follows (similar to [I-D.ietf-idr-ls-distribution]):

- o Autonomous System: opaque value (32 Bit AS Number)
- o BGP-LS Identifier: opaque value (32 Bit ID). In conjunction with ASN, uniquely identifies the BGP-LS domain as described in [I-D.ietf-idr-ls-distribution]. This sub-TLV is present only if the node implements BGP-LS and the ID is set by the operator.
- o Area ID: It is used to identify the 32 Bit area to which the TE object belongs. Area Identifier allows the different TE objects of the same router to be discriminated.
- o Router ID: opaque value. Usage is described in [I-D.ietf-idr-ls-distribution] for IGP Router ID. In case only local TE information is transported and PCE learns TED only from PCEP, it contain the unique local TE IPv4 or IPv6 router ID.
- o There can be at most one instance of each sub-TLV type present in any Node Descriptor.

9.2.5. TE Link Descriptors TLV

The TE Link Descriptors TLV contains Link Descriptors for each TE link. This TLV MUST be included in the TE Report when during a given PCEP session a TE link is first reported to a PCE. A PCC sends to a PCE the first TE Report either during State Synchronization, or when a new TE link is learned at the PCC. The length of this TLV is variable. The value contains one or more TE Link Descriptor Sub-TLVs

The 'TE Link descriptor' TLVs uniquely identify a link among multiple parallel links between a pair of anchor routers similar to [I-D.ietf-idr-ls-distribution].

```
  0                   1                   2                   3
  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |           Type=[TBD14]        |             Length            |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                                                               |
 //              Link Descriptor Sub-TLVs (variable)            //
 |                                                               |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The Link Descriptor Sub-TLV type and lengths are listed in the following table:

Sub-TLV	Description	IS-IS TLV /Sub-TLV	Value defined in:
TBD15	Link Local/Remote Identifiers	22/4	[RFC5307]/1.1
TBD16	IPv4 interface address	22/6	[RFC5305]/3.2
TBD17	IPv4 neighbor address	22/8	[RFC5305]/3.3
TBD18	IPv6 interface address	22/12	[RFC6119]/4.2
TBD19	IPv6 neighbor address	22/13	[RFC6119]/4.3

The format and semantics of the 'value' fields in most 'Link Descriptor' sub-TLVs correspond to the format and semantics of value fields in IS-IS Extended IS Reachability sub-TLVs, defined in [RFC5305], [RFC5307] and [RFC6119]. Although the encodings for 'Link Descriptor' TLVs were originally defined for IS-IS, the TLVs can carry data sourced either by IS-IS or OSPF or direct.

The information about a link present in the LSA/LSP originated by the local node of the link determines the set of sub-TLVs in the Link Descriptor of the link as described in [I-D.ietf-idr-ls-distribution].

9.2.6. TE Node Attributes TLV

This is an optional, non-transitive attribute that is used to carry TE node attributes. The TE node attribute TLV may be encoded in the TE node Object.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           Type=[TBD20]          |             Length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
```

```
   //              Node Attributes Sub-TLVs (variable)              //
   |                                                               |
   +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The Node Attributes Sub-TLV type and lengths are listed in the following table:

Sub TLV	Description	Length	Value defined in:
TBD21	Node Flag Bits	1	[I-D.ietf-idr-ls-distribution]/3.3.1.1
TBD22	Opaque Node Properties	variable	[I-D.ietf-idr-ls-distribution]/3.3.1.5
TBD23	Node Name	variable	[I-D.ietf-idr-ls-distribution]/3.3.1.3
TBD24	IS-IS Area Identifier	variable	[I-D.ietf-idr-ls-distribution]/3.3.1.2
TBD25	IPv4 Router-ID of Local Node	4	[RFC5305]/4.3
TBD26	IPv6 Router-ID of Local Node	16	[RFC6119]/4.1

9.2.7. TE Link Attributes TLV

TE Link attribute TLV may be encoded in the TE Link Object. The format and semantics of the 'value' fields in some 'Link Attribute' sub-TLVs correspond to the format and semantics of value fields in IS-IS Extended IS Reachability sub-TLVs, defined in [RFC5305], [RFC5307] and [I-D.ietf-idr-ls-distribution]. Although the encodings for 'Link Attribute' TLVs were originally defined for IS-IS, the TLVs can carry data sourced either by IS-IS or OSPF or direct.

```
    0                   1                   2                   3
```

```
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|           Type=[TBD27]        |             Length            |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                                                               |
//              Link Attributes Sub-TLVs (variable)            //
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

The following 'Link Attribute' sub-TLVs are are valid :

Sub-TLV	Description	IS-IS TLV /Sub-TLV BGP-LS TLV	Defined in:
TBD28	IPv4 Router-ID of Local Node	134/---	[RFC5305]/4.3
TBD29	IPv6 Router-ID of Local Node	140/---	[RFC6119]/4.1
TBD30	IPv4 Router-ID of Remote Node	134/---	[RFC5305]/4.3
TBD31	IPv6 Router-ID of Remote Node	140/---	[RFC6119]/4.1
TBD32	Link Local/Remote Identifiers	22/4	[RFC5307]/1.1
TBD33	Administrative group (color)	22/3	[RFC5305]/3.1
TBD34	Maximum link bandwidth	22/9	[RFC5305]/3.3
TBD35	Max. reservable link bandwidth	22/10	[RFC5305]/3.5
TBD36	Unreserved bandwidth	22/11	[RFC5305]/3.6
TBD37	TE Default Metric	22/18	[I-D.ietf-idr-ls-distribution] /3.3.2.3

TBD38	Link Protection Type	22/20	[RFC5307]/1.2
TBD39	MPLS Protocol Mask	1094	[I-D.ietf-idr-ls-distribution]/3.3.2.2
TBD40	IGP Metric	1095	[I-D.ietf-idr-ls-distribution]/3.3.2.4
TBD41	Shared Risk Link Group	1096	[I-D.ietf-idr-ls-distribution]/3.3.2.5
TBD42	Opaque link attributes	1097	[I-D.ietf-idr-ls-distribution]/3.3.2.6
TBD43	Link Name attribute	1098	[I-D.ietf-idr-ls-distribution]/3.3.2.7

10. Other Considerations

10.1. Inter-AS Links

The main source of TE information is the IGP, which is not active on inter-AS links. In some cases, the IGP may have information of inter-AS links ([RFC5392], [RFC5316]). In other cases, an implementation SHOULD provide a means to inject inter-AS links into PCEP. The exact mechanism used to provision the inter-AS links is outside the scope of this document.

11. Security Considerations

This document extends PCEP to support TED population including a new TERpt message with new object and TLVs. Procedures and protocol extensions defined in this document do not effect the overall PCEP security model. See [RFC5440], [I-D.ietf-pce-pceps]. Tampering with the TERpt message may have an effect on path computations at PCE. It also provides adversaries an opportunity to eavesdrop and learn sensitive information and plan sophisticated attacks on the network infrastructure. The PCE implementation SHOULD provide mechanisms to prevent strains created by network flaps and amount of TED information. Thus it is suggested that any mechanism used for securing the transmission of other PCEP message be applied here as well. As a general precaution, it is RECOMMENDED that these PCEP extensions only be activated on authenticated and encrypted sessions belonging to the same administrative authority.

12. Manageability Considerations

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

12.1. Control of Function and Policy

In addition to configuring specific PCEP session parameters, as specified in section 8.1 of [RFC5440], a PCE or PCC implementation MUST allow configuring the TED PCEP capability. A PCC SHOULD allow the operator to specify an TED population policy where TERpt are sent to which PCE.

12.2. Information and Data Models

PCEP session configuration and information in the PCEP MIB module SHOULD be extended to include advertised TED capabilities, TED synchronization status and TED etc.

12.3. Liveness Detection and Monitoring

PCEP protocol extensions defined in this document do not require any new mechanisms beyond those already defined in section 8.3 of [RFC5440].

12.4. Verify Correct Operations

Mechanisms defined in section 8.4 of [RFC5440] also apply to PCEP protocol extensions defined in this document. In addition to monitoring parameters defined in [RFC5440], a PCEP implementation with TED SHOULD provide the following parameters:

o Total number of TE Reports

o Number of TE nodes and links

o Number of dropped TERpt messages

12.5. Requirements On Other Protocols

PCEP protocol extensions defined in this document do not put new requirements on other protocols.

12.6. Impact On Network Operations

Mechanisms defined in section 8.6 of [RFC5440] 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 TERpt messages sent by a PCEP speaker and processed by the peer. It SHOULD also allow sending a notification when a rate threshold is reached.

13. IANA Considerations

14. Acknowledgments

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

15. References

15.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009.

15.2. Informative References

[RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", RFC 3630, September 2003.

[RFC4203] Kompella, K. and Y. Rekhter, "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, October 2005.

[RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, August 2006.

[RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, October 2008.

[RFC5307] Kompella, K. and Y. Rekhter, "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 5307, October 2008.

[RFC5316] Chen, M., Zhang, R., and X. Duan, "ISIS Extensions in Support of Inter-Autonomous System (AS) MPLS and GMPLS Traffic Engineering", RFC 5316, December 2008.

[RFC5392] Chen, M., Zhang, R., and X. Duan, "OSPF Extensions in Support of Inter-Autonomous System (AS) MPLS and GMPLS Traffic Engineering", RFC 5392, January 2009.

[RFC6119] Harrison, J., Berger, J., and M. Bartlett, "IPv6 Traffic Engineering in IS-IS", RFC 6119, February 2011.

[RFC6549] Lindem, A., Roy, A., and S. Mirtorabi, "OSPFv2 Multi-Instance Extensions", RFC 6549, March 2012.

[RFC6822] Previdi, S., Ginsberg, L., Shand, M., Roy, A., and D. Ward, "IS-IS Multi-Instance", RFC 6822, December 2012.

[I-D.ietf-pce-stateful-pce]
Crabbe, E., Minei, I., Medved, J., and R. Varga, "PCEP Extensions for Stateful PCE", draft-ietf-pce-stateful-pce-10 (work in progress), October 2014.

[I-D.ietf-pce-pceps]
Lopez, D., Dios, O., Wu, W., and D. Dhody, "Secure Transport for PCEP", draft-ietf-pce-pceps-02 (work in progress), October 2014.

[I-D.ietf-idr-ls-distribution]
Gredler, H., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and TE Information using BGP", draft-ietf-idr-ls-distribution-10 (work in progress), January 2015.

[I-D.lee-pce-transporting-te-data]
Lee, Y. and z. zhenghaomian@huawei.com, "PCE in Support of Transporting Traffic Engineering Data", draft-lee-pce-transporting-te-data-01 (work in progress), October 2014.

Appendix A. Contributor Addresses

Udayasree Palle
Huawei Technologies
Divyashree Techno Park, Whitefield
Bangalore, Karnataka 560037
India

EMail: udayasree.palle@huawei.com

Sergio Belotti
Alcatel-Lucent
Italy

EMail: sergio.belotti@alcatel-lucent.com

Authors' Addresses

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

EMail: dhruv.ietf@gmail.com

Young Lee
Huawei Technologies
5340 Legacy Drive, Building 3
Plano, TX 75023
USA

EMail: leeyoung@huawei.com

Daniele Ceccarelli
Ericsson
Torshamnsgatan,48
Stockholm
Sweden

EMail: daniele.ceccarelli@ericsson.com