

PCE Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 24, 2016

D. Dhody
Y. Lee
Huawei Technologies
D. Ceccarelli
Ericsson
September 21, 2015

**PCEP Extension for Distribution of Link-State and TE Information.
draft-dhodylee-pce-pcep-ls-00**

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 (LS) routing protocol supporting traffic engineering extensions.

This document extends the Path Computation Element Communication Protocol (PCEP) with Link-State and TE Information.

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 <http://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 March 24, 2016.

Copyright Notice

Copyright (c) 2015 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 (<http://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	4
2.	Terminology	4
3.	Applicability	4
4.	Requirements for PCEP extension	5
5.	New Functions to distribute link-state (and TE) via PCEP	5
6.	Overview of Extension to PCEP	6
6.1.	New Messages	6
6.2.	Capability Advertisement	6
6.3.	Initial Link-State (and TE) Synchronization	7
6.3.1.	Optimizations for LS Synchronization	9
6.4.	LS Report	10
7.	Transport	10
8.	PCEP Messages	10
8.1.	LS Report Message	10
8.2.	The PCErr Message	11
9.	Objects and TLV	11
9.1.	Open Object	12
9.1.1.	LS Capability TLV	12
9.2.	LS Object	12
9.2.1.	Routing Universe TLV	14
9.2.2.	Local Node Descriptors TLV	15
9.2.3.	Remote Node Descriptors TLV	16
9.2.4.	Node Descriptors Sub-TLVs	16
9.2.5.	Multi-Topology ID TLV	17
9.2.6.	Link Descriptors TLV	17
9.2.7.	Prefix Descriptors	19
9.2.8.	PCEP-LS Attributes TLV	19
9.2.8.1.	Node Attributes TLV	19
9.2.8.2.	Link Attributes TLV	20
9.2.8.3.	Prefix Attributes TLV	22
10.	Other Considerations	23
10.1.	Inter-AS Links	23
11.	Security Considerations	23
12.	Manageability Considerations	23
12.1.	Control of Function and Policy	23
12.2.	Information and Data Models	23
12.3.	Liveness Detection and Monitoring	23
12.4.	Verify Correct Operations	23
12.5.	Requirements On Other Protocols	24

12.6.	Impact On Network Operations	24
13.	IANA Considerations	24
14.	TLV/Sub-TLV Code Points Summary	24
15.	Acknowledgments	26
16.	References	26
16.1.	Normative References	27
16.2.	Informative References	27
Appendix A.	Contributor Addresses	29
	Authors' Addresses	29

[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 describes a mechanism by which Link State and TE information can be collected from networks and shared with PCE using the PCEP itself. This is achieved using a new PCEP message format. The mechanism is applicable to physical and virtual links as well as further subjected to various policies.

A router maintains one or more databases for storing link-state and TE information about nodes and links in any given area. Link attributes stored in these databases include: local/remote IP addresses, local/ remote interface identifiers, link metric and TE metric, link bandwidth, reservable bandwidth, per CoS class reservation state, preemption and Shared Risk Link Groups (SRLG). The router's PCEP process can retrieve topology from these databases and distribute it to a PCE, either directly or via another PCEP Speaker, using the encoding specified in this document.

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 or BGP-LS running at the PCE to learn link-state and TE information.
- o Where there is IGP or BGP-LS running but with a need for a faster TE and link-state population and convergence at the PCE.
 - * A PCE may receive partial information (say basic TE, link-state) from IGP and other information (optical and impairment) from PCEP.
 - * A PCE may receive full information from both existing mechanism (IGP or BGP) and PCEP.

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

How a PCE manages the link-state (and TE) information is implementation specific and thus out of scope of this document.

4. Requirements for PCEP extension

Following key requirements associated with link-state (and TE) distribution are identified for PCEP:

1. The PCEP speaker supporting this draft **MUST** be a mechanism to advertise the Link-State (and TE) distribution capability.
2. PCC supporting this draft **MUST** have the capability to report the link-state (and TE) information to the PCE. This includes self originated information and remote 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 after.
3. A PCE **MAY** learn link-state (and TE) from PCEP as well as from existing mechanism like IGP/BGP-LS. PCEP extension **MUST** have a mechanism to link the information learned via other means. There **MUST NOT** be any changes to the existing link-state (and TE) population mechanism via IGP/BGP-LS. PCEP extension **SHOULD** keep the properties in a protocol (IGP or BGP-LS) neutral way, such that an implementation maynot need to know about any OSPF or IS-IS or BGP protocol specifics.
4. It **SHOULD** be possible to encode only the changes in link-state (and TE) 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 distribute link-state (and TE) via PCEP

Several new functions are required in PCEP to support distribution of link-state (and TE) information. 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 distribution of link-state (and TE) information defined in this document.

- o Link-State (and TE) synchronization (C-E): after the session between the PCC and a PCE is initialized, the PCE must learn Link-State (and TE) information before it can perform path computations. In case of stateful PCE it is RECOMENDED that this operation be done before LSP state synchronization.
- o Link-State (and TE) Report (C-E): a PCC sends a LS report to a PCE whenever the Link-State (and TE) information changes.

6. Overview of Extension to PCEP

6.1. New Messages

In this document, we define a new PCEP messages called LS Report (LSRpt), a PCEP message sent by a PCC to a PCE to report link-state (and TE) information. Each LS Report in a LSRpt message can contain the node or link properties. An unique PCEP specific LS identifier (LS-ID) is also carried in the message to identify a node or link and that remains constant for the lifetime of a PCEP session. This identifier on its own is sufficient when no IGP or BGP-LS running in the network for PCE to learn link-state (and TE) information. Incase PCE learns some information from PCEP and some from the existing mechanism, the PCC SHOULD include the mapping of IGP or BGP-LS identifier to map the information populated via PCEP with IGP/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 LS (and TE) distribution via PCEP extensions. A PCEP Speaker includes the "LS Capability" TLV, described in [Section 9.1.1](#), in the OPEN Object to advertise its support for PCEP-LS extensions. The presence of the LS Capability TLV in PCC's OPEN Object indicates that the PCC is willing to send LS Reports whenever local link-state (and TE) information changes. The presence of the LS Capability TLV in PCE's OPEN message indicates that the PCE is interested in receiving LS Reports whenever local link-state (and TE) information changes.

The PCEP protocol extensions for LS (and TE) distribution MUST NOT be used if one or both PCEP Speakers have not included the LS 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 LSRpt message from the PCC, it SHOULD generate a PCErr with error-type 19 (Invalid Operation), error-value TBD1 (Attempted LS Report if LS capability was not advertised) and it will terminate the PCEP session.

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

6.3. Initial Link-State (and TE) Synchronization

The purpose of LS Synchronization is to provide a checkpoint-in- time state replica of a PCC's link-state (and TE) data base 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 LS synchronization should be done before LSP state synchronization.

During LS Synchronization, a PCC first takes a snapshot of the state of its database, then sends the snapshot to a PCE in a sequence of LS Reports. Each LS Report sent during LS Synchronization has the SYNC Flag in the LS Object set to 1. The end of synchronization marker is a LSRpt message with the SYNC Flag set to 0 for an LS Object with LS-ID equal to the reserved value 0. If the PCC has no link-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 LS synchronization, it MUST send a PCErr message with error-type TBD2 (LS 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 LS synchronization messages. It MUST respond with a PCErr message with error-type TBD2 (LS Synchronization Error) and error-value 1 (indicating an error in processing the LSRpt) if it encounters a problem with the LS Report it received from the PCC and it MUST terminate the session.

The LS reports can carry local as well as remote link-state (and TE) information depending on the R flag in LS capability TLV.

The successful LS Synchronization sequences is shown in Figure 1.

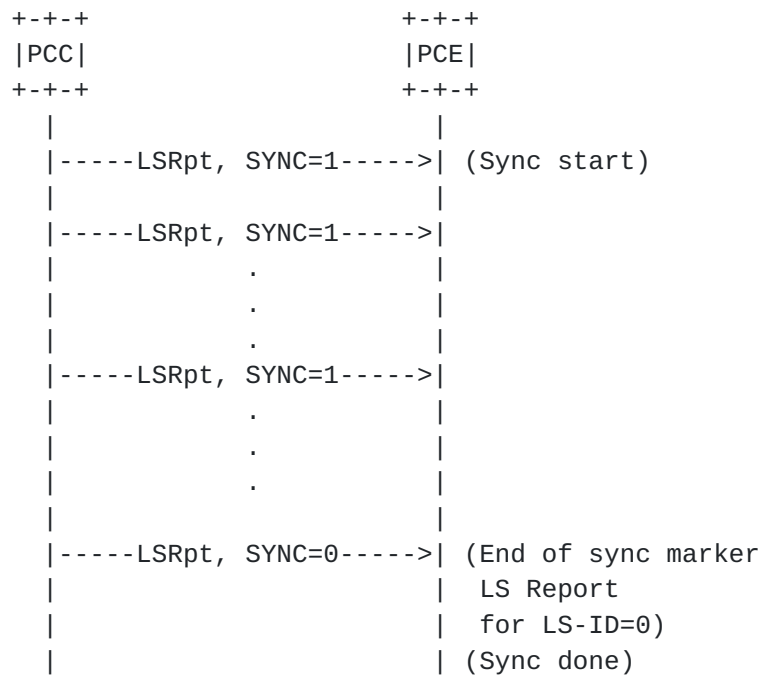


Figure 1: Successful LS synchronization

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

6.4. LS Report

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

More details about LSRpt message are in [Section 8.1](#).

7. Transport

A permanent PCEP session MUST be established between a PCE and PCC supporting link-state (and TE) distribution 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. 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 link-state (and TE) information. A LSRpt message can carry more than one LS Reports. The Message-Type field of the PCEP common header for the LSRpt message is set to [TBD3].

The format of the LSRpt message is as follows:

```
<LSRpt Message> ::= <Common Header>  
                    <ls-report-list>
```

Where:

```
<ls-report-list> ::= <LS>[<ls-report-list>]
```


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

A PCE may choose to implement a limit on the LS information a single PCC can populate. If a LSRpt 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 LSRpt message triggering this condition and SHOULD terminate the session.

8.2. The PCErr Message

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

The format of a PCErr message from [\[RFC5440\]](#) is extended as follows:

The format of the PCErr 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> | <ls-id-list>]
           <error-obj-list>
```

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

```
<ls-id-list> ::= <LS> [<ls-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. LS Capability TLV

The LS-CAPABILITY TLV is an optional TLV for use in the OPEN Object for link-state (and TE) distribution via PCEP capability advertisement. Its format is shown in the following figure:



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 LS information learned via other means like IGP and BGP-LS; if set to 1 by a PCE, the R Flag indicates that the PCE is capable of receiving remote LS information (from the PCC point of view). The R Flag must be advertised by both a PCC and a PCE for LSRpt messages to report remote as well as local LS information on a PCEP session. The TLVs related to IGP/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 LS capability implies support of local link-state (and TE) distribution, as well as the objects, TLVs and procedures defined in this document.

9.2. LS Object

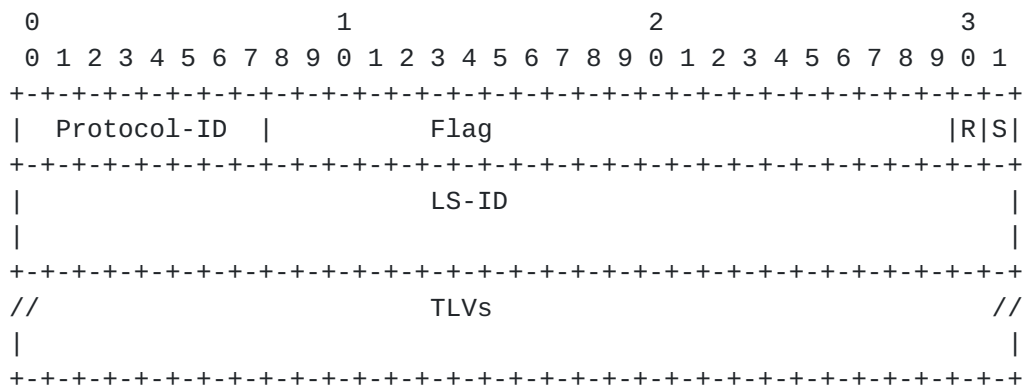
The LS (link-state) object MUST be carried within LSRpt messages and MAY be carried within PCerr messages. The LS object contains a set of fields used to specify the target node or link. It also contains a flag indicating to a PCE that the LS synchronization is in progress. The TLVs used with the LS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

LS Object-Class is [TBD6].

Four Object-Type values are defined for the LS object so far:

- o LS Node: LS Object-Type is 1.
- o LS Link: LS Object-Type is 2.
- o LS IPv4 Topology Prefix: LS Object-Type is 3.
- o LS IPv6 Topology Prefix: LS Object-Type is 4.

The format of all types of LS object is as follows:



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 LSRpt sent from a PCC during LS Synchronization. The S Flag MUST be set to 0 in other LSRpt messages sent from the PCC.

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

LS-ID(64-bit): A PCEP-specific identifier for the node or link or prefix information. A PCC creates an unique LS-ID for each node/link/prefix that is constant for the lifetime of a PCEP session. The PCC will advertise the same LS-ID on all PCEP sessions it maintains at a given times. All subsequent PCEP messages then address the node/link/prefix by the LS-ID. The values of 0 and 0xFFFFFFFFFFFFFFFF are reserved.

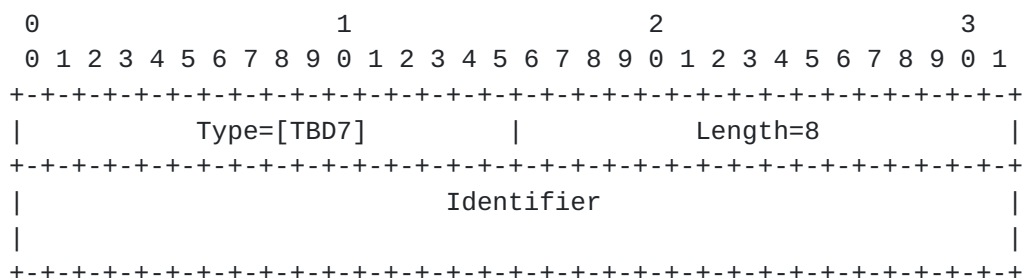
Unassigned bits are considered reserved. They MUST be set to 0 on transmission and MUST be ignored on receipt.

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

9.2.1. Routing Universe TLV

In case of remote link-state (and TE) population when existing IGP/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 [\[RFC6549\]](#) for more information. These instances define independent "routing universes". The 64-Bit 'Identifier' field is used to identify the "routing universe" where the LS object belongs. The LS objects representing IGP objects (nodes or links or prefix) from the same routing universe MUST have the same 'Identifier' value; LS 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:



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	Default Layer 3 Routing topology
1-31	Reserved

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

9.2.2. Local 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 link attribute TLV (see [Section 9.2.8.2](#)).

It is desirable that the Router-ID assignments inside the Node Descriptor are globally unique. Some considerations for globally unique Node/Link/Prefix identifiers are described in [[I-D.ietf-idr-ls-distribution](#)].

The Local Node Descriptors TLV contains Node Descriptors for the node anchoring the local end of the link. This TLV MUST be included in the LS Report when during a given PCEP session a node/link/prefix is first reported to a PCE. A PCC sends to a PCE the first LS Report either during State Synchronization, or when a new node/link/prefix 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/prefix information such that global uniqueness of the node/link/prefix 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								
+-+-+-----+-+-																																							

The value contains one or more Node Descriptor Sub-TLVs defined in [Section 9.2.4](#).

9.2.3. Remote Node Descriptors TLV

The Remote Node Descriptors contains Node Descriptors for the node anchoring the remote end of the link. This TLV MUST be included in the LS Report when during a given PCEP session a link is first reported to a PCE. A PCC sends to a PCE the first LS Report either during State Synchronization, or when a new 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. 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	IGP Router-ID	Variable
TBD44	Multi-Topology-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 OSPF Area ID: It is used to identify the 32 Bit area to which the LS object belongs. Area Identifier allows the different LS objects of the same router to be discriminated.

- o IGP Router ID: opaque value. Usage is described in [\[I-D.ietf-idr-ls-distribution\]](#) for IGP Router ID. In case only local information is transported and PCE learns link-state (and TE) information only from PCEP, it contain the unique local TE IPv4 or IPv6 router ID.
- o Multi-Topology-ID: Usage is described in [\[I-D.ietf-idr-ls-distribution\]](#) for MT-ID.
- o There can be at most one instance of each sub-TLV type present in any Node Descriptor.

9.2.5. Multi-Topology ID TLV

The Multi-Topology ID (MT-ID) TLV carries one or more IS-IS or OSPF Multi-Topology IDs for a link, node or prefix. The semantics of the IS-IS MT-ID are defined in [Section 7.2 of \[RFC5120\]](#). The MT-ID TLV MAY be present in a Link Descriptor, a Prefix Descriptor, or in the attribute of a node (Node Attributes TLV) in LS object.

The format and handling of the MT-ID TLV is as defined in [\[I-D.ietf-idr-ls-distribution\]](#).

In a Link or Prefix Descriptor, only a single MT-ID TLV containing the MT-ID of the topology where the link or the prefix is reachable is allowed. In case one wants to advertise multiple topologies for a given Link Descriptor or Prefix Descriptor, multiple reports need to be generated where each LS object contains an unique MT-ID. In the attribute of a node (Node Attributes TLV) in LS object, one MT-ID TLV containing the array of MT-IDs of all topologies where the node is reachable is allowed.

9.2.6. Link Descriptors TLV

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

The '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
TBD44	Multi-Topology identifier	-	[I-D.ietf-idr- ls-distribution] /3.2.1.5

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](#)].

9.2.7. Prefix Descriptors

The 'Prefix Descriptor' field is a set of Type/Length/Value (TLV) triplets. 'Prefix Descriptor' TLVs uniquely identify an IPv4 or IPv6 Prefix originated by a Node. The following TLVs are valid as Prefix Descriptors in the IPv4/IPv6 Prefix-

TLV Code Point	Description	Length	Value defined in:
TBD44	Multi-Topology Identifier	variable	[I-D.ietf-idr-ls-distribution] /3.2.1.5
TBD45	OSPF Route Type	1	[I-D.ietf-idr-ls-distribution] /3.2.3.1
TBD46	IP Reachability Information	variable	[I-D.ietf-idr-ls-distribution] /3.2.3.2

9.2.8. PCEP-LS Attributes TLV

9.2.8.1. Node Attributes TLV

This is an optional attribute that is used to carry node attributes. The node attribute TLV may be encoded in the LS 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								
+-+-+-----+-+-																																							

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

Sub TLV	Description	Length	Value defined in:
TBD44	Multi-Topology Identifier	variable	[I-D.ietf-idr-ls-distribution]/3.2.1.5
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

9.2.8.2. Link Attributes TLV

Link attribute TLV may be encoded in the LS 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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |                               |
//                               Link Attributes Sub-TLVs (variable)                               //
|                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

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

Sub-TLV	Description	IS-IS TLV /Sub-TLV BGP-LS TLV	Defined in:
TBD25	IPv4 Router-ID of Local Node	134/---	[RFC5305]/4.3
TBD26	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

9.2.8.3. Prefix Attributes TLV

Prefix attribute TLV may be encoded in the LS Prefix Object. Prefixes are learned from the IGP (IS-IS or OSPF) or BGP topology with a set of IGP attributes (such as metric, route tags, etc.). This section describes the different attributes related to the IPv4/IPv6 prefixes. Prefix Attributes TLVs SHOULD be encoded in the LS Prefix 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=[TBD53]              |          Length                  |
+-----+-----+-----+-----+-----+-----+-----+-----+
|                                     |                                     |
|//          Prefix Attributes Sub-TLVs (variable)          //
|                                     |                                     |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

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

Sub-TLV	Description	BGP-LS TLV	Defined in:
TBD47	IGP Flags	1152	[I-D.ietf-idr- ls-distribution] /3.3.3.1
TBD48	Route Tag	1153	[I-D.ietf-idr- ls-distribution] /3.3.3.2
TBD49	Extended Tag	1154	[I-D.ietf-idr- ls-distribution] /3.3.3.3
TBD50	Prefix Metric	1155	[I-D.ietf-idr- ls-distribution] /3.3.3.4
TBD51	OSPF Forwarding Address	1156	[I-D.ietf-idr- ls-distribution] /3.3.3.5
TBD52	Opaque Prefix Attribute	1157	[I-D.ietf-idr- ls-distribution] /3.3.3.6

10. Other Considerations

10.1. Inter-AS Links

The main source of LS (and 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 for LS (and TE) distribution including a new LSRpt 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 LSRpt 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 LS (and TE) 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

12.1. Control of Function and Policy

TBD.

12.2. Information and Data Models

TBD.

12.3. Liveness Detection and Monitoring

TBD.

12.4. Verify Correct Operations

TBD.

[12.5.](#) Requirements On Other Protocols

TBD.

[12.6.](#) Impact On Network Operations

TBD.

[13.](#) IANA Considerations

[14.](#) TLV/Sub-TLV Code Points Summary

This section contains the global table of all TLVs/Sub-TLVs in LS object defined in this document.

Sub-TLV	Description	Ref TLV /Sub-TLV	Value defined in:
TBD7	Routing Universe	--	Sec 9.2.1
TBD8	Local Node Descriptors	256	[I-D.ietf-idr-ls-distribution]/3.2.1.2
TBD9	Remote Node Descriptors	257	[I-D.ietf-idr-ls-distribution]/3.2.1.3
TBD10	Autonomous System	512	[I-D.ietf-idr-ls-distribution]/3.2.1.4
TBD11	BGP-LS Identifier	513	[I-D.ietf-idr-ls-distribution]/3.2.1.4
TBD12	OSPF Area-ID	514	[I-D.ietf-idr-ls-distribution]/3.2.1.4
TBD13	IGP Router-ID	515	[I-D.ietf-idr-ls-distribution]/3.2.1.4
TBD44	Multi-Topology-ID	263	[I-D.ietf-idr-ls-distribution]/3.2.1.5
TBD14	Link Descriptors	--	Sec 9.2.6
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
TBD45	OSPF Route Type	264	[I-D.ietf-idr-ls-distribution]/3.2.3.1
TBD46	IP Reachability Information	265	[I-D.ietf-idr-ls-distribution]/3.2.3.2
TBD20	Node Attributes	--	Sec 9.2.8.1
TBD21	Node Flag Bits	1024	[I-D.ietf-idr-ls-distribution]/3.3.1.1
TBD22	Opaque Node Properties	1025	[I-D.ietf-idr-ls-distribution]/3.3.1.5
TBD23	Node Name	1026	[I-D.ietf-idr-ls-distribution]/3.3.1.3
TBD24	IS-IS Area Identifier	1027	[I-D.ietf-idr-ls-distribution]/3.3.1.2
TBD25	IPv4 Router-ID of Local Node	134/--	[RFC5305]/4.3
TBD26	IPv6 Router-ID of Local Node	140/--	[RFC6119]/4.1
TBD27	Link Attributes	--	Sec 9.2.8.2
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	22/20	[RFC5307]/1.2

	Type		
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
TBD53	Prefix Attributes	--	Sec 9.2.8.3
TBD47	IGP Flags	1152	[I-D.ietf-idr- ls-distribution] /3.3.3.1
TBD48	Route Tag	1153	[I-D.ietf-idr- ls-distribution] /3.3.3.2
TBD49	Extended Tag	1154	[I-D.ietf-idr- ls-distribution] /3.3.3.3
TBD50	Prefix Metric	1155	[I-D.ietf-idr- ls-distribution] /3.3.3.4
TBD51	OSPF Forwarding Address	1156	[I-D.ietf-idr- ls-distribution] /3.3.3.5
TBD52	Opaque Prefix Attribute	1157	[I-D.ietf-idr- ls-distribution] /3.3.3.6
+-----+-----+-----+-----+			

15. Acknowledgments

This document borrows some of the structure and text from the [\[I-D.ietf-idr-ls-distribution\]](#).

16. References

16.1. 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, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", [RFC 5305](#), DOI 10.17487/RFC5305, October 2008, <<http://www.rfc-editor.org/info/rfc5305>>.
- [RFC5307] Kompella, K., Ed. and Y. Rekhter, Ed., "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", [RFC 5307](#), DOI 10.17487/RFC5307, October 2008, <<http://www.rfc-editor.org/info/rfc5307>>.
- [RFC5440] Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", [RFC 5440](#), DOI 10.17487/RFC5440, March 2009, <<http://www.rfc-editor.org/info/rfc5440>>.
- [RFC6119] Harrison, J., Berger, J., and M. Bartlett, "IPv6 Traffic Engineering in IS-IS", [RFC 6119](#), DOI 10.17487/RFC6119, February 2011, <<http://www.rfc-editor.org/info/rfc6119>>.
- [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-11](#) (work in progress), June 2015.

16.2. Informative References

- [RFC3630] Katz, D., Kompella, K., and D. Yeung, "Traffic Engineering (TE) Extensions to OSPF Version 2", [RFC 3630](#), DOI 10.17487/RFC3630, September 2003, <<http://www.rfc-editor.org/info/rfc3630>>.
- [RFC4203] Kompella, K., Ed. and Y. Rekhter, Ed., "OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", [RFC 4203](#), DOI 10.17487/RFC4203, October 2005, <<http://www.rfc-editor.org/info/rfc4203>>.
- [RFC4655] Farrel, A., Vasseur, J., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", [RFC 4655](#), DOI 10.17487/RFC4655, August 2006, <<http://www.rfc-editor.org/info/rfc4655>>.

- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", [RFC 5120](#), DOI 10.17487/RFC5120, February 2008, <<http://www.rfc-editor.org/info/rfc5120>>.
- [RFC5316] Chen, M., Zhang, R., and X. Duan, "ISIS Extensions in Support of Inter-Autonomous System (AS) MPLS and GMPLS Traffic Engineering", [RFC 5316](#), DOI 10.17487/RFC5316, December 2008, <<http://www.rfc-editor.org/info/rfc5316>>.
- [RFC5392] Chen, M., Zhang, R., and X. Duan, "OSPF Extensions in Support of Inter-Autonomous System (AS) MPLS and GMPLS Traffic Engineering", [RFC 5392](#), DOI 10.17487/RFC5392, January 2009, <<http://www.rfc-editor.org/info/rfc5392>>.
- [RFC6549] Lindem, A., Roy, A., and S. Mirtorabi, "OSPFv2 Multi-Instance Extensions", [RFC 6549](#), DOI 10.17487/RFC6549, March 2012, <<http://www.rfc-editor.org/info/rfc6549>>.
- [RFC6822] Previdi, S., Ed., Ginsberg, L., Shand, M., Roy, A., and D. Ward, "IS-IS Multi-Instance", [RFC 6822](#), DOI 10.17487/RFC6822, December 2012, <<http://www.rfc-editor.org/info/rfc6822>>.
- [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-11](#) (work in progress), April 2015.
- [I-D.ietf-pce-pceps]
Lopez, D., Dios, O., Wu, W., and D. Dhody, "Secure Transport for PCEP", [draft-ietf-pce-pceps-04](#) (work in progress), May 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

EEmail: udayasree.palle@huawei.com

Sergio Belotti
Alcatel-Lucent
Italy

EEmail: sergio.belotti@alcatel-lucent.com

Authors' Addresses

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

EEmail: dhruv.ietf@gmail.com

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

EEmail: leeyoung@huawei.com

Daniele Ceccarelli
Ericsson
Torshamnsgatan, 48
Stockholm
Sweden

EEmail: daniele.ceccarelli@ericsson.com

