Workgroup: Path Computation Element
Internet-Draft:
draft-chen-pce-pce-initiated-ip-tunnel-04
Published: 11 January 2024
Intended Status: Standards Track
Expires: 14 July 2024
Authors: X. Chen
Huawei Technologies
H. Shi, Ed.
Huawei Technologies
Z. Li
Huawei Technologies

# PCE-initiated IP Tunnel

## Abstract

This document specifies a set of extensions to PCEP to support PCE-initiated IP Tunnel to satisfy the requirement which is introduced in [I-D.li-spring-tunnel-segment]. The extensions include the setup, maintenance and teardown of PCE-initiated IP Tunnels, without the need for local configuration on the PCC.

## Discussion Venues

This note is to be removed before publishing as an RFC.

Discussion of this document takes place on the Path Computation Element Working Group mailing list (pce@ietf.org), which is archived at https://mailarchive.ietf.org/arch/browse/pce/.

Source for this draft and an issue tracker can be found at https://github.com/VMatrix1900/draft-pce-initiated-ip-tunnel.

## Status of This Memo

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

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 14 July 2024.

## Copyright Notice



## Table of Contents

## 1. Introduction

[I-D.li-spring-tunnel-segment] introduces a new type of segment, Tunnel Segment, for the segment routing. Tunnel segment can be used to reduce SID stack depth of SR path, span the non-SR domain or provide differentiated services. The tunnel segment can be allocated for MPLS RSVP-TE tunnel, SR-TE tunnel or IP tunnel. IP tunnel is also useful in SD-WAN scenario.

[I-D.li-spring-tunnel-segment] introduces two ways to set up the tunnel which is used as tunnel segment: one is to configure tunnel on the device, the other is PCE-initiated tunnel.

[RFC8231], [RFC8281] and [RFC8664] has defined how to set up the PCE initiated RSVP-TE LSP and SR-TE LSP. This document specifies a set of extensions to PCEP to support PCE-initiated IP Tunnel. The extensions include the setup, maintenance and teardown of PCE-initiated IP Tunnels, without the need for local configuration on the PCC.

## 2. Terminology

*SR: Segment Routing

*SR-TE: Segment Routing Traffic Engineering

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

The following terms are defined in [RFC8281]:

*PCE-initiated LSP: LSP that is instantiated as a result of a request from the PCE.

The following terms are defined in this document:

*IP Tunnel: Tunnel that uses IP encapsulation.

*PCE-initiated IP Tunnel: IP Tunnel that is instantiated as a result of a request from the PCE.

The message formats in this document are specified using Routing Backus-Naur Format (RBNF) encoding as specified in [RFC5511].

### 2.1. Requirements Language

The key words "**MUST**", "**MUST NOT**", "**REQUIRED**", "**SHALL**", "**SHALL NOT**", "**SHOULD**", "**SHOULD NOT**", "**RECOMMENDED**", "**NOT RECOMMENDED**", "**MAY**", and "**OPTIONAL**" in this document are to be interpreted as described in

BCP14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

## 3. Procedures for PCE-initiated IP Tunnel

### 3.1. Overview of Procedures

A PCC or PCE indicates its ability to support PCE Initiated dynamic tunnel during the PCEP Initialization Phase via "PCE Initiated Tunnel Capability" TLV (see details in Section 5.1).

In this document, the procedure is only about PCE Initiated dynamic IP Tunnel. The decision of when to instantiate or delete a PCE-initiated IP Tunnel is out of the scope of this document.

This section introduces the procedure to support PCE provisioned IP Tunnel as follows:

Firstly both the PCC and the PCE negotiate the PCE Initiated Tunnel Capability for tunnel types during the PCE session initiation phase. On the PCEP session with PCE Initiated Tunnel Capability PCE communicates with PCC to set up, maintain and tear down PCE-initiated IP Tunnels.

The procedure about tunnel state synchronization, PCC local policy and timeout process, the session failure process, etc. will be specified in the future version.

### 3.2. Capability Advertisement

During PCEP session establishment, both the PCC and the PCE must announce their support of PCEP extensions defined in this document. A PCEP Speaker (PCE or PCC) includes the "PCE Initiated Tunnel Capability" TLV, described in Section 5.1, in the OPEN Object to advertise its support for PCEP extensions for PCE Initiated IP Tunnel Capability.

The PCE Initiated Tunnel Capability TLV includes the tunnel types that are supported by PCEP Speaker. Each tunnel type is indicated by one bit.

The presence of the PCE Initiated Tunnel Capability TLV in PCE's OPEN message indicates that the PCE can support the instantiation of PCE- initiated Tunnels and the types of the tunnels which PCE can initiate.

The presence of such Capability TLV in PCC's OPEN Object indicates that the PCC can support to instantiate the tunnel according to the PCE's indication and the types of the tunnels which PCC can setup automatically according to the PCE's request.

If PCE has such capability TLV and PCC has no such capability TLV PCE **MUST NOT** send the PCE messages for procedure of PCE initiated IP Tunnel. And if PCC receives such messages it should send PCErr message to PCE.

If both PCE and PCC have such capability TLV they only negotiate the types of the tunnels both PCE and PCC can support. PCE **MUST** only initiate the specific tunnel which both PCE and PCC can support. Otherwise PCC **MUST** send the PCErr message.

### 3.3. Tunnel Operations

#### 3.3.1. PCE IP Tunnel Instantiation

To instantiate a tunnel, the PCE sends a Path Computation Tunnel Initiate (PCTunnelInitiate) message to the PCC. The PCTunnelInitiate message **MUST** include the SRP object (see details in Section 5.2) and TUNNEL object (see details in Section 5.3) . The TUNNEL object **MUST** have a PTUNNEL-ID of 0 and **MUST** include the Tunnel Identifier TLV with the TUNNEL-ID 0 and the Tunnel Name TLV. The TUNNEL object **MAY** have the Tunnel Parameter TLV.

The PCC creates the different type of tunnel using the end point address carried in Tunnel Identifier TLV and sends the Path Computation Tunnel State Report (PCTunneRpt) message to PCE. The PCTunneRpt message **MUST** include the SRP object and TUNNEL object. PCC assigns a unique PTUNNEL-ID carried via TUNNEL object and a unique TUNNEL-ID carried via Tunnel Identifier TLV(see details in Section 5.3) in TUNNEL object for the tunnel. PCC indicates the operational state in the TUNNEL object.

The PCTunneRpt message **MUST** include the SRP object, with the SRP-ID-NUMBER used in the SRP object of the PCTunnelInitiate message.

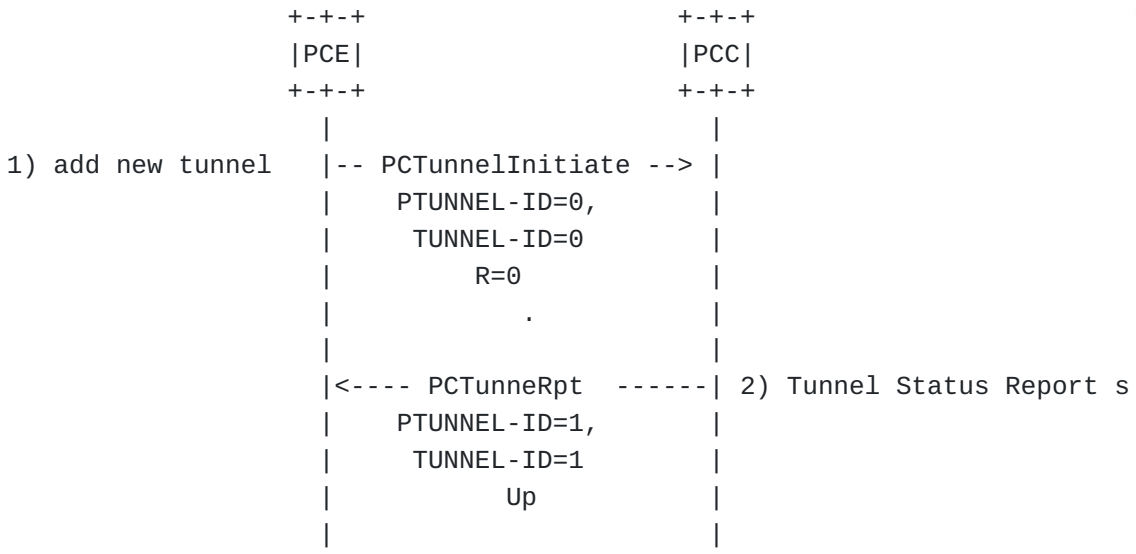

### 3.3.2. PCE IP Tunnel Update

To update the parameters used to create a tunnel, the PCE sends a Path Computation Tunnel Update (PCTunnelUpd) message to the PCC. The PCTunnelUpd message **MUST** include the SRP object and TUNNEL object. The TUNNEL object **MUST** have specific PTUNNEL-ID and **MUST** have specific Tunnel Identifier TLV. The TUNNEL object **MUST** carry any of the Tunnel Parameter TLV and Tunnel Attribute TLV.

The PCC updates the encapsulation parameters and/or attributes of the tunnel and PCC sends the PCTunneRpt message to PCE to report updated state.

The PCTunneRpt message **MUST** include the SRP object, with the SRP-ID-NUMBER used in the SRP object of the PCTunnelUpd message.

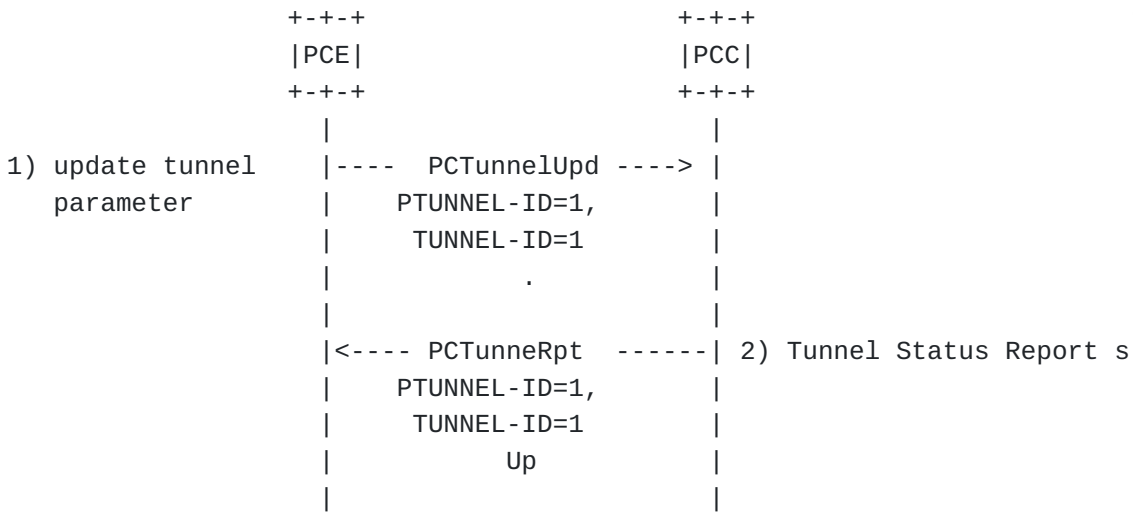


### 3.3.3. PCE IP Tunnel Deletion

To delete a tunnel, the PCE sends a Path Computation Tunnel Initiate (PCTunnelInitiate) message to the PCC. The PCTunnelInitiate message **MUST** include the SRP object and TUNNEL object and the 'R' flag in SRP object **SHOULD** be set. The TUNNEL object **MUST** have specific PTUNNEL- ID and **MUST** have specific Tunnel Identifier TLV.

The PCC delete the tunnel specified by PTUNNEL-ID and PCC sends the PCTunneRpt message to PCE to report updated state.

The PCTunneRpt message **MUST** include the SRP object, with the SRP-ID-NUMBER used in the SRP object of the PCTunnelInitiate message.

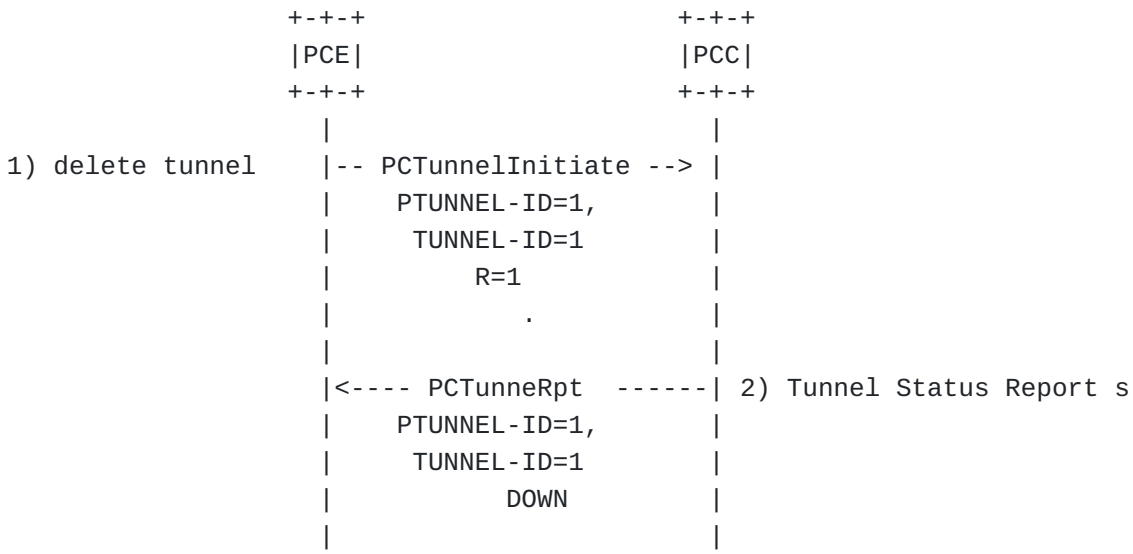


## 4. PCEP Messages

To initiate a tunnel, a PCE sends a PCTunnelInitiate message to a PCC.

To report the state of a tunnel, a PCC sends a PCTunnelRpt message to a PCE.

To modify the parameters of a tunnel, a PCE sends a PCTunnelUpd message to a PCC.

The message format, objects and TLVs are discussed separately below for the creation and the deletion cases.

### 4.1. PCTunnelInitiate Message

A Path Computation Tunnel Initiate message which is also referred to as PCTunnelInitiate message is a PCEP message sent by a PCE to a PCC to trigger tunnel instantiation or deletion.

The Message-Type field of the PCEP common header for the PCTunnelInitiate message is to be assigned by IANA. The PCTunnelInitiate message **MUST** include the SRP and the TUNNEL objects. If the SRP object is missing, the PCC **MUST** send a PCErr with error- type 6 (Mandatory Object missing) and error-value=10 (SRP Object missing) (per [RFC8231]). If the TUNNEL object is missing, the PCC **MUST** send a PCErr with error-type 6 (Mandatory Object missing) and error-value which means TUNNEL Object missing.

Tunnel instantiation is done by sending an Tunnel Initiate Message with an TUNNEL object with the reserved PTUNNEL-ID 0. Tunnel deletion is done by sending an Tunnel Initiate Message with an TUNNEL object carrying the PTUNNEL-ID of the tunnel to be removed and an SRP object with the R flag set.

The format of a PCTunnelInitiate message for tunnel instantiation is as follows:

```
<PCTunnelInitiate Message> ::= <Common Header>
                               <PCE-initiated-tunnel-list>
Where:
<PCE-initiated-tunnel-list> ::= <PCE-initiated-tunnel-request>
                                [<PCE-initiated-tunnel-request>]
<PCE-initiated-tunnel-request> ::= (<PCE-initiated-tunnel-instantiati
                                    |<PCE-initiated-tunnel-deletion>)
<PCE-initiated-tunnel-instantiation> ::= <SRP>
                                         <TUNNEL>
<PCE-initiated-tunnel-deletion> ::= <SRP>
                                    <TUNNEL>
```

The SRP object defined in [RFC8231] can be used in this document to correlate tunnel initiate requests and update requests sent by the PCE with the error reports and tunnel state reports sent by the PCC. Every request from the PCE sends a new SRP- ID-NUMBER. This number is unique per PCEP session and is incremented each time an operation (initiation, update, etc) is requested from the PCE. The value of the SRP-ID-NUMBER **MUST** be echoed back by the PCC in PCErr and PCTunnelRpt messages to allow for correlation between requests made by the PCE and errors or state reports generated by the PCC. Procedure of PCE-initiated IP Tunnel share the same number space of the SRP-ID-NUMBER with procedure of stateful PCE.

The <TUNNEL> object is an new object introduced in this document. <TUNNEL> object in PCTunnelInitiate message **MUST** include Tunnel Identifier TLV and Tunnel Name TLV. Tunnel Parameter TLV is optionally included.

The Tunnel Initiate message for tunnel instantiation has the TUNNEL object with the TUNNEL-ID in Tunnel Identifier TLV 0. The Tunnel Initiate message for tunnel deletion has the TUNNEL object carrying the TUNNEL-ID of the TUNNEL to be removed.

### 4.2. PCTunnelUpd Message

A Path Computation Tunnel Update Request message (also referred to as PCTunnelUpd message) is a PCEP message sent by a PCE to a PCC to update the encapsulation parameters and/or attributes of a tunnel. A PCTunnelUpd message can carry more than one Tunnel Update Request.

The Message-Type field of the PCEP common header for the PCUpd message is to be assigned by IANA.

The PCTunnelUpd message **MUST** include the SRP and the TUNNEL objects. If the SRP object is missing, the PCC **MUST** send a PCErr with error-type 6 (Mandatory Object missing) and error-value=10 (SRP Object

missing) (per [RFC8231]). If the TUNNEL object is missing, the PCC **MUST** send a PCErr with error-type 6 (Mandatory Object missing) and error-value which means TUNNEL Object missing.

The format of a PCTunnelUpd message for tunnel parameter update is as follows:

```
   <PCETunnelUpd Message> ::= <Common Header>
                              <tunnel-update-request-list>
Where:
   <tunnel-update-request-list> ::= <tunnel-update-request>
                                    [<tunnel-update-request-list>]
   <tunnel-update-request> ::= <SRP>
                               <TUNNEL>
```

<TUNNEL> object in PCTunnelUpd message **MUST** include Tunnel Identifier TLV and any of Tunnel Parameter TLV and Tunnel Attribute TLV. Tunnel Name TLV is not included.

### 4.3. PCTunnelRpt Message

A Path Computation Tunnel State Report message which is also referred to as PCTunnelRpt message is a PCEP message sent by a PCC to a PCE to report the current state of a tunnel. A PCTunnelRpt message can carry more than one Tunnel State Reports. A PCC sends an Tunnel State Report in response to a Tunnel Initiate Request for creation or a Tunnel Update Request from a PCE.

The Message-Type field of the PCEP common header for the PCTunnelRpt message is to be assigned by IANA. The PCTunnelRpt message **MUST** include the SRP and the TUNNEL objects. If the SRP object is missing, the PCE **MUST** send a PCErr with error-type 6 (Mandatory Object missing) and error-value=10 (SRP Object missing) (per [RFC8231]). If the TUNNEL object is missing, the PCE **MUST** send a PCErr with error-type 6 (Mandatory Object missing) and error-value which means TUNNEL Object missing.

The format of a PCTunnelRpt message for tunnel instantiation is as follows:

```
   <PCTunnelRpt Message> ::= <Common Header>
                             <tunnel-state-report-list>
Where:
   <tunnel-state-report-list> ::= <tunnel-state-report>
                                  [<tunnel-state-report-list>]
   <tunnel-state-report> ::= <SRP>
                             <TUNNEL>
```

<TUNNEL> object in PCTunnelRpt message **MUST** include Tunnel Identifier TLV.

Tunnel Parameter TLV and Tunnel Attribute TLV is optionally included in PCTunnelRpt message. In the first PCTunnelRpt message in response to the PCTunnelInitiate message Tunnel Name TLV **MUST** be included. And in the subsequent PCTunnelRpt message Tunnel Name TLV is optionally included.

# 5. PCEP Objects

## 5.1. OPEN Object

### 5.1.1. PCE Initiated Tunnel Capability TLV

The PCE-INITIATE-TUNNEL-CAPABILITY TLV is an optional TLV associated with the OPEN Object [RFC5440] to exchange PCE-initiated tunnel capability of PCEP speakers.

Its format is shown in the following figure:

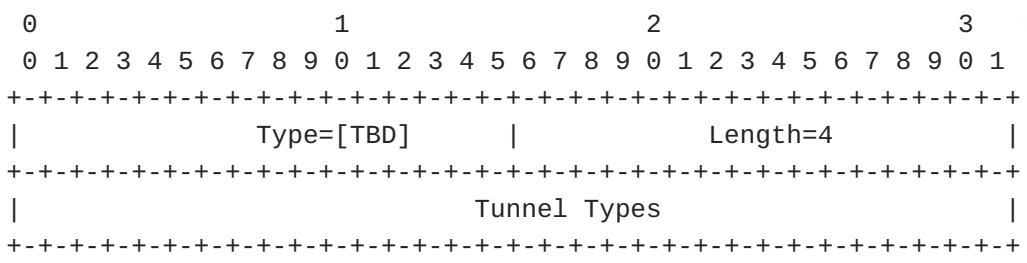


The type of the TLV is to be assigned by IANA and it has a fixed length of 4 octets.

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

Each bit indicates one kind of tunnel. Each bit from right to left successively represents the value of tunnel type which is 0 to 31. The value of tunnel types refer to the registry for "BGP Tunnel Encapsulation Attribute Tunnel Types" [RFC9012] assigned by IANA. This document only use the IP tunnel type.

The assignments used by this document are as follows:

| Tunnel Type | Value |
|---|---|
| Reserved | 0 |
| GRE | 2 |
| VXLAN | 8 |
| NVGRE | 9 |
| MPLS in GRE | 11 |
| MPLS in UDP | 13 |

Table 1: Tunnel Type

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

### 5.2. SRP Object

<SRP> object is defined in [RFC8231]. In this document <SRP> is used to correlate PCTunnelInitiate and PCTunnelRpt or PCErr message.

'R' Flag in <SRP> object is defined in [RFC8281]. When PCE requests PCC to create the IP tunnel 'R' Flag in <SRP> is set to 0. When PCE requests PCC to delete the IP tunnel 'R' Flag in <SRP> is set to 1.

Other flags must be set to 0 and if PCC receive the PCTunnelInitiate message with other reserved flags in <SRP> set to 1 PCC will send the PCErr message.

In procedure of PCE-initiated IP tunnel <SRP> object carries no optional TLVs.

### 5.3. TUNNEL Object

The TUNNEL object **MUST** be present within PCTunnelInitiate, PCTunnelRpt and PCTunnelUpd messages. The TUNNEL object contains a set of fields used to specify the target tunnel, the flags to indicate the state of the tunnel or operation to be performed on the tunnel and TLVs.

TUNNEL Object-Class is to be assigned by IANA.

TUNNEL Object-Type is 1.

The format of the TUNNEL object body is shown in following Figure:

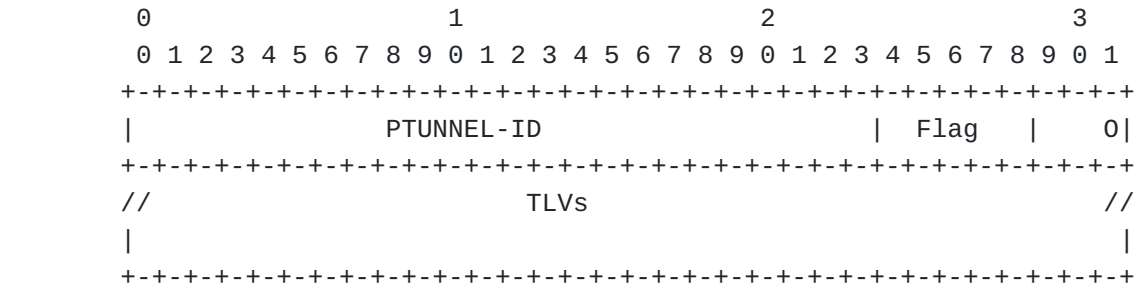


PTUNNEL-ID (24 bits): A PCEP-specific identifier for the tunnel. A PCC creates a unique PTUNNEL-ID for each tunnel that is constant for the lifetime of a PCEP session. The PCC will advertise the same PTUNNEL-ID on all PCEP sessions. The mapping of the Tunnel Name to PTUNNEL-ID is communicated to the PCE by sending a PCTunnelRpt message containing the TUNNEL-NAME TLV. All subsequent PCEP messages then address the tunnel by the PTUNNEL-ID. The values of 0 and 0xFFFFFF are reserved.

Flags (8 bits):

O(Operational - 3 bits): On PCTunnelRpt messages, the O Field represents the operational status of the tunnel.

The following values are defined:

0 - DOWN: The tunnel can't carry the traffic.

1 - UP: The tunnel can carry the traffic.

2-7 - Reserved: these values are reserved for future use.

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 TUNNEL Object are described in the following sections.

### 5.3.1. Tunnel Identifier TLV

The Tunnel Identifier TLV **MUST** be included in the TUNNEL object in PCTunnelInitiate, PCTunnelRpt and PCTunnelUpd messages for PCE-initiated IP Tunnels. If the TLV is missing, the PCE will generate an error with error-type 6 (mandatory object missing) and error-value which means Tunnel Identifier TLV missing and close the session. There are two Tunnel Identifier TLVs, one for IPv4 and one for IPv6.

The format of the IPV4-TUNNEL-Identifier TLV is shown in the following figure:

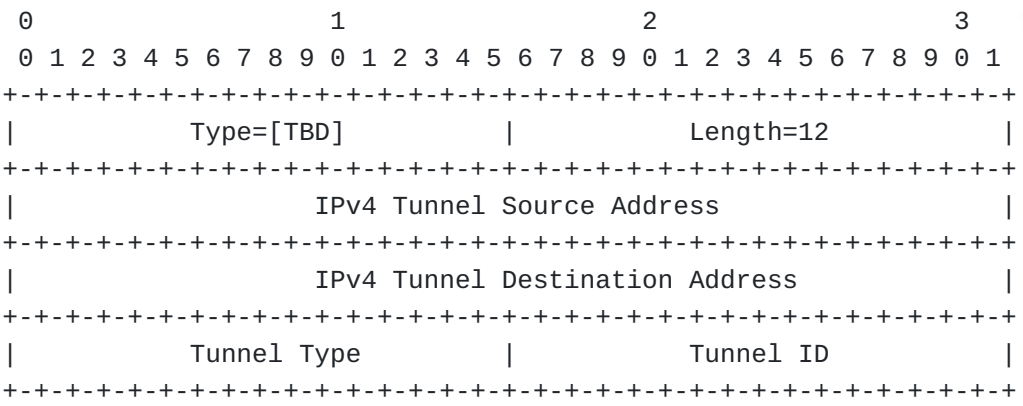


The type of the TLV is to be assigned by IANA and it has a fixed length of 12 octets. The value contains the following fields:

IPv4 Tunnel Source Address: contains the source IPv4 address of the ingress node of the tunnel.

IPv4 Tunnel Destination Address: contains the destination IPv4 address of the egress node of the tunnel.

Tunnel Type: contains the type of tunnel (see Table 1).

Tunnel ID: Tunnel ID remains constant over the life time of a tunnel. A PCC creates a unique Tunnel ID for each tunnel. Each tunnel type has individual identifier space. The Tunnel ID is allocated on id space of the tunnel type and is unique in the same id space.

The PCC will advertise the same Tunnel ID on all PCEP sessions. The mapping of the Tunnel Name to Tunnel ID is communicated to the PCE by sending a PCTunnelRpt message containing the TUNNEL-NAME TLV. The values of 0 and 0xFFFF are reserved.

The format of the IPV6-TUNNEL-Identifier TLV is shown in following figure:

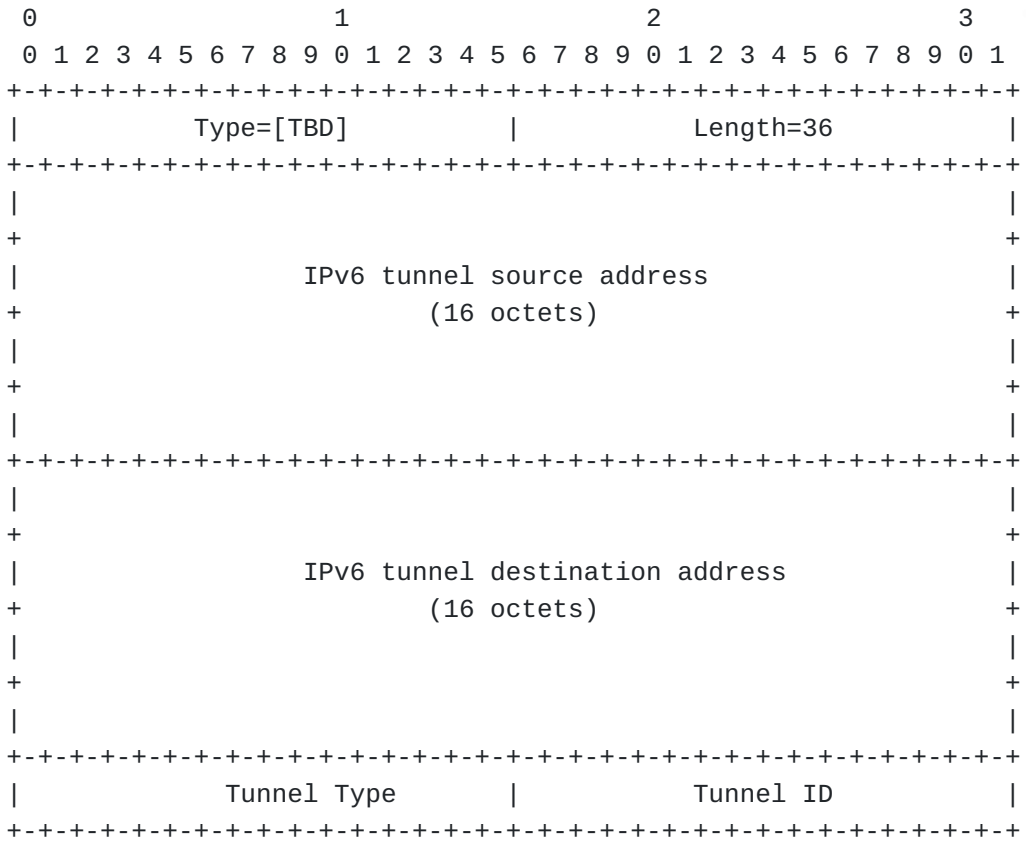


The type of the TLV is to be assigned by IANA and it has a fixed length of 36 octets. The value contains the following fields:

IPv6 Tunnel Source Address: contains the source IPv6 address of the ingress node of the tunnel.

IPv6 Tunnel Destination Address: contains the destination IPv6 address of the egress node of the tunnel.

Tunnel Type: contains the type of tunnel (see Table 1).

Tunnel ID: Tunnel ID remains constant over the life time of a tunnel. A PCC creates a unique Tunnel ID for each TUNNEL. Each tunnel type has individual identifier space. The tunnel ID is allocated on id space of the tunnel type and is unique in the same id space.

The PCC will advertise the same Tunnel ID on all PCEP sessions. The mapping of the Tunnel Name to Tunnel ID is communicated to the PCE by sending a PCTunnelRpt message containing the TUNNEL-NAME TLV. The values of 0 and 0xFFFF are reserved.

### 5.3.2. Tunnel Name TLV

The Tunnel Name TLV **MUST** be included in the TUNNEL object in PCTunnelInitiate messages for PCE-initiated IP Tunnels. If the TLV is missing, the PCE will generate an error with error-type 6 (mandatory object missing) and error-value which means Tunnel Name TLV missing and close the session.

Each tunnel **MUST** have a tunnel name that is unique in the PCC. This tunnel name **MUST** remain constant throughout a tunnel's lifetime.

The TUNNEL-NAME TLV **MUST** be included in the PCTunnelRpt message when a tunnel is first reported to a PCE in response to the PCTunnelInitiate message to create the tunnel. The tunnel name **MAY** be included in subsequent PCTunnelRpt messages for the tunnel.

The format of the TUNNEL-NAME TLV is shown in the following figure:

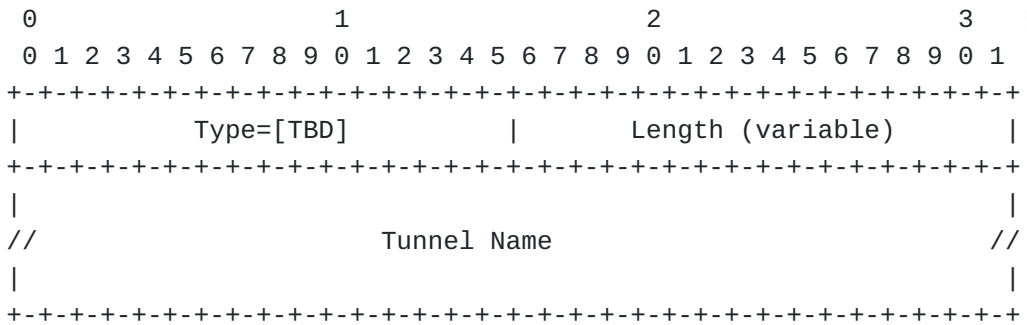


The type of the TLV is to be assigned by IANA and it has a variable length, which **MUST** be greater than 0.

### 5.3.3. Tunnel Parameter TLV

The Tunnel Parameter TLV and/or Tunnel Attribute TLV(see details in following section) **MUST** be included in the TUNNEL object in PCTunnelUpd messages for PCE-initiated IP Tunnels. If both of the TLVs are missing, the PCE will generate an error with error-type 6 (mandatory object missing) and error-value which means Tunnel Parameter TLV and Tunnel Attribute TLV missing and close the session.

The Tunnel Parameter TLV **MAY** be included in the TUNNEL object in PCTunnelInitiate and PCTunnelRpt messages for PCE-initiated IP Tunnels.

Tunnel Parameter TLV specifies information needed to construct the encapsulation header when sending packets through that tunnel.

The tunnel with different type has different encapsulation mode and each tunnel with same type **MAY** has different encapsulation parameters. When PCE initiate setup of the tunnel PCE can specify the encapsulation parameter of the tunnel and PCC will setup the tunnel and encapsulate the packet according to the parameters.

After the tunnel has been triggered to instantiate PCE can send PCTunnelUpd message to modify the encapsulation parameter.

The format of the TUNNEL-PARAMETER TLV is shown in following figure:

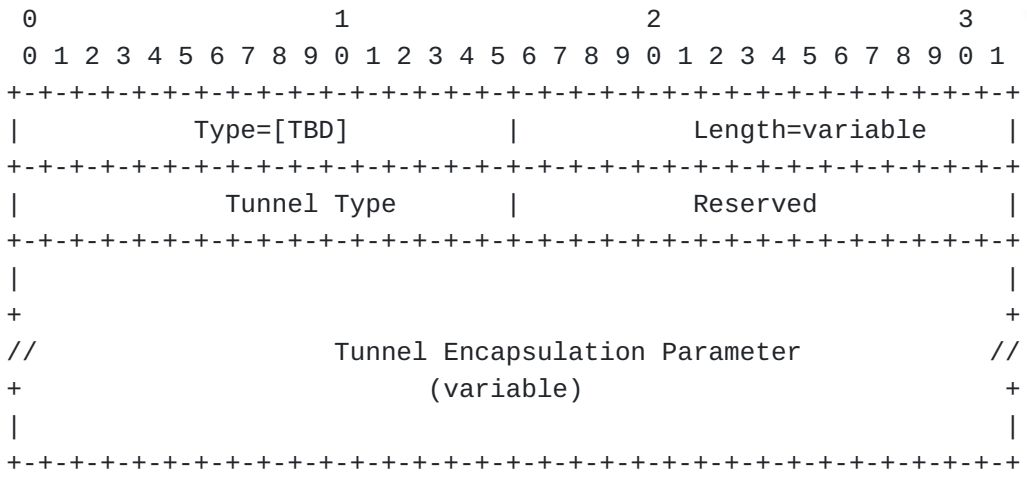


The type of the TLV is to be assigned by IANA and it has a variable length, which **MUST** be greater than 0. The minimum value of length is 4 without any parameter. The value contains the following fields:

Tunnel Type: contains the type of tunnel (see Table 1).

#### 5.3.3.1. GRE

When the tunnel type of the TLV is GRE, the following is the structure of the value field of Tunnel Encapsulation Parameter:

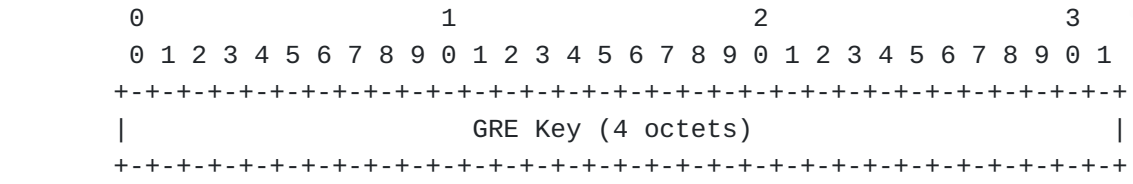


*GRE Key: 4-octet field [RFC2890]. The actual method by which the key is obtained by PCE is beyond the scope of this document. The key is inserted into the GRE encapsulation header of the payload packets sent by ingress router to the egress router. It is intended to be used for identifying extra context information about the received payload.

Note that the key is optional. Unless a key value is being used, the GRE encapsulation **MUST NOT** be present. If GRE tunnel didn't use the GRE key the PCTunnelInitiate message needn't carry the TUNNEL-PARAMETER TLV. If GRE tunnel firstly use the GRE key the PCTunnelInitiate message need carry the TUNNEL-PARAMETER TLV. Then if the GRE tunnel quit using the GRE key the PCTunnelUpd message can carry the TUNNEL-PARAMETER TLV without GRE key to delete the parameter previously set.

MPLS in GRE has the same encapsulation with GRE.

#### 5.3.3.2. VXLAN

When the tunnel type of the TLV is VXLAN, the following is the structure of the value field of Tunnel Encapsulation Parameter:

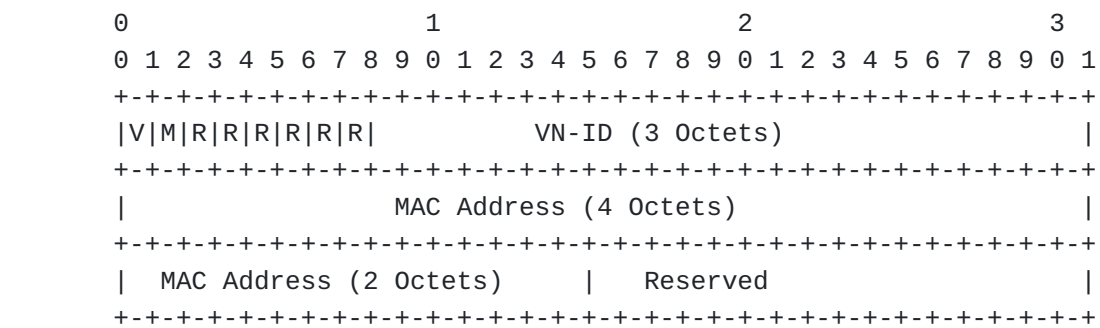


The definition of the fields refer to [RFC9012].

#### 5.3.3.3. NVGRE

When the tunnel type of the TLV is NVGRE, the following is the structure of the value field of Tunnel Encapsulation Parameter:

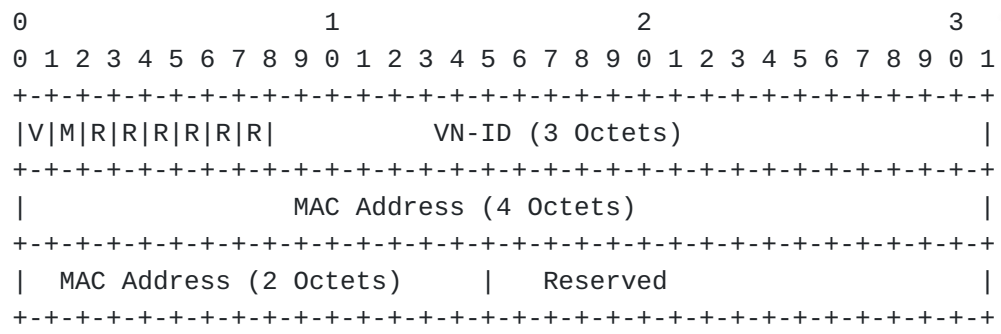


The definition of the fields refer to [RFC9012].

#### 5.3.3.4. MPLS-in-UDP

When the tunnel type of the TLV is MPLS-in-UDP, the following is the structure of the value field of Tunnel Encapsulation Parameter:

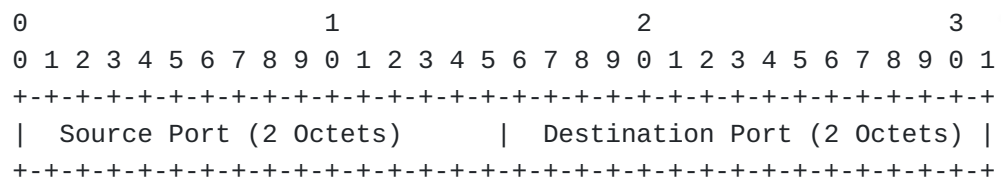


Source Port: UDP source port.

Destination Port: UDP destination port.

### 5.3.4. Tunnel Attribute TLV

The Tunnel Attribute TLV **MAY** be included in the TUNNEL object in PCTunnelInitiate, PCTunnelUpd, PCTunnelRpt messages for PCE-initiated IP Tunnels.

Tunnel Attribute TLV specifies some of the information of the tunnel such as metric or TE metric which are carried in sub-TLVs.

The format of the TUNNEL-ATTRIBUTE TLV is shown in following figure:

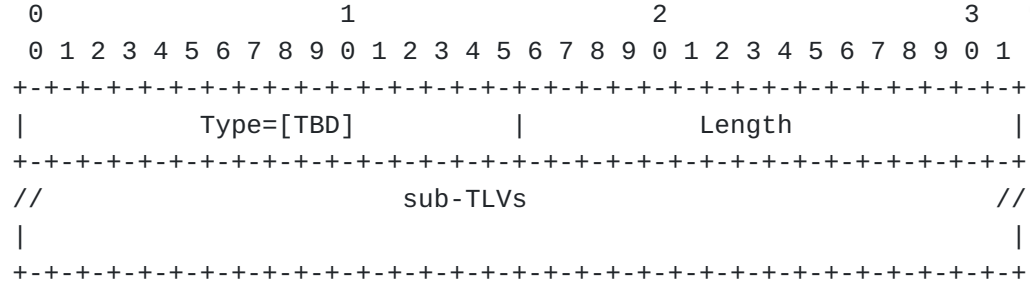


The type of the TLV is to be assigned by IANA and it has a variable length. The minimum value of length is 0 without any parameter. The value contains the following fields:

sub-TLVs: Each sub-TLV has the Type (two octets), Length (two octets), Value. The length is the length of the value of the sub-TLV. Unknown sub-TLVs are to be ignored and skipped upon receipt.

This document defines the following sub-TLVs.

#### 5.3.4.1. Metric Sub-TLV

The following is the structure of the sub-TLV of metric:

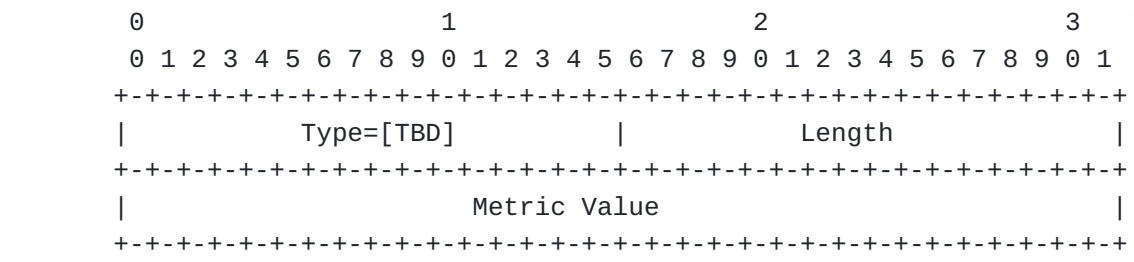


The type of the sub-TLV is to be assigned by IANA and it has a fixed length of 4 octets.

The value comprises a single field - Metric Value (32 bits): value of metric.

#### 5.3.4.2. TE Metric Sub-TLV

The following is the structure of the sub-TLV of traffic engineering metric:

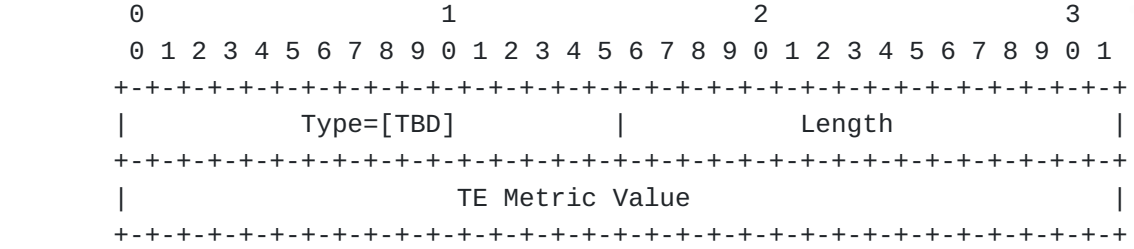


The type of the sub-TLV is to be assigned by IANA and it has a fixed length of 4 octets.

The value comprises a single field - TE Metric Value (32 bits): value of traffic engineering metric.

## 6. Security Considerations

TBD.

## 7. IANA Considerations

TBD.

## 8. References

### 8.1. Normative References


**[I-D.li-spring-tunnel-segment]** Li, Z., Li, C., and N. Wu, "Tunnel Segment in Segment Routing", Work in Progress, Internet-Draft, draft-li-spring-tunnel-segment-02, 14 April 2021, <https://datatracker.ietf.org/doc/html/draft-li-spring-tunnel-segment-02>.

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

**[RFC2890]** Dommety, G., "Key and Sequence Number Extensions to GRE", RFC 2890, DOI 10.17487/RFC2890, September 2000, <https://www.rfc-editor.org/info/rfc2890>.

**[RFC5440]** Vasseur, JP., Ed. and JL. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, DOI 10.17487/RFC5440, March 2009, <https://www.rfc-editor.org/info/rfc5440>.

**[RFC5511]** Farrel, A., "Routing Backus-Naur Form (RBNF): A Syntax Used to Form Encoding Rules in Various Routing Protocol Specifications", RFC 5511, DOI 10.17487/RFC5511, April 2009, <https://www.rfc-editor.org/info/rfc5511>.

**[RFC8174]** Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

**[RFC9012]** Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", RFC 9012, DOI 10.17487/RFC9012, April 2021, <https://www.rfc-editor.org/info/rfc9012>.


### 8.2. Informative References


**[RFC8231]** Crabbe, E., Minei, I., Medved, J., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for Stateful PCE", RFC 8231, DOI 10.17487/RFC8231, September 2017, <https://www.rfc-editor.org/info/rfc8231>.

**[RFC8281]** Crabbe, E., Minei, I., Sivabalan, S., and R. Varga, "Path Computation Element Communication Protocol (PCEP) Extensions for PCE-Initiated LSP Setup in a Stateful PCE

Model", RFC 8281, DOI 10.17487/RFC8281, December 2017, <https://www.rfc-editor.org/info/rfc8281>.

[RFC8664] Sivabalan, S., Filsfils, C., Tantsura, J., Henderickx, W., and J. Hardwick, "Path Computation Element Communication Protocol (PCEP) Extensions for Segment Routing", RFC 8664, DOI 10.17487/RFC8664, December 2019, <https://www.rfc-editor.org/info/rfc8664>.

## Authors' Addresses

Xia Chen
Huawei Technologies
China

Email: jescia.chenxia@huawei.com

Hang Shi (editor)
Huawei Technologies
China

Email: shihang9@huawei.com

Zhenbin Li
Huawei Technologies
China

Email: lizhenbin@huawei.com