

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 22 June 2022

H. Chen
M. McBride
Futurewei
R. Chen
ZTE Corporation
G. Mishra
Verizon Inc.
A. Wang
China Telecom
Y. Liu
China Mobile
Y. Fan
Casa Systems
B. Khasanov
Yandex LLC
L. Liu
Fujitsu
X. Liu
Volta Networks
19 December 2021

BGP for BIER-TE Path
draft-chen-idr-bier-te-path-03

Abstract

This document describes extensions to Border Gateway Protocol (BGP) for distributing a Bit Index Explicit Replication Traffic/Tree Engineering (BIER-TE) path. A new Tunnel Type for BIER-TE path is defined to encode the information about a BIER-TE path.

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](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

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-Draft

BIER-TE Path

December 2021

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 22 June 2022.

Copyright Notice

Copyright (c) 2021 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the [Trust Legal Provisions](#) and are provided without warranty as described in the Revised BSD License.

Table of Contents

1.	Introduction	3
1.1.	Terminologies	3
2.	Overview of BGP for BIER-TE Path	4
2.1.	Example BIER-TE Topology with BGP	4
2.2.	Distributing Path to Ingress	5
3.	Extensions to BGP	6
3.1.	New SAFI and NLRI	6
3.2.	New Tunnel Type for BIER-TE	7
3.3.	Path BitPositions Sub-TLV	8
3.4.	Path Name Sub-TLV	9
3.5.	Traffic Description Sub-TLVs	10
4.	Security Considerations	11
5.	Acknowledgements	11
6.	IANA Considerations	11
6.1.	Existing Registry: SAFI Parameters	11
6.2.	Existing Registry: BGP TEA Tunnel Types	12
6.3.	Existing Registry: BGP TEA sub-TLVs	12
7.	References	12
7.1.	Normative References	12
7.2.	Informative References	13

Appendix A.	Extensions to PMSI_TUNNEL Attribute	13
A.1.	New Tunnel Type for BIER-TE	14
Authors' Addresses	14

[1.](#) Introduction

[I-D.ietf-bier-te-arch] introduces Bit Index Explicit Replication (BIER) Tree Engineering (BIER-TE). It is an architecture for per-packet stateless explicit point to multipoint (P2MP) multicast path/tree, which is called BIER-TE path, and based on the BIER architecture defined in [[RFC8279](#)].

A Bit-Forwarding Router (BFR) in a BIER-TE domain has a BIER-TE Bit Index Forwarding Table (BIFT). A BIER-TE BIFT on a BFR comprises a forwarding entry for a BitPosition (BP) assigned to each of the adjacencies of the BFR. If the BP represents a forward connected adjacency, the forwarding entry for the BP forwards the multicast packet with the BP to the directly connected BFR neighbor of the adjacency. If the BP represents a BFER (i.e., egress node) or say a local decap adjacency, the forwarding entry for the BP decapsulates the multicast packet with the BP and passes a copy of the payload of the packet to the packet's NextProto within the BFR.

A Bit-Forwarding Ingress Router (BFIR) in a BIER-TE domain receives the information or instructions about which multicast flows/packets are mapped to which BIER-TE paths that are represented by BitPositions or say BitStrings. After receiving the information or instructions, the ingress node/router encapsulates the multicast packets with the BitPositions for the corresponding BIER-TE paths, replicates and forwards the packets with the BitPositions along the BIER-TE paths.

This document proposes some procedures and extensions to BGP for distributing a BIER-TE path to the Bit-Forwarding Ingress Router (BFIR) of the path. It specifies a way of encoding the information about a BIER-TE path in a BGP UPDATE message, which can be distributed to the BFIR of the path.

[1.1.](#) Terminologies

The following terminologies are used in this document.

BIER: Bit Index Explicit Replication.

BIER-TE: BIER Tree Engineering.

BFR: Bit-Forwarding Router.

BFIR: Bit-Forwarding Ingress Router.

BFER: Bit-Forwarding Egress Router.

BFR-id: BFR Identifier. It is a number in the range [1,65535].

BFR-NBR: BFR Neighbor.

BFR-prefix: An IP address (either IPv4 or IPv6) of a BFR.

BIRT: Bit Index Routing Table. It is a table that maps from the BFR-id (in a particular sub-domain) of a BFER to the BFR-prefix of that BFER, and to the BFR-NBR on the path to that BFER.

BIFT: Bit Index Forwarding Table.

P-tunnel: A multicast tunnel through the network of one or more SPs.

PMSI: Provider Multicast Service Interface. PMSI is an abstraction that represents a multicast service for carrying packets. A PMSI is instantiated via one or more P-tunnels.

I-PMSI A-D Route: Inclusive PMSI Auto-Discovery route.

S-PMSI A-D route: Selective PMSI Auto-Discovery route.

x-PMSI A-D route: A route that is either an I-PMSI A-D route or an S-PMSI A-D route.

[2.](#) Overview of BGP for BIER-TE Path

This section briefs the BGP for BIER-TE path and illustrates some details through a simple example BIER-TE topology.

2.1. Example BIER-TE Topology with BGP

An example BIER-TE domain topology using SDN controller with a BGP to distribute BIER-TE path is shown in Figure 1. There are 8 nodes/BFRs A, B, C, D, E, F, G and H in the domain. Nodes/BFRs A, H, E, F and D are BFIRs (i.e., ingress nodes) or BFERs (i.e., egress nodes). The controller has a BGP session with each of the edge nodes in the domain, including BFIRs (i.e., ingress nodes A, H, E, F and D), and each of the non edge nodes in the domain (i.e., nodes B, C and G). Note that some of connections and the BGP on each edge node are not shown in the figure.




Figure 1: Example BIER-TE Topology with Controller

Nodes/BFRs D, F, E, H and A are BFERs (or BFIRs) and have local decap

adjacency BitPositions 1, 2, 3, 4, and 5 respectively.

The BitPositions for the forward connected adjacencies are represented by i' , where i is from 1 to 20.

[2.2.](#) Distributing Path to Ingress

This section describes how the SDN controller distributes a BIER-TE path to its ingress node.

There are two scenarios for distributing the BIER-TE path information. In the first scenario, the ingress node is directly connected to the controller. The path information should not be propagated beyond the ingress node. In the second scenario, the ingress node is not directly connected to the controller. The path information should be propagated throughout the domain until it reaches the ingress node.

Suppose that node A in Figure 1 wants to have a BIER-TE path from ingress node A to egress nodes H and F. The path satisfies a set of constraints. The controller obtains the request from an application or user configuration. It finds a BIER-TE path satisfying the constraints and distributes the path to ingress node A.

If A is directly connected to the controller (e.g., as the example network in Figure 1), then the controller sends A the information about the path in a Update message in one of two ways. In one way, the controller sends each of its BGP peers, including the BGP peer running on node A, a Update message about the explicit path, with a route target matching the BGP identifier of ingress node A, and NO_ADVERTISE community. Ingress node A accepts this message from the controller and installs a forwarding entry for the BIER-TE path, but will not advertise it to any peer. All the other peers do not accept the message.

In another way, the controller sends A a Update message directly through the local session between them, but does not send the message to any other peers. The message contains the information about the path, a route target matching the BGP identifier of ingress node A

and the NO_ADVERTISE community. Ingress node A accepts this message from the controller and installs a forwarding entry for the BIER-TE path, but will not advertise it.

If A is not directly connected to the controller, then the controller distributes the information about the explicit path to the ingress node A across the network. To achieve this, the controller advertises a BGP Update message to all its BGP peers, where the message contains the information about the path, a route target matching the BGP identifier of ingress node A, but does not have NO_ADVERTISE community. Each of the BGP peers advertises the received Update to its BGP neighbors according to normal BGP propagation rules. Eventually, ingress node A accepts this message and installs a forwarding entry for the BIER-TE path, which imports the packets to be transported by the path into the path.

[3.](#) Extensions to BGP

This section defines a new Tunnel Type (or say TLV) for BIER-TE path/tunnel under Tunnel Encapsulation Attribute and a new SAFI. This new SAFI and the existing AFI for IPv4/IPv6 pair uses a new NLRI for indicating a BIER-TE Path.

[3.1.](#) New SAFI and NLRI

A new SAFI, called BIER-TE path SAFI, is defined. Its codepoint (TBD1) is to be assigned by IANA. This new SAFI and the existing AFI for IPv4/IPv6 pair uses a new NLRI, which is defined 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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| NLRI Length |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Distinguisher (4 octets)                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               Tunnel Identifier (11/23 octets)                               ~

```


Tunnel Encaps Attribute (23)
 Tunnel Type (TBD2): BIER-TE Path
 Path BitPositions sub-TLV
 Path Name sub-TLV
 Traffic Description sub-TLV

Where:

- * Tunnel Type (TBD2) is to be assigned by IANA.
- * Path BitPositions sub-TLV encodes the bit positions of the BIER-TE path.
- * Path Name sub-TLV encodes the name of a BIER-TE path.
- * Traffic Description sub-TLV encodes the multicast traffic that is transported by the BIER-TE path.

3.3. Path BitPositions Sub-TLV

The bit positions of a BIER-TE path are encoded in a Path BitPositions sub-TLV. The format of the sub-TLV is illustrated below.

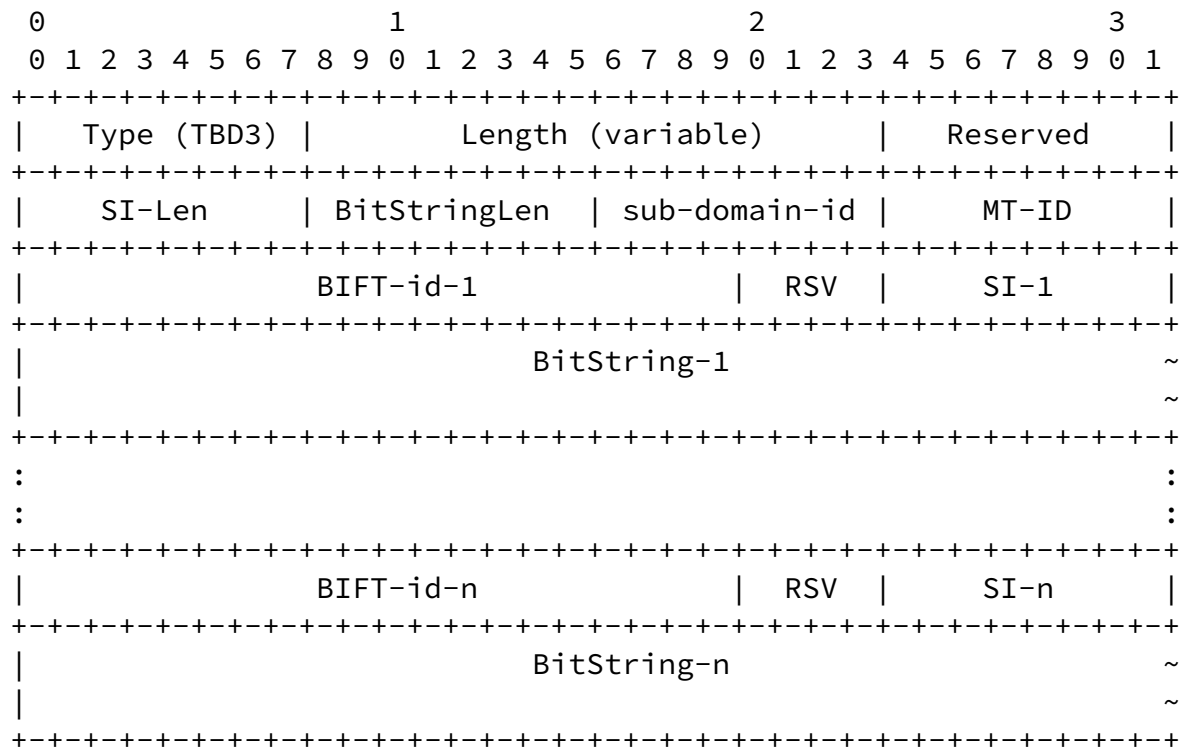


Figure 3: Path BitPositions Sub-TLV Format

Type: Its value (TBD3) is to be assigned by IANA.

Length: It is variable.

Reserved/RSV: MUST be set to zero by the sender and MUST be ignored by the receiver.

SI-Len (SI Length) - 8 bits: The length in bits of the SI field.

BitStringLen (Bit String Length) - 8 bits: The length in bits of the BitString field according to [\[RFC8296\]](#). If k is the length of the BitString, the value of BitStringLen is $\log_2(k)-5$. For example, BitStringLen = 1 indicates $k = 64$, BitStringLen = 7 indicates $k = 4096$.

sub-domain-id: Unique value identifying the BIER sub-domain within the BIER domain.

MT-ID: Multi-Topology ID identifying the topology that is associated with the BIER sub-domain.

<BIFT-id, SI, BitString> tuple: Each tuple <BIFT-id- i , SI- i , BitString- i > ($i = 1, 2, \dots, n$) represents/encodes a set of bit positions on the BIER-TE path with a BIFT ID. All the tuples in the sub-TLV represent/encode the BIER-TE path (i.e., all the bit positions of the BIER-TE path).

[3.4.](#) Path Name Sub-TLV

The name of a BIER-TE path is encoded in a Path Name sub-TLV. The format of the sub-TLV is illustrated below.

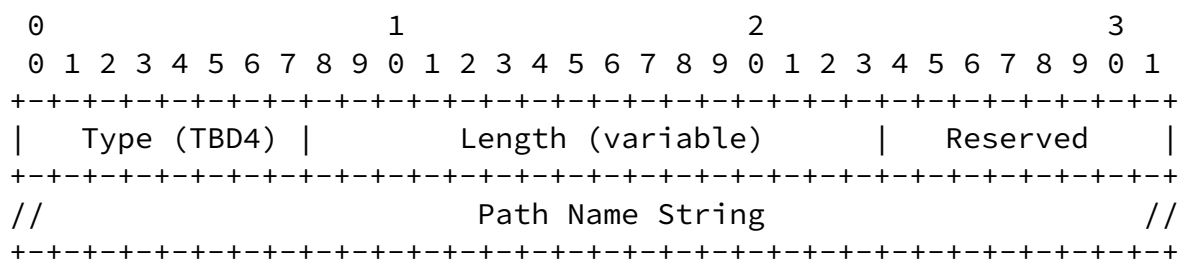


Figure 4: Path Name Sub-TLV Format

Type: Its value (TBD4) is to be assigned by IANA.

Length: It is variable.

Reserved: MUST be set to zero by the sender and MUST be ignored by

the receiver.

Path Name String: It represents/encodes the name of the BIER-TE path in a string of chars.

3.5. Traffic Description Sub-TLVs

A Traffic Description Sub-TLV describes the traffic to be imported into a BIER-TE path. Two Traffic Description Sub-TLVs are defined. They are multicast traffic sub-TLVs for IPv4 and IPv6.

The multicast traffic sub-TLVs for IPv4 and IPv6 are shown in Figure 5 and Figure 6 respectively.

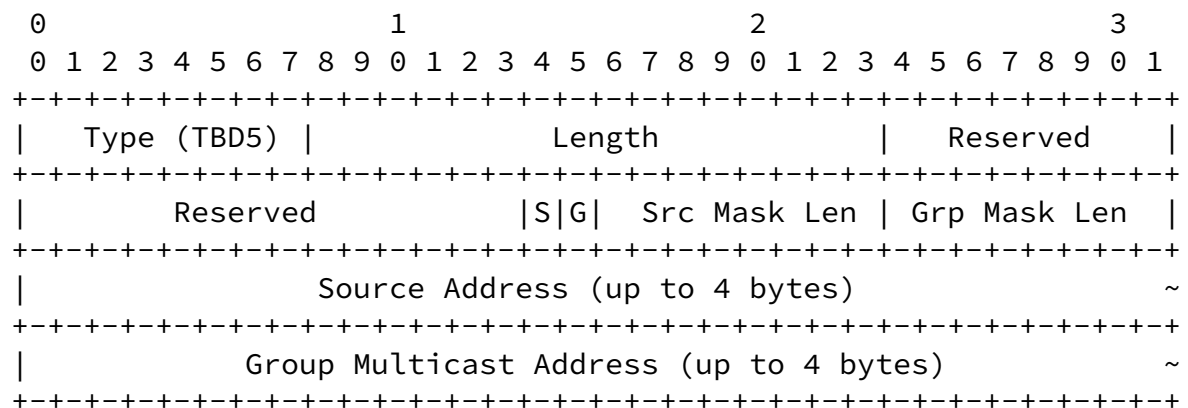
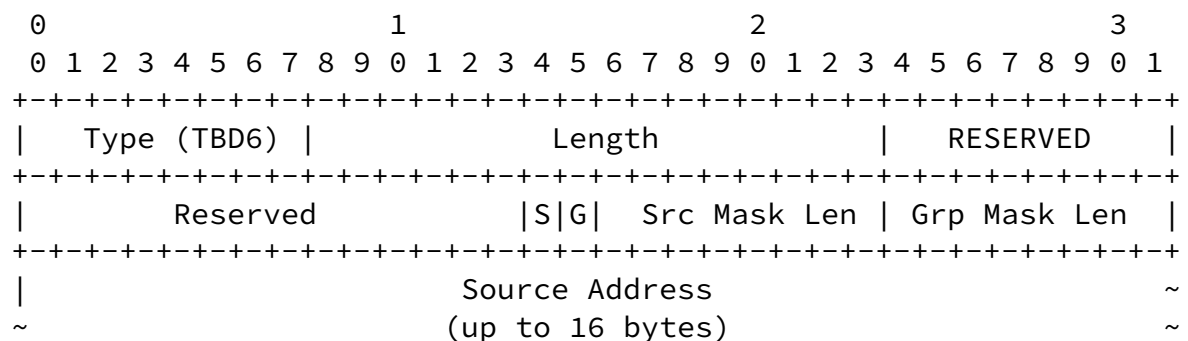


Figure 5: Multicast Traffic for IPv4 Sub-TLV






Figure 6: Multicast Traffic for IPv6 Sub-TLV

The address fields and address mask lengths of the two Multicast Traffic sub-TLVs contain source and group prefixes for matching against packets noting that the two address fields are up to 32 bits for an IPv4 Multicast Traffic and up to 128 bits for an IPv6 Multicast Traffic.

The Reserved field MUST be set to zero and ignored on receipt.

Two bit flags (S and G) are defined to describe the multicast wilddarding in use. If the S bit is set, then source wilddarding is in use and the values in the Source Mask Length and Source Address fields MUST be ignored. If the G bit is set, then group wilddarding is in use and the values in the Group Mask Length and Group multicast Address fields MUST be ignored. The G bit MUST NOT be set unless the S bit is also set: if a Multicast Traffic sub-TLV is received with S bit = 0 and G bit = 1 the receiver MUST respond with an error (Malformed Multicast Traffic).

The three multicast mappings may be achieved as follows:

- (S, G): S bit = 0, G bit = 0, the Source Address and Group multicast Address prefixes are both used to define the multicast traffic.
- (*, G): S bit = 1, G bit = 0, the Group multicast Address prefix is used to define the multicast traffic, but the Source Address prefix is ignored.
- (*, *): S bit = 1, G bit = 1, the Source Address and Group multicast Address prefixes are both ignored.

4. Security Considerations

Protocol extensions defined in this document do not affect the BGP security other than those as discussed in the Security Considerations section of [RFC9012].

5. Acknowledgements

The authors of this document would like to thank Tony Przygienda, Susan Hares, and Jeffrey Zhang for their comments.

6. IANA Considerations

6.1. Existing Registry: SAFI Parameters

This document requests assigning a new SAFI in the registry "Subsequent Address Family Identifiers (SAFI) Parameters" as follows:

Chen, et al.

Expires 22 June 2022

[Page 11]

Internet-Draft

BIER-TE Path

December 2021

Code Point	Description	Reference
TBD1(179 suggested)	BIER-TE Policy SAFI	This document

6.2. Existing Registry: BGP TEA Tunnel Types

This document requests assigning a new Tunnel-Type in the registry "BGP Tunnel Encapsulation Attribute Tunnel Types" as follows:

Code Point	Description	Reference
TBD2(16 suggested)	BIER-TE Tunnel/Path	This document

6.3. Existing Registry: BGP TEA sub-TLVs

This document requests assigning a few of new sub-TLVs in the registry "BGP Tunnel Encapsulation Attribute sub-TLVs" as follows:

Code Point	Description	Reference
------------	-------------	-----------

TBD3(16 suggested)	Path BitPositions	This document
+=====+	+=====+	+=====+
TBD4(17 suggested)	Path Name	This document
+=====+	+=====+	+=====+
TBD5(18 suggested)	IPv4 Multicast Traffic	This document
+=====+	+=====+	+=====+
TBD6(19 suggested)	IPv6 Multicast Traffic	This document
+=====+	+=====+	+=====+

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", [RFC 6514](#), DOI 10.17487/RFC6514, February 2012, <<https://www.rfc-editor.org/info/rfc6514>>.

Chen, et al.

Expires 22 June 2022

[Page 12]

Internet-Draft

BIER-TE Path

December 2021

- [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>>.
- [RFC8279] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast Using Bit Index Explicit Replication (BIER)", [RFC 8279](#), DOI 10.17487/RFC8279, November 2017, <<https://www.rfc-editor.org/info/rfc8279>>.
- [RFC8296] Wijnands, IJ., Ed., Rosen, E., Ed., Dolganow, A., Tantsura, J., Aldrin, S., and I. Meilik, "Encapsulation for Bit Index Explicit Replication (BIER) in MPLS and Non-MPLS Networks", [RFC 8296](#), DOI 10.17487/RFC8296, January 2018, <<https://www.rfc-editor.org/info/rfc8296>>.
- [RFC9012] Patel, K., Van de Velde, G., Sangli, S., and J. Scudder, "The BGP Tunnel Encapsulation Attribute", [RFC 9012](#),

7.2. Informative References

- [I-D.ietf-bier-te-arch]
Eckert, T., Cauchie, G., and M. Menth, "Tree Engineering for Bit Index Explicit Replication (BIER-TE)", Work in Progress, Internet-Draft, [draft-ietf-bier-te-arch-11](#), 15 November 2021, <<https://www.ietf.org/archive/id/draft-ietf-bier-te-arch-11.txt>>.
- [RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", [RFC 5226](#), DOI 10.17487/RFC5226, May 2008, <<https://www.rfc-editor.org/info/rfc5226>>.
- [RFC5575] Marques, P., Sheth, N., Raszuk, R., Greene, B., Mauch, J., and D. McPherson, "Dissemination of Flow Specification Rules", [RFC 5575](#), DOI 10.17487/RFC5575, August 2009, <<https://www.rfc-editor.org/info/rfc5575>>.

Appendix A. Extensions to PMSI_TUNNEL Attribute

This section defines a new Tunnel Type (or TLV) for BIER-TE path/tunnel under the PMSI_TUNNEL Attribute (PTA) defined in [[RFC6514](#)]. It describes a couple of new sub-TLVs encoding the information about a BIER-TE path.

A.1. New Tunnel Type for BIER-TE

The PMSI Tunnel attribute carried by an x-PMSI A-D route identifies P-tunnel for PMSI. For the PTA with Tunnel Type BIER-TE, the PTA is constructed by the SDN controller and distributed to the ingress node of the BIER-TE tunnel.

The format of the PMSI_TUNNEL Attribute with the new Tunnel Type (TBD) for BIER-TE is shown in Figure 7.

Email: huaimo.chen@futurewei.com

Mike McBride
Futurewei

Email: michael.mcbride@futurewei.com

Ran Chen
ZTE Corporation

Email: chen.ran@zte.com.cn

Gyan S. Mishra
Verizon Inc.
13101 Columbia Pike
Silver Spring, MD 20904
United States of America

Phone: 301 502-1347
Email: gyan.s.mishra@verizon.com

Aijun Wang
China Telecom
Beiqijia Town, Changping District
Beijing
102209
China

Email: wangaj3@chinatelecom.cn

Yisong Liu
China Mobile

Email: liuyisong@chinamobile.com

Yanhe Fan
Casa Systems
United States of America

Email: yfan@casa-systems.com

Boris Khasanov
Yandex LLC
Moscow

Email: bhassanov@yahoo.com

Lei Liu
Fujitsu
United States of America

Email: liulei.kddi@gmail.com

Xufeng Liu
Volta Networks
McLean, VA
United States of America

Email: xufeng.liu.ietf@gmail.com

