

Networking Working Group
Internet-Draft
Intended status: Standards Track
Expires: 1 September 2023

R. Chen
Zh. Zhang
ZTE Corporation
H. Chen
S. Dhanaraj
Futurewei
F. Qin
China Mobile
A. Wang
China Telecom
28 February 2023

PCEP Extensions for BIER-TE
draft-chen-pce-bier-10

Abstract

Bit Index Explicit Replication (BIER)-TE shares architecture and packet formats with BIER as described in [RFC8279]. BIER-TE forwards and replicates packets based on a BitString in the packet header, but every BitPosition of the BitString of a BIER-TE packet indicates one or more adjacencies as described in [RFC9262]. BIER-TE Path can be derived from a Path Computation Element (PCE).

This document specifies extensions to the Path Computation Element Protocol (PCEP) that allow a PCE to compute and initiate the path for the BIER-TE.

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 1 September 2023.

Copyright Notice

Copyright (c) 2023 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](#)
- [2.](#) Conventions used in this document [3](#)
- [3.](#) Overview of PCEP Operation in BIER Networks [3](#)
- [4.](#) LSP Operations [3](#)
- [5.](#) PCEP Messages [4](#)
- [6.](#) Object Formats [4](#)
 - [6.1.](#) The OPEN Object [4](#)
 - [6.1.1.](#) The BIER-TE PCE Capability sub-TLV [4](#)
 - [6.2.](#) the LSP Object [5](#)
 - [6.3.](#) The RP/SRP Object [6](#)
 - [6.4.](#) END-POINTS object [6](#)
 - [6.5.](#) Objective Functions [7](#)
 - [6.6.](#) ERO Object [7](#)
 - [6.6.1.](#) BIER-TE-ERO Subobject [7](#)
 - [6.7.](#) RRO Object [9](#)
- [7.](#) Procedures [9](#)
 - [7.1.](#) Exchanging the BIER-TE Capability [9](#)
 - [7.2.](#) BIER-TE-ERO Processing [10](#)
 - [7.3.](#) BIER-TE-RRO Processing [10](#)
- [8.](#) IANA Considerations [10](#)
 - [8.1.](#) New Path Setup Type [11](#)
 - [8.2.](#) BIER-TE-PCE-CAPABILITY Sub-TLV Type Indicators [11](#)
 - [8.3.](#) PCEP TLV Type Indicators [11](#)
 - [8.4.](#) Objective Functions [12](#)
 - [8.5.](#) BIER-TE-ERO and RRO Subobjects [12](#)
 - [8.6.](#) PCEP-Error Objects and Types [12](#)
- [9.](#) Security Considerations [13](#)
- [10.](#) Acknowledgements [13](#)
- [11.](#) Normative references [13](#)
- Authors' Addresses [15](#)

1. Introduction

Bit Index Explicit Replication (BIER)-TE shares architecture and packet formats with BIER as described in [RFC8279]. BIER-TE forwards and replicates packets based on a BitString in the packet header, but every BitPosition of the BitString of a BIER-TE packet indicates one or more adjacencies as described in [RFC9262]. BIER-TE Path can be derived from a Path Computation Element (PCE).

[RFC8623] specifies a set of extensions to PCEP that allow a PCE to compute and recommend network paths in compliance with [RFC4657] and defines objects and TLVs for P2MP TE LSPs.

This document uses a PCE for computing one or more BIER-TE paths taking into account various constraints and objective functions.

2. Conventions used in this document

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.

3. Overview of PCEP Operation in BIER Networks

BIER-TE forwards and replicates packets based on a BitString in the packet header, and every BitPosition of the BitString of a BIER-TE packet indicates one or more adjacencies as described in [RFC9262]. In a PCEP session, An ERO object specified in [RFC5440] can be extended to carry a BIER-TE path consists of one or more BIER-TE-ERO subobject(s). BIER-TE computed by a PCE can be represented in the following forms:

- * An ordered set of adjacencies BitString(s) in which each bit represents that the adjacencies to which the BFR should replicate packets to in the domain.

In this document, we define a set of PCEP protocol extensions, including a new PCEP capability, a new Path Setup Type (PST), reuse BIER END-POINT Object, a new Objective Functions subobjects, a new ERO subobjects, a new RRO subobjects, a new PCEP error codes and procedures.

4. LSP Operations

LSP operations for active and passive stateful PCE operations and on P2MP TE LSPs (described in [RFC8623]) are applicable for BIER-TE LSPs as well.

5. PCEP Messages

The PCEP Message of P2MP TE LSPs(defined in [[RFC8623](#)]) are applicable for BIER-TE LSPs as well.

The PCReq message, PCRep message and PCRpt message may be extended to support encoding of OF object so that to indicate the required/ desired objective function to be applied by the PCE during path computation or within a PCRep/PCRpt message so as to indicate the objective function that was used by the PCE during path computation.

6. Object Formats

6.1. The OPEN Object

This document defines one new optional TLV for use in the OPEN object.

6.1.1. The BIER-TE PCE Capability sub-TLV

[[RFC8408](#)] defines the PATH-SETUP-TYPE-CAPABILITY TLV for use in the OPEN object. The PATH-SETUP-TYPE-CAPABILITY TLV contains an optional list of sub-TLVs which are intended to convey parameters that are associated with the path setup types supported by a PCEP speaker.

This document defines a new Path Setup Type (PST) for BIER-TE as follows:

* PST = TBD1: Path is setup using BIER-TE technique.

A PCEP speaker MUST indicate its support of the function described in this document by sending a PATH-SETUP-TYPE-CAPABILITY TLV in the OPEN object with this new PST included in the PST list.

This document also defines the BIER-TE-PCE-CAPABILITY sub-TLV. PCEP speakers use this sub-TLV to exchange BIER-TE capability. If a PCEP speaker includes PST=TBD1 in the PST List of the PATH-SETUP-TYPE-CAPABILITY TLV then it MUST also include the BIER-TE-PCE-CAPABILITY sub-TLV inside the PATH-SETUP-TYPE-CAPABILITY TLV.

The format of the BIER-TE-PCE-CAPABILITY sub-TLV is shown in the following figure:




Figure 1: Figure 1 BIER-TE-PCE-CAPABILITY sub-TLV format

The code point for the TLV type is to be defined by IANA.

Length: 4 bytes.

Flags: A single flag is defined (as per setion 7.1.1 of [\[RFC8296\]](#)):

- * U (1bit):if set to 1 by a PCC, the U flag indicates that the PCC allows modification of LSP parameters; if set to 1 by a PCE, the U flag indicates that the PCE is capable of updating LSP parameters. The flag must be advertised by both a PCC and a PCE for PCUpd messages to be allowed on a PCEP session.
- * The remaining "Flags" fields are currently unused, and MUST be set to zero on transmission and ignored on reception.

6.2. the LSP Object

[RFC8623] specifies the IPv4 and IPv6 P2MP-LSP-IDENTIFIERS TLVs to be included in the LSP object. For BIER-TE LSP, this document defines BIER-TE-IDENTIFIERS TLVs for the LSP object. The BIER-TE-IDENTIFIERS TLV MUST be included in the LSP object in a PCRpt message for BIER-TE LSP. If the P2MP-LSP-IDENTIFIER TLV is missing, the PCE MUST respond with a PCErr message carrying error-type 6 ("mandatory object missing") and error-value TBD3 ("BIER-TE-IDENTIFIERS TLV missing") and close the PCEP session.

The BIER-TE-IDENTIFIERS TLV MAY optionally be included in the LSP object in the PCUpd, the PCReq and the PCRep message for BIER-TE LSPs and the BIER-TE-IDENTIFIERS TLV SHOULD NOT be included in a PCInitiate message.

The format of the BIER-TE-IDENTIFIERS TLV is shown in Figure 2:




Figure 2: Figure 2

Type: TBD4

Length: 1 octet. Contains the total length of the TLV in octets.

Tunnel Identifier(as per [I-D.ietf-idr-bier-te-path]): 11/23 octet value contains:

- * sub-domain-id (1 octet): It is id of the sub domain through which the BIER-TE tunnel crosses.
- * BFR-id (2 octets): It is the BFR-id of the BFIR of the BIER-TE tunnel.
- * Tunnel-ID (4 octets): It is a number uniquely identifying a BIER-TE tunnel within the BFIR and sub domain.
- * BFR-prefix (4/16 octets): It is a BFR-prefix of the BFIR of the BIER-TE tunnel. It occupies 4 octets for IPv4 and 16 octets for IPv6.

6.3. The RP/SRP Object

In order to setup an BIER-TE, a new PATH-SETUP-TYPE TLV MUST be contained in RP/SRP object. This document defines a new Path Setup Type (PST=TBD1) for BIER-TE.

6.4. END-POINTS object

The END-POINTS object which is defined in [RFC8306]is used in a PCReq message to specify the BIER information of the path for which a path computation is requested. To represent the end points for a BIER path efficiently, we reuse the P2MP END-POINTS object body for IPv4(Object-Type 3) and END-POINTS object body for IPv6 (Object-Type 4) which is defined in [RFC8306].

6.5. Objective Functions

[RFC5541] defines a mechanism to specify an objective function (OF) that is used by a PCE when it computes a path. For a BIER-TE path, a new OF is defined.

Objective Function Code: TBD5

Name: Minimum Bit Sets (MBS)

Description: Find a path represented by BitPositions that has the minimum number of bit sets.

For each bit set that represents a part of the BIER-TE path, the ingress of the path constructs a copy of the packet containing the bit set and applies the BIER-TE forwarding procedure to forward the packet copy. When a path is computed to have the minimum number of bit sets, the ingress of the path generates the minimum number of the packet copies and applies the BIER-TE forwarding procedure in the minimum number of times. The number of packet copies generated and transmitted in the network along the path may be minimum.

6.6. ERO Object

BIER-TE consists of one or more adjacencies BitStrings where every BitPosition of the BitString indicates one or more adjacencies, as described in([RFC8279]).

The ERO object specified in [RFC5440] is used to encode the path of a TE LSP through the network. The ERO is carried within a PCRep message to provide the computed TE LSP if the path computation was successful. In order to carry BIER-TE explicit paths, this document defines a new ERO subobjects referred to as "BIER-TE-ERO subobjects" whose formats are specified in the following section. An BIER-TE-ERO subobjects carrying a adjacencies BitStrings consists of one or more BIER-TE-ERO subobject(s).

6.6.1. BIER-TE-ERO Subobject

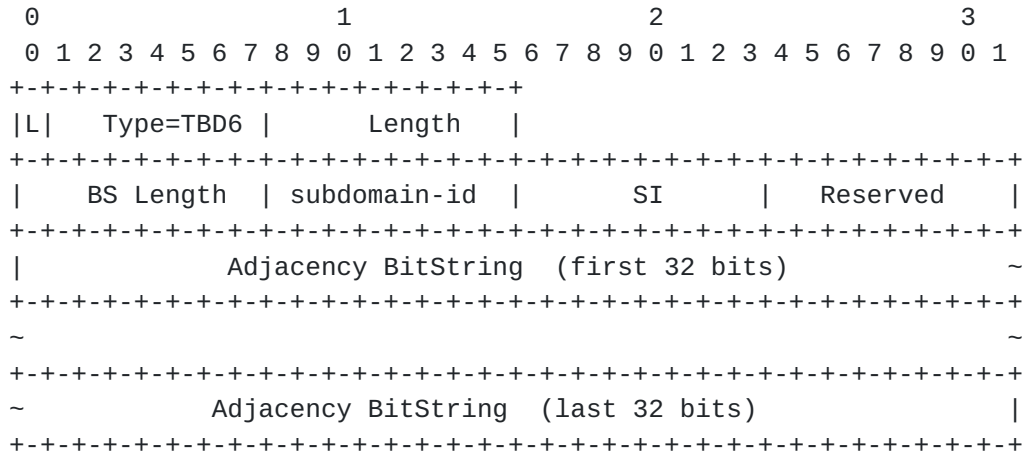


Figure 3: Figure 3

The 'L' Flag: Indicates whether the subobject represents a loose-hop in the LSP[RFC3209]. If the bit is not set, the subobject represents a strict hop in the explicit route.

Type: TBD6

Length: 1 octet ([RFC3209]). Contains the total length of the subobject in octets. The Length MUST be at least 8, and MUST be a multiple of 4.

BS Length: A 1 octet field encodes the length in bits of the BitString as per [RFC8296], the maximum length of the BitString is 5, it indicates the length of BitString is 1024.It is used to refer to the number of bits in the BitString.

subdomain-id: Unique value identifying the BIER subdomain. 1 octet.

SI: Set Identifier ([Section 1 of \[RFC8279\]](#) used in the encapsulation for this BIER subdomain for this BitString length, 1 octet.

The "Reserved" (1 octets) fields are currently unused, and MUST be set to zero on transmission and ignored on reception.

Adjacency BitString: a variable length field encoding the Adjacency BitString where every BitPosition of the BitString indicates one or more adjacencies.the length of this field is according the BS length. The minimum value of this field is 64 bits, and the maximum value of this field is 1024 bits.

Notice:

The maximum value of BS Length is limited to the 1024 bits, in case the BIER-TE-ERO Subobject is too long.

6.7. RRO Object

An RRO contains one or more subobjects called "BIER-TE-RRO subobjects", whose format is shown below:

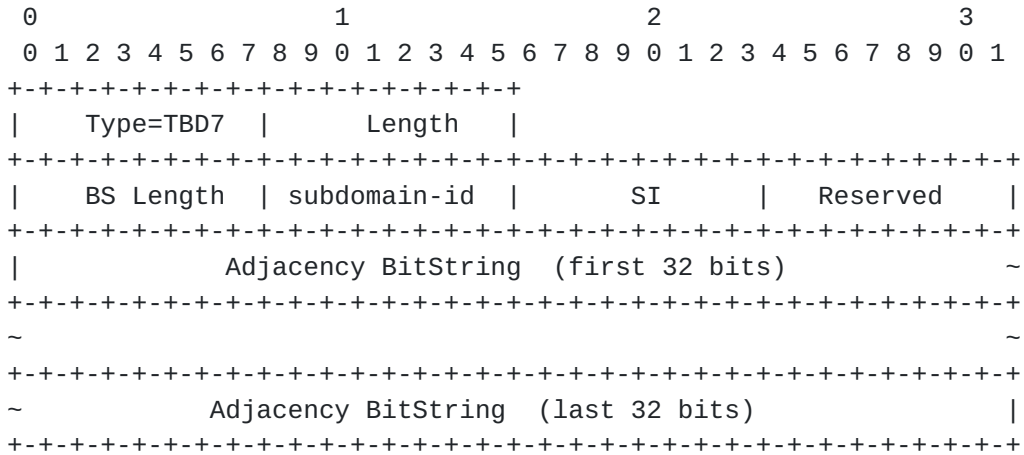


Figure 4: Figure 4

The format of the BIER-TE-RRO subobject is the same as that of the BIER-TE-ERO subobject, but without the L-Flag.

For the integrity of the protocol, we define a new BIER-TE-RRO object, but its actual value is consistent with ERO. The PCC reports an BIER-TE to a PCE by sending a PCRpt message with RRO object.

7. Procedures

7.1. Exchanging the BIER-TE Capability

A PCC indicates that it is capable of supporting the head-end functions for BIER-TE by including the BIER-TE-PCE-CAPABILITY sub-TLV in the Open message that it sends to a PCE. A PCE indicates that it is capable of computing BIER-TE by including the BIET-TE-PCE-CAPABILITY sub-TLV in the Open message that it sends to a PCC.

If a PCEP speaker receives a PATH-SETUP-TYPE-CAPABILITY TLV with a PST list containing PST=TBD1, and supports that path setup type, then it checks for the presence of the SR-PCE-CAPABILITY sub-TLV. If that sub-TLV is absent, then the PCEP speaker MUST send a PCErr message

with Error-Type = 10 ("Reception of an invalid object") and Error-value = TBD8("Missing PCE-BIER-TE-CAPABILITY sub-TLV") and MUST then close the PCEP session. If a PCEP speaker receives a PATH-SETUP-TYPE-CAPABILITY TLV with a BIER-TE-PCE-CAPABILITY sub-TLV, but the PST list does not contain PST=TBD1, then the PCEP speaker MUST ignore the BIER-TE-PCE-CAPABILITY sub-TLV.

7.2. BIER-TE-ERO Processing

If a PCC does not support the BIER-TE PCE Capability and thus cannot recognize the BIER-TE-ERO or BIER-TE-RR0 subobjects, The ERO and BIER-TE-ERO subobject processing remains as per [[RFC5440](#)].

If a PCC receives an BIER-TE-ERO subobject in which either BitStringLength or Adjacency BitString or SI is absent, it MUST consider the entire BIER-TE-ERO subobject invalid and send a PCerr message with Error-Type = 10 ("Reception of an invalid object"), Error-Value = TBD9 ("BitStringLength is absent ") or Error-Value = TBD10 ("Adjacency BitString is absent") or Error-Value = TBD11("SI is absent").

If a PCC receives an BIER-TE-ERO subobject in which BitStringLength values are not chosen from: 64, 128, 256, 512, 1024, as it described in ([[RFC8279](#)]). The PCC MUST send a PCerr message with Error-Type = 10 ("Reception of an invalid object") and Error-Value = TBD12 ("Invalid BitStringLength").

When a PCEP speaker detects that all subobjects of ERO are not of type TBD6, and if it does not handle such ERO, it MUST send a PCerr message with Error-Type = 10 ("Reception of an invalid object") and Error-Value = TBD13 ("Non-identical ERO subobjects") as per [[RFC8664](#)].

7.3. BIER-TE-RR0 Processing

The syntax checking rules that apply to the BIER-TE-RR0 subobject are identical to those of the BIER-TE-ERO subobject

The actual value of BIER-TE-RR0 subobject is consistent with ERO. The PCC reports an BIER-TE to a PCE by sending a PCRpt message with RR0 object.

8. IANA Considerations

IANA has registered the code points for the protocol elements defined in this document.

8.1. New Path Setup Type

A sub-registry within the "Path Computation Element Protocol (PCEP) Numbers" registry called "PCEP Path Setup Types" was created in [RFC8408]. The document requests a new codepoint within this registry, as follows:

value	Meaning	Reference
TBD1	Path is setup using BIER-TE technique	This Document

8.2. BIER-TE-PCE-CAPABILITY Sub-TLV Type Indicators

IANA has created a new sub-registry, named "PATH-SETUP-TYPE-CAPABILITY Sub-TLV Type Indicators", within the "Path Computation Element Protocol (PCEP) Numbers" registry to manage the type indicator space for sub-TLVs of PATH-SETUP-TYPE-CAPABILITY TLV. This document defines a new sub-TLV type.

value	Meaning	Reference
TBD2	BIER-TE-PCE-CAPABILITY	This Document

8.3. PCEP TLV Type Indicators

The document requests a new code point in the existing "PCEP TLV Type Indicators" registry as follows:

value	Meaning	Reference
TBD4	BIER-TE-IDENTIFIERS TLV	This Document

8.4. Objective Functions

This document requests a new objective functions from the "Objective Function" subregistry within the "Path Computation Element Protocol (PCEP) Numbers" registry:

Value	Meaning	Reference
TBD5	Minimum Bit Sets (MBS)	This Document

8.5. BIER-TE-ERO and RRO Subobjects

This document defines a new subobject type for the PCEP explicit route object (ERO) and a new subobject type for the PCEP RRO. The code points for subobject types of these objects are maintained in the RSVP parameters registry, under the EXPLICIT_ROUTE and ROUTE_RECORD objects, respectively.

Object	Subobject	Subobject Type
EXPLICIT_ROUTE	BIER-TE-ERO (PCEP specific)	TBD6
ROUTE_RECORD	BIER-TE-RRO (PCEP specific)	TBD7

8.6. PCEP-Error Objects and Types

IANA is requested to allocate code-points in the "PCEP-ERROR Object Error Types and Values" subregistry for the following new error-types and error-values:

Error-Type	Meaning	Error-value
6	mandatory object missing	
		TBD3:BIER-TE-IDENTIFIERS TLV missing
10	Reception of an invalid object	
		TBD8: Missing PCE-BIER-TE-CAPABILITY subobjects
		TBD9: BitStringLength is absent
		TBD10: Adjacency BitString is absent
		TBD11: SI is absent
		TBD12: Invalid BitStringLength
		TBD13: Non-identical ERO subobjects

9. Security Considerations

The security considerations described in [RFC5440], [RFC8231], [RFC8281] and[RFC8408]are applicable to this specification. No additional security measures are required.

10. Acknowledgements

The authors thank Dhruv Dhody, Benchong Xu, Chun Zhu, and Zhaohui Zhang and many others for their suggestions and comments.

11. Normative references

[I-D.ietf-idr-bier-te-path]
 Chen, H., McBride, M., Chen, R., Mishra, G. S., Wang, A., Liu, Y., Fan, Y., Khasanov, B., Liu, L., and X. Liu, "BGP for BIER-TE Path", Work in Progress, Internet-Draft, [draft-ietf-idr-bier-te-path-01](https://datatracker.ietf.org/doc/html/draft-ietf-idr-bier-te-path-01), 7 January 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-idr-bier-te-path-01>>.

- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", [RFC 3209](#), DOI 10.17487/RFC3209, December 2001, <<https://www.rfc-editor.org/info/rfc3209>>.
- [RFC4657] Ash, J., Ed. and J.L. Le Roux, Ed., "Path Computation Element (PCE) Communication Protocol Generic Requirements", [RFC 4657](#), DOI 10.17487/RFC4657, September 2006, <<https://www.rfc-editor.org/info/rfc4657>>.
- [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>>.
- [RFC5541] Le Roux, JL., Vasseur, JP., and Y. Lee, "Encoding of Objective Functions in the Path Computation Element Communication Protocol (PCEP)", [RFC 5541](#), DOI 10.17487/RFC5541, June 2009, <<https://www.rfc-editor.org/info/rfc5541>>.
- [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>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC8306] Zhao, Q., Dhody, D., Ed., Palleti, R., and D. King, "Extensions to the Path Computation Element Communication Protocol (PCEP) for Point-to-Multipoint Traffic

Engineering Label Switched Paths", [RFC 8306](#),
DOI 10.17487/RFC8306, November 2017,
<<https://www.rfc-editor.org/info/rfc8306>>.

- [RFC8408] Sivabalan, S., Tantsura, J., Minei, I., Varga, R., and J. Hardwick, "Conveying Path Setup Type in PCE Communication Protocol (PCEP) Messages", [RFC 8408](#), DOI 10.17487/RFC8408, July 2018, <<https://www.rfc-editor.org/info/rfc8408>>.
- [RFC8623] Palle, U., Dhody, D., Tanaka, Y., and V. Beeram, "Stateful Path Computation Element (PCE) Protocol Extensions for Usage with Point-to-Multipoint TE Label Switched Paths (LSPs)", [RFC 8623](#), DOI 10.17487/RFC8623, June 2019, <<https://www.rfc-editor.org/info/rfc8623>>.
- [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>>.
- [RFC9262] Eckert, T., Ed., Menth, M., and G. Cauchie, "Tree Engineering for Bit Index Explicit Replication (BIER-TE)", [RFC 9262](#), DOI 10.17487/RFC9262, October 2022, <<https://www.rfc-editor.org/info/rfc9262>>.

Authors' Addresses

Ran Chen
ZTE Corporation
Email: chen.ran@zte.com.cn

Zheng Zhang
ZTE Corporation
Email: zhang.zheng@zte.com.cn

Huaimo Chen
Futurewei
Email: huaimo.chen@futurewei.com

Senthil Dhanaraj
Futurewei
Email: senthil.dhanaraj.ietf@gmail.com

Fengwei Qin
China Mobile
Email: qinfengwei@chinamobile.com

Aijun Wang
China Telecom
Email: wangaj3@chinatelecom.cn