

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: October 17, 2021

H. Chen
M. McBride
Futurewei
A. Wang
China Telecom
G. Mishra
Verizon Inc.
Y. Liu
China Mobile
Y. Fan
Casa Systems
L. Liu
Fujitsu
X. Liu
Volta Networks
April 15, 2021

**PCE for BIER-TE Path
draft-chen-pce-bier-te-path-01**

Abstract

This document describes extensions to Path Computation Element (PCE) communication Protocol (PCEP) for supporting Bit Index Explicit Replication (BIER) Traffic Engineering (TE) paths.

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 [RFC 2119](#) [[RFC2119](#)].

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 October 17, 2021.

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 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.](#) Terminologies [4](#)
- [2.](#) Overview of PCE for BIER-TE [5](#)
- [2.1.](#) Example BIER-TE Topology with PCE [5](#)
- [2.2.](#) A Brief Flow of PCEP Messages for a BIER-TE Path [6](#)
- [2.3.](#) Procedures on Ingress [8](#)
- [3.](#) Extensions to PCEP [9](#)
- [3.1.](#) BIER-TE Path Capability [9](#)
- [3.2.](#) Extensions to SRP [10](#)
- [3.2.1.](#) SRP Object Flag Field [10](#)
- [3.2.2.](#) Reuse of Multicast Flow Specification TLV [11](#)
- [3.3.](#) Ingress Node Object [11](#)
- [3.4.](#) Objective Functions [13](#)
- [3.5.](#) BIER-TE Path Subobject [14](#)
- [3.6.](#) BIER-TE Path Subobject in ERO [15](#)
- [3.7.](#) BIER-TE Path Subobject in RRO [15](#)
- [4.](#) Procedures [16](#)
- [4.1.](#) BIER-TE Path Creation [16](#)
- [4.2.](#) BIER-TE Path Update [17](#)
- [4.3.](#) BIER-TE Path Deletion [17](#)
- [5.](#) The PCEP Messages [17](#)
- [5.1.](#) The PCRpt Message [17](#)
- [5.2.](#) The PCUpd Message [18](#)
- [5.3.](#) The PCInitiate Message [18](#)
- [5.4.](#) The PCReq Message [18](#)
- [5.5.](#) The PCRep Message [18](#)
- [6.](#) IANA Considerations [19](#)
- [6.1.](#) PST for BIER-TE Path [19](#)
- [6.2.](#) PCE-BIER-TE-Path Capability sub-TLV [19](#)

6.3.	SRP Object Flag Field	19
6.4.	Ingress Node Object	20
6.5.	OF Code Points	20
6.6.	PCEP BIER-TE Path Subobjects	20
7.	Security Considerations	21
8.	Acknowledgements	21
9.	References	21
9.1.	Normative References	21
9.2.	Informative References	22
	Authors' Addresses	23

[1.](#) Introduction

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

A Bit-Forwarding Ingress Router (BFIR) in a BIER-TE domain receives the information or instructions from a controller such as a stateful PCE about which multicast flows/packets are mapped to which P2MP paths. The multicast flows/packets are indicated by multicast and source addresses. The paths 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 P2MP paths, replicates and forwards the packets with the BitPositions along the P2MP paths.

[RFC8231] describes a set of extensions to PCEP to provide stateful control. A stateful PCE has access to not only the information carried by the network's Interior Gateway Protocol (IGP) but also the set of active paths and their reserved resources. The additional state allows the PCE to compute constrained paths while considering individual paths and their interactions.

To compute and initiate BIER-TE P2MP paths, the stateful PCE needs to be extended. For a BIER-TE P2MP path, some new state information will be stored and maintained, which includes the BitPositions, multicast group and multicast source for the path. The PCE gets the egresses of the path, the same multicast group and source from the egresses when each of the egresses reports to the PCE that it receives a multicast join with the multicast group and source. With this information, the PCE finds an ingress for the path, computes the path from the ingress to the egresses that has the optimal BitPositions and satisfies the constraints, and then initiates the BIER-TE path at the ingress of the path through sending the ingress the BitPositions of the path, multicast group and source in a PCEP message such as PCInitiate. After receiving the message, the ingress

creates a forwarding entry that imports the packets with the multicast group/address and source into the BIER-TE path (i.e., encapsulates the packets with a BIER-TE header having the BitPositions of the path), and then reports the status of the path to the PCE in a PCEP message such as PCRpt.

[I-D.chen-pce-bier] describes part of the solution for this, which is mainly the BIER-ERO subobject used for P2MP paths.

This document proposes a comprehensive solution for computing and establishing BIER-TE P2MP paths.

1.1. Terminologies

The following terminologies are used in this document.

PCE: Path Computation Element

PCEP: PCE communication Protocol

PCC: Path Computation Client

CE: Customer Edge

PE: Provider Edge

BIER: Bit Index Explicit Replication.

BIER-TE: BIER Traffic/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.

LSP-DB: Label Switching Path DataBase.

TED: Traffic/Tree Engineering DataBase.

2. Overview of PCE for BIER-TE

This section briefly describes PCE for BIER-TE and illustrates some details through a simple example BIER-TE topology.

2.1. Example BIER-TE Topology with PCE

An example BIER-TE topology for a BIER-TE domain with a PCE 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). There is a connection (i.e., PCE session) between the PCE and the PCC running on each of the possible ingress and egress nodes in the domain. Note that some of connections and the PCC on each node are not shown in the figure.

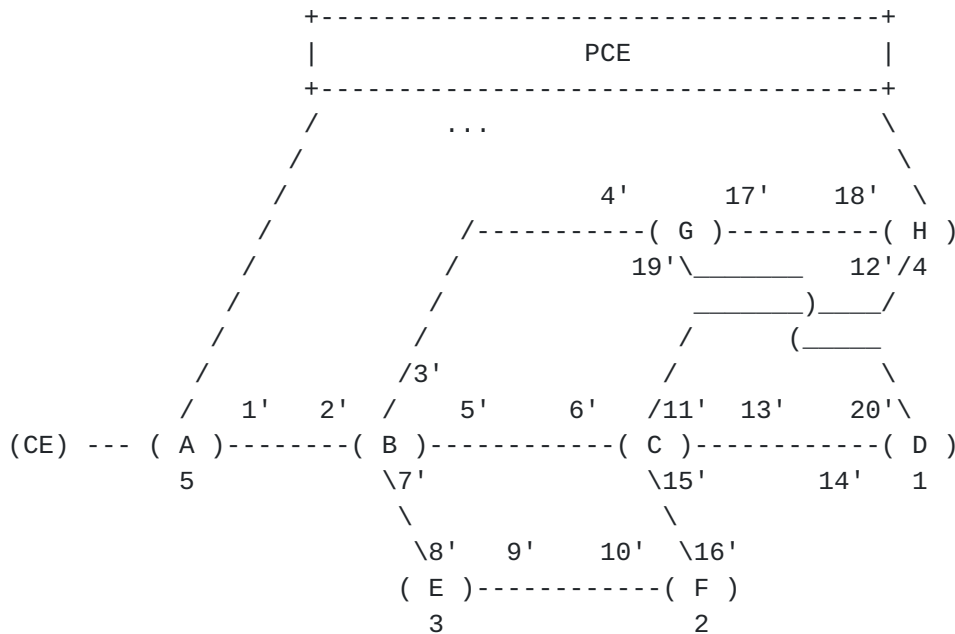


Figure 1: Example BIER-TE Topology with PCE

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. For simplicity, these BPs are represented by (SI:BitString), where SI = 0 and BitString is of 8 bits. BPs 1, 2, 3, 4, and 5 are represented by 1 (0:00000001), 2 (0:00000010), 3 (0:00000100), 4 (0:00001000) and 5 (0:00010000) respectively.

The BitPositions for the forward connected adjacencies are represented by i' , where i is from 1 to 20. In one option, they are encoded as $(n+i)$, where n is a power of 2 such as 32768. For simplicity, these BitPositions are represented by $(SI:BitString)$, where $SI = (6 + (i-1)/8)$ and $BitString$ is of 8 bits. BitPositions i' (i from 1 to 20) are represented by $1'(6:00000001)$, $2'(6:00000010)$, $3'(6:00000100)$, $4'(6:00001000)$, $5'(6:00010000)$, $6'(6:00100000)$, $7'(6:01000000)$, $8'(6:10000000)$, $9'(7:00000001)$, $10'(7:00000010)$, . . . , $16'(7:10000000)$, $17'(8:00000001)$, $18'(8:00000010)$, . . . , $20'(8:00001000)$.

For a link between two nodes X and Y, there are two BitPositions for two forward connected adjacencies. These two forward connected adjacency BitPositions are assigned on nodes X and Y respectively. The BitPosition assigned on X is the forward connected adjacency of Y. The BitPosition assigned on Y is the forward connected adjacency of X.

For example, for the link between nodes B and C in the figure, two forward connected adjacency BitPositions 5' and 6' are assigned to two ends of the link. BitPosition 5' is assigned on node B to B's end of the link. It is the forward connected adjacency of node C. BitPosition 6' is assigned on node C to C's end of the link. It is the forward connected adjacency of node B.

2.2. A Brief Flow of PCEP Messages for a BIER-TE Path

For a BIER-TE Path to transport the packets with a given multicast group/address and source in a BIER-TE domain, a sequence of PCEP messages are exchanged between the PCE for the domain and the PCEs for the domains containing the source, and between the PCE for the domain and the PCCs running on the BFERs/BFIRs of the domain.

Suppose that each of nodes H, D and F receives a multicast join with a same multicast group/address and source, which are MGa and MSa respectively. For simplicity, assume that the multicast source MSa is in the left domain containing the CE in Figure 1. The following is a brief flow of PCEP messages for computing and creating a BIER-TE Path to transport the packets to H, D and F.

At first, the PCC running on each of nodes H, D and F sends the PCE a PCEP message such as PCRpt. The message contains the multicast group and source (i.e., MGa and MSa), which reports to the PCE that the node receives a multicast join with MGa and MSa. Note that a PCEP message is sent to the PCE from the PCC on a node to report that the node leaves when the node receives a multicast leave with MGa and MSa.

After receiving the PCEP messages from nodes H, D and F reporting multicast join with MGa and MSa, the PCE for the domain containing these nodes determines that nodes H, D and F are the egress nodes of a BIER-TE path since they have the same multicast group and source.

Second, the PCE for the domain sends a PCEP message such as PCReq to each of the PCEs for the domains that may contain the multicast source. This message requests the PCE (that may contain the source) to find an ingress node for the BIER-TE path having egress nodes H, D and F. The message contains the multicast group and source (i.e., MGa and MSa). For example, the PCE for the BIER-TE domain sends the PCEP message to the PCE (called PCE-L) for the left domain containing CE (note that this PCE is not shown in the figure).

After receiving the PCEP message requesting to find an ingress node, the PCE (e.g., PCE-L) for the domain containing the multicast source computes the ingress node that is reachable from the source with minimum cost (e.g., ingress node A). The PCE for the domain without the source can not find any ingress node.

Third, the PCE for the domain with the source sends the PCE for the BIER-TE domain a PCEP message such as PCRep with the ingress node. The PCE for the domain without the source sends the PCE for the BIER-TE domain a PCEP message such as PCRep with NO INGRESS FOUND.

After receiving the PCEP message with the ingress node, the PCE for the BIER-TE domain computes a P2MP path from the ingress node (e.g., A) to the egress nodes (e.g., H, D and F). The path has the optimal BitPositions and satisfies the constraints. The optimal BitPositions means the BitPositions for the path has the minimum number of bit sets and the minimum bit distance.

Fourth, the PCE for the BIER-TE domain sends a PCEP message such as PCInitiate to the PCC on the ingress node (e.g., A) for the ingress to create a BIER-TE path to transport the packets for the given multicast group and source. The message contains the BitPositions for the path, the multicast group and source.

After receiving the PCEP message with the path, the PCC on the ingress (e.g., A) creates the BIER-TE path, i.e., a forwarding entry that imports the packets with the multicast group/address and source into the BIER-TE path (i.e., encapsulates the packets with a BIER-TE header having the BitPositions of the path).

And then the PCC on the ingress sends the PCE a PCEP message such as PCRpt reporting the status of the path to the PCE.

After receiving the PCEP message with the status of the path, the PCE for the domain updates the information about the path accordingly.

2.3. Procedures on Ingress

This section introduces the procedures for the ingress node of a P2MP path to get the BitPositions representing the explicit P2MP path from the ingress node to its egress nodes from the PCE.

Suppose that node A in Figure 1 wants to have an explicit P2MP path from ingress node A to egress nodes H and F. The path satisfies a set of constraints. In one case, the PCC running on ingress node A sends a request for the path to the PCE. The request contains the set of constraints, objective functions, the ingress node and the egress nodes. After receiving the request, the PCE computes an explicit P2MP path, which satisfies the constraints and is from the given ingress node to the egress nodes. While computing the path, the PCE will optimize the BitPositions of the path. That is that, for a given length of BitString, the path computed uses the minimum number of BitStrings (i.e., bit sets) and satisfies the constraints. The length is given by the value in BitStrLen field in the PCE-BIER-TE-Path-Capability sub-TLV. The PCE sends a reply with the path to the PCC. The reply contains the BitPositions representing the explicit P2MP path.

For example, assume that the explicit P2MP path computed by the PCE traverses the link/adjacency from A to B (indicated by BP 2'), the link/adjacency from B to G (indicated by BP 4') and the link/adjacency from B to C (indicated by BP 6'), the link/adjacency from G to H (indicated by BP 18'), and the link/adjacency from C to F (indicated by BP 16'). This path is represented by {2', 4', 6', 16', 18', 2, 4}, where BitPositions 2 and 4 indicate egress nodes F and H respectively. The reply sent to the PCC on node A by the PCE contains the path represented by {2', 4', 6', 16', 18', 2, 4}.

In another case, a request for a P2MP path is from a user or application. After receiving the request, the PCE finds an ingress node if no ingress is given, and computes an explicit P2MP path from the ingress node to the egress nodes and sends the path to the PCC running on the ingress node.

After receiving the P2MP path, for any packet from CE to be transported by the path, such as the packet with the multicast address, the ingress node encapsulates the packet with the BitPositions representing the path and forwards the packet according to its BIFT.

For example, when ingress node A receives the path represented by BitPositions {2', 4', 6', 16', 18', 2, 4}, it encapsulates every packet from CE with the multicast address with the BitPositions and then forwards the packet along the P2MP path according to its BIFT.

A forwards the packet to B according to the forwarding entry for BP 2' in its BIFT.

After receiving the packet from A, B forwards the packet to G and C according to the forwarding entries for BPs 4' and 6' in B's BIFT respectively. The packet received by G has path {16', 18', 2, 4}. The packet received by C has path {16', 18', 2, 4}.

After receiving the packet from B, G sends the packet to H according to the forwarding entry for BP 18' in G's BIFT.

After receiving the packet from B, C sends the packet to F according to the forwarding entry for BP 16' in C's BIFT.

Egress node H of the P2MP path receives the packet with BitPosition 4. It decapsulates the packet and pass the payload of the packet to the packet's NextProto.

Egress node F of the P2MP path receives the packet with BitPosition 2. It decapsulates the packet and pass the payload of the packet to the packet's NextProto.

3. Extensions to PCEP

This section describes extensions to PCEP.

3.1. BIER-TE Path Capability

During a PCEP session establishment, PCEP Speakers (PCE or PCC) indicate their ability to support BIER-TE paths. The OPEN object in the Open message contains the PATH-SETUP-TYPE-CAPABILITY TLV, which is defined in [[RFC8408](#)]. The TLV contains a list of Path Setup Types (PSTs) and optional sub-TLVs associated with the PSTs. The sub-TLVs convey the parameters that are associated with the PSTs supported by a PCEP speaker.

This document defines a new PST value:

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

A new sub-TLV associated with this new PST is defined, which is called PCE-BIER-TE-Path-Capability sub-TLV. The format of this new sub-TLV is illustrated in the figure below.




Figure 2: PCE-BIER-TE-Path-Capability sub-TLV

Type - 16 bits: TBD2 is to be assigned by IANA.

Length - 16 bits: 4 is the total length in bytes of the remainder of the TLV, excluding the Type and Length fields.

SILen (SI Length) - 5 bits: The length in bits of the SI field.

BitStrLen (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 BitStrLen is $\log_2(k)-5$. For example, BitStrLen = 1 indicates k = 64, BitStrLen = 7 indicates k = 4096.

Reserved - 19 bits: MUST be set to zero by the sender and MUST be ignored by the receiver.

A PCEP speaker supporting BIER-TE paths includes the new PST and sub-TLV in the PATH-SETUP-TYPE-CAPABILITY TLV.

3.2. Extensions to SRP

For a PCEP message, when it is used for a BIER-TE path, the SRP (Stateful PCE Request Parameters) object in the message MUST include the PATH-SETUP-TYPE TLV defined in [RFC8408]. The TLV MUST contain the PST = TBD1 for path setup using BIER-TE.

Three contiguous bits in SRP Object Flag Field are defined to indicate one of the assistant operations for a BIER-TE path. This three bits field is called AOP (Assistant Operations). In addition, the Multicast Flow Specification TLV defined in [I-D.ietf-pce-pcep-flowspec] is re-used in the SRP object for indicating Multicast Traffic.

3.2.1. SRP Object Flag Field

The three bits for AOP are bits 27 to 29 (the exact bits to be assigned by IANA) in the SRP Object Flag Field. The values of AOP are defined as follows:

AOP Value	Meaning (Assistant Operation)
0x001 (J):	Join with Multicast Group and Source
0x010 (L):	Leave from Multicast Group and Source
0x011 (I):	Ingress node computation

The value of AOP indicates one of the three operations above. When any of the other values is received, an error MUST be reported.

When the PCC running on an edge node of a BIER-TE domain sends the PCE for the domain a PCEP message such as PCRpt to report that the edge node receives a multicast join, the message MUST include a SRP object with AOP == 0x001 (J).

When the PCC running on an edge node of a BIER-TE domain sends the PCE for the domain a PCEP message such as PCRpt to report that the edge node receives a multicast leave, the message MUST include a SRP object with AOP == 0x010 (L).

When the PCE for the domain sends a PCEP message such as PCReq to another PCE for requesting to find an ingress node for a BIER-TE path, the message MUST include a SRP object with AOP == 0x011 (I).

3.2.2. Reuse of Multicast Flow Specification TLV

For a PCE-Initiated BIER-TE path, when a PCE sends a PCC a message such as PCInitiate message to create a BIER-TE path in a BIER-TE domain, the message MUST contain a Multicast Flow Specification TLV in the SRP object. The TLV indicates the multicast traffic that the BIER-TE path will carry.

When the PCC running on an edge node of a BIER-TE domain sends the PCE for the domain a PCEP message to report that the edge node receives a multicast join or leave with a multicast group/address and source, the message MUST contain a Multicast Flow Specification TLV in the SRP object. The TLV indicates the multicast group by the multicast group address and/or multicast source address.

When the PCE for a BIER-TE domain sends another PCE a PCEP message to request for finding an ingress node of a BIER-TE path, the message MUST contain a Multicast Flow Specification TLV in the SRP object. The TLV indicates the multicast source.

3.3. Ingress Node Object

To represent an ingress node, a new ingress node object is defined. The format of the new object for IPv4 (OT = 1) is as follows:

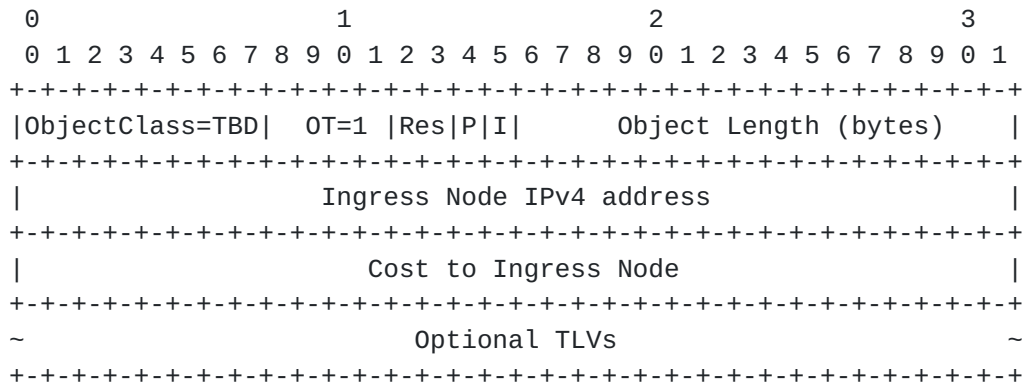


Figure 3: Ingress Node Object for IPv4

ObjectClass: TBD is to be assigned by IANA.

OT: 1 for IPv4.

Res, P, I and Object Length: Same as those defined in Common Object Header in [\[RFC5440\]](#).

Ingress Node IPv4 address: Indicates an IPv4 address of an ingress node.

Cost to Ingress Node: Indicates the cost from the multicast source to the ingress node.

No optional TLV is defined so far.

The format of the new object for IPv6 (OT = 2) is as follows:

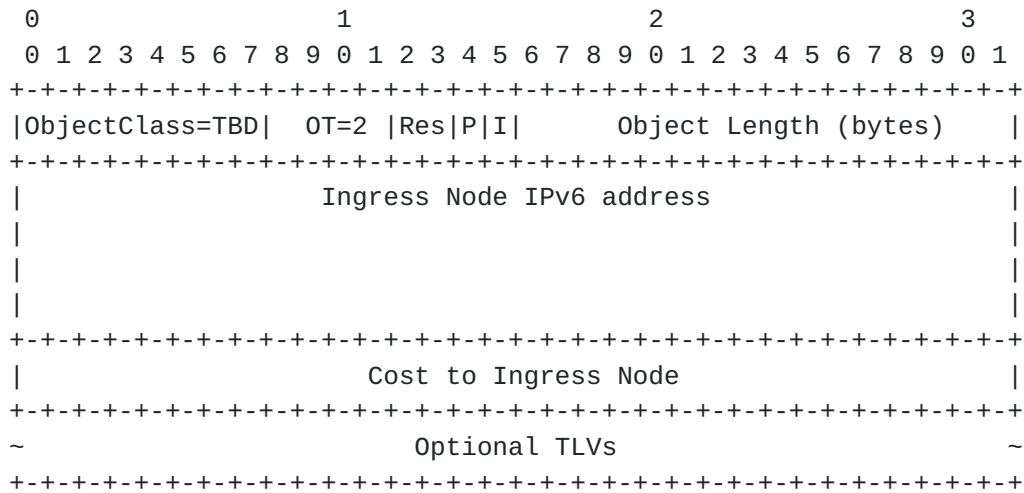


Figure 4: Ingress Node Object for IPv6

TBD, Res, P, I, Object Length, and Cost to Ingress Node:
Same as those defined in Ingress Node Object for IPv4.

OT: 2 for IPv6.

Ingress Node IPv6 address: Indicates an IPv6 address of an ingress node.

No optional TLV is defined so far.

3.4. 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, the following new OF is defined.

Objective Function Code: TBD8
Name: Minimum Bit Sets (MBS)
Description: Find a path represented by BitPositions that has the minimum number of bit sets.

Objective Function Code: TBD9
Name: Minimum Bits (MB)
Description: Find a path represented by BitPositions that has the minimum bit distance. The bit distance of BitPositions is the distance from the lowest bit to the highest bit in BitPositions.

3.5. BIER-TE Path Subobject

A BIER-TE path is represented by the BitPositions for the adjacencies through which the path traverses. A BitPosition is represented by a SI:BitString or a number.

A new subobject, called BIER-TE Path subobject (or BIER-TE-ERO subobject), is defined to contain the information about one or more BitPositions.

The format of a BIER-TE Path subobject in a ERO is shown in the figure below.




Figure 5: BIER-TE Path Subobject in ERO

L Flag (1 bit): It indicates whether the subobject represents a loose-hop in the path.

Type (7 bits): It is to be assigned by IANA. It identifies the BIER subobject type.

Length (8 bits): It contains the total length of the subobject in octets. The Length MUST be at least 4.

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.

BitPositions: It MUST be at least one BitPosition.

For the subobject in a message received from a PCEP session, the format of the BitPositions in the subobject is determined by the values of SILen and BitStrLen in the PCE-BIER-TE-Path-Capability sub-TLV exchanged during the establishment of the session. When both SILen and BitStrLen are greater than zero, each of the BitPositions has two parts SI and BitString, where SI occupies SILen bits and

BitString occupies BitStrLen bits. When both SILEn and BitStrLen are zeros, each of the BitPositions is a number of 16 bits.

For example, when SILEn = 8 and BitStrLen = 1 (indicating BitString is of 64 bits), each BitPosition has a SI of 8 bits and a BitString of 64 bits. For simplicity, BitString of 8 bits is used below. The BitPositions for a BIER-TE path are sorted in descending order before they are put into a BIER-TE path subobject. For BIER-TE path {2', 4', 6', 16', 18', 2, 4}, when its BitPositions are sorted, it is {18', 16', 6', 4', 2', 4, 2}, which is {18'(8:00000010), 16'(7:10000000), 6'(6:00100000), 4'(6:00001000), 2'(6:00000010), 4 (0:00001000), 2 (0:00000010)}. The BitPositions with the same SI are stored in one BitString. For example, 6'(6:00100000), 4'(6:00001000) and 2'(6:00000010) are stored in (SI:BitString) = (6:00101010), where SI = 6. BIER-TE path {18', 16', 6', 4', 2', 4, 2} is encoded in the BIER-TE path subobject in the figure below. The path uses four BitStrings of 8 bits.

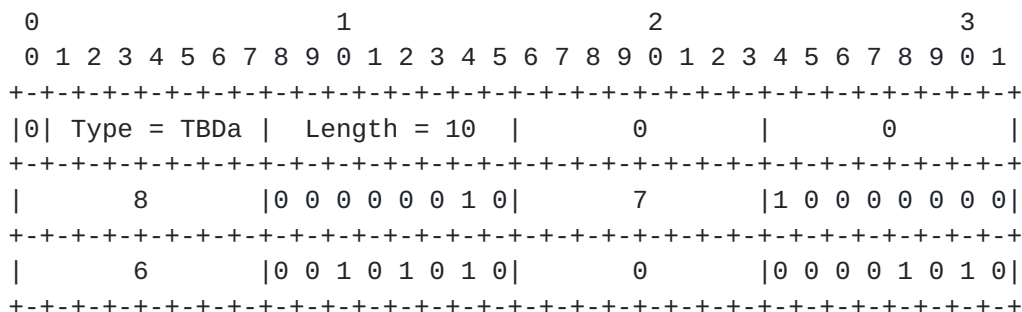


Figure 6: BIER-TE Path Subobject for a Path

3.6. BIER-TE Path Subobject in ERO

The ERO defined in [RFC5440] may contain a BIER-TE Path subobject for the BitPositions of a BIER-TE path. The BitPositions in the BIER-TE Path subobject for the BIER-TE path MUST be in descending order. When an ERO contains one or more BIER-TE Path subobjects for a BIER-TE path, the ERO MUST NOT include any other type of subobjects (i.e., it MUST include only BIER-TE Path subobjects). The first one is used and the others are ignored.

3.7. BIER-TE Path Subobject in RRO

A BIER-TE Path Subobject in a RRO (Record Route Object) has the same format as a BIER-TE Path subobject in a ERO except for L flag. The former does not have L flag. The format of a BIER-TE Path Subobject in a RRO is shown in the figure below.

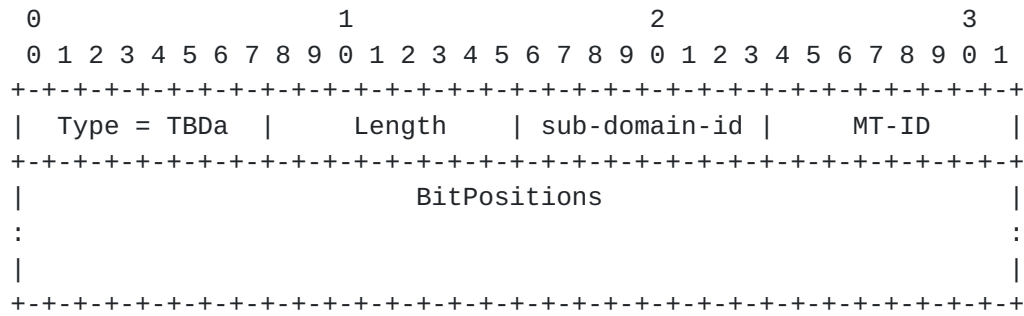


Figure 7: BIER-TE Path Subobject in RRO

A PCC may send a PCE a message such as a PCRpt message defined in [RFC8231]. The message contains a RRO with one BIER-TE Path subobject having the BitPositions for the actual BIER-TE path that is used to transport the traffic in the BIER-TE domain. The BitPositions in the BIER-TE Path subobject for the BIER-TE path MUST be in descending order.

4. Procedures

This section describes the procedures related to a BIER-TE path.

4.1. BIER-TE Path Creation

For PCC-Initiated BIER-TE path, a PCC MUST delegate the path by sending a path computation report (PCRpt) message with its demanded resources to a stateful PCE. Note the PCC MAY use the PCReq/PCRep before delegating.

Upon receiving the delegation via PCRpt message, the stateful PCE MUST compute a path based on the network resource availability stored in the TED.

The stateful PCE will send a PCUpd message for the BIER-TE path to the PCC. The stateful PCE MUST update its local LSP-DB and TED and would need to synchronize the information with other PCEs in the domain.

For PCE-Initiated BIER-TE path, the stateful PCE MUST compute a BIER-TE path per request from network management systems or applications automatically based on the network resource availability in the TED and send a PCInitiate message with the path information to the PCC. After receiving the PCInitiate message, the PCC creates the BIER-TE path.

For both PCC-Initiated and PCE-Initiated BIER-TE paths:

- o The stateful PCE MUST update its local LSP-DB and TED with the paths.
- o Upon receiving the PCUpd message or PCInitiate message for the path from the PCE with a found path, the PCC determines that it is a BIER-TE path by the PST = TBD1 for path setup using BIER-TE in the PATH-SETUP-TYPE TLV of the SRP object in the message.

4.2. BIER-TE Path Update

After a BIER-TE path is created in a BIER-TE domain, when some network events such as a node failure happen on the path (called old path) or a leaf/egress joins/leaves, the PCE computes a new BIER-TE path and replaces the old path with the new path. The new path satisfies the same constraints as the old path.

The PCE sends a PCUpd message to the PCC running on the ingress node. The message contains the information about the new BIER-TE path. After receiving the message, the PCC overwrites (or replaces) the BIER-TE path with the new BIER-TE path.

4.3. BIER-TE Path Deletion

For a BIER-TE path that has been created in a BIER-TE domain, after receiving a request for deleting the path from a user or application, the PCE MUST send a PCInitiate or PCUpd message to the PCC running on the ingress node of the path to remove the path.

5. The PCEP Messages

5.1. The PCRpt Message

The Path Computation State Report (PCRpt) message is a PCEP message sent by a PCC to a PCE to report the status of one or more LSPs, as per [[RFC8281](#)]. Each LSP State Report in a PCRpt message contains the actual LSP's path, bandwidth, operational and administrative status, etc. An LSP Status Report carried in a PCRpt message is also used in delegation or revocation of control of an LSP to/from a PCE.

In the case of a BIER-TE path, a PATH-SETUP-TYPE TLV with PST = TBD1 for path setup using BIER-TE MUST be carried in the SRP object in the PCRpt message. A BIER-TE path in the message is represented by a BIER-TE path subobject.

In addition, a PCRpt message is sent from the PCC running on an edge node to the PCE to report that the edge node as leaf/egress joins/leaves to/from a multicast group and source.

5.2. The PCUpd Message

The Path Computation Update Request (PCUpd) message is a PCEP message sent by a PCE to a PCC to update LSP parameters on one or more LSPs, as per [RFC8281]. In the case of a BIER-TE path, a PATH-SETUP-TYPE TLV with PST = TBD1 for path setup using BIER-TE MUST be carried in the SRP object in the PCUpd message. Each BIER-TE path Update Request in a PCUpd message contains all parameters that a PCE wishes to be set for a given BIER-TE path. A BIER-TE path in the message is represented by a BIER-TE path subobject.

5.3. The PCInitiate Message

The LSP Initiate Request (PCInitiate) message is a PCEP message sent by a PCE to a PCC to trigger LSP instantiation or deletion, as per [RFC8281]. In the case of a BIER-TE path, a PATH-SETUP-TYPE TLV with PST = TBD1 for path setup using BIER-TE MUST be carried in the SRP object in the PCInitiate message. A BIER-TE path in the message is represented by a BIER-TE path subobject. The multicast packets to be transported by the BIER-TE path is specified by the Multicast Flow Specification TLV included in the SRP object.

5.4. The PCReq Message

The Path Computation Request (PCReq) message is a PCEP message sent by a PCC to a PCE to request a path computation [RFC5440], and it may contain the LSP object [RFC8231] to identify the LSP for which the path computation is requested. In the case of a BIER-TE path, a PATH-SETUP-TYPE TLV with PST = TBD1 for path setup using BIER-TE MUST be carried in the SRP object in the PCReq message.

In addition, a PCReq message is sent from the PCE (as a PCC) for the BIER-TE domain to another PCE for the domain that may contain the multicast source for a BIER-TE path in order to find an ingress node for the BIER-TE path.

5.5. The PCRep Message

The Path Computation Reply (PCRep) message is a PCEP message sent by a PCE to a PCC in reply to a path computation request [RFC5440], and it may contain the LSP object [RFC8231] to identify the LSP for which the path is computed. A PCRep message can contain either a set of computed paths if the request can be satisfied or a negative reply if not. A negative reply may indicate the reason why no path could be found. In the case of a BIER-TE path, a PATH-SETUP-TYPE TLV with PST = TBD1 for path setup using BIER-TE MUST be carried in the SRP object in the PCRep message. Each of the computed paths in the message is represented by a BIER-TE path subobject.

In addition, a PCRep message is sent from the PCE for the domain that may contain the multicast source for a BIER-TE path to the PCC (i.e., the PCE for the BIER-TE domain) in response to the request for finding an ingress node for the BIER-TE path. A PCRep message can contain either a set of ingress nodes represented by ingress node objects if the request can be satisfied or a negative reply if not.

6. IANA Considerations

6.1. PST for BIER-TE Path

IANA is requested to allocate a new code point within registry "PCEP Path Setup Types" under "Path Computation Element Protocol (PCEP) Numbers" as follows:

Value	Description	Reference
TBD1 (2)	Path is setup using BIER-TE	This document

6.2. PCE-BIER-TE-Path Capability sub-TLV

IANA is requested to allocate a new code point within registry "PATH-SETUP-TYPE-CAPABILITY Sub-TLV Type Indicators" under "Path Computation Element Protocol (PCEP) Numbers" as follows:

Value	Meaning	Reference
TBD2 (1)	PCE-BIER-TE-Path Capability	This document

6.3. SRP Object Flag Field

IANA is requested to allocate the following bits in the "SRP Object Flag Field" subregistry under the "Path Computation Element Protocol (PCEP) Numbers" registry:

Value	Description	Reference
27-29	Assistant Operations for Path	This document

6.4. Ingress Node Object

IANA is requested to allocate the following Object-Class Value in the "PCEP Objects" subregistry under the "Path Computation Element Protocol (PCEP) Numbers" registry:

Object-Class Value	Name	Object-Type	Reference
TBD (45)	INGRESS	0: Reserved	This document
		1: IPv4 Address	This document
		2: IPv6 Address	This document
		3-15: Unassigned	

6.5. OF Code Points

IANA is requested to allocate the following Objective Function Code Points in the "Objective Function" subregistry under the "Path Computation Element Protocol (PCEP) Numbers" registry:

Code Point	Name	Reference
TBD8 (18)	Minimum Bit Sets (MBS)	This document
TBD9 (19)	Minimum Bit Distance (MBD)	This document

6.6. PCEP BIER-TE Path Subobjects

This document defines a new subobject, called PCE BIER-TE Path (or BIER-TE-ERO) subobject, for PCEP ERO object. It also defines a new subobject, called PCE BIER-TE Path (or BIER-TE-RR0) subobject, for PCEP RR0 object. The code points of the subobjects for the objects are maintained under ERO and RR0 objects in the RSVP Parameters registry.

IANA is requested to allocate a code point under "Subobject type - 20 EXPLICIT_ROUTE - Type 1 Explicit Route" within registry "Resource Reservation Protocol (RSVP) Parameters" for PCEP BIER-TE path subobject as follows:

Value	Name	Reference
TBDa (63)	PCEP BIER-TE Path	This document

IANA is requested to allocate a code point under "Subobject type - 21 ROUTE_RECORD - Type 1 Explicit Route" within registry "Resource Reservation Protocol (RSVP) Parameters" for PCEP BIER-TE path subobject as follows:

Value	Name	Reference
TBDa (63)	PCEP BIER-TE Path	This document

7. Security Considerations

TBD

8. Acknowledgements

The authors would like to thank Dhruv Dhody, and others for their comments to this work.

9. References

9.1. Normative References

[I-D.ietf-bier-te-arch] Eckert, T., Cauchie, G., and M. Menth, "Tree Engineering for Bit Index Explicit Replication (BIER-TE)", [draft-ietf-bier-te-arch-09](#) (work in progress), October 2020.

[I-D.ietf-pce-pcep-flowspec] Dhody, D., Farrel, A., and Z. Li, "PCEP Extension for Flow Specification", [draft-ietf-pce-pcep-flowspec-12](#) (work in progress), October 2020.

[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>>.

[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, J.L., Vasseur, J.P., and Y. Lee, "Encoding of Objective Functions in the Path Computation Element Communication Protocol (PCE)", [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 (PCE) Extensions for Stateful PCE", [RFC 8231](#), DOI 10.17487/RFC8231, September 2017, <<https://www.rfc-editor.org/info/rfc8231>>.
- [RFC8232] Crabbe, E., Minei, I., Medved, J., Varga, R., Zhang, X., and D. Dhody, "Optimizations of Label Switched Path State Synchronization Procedures for a Stateful PCE", [RFC 8232](#), DOI 10.17487/RFC8232, September 2017, <<https://www.rfc-editor.org/info/rfc8232>>.
- [RFC8279] Wijnands, I.J., 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 (PCE) 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, I.J., 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>>.
- [RFC8408] Sivabalan, S., Tantsura, J., Minei, I., Varga, R., and J. Hardwick, "Conveying Path Setup Type in PCE Communication Protocol (PCE) Messages", [RFC 8408](#), DOI 10.17487/RFC8408, July 2018, <<https://www.rfc-editor.org/info/rfc8408>>.

9.2. Informative References

- [I-D.chen-pce-bier]
Chen, R., Zhang, Z., Dhanaraj, S., and F. Qin, "PCEP Extensions for BIER-TE", [draft-chen-pce-bier-08](#) (work in progress), November 2020.

[RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", [RFC 8402](#), DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.

Authors' Addresses

Huaimo Chen
Futurewei
Boston, MA
USA

Email: Huaimo.chen@futurewei.com

Mike McBride
Futurewei

Email: michael.mcbride@futurewei.com

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

Email: wangaj3@chinatelecom.cn

Gyan S. Mishra
Verizon Inc.
13101 Columbia Pike
Silver Spring MD 20904
USA

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

Yisong Liu
China Mobile

Email: liuyisong@chinamobile.com

Yanhe Fan
Casa Systems
USA

Email: yfan@casa-systems.com

Lei Liu
Fujitsu

USA

Email: liulei.kddi@gmail.com

Xufeng Liu
Volta Networks

McLean, VA
USA

Email: xufeng.liu.ietf@gmail.com

