

Networking Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 15, 2016

Ran. Chen
Zheng. Zhang
ZTE Corporation
October 13, 2015

PCEP Extensions for BIER
draft-chen-pce-bier-00

Abstract

Bit Index Explicit Replication (BIER) is an architecture that provides optimal multicast forwarding through a "BIER domain" without requiring intermediate routers to maintain any multicast related per-flow state. BIER forwards and replicates packets based on a BitString in the packet header. A BIER Path can be derived from a variety of mechanisms, including an IGP Shortest Path Tree (SPT), explicit configuration, or a Path Computation Element (PCE).

This document specifies extensions to the Path Computation Element Protocol (PCEP) to handle requests and responses for the computation of paths for BIER TE LSPs.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

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

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

This Internet-Draft will expire on April 15, 2016.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of

publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- [1. Introduction](#) [2](#)
- [2. Conventions used in this document](#) [3](#)
- [3. Overview of PCEP Operation in BIER Networks](#) [3](#)
- [4. BIER PCEP Message Extensions](#) [3](#)
 - [4.1. BIER Capability Advertisement](#) [3](#)
 - [4.1.1. The OPEN Object](#) [3](#)
 - [4.1.1.1. The BIER PCE Capability TLV](#) [3](#)
 - [4.2. Path Computation Request/Reply Message Extensions](#) [4](#)
 - [4.2.1. The RP/SPR Object](#) [4](#)
 - [4.2.2. The New BIER END-POINT Object](#) [5](#)
 - [4.2.3. ERO Object](#) [5](#)
 - [4.2.3.1. BIER-ERO Subobject](#) [5](#)
 - [4.2.4. RRO Object](#) [7](#)
 - [4.2.4.1. RRO Processing](#) [7](#)
- [5. Security Considerations](#) [7](#)
- [6. IANA Considerations](#) [7](#)
 - [6.1. PCEP Objects](#) [7](#)
 - [6.2. PCEP-Error Objects and Types](#) [8](#)
 - [6.3. PCEP TLV Type Indicators](#) [8](#)
 - [6.4. New Path Setup Type](#) [8](#)
- [7. Normative references](#) [9](#)
- [Authors' Addresses](#) [9](#)

1. Introduction

Bit Index Explicit Replication (BIER) is an architecture that provides optimal multicast forwarding through a "BIER domain" without requiring intermediate routers to maintain any multicast related per-flow state. BIER forwards and replicates packets based on a BitString in the packet header. A BIER Path can be derived from a variety of mechanisms, including an IGP Shortest Path Tree (SPT), explicit configuration, or a Path Computation Element (PCE).

This document specifies extensions to the Path Computation Element Protocol (PCEP) to handle requests and responses for the computation of paths for BIER TE LSPs.

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 forwards and replicates packets based on a BitString in the packet header. In a PCEP session, An ERO object specified in [[RFC5440](#)] carrying a BIER-TE path consists of one or more BIER-subobject(s). BIER-TE LSPs computed by a PCE can be represented in one of the following forms:

- o An ordered set of BitString(s) in which each bit represents exactly one router that the BIER-TE paths pass through in the domain.
- o An ordered set of BFR-id(s). In this case, the PCC needs to convert the BFR-ID into the corresponding BitString(s).
- o An ordered set of BFR-prefix(es). In this case, the PCC needs to convert the BFR-ID into the corresponding B BitString(s).

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

4. BIER PCEP Message Extensions

The following section describes the protocol extensions required to support BIER-TE path.

4.1. BIER Capability Advertisement

4.1.1. The OPEN Object

This document defines a new optional TLV for use in the OPEN Object.

4.1.1.1. The BIER PCE Capability TLV

The BIER-PCE-CAPABILITY TLV is an optional TLV associated with the OPEN Object to exchange BIER capability of PCEP speakers. The format of the BIER-PCE-CAPABILITY TLV is shown in the following figure:




Figure 1

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

The "Reserved" (2 octets) and "Flags" (2 octet) fields are currently unused, and MUST be set to zero on transmission and ignored on reception.

4.1.1.1.1. Exchanging BIER Capability

This document defines a new optional BIER-PCE-CAPABILITY TLV for use in the OPEN object to negotiate the BIER capability. The inclusion of this TLV in the OPEN message destined to a PCC indicates the PCE's capability to perform BIER-TE path computations, and the inclusion of this TLV in the OPEN message destined to a PCE indicates the PCC's capability to support BIER-TE Path.

A PCE that is able to support the BIER extensions defined in this document SHOULD include the BIER-PCE-CAPABILITY TLV on the OPEN message. If the PCE does not include the BIER-PCE-CAPABILITY TLV in the OPEN message and PCC does include the TLV, it is RECOMMENDED that the PCC indicates a mismatch of capabilities.

4.2. Path Computation Request/Reply Message Extensions

4.2.1. The RP/SPR Object

In order to setup an BIER-TE LSP, a new PATH-SETUP-TYPE TLV([I-D.ietf-pce-lsp-setup-type]) MUST be contained in RP or SRP object. This document defines a new Path Setup Type (PST) for BIER as follows:

- o PST = 2: Path is setup using BIER Traffic Engineering technique.

If a PCEP speaker does not recognize the PATH-SETUP-TYPE TLV, it MUST ignore the TLV in accordance with [RFC5440]. If a PCEP speaker recognizes the TLV but does not support the TLV, it MUST send PCErr with Error-Type = 2 (Capability not supported).

4.2.2. The New BIER END-POINT Object

The END-POINTS object 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 define a new END-POINT Object for the BIER path:

The format of the new END-POINTS Object is as follows:

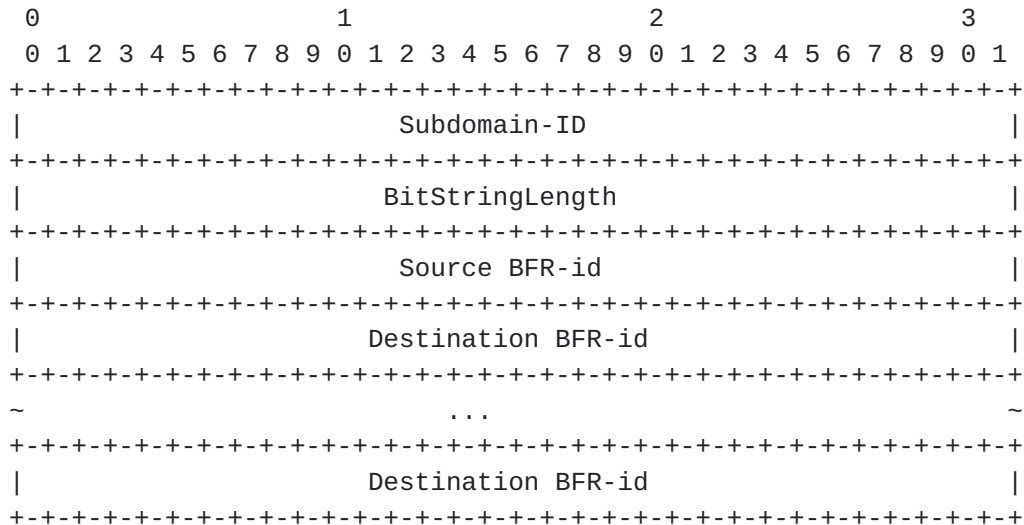


Figure 2

4.2.3. ERO Object

An BIER-TE path consists of one or more BitString/BFR-id/BFR-prefix where each BitString/BFR-id/BFR-prefix MAY be associated with the identifier that represents the node.

The ERO object specified in [RFC5440] is used to carry BIER-TE path information. In order to carry BitString/BFR-id/BFR-prefix, this document defines three new ERO subobjects referred to as "BIER-ERO subobjects" whose formats are specified in the following section. An ERO object carrying a BIER-TE path consists of one or more BIER-ERO subobject(s).

4.2.3.1. BIER-ERO Subobject




Figure 3




Figure 4




Figure 5

4.2.3.1.1. ERO Processing

If a PCC finds a non-recognize the SR-ERO subobject, the PCC MUST respond with a PCErr message with Error-Type=3 ("Unknown Object") and Error-Value=2 ("Unrecognized object Type") or Error-Type=4 ("Not supported object") and Error-Value=2 ("Not supported object Type") as described in [RFC5440] .

If a PCC receives an BIER-ERO subobject format 1 in which either BitStringLength or BitString is absent, it MUST consider the entire

ERO object invalid and send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-Value = TBD ("BitStringLength is absent ") and Error-Value = TBD ("BitString is absent ")

If a PCC detects that all subobjects of ERO are not identical, it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-Value = TBD ("Non-identical ERO subobjects"), as described in ([I-D.ietf-pce-segment-routing])[pce-].

If a PCC receives an SR-ERO subobject in which BFR-ID=0, it MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-Value = TBD ("Invalid BFR-ID").

If a PCC receives an SR-ERO subobject in which BitStringLength values are not chosen from: 64, 128, 256, 512, 1024, 2048, and 4096, as it described in ([I-D.ietf-bier-architecture]). The PCC MUST send a PCErr message with Error-Type = 10 ("Reception of an invalid object") and Error-Value = TBD ("Invalid BitStringLength").

4.2.4. RRO Object

A PCC can record BIER-TE LSP and report the LSP to a PCE via RRO. An RRO object contains one or more subobjects called "BIER-RRO subobjects" whose formats are the same as that of SR-ERO subobject.

4.2.4.1. RRO Processing

Processing rules of SR-RRO subobject are identical to those of SR-ERO subobject defined in [section 4.2.3.1.1](#) in this document.

5. Security Considerations

TBD.

6. IANA Considerations

6.1. PCEP Objects

As discussed in [Section 4.2.2](#), a new END-POINTS Object-Type is defined. IANA has made the following Object-Type allocations from the "PCEP Objects" sub-registry:

| Object | Object-Class Value |
|-----------------------|--------------------|
| ----- | ----- |
| BIER END-POINT Object | TBD |

As discussed in [Section 4.2.3](#) and 4.2.4, a new sub-object type for the PCEP explicit route object (ERO), and a new sub-object type for the PCEP record route object (RRO) are defined.

IANA has made the following sub-objects allocation from the RSVP Parameters registry:

| Object | Sub-Object | Sub-Object Type |
|----------------|--------------------------|-----------------|
| EXPLICIT_ROUTE | BIER-ERO (PCEP-specific) | TBD |
| ROUTE_RECORD | BIER-RRO (PCEP-specific) | TBD |

6.2. PCEP-Error Objects and Types

As described in [Section 4.2.3.1.1](#), a number of new PCEP-ERROR Object Error Values have been defined.

| Reference | Error-Type | Meaning |
|-------------------------|------------|--|
| | 10 | Reception of an invalid object. |
| RFC5540 | | Error-value = TBD: BitStringLength is absent |
| This document | | Error-value = TBD: BitString is absent |
| This document | | Error-value = TBD: invalid BFR-ID |
| This document | | Error-value = TBD: Invalid BitStringLength |

6.3. PCEP TLV Type Indicators

IANA is requested to allocate a new code point in the PCEP TLV Type Indicators registry, as follows:

| Value | Meaning | Reference |
|-------|-------------------------|---------------|
| TBD | BIER-PCE-CAPABILITY TLV | This document |

6.4. New Path Setup Type

IANA is requested to allocate a new code point in the PCEP PATH_SETUP_TYPE TLV PST field registry, as follows:

| Value | Description | Reference |
|-------|--|---------------|
| 2 | Path is setup using BIER Traffic Engineering technique | This document |

7. Normative references

[I-D.ietf-bier-architecture]

Wijnands, I., Rosen, E., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast using Bit Index Explicit Replication", [draft-ietf-bier-architecture-02](#) (work in progress), July 2015.

[I-D.ietf-pce-lsp-setup-type]

Sivabalan, S., Medved, J., Minei, I., Crabbe, E., Varga, R., Tantsura, J., and J. Hardwick, "Conveying path setup type in PCEP messages", [draft-ietf-pce-lsp-setup-type-03](#) (work in progress), June 2015.

[I-D.ietf-pce-segment-routing]

Sivabalan, S., Medved, J., Filsfils, C., Crabbe, E., Lopez, V., Tantsura, J., Henderickx, W., and J. Hardwick, "PCEP Extensions for Segment Routing", [draft-ietf-pce-segment-routing-06](#) (work in progress), August 2015.

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

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

Authors' Addresses

Ran Chen
 ZTE Corporation
 No.50 Software Avenue, Yuhuatai District
 Nanjing, Jiangsu Province 210012
 China

Phone: +86 025 88014636
 Email: chen.ran@zte.com.cn

Zheng Zhang
ZTE Corporation
No.50 Software Avenue, Yuhuatai District
Nanjing, Jiangsu Province 210012
China

Email: zhang.zheng@zte.com.cn