

OSPF
Internet-Draft
Intended status: Standards Track
Expires: March 18, 2017

P. Psenak, Ed.
N. Kumar
IJ. Wijnands
Cisco
A. Dolganow
Alcatel-Lucent
T. Przygienda
J. Zhang
Juniper Networks, Inc.
S. Aldrin
Google, Inc.
September 14, 2016

OSPF Extensions for BIER
draft-ietf-bier-ospf-bier-extensions-03.txt

Abstract

Bit Index Explicit Replication (BIER) is an architecture that provides multicast forwarding through a "BIER domain" without requiring intermediate routers to maintain multicast related per-flow state. Neither does BIER require an explicit tree-building protocol for its operation. A multicast data packet enters a BIER domain at a "Bit-Forwarding Ingress Router" (BFIR), and leaves the BIER domain at one or more "Bit-Forwarding Egress Routers" (BFERs). The BFIR router adds a BIER header to the packet. Such header contains a bit-string in which each bit represents exactly one BFER to forward the packet to. The set of BFERs to which the multicast packet needs to be forwarded is expressed by the according set of bits switched on in BIER packet header.

This document describes the OSPF protocol extension required for BIER with MPLS encapsulation.

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 March 18, 2017.

Copyright Notice

Copyright (c) 2016 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. | Flooding of the BIER Information in OSPF | 3 |
| 2.1. | The BIER Sub-TLV | 3 |
| 2.2. | The BIER MPLS Encapsulation Sub-TLV | 4 |
| 2.3. | Optional BIER Tree Type Sub-TLV | 5 |
| 2.4. | Flooding scope of BIER Information | 6 |
| 3. | Security Considerations | 7 |
| 4. | IANA Considerations | 7 |
| 5. | Acknowledgments | 7 |
| 6. | Normative References | 7 |
| | Authors' Addresses | 8 |

[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. Neither does BIER explicitly require a tree-building protocol for its operation. A multicast data packet enters a BIER domain at a "Bit-Forwarding Ingress Router" (BFIR), and leaves the BIER domain at one or more "Bit-Forwarding Egress Routers" (BFERs). The BFIR router adds a BIER header to the packet. The BIER header contains a bit-string in which each bit represents exactly one BFER to forward the packet to. The set of BFERs to which the multicast packet needs to be forwarded is expressed by setting the bits that correspond to those routers in the BIER header.

BIER architecture requires routers participating in BIER to exchange BIER related information within a given domain. BIER architecture permits link-state routing protocols to perform distribution of such information. This document describes extensions to OSPF necessary to carry BIER specific information in the case where BIER uses MPLS encapsulation as described in [[I-D.wijnands-mpls-bier-encapsulation](#)].

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](#)].

2. Flooding of the BIER Information in OSPF

All BIER specific information that a BFR needs to advertise to other BFRs is associated with a BFR-Prefix. BFR prefix is a unique (within a given BIER domain), routable IP address that is assigned to each BFR as described in more detail in section 2 of [[I-D.wijnands-bier-architecture](#)].

Given that BIER information must be associated with a BFR prefix, the OSPF Extended Prefix Opaque LSA [[I-D.ietf-ospf-prefix-link-attr](#)] has been chosen to flood it.

2.1. The BIER Sub-TLV

A new Sub-TLV of the Extended Prefix TLV (defined in [[I-D.ietf-ospf-prefix-link-attr](#)]) is defined for distributing BIER information. The new Sub-TLV is called BIER Sub-TLV. Multiple BIER Sub-TLVs may be included in the Extended Prefix TLV.

BIER Sub-TLV has the following format:

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Type               |               Length               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Sub-domain-ID |      MT-ID      |               BFR-id               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               Sub-TLVs (variable)               |
+-+-----+
|
```

Type: TBD

Length: variable

Sub-domain-ID: Unique value identifying the BIER sub-domain within the BIER domain, as described in section 1 of [\[I-D.wijnands-bier-architecture\]](#).

MT-ID: Multi-Topology ID (as defined in [\[RFC4915\]](#)) that identifies the topology that is associated with the BIER sub-domain.

BFR-id: A 2 octet field encoding the BFR-id, as documented in [section 2 \[I-D.wijnands-bier-architecture\]](#). If the BFR is not locally configured with a valid BFR-id, the value of this field is set to invalid BFR-id per [\[I-D.wijnands-bier-architecture\]](#).

Each BFR sub-domain MUST be associated with one and only one OSPF topology that is identified by the MT-ID. If the association between BIER sub-domain and OSPF topology advertised in the BIER sub-TLV by other BFRs is in conflict with the association locally configured on the receiving router, whole BIER sub-TLV of the advertising routers MUST be ignored.

If a BFR advertises the same Sub-domain-ID in multiple BIER sub-TLVs, the BRF MUST be treated as if it did not advertise a BIER sub-TLV for such sub-domain.

All BFRs MUST detect advertisement of duplicate valid BFR-IDs for a given MT-ID and Sub-domain-ID. When such duplication is detected all BFRs advertising duplicates MUST be treated as if they did not advertise a valid BFR-id.

[2.2.](#) The BIER MPLS Encapsulation Sub-TLV

BIER MPLS Encapsulation Sub-TLV is a sub-TLV of the BIER Sub-TLV. BIER MPLS Encapsulation Sub-TLV is used in order to advertise MPLS specific information used for BIER. It MAY appear multiple times in the BIER Sub-TLV.

BIER MPLS Encapsulation Sub-TLV has the following format:

```

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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               |                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Lbl Range Size |                               Label Range Base |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| BS Length |                               Reserved |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Type: TBD

Length: 4 bytes

Label Range Size: A 1 octet field encoding the label range size of the label range. It MUST be greater than 0, otherwise the advertising router MUST be treated as if it did not advertise a BIER sub-TLV.

Label Range Base: A 3 octet field, where the 20 rightmost bits represent the first label in the label range.

BS Length: A 1 octet field encoding the supported BitString length associated with this BFR-prefix. The values allowed in this field are specified in section 3 of [\[I-D.wijnands-mpls-bier-encapsulation\]](#).

The "label range" is the set of labels beginning with the label range base and ending with (label range base)+(label range size)-1. A unique label range is allocated for each BitStream length and Sub-domain-ID. These labels are used for BIER forwarding as described in [\[I-D.wijnands-bier-architecture\]](#) and [\[I-D.wijnands-mpls-bier-encapsulation\]](#).

The size of the label range is determined by the number of Set Identifiers (SI) (section 2 of [\[I-D.wijnands-bier-architecture\]](#)) that are used in the network. Each SI maps to a single label in the label range. The first label is for SI=0, the second label is for SI=1, etc.

If same BS length is repeated in multiple BIER MPLS Encapsulation Sub-TLV inside the same BIER Sub-TLV, the advertising router MUST be treated as if it did not advertise a BIER sub-TLV.

Label ranges within all BIER MPLS Encapsulation Sub-TLV inside the same BIER Sub-TLV MUST NOT overlap. If the overlap is detected, the advertising router MUST be treated as if it did not advertise a BIER sub-TLV.

All advertised labels MUST be valid, otherwise the advertising router MUST be treated as if it did not advertise a BIER sub-TLV.

[2.3.](#) Optional BIER Tree Type Sub-TLV

This Sub-TLV carries the information associated with the supported BIER tree type for a subdomain. This Sub-TLV is optional and its absence has the same semantics as its presence with Tree Type value 0 (SPF). When Tree Type 0 is used it is recommended that this Sub-TLV is omitted in order to reduce the space consumed in the parent TLV.

This Sub-TLV MAY occur no more than once in a BIER sub-TLV. If multiple occurrences of this Sub-TLV are present in a single BIER Sub-TLV the advertising router MUST be treated as if it did not advertise a BIER sub-TLV.

If the tree type (implied or explicitly advertised) does not match the locally configured tree type associated with the matching subdomain the advertising router MUST be treated as if it did not advertise a BIER sub-TLV.

```

      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
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|   Type           |   Length           |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| Tree Type        |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Type: value of 1 indicating BIER Tree Type.

Length: 1 octet.

Tree Type: 1 octet

2.4. Flooding scope of BIER Information

Flooding scope of the OSPF Extended Prefix Opaque LSA [[I-D.ietf-ospf-prefix-link-attr](#)] that is used for advertising BIER Sub TLV is set to area. To allow BIER deployment in a multi-area environment, OSPF must propagate BIER information between areas. The following procedure is used in order to propagate BIER related information between areas:

When an OSPF ABR advertises a Type-3 Summary LSA from an intra-area or inter-area prefix to all its connected areas, it will also originate an Extended Prefix Opaque LSA, as described in [[I-D.ietf-ospf-prefix-link-attr](#)]. The flooding scope of the Extended Prefix Opaque LSA type will be set to area-scope. The route-type in the OSPF Extended Prefix TLV is set to inter-area. When determining whether a BIER Sub-TLV should be included in this LSA ABR will:

- look at its best path to the prefix in the source area and find the advertising router associated with the best path to that prefix.
- determine if such advertising router advertised a BIER Sub-TLV for the prefix. If yes, ABR will copy the information from

such BIER MPLS Sub-TLV when advertising BIER MPLS Sub-TLV to each connected area.

3. Security Considerations

Implementations must assure that malformed TLV and Sub-TLV permutations do not result in errors which cause hard OSPF failures.

4. IANA Considerations

The document requests three new allocations from the OSPF Extended Prefix sub-TLV registry as defined in [\[I-D.ietf-ospf-prefix-link-attr\]](#).

BIER Sub-TLV: TBD

BIER MPLS Encapsulation Sub-TLV: TBD

BIER Tree Type Sub-TLV: TBD

5. Acknowledgments

The authors would like to thank Rajiv Asati, Christian Martin, Greg Shepherd and Eric Rosen for their contribution.

6. Normative References

[I-D.ietf-ospf-prefix-link-attr]

Psenak, P., Gredler, H., rjs@rob.sh, r., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute Advertisement", [draft-ietf-ospf-prefix-link-attr-13](#) (work in progress), August 2015.

[I-D.wijnands-bier-architecture]

Wijnands, I., Rosen, E., Dolganow, A., and T. Przygienda, "Multicast using Bit Index Explicit Replication", [draft-wijnands-bier-architecture-00](#) (work in progress), September 2014.

[I-D.wijnands-mpls-bier-encapsulation]

Wijnands, I., Rosen, E., Dolganow, A., and J. Tantsura, "Encapsulation for Bit Index Explicit Replication in MPLS Networks", [draft-wijnands-mpls-bier-encapsulation-00](#) (work in progress), September 2014.

[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>>.

[RFC4915] Psenak, P., Mirtorabi, S., Roy, A., Nguyen, L., and P. Pillay-Esnault, "Multi-Topology (MT) Routing in OSPF", [RFC 4915](#), DOI 10.17487/RFC4915, June 2007, <<http://www.rfc-editor.org/info/rfc4915>>.

Authors' Addresses

Peter Psenak (editor)
Cisco
Apollo Business Center
Mlynske nivy 43
Bratislava 821 09
Slovakia

Email: ppsenak@cisco.com

Nagendra Kumar
Cisco
7200 Kit Creek Road
Research Triangle Park, NC 27709
US

Email: naikumar@cisco.com

IJsbrand Wijnands
Cisco
De Kleetlaan 6a
Diegem 1831
Belgium

Email: ice@cisco.com

Andrew Dolganow
Alcatel-Lucent
600 March Rd.
Ottawa, Ontario K2K 2E6
Canada

Email: andrew.dolganow@alcatel-lucent.com

Tony Przygienda
Juniper Networks, Inc.
10 Technology Park Drive
Westford, MA 01886
USA

Email: prz@juniper.net

Jeffrey Zhang
Juniper Networks, Inc.
10 Technology Park Drive
Westford, MA 01886
USA

Email: zzhang@juniper.net

Sam Aldrin
Google, Inc.
1600 Amphitheatre Parkway
Mountain View, CA
USA

Email: aldrin.ietf@gmail.com

