

OSPF
Internet-Draft
Intended status: Standards Track
Expires: April 6, 2018

P. Psenak, Ed.
N. Kumar
IJ. Wijnands
Cisco
A. Dolganow
Nokia
T. Przygienda
J. Zhang
Juniper Networks, Inc.
S. Aldrin
Google, Inc.
October 3, 2017

OSPF Extensions for BIER
draft-ietf-bier-ospf-bier-extensions-08.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 set 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 <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 April 6, 2018.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](https://trustee.ietf.org/license-info) 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	2
2.	Flooding of the BIER Information in OSPF	3
2.1.	BIER Sub-TLV	3
2.2.	BIER MPLS Encapsulation Sub-TLV	4
2.3.	Optional BIER Tree Type Sub-TLV	5
2.4.	Optional BIER sub-domain BSL conversion Sub-TLV	6
2.5.	Flooding scope of BIER Information	7
3.	Security Considerations	8
4.	IANA Considerations	8
4.1.	BIER Tree Type Registry	8
5.	Acknowledgments	9
6.	Normative References	9
	Authors' Addresses	10

[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

Type: TBD1

[illegible]

Type: TBD2

Length: 4 octets

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. The 4 leftmost bits MUST be ignored.

Bit String Length: A 4 bits field encoding the supported BitString length associated with this BFR-prefix. The values allowed in this field are specified in section 2 of [\[I-D.ietf-bier-mpls-encapsulation\]](#).

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 [\[I-D.ietf-bier-architecture\]](#) and [\[I-D.ietf-bier-mpls-encapsulation\]](#).

The size of the label range is determined by the number of Set Identifiers (SI) (section 1 of [\[I-D.ietf-bier-architecture\]](#)) 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 BIER sub-TLV MUST be ignored.

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 BIER sub-TLV MUST be ignored.

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

The BIER Tree Type Sub-TLV is a Sub-TLV of the BIER Sub-TLV. This Sub-TLV carries the information associated with the supported BIER tree type for a sub-domain. This Sub-TLV is optional and its absence

The BIER sub-domain BSL conversion Sub-TLV is optional and its absence indicates that the router is NOT capable of imposing different BSLs but will always forward the packet with the BSL unchanged. This sub-TLV MAY occur at most once in a given BIER sub-TLV. If multiple occurrences of this sub-TLV are received in a given BIER sub-TLV, the BIER sub-TLV MUST be ignored.

The BIER sub-domain BSL conversion Sub-TLV has 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: TBD4.

Length: 0 octets.

2.5. Flooding scope of BIER Information

The flooding scope of the OSPF Extended Prefix Opaque LSA [[RFC7684](#)] that is used for advertising the BIER Sub-TLV is set to area-local. To allow BIER deployment in a multi-area environment, OSPF must propagate BIER information between areas.

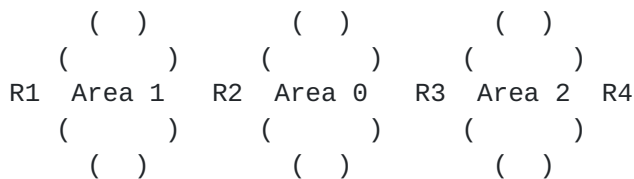


Figure 1: BIER propagation between areas

The following procedure is used in order to propagate BIER related information between areas:

When an OSPF Area Border Router (ABR) advertises a Type-3 Summary LSA from an intra-area or inter-area prefix to all its attached areas, it will also originate an Extended Prefix Opaque LSA, as described in [[RFC7684](#)]. The flooding scope of the Extended Prefix Opaque LSA type will be set to area-local. 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, an OSPF ABR will:

- Examine 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, the ABR will copy the information

from such BIER Sub-TLV when advertising BIER Sub-TLV to each attached area.

In the Figure 1, R1 advertises a prefix 192.0.0.1/32 in Area 1. It also advertises Extended Prefix Opaque LSA for prefix 192.0.0.1/32 and includes BIER Sub-TLV in it. Area Border Router (ABR) R2 calculates the reachability for prefix 192.0.0.1/32 inside Area 1 and propagates it to Area 0. When doing so, it copies the entire BIER Sub-TLV (including all its Sub-TLVs) it received from R1 in Area 1 and includes it in the Extended Prefix Opaque LSA it generates for 192.0.0.1/32 in Area 0. ABR R3 calculates the reachability for prefix 192.0.0.1/32 inside Area 0 and propagates it to Area 2. When doing so, it copies the entire BIER Sub-TLV (including all its Sub-TLVs) it received from R2 in Area 0 and includes it in the Extended Prefix Opaque LSA it generates for 192.0.0.1/32 in Area 2.

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 [[RFC7684](#)].

BIER Sub-TLV: TBD1

BIER MPLS Encapsulation Sub-TLV: TBD2

BIER Tree Type Sub-TLV: TBD3

BIER sub-domain BSL conversion Sub-TLV

4.1. BIER Tree Type Registry

IANA is requested to set up a registry called "BIER Tree Type". The registration policy for this registry is "Standards Action" ([[RFC8126](#)] and [[RFC7120](#)]).

Values in this registry must come from the range 0-255.

The initial value in the BIER Tree Type registry is:

0: Shortest Path First (SPF) algorithm based on standard IGP link metric.

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-bier-architecture]

Wijnands, I., Rosen, E., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast using Bit Index Explicit Replication", [draft-ietf-bier-architecture-06](#) (work in progress), April 2017.

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

Wijnands, I., Rosen, E., Dolganow, A., Tantsura, J., Aldrin, S., and I. Meilik, "Encapsulation for Bit Index Explicit Replication in MPLS and non-MPLS Networks", [draft-ietf-bier-mpls-encapsulation-07](#) (work in progress), June 2017.

[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>>.

[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, <<https://www.rfc-editor.org/info/rfc4915>>.

[RFC7120] Cotton, M., "Early IANA Allocation of Standards Track Code Points", [BCP 100](#), [RFC 7120](#), DOI 10.17487/RFC7120, January 2014, <<https://www.rfc-editor.org/info/rfc7120>>.

[RFC7684] Psenak, P., Gredler, H., Shakir, R., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute Advertisement", [RFC 7684](#), DOI 10.17487/RFC7684, November 2015, <<https://www.rfc-editor.org/info/rfc7684>>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", [BCP 26](#), [RFC 8126](#), DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.

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
Nokia
750 Chai Chee Rd
06-06 Viva Business Park
Singapore 469004

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

