

Internet Engineering Task Force
Internet-Draft
Updates: [8296](#) (if approved)
Intended status: Standards Track
Expires: November 14, 2019

S. Dhanaraj, Ed.
Huawei
IJ. Wijnands
P. Psenak
Cisco Systems, Inc.
Z. Zhang
Juniper Networks.
G. Yan
J. Xie
Huawei
May 13, 2019

LSR Extensions for BIER over Ethernet
draft-ietf-bier-lsr-ethernet-extensions-00

Abstract

Bit Index Explicit Replication (BIER) [[RFC8279](#)] is an architecture that provides multicast forwarding through a "BIER domain" without requiring intermediate routers to maintain multicast related per-flow state. BIER can be supported in MPLS and non-MPLS networks. The common BIER header format and encapsulation for MPLS and non-MPLS networks is specified in [[RFC8296](#)].

This document specifies the required extensions to the IS-IS [[RFC1195](#)] and OSPFv2 [[RFC2328](#)] protocol for supporting BIER in non-MPLS networks using BIER in Ethernet 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 November 14, 2019.

Copyright Notice

Copyright (c) 2019 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	2
2.	Terminology	4
2.1.	Requirements Language	5
3.	Specification	5
3.1.	IS-IS BIER Ethernet Encapsulation Sub-sub TLV	5
3.2.	OSPFv2 BIER Ethernet Encapsulation Sub-TLV	6
4.	Security Considerations	8
5.	IANA Considerations	8
5.1.	IS-IS sub-sub-TLVs for BIER Info sub-TLV Registry	8
5.2.	OSPFv2 Extended Prefix TLV Sub-TLVs Registry	8
6.	Acknowledgments	9
7.	References	9
7.1.	Normative References	9
7.2.	Informative References	9
	Authors' Addresses	11

[1.](#) Introduction

Bit Index Explicit Replication (BIER) [[RFC8279](#)] is an architecture that provides multicast forwarding through a "BIER domain" without requiring intermediate routers to maintain multicast related per-flow state. BIER can be supported in MPLS and non-MPLS networks.

[RFC8296] specifies a common BIER header format for both MPLS and non-MPLS networks, though the first 20-bits of the BIER header (referred as BIFT-id) is a "MPLS Label" in case of MPLS networks and

is a "domain-wide-unique-value" representing the combination of SD-BSL-SI in case of non-MPLS networks.

[[I-D.ietf-bier-non-mpls-bift-encoding](#)] specifies two optional ways of statically assigning domain-wide-unique mapping between BIFT-IDs and SD-BSL-SI combination.

However, BIER architecture [[RFC8279](#)] does NOT require domain-wide-unique BIFT-IDs to be used (even for non-MPLS encapsulation). As discussed in [[I-D.zzhang-bier-rift](#)], the BIFT-ID in case of non-MPLS encapsulation can also just be a local 20-bit opaque value and signaled just like in MPLS case.

As an example, suppose a particular BIER domain contains a SD (SD 0), supports two BSLs (256 and 512), and contains 1024 BFRs. A BFR that is provisioned for above SD, and that supports both BSLs, could advertise the following set of BIFT-id's:

BIFT-id 1: corresponding to SD 0, BSL 256, SI 0.

BIFT-id 2: corresponding to SD 0, BSL 256, SI 1.

BIFT-id 3: corresponding to SD 0, BSL 256, SI 2.

BIFT-id 4: corresponding to SD 0, BSL 256, SI 3.

BIFT-id 5: corresponding to SD 0, BSL 512, SI 0.

BIFT-id 6: corresponding to SD 0, BSL 512, SI 1.

Notice that the example uses ranges of continuous BIFT-ids:

BIFT-id range [1 to 4] correspond to <SD 0, BSL 256>. The first BIFT-id in the range correspond to SI=0, the second correspond to SI=1, and so on.

BIFT-id range [5 to 6] correspond to <SD 0, BSL 512>. The first BIFT-id in the range correspond to SI=0, the second correspond to SI=1.

Strictly speaking, using contiguous range is not required, but it is done for the purpose of simplified signaling similar to MPLS label blocks (notice that locally assigning BIFT-ID ranges requires no

manual processing just like in the case of MPLS label block allocation).

Processing and forwarding of BIER packets requires special software and hardware capabilities. The BFRs supporting a BIER encapsulation type MUST advertise this capability (along with the other required parameters specific to the encapsulation) to the other routers in BIER domain. This advertisement, for example, will enable the other BFRs in the BIER domain in deciding, whether to include or exclude the advertising router from the BAR and/or IPA algorithm while computing the multicast path for a specific encapsulation type.

[RFC8401] and [RFC8444] specifies the required extensions to the IS-IS [RFC1195] and OSPFv2 [RFC2328] protocol respectively for the distribution of BIER sub-domain information including the Sub-sub-TLV required to support BIER in MPLS encapsulation for MPLS networks.

This document specifies the required extensions to the IS-IS [RFC1195] and OSPFv2 [RFC2328] protocol for supporting BIER using BIER in Ethernet encapsulation with dynamically and locally assigned BIFT-IDs.

Support for other encapsulation types are outside the scope of this document.

[2.](#) Terminology

Some of the terminology specified in [RFC8279] is replicated here and extended by necessary definitions:

BIER: Bit Index Explicit Replication

(The overall architecture of forwarding multicast using a Bit Position).

BIER-MPLS: BIER in MPLS encapsulation.

(Encapsulation of BIER header inside MPLS header in MPLS networks).

BIER-ETH: BIER in Ethernet encapsulation.

(Encapsulation of BIER header inside Ethernet header (EtherType=0xAB37) in non-MPLS networks).

BFR: Bit Forwarding Router (A router that participates in Bit Index Multipoint Forwarding). A BFR is identified by a unique BFR-prefix in a BIER domain.

BIFT: Bit Index Forwarding Table used to forward the BIER packets in a domain.

BAR: BIER Algorithm. Used to calculate underlay nexthops as defined by the BAR value.

IPA: IGP Algorithm. May be used to modify, enhance or replace the calculation of underlay paths as defined by the BAR value

[2.1.](#) Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14](#) [[RFC2119](#)] [[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

[3.](#) Specification

A BIER sub-domain MAY support multiple BIER encapsulation types like BIER-MPLS, BIER-ETH. The different encapsulation types supported by a BFR in a sub-domain MUST share the same BFR-id. This would allow the BFR's in transit to translate the encapsulation from one type to the other while forwarding the packet in the BIER sub-domain.

When a BFIR/BFR supports multiple BIER encapsulation types, when sending to a BIER neighbor it MUST use a type that the neighbor also supports. If the neighbor also supports more than one encapsulation type that this BFIR/BFR supports, the type selection could be a matter of local policy and is outside the scope of this document.

3.1. IS-IS BIER Ethernet Encapsulation Sub-sub TLV

BIER Info sub-TLV defined in [RFC8401] is used to advertise the sub-domain id, and other associated parameters of the sub-domain like BFR-id, MT, BAR, IPA.

This document introduces new sub-sub-TLV under BIER Info sub-TLV to advertise the ethernet encapsulation capability and other associated parameters of the encapsulation.

This sub-sub-TLV carries the information for the BIER Ethernet encapsulation including the BitString length supported for a certain <MT,SD> pair.

It is advertised within the BIER Info sub-TLV defined in [RFC8401] which in-turn is carried within the TLVs 235, 237 [RFC5120] or TLVs 135 [RFC5305], or TLV 236 [RFC5308].

This sub-sub-TLV MAY appear multiple times within a single BIER Info sub-TLV. If the same BitString length is repeated in multiple BIER Ethernet encapsulation sub-sub-TLVs inside the same BIER Info sub-TLV, the BIER Info sub-TLV MUST be ignored.

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	
+-----+			
Max SI		BS Len BIFT-id	
+-----+			

Type: 2 (suggested value - To be assigned by IANA).

Length: 4

Max SI: A 1 octet field encoding the Maximum Set Identifier (Section 1 of [RFC8279]) used in the encapsulation for this BIER subdomain for this BitString length. The first BIFT-id is for SI=0,

the second BIFT-id is for SI=1, etc. If the BIFT-id associated with the Maximum Set Identifier exceeds the 20-bit range, the sub-sub-TLV MUST be ignored.

Local BitString Length (BS Len): A 4 bit field encoding the bitstring length (as per [[RFC8296](#)]) supported for the encapsulation.

BIFT-id: A 20 bit field encoding the first BIFT-id of the BIFT-id range.

The "BIFT-id range" is the set of 20-bit values beginning with the BIFT-id and ending with (BIFT-id + (Max SI)). A unique BIFT-id range is allocated for each BitString length and sub-domain-id. These BIFT-id's are used for BIER forwarding as described in [[RFC8279](#)] and [[RFC8296](#)].

The size of the BIFT-id range is determined by the number of SI's ([Section 1 of \[RFC8279\]](#)) that are used in the network. Each SI maps to a single BIFT-id in the BIFT-id range: the first BIFT-id is for SI=0, the second BIFT-id is for SI=1, etc.

If the BIFT-id associated with the Maximum Set Identifier exceeds the 20-bit range, the BIER Ethernet Encapsulation Sub-sub-TLV containing the error MUST be ignored.

[3.2.](#) OSPFv2 BIER Ethernet Encapsulation Sub-TLV

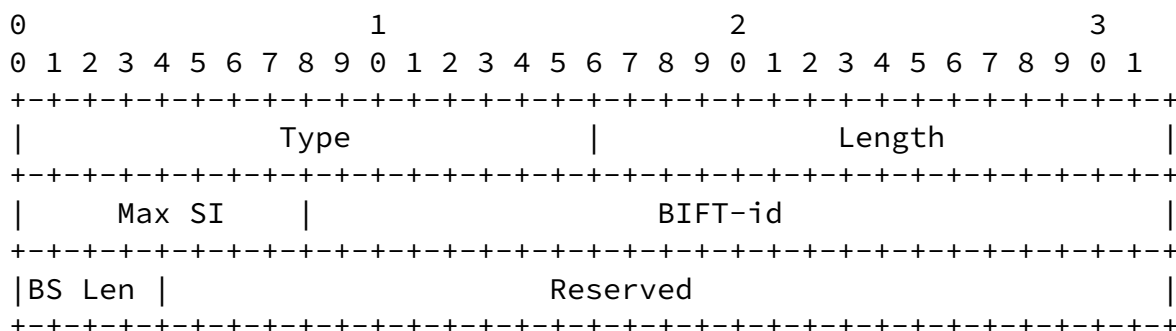
BIER Sub-TLV defined in [[RFC8444](#)] is used to advertise the sub-domain id, and other associated parameters of the sub-domain like BFR-id, MT, BAR, IPA.

This document introduces new Sub-TLV under BIER Sub-TLV to advertise the ethernet encapsulation capability and other associated parameters of the encapsulation.

This Sub-TLV carries the information for the BIER Ethernet encapsulation including the BitString length supported for a certain <MT,SD> pair.

It is advertised within the BIER Sub-TLV defined in [RFC8444] which in-turn is carried within the OSPFv2 Extended Prefix TLV defined in [RFC7684].

This Sub-TLV MAY appear multiple times within a single BIER Sub-TLV. If the same BitString length is repeated in multiple BIER Ethernet encapsulation Sub-TLVs inside the same BIER Sub-TLV, the BIER Sub-TLV MUST be ignored.



Type: 11 (suggested value - To be assigned by IANA).

Length: 8

Max SI: A 1 octet field encoding the Maximum Set Identifier ([Section 1 of \[RFC8279\]](#)) used in the encapsulation for this BIER subdomain for this BitString length. The first BIFT-id is for SI=0, the second BIFT-id is for SI=1, etc. If the BIFT-id associated with the Maximum Set Identifier exceeds the 20-bit range, the sub-sub-TLV MUST be ignored.

BIFT-id: A 3-octet field, where the 20 rightmost bits represent the first BIFT-id in the BIFT-id range. The 4 leftmost bits MUST be ignored.

The "BIFT-id range" is the set of 20-bit values beginning with the BIFT-id and ending with (BIFT-id + (Max SI)). A unique BIFT-id range is allocated for each BitString length and sub-domain-id. These BIFT-id's are used for BIER forwarding as described in [RFC8279] and [RFC8296].

([Section 1 of \[RFC8279\]](#)) that are used in the network. Each SI maps to a single BIFT-id in the BIFT-id range: the first BIFT-id is for SI=0, the second BIFT-id is for SI=1, etc.

If the BIFT-id associated with the Maximum Set Identifier exceeds the 20-bit range, the BIER Ethernet Encapsulation Sub-sub-TLV containing the error MUST be ignored.

Local BitString Length (BS Len): A 4 bit field encoding the bitstring length (as per [\[RFC8296\]](#)) supported for the encapsulation.

Reserved: SHOULD be set to 0 on transmission and MUST be ignored on reception.

[4.](#) Security Considerations

Security concerns for IS-IS are addressed in [\[RFC5304\]](#) and [\[RFC5310\]](#) and the security concerns for IS-IS extensions for BIER are addressed in [\[RFC8401\]](#). This document introduces new sub-sub-TLV for the already existing IS-IS TLVs defined for distributing the BIER sub-domain information in [\[RFC8401\]](#). It does not introduce any new security risks to IS-IS.

Security concerns and required extensions for OSPFv2 are addressed in [\[RFC2328\]](#) and [\[RFC7684\]](#) and the security concerns for OSPFv2 extensions for BIER are addressed in [\[RFC8444\]](#). This document introduces new Sub-TLV for the already existing OSPFv2 TLV defined for distributing the BIER sub-domain information in [\[RFC8444\]](#). It does not introduce any new security risks to OSPFv2.

[5.](#) IANA Considerations

The document requests new allocations from the IANA registries as follows

[5.1.](#) IS-IS sub-sub-TLVs for BIER Info sub-TLV Registry

BIER Ethernet Encapsulation sub-sub-TLV: 2 (suggested)

[5.2.](#) OSPFv2 Extended Prefix TLV Sub-TLVs Registry

BIER Ethernet Encapsulation Sub-TLV: 11 (suggested)

6. Acknowledgments

The author wants to thank Antonie Przygienda for his comments and suggestions.

7. References

7.1. Normative References

- [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>>.
- [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>>.
- [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>>.
- [RFC8401] Ginsberg, L., Ed., Przygienda, T., Aldrin, S., and Z. Zhang, "Bit Index Explicit Replication (BIER) Support via IS-IS", [RFC 8401](#), DOI 10.17487/RFC8401, June 2018, <<https://www.rfc-editor.org/info/rfc8401>>.
- [RFC8444] Psenak, P., Ed., Kumar, N., Wijnands, IJ., Dolganow, A., Przygienda, T., Zhang, J., and S. Aldrin, "OSPFv2 Extensions for Bit Index Explicit Replication (BIER)", [RFC 8444](#), DOI 10.17487/RFC8444, November 2018, <<https://www.rfc-editor.org/info/rfc8444>>.

7.2. Informative References

- [I-D.ietf-bier-non-mpls-bift-encoding] Wijnands, I., Xu, X., and H. Bidgoli, "An Optional Encoding of the BIFT-id Field in the non-MPLS BIER Encapsulation", [draft-ietf-bier-non-mpls-bift-encoding-01](#) (work in progress), October 2018.

[I-D.zzhang-bier-rift]

Zhang, Z., Ma, S., and Z. Zhang, "Supporting BIER with RIFT", [draft-zzhang-bier-rift-00](#) (work in progress), March 2018.

[RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", [RFC 1195](#), DOI 10.17487/RFC1195, December 1990, <<https://www.rfc-editor.org/info/rfc1195>>.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, [RFC 2328](#), DOI 10.17487/RFC2328, April 1998, <<https://www.rfc-editor.org/info/rfc2328>>.

[RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", [RFC 5120](#), DOI 10.17487/RFC5120, February 2008, <<https://www.rfc-editor.org/info/rfc5120>>.

[RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", [RFC 5304](#), DOI 10.17487/RFC5304, October 2008, <<https://www.rfc-editor.org/info/rfc5304>>.

[RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", [RFC 5305](#), DOI 10.17487/RFC5305, October 2008, <<https://www.rfc-editor.org/info/rfc5305>>.

[RFC5308] Hopps, C., "Routing IPv6 with IS-IS", [RFC 5308](#), DOI 10.17487/RFC5308, October 2008, <<https://www.rfc-editor.org/info/rfc5308>>.

[RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", [RFC 5310](#), DOI 10.17487/RFC5310, February 2009, <<https://www.rfc-editor.org/info/rfc5310>>.

[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>>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in [RFC 2119](#) Key Words", [BCP 14](#), [RFC 8174](#), DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

Dhanaraj, et al.

Expires November 14, 2019

[Page 10]

Internet-Draft LSR Extensions for BIER over Ethernet

May 2019

Authors' Addresses

Senthil Dhanaraj (editor)
Huawei

Email: senthil.dhanaraj.ietf@gmail.com

IJsbrand Wijnands
Cisco Systems, Inc.

Email: ice@cisco.com

Peter Psenak
Cisco Systems, Inc.

Email: ppsenak@cisco.com

Zhaohui Zhang
Juniper Networks.

Email: zzhang@juniper.net

Gang Yan
Huawei

Email: yangang@huawei.com

Jingrong Xie
Huawei

Email: xiejingrong@huawei.com