

BIER Working Group
Internet-Draft
Intended status: Standards Track
Expires: 4 December 2025

R. Chen
Z. Zhang
ZTE Corporation
V. Govindan
Cisco
IJ. Wijnands
Individual
Z. Zhang
Juniper Networks
2 June 2025

BGP Link-State extensions for BIER
draft-ietf-bier-bgp-ls-bier-ext-20

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 also does not require any 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. The BIER header contains a bitstring 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.

BGP Link-State (BGP-LS) enables the collection of various topology informations from the network, and the topology informations are used by the controller to calculate the forwarding tables and then propagate them onto the BFRs (instead of having each node to calculate on its own) and that can be for both inter-as and intra-as situations.

This document specifies extensions to the BGP Link-state address-family in order to advertise the BIER informations.

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 4 December 2025.

Copyright Notice

Copyright (c) 2025 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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
1.1. Requirements Language	3
2. BGP-LS Extensions for BIER	3
2.1. Prefix Attributes TLVs	4
2.2. The BIER information TLV	4
2.3. The BIER MPLS Encapsulation TLV	5
2.4. The BIER non-MPLS Encapsulation TLV	6
3. Equivalent IS-IS BIER TLVs/Sub-TLVs	7
4. Equivalent OSPFv2/OSPFv3 BIER TLVs/Sub-TLVs	7
5. IANA Considerations	8
6. Security Considerations	8
7. Acknowledgments	9
8. References	9
8.1. Normative References	9
8.2. Informative References	10
Authors' Addresses	11

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 also does not require any 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. The BIER header contains a bitstring 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 are expressed by setting the bits that correspond to those routers in the BIER header.

The BGP-LS address-family/sub-address-family have been defined to allow BGP to carry Link-State informations. This document specifies extensions to the BGP Link-state address-family in order to advertise BIER-specific informations, Similar to BGP-LS Advertisement of IGP Traffic Engineering Performance Metric Extensions([RFC8571]). An external component (e.g., a controller/a PCE(see [RFC4655] for PCE-Based Architecture ,[RFC5440] for PCEP and [RFC5376] for Inter-AS Requirements for the PCEP.))then can learn the BIER informations in the "northbound" direction and calculate BIRT/BIFT and then propagate them onto BFRs (instead of having each BFR to calculate on its own), and that can be for both inter-as and intra-as situations.

1.1. 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.

2. BGP-LS Extensions for BIER

[RFC8279] defines the BFR - A router that supports BIER is known as a "Bit-Forwarding Router"(BFR), and each BFR MUST be assigned a "BFR-Prefix". A BFR's Prefix MUST be an IP address (either IPv4 or IPv6) of the BFR, and MUST be unique and routable within the BIER domain as described in section 2 of [RFC8279], and then external component (e.g., a controller) need to collect BIER informations of BIER routers are associated with the BFR-Prefix in the "northbound" direction within the BIER domain.

Given that the BIER informations are associated with the prefix, the Prefix Attribute TLV [RFC9552] can be used to carry the BIER informations. The new Prefix Attribute TLVs are defined for the encoding of BIER informations.

2.1. Prefix Attributes TLVs

The following Prefix Attribute TLVs are defined:

Type	Description	Section
TBD1	BIER information	section 3.2
TBD2	BIER MPLS Encapsulation	section 3.3
TBD3	BIER non-MPLS Encapsulation	section 3.4

Table 1: The new Prefix Attribute TLVs

2.2. The BIER information TLV

A new Prefix Attribute TLV (defined in [RFC9552] is defined for distributing BIER informations. The new TLV is called the BIER information TLV. The BIER information TLV may appear multiple times.

The following BIER information TLV is defined:




Figure 1: The BIER information TLV

Type: A 2-octet field with value TBD, see IANA Considerations section.

Length: 2 octets.

Subdomain-id: Unique value identifying the BIER sub-domain, 1 octet.

BFR-ID: A 2-octet field encoding the BFR-ID, as documented in [RFC8279]. If the BFR-ID is zero, it means, the advertising router is not advertising any BIER-id. In some environment, BFR-ID can be configured by NMS, The BFR-ID should be sent to a controller.

BAR: A 1-octet field encoding the BIER Algorithm, used to calculate underlay paths to reach BFRs. Values are allocated from the "BIER Algorithms" registry which are defined in [RFC8401], [RFC8444] and [I-D.ietf-bier-ospfv3-extensions].

IPA: A 1-octet field encoding the IGP Algorithm, used to either modify, enhance, or replace the calculation of underlay paths to reach BFRs as defined by the BAR value. Values are from the "IGP Algorithm" registry.

Reserved: MUST be 0 on transmission, ignored on reception. May be used in future versions.

2.3. The BIER MPLS Encapsulation TLV

The BIER MPLS Encapsulation TLV is used in order to advertise MPLS specific informations used for BIER. It MAY appear multiple times.

In some environment, each router allocates its labels, and advertises it to the controller. That solution is simpler as the controller does not need to deal with label allocation. If the controller has to deal with Label allocation, there needs to be a (global) range carved out such there are no conflicts. We can avoid all that by having the router allocate the BIER Label range and advertise it to the controller.

The following the BIER MPLS Encapsulation Sub-TLV is defined:




Figure 2: The BIER MPLS Encapsulation TLV

Type: A 2-octet field with value TBD, see IANA Considerations section.

Length: 2 octets.

Max SI: A 1-octet field encoding the maximum Set Identifier(as defined in [RFC8279]), used in the encapsulation for this BIER subdomain for this BitString length.

Label: First label of the range, 20 bits. The labels are as defined in [RFC8296].

BS Len: A 4-bit field encoding the Bitstring length as per [RFC8296].

BS length in multiple BIER MPLS Encapsulation Sub-TLV associated with the same BIER Sub-TLV MUST NOT repeat, otherwise only the first BIER MPLS Encapsulation Sub-TLV with such BS length MUST be used and any subsequent BIER MPLS Encapsulation Sub-TLVs with the same BS length MUST be ignored.

2.4. The BIER non-MPLS Encapsulation TLV

The BIER non-MPLS Encapsulation TLV is used in order to advertise non-MPLS encapsulation(e.g. ethernet encapsulation) capability and other associated parameters of the encapsulation.It MAY appear multiple times.

The following the BIER non-MPLS Encapsulation Sub-TLV is defined:




Figure 3: The BIER non-MPLS Encapsulation TLV

Type: A 2-octet field with value TBD, see IANA Considerations section.

Length: 2 octets.

Max SI: A 1-octet field encoding the maximum Set Identifier(as defined in [RFC8279]), used in the encapsulation for this BIER subdomain for this BitString length.

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].

Local BitString Length (BS Len): A 4-bit field encoding the Bitstring length as per [RFC8296].

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

3. Equivalent IS-IS BIER TLVs/Sub-TLVs

This section illustrates the IS-IS BIER Extensions Sub-TLVs/Sub-Sub-TLVs mapped to the ones defined in this document.

The following table illustrates for each BGP-LS TLV, and its equivalence in IS-IS.

Description	IS-IS TLV/ Sub-TLV	Reference
BIER information	BIER info Sub-TLV	[RFC8401]
BIER MPLS Encapsulation	BIER MPLS Encapsulation Sub-Sub-TLV	[RFC8401]
BIER non-MPLS Encapsulation	BIER non-MPLS Encapsulation Sub-Sub-TLV	[I-D.ietf-bier-lsr-non-mpls-extensions]

Table 2: IS-IS BIER Extensions Sub-TLVs/Sub-Sub-TLVs

4. Equivalent OSPFv2/OSPFV3 BIER TLVs/Sub-TLVs

This section illustrates the BIER Extensions TLVs/Sub-TLVs mapped to the ones defined in this document.

The following table illustrates for each BGP-LS TLV, and its equivalence in OSPFv2/OSPFV3.

Description	OSPFv2/OSPFV3 sub-TLV/Sub- Sub-TLV	Reference
BIER information	BIER Sub-TLV	[RFC8444], [I-D.ietf-bier-ospfv3-extensions]
BIER MPLS Encapsulation	BIER MPLS Encapsulation Sub-TLV	[RFC8444], [I-D.ietf-bier-ospfv3-extensions]
BIER non-MPLS Encapsulation	BIER non-MPLS Encapsulation Sub-TLV	[I-D.ietf-bier-lsr-non-mpls-extensions]

Table 3: OSPFv2/OSPFV3 BIER TLVs/Sub-TLVs

5. IANA Considerations

IANA maintains a registry group called "Border Gateway Protocol - Link State (BGP-LS) Parameters" with a registry called "BGP-LS NLRI and Attribute TLVs". The following TLV codepoints are suggested (for early allocation by IANA):

Code Point	Description	Value defined in
TBD1	BIER information	this document
TBD2	BIER MPLS Encapsulation	this document
TBD3	BIER non-MPLS Encapsulation	this document

Table 4: The new Prefix Attribute TLVs

6. Security Considerations

Procedures and protocol extensions defined in this document do not affect the BGP security model. See the "Security Considerations" section of [RFC4271] for a discussion of BGP security. Security considerations for acquiring and distributing BGP-LS information are discussed in [RFC9552].

The TLVs introduced in this document are used to propagate the Bit Index Explicit Replication (BIER) defined in [RFC8401], [RFC8444] , [I-D.ietf-bier-ospfv3-extensions], and [I-D.ietf-bier-lsr-non-mpls-extensions]. These TLVs represent the bier information associated with the prefix. It is assumed that the IGP instances originating these TLVs will support all the required security and authentication mechanisms in [RFC8401], [RFC8444], [I-D.ietf-bier-ospfv3-extensions] ,and [I-D.ietf-bier-lsr-non-mpls-extensions] in order to prevent any security issues when propagating the TLVs into BGP-LS. The advertisement of the link attribute information defined in this document present no additional risk beyond that associated with the existing link attribute informations already supported in [RFC9552].

7. Acknowledgments

The authors thank Peter Psenak, Ketan Talaulikar, Gyan Mishra and Benchong Xu and many others for their suggestions and comments.

8. References

8.1. Normative References

- [I-D.ietf-bier-lsr-non-mpls-extensions]
Dhanaraj, S., Yan, G., Wijnands, I., Psenak, P., Zhang, Z. J., and J. Xie, "LSR Extensions for BIER non-MPLS Encapsulation", Work in Progress, Internet-Draft, draft-ietf-bier-lsr-non-mpls-extensions-03, 7 February 2024, <<https://datatracker.ietf.org/doc/html/draft-ietf-bier-lsr-non-mpls-extensions-03>>.
- [I-D.ietf-bier-ospfv3-extensions]
Psenak, P., Nainar, N. K., and I. Wijnands, "OSPFv3 Extensions for BIER", Work in Progress, Internet-Draft, draft-ietf-bier-ospfv3-extensions-07, 1 December 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-bier-ospfv3-extensions-07>>.
- [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>>.
- [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>>.

- [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>>.
- [RFC8571] Ginsberg, L., Ed., Previdi, S., Wu, Q., Tantsura, J., and C. Filsfils, "BGP - Link State (BGP-LS) Advertisement of IGP Traffic Engineering Performance Metric Extensions", RFC 8571, DOI 10.17487/RFC8571, March 2019, <<https://www.rfc-editor.org/info/rfc8571>>.
- [RFC9552] Talaulikar, K., Ed., "Distribution of Link-State and Traffic Engineering Information Using BGP", RFC 9552, DOI 10.17487/RFC9552, December 2023, <<https://www.rfc-editor.org/info/rfc9552>>.

8.2. Informative References

- [RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <<https://www.rfc-editor.org/info/rfc4271>>.
- [RFC4655] Farrel, A., Vasseur, J.-P., and J. Ash, "A Path Computation Element (PCE)-Based Architecture", RFC 4655, DOI 10.17487/RFC4655, August 2006, <<https://www.rfc-editor.org/info/rfc4655>>.

- [RFC5376] Bitar, N., Zhang, R., and K. Kumaki, "Inter-AS Requirements for the Path Computation Element Communication Protocol (PCECP)", RFC 5376, DOI 10.17487/RFC5376, November 2008, <<https://www.rfc-editor.org/info/rfc5376>>.
- [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>>.

Authors' Addresses

Ran Chen
ZTE Corporation
Nanjing
China
Email: chen.ran@zte.com.cn

Zheng Zhang
ZTE Corporation
Nanjing
China
Email: zhang.zheng@zte.com.cn

Vengada Prasad Govindan
Cisco
Email: venggovi@cisco.com

IJsbrand Wijnands
Individual
Email: ice@braindump.be

Zhaohui Zhang
Juniper Networks
Email: zzhang@juniper.net