

BIER WG
Internet-Draft
Intended status: Standards Track
Expires: 26 July 2024

Z. Zhang
ZTE Corporation
A. Przygienda
Juniper Networks
23 January 2024

BIER in BABEL
draft-zhang-bier-babel-extensions-10

Abstract

BIER introduces a novel multicast architecture. It does not require a signaling protocol to explicitly build multicast distribution trees, nor does it require intermediate nodes to maintain any per-flow state.

Babel defines a distance-vector routing protocol that operates in a robust and efficient fashion both in wired as well as in wireless mesh networks. This document defines a way to carry necessary BIER signaling information in Babel.

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 26 July 2024.

Copyright Notice

Copyright (c) 2024 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	2
2. Terminology	2
3. Conventions Used in This Document	3
4. Advertisement of BIER information	3
4.1. BIER BFR-prefix and BIER sub-TLV	3
4.1.1. BIER sub-TLV	3
4.2. BIER MPLS Encapsulation sub-sub-TLV	4
4.3. BIER non-MPLS Encapsulation sub-sub-TLV	5
4.3.1. BIER IPv6 transportation sub-sub-TLV	6
5. Tree types and tunneling	7
6. Security Considerations	7
7. IANA Considerations	7
8. References	7
8.1. Normative References	7
8.2. Informative References	8
Authors' Addresses	8

1. Introduction

[RFC8279] introduces a novel multicast architecture. It does not require a signaling protocol to explicitly build multicast distribution trees, nor does it require intermediate nodes to maintain any per-flow state. All procedures necessary to support BIER are abbreviated by the "BIER architecture" moniker in this document.

[RFC8966] define a distance-vector routing protocol under the name of "Babel". Babel operates in a robust and efficient fashion both in ordinary wired as well as in wireless mesh networks.

2. Terminology

The terminology of this documents follows [RFC8279] and [RFC8966].

3. Conventions Used in This Document

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.

4. Advertisement of BIER information

In case a router is configured with BIER information, and Babel is the routing protocol used, such a router MAY use Babel protocol to announce the BIER information using the BIER sub-TLV specified below.

4.1. BIER BFR-prefix and BIER sub-TLV

BFR-prefix and according information is carried in a Babel Update TLV per [RFC8966]. A new sub-TLV is defined to convey further BIER information such as BFR-id, sub-domain-id and BSL. Two sub-sub-TLVs are carried as payload of BIER sub-TLV.

The mandatory bit of BIER sub-TLV should be set to 0. If a router cannot recognize a sub-TLV, the router MUST ignore this unknown sub-TLV.

4.1.1. BIER sub-TLV

The BIER sub-TLV format aligns exactly with the definition and restrictions in [RFC8401], [RFC8444] and [I-D.ietf-bier-ospfv3-extensions]. It is a sub-TLV of Babel update TLV. The prefix MUST NOT be summarized and the according sub-TLV MUST be treated as optional and transitive.

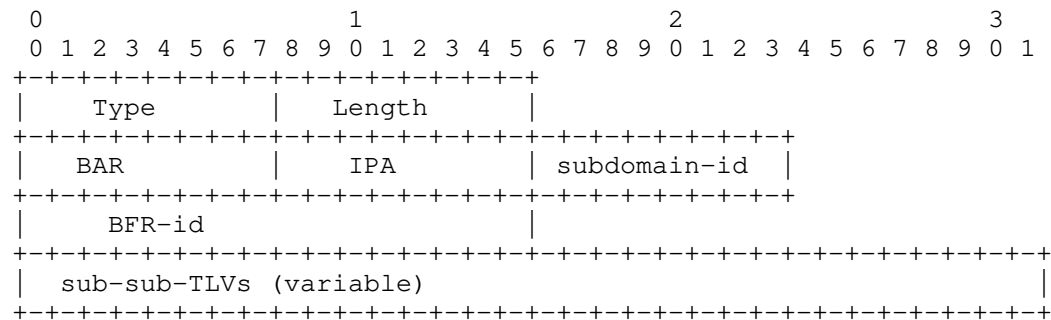


Figure 1: BIER sub-TLV

- * Type: as indicated in IANA section.
 - * Length: 1 octet. Include the length of BIER sub-TLV and potential length of the sub-sub-TLVs.
 - * BAR: BIER Algorithm. Specifies a BIER-specific algorithm used to calculate underlay paths to reach BFERs. Values are allocated from the "BIER Algorithms" registry. 1 octet.
 - * IPA: IGP Algorithm. Specifies an IGP Algorithm to either modify, enhance, or replace the calculation of underlay paths to reach BFERs as defined by the BAR value. Values are from the IGP Algorithm registry. 1 octet.
 - * 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 no BFR-id has been assigned this field is set to the invalid BFR-id.
- 4.2. BIER MPLS Encapsulation sub-sub-TLV

The BIER MPLS Encapsulation sub-sub-TLV can be carried by BIER sub-TLV. The format and restrictions are aligned with [RFC8401], [RFC8444] and [I-D.ietf-bier-ospfv3-extensions]. This sub-sub-TLV carries the information for the BIER MPLS encapsulation including the label range for a specific BSL for a certain <MT,SD> pair.




Figure 2: MPLS Encapsulation sub-sub-TLV

- * Type: value of 1 indicating MPLS encapsulation.
 - * Length: 1 octet
 - * Max SI: Maximum Set Identifier (Section 1 of [RFC8279]) used in the encapsulation for this BIER subdomain for this BitString length, 1 octet. 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 the label associated with the Maximum Set Identifier exceeds the 20-bit range, the sub-sub-TLV MUST be ignored.
 - * Local BitString Length (BS Len): Encoded BitString length as per [RFC8296]. 4 bits.
 - * Label: First label, 20 bits. The labels are as defined in [RFC8296].
- 4.3. BIER non-MPLS Encapsulation sub-sub-TLV

The BIER non-MPLS Encapsulation sub-sub-TLV can be carried by BIER sub-TLV. The format and restrictions are aligned with [I-D.ietf-bier-lsr-non-mpls-extensions]. This sub-sub-TLV carries the information for the BIER MPLS encapsulation including the label range for a specific BSL for a certain <MT,SD> pair.




Figure 3: non-MPLS Encapsulation sub-sub-TLV

- * Type: value of 2 indicating non-MPLS encapsulation.
- * Length: 1 octet
- * Max SI: Maximum Set Identifier (Section 1 of [RFC8279]) used in the encapsulation for this BIER subdomain for this BitString length, 1 octet. 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)). These BIFT-id's are used for BIER forwarding as described in [RFC8279] and [RFC8296].
- * Local BitString Length (BS Len): Encoded BitString length as per [RFC8296]. 4 bits.

4.3.1. BIER IPv6 transportation sub-sub-TLV

The BIER IPv6 transportation sub-sub-TLV can be carried by BIER non-MPLS Encapsulation sub-sub-TLV. The format and restrictions are aligned with [I-D.ietf-bier-bierin6]. A node that requires IPv6 encapsualtion MUST advertise the BIER IPv6 transportation sub-sub-TLV according to local configuration or policy in the BIER domain to request other BFRs to always use IPv6 encapsulation.

The format is the same with the definition in section 4.1, [I-D.ietf-bier-bierin6].

5. Tree types and tunneling

Since Babel is performing a diffusion computation, support for different tree types is not as natural as with link-state protocols. Hence this specification is assuming that normal Babel reachability computation is performed without further modifications.

BIER architecture does not rely on all routers in a domain performing BFR procedures. How to support tunnels that will allow to tunnel BIER across such routers in Babel is for further study.

6. Security Considerations

Security considerations discussed in [RFC8296], [RFC8966] apply to this document.

7. IANA Considerations

A new type of Babel update sub-TLV needs to be defined for BIER information advertisement.

8. References

8.1. 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>>.
- [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>>.
- [RFC8966] Chroboczek, J. and D. Schinazi, "The Babel Routing Protocol", RFC 8966, DOI 10.17487/RFC8966, January 2021, <<https://www.rfc-editor.org/info/rfc8966>>.

8.2. Informative References

- [I-D.ietf-bier-bierin6]
Zhang, Z., Zhang, Z. J., Wijnands, I., Mishra, M. P., Bidgoli, H., and G. S. Mishra, "Supporting BIER in IPv6 Networks (BIERin6)", Work in Progress, Internet-Draft, draft-ietf-bier-bierin6-08, 18 September 2023, <<https://datatracker.ietf.org/doc/html/draft-ietf-bier-bierin6-08>>.
- [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-01, 19 September 2022, <<https://datatracker.ietf.org/doc/html/draft-ietf-bier-lsr-non-mpls-extensions-01>>.
- [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>>.

Authors' Addresses

Zheng (Sandy) Zhang
ZTE Corporation
Email: zhang.zheng@zte.com.cn

Tony Przygienda
Juniper Networks
Email: prz@juniper.net