

Workgroup: BIER Working Group
Internet-Draft: draft-xzlnp-bier-ioam-03
Published: 12 January 2022
Intended Status: Standards Track
Expires: 16 July 2022
Authors: X. Min Z. Zhang Y. Liu
 ZTE Corp. ZTE Corp. China Mobile
 N. Nainar C. Pignataro
 Cisco Systems, Inc. Cisco Systems, Inc.

Bit Index Explicit Replication (BIER) Encapsulation for In-situ OAM (IOAM) Data

Abstract

In-situ Operations, Administration, and Maintenance (IOAM) records operational and telemetry information in the packet while the packet traverses a path in the network. Bit Index Explicit Replication (BIER) is an architecture that provides optimal multicast forwarding through a "multicast domain", without requiring intermediate routers to maintain any per-flow state or to engage in an explicit tree-building protocol. The BIER header contains a bit-string in which each bit represents exactly one egress router to forward the packet to. This document outlines the requirements to carry IOAM data in BIER header and specifies how IOAM data fields are encapsulated in BIER header.

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 16 July 2022.

Copyright Notice

Copyright (c) 2022 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. Conventions Used in This Document](#)
 - [2.1. Requirements Language](#)
 - [2.2. Abbreviations](#)
- [3. Requirements to carry IOAM data](#)
- [4. IOAM data fields encapsulation in BIER header](#)
- [5. Considerations](#)
 - [5.1. Selecting the encapsulation approach](#)
 - [5.2. Interaction with the BIER OAM field](#)
- [6. Security Considerations](#)
- [7. IANA Considerations](#)
- [8. Acknowledgements](#)
- [9. References](#)
 - [9.1. Normative References](#)
 - [9.2. Informative References](#)
- [Authors' Addresses](#)

1. Introduction

In-situ Operations, Administration, and Maintenance (IOAM) records operational and telemetry information in the packet while the packet traverses a path in the network. [[I-D.ietf-ippm-ioam-data](#)] defines multiple IOAM options with different IOAM data fields used to record various telemetry data from the transit nodes. [[I-D.ietf-ippm-ioam-direct-export](#)] defines IOAM Direct Export option with IOAM data fields, which indicate telemetry data to be collected without being embedded in data packets. The term "in-situ" refers to the fact that the OAM data is added to the data packets rather than being sent within packets specifically dedicated to OAM.

Bit Index Explicit Replication (BIER), as defined in [[RFC8279](#)], is an architecture that provides optimal multicast forwarding through a "multicast domain", without requiring intermediate routers to maintain any per-flow state or to engage in an explicit tree-building protocol. The BIER header, as defined in [[RFC8296](#)], contains a bit-string in which each bit represents exactly one egress router to forward the packet to.

This document outlines the requirements to carry IOAM data in BIER header and specifies how IOAM data fields are encapsulated in BIER header.

2. Conventions Used in This Document

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.

2.2. Abbreviations

Abbreviations used in this document:

BFER: Bit Forwarding Egress Router

BFIR: Bit Forwarding Ingress Router

BIER: Bit Index Explicit Replication

GRE: Generic Routing Encapsulation

IOAM: In-situ Operations, Administration, and Maintenance

OAM: Operations, Administration, and Maintenance

3. Requirements to carry IOAM data

[[I-D.ietf-bier-use-cases](#)] lists many use cases for BIER. Usually there are many multicast flows within one network domain, and some of the multicast flows, such as live video and real-time meeting, are sensitive to packet loss, delay and other factors. The network operator wants to know the real-time statistics for these flows, such as delay, sequence, the ingress/egress interface, and the usage of buffer.

So methods are needed for measuring the real-time transportation guarantee of BIER packets. This document attempts to provide a way to record operational and telemetry information in the BIER packets through in-situ OAM.

4. IOAM data fields encapsulation in BIER header

The BIER header is defined in [[RFC8279](#)]. The BIER OAM header that follows BIER header is defined in [[I-D.ietf-bier-ping](#)]. IOAM-Data-Fields can either be carried in BIER using a new type of OAM message

which follows the BIER OAM header (referred to as option 1), or be carried in BIER using a new next protocol header which immediately follows the BIER header (referred to as option 2). In this document, option 2 is selected and the reason is discussed in Section 5.1. An IOAM header is added containing different IOAM-Data-Fields defined in [I-D.ietf-ippm-ioam-data] and [I-D.ietf-ippm-ioam-direct-export].

In a BIER domain where IOAM is applied, inserting the IOAM header into BIER packets is enabled at the BFIRs, which also serve as IOAM encapsulating nodes by means of configuration, and deleting the IOAM header from BIER packets is enabled at the BFERs, which also serve as IOAM decapsulating nodes by means of configuration.

The Encapsulation format for IOAM over BIER is defined as follows:

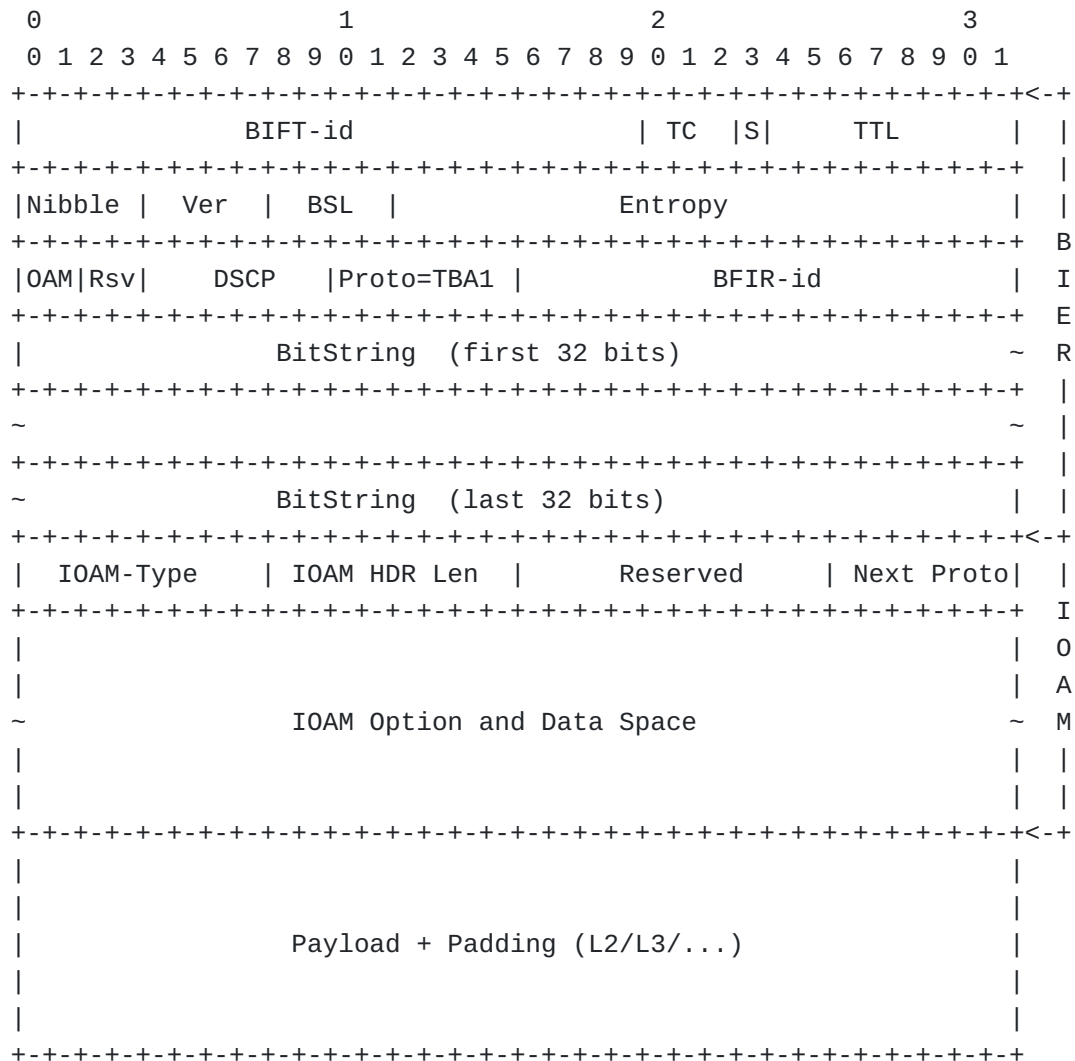


Figure 1: IOAM Encapsulation Format within BIER

The BIER header and fields are defined in [[RFC8296](#)]. Within the BIER header, a 6-bit field as "Proto" (Next Protocol) is used to identify the type of the payload immediately following the BIER header, The "Proto" value is set to TBA1 when the IOAM header is present.

The IOAM related fields in BIER are defined as follows:

IOAM-Type: An 8-bit field defining the IOAM option type, as defined in Section 8.1 of [[I-D.ietf-ippm-ioam-data](#)] and Section 4.1 of [[I-D.ietf-ippm-ioam-direct-export](#)].

IOAM HDR Len: An 8-bit unsigned integer. Length of the IOAM header in 4-octet units.

Reserved: A 10-bit reserved field MUST be set to zero upon transmission and ignored upon receipt.

Next Proto: A 6-bit unsigned integer that identifies the type of payload immediately following this IOAM option. The semantics of this field are identical to the "Proto" field in [[RFC8296](#)].

IOAM Option and Data Space: IOAM option header and data is present as specified by the IOAM-Type field, and is defined in Section 5 of [[I-D.ietf-ippm-ioam-data](#)] and Section 3 of [[I-D.ietf-ippm-ioam-direct-export](#)].

Multiple IOAM options MAY be included within a BIER encapsulation. For example, if a BIER encapsulation contains two IOAM options preceding a data payload, the "Next Proto" field of the first IOAM option would be set to the value of TBA1 that indicates a second IOAM option follows, while the "Next Proto" field of the second IOAM option would be set to the value of "BIER Next Protocol" indicating the type of the data payload. Each type of IOAM option MUST occur at most once within a BIER encapsulation.

5. Considerations

This section summarizes a set of considerations on the overall approach taken for IOAM data encapsulation in BIER, as well as deployment considerations.

5.1. Selecting the encapsulation approach

Both the options described in Section 4 are supposed to be feasible, nevertheless this document needs to select one as standardized encapsulation for IOAM over BIER. Considering the fact that the encapsulation format option 2 using a new next protocol header is more concise than option 1 using a new type of OAM message, and many other transport protocols, e.g., GRE, use a new next protocol header

to encapsulate IOAM data, the encapsulation format option 2 is selected as the standardized one.

5.2. Interaction with the BIER OAM field

[RFC8296] defines a two-bit field, referred to as OAM. [I-D.ietf-bier-pmmm-oam] describes how to use the two-bit OAM field for alternate marking performance measurement method. This document would not change the semantics of the two-bit OAM field. The BIER IOAM header and the BIER OAM field are orthogonal and they can co-exist in one packet, i.e., a BIER packet with IOAM data can set the OAM field and a BIER packet with OAM field set can carry IOAM data too.

6. Security Considerations

This document describes the encapsulation of IOAM data fields in BIER. Security considerations of the specific IOAM data fields are described in [I-D.ietf-ippm-ioam-data] and [I-D.ietf-ippm-ioam-direct-export].

IOAM is considered a "per domain" feature, where one or several operators decide on configuring IOAM according to their needs. IOAM is intended for deployment in limited domains [RFC8799]. As such, it assumes that a node involved in IOAM operation has previously verified the integrity of the path. Still, operators need to properly secure the IOAM domain to avoid malicious configuration and use, which could include injecting malicious IOAM packets into the domain.

As this document describes new protocol fields within the existing BIER encapsulation, these are similar to the security considerations of [RFC8296].

7. IANA Considerations

In the "BIER Next Protocol Identifiers" registry created for [RFC8296], a new Next Protocol Value for IOAM is requested from IANA as follows:

BIER Next Protocol Identifier	Description	Semantics Definition	Reference
TBA1	In-situ OAM (IOAM)	Section 4	This Document

Table 1: New BIER Next Protocol Identifier

8. Acknowledgements

The authors would like to acknowledge Greg Mirsky for his thorough review and very helpful comments.

9. References

9.1. Normative References

[I-D.ietf-ippm-ioam-data] Brockners, F., Bhandari, S., and T. Mizrahi, "Data Fields for In-situ OAM", Work in Progress, Internet-Draft, draft-ietf-ippm-ioam-data-17, 13 December 2021, <<https://www.ietf.org/archive/id/draft-ietf-ippm-ioam-data-17.txt>>.

[I-D.ietf-ippm-ioam-direct-export] Song, H., Gafni, B., Zhou, T., Li, Z., Brockners, F., Bhandari, S., Sivakolundu, R., and T. Mizrahi, "In-situ OAM Direct Exporting", Work in Progress, Internet-Draft, draft-ietf-ippm-ioam-direct-export-07, 13 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-ippm-ioam-direct-export-07.txt>>.

[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>>.

9.2. Informative References

[I-D.ietf-bier-ping] Kumar, N., Pignataro, C., Akiya, N., Zheng, L., Chen, M., and G. Mirsky, "BIER Ping and Trace", Work in Progress, Internet-Draft, draft-ietf-bier-ping-07, 11 May

2020, <<https://www.ietf.org/archive/id/draft-ietf-bier-ping-07.txt>>.

[I-D.ietf-bier-pmmm-oam] Mirsky, G., Zheng, L., Chen, M., and G. Fioccola, "Performance Measurement (PM) with Marking Method in Bit Index Explicit Replication (BIER) Layer", Work in Progress, Internet-Draft, draft-ietf-bier-pmmm-oam-11, 4 October 2021, <<https://www.ietf.org/archive/id/draft-ietf-bier-pmmm-oam-11.txt>>.

[I-D.ietf-bier-use-cases]

Kumar, N., Asati, R., Chen, M., Xu, X., Dolganow, A., Przygienda, T., Gulko, A., Robinson, D., Arya, V., and C. Bestler, "BIER Use Cases", Work in Progress, Internet-Draft, draft-ietf-bier-use-cases-12, 10 September 2020, <<https://www.ietf.org/archive/id/draft-ietf-bier-use-cases-12.txt>>.

[RFC8799] Carpenter, B. and B. Liu, "Limited Domains and Internet Protocols", RFC 8799, DOI 10.17487/RFC8799, July 2020, <<https://www.rfc-editor.org/info/rfc8799>>.

Authors' Addresses

Xiao Min
ZTE Corp.
Nanjing
China

Email: xiao.min2@zte.com.cn

Zheng(Sandy) Zhang
ZTE Corp.
Nanjing
China

Email: zhang.zheng@zte.com.cn

Yisong Liu
China Mobile
Beijing
China

Email: liuyisong@chinamobile.com

Nagendra Kumar Nainar
Cisco Systems, Inc.
7200-11 Kit Creek Road
Research Triangle Park, NC 27709
United States of America

Email: naikumar@cisco.com

Carlos Pignataro
Cisco Systems, Inc.
7200-11 Kit Creek Road
Research Triangle Park, NC 27709
United States of America

Email: cpignata@cisco.com