

Network Working Group
Internet-Draft
Intended status: Experimental
Expires: 27 August 2025

T. Przygienda
J. Zhang
Juniper Networks
H. Bidgoli
Nokia
I. Wijnands
Individual
23 February 2025

Unmasked BIER Mode
draft-zzhang-bier-unmasked-bier-01

Abstract

The document introduces a new mode of interpretation of the bitmask field in the BIER encoding, called unmasked BIER, that solves the problem of BIER originator targeting receivers across many different sets and hence, in worst case, degrading into ingress replication.

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 27 August 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	2
2. Unmasked BIER	3
3. U-BIER Interpretation of the bitmask field	3
4. Signalling for MPLS in ISIS via U-Mode BIER MPLS Encapsulation Sub-sub-TLV	3
5. Further Considerations	4
6. Security Considerations	4
7. IANA Section	4
8. Contributors	5
9. Acknowledgement	5
10. References	5
10.1. Normative References	5
10.2. Informative References	5
Authors' Addresses	5

1. Introduction

BIER technology [RFC8296] in its standard form may suffer from a degradation towards ingress replication in a scenario where the total number of involved BFRs is very large but most packets address varying combination of relatively few receivers spread across different sets.

More explicitly, BIER carries in its standard form a bitmask of successive BFR receivers represented as bits in a set and thus on large numbers of involved BFRs and few receivers w/o customized allocation of BFR IDs or application awareness of the address locality each receiver on a packet may be found in a different set in most extreme cases. Obviously, when receivers are in different sets, multiple packets must be sent, one per involved set. Such a scenario can become common in large distributed computation environments based on consensus building algorithms (or building closures by successive folding of results), e.g. multi-CPU HPC or data centers and RAID storage where storing data while addressing large number of always changing stripes can lead to this kind of anomaly.

2. Unmasked BIER

To deal with the problem of small number of receivers spread across many sets we introduce into BIER a special interpretation of the bitmask field called unmasked BIER or U-BIER for brevity's sake. U-BIER bitmask is interpreted as sequence of BFR-IDs that can belong to different sets and be interspersed with "holes" consisting of illegal BFR-IDs to increase hardware processing efficiency.

Support for processing of unmasked BIER is indicated by signalling special label or its equivalent analogous to BIER Encapsulation signalling. Based on this info the upstream node can decide on a hop-by-hop basis whether it compresses the same BIER frames it received for different sets (how the recognition of same frames is performed is outside the scope of this document) into a single U-BIER frame or propagates a U-BIER frame with according changes, i.e. only providing the BFR-IDs for which the receiver is responsible. The removal of BFR IDs from the packet before forwarding it is completely analogous to the normal bitmask processing in a sense but here spanning possibly many sets in the lookup engine. In case the downstream receiver does not support U-BIER the frame needs to be replicated into according standard BIER frames with usual bitmask interpretation.

3. U-BIER Interpretation of the bitmask field

The bitmask field MUST be interpreted in U-BIER mode as sequence of BFR-IDs (as an example, within 256 bits maximum 16 receivers can be encoded) whereas illegal BFR-IDs are allowed in any position and MUST be ignored. Duplicates are also allowed and MUST be treated as a single occurrence. No guarantees are provided in terms of sorting of the BFR-IDs in U-BIER.

4. Signalling for MPLS in ISIS via U-Mode BIER MPLS Encapsulation Sub-sub-TLV

The signalling follows the BIER MPLS Encapsulation Sub-sub-TLV, i.e. it is <MT,SD,BML> specific but obviously does not include the concept of a set.

```

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 |
+-----+-----+-----+-----+
|  Reserved  | BS Len | Label |
+-----+-----+-----+-----+
| ...

```

Type: TBD1

Length: 4 * of U-Mode BitString Length specific labels.

Reserved: MUST be 0 on send and ignored on reception.

Label: Frames received on this label interpret the BitMask as U-BIER.

In case multiple values are advertised for same bitstring length, the sub-sub-TLV MUST be ignored. In case same label value is advertised for different bitstring lengths, the sub-sub-TLV MUST be ignored.

Label values MUST NOT match any of the reserved values defined in [RFC3032].

5. Further Considerations

U-BIER can be deployed within existing BIER node by node, either as addition or the only mode supported. Observe that in an environment where only some nodes support U-BIER those may have to disaggregate U-BIER into normal BIER and hence fall back to average replication. In a network where some nodes support only U-BIER normal BIER nodes may consider them not supporting BIER at all or may perform translation from normal BIER into U-BIER frames, however, it must be expected that they will use a single frame and basically send U-BIER with the BFR IDs of a single set only, not rendering any gains unless a mechanism is in place the same frame is detected sent on multiple sets and coalesced into a single U-BIER frame. This can be achieved by many mechanisms out the scope of this draft, amongst them a shim after the BIER encapsulation [I-D.zhang-bier-extension-headers] in way similar to inband telemetry or fragmentation support.

Mixing U-BIER and normal BIER within a subdomain can lead to difficulties of decoding the according packets in the middle in terms of destinations targeted if the involved parties in the middle do not possess the correct binding information. An obvious method to deal with this problem is to deploy U-BIER specifically in its own sub-domain where the interpretation of the bitmask is unambiguous.

6. Security Considerations

TBD

7. IANA Section

TBD

8. Contributors

TBD

9. Acknowledgement

Michael Menth and Toerless Eckert mentioned the problem and approaches to address it in open forums for the first time.

10. References

10.1. Normative References

- [RFC3032] Rosen, E., Tappan, D., Fedorkow, G., Rekhter, Y., Farinacci, D., Li, T., and A. Conta, "MPLS Label Stack Encoding", RFC 3032, DOI 10.17487/RFC3032, January 2001, <<https://www.rfc-editor.org/info/rfc3032>>.
- [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>>.

10.2. Informative References

- [I-D.zzhang-bier-extension-headers] Zhang, Z. J., Min, X., Liu, Y., and H. Bidgoli, "BIER Extension Headers", Work in Progress, Internet-Draft, draft-zzhang-bier-extension-headers-03, 25 February 2024, <<https://datatracker.ietf.org/doc/html/draft-zzhang-bier-extension-headers-03>>.

Authors' Addresses

Tony Przygienda
Juniper Networks
Email: prz@juniper.net

Jeffrey (Zhaohui) Zhang
Juniper Networks
Email: zzhang@juniper.net

Hooman Bigdoli
Nokia
Email: hooman.bigdoli@nokia.com

IJsbrand Wijnands
Individual
Email: ice@braindump.be