

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 9 August 2025

H. Bidgoli, Ed.
Nokia
I. Wijnands
M. Mishra
Cisco System
Z. Zhang
Juniper Networks
E. Leyton
Verizon
5 February 2025

M-LDP Signaling Through BIER Core
draft-ietf-bier-mlbp-signaling-over-bier-04

Abstract

Consider an end-to-end Multipoint LDP (mLDP) network, where it is desirable to deploy BIER in a segment of this network. It might be desirable to deploy BIER with minimal disruption to the mLDP network or a redesign of the network.

This document describes a procedure needed for mLDP tunnels to be signaled over and stitched through a BIER core, allowing LDP routers to run traditional mLDP services through a BIER core.

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 9 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. Conventions used in this document	3
2.1. Definitions	3
3. mLDP Signaling Through BIER domain	4
3.1. Ingress BBR procedure	5
3.1.1. Automatic tLDP session Creation	5
3.1.2. ECMP Method on IBBR	5
3.2. Egress BBR procedure	5
3.2.1. IBBR procedure for arriving upstream assigned label	6
3.2.2. BIER Interface ID SUB-TLVs	6
4. Datapath Forwarding	7
4.1. Datapath traffic flow	7
5. Recursive FEC	7
6. IANA Consideration	7
7. Security Considerations	8
8. Acknowledgments	8
9. References	8
9.1. Normative References	8
9.2. Informative References	9
Authors' Addresses	9

1. Introduction

Some operators that are using mLDP P2MP LSPs for their multicast transport would like to deploy BIER technology in some segments of their network. As an example, if a operator is simplifying their core network with IGP underlay and BGP overlay and migrating from LDP to Segment Routing, they might prefer BIER in this segment of network for Multicast technology as BIER uses IGP underlay. The operator might still want to legacy multicast protocols like mLDP on the access of the network and offer end to end multicast services. This draft explains a method to signal mLDP services through a BIER domain, with minimal disruption and operational impact to the mLDP domain.

2. Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

2.1. Definitions

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.

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.

BFIR:

Bit-Forwarding Ingress Router, The ingress border router that inserts the Bit Map into the packet. Each BFIR must have a valid BFR-id assigned. BFIR is a term used for data plane packet forwarding.

BFER:

Bit-Forwarding Egress Router, A router that participates in Bit Index Forwarding as leaf. Each BFER must have a valid BFR-id assigned. BFER is a term used for data plane packet forwarding.

BBR:

BIER Boundary router. The router between the LDP domain and BIER domain.

IBBR:

Ingress BIER Boundary Router. The ingress router from a signaling point of view. It maintains mLDP adjacency toward the LDP domain and determines if the Multipoint LDP FEC needs to be signaled across the BIER domain via Targeted LDP.

EBBR:

Egress BIER Boundary Router. The egress router in a BIER domain from signaling point of view. It terminates the targeted LDP signaling through the BIER domain. It also keeps track of all IBBRs that are part of this P2MP tree

BIFT:

Bit Index Forwarding Table.

BIER sub-domain:

A further distinction within a BIER domain identified by its unique sub-domain identifier. A BIER sub-domain can support multiple BitString Lengths.

BFR-id:

An optional, unique identifier for a BFR within a BIER sub-domain, all BFERS and BFIRS need to be assigned a BFR-id.

ILM:

MPLS Incoming Label Map.

3. mLDP Signaling Through BIER domain

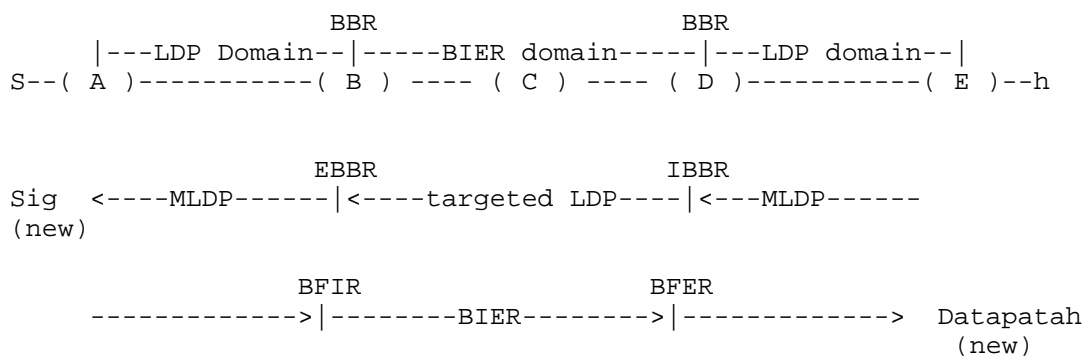


Figure 1: BIER boundary router

As per above figure, point-to-multipoint and multipoint-to-multipoint LSPs established via mLDP [RFC6388] can be signaled through a bier domain via Targeted LDP sessions. The procedure of using LDP Multipoint Extensions on Targeted LDP session is explained in [RFC7060].

This document provides additional details and defines some needed procedures to use LDP Multipoint Extensions on Targeted LDP as per [RFC7060] to connect two or more mLDP domains over a BIER domain.

3.1. Ingress BBR procedure

In above Figure, the Ingress BBR (IBBR) is connected to the mLDP domain on downstream and a BIER domain on the upstream. To connect the LDP domains via the BIER domain, IBBR needs to establish a targeted LDP session with EBBR closest to the root of the P2MP or MP2MP LSP. To do so IBBR will follow procedures in [RFC7060] in particular section 6 "Targeted mLDP with Multicast Tunneling".

The Target LDP session can be established manually via configuration or via automated mechanism.

3.1.1. Automatic tLDP session Creation

tLDP sessions can be signaled automatically from every IBBR to the appropriate EBBR. When mLDP FEC arrives at the IBBR from LDP domain. IBBR can automatically start a tLDP session to the EBBR closest to the root node. Both IBBR and EBBR should be in auto-discovery mode and react to the arriving tLDP signaling packets (i.e. targeted hellos, keep-alives etc...) to establish the session automatically.

The root node address in the mLDP FEC can be used to find the EBBR. To identify the EBBR, the same procedures as [RFC7060] section 2.1 can be used or the procedures as explained in the [draft-ietf-bier-pim-signaling] appendix A.

3.1.2. ECMP Method on IBBR

If the IBBR finds multiple equal cost EBBRs on the path to the root, it can use a vendor specific algorithm to choose between the EBBRs. These algorithms are beyond the scope of this draft. As an example the IBBR can use the lowest EBBR IP address to establish its mLDP signaling to.

3.2. Egress BBR procedure

The Egress BBR (EBBR) is connected to the upstream mLDP domain. The EBBR should accept the tLDP session generated from the IBBR. It should assign a unique "upstream assigned label" for each arriving FEC generated by the IBBRs.

The EBBR should follow the [RFC7060] procedures with the following modifications:

- * The label assigned by the EBBR cannot be Implicit Null. This is to ensure that the identity of each p2mp and/or mp2mp tunnel in the BIER domain is uniquely distinguished.
- * The label can be assigned from a domain-wide Common Block (DCB) [RFC9573]
- * The Interface ID TLV, as per [RFC6389] should includes a new BIER sub-tlv as tunnel identifier.

The EBBR will also generate a new label and FEC toward the root in the LDP domain. The EBBR Should stitch this generated label with the "upstream assigned label" to complete the P2MP or MP2MP LSP.

With the same token the EBBR should track all the arriving FECs and the IBBRs that are generating these FECs. The EBBR will use this information to build the BIER header for each set of common FEC arriving from the IBBRs. The EBBR will use the BFR-ID of the IBBRs to build the BIER header.

3.2.1. IBBR procedure for arriving upstream assigned label

Upon receiving the "upstream assigned label", the IBBR should create its own stitching instruction between the "upstream assigned label" and the downstream signaled label.

3.2.2. BIER Interface ID SUB-TLVs

As per [RFC6389] when LDP is used for upstream label assignment, the Interface ID TLV is used for signaling the Tunnel Identifier and it carries sub-TLVs. This document defines two new Interface ID sub-TLVs for BIER.

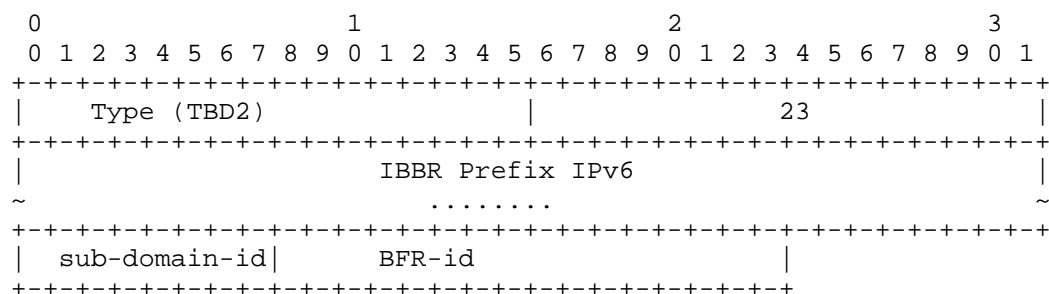
Below is the Interface ID BIER sub-TLV for IPv4 BIER prefix:

```

      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 (TBD1)      |      15      |
+-----+-----+-----+-----+-----+-----+-----+-----+
|      IBBR Prefix IPv4      |
+-----+-----+-----+-----+-----+-----+-----+-----+
| sub-domain-id | BFR-id |
+-----+-----+-----+-----+-----+-----+-----+

```

Below is the Interface ID BIER sub-TLV for IPv6 BIER prefix:



4. Datapath Forwarding

4.1. Datapath traffic flow

On the BFIR when the MPLS label for P2MP/MP2MP LSP arrives from upstream, a lookup in the ILM table is done and the label is swapped with the tLDP upstream assigned label. The BFIR will note all the BFERs that are interested in specific P2MP/MP2MP LSP (as per section 3.2). The BFIR will put the corresponding BIER header with the bit index set for all the IBBRs interested in this stream. The BFIR builds this bit index from the BFR-ID of the interested IBBRs. The BFIR will set the BIERHeader.Proto = MPLS and will forward the BIER packet into the BIER domain.

In the BIER domain, normal BIER forwarding procedure will be done, as per [RFC8279]

The BFERs will receive the BIER packet and will look at the protocol field of BIER header, indicating MPLS protocol. The BFER will remove the BIER header and will do a lookup in the ILM table for the upstream assigned label and perform its corresponding action.

It should be noted that these procedures are also valid if BFIR is the ILER and/or BFER is the ELER as per [RFC7060]

5. Recursive FEC

The procedures above are also valid for mLDP recursive FEC. The root used to determine the EBBR is the outer root of the FEC. The entire recursive FEC needs to be preserved when it is forwarded via tLDP and the label request.

6. IANA Consideration

IANA maintains a registry of Interface ID Types for use in GMPLS in the registry "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Parameters" and sub-registry "Interface_ID Types"

This document defines and requests two new Interface_ID Type for BIER from the Interface_ID Types space,

- * BIER IPv4 prefix TLV (Value TBD)

- * BIER IPv6 prefix TLV (Value TBD)

7. Security Considerations

While in the BIER domain the security considerations of [RFC8279] are relevant to this document.

The implementation should also take into account the security recommendations of [RFC6389].

8. Acknowledgments

Authors would like to acknowledge Jingrong Xie and Nabeel Cocker for his comments and help on this draft.

9. References

9.1. Normative References

- [draft-ietf-bier-pim-signaling]
"H. Bidgoli, F. Xu, J. Kotalwar, I. Wijnands, M. Mishra, Z. Zhang", 29 July 2020.
- [RFC2119] "S. Bradner, "Key words for use in RFCs to Indicate Requirement Levels"", March 1997.
- [RFC6388] "I.J. Wijnands, I. Minei, K. Kompella, B.Thomas "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint LSP"". "
- [RFC6389] "R Aggarwal, JL. Le Roux, "MPLS Upstream Label Assignment for LDP"", November 2011.
- [RFC7060] "M. Napierala, E. Rosen, I. Wijnands", November 2013.
- [RFC8279] "I. Wijnands, E. Rosen, A. ADolganow, T. Prygienda, S. Aldrin", November 2017.
- [RFC9573] "Z. Zhang, E. Rosen, W. Lin, Z. Li, I. Wijnands, "MVPN/ EVPN Tunnel Aggregation with Common Labels"", 27 April 2018.

9.2. Informative References

- [RFC8401] "Ginsberg, L., Przygienda, T., Aldrin, S., and Z. Zhang, "BIER Support via ISIS", June 2018.
- [RFC8444] "Psenak, P., Kumar, N., Wijnands, IJ., Dolganow, A., Przygienda, T., Zhang, Z., and S. Aldrin, "OSPF Extensions for Bit Index Explicit Replication", June 2018.
- [RFC8556] "Rosen, E., Ed., Sivakumar, M., Wijnands, IJ., Aldrin, S., Dolganow, A., and T. Przygienda, "Multicast VPN Using Index Explicit Replication (BIER)", April 2018.

Authors' Addresses

Hooman Bidgoli (editor)
Nokia
Ottawa
Canada
Email: hooman.bidgoli@nokia.com

IJsbrand Wijnands
Cisco System
Diegem
Belgium
Email: ice@cisco.com

Mankamana Mishra
Cisco System
Milpitas,
United States of America
Email: mankamis@cisco.com

Zhaohui Zhang
Juniper Networks
Boston,
United States of America
Email: zzhang@juniper.com

Eddie Leyton
Verizon
Email: Edward.leyton@verizonwireless.com