

**Multipoint LDP Extension for P2MP-based BIER**  
**draft-xie-mpls-ldp-bier-extension-01**

Abstract

Bit Index Explicit Replication (BIER) is a new architecture that provides optimal multicast forwarding without requiring intermediate routers to maintain any per-flow state by using a multicast-specific BIER header. An extension to the Label Distribution Protocol (LDP) defined in [\[RFC5036\]](#) for the setup of point-to-multipoint (P2MP) is described in [\[RFC6388\]](#) is called mLDP, and is used for multicast-specific tree building. This document describes mLDP extensions to setup a P2MP with BIER information, which is called P2MP based BIER in [\[I-D.xie-bier-mvpn-mpls-p2mp\]](#).

Requirements Language

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 [\[RFC2119\]](#).

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 March 9, 2019.

## Copyright Notice

Copyright (c) 2018 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

## Table of Contents

<a href="#">1.</a>	<a href="#">Introduction</a>	<a href="#">2</a>
<a href="#">2.</a>	<a href="#">Terminology</a>	<a href="#">3</a>
<a href="#">3.</a>	<a href="#">MLDP P2MP BIER Signalling</a>	<a href="#">4</a>
<a href="#">3.1.</a>	<a href="#">Definitions</a>	<a href="#">4</a>
<a href="#">3.2.</a>	<a href="#">Example</a>	<a href="#">4</a>
<a href="#">3.3.</a>	<a href="#">Signaling the P2MP BIER</a>	<a href="#">5</a>
<a href="#">3.4.</a>	<a href="#">The P2MP BIER LSP Identifier</a>	<a href="#">5</a>
<a href="#">3.5.</a>	<a href="#">BIER TLV</a>	<a href="#">6</a>
<a href="#">3.6.</a>	<a href="#">Make Before Break (MBB)</a>	<a href="#">7</a>
<a href="#">4.</a>	<a href="#">Capability and Error handling</a>	<a href="#">7</a>
<a href="#">4.1.</a>	<a href="#">BIER Capability</a>	<a href="#">7</a>
<a href="#">4.2.</a>	<a href="#">BIER Status</a>	<a href="#">8</a>
<a href="#">4.3.</a>	<a href="#">Check Capability and Use Status for Error Handling</a>	<a href="#">9</a>
<a href="#">5.</a>	<a href="#">IANA Considerations</a>	<a href="#">11</a>
<a href="#">6.</a>	<a href="#">Security Considerations</a>	<a href="#">11</a>
<a href="#">7.</a>	<a href="#">Acknowledgements</a>	<a href="#">12</a>
<a href="#">8.</a>	<a href="#">References</a>	<a href="#">12</a>
<a href="#">8.1.</a>	<a href="#">Normative References</a>	<a href="#">12</a>
<a href="#">8.2.</a>	<a href="#">Informative References</a>	<a href="#">13</a>
	<a href="#">Authors' Addresses</a>	<a href="#">13</a>

## [1. Introduction](#)

Bit Index Explicit Replication (BIER) is a new architecture that provides optimal multicast forwarding without requiring intermediate routers to maintain any per-flow state by using a multicast-specific BIER header. An extension to the Label Distribution Protocol (LDP) defined in [\[RFC5036\]](#) for the setup of point-to-multipoint (P2MP) is described in [\[RFC6388\]](#) is called mLDP, and is used for multicast-specific tree building. This document describes mLDP extensions to



setup a P2MP with BIER information, which is called a P2MP based BIER in [[I-D.xie-bier-mvpn-mpls-p2mp](#)] .

Related documents that may be of interest include [[RFC5561](#)], and [[RFC3988](#)].

## 2. Terminology

Readers of this document are assumed to be familiar with the terminology and concepts of the documents listed as Normative References. For convenience, some of the more frequently used terms and new terms list below.

- o mLDP: Multipoint extensions for LDP.
- o P2MP LSP: An LSP that has one Ingress LSR and one or more Egress LSRs.
- o Ingress LSR: An Ingress LSR for a particular LSP is an LSR that can send a data packet along the LSP. MP2MP LSPs can have multiple Ingress LSRs, P2MP LSPs have just one, and that node is often referred to as the "root node".
- o Egress LSR: An Egress LSR for a particular LSP is an LSR that can remove a data packet from that LSP for further processing. P2P and MP2P LSPs have only a single egress node, but P2MP and MP2MP LSPs can have multiple egress nodes.
- o Transit LSR: An LSR that has reachability to the root of the MP LSP via a directly connected upstream LSR and one or more directly connected downstream LSRs.
- o Bud LSR: An LSR that is an egress but also has one or more directly connected downstream LSRs.
- o Leaf node: A leaf node can be either an Egress or Bud LSR when referred to in the context of a P2MP LSP.
- o FEC: Forwarding Equivalence Class
- o P2MP FEC: defined in [RFC6388](#).
- o F-BM: Forwarding Bit Mask
- o BSL: Bit String Length, that is 64, 128, 256, etc (per [[RFC8279](#)]).

### 3. MLDP P2MP BIER Signalling

#### 3.1. Definitions

F-BM: Forwarding Bit Mask, An array of Bit, which is defined in [\[RFC8279\]](#).

Peer F-BM: For LSR A and P2MP FEC<Root,N>, this is the F-BM that included in Label Mapping from a downstream LSR for P2MP FEC<Root,N> to A.

Downstream F-BM: For LSR A and P2MP FEC<Root,N>, this is the OR'ing result of each of the F-BM included in Label Mapping from downstream LSR for P2MP FEC<Root,N> to A. A use this Downstream F-BM in its Label Mapping to upstream node for P2MP FEC<Root,N>. In other words, A's Downstream F-BM is A's upstream node's Peer F-BM.

#### 3.2. Example

Consider LSRs A - F, interconnected as follows:




Figure 1: P2MP-based BIER Topology

Say that the node D has a BFR-id 1, F has a BFR-id 2, and E has a BFR-id 3, and we use a Bit String Length 4 (which is not valid per [\[RFC8296\]](#)) as an example.

Consider an P2MP FEC<Root=A,N=10> for which A is the Root, and say that D,E,F are all the Leafs of this P2MP FEC<Root=A, N=10>.

While D send LDP Mapping to C, it includes a F-BM 0001. Say that C got a Peer<D> F-BM 0001, and then C form a Downstream F-BM 0001.

While F send LDP Mapping to C, it includes a F-BM 0010. Say that C got a Peer<F> F-BM 0010, and then C form a Downstream F-BM 0011.

While C send LDP Mapping to B, it includes a F-BM 0010. Say that B got a Peer<C> F-BM 0011, and then B form a Downstream F-BM 0011.

While E send LDP Mapping to B, it includes a F-BM 0100. Say that B got a Peer<E> F-BM 0100, and then B form a Downstream F-BM 0111.

While B send LDP Mapping to A, it includes a F-BM 0111. Say that A got a Peer<B> F-BM 0111, and then A form a Downstream F-BM 0111.

This memo describes how every nodes in a P2MP tree get Peer F-BM from every downstream LSR, form a Downstream F-BM by OR'ing all it's Downstream Peer F-BM, and send a Mapping to it's upstream node using the Downstream F-BM.

### **3.3. Signaling the P2MP BIER**

The procedure for signalling the P2MP BIER is performed hop-by-hop by each LSR L along an P2MP LSP for a given P2MP FEC<Root,N>. The steps are as follows:

1. First, L computes its Downstream F-BM for P2MP FEC<Root,N>:

If L is a leaf for P2MP FEC<Root,N>, L sets the F-BM with it's BFRID's BitPosition to 1.

Otherwise (L is not a Leaf), L computes the Downstream F-BM by OR'ing all it's downstream Node's F-BM, as described above.

2. For each LDP neighbor of L to which L decides to send a Mapping for FEC F, L attaches an BIER TLV with the F-BM that it computed for this P2MP FEC.

3. When a new BIER TLV is received for P2MP FEC<Root,N> from a downstream LSR or the set of downstream LSRs, L returns to step 1. If the newly computed Downstream F-BM is unchanged, L SHOULD NOT advertise new information to its upstream neighbor. Otherwise, L readvertises its Mappings to its upstream neighbor with an updated BIER TLV.

This behavior is standard for attributes such as path vector, hop count, and MTU, and the same rules apply, as specified in [[RFC5036](#)].

If the Downstream F-BM changes, L MAY readvertise the new F-BM immediately, or hold down the readvertisement for a while.

### **3.4. The P2MP BIER LSP Identifier**

[RFC6388] defined the P2MP FEC Element, which include a LDP MP Opaque Value Element. It also defined a Generic LSP Identifier as the LDP MP Opaque Value Element, with a TLV of Type 1. This document defined a new type of LDP MP Opaque Value Element, called the P2MP BIER LSP Identifier.

The encoding for the P2MP BIER LSP Identifier is as follows:



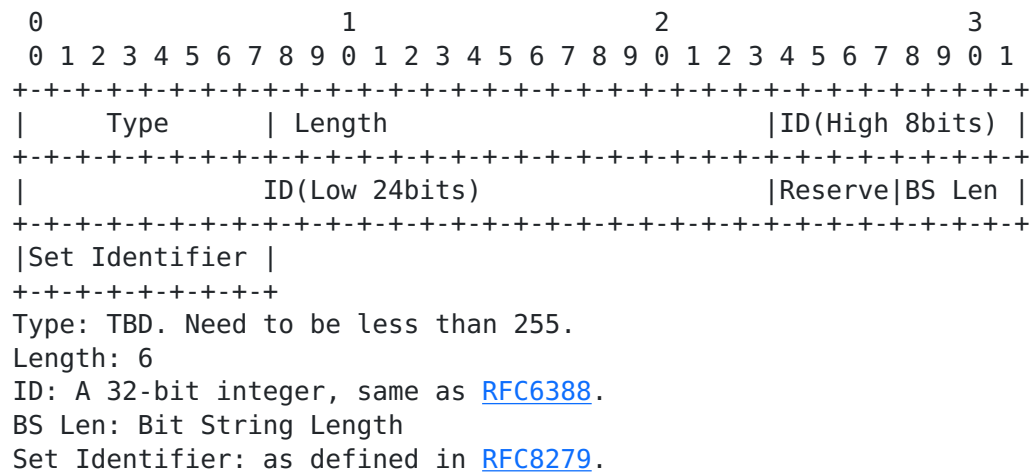


Figure 2: P2MP BIER LSP Identifier

### 3.5. BIER TLV

The BIER TLV encodes information on the F-BM for an LSP, from the advertising LSR to the egress(es) over all valid paths.

The encoding for the BIER TLV is as follows:

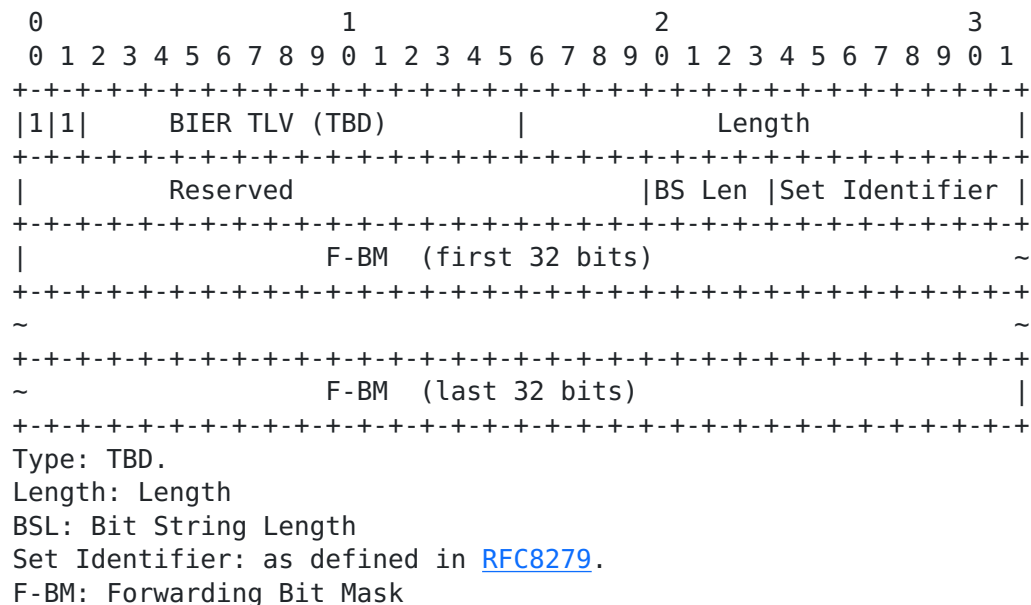


Figure 3: BIER TLV





### 3.6. Make Before Break (MBB)

The Make Before Break (MBB) mechanism for mLDP P2MP defined in [RFC6388](#) also applies.

## 4. Capability and Error handling

The extensions defined in this document utilize the existing LDP error handling defined in [[RFC5036](#)]. If an LSR receives an error notification from a peer for a session, it terminates the LDP session by closing the TCP transport connection for the session and discarding all multi-topology label mappings learned via the session.

An LSR should respond with an LDP MP Status in LDP Notification Messages when it receives an LDP Label Mapping message with a P2MP FEC element specifying an BIER TLV that is not locally known or not supported. The LSR MUST also discard the entire message before sending the Notification message.

### 4.1. BIER Capability

A new optional capability parameter TLV, RMR Capability, is defined. Following is the format of the RMR Capability Parameter:

```

      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
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|U|F| BIER Capability (IANA) | Length (= 2) |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|S| Reserved |P|D|I|R| Rsv |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
BIER Capability TLV Format
U: set to 1. Ignore, if not known.
F: Set to 0. Do not forward.
S: MUST be set to 1 to advertise the BIER TLV.
P: The node has BIER P-Capability.
D: The node has BIER D-Capability.
I: The node Ignore the BIER Header except the Label.
R: The node Require a packet without BIER Header except the Label.
```

Figure 4: BIER Capability

The BIER Capability TLV MUST be advertised in the LDP Initialization message. If the peer has not advertised the BIER capability, then label messages including a BIER TLV MUST NOT be sent to the peer.

If an LSR has not advertised that it is BIER capable, its LDP peers MUST NOT send it messages that include BIER TLV. If an LSR receives



a Label Mapping message with an BIER TLV from downstream LSR-D and its upstream LSR-U has not advertised that it is BIER capable, the LSR MUST send an BIER notification immediately to LSR-D. If this happens, an P2MP BIER LSP will not be established, a normal P2MP LSP will not be established either.

P-Capability indicate a complete BIER function, which includes P-Capability and D-Capability. If a node support P-Capability, then it support the whole BIER function, which means it support both P-capability and D-capability.

D-Capability indicate a subset of BIER function, to Disposition some length of a packet from some offset. For example, from BIER Label and a whole ( BIER Header Length ) , or from the position after BIER Label and a length of ( BIER Header Length - BIER Label Length ) . If a node don't support P-Capability, it may still support D-Capability. If a node don't support D-capability, it must not support P-Capability.

If a Node doesn't have P-Capability, then P flag MUST be cleared. Whether the node will be a Branch or BUD or Leaf, the I flag SHOULD be set.

if a node doesn't have D-Capability, then P and D flag MUST be cleared. If the node will be a BUD or Leaf then R flag SHOULD be set. if the node will be a Branch then R flag MAY not be set.

if a node doesn't have P-Capability but does have D-Capability, then D flag SHOULD be set, but R flag MAY be set or not be set.

#### **4.2. BIER Status**

A new optional LDP MP Status Value Element ([RFC6388](#)), BIER Status, is defined. Following is the format of the BIER Status:

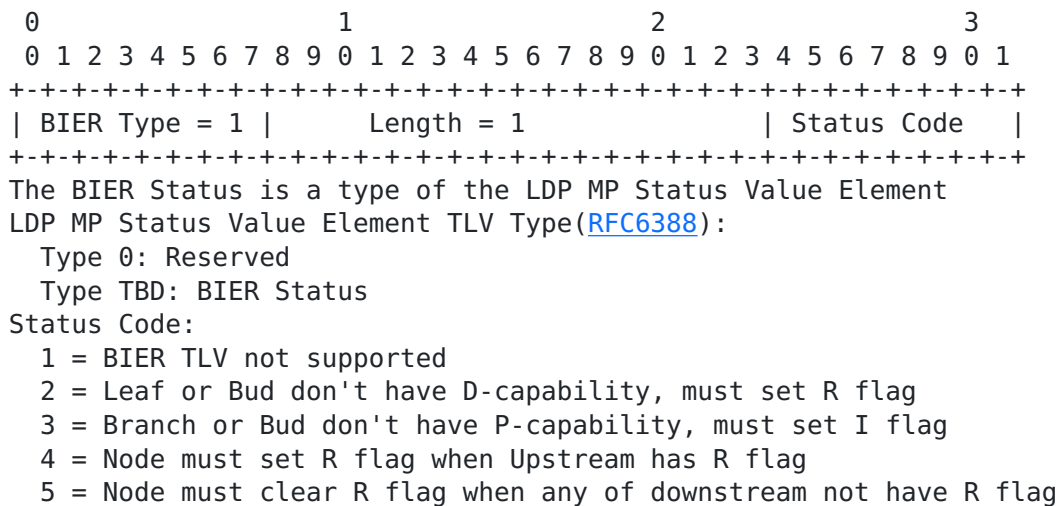


Figure 5: BIER Status TLV

#### 4.3. Check Capability and Use Status for Error Handling

In order to deploy P2MP BIER on a network with some legacy nodes, which may not support P-capability or D-capability, some Capability flags need to be checked and notification messages may be generated.

If all edge nodes support D-capability, but some nodes don't support P-capability and they set a I flag as required, then deployment of P2MP BIER is fine, and a BIER packet can walk from Root to all Leaf(s) without any change except the BIER Label.

If an LSR support P-capability, and it's upstream node also support P-capability, and when the LSR receives a Label Mapping message with an BIER TLV and R flag set from downstream LSR-D, the LSR will accept such Label Mapping message. If receives a Label Mapping wiht an BIER TLV and R flag cleaned another downstream LSR-D', the LSR will accept too.

If an LSR receives a Label Mapping message with an BIER TLV from downstream LSR-D with a R flag, and its upstream LSR-U has no P-capability, the LSR MUST send an BIER notification immediately to LSR-D. If this happens, an P2MP BIER LSP will not be established, a normal P2MP LSP will not be established either.

When an LSR receives a Label Mapping message, it need to do some check before process and build P2MP BIER forwarding table. Such check includes:

- 1) Check if the node's P-capability and D-capability conflict with it's I flag and R flag.



The following table list the whole check matrix.

Node's Role	Node's PDIR-flag	Check Result	Rule
Leaf	[PDxx]	OK	(*) Comment 1
Leaf	[-Dxx]	OK	
Leaf	[--xR]	OK	Rule 1
Leaf	[--x-]	ERR	Rule 1
Branch	[PDxx]	OK	(*) Comment 1
Branch	[-DIx]	OK	Rule 2
Branch	[-D-x]	ERR	Rule 2
Branch	[--Ix]	OK	(*) Comment 2
Branch	[---x]	ERR	Rule 2
BUD	[PDxx]	OK	(*) Comment 1
BUD	[-DIx]	OK	Rule 2
BUD	[-D-x]	ERR	Rule 2
BUD	[--IR]	OK	
BUD	[--I-]	ERR	Rule 1
BUD	[---R]	ERR	Rule 2
BUD	[----]	ERR	Rule 1 & 2

(\*) Comment 1: In some cases, set a R flag is useful  
 (\*) Comment 2: In some cases, Clear R flag on a Branch node is useful  
 Rule 1: Leaf don't have D-capability, must set R flag  
 Rule 2: Branch don't have P-capability, must set I flag

Figure 6: BIER Self Check Matrix

2) Check the node's R flag, the node's upstream R flag, the node's downstream R flag.

The following table list the whole check martrix about the R flag of node, the upstream node, the downstream nodes.









## 7. Acknowledgements

TBD

## 8. References

### 8.1. Normative References

- [I-D.ietf-bier-mvpn]  
Rosen, E., Sivakumar, M., Aldrin, S., Dolganow, A., and T. Przygienda, "Multicast VPN Using BIER", [draft-ietf-bier-mvpn-11](#) (work in progress), March 2018.
- [I-D.xie-bier-mvpn-mpls-p2mp]  
Xie, J., McBride, M., Chen, M., and L. Geng, "Multicast VPN Using MPLS P2MP and BIER", [draft-xie-bier-mvpn-mpls-p2mp-02](#) (work in progress), July 2018.
- [RFC3988] Black, B. and K. Kompella, "Maximum Transmission Unit Signalling Extensions for the Label Distribution Protocol", [RFC 3988](#), DOI 10.17487/RFC3988, January 2005, <<https://www.rfc-editor.org/info/rfc3988>>.
- [RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", [RFC 5036](#), DOI 10.17487/RFC5036, October 2007, <<https://www.rfc-editor.org/info/rfc5036>>.
- [RFC5561] Thomas, B., Raza, K., Aggarwal, S., Aggarwal, R., and JL. Le Roux, "LDP Capabilities", [RFC 5561](#), DOI 10.17487/RFC5561, July 2009, <<https://www.rfc-editor.org/info/rfc5561>>.
- [RFC6388] Wijnands, IJ., Ed., Minei, I., Ed., Kompella, K., and B. Thomas, "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", [RFC 6388](#), DOI 10.17487/RFC6388, November 2011, <<https://www.rfc-editor.org/info/rfc6388>>.
- [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>>.

## 8.2. Informative 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>>.

### Authors' Addresses

Jingrong Xie  
Huawei Technologies  
Q15 Huawei Campus, No.156 Beiqing Rd.  
Beijing 100095  
China

Email: xiejingrong@huawei.com

Mach Chen  
Huawei

Email: mach.chen@huawei.com

Robin Li  
Huawei

Email: lizhenbin@huawei.com