

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 6, 2018

J. Xie
Y. Liu
M. McBride
Huawei Technologies
March 5, 2018

PIM Extensions for P2MP-based BIER
draft-xie-pim-bier-extension-00

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. PIM is a well-known multicast-specific routing protocol which is widely deployed either in a VRF context or in a Non-VRF context. This document describes PIM extensions to signal a P2MP Tree with BIER information, which is called a P2MP based BIER in [I.D.xie-bier-mvpn-mpls-p2mp]. PIM is required to alloc Label to build a P2MP tree hop-by-hop, and build a P2MP based BIER forwarding table further. This requires a BitMask being carried as a PIM Join Attribute by downstream node to upstream node hop-by-hop, and the behavior is like precedures as specified in [[RFC6807](#)].

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 September 6, 2018.

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

1.	Introduction	2
2.	Terminology	3
3.	PIM Signaling a P2MP BIER tree	3
3.1.	Example of signaling the P2MP-BIER	3
3.2.	BIER-Supported Hello Option	5
3.3.	New BIER F-BM Join Attribute Format	6
4.	How to Use BIER F-BM Join Attribute	7
5.	Capability and Error Handling	8
6.	IANA Considerations	8
7.	Security Considerations	8
8.	Acknowledgements	8
9.	References	8
9.1.	Normative References	8
9.2.	Informative References	9
	Authors' Addresses	9

[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. PIM defined in [[RFC7761](#)] is a well-known multicast-specific routing protocol which is widely deployed either in a VRF context or in a Non-VRF context. This document describes PIM extensions to signal a P2MP Tree with BIER information, which is called a P2MP based BIER in [[I-D.xie-bier-mvpn-mpls-p2mp](#)], in which

PIM is required to alloc Label to build a P2MP tree hop-by-hop, and build a P2MP based BIER forwarding table further. This requires a BitMask being carried as a PIM Join Attribute by downstream node to upstream node hop-by-hop, and the behavior is like precedures as specified in [[RFC6807](#)].

This document defines support for MPLS encapsulation as specified in [[RFC8296](#)]. Support for other encapsulation types is outside the scope of this document. The use of multiple encapsulation types is outside the scope of this document.

[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 BFR: BIER Forwarding Router
- o BFR-ID: BIER Forwarding Router IDentify.
- o P2MP: Point to Multi-point
- o P2MP based BIER: BIER using P2MP as topology
- o F-BM: Forwarding Bit Mask
- o BSL: Bit String Length, that is 64, 128, 256, etc (per [[RFC8279](#)]).

[3.](#) PIM Signaling a P2MP BIER tree

[3.1.](#) Example of signaling the P2MP-BIER

Consider LSRs A - F, interconnected as follows:

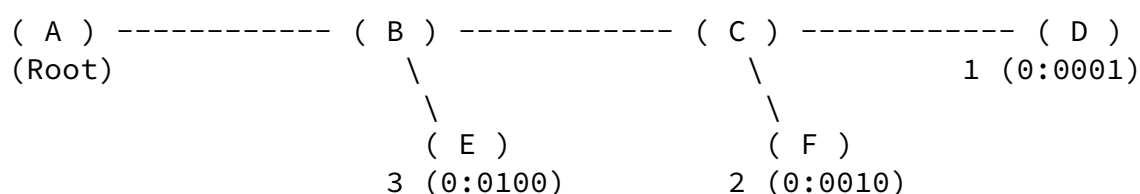


Figure 1: P2MP-based BIER Topology

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 a target PIM tree identified by $\langle S=\text{RootAddress}, G \rangle$, for which A is the Root, and say that D,E,F are all the Leafs of this PIM tree.

When D join the PIM tree, it alloc a Label, and bring this label with a F-BM of 0001 in a PIM Join attribute, and send it to the upstream node C.

When F join the PIM tree, it alloc a Label, and bring this label with a F-BM of 0010 in a PIM Join attribute, and send it to the upstream node C.

When E join the PIM tree, it alloc a Label, and bring this label with a F-BM of 0100 in a PIM Join attribute, and send it to the upstream node B.

When C get the PIM join messages from D and F, then C will establish a PIM state (S,G) and one or many downstream states, C also establish a PIM upstream state and send PIM Join to its upstream neighbor B, with a new allocated Label and a F-BM of 0011.

When B get the PIM join messages from E and C, then B will establish a PIM state (S,G) and one or many downstreams, B also establish a PIM upstream state and send PIM Join to it's upstream neighbor A, with a new allocated Label and a F-BM of 0111.

When A get the PIM join message from B, A will establish a PIM state (S,G) and the downstream(s).

Each node of the PIM tree will establish a routing state of PIM (S,G), and a forwarding state of P2MP based BIER. Here we list the forwarding state of P2MP based BIER on every node.

A:

NHLFE (TreeID, OutInterface<to B>, OutLabel<alloc by B>,
F-BM=0111)

B:

ILM(inLabel<alloc by B>, action<Replication to TreeID>,
flag=CheckBS|Branch, BSL)

NHLFE (TreeID, OutInterface<to C>, OutLabel<alloc by C>,
F-BM=0011)

NHLFE (TreeID, OutInterface<to E>, OutLabel<alloc by E>,
F-BM=0100)

C:

ILM(inLabel<alloc by C>, action<Replication to TreeID>,
flag=CheckBS|Branch, BSL)

NHLFE (TreeID, OutInterface<to D>, OutLabel<alloc by D>,
F-BM=0001)

NHLFE (TreeID, OutInterface<to F>, OutLabel<alloc by F>,
F-BM=0100)

E:

ILM(inLabel<alloc by E>, action<Replication to TreeID>,
flag=CheckBS|Leaf, BSL)

LEAF(TreeID, F-BM=0100, Flag=PopBIERincluding)

D:

ILM(inLabel<alloc by D>, action<Replication to TreeID>,
flag=CheckBS|Leaf, BSL)

LEAF(TreeID, F-BM=0001, Flag=PopBIERincluding)

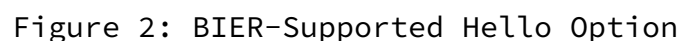
```
ILM(inLabel<alloc by F>, action<Replication to TreeID>,
flag=CheckBS|Leaf, BSL)
```

3.2. BIER-Supported Hello Option

Xie, et al.

[Page 5]

March 2018



```
OptionType = TBD, OptionLength = 4.
```

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 Disposit BIER Header of a packet including or excluding the BIER Label. If a node doesn't support P-Capability, it may still support D-Capability. If a node don't support D-capability, it is supposed not to 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.

3.3. New BIER F-BM Join Attribute Format

When a PIM router supports this mechanism and has determined from a received Hello that the neighbor supports this mechanism, and also that all the neighbors on the interface support the use of join attributes, it will send Join/Prune messages that MAY include a BIER F-BM Join Attribute. The mechanism to process a PIM Join Attribute is described in [\[RFC5384\]](#). The format of the new attribute is specified in the following.

```

+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
~                               F-BM  (last 32 bits)                               |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

Figure 3: BIER F-BM Join Attribute

F bit: 0 (Non-Transitive Attribute).

E bit: As specified by [\[RFC5384\]](#).

Attr_Type: TBD.

Length: The minimum length is 10, that is a 2 octets BSL and a minimum of 8 octets (64 bits).

BS Len: Bit String Length, that is the Length of F-BM

Set Identifier: As defined in [\[RFC8279\]](#).

F-BM: As defined in [\[RFC8279\]](#).

4. How to Use BIER F-BM Join Attribute

A router supporting this mechanism MUST, unless administratively disabled, include the PIM Join Attribute option in its PIM Hellos. See [\[RFC5384\]](#) and "PIM-Hello Options" on [PIM-REG] for details.

It is RECOMMENDED that implementations allow for administrative control of whether to make use of this mechanism. Implementations MAY also allow further control of what information to store and send upstream, by configuring whether the node require a packet without BIER header.

It is important to note that when a node's downstream F-BM OR'ing result changed, it SHOULD trigger a new Join message with an updated BIER F-BM Join Attribute.

When a router removes a link from an oif-list, it needs to be able to reevaluate the BIER F-BM that it will advertise upstream. This happens when an oif-list entry is timed out or a Prune is received.

used only for entries in the join-list part of the Join/Prune message. If the attribute is used in the prune-list, an implementation MUST ignore it and process the Prune as if the attribute were not present.

It is also RECOMMENDED that join suppression be disabled on a LAN when BIER F-BM is used.

It is RECOMMENDED that, when triggered Join/Prune messages are sent by a downstream router, the BIER F-BM information not be included in the message. This way, when convergence is important, avoiding the processing time to build an BIER F-BM record in a downstream router and processing time to parse the message in the upstream router will help reduce convergence time. If an upstream router receives a Join/Prune message with no BIER F-BM data, it SHOULD NOT interpret the message as a trigger to clear or reset the BIER F-BM data it has cached.

[5.](#) Capability and Error Handling

TBD.

[6.](#) IANA Considerations

Allocation is expected from IANA for codepoints from the "PIM-Hello Options" registry and the "PIM Join Attribute Types" registry.

[7.](#) Security Considerations

TBD

[8.](#) Acknowledgements

TBD

[9.](#) References

[9.1.](#) 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-10](#) (work in progress), February 2018.

[I-D.xie-bier-mvpn-mpls-p2mp]

Xie, J., "Multicast VPN Using MPLS P2MP and BIER", [draft-xie-bier-mvpn-mpls-p2mp-00](#) (work in progress), October 2017.

[RFC5384] Boers, A., Wijnands, I., and E. Rosen, "The Protocol Independent Multicast (PIM) Join Attribute Format", [RFC 5384](#), DOI 10.17487/RFC5384, November 2008, <<https://www.rfc-editor.org/info/rfc5384>>.

[RFC6807] Farinacci, D., Shepherd, G., Venaas, S., and Y. Cai, "Population Count Extensions to Protocol Independent Multicast (PIM)", [RFC 6807](#), DOI 10.17487/RFC6807, December 2012, <<https://www.rfc-editor.org/info/rfc6807>>.

[RFC7761] Fenner, B., Handley, M., Holbrook, H., Kouvelas, I., Parekh, R., Zhang, Z., and L. Zheng, "Protocol Independent Multicast - Sparse Mode (PIM-SM): Protocol Specification (Revised)", STD 83, [RFC 7761](#), DOI 10.17487/RFC7761, March 2016, <<https://www.rfc-editor.org/info/rfc7761>>.

[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

[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

Internet-Draft

PIM Extensions for P2MP-based BIER

March 2018

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

Email: xiejingrong@huawei.com

Yisong Liu
Huawei Technologies

Email: liuyisong@huawei.com

Mike McBride
Huawei Technologies

Email: mmcbride7@gmail.com

Xie, et al.

Expires September 6, 2018

[Page 10]