

Network Working Group  
Internet-Draft  
Intended status: Standards Track  
Expires: April 30, 2018

J. Xie  
Huawei  
October 27, 2017

Multicast VPN Using MPLS P2MP and BIER  
draft-xie-bier-mvpn-mpls-p2mp-00

Abstract

The Multicast Virtual Private Network (MVPN) specifications defines P-tunnels for carrying multicast traffic across the backbone. A variety of P-tunnel types are supported. Bit Index Explicit Replication (BIER) is a new architecture that provides optimal multicast forwarding through a "multicast domain", without requiring intermediate routers to maintain any per-flow state. The purpose of the current document is to specify one way of carrying multicast traffic over an SP MPLS backbone network using compatible method and encapsulation of BIER. It uses a pre-build P2MP as the BIER topology, and uses mLDP/RSVP-TE protocol extension to build BIER-related underlay routing and forwarding information in-band when building the P2MP topology.

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 April 30, 2018.

## Copyright Notice

Copyright (c) 2017 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. Use of the PTA in MVPN Routes . . . . .	4
3.1. Overview . . . . .	4
3.2. Use of the PTA in x-PMSI A-D Routes . . . . .	4
3.3. Use of the PTA in Leaf A-D routes . . . . .	6
4. P2MP LSP based BIER Forwarding Procedures . . . . .	7
4.1. Overview . . . . .	7
4.2. Building P2MP LSP based BIER forwarding state . . . . .	8
4.3. Live-Live protection . . . . .	9
5. Provisioning Considerations . . . . .	9
6. IANA Considerations . . . . .	10
7. Security Considerations . . . . .	11
8. Acknowledgements . . . . .	11
9. References . . . . .	11
9.1. Normative References . . . . .	11
9.2. Informative References . . . . .	13
Author's Address . . . . .	13

## 1. Introduction

[RFC6513] and [RFC6514] specify the protocols and procedures that a Service Provider (SP) can use to provide Multicast Virtual Private Network (MVPN) service to its customers. Multicast tunnels are created through an SP's backbone network; these are known as "P-tunnels". The P-tunnels are used for carrying multicast traffic across the backbone. The MVPN specifications allow the use of several different kinds of P-tunnel technology. In an MPLS network, such P-tunnel can be mLDP P2MP or RSVP-TE P2MP.

Bit Index Explicit Replication (BIER) [I-D.ietf-bier-architecture] is an architecture that provides optimal multicast forwarding through a "multicast domain", without requiring intermediate routers to maintain any per-flow state.

BIER architecture requires routers participating in BIER to exchange BIER related information within a given domain. BIER architecture permits IGP/BGP or any other routing protocols to perform distribution of such information. Such routing protocols are defined as Underlay protocols.

In an MPLS network, [I-D.ietf-bier-mpls-encapsulation] define a BIER Header within it an initial 4 octets MPLS-Label, to encapsulate Multicast packet and transport through the MPLS network.

The purpose of the current document is to specify one way of carrying multicast traffic over an SP MPLS backbone network, using compatible method and encapsulation described in the above BIER documents. It uses a pre-build P2MP as the BIER topology, and uses mLDP/RSVP-TE protocol extension to build BIER-related underlay routing and forwarding information in-band when building the P2MP topology.

## 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 LSP: Label Switch Path
- o LSR: Label Switching Router
- o P2MP: Point to Multi-point
- o P-tunnel: A multicast tunnel through the network of one or more SPs. P-tunnels are used to transport MVPN multicast data.
- o PMSI: Provider Multicast Service Interface
- o x-PMSI A-D route: a route that is either an I-PMSI A-D route or an S-PMSI A-D route.
- o PTA: PMSI Tunnel attribute. A type of BGP attribute known as the PMSI Tunnel attribute.
- o P2MP LSP based BIER: BIER using P2MP LSP as topology

### 3. Use of the PTA in MVPN Routes

#### 3.1. Overview

According to [I-D.ietf-bier-architecture], the P2MP LSP based BIER is a REAL BIER, which using a P2MP LSP as the underlay topology. The P2MP LSP is not only a LSP, but also a topology as the BIER underlay. The P2MP LSP based BIER is P-tunnel, which is used for bearing multicast flows. Every flow can think as binding to an independent tunnel, which is constructed by the BitString in the BIER header of every packet of the flow. Multicast flows are transported in SPMSI-only mode, on P2MP LSP based BIER tunnels, and never directly on P2MP LSP tunnel.

#### 3.2. Use of the PTA in x-PMSI A-D Routes

As defined in [RFC6514], the PMSI Tunnel attribute (PTA) carried by an x-PMSI A-D route identifies the P-tunnel that is used to instantiate a particular PMSI. If a PMSI is to be instantiated by P2MP LSP based BIER, the PTA is constructed by a BFIR, which is also a Ingress LSR. This document defines the following Tunnel Types:

+ TBD - RSVP-TE P2MP LSP based BIER

+ TBD - mLDP P2MP LSP based BIER

Allocation is expected from IANA for two new tunnel type codepoints from the "P-Multicast Service Interface Tunnel (PMSI Tunnel) Tunnel Types" registry. These codepoints will be used to indicate that the PMSIs is instantiated by MLDP or RSVP-TE extension with support of BIER.

When the Tunnel Type is set to RSVP-TE P2MP LSP based BIER, the Tunnel Identifier include two parts, as follows:

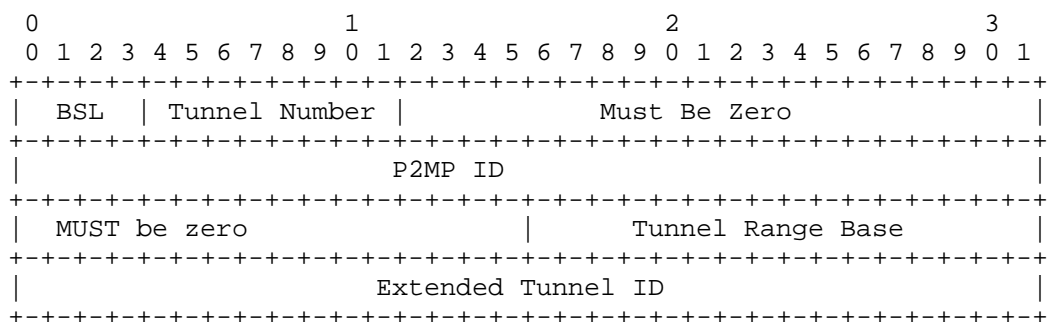


Figure 1: PTA of RSVP-TE P2MP LSP based BIER

BSL: A 4 bits field. The values allowed in this field are specified in section 2 of [I-D.ietf-bier-mpls-encapsulation].

Tunnel Number: A 1 octet field encoding the Number of the Tunnel range. It MUST be greater than 0.

<Extended Tunnel ID, Reserved, Tunnel Range Base, P2MP ID>: A ID as carried in the RSVP-TE P2MP LSP SESSION Object defined in [RFC4875].

The "Tunnel Range" is the set of P2MP LSPs beginning with the Tunnel Range base and ending with ((Tunnel Range base)+(Tunnel Number)- 1). A unique Tunnel Range is allocated for the BSL and a Sub-domain-ID implicated by the P2MP.

The size of the Tunnel Range is determined by the number of Set Identifiers (SI) (section 1 of [I-D.ietf-bier-architecture]) that are used in the topology of the P2MP-LSP. Each SI maps to a single Tunnel in the Tunnel Range. The first Tunnel is for SI=0, the second Tunnel is for SI=1, etc.

When the Tunnel Type is set to mLDP P2MP LSP based BIER, the Tunnel Identifier include two parts, as follows:

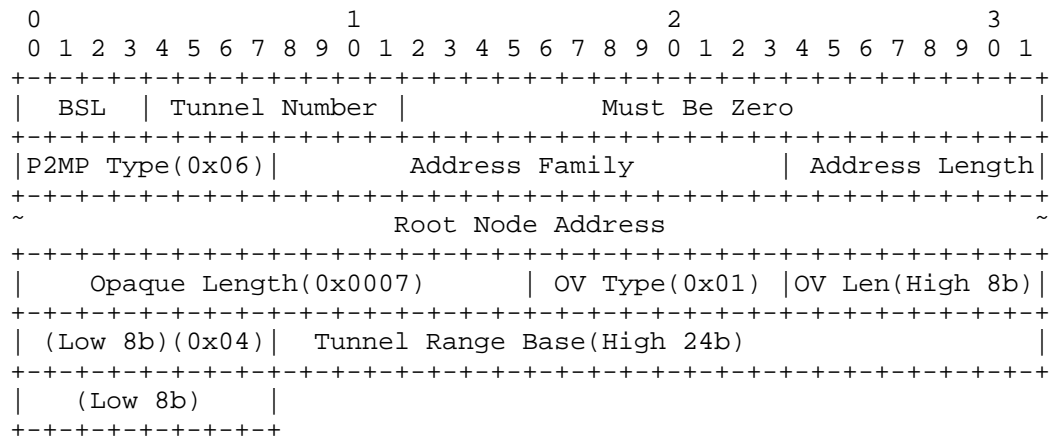


Figure 2: PTA of MLDP P2MP LSP based BIER

BSL: A 4 bits field. The values allowed in this field are specified in section 2 of [I-D.ietf-bier-mpls-encapsulation].

Tunnel Number: A 1 octet field encoding the Number of the Tunnel range. It MUST be greater than 0.

<Type=0x06, AF, AL, RootNodeAddr, Opqeue Length=0x0007, OV Type=0x01, OV Len=0x04, Tunnel Range Base>: A P2MP Forwarding Equivalence Class (FEC) Element, with a Generic LSP Identifier TLV as the opaque value element, defined in [RFC6388].

The "Tunnel Range" is the set of P2MP LSPs beginning with the Tunnel Range base and ending with ((Tunnel Range base)+(Tunnel Number)- 1). A unique Tunnel Range is allocated for the BSL and a Sub-domain-ID implicated by the P2MP.

The size of the Tunnel Range is determined by the number of Set Identifiers (SI) (section 1 of [I-D.ietf-bier-architecture]) that are used in the topology of the P2MP-LSP. Each SI maps to a single Tunnel in the Tunnel Range. The first Tunnel is for SI=0, the second Tunnel is for SI=1, etc.

When the Tunnel Type is any of the above two, The "MPLS label" field OPTIONAL contain an upstream-assigned non-zero MPLS label. It is assigned by the router (a BFIR) that constructs the PTA. Absence of an MPLS Label is indicated by setting the MPLS Label field to zero.

When the Tunnel Type is any of the above two, two of the flags, LIR and LIR-pF, in the PTA "Flags" field are meaningful. Details about the use of these flags can be found in [RFC6513], [I-D.ietf-bess-mvpn-expl-track] and [I-D.ietf-bier-mvpn]].

### 3.3. Use of the PTA in Leaf A-D routes

Before an egress PE can receive a (C-S,C-G) flow from a given ingress PE via RSVP-TE/MLDP P2MP LSP based BIER, the egress PE must have received one of the following x-PMSI A-D routes from the ingress PE:

- o A "more specific" x-PMSI A-D route, (C-S,C-G) S-PMSI A-D route.
- o A "less specific" x-PMSI A-D route, (C-\*,C-\*), (C-\*,C-G), or (C-S,C-\*) S-PMSI A-D route.

In which, the PTA tunnel Type is "RSVP-TE P2MP LSP based BIER" or "MLDP P2MP LSP based BIER".

The rules for determining which x-PMSI A-D route is the match for reception are given in [RFC6625]. If such a route is found, we refer to it as the "matching x-PMSI A-D route." If no matching x-PMSI A-D route for (C-S,C-G) is found, the egress PE cannot receive the (C-S,C-G) flow from the ingress PE via RSVP-TE/MLDP P2MP LSP based BIER until such time as a matching route is received.

When an egress PE determines that it needs to receive a (C-S,C-G) flow from a particular ingress PE via RSVP-TE/MLDP P2MP LSP based BIER, it originates a Leaf A-D route. Construction of the Leaf A-D route generally follows the procedures specified in [RFC6514], or optionally, the procedures specified in [I-D.ietf-bess-mvpn-expl-track]. However, when RSVP-TE/MLDP P2MP LSP based BIER is being used, the Leaf A-D route MUST carry a PTA that is constructed as follows:

1. The tunnel type MUST be set to RSVP-TE/MLDP P2MP LSP based BIER, corresponding to the PTA of the matching x-PMSI A-D route.
2. The MPLS label field SHOULD be set to zero.
3. The BFR-Prefix field of the Tunnel Identifier field MUST be set to the egress PE's IP-Address. This IP-Address is the same as the Originating Router's IP Addr field of the NLRI of the Leaf A-D route.

When an ingress PE receives such a Leaf A-D route, it learns the BFR-Prefix of the egress PE from the PTA. The ingress PE does not make any use the value of the PTA's MPLS label field.

Failure to properly construct the PTA cannot always be detected by the protocol, and will cause improper delivery of the data packets.

#### 4. P2MP LSP based BIER Forwarding Procedures

The MVPN application plays the role of the "multicast flow overlay" as described in [I-D.ietf-bier-architecture].

This section specifies some OPTIONAL rules for forwarding a BIER-encapsulated data packet within a P2MP topology underlay.

These rules will produce the same results as the procedures in [I-D.ietf-bier-architecture], on condition that the underlay topology is a P2MP.

##### 4.1. Overview

As [I-D.ietf-bier-architecture] describes:

1. BIER support using the default topology of the unicast IGP as the routing underlay. To quote from [I-D.ietf-bier-architecture]:  
"By default, each sub-domain uses the default topology of the unicast IGP as the routing underlay."

2. BIER also support using other topologies as the routing underlay, including a tree topology. To quote from [I-D.ietf-bier-architecture]: "Alternatively, one could deploy a routing underlay that creates a multicast-specific tree of some sort. Then BIER could be used to forward multicast data packets along the multicast-specific tree, while unicast packets follow the 'ordinary' OSPF best path."

This document specifies one OPTIONAL Forwarding Procedure of BIER encapsulation packet, on the condition that the BIER underlay topology is P2MP LSP, as describes in the above sections. Comparing to the Forwarding Procedure, which is described in [I-D.ietf-bier-architecture], and which is on a underlay of unicast IGP topology, there is some simplification:

1. Not need to Edit the BitString when forwarding packet to Neighbor, for the underlay P2MP topology is already loop-free.
2. Not need to use Entropy in the BIER Header, for current P2MP topology is already ECMP-eliminate.

The optional BIER forwarding procedure is, on the basis of P2MP forwarding procedure according to the BIER-MPLS label, and use the BitString to prune the undesired P2MP downstream.

The enhancement to the P2MP forwarding is to add a Forwarding BitMask to existing NHLFE defined in [RFC3031], for checking with the BitString in a packet, to determin whether the packet is to be forwarded or pruned. If the checking result by AND'ing a packet's BitString with the F-BM of the NHLFE (i.e., Packet->BitString &= F-BM) is non-zero, then forward the packet to the next-hop indicated by the NHLFE entry, and the Label is switched to the proper one in the NHLFE. If the result is zero, then do not forward the packet to the next-hop indicated by the NHLFE entry.

#### 4.2. Building P2MP LSP based BIER forwarding state

When RSVP-TE/MLDP P2MP LSP based BIER are used, then it is not nessary to use IGP or BGP to build the BIER routing table and forwarding table. Instead, the BIER layer information is carried by MLDP or RSVP-TE, and when MLDP or RSVP-TE build P2MP LSP, it build the BIER forwarding state in-band.

The procedure for building RSVP-TE/MLDP P2MP LSP based BIER forwarding state using mLDP or RSVP-TE is outside the scope of this document.



#### 4.3. Live-Live protection

As described above, loop and redundancy, ECMP and Entropy, are all not supported in current P2MP LSP underlay. There will be extra P2MP LSP convergence, after IGP convergence, in the case of link or node failure.

On the other hand, Multicast has special Service Level Agreement (SLA), especially when multicast service is compressed or uncompressed video. Accordingly, there are some multicast-specific methods of protection, such as Live-Live. [RFC7431] defines a method of detecting failure locally by comparing the packets received from live-live paths. [I-D.ietf-bess-mvpn-fast-failover] defines a Live-Live method for protecting Multicast in MVPN.

This document specifies one OPTIONAL extension to enhance Live-Live protection, re-using the Entropy field of BIER header as a Sequence number of multicast packet, on the condition that the field is not used for ECMP, such as in the P2MP LSP topology described above.

This is an optional function of BIER Layer. If this function is enabled, every BFR of the domain is required to support, which means:

1. The BFIR (and Ingress LSR) will push a sequence-number in the Entropy field, per-flow per-packet.
2. The middle BFR will ignore the Entropy field, and not do the selection of multi-tables.
3. The BFER (and Egress LSR) will do packet check from live-live paths, and do forward packet with zero packet loss, on a per-flow basis.

#### 5. Provisioning Considerations

P2MP LSP based BIER use concepts of both RSVP-TE/MLDP and BIER. Some provisioning considerations list below:

Sub-domain:

In P2MP LSP based BIER, every P2MP LSP is a specific BIER underlay topology, and an implicit Sub-domain. RSVP-TE/MLDP build the BIER information of the implicit sub-domain when build P2MP LSP. MVPN get the implicit sub-domain when specified with which RSVP-TE/MLDP P2MP LSP.

In the following conditions, there may be requirements to configure an explicit sub-domain ID for P2MP LSP based BIER:

1. P2MP LSP based BIER, use the native procedure of forwarding described in [I-D.ietf-bier-architecture], which require Consistent Per-Sub-domain BIFT.
2. P2MP LSP based BIER is shared by multiple VPNs, and an explicit sub-domain ID is configured as anchor for using by these VPNs.

When explicitly configing a sub-domain ID for P2MP LSP based BIER, the ID should be great than 255. For the [0-255] has been defined to use by IGP, BGP and MVPN, as specified in [I-D.ietf-bier-ospf-bier-extensions], [I-D.ietf-bier-isis-extensions], [I-D.ietf-bier-idr-extensions] and [I-D.ietf-bier-mvpn].

BFR-prefix:

In P2MP LSP based BIER, every BFR is also a LSR. So the BFR-prefix in the sub-domain is by default identified by LSR-id. Additionally, When BFR/LSR is also a MVPN PE, BFR-prefix is also the same as Originating Router's IP Address of x-PMSI A-D route or Leaf A-D route.

BFR-id:

When using protocols like RSVP-TE, which initializes P2MP LSP from a specific Ingress Node, BFR-id which is unique in P2MP LSP scope, can be auto-provisioned by Ingress Node, or conventionally configure on every Egress Nodes.

BSL and BIER-MPLS Label Block Size:

In P2MP LSP based BIER, Every P2MP LSP or implicit sub-domain requires a single BSL, and a specific BIER-MPLS Label block size for this BSL.

VPN-Label:

In P2MP LSP based BIER, a P2MP LSP based BIER 'P-tunnel' can be shared by multiple VPNs or a single VPN. When a P2MP LSP based BIER being shared by multiple VPNs, an Upstream-assigned VPN-Label is required. It can be auto-provisioned or manual configured by the BFIR or Ingress LSR.

## 6. IANA Considerations

Allocation is expected from IANA for two new tunnel type codepoints for "RSVP-TE P2MP LSP based BIER" and "MLDP P2MP LSP based BIER" from

the "P-Multicast Service Interface Tunnel (PMSI Tunnel) Tunnel Types" registry.

## 7. Security Considerations

This document does not introduce any new security considerations other than already discussed in [I-D.ietf-bier-architecture].

## 8. Acknowledgements

TBD

## 9. References

### 9.1. Normative References

[I-D.ietf-bess-mvpn-expl-track]

Dolganow, A., Kotalwar, J., Rosen, E., and Z. Zhang, "Explicit Tracking with Wild Card Routes in Multicast VPN", draft-ietf-bess-mvpn-expl-track-03 (work in progress), September 2017.

[I-D.ietf-bess-mvpn-fast-failover]

Morin, T. and R. Kebler, "Multicast VPN fast upstream failover", draft-ietf-bess-mvpn-fast-failover-02 (work in progress), March 2017.

[I-D.ietf-bier-architecture]

Wijnands, I., Rosen, E., Dolganow, A., Przygienda, T., and S. Aldrin, "Multicast using Bit Index Explicit Replication", draft-ietf-bier-architecture-08 (work in progress), September 2017.

[I-D.ietf-bier-idr-extensions]

Xu, X., Chen, M., Patel, K., Wijnands, I., and T. Przygienda, "BGP Extensions for BIER", draft-ietf-bier-idr-extensions-03 (work in progress), August 2017.

[I-D.ietf-bier-isis-extensions]

Ginsberg, L., Przygienda, T., Aldrin, S., and Z. Zhang, "BIER support via ISIS", draft-ietf-bier-isis-extensions-06 (work in progress), October 2017.

- [I-D.ietf-bier-mpls-encapsulation]  
Wijnands, I., Rosen, E., Dolganow, A., Tantsura, J., Aldrin, S., and I. Meilik, "Encapsulation for Bit Index Explicit Replication in MPLS and non-MPLS Networks", draft-ietf-bier-mpls-encapsulation-11 (work in progress), October 2017.
- [I-D.ietf-bier-mvpn]  
Rosen, E., Sivakumar, M., Aldrin, S., Dolganow, A., and T. Przygienda, "Multicast VPN Using BIER", draft-ietf-bier-mvpn-08 (work in progress), October 2017.
- [I-D.ietf-bier-ospf-bier-extensions]  
Psenak, P., Kumar, N., Wijnands, I., Dolganow, A., Przygienda, T., Zhang, Z., and S. Aldrin, "OSPF Extensions for BIER", draft-ietf-bier-ospf-bier-extensions-09 (work in progress), October 2017.
- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, DOI 10.17487/RFC3031, January 2001, <<https://www.rfc-editor.org/info/rfc3031>>.
- [RFC4875] Aggarwal, R., Ed., Papadimitriou, D., Ed., and S. Yasukawa, Ed., "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, DOI 10.17487/RFC4875, May 2007, <<https://www.rfc-editor.org/info/rfc4875>>.
- [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>>.
- [RFC6513] Rosen, E., Ed. and R. Aggarwal, Ed., "Multicast in MPLS/BGP IP VPNs", RFC 6513, DOI 10.17487/RFC6513, February 2012, <<https://www.rfc-editor.org/info/rfc6513>>.
- [RFC6514] Aggarwal, R., Rosen, E., Morin, T., and Y. Rekhter, "BGP Encodings and Procedures for Multicast in MPLS/BGP IP VPNs", RFC 6514, DOI 10.17487/RFC6514, February 2012, <<https://www.rfc-editor.org/info/rfc6514>>.

- [RFC6625] Rosen, E., Ed., Rekhter, Y., Ed., Hendrickx, W., and R. Qiu, "Wildcards in Multicast VPN Auto-Discovery Routes", RFC 6625, DOI 10.17487/RFC6625, May 2012, <<https://www.rfc-editor.org/info/rfc6625>>.
- [RFC7431] Karan, A., Filsfils, C., Wijnands, IJ., Ed., and B. Decraene, "Multicast-Only Fast Reroute", RFC 7431, DOI 10.17487/RFC7431, August 2015, <<https://www.rfc-editor.org/info/rfc7431>>.

## 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>>.

## Author's Address

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

Email: [xiejingrong@huawei.com](mailto:xiejingrong@huawei.com)