Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: May 6, 2021

H. Chen
Futurewei
A. Wang
China Telecom
G. Mishra
Verizon Inc.
Y. Fan
Casa Systems
L. Liu
Fujitsu
X. Liu
Volta Networks

November 2, 2020

BIER Fast ReRoute
draft-chen-bier-frr-00

Abstract

This document describes a mechanism for fast re-route (FRR) protection against the failure of a node or link in the core of a "Bit Index Explicit Replication" (BIER) domain. It does not have any per-flow state in the core. For a multicast packet to traverse a node in the domain, when the node fails, its upstream hop as a PLR reroutes the packet around the failed node once it detects the failure.

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] [RFC8174] when, and only when, they appear in all capitals, as shown here.

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 May 6, 2021.

Table of Contents

1. Introduction 2
1.1. Terminology 3
2. BIER FRR Solution 4
2.1. Overview of BIER forwarding 4
2.2. FRR Bit Index Routing Tables 5
2.3. FRR Bit Index Forwarding Tables 6
2.4. Updated Forwarding Procedure 7
2.5. Switching between FRR and Normal Forwarding 7
3. Example Application of BIER FRR 8
3.1. Example BIER Topology 8
3.2. BIRT and BIFT on a BFR 8
3.3. FRR-BIRTs and FRR-BIFTs on a BFR 10
3.4. Forwarding using FRR-BIFT 12
4. Security Considerations 13
5. IANA Considerations 13
6. Acknowledgements 13
7. References 13
7.1. Normative References 13
7.2. Informative References 15
Authors' Addresses 15

1. Introduction

[RFC8279] specifies "Bit Index Explicit Replication" (BIER). It provides optimal forwarding of multicast data packets through a "multicast/BIER domain". It does not require the use of a protocol

for explicitly building multicast distribution trees, and it does not require intermediate nodes to maintain any per-flow state.

This document describes a mechanism for fast re-route (FRR) protection against the failure of a node or link in the core of a BIER domain, which is called BIER-FRR. For a multicast packet with a BIER header to traverse a node in the domain, when the node fails, its upstream hop as a point of local repair (PLR) reroutes the packet around the failed node once it detects the failure.

This BIER-FRR does not require intermediate nodes to maintain any per-flow state for FRR protection against the failure of a node or link along the flow.

1.1. Terminology

BFR: Bit-Forwarding Router.

BFIR: Bit-Forwarding Ingress Router.

BFER: Bit-Forwarding Egress Router.

BFR-id: BFR Identifier. It is a number in the range [1,65535].

BFR-NBR: BFR Neighbor.

F-BM: Forwarding Bit Mask.

BFR-prefix: An IP address (either IPv4 or IPv6) of a BFR.

BIRT: Bit Index Routing Table. It is a table that maps from the BFR-id (in a particular sub-domain) of a BFER to the BFR-prefix of that BFER, and to the BFR-NBR on the path to that BFER.

BIFT: Bit Index Forwarding Table.

FRR: Fast Re-Route.

PLR: Point of Local Repair.

LFA: Loop-Free Alternate.

RLFA: Remote LFA.

DLFA: Remote LFA with Directed forwarding.

IGP: Interior Gateway Protocol.

LSDB: Link State DataBase.

SPF: Shortest Path First.

SPT: Shortest Path Tree.

SPT-old(R): The SPT rooted at node R using LSDB before X fails (i.e., old LSDB).

SPT-new(R, X): The SPT rooted at node R using LSDB without X after X fails (i.e., new LSDB).

P-Space P(R,X): The set of nodes that are reachable from R without going through X. In other words, it is the set of nodes that are not downstream of X in SPT-old(R).

Extended P-Space P'(R,X): The set of nodes that are reachable from R or a neighbor of R, without going through X.

Q-Space Q(D,X): The set of nodes that do not use X to reach destination D using the old LSDB.

PQ node(R,X): A member of both the P-Space P(R, X) (or the extended P-Space P'(R, X)) and the Q-Space (D, X).

2. BIER FRR Solution

A Bit-Forwarding Router (BFR) in a BIER sub-domain builds and maintains a "FRR Bit Index Routing Table" (FRR-BIRT) for each of its BFR Neighbors (BFR-NBRs) to provide BIER-FRR. The BFR builds each FRR-BIRT based on a BIRT defined in [RFC8279]. A "FRR Bit Index Forwarding Table" (FRR-BIFT) is derived from a FRR-BIRT in the same way as a BIFT is derived from a BIRT, which is defined in [RFC8279].

The forwarding procedure defined in [RFC8279] is enhanced/updated for FRR-BIFTs. Once the BFR as a PLR detects the failure of its BFR-NBR X, it uses the FRR-BIFT for X to forward packets with BIER headers to get around failed X according to the updated/enhanced forwarding procedure.

2.1. Overview of BIER forwarding

This section briefs the BIRT, BIFT and forwarding procedure defined in [RFC8279].

There is a "Bit Index Routing Table" (BIRT) for a BIER sub-domain on a BFR. The BIRT maps the BFR Identifier (BFR-id) (in the sub-domain) of a Bit-Forwarding Egress Router (BFER) to the BFR-prefix of that

BFER, and to the BFR-NBR on the shortest path to that BFER. In other words, the BIRT has a route or say a next hop (i.e., BFR-NBR on the path) to every BFER.

From the BIRT on the BFR, a "Bit Index Forwarding Table" (BIFT) is derived. In addition to having a route to a BFER in each row of the BIFT which is the same as the BIRT, it has a "Forwarding Bit Mask" (F-BM) in its each row. For the rows in the BIRT that have the same SI and the same BFR-NBR, the F-BM for each of these rows in the BIFT is the logical OR of the BitStrings of these rows.

This BIFT is programmed into the data plane and used to forward a packet with a BIER header. The header contains SI, BitString, BitStringLength, and sub-domain.

When a BFR receives a packet, for each BFER k (from the rightmost to the leftmost) represented in the SI and BitString of the packet, if BFER k is the BFR itself, the BFR copies the packet, sends the copy to the multicast flow overlay and clears bit k in the original packet; otherwise the BFR finds the row (i.e., forwarding entry) in the BIFT for the sub-domain using the SI and BitString as the key or say index, and then copies, updates and forwards the packet to the BFR-NBR (i.e., the next hop) indicated by the row (i.e., forwarding entry).

After copying the packet and before forwarding it to the BFR-NBR, the packet's BitString is updated by ANDing it with the F-BM in the forwarding entry (i.e., PacketCopy->BitString &= F-BM).

After forwarding the updated packet to a BFR-NBR and before forwarding the original packet to another BFR-NBR, the original packet's BitString is changed by ANDing it with the INVERSE of the F-BM (i.e., Packet->BitString &= ~F-BM).

2.2. FRR Bit Index Routing Tables

Each BFR in a BIER sub-domain builds and maintains a number of "FRR Bit Index Routing Tables" (FRR-BIRTs). There is a FRR-BIRT for each BFR-NBR of the BFR. The BFR builds each FRR-BIRT based on its BIRT. It has the same format as the BIRT.

The FRR-BIRT for BFR-NBR X of the BFR considers the failure of X and maps the BFR-id (in the sub-domain) of a BFER to the BFR-prefix of that BFER, and to BFR-NBR N on the path to that BFER. In other words, the FRR-BIRT has a route or say a next hop (i.e., BFR-NBR N on the path, where N is not X) to every BFER when BFR-NBR X fails.

The BFR may build the FRR-BIRT for BFR-NBR X by copying its BIRT to the FRR-BIRT first, and then change the next hop with value BFR-NBR X in the FRR-BIRT to a backup next hop (BNH) to protect against the failure of X. In other wards, for the BFR-id of a BFER in the FRR-BIRT for BFR-NBR X, if the next hop BFR-NBR on the path to the BFER is X, it is changed to a BNH when there is a BNH on a backup path to the BFER without going through X and the link from the BFR to X.

If there is not any BNH to a BFER to protect against the failure of X, the next hop BFR-NBR X to the BFER in the FRR-BIRT for BFR-NBR X is changed to NULL. For a multicast packet having the BFER as one of its destinations, if the next hop BFR-NBR to the BFER is NULL, the BFR does not send the packet to the next hop BFR-NBR NULL but drops it when X fails.

Note: In another option, the next hop BFR-NBR X to the BFER in the FRR-BIRT for BFR-NBR X keeps unchanged when there is not any BNH to the BFER to protect against the failure of X. In this case, for a multicast packet having the BFER as one of its destinations, the BFR sends the packet to X when X fails.

In one implementation, the BNH is the Loop-Free Node-Protecting Alternate defined in [RFC5286] to protect against the failure of X and link from the BFR to X. In another implementation, the BNH is the virtual Loop-Free Alternate (LFA), i.e., PQ node, defined in [RFC7490]. In a special case, a PQ node is a Loop-Free Node-Protecting Alternate defined in [RFC5286].

2.3. FRR Bit Index Forwarding Tables

From each FRR-BIRT on the BFR, a "FRR Bit Index Forwarding Table" (FRR-BIFT) is derived. In addition to having a route to a BFER in each row of the FRR-BIFT which is the same as the FRR-BIRT, it has a "Forwarding Bit Mask" (F-BM) in its each row. For the rows in the FRR-BIRT that have the same SI and the same BFR-NBR, the F-BM for each of these rows in the FRR-BIFT is the logical OR of the BitStrings of these rows.

This FRR-BIFT is programmed into the data plane and is not used to forward any packet in normal operations. It is activated to forward a packet with a BIER header once the BFR detects the failure of BFR-NBR. The header contains SI, BitString, BitStringLength, and sub-domain.

2.4. Updated Forwarding Procedure

The forwarding procedure defined in [RFC8279] is updated/enhanced for a FRR-BIFT to consider the case where the next hop BFR-NBR to a BFER is NULL. For a multicast packet with the BitString indicating a BFER as one of its destinations, the updated forwarding procedure checks whether the next hop BFR-NBR to the BFER in the FRR-BIFT is NULL. If it is NULL, the procedure will not send the packet to this next hop BFR-NBR NULL but drop the packet.

The updated procedure is described in Figure 1. It is used with a FRR-BIFT for BFR-NBR X on a BFR to forward multicast packets when X fails. It can also be used with a BIFT on the BFR to forward multicast packets in normal operations.

```
  Packet = the packet received by BFR;
  FOR each BFER k (from the rightmost in Packet's BitString) {
      IF BFER k is the BFR itself {
          copies Packet, sends the copy to the multicast
          flow overlay and clears bit k in Packet's BitString
      } else {
          finds the row in the FRR-BIFT for the sub-domain using
          Packet's SI and BitString as the key/index
          IF BFR-NBR in the row is not NULL {
              Copies Packet, updates the copy's BitString by ANDing
              it with F-BM in the row, sends updated copy to BFR-NBR
          } // BFR-NBR == NULL, not sent Packet to BFR-NBR
          updates Packet's BitString by ANDing it with the INVERSE
          of the F-BM in the row
      }
  }
```

Figure 1: Updated Forwarding Procedure

2.5. Switching between FRR and Normal Forwarding

The FRR-BIFTs will be pre-computed and installed ready for activation when a failure is detected. Once the BFR detects the failure of its BFR-NBR X, it activates the FRR-BIFT for X to forward packets with BIER headers and de-activates its BIFT. After activation of the FRR-BIFT, it remains in effect until it is no longer required.

In general, when the routing protocol has re-converged on the new topology taking into account the failure of X, the BIRT is re-computed using the updated LSDB and the BIFT is re-derived from the BIRT. Once the BIFT is installed ready for activation, it is activated to forward packets with BIER headers and the FRR-BIFT for X is de-activated.

From the new topology, the BFR computes/re-computes the FRR-BIRT for each BFR-NBR Y of the BFR and the FRR-BIFT for Y is derived/re-derived from the FRR-BIRT for Y. The FRR-BIFT is installed/re-installed ready for activation when Y fails.

3. Example Application of BIER FRR

This section illustrates an example application of BIER FRR on a BFR in a BIER topology in Figure 2.

3.1. Example BIER Topology

An example BIER topology for a BIER sub-domain is shown in Figure 2. It has 8 nodes/BFRs A, B, C, D, E, F, G and H. Each of the links connecting these nodes/BFRs has a cost. The link cost of 1 is default and is not indicated in the figure. The link cost of other value such as 2 is indicated in the figure.

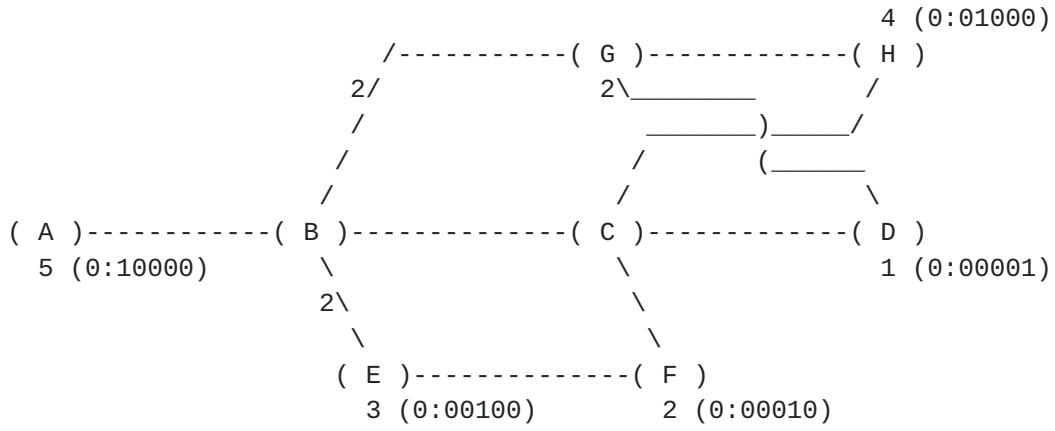

Figure 2: Example BIER Topology

Nodes/BFRs D, F, E, H and A are BFERs and have BFR-ids 1, 2, 3, 4, and 5 respectively. For simplicity, these BFR-ids are represented by (SI:BitString), where SI = 0 and BitString is of 5 bits. BFR-ids 1, 2, 3, 4, and 5 are represented by (0:00001), (0:00010), (0:00100), (0:01000) and (0:10000) respectively.

3.2. BIRT and BIFT on a BFR

Every BFR in a BIER sub-domain/topology builds and maintains a Bit Index Routing Table (BIRT). For the BIER topology in Figure 2, each of 8 nodes/BFRs A, B, C, D, E, F, G and H builds and maintains a BIRT using the LSDB for the topology.

The BIRT built on BFR B (i.e. node B) is shown in Figure 3.

BFR-id (SI:BitString)	BFR-Prefix of Dest BFER	BFR-NBR (Next Hop)
1 (0:00001)	D	C
2 (0:00010)	F	C
3 (0:00100)	E	E
4 (0:01000)	H	C
5 (0:10000)	A	A

Figure 3: Bit Index Routing Table on BFR B

The 1st row in the BIRT indicates that the next hop BFR-NBR on the shortest path to BFER D with BFR-id 1 is BFR C.

The 2nd row in the BIRT indicates that the next hop BFR-NBR on the shortest path to BFER F with BFR-id 2 is BFR C.

The 3rd row in the BIRT indicates that the next hop BFR-NBR on the shortest path to BFER E with BFR-id 3 is BFR E.

The 4-th row in the BIRT indicates that the next hop BFR-NBR on the shortest path to BFER H with BFR-id 4 is BFR C.

The 5-th row in the BIRT indicates that the next hop BFR-NBR on the shortest path to BFER A with BFR-id 5 is BFR A.

From this BIRT on BFR B, a Bit Index Forwarding Table (BIFT) is derived. This BIFT is shown in Figure 4.

The 1st, 2nd and 4-th rows in the BIRT have the same SI = 0 and next hop BFR-NBR = C. The F-BM for each of these three rows in the BIFT is the logical OR of the BitStrings of these rows, which is 01011 (00001 OR 00010 OR 01000 = 01011).

The F-BM for 3rd row in the BIFT is 00100. The F-BM for 5-th row in the BIFT is 10000.

BFR-id (SI:BitString)	F-BM	BFR-NBR (Next Hop)
1 (0:00001)	01011	C
2 (0:00010)	01011	C
3 (0:00100)	00100	E
4 (0:01000)	01011	C
5 (0:10000)	10000	A

Figure 4: Bit Index Forwarding Table on BFR B

3.3. FRR-BIRTs and FRR-BIFTs on a BFR

Every BFR in a BIER sub-domain/topology builds and maintains a number of FRR Bit Index Routing Tables (FRR-BIRTs). For the BIER topology in Figure 2, each of 8 nodes/BFRs A, B, C, D, E, F, G and H builds and maintains a number of FRR-BIRTs using the LSDB for the topology for its every BFR-NBR.

For example, BFR B (i.e., node B) in the BIER topology builds and maintains four FRR-BIRTs for its four BFR-NBRs (BFR C, BFR E, BFR A and BFR G) respectively. The FRR-BIRT for BFR C built by BFR B is shown in Figure 5.

BFR-id (SI:BitString)	BFR-Prefix of Dest BFER	BFR-NBR (Next Hop)
1 (0:00001)	D	G
2 (0:00010)	F	E
3 (0:00100)	E	E
4 (0:01000)	H	G
5 (0:10000)	A	A

Figure 5: FRR Bit Index Routing Table for BFR C on BFR B

The 1st row in the FRR-BIRT indicates that the next hop BFR-NBR on the path to BFER D with BFR-id 1 is BFR G. G is the Loop-Free Node-Protecting Alternate defined in [RFC5286] to protect against the failure of C and link from B to C.

The 2nd row in the FRR-BIRT indicates that the next hop BFR-NBR on the path to BFER F with BFR-id 2 is BFR E. E is the Loop-Free Node-Protecting Alternate defined in [RFC5286] to protect against the failure of C and link from B to C.

The 3rd row in the FRR-BIRT indicates that the next hop BFR-NBR on the path to BFER E with BFR-id 3 is BFR E.

The 4-th row in the FRR-BIRT indicates that the next hop BFR-NBR on the path to BFER H with BFR-id 4 is BFR G. G is the Loop-Free Node-Protecting Alternate defined in [RFC5286] to protect against the failure of C and link from B to C.

The 5-th row in the FRR-BIRT indicates that the next hop BFR-NBR on the path to BFER A with BFR-id 5 is BFR A.

From this FRR-BIRT for BFR C on BFR B, a FRR Bit Index Forwarding Table (FRR-BIFT) is derived. This FRR-BIFT for BFR C is shown in Figure 6.

The 1st and 4-th rows in the FRR-BIRT have the same SI = 0 and next hop BFR-NBR = G. The F-BM for each of these two rows in the FRR-BIFT is the logical OR of the BitStrings of these rows, which is 01001 (00001 OR 01000 = 01001).

BFR-id (SI:BitString)	F-BM	BFR-NBR (Next Hop)
1 (0:00001)	01001	G
2 (0:00010)	00110	E
3 (0:00100)	00110	E
4 (0:01000)	01001	G
5 (0:10000)	10000	A

Figure 6: FRR Bit Index Forwarding Tablefor BFR C on BFR B

The 2nd and 3rd rows in the FRR-BIRT have the same SI = 0 and next hop BFR-NBR = E. The F-BM for each of these two rows in the FRR-BIFT is the logical OR of the BitStrings of these rows, which is 00110 (00010 OR 00100 = 00110).

The F-BM for 5-th row in the FRR-BIFT is 10000.

3.4. Forwarding using FRR-BIFT

Suppose that there is a multicast traffic from BFR A as ingress/BFIR to egresses/BFERs D, F, E and H. For every packet of the traffic, after receiving it, BFR A adds a BIER header into the packet and sends the packet with the BIER header to BFR B. The BIER header contains (SI:BitString) = (0:01111) for egresses/BFERs D, F, E and H.

In normal operations, after receiving the packet from BFR A, BFR B copies, updates and sends the packet to BFR C and BFR E using the BIFT on BFR B according to the forwarding procedure defined in [RFC8279].

Once BFR B detects the failure of its BFR-NBR X, after receiving the packet from BFR A, BFR B copies, updates and sends the packet using the FRR-BIFT for X on BFR B to avoid X and link from B to X according to the forwarding procedure defined in [RFC8279].

For example, once BFR B detects the failure of its BFR-NBR C, after receiving the packet from BFR A, BFR B copies, updates and sends the packet to BFR G and BFR E using the FRR-BIFT for BFR C on BFR B to avoid C and link from B to C.

The packet received by BFR B from BFR A contains (SI:BitString) = (0:01111). The rightmost one bit in BitString is bit 1. For BFER 1 (0:00001) (i.e., BFR D as BFER), BFR B gets the 1st row (i.e., forwarding entry) in the FRR-BIFT for BFR C. The next hop BFR-NBR is G in the row. BFR B copies, updates and forwards the packet to G.

The packet sent to G contains the updated BitString = 01001, which is 01111 & F-BM in the row = 01111 & 01001.

After sending the packet to G, BFR B updates the original packet by ANDing its BitString with the INVERSE of the F-BM in the row. The updated BitString = 00110, which is 01111 & ~F-BM in the row = 01111 & 00110.

For the packet containing BitString = 00110, the rightmost one bit in BitString is bit 2. For BFER 2 (0:00010) (i.e., BFR F as BFER), BFR B gets the 2nd row (i.e., forwarding entry) in the FRR-BIFT for BFR

C. The next hop BFR-NBR is E in the row. BFR B copies, updates and forwards the packet to E.

The packet sent to E contains the updated BitString = 00110, which is 00110 & F-BM in the 2nd row = 00110 & 00110.

After sending the packet to E, BFR B updates the original packet by ANDing its BitString with the INVERSE of the F-BM in the 2nd row. The updated BitString = 00000, which is 00110 & ~F-BM in the row = 00110 & 11001.

The updated packet has BitString without any one bit. BFR B finishes forwarding the packet from A to D, F, E and H.

4. Security Considerations

TBD.

5. IANA Considerations

No requirements for IANA.

6. Acknowledgements

The authors would like to thank people for their comments to this work.

7. References

7.1. Normative 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>.

[RFC5226] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 5226, DOI 10.17487/RFC5226, May 2008, <https://www.rfc-editor.org/info/rfc5226>.

[RFC5250] Berger, L., Bryskin, I., Zinin, A., and R. Coltun, "The OSPF Opaque LSA Option", RFC 5250, DOI 10.17487/RFC5250, July 2008, <https://www.rfc-editor.org/info/rfc5250>.

[RFC5286] Atlas, A., Ed. and A. Zinin, Ed., "Basic Specification for IP Fast Reroute: Loop-Free Alternates", RFC 5286, DOI 10.17487/RFC5286, September 2008, <https://www.rfc-editor.org/info/rfc5286>.

[RFC5714] Shand, M. and S. Bryant, "IP Fast Reroute Framework", RFC 5714, DOI 10.17487/RFC5714, January 2010, <https://www.rfc-editor.org/info/rfc5714>.

[RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, DOI 10.17487/RFC5880, June 2010, <https://www.rfc-editor.org/info/rfc5880>.

[RFC7356] Ginsberg, L., Previdi, S., and Y. Yang, "IS-IS Flooding Scope Link State PDUs (LSPs)", RFC 7356, DOI 10.17487/RFC7356, September 2014, <https://www.rfc-editor.org/info/rfc7356>.

[RFC7490] Bryant, S., Filsfils, C., Previdi, S., Shand, M., and N. So, "Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)", RFC 7490, DOI 10.17487/RFC7490, April 2015, <https://www.rfc-editor.org/info/rfc7490>.

[RFC7684] Psenak, P., Gredler, H., Shakir, R., Henderickx, W., Tantsura, J., and A. Lindem, "OSPFv2 Prefix/Link Attribute Advertisement", RFC 7684, DOI 10.17487/RFC7684, November 2015, <https://www.rfc-editor.org/info/rfc7684>.

[RFC7770] Lindem, A., Ed., Shen, N., Vasseur, JP., Aggarwal, R., and S. Shaffer, "Extensions to OSPF for Advertising Optional Router Capabilities", RFC 7770, DOI 10.17487/RFC7770, February 2016, <https://www.rfc-editor.org/info/rfc7770>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[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>.

[RFC8556] Rosen, E., Ed., Sivakumar, M., Przygienda, T., Aldrin, S., and A. Dolganow, "Multicast VPN Using Bit Index Explicit Replication (BIER)", RFC 8556, DOI 10.17487/RFC8556, April 2019, <https://www.rfc-editor.org/info/rfc8556>.

7.2. Informative References

[I-D.ietf-rtgwg-segment-routing-ti-lfa]
Litkowski, S., Bashandy, A., Filsfils, C., Decraene, B., Francois, P., Voyer, D., Clad, F., and P. Camarillo, "Topology Independent Fast Reroute using Segment Routing", draft-ietf-rtgwg-segment-routing-ti-lfa-04 (work in progress), August 2020.

[I-D.ietf-spring-segment-protection-sr-te-paths]
Hegde, S., Bowers, C., Litkowski, S., Xu, X., and F. Xu, "Segment Protection for SR-TE Paths", draft-ietf-spring-segment-protection-sr-te-paths-00 (work in progress), September 2020.

[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>.

[RFC8401] Ginsberg, L., Ed., Przygienda, T., Aldrin, S., and Z. Zhang, "Bit Index Explicit Replication (BIER) Support via IS-IS", RFC 8401, DOI 10.17487/RFC8401, June 2018, <https://www.rfc-editor.org/info/rfc8401>.

[RFC8444] Psenak, P., Ed., Kumar, N., Wijnands, IJ., Dolganow, A., Przygienda, T., Zhang, J., and S. Aldrin, "OSPFv2 Extensions for Bit Index Explicit Replication (BIER)", RFC 8444, DOI 10.17487/RFC8444, November 2018, <https://www.rfc-editor.org/info/rfc8444>.

Authors' Addresses

Huaimo Chen
Futurewei
Boston, MA
USA

Email: Huaimo.chen@futurewei.com

Aijun Wang
China Telecom
Beiqijia Town, Changping District
Beijing, 102209
China

Email: wangaj3@chinatelecom.cn

Gyan S. Mishra
Verizon Inc.
13101 Columbia Pike
Silver Spring MD 20904
USA

Phone: 301 502-1347
Email: gyan.s.mishra@verizon.com

Yanhe Fan
Casa Systems
USA

Email: yfan@casa-systems.com

Lei Liu
Fujitsu

USA

Email: liulei.kddi@gmail.com

Xufeng Liu
Volta Networks

McLean, VA
USA

Email: xufeng.liu.ietf@gmail.com