

Network Working Group
Internet-Draft
Intended status: Informational
Expires: July 6, 2022

H. Chen, Ed.
M. McBride
Futurewei
S. Lindner
M. Menth
University of Tuebingen
A. Wang
China Telecom
G. Mishra
Verizon Inc.
Y. Liu
China Mobile
Y. Fan
Casa Systems
L. Liu
Fujitsu
X. Liu
Volta Networks
January 2, 2022

**BIER Fast ReRoute
draft-chen-bier-frr-04**

Abstract

BIER is a scalable multicast overlay [[RFC8279](#)] that utilizes a routing underlay, e.g., IP, to build up its Bit Index Forwarding Tables (BIFTs). This document proposes Fast Reroute Extensions for BIER (BIER-FRR). It protects BIER traffic after detecting the failure of a link or node in the core of a BIER domain until affected BIFT entries are recomputed after reconvergence of the routing underlay. The BIER-FRR extensions are applied locally at the point of local repair (PLR) and do not introduce any per-flow state. The document specifies nomenclature for BIER-FRR and gives examples for its integration in BIER forwarding. Furthermore, it presents operation modes for BIER-FRR. Link and node protection may be chosen as protection level. Moreover, the backup strategies tunnel-based BIER-FRR and LFA-based BIER-FRR are defined and compared.

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 July 6, 2022.

Copyright Notice

Copyright (c) 2022 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	3
2.	Extensions for BIER-FRR	5
2.1.	Definition of Forwarding Actions	5
2.2.	Definition of Backup Forwarding Entries	5
2.3.	Activating and Deactivating Backup Forwarding Entries	6
2.4.	Computation of the Backup F-BM	7
3.	Representations for BIER-FRR Forwarding Data	7
3.1.	Potential Emergence of Redundant Packets	7
3.2.	Primary BIFT and Single Backup BIFT	9
3.3.	Primary BIFT and Failure-Specific Backup BIFTs	10
4.	Protection Levels	11
4.1.	Link Protection	11
4.2.	Node Protection	12
4.3.	Example	12

5.	Backup Strategies	12
5.1.	Tunnel-Based BIER-FRR	12
5.1.1.	Tunnel-Based BIER-FRR with Link Protection	13
5.1.2.	Tunnel-Based BIER-FRR with Node Protection	14
5.1.3.	Implementation Experience	16
5.2.	LFA-based BIER-FRR	16
5.2.1.	Relation of BIER-LFAs to IP-LFAs and Prerequisites	16
5.2.2.	Definition of BIER-LFAs	16
5.2.3.	Protection Coverage of BIER-LFA Types	17
5.2.4.	Sets of Supported BIER-LFAs	18
5.2.5.	Link Protection	18
5.2.6.	Node Protection	20
5.2.7.	Optimization Potential to Reduce Redundant BIER Packets in Failure Cases	22
6.	Comparison	22
6.1.	Comparison of LFA-Based Protection for IP-FRR and BIER- FRR	22
6.2.	Advantages and Disadvantages of Tunnel-Based BIER-FRR	23
6.2.1.	Advantages	23
6.2.2.	Disadvantages	23
6.3.	Advantages and Disadvantages of LFA-Based BIER-FRR	24
6.3.1.	Advantages	24
6.3.2.	Disadvantages	24
7.	Security Considerations	24
8.	IANA Considerations	24
9.	Contributors	25
10.	Acknowledgements	25
11.	References	25
11.1.	Normative References	25
11.2.	Informative References	26
Appendix A.	Specific Backup Strategy Examples	26
A.1.	LFA-based BIER-FRR using Single BIFT	26
A.2.	LFA-based BIER-FRR using Multiple Backup BIFTs	28
	Authors' Addresses	30

[1.](#) Introduction

With BIER [[RFC8279](#)], a Bit-Forwarding Router (BFR) forwards BIER packets based on a bitstring in the BIER header using the information in the Bit Index Forwarding Table (BIFT). Its entries are locally derived from a routing underlay or set by a controller. In case of a persistent link or node failure, BIER traffic may not be delivered until the BIFT has been updated based on the reconverged routing underlay or by the controller.

BIER packets are usually forwarded without an outer IP header. If a link or node fails, the corresponding BFR neighbor (BFR-NBR) is no longer reachable. Fast reroute (FRR) mechanisms in the routing

underlay, e.g., IP-FRR, apply only to IP packets so that BIER traffic would be dropped. BIER traffic can be delivered again only after reconvergence of the routing underlay and recalculation of the BIFT. Thus, tunneling BIER packets can help to reach the BFR-NBR in case of a link failure by leveraging FRR capabilities of the routing underlay if such mechanisms are available. However, this does not help in case of a node failure. Then, all destinations having the failed node as BFR-NBR cannot be reached anymore. As BIER carries multicast traffic which has often realtime requirements, there is a particular need to protect BIER traffic against too long outages after failures.

In this document we propose nomenclature for Fast Reroute Extensions for BIER (BIER-FRR). As soon as a BFR detects a BFR-NBR is unreachable, BIER-FRR enables a BFR to quickly reroute affected BIER packets with the help of backup forwarding entries. To avoid redundant packets, backup forwarding entries should be processed prior to normal forwarding entries. To achieve that goal, two possible representations for backup forwarding entries are proposed.

The protection level can be either link protection or node protection. Link protection protects only the failure of a link. It is simple but may not work if a BFR fails. Node protection is more complex but also protects against the failure of BFRs. The backup strategy determines the selection of the backup forwarding entries.

Examples for backup strategies are tunnel-based BIER-FRR and LFA-based BIER-FRR

- o Tunnel-based BIER-FRR leverages mechanisms of the routing underlay for FRR purposes. The routing underlay restores connectivity faster than BIER as a reconverged routing underlay is prerequisite for recalculation of the BIFT. If the routing underlay leverages FRR mechanisms, its forwarding ability is restored long before reconvergence is completed. To leverage fast restoration of the routing underlay, BIER traffic affected by a failure is tunneled over the routing underlay.
- o LFA-based BIER-FRR reroutes BIER traffic to alternative neighbors in case of a failure. It utilizes the principles of IP-FRR but requires that LFAs are BFRs. Normal BIER-LFAs can be reached without tunneling, remote BIER-LFAs utilize a tunnel, and topology-independent BIER-LFAs leverage explicit paths to reach the backup BFR-NBR. In contrast to tunnel-based FRR, LFA-based BIER-FRR does not require fast reroute mechanisms in the routing underlay.

BIER-FRR as presented in this document follows a primary/backup path principle, also known as 1:1 protection. It is opposite to 1+1

protection which denotes a live-live protection principle. This has been considered for BIER in [[BrA117](#)].

2. Extensions for BIER-FRR

In this section, forwarding actions and backup forwarding entries are defined. Then, the computation of the backup F-BM and the BIER forwarding process with BIER-FRR are explained.

2.1. Definition of Forwarding Actions

A BFR-NBR is directly connected if it is a next hop on the network layer, i.e., if it can be reached via the link layer technology. Otherwise, the BFR-NBR is indirectly connected.

We define the following forwarding actions.

- o Plain: Sends the mere BIER packet to a BFR-NBR via a direct link and without a tunnel header. That means, the packet is not sent over the routing underlay.
- o Tunnel: Encapsulates the BIER packet with a tunnel header towards a BFR-NBR and sends it over the routing underlay.
- o Explicit: Forwards the packet over an explicit path to a BFR-NBR. The path information must be given. If segment routing is used for this purpose, the segment IDs (SIDs) must be given. Two forwarding actions of type Explicit are equal only if they share the same explicit path.

The forwarding actions in the BIFT as proposed in [[RFC8279](#)] are given implicitly as they are derived from the connectedness of the BFR-NBR. If the BFR-NBR is directly connected, the forwarding action is Plain. If the BFR-NBR is not directly connected, the forwarding action is Tunnel.

2.2. Definition of Backup Forwarding Entries

The BIFT as proposed in [[RFC8279](#)] contains a F-BM and a BFR-NBR for a specific BFER. They constitute a primary forwarding entry. BIER-FRR extends this regular BIFT by additional columns containing backup forwarding entries. A backup forwarding entry contains

- o a backup F-BM (BF-BM),
- o a backup BFR-NBR (BBFR-NBR),
- o a backup forwarding action (BFA), and

- o a backup entry active (BEA) flag.

Backup F-BM and backup BFR-NBR have the same structure as their primary counterparts. The backup forwarding action is a forwarding action as defined in [Section 2.1](#). The BEA flag indicates whether the backup forwarding entry is active. When it is active, the backup F-BM, backup BFR-NBR, and the backup forwarding action are used for the forwarding of BIER packets instead of the primary forwarding entry. The structure of the extended BIFT is given in Figure 1.

+-----+	+-----+	+-----+	+-----+	+-----+	+-----+	+-----+
BFR-id	F-BM	BFR-NBR	BF-BM	BBFR-NBR	BFA	BEA
+=====+	+=====+	+=====+	+=====+	+=====+	+=====+	+=====+
...
+-----+	+-----+	+-----+	+-----+	+-----+	+-----+	+-----+

Figure 1: Structure of an extended BIFT with backup forwarding entries.

The primary action is not given in the BIFT as it is derived from the BFR-NBR. In contrast, the backup forwarding action is given in the extended BIFT. Moreover, an explicit path must be indicated in case of forwarding action Explicit. However, explicit paths are implementation-specific and, therefore, this information is not indicated in the table. The values for the backup BFR-NBR and the backup action depend on the desired protection level and the backup strategy. Examples for them are described in [Section 5.1](#) and [Section 5.2](#). The backup F-BM depends on the backup BFR-NBR. Its computation is explained in [Section 2.4](#).

2.3. Activating and Deactivating Backup Forwarding Entries

When a primary BFR-NBR is not reachable over the implicit primary action, a failure is observed. Then, the BEA flag of the corresponding backup forwarding entry is set.

If the primary BFR-NBR is directly connected, the information about the failed interface is sufficient to detect its unreachability. If the primary BFR-NBR is indirectly connected, a BFD session between the BFR as PLR and the BFR-NBR may be used to monitor its reachability.

If the primary BFR-NBR is reachable again, the BEA flag is deactivated. This may be caused by the disappearance of the failure or by a change of the primary BFR-NBR due to a reconfiguration of the BIFT.

2.4. Computation of the Backup F-BM

The primary F-BM of a specific BFER indicates all BFERs that share the same primary BFR-NBR. The backup F-BM of a specific BFER indicates

- o all BFERs that share the primary and backup BFR-NBR of the specific BFER and
- o all BFERs that have the backup BFR-NBR of the specific BFER as primary BFR-NBR.

3. Representations for BIER-FRR Forwarding Data

We show that backup entries need to be used first to reduce the number of redundant packets in the single extended BIFT (presented in [Section 2.2](#)). This may be hard or cannot be achieved on some hardware platforms. Therefore, two alternate representations of forwarding entries are proposed. The first is a primary BIFT and single backup BIFT (SBB). The second is a primary BIFT and multiple failure-specific backup BIFTs (FBB).

3.1. Potential Emergence of Redundant Packets

The BIER forwarding procedure in failure-free scenarios avoids redundant packets, i.e., it ensures that at most a single copy is sent per link for every BIER packet. However, this property might be violated when BIER-FRR as presented in [Section 2](#) is applied to protect against a failure.

Figure 2 shows an example of a BIER network. BFRs have the prefix "B" and are numbered by their BFR-ids. To simplify the example, every BFR is a BFER and its bit position in the bitstring equals its BFR-id. The number on a link is its cost which is used by the routing underlay for computing the shortest paths.

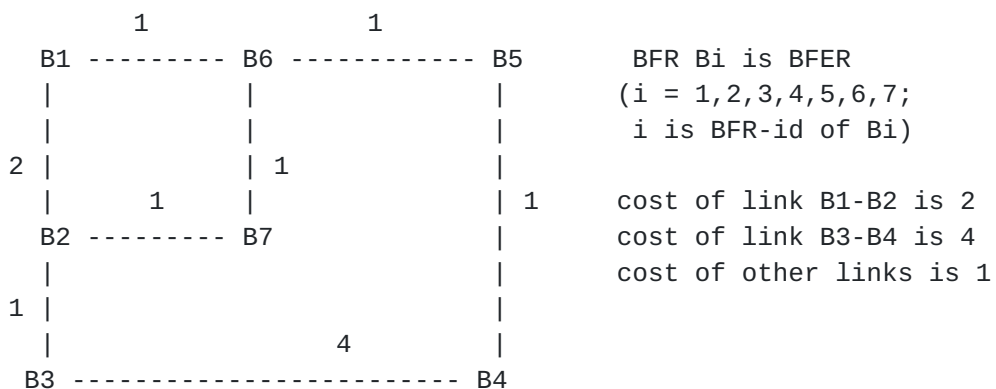


Figure 2: BIER network example.

The extended BIFT with backup forwarding entries for LFA-based BIER-FRR with link protection built by BFR B1 is illustrated in Figure 3.

BFR-id	F-BM	BFR-NBR	BF-BM	BBFR-NBR	BFA	BEA
2	0000110	B2	1111110	B6	Plain	
3	0000110	B2	1111110	B6	Plain	
4	1111000	B6	1111110	B2	Plain	
5	1111000	B6	1111110	B2	Plain	
6	1111000	B6	1111110	B2	Plain	
7	1111000	B6	1111110	B2	Plain	

Figure 3: B1's extended BIFT for LFA-based FRR with link protection.

We show how redundant packets can occur in case of a failure. To that end, we consider the extended BIFT for BFR 1 in Figure 3. It has backup forwarding entries for LFA-based FRR and link protection. For a BIER packet with destinations B2 and B6 (i.e., bitstring 0100010), BFR B1 sends a single packet copy on link B1-B2 and on link B1-B6 in the absence of a failure.

When the link B1-B6 fails, B1 as a PLR detects the failure. Therefore, B1 sets the BEA flag for all destinations that have B6 as BFR-NBR. We consider again that B1 sends a BIER packet to B2 and B6. At first, it sends a copy with bitstring 0000010 to B2 using the corresponding primary forwarding entry in the extended BIFT in Figure 3.

Then, B1 sends another copy of the packet with bitstring 0100000 for B6 to B2 using the backup forwarding entry since the BEA flag is activated.

This is a second (redundant) copy over the same link B1-B2. It can be prevented if the backup forwarding entry is used first. When using the backup forwarding entry, B1 sends only a single copy of the packet with bitstring 0100010 to B2. It will not send any copy of the packet to B2 again since the bitstring in the packet will be all cleaned by the BF-BM 1111110. Thus, prioritized processing of BFERs with unreachable BFR-NBRs helps to reduce redundant packet copies.

3.2. Primary BIFT and Single Backup BIFT

The extended BIFT may be separated into two BIFTs. One is a primary BIFT and the other is a single backup BIFT. The primary BIFT is the same as the regular BIFT. The backup BIFT contains the backup forwarding entries, including BF-BM, BBFR-NBR, BFA and BEA in the extended BIFT. When a BFR as a PLR detects that BFR-NBR N is unreachable, it activates the BEA flag for all BFERs in the backup BIFT that have BFR-NBR as primary BFR-NBR. When a BFR forwards a BIER packet, it processes the packet first using the backup BIFT and then using the primary BIFT. With this prioritization, the number of redundant packet copies can be reduced.

B1's extended BIFT in Figure 3 is separated into the primary BIFT in Figure 4 and the single backup BIFT in Figure 5.

BFR-id	F-BM	BFR-NBR
2	0000110	B2
3	0000110	B2
4	1111000	B6
5	1111000	B6
6	1111000	B6
7	1111000	B6

Figure 4: B1's primary BIFT for the BIER network example.

BFR-id	BF-BM	BBFR-NBR	BFA	BEA	Comment: protects
					failure of
2	1111110	B6	Plain		Link B1->B2
3	1111110	B6	Plain		Link B1->B2
4	1111110	B2	Plain		Link B1->B6
5	1111110	B2	Plain		Link B1->B6
6	1111110	B2	Plain		Link B1->B6
7	1111110	B2	Plain		Link B1->B6

Figure 5: B1's backup BIFT for the BIER network example.

Each forwarding entry in the backup BIFT contains BF-BM, BBFR-NBR, BFA and BEA. When a BFR-NBR fails, the BEA flag is activated for all BFRs in the backup BIFT that have BFR-NBR as primary BFR-NBR. For example, BFRs B4, B5, B6 and B7 have BFR-NBR B6 as their primary BFR-NBR. When BFR-NBR B6 fails, the BEA flag for BFRs B4, B5, B6 and B7 is activated, i.e., the BEA in the last four entries in the backup BIFT is set to one.

3.3. Primary BIFT and Failure-Specific Backup BIFTs

As an alternative, the information in the extended BIFT may be represented in a primary BIFT and several, failure-specific backup BIFTs. A failure-specific backup BIFT is a backup BIFT for the unreachability of BFR-NBR N. A backup BIFT for the failure of N is simply called a backup BIFT for N. It has the same structure as the regular BIFT but has an entry for a backup forwarding action. Thus, a BFR has a primary BIFT, which is the same as the regular BIFT, and a backup BIFT for each of its BFR-NBRs.

The BFR uses the primary BIFT to forward BIER packets under failure-free conditions. When the BFR as a PLR detects that BFR-NBR N is unreachable, it uses the backup BIFT for N to forward all BIER packets. After the routing underlay has re-converged on the new network topology, the primary BIFT is re-computed. Once the re-computed primary BIFT is installed, it is used to forward all BIER packets.

We illustrate the concept using the example from extended BIFT in Figure 3. Figure 4 shows the primary BIFT of B1 in this context.

BFR B1 in Figure 2 has two neighbors: B6 and B2. B1 has two backup BIFTs with link protection: one for B6 and another for B2. B1 has also two backup BIFTs with node protection. Figure 6 is B1's backup BIFT for B6 to react to the unreachability of B1 in a similar way as with the extended BIFT in Figure 3.

BFR-id	F-BM	BFR-NBR	Forwarding Action	Comment: protects
				failure of
2	1111110	B2	Plain	
3	1111110	B2	Plain	
4	1111110	B2	Plain	Link B1->B6
5	1111110	B2	Plain	Link B1->B6
6	1111110	B2	Plain	Link B1->B6
7	1111110	B2	Plain	Link B1->B6

Figure 6: B1's backup BIFT for B6 for LFA-based BIER FRR with link protection.

Once B1 as a PLR detects that B6 is unreachable through the link to B6, it uses the backup BIFT for B6 to forward all BIER packets. As this representation is equivalent to the concept of single primary and single backup BIFT, redundant packets for the same forwarding action are avoided.

4. Protection Levels

Link and node protection may be supported. Link protection protects against the failure of an adjacent link while node protection protects against the failure of a neighboring node and the path towards that node. Depending on the supported service, link protection or node protection may be relevant. Both protection levels can be combined with any backup strategy in [Section 5](#).

4.1. Link Protection

With link protection the backup path avoids the failed link (i.e., the failed primary path from the PLR to the primary BFR-NBR, excluding the primary BFR-NBR), but the backup path may include the primary BFR-NBR. Therefore, the backup path is still operational if the primary path fails. The disadvantage of link protection is that

it fails if the primary BFR-NBR itself is not operational. However, link protection has also advantages. It often leads to shorter backup paths than node protection. In case of tunnel-based BIER-FRR, link protection causes at most one redundant packet while node protection can cause more redundant packets. In case of LFA-based BIER-FRR, link protection can protect more BFERs with normal BIER-LFAs than node protection.

4.2. Node Protection

With node protection, the backup path avoids the failed node and the link to the node (i.e., the failed primary path from the PLR to the primary BFR-NBR, including the primary BFR-NBR). Therefore, the backup path must not include the primary path or the primary BFR-NBR so that the backup path is still operational if these elements fail. If a BFER and its primary BFR-NBR are the same, only link protection is possible for that BFER. An advantage of node protection is the improved protection quality compared to link protection. However, node protection has also disadvantages. It often leads to longer backup paths than link protection. For tunnel-based BIER-FRR, possibly more redundant packets are transmitted over a link than with link protection. For LFA-based BIER-FRR, possibly fewer BFERs can be protected with normal BIER-LFAs so that more remote BIER-LFAs or topology-independent BIER-LFAs are needed which are more complex.

4.3. Example

In Figure 2, B1's primary path towards BFER B5 is B1-B6-B5. Node protection for BFER B5 can be achieved only via the backup path B1-B2-B3-B4-B5. Link protection for BFER 5 is achieved via the backup path B1-B2-B7-B6 and in addition via the backup path B1-B2-B3-B4-B5-B6. The backup entries depend on the protection level and on the backup strategy. Example BIFTs for link and node protection are given in [Section 5](#).

5. Backup Strategies

The backup strategies determine the selection of the backup forwarding entries. They have an impact on the backup BFR-NBR and on the backup action, and thereby on the backup path. In the following, tunnel-based BIER-FRR and LFA-based BIER-FRR are presented.

5.1. Tunnel-Based BIER-FRR

The routing underlay may be able to forward packets towards their destinations despite an existing failure. This may be achieved, e.g., due to FRR mechanisms in the routing underlay. In that case,

the primary BFR-NBR is not reachable via the primary action (Plain), but it may be reachable via a backup action (Tunnel).

Tunnel-based BIER-FRR encapsulates BIER packets affected by a failure in the routing underlay to leverage its fast restoration capability. The affected BIER packets can be delivered towards their destinations as soon as the connectivity in the routing underlay is restored. The appropriate backup forwarding entries in a BIFT for BIER-FRR depend on the desired protection level.

5.1.1. Tunnel-Based BIER-FRR with Link Protection

With link protection, the backup BFR-NBRs equal the primary BFR-NBRs. If a primary BFR-NBR is directly connected to the BFR as a PLR, the corresponding backup forwarding action is Tunnel. As a result, the BIER packets affected by a failure are tunneled over the routing underlay to their BFR-NBR instead of being sent directly as plain BIER packets to the BFR-NBR. If a primary BFR-NBR is not directly connected to the BFR as a PLR (i.e., the implicit, primary action is Tunnel), the corresponding backup action is also Tunnel. The backup F-BMs are the same as the primary F-BMs, which is in line with the computation of the backup F-BMs in [Section 2.4](#).

BFR-id	BF-BM	BBFR-NBR	BFA	BEA	Comment: protects
					failure of
2	0000110	B2	Tunnel		Link B1->B2
3	0000110	B2	Tunnel		Link B1->B2
4	1111000	B6	Tunnel		Link B1->B6
5	1111000	B6	Tunnel		Link B1->B6
6	1111000	B6	Tunnel		Link B1->B6
7	1111000	B6	Tunnel		Link B1->B6

Figure 7: B1's backup BIFT for tunnel-based BIER-FRR with link protection.

Figure 7 shows B1's backup BIFT for tunnel-based BIER-FRR with link protection for the BIER network example of Figure 2. The backup BFR-NBRs and backup F-BMs in this backup BIFT are the same as the primary BFR-NBRs and primary F-BMs in the primary BIFT in Figure 4, but the backup actions in this backup BIFT are Tunnel while the primary

actions in the primary BIFT are Plain (which are not shown, but implied).

When B1 as a PLR detects failure of its link to B6, a BIER packet with bitstring 01000000 for B6 is tunneled by B1 towards B6 via the routing underlay. The exact path of the backup tunnel depends on the routing underlay. It may be B1-B2-B7-B6 or B1-B2-B3-B4-B5-B6.

If a BIER packet is destined to {B2, B5, B7}, first an encapsulated packet copy is forwarded via link B1-B2 to backup BFR-NBR B6 with backup action Tunnel to deliver packet copies to BFER B5 and B7. Then, a non-encapsulated packet copy is forwarded via link B1-B2 to BFR-NBR B2 with primary action Plain to deliver a packet copy to BFER B2. Thus, with tunnel-based BIER-FRR, a single redundant packet copy can occur in case of a failure because an encapsulated packet copy and a non-encapsulated packet copy are forwarded over the same link. This happens although BIER packets affected by failures are forwarded before BIER packets not affected by failures.

A BIER packet with bitstring 10000000 for B7 is forwarded on the backup path B1-B2-B7-B6-B7 as it is first delivered to the backup BFR-NBR B6. Thus, the backup path can be unnecessarily long. This phenomenon is known from facility backup method in [[RFC4090](#)] which takes similar paths as tunnel-based BIER-FRR.

5.1.2. Tunnel-Based BIER-FRR with Node Protection

To determine the backup forwarding entries with node protection, a case analysis for the BFER to protect is needed. If the BFER is the same as its primary BFR-NBR, node protection is not possible for that BFER. Therefore, link protection is applied, i.e., the backup BFR-NBR is set to the primary BFR-NBR. If that level of protection is not sufficient, egress protection in [[I-D.chen-bier-egress-protect](#)] may be applied. Otherwise (i.e., the BFER is different from its primary BFR-NBR), the backup BFR-NBR is set to the primary BFR-NBR's primary BFR-NBR for that BFER, i.e., the backup BFR-NBR is a next hop BFR. In all cases, the backup action is Tunnel. In the first case, the backup F-BM is set to all zeroes plus the bit enabled for the BFER to protect. In the second case, the backup F-BM is computed in the way described in [Section 2.4](#).

BFR-id	BF-BM	BBFR-NBR	BFA	BEA	Comment: protects
					failure of
2	0000010	B2	Tunnel		Link B1->B2
3	0000100	B3	Tunnel		BFR-NBR B2
4	0011000	B5	Tunnel		BFR-NBR B6
5	0011000	B5	Tunnel		BFR-NBR B6
6	0100000	B6	Tunnel		Link B1->B6
7	1000000	B7	Tunnel		BFR-NBR B6

Figure 8: B1's backup BIFT for tunnel-based BIER-FRR with node protection.

Figure 8 shows B1's backup BIFT for tunnel-based BIER-FRR with node protection for the BIER network example in Figure 2. BFRs B2 and B6 are direct neighbors of B1. To protect them, only link protection is applied as B1's primary BFR-NBR for them are those nodes themselves. According to the description above, only the bit for B2 is set in the backup F-BM of B2. The same holds for B6. For BFER B5, the backup BFR-NBR is B5 as it is B1's next next hop BFR towards B5. Similarly, for BFER B7, the backup BFR-NBR is B7. When B1 as a PLR detects the failure of its BFR-NBR B6, a BIER packet with bitstring 1010010 for {B2, B5, B7} is processed as follows. An encapsulated copy of the packet is sent via tunnel B1-B2-B3-B4-B5, another encapsulated copy is sent via tunnel B1-B2-B7, and a non-encapsulated copy is sent via link B1-B2. In this example, two redundant packets are sent on link B1-B2. Thus, with node protection, more redundant packets copies may be sent than with link protection.

Caveat: If the routing underlay does not provide node protection, tunnel-based BIER-FRR cannot provide node protection, either. This is shown by the following example. The underlay in the networking example of Figure 2 offers only link protection. B6 fails and B1 must forward a packet to B5. According to the backup BIFT in Figure 8 the packet is tunneled towards B5 and the tunnel path B1-B2-B7-B6-B5 may be taken for this purpose by the underlay due to FRR with link protection. However, B6 is also unreachable at B7 so that the packet is returned to B2 and the packet loops between B2 and B7.

5.1.3. Implementation Experience

Tunnel-based BIER-FRR has been implemented in P4 for the software-switch bmv2 [[MeLi20b](#)] and the hardware switching ASIC Tofino [[MeLi21](#)]. Performance results have been provided.

5.2. LFA-based BIER-FRR

LFA-based BIER-FRR leverages alternate BFRs to deliver BIER packets to BFRs for which the primary BFR-NBR is unreachable. It does not rely on any fast restoration/protection mechanisms in the underlay. First, some prerequisites for LFA-based BIER-FRR are clarified, BIER-LFAs are defined, and then link and node protection for LFA-based BIER-FRR are discussed using a single backup BIFT.

5.2.1. Relation of BIER-LFAs to IP-LFAs and Prerequisites

A loop-free alternate (LFA) for a specific destination is an alternate node to which a packet is sent if the primary next hop for this destination is not reachable. This alternate node should be able to forward the packet without creating a forwarding loop. LFAs have been defined for IP networks in [[RFC5286](#)], [[RFC7490](#)] and [[I-D.ietf-rtgwg-segment-routing-ti-lfa](#)]. We denote such LFAs as IP-LFAs. BIER-LFAs are very similar to IP-LFAs, but a BIER-LFA node must be a BFR. If only a subset of the nodes in the routing underlay are BFRs, some IP-LFAs in the routing underlay may not be usable as BIER-LFAs. To compute BIER-LFAs, network topology and link cost information from the routing underlay are needed. This is a difference to tunnel-based BIER-FRR where knowledge about the primary BIFTs of a PLR and its BFR-NBRs is sufficient.

LFA-based BIER-FRR may reuse IP-LFAs in the following sense as BIER-LFAs. If an IP-LFA node for the destination of a specific BFER is a BFR, it may be reused as backup BFR-NBR for that BFER together with the backup action that is applied for that IP-LFA on the IP layer. A normal IP-LFA corresponds to backup action plain, a remote IP-LFA to Tunnel, and a TI-IP-LFA to Explicit.

5.2.2. Definition of BIER-LFAs

As for IP-LFAs, there are several, different types of BIER-LFAs:

- o A BFR is a normal BIER-LFA for a specific BFER if it is directly connected to the PLR and
 1. the BFER can be reached from it through the BIER domain

2. both the path from the PLR to it and the path from it to the BFER are disjoint with the primary path from the PLR to the primary BFR-NBR. These paths
 - + may contain the primary BFR-NBR for link protection, and
 - + must not contain the primary BFR-NBR for node protection.
- o A BFR is a remote BIER-LFA for a specific BFER if it is not directly connected to the PLR, if it can be reached via a tunnel from the PLR, and if it also satisfies the aforementioned conditions 1 and 2.
- o A BFR is a TI-BIER-LFA for a specific BFER if it is not directly connected to the PLR, if it cannot be reached via a tunnel from the PLR, if it is reachable from the PLR via an explicit path (i.e., with the help of a SR header), and if it also satisfies the aforementioned conditions 1 and 2.

For some BFERs, one or more normal BIER-LFAs are available at a specific PLR. For other BFERs, only remote and TI-LFAs are available. And there may be some BFERs, for which only TI-LFAs are available.

The backup actions to reroute BIER packets depending on the BIER-LFA types are:

- o For normal BIER-LFA: Plain
- o For remote BIER-LFA: Tunnel
- o For TI-BIER-LFA: Explicit

5.2.3. Protection Coverage of BIER-LFA Types

The protection coverage is the set of BFERs that can be protected with a desired protection level by a certain BIER-LFA type. The BIER-LFA types have the following properties:

- o Normal BIER-LFAs
 - * The protection coverage is the least because some or many BFERs cannot be protected with the desired protection level or even not at all.
 - * Redundant packet copies are avoided.
 - * No encapsulation overhead.

- o Remote BIER-LFAs
 - * They increase the protection coverage of normal BIER-LFAs.
 - * Redundant packet copies may occur on a link similar to tunnel-based BIER-FRR.
 - * Same encapsulation overhead as with tunnel-based BIER-FRR.
- o TI-BIER-LFAs
 - * They complement the protection coverage of normal and remote BIER-LFAs to 100%.
 - * Redundant packets may occur on a link similar to tunnel-based BIER-FRR.
 - * Same or similar encapsulation overhead as with tunnel-based BIER-FRR depending on the FRR mechanism in the routing underlay.

5.2.4. Sets of Supported BIER-LFAs

Normal BIER-LFAs are simplest, as they require neither tunneling nor explicit paths. Remote BIER-LFAs are more powerful, but entail more header overhead and require more functionality from the PLR. TI-BIER-LFAs are most complex as they require the use of explicit paths. When LFA-based BIER-FRR is utilized, the set of supported BIER-LFAs must be indicated. The following options are available:

- o Option 1: only normal BIER-LFAs are supported
- o Option 2: normal and remote BIER-LFAs are supported
- o Option 3: all BIER-LFA types are supported

5.2.5. Link Protection

With link protection, normal BIER-LFAs are preferred over remote LFAs and remote BIER-LFAs are preferred over TI-BIER-LFAs. Depending on the set of supported BIER-LFAs, a BFER may not be protectable.

Figure 5 illustrates B1's backup BIFT for LFA-based BIER-FRR with link protection in the networking example of Figure 2.

If the link B1-B6 fails, B1 cannot reach the BFERs B4, B5, B6, and B7 over their primary BFR-NBR. Therefore, B1 sends their traffic via the backup BFR-NBR B2 together with the traffic for B2 and B3 as B2

is their primary BFR-NBR. As a consequence, the backup F-BM is 1111110 in that case. Likewise, if the link B1-B2 fails, B1 sends all traffic to B6, and the backup F-BM is 1111110 also in that case.

B1 requires only normal BIER-LFAs to protect all BFERs. This can be substantially different for other BFRs. Figure 9 and Figure 10 show the backup BIFTs for B7 and B5 respectively. BFR B7 requires one normal BIER-LFA, three remote BIER-LFAs, and two TI-BIER-LFAs to protect all BFERs. And BFR B5 requires even one normal BIER-LFA, one remote BIER-LFA, and four TI-BIER-LFAs as backup BFR-NBRs. Thus, depending on the set of supported BIER-LFAs, a BFER cannot be protected by BIER-FRR.

We now assume B7 has a BIER packet with destinations {B1, B4, B5, B6}. When link B7-B6 fails, the packet copy for B1 is sent to B2 using forwarding action Plain, the packet copy to B4 is tunneled via B2 to B3, and the packet copies towards B5 and B6 are sent via explicit paths towards B4 and B1 respectively. As these packet copies have different headers, they all need to be sent. Hence, we observe three redundant copies.

+-----+-----+-----+-----+-----+-----+						
BFR-id	BF-BM	BBFR-NBR	BFA	BEA	Comment: protects	
					failure of	
+=====+=====+=====+=====+=====+=====+						
1	0000111	B2	Plain		Link B7->B6	
+-----+-----+-----+-----+-----+-----+						
2	0000110	B1	Tunnel		Link B1->B2	
+-----+-----+-----+-----+-----+-----+						
3	0000110	B1	Tunnel		Link B1->B2	
+-----+-----+-----+-----+-----+-----+						
4	0001000	B3	Tunnel		Link B1->B6	
+-----+-----+-----+-----+-----+-----+						
5	0010000	B4	Explicit		Link B1->B6	
+-----+-----+-----+-----+-----+-----+						
6	0100000	B1	Explicit		Link B1->B6	
+-----+-----+-----+-----+-----+-----+						

Figure 9: B7's backup BIFT with link protection.

BFR-id	BF-BM	BBFR-NBR	BFA	BEA	Comment: protects
					failure of
1	1100011	B3	Explicit		Link B5->B6
2	1100011	B3	Explicit		Link B5->B6
3	0000100	B4	Plain		Link B5->B6
4	0001000	B3	Tunnel		Link B5->B4
6	1100011	B3	Explicit		Link B5->B6
7	1100011	B3	Explicit		Link B5->B6

Figure 10: B5's backup BIFT with link protection.

5.2.6. Node Protection

To determine the backup forwarding entries with node protection, a case analysis for the BFER to protect is needed again. If the BFER is the same as its primary BFR-NBR, node protection is not possible for that BFER. In this case, link protection is applied. Otherwise, the BFER must be protected by a node-protecting BIER-LFA. Thereby, normal BIER-LFAs are preferred over remote BIER-LFAs and remote BIER-LFAs are preferred over TI-BIER-LFAs. Depending on the set of allowed BIER-LFAs, a BFER may not be protectable.

Figure 11 illustrates B1's backup BIFT for the LFA-based BIER-FRR with node protection in the networking example of Figure 2.

BFR-id	BF-BM	BBFR-NBR	BFA	BEA	Comment: protects
					failure of
2	1111010	B6	Plain		BFR-NBR B2
3	0000100	B4	Tunnel		BFR-NBR B2
4	0001000	B3	Tunnel		BFR-NBR B6
5	0010000	B4	Explicit		BFR-NBR B6
6	1100100	B2	Plain		BFR-NBR B6
7	1100100	B2	Plain		BFR-NBR B6

Figure 11: B1's backup BIFT with node protection.

As the primary BFR-NBR of B1 for BFER B6 is B6 itself, only link protection can be applied. Therefore, B2 is used as normal, link-protection BIER-LFA to protect B6. Likewise, the primary BFR-NBR of B1 for BFER B2 is B2, and therefore, B2 is protected with B6 as normal, link-protecting BIER-LFA. BFER B7 is protected against the failure of node B6 with B2 as normal, node-protecting BIER-LFA as B2 has a shortest path towards B7 that does not traverse B6. The backup F-BMs for BFER 6 and BFER 7 are {B2, B6, B7} because if B6 is unreachable, the traffic for these BFERs is sent via link B1-B2 with forwarding action Plain.

BFER B4 is not reachable through a normal LFA when BFR B6 fails. However, B3 is a remote, node-protecting BIER-LFA for BFER B4 because B3 has a shortest path towards B4, and B3 is reachable through a shortest path from B1, and the resulting backup path from B1 to B4 does not traverse B6. Likewise, B4 is a remote LFA for BFER B3 if BFR B2 fails.

BFER B5 is neither reachable through a normal BIER-LFA nor through a remote BIER-LFA when BFR B6 fails. However, B4 is a node-protecting TI-LFA for BFER B5 because B4 has a shortest path towards B5 that does not traverse B6. Moreover, B4 is reachable through the explicit path B1-B2-B3-B4.

5.2.7. Optimization Potential to Reduce Redundant BIER Packets in Failure Cases

Redundant packets occur with LFA-based BIER-FRR if BIER packets are sent over a specific link in different forms. These forms are

- o plain BIER packets (plain primary transmission or reroute to normal BIER-LFA)
- o BIER packets encapsulated to a specific BFR-NBR (tunneled primary transmission or reroute to remote BIER-LFA)
- o BIER packets with an encoded explicit path (reroute to TI-LFA)

When different remote LFAs are addressed, even multiple redundant packets can be caused through remote LFAs. The same can happen with TI-LFAs. Some redundant packets can be avoided if remote LFAs or TI-LFAs are chosen such that they can protect several BFERs and thereby avoid the need for another remote LFA or TI-LFA. However, this may lead to longer backup paths. This is a new, potential optimization objective for the choice of remote or TI-BIER-LFAs which does not exist for IP-FRR. Its relevance may depend on the use case.

We illustrate this optimization potential. We consider LFA-based BIER-FRR with link protection for B7. Its backup BIFT is given in Figure 9. As observed in [Section 5.2.5](#), B7 needs to send four copies to forward a packet to {B1, B4, B5, B6}. If we choose the more complex TI-BIER-LFA B4 to protect BFER B4 instead of the remote BIER-LFA B3, then only two redundant copies need to be sent.

6. Comparison

This section first discusses the difference of IP-LFAs for IP-FRR and BIER-LFAs for BIER-FRR. Then it discusses advantages and disadvantages of tunnel-based and LFA-based BIER-FRR.

6.1. Comparison of LFA-Based Protection for IP-FRR and BIER-FRR

LFAs have been first proposed for IP networks. They are simple in the sense that they do not require any tunneling overhead. However, some destinations cannot be protected against some link failures and even more destinations cannot be protected against some node failures. Therefore, remote LFAs (R-LFAs) have been defined to improve that coverage by tunneling the affected traffic to another node from where the traffic can reach the destination via normal forwarding. Nevertheless, there may be still some destinations that cannot be protected against link or node failures. Therefore, topology-independent LFAs (TI-LFAs) have been defined where affected

traffic is tunneled via an explicit path (preferably using segment routing headers) to another node from where the traffic can reach its destination via normal IP forwarding. With TI-LFAs, all destinations can be protected against any failures as long as connectivity exists.

LFA-based BIER-FRR adopts the idea of LFAs. It differs from IP-FRR as the LFA target node, i.e., the node to which the traffic is deviated, must be a BFR. If an IP-LFA target is a BFR, it can be utilized as a BIER-LFA; otherwise it cannot serve as BIER-LFA. Thus, if only some nodes of the underlay are BFRs, the BIER-LFAs will be substantially different from IP-LFAs. Moreover, this makes it more difficult to find normal LFAs for which tunneling is not needed. That means, LFA-based BIER-FRR is likely to require more remote LFAs and TI-LFAs than IP-FRR under such conditions.

6.2. Advantages and Disadvantages of Tunnel-Based BIER-FRR

6.2.1. Advantages

- o Computation of backup forwarding entries is very simple. Only primary BIFTs of a PLR and its BFR-NBRs are needed to compute the backup forwarding entries. Routing information from the routing underlay is not needed.
- o The forwarding action Explicit is not needed. However, depending on the underlay, explicit forwarding may be used to achieve FRR in the underlay.

6.2.2. Disadvantages

- o It requires a FRR mechanism on the underlay.
- o It is limited to the protection level of the underlay. E.g., if the underlay supports only link protection, tunnel-based BIER-FRR cannot provide node protection.
- o Redundant packet copies may occur in tunnel-based BIER-FRR.
- o In case of node protection, backup paths may be extended more than needed.
- o Requires a tunneling header for any rerouting, which creates header overhead.

6.3. Advantages and Disadvantages of LFA-Based BIER-FRR

6.3.1. Advantages

- o Does not rely on any fast protection of the underlay.
- o Can provide better protection on the BIER layer than on the IP layer; this is possible if LFA-based BIER-FRR utilizes BIER-LFAs with better protection level than LFA-based IP-FRR. E.g., the underlay may provide only FRR with link protection while BIER-FRR may provide node protection for BIER traffic.
- o Avoids header overhead for normal BIER-LFAs.

6.3.2. Disadvantages

- o Computation of backup forwarding entries requires routing information from the underlay.
- o Computation of backup forwarding entries more complex if some nodes of the underlay are not BFRs.
- o Need for forwarding action Tunnel to protect some BFRs, which adds header overhead.
- o Need for forwarding action Explicit to achieve full protection coverage for some topologies; otherwise only partial protection coverage. This requires support for explicit paths, e.g., segment routing.
- o More remote and TI-LFAs needed than for IP-FRR if some nodes in the routing underlay are not BFRs.
- o Redundant packet copies may occur in LFA-based BIER-FRR (but it's less than with tunnel-based BIER-FRR).

7. Security Considerations

TBD.

8. IANA Considerations

No requirements for IANA.

9. Contributors

Daniel Merling
Germany
Email: daniel.merling@uni-tuebingen.de

Xuesong Geng
China
Email: gengxuesong@huawei.com

10. Acknowledgements

The authors would like to thank Jeffrey Zhang, Tony Przygienda and Shaofu Peng for their comments to this work.

11. References

11.1. 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>>.
- [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>>.
- [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>>.
- [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>>.

11.2. Informative References

- [BrAl17] Braun, W., Albert, M., Eckert, T., and M. Menth, "Performance Comparison of Resilience Mechanisms for Stateless Multicast Using BIER", May 2017.
- [I-D.chen-bier-egress-protect]
Chen, H., McBride, M., Wang, A., Mishra, G. S., Liu, Y., Menth, M., Khasanov, B., Geng, X., Fan, Y., Liu, L., and X. Liu, "BIER Egress Protection", [draft-chen-bier-egress-protect-03](#) (work in progress), October 2021.
- [I-D.ietf-rtgwg-segment-routing-ti-lfa]
Litkowski, S., Bashandy, A., Filsfils, C., Francois, P., Decraene, B., and D. Voyer, "Topology Independent Fast Reroute using Segment Routing", [draft-ietf-rtgwg-segment-routing-ti-lfa-07](#) (work in progress), June 2021.
- [MeLi20b] Merling, D., Lindner, S., and M. Menth, "P4-Based Implementation of BIER and BIER-FRR for Scalable and Resilient Multicast", November 2020.
- [MeLi21] Merling, D., Lindner, S., and M. Menth, "Hardware-based Evaluation of Scalable and Resilient Multicast with BIER in P4", March 2020.
- [RFC4090] Pan, P., Ed., Swallow, G., Ed., and A. Atlas, Ed., "Fast Reroute Extensions to RSVP-TE for LSP Tunnels", [RFC 4090](#), DOI 10.17487/RFC4090, May 2005, <<https://www.rfc-editor.org/info/rfc4090>>.

Appendix A. Specific Backup Strategy Examples

This appendix demonstrates the computations of some specific backup strategy options in details.

A.1. LFA-based BIER-FRR using Single BIFT

In the LFA-based BIER-FRR using single BIFT, every BFR has a single BIFT for a level of protection. Its structure is the same as the one in Figure 1.

The following presents the details in BFR B1 in Figure 2 for building the BIFT for BIER-FRR link protection.

At first, BFR B1 obtains a BIER-LFA as BBFR-NBR for each BFER. B6 is normal BIER-LFA as BBFR-NBR for BFER B2 and B3. B2 is normal BIER-LFA as BBFR-NBR for BFER B4, B5, B6 and B7. Figure 12 illustrates

B1's intermediate BIFT for link protection filled with values for BBFR-NBRs and BFAs.

BFR-id	F-BM	BFR-NBR	BF-BM	BBFR-NBR	BFA	BEA
2	0000110	B2		B6	Plain	
3	0000110	B2		B6	Plain	
4	1111000	B6		B2	Plain	
5	1111000	B6		B2	Plain	
6	1111000	B6		B2	Plain	
7	1111000	B6		B2	Plain	

Figure 12: B1's intermediate BIFT for link protection.

From the intermediate BIFT, BFRs B2 and B3 have the same BFR-NBR B2 and BBFR-NBR B6, BFRs B4 to B7 have the same BFR-NBR B6 as the BBFR-NBR B6 for BFER B2. According to [Section 2.4](#), the BF-BM for BFER B2 has the bits for B2 and B3 as well as the bits for B4 to B7, which is 1111110. The BF-BM for BFER B3 is also 1111110. Similarly, the BF-BM for each of BFRs B3 to B7 is computed, which is 1111110.

With the BF-BMs, BFR B1 has the BIFT for link protection, which is illustrated in Figure 13.

BFR-id	F-BM	BFR-NBR	BF-BM	BBFR-NBR	BFA	BEA
2	0000110	B2	1111110	B6	Plain	
3	0000110	B2	1111110	B6	Plain	
4	1111000	B6	1111110	B2	Plain	
5	1111000	B6	1111110	B2	Plain	
6	1111000	B6	1111110	B2	Plain	
7	1111000	B6	1111110	B2	Plain	

Figure 13: B1's BIFT for BIER-FRR link protection.

A.2. LFA-based BIER-FRR using Multiple Backup BIFTs

For the LFA-based BIER-FRR using multiple backup BIFTs, in addition to a primary BIFT, a BFR has a backup BIFT for each of its BFR-NBRs with a level of protection. The backup BIFT for BFR-NBR N with link protection (or simply called the backup BIFT for link to N) assumes that the link to N failed. The BFR uses it to protect against the failure of its link to N. The backup BIFT for N with node protection (or simply called the backup BIFT for N) assumes that node N failed. The BFR uses it to protect against the failure of N. Once the BFR as a PLR detects the failure of its link to N, it uses the backup BIFT for link to N to forward all BIER packets. When the BFR as a PLR detects the failure of its BFR-NBR N, it uses the backup BIFT for N to forward all BIER packets.

Even though a BFR has multiple backup BIFTs, the LFA-based BIER-FRR using multiple backup BIFTs is scalable. Both the size of a backup BIFT and the number of backup BIFTs on the BFR are small. Assume a BIER network has 1000 BFRs and 100 BFRs, and each BFR has 10 BFR-NBRs on average. The size of a backup BIFT is 100 forwarding entries. The number of backup BIFTs on the BFR is 20 on average. The total size of all backup BIFTs is $100 * 20 = 2000$ forwarding entries.

The following presents the details in BFR B1 in Figure 2 for building the backup BIFT for link to B2 to support BIER-FRR link protection.

To support link protection, BFR B1 in Figure 2 has two backup BIFTs: one for link to B2 and the other for link to B6. The backup BIFT for link to B2 is illustrated in Figure 14.

BFR-id	F-BM	BFR-NBR	Forwarding Action	Comment: protects failure of
2	1111110	B6	Plain	Link B1->B2
3	1111110	B6	Plain	Link B1->B2
4	1111110	B6	Plain	
5	1111110	B6	Plain	
6	1111110	B6	Plain	
7	1111110	B6	Plain	

Figure 14: B1's backup BIFT for link to B2.

BFR B1 builds the backup BIFT for link to B2 in two steps. In the first step, it builds the backup BIRT for link to B2 through copying its regular BIRT to the backup BIRT and then changing BFR-NBR B2 in the backup BIRT to a backup BFR-NBR to protect against the failure of the link to B2. The backup BIRT for link to B2 built by B1 is illustrated in Figure 15.

BFR-id	BFR's Prefix	BFR-NBR	Forwarding Action	Comment: protects failure of
2	B2	B6	Plain	Link B1->B2
3	B3	B6	Plain	Link B1->B2
4	B4	B6	Plain	
5	B5	B6	Plain	
6	B6	B6	Plain	
7	B7	B6	Plain	

Figure 15: B1's backup BIRT for link to B2.

The BFR-NBR in each of the first two routing entries with BFR-NBR B2 originally from the BIRT is changed to its corresponding backup BFR-NBR. The BFR-NBR B2 in the first entry is changed to backup BFR-NBR

BIER-LFA B6. The BFR-NBR B2 in the second entry is changed to backup BFR-NBR BIER-LFA B6.

In the second step, BFR B1 derives the backup BIFT for link to B2 from the backup BIRT for link to B2 in the same way as it derives its regular BIFT from its BIRT defined in [[RFC8279](#)]. The result of the backup BIFT for link to B2 is the one shown in Figure 14.

When BFR B1 as a PLR detects the failure of its link to B2, it forwards all the BIER packets using the FRR-BIFT for link to B2. There is no redundant packet. For example, for a BIER packet with destinations B2 and B6 (i.e., bitstring 0100010), BFR B1 sends a single packet copy on the link to B6 using the backup BIFT for link to B2 after detecting the failure of its link to B2. It will not send any copy of the packet to B6 again since the bitstring in the packet will be all cleaned by the F-BM 1111110 after sending the packet to B6 via its link to B6. Similarly, for a BIER packet with destinations B2, B5 and B7 (i.e., bitstring 1010010), BFR B1 sends a single packet copy on its link to B6 using the backup BIFT for link to B2 after detecting the failure of its link to B2.

Authors' Addresses

Huaimo Chen (editor)
Futurewei
Boston, MA
USA

Email: Huaimo.chen@futurewei.com

Mike McBride
Futurewei

Email: michael.mcbride@futurewei.com

Steffen Lindner
University of Tuebingen

Email: steffen.lindner@uni-tuebingen.de

Michael Menth
University of Tuebingen

Email: menth@uni-tuebingen.de

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

Yisong Liu
China Mobile

Email: liuyisong@chinamobile.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

