Network Working Group
Internet-Draft
Intended status: Informational
Expires: 3 January 2026

H. Chen
M. McBride
Futurewei
S. Lindner
M. Menth
University of Tuebingen
A. Wang
China Telecom
G. Mishra
Verizon Inc.

2 July 2025

BIER Fast Reroute (BIER-FRR)
draft-ietf-bier-frr-10

Abstract

This document describes BIER Fast Reroute (BIER-FRR) as a mechanism for Fast Reroute (FRR) in Bit Index Explicit Replication (BIER) networks. It enhances the resiliency of BIER by quickly rerouting BIER traffic in the event of a link or node failure. This ensures that multicast traffic continues to reach its intended destinations, thereby minimizing packet loss and service disruption. BIER-FRR is designed to integrate seamlessly with existing BIER operations without requiring per-flow state or additional signaling. The document suggests additional structures for BIER to hold information for backup forwarding entries and to enable them in case of detected failures. BIER-FRR can implement different protection levels, e.g., link protection or node protection, and different protection strategies. Tunnel-based BIER-FRR and LFA-based BIER-FRR are introduced as protection strategies and their implementation with the proposed extensions. A comparison highlights the differences between both approaches. This document serves as an introductory primer to support future, more comprehensive, BIER Fast Reroute analysis and solution development.

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.

This Internet-Draft will expire on 3 January 2026.

Copyright Notice

Table of Contents

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 ([RFC8279] Section 4.1) 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 a controller.

Typically, BIER packets are forwarded without an outer IP header. Consequently, if a link or node failure occurs, the corresponding BFR Neighbor (BFR-NBR) becomes unreachable. Fast Reroute (FRR) mechanisms in the routing underlay, such as IP-FRR [RFC5286], apply

exclusively to IP packets, leading to potential loss of BIER traffic. BIER traffic can only be restored after the routing underlay has reconverged and the BIFT has been recalculated. Tunneling BIER packets can serve as a solution to reach the BFR-NBR in the case of a link failure by leveraging the FRR capabilities of the routing underlay, provided such mechanisms are available. However, tunneling a single BIER packet does not help in the case of node failures because many next-next-hops on the way to destinations need a packet copy when the next-hop becomes unreachable. Given that BIER may carry multicast traffic with real-time requirements, there is a particular need to protect BIER traffic against prolonged outages following failures.

This document introduces a nomenclature for Fast Reroute in BIER (BIER-FRR). Upon detecting that a BFR-NBR is unreachable, BIER-FRR enables a BFR to quickly reroute affected BIER packets using backup forwarding entries. To avoid the generation of redundant packets, backup forwarding entries should be processed before normal forwarding entries.

The protection level offered by BIER-FRR can be either link protection or node protection. Link protection is limited to safeguarding against link failures and is simpler to implement but may not be effective if a BFR itself fails. Node protection, while more complex, also guards against the failure of BFRs. The choice of backup strategy determines the selection of backup forwarding entries. Examples of backup strategies include tunnel-based BIER-FRR and LFA-based BIER-FRR:

* Tunnel-based BIER-FRR: This approach leverages the mechanisms of the routing underlay for FRR purposes. The routing underlay typically restores connectivity faster than BIER, as the reconvergence of the routing underlay is a prerequisite for the recalculation of the BIFT. When the routing underlay utilizes FRR mechanisms, its forwarding capabilities are restored well before reconvergence is completed. To benefit from the rapid restoration of the routing underlay, BIER traffic affected by a failure is tunneled over the routing underlay.

* LFA-based BIER-FRR: This approach reroutes BIER traffic to alternative neighbors in the event of a failure. It applies the principles of IP-FRR, requiring that LFAs are also BFRs. Normal (ie, non-tunneled or direct) BIER-LFAs can be reached without tunneling, remote BIER-LFAs use a tunnel and topology-independent BIER-LFAs use explicit paths to reach the backup BFR-NBR. Unlike tunnel-based FRR, LFA-based BIER-FRR does not depend on fast reroute mechanisms in the routing underlay.

BIER-FRR describes extensions to BIER so that both strategies can be implemented, but it does not mandate a specific one. The BIER-FRR mechanisms described in this document adhere to a primary/backup path model, also known as 1:1 protection where traffic is forwarded either over a primary path or over a backup path. It is in contrast to a 1+1 protection model, where traffic is duplicated across both primary and backup paths. That principle has been implemented by Multicast-only Fast Reroute (MoFRR) [RFC7431] and was explored for BIER in [BrAl17].

2. Terminology

This document uses the following definitions:

BIER: Bit Index Explicit Replication

BIER-FRR: Bit Index Explicit Replication Fast ReRoute

BFR: Bit-Forwarding Router

BFR-NBR: Bit-Forwarding Neighbor

BFIR: Bit-Forwarding Ingress Router

BFER: Bit-Forwarding Egress Router

BIFT: Bit Index Forwarding Table

F-BM: Forwarding Bit Mask

PLR: Point of Local Repair

LFA: Loop Free Alternate

BF-BM: Backup F-BM

BBFR-NBR: Backup BFR-NBR

BFA: Backup Forwarding Action

BEA: Backup Entry Active

3. Definition of BIER-FRR

BIER-FRR proposes a backup BIFT that comprises backup forwarding entries. They are executed before the primary forwarding entries in the normal BIFT which is also denoted primary BIFT in this context. In the following, forwarding actions are defined and the structure of the backup BIFT is introduced. Then activation and deactivation of backup forwarding entries as well as the derivation of the backup F-BM (BF-BM) are explained.

3.1. Definition of Forwarding Actions

A BFR-NBR is considered directly connected if it is a link layer next-hop. Conversely, if the BFR-NBR cannot be reached directly through the link layer, it is regarded as indirectly connected.

The following forwarding actions are defined:

* Plain: The BIER packet is sent directly to a BFR-NBR via a direct link without encapsulation in a tunnel header. This indicates that the packet is not forwarded through the underlying network.

* Tunnel: The BIER packet is encapsulated with a tunnel header and forwarded to a BFR-NBR over the routing underlay.

* Explicit: The packet is forwarded along an explicit path to a BFR-NBR. The specific path information must be provided. If segment routing is employed for this purpose, the segment IDs (SIDs) must be specified. Two forwarding actions of type Explicit are considered equivalent only if they utilize the same explicit path.

In the BIFT as outlined in [RFC8279], the forwarding actions are implicitly determined by the connectivity status 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.

3.2. Backup BIFT

The structure of the backup BIFT is given in Figure 1.

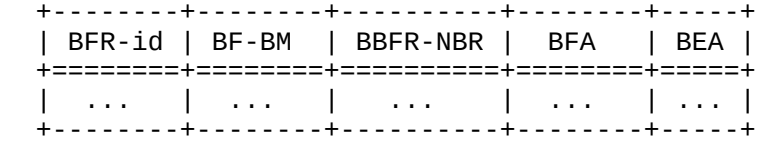

Figure 1: Structure of the backup BIFT.

The columns refer to:

* BFR-id: the bit position of a BFER for which this row in the backup BIFT applies.

* BF-BM: the Backup F-BM used for forwarding, used like the primary F-BM.

* BBFR-NBR: the Backup BFR-NBR used for forwarding, used like the primary BFR-NBR.

* BFA: the Backup Forwarding Action takes values as introduced in Section 3.1 and indicates how the packet is forwarded to the BBFR-NBR.

* BEA: the Backup Entry Active flag indicates if the backup forwarding entry of this row is active.

The structure and semantics of the first three fields are identical to the entries of the primary BIFT, as defined in Figure 3 of [RFC8279], and they are used in a very similar way. The BEA indicates if the backup forwarding entry is executed. In that case, the BFA indicates the forwarding action for the packet.

3.3. 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 Bidirectional Forwarding Detection (BFD) [RFC5880]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.

3.4. Usage of the Backup BIFT

An incoming packet is first matched against the backup BIFT. A row in the backup BIFT matches a packet if the BEA flag in the backup BIFT is set and if the BFR-id is set in the packet's bitstring. Then, the BF-BM of the matching backup forwarding entry is applied to the packet's bitstring. That means, the packet is copied and in its bitstring the bits other than those set in BF-BM are cleared before the packet is forwarded to the BBFR-NBR with the indicated BFA. Finally, the bits of the BF-BM are cleared in the bitstring of the remaining packet. In the absence of a match of the remaining packet, the normal forwarding procedure continues, i.e., forwarding based on the primary BIFT as described in [RFC8279].

Note: If a BFR-id matches in the primary or backup BIFT, and the transmission is not successful, the F-BM or BF-BM is still applied to the bitstring of the remaining packet.

3.5. Computation of the Backup F-BM

The primary F-BM of a specific BFER identifies all BFERs that share the same primary BFR-NBR. The backup F-BM for a specific BFER is computed to indicate:

* All BFERs that share both the primary and backup BFR-NBRs of the specific BFER, and

* All BFERs for which the backup BFR-NBR of the specific BFER serves as the primary BFR-NBR.

3.6. Alternative Representations of Backup Forwarding Entries

Alternative representations of backup forwarding entries are possible as long as the same behavior is ensured. Two other variants are introduced in the following sections.

3.7. Single Extended BIFT

The information of the primary BIFT and the backup BIFT may be combined in a single extended BIFT. Its structure is illustrated in Figure 2

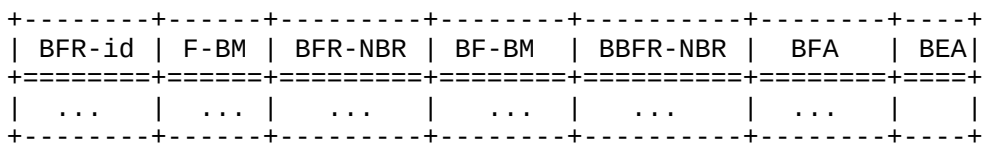

Figure 2: Structure of a single extended BIFT including backup forwarding entries.

To ensure the same behavior, the BEA flag must be set like in the backup BIFT. Furthermore, two matching passes through the extended BIFT are needed. A first one matches the bitstring combined with BEA=1. If no further match is possible, then another pass with the remaining bitstring combined with BEA=0 is performed.

3.8. Primary BIFT and Failure-Specific Backup BIFTs

To avoid two distinct passes through a BIFT, the information of the primary BIFT and backup BIFT may be combined into a primary BIFT and multiple failure-specific BIFTs. Each failure-specific BIFT corresponds to a specific failure scenario. Failure-specific backup BIFTs are structured like normal backup BIFTs, but do not have a BEA flag as they are enabled or disabled as a whole.

In the absence of a failure, packets are processed using the primary BIFT. In case of a failure, packets are processed using a failure-specific BIFT that matches the occurred failure. That means, there should be failure-specific BIFTs for at least any adjacent link to protect against all single-link failures. To support multiple failures, even more failure-specific BIFTs are needed. If failure-specific BIFTs are provided for only single-link failures, the BIFT should be taken that covers the most relevant single failure.

4. Illustration and the Need for Prioritized Backup Forwarding Entries

In this section, BIER-FRR is illustrated using a small example. It is pointed out that unnecessary redundant packets may occur if primary forwarding entries are erroneously applied before backup forwarding entries. Therefore, it is important that the backup BIFT is applied before the primary BIFT.

4.1. Example

Figure 3 presents an example of a BIER network. In this example, BFRs are identified by the prefix "B" followed by their BFR-ids. For simplicity, each BFR also serves as a BFER, and its bit position in the bitstring corresponds to its BFR-id. The number assigned to each link represents its cost, which the routing underlay uses to compute the shortest paths.

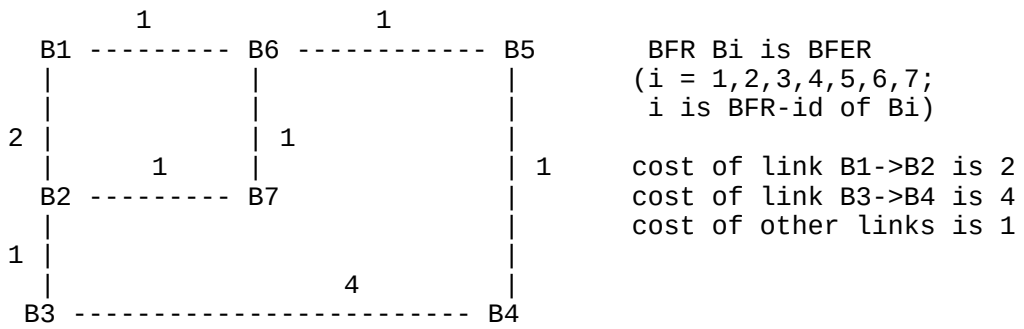

Figure 3: BIER network example.

In the absence of a failure, traffic for BFR-id 2 and 3 is forwarded via BFR-NBR B2 and traffic to BFR-id 4, 5, 6, and 7 is forwarded to BFR-NBR B6. If a packet with bitstring 0001100 (destinations B3 and B4) is forwarded, the row for BFR-id B3 matches first. A packet with bitstring 0000100 is sent to B2 and the bitstring of the remaining packet is also processed with F-BM 0001100, i.e., the remaining bitstring is 0001000. Then the remaining bitstring is matched again so that BFR-id B4 yields a match. A packet copy with bitstring 0001000 is sent to B6 and the application of the F-BM 1111000 to the bitstring of the remaining packet results in 0000000, which terminates the forwarding process. This BIER forwarding process avoids redundant packet copies.

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.

A backup BIFT for B1 in the example of Figure 2 is given in Figure 4. It implements LFA-based FRR as a protection strategy and link protection.

If B1 cannot reach B2 or B6, BEA will be set to 1 in the rows for the backup BIFT for which B2 or B6 is the BFR-NBR in the primary BIFT. Thus, if B1 cannot reach B2, traffic for BFR-id 2 and 3 will be forwarded over B6 and 1111110 is applied as BF-BM. This mask also includes all the BFR-ids that have B6 as their primary BFR-NBR. Likewise, if B1 cannot reach B6, traffic for BFR-id 4, 5, 6, and 7 will be forwarded over B2 and again 1111110 is applied as BF-BM for the same reason.

4.2. B1's backup BIFT for LFA-based FRR with link protection

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 5: B1's backup BIFT for LFA-based FRR with link protection.

We now consider that the link B1->B2 failed and that B1 needs to forward a packet with bitstring 0001100. Therefore, the BEA is set for BFR-id 2 and 3 in the backup BIFT. If B1 needs to forward a packet with bitstring 0001100 (destinations B3 and B4), the row for BFR-id B3 in the backup BIFT matches first. Therefore, a packet with bitstring 0001100 is sent to B6 and the bitstring of the remaining packet is also processed with BF-BM 1111110 so that the remaining bitstring is 0000000, which terminates the forwarding process. That is, only a single packet copy is sent to B6.

5. Prioritization of Backup Forwarding Entries over Primary Forwarding Entries

BIER-FRR defines that the backup BIFT is applied before the primary BIFT. The reason for that is twofold. First, applying the primary BIFT first may erase the forwarding information for BFERs whos primary BFR-NBR is unreachable. Second, if that can be fixed, redundant packets can occur if the primary BIFT is applied before the backup BIFT. These issues are demonstrated in the above example when the link B1->B2 has failed and B1 applies the primary BIFT before the backup BIFT when forwarding a packet with bitstring 0011000 (B3 and B4 as destinations).

We first assume that B1 just ignores the failed interface when forwarding the packet with the primary BIFT but processes the bitstring of remaining packet like if the transmission was successful. That means, when BFR-id 3 matches first in the primary BIFT, no packet is sent to B2, but the bits in the bitstring are still cleared, leading to a remaining bitstring of 0001000. Another pass through the primary BIFT forwards a packet copy to B6 and clears the remaining bitstring to 0000000, which terminates the forwarding process. However, no packet will reach B3 as the bitstring information was lost during the unsuccessful transmission.

We now assume a feature that saves the bitstring information when the transmission to a specific BFR-id was not successful. This can be done by AND-ing the remaining bitstring and the F-BM and OR-ing the result with a remaining backup bitstring which was initially zero. Only then the bits of the F-BM are cleared from the remaining bitstring. When B1 is to forward a packet with bitstring 0001100, the first match in the primary BIFT is for BFR-id 3. As the transmission is not successful, 00000100 is saved in the remaining backup bitstring and the remaining bitstring is 0001000. Therefore, a second match in the primary BIFT is for BFR-id 4, which sends a packet copy with bitstring 0001000 to B6. Then, the remaining backup bitstring is processed with the backup BIFT. As there is a match for BFR-id 3, another packet is sent to B6, now with bitstring 0000100. This can be considered redundant.

Below the line, it is important to first process backup forwarding entries before backup forwarding entries. This avoids additions to the forwarding process with the primary BIFT and avoids redundant packets.

6. Protection Levels

Both link protection and node protection may be supported. Link protection is designed to safeguard against the failure of an adjacent link, whereas node protection addresses the failure of a neighboring node and the associated path leading to that node. The relevance of link or node protection depends on the specific service being supported. Additionally, both protection levels can be combined with any of the backup strategies outlined in Section 7.

6.1. Link Protection

In link protection, the backup path is designed to circumvent the failed link, i.e., the failed primary path from the PLR to the primary BFR-NBR, while still potentially including the primary BFR-NBR itself. Consequently, the backup path with link protection cannot protect against the failure of the primary BFR-NBR..

6.2. Node Protection

In node protection, the backup path is designed to avoid both the failed node and the link to that node, i.e., the failed primary path from the PLR to the primary BFR-NBR, including the primary BFR-NBR. Consequently, the backup path with link protection also protects against the failure of the primary BFR-NBR. If a BFER and its primary BFR-NBR are the same, only link protection is feasible for that BFER.

6.3. Example

In the network depicted in Figure 3, the primary path from BFR B1 to BFER B5 is B1->B6->B5. Protecting BFER B5 from a BFR-NBR B6 node failure can only be provided through the backup path B1->B2->B3->B4->B5. Link protection for BFER B5 is achieved via the backup path B1->B2->B7->B6, and additionally through the backup path B1->B2->B3->B4->B5->B6. The specific backup entries are determined by the selected protection level and backup strategy. Example BIFTs illustrating link and node protection are provided in Section 7.

7. Backup Strategies

Backup strategies determine the selection of backup forwarding entries, influencing both the choice of the backup BFR-NBR and the backup forwarding action, and consequently, the backup path. The following sections present tunnel-based BIER-FRR and LFA-based BIER-FRR as potential strategies. Both can be implemented with BIER-FRR presented in Section 3.

7.1. Tunnel-Based BIER-FRR

The routing underlay may possess the capability to forward packets to their destinations even in the presence of a failure, potentially due to FRR mechanisms within the routing underlay. In such scenarios, while the primary BFR-NBR may no longer be reachable via the primary action (Direct), it could still be accessible through a backup action (Tunnel).

Tunnel-based BIER-FRR encapsulates BIER packets impacted by a failure within the routing underlay, thereby leveraging the routing underlay's fast restoration capabilities. As soon as connectivity in the routing underlay is reestablished, the affected BIER packets can be forwarded to their intended destinations. The appropriate backup forwarding entries in a BIFT for BIER-FRR are determined by the desired protection level.

7.1.1. Tunnel-Based BIER-FRR with Link Protection

In the context of link protection, the backup BFR-NBRs are identical to the primary BFR-NBRs. If a primary BFR-NBR is directly connected to the BFR acting as the Point of Local Repair (PLR), the corresponding backup forwarding action is Tunnel. Consequently, BIER packets affected by a failure are tunneled through the routing underlay to their BFR-NBR, rather than being directly sent as pure BIER packets. If the 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 identical to the primary F-BMs, which is consistent with the computation of backup F-BMs described in Section 3.5.

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 6: B1’s backup BIFT for tunnel-based BIER-FRR with link protection.

Figure 6 illustrates B1’s backup BIFT for tunnel-based BIER-FRR with link protection in the BIER network example depicted in Figure 3. The backup BFR-NBRs and backup F-BMs in this backup BIFT correspond to the primary BFR-NBRs and primary F-BMs in the primary BIFT. However, the backup actions in this backup BIFT are Tunnel, while the primary forwarding actions in the primary BIFT are Direct (which are not explicitly shown but are implicit).

When B1, acting as the PLR, detects a failure of its link to B6, a BIER packet with the bitstring 0100000 destined for B6 is tunneled by B1 through the routing underlay towards B6. The specific path of the backup tunnel depends on the routing underlay and could be B1->B2->B7->B6 or B1->B2->B3->B4->B5->B6.

If a BIER packet is destined for {B2, B5, B7}, an encapsulated packet copy is first forwarded via link B1->B2 to backup BFR-NBR B6 using the backup forwarding action Tunnel to deliver packet copies to BFERs B5 and B7. Subsequently, a non-encapsulated packet copy is forwarded via link B1->B2 to BFR-NBR B2 using the primary forwarding action Direct to deliver a packet copy to BFER B2. Therefore, with tunnel-based BIER-FRR, and link protection, a single redundant packet copy may occur in the event of a failure because an encapsulated and a non-encapsulated packet copy are forwarded over the same link. This redundancy occurs even though BIER packets affected by failures are forwarded before those unaffected by failures. The redundant packet is rather caused by the fact that two packet copies are sent over the link with different next-hops on the BIER layer, namely B2 and B6.

A BIER packet with the bitstring 1000000 destined for B7 is forwarded along the backup path B1->B2->B7->B6->B7, as it is first delivered to the backup BFR-NBR B6. Consequently, the backup path may be unnecessarily long. This phenomenon is similar to the facility backup method described in [RFC4090] which employs paths analogous to those in tunnel-based BIER-FRR..

7.1.2. Tunnel-Based BIER-FRR with Node Protection

To determine the backup forwarding entries for node protection, two cases need to be distinguished. If the BFER is the same as its primary BFR-NBR, node protection is not feasible for that BFER. Therefore, link protection is applied, meaning the backup BFR-NBR is set to the primary BFR-NBR. If 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, making the backup BFR-NBR a next-next-hop BFR. In both cases, the backup forwarding action is Tunnel. In

the first case, the backup F-BM is set to all zeros with the bit for the BFER to be protected enabled. In the second case, the backup F-BM is computed as described in Section 3.5.

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 7: B1’s backup BIFT for tunnel-based BIER-FRR with node protection.

Figure 7 illustrates B1’s backup BIFT for tunnel-based BIER-FRR with node protection in the BIER network example provided in Figure 3. BFERs B2 and B6 are direct neighbors of B1. To protect them, only link protection is applied, as B1’s primary BFR-NBRs for these nodes are the nodes themselves. As described above, only the bit for B2 is set in the backup F-BM of B2, and similarly 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, acting as the PLR, detects the failure of its BFR-NBR B6, a BIER packet with bitstring 1010010 destined 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 over link B1->B2. Therefore, node protection may result in more redundant packet copies than link protection.

Caveat: If the routing underlay does not support node protection, tunnel-based BIER-FRR will similarly be unable to provide node protection. This limitation is illustrated in the following example. In the network depicted in Figure 3, the underlay offers only link protection. If BFR-NBR B6 fails and B1 must forward a packet to B5, according to the backup BIFT in Figure 7 the packet is tunneled towards B5. The underlay may route the packet along the path

B1->B2->B7->B6->B5 due to FRR with link protection. However, since B6 is also unreachable from B7, the packet is returned to B2, resulting in a loop between B2 and B7.

7.2. LFA-based BIER-FRR

LFA-based BIER-FRR leverages alternate BFRs to deliver BIER packets to BFERs if their primary BFR-NBR is unreachable. This approach does not rely on any fast restoration or protection mechanisms in the underlying routing infrastructure. First, the prerequisites for LFA-based BIER-FRR are clarified, followed by the definition of BIER-LFAs. Subsequently, link and node protection for LFA-based BIER-FRR are discussed using a single backup BIFT.

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

An LFA for a specific destination is an alternate node to which a packet is sent if the primary next-hop for that destination is unreachable. This alternate node should be capable of forwarding 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], and such LFAs are referred to as IP-LFAs. BIER-LFAs are 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 required. This differs from tunnel-based BIER-FRR, where knowledge of the primary BIFTs of a PLR and its BFR-NBRs is sufficient.

LFA-based BIER-FRR may reuse IP-LFAs as BIER-LFAs under the following conditions: if an IP-LFA node for the destination of a specific BFER is a BFR, it may be reused as the backup BFR-NBR for that BFER, along with the backup action applied for that IP-LFA at the IP layer. A normal IP-LFA corresponds to the backup forwarding action Direct, a remote IP-LFA to Tunnel, and a TI-IP-LFA to Explicit.

7.2.2. Definition of BIER-LFAs

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

* A BFR is considered 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 the BFR and the path from the BFR to the BFER are disjoint from the primary path from the PLR to the primary BFR-NBR. These paths:

   - may include the primary BFR-NBR for link protection.

   - must not include the primary BFR-NBR for node protection.

* A BFR is considered a remote BIER-LFA for a specific BFER if it is not directly connected to the PLR, can be reached via a tunnel from the PLR, and satisfies the aforementioned conditions 1 and 2.

* A BFR is considered a TI-BIER-LFA for a specific BFER if it is not directly connected to the PLR, cannot be reached via a tunnel from the PLR, but is reachable from the PLR via an explicit path (e.g., with the assistance of a Segment Routing (SR) header), and satisfies the aforementioned conditions 1 and 2.

For the protection of some BFERs, one or more normal BIER-LFAs may be available at a specific PLR. For the protection of other BFERs, only remote or TI-BIER-LFAs may be available. There may also be BFERs which can be protected only through TI-BIER-LFAs.

The backup forwarding actions for rerouting BIER packets depending on the type of BIER-LFA are:

* For normal BIER-LFA: Direct

* For remote BIER-LFA: Tunnel

* For TI-BIER-LFA: Explicit

7.2.3. Protection Coverage of BIER-LFA Types

Protection coverage refers to the set of BFERs that can be protected with a desired level of protection by a particular type of BIER-LFA. The BIER-LFA types exhibit the following characteristics:

* Normal BIER-LFAs

   - The protection coverage is the least as some or many BFERs may not be protected at the desired protection level or at all.

   - Redundant packet copies are avoided.

   - There is no encapsulation overhead.

* Remote BIER-LFAs

- They enhance the protection coverage of normal BIER-LFAs.

- Redundant packet copies may occur on a link, similar to tunnel-based BIER-FRR.

- The encapsulation overhead is similar to that of tunnel-based BIER-FRR.

* TI-BIER-LFAs

- They complement the protection coverage of normal and remote BIER-LFAs to achieve 100% coverage.

- Redundant packets may occur on a link, similar to tunnel-based BIER-FRR.

- The encapsulation overhead is similar or equivalent to that of tunnel-based BIER-FRR, depending on the FRR mechanism employed in the routing underlay.

- There is increased complexity as segment routing, or some other forms of explicit tunnels, needs to be supported by the routing underlay.

7.2.4. Sets of Supported BIER-LFAs

Normal BIER-LFAs are the simplest option, as they do not require tunneling or explicit paths. Remote BIER-LFAs offer greater capabilities but introduce additional header overhead and require more functionality from the PLR. TI-BIER-LFAs are the most complex BIER-LFAs, necessitating the use of explicit paths. When implementing LFA-based BIER-FRR, it is essential to specify the set of supported BIER-LFAs. The available options are as follows:

* Option 1: Only normal BIER-LFAs are supported.

* Option 2: Both normal and remote BIER-LFAs are supported.

* Option 3: All types of BIER-LFAs are supported.

Options 1 and 2 may not be able to protect the reachability of all BFERs against all single link failures and all single node failures.

7.2.5. Link Protection

In the following, LFA-based BIER-FRR with link protection is illustrated. Thereby, normal BIER-LFAs are prioritized over remote LFAs, and remote BIER-LFAs are preferred over TI-BIER-LFAs. Depending on the specific PLR, simple BIER-LFAs are sufficient, remote BIER-LFAs are needed, or even TI-BIER-LFAs to protect the reachability of all BFERs against single link failures..

If the link between B1 and B6 fails, B1 cannot reach the BFERs B4, B5, B6, and B7 via their primary BFR-NBR. Consequently, B1 forwards their traffic via the backup BFR-NBR B2, along with the traffic for B2 and B3, as B2 is their primary BFR-NBR. In this scenario, the backup F-BM is set to 1111110. Similarly, if the link between B1 and B2 fails, B1 routes all traffic to B6, with the backup F-BM also set to 1111110.

B1 requires only normal BIER-LFAs to protect all BFERs. However, this situation can vary significantly for other BFRs. Figure 8 and Figure 9 present 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. BFR B5 requires 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, it may not be possible to protect all BFERs using BIER-FRR.

Consider a scenario where B7 holds a BIER packet with destinations {B1, B4, B5, B6}. If the link between B7 and B6 fails, the packet copy for B1 is sent to B2 using the backup forwarding action Direct, the packet copy for B4 is tunneled via B2 to B3, and the packet copies for B5 and B6 are sent via explicit paths to B4 and B1, respectively. Since these packet copies have different next-hops on the BIER layer, all of them must be transmitted, resulting in three redundant copies.

BFR-id	BF-BM	BBFR-NBR	BFA	BEA	Comment: protects failure of
1	0000111	B2	Direct		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 8: 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	Direct		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 9: B5’s backup BIFT with link protection.

7.2.6. Node Protection

To determine the backup forwarding entries for node protection, it is necessary to conduct a case-by-case analysis of the BFER to be protected. If the BFER is the same as its primary BFR-NBR, node protection is not feasible for that BFER, and link protection must be applied instead. In all other cases, the BFER should be protected by a node-protecting BIER-LFA. In this context, normal BIER-LFAs are prioritized over remote BIER-LFAs, and remote BIER-LFAs are preferred over TI-BIER-LFAs. Depending on the set of supported BIER-LFAs, it

may not be possible to protect certain BFERs.

Figure 10 illustrates B1’s backup BIFT for LFA-based BIER-FRR with node protection, using the network example provided in Figure 3.

BFR-id	BF-BM	BBFR-NBR	BFA	BEA	Comment: protects failure of
2	1111010	B6	Direct		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	Direct		BFR-NBR B6
7	1100100	B2	Direct		BFR-NBR B6

Figure 10: B1’s backup BIFT with node protection.

As B6 serves as the primary BFR-NBR for BFER B6, only link protection can be applied. Consequently, B2 is utilized as a normal, link-protecting BIER-LFA to safeguard B6. Similarly, as B2 is the primary BFR-NBR for BFER B2, B2 is protected with B6 as its normal, link-protecting BIER-LFA. BFER B7 is protected against the failure of node B6 by using B2 as its normal, node-protecting BIER-LFA, as B2 has a shortest path to B7 that does not traverse B6. The backup F-BMs for BFERs B6 and B7 are set to {B2, B6, B7}, as traffic for these BFERs is routed via link B1->B2 with the backup forwarding action Direct when B6 is unreachable.

BFER B4 cannot be reached via a normal LFA when BFR B6 fails. However, B3 serves as a remote, node-protecting BIER-LFA for BFER B4, as B3 has a shortest path to B4, is reachable from B1 via a shortest path, and the resulting backup path from B1 to B4 does not traverse B6. Similarly, B4 serves as 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 acts as a node-protecting TI-BIER-LFA for BFER B5 as B4 is reachable through the explicit path B1->B2->B3->B4 and has a shortest path to B5 that does not traverse B6.

Consider a scenario where B1 holds a BIER packet with destinations {B4, B5, B6}. If the link between B1 and B2 fails, the packet copy for B1 is sent to B2 using the backup forwarding action Direct, a packet copy for B4 is tunneled via B2, and a packet copy for B5 is sent via an explicit path to B4. Since these packet copies have different next-hops on the BIER layer, all of them must be transmitted, resulting in two redundant copies.

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

Redundant packets can occur with LFA-based BIER-FRR when BIER packets are transmitted over a specific link in different forms, including:

* Directly sent BIER packets (either primary transmission or reroute to a normal BIER-LFA).

* BIER packets encapsulated for transmission to a specific BFR-NBR (either tunneled primary transmission or reroute to a remote BIER-LFA).

* BIER packets routed with an encoded explicit path (reroute to a TI-LFA).

When different remote BIER-LFAs are utilized, multiple redundant packets may be generated. A similar situation can arise with TI-BIER-LFAs. However, some redundant packets can be mitigated if remote BIER-LFAs or TI-BIER-LFAs are selected such that they can protect multiple BFERs, thereby reducing the need for additional remote BIER-LFAs or TI-BIER-LFAs. This approach, while potentially leading to longer backup paths, introduces a new optimization objective for the selection of remote or TI-BIER-LFAs, which does not exist in IP-FRR. The relevance of this optimization may vary depending on the specific use case.

To illustrate this optimization potential, consider LFA-based BIER-FRR with link protection for B7, as described in its backup BIFT in Figure 8. As noted in Section 7.2.5, B7 needs to transmit four copies to forward a packet to {B1, B4, B5, B6}. If the more complex TI-BIER-LFA B4 is chosen to protect BFER B4 instead of the remote BIER-LFA B3, only two redundant copies need to be transmitted.

8. Comparison

This section first addresses the differences between IP-LFAs for IP-FRR and BIER-LFAs for BIER-FRR. It then examines the advantages and disadvantages of tunnel-based and LFA-based BIER-FRR.

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

LFAs were initially proposed for IP networks. They are straightforward in that they do not require any tunneling overhead. However, certain destinations cannot be protected against specific link failures, and even more destinations may be unprotectable against certain node failures. To improve coverage, remote LFAs (R-LFAs) were introduced, which tunnel affected traffic to another node from which the traffic can reach the destination through normal forwarding. Despite this, there may still be destinations that remain unprotected against link or node failures. To address this, topology-independent LFAs (TI-LFAs) were developed, wherein affected traffic is tunneled via an explicit path (preferably using segment routing headers) to another node from which the traffic can reach its destination through standard IP forwarding. With TI-LFAs, all destinations can be protected against any failures as long as connectivity exists.

LFA-based BIER-FRR adopts the principles of LFAs but differs from IP-FRR in that the LFA target node, i.e., the next-hop on the BIER layer to which traffic is diverted, 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 a BIER-LFA. Consequently, if only a subset of nodes in the underlay are BFRs, the BIER-LFAs will differ substantially from IP-LFAs. Furthermore, this makes it more challenging to find normal BIER-LFAs which do not require tunneling. As a result, LFA-based BIER-FRR is likely to require more remote BIER-LFAs and TI-BIER-LFAs than IP-FRR under such conditions.

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

8.2.1. Advantages

* The computation of backup forwarding entries for tunnel-based BIER-FRR is straightforward, requiring only the primary BIFTs of a PLR and its BFR-NBRs. No routing information from the routing underlay is needed.

* The forwarding action "Explicit" is not required for tunnel-based BIER-FRR. However, depending on the underlay, explicit forwarding may still be utilized to achieve FRR in the underlay.

8.2.2. Disadvantages

* Tunnel-based BIER-FRR relies on the presence of a FRR mechanism in the underlay.

* Its protection level is constrained by the protection level provided by the underlay. For instance, if the underlay supports only link protection, tunnel-based BIER-FRR cannot offer node protection.

* Redundant packet copies may occur in tunnel-based BIER-FRR.

* Backup paths may be longer than with LFA-based BIER-FRR.

* A tunneling header is required for any rerouting, resulting in additional header overhead.

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

8.3.1. Advantages

* LFA-based BIER-FRR does not depend on any fast protection mechanisms in the underlay.

* Therefore, it can provide superior protection at the BIER layer compared to the IP layer, particularly if LFA-based BIER-FRR utilizes BIER-LFAs with a higher protection level than those used in LFA-based IP-FRR. For example, the underlay may only offer FRR with link protection, while BIER-FRR can provide node protection for BIER traffic.

* LFA-based BIER-FRR avoids header overhead for normal BIER-LFAs.

8.3.2. Disadvantages

* The computation of backup forwarding entries requires routing information from the underlay.

* The computation of backup forwarding entries is more complex when some nodes in the underlay are not BFRs because then BIER-LFAs differ from IP-LFAs..

* The "Tunnel" forwarding action is required to protect certain BFERs, which adds header overhead.

* The "Explicit" forwarding action is necessary to achieve full protection coverage in some topologies; without it, only partial protection coverage is possible. This requires support for explicit paths, such as Segment Routing.

* More remote BIER-LFAs and TI-BIER-LFAs are needed compared to IP-FRR if some nodes in the routing underlay are not BFRs.

* Redundant packet copies may occur in LFA-based BIER-FRR, though this is less frequent than with tunnel-based BIER-FRR as simple BIER-LFAs do not require a tunnel.

9. Security Considerations

This specification does not introduce additional security concerns beyond those already discussed in the BIER architecture [RFC8279] along with the IP FRR [RFC5286] and LFA [RFC7490] specifications.

10. IANA Considerations

No requirements for IANA.

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>.

[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>.

[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>.

[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>.

[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>.

11.2. Informative References

[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>.

[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", Work in Progress, Internet-Draft, draft-chen-bier-egress-protect-07, 28 March 2024, <https://datatracker.ietf.org/doc/html/draft-chen-bier-egress-protect-07>.

[I-D.ietf-rtgwg-segment-routing-ti-lfa] Bashandy, A., Litkowski, S., Filsfils, C., Francois, P., Decraene, B., and D. Voyer, "Topology Independent Fast Reroute using Segment Routing", Work in Progress, Internet-Draft, draft-ietf-rtgwg-segment-routing-ti-lfa-21, 12 February 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-rtgwg-segment-routing-ti-lfa-21>.

[BrAl17] Braun, W., Albert, M., Eckert, T., and M. Menth, "Performance Comparison of Resilience Mechanisms for Stateless Multicast Using BIER", May 2017.

Acknowledgments

The authors would like to thank Daniel Merling, Jeffrey Zhang, Tony Przygienda, Shaofu Peng and Toerless Eckert for their comments to this work. A special thank you to Gunter van de Velde for his extensive editing to help bring this document to publication.

Contributors

Yisong Liu
China Mobile
Email: liuyisong@chinamobile.com

Yanhe Fan
Casa Systems
United States of America
Email: yfan@casa-systems.com

Lei Liu
Fujitsu
United States of America
Email: liulei.kddi@gmail.com

Xufeng Liu
Alef Edge
United States of America
Email: xufeng.liu.ietf@gmail.com

Xuesong Geng
China
Email: gengxuesong@huawei.com

Authors' Addresses

Huaimo Chen
Futurewei
Email: hchen.ietf@gmail.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
United States of America
Phone: 301 502-1347
Email: gyan.s.mishra@verizon.com