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

H. Chen
M. McBride
Futurewei
S. Lindner
M. Menth
University of Tuebingen
T. Eckert
Futurewei

2 March 2026

A Framework for Fast Reroute with Bit Index Explicit Replication (BIER-FRR)
draft-ietf-bier-frr-12

Abstract

This document provides a framework for the development of Fast Reroute (FRR) mechanisms for Bit Index Explicit Replication forwarding (BIER-FRR). BIER-FRR can provide protection against link or BFR failure by invoking locally pre-determined repair paths that can react in the same time-scales as (unicast) FRR for MPLS or IP networks - "sub 50msec", and without the creation of additional per-path or per-flow state coordinated across multiple routers/LSR.

BIER-FRR can be implemed locally within a router/LSR with minimal interoperability requirements against other router/LSR. It can therefore easily be introduced incrementally or selectively where needed. BIER-FRR implementing nodes only need to understand the routing topology of the network for calculation of repair paths and know what type of unicast encapsulation can be used to send ("tunnel") BIER packets to remote BFR.

This document proposes and discusses different options for BIER forwardng (BIFT) extensions to support BIER-FRR. These are exemplary and non-normative. This document does not specify any standards or experiments but aims to support such efforts.

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/.

Copyright Notice

Table of Contents

1. Terminology 3
2. Overview 4
  2.1. Benefits 5
    2.1.1. BIER versus IP multicast 5
    2.1.2. FRR for BIER versus IP/MPLS unicast and multicast . . 6
  2.2. Introduction 9
    2.2.1. Tunnel-based BIER-FRR 9
    2.2.2. LFA-based BIER-FRR 11
    2.2.3. Re-use of routing underlay FRR adjacencies 15
    2.2.4. Conceptual BIER forwarding with FRR 16
3. Definition of BIER-FRR 18
  3.1. Definition of Forwarding Actions 18
  3.2. Backup BIFT 18
  3.3. Activating and Deactivating Backup Forwarding Entries . . 19
  3.4. Usage of the Backup BIFT 20
  3.5. Computation of the Backup F-BM 20
  3.6. Alternative Representations of Backup Forwarding Entries 20
  3.7. Single Extended BIFT 20
  3.8. Primary BIFT and Failure-Specific Backup BIFTs 21
4. Illustration and the Need for Prioritized Backup Forwarding Entries 21
  4.1. Example 21
  4.2. B1's backup BIFT for LFA-based FRR with link protection 23

5. Prioritization of Backup Forwarding Entries over Primary Forwarding Entries 23
6. Protection Levels 24
6.1. Link Protection 25
6.2. Node Protection 25
6.3. Example 25
7. Backup Strategies 25
7.1. Tunnel-Based BIER-FRR 25
7.1.1. Tunnel-Based BIER-FRR with Link Protection 26
7.1.2. Tunnel-Based BIER-FRR with Node Protection 27
7.2. LFA-based BIER-FRR 29
7.2.1. Relation of BIER-LFAs to IP-LFAs and Prerequisites . 29
7.2.2. Definition of BIER-LFAs 29
7.2.3. Protection Coverage of BIER-LFA Types 30
7.2.4. Sets of Supported BIER-LFAs 31
7.2.5. Link Protection 31
7.2.6. Node Protection 33
7.2.7. Optimization Potential to Reduce Redundant BIER Packets in Failure Cases 35
8. Comparison 35
8.1. Comparison of LFA-Based Protection for IP-FRR and BIER-FRR 36
8.2. Advantages and Disadvantages of Tunnel-Based BIER-FRR . . 36
8.2.1. Advantages 36
8.2.2. Disadvantages 36
8.3. Advantages and Disadvantages of LFA-Based BIER-FRR . . . 37
8.3.1. Advantages 37
8.3.2. Disadvantages 37
9. Security Considerations 38
10. IANA Considerations 38
Acknowledgments 38
References 38
Normative References 38
Informative References 39
Appendix A. Non-working FRR options 40
A.1. BIER-in-BIER encapsulation 40
Appendix B. Changelog 41
B.1. rev 11 - sent back from IESG to WG 41
B.2. rev 11 - sent back from IESG to WG 41
B.3. Resolved IESG discuss / comments before rev 11. 42
B.4. TBD 42
Contributors 44
Authors' Addresses 45

1. 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

2. Overview

BIER-FRR describes how IP FRR style sub-50 msec protection can be done for BIER. 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 1+1 principle has been described by Multicast-only Fast Reroute (MoFRR) [RFC7431] and was explored for BIER in [BrAl17].

This memo is informational because it is a technology primer explaining the benefits of and mechanisms for Fast ReRoute (FRR) with Bit Inexed Explicit Replication (BIER) stateless multicast forwarding. This document is not a standards track document because

- Most if not all mechanisms possible can be implemented solely on single routers supporting BIER without the same or "interoperable" new mechanism to be supported by other routers supporting BIER.

- At this point in time, it is unclear which of the advanced mechanisms presented are most feasible for implementation adoption in different type of routers supporting BIER, because the feasibility of implemeting them depend on the specifiec abilities of the routers forwarding plane.

2.1. Benefits

BIER is a novel method that allows to simplify and scale the deployments of multicast services significantly over widely deployed technologies like [PIM-SM] for IP networks including SRv6 networks or [mLDP] for MPLS or SR-MPLS networks. The key novelty of BIER is that it is stateless whereas the prior mechanisms are stateful.

BIER-FRR allows to achieve fundamentally the same so-called "sub 50msec" recovery from link or node failure that [IP-FRR] achieves the same networks for unicast traffic. Stateful multicast mechanisms including [PIM-SM] and [mLDP] can not support FRR directly. Instead, they must be combined with prior mechanisms to achieve link-protection (such as [RSVP-TE] with [RFC4090] or explicit paths with SR). In summary, link-protection with these pre-existing mechanisms is more complex and less efficient, and node-protection is mostly considered infeasible because of the involved complexity and excess traffic it causes.

BIER-FRR likely allows for the most efficient and simple multicast FRR because it fundamentally operates like unicast, except that a BIER packet does not indicate only one destination but a list of multiple packets, allowing for each router to perform unicast like forwarding plus whenever necessary replication of the packet to any outgoing interface through which one or more of these destinations need to be reached.

The following sections detail these summaries.

2.1.1. BIER versus IP multicast

BIER [RFC8279] is a novel method for so-called "stateless" forwarding and replication of traffic, especially IP multicast, across networks such as Service Provider network. BIER is intended to replace prior, so-called "stateful" methods of IP multicast, such as the predominantly deployed [PIM-SM] in these networks.

In IP (unicast), variations in user traffic across such a Service Provider network do only impact capacity utilization of the network (links and nodes), but not the control-plane activity and forwarding plane state scalability: IGP and BGP routing protocol operations and scalability and derived forwarding plane scalability and change performance. Only changes in the network topology through failures, recovery and network expansion cause the need for high-performance control-plane activity and/or increased scale requirements. This makes IP unicast Service Provider network designs arguably very easily scalable to arbitrary levels of throughput, a core requirement for successful Service Provider network deployments.

In stateful IP multicast, this is not the case. Every individual new IP multicast application, with new senders and/or receivers can create new control-plane state (from e.g.: PIM-SM) hop-by-hop across such a Service Provider network. Every failure or recovery of links or nodes in the network require control-plane and forwarding-plane re-convergence that needs to be able to scale not with the total size of the network (unicast routing table), but with the total number of IP multicast applications. Denial of Service attacks in IP multicast are equally not limited to attempts of overloading network bandwidth, but can more easily attempt to overload this control-plane and forwarding-plane state through the creation of new IP multicast application state. Because this does not even need for the attackers to have significant bandwidth available, this attack is a lot easier to achieve.

BIER resolves these issues by eliminating the per-IP-multicast application state across Service Provider networks by utilizing fundamentally very much the same packet forwarding paradigm as IP (unicast). The packet simply contains a list of destinations in it’s packet header (encoded very efficiently as a bitstring), and on every hop the packet is replicated and forwarded to all the next-hop routers towards one or more of those destinations. No IP-multicast group state is required for this forwarding.

2.1.2. FRR for BIER versus IP/MPLS unicast and multicast

In IP-FRR, which is also applicable to FRR for SR-MPLS (and hence MPLS networks), unicast packets are forwarded by every router/LSR (Label Switched Router) based

In IP (IPv4 and IPv6), packets are forwarded hop-by-hop in routers based on the destination address of the packet. for each destination (prefix) the forwarding entry indicates a next-hop which may be an outgoing interface or an outgoing intercace with a particular next-hop-neighbor if the outgoing interface is a multi-access subnet (like ethernet with multicast routers connected to it).

IP-FRR amends the per-packet forwarding with a check whether the outgoing interface or the next-hop-neighbor are failed. If so, then the packet is instead forwarded to a so-called FRR adjacency, which is primarily a different interface/next-hop-neighbor, but can also involve an additional 'steering' header, such as from SR-MPLS or [SRH]. These FRR adjacencies are pre-calculated by the routing control plane.

Note too, that these IP-FRR mechanisms are not limited to IP networks including SRv6 networks, but they are equally applicable to SR-MPLS networks.

In BIER-FRR, exactly the same principles of IP-FRR can be used, except that when a BIER packet is processed by the forwarding plane, a forwarding decision needs to be taken not only once (as in unicast), but once for every bit (destination) in the bitstring. For each bit, an interface/next-hop-neighbor is determined and a copy of the packet is sent to it. Unless such a copy was already made for a prior examined bit in the packets bitstring, hence avoiding more than one copy of the packet to each interface/next-hop-neighbor. Each packet will also receive a modified version of the bitstring which indicates the subset of bits in the received packets bitstring which are reachable across this same interface/next-hop-neighbor.

Beside being able to apply all the same mechanisms to BIER-FRR as available in IP-FRR, BIER-FRR specifically maintains the benefit of the most efficient hop-by-hop replication, even when a packet has to be put onto a BIER-FRR adjacency.

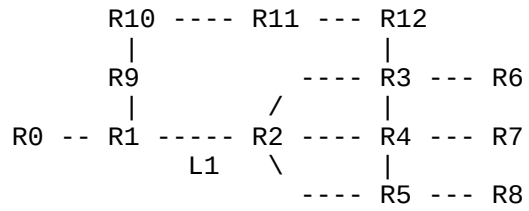

Figure 1: Example 1

pre-existing unicast LFA algorithms for LFA from R1 to R6 in case of R2 failure would for example calculate that it is sufficient to encapsulate a packet up to R10 but towards R8, that encapsulation endpoint would need to be R11 because R10 would send back packets to R8 to R9 because that is its shorted path (before reconvernce).

Consider some multicast packets from R0 would need to be received by a subset of (R6,R7,R8). They would be forwarded from R0 to R1, by R1 to R2, and from R2 to R3, R4, R5 to then reach R6, R7, R8. Assume

one multicast group, G1 would only need to go to R6. R2 would then have multicast forwarding state G1 -> (R3) to copy the packet only to R3. Another group G2 would need to go to R6, R7, R8. R2 would then have G2 -> (R3, R4, R5) meaning it would replicate packets to G1 to R3, R4 and R5.

Assumme R2 fails. FRR in R1 would recognize this as a failure of the link to R2 and be configured to assume this indicates failure of R2 and hence pre-established R2 node-protection FRR would kich in. Instead of sending a packet to R2, R1 it would send the packet onto the FRR packup path. With existing stateful multicast, FRR in R1 would have needed to pre-calculate and install three unicast backup path ("tunnels"). R1-...->R3, R1-...->R4 and R1-...->R5. R1 would then have to send three copies of each multicast packet, one into each backup path. If this was a packet for G1, then two of these copies would be in wain and be discarded by R4 and R5.

If a stateful multicast FRR solution would want to avoid such unnecessary excess traffic, it would need to install per group FRR state - a separate state for G1 and G2. In real networks there are not just 2 groups but thousands (#G). In the picture, R1 only has one outgoing interface to which to replicate. In a real network, it would have on average N outgoing interfaces/next-hops and this type of traffic optimized FRR for stateful multicast would create #G*N additional states on each hop. In addition, new protocol mechanisms to signal all this information would need to be developed, because today, R1 would not know which group would need to go to which subset of R3,R4,R5. Such protocol mechanisms where never developed and hence this type of stateful FRR for "node protection" (in this case for the failure of R2) is just theoretical.

With BIER-FRR, the solution becomes extremely simple for this example. The BIER-FRR control plane in R1 can calculate that the FRR adjacency for the case of R2 failing for all three possible destinations R6, R7, R8 can be R11. When R1 recognizes failure of (link to) R2, it would send the BIER packet with exactly the same bits encapsulated towards R11 (via R9, R10) and there normal BIER forwarding proceeds. According to the subset of bits for R6, R7, R8 set in the bitstring, R3,R4,R5 would make according copies. Like in FRR for unicast, it does not matter to BIER that the packet did arrive from R10 as opposed to R2, it is only revelant what destinations it has (in its bitstring).

R2 could not simply send the packt to R9, because it is not aware of R2 having failed, so its routing table would send packet packets for R6,R7,R8 towards R1 because that is the shortest path from R9 to R6,R7,R8. Likewise, R10 too would still send back packets towards R8 to R1. Only for R11 are paths not across the failed R2 the shortest path towards R6,R7,R8.

Note that the forwarding described for BIER-FRR here is LFA-based BIER-FRR in the terminology introduced later in this document.

2.2. Introduction

The protection level offered by BIER-FRR can be either link protection or node protection. This is like in IP FRR and both of them can co-exist.

Link protection is limited to safeguarding against link failures and is simpler to implement but will not be effective if a BFR neighbor 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.

As in IP FRR, one would typically use node protection and apply link-protection only when the next-hop node is the actual destination (BFER), so that only link-protection makes sense.

The followin introductory text explains the two fundamental approaches to BIER-FRR, one being tunnel-based BIER-FRR, and the other being LFA-based BIER-FRR.

2.2.1. Tunnel-based BIER-FRR

In this approach, a BIER forwarding entry to a BFR-NBR is replaced by an (IP)-FRR protected adjacency which encapsulates the BIER packet into a unicast packet towards that BFR-NBR. This is exactly what classicial (non IP) FRR in MPLS was and is doing with RSVP-TE backup tunnels. Hence the name tunnel-based.

For example, assuming the topology of Figure 1 and R1 as the BFR where this form of FRR is to be installed. Usually a BIER forwarding entry towards R2 is simply an L2 adjacency for BIER packets. Instead, that L2 adjacency towards R2 is set to be a unicast encapsulation (IPv4, IPv6 or SR-MPLS) with a unicast IPv4/IPv6 destination of R2. In the case of SR-MPLS, the actual encapsulation will have an MPLS SID for the IPv4 and/or IPv6 address of R2.

Absent a failure of link L1, this change means simply an "unnecessary" encapsulation of BIER packets from R1 to R2 and decapsulation there.

If instead, the link L1 fails, the BIER packet on R1 will get encapsulated into the unicast packet, and once it is to be forwarded as unicast, the unicast forwarding will determine that the link L1 has failed and will hence forward the packet on the pre-calculated unicast FRR adjacency towards R2. In the Figure 1 topology this means that so-called TI-LFA (Topology Independent LFA) has to be used which can calculate that packets towards R2 (in case of L1 link failures) need to be encapsulated with a destination address of R11 - because sending the FRR packet only towards R9 or R10 would have it be returned towards R1. In other topologies, simple (non TI) LFA may suffice. In those cases, the FRR adjacency would simply be to send the unicast packet to another next-hop without additional encapsulation.

2.2.1.1. Limitations

Tunnel-based BIER-FRR is only applicable to link-protection but not node-protection. This is because Tunnel-based BIER-FRR is based on only having to replace BIER (L2) adjacencies with already existing FRR protected unicast adjacencies. Tunnel-based BIER-FRR is not applicable to node-protection because a failure of the next-hop means that some other router in the network would need to replace the replications of BIER packet that this failed router could have done, and there is no unicast FRR function that would do this.

The traffic efficiency of Tunnel-based BIER-FRR is typically lower than that of LFA-based BIER-FRR because the traffic needs to be tunneled towards the BIER-NH and from there back as a regular BIER packet. In the topology of Figure 1 a Tunnel-based BIER-FRR packet from R1 to R6 would go R1-R9-R10-R11-R11-R12-R3-R2(decap)-R3-R6, whereas as described in before, LFA-based would have it simply go R1-R9-R10-R11-R11-R12-R3-R6 (not via R2).

Because Tunnel-based BIER-FRR requires FRR protected adjacencies to a next-hop, unicast FRR needs to be configured to not only provide node-protection adjacencies but also link-protection adjacencies to all direct neighbors that are also BIER neighbors. This should be standard unicast FRR deployment though when node-protection is used.

2.2.1.2. Benefits

Tunnel-based BIER-FRR is least demanding to the BIER forwarding plane and should be most easy to implement on any BFR already supporting unicast LFA. No changes to the BIER forwarding tables (BIFT) are required other than replacing the BFR-NBR entries.

For both unicast and BIER, Tunnel-based FRR is logically something that can happen at the link-layer, not impacting the forwarding table, and hence not requiring to scale with the size of the forwarding table. Instead of sending a packet to a neighbor in a connected subnet, FRR protection adds a check whether the interface to the subnet is up and/or the neighbor is alive, and if not, then packet is instead redirected to another interface (or neighbor on the same interface) with an additional encapsulation.

The interoperability requirement for Tunnel-based is solely that the tailends of pre-calculated unicast FRR adjacencies are also BFR and will accept to forward unicast FRR encapsulated BIER packets as BIER again after decapsulation.

When the total amount of multicast/BIER traffic is insignificant compared to the overall bandwidth in the network, the traffic inefficiencies of Tunnel-based BIER-FRR may be irrelevant, but the inability to provide node-protection may be insufficient.

Tunnel-based BIER-FRR does not necessarily have to be built from (IP)-LFA based unicast adjacencies. It can equally be built from RSVP-TE tunnels with FRR protection (or any other unicast FRR protected tunnels), which may be attractive if those already exist in the network.

2.2.2. LFA-based BIER-FRR

In LFA-based BIER-FRR as in (unicast) IP FRR, the goal is to directly get packets in the failure situation as directly as possible to the ultimate destinations insted of back to the next-hop router as in Tunnel-Mode.

The mechanisms are exactly the same as in unicast IP FRR, specifically that it requires to establish for every destination in the unicast or BIER forwarding table (BIFT) new "FRR protected" adjacencies", and hence it requires additional forwarding resources in the order of forwarding table entries. These can often be optimized, but it does cost overall more forwarding plane resources than Tunnel-based BIER-FRR.

BIER-FRR will often be able to re-use pre-existing (unicast) IP FRR adjcancies in the BIER forwarding table (BIFT), but there are some optimizations and differences to observe. These are described here.

2.2.2.1. Principles of (unicast) LFA

The principles of unicast LFA are to determine a path from the node determining an adjacent link or node failure towards some desting assuming that the intermediate hops on this alternative path still have the routing tables in which there is no failed link or node. The problem with this principle is that nodes along the best alternative path would not send packets along the intended path but back to where they come from because without a failure, that could be the best path for them.

Using the example of Figure 2, and examining the case where R1 tries to calculate unicast link-protection FRR for the failure of link L1. It calculates for each of the destinations of interest, R6, R7 and R8 the shortest path. Each of these path goes along

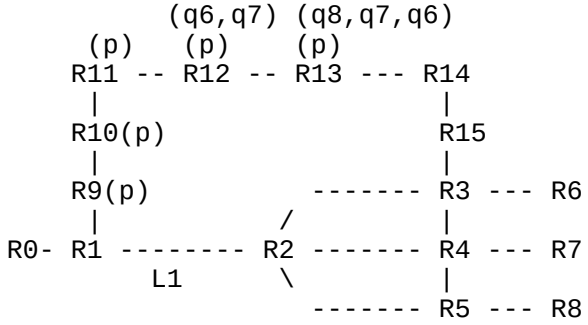

Figure 2: Example 2

LFA characterizes routers along an FRR path by whether they are part of the so-called P-space and/or Q-space [RFC7490].

The p)re-converge space are all routers that can be reached from the router of interest (R1) in the pre-convergence routing table without going through the failed link (or node). For example from R1 under failure of L1 or R2, the routers R9, R10, R11, R12 and R13 are in p-space. R14 is not in p-space anymore because assuming all links have equal cost, the path from R9 to R14 is ECMP through either R1 and the failed link/node or R10. And R1 can not know which route R9 is using. So packets from R1 to R14 via R9 might be returned by R9. And R15 is most easily not in p-space anymore, because R9 would be guaranteed to send packets from R1 towards R15 back to R1.

q-space is the reverse p-space (as q is typographically a reverse p), indicating whether a router can reach the ultimate destination without going through the failed link/node in the pre-converged routing tables.

R13 is in the q-space for R8, whereas R12 is not. It has ECMP paths via R13 and R11, and the path via R11 would go via the failed link/ node (L1/R2). However, R12 is in q-space for R6 and R7 though.

The FRR logic is now simply that if packets can not be sent directly to an alternative next-hop without encapsulation, then they need to be encapsulated and sent to a router that is both in p-space and q-space for the final destination. Because R9 is not in q-space for any of the destinations of interest in BIER, directly sending packets to R9 does not work - whether link failure of L1 is assumed, or node failure of R2.

Instead, R1 would create for R6 an FRR adjacency encapsulating the packet with an appropriate unicast encapsulation destined to R12 or R13 and make R9 the next-hop for this encapsulated packet. Likewise for R7. For R8, it would need to use R13 as the encapsulation destination. Most likely R1 would pick R12 for R6 and R7 because the path with encapsulation is typically choosen to be as short as possible but with most often insignificant differences.

2.2.2.2. Adjustments for BIER-FRR

If the same logic as for unicast is choosen for BIER-FRR, it does make a significant difference whether R12 or R13 are picked as FRR for R6,R7 or if R13 is picked. For BIER, the same FRR adjacency "encapsualte towards R13" would mean that a BIER packet destined to R6 and R8 could be sent once across that FRR adjcency. If instead R6,R7 had the "encapsulate towards R12" FRR adjacency, and R8 "encapsulate towards R13", then BIER would need to send two copies of that BIER packet to R6,R8: once encapsulated for R6 towards R12 and once encapsulated for R8 towards R13. In this case, if L1 or R2 are failling, R1 would need to send two copies of the same packet instead of one if R13 was choosen as the pq-router for R6, R7 and R8.

So, while the principle of unicast FRR holds true, BIER makes it more beneficial to choose longer encapsulated paths to minimize unnecessary replications. On the other hand (not shown by example topology), longer encapsulation paths may also result in longer paths from the point of decapsulation, and there may be a more complex weighting of multiple copies across the same encapsulated path versus such longer paths after the decapsulation.

2.2.2.3. Topology Independent FRR

A simple encapsulation towards a router in pq-space may not suffice. This is the same for BIER-FRR as it is for unicast FRR.

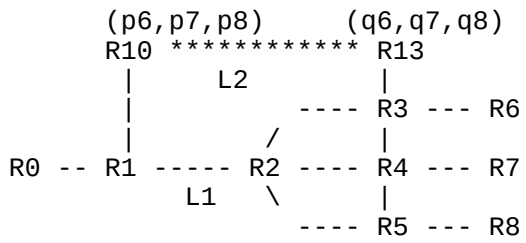

Figure 3: Example 3

Consider in Figure 3 that a link between R10 and R13 has a significantly higher cost than the other links in the topology. In result, R10 is in p-space, R13 in q-space and hence there is no single router in pq-space that R1 could encapsulate packets to to reach R6, R6, R8 under failure of L1 or R2. If R1 would want to send packets to ANY node in the topology across R10 under failure of L1 or R2, R10 would simply send them back because of the high cost of L2.

To get packets across L2, a steering encapsulation is required that does ignore the actual IGP routing table. For example with SR-MPLS or SRv6 this could be a two-hop SID list, once to R10 and then to an adjacency-SID for R13. An adjacency SIP would only be possible to interpret by R10 and indicate to R10 to forward the packet to the interface of R13 connecting to R10. This type of FRR solutions are documented in [RFC9855]. There are no additional considerations for BIER-FRR.

2.2.2.4. Partitioned q-space

Even though the big advantage of BIER-FRR with node protection is that it allows most often and as shown in the prior example Figure 2 to send out just one copy of a packet to an FRR adjacency instead of to the primary adjacency, there are also cases that more than one copy need to be sent. This is always true in BIER-FRR node protection, when a BIER packet though the primary adjacency would have to reach multiple destinations (BFER) and the q-spaces for those destinations are non-overlapping (partitioned).

In Figure 4, R1 is the PLR and R2 is the node assumed to fail. R1 would have a BIFT in which R5, R6, R7, R8 are reachable via R2, but in case of a node failure of R2, there is no single node which is in q-space for all of R5, R6, R7, R8 (shown as p5, p6, p7, p8). R12 for example is only in p-space for R1 but not in q-space for any destination because it would go via R1 to reach any of R5, R6, R7, R8.

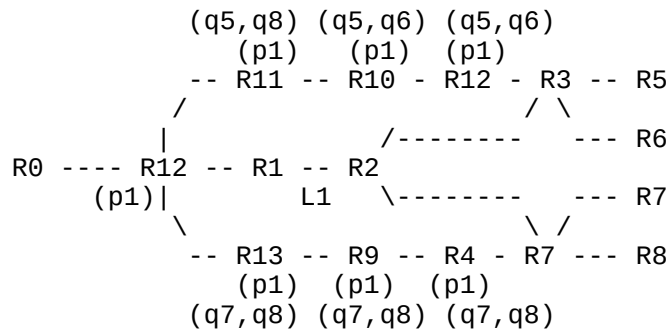

Figure 4: Example 4

In result, R1 does need to create an FRR adjacencies for R5 and R6 towards R11 and one for R7 and R8 towards R13. If then a packet comes towards all of R5, R6, R7, R8, R1 will create one copy towards R11 for R5 and R6, and another copy towards R13 for R7 and R8. This creates double the amount of originally received traffic from R12 back towards R12, but from there on, it again has only one copy on each link as opposed to any non-BIER FRR mechanism which would still carry two copies for example one for R5 and another for R6.

See Appendix A.1 for discussions of non-working enhancements.

2.2.3. Re-use of routing underlay FRR adjacencies

If the routes in the network are the same for BIER as they are for unicast routing from which FRR is being calculated, then the FRR adjacencies created for unicast can equally be re-used for BIER.

For example, when creating BIER-FRR entries in the BIFT, the BIER control plane could simply look up for every BFER the IP FRR information for its BFR-prefix, which is the IP address of the BFER. The FRR adjacency is then the same that could be used in the BIFT. This does of course not take care of the optimization opportunities for BIER discussed in Section 2.2.2.2. But it equally works for the partitioned q-space (Section 2.2.2.4).

What always needs to be set up explicitly for BIER with LFA-based BIER-FRR is triggering of the FRR action from the BIER forwarding (BIFT), because no unicast packet is involved at that point.

2.2.4. Conceptual BIER forwarding with FRR

[RFC8279], Figure 3, specifies the BIER forwarding conceptually through the structure of the Bit Index Forwarding Table (BIFT).

Figure 5 shows an excerpt of the BIFT of R1 in Figure 4, using slightly simplified representations to make it easiert to understand. Every BFER has a BIFT forwarding entry (row). The forwarding 'address' is the BIFT-id and is the BFRs bit in the packets bitstring. The BFR-NBR (BFR Neighbor) is the adjacency to which a copy for the BFR-id needs to be sent.

The F-BM is the mechanism in BIER by which it avoid sending duplicates. It is a bitstring including the bit of all BFR-id that can be reached when a copy to this lines BFR-NBR is made. o

Assume a BIER packet to R5,R7,R13. Ithas a bitstring with the appropriate bit for each of these destinations set. It first is matched against the one for R5 (first forwarding entry). It matches. A copy of the packet is made and sent towards R1. This copy will have all bits in its bitstring removed which are for R5,R6,R7,R8 - because these are the BFER that can be reached (shortest path). The processing of the remaining forwarding plane entries continues with a packet where those four bits are removed. Hence the rows for R6-R12 will not match, but the forwarding entry for R13 will match again - and a second packet copy will be made following the same rules.

BFR-id	F-BM	BFR-NBR
R5	R5,R6,R7,R8	R2
R6	R5,R6,R7,R8	R2
R7	R5,R6,R7,R8	R2
R8	R5,R6,R7,R8	R2
R11	R11,R12,R13	R12
R12	R11,R12,R13	R12
R13	R11,R12,R13	R12

Figure 5: BIFT Example 4 R1 (exerpt)

For FRR, logically the BIFT will need to have multiple forwarding entries that need to behave differently under the event of one specific link or node failure. For example Figure 6 shows the R1 BIFT excerpt under failure of R2.

BFR-id	F-BM	BFR-NBR
R5	R5,R6	encap(dest:R11) NH:R12
R6	R5,R6	encap(dest:R11) NH:R12
R7	R7,R8	encap(dest:R13) NH:R12
R8	R7,R8	encap(dest:R13) NH:R12
R11	R11,R12,R13	R12
R12	R11,R12,R13	R12
R13	R11,R12,R13	R12

Figure 6: BIFT Example 4 R1 under R2 failure

In unicast LFA, there is alwys only one forwarding entry for a packet that needs to be looked up and processed. And in result also only the FRR adjacency to be considered: The one for the next-hop of the forwarding entry.

In BIER-FRR with LFA-mode, this is not the case. Not only does a node failure impact potentially multiple BIFT entries, but one and the same BIFT entry may have multiple conflicting F-BM and BFR-NBR depending on which BFR-NBR has just failed.

TBD: shuold make an example of this problem.

Therefore, this memo will discuss some possible optimizations to support this complexity.

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 this subsection, 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. 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 7.

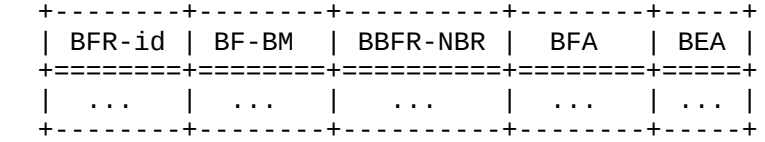

Figure 7: Structure of the backup BIFT.

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 8.

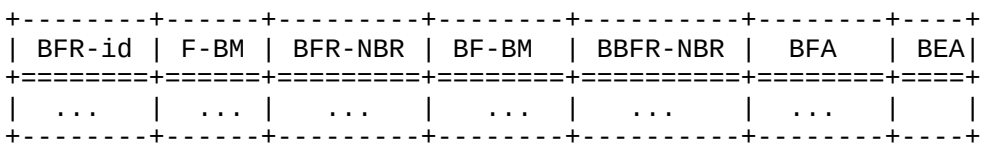

Figure 8: Structure of a single extended BIFT including backup forwarding entries.

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 9 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.

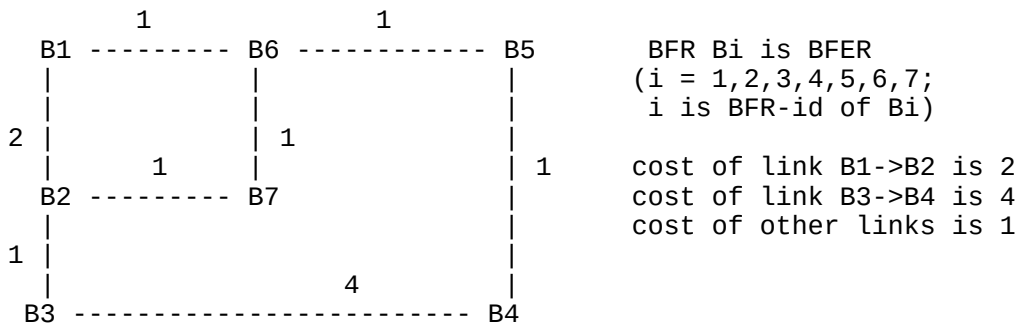

Figure 9: BIER network example.

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 10: B1’s primary BIFT.

A backup BIFT for B1 in the example of Figure 2 is given in Figure 10. 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 11: 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 9, 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 12: B1's backup BIFT for tunnel-based BIER-FRR with link protection.

Figure 12 illustrates B1's backup BIFT for tunnel-based BIER-FRR with link protection in the BIER network example depicted in Figure 9. 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 13: B1's backup BIFT for tunnel-based BIER-FRR with node protection.

Figure 13 illustrates B1's backup BIFT for tunnel-based BIER-FRR with node protection in the BIER network example provided in Figure 9. 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 9, 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 13 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 14 and Figure 15 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 14: 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 15: 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 16 illustrates B1’s backup BIFT for LFA-based BIER-FRR with node protection, using the network example provided in Figure 9.

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 16: 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 14. 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.

Acknowledgments

The authors would like to thank Daniel Merling, Jeffrey Zhang, Tony Przygienda and Shaofu Peng 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.

References

Normative References

[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/rfc/rfc5286>.

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

[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/rfc/rfc7431>.

[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/rfc/rfc7490>.

[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/rfc/rfc8279>.

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", 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>.

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

[mLDP] Wijnands, IJ., Ed., Minei, I., Ed., Kompella, K., and B. Thomas, "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", RFC 6388, DOI 10.17487/RFC6388, November 2011, <https://www.rfc-editor.org/rfc/rfc6388>.

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

[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/rfc/rfc4090>.

[RFC9855] Bashandy, A., Litkowski, S., Filsfils, C., Francois, P., Decraene, B., and D. Voyer, "Topology Independent Fast Reroute Using Segment Routing", RFC 9855, DOI 10.17487/RFC9855, October 2025, <https://www.rfc-editor.org/rfc/rfc9855>.

[RSVP-TE] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <https://www.rfc-editor.org/rfc/rfc3209>.

[SRH] Filsfils, C., Ed., Dukes, D., Ed., Previdi, S., Leddy, J., Matsushima, S., and D. Voyer, "IPv6 Segment Routing Header (SRH)", RFC 8754, DOI 10.17487/RFC8754, March 2020, <https://www.rfc-editor.org/rfc/rfc8754>.

Appendix A. Non-working FRR options

A.1. BIER-in-BIER encapsulation

Figure 17 is again the example shown in Figure 4. One option not discussed in before - because it can reasonably not be made to work - is to attempt using BIER-in-BIER encapsulation to improve over the solution described.

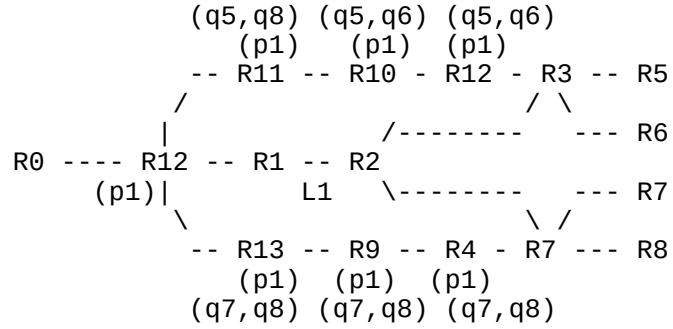

Figure 17: Example 4

One way how one could attempt to avoid having to send out two copies from R1 under the failure of R2, one towards R11 to then reach R5, R6 and one towards R13 to then reach R7, R8 is to consider encapsulating the packet into a new BIER header which for example has as

destinations (BFER) R11 and R13. When R11 and R13 respectively receive this packet, they decapsulate it, encounter the original BIER packet (that was FRR redirected by R1), and continue to forward it to it’s ultimate destinations.

Unfortunately, this topology also shows how this would fail. When R11 would process the original BIER packet and encounter the bits for either R7 and.or R8, then it would send a copy back to R12 to reach R7, R8. That copy would again reach R1, and in result in another FRR’ed copy of the same packet. This duplication would only stop due to TTL expiry.

Arguably, if the BIER encapsulation was choosen not to send copies towards R11, R13, but instead R12, R4, then this looping would not happen because the copy for R7, R8 from R12 would instead be sent towards R3, which would send it towards R2 which assumably has failed and hence the packet copy from R3 would be dropped.

However, in reality FRR would not only be set up on R1 towards R2, but equally on R3 and R7, so both of them would equally create FRR packets when they recognize R2 to be down.

Appendix B. Changelog

[RFC-Editor: Please remove this section].

B.1. rev 11 - sent back from IESG to WG

Intermediate version. Added primarily the Overview Chapter to better introduce the concepts and benefits of BIER-FRR to readers (customers et all) only aware of unicast FRR, and to better explain unicast FRR concepts (p, q space) to readers primarily aware of BIER.

Changes / optimizations in further parts of the document yet.

B.2. rev 11 - sent back from IESG to WG

Triage of IESG review feedback. Fixed the following core / simple feedback. See TBD section below for the missing IESG and directorate review feedback that will need to be folded into the next rev’s.

Brought document into kramdown format for easier editing.

Rewrote abstract to answer Roman Danyliw / テ詠ic Vyncke / Brian Haberman questions about scope (framework) and intended status (informational) of document.

Changed set of authors to meeth 5-authors max requirements. Changed authors to contributors.

Removed RFC2119 boilerplate because as a framework, this document does not use RFC2119 language (Mike Bishop).

Resolved Eric Vyncke Abstract text convern (removed). But see TBD for more work on refining text required.

Resolved expanding LFA on first use.

Ketan: Please remove extra "." I saw a few other similar instances in the document.

Ketan: minor: perhaps s/reconvergence/control plane reconvergence .

Ketan: fixed "persistent failure" text.

Ketan: In the following ? ... perhaps "In this subsection," ?

B.3. Resolved IESG discuss / comments before rev 11.

Added text for BFD referring to RFC5880 (prior no use of RFC5880 reference).

EVyncke: s/without encapsulation in a tunnel header/without encapsulation in a tunnel/

EVyncke: s/link layer technology/link-layer technology/

B.4. TBD

Fold in RTGDIR feedback (eckert).

Fold in unanswered questions from INTDIR review (Haberman): https://datatracker.ietf.org/doc/review-ietf-bier-frr-08-intdir-telechat-haberman-2025-06-03/)

Section 6.1: Add text about PMTU when using tunnels (Evyncke DISCUSS). Although: RFC7490 which explicitly require stunnels also does not address tunnels MTU issues. Maybe attempt to declare MTU problem out of scope given how we’re "just" doing something similar for BIER that several unicast RFCs are doing - without addressing MTU.

Check/remove unused references. Add text explaining benefits of reading reference "Performance Comparison of Resilience Mechanisms for Stateless Multicast Using BIER" (aka: which pieces relevant to this draft does this research paper cover).

Add text about egress protection (Aka: node protection against BFER failure), reference I-D.chen-bier-egress-protect (Roman Danilyv).

EVyncke: I fail to see the logical link between Typically, BIER packets are forwarded without an outer IP header. and the consequence if a link or node failure occurs, the corresponding BFR Neighbor (BFR-NBR) becomes unreachable. Strongly suggest adding some explanations. Answer: linkage is that BIER can not automatically use IP FRR but has to deal with unreachability events itself. But even if BIER was using per-hop IP/MPLS encap to rely on IP/MPLS FRR, then the result would not be as good as "direct" BIER-FRR. Text rewrite requires some restructuring.

EVyncke: Should a reference be provide for SR in If segment routing is employed ?

Evyncke: The last paragraph does not mention the ’explicit’ forwarding action, is it on purpose ? If so, the read will welcome an explanation.

Ketan: major: Please provide a reference or explanation of "normal BIER-LFAs". Did you mean RFC5286? Same goes for the other types - please provide references. TBD because text needs more structured rewrite of aligning the BIER behavior with the pre-defined unicast FRR terminology / cases.

Ketan: Is that IP-TI-LFA ? Same Q for TI-BIER-LFA. Need more structured rewrite of text...

Ketan: Isn’t that 100% theoretical? Practically, there are limits of platform and implementations. Also, all routers should be BFRs. Answer: No, should be as practically applicable as unicast TI-LFA is. There may be othrer platform limitations for BIER-FRR though.

DebCooley: The word ’tunnel’ is used many times in this draft. There is no definition of what is meant by tunnel(s), I have to assume that they are not for security purposes. If they are specific types of tunnels, e.g. MPLS or other security tunnel options (IPsec), then it would be nice to have that defined. Yes: Need to define "tunnel" for the purpose of this ocument as an encapsulation of BIER packets into some unicast header that allows forwarding of the packet to a remote BFR. On the other hand, RFC like RFC7490 (RLFA in unicast) uses "tunnel" without explaininf/defining it.

Ketan: References to respective RFCs related to different types of LFA/FRR unicast mechanisms would be helpful in this section as well. Yes!

All of Mohammeds IESG review (sorry, ran out of time).

Contributors

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

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

Toerless Eckert
Futurewei
Email: tte@cs.fau.de