
Stream:	Internet Engineering Task Force (IETF)				
RFC:	9860				
Category:	Informational				
Published:	September 2025				
ISSN:	2070-1721				
Authors:	Y. Liu	M. McBride	Z. Zhang	J. Xie	C. Lin
	<i>China Mobile</i>	<i>Futurewei</i>	<i>ZTE</i>	<i>Huawei</i>	<i>New H3C Technologies</i>

RFC 9860

Multicast-Only Fast Reroute (MoFRR) Based on Topology Independent Loop-Free Alternate (TI-LFA) Fast Reroute

Abstract

This document specifies the use of Topology Independent Loop-Free Alternate (TI-LFA) mechanisms with Multicast-only Fast Reroute (MoFRR) for Protocol Independent Multicast (PIM). The TI-LFA provides Fast Reroute (FRR) protection for unicast traffic in IP networks by precomputing backup paths that avoid potential failures. By integrating TI-LFA with MoFRR, this document extends the benefits of FRR mechanisms to multicast traffic, enabling enhanced resilience and minimized packet loss in multicast networks. The document outlines an optional approach to implement TI-LFA in conjunction with MoFRR for PIM, ensuring that multicast traffic is rapidly rerouted in the event of a failure.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are candidates for any level of Internet Standard; see Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc9860>.

Copyright Notice

Copyright (c) 2025 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
1.1. Terminology	3
2. Problem Statement	4
2.1. LFA for MoFRR	4
2.2. TI-LFA for MoFRR	5
3. A Solution	5
3.1. Overview	5
3.2. Procedure	6
4. Illustration	7
5. Management and Operational Considerations	10
6. IANA Considerations	10
7. Security Considerations	10
8. References	11
8.1. Normative References	11
8.2. Informative References	13
Contributors	13
Authors' Addresses	13

1. Introduction

Multicast-only Fast Reroute (MoFRR), as defined in [RFC7431], offers a mechanism to reduce multicast packet loss in the event of node or link failures by introducing simple enhancements to multicast routing protocols, such as Protocol Independent Multicast (PIM) [RFC7761].

However, the MoFRR mechanism [RFC7431], which selects the secondary multicast next hop based solely on the Loop-Free Alternate (LFA) FRR defined in [RFC7431], has limitations in certain multicast deployment scenarios (see Section 2).

This document introduces a new mechanism for MoFRR using FRR for Topology Independent Loop-Free Alternate (TI-LFA) [RFC9855]. Unlike conventional methods, TI-LFA is independent of network topology, enabling broader coverage across diverse network environments. This mechanism is applicable to PIM networks, including cases where PIM operates directly over IP in Segment Routing (SR) networks.

The TI-LFA mechanism is designed for standard link-state Interior Gateway Protocol (IGP) shortest path and SR scenarios. For each destination advertised by the IGP in a network, TI-LFA pre-installs a backup forwarding entry for the protected destination, which is ready to be activated upon the detection of a link failure used to reach that destination. This document leverages the backup path computed by TI-LFA through the IGP as a secondary Upstream Multicast Hop (UMH) for PIM. By sending PIM secondary Join messages hop by hop on the TI-LFA backup path, a FRR backup path can be created for PIM multicast.

The techniques described in this document are limited to protecting links and nodes within a link-state IGP area. Protecting domain exit routers and/or links attached to other routing domains is beyond the scope of this document. All the Segment Identifiers (SIDs) required are contained within the Link State Database (LSDB) of the IGP.

The approach does not alter the existing management and operation of LFA, TI-LFA, and Remote LFA (RLFA) [RFC7916] [RFC8102] [RFC9855]. Additionally, during post-failure reconvergence, micro-loops [RFC5715] may form due to transient forwarding inconsistencies across routers. PIM micro-loop prevention is out of scope for this document.

Note that this document introduces an optional approach for backup Join paths, designed to enhance the protection scope of existing multicast systems. It is fully compatible with current protocol implementations and does not necessitate any changes to the protocols or forwarding functions on intermediate nodes. All nodes along the backup Join paths must support the Reverse Path Forwarding (RPF) Vector Attribute as defined in [RFC5496] and [RFC7891]. If there is a choice between vector and non-vector Join messages on the intermediate nodes, the non-vector option should be prioritized, which implies that protection paths will remain inactive. This document does not modify the handling of conflicts in PIM Join messages. For guidance on conflicts in Join attributes, please refer to Section 3.3.3 of [RFC5384].

1.1. Terminology

This document utilizes the terminology as defined in [RFC7431] and incorporates the concepts established in [RFC7490]. The definitions of individual terms are not reiterated within this document.

2. Problem Statement

2.1. LFA for MoFRR

Section 3 of [RFC7431] specifies that a secondary UMH in PIM for MoFRR is a Loop-Free Alternate (LFA). However, the conventional LFA mechanism requires that at least one neighbor's next hop to the destination node is a loop-free next hop. Due to existing limitations of the LFA mechanism in network deployments, such as topology dependency and incomplete destination coverage, the LFA mechanism can only be deployed in certain network topology environments. In specific network topologies, the secondary UMH cannot be computed in PIM for MoFRR, preventing PIM from establishing a standby multicast tree, and thus preventing the implementation of MoFRR protection. Consequently, the MoFRR functionality [RFC7431] in PIM is applicable only in network topologies where LFA is feasible.

The limitations of the MoFRR applicability [RFC7431] can be illustrated using the example network depicted in Figure 1. In this example, the metric of the R1-R4 link is 20, the metric of the R5-R6 link is 100, and the metrics of the other links are 10. All link metrics are bidirectional.

For multicast source S1 and receiver R, the primary path of the PIM Join packet is R3->R2->R1, and the secondary path is R3->R4->R1, which corresponds to the LFA path of unicast routing. In this scenario, MoFRR [RFC7431] operates effectively.

For multicast source S2 and receiver R, the primary path of the PIM Join packet is R3->R2. However, an LFA does not exist. If R3 sends the packet to R4, R4 will forward it back to R3 because the IGP shortest path from R4 to R1 is R4->R3->R2. In this case, MoFRR [RFC7431] cannot calculate a secondary UMH. Similarly, for multicast source S3 and receiver R, the MoFRR mechanism [RFC7431] is ineffective.




Figure 1: Example Network Topology

2.2. TI-LFA for MoFRR

The alternate path provided by the TI-LFA mechanism is represented as a segment list, which includes the Node SID of the P-space node and the Adjacency SIDs of the links between the P-space and Q-space nodes. When a remote PQ node exists in both P-space and Q-space, the segment list requires only the PQ node's Node SID.

PIM can look up the corresponding node's IP address in the unicast route base according to the Node SID and the IP addresses of the endpoints of the corresponding link in the unicast route base according to the Adjacency SIDs. However, multicast protocol packets cannot be directly forwarded along the path of the segment list.

To establish a standby multicast tree, PIM Join messages need to be transmitted hop by hop. However, not all nodes and links on the unicast alternate path are included in the segment list. If PIM protocol packets are transmitted solely in unicast mode, they effectively traverse the unicast tunnel like unicast traffic and do not pass through the intermediate nodes of the tunnel. Consequently, the intermediate nodes on the alternate path cannot forward multicast traffic because they lack PIM state entries. PIM must create entries on each device hop by hop, generating an incoming interface and an outgoing interface list, to form a complete end-to-end multicast tree for forwarding multicast traffic. Therefore, simply sending PIM Join packets using the segment list, as done with unicast traffic, is insufficient to establish a standby multicast tree.

3. A Solution

3.1. Overview

A secondary UMH serves as a candidate next hop that can be used to reach the root of a multicast tree. In this document, the secondary UMH is derived from unicast routing, utilizing the segment list computed by TI-LFA.

The path information from the segment list is incorporated into the PIM packets to guide hop-by-hop RPF selection. The IP address corresponding to the Node SID can be used as the segmented root node, while the IP addresses of the interfaces at both endpoints of the link associated with the Adjacency SID can be used as the local upstream interface and upstream neighbor.

[RFC5496] defines the PIM RPF Vector Attribute, which can carry the node's IP address corresponding to the Node SID. Additionally, [RFC7891] defines the Explicit RPF Vector, which can carry the peer's IP address corresponding to the Adjacency SID.

For instance, in the network illustrated in [Figure 1](#), the secondary path for the PIM Join packet towards multicast source S2 cannot be computed by MoFRR [RFC7431], as previously described. Using TI-LFA, R3 sends the packet to R4 while including an RPF Vector that contains the IP address of R1, which serves as R3's PQ node for the protected R3-R2 link. R4 then forwards the packet to R1 via the R1-R4 link according to the unicast route associated with the RPF Vector. R1 subsequently forwards the packet to R2, thus establishing the secondary path R3->R4->R1->R2.

Additionally, for multicast source S3 and receiver R, the primary path of the PIM Join packet is R3->R2->R5. Using TI-LFA, R3 sends the PIM Join packet to R7 while including two RPF Vectors:

- The first RPF Vector contains the IP address of R6, which serves as R3's P-node for the protected R3-R2 link.
- The second RPF Vector contains the interface IP addresses of R6 and R5, which serve as R3's Q-node for the protected R3-R2 link.

The first RPF Vector guides the PIM Join path through R3->R7->R6, while the second RPF Vector guides the PIM Join path through R6->R5, thereby establishing the secondary path R3->R7->R6->R5.

This document leverages the above RPF Vector standards, obviating the need for PIM protocol extensions. This approach allows the establishment of a standby multicast tree based on the segment list calculated by TI-LFA, thereby providing comprehensive MoFRR protection for multicast services across diverse network environments.

3.2. Procedure

Consider a segment list calculated by TI-LFA as (NodeSID(A), AdjSID(A-B)). Node A belongs to the P-space, and node B belongs to the Q-space. The IP address corresponding to NodeSID(A) can be retrieved from the local LSDB of the IGP and assumed to be IP-a. Similarly, the IP addresses of the two endpoints of the link corresponding to AdjSID(A-B) can also be retrieved from the local LSDB and assumed to be IP-La and IP-Lb.

Within the PIM process, IP-a is treated as the standard RPF Vector Attribute and added to the PIM Join packet. IP-La is considered the local address of the incoming interface, and IP-Lb is regarded as the address of the upstream neighbor. Consequently, IP-Lb can be included in the PIM Join packet as the Explicit RPF Vector Attribute.

The PIM protocol initially selects the RPF incoming interface and upstream neighbor towards IP-a and proceeds hop by hop to establish the PIM standby multicast tree until reaching node A. At node A, IP-Lb is treated as the PIM upstream neighbor. Node A identifies the incoming interface in the unicast routing table based on IP-Lb, and IP-Lb is used as the RPF upstream address for the PIM Join packet directed towards node B.

Upon receiving the PIM Join packet at node B, the PIM protocol, finding no additional RPF Vector Attributes, selects the RPF incoming interface and upstream neighbor towards the multicast source directly. The protocol then continues hop by hop to establish the PIM standby multicast tree, extending to the router directly connected to the source.

When a remote PQ node exists in both P-space and Q-space, the processing can be simplified to involve only node A.

The primary path of the PIM Join packet is R6->R2->R1, and the secondary path is R6->R5->R4->R3->R2->R1. However, the conventional LFA does not function properly for the secondary path because the shortest path to R2 from R5 (or from R4) still traverses the R6-R2 link. In this scenario, R6 must calculate the secondary UMH using the proposed MoFRR method based on TI-LFA.

According to the TI-LFA algorithm, the P-space and Q-space are illustrated in Figure 3. The TI-LFA repair path consists of the Node SID of R4 and the Adjacency SID of R4->R3. On the Segment Routing over MPLS (SR-MPLS) data plane, the repair segment list is (Label-R4, Label-R4-R3). On the Segment Routing over IPv6 (SRv6) data plane, the repair segment list is (SID-R4, SID-R4-R3).




Figure 3: P-Space and Q-Space

In the PIM process, the IP addresses associated with the repair segment list are retrieved from the IGP LSDB.

On the IPv4 data plane, the Node SID Label-R4 corresponds to IP4-R4, which will be carried in the RPF Vector Attribute. The Adjacency SID Label-R4-R3 corresponds to the local address IP4-R4-R3 and the remote peer address IP4-R3-R4, with IP4-R3-R4 carried in the Explicit RPF Vector Attribute.

On the IPv6 data plane, the End SID SID-R4 corresponds to IP6-R4, which will be carried in the RPF Vector Attribute. The End.X SID SID-R4-R3 corresponds to the local address IP6-R4-R3 and the remote peer address IP6-R3-R4, with IP6-R3-R4 carried in the Explicit RPF Vector Attribute.

Subsequently, R6 installs the secondary UMH using these RPF Vectors.




Figure 4: Forwarding PIM Join Packet Along Secondary Path (IPv4)

On the IPv4 data plane, the forwarding of the PIM Join packet along the secondary path is shown in [Figure 4](#).

R6 inserts two RPF Vector Attributes in the PIM Join packet: IP4-R4 of Type 0 (RPF Vector Attribute) and IP4-R3-R4 of Type 4 (Explicit RPF Vector Attribute). R6 then forwards the packet along the secondary path.

When R5 receives the packet, it performs a unicast route lookup of the first RPF Vector IP4-R4 and sends the packet to R4.

R4, being the owner of IP4-R4, removes the first RPF Vector from the packet and forwards it according to the next RPF Vector. R4 sends the packet to R3 based on the next RPF Vector IP4-R3-R4, as its PIM neighbor R3 corresponds to IP4-R3-R4.

When R3 receives the packet, as the owner of IP4-R3-R4, it removes the RPF Vector. The packet, now devoid of RPF Vectors, is forwarded to the source through R3->R2->R1 based on unicast routes.

After the PIM Join packet reaches R1, a secondary multicast tree, R1->R2->R3->R4->R5->R6, is established hop by hop for (S, G) using MoFRR based on TI-LFA.




Figure 5: Forwarding PIM Join Packet Along Secondary Path (IPv6)

On the IPv6 data plane, the forwarding of the PIM Join packet along the secondary path is illustrated in [Figure 5](#). The procedure is analogous to that of the IPv4 data plane.

When a remote PQ node exists in both P-space and Q-space, the processing can be simplified to involve only the PQ node. In this case, only a single RPF Vector needs to be carried, and all other processing steps remain unchanged.

5. Management and Operational Considerations

The management of the proposed approach is consistent with [\[RFC7916\]](#). However, in the operation of this approach, the node on the backup Join paths must have an independent configuration strategy for installing RPF Vector Attributes in the PIM Join packet and controlling the sending of this PIM Join message.

All nodes on the backup Join paths must be able to parse the PIM Join message with the RPF Vector Attribute. If the nodes do not understand the RPF Vector Attribute in the PIM Join packet, then they must discard the RPF Vector Attribute because failing to remove the RPF Vectors could cause upstream nodes to send the Join packet back towards these nodes causing loops.

If an administrator is manually specifying the path that the Join messages need to be sent on, it is recommended that the administrator computes the path to include nodes that support the Explicit RPF Vector and check that the state is created correctly on each node along the path. Tools like Mtrace [\[RFC8487\]](#) can be used for debugging and to ensure that the Join state is set up correctly.

6. IANA Considerations

This document has no IANA actions.

7. Security Considerations

This document does not introduce additional security concerns. It does not change the security properties of PIM. For general PIM - Sparse Mode (PIM-SM) protocol security considerations, see [\[RFC7761\]](#). The security considerations of LFA, RLFA, and MoFRR described in [\[RFC5286\]](#), [\[RFC7490\]](#), and [\[RFC7431\]](#) should apply to this document.

When deploying TI-LFA, packets may be sent over nodes and links they were not transported through before, potentially raising the following security issues:

1. Spoofing and false route advertisements

- Dependencies of LFA/RLFA/TI-LFA on routing information
 - LFAs depend on accurate routing information to determine alternate paths. If an attacker can inject false routing information (e.g., by spoofing link-state advertisements), it could cause the network to select suboptimal or malicious paths for LFAs.
 - RLFA and TI-LFA also depend on accurate routing information, particularly for determining the tunneling paths or explicit paths. False route advertisements could mislead the network into using insecure or compromised paths.

2. On-path attacks

- Use of alternate paths
 - By rerouting traffic through alternate paths, especially those that traverse multiple hops (as in RLFA and TI-LFA), the risk of on-path attacks increases if any of the intermediate routers on the alternate path are compromised.
 - TI-LFA, which uses explicit paths, might expose traffic to routers that were not originally part of the primary path, potentially allowing for interception or alteration of the traffic.

3. Confidentiality and integrity

- Traffic encapsulation
 - RLFA and TI-LFA involve encapsulating traffic, which may expose it to vulnerabilities if the encapsulation mechanisms are not secure. For instance, if IPsec or another secure encapsulation method is not used, an attacker might be able to intercept or alter the traffic in transit.
- Protection of explicit paths
 - TI-LFA relies on explicit paths that are typically defined using SR. If these paths are not properly protected, an attacker could manipulate the segment list to reroute traffic through malicious nodes.

4. Increased attack surface

- Extended topology
 - By introducing LFA, RLFA, and TI-LFA, the network increases its reliance on additional routers and links, thereby expanding the potential attack surface. Compromise of any router in these alternate paths could expose traffic to unauthorized access or disruption.

To address security issues [1](#) and [2](#) mentioned above, control plane protocols need to provide security protection. To mitigate the risks associated with false route advertisements and on-path attacks, it is recommended to use secure routing protocols (e.g., OSPFv3 with IPsec, IS-IS HMAC-SHA256, or PIM with IPsec) that provide authentication and integrity protection for routing updates.

To address security issues [3](#) and [4](#) mentioned above, these mechanisms need to run within a trusted network. The security of LFA, RLFA, and TI-LFA mechanisms heavily relies on the trustworthiness of the underlying routing infrastructure. As the solution described in the document is based on SR technology, readers should be aware of the security considerations related to this technology (see [\[RFC8402\]](#)) and its data plane instantiations (see [\[RFC8660\]](#), [\[RFC8754\]](#), and [\[RFC8986\]](#)).

8. References

8.1. 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/info/rfc5286>>.
- [RFC5384] Boers, A., Wijnands, I., and E. Rosen, "The Protocol Independent Multicast (PIM) Join Attribute Format", RFC 5384, DOI 10.17487/RFC5384, November 2008, <<https://www.rfc-editor.org/info/rfc5384>>.
- [RFC5496] Wijnands, IJ., Boers, A., and E. Rosen, "The Reverse Path Forwarding (RPF) Vector TLV", RFC 5496, DOI 10.17487/RFC5496, March 2009, <<https://www.rfc-editor.org/info/rfc5496>>.
- [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>>.
- [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>>.
- [RFC7761] 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/info/rfc7761>>.
- [RFC7891] Asghar, J., Wijnands, IJ., Ed., Krishnaswamy, S., Karan, A., and V. Arya, "Explicit Reverse Path Forwarding (RPF) Vector", RFC 7891, DOI 10.17487/RFC7891, June 2016, <<https://www.rfc-editor.org/info/rfc7891>>.
- [RFC7916] Litkowski, S., Ed., Decraene, B., Filsfils, C., Raza, K., Horneffer, M., and P. Sarkar, "Operational Management of Loop-Free Alternates", RFC 7916, DOI 10.17487/RFC7916, July 2016, <<https://www.rfc-editor.org/info/rfc7916>>.
- [RFC8402] Filsfils, C., Ed., Previdi, S., Ed., Ginsberg, L., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", RFC 8402, DOI 10.17487/RFC8402, July 2018, <<https://www.rfc-editor.org/info/rfc8402>>.
- [RFC8660] Bashandy, A., Ed., Filsfils, C., Ed., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing with the MPLS Data Plane", RFC 8660, DOI 10.17487/RFC8660, December 2019, <<https://www.rfc-editor.org/info/rfc8660>>.
- [RFC8754] 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/info/rfc8754>>.
- [RFC8986] Filsfils, C., Ed., Camarillo, P., Ed., Leddy, J., Voyer, D., Matsushima, S., and Z. Li, "Segment Routing over IPv6 (SRv6) Network Programming", RFC 8986, DOI 10.17487/RFC8986, February 2021, <<https://www.rfc-editor.org/info/rfc8986>>.
-

- [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, September 2025, <<https://www.rfc-editor.org/info/rfc9855>>.

8.2. Informative References

- [RFC5715] Shand, M. and S. Bryant, "A Framework for Loop-Free Convergence", RFC 5715, DOI 10.17487/RFC5715, January 2010, <<https://www.rfc-editor.org/info/rfc5715>>.
- [RFC8102] Sarkar, P., Ed., Hegde, S., Bowers, C., Gredler, H., and S. Litkowski, "Remote-LFA Node Protection and Manageability", RFC 8102, DOI 10.17487/RFC8102, March 2017, <<https://www.rfc-editor.org/info/rfc8102>>.
- [RFC8487] Asaeda, H., Meyer, K., and W. Lee, Ed., "Mtrace Version 2: Traceroute Facility for IP Multicast", RFC 8487, DOI 10.17487/RFC8487, October 2018, <<https://www.rfc-editor.org/info/rfc8487>>.

Contributors

Mengxiao Chen

New H3C Technologies

China

Email: chen.mengxiao@h3c.com

Authors' Addresses

Yisong Liu

China Mobile

China

Email: liuyisong@chinamobile.com

Mike McBride

Futurewei Inc.

United States of America

Email: michael.mcbride@futurewei.com

Zheng (Sandy) Zhang

ZTE Corporation

China

Email: zhang.zheng@zte.com.cn

Jingrong Xie

Huawei Technologies

China

Email: xiejingrong@huawei.com

Changwang Lin

New H3C Technologies

China

Email: linchangwang.04414@h3c.com