PIM Working Group
Internet Draft
Intended status: Standards Track
Expires: August 21, 2021

Yisong Liu
China Mobile
M. McBride
Futurewei
Z. Zhang
ZTE
J. Xie
Huawei

Feb 21, 2021

Multicast-only Fast Reroute Based on Topology Independent Loop-free Alternate Fast Reroute
draft-liu-pim-mofrr-tilfa-03

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

The list of current Internet-Drafts can be accessed at http://www.ietf.org/ietf/1id-abstracts.txt

The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html

This Internet-Draft will expire on August 21, 2021.

Copyright Notice

Copyright (c) 2021 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 (http://trustee.ietf.org/license-info) in effect on the date of

publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Abstract

Multicast-only Fast Reroute (MoFRR) has been defined in [RFC7431], but the selection of the secondary multicast next hop only according to the loop-free alternate fast reroute, which has restrictions in multicast deployments. This document describes a mechanism for Multicast-only Fast Reroute by using Topology Independent Loop-free Alternate fast reroute, which is independent of network topology and can achieve covering more network environments.

Table of Contents

1. Introduction

As the deployment of video services, operators are paying more and more attention to solutions that minimize the service disruption due to faults in the IP network carrying the packets for these services. Multicast-only Fast Reroute (MoFRR) has been defined in [RFC7431], which can minimize multicast packet loss in a network when node or link failures occur by making simple enhancements to multicast routing protocols such as Protocol Independent Multicast (PIM). But the selection of the secondary multicast next hop only according to the loop-free alternate fast reroute in [RFC7431], and there are limitations in multicast deployments for this mechanism. This

document describes a new mechanism for Multicast-only Fast Reroute using Topology Independent Loop-free Alternate (TILFA) fast reroute, which is independent of network topology and can achieve covering more network environments.

1.1. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

1.2. Terminology

This document use the terms defined in [RFC7431], and also use the concepts defined in [RFC7490]. The specific content of each term is not described in this document.

2. Problem Statement

In [RFC7431] section 3, the secondary Upstream Multicast Hop (UMH) of PIM for MoFRR is a loop-free alternate (LFA). However, the traditional LFA mechanism needs to satisfy at least one neighbor whose next hop to the destination node is an acyclic next hop, existing limitations in network deployments, and can only cover part of the network topology environments. In some network topology, the corresponding secondary UMH cannot be calculated, so PIM cannot establish a standby multicast tree and cannot implement MoFRR protection. Therefore, the current MoFRR of PIM is only available in the network topology applicable to LFA.

The remote loop-free alternate (RLFA) defined in [RFC7490] is extended from the LFA and can cover more network deployment scenarios through the tunnel as an alternate path. The RLFA mechanism needs to satisfy at least one node assumed to be N in the network that the fault node is neither on the path from the source node to the N node, nor on the path from the N node to the destination node. RLFA only has enhancement compared to LFA but still has limitations in network deployments.

[I-D.ietf-rtgwg-segment-routing-ti-lfa] defined a unicast FRR solution based on the TILFA mechanism. The TILFA mechanism can express the backup path with an explicit path, and has no constraint on the topology, providing a more reliable FRR mechanism. The unicast traffic can be forwarded according to the explicit path list as an alternate path to implement unicast traffic protection, and can achieve full coverage of various networking environments.

The alternate path provided by the TILFA mechanism is actually a Segment List, including one or more Adjacency SIDs of one or more links between the P space and the Q space, and the NodeSID of P space node. PIM can look up the corresponding node IP address in the unicast route according to the NodeSID, and the IP addresses of the two endpoints of the corresponding link in the unicast route according to the Adjacency SIDs, but the multicast protocol packets cannot be directly sent along the path of the Segment List.

PIM join message need to be sent hop-by-hop to establish a standby multicast tree. However, not all of the nodes and links on the unicast alternate path are included in the Segment List. If the PIM protocol packets are transmitted only in unicast mode, then equivalently the PIM packets are transmitted through the unicast tunnel like unicast traffic, and cannot pass through the intermediate nodes of the tunnel. The intermediate nodes of the alternate path cannot forward multicast traffic because there are no PIM state entries on the nodes. PIM needs to create entries on the device hop-by-hop and generate an incoming interface and an outgoing interface list. So it can form an end-to-end complete multicast tree for forwarding multicast traffic. Therefore, it is not possible to send PIM packets like unicast traffic according to the Segment List path and can only establish a standby multicast tree.

It is available in principle that the path information of the Segment List is added to the PIM packets to guide the hop-by-hop RPF selection. The IP address of the node corresponding to the NodeSID can be used as the segmented root node, and the IP addresses of the interfaces at both endpoints of the link corresponding to the Adjacency SID can be used directly as the local upstream interface and upstream neighbor, but there is currently no field in protocol packet to carry the explicit path specified by the Segment List. For the PIM protocol, the PIM RPF Vector attribute was defined in [RFC5496], which can carry the node IP address corresponding to the NodeSID. The Explicit RPF Vector was defined in [RFC7891], which can carry the peer IP address corresponding to the Adjacency SID, but if there are multiple same peer IP addresses corresponding to the Adjacency SID (i.e. anycast IP address), the upstream neighbor of RPF selection may be different from the actual upstream link corresponding to the Adjacency SID, which can make the PIM join path and the TILFA calculation path inconsistent.

3. Solution

An Upstream Multicast Hop (UMH) is a candidate next-hop that can be used to reach the root of the tree. In This document the secondary

UMH is based on unicast routing to find Segment List calculated by TILFA.

This document extends the PIM protocol, to establish the standby multicast tree according to the Segment List calculated by TILFA, and can achieve full coverage of various networking environments for MoFRR protection of multicast services.

Assume that the Segment List calculated by TILFA is (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 looked up in the local link state database of the IGP protocol, and can be assumed to be IP-a. The IP addresses of the two endpoints of the link corresponding to AdjSID(A-B) can also be looked up in the local link state database of the IGP protocol, and can be assumed to be IP-La and IP-Lb.

3.1. Secondary UMH Selection

In the procedure of PIM, IP-a can be looked as the normal RPF vector attribute and added to the PIM join packet. IP-La and IP-Lb can be looked as the RPF Vector attribute of the adjacency relationship, called Adjacency RPF Vector, which is a new type of PIM join attributes, and added to the PIM join packet too.

The PIM protocol firstly can select the RPF incoming interface and upstream towards IP-a, and can join hop-by-hop to establish the PIM standby multicast tree until the node A. On the node A, IP-Lb can be looked as one PIM neighbor. If there are multiple PIM neighbors with the same address IP-Lb, all of the corresponding local interfaces on the node A need to be checked. The interface that is the only one with the IP address IP-La can be looked as the RPF incoming interface. The node A can send the PIM join packet to the node B on the interface of IP-La, and IP-Lb is used as the RPF upstream address of the PIM join.

After the PIM join packet is received on the node B, the PIM protocol can find no more join attributes and select the RPF incoming interface and upstream towards the multicast source directly, and then can continue to join hop-by-hop to establish the PIM standby multicast tree until the router directly connected the source.

3.2. Extension Protocol Fields Conflict

PIM Adjacency RPF Vector attribute is newly defined in join attributes. If there are conflicts from multiple downstream PIM neighbors, the mechanism in [RFC5384] Section 3.3.3 can be used to

select a PIM downstream neighbor with a numerically smallest IP address. If at least two neighbors have the same IP address, the interface index MUST be used as a tie breaker.

4. Packet Format

This section describes the format of PIM join message extension introduced by this document.

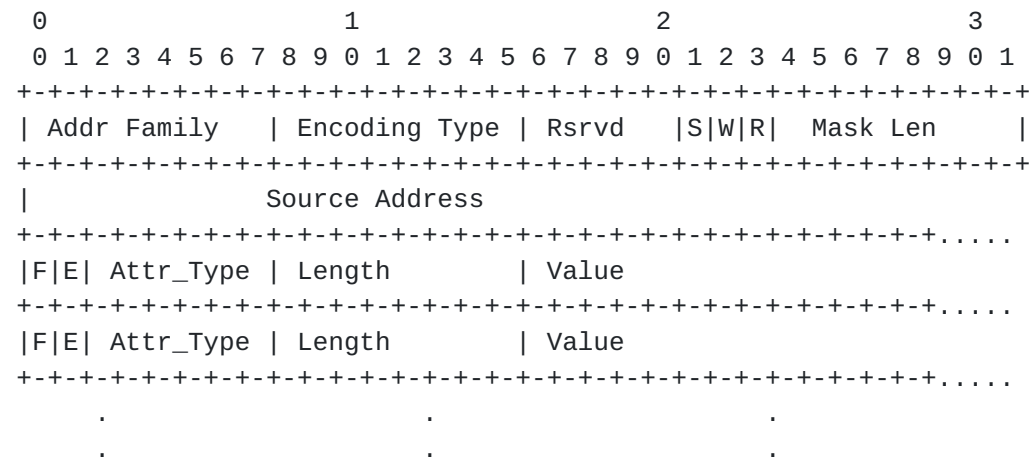

The original PIM join attribute already has been defined in [RFC5384]

Attr_Type:

0- Vector ;

4- Explicit RPF Vector ;

Other existing definitions are not related to RPF Vector Attribute.

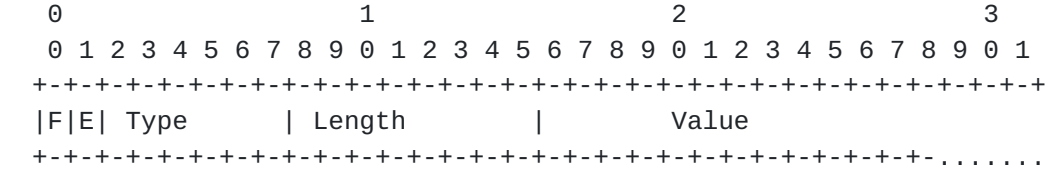

The definition of Adjacency RPF Vector attribute

F bit: 0, indicating that the unrecognized device does not forward the attribute

E bit: indicates the last join attribute

Type: TBD

Length: depends on the address family of Encoded-Unicast address, including the length of 2 addresses.

Value: Encoded-Unicast Address format defined in [RFC7761] Section 4.9.1, including 2 addresses. The first one indicates the address of the local interface, and the second one indicates the address of the peer interface. Only the case of the same address family is supported.

5. IANA Considerations

This document requests IANA to assign a registry for Adjacency RPF Vector in PIM Join Attribute and the Explicit Path TLV Node Address Sub TLV. The assignment is requested permanent for IANA when this document is published as an RFC. The string TBD should be replaced by the assigned values accordingly.

6. Security Considerations

For general PIM-SM protocol Security Considerations, see [RFC7761].

TBD

7. References

7.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC5384] Boers, A., Wijnands, I., and E. Rosen, "The Protocol Independent Multicast (PIM) Join Attribute Format", RFC 5384, November 2008

[RFC5496] Wijnands, IJ., Boers, A., and E. Rosen, "The Reverse Path Forwarding (RPF) Vector TLV", RFC 5496, March 2009

[RFC7431] Karan, A., Filsfils, C., Wijnands, IJ., Ed., and B. Decraene, "Multicast-Only Fast Reroute", RFC 7431, August 2015

[RFC7490] Bryant, S., Filsfils, C., Previdi, S., Shand, M., and N. So, "Remote Loop-Free Alternate (LFA) Fast Reroute (FRR)", RFC 7490, April 2015

[RFC7761] Fenner, B., Handley, M., Holbrook, H., Kouvelas, I.,Parekh, R., Zhang, Z., and L. Zheng, "Protocol IndependentMulticast - Sparse Mode (PIM-SM): Protocol Specification(Revised)", RFC 7761, March 2016

[RFC7891] Asghar, J., Wijnands, IJ., Ed., Krishnaswamy, S., Karan, A., and V. Arya, "Explicit Reverse Path Forwarding (RPF) Vector", RFC 7891, June 2016

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, May 2017

[I-D.ietf-rtgwg-segment-routing-ti-lfa] Litkowski, S., Bashandy, A., Filsfils, C., Francois, P., Decraene, B., and D. Voyer, "Topology Independent Fast Reroute using Segment Routing", draft-ietf-rtgwg-segment-routing-ti-lfa-06, work-in-progress, February 2021

7.2. Informative References

TBD

8. Acknowledgments

The authors would like to thank the following for their valuable contributions of this document:

TBD

Authors' Addresses

Yisong Liu
China Mobile
China
Email: liuyisong@chinamobile.com

Mike McBride
Futurewei Inc.
USA
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